

---

# Accommodating Picky Customers: Regret Bound and Exploration Complexity for Multi-Objective Reinforcement Learning

---

**Jingfeng Wu**

Department of Computer Science  
Johns Hopkins University  
Baltimore, MD 21218  
uuujf@jhu.edu

**Vladimir Braverman**

Department of Computer Science  
Johns Hopkins University  
Baltimore, MD 21218  
vova@cs.jhu.edu

**Lin F. Yang**

Department of Electrical and Computer Engineering  
University of California, Los Angeles  
Los Angeles, CA 90095  
linyang@ee.ucla.edu

## Abstract

In this paper we consider multi-objective reinforcement learning where the objectives are balanced using preferences. In practice, the preferences are often given in an adversarial manner, e.g., customers can be picky in many applications. We formalize this problem as an episodic learning problem on a Markov decision process, where transitions are unknown and a reward function is the inner product of a preference vector with pre-specified multi-objective reward functions. We consider two settings. In the online setting, the agent receives a (adversarial) preference every episode and proposes policies to interact with the environment. We provide a model-based algorithm that achieves a nearly minimax optimal regret bound  $\tilde{O}(\sqrt{\min\{d, S\} \cdot H^2 S A K})$ , where  $d$  is the number of objectives,  $S$  is the number of states,  $A$  is the number of actions,  $H$  is the length of the horizon, and  $K$  is the number of episodes. Furthermore, we consider preference-free exploration, i.e., the agent first interacts with the environment without specifying any preference and then is able to accommodate arbitrary preference vector up to  $\epsilon$  error. Our proposed algorithm is provably efficient with a nearly optimal trajectory complexity  $\tilde{O}(\min\{d, S\} \cdot H^3 S A / \epsilon^2)$ . This result partly resolves an open problem raised by [Jin et al. \[2020\]](#).

## 1 Introduction

In single-objective reinforcement learning (RL), a *scalar reward* is pre-specified and an agent learns a policy to maximize the long-term cumulative reward [[Sutton and Barto, 2018](#)]. However, in many real-world applications, we need to optimize multiple objectives for the same (unknown) environment, even when these objectives are possibly contradicting [[Roijers et al., 2013](#)]. For example, in an autonomous driving application, each passenger may have a different preference of driving styles: some of the passengers prefer a very steady riding experience while other passengers enjoy the fast acceleration of the car. Therefore, traditional single-objective RL approach may fail to be applied in such scenarios. One way to tackle this issue is the *multi-objective reinforcement learning* (MORL) [[Roijers et al., 2013](#), [Yang et al., 2019](#), [Natarajan and Tadepalli, 2005](#), [Abels et al., 2018](#)]

method, which models the multiple objectives by a *vectorized reward*, and an additional *preference vector* to specify the relative importance of each objective. The agent of MORL needs to find policies to optimize the cumulative preference-weighted rewards.

If the preference vector is fixed or drawn from a fixed distribution, MORL is no more challenging than single-objective RL, as we can predict the objective to be optimized and apply (variants of) single-objective RL algorithms (e.g., [Azar et al., 2017]). However, more often in practice, the preference vector (under which the weighted objective needs to be optimized) is:

- (i) adversarially provided,
- (ii) or even not available in the learning phase.

Once more taking the autonomous driving application as example: (i) the intelligent system needs to adapt driving configurations to accommodate every customer, even though the next customer can have picky preference that is unpredictable; (ii) when the intelligent system is under development, the future customer cannot be known in advance, nonetheless, when the system is deployed, it has to be capable to accommodate any potential customer who rides the car. Such MORL examples are common, to name a few beyond the autonomous driving one: medical treatment must take care of every patient even in very rare health conditions; an education system should accommodate every student according to his/her own characteristics; and emergency response systems have to be responsible in all extreme cases. Due to these exclusive challenges that have not appeared in single-objective RL, new *sample-efficient* algorithms for MORL need to be developed, as well as their *theoretical grounds* need to be established.

In this work, we study *provable* sample-efficient algorithms for MORL that resolve the aforementioned issues. Specifically, we consider MORL on a finite-horizon Markov decision process (MDP) with an unknown transition kernel,  $S$  states,  $A$  actions,  $H$  steps, and  $d$  reward functions that represent the  $d$  difference objectives. We investigate MORL problems in two paradigms: (i) *online MORL* (where the preferences are adversarially presented), and (ii) *preference-free exploration* (where the preferences are not available in the learning/exploration phase). The two settings and our contributions are explained respectively in the following.

**Setting (i): Online MORL.** We first consider an online learning problem to capture the challenge that preference vectors could be adversarially provided in MORL. In the beginning of each episode, the MORL agent is provided (potentially adversarially) with a preference vector, and the agent interacts with the unknown environment to collect data and rewards (that is specified by the provided preference). The performance of an algorithm is measured by the *regret*, i.e., the difference between the rewards collected by the agent and those would be collected by a theoretically optimal agent (who could use varying policies that adapt to the preferences). This setting generalizes the classical online single-objective RL problem [Azar et al., 2017] (where  $d = 1$ ).

**Contribution (i).** For online MORL, we provide a provably efficient algorithm with a regret upper bound  $\tilde{O}(\sqrt{\min\{d, S\} \cdot H^2 SAK})^1$ , where  $K$  is the number of episodes for interacting with the environment. Furthermore, we show an information-theoretic lower bound  $\Omega(\sqrt{\min\{d, S\} \cdot H^2 SAK})$  for online MORL. These bounds together show that, ignoring logarithmic factors, our algorithm resolves the online MORL problems optimally.

**Setting (ii): Preference-Free Exploration.** We further consider an unsupervised MORL problem to capture the issue that preferences could be hard to obtain in the training phase. The MORL agent first interacts with the unknown environment in the absence of preference vectors; afterwards, the agent is required to use the collected data to compute near-optimal policies for an arbitrarily specified preference vector. The performance is measured by the *sample complexity*, i.e., the minimum amount of trajectories that an MORL agent needs to collect during exploration in order to be near-optimal during planning. This setting extends the recent proposed *reward-free exploration* problem [Jin et al., 2020] to MORL.

**Contribution (ii).** For preference-free exploration, we show that a simple variant of the proposed online algorithm can achieve nearly optimal sample complexity. In particular, the algorithm achieves a sample complexity upper bound  $\tilde{O}(\min\{d, S\} \cdot H^3 SA/\epsilon^2)$  where  $\epsilon$  is the tolerance of the planning error; and we also show a sample complexity lower bound,  $\Omega(\min\{d, S\} \cdot H^2 SA/\epsilon^2)$ , for any

---

<sup>1</sup>We use  $\tilde{O}(\cdot)$  to hide (potential) polylogarithmic factors in  $\mathcal{O}(\cdot)$ , i.e.,  $\tilde{O}(n) := \mathcal{O}(n \log^k n)$  for sufficiently large  $n$  and some absolute constant  $k > 0$ .

algorithm. These bounds suggest that our algorithm is optimal in terms of  $d, S, A, \epsilon$  up to logarithmic factors. It is also worth noting that our results for preference-free exploration partly answer an open question raised by Jin et al. [2020]: reward-free RL is easier when the unknown reward functions enjoy low-dimensional representations (as in MORL).

**Paper Layout.** The remaining paper is organized as follows: the preliminaries are summarized in Section 2; then in Section 3, we formally introduce the problem of online MORL, our algorithm and its regret upper and lower bounds; then in Section 4, we turn to study the preference-free exploration problem in MORL, where we present an exploration algorithm with sample complexity analysis (with both upper bound and lower bound), and compare our results with existing results for related problems; finally, the related literature is reviewed in Section 5 and the paper is concluded by Section 6.

## 2 Preliminaries

We specify a finite-horizon Markov decision process (MDP) by a tuple of  $(\mathcal{S}, \mathcal{A}, H, \mathbb{P}, \mathbf{r}, \mathcal{W})$ .  $\mathcal{S}$  is a finite state set where  $|\mathcal{S}| = S$ .  $\mathcal{A}$  is a finite action set where  $|\mathcal{A}| = A$ .  $H$  is the length of the horizon.  $\mathbb{P}(\cdot | x, a)$  is a stationary, unknown transition probability to a new state for taking action  $a$  at state  $x$ .  $\mathbf{r} = \{\mathbf{r}_1, \dots, \mathbf{r}_H\}$ , where  $\mathbf{r}_h : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]^d$  represents a  $d$ -dimensional vector rewards function that captures the  $d$  objectives<sup>2</sup>.  $\mathcal{W} := \{w \in [0, 1]^d, \|w\|_1 = 1\}$  specifies the set of all possible preference vectors<sup>3</sup>, where each preference vector  $w \in \mathcal{W}$  induces a scalar reward function by<sup>4</sup>  $r_h(\cdot, \cdot) = \langle w, \mathbf{r}_h(\cdot, \cdot) \rangle$  for  $h = 1, \dots, H$ . A policy is represented by  $\pi := \{\pi_1, \dots, \pi_H\}$ , where each  $\pi_h(\cdot)$  maps a state to a distribution over the action set. Fixing a policy  $\pi$ , we will consider the following generalized  $Q$ -value function and generalized value function:

$$Q_h^\pi(x, a; w) := \mathbb{E}[\sum_{j=h}^H \langle w, \mathbf{r}_j(x_j, a_j) \rangle | x_h = x, a_h = a], \quad V_h^\pi(x; w) := Q_h^\pi(x, \pi_h(x); w),$$

where  $x_j \sim \mathbb{P}(\cdot | x_{j-1}, a_{j-1})$  and  $a_j \sim \pi_j(x_j)$  for  $j > h$ . Note that compared with the typical  $Q$ -value function (or the value function) used in single-objective RL, here the generalized  $Q$ -value function (or the generalized value function) takes the preference vector as an additional input, besides the state-action pair. Fixing a preference  $w \in \mathcal{W}$ , the optimal policy under  $w$  is defined as  $\pi_w^* := \arg \max_\pi V_1^\pi(x_1; w)$ . For simplicity, we denote

$$Q_h^*(x, a; w) := Q_h^{\pi_w^*}(x, a; w), \quad V_h^*(x; w) := V_h^{\pi_w^*}(x; w) = \max_a Q_h^*(x, a; w).$$

The following abbreviation is adopted for simplicity:

$$\mathbb{P}V_h^\pi |_{x,a,w} := \sum_{y \in \mathcal{S}} \mathbb{P}(y|x, a) V_h^\pi(y; w),$$

and similar abbreviations will be adopted for variants of probability transition (e.g., the empirical transition probability) and variants of value function (e.g., the estimated value function). Finally, we remark that the following Bellman equations hold for the generalized  $Q$ -value functions

$$\begin{cases} Q_h^\pi(x, a; w) = \langle w, \mathbf{r}_h(x, a) \rangle + \mathbb{P}V_{h+1}^\pi |_{x,a,w}, & \begin{cases} Q_h^*(x, a; w) = \langle w, \mathbf{r}_h(x, a) \rangle + \mathbb{P}V_{h+1}^* |_{x,a,w}, \\ V_h^*(x; w) = \max_a Q_h^*(x, a; w). \end{cases} \\ V_h^\pi(x; w) = Q_h^\pi(x, \pi_h(x); w); \end{cases}$$

With the above preparations, we are ready to discuss our algorithms and theory for online MORL (Section 3) and preference-free exploration (Section 4).

## 3 Online MORL

**Problem Setups.** The online setting captures the first difficulty in MORL, where the preference vectors can be adversarially provided to the MORL agent. Formally, the MORL agent interacts with

<sup>2</sup>For the sake of presentation, we discuss bounded deterministic reward functions in this work. The techniques can be readily extended to stochastic reward settings.

<sup>3</sup>The condition that  $w \in [0, 1]^d$  is only assumed for convenience. Our results naturally generalize to preference vectors that are entry-wisely bounded by absolute constants.

<sup>4</sup>The linear scalarization method can be generalized. See more discussions in Section 3.2, Remark 2.

an unknown environment through Protocol 1: at the beginning of the  $k$ -th episode, an adversary selects a preference vector  $w^k$  and reveals it to the agent; then starting from a fixed initial state<sup>5</sup>  $x_1$ , the agent draws a trajectory from the environment by recursively taking an action and observing a new state, and collects rewards from the trajectory, where the rewards are scalarized from the vector rewards by the given preference. The agent’s goal is to find the policies that maximize the cumulative rewards. Its performance will be measured by the following *regret*: suppose the MORL agent has interacted with the environment through Protocol 1 for  $K$  episodes, where the provided preferences are  $\{w^k\}_{k=1}^K$  and the adopted policies are  $\{\pi^k\}_{k=1}^K$  correspondingly, we consider the regret of the collected rewards (in expectation) competing with the theoretically maximum collected rewards (in expectation):

$$\text{regret}(K) := \sum_{k=1}^K V_1^*(x_1; w^k) - V_1^{\pi^k}(x_1; w^k). \quad (1)$$

Clearly, the regret (1) is always non-negative, and a smaller regret implies a better online performance. We would like to highlight that the regret (1) allows the theoretically optimal agent to adopt varying policies that adapt to the preferences.

---

#### Protocol 1 Environment Interaction Protocol

---

**Require:** MDP( $\mathcal{S}, \mathcal{A}, H, \mathbb{P}, r, \mathcal{W}$ ) and online agent

- 1: agent observes ( $\mathcal{S}, \mathcal{A}, H, r, \mathcal{W}$ )
  - 2: **for** episode  $k = 1, 2, \dots, K$  **do**
  - 3:   agent receives an initial state  $x_1^k = x_1$ , and a preference  $w^k$  (from an adversary)
  - 4:   **for** step  $h = 1, 2, \dots, H$  **do**
  - 5:     agent chooses an action  $a_h^k$ , and collects reward  $\langle w^k, r_h(x_h^k, a_h^k) \rangle$
  - 6:     agent transits to a new state  $x_{h+1}^k \sim \mathbb{P}(\cdot | x_h^k, a_h^k)$
  - 7:   **end for**
  - 8: **end for**
- 

#### Connections to Online Single-Objective RL.

The online regret minimization problems are well investigated in the context of single-objective RL (see, e.g., [Azar et al., 2017, Jin et al., 2018]). In both online single-objective RL and our studied online MORL, it is typical to assume that the transition probability  $\mathbb{P}$  is the only unknown information about the environment since estimating a stochastic reward is relatively easy (see, e.g., [Azar et al., 2017, Jin et al., 2018]). In particular, single-objective RL is a special case of MORL in the online setting, where the preference is fixed during the entire learning process (i.e.,  $w^k := w$  for all  $k$ ) — therefore an online MORL algorithm naturally applies to single-objective RL. However, the reverse is not true as that in MORL the preference vectors can change over time and are potentially adversarial.

#### Comparison with Single-Objective Stochastic and Adversarial Reward Setting.

The essential difficulty of MORL is further reflected in the regret (1). Specifically, the regret (1) compares the performance of the agent’s policies to a sequence of *optimal* policy under *each* given preference, which could vary over time. However, in online single-objective RL, either with known/stochastic rewards [Azar et al., 2017, Jin et al., 2018] or adversarial rewards [Rosenberg and Mansour, 2019, Jin et al., 2019], the benchmark policy is supposed to be *fixed* over time (the best policy in the hindsight). This difference suggests that online MORL could be more challenging than online single-objective RL, as a harder performance measurement is adopted. More specifically, when measured by regret




Figure 1: A regret comparison of MO-UCBVI vs. the best-in-hindsight policy in a simulated random multi-objective MDP. Note that the best-in-hindsight policy is the optimal policy for single-objective RL. The plots show that the best-in-hindsight policy will incur linear regret in online MORL, and the proposed MO-UCBVI achieves sub-linear regret as predicted by Theorem 1. See Appendix A for details.

<sup>5</sup>Without loss of generality, we fix the initial state; otherwise we may as well consider an MDP with an external initial state  $x_0$  with zero reward for all actions, and a transition  $\mathbb{P}_0(\cdot | x_0, a) = \mathbb{P}_0(\cdot)$  for all action  $a$ . This is equivalent to our setting by letting the horizon length  $H$  be  $H + 1$ .

(1), existing algorithms for single-objective RL (e.g., [Azar et al., 2017]) easily suffer a  $\propto \Theta(K)$  regret when the sequence of preferences is adversarially designed; in contrast, we will show an MORL algorithm that experiences at most  $\propto \tilde{O}(\sqrt{K})$  regret under any sequence of preferences. A numerical simulation of this issue is presented in Figure 1.

### 3.1 A Sample-Efficient Online Algorithm

In this part we introduce an algorithm for online MORL, called *multi-objective upper confidence bound value iteration* (**MO-UCBVI**). A simplified version of **MO-UCBVI** is presented as Algorithm 1, and a more advanced version (Algorithm 4) can be found in Appendix B. Our algorithm is inspired by UCBVI [Azar et al., 2017] that achieves minimax regret bound in online single-objective RL.

---

#### Algorithm 1 MO-UCBVI

---

```

1: initialize history  $\mathcal{H}^0 = \emptyset$ 
2: for episode  $k = 1, 2, \dots, K$  do
3:    $N^k(x, a), \hat{\mathbb{P}}^k(y | x, a) \leftarrow \text{Empi-Prob}(\mathcal{H}^{k-1})$ 
4:   compute bonus  $b^k(x, a) := c \cdot \sqrt{\frac{\min\{d, S\} H^{2\iota}}{N^k(x, a)}}$ , where  $\iota = \log(HSAK/\delta)$  and  $c$  is a constant
5:   receive a preference  $w^k$ 
6:    $\{Q_h^k(x, a; w^k)\}_{h=1}^H \leftarrow \text{UCB-Q-Value}(\hat{\mathbb{P}}^k, w^k, b^k)$ 
7:   receive initial state  $x_1^k = x_1$ 
8:   for step  $h = 1, 2, \dots, H$  do
9:     take action  $a_h^k = \arg \max_a Q_h^k(x_h^k, a; w^k)$ , and obtain a new state  $x_{h+1}^k$ 
10:  end for
11:  update history  $\mathcal{H}^k = \mathcal{H}^{k-1} \cup \{x_h^k, a_h^k\}_{h=1}^H$ 
12: end for

13: Function Empi-Prob
14: Require: history  $\mathcal{H}^{k-1}$ 
15: for  $(x, a, y) \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$  do
16:    $N^k(x, a, y) := \#\{(x, a, y) \in \mathcal{H}^{k-1}\}$ , and  $N^k(x, a) := \sum_y N^k(x, a, y)$ 
17:   if  $N^k(x, a) > 0$  then
18:      $\hat{\mathbb{P}}^k(y | x, a) = N^k(x, a, y) / N^k(x, a)$ 
19:   else
20:      $\hat{\mathbb{P}}^k(y | x, a) = 1/S$ 
21:   end if
22: end for
23: return  $N^k(x, a)$  and  $\hat{\mathbb{P}}^k(y | x, a)$ 

24: Function UCB-Q-Value
25: Require: empirical transition  $\hat{\mathbb{P}}^k$ , preference  $w^k$ , and bonus  $b^k(x, a)$ 
26: set  $V_{H+1}^k(x; w^k) = 0$ 
27: for step  $h = H, H-1, \dots, 1$  do
28:   for  $(x, a) \in \mathcal{S} \times \mathcal{A}$  do
29:      $Q_h^k(x, a; w^k) = \min \{H, \langle w^k, \mathbf{r}_h(x, a) \rangle + b^k(x, a) + \hat{\mathbb{P}}_h^k V_{h+1}^k |_{x, a, w^k}\}$ 
30:      $V_h^k(x; w^k) = \max_{a \in \mathcal{A}} Q_h^k(x, a; w^k)$ 
31:   end for
32: end for
33: return  $\{Q_h^k(x, a; w^k)\}_{h=1}^H$ 

```

---

In Algorithm 1, the agent interacts with the environment according to Protocol 1, and use an *optimistic policy* to explore the unknown environment and collect rewards. The optimistic policy is a greedy policy with respect to an optimistic estimation to the value function, UCB-Q-Value (lines 6 and 9). Specifically, UCB-Q-Value is constructed to maximize the cumulative reward, which is scalarized by the current preference, through dynamic programming over an empirical transition probability (line 24). The empirical transition probability is inferred from the data collected so far (lines 3 and

13), which might not be accurate if a state-action pair has not yet been visited for sufficient times. To mitigate this inaccuracy, UCB-Q-Value utilizes an extra *exploration bonus* (lines 4 and 29) so that: (i) UCB-Q-Value never under-estimates the optimal true value function, for *whatever preference vector* (and with high probability); and (ii) the added bonus *shrinks quickly* enough (as the corresponding state-action pair continuously being visited) so that the over-estimation is under control. The overall consequence is that: **MO-UCBVI** explores the environment via a sequence of optimistic policies, in order to collect rewards under a sequence of adversarially provided preferences; since the policies are optimistic for any preference, the incurred regret would not exceed the total amount of added bonus; and since the bonus decays fast, their sum up would be sublinear (with respect to the number of episodes played). Therefore **MO-UCBVI** only suffers a sublinear regret, even when the preferences are adversarially presented. The above intuition is formalized in the next part.

### 3.2 Theoretical Analysis

The next two theorems justify regret upper bounds for **MO-UCBVI** and a regret lower bound for online MORL problems, respectively.

**Theorem 1** (Regret bounds for **MO-UCBVI**). *Suppose  $K$  is sufficiently large. Then for any sequence of the incoming preferences  $\{w^k\}_{k=1}^K$ , with probability at least  $1 - \delta$ :*

- the regret (1) of **MO-UCBVI** (see Algorithm 1) satisfies

$$\text{regret}(K) \leq \mathcal{O}\left(\sqrt{\min\{d, S\}} \cdot H^3 SAK \log(HSAK/\delta)\right);$$

- and the regret (1) of a Bernstein-variant **MO-UCBVI** (see Algorithm 4 in Appendix B) satisfies

$$\text{regret}(K) \leq \mathcal{O}\left(\sqrt{\min\{d, S\}} \cdot H^2 SAK \log^2(HSAK/\delta)\right).$$

**Theorem 2** (A regret lower bound for MORL). *There exist some absolute constants  $c, K_0 > 0$ , such that for any  $K > K_0$ , any MORL algorithm that runs  $K$  episodes, there is a set of MOMDPs and a sequence of (necessarily adversarially chosen) preferences vectors such that*

$$\mathbb{E}[\text{regret}(K)] \geq c \cdot \sqrt{\min\{d, S\}} \cdot H^2 SAK,$$

where the expectation is taken with respect to the randomness of choosing MOMDPs and the randomness of the algorithm for collecting dataset.

*Remark 1.* When  $d = 1$ , MORL recovers the single-objective RL setting, and Theorem 1 recovers existing nearly minimax regret bounds for single-objective RL [Azar et al., 2017, Zanette and Brunskill, 2019]. Moreover, the lower bound in Theorem 2 implies that our upper bound in Theorem 1 is tight ignoring logarithm terms. Interestingly, the lower bound suggests MORL with  $d > 2$  is truly harder than single objective RL (corresponding to  $d = 1$ ) as the sequence of preferences can be adversarially chosen.

*Remark 2.* Theorem 1 (as well as our other theorems) applies to general scalarization methods besides the linear one as adopted by Algorithm 1 and many other MORL papers [Yang et al., 2019, Roijers et al., 2013, Natarajan and Tadepalli, 2005, Abels et al., 2018]. In particular, our results apply to scalarization functions  $r_h(\cdot, \cdot) = f(r_h(\cdot, \cdot); w)$  that are (1) deterministic, (2) Lipschitz continuous for  $w$ , and (3) bounded between  $[0, 1]$  (which can be relaxed). This will be clear from proofs in Appendix B, where we treat the potentially adversarially given preferences by a covering argument and union bound, and these techniques are not dedicated to linear scalarization function and can be easily extended to more general cases.

The proof of Theorem 1 leverages standard analysis procedures for single-objective RL [Azar et al., 2017, Zanette and Brunskill, 2019], and a covering argument with an union bound to tackle the challenge of adversarial preferences. Th rigorous proof is included in Appendix B.

Specifying an adversarial process of providing preferences is the key challenge for proving Theorem 2. To handle this issue, we use reduction techniques and utilize a lower bound that we will show shortly for preference-free exploration problems. We refer the readers to Theorem 4 and Appendix E for more details.

## 4 Preference-Free Exploration

**Problem Setups.** *Preference-free exploration* (PFE) captures the second difficulty in MORL: the preference vector might not be observable when the agent explores the environment. Specifically, PFE

consists of an exploration phase and a planning phase. Similarly as in the online setting, the transition probability is hidden from the MORL agent. In the exploration phase, the agent interacts with the unknown environment to collect samples, however the agent has no information about the preference vectors at this point. Afterwards PFE switches to the planning phase, where the agent is prohibited to obtain new data, and is required to compute near-optimal policy for any preference-weighted reward functions. Since this task is no longer in an online fashion, we turn to measure the performance of a PFE algorithm by the minimum number of required trajectories (in the exploration phase) so that the algorithm can behave near-optimally in the planning phase. This is made formal as follows: a PFE algorithm is called  $(\epsilon, \delta)$ -PAC (Probably Approximately Correct), if

$$\mathbb{P}\{\forall w \in \mathcal{W}, V_1^*(x_1; w) - V_1^{\pi_w}(x_1; w) \leq \epsilon\} \geq 1 - \delta,$$

where  $\pi_w$  is the policy outputted by the PFE algorithm for input preference  $w$ , then the *sample complexity* of a PFE algorithm is defined by the least amount of trajectories it needs to collect in the exploration phase for being  $(\epsilon, \delta)$ -PAC in the planning phase.

**Connections to Reward-Free Exploration.** PFE problem is a natural extension to the recent proposed *reward-free exploration* (RFE) problem [Jin et al., 2020, Kaufmann et al., 2020, Wang et al., 2020, Ménard et al., 2020, Zhang et al., 2020b]. Both problems consist of an exploration phase and a planning phase; the difference is in the planning phase: in RFE, the agent needs to be able to compute near-optimal policies for *all reward functions*, while in PFE, the agent only needs to achieve that for *all preference-weighted reward functions*, i.e., the reward functions that can be represented as the inner product of a  $d$ -dimensional preference vectors and the  $d$ -dimensional vector rewards functions (i.e., the  $d$  objectives in MORL). A PFE problem reduces to a RFE problem if  $d = SA$  such that every reward function can be represented as a preference-weighted reward function. However, if  $d \ll SA$ , it is conjectured by Jin et al. [2020] that PFE can be solved with a much smaller sample complexity than RFE. Indeed, in the following part we show an algorithm that improves a  $\propto \tilde{O}(S^2)$  dependence for RFE to  $\propto \tilde{O}(\min\{d, S\} \cdot S)$  dependence for PFE, in terms of sample complexity.

#### 4.1 A Sample-Efficient Exploration Algorithm

We now present a simple variant of **MO-UCBVI** that is sample-efficient for PFE. The algorithm is called *preference-free upper confidence bound exploration* (**PF-UCB**), and is discussed separately as in the exploration phase (Algorithm 2) and in the planning phase (Algorithm 3) in below.

---

##### Algorithm 2 **PF-UCB** (Exploration)

---

- 1: initialize history  $\mathcal{H}^0 = \emptyset$
  - 2: **for** episode  $k = 1, 2, \dots, K$  **do**
  - 3:    $N^k(x, a), \hat{\mathbb{P}}^k(y | x, a) \leftarrow \text{Empi-Prob}(\mathcal{H}^{k-1})$
  - 4:   compute bonus  $c^k(x, a) := \frac{H^2 S}{2N^k(x, a)} + 2b^k(x, a)$  for  $b^k(x, a)$  defined in Algorithms 1 or 3
  - 5:    $\{\bar{Q}_h^k(x, a)\}_{h=1}^H \leftarrow \text{UCB-Q-Value}(\hat{\mathbb{P}}^k, w = 0, c^k)$
  - 6:   receive initial state  $x_1^k = x_1$
  - 7:   **for** step  $h = 1, 2, \dots, H$  **do**
  - 8:     take action  $a_h^k = \arg \max_a \bar{Q}_h^k(x_h^k, a)$ , and obtain a new state  $x_{h+1}^k$
  - 9:   **end for**
  - 10:   update history  $\mathcal{H}^k = \mathcal{H}^{k-1} \cup \{x_h^k, a_h^k\}_{h=1}^H$
  - 11: **end for**
- 

Algorithm 2 presents our PFE algorithm in the exploration phase. Indeed, Algorithm 2 is a modified **MO-UCBVI** (Algorithm 1) by setting the preference to be zero (Algorithm 2, line 5), and slightly enlarging the exploration bonus (Algorithm 2, line 4). The intention of an increased exploration bonus will be made clear later when we discuss the planning phase. With a zero preference vector, the **UCB-Q-Value** in Algorithm 2 will identify a trajectory along which the cumulative bonus (instead of the cumulative rewards) is maximized (with respect to the empirical transition probability). Also note that the bonus function (Algorithm 2, line 4) is negatively correlated with the number of visits to a state-action pair. Hence the greedy policy with respect to the zero-preference **UCB-Q-Value** tends to visit the state-actions pairs that are associated with large bonus, i.e., those that have been visited

---

**Algorithm 3** PF-UCB (Planning)

---

**Require:** history  $\mathcal{H}^K$ , preference vector  $w$

- 1: **for**  $k = 1, 2, \dots, K$  **do**
  - 2:    $N^k(x, a), \widehat{\mathbb{P}}^k(y|x, a) \leftarrow \text{Empi-Prob}(\mathcal{H}^{k-1})$
  - 3:   compute bonus  $b^k(x, a) := c \cdot \sqrt{\frac{\min\{d, S\} H^2 \iota}{N^k(x, a)}}$  where  $\iota = \log(HSAK/\delta)$  and  $c$  is a constant
  - 4:    $\{Q_h^k(\cdot, \cdot; w)\}_{h=1}^H \leftarrow \text{UCB-Q-Value}(\widehat{\mathbb{P}}^k, w, b^k)$
  - 5:   infer greedy policy  $\pi_h^k(x) = \arg \max_a Q_h^k(x, a; w)$
  - 6: **end for**
  - 7: **return**  $\pi$  drawn uniformly from  $\{\pi^k\}_{k=1}^K$
- 

for less times. In sum, Algorithm 2 explores the unknown environment “uniformly”, without the guidance of preference vectors.

Then Algorithm 3 shows our PFE algorithm in the planning phase. Given any preference vector, Algorithm 3 computes a sequence of optimistically estimated value functions based on the data collected from the exploration phase, and then outputs a greedy policy with respect to a randomly drawn optimistic value estimation. Note that the bonus in Algorithm 3 is set as the one in Algorithm 1, and recall that the bonus in Algorithm 2 is an enlarged one. The relatively large bonus in the exploration phase guarantees that the regret in the planning phase never exceeds that in the exploration phase. On the other hand, based on Theorem 1 for MO-UCBVI (Algorithm 1), we have already known that the exploration algorithm (Algorithm 2), a modified Algorithm 1, suffers at most  $\tilde{\mathcal{O}}(\sqrt{K})$  regret, hence the planning algorithm (Algorithm 3) experiences at most  $\tilde{\mathcal{O}}(1/\sqrt{K})$  error. The next part rigorously justifies these discussions.

## 4.2 Theoretic Analysis

We first provide Theorem 3 to justify the trajectory complexity of Algorithms 2 and 3 in the PFE setting; then we present Theorem 4 that gives an information-theoretic lower bound on the trajectory complexity for any PFE algorithm.

**Theorem 3** (A trajectory complexity of PF-UCB). *Suppose  $\epsilon > 0$  is sufficiently small. Then for PF-UCB (Algorithm 2) run for*

$$K = \mathcal{O}(\min\{d, S\} \cdot H^3 SA \iota / \epsilon^2), \quad \text{where } \iota := \log(HSA/(\delta\epsilon)),$$

*episodes, PF-UCB (Algorithm 3) is  $(\epsilon, \delta)$ -PAC.*

**Theorem 4** (A lower bound for PFE). *There exist absolute constants  $c, \epsilon_0 > 0$ , such that for any  $0 < \epsilon < \epsilon_0$ , there exists a set of MOMDPs such that any PFE algorithm that is  $(\epsilon, 0.1)$ -PAC on them, it needs to collect at least*

$$K \geq \Omega(\min\{d, S\} \cdot H^2 SA / \epsilon^2)$$

*trajectories in expectation (with respect to the randomness of choosing MOMDPs and the exploration algorithm).*

*Remark 3.* According to Theorems 3 and 4, the trajectory complexity of PF-UCB is optimal for  $d, S, A, \epsilon$  ignoring logarithmic factors, but is an  $H$  factor loose compared with the lower bound. This is because the current Algorithm 2 utilizes a *preference-independent*, Hoeffding-type bonus since the preference vector is not available during exploration. We leave it as an open problem to further remove this gap about  $H$ .

**Proof Sketch of Theorem 3.** Theorem 3 is obtained in three procedures. (1) We first observe the total regret incurred by Algorithm 2 is  $\tilde{\mathcal{O}}(\sqrt{\min\{d, S\} H^3 SAK})$  according to Theorem 1. (2) Then utilizing the enlarged exploration bonus, we show that in each episode, the planning error is at most constant times of the incurred exploration error. (3) Thus the averaged planning error is at most  $\tilde{\mathcal{O}}(\sqrt{\min\{d, S\} H^3 SA/K})$  as claimed. A complete proof is deferred to Appendix C.

Note that the second argument is motivated by [Zhang et al., 2020a, Wang et al., 2020]. However in their original paper, a brute-force union bound over all possible value functions are required to obtain similar effect, due to the limitation of model-free algorithm [Zhang et al., 2020a] (see Appendix F



for more details) or linear function approximation [Wang et al., 2020]. This will cause a loose,  $\propto \tilde{\mathcal{O}}(S^2)$  complexity in the obtained bound. Different from their approach, we carefully manipulate a lower order term to avoid union bounding all value functions during the second argument. As a consequence we obtain a near-tight  $\propto \tilde{\mathcal{O}}(\min\{d, S\} \cdot S)$  dependence in the final bound. We believe this observation has broader application in the analysis of similar RL problems.

**Proof Sketch of Theorem 4.** We next introduce the idea of constructing the hard instance that witnesses the lower bound in Theorem 4. A basic ingredient is the hard instance given by Jin et al. [2020]. However, this hard instance is invented for RFE, where the corresponding lower bound is  $K \geq \Omega(H^2 S^2 A / \epsilon^2)$ . Note this lower bound cannot match the upper bound in Theorem 3 in terms of  $d$  and  $S$ . In order to develop a dedicated lower bound for PFE, we utilize Johnson–Lindenstrauss Lemma [Johnson and Lindenstrauss, 1984] to refine the hard instance in [Jin et al., 2020], and successfully reduce a factor  $S$  to  $\min\{d, S\}$  in their lower bound, which gives the result in Theorem 4. We believe that the idea to refine RL hard instance by Johnson–Lindenstrauss Lemma is of broader interests. A rigorous proof is deferred to Appendix D.

**Application in Reward-Free Exploration.** By setting  $d = SA$  and allowing arbitrary reward functions, PFE problems reduce to RFE problems. Therefore as a side product, Theorem 3 implies the following results for RFE problems on stationary or non-stationary MDPs<sup>6</sup>:

**Corollary 5.** *Suppose  $\epsilon > 0$  is sufficiently small. Consider the reward-free exploration problems on a stationary MDP. Suppose PF-UCB (Algorithm 2) is run for*

$$K = \mathcal{O}(H^3 S^2 A \iota / \epsilon^2), \quad \text{where } \iota := \log(HSA / (\delta\epsilon)),$$

*episodes, then PF-UCB (Algorithm 3) is  $(\epsilon, \delta)$ -PAC. Moreover, if the MDP is non-stationary, the above bound will be revised to  $K = \mathcal{O}(H^4 S^2 A \iota / \epsilon^2)$ .*

*Remark 4.* When applied to RFE, PF-UCB matches the rate shown in [Kaufmann et al., 2020] and improves an  $H$  factor compared with [Jin et al., 2020] (for both stationary and non-stationary MDPs), but is an  $H$  factor loose compared with the current best rates, [Ménard et al., 2020] (for non-stationary MDPs) and [Zhang et al., 2020b] (for stationary MDPs). However we highlight that our results are superior in the context of PFE, since PF-UCB adapts with the structure of the rewards. In specific, if rewards admit a  $d$ -dimensional feature for  $d < S$ , PF-UCB only needs  $\propto \tilde{\mathcal{O}}(\min\{d, S\} \cdot S)$  samples, but the above methods must explore the whole environment with a high precision which consumes  $\propto \tilde{\mathcal{O}}(S^2)$  samples.

**Application in Task-Agnostic Exploration.** PFE is also related to *task-agnostic exploration* (TAE) [Zhang et al., 2020a]: in PFE, the agent needs to plan for an arbitrary reward function from a fixed and bounded  $d$ -dimensional space; and in TAE, the agent needs to plan for  $N$  fixed reward functions. Due to the nature of the problem setups, PFE algorithms (that do not exploit the given reward basis  $\mathbf{r}$  during exploration, e.g., ours) and TAE algorithms can be easily applied to solve the other problem through a covering argument and a union bound, and with a modification of  $\min\{d, S\} \leftrightarrow \log(N)$  in the obtained trajectory complexity bounds. For TAE on a non-stationary MDP, Zhang et al. [2020a] show an algorithm which takes  $\tilde{\mathcal{O}}(\log(N) \cdot H^5 SA / \epsilon^2)$  episodes for TAE<sup>7</sup>. In comparison, when applied to TAE on a non-stationary MDP, Theorem 3 implies PF-UCB only takes  $\tilde{\mathcal{O}}(\log(N) \cdot H^4 SA / \epsilon^2)$  episodes<sup>8</sup>, which improves [Zhang et al., 2020a].

<sup>6</sup>Our considered MDP is *stationary* as the transition probability  $\mathbb{P}$  is fixed (across different steps). An MDP is called *non-stationary*, if the transition probability varies at different steps, i.e., replacing  $\mathbb{P}$  by  $\{\mathbb{P}_h\}_{h=1}^H$ .

<sup>7</sup>This bound is copied from [Zhang et al., 2020a], which is erroneously stated due to a technical issue in the proof of Lemma 2 in their original paper. The issue can be fixed by a covering argument and union bound on the value functions, but then the obtained bound should be  $\tilde{\mathcal{O}}(\log(N) H^5 S^2 A / \epsilon^2)$ . See Appendix F for details.

<sup>8</sup>The conversion holds as follows. First set  $d = 1$  in our algorithm to yield an algorithm for TAE with a single agnostic task, where we have  $\min\{d, S\} = 1$ . Then one can extend this algorithm to TAE with  $N$  agnostic tasks using a union bound to have the algorithm succeed simultaneously for all  $N$  tasks, which adds a  $\log N$  multiplicative factor in the sample complexity bound. In this way, we obtain a TAE algorithm with a sample complexity bound in Theorem 3 where  $\min\{d, S\}$  is replaced with  $\log N$ .

## 5 Related Works

**MORL.** MORL receives extensive attention from RL applications [Yang et al., 2019, Natarajan and Tadepalli, 2005, Mossalam et al., 2016, Abels et al., 2018, Roijers et al., 2013]. However, little theoretical results are known, especially in terms of the sample efficiency and adversarial preference settings. A large amount of MORL works focus on identifying (a cover for) the policies belonging to the *Pareto front* [Yang et al., 2019, Roijers et al., 2013, Cheung, 2019], but the correspondence between a preference and an optimal policy is ignored. Hence they cannot “accommodate picky customers” as our algorithms do. To our knowledge, this paper initiates the theoretical study of the sample efficiency for MORL in the setting of adversarial preferences.

**Adversarial MDP.** Similarly to the online MORL problem studied in this paper, the adversarial MDP problem [Neu et al., 2012, Rosenberg and Mansour, 2019, Jin et al., 2019] also allows the reward function to change adversarially over time. However, we study a totally different regret (see (1)). In the adversarial MDP problem, the regret is measured against a *fixed* policy that is best in the hindsight; but in online MORL problem, we study a regret against a sequence of optimal policies with respect to the sequence of incoming preferences, i.e., our benchmark policy may *vary* over time. Therefore our regret bound for online MORL problems is orthogonal to those for adversarial MDP problems [Neu et al., 2012, Rosenberg and Mansour, 2019, Jin et al., 2019].

**Constrained MDP.** In the problem of the *constrained MDP*,  $d$  constraints are enforced to restrict the policy domain, and the agent aims to find a policy that belongs to the domain and maximizes the cumulative rewards [Achiam et al., 2017, Efroni et al., 2020, El Chamie et al., 2018, Fisac et al., 2018, Wachi and Sui, 2020, Garcia and Fernández, 2015, Brantley et al., 2020]. Constrained MDP is related to MORL as when the constraints are soft, they can be formulated as “objectives” with negative weights. However, there is a fundamental difference: constrained MDP aims to optimize only one (and known) objective, where in MORL studied in this paper, we aim to be able to find near optimal policies for any preference-weighted objective (to accommodate picky customers).

**Reward-Free Exploration.** The proposed preference-free exploration problem generalizes the problems of reward-free exploration [Jin et al., 2020]. Compared with existing works for RFE [Jin et al., 2020, Kaufmann et al., 2020, Ménard et al., 2020, Zhang et al., 2020b], our method has the advantage of adapting with rewards that admit low-dimensional structure — this partly answers an open question raised by Jin et al. [2020]. We also note that Wang et al. [2020] study RFE with linear function approximation on the value functions; in contrast our setting can be interpreted as RFE with linear function approximation on the reward functions, which is orthogonal to their setting.

## 6 Conclusion

In this paper we study provably sample-efficient algorithms for multi-objective reinforcement learning in both online and preference-free exploration settings. For both settings, sample-efficient algorithms are proposed and their sample complexity analysis is provided; moreover, two information-theoretic lower bounds are proved to justify the near-tightness of the proposed algorithms, respectively. Our results extend existing theory for single-objective RL and reward-free exploration, and resolve an open question raised by Jin et al. [2020].

## Acknowledgement

This research was supported in part by NSF CAREER grant 1652257, ONR Award N00014-18-1-2364, the Lifelong Learning Machines program from DARPA/MTO, and NSF HDR TRIPODS grant 1934979.

## References

- Axel Abels, Diederik M Roijers, Tom Lenaerts, Ann Nowé, and Denis Steckelmacher. Dynamic weights in multi-objective deep reinforcement learning. *arXiv preprint arXiv:1809.07803*, 2018.
- Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained policy optimization. In *International Conference on Machine Learning*, pages 22–31. PMLR, 2017.

- Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 263–272. JMLR. org, 2017.
- Kianté Brantley, Miroslav Dudik, Thodoris Lykouris, Sobhan Miryoosefi, Max Simchowitz, Aleksandrs Slivkins, and Wen Sun. Constrained episodic reinforcement learning in concave-convex and knapsack settings. *arXiv preprint arXiv:2006.05051*, 2020.
- Wang Chi Cheung. Exploration-exploitation trade-off in reinforcement learning on online markov decision processes with global concave rewards. *arXiv preprint arXiv:1905.06466*, 2019.
- Christoph Dann, Tor Lattimore, and Emma Brunskill. Unifying pac and regret: Uniform pac bounds for episodic reinforcement learning. *arXiv preprint arXiv:1703.07710*, 2017.
- Yonathan Efroni, Shie Mannor, and Matteo Pirota. Exploration-exploitation in constrained mdps. *arXiv preprint arXiv:2003.02189*, 2020.
- Mahmoud El Chamie, Yue Yu, Behçet Açıkmeşe, and Masahiro Ono. Controlled markov processes with safety state constraints. *IEEE Transactions on Automatic Control*, 64(3):1003–1018, 2018.
- Jaime F Fisac, Anayo K Akametalu, Melanie N Zeilinger, Shahab Kaynama, Jeremy Gillula, and Claire J Tomlin. A general safety framework for learning-based control in uncertain robotic systems. *IEEE Transactions on Automatic Control*, 64(7):2737–2752, 2018.
- Javier Garcia and Fernando Fernández. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research*, 16(1):1437–1480, 2015.
- Chi Jin, Zeyuan Allen-Zhu, Sebastien Bubeck, and Michael I Jordan. Is q-learning provably efficient? In *Advances in Neural Information Processing Systems*, pages 4863–4873, 2018.
- Chi Jin, Tiancheng Jin, Haipeng Luo, Suvrit Sra, and Tiancheng Yu. Learning adversarial mdps with bandit feedback and unknown transition, 2019.
- Chi Jin, Akshay Krishnamurthy, Max Simchowitz, and Tiancheng Yu. Reward-free exploration for reinforcement learning. *arXiv preprint arXiv:2002.02794*, 2020.
- William B Johnson and Joram Lindenstrauss. Extensions of lipschitz mappings into a hilbert space. *Contemporary mathematics*, 26(189-206):1, 1984.
- Emilie Kaufmann, Pierre Ménard, Omar Darwiche Domingues, Anders Jonsson, Edouard Leurent, and Michal Valko. Adaptive reward-free exploration. *arXiv preprint arXiv:2006.06294*, 2020.
- Andreas Maurer and Massimiliano Pontil. Empirical bernstein bounds and sample variance penalization. *arXiv preprint arXiv:0907.3740*, 2009.
- Pierre Ménard, Omar Darwiche Domingues, Anders Jonsson, Emilie Kaufmann, Edouard Leurent, and Michal Valko. Fast active learning for pure exploration in reinforcement learning. *arXiv preprint arXiv:2007.13442*, 2020.
- Hossam Mossalam, Yannis M Assael, Diederik M Roijers, and Shimon Whiteson. Multi-objective deep reinforcement learning. *arXiv preprint arXiv:1610.02707*, 2016.
- Sriraam Natarajan and Prasad Tadepalli. Dynamic preferences in multi-criteria reinforcement learning. In *Proceedings of the 22nd international conference on Machine learning*, pages 601–608, 2005.
- Gergely Neu, Andras Gyorgy, and Csaba Szepesvári. The adversarial stochastic shortest path problem with unknown transition probabilities. In *Artificial Intelligence and Statistics*, pages 805–813, 2012.
- Diederik M Roijers, Peter Vamplew, Shimon Whiteson, and Richard Dazeley. A survey of multi-objective sequential decision-making. *Journal of Artificial Intelligence Research*, 48:67–113, 2013.
- Aviv Rosenberg and Yishay Mansour. Online convex optimization in adversarial markov decision processes. *arXiv preprint arXiv:1905.07773*, 2019.

- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- Akifumi Wachi and Yanan Sui. Safe reinforcement learning in constrained markov decision processes. In *International Conference on Machine Learning*, pages 9797–9806. PMLR, 2020.
- Ruosong Wang, Simon S Du, Lin F Yang, and Ruslan Salakhutdinov. On reward-free reinforcement learning with linear function approximation. *arXiv preprint arXiv:2006.11274*, 2020.
- Runzhe Yang, Xingyuan Sun, and Karthik Narasimhan. A generalized algorithm for multi-objective reinforcement learning and policy adaptation. In *Advances in Neural Information Processing Systems*, pages 14610–14621, 2019.
- Andrea Zanette and Emma Brunskill. Tighter problem-dependent regret bounds in reinforcement learning without domain knowledge using value function bounds. In *International Conference on Machine Learning*, pages 7304–7312. PMLR, 2019.
- Xuezhou Zhang, Yuzhe Ma, and Adish Singla. Task-agnostic exploration in reinforcement learning, 2020a.
- Zihan Zhang, Simon S Du, and Xiangyang Ji. Nearly minimax optimal reward-free reinforcement learning. *arXiv preprint arXiv:2010.05901*, 2020b.

## Checklist

1. For all authors...
  - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [\[Yes\]](#)
  - (b) Did you describe the limitations of your work? [\[Yes\]](#)
  - (c) Did you discuss any potential negative societal impacts of your work? [\[N/A\]](#)
  - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [\[Yes\]](#)
2. If you are including theoretical results...
  - (a) Did you state the full set of assumptions of all theoretical results? [\[Yes\]](#)
  - (b) Did you include complete proofs of all theoretical results? [\[Yes\]](#)
3. If you ran experiments...
  - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [\[Yes\]](#)
  - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [\[Yes\]](#)
  - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [\[N/A\]](#)
  - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [\[N/A\]](#)
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
  - (a) If your work uses existing assets, did you cite the creators? [\[Yes\]](#)
  - (b) Did you mention the license of the assets? [\[N/A\]](#)
  - (c) Did you include any new assets either in the supplemental material or as a URL? [\[N/A\]](#)
  - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [\[N/A\]](#)
  - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [\[N/A\]](#)
5. If you used crowdsourcing or conducted research with human subjects...
  - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [\[N/A\]](#)

- (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
- (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]




Figure 2: **MO-UCBVI** vs. Q-learning [Jin et al., 2018] in a simulated random multi-objective MDP. The plot suggests that **MO-UCBVI** achieves sublinear regret but Q-learning incurs linear regret.

## A Numerical Simulations

Code for simulations is available at <https://github.com/uuujf/MORL>.

**Comparison with the Optimal Single-Objective RL Algorithm.** Figure 1 shows a regret comparison between the proposed **MO-UCBVI** and the best-in-hindsight policy. In the experiment we simulate a random multi-objective MDP with  $S = 20$ ,  $A = 5$ ,  $H = 10$ ,  $d = 15$ , and run the algorithms for  $K = 5000$  episodes. The best-in-hindsight policy refers to a policy that is fixed across episodes and achieves maximum cumulative rewards in the whole game, i.e.,  $\arg \max_{\pi} \sum_{k=1}^K V_1^k(x_1; w^k)$ . Note that the best-in-hindsight policy is the optimal policy in the single-objective RL setting; however it could be much worse than a time-varying policy in the multi-objective RL setting. Since the information-theoretically adversarial preferences are computationally infeasible to compute, in Figure 1 we simply use a randomly generated set of preferences, for which the “best-in-hindsight” policy already performs poorly. In sum, the plot shows that our algorithm achieves sublinear regret in online MORL, but the best single-objective RL algorithm will incur linear regret in online MORL.

**Comparison with Q-Learning.** Since MORL is connected to task-agnostic exploration [Zhang et al., 2020a] in the exploration setting (see discussions in Section 4), it is tempted to think that the Q-learning method studied in Jin et al. [2018], Zhang et al. [2020a] could also work in the setting of online MORL. However this is refuted by Figure 2. We simulate a random multi-objective MDP with  $S = 20$ ,  $A = 5$ ,  $H = 10$ ,  $d = 15$ , and run the algorithms for  $K = 5000$  episodes. The Q-learning algorithm is specified by Algorithm 1 in Jin et al. [2018], excepted that the reward is now a linear scalarization of a preference vector and a multi-objective reward vector. Figure 2 shows that Q-learning cannot achieve a sublinear regret in the setting of online MORL.

**The Effect of Number of Objectives.** Figure 3 shows the performance of **MO-UCBVI** with different number of objectives. In the experiment we simulate a random multi-objective MDP with  $S = 20$ ,  $A = 5$ ,  $H = 10$ , and number of objectives  $d \in \{1, 5, 15, 20, 30\}$ , and run **MO-UCBVI** for  $K = 5000$  episodes. As shown in Figure 3, the regret will increase as the number of objectives increases; moreover, in all settings **MO-UCBVI** achieves a sublinear regret.

## B Proof of Theorem 1

Our proof utilizes techniques from [Zanette and Brunskill, 2019].

**Notations.** For two functions  $f(x) \geq 0$  and  $g(x) \geq 0$  defined for  $x \in [0, \infty)$ , we write  $f(x) \lesssim g(x)$  if  $f(x) \leq c \cdot g(x)$  for some absolute constant  $c > 0$ ; we write  $f(x) \gtrsim g(x)$  if  $g(x) \lesssim f(x)$ ; and we write  $f(x) \approx g(x)$  if  $f(x) \lesssim g(x) \lesssim f(x)$ . Moreover, we write  $f(x) = \mathcal{O}(g(x))$  if  $\lim_{x \rightarrow \infty} f(x)/g(x) < c$  for some absolute constant  $c > 0$ ; we write  $f(x) = \Omega(g(x))$  if  $g(x) = \mathcal{O}(f(x))$ ; and we write  $f(x) = \Theta(g(x))$  if  $f(x) = \mathcal{O}(g(x))$  and  $g(x) = \mathcal{O}(f(x))$ . To hide




Figure 3: The effect of number of objectives. The plot shows the performance of **MO-UCBVI** in a simulated random multi-objective MDP with different number of objectives. The plot suggests that regret increases as number of objectives increases, but **MO-UCBVI** achieves sublinear regret in all cases, as predicted by Theorem 1.

the logarithmic factors, we write  $f(x) = \tilde{\mathcal{O}}(g(x))$  if  $f(x) = \mathcal{O}(g(x) \log^d x)$  for some absolute constant  $d > 0$ . For  $a, b \in \mathbb{R}$ , we write  $a \wedge b := \min\{a, b\}$  and  $a \vee b := \max\{a, b\}$ . We will use  $\iota$  to denote a general logarithmic factor. Let  $\pi^k$  be the planning policy at the  $k$ -th episode, i.e., a greedy policy that maximizes  $Q_h^k(x, a)$ . Let  $w_h^k(x, a) := \mathbb{P}\{(x_h, a_h) = (x, a) \mid \pi^k, \mathbb{P}\}$  and  $w^k(x, a) := \sum_h w_h^k(x, a)$ .

### B.1 Good Events and Useful Lemmas

Fix  $\epsilon$  to be a constant to be determined later (we will set  $\epsilon = 1/K$ ). Consider the following events:

$$G_H := \left\{ \forall x, a, h, k, w, \left| (\hat{\mathbb{P}}^k - \mathbb{P})V_{h+1}^* \Big|_{x, a, w} \right| \leq 2\epsilon + \sqrt{\frac{(d \wedge S)H^2}{2N^k(x, a)} \log \frac{6H^2SAK}{\delta\epsilon}} \right\}, \quad (G_H)$$

$$G_{\hat{\mathbb{P}}} := \left\{ \forall x, a, h, k, w, \left| (\hat{\mathbb{P}}^k - \mathbb{P})V_{h+1}^* \Big|_{x, a, w} \right| \leq 2\epsilon + \sqrt{\frac{2(d \wedge S) \|V_{h+1}^*\|_{\hat{\mathbb{P}}^k}^2}{N^k(x, a)} \log \frac{6H^2SAK}{\delta\epsilon}} + \frac{7(d \wedge S)H}{3N^k(x, a)} \log \frac{6H^2SAK}{\delta\epsilon} \right\}, \quad (G_{\hat{\mathbb{P}}})$$

$$G_V := \left\{ \forall x, a, h, k, w, \left| \|V_h^*\|_{\hat{\mathbb{P}}^k} - \|V_h^*\|_{\mathbb{P}} \Big|_{x, a, w} \right| \leq 2\epsilon + \sqrt{\frac{4(d \wedge S)H^2}{N^k(x, a)} \log \frac{6H^2SAK}{\delta\epsilon}} \right\}, \quad (G_V)$$

$$G_P := \left\{ \forall x, a, y, k, \left| \hat{\mathbb{P}}^k(y \mid x, a) - \mathbb{P}(y \mid x, a) \right| \leq \sqrt{\frac{2\mathbb{P}(y \mid x, a)}{N^k(x, a)} \log \frac{2S^2AK}{\delta}} + \frac{2}{3N^k(x, a)} \log \frac{2S^2AK}{\delta} \right\}, \quad (G_P)$$

$$G_{\hat{\mathbb{P}}} := \left\{ \forall x, a, y, k, \left| \hat{\mathbb{P}}^k(y \mid x, a) - \mathbb{P}(y \mid x, a) \right| \leq \sqrt{\frac{2\hat{\mathbb{P}}^k(y \mid x, a)}{N^k(x, a)} \log \frac{2S^2AK}{\delta}} + \frac{7}{3N^k(x, a)} \log \frac{2S^2AK}{\delta} \right\}, \quad (G_{\hat{\mathbb{P}}})$$

$$G_N := \left\{ \forall x, a, k, N^k(x, a) \geq \frac{1}{2} \sum_{j < k} w^j(x, a) - H \log \frac{HSA}{\delta} \right\}. \quad (G_N)$$

**Lemma 1** (Probability of good events.). *Each of the good event holds with probability at least  $1 - \delta$ . In particular, they hold simultaneously with probability at least*

$$\mathbb{P} \{G_H \cap G_{\hat{B}} \cap G_V \cap G_P \cap G_{\hat{P}} \cap G_N\} \geq 1 - 6\delta.$$

*Proof.* We only need to prove that each of the good event holds with probability at least  $1 - \delta$ .

We first show that  $\mathbb{P} \{G_H\} \geq 1 - \delta$  by Hoeffding's inequality and a covering argument. In particular by Hoeffding's inequality we have that for fixed  $x, a, h, k, w$ ,

$$\left| (\hat{\mathbb{P}}^k - \mathbb{P})V_{h+1}^* \Big|_{x,a,w} \right| \leq \sqrt{\frac{H^2}{2N^k(x,a)}} \log \frac{2}{\delta}$$

holds with probability at least  $1 - \delta$ . We first apply a covering argument for the preferences set. We consider an  $\frac{\epsilon}{H}$ -covering  $\mathcal{C}$  for the unit ball  $\{w \in \mathbb{R}^d : \|w\|_1 \leq 1\}$ , then  $|\mathcal{C}| \leq \left(\frac{3H}{\epsilon}\right)^d$ , and for any  $w$  in the ball, there exists  $w' \in \mathcal{C}$  such that  $\|w - w'\|_1 \leq \frac{\epsilon}{H}$ . Then a union bound on  $w \in \mathcal{C}$  and  $(x, a, k, h) \in \mathcal{S} \times \mathcal{A} \times [K] \times [H]$  yields that with probability at least  $1 - \delta$ , the following holds for every  $(x, a, k, h, w) \in \mathcal{S} \times \mathcal{A} \times [K] \times [H] \times \mathcal{C}$ :

$$\left| (\hat{\mathbb{P}}^k - \mathbb{P})V_{h+1}^* \Big|_{x,a,w} \right| \leq \sqrt{\frac{dH^2}{2N^k(x,a)}} \log \frac{6H^2SAK}{\delta\epsilon}$$

Now consider an arbitrary  $w \in \mathcal{W}$  (hence in the unit ball), let  $w' \in \mathcal{C}$  be such that  $\|w - w'\|_1 \leq \frac{\epsilon}{H}$ , then by Lemma 2 we have that with probability at least  $1 - \delta$ ,

$$\begin{aligned} \left| (\hat{\mathbb{P}}^k - \mathbb{P})V_{h+1}^* \Big|_{x,a,w} \right| &\leq \left| \hat{\mathbb{P}}^k V_{h+1}^* \Big|_{x,a,w} - \hat{\mathbb{P}}^k V_{h+1}^* \Big|_{x,a,w'} \right| \\ &\quad + \left| (\hat{\mathbb{P}}^k - \mathbb{P})V_{h+1}^* \Big|_{x,a,w'} \right| + \left| \mathbb{P}V_{h+1}^* \Big|_{x,a,w'} - \mathbb{P}V_{h+1}^* \Big|_{x,a,w} \right| \\ &\leq 2\epsilon + \sqrt{\frac{dH^2}{2N^k(x,a)}} \log \frac{6H^2SAK}{\delta\epsilon} \end{aligned}$$

holds for every  $(x, a, k, h, w) \in \mathcal{S} \times \mathcal{A} \times [K] \times [H] \times \mathcal{W}$ . Similarly, we can apply the above covering argument for the value function set by considering an  $\epsilon$ -covering  $\mathcal{C}$  for the  $H$ -ball  $\{v \in \mathbb{R}^S : \|v\|_\infty \leq H\}$ , then  $|\mathcal{C}| \leq \left(\frac{3H}{\epsilon}\right)^S$ , then we obtain that with probability at least  $1 - \delta$ ,

$$\left| (\hat{\mathbb{P}}^k - \mathbb{P})V_{h+1}^* \Big|_{x,a,w} \right| \leq 2\epsilon + \sqrt{\frac{SH^2}{2N^k(x,a)}} \log \frac{6H^2SAK}{\delta\epsilon}$$

holds for every  $x, a, k, h$  and every value functions, hence for every preference-induced value functions. The two inequalities together show that  $\mathbb{P} \{G_H\} \geq 1 - \delta$ .

Similarly, we can apply the covering arguments with the empirical Bernstein's inequalities [Maurer and Pontil, 2009] to show that  $\mathbb{P} \{G_{\hat{B}}\} \geq 1 - \delta$ .

$\mathbb{P} \{G_V\} \geq 1 - \delta$  is proved by the covering arguments with Theorem 10 from [Maurer and Pontil, 2009].

$\mathbb{P} \{G_P\} \geq 1 - \delta$  and  $\mathbb{P} \{G_{\hat{P}}\} \geq 1 - \delta$  are proved by Bernstein and empirical Bernstein's inequalities [Maurer and Pontil, 2009], respectively.

Finally,  $\mathbb{P} \{G_N\} \geq 1 - \delta$  is due to Lemma F.4 from [Dann et al., 2017].

□

**Lemma 2** (Continuity). *For every  $h$ , we have  $|V_h^*(x; w) - V_h^*(x; w')| \leq (H - h + 1) \|w - w'\|_1$ .*

*Proof.* We prove it by induction. For  $H + 1$ , we have  $V_{H+1}^*(x; w) = V_{H+1}^*(x; w') = 0$ . Now suppose

$$\left| V_{h+1}^*(x; w) - V_{h+1}^*(x; w') \right| \leq (H - h) \|w - w'\|_1,$$



and consider  $h$ . Without loss of generality, suppose  $V_h^*(x; w) \geq V_h^*(x; w')$ . Then

$$\begin{aligned}
|V_h^*(x; w) - V_h^*(x; w')| &= V_h^*(x; w) - V_h^*(x; w') \\
&= \max_a Q_h^*(x, a; w) - \max_{a'} Q_h^*(x, a; w') \\
&\leq Q_h^*(x, a; w) - Q_h^*(x, a; w') \quad (\text{let } a = \arg \max_a Q_h^*(x, a; w)) \\
&\leq |(w - w', \mathbf{r}_h(x, a))| + \sum_{y \in \mathcal{S}} \mathbb{P}(y | x, a) (V_{h+1}^*(y; w) - V_{h+1}^*(y; w')) \\
&\leq \|w - w'\|_1 + (H - h) \|w - w'\|_1 \\
&= (H - h + 1) \|w - w'\|_1.
\end{aligned}$$

Hence the conclusion holds.  $\square$

**Lemma 3** (Bounds for lower order terms). *If events  $(\mathbf{G}_P)$ ,  $(\mathbf{G}_{\hat{P}})$  hold, then for every  $V_1, V_2$  such that  $0 \leq V_1(x; w) \leq V_2(x; w) \leq H$ , and for every  $x, a, h, k, w$ , the following inequalities hold:*

1.  $\left| (\hat{\mathbb{P}}^k - \mathbb{P})(V_2 - V_1) \Big|_{x, a, w} \right| \leq \frac{1}{H} \mathbb{P}(V_2 - V_1) \Big|_{x, a, w} + \frac{2H^2 S_\iota}{N^k(x, a)};$
2.  $\left| (\hat{\mathbb{P}}^k - \mathbb{P})(V_2 - V_1) \Big|_{x, a, w} \right| \leq \frac{1}{H} \hat{\mathbb{P}}(V_2 - V_1) \Big|_{x, a, w} + \frac{3H^2 S_\iota}{N^k(x, a)};$
3.  $\left| (\hat{\mathbb{P}}^k - \mathbb{P})(V_2 - V_1)^2 \Big|_{x, a, w} \right| \leq \mathbb{P}(V_2 - V_1)^2 \Big|_{x, a, w} + \frac{2H^2 S_\iota}{N^k(x, a)};$
4.  $\|V_2 - V_1\|_{\hat{\mathbb{P}}^k}^2 \Big|_{x, a, w} \leq 2 \mathbb{P}(V_2 - V_1)^2 \Big|_{x, a, w} + \frac{2H^2 S_\iota}{N^k(x, a)};$

*Proof.* For simplicity in this proof we denote  $p(y) := \mathbb{P}(y | x, a)$  and  $\hat{p}(y) := \hat{\mathbb{P}}^k(y | x, a)$ .

For the first inequality,

$$\begin{aligned}
&\left| (\hat{\mathbb{P}}^k - \mathbb{P})(V_2 - V_1) \Big|_{x, a, w} \right| \leq \sum_{y \in \mathcal{S}} |\hat{p}^k(y) - p(y)| (V_2(y; w) - V_1(y; w)) \\
&\leq \sum_{y \in \mathcal{S}} \left( \sqrt{\frac{2p(y)\iota}{N^k(x, a)}} + \frac{2\iota}{3N^k(x, a)} \right) (V_2(y; w) - V_1(y; w)) \quad (\text{since } (\mathbf{G}_P) \text{ holds}) \\
&\leq \sum_{y \in \mathcal{S}} \left( \frac{p(y)}{H} + \frac{H\iota}{2N^k(x, a)} + \frac{2\iota}{3N^k(x, a)} \right) (V_2(y; w) - V_1(y; w)) \quad (\text{use } \sqrt{ab} \leq \frac{1}{2}(a + b)) \\
&\leq \sum_{y \in \mathcal{S}} \frac{p(y)}{H} (V_2(y; w) - V_1(y; w)) + \frac{2H^2 S_\iota}{N^k(x, a)} \\
&= \frac{1}{H} \mathbb{P}(V_2 - V_1) \Big|_{x, a, w} + \frac{2H^2 S_\iota}{N^k(x, a)}.
\end{aligned}$$

The second inequality is proved in a same way as the first one, except that in the second step we use event  $(\mathbf{G}_{\hat{P}})$  instead of  $(\mathbf{G}_P)$ .

For the third inequality,

$$\begin{aligned}
& \left| (\widehat{\mathbb{P}}^k - \mathbb{P})(V_2 - V_1)^2 \Big|_{x,a,w} \right| \leq \sum_y |\hat{p}(y) - p(y)| (V_2(y; w) - V_1(y; w))^2 \\
& \leq \sum_y \left( \sqrt{\frac{2p(y)\iota}{N^k(x, a)}} + \frac{2\iota}{3N^k(x, a)} \right) (V_2(y; w) - V_1(y; w))^2 \quad (\text{by } (\mathbf{G_P})) \\
& \leq \sum_y \left( p(y) + \frac{\iota}{2N^k(x, a)} + \frac{2\iota}{3N^k(x, a)} \right) (V_2(y; w) - V_1(y; w))^2 \quad (\text{use } \sqrt{ab} \leq a/2 + b/2) \\
& \leq \mathbb{P}(V_2 - V_1)^2 \Big|_{x,a,w} + \frac{2H^2 S \iota}{N^k(x, a)}.
\end{aligned}$$

For the fourth inequality, notice the following

$$\|V_2 - V_1\|_{\widehat{\mathbb{P}}^k}^2 \leq \widehat{\mathbb{P}}^k(V_2 - V_1)^2 \leq \mathbb{P}(V_2 - V_1)^2 + (\widehat{\mathbb{P}}^k - \mathbb{P})(V_2 - V_1)^2,$$

and use the third inequality. □

## B.2 Proof of the Hoeffding Variant.

In this part, we follow notations in Algorithm 1 and prove the first claim in Theorem 1.

**Bonus.** We set the bonus function in Algorithm 1 to be

$$b^k(x, a) := 2\epsilon + \sqrt{\frac{(d \wedge S)H^2 \iota}{2N^k(x, a)}}, \quad \iota := \log \frac{6H^2 SAK}{\delta \epsilon} \quad (2)$$

where  $\epsilon$  will be set as  $\epsilon = 1/K$ .

**Lemma 4** (Optimistic value function estimation). *If event  $(\mathbf{G_H})$  holds, then  $V_h^*(x; w) \leq V_h^k(x; w)$ .*

*Proof.* We prove the lemma by induction over  $h \in [H]$ . For  $H+1$ ,  $V_{H+1}^k(x; w) = V_{H+1}^*(x; w) = 0$ . Now suppose  $V_{h+1}^k(x; w) \geq V_{h+1}^*(x; w)$  holds for all  $(x, k, w) \in \mathcal{S} \times [K] \times \mathcal{W}$ , and consider  $h$ . Let  $a = \arg \max_a Q_h^*(x, a; w)$ , then by definition we have

$$\begin{aligned}
V_h^*(x; w) &= Q_h^*(x, a; w) = \langle w, \mathbf{r}_h(x, a) \rangle + \mathbb{P}V_{h+1}^* \Big|_{x,a,w}, \\
V_h^k(x; w) &\geq Q_h^k(x, a; w) = H \wedge \left( b^k(x, a) + \langle w, \mathbf{r}_h(x, a) \rangle + \widehat{\mathbb{P}}^k V_{h+1}^k \Big|_{x,a,w} \right).
\end{aligned}$$

If  $Q_h^k(x, a; w) = H$ , then  $V_h^k(x; w) \geq H \geq V_h^*(x; w)$ ; otherwise we have

$$\begin{aligned}
V_h^k(x; w) - V_h^*(x; w) &\geq b^k(x, a) + \widehat{\mathbb{P}}^k V_{h+1}^k \Big|_{x,a,w} - \mathbb{P}V_{h+1}^* \Big|_{x,a,w} \\
&\geq b^k(x, a) + \widehat{\mathbb{P}}^k V_{h+1}^* \Big|_{x,a,w} - \mathbb{P}V_{h+1}^* \Big|_{x,a,w} \quad (\text{by induction hypothesis}) \\
&\geq 0, \quad (\text{by } (\mathbf{G_H}))
\end{aligned}$$

which completes the induction and finishes the proof. □

**Lemma 5** (Per-episode regret decomposition). *If events  $(\mathbf{G_H})$ ,  $(\mathbf{G_P})$  hold, then for every  $k$ , it holds:*

$$V_1^*(x_1; w^k) - V_1^{\pi^k}(x_1; w^k) \lesssim \mathbb{E}_{\mathbb{P}, \pi^k} \sum_{h=1}^H H \wedge \left( \epsilon + \sqrt{\frac{(d \wedge S)H^2 \iota}{N^k(x, a)}} + \frac{H^2 S \iota}{N^k(x_h, a_h)} \right).$$

*Proof.* We first note the following recursion,

$$V_h^k(x; w) - V_h^{\pi^k}(x; w) = Q_h^k(x, a; w) - Q_h^{\pi^k}(x, a; w) \quad (\text{set } a = \pi_h^k(x))$$

$$\begin{aligned}
&\leq b^k + \widehat{\mathbb{P}}^k V_{h+1}^k - \mathbb{P} V_{h+1}^{\pi^k} \Big|_{x,a,w} \\
&= b^k + (\widehat{\mathbb{P}}^k - \mathbb{P}) V_{h+1}^* + (\widehat{\mathbb{P}}^k - \mathbb{P})(V_{h+1}^k - V_{h+1}^*) + \mathbb{P}(V_{h+1}^k - V_{h+1}^{\pi^k}) \Big|_{x,a,w} \\
&\leq 2b^k + (\widehat{\mathbb{P}}^k - \mathbb{P})(V_{h+1}^k - V_{h+1}^*) + \mathbb{P}(V_{h+1}^k - V_{h+1}^{\pi^k}) \Big|_{x,a,w} \quad (\text{by (G}_H\text{)}) \\
&\leq 2b^k + \frac{2H^2 S \iota}{N^k} + \left(1 + \frac{1}{H}\right) \mathbb{P}(V_{h+1}^k - V_{h+1}^{\pi^k}) \Big|_{x,a,w} \quad (\text{by Lemma 3}),
\end{aligned}$$

and by solving the recursion and noting that  $(1 + 1/H)^H \leq e$  we have

$$\begin{aligned}
V_1^k(x_1; w^k) - V_1^{\pi^k}(x_1; w^k) &\leq e \cdot \mathbb{E}_{\mathbb{P}, \pi^k} \sum_{h=1}^H \left( 2b^k(x_h, a_h) + \frac{2H^2 S \iota}{N^k(x_h, a_h)} \right) \\
&\lesssim \mathbb{E}_{\mathbb{P}, \pi^k} \sum_{h=1}^H \left( \epsilon + \sqrt{\frac{(d \wedge S) H^2 \iota}{N^k(x, a)}} + \frac{H^2 S \iota}{N^k(x_h, a_h)} \right).
\end{aligned}$$

The claim is proved by using Lemma 4 and noting that the value functions are bounded by  $H$ .  $\square$

Now we are ready to prove the first part of Theorem 1.

**Theorem 6** (Restatement of Theorem 1, Hoeffding part). *Consider Algorithm 1 with bonus function (2) and  $\epsilon = 1/K$ , then for any sequence of the incoming preferences  $\{w^1, w^2, \dots, w^K\}$ , the regret (1) of Algorithm 1 is bounded by*

$$\text{regret}(K) \lesssim \sqrt{(d \wedge S) \cdot H^3 \text{SAK} \cdot \log(H \text{SAK} / \delta)} + H^2 S^2 A \log^2(H \text{SAK} / \delta)$$

with probability at least  $1 - \delta$ .

*Proof.* First by Lemma 5 we have

$$\begin{aligned}
\text{regret}(K) &= \sum_{k=1}^K V_1^*(x_1; w^k) - V_1^{\pi^k}(x_1; w^k) \\
&\lesssim \sum_{k=1}^K \sum_{h=1}^H \sum_{x,a} w_h^k(x, a) \cdot H \wedge \left( \epsilon + \frac{H^2 S \iota}{N^k(x, a)} + \sqrt{\frac{(d \wedge S) H^2 \iota}{N^k(x, a)}} \right) \\
&\lesssim \underbrace{\sum_{k=1}^K \sum_{h=1}^H \sum_{x, a \notin L^k} w_h^k(x, a) \cdot H}_{(i)} + \underbrace{\sum_{k=1}^K \sum_{h=1}^H \sum_{x, a \in L^k} w_h^k(x, a) \cdot \epsilon}_{(ii)} \\
&\quad + \underbrace{\sum_{k=1}^K \sum_{h=1}^H \sum_{x, a \in L^k} w_h^k(x, a) \cdot \frac{H S \iota}{N^k(x, a)}}_{(iii)} + \underbrace{\sum_{k=1}^K \sum_{h=1}^H \sum_{x, a \in L^k} w_h^k(x, a) \cdot \sqrt{\frac{(d \wedge S) H^2 \iota}{N^k(x, a)}}}_{(iv)},
\end{aligned}$$

where we set

$$L^k := \{(x, a) : \sum_{j < k} w^j(x, a) \geq 2H\iota\}. \quad (3)$$

We next bound each terms separately.

*Term (i).* By (3) we have that (i) =  $\sum_{k=1}^K \sum_{x, a \notin L^k} w^k(x, a) \cdot H \lesssim H^2 S A \iota$ .

*Term (ii).* By definition we have that (ii) =  $\sum_{k=1}^K \sum_{x, a \in L^k} w^k(x, a) \cdot \epsilon \lesssim H \text{SAK} \epsilon$ .

*Term (iii).* The third term is bounded as follows:

$$(iii) = \sum_{k=1}^K \sum_{x, a} w^k(x, a) \cdot \frac{H^2 S \iota}{N^k(x, a)} \cdot \mathbb{1} \left[ \sum_{j < k} w^j(x, a) \geq 2H\iota \right] \quad (\text{by (3)})$$

$$\begin{aligned}
&\lesssim \sum_{x,a} \sum_{k=1}^K w^k(x,a) \cdot \frac{H^2 S \iota}{\sum_{j<k} w^j(x,a) - H \iota} \cdot \mathbb{1} \left[ \sum_{j<k} w^j(x,a) \geq 2H \iota \right] \quad (\text{by } (\mathbf{G}_N)) \\
&\lesssim H^2 S \iota \cdot SA \cdot \iota = H^2 S^2 A \iota^2. \quad (\text{integration trick})
\end{aligned}$$

Term (iv). The fourth term is bounded as follows:

$$\begin{aligned}
(\text{iv}) &= \sum_{k=1}^K \sum_{x,a} w^k(x,a) \cdot \sqrt{\frac{(d \wedge S) H^2 \iota}{N^k(x,a)}} \cdot \mathbb{1} \left[ \sum_{j<k} w^j(x,a) \geq 2H \iota \right] \quad (\text{by } (3)) \\
&\lesssim \sum_{x,a} \sum_{k=1}^K w^k(x,a) \cdot \sqrt{\frac{(d \wedge S) H^2 \iota}{\sum_{j<k} w^j(x,a) - H \iota}} \cdot \mathbb{1} \left[ \sum_{j<k} w^j(x,a) \geq 2H \iota \right] \quad (\text{by } (\mathbf{G}_N)) \\
&\lesssim \sqrt{(d \wedge S) H^2 \iota} \cdot SA \cdot \sqrt{HK} = \sqrt{(d \wedge S) H^3 S A K \iota}. \quad (\text{integration trick})
\end{aligned}$$

Summing up these terms, setting  $\epsilon = 1/K$  and applying Lemma 1, we complete the proof.  $\square$

---

#### Algorithm 4 MO-UCBVI (Bernstein Variant)

---

```

1: initialize history  $\mathcal{H}^0 = \emptyset$ ,  $\iota = \log(HSAK/\delta)$ 
2: for episode  $k = 1, 2, \dots, K$  do
3:    $N^k(x, a), \widehat{\mathbb{P}}^k(y | x, a) \leftarrow \text{Empi-Prob}(\mathcal{H}^{k-1})$ 
4:   receive a preference  $w^k$ 
5:   set  $V_{H+1}^k(x; w^k) = \underline{V}_{H+1}^k(x; w^k) = 0$ 
6:   for step  $h = H, H-1, \dots, 1$  do
7:     for  $(x, a) \in \mathcal{S} \times \mathcal{A}$  do
8:       compute bonus  $b_h^k(x, a) := \sqrt{\frac{(d \wedge S) \iota}{N^k(x, a)}} \cdot \left( \|V_{h+1}^k\|_{\widehat{\mathbb{P}}^k} + \|V_{h+1}^k - \underline{V}_{h+1}^k\|_{\widehat{\mathbb{P}}^k} \right) + \frac{(d \wedge S) H \iota}{N^k(x, a)}$ 
9:       compute bonus  $a_h^k(x, a) := \sqrt{\frac{(d \wedge S) \iota}{N^k(x, a)}} \cdot \left( \|\underline{V}_{h+1}^k\|_{\widehat{\mathbb{P}}^k} + \|V_{h+1}^k - \underline{V}_{h+1}^k\|_{\widehat{\mathbb{P}}^k} \right) + \frac{(d \wedge S) H \iota}{N^k(x, a)}$ 
10:       $Q_h^k(x, a; w^k) = \min \left\{ H, \langle w^k, \mathbf{r}_h(x, a) \rangle + b_h^k(x, a) + \widehat{\mathbb{P}}_h^k V_{h+1}^k|_{x, a, w^k} \right\}$ 
11:       $V_h^k(x; w^k) = \max_{a \in \mathcal{A}} Q_h^k(x, a; w^k)$ 
12:       $\pi_h^k(x) = \arg \max_a Q_h^k(x, a; w^k)$ 
13:       $\underline{Q}_h^k(x, a; w^k) = \max \left\{ 0, \langle w^k, \mathbf{r}_h(x, a) \rangle - a_h^k(x, a) + \widehat{\mathbb{P}}_h^k \underline{V}_{h+1}^k|_{x, a, w^k} \right\}$ 
14:       $\underline{V}_h^k(x; w^k) = \underline{Q}_h^k(x, \pi_h^k(x); w^k)$ 
15:     end for
16:   end for
17:   receive initial state  $x_1^k = x_1$ 
18:   for step  $h = 1, 2, \dots, H$  do
19:     take action  $a_h^k = \pi_h^k(x_h^k)$ , and obtain a new state  $x_{h+1}^k$ 
20:   end for
21:   update history  $\mathcal{H}^k = \mathcal{H}^{k-1} \cup \{x_h^k, a_h^k\}_{h=1}^H$ 
22: end for

```

---

### B.3 Proof of Bernstein Variant

In this part, we follow notations in Algorithm 4 and prove the second claim in Theorem 1.

For a (estimated) transition kernel  $\mathbb{P}$  and a (estimated) value function  $V_h(y; w)$ , define its one-step variance as

$$\|V_h\|_{\mathbb{P}}^2 \Big|_{x, a, w} := \sum_{y \in \mathcal{S}} \mathbb{P}(y | x, a) \left( V_h(y; w) - \mathbb{P} V_h|_{x, a, w} \right)^2.$$

Moreover, we often omit  $|_{x, a, w}$  and simply write  $\|V_h\|_{\mathbb{P}}^2$  when  $x, a, w$  are clear from the context.

**Bonus.** Recall that

$$\begin{aligned} a_h^k(x, a, w) &:= 2\epsilon + \sqrt{\frac{2(d \wedge S)\iota}{N^k(x, a)}} \cdot \left( \|\underline{V}_h^k\|_{\widehat{\mathbb{P}}^k} + \|V_h^k - \underline{V}_h^k\|_{\widehat{\mathbb{P}}^k} \right) + \frac{7(d \wedge S)H\iota}{3N^k(x, a)}, \\ b_h^k(x, a, w) &:= 2\epsilon + \sqrt{\frac{2(d \wedge S)\iota}{N^k(x, a)}} \cdot \left( \|V_h^k\|_{\widehat{\mathbb{P}}^k} + \|V_h^k - \underline{V}_h^k\|_{\widehat{\mathbb{P}}^k} \right) + \frac{7(d \wedge S)H\iota}{3N^k(x, a)}, \end{aligned} \quad (4)$$

where  $\epsilon$  is a parameter to be determined (we will set  $\epsilon = 1/K$ ).

**Lemma 6** (Optimistic value estimation). *Under event  $(\mathbf{G}_{\widehat{\mathbb{B}}})$ , we have that for every  $x, a, k, h, w$ ,*

- $\left| (\widehat{\mathbb{P}}^k - \mathbb{P})V_{h+1}^* \Big|_{x, a, w} \right| \leq b_h^k(x, a, w) \wedge a_h^k(x, a, w);$
- $\underline{V}_h^k(x; w) \leq V_h^*(x; w) \leq V_h^k(x; w)$

*Proof.* We prove the claims by induction. For  $H + 1$ ,  $V_{H+1}^*(x, a; w) = \overline{V}_{H+1}^k(x, a; w) = \underline{V}_{H+1}^k(x, a; w) = 0$ , hence the hypotheses hold. Now suppose the hypotheses hold for  $h + 1$ , i.e.,

$$\left| (\widehat{\mathbb{P}}^k - \mathbb{P})V_{h+1}^* \Big|_{x, a, w} \right| \leq b_h^k(x, a, w) \wedge a_h^k(x, a, w), \quad (5)$$

$$\underline{V}_{h+1}^k(x; w) \leq V_{h+1}^*(x; w) \leq \overline{V}_{h+1}^k(x; w) \quad (6)$$

and consider  $h$ . First we have the following two sets of inequalities:

$$\begin{aligned} \overline{V}_h^k(x; w) - V_h^*(x; w) &\geq \overline{Q}_h^k(x, a; w) - Q_h^*(x, a; w) \quad (\text{set } a = \pi_h^*(x)) \\ &\geq b_h^k(x, a, w) + \widehat{\mathbb{P}}^k \overline{V}_{h+1}^k \Big|_{x, a, w} - \mathbb{P} V_{h+1}^* \Big|_{x, a, w} \\ &\geq b_h^k(x, a, w) + (\widehat{\mathbb{P}}^k - \mathbb{P})V_{h+1}^* \Big|_{x, a, w} \quad (\text{by (6)}) \\ &\geq 0, \quad (\text{by (5)}) \end{aligned}$$

and

$$\begin{aligned} V_h^*(x; w) - \underline{V}_h^k(x; w) &\geq Q_h^*(x, a; w) - \underline{Q}_h^k(x, a; w) \quad (\text{set } a = \underline{\pi}_h^*(x)) \\ &\geq a_h^k(x, a, w) + \mathbb{P} V_{h+1}^* \Big|_{x, a, w} - \widehat{\mathbb{P}}^k \underline{V}_{h+1}^k \Big|_{x, a, w} \\ &\geq a_h^k(x, a, w) + (\mathbb{P} - \widehat{\mathbb{P}}^k)V_{h+1}^* \Big|_{x, a, w} \quad (\text{by (6)}) \\ &\geq 0. \quad (\text{by (5)}) \end{aligned}$$

These justify that

$$\underline{V}_h^k(x; w) \leq V_h^*(x; w) \leq \overline{V}_h^k(x; w). \quad (7)$$

Second, the above inequalities imply that

$$\|V_h^*\|_{\widehat{\mathbb{P}}^k} \leq \|\overline{V}_h^k\|_{\widehat{\mathbb{P}}^k} + \|\overline{V}_h^k - V_h^*\|_{\widehat{\mathbb{P}}^k} \leq \|\overline{V}_h^k\|_{\widehat{\mathbb{P}}^k} + \|\overline{V}_h^k - \underline{V}_h^k\|_{\widehat{\mathbb{P}}^k}, \quad (8)$$

$$\|V_h^*\|_{\widehat{\mathbb{P}}^k} \leq \|\underline{V}_h^k\|_{\widehat{\mathbb{P}}^k} + \|V_h^* - \underline{V}_h^k\|_{\widehat{\mathbb{P}}^k} \leq \|\underline{V}_h^k\|_{\widehat{\mathbb{P}}^k} + \|\overline{V}_h^k - \underline{V}_h^k\|_{\widehat{\mathbb{P}}^k}, \quad (9)$$

therefore,

$$\begin{aligned} \left| (\widehat{\mathbb{P}}^k - \mathbb{P})V_h^* \Big|_{x, a, w} \right| &\leq 2\epsilon + \sqrt{\frac{2(d \wedge S)\iota}{N^k(x, a)}} \cdot \|V_h^*\|_{\widehat{\mathbb{P}}^k} + \frac{7(d \wedge S)H\iota}{3N^k(x, a)} \quad (\text{by } (\mathbf{G}_{\widehat{\mathbb{B}}})) \\ &\leq a_{h-1}^k(x, a; w) \wedge b_{h-1}^k(x, a; w). \quad (\text{by (8), (9) and (4)}) \end{aligned}$$

This completes our induction.  $\square$

**Lemma 7** (Bonus upper bound). *If events  $(G_{\hat{P}})$ ,  $(G_V)$  hold, then we have that*

$$b_h^k \vee a_h^k \Big|_{x,a,w} \leq 3\epsilon + \sqrt{\frac{2(d \wedge S)\iota}{N^k}} \left( \|V_{h+1}^*\|_{\mathbb{P}} + 2\sqrt{\mathbb{P}(V_{h+1}^k - \underline{V}_{h+1}^k)^2} \right) + \frac{15HS\iota}{N^k} \Big|_{x,a,w}.$$

*Proof.* We first prove the bound for  $b_h^k(x, a, w)$ . Note that

$$\begin{aligned} & \|V_{h+1}^k\|_{\hat{\mathbb{P}}^k} + \|V_{h+1}^k - \underline{V}_{h+1}^k\|_{\hat{\mathbb{P}}^k} \leq \|V_{h+1}^*\|_{\hat{\mathbb{P}}^k} + \|V_{h+1}^k - V_{h+1}^*\|_{\hat{\mathbb{P}}^k} + \|V_{h+1}^k - \underline{V}_{h+1}^k\|_{\hat{\mathbb{P}}^k} \\ & \leq \|V_{h+1}^*\|_{\hat{\mathbb{P}}^k} + 2\|V_{h+1}^k - \underline{V}_{h+1}^k\|_{\hat{\mathbb{P}}^k} \quad (\text{by Lemma 6}) \\ & \leq 2\epsilon + \|V_{h+1}^*\|_{\mathbb{P}} + \sqrt{\frac{4(d \wedge S)H^2\iota}{N^k(x, a)}} + 2\|V_{h+1}^k - \underline{V}_{h+1}^k\|_{\hat{\mathbb{P}}^k} \quad (\text{by } (G_V)) \\ & \leq 2\epsilon + \|V_{h+1}^*\|_{\mathbb{P}} + \sqrt{\frac{4(d \wedge S)H^2\iota}{N^k}} + 2\sqrt{2\mathbb{P}(V_{h+1}^k - \underline{V}_{h+1}^k)^2 + \frac{2H^2S\iota}{N^k}} \Big|_{x,a,w} \quad (\text{by Lemma 3}) \\ & \leq 2\epsilon + \|V_{h+1}^*\|_{\mathbb{P}} + 2\sqrt{\mathbb{P}(\bar{V}_{h+1}^k - \underline{V}_{h+1}^k)^2} + \sqrt{\frac{50H^2S\iota}{N^k}} \Big|_{x,a,w}. \end{aligned}$$

Then the bound for  $b_h^k(x, a, w)$  is obtained by substituting the above into (4) and use

$$\sqrt{\frac{2(d \wedge S)\iota}{N^k(x, a)}} \cdot 2\epsilon \leq \epsilon^2 + \frac{2(d \wedge S)\iota}{N^k(x, a)} \leq \epsilon + \frac{2HS\iota}{N^k(x, a)}.$$

Similarly we can obtain the bound for  $a_h^k(x, a, w)$ .  $\square$

**Lemma 8** (Per-episode regret decomposition). *If events  $(G_{\hat{P}})$ ,  $(G_V)$ ,  $(G_P)$  hold, then*

$$\begin{aligned} & V_1^*(x_1; w^k) - V_1^{\pi^k}(x_1; w^k) \lesssim \\ & \mathbb{E}_{\mathbb{P}, \pi^k} \sum_{h=1}^H H \wedge \left( \epsilon + \sqrt{\frac{(d \wedge S)\iota}{N^k}} \left( \|V_{h+1}^*\|_{\mathbb{P}} + \sqrt{\mathbb{P}(V_{h+1}^k - \underline{V}_{h+1}^k)^2} \right) + \frac{H^2S\iota}{N^k} \Big|_{x_h, a_h, w^k} \right). \end{aligned}$$

*Proof.* We first note the following recursion,

$$\begin{aligned} & V_h^k(x; w) - V_h^{\pi^k}(x; w) = Q_h^k(x, a; w) - Q_h^{\pi^k}(x, a; w) \quad (\text{set } a = \pi_h^k(x)) \\ & \leq b_h^k + \hat{\mathbb{P}}^k V_{h+1}^k - \mathbb{P} V_{h+1}^{\pi^k} \Big|_{x,a,w} \\ & = b_h^k + (\hat{\mathbb{P}}^k - \mathbb{P})V_{h+1}^* + (\hat{\mathbb{P}}^k - \mathbb{P})(V_{h+1}^k - V_{h+1}^*) + \mathbb{P}(V_{h+1}^k - V_{h+1}^{\pi^k}) \Big|_{x,a,w} \\ & \leq 2b_h^k + (\hat{\mathbb{P}}^k - \mathbb{P})(V_{h+1}^k - V_{h+1}^*) + \mathbb{P}(V_{h+1}^k - V_{h+1}^{\pi^k}) \Big|_{x,a,w} \quad (\text{by Lemma 6}) \\ & \leq 2b_h^k + \frac{2H^2S\iota}{N^k} + \left(1 + \frac{1}{H}\right) \mathbb{P}(V_{h+1}^k - V_{h+1}^{\pi^k}) \Big|_{x,a,w} \quad (\text{by Lemma 3}), \end{aligned}$$

and by solving the recursion and noting that  $(1 + 1/H)^H \leq e$  we have

$$\begin{aligned} & V_1^k(x_1; w^k) - V_1^{\pi^k}(x_1; w^k) \leq e \cdot \mathbb{E}_{\mathbb{P}, \pi^k} \sum_{h=1}^H \left( 2b_h^k(x_h, a_h) + \frac{2H^2S\iota}{N^k(x_h, a_h)} \right) \\ & \lesssim \mathbb{E}_{\mathbb{P}, \pi^k} \sum_{h=1}^H \left( \epsilon + \sqrt{\frac{(d \wedge S)\iota}{N^k}} \cdot \left( \|V_{h+1}^*\|_{\mathbb{P}} + \sqrt{\mathbb{P}(V_{h+1}^k - \underline{V}_{h+1}^k)^2} \right) + \frac{H^2S\iota}{N^k} \Big|_{x_h, a_h, w^k} \right), \end{aligned}$$

where the last inequality is by Lemma 7. The claim is proved by using Lemma 6 and noting that the value functions are bounded by  $H$ .  $\square$

**Lemma 9** (Lower order regrets). *By setting  $\epsilon = 1/K$ , we have that*

$$\sum_{k=1}^K \mathbb{E}_{\mathbb{P}, \pi^k} \left[ \sum_{h=1}^H \mathbb{P}(V_{h+1}^k - \underline{V}_{h+1}^k)^2 \Big|_{x_h, a_h, w^k} \right] \lesssim H^5 S^2 A \iota^2.$$

*Proof.* We first note the following recursion,

$$\begin{aligned} V_h^k(x; w) - \underline{V}_h^k(x; w) &= Q_h^k(x, a; w) - \underline{Q}_h^k(x, a; w) \quad (\text{set } a = \pi_h^k(x)) \\ &\leq a_h^k + b_h^k + \widehat{\mathbb{P}}^k(V_{h+1}^k - \underline{V}_{h+1}^k) \Big|_{x, a, w} \\ &= a_h^k + b_h^k + (\widehat{\mathbb{P}}^k - \mathbb{P})(V_{h+1}^k - \underline{V}_{h+1}^k) + \mathbb{P}(V_{h+1}^k - \underline{V}_{h+1}^k) \Big|_{x, a, w} \\ &\leq a_h^k + b_h^k + \frac{2H^2 S \iota}{N^k} + \left(1 + \frac{1}{H}\right) \mathbb{P}(V_{h+1}^k - \underline{V}_{h+1}^k) \Big|_{x, a, w} \quad (\text{by Lemma 3}) \\ &\leq 6\epsilon + 6H \sqrt{\frac{2(d \wedge S)\iota}{N^k}} + \frac{30HS\iota}{N^k} + \frac{2H^2 S \iota}{N^k} + \left(1 + \frac{1}{H}\right) \mathbb{P}(V_{h+1}^k - \underline{V}_{h+1}^k) \Big|_{x, a, w} \\ &\quad (\text{by Lemmas 6 and 7}) \\ &\leq 6\epsilon + \sqrt{\frac{100H^2 S \iota}{N^k}} + \frac{32H^2 S \iota}{N^k} + \left(1 + \frac{1}{H}\right) \mathbb{P}(V_{h+1}^k - \underline{V}_{h+1}^k) \Big|_{x, a, w}, \end{aligned}$$

by solving which and noting that  $(1 + 1/H)^H \leq e$  and that the value functions are bounded by  $H$ , we obtain that

$$V_h^k(x; w) - \underline{V}_h^k(x; w) \lesssim \mathbb{E}_{\mathbb{P}, \pi^k} \sum_{t \geq h} H \wedge \left( \epsilon + \sqrt{\frac{100H^2 S \iota}{N^k(x_h, a_h)}} + \frac{H^2 S \iota}{N^k(x_h, a_h)} \right),$$

therefore

$$\begin{aligned} \left( V_h^k(x; w) - \underline{V}_h^k(x; w) \right)^2 &\lesssim \left( \mathbb{E}_{\mathbb{P}, \pi^k} \sum_{t \geq h} H \wedge \left( \epsilon + \sqrt{\frac{H^2 S \iota}{N^k(x_h, a_h)}} + \frac{H^2 S \iota}{N^k(x_h, a_h)} \right) \right)^2 \\ &\lesssim H \cdot \mathbb{E}_{\mathbb{P}, \pi^k} \sum_{t \geq h} \left( H \wedge \left( \epsilon + \sqrt{\frac{H^2 S \iota}{N^k(x_h, a_h)}} + \frac{H^2 S \iota}{N^k(x_h, a_h)} \right) \right)^2 \\ &\lesssim H \cdot \mathbb{E}_{\mathbb{P}, \pi^k} \sum_{t \geq h} H^2 \wedge \left( \epsilon^2 + \frac{H^2 S \iota}{N^k(x_h, a_h)} + \frac{H^4 S^2 \iota^2}{(N^k(x_h, a_h))^2} \right). \end{aligned}$$

Then

$$\begin{aligned} \text{LHS} &:= \sum_{k=1}^K \mathbb{E}_{\mathbb{P}, \pi^k} \sum_{h=1}^H \mathbb{P}(V_{h+1}^k - \underline{V}_{h+1}^k)^2 \Big|_{x_h, a_h, w^k} \\ &= \sum_{k=1}^K \sum_{h=1}^H \sum_{x, a} w_h^k(x, a) \mathbb{P}(V_{h+1}^k - \underline{V}_{h+1}^k)^2 \Big|_{x, a, w^k} \\ &\lesssim H \cdot \sum_{k=1}^K \sum_{h=1}^H \sum_{x, a} w_h^k(x, a) \cdot \mathbb{E}_{\mathbb{P}, \pi^k} \sum_{t \geq h} H^2 \wedge \left( \epsilon^2 + \frac{H^2 S \iota}{N^k(x_h, a_h)} + \frac{H^4 S^2 \iota^2}{(N^k(x_h, a_h))^2} \right) \\ &\lesssim H^2 \cdot \sum_{k=1}^K \sum_{h=1}^H \sum_{x, a} w_h^k(x, a) \cdot H^2 \wedge \left( \epsilon^2 + \frac{H^2 S \iota}{N^k(x_h, a_h)} + \frac{H^4 S^2 \iota^2}{(N^k(x_h, a_h))^2} \right) \\ &\lesssim \sum_{k, h} \sum_{x, a \notin M^k} w_h^k(x, a) H^4 + \sum_{k, h} \sum_{x, a \in M^k} w_h^k(x, a) \left( H^2 \epsilon^2 + \frac{H^4 S \iota}{N^k(x, a)} + \frac{H^6 S^2 \iota^2}{(N^k(x, a))^2} \right), \end{aligned}$$

where we set

$$M^k := \{(x, a) : \sum_{j < k} w^j(x, a) \geq 2HS\iota\}.$$

Then by (G<sub>N</sub>) and the integration tricks (see the proof of Theorem 7), we obtain

$$\begin{aligned} \text{LHS} &\lesssim H^4 \cdot SA \cdot HS\iota + H^2\epsilon^2 \cdot SA \cdot HK + H^4S\iota \cdot SA \cdot \iota + H^6S^2\iota^2 \cdot SA/(HS\iota) \\ &\lesssim H^5S^2A\iota^2 + H^3SAK\epsilon^2 \lesssim H^5S^2A\iota^2, \end{aligned}$$

where the last inequality is because  $\epsilon = 1/K$ .  $\square$

Now we are ready to prove the second part of Theorem 1.

**Theorem 7** (Restatement of Theorem 1, Bernstein part). *Consider Algorithm 4 with bonus function (4) and  $\epsilon = 1/K$ , then for any sequence of the incoming preferences  $\{w^1, w^2, \dots, w^K\}$ , the regret (1) of Algorithm 4 is bounded by*

$$\text{regret}(K) \lesssim \sqrt{(d \wedge S)H^2SAK\iota^2} + H^{2.5}S^2A\iota^2, \quad \iota := \log(HSAK/\delta),$$

with probability at least  $1 - \delta$ .

*Proof.* First by Lemma 8 we have

$$\begin{aligned} \text{regret}(K) &= \sum_{k=1}^K V_1^*(x_1; w^k) - V_1^{\pi^k}(x_1; w^k) \\ &\lesssim \sum_{k=1}^K \sum_{h=1}^H \sum_{x,a} w_h^k(x, a) H \wedge \left( \epsilon + \frac{H^2S\iota}{N^k} + \sqrt{\frac{(d \wedge S)\iota}{N^k}} \left( \sqrt{\mathbb{P}(V_{h+1}^k - \underline{V}_{h+1}^k)^2} + \|V_{h+1}^*\|_{\mathbb{P}} \right) \right) \\ &\lesssim \underbrace{\sum_{k=1}^K \sum_{h=1}^H \sum_{x,a \notin L^k} w_h^k(x, a) H}_{(i)} + \underbrace{\sum_{k=1}^K \sum_{h=1}^H \sum_{x,a \in L^k} w_h^k(x, a) \epsilon}_{(ii)} + \underbrace{\sum_{k=1}^K \sum_{h=1}^H \sum_{x,a \in L^k} w_h^k(x, a) \frac{H^2S\iota}{N^k(x, a)}}_{(iii)} \\ &\quad + \underbrace{\sum_{k=1}^K \sum_{h=1}^H \sum_{x,a \in L^k} w_h^k(x, a) \cdot \sqrt{\frac{(d \wedge S)\iota}{N^k(x, a)}} \cdot \sqrt{\mathbb{P}(V_{h+1}^k - \underline{V}_{h+1}^k)^2}}_{(iv)} \Big|_{x,a,w^k} \\ &\quad + \underbrace{\sum_{k=1}^K \sum_{h=1}^H \sum_{x,a \in L^k} w_h^k(x, a) \cdot \sqrt{\frac{(d \wedge S)\iota}{N^k(x, a)}} \cdot \|V_{h+1}^*\|_{\mathbb{P}}}_{(v)} \Big|_{x,a,w^k}, \end{aligned}$$

where we set

$$L^k := \{(x, a) : \sum_{j < k} w^j(x, a) \geq 2H\iota\}. \quad (10)$$

We next bound each of these terms separately.

*Term (i).* By (10) we have that (i) =  $\sum_{k=1}^K \sum_{x,a \notin L^k} w^k(x, a) H \lesssim H^2SA\iota$ .

*Term (ii).* By definition we have that (ii) =  $\sum_{k=1}^K \sum_{x,a \in L^k} w^k(x, a) \epsilon \lesssim HSAK\epsilon$ .

*Term (iii).* The third term is bounded as follows:

$$\begin{aligned} (iii) &= \sum_{k=1}^K \sum_{x,a} w^k(x, a) \cdot \frac{H^2S\iota}{N^k(x, a)} \cdot \mathbb{1} \left[ \sum_{j < k} w^j(x, a) \geq 2H\iota \right] \quad (\text{by (10)}) \\ &\lesssim \sum_{x,a} \sum_{k=1}^K w^k(x, a) \cdot \frac{H^2S\iota}{\sum_{j < k} w^j(x, a) - H\iota} \cdot \mathbb{1} \left[ \sum_{j < k} w^j(x, a) \geq 2H\iota \right] \quad (\text{by (G}_N\text{)}) \\ &\lesssim H^2S\iota \cdot SA \cdot \iota = H^2S^2A\iota^2. \quad (\text{integration trick}) \end{aligned}$$



Term (iv). By Lemma 9 and the integration tricks, we bound the fourth term as follows:

$$\begin{aligned} \text{(iv)} &\leq \sqrt{(d \wedge S)\iota} \cdot \sqrt{\sum_{k=1}^K \sum_{h=1}^H \sum_{x,a \in L^k} \frac{w_h^k(x,a)}{N^k(x,a)}} \cdot \sqrt{\sum_{k=1}^K \sum_{h=1}^H \sum_{x,a \in L^k} w_h^k(x,a) \cdot \mathbb{P}(V_{h+1}^k - \underline{V}_{h+1}^k)^2} \\ &\lesssim \sqrt{S\iota} \cdot \sqrt{SA\iota} \cdot \sqrt{H^5 S^2 A\iota^2} \lesssim H^{2.5} S^2 A\iota^2. \end{aligned}$$

Term (v). The fifth term is the leading term. We proceed to bound this term by noting

$$\begin{aligned} \text{(v)} &\leq \sum_{k=1}^K \sum_{h=1}^H \sum_{x,a \in L^k} w_h^k(x,a) \cdot \left( \sqrt{\frac{(d \wedge S)\iota}{N^k(x,a)}} \cdot \left( \|V_{h+1}^{\pi^k}\|_{\mathbb{P}} + \|V_{h+1}^* - V_{h+1}^{\pi^k}\|_{\mathbb{P}} \right) \right) \\ &\leq \sqrt{(d \wedge S)\iota} \cdot \sqrt{\sum_{k=1}^K \sum_{h=1}^H \sum_{x,a \in L^k} \frac{w_h^k(x,a)}{N^k(x,a)}} \\ &\quad \left( \underbrace{\sum_{k=1}^K \sum_{h=1}^H \sum_{x,a \in L^k} w_h^k(x,a) \cdot \|V_{h+1}^{\pi^k}\|_{\mathbb{P}}^2}_{\text{(v1)}} + \underbrace{\sum_{k=1}^K \sum_{h=1}^H \sum_{x,a \in L^k} w_h^k(x,a) \cdot \|V_{h+1}^* - V_{h+1}^{\pi^k}\|_{\mathbb{P}}^2}_{\text{(v2)}} \right), \end{aligned}$$

where

$$\text{(v1)} = \sum_{k=1}^K \text{Var}_{\pi^k, \mathbb{P}} \left[ \sum_{h=1}^H \langle w^k, \mathbf{r}_h(x_h, a_h) \rangle \right] \leq H^2 K$$

by the law of total variance and that the cumulative reward cannot exceed  $H$ , and

$$\begin{aligned} \text{(v2)} &\leq \sum_{k=1}^K \sum_{h=1}^H \sum_{x,a \in L^k} w_h^k(x,a) \cdot \mathbb{P}(V_{h+1}^* - V_{h+1}^{\pi^k})^2 \Big|_{x,a,w^k} \\ &\leq H \cdot \sum_{k=1}^K \sum_{h=1}^H \sum_{x,a \in L^k} w_h^k(x,a) \cdot \mathbb{P}(V_{h+1}^* - V_{h+1}^{\pi^k}) \Big|_{x,a,w^k} \\ &\leq H^2 \cdot \sum_{k=1}^K \left( V_1^*(x_1; w^k) - V_1^{\pi^k}(x_1; w^k) \right) = H^2 \cdot \text{regret}(K). \end{aligned}$$

Hence,

$$\begin{aligned} \text{(v)} &\lesssim \sqrt{(d \wedge S)\iota} \cdot \sqrt{SA\iota} \cdot \left( \sqrt{H^2 K} + H \cdot \sqrt{\text{regret}(K)} \right) \\ &\lesssim \sqrt{(d \wedge S)H^2 SAK\iota^2} + \sqrt{(d \wedge S)H^2 SA\iota^2} \cdot \sqrt{\text{regret}(K)}. \end{aligned}$$

Summing up terms (i) to (v) and choosing  $\epsilon = 1/K$ , we obtain

$$\begin{aligned} \text{regret}(K) &\lesssim H^2 SA\iota + HSAK\epsilon + H^2 S^2 A\iota^2 + H^{2.5} S^2 A\iota^2 \\ &\quad + \sqrt{(d \wedge S)H^2 SAK\iota^2} + \sqrt{(d \wedge S)H^2 SA\iota^2} \cdot \sqrt{\text{regret}(K)} \\ &\lesssim H^{2.5} S^2 A\iota^2 + \sqrt{(d \wedge S)H^2 SAK\iota^2} + \sqrt{(d \wedge S)H^2 SA\iota^2} \cdot \sqrt{\text{regret}(K)}, \end{aligned}$$

solving which we obtain

$$\text{regret}(K) \lesssim \sqrt{(d \wedge S)H^2 SAK\iota^2} + H^{2.5} S^2 A\iota^2.$$

Applying Lemma 1 completes the proof.  $\square$

## C Proof of Theorem 3 and Corollary 5

In this section, we follow notations in Algorithm 2 and Algorithm 3 and prove Theorem 3. Let  $\bar{V}_h^k(x; w) := \max_a \bar{Q}_h^k(x, a; w)$ . Let  $\pi^k$  be the planning policy at the  $k$ -th episode and  $\bar{\pi}^k$  be the exploration policy at the  $k$ -th episode.

**Bonus.** We set the bonus functions in Algorithm 3 and Algorithm 2 to be

$$b^k(x, a) := 2\epsilon + \sqrt{\frac{(d \wedge S)H^2\iota}{2N^k(x, a)}}, \quad c^k(x, a) := \frac{3H^2S\iota}{N^k(x, a)} + 2b^k(x, a), \quad \iota := \log \frac{6H^2SAK}{\delta\epsilon}, \quad (11)$$

where  $\epsilon$  is a parameter to be decided later (we will set  $\epsilon = 1/K$ ).

We note Lemma 4 applies for the planning phase. The follows lemma relates planning error with exploration regret.

**Lemma 10** (Planning error). *If events  $(G_H)$ ,  $(G_{\hat{P}})$  hold, then for every  $x, a, h, k, w$ ,*

$$V_h^k(x, w) - V_h^{\pi^k}(x, w) \leq \left(1 + \frac{1}{H}\right)^{H-h+1} \cdot \bar{V}_h^k(x).$$

In particular  $V_1^k(x_1; w) - V_1^{\pi^k}(x_1; w) \leq e \cdot \bar{V}_1^k(x)$ .

*Proof.* We prove the conclusion by induction. For  $H + 1$ , the conclusion holds since both the left high side and the right hand side are zero. Next we assume the conclusion holds for  $h + 1$ , i.e.,

$$V_{h+1}^k(x, w) - V_{h+1}^{\pi^k}(x, w) \leq \left(1 + \frac{1}{H}\right)^{H-h} \cdot \bar{V}_{h+1}^k(x), \quad (12)$$

and consider  $h$ . Recall the definitions in Algorithm 2 and Algorithm 3. Set  $a = \pi_h^k(x)$ , then we have

$$V_h^k(x; w) = Q_h^k(x, a; w) = H \wedge \left( \langle w, \mathbf{r}_h(x, a) \rangle + \hat{\mathbb{P}}^k V_{h+1}^k \Big|_{x, a, w} + b^k(x, a) \right), \quad (13)$$

$$V_h^{\pi^k}(x; w) = Q_h^{\pi^k}(x, a; w) = \langle w, \mathbf{r}_h(x, a) \rangle + \mathbb{P} V_{h+1}^{\pi^k} \Big|_{x, a, w}, \quad (14)$$

and

$$\bar{V}_h^k(x) = \max_{a'} \bar{Q}_h^k(x, a') \geq \bar{Q}_h^k(x, a) = H \wedge \left( \hat{\mathbb{P}}^k \bar{V}_{h+1}^k \Big|_{x, a} + c^k(x, a) \right).$$

If  $\bar{Q}_h^k(x, a) = H$ , clearly we have  $V_h^k(x; w) - V_h^{\pi^k}(x; w) \leq H \leq \bar{V}_h^k(x)$ , thus the conclusion holds. In the following suppose

$$\bar{V}_h^k(x) = \max_{a'} \bar{Q}_h^k(x, a') \geq \bar{Q}_h^k(x, a) = \hat{\mathbb{P}}^k \bar{V}_{h+1}^k \Big|_{x, a} + c^k(x, a). \quad (15)$$

Then (13) and (14) yield

$$\begin{aligned} & V_h^k(x, w) - V_h^{\pi^k}(x, w) \leq \hat{\mathbb{P}}^k V_{h+1}^k \Big|_{x, a, w} - \mathbb{P} V_{h+1}^{\pi^k} \Big|_{x, a, w} + b^k(x, a) \\ &= \hat{\mathbb{P}}^k (V_{h+1}^k - V_{h+1}^{\pi^k}) + (\mathbb{P} - \hat{\mathbb{P}}^k) (V_{h+1}^* - V_{h+1}^{\pi^k}) + (\hat{\mathbb{P}}^k - \mathbb{P}) V_{h+1}^* \Big|_{x, a, w} + b^k(x, w) \\ &\leq \hat{\mathbb{P}}^k (V_{h+1}^k - V_{h+1}^{\pi^k}) + (\mathbb{P} - \hat{\mathbb{P}}^k) (V_{h+1}^* - V_{h+1}^{\pi^k}) \Big|_{x, a, w} + 2b^k(x, a). \quad (\text{by Lemma 4}) \\ &\leq \left(1 + \frac{1}{H}\right) \hat{\mathbb{P}}^k (V_{h+1}^k - V_{h+1}^{\pi^k}) \Big|_{x, a, w} + \frac{3H^2S\iota}{N^k(x, a)} + 2b^k(x, a) \quad (\text{by Lemma 3}) \\ &\leq \left(1 + \frac{1}{H}\right) \hat{\mathbb{P}}^k \left( \left(1 + \frac{1}{H}\right)^{H-h} \bar{V}_{h+1}^k \right) \Big|_{x, a} + \frac{3H^2S\iota}{N^k(x, a)} + 2b^k(x, a) \quad (\text{by (12)}) \\ &= \left(1 + \frac{1}{H}\right)^{H-h+1} \hat{\mathbb{P}}^k \bar{V}_{h+1}^k \Big|_{x, a} + c^k(x, a) \quad (\text{by (11)}) \\ &\leq \left(1 + \frac{1}{H}\right)^{H-h+1} \bar{Q}_h^k(x, a) \leq \left(1 + \frac{1}{H}\right)^{H-h+1} \bar{V}_h^k(x) \quad (\text{by (15)}). \end{aligned}$$

Thus the conclusion also holds for  $h$ . By induction we complete the proof.  $\square$

The following lemma is a consequence of Theorem 6 with zero reward.

**Lemma 11** (Zero reward regret). *In (11) set  $\epsilon = 1/K$ , then with probability at least  $1 - \delta$ , the total value (regret) of Algorithm 2 is bounded by*

$$\sum_{k=1}^K \bar{V}_1^k(x_1) \lesssim \sqrt{(d \wedge S) \cdot H^3 S A K \iota} + H^2 S^2 A \iota^2, \quad \iota := \log(H S A K / \delta).$$

*Proof.* Use Theorem 6 with preference/reward set to be zero.  $\square$

We are ready to prove Theorem 3.

**Theorem 8** (Restatement of Theorem 3). *Consider Algorithm 2 and Algorithm 3 with bonus function (11) and  $\epsilon = 1/K$ , and suppose that Algorithm 2 is run for*

$$K \gtrsim \frac{(d \wedge S) \cdot H^3 S A \iota}{\epsilon^2} + \frac{H^2 S^2 A \iota^2}{\epsilon}, \quad \iota := \log \frac{H S A}{\delta \epsilon}$$

*episodes, then Algorithm 3 outputs an  $(\epsilon, \delta)$ -PAC policy for preference-free exploration.*

*Proof.* Recall that in the planning phase, we uniformly sample one of the  $K$  policies as our final policy. Thus, with probability at least  $1 - \delta$ ,

$$\begin{aligned} V_1^*(x_1; w) - V_1^\pi(x_1; w) &= \frac{1}{K} \sum_{k=1}^K \left( V_1^*(x_1; w) - V_1^{\pi^k}(x_1; w) \right) \\ &\leq \frac{1}{K} \sum_{k=1}^K \left( V_1^k(x_1; w) - V_1^{\pi^k}(x_1; w) \right) \quad (\text{by Lemma 4}) \\ &\leq \frac{e}{K} \sum_{k=1}^K \bar{V}_1^k(x_1) \quad (\text{by Lemma 10}) \\ &\lesssim \sqrt{\frac{(d \wedge S) \cdot H^3 S A \iota}{K}} + \frac{H^2 S^2 A \iota^2}{K}, \quad (\text{by Lemma 11}) \end{aligned}$$

where  $\iota = \log(H S A K / \delta)$ , for the right hand side to be bounded by  $\epsilon$ , it suffices to set

$$K \gtrsim \frac{(d \wedge S) \cdot H^3 S A}{\epsilon^2} \log \frac{H S A}{\delta \epsilon} + \frac{H^2 S^2 A}{\epsilon} \log^2 \frac{H S A}{\delta \epsilon}.$$

$\square$

### C.1 Proof of Corollary 5

For the first conclusion, we simply set  $d = S A$  and apply Theorem 3.

For the second conclusion about non-stationary MDPs, the proof logic follows with small revisions.

First, as the MDP is non-stationary, we replace  $\mathbb{P}$  by  $\{\mathbb{P}_h\}_{h=1}^H$ . Similar in the algorithms, we need to replace  $N^k(x, a)$  by  $\{N_h^k(x, a)\}_{h=1}^H$ , which represents the number of visits to  $(x, a)$  at step  $h$  upto episode  $k$ . Then the empirical transition probability will be estimated by  $\hat{\mathbb{P}}_h^k(y | x, a) = N_h^k(x, a, y) / N_h^k(x, a)$ , for  $N_h^k(x, a) \geq 1$ .

Second, note that for stationary MDP, we have  $\sum_{(x,a)} N^K(x, a) = H K$ , but for non-stationary MDP, this needs to be replaced by  $\sum_{(x,a)} N_h^K(x, a) = K$  for every  $h$ . Then for non-stationary MDP, the integration tricks used in the proof of Theorem 6 needs to be revised accordingly. Let us take

term (iv) as an example, the revised analysis should be:

$$\begin{aligned}
\text{(iv)} &= \sum_{k=1}^K \sum_{h=1}^H \sum_{x,a} w_h^k(x,a) \cdot \sqrt{\frac{(d \wedge S)H^2\iota}{N_h^k(x,a)}} \cdot \mathbf{1} \left[ \sum_{j < k} w_h^j(x,a) \geq 2\iota \right] \quad (\text{by a revised (3)}) \\
&\lesssim \sum_{h=1}^H \sum_{x,a} \sum_{k=1}^K w_h^k(x,a) \cdot \sqrt{\frac{(d \wedge S)H^2\iota}{\sum_{j < k} w_h^j(x,a) - \iota}} \cdot \mathbf{1} \left[ \sum_{j < k} w_h^j(x,a) \geq 2\iota \right] \quad (\text{by a revised (G}_N\text{)}) \\
&\lesssim \sqrt{(d \wedge S)H^2\iota} \cdot HSA \cdot \sqrt{K} = \sqrt{(d \wedge S)H^4SAK\iota}. \quad (\text{integration trick})
\end{aligned}$$

Therefore the obtained bound has an enlarged  $H$  dependence.

Last, repeating the proof of Theorem 3 with the revised notations and the above revised inequalities, we obtain the exploration trajectory complexity for non-stationary MDP, where the order of  $H$  is enlarged due to the revised Theorem 6.

## D Proof of Theorem 4

In this section, we follow notations in Algorithm 2 and Algorithm 3 and prove Theorem 4. Our construction of hard instance is based on the results from Jin et al. [2020].

The following two definitions are migrated from reward-free exploitation (RFE) [Jin et al., 2020] to preference-free exploitation (PFE) in MORL.

**Definition 1** (Set of MOMDPs). Fix  $\mathcal{S}, \mathcal{A}, H, \mathbf{r}, \mathcal{W}$  as the state sets, action sets, length of horizon, reward vector, preferences set, respectively. Then a set of transition probabilities  $\mathcal{P}$  induces a set of MOMDPs, denoted as  $(\mathcal{S}, \mathcal{A}, H, \mathcal{P}, \mathbf{r}, \mathcal{W}) := \{(\mathcal{S}, \mathcal{A}, H, \mathbb{P}, \mathbf{r}, \mathcal{W}) : \mathbb{P} \in \mathcal{P}\}$ .

**Definition 2** ( $(\epsilon, p)$ -correctness). We say a PFE algorithm is  $(\epsilon, p)$ -correct for a set of MOMDPs  $(\mathcal{S}, \mathcal{A}, H, \mathcal{P}, \mathbf{r}, \mathcal{W})$ , if with probability at least  $1 - p$ ,

$$V_1^{\pi_w}(x_1; w; \mathbb{P}) \geq V_1^*(x_1; w; \mathbb{P}) - \epsilon, \quad \text{for all } w \in \mathcal{W} \text{ and } \mathbb{P} \in \mathcal{P},$$

where  $\pi_w$  is the policy given by the PFE algorithm for preference  $w \in \mathcal{W}$ .

**Basic hard instance.** We specify the following set of MOMDPs  $(\mathcal{S}, \mathcal{A}, 2, \mathcal{P}_{\text{basic}}(s, \epsilon), \mathbf{r}, \mathcal{W})$  as a basic hard instance:

- A state set  $\mathcal{S} := \{s\} \cup [d]$ , where  $s$  is the initial state, and  $|\mathcal{S}| = d + 1$ .
- An action set  $\mathcal{A} := [A]$ .
- A horizon length  $H = 2$ .
- A  $d$ -dimensional action-independent reward vector, described by a matrix  $\mathbf{r} := (0, I_{d \times d}) \in \mathbb{R}^{d \times (d+1)}$ , i.e., the  $i$ -th row  $\mathbf{r}^{(i)}(\cdot) \in \mathbb{R}^{d+1}$  sets reward 1 at state  $i$ , and reward 0 elsewhere.
- A preference set  $\mathcal{W} := \{w \in \mathbb{R}^d, \|w\|_1 = 1\}$ . For  $w \in \mathcal{W}$ ,  $w^\top \mathbf{r} \in \mathbb{R}^{d+1}$  gives a scalarized reward.
- A set of transition probabilities

$$\mathcal{P}_{\text{basic}}(s, \epsilon) := \left\{ \mathbb{P} : \forall a \in [A], i \in [d], \left| \mathbb{P}(i | s, a) - \frac{1}{d} \right| \leq \frac{\epsilon}{d}, \mathbb{P}(i | i, a) = 1 \right\}.$$

For  $\mathbb{P} \in \mathcal{P}_{\text{basic}}(s, \epsilon)$ , the initial state is always set to be  $s$ ; then the agent takes an action and transits to  $[d]$  with a near uniform distribution; finally,  $[d]$  are all absorbing states.

For the above basic hard instance, Jin et al. [2020] gives the following lemma to characterizes its learning complexity.

**Lemma 12** (Lower bound on the sample complexity of PFE for the basic hard instance). Fix  $\epsilon \leq 1, p \leq \frac{1}{2}, A \geq 2$ , suppose  $d \geq \mathcal{O}(\log A)$ . There exists a distribution  $\mathcal{D}$  over  $\mathcal{P}_{\text{basic}}(s, \epsilon)$ , such that any  $(\epsilon/12, p)$ -correct PFE algorithm  $\text{alg}$  for  $(\mathcal{S}, \mathcal{A}, 2, \mathcal{P}_{\text{basic}}, \mathbf{r}, \mathcal{W})$  must satisfy

$$\mathbb{E}_{\mathbb{P} \sim \mathcal{D}} \mathbb{E}_{\mathbb{P}, \text{alg}} [K] \geq \Omega\left(\frac{dA}{\epsilon^2}\right),$$

where  $K$  is the number of trajectories collected by  $\text{alg}$  in the exploration phase.

*Proof.* See Jin et al. [2020], Lemma D.2. Note that by the construction of the basic hard instance, the preference-weighted reward recovers any reward that is action-independent, thus for the basic hard instance defined in the language of PFE, it equals to the basic hard instance studied by Jin et al. [2020], Lemma D.2 for RFE.  $\square$

**Hard instance in full version.** Based on the basic hard instance, we next build the hard instance in full version that witnesses the lower bound in Theorem 4. Let  $n = 2^{\ell_0}$ . The hard instance is a set of MOMDPs  $(\mathcal{S}, \mathcal{A}, H, \mathcal{P}, \mathbf{r}, \mathcal{W})$  specified as follows:

- A state set

$$\mathcal{S} := \{(s, \ell) : s \in [2^\ell], \ell = 0, \dots, \ell_0\} \cup \{(s, \ell_0 + 1) : s \in [d]\},$$

where  $(0, 0)$  is the initial state. The states can be viewed as a  $\ell_0$ -layer binary tree with  $d$  leaves attached to the last layer. Clearly,  $|\mathcal{S}| = 2^{\ell_0+1} - 1 + d = 2n - 1 + d = \Theta(n)$ .

- An action set  $\mathcal{A} := [A]$ .
- A horizon length  $H$ .
- A  $2d$ -dimensional action-independent reward vector, described by a matrix

$$\mathbf{r} := \begin{pmatrix} 0_{d \times (n-1)} & A_{d \times n} & 0_{d \times d} \\ 0_{d \times (n-1)} & 0_{d \times n} & I_{d \times d} \end{pmatrix} \in \mathbb{R}^{2d \times (2n-1+d)},$$

where  $A \in \mathbb{R}^{d \times n}$  is a fixed matrix that we will specify later.

- A preference set  $\mathcal{W} := \mathcal{W}_A \oplus \mathcal{W}_{\text{basic}} \subset \mathbb{R}^{2d}$ . Here  $\mathcal{W}_{\text{basic}} \subset \mathbb{R}^d$  is the preference set for the basic hard instance and gives the set of the last  $d$ -dimensions of a preference vector. And we define  $\mathcal{W}_A := \{Ae_i, e_i \in \mathbb{R}^n, i = 1, \dots, n\} \subset \mathbb{R}^d$ , which gives the set of the first  $d$ -dimensions of a preference vector. Here  $e_i$  is the  $i$ -th standard coordinate vector for  $\mathbb{R}^n$ . Then  $w^\top \mathbf{r} \in \mathbb{R}^{2n-1+d}$  gives a scalarized reward. The first  $2n - 1$  dimensions of the scalarized reward  $w^\top \mathbf{r}$  are always zero, i.e., the reward for the states in the  $0, \dots, \ell_0$ -th layers is always zero. The last  $d$  dimensions of the scalarized reward is sampled from  $\{A^\top Ae_i, i = 1, \dots, n\}$ , which specifies the reward for states in the  $(\ell_0 + 1)$ -th layer.
- A set of transition probabilities

$$\begin{aligned} \mathcal{P} := \{ & \mathbb{P}(x_{\ell+1} = (s, \ell + 1) \mid x_\ell = (s, \ell), a = 1) = 1, \ell = 0, \dots, \ell_0 - 1, \\ & \mathbb{P}(x_{\ell+1} = (2^\ell + s, \ell + 1) \mid x_\ell = (s, \ell), a > 1) = 1, \ell = 0, \dots, \ell_0 - 1; \\ & \mathbb{P}(x_{\ell_0+1} \mid x_{\ell_0} = (s, \ell_0), a) = \mathbb{P}_{\text{basic}}^s, \mathbb{P}_{\text{basic}}^s \in \mathcal{P}_{\text{basic}}(s, \epsilon)\}. \end{aligned}$$

Notice the transition probabilities up to the  $\ell_0$ -th layer are fixed. And the transition probability from the  $\ell_0$ -th layer to the  $(\ell_0 + 1)$ -th layer is specified by  $n$  multiples of the transition probabilities that are defined earlier for the basic hard instance. In sum, a transition probability  $\mathbb{P} \in \mathcal{P}$  corresponds to  $n$  multiples of transition probabilities  $\mathbb{P}_{\text{basic}}^s \in \mathcal{P}_{\text{basic}}(s, \epsilon)$ ,  $s \in [n]$ .

The following lemma is the key ingredient for obtaining lower bound for PFE.

**Lemma 13.** *Let  $e_1, \dots, e_n$  be the standard coordinate vector for  $\mathbb{R}^n$ . Let  $e_0$  be the zero vector. Suppose  $d \geq \frac{200}{\epsilon^2} \log(n + 1)$ . Then there exists a matrix  $A \in \mathbb{R}^{d \times n}$  such that*

$$\|A^\top Ae_i - e_i\|_\infty \leq \epsilon_1, \quad i = 1, \dots, n.$$

*Proof.* For  $d \geq \frac{8}{\epsilon^2} \log(n + 1)$ , by Johnson-Lindenstrauss lemma [Johnson and Lindenstrauss, 1984], there exists matrix  $A \in \mathbb{R}^{d \times n}$ , such that

$$(1 - \epsilon) \|e_i - e_j\|_2 \leq \|Ae_i - Ae_j\|_2 \leq (1 + \epsilon) \|e_i - e_j\|_2, \quad \forall i, j \in \{0, 1, \dots, n\}.$$

Thus for  $i \neq j$ ,

$$\begin{aligned} \langle Ae_i, Ae_j \rangle &= \frac{1}{2} \left( \|A(e_i - e_0)\|_2^2 + \|A(e_j - e_0)\|_2^2 - \|A(e_i - e_j)\|_2^2 \right) \\ &\begin{cases} \leq \frac{1}{2} (2(1 + \epsilon)^2 - 2(1 - \epsilon)^2) = 4\epsilon, \\ \geq \frac{1}{2} (2(1 - \epsilon)^2 - 2(1 + \epsilon)^2) = -4\epsilon. \end{cases} \end{aligned}$$

And for  $i = j$ ,

$$\langle Ae_i, Ae_i \rangle = \|A(e_i - e_0)\|_2^2 \begin{cases} \leq (1 + \epsilon)^2 \leq 1 + 3\epsilon, \\ \geq (1 - \epsilon)^2 \geq 1 - 3\epsilon. \end{cases}$$

In sum

$$\delta_{i,j} - 4\epsilon \leq \langle A^\top Ae_i, e_j \rangle \leq \delta_{i,j} + 4\epsilon.$$

Note that  $A^\top Ae_i = \sum_{j=1}^n \langle A^\top Ae_i, e_j \rangle e_j$ . Thus the above inequality implies  $\|A^\top Ae_i - e_i\|_\infty \leq 4\epsilon$ . A rescale of  $\epsilon$  completes the proof.  $\square$

We can then extend results from Jin et al. [2020] to PFE using Lemma 13.

**Lemma 14** (Generalized Lemma D.3 of Jin et al. [2020]). *Fix  $\epsilon < \frac{1}{4}$ ,  $\epsilon_0 < \frac{1}{16H}$  and  $\epsilon_1 < \frac{1}{16}$ . Consider  $\mathbb{P} \in \mathcal{P}$  and  $w = (Ae_s, v) \in \mathcal{W}$ , where  $v \in \mathbb{R}^d$  is nearly uniform, i.e.,*

$$\left\| v - \frac{1}{d} \mathbf{1} \right\|_\infty \leq \frac{\epsilon_0}{d}.$$

*Then if policy  $\pi$  is  $\epsilon$ -optimal, i.e.,  $V_1^\pi(\mathbb{P}, w) \geq V_1^*(\mathbb{P}, w) - \epsilon$ , it must visit state  $(s, \ell_0)$  with probability at least  $\frac{1}{2}$ , i.e.,*

$$\mathbb{P}^\pi[x_{\ell_0} = (s, \ell_0)] \geq \frac{1}{2}.$$

*Proof.* From the preference vector we compute the scalarized reward as  $r = (0, A^\top Ae_s, v)$ , i.e., states in the  $0, \dots, (\ell_0 - 1)$ -th layers has zero reward, states in the  $\ell_0$ -th layer has reward as  $A^\top Ae_s$ , which takes value nearly 1 at state  $(s, \ell_0)$  and value nearly 0 at other states, and states in the last layer has nearly uniform reward  $v$ .

Let  $\bar{v} := \frac{1}{d} \sum_{y=1}^d v[y]$  be the expected reward in the final layer under a uniform visiting distribution. Then by construction we have  $|v(y) - \bar{v}| \leq \epsilon_0$ , which implies  $|v(y) - v(y')| \leq 2\epsilon_0$ . Therefore,

$$|\mathbb{P}v(\cdot) - \mathbb{P}'v(\cdot)| \leq 2\epsilon_0, \text{ for any two transition probability } \mathbb{P} \text{ and } \mathbb{P}'.$$

Now we can compute the value of a policy  $\pi$ :

$$\begin{aligned} V_1^\pi &= \langle \mathbb{P}^\pi[(\cdot, \ell_0)], A^\top Ae_s(\cdot) \rangle + (H - \ell_0 - 1) \langle \mathbb{P}^\pi[(\cdot, \ell_0 + 1)], v(\cdot) \rangle \\ &\quad \text{(since the last layer is absorbing)} \\ &\leq (1 + \epsilon_1) \mathbb{P}^\pi[(s, \ell_0)] + \epsilon_1 (1 - \mathbb{P}^\pi[(s, \ell_0)]) + (H - \ell_0 - 1) \langle \mathbb{P}^\pi[(\cdot, \ell_0 + 1)], v(\cdot) \rangle \\ &\quad \text{(by Lemma 13)} \\ &\leq \epsilon_1 + \mathbb{P}^\pi[(s, \ell_0)] + (H - \ell_0 - 1) \langle \mathbb{P}^\pi[(\cdot, \ell_0 + 1)], v(\cdot) \rangle. \end{aligned}$$

On the other hand we can lower bound the value of the optimal policy by a policy such that  $x_{\ell_0} = (s, \ell_0)$  is taken with probability 1 (this is doable since the transition probability before the  $\ell_0$ -th layer is deterministic):

$$V_1^* \geq A^\top Ae_s(s) + (H - \ell_0 - 1) \langle \mathbb{P}^*[(\cdot, \ell_0 + 1)], v(\cdot) \rangle \geq 1 - \epsilon_1 + (H - \ell_0 - 1) \langle \mathbb{P}^*[(\cdot, \ell_0 + 1)], v(\cdot) \rangle.$$

In sum

$$\begin{aligned} \frac{1}{4} &\geq \epsilon \geq V_1^* - V_1^\pi \geq 1 - \mathbb{P}^\pi[(s, \ell_0)] - 2\epsilon_1 - 2(H - \ell_0 - 1)\epsilon_0 \\ &\geq 1 - 2\epsilon_1 - 2H\epsilon_0 - \mathbb{P}^\pi[(s, \ell_0)] \geq \frac{3}{4} - \mathbb{P}^\pi[(s, \ell_0)]. \end{aligned}$$

$\square$

**Lemma 15** (Lemma D.4 of Jin et al. [2020]). *Suppose  $H \geq 2(\ell_0 + 1)$ . Then a PFE algorithm  $\text{alg}$  that is  $(\epsilon, p)$ -correct for the hard instance induces  $n$  PFE algorithms  $\text{alg}^s$ ,  $s \in [n]$ , which are all  $(\frac{\epsilon}{4H}, p)$ -correct for the basic hard instance.*

*Proof.* Suppose policy  $\pi$  satisfies  $V_1^\pi(\mathbb{P}, w) \geq V_1^*(\mathbb{P}, w) - \epsilon$ , for all  $\mathbb{P} \in \mathcal{P}$  and  $w \in \mathcal{W}$ . Then by setting  $w = (0, A^T A e_s, v)$  for  $s = 1, \dots, n$ , we obtain  $n$  sub-policies, each of which visits a corresponding state from  $(1, \ell_0), \dots, (n, \ell_0)$  with probability at least  $\frac{1}{2}$  by Lemma 14 and induces a near-optimal policy for the basic hard instance. The last claim is since

$$\begin{aligned} \epsilon &\geq V_1^*(\mathbb{P}, w) - V_1^\pi(\mathbb{P}, w) \\ &\geq \mathbb{P}^\pi [x_{\ell_0} = (s, \ell_0)] \cdot (H - \ell_0 - 1) \cdot (V_1^*(s; \mathbb{P}_{\text{basic}}^s, v) - V_1^\pi(s; \mathbb{P}_{\text{basic}}^s, v)) \\ &\geq \frac{1}{2}(H - \ell_0 - 1) \cdot (V_1^*(s; \mathbb{P}_{\text{basic}}^s, v) - V_1^\pi(s; \mathbb{P}_{\text{basic}}^s, v)) \quad (\text{by Lemma 14}) \\ &\geq \frac{H}{4} \cdot (V_1^*(s; \mathbb{P}_{\text{basic}}^s, v) - V_1^\pi(s; \mathbb{P}_{\text{basic}}^s, v)), \quad (\text{since } H \geq 2(\ell_0 + 1)) \end{aligned}$$

which implies  $V_1^\pi(s; \mathbb{P}_{\text{basic}}^s, v) \geq V_1^*(s; \mathbb{P}_{\text{basic}}^s, w) - \frac{\epsilon}{4H}$  holds with probability  $p$ . □

We are now ready to state Theorem 4 formally and deliver the proof.

**Theorem 9** (Restatement of Theorem 4). *Fix  $\epsilon, p$ . There exists a set of MOMDPs induced by a set of transition probabilities  $\mathcal{P}$ , and a distribution  $\mathcal{D}$  over  $\mathcal{P}$ , such that if a PFE algorithm (alg) is  $(\epsilon, p)$ -correct for the set of MOMDPs, then the number of trajectories  $K$  that alg needs to collect in the exploitation phase must satisfy*

$$\mathbb{E}_{\mathbb{P} \sim \mathcal{D}} \mathbb{E}_{\text{alg}} [K] \geq \Omega \left( \frac{\min \{d, S\} \cdot H^2 S A}{\epsilon^2} \right).$$

*Proof.* Let  $K^s$  be the number of visits to state  $(s, \ell_0)$  by alg. Then according to Lemma 15, there are  $n$  induced algorithms  $\text{alg}^s$ ,  $s \in [n]$  that is  $(\frac{\epsilon}{4H}, p)$ -correct for the basic hard instance, and each of them collects  $k^s$  number of trajectories.

However, by Lemma 12 there exists a distribution  $\mathcal{D}^s$  over  $\mathcal{P}_{\text{basic}}(s, \frac{\epsilon}{4H})$  such that

$$\mathbb{E}_{\mathbb{P} \sim \mathcal{D}^s} \mathbb{E}_{\text{alg}^s} [K^s] \geq \Omega \left( \frac{d A H^2}{\epsilon^2} \right).$$

Notice that during each episode, alg can and only can visit one of  $(s, \ell_0)$  in the hard instance, thus we have  $K = \sum_{s \in [n]} K^s$ . Making a summation we obtain

$$\mathbb{E} [K] \geq \Omega \left( \frac{d \cdot n A H^2}{\epsilon^2} \right) = \Omega \left( \frac{d S A H^2}{\epsilon^2} \right).$$

On the other hand, clearly in our construction, we can set  $d = n$  where Lemma 13 holds trivially. Then by the same procedure we obtain

$$\mathbb{E} [K] \geq \Omega \left( \frac{n \cdot n A H^2}{\epsilon^2} \right) = \Omega \left( \frac{S^2 A H^2}{\epsilon^2} \right).$$

In sum we have

$$\mathbb{E} [K] \geq \Omega \left( \frac{\min \{d, S\} \cdot S A H^2}{\epsilon^2} \right).$$

□

## E Proof of Theorem 2

We now prove Theorem 2 based on Theorem 9.

**Theorem 10** (Restatement of Theorem 2). *Fix  $S, A, H$ . Suppose  $K > 0$  is sufficiently large. Then for any algorithm that runs for  $K$  episodes, there exists a set of MDPs  $\mathcal{P}$  and a sequence of (necessarily adversarially chosen) preferences  $\{w^1, \dots, w^K\}$ , such that*

$$\mathbb{E}_{\mathbb{P} \sim \mathcal{P}, \text{alg}} [\text{regret}(K)] \geq \frac{1}{c} \cdot \sqrt{\min \{d, S\} H^2 S A K}$$

for some absolute constant  $c > 1$ , where the expectation is taken with respect to the randomness of drawing an MDP and an algorithm collecting a dataset from the chosen MDP during the exploration phase.

*Proof.* Let us fix  $d, S, A, H$  and a MDP structure as described in the proof of Theorem 9. Let  $K$  to be a large number, and  $\epsilon := \sqrt{\min\{d, S\} H^2 SA / K_0}$  to be a small number.

We only consider the randomness of (i) choosing an MDP with transition  $\mathbb{P}$  and (ii) an algorithm collecting a dataset  $\mathcal{H}^K$  from the chosen MDP during the exploration phase. In other words, we will take expectation over all the randomness during the planning phase, in particular, the considered regret is understood as

$$\text{regret}(K) := \sum_{k=1}^K \mathbb{E} \left[ V_1^*(x_1; w^k) - V_1^{\pi^k}(x_1; w^k) \mid \mathbb{P}, \mathcal{H}^K \right].$$

With this in mind, let us consider the following probability measures:

- Let  $\mathcal{P}_0(\cdot)$  be the probability measure induced by the randomness of drawing a transition kernel  $\mathbb{P}$  from the set of transitions  $\mathcal{P}$ .
- Let  $\mathcal{P}(\cdot \mid \mathbb{P})$  to denote the conditional probability measure induced by the randomness of running an algorithm `alg` on a MDP with a fixed transition kernel  $\mathbb{P}$ , i.e., the probability measure of collecting a dataset.
- Let  $\mathcal{P}(\cdot)$  be the joint probability measure induced by the randomness of drawing a transition kernel  $\mathbb{P}$ , and the randomness of collecting a dataset.

*Converting online-MORL algorithms into PFE algorithms.* Let `alg` be an MORL algorithm determined by any fixed rules (but could contain random bits), and let `adv` be an (arbitrary) adversary that provides preferences for `alg` in the online MORL game. Then we inductively define a preference-free exploration algorithm  $\text{ALG}(K, \mathbb{P})$ , which runs for at most  $K$  episodes on an MDP with transition kernel  $\mathbb{P}$ :

- At episode 1, `adv` chooses a preference  $w^1$  based on the fixed rule of `alg`; and `alg` interacts with  $\mathbb{P}$  under the guidance of  $w^1$  to collect dataset  $\mathcal{H}^1$ ;
- At episode  $k \leq K$ , `adv` chooses a preference  $w^k$  based on the fixed rule of `alg` and the history  $\mathcal{H}^{k-1}$ ; and `alg` interacts with  $\mathbb{P}$  under the guidance of  $w^k$  and collect dataset  $\mathcal{H}^k = \mathcal{H}^{k-1} \cup \{\text{new empirical observations}\}$ .
- At episode  $k \leq K$ , if the agent receives a sequence of planning preferences  $w_1, \dots, w_n$ , the agent outputs a sequence of policies  $\pi_1, \dots, \pi_n$  respectively according to the planning rule that `alg` adopts at the current episode (which is based on history  $\mathcal{H}^{k-1}$ ). For simplicity, we denote  $\text{ALG}(\mathbb{P}, k)[\cdot] : \{w\} \rightarrow \{\pi\}$  be such a planning process, i.e.,  $\text{ALG}(\mathbb{P}, k)[w_i] = \pi_i$ .

We now extend  $\text{ALG}$  into a ‘‘correct’’ PFE algorithm that stops at finite time:

$$\widetilde{\text{ALG}}(\mathbb{P}, k)[\cdot] := \begin{cases} \text{ALG}(\mathbb{P}, k)[\cdot], & 1 \leq k \leq K, \\ \text{PF-UCB}(\mathbb{P}, k - K)[\cdot], & K < k \leq 2K, \\ \text{output random policy for any preference,} & k > 2K. \end{cases}$$

To justify the correctness of this augmented algorithm, consider two error random variables:

$$\begin{aligned} \text{error}(\mathbb{P}, k, \widetilde{\text{ALG}})[w] &:= \mathbb{E}[V_1^*(w; \mathbb{P}) - V_1^\pi(w; \mathbb{P}) \mid \mathbb{P}, \mathcal{H}^{k-1}], \quad \text{for } \pi = \widetilde{\text{ALG}}(\mathbb{P}, k)[w], \\ \text{error}(\mathbb{P}, k, \widetilde{\text{ALG}}) &:= \max_w \text{error}(\mathbb{P}, k, \widetilde{\text{ALG}})[w], \end{aligned}$$

and a stopping time

$$\mathring{K} := (2K) \wedge \min \left\{ k : \text{error}(\mathbb{P}, k, \widetilde{\text{ALG}}) \leq \epsilon \right\}.$$

Clearly  $\mathring{K} \leq 2K$  is bounded.



The augmented PFE algorithm  $\widetilde{\text{ALG}}(\mathbb{P}, \mathring{K})[\cdot]$  is correct. By construction, we have that for any MDP with fixed transition kernel, if  $\widetilde{\text{ALG}}$  is run for  $2K$  episodes, its output becomes the same output of **PF-UCB** that runs for  $K$  episodes, which is  $(\epsilon, \delta \leq 0.9)$ -correct for any preference according to Theorem 8. Mathematically, for any  $\mathbb{P}$ ,

$$\mathbf{P} \left\{ \text{error}(\mathbb{P}, 2K, \widetilde{\text{ALG}}) > \epsilon \mid \mathbb{P} \right\} < \delta \leq 0.9,$$

then over the randomness of choosing a transition kernel  $\mathbb{P} \sim \mathbf{P}_0$ , we have that

$$\mathbf{P} \left\{ \text{error}(\mathbb{P}, 2K, \widetilde{\text{ALG}}) > \epsilon \right\} = \mathbb{E}_{\mathbb{P} \sim \mathbf{P}_0} \left[ \mathbf{P} \left\{ \text{error}(\mathbb{P}, 2K, \widetilde{\text{ALG}}) > \epsilon \mid \mathbb{P} \right\} \right] < 0.9.$$

By discussing whether  $\mathring{K} < 2K$  (where the error is smaller than  $\epsilon$  with probability 1) or  $\mathring{K} = 2K$  (where we apply the above inequality), we have that

$$\mathbf{P} \left\{ \text{error}(\mathbb{P}, \mathring{K}, \widetilde{\text{ALG}}) > \epsilon \right\} < 0.9.$$

This implies that  $\widetilde{\text{ALG}}$  that runs for  $\mathring{K}$  episodes is  $(\epsilon, 0.9)$ -correct for the set of MDP described in the proof of Theorem 9. Then by Theorem 9, we have that

$$\mathbb{E}[\mathring{K}] \geq \frac{1}{c} \cdot \frac{\min\{d, S\} H^2 S A}{\epsilon^2} = \frac{1}{c} \cdot K,$$

for some absolute constant  $c > 1$ . On the other hand,

$$\frac{K}{c_1} \leq \mathbb{E}[\mathring{K}] = \mathbb{E} \left[ \mathbb{1} \left[ \mathring{K} > \frac{K}{2c} \right] \cdot \mathring{K} \right] + \mathbb{E} \left[ \mathbb{1} \left[ \mathring{K} \leq \frac{K}{2c} \right] \cdot \mathring{K} \right] \leq 2K \cdot \mathbf{P} \left\{ \mathring{K} > \frac{K}{2c} \right\} + \frac{K}{2c},$$

which implies that

$$\mathbf{P} \left\{ \mathring{K} > \frac{K}{2c} \right\} \geq \frac{1}{4c},$$

then by the definition of  $\mathring{K}$ , we have that

$$\mathbf{P} \left\{ \text{for each } 1 \leq k \leq \frac{K}{2c}, \text{error}(\mathbb{P}, k, \widetilde{\text{ALG}}) > \epsilon \right\} \geq \mathbf{P} \left\{ \mathring{K} > \frac{K}{2c} \right\} \geq \frac{1}{4c},$$

which further yields that for each  $1 \leq k \leq \frac{K}{2c_1}$ ,

$$\mathbb{E} \left[ \text{error}(\mathbb{P}, k, \widetilde{\text{ALG}}) \right] \geq \frac{1}{4c} \cdot \epsilon, \quad (16)$$

and

$$\mathbb{E} \left[ \sum_{k=1}^{\lfloor K/(2c) \rfloor} \text{error}(\mathbb{P}, k, \widetilde{\text{ALG}}) \right] \geq \frac{1}{4c} \cdot \epsilon \cdot \frac{K}{2c} = \frac{1}{8c^2} \cdot \sqrt{\min\{d, S\} H^2 S A K}. \quad (17)$$

*Regret lower bound.* Recall the definition of  $\text{error}(\mathbb{P}, k, \widetilde{\text{ALG}})$  and note that for  $k \leq K/(2c) \leq K$ ,  $\widetilde{\text{ALG}}[\mathbb{P}, k]$  is the same as  $\text{ALG}(\mathbb{P}, k)$ , which plans in the same way as  $\text{alg}$  in the  $k$ -th episode.

We note that (17) holds for any adversary that chooses its preferences inductively based on the online MORL algorithm. Now we specify the adversary  $\text{adv}$  as follows: for  $k \leq K/(2c)$ ,  $\text{adv}$  specifies a preference  $w^{\text{adv}}$  based on the rule of  $\text{alg}$ , the dataset  $\mathcal{H}^{k-1}$ , and the transition kernel  $\mathbb{P}$ , such that

$$(w^{\text{adv}})^k := \arg \max_w \mathbb{E}[V_1^*(w; \mathbb{P}) - V_1^\pi(w; \mathbb{P}) \mid \mathbb{P}, \mathcal{H}^{k-1}], \quad \text{for } \pi = \widetilde{\text{ALG}}(\mathbb{P}, k)[w],$$

then

$$\text{error}(\mathbb{P}, k, \widetilde{\text{ALG}})[(w^{\text{adv}})^k] = \text{error}(\mathbb{P}, k, \widetilde{\text{ALG}}).$$

As a consequence, we have that

$$\begin{aligned} \mathbb{E}[\text{regret}(K)] &\geq \mathbb{E}[\text{regret}(K/(2c))] = \mathbb{E} \left[ \sum_{k=1}^{\lfloor K/(2c) \rfloor} \text{error}(\mathbb{P}, k, \widetilde{\text{ALG}})[(w^{\text{adv}})^k] \right] \\ &= \mathbb{E} \left[ \sum_{k=1}^{\lfloor K/(2c) \rfloor} \text{error}(\mathbb{P}, k, \widetilde{\text{ALG}}) \right] \geq \frac{1}{8c^2} \cdot \sqrt{\min\{d, S\} H^2 S A K}, \end{aligned}$$

and a rescaling of the constant completes our proof.  $\square$

## F Discussion on [Zhang et al., 2020a]

There is a technical error in the proof of Lemma 2 in [Zhang et al., 2020a]. In particular, the sequence considered in Page 11, between equations (16) and (17),

$$\left\{ \mathbb{1} [k_i \leq K] \cdot \left[ \left[ \widehat{\mathbb{P}}^{k_1} - \mathbb{P}_h \right] \left[ \overline{V}_{h+1}^{\pi_k} - V_{h+1}^{\pi_k} \right] (s, a) + \left( r_h^{k_1} - \mathbb{E}[r_h](s_h^{k_i}, a_h^{k_i}) \right) \right] \right\}_{i=1}^{\tau},$$

is not a martingale difference sequence, since  $\pi_k$  depends on the randomness upto episode  $k$ , however  $k_1, \dots, k_{\tau}$  are all no larger than  $k$ . Thus Azuma-Hoeffding's inequality cannot be applied for this sequence.

To fix this error, one may consider applying a covering argument and union bound over the value functions, but the obtained bound in their Theorem 1 should be revised to  $\tilde{\mathcal{O}}(\log(N)H^5S^2A/\epsilon^2)$ .