
When a Reinforcement Learning Agent Encounters Unknown Unknowns

Juntian Zhu*

School of Artificial Intelligence and Data Science
University of Science and Technology of China
zjt1229@mail.ustc.edu.cn

Miguel de Carvalho

School of Mathematics
University of Edinburgh
Miguel.deCarvalho@ed.ac.uk

Zhouwang Yang

School of Mathematical Sciences
University of Science and Technology of China
yangzw@ustc.edu.cn

Fengxiang He

School of Informatics
University of Edinburgh
F.He@ed.ac.uk

Abstract

An AI agent might surprisingly find she has reached an unknown state which she has never been aware of – an unknown unknown. We mathematically ground this scenario in reinforcement learning: an agent, after taking an action calculated from value functions Q and V defined on the *aware domain*, reaches a state out of the domain. To enable the agent to handle this scenario, we propose an *episodic Markov decision process with growing awareness* (EMDP-GA) model, taking a new *noninformative value expansion* (NIVE) approach to expand value functions to newly aware areas: when an agent arrives at an unknown unknown, value functions Q and V whereon are initialised by noninformative beliefs – the averaged values on the aware domain. This design is out of respect for the complete absence of knowledge in the newly discovered state. The upper confidence bound momentum Q-learning is then adapted to the growing awareness for training the EMDP-GA model. We prove that (1) the regret of our approach is asymptotically consistent with the state of the art (SOTA) without exposure to unknown unknowns in an extremely uncertain environment, and (2) our computational complexity and space complexity are comparable with the SOTA – these collectively suggest that though an unknown unknown is surprising, it will be asymptotically properly discovered with decent speed and an affordable cost.

1 Introduction

Reinforcement learning aims to train an agent that takes sequential actions, moving her state in an environment, for maximising her cumulative rewards received from interactions with the environment [30]. A major model in reinforcement learning is Markov decision process (MDP), assuming the state transitions and rewards from actions do not depend on the history given the current state [35]. Further, episodic MDP (EMDP) allows varying rewards from executing the same action at the same state but in different timings [41]. In parallel, value-based algorithms learn value functions Q and V over the space of (state, action)-pairs and the state space, respectively, which then deliver a policy model for the agent, taking the actions that maximise Q [30]. The domain of value functions is naturally restricted by the awareness of all possible states – mathematically, they are defined over all aware states, termed as *aware domain* in this paper. Correspondingly, the exterior of the aware domain is termed as *unaware domain*.

*This work was completed when J. Zhu was visiting the University of Edinburgh.




Figure 1: Growing awareness

Consider a scenario that is completely untouched in the literature of reinforcement learning – an agent might surprisingly find she has reached an unknown state where she is even unaware of. The former US Secretary of Defense Donald Rumsfeld referred to this scenario as an unknown unknown in a news conference [27]. Unknown unknowns broadly exist in security-critical applications, such as autonomous vehicles [33], quantitative finance [22] and economics [5], when the environment is extremely under-explored and uncertain.

This paper proposes the first mathematical model to ground unknown unknowns in reinforcement learning. Although actions are calculated by maximising value functions Q and V that are *defined* on the aware domain, some of them could move the agent to unknown unknown states in the unaware domain. Based upon this, this paper designs a novel model, *episodic Markov decision process with growing awareness* (EMDP-GA), which employs a newly designed *noninformative value expansion* (NIVE) approach to expand value functions to the newly aware states, when the agent surprisingly reaches there. Figure 1 illustrates the expansion of the aware domain, where S is the whole state space and S_t is the aware domain after exploring t episodes. The initial values are chosen as a noninformative prior [32], such as the averaged values on the previous aware domain, out of respect for the extreme lack of knowledge in this unknown unknown state. We further adapt upper confidence bound momentum Q-learning (UCBMQ) [24] to the growing awareness for training the EMDP-GA, termed as UCBMQ-GA. UCBMQ employs a momentum term to correct the bias in Q-learning, but relies on the upper bounds of value functions Q and V . We prove that our NIVE can preserve this ‘upper-bound’ property when expanding the value functions, which will be invalidated by multiplying NIVE’s expansion by any constant $d < 1$. This confirms the efficiency of NIVE as well as our adaptation made in EMDP-GA and UCBMQ-GA.

Extensive theoretical analysis is conducted for verifying our approach. We analyse the regret over the whole course, which is defined as the cumulative reward loss over the whole learning process against the hidden optimal policy. The regret has an upper bound in order of $\tilde{O}(\sqrt{H^3 S A T} + H^4 S A + H^2 S \sqrt{T})$, where S is the size of the state space, A is the size of the action space, T is the number of episodes, and H is the length of an episode. This regret bound is sublinear with respect to T , and comparable with the state of the art (SOTA), including upper confidence bound-advantage [40], upper confidence bound momentum Q-learning [24] and monotonic value propagation [38], despite our exposure to unknown unknowns, while they are not. The computing complexity and space complexity of employing the UCBMQ-GA approach for training the EMDP-GA model are of order $\mathcal{O}(H(S + A)T)$ and $\mathcal{O}(H S^2 A)$, respectively, which are comparable with the SOTA [24]. These results collectively suggest that although unknown unknowns are surprising, they can nonetheless be discovered asymptotically properly with a decent speed and an affordable cost.

2 Related Works

Episodic Markov decision process (EMDP) A large volume of literature has been on developing EMDP over the last decade. Neu et al. study the adversarial stochastic shortest path problem in

EMDPs [26]. Zimin and Neu investigate EMDPs with a layered structure [41]. Dann and Brunskill study EMDPs from the perspective of probably approximately correct (PAC) learning [6]. EMDP algorithms can be roughly classified into model-based and model-free approaches. Model-based approaches learn the EMDP model prior to or concurrently with policy optimisation and require $\Omega(S^2AH)$ memory space to store the value function estimators [25], in which recent developments have achieved lower regret and burn-in costs [2, 7, 36, 11, 40, 39, 38, 21]. Here, burn-in costs are the minimum sample size needed for an algorithm to operate sample-optimally. In contrast, model-free approaches estimate the optimal value function or optimise the policy directly, without explicit model estimation, and require only $O(SAH)$ memory to store these estimators. Extensive theoretical analyses of model-free approaches have also been conducted [40, 24, 23, 8, 31, 20, 1]. Amongst them, a remarkable result is a minimax regret lower bound, $\min\{\sqrt{SAH^3T}, HT\}$, established by Domingues et al. [10], offering a theoretical benchmark.

Awareness and reverse Bayes Similar terms are seen in the literature of choice theory [13, 28, 19]. Reverse Bayesianism is proposed to model the expanding decision universe as awareness grows [15, 16, 17]. Belief formation and responses to unforeseen events, as predicted by reverse Bayesianism, are experimentally studied [3]. This framework was then applied in economic theory [34, 4, 14] and epistemology [37]. However, no practical approach has been proposed to train an agent to be able to handle unknown unknowns.

3 Notations and Preliminaries

Episodic Markov decision process (EMDP) An EMDP is defined by a tuple $M = (\mathcal{S}, \mathcal{A}, P, r, H, T)$ [12, 18], described below. \mathcal{S} and \mathcal{A} are finite state space and finite action space, respectively; mathematically, they are subsets of Euclidean spaces. If an agent at a state $s \in \mathcal{S}$ takes an action $a \in \mathcal{A}$, it will be transited to a new state $s' \in \mathcal{S}$. An EMDP agent takes *episodic* actions for learning – after an episode, the agent comes back to the initial state s_1 , and starts the next episode. Usually, episodes involve the same number of time steps, termed as *horizon* $H \in \mathbb{N}$. $T \in \mathbb{N}$ is the number of episodes in a learning process. We can thus denote the episode t by the trajectory $(s_1^t, a_1^t, \dots, s_H^t, a_H^t, s_{H+1}^t)$. For any episode, at time step $h \in [H] \triangleq \{1, \dots, H\}$, a transition function $P_h : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ defines the probability $P_h(s'|s, a)$ of transition from state s to state s' when taking action a . This interaction leads to an immediate reward $r_h : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$. Let $\pi_h(a|s)$ denote the probability of taking action a at state s . A policy model $\pi = (\pi_1, \dots, \pi_H) \in \Pi$ characterises the agent’s action course. All possible policy models constitute a policy space $\Pi = \{\pi = (\pi_1, \dots, \pi_H) | \pi_h : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]\}$. We denote the policy used to get episode t by $\pi^t \in \Pi$.

Value functions and regret At time step $h \in [H]$, for any state $s \in \mathcal{S}$ and action $a \in \mathcal{A}$, the Q-value function of a policy π is defined as $Q_h^\pi(s, a) = r_h(s, a) + \mathbb{E}_{p, \pi} \left[\sum_{l=h+1}^H r_l(s_l, a_l) | s_h = s, a_h = a \right]$, and the value function is defined as $V_h^\pi(s) = \mathbb{E}_{p, \pi} \left[\sum_{l=h}^H r_l(s_l, a_l) | s_h = s \right]$. Their relationship is governed by the Bellman equation

$$V_h^\pi(s) = \sum_{a \in \mathcal{A}} \pi(a|s) Q_h^\pi(s, a), \quad Q_h^\pi(s, a) = r_h(s, a) + \sum_{s' \in \mathcal{S}} p(s'|s, a) V_h^\pi(s').$$

The objective of reinforcement learning is to learn a policy that maximises the cumulative rewards, denoted by $V_1^\pi(s)$. Shreve and Bertsekas prove that, under certain conditions, there exists an optimal deterministic policy π^* satisfy $V_h^*(s) \triangleq V_h^{\pi^*}(s) = \max_{\pi} V_h^\pi(s)$, $Q_h^*(s, a) \triangleq Q_h^{\pi^*}(s, a) = \max_{\pi} Q_h^\pi(s, a)$ [29]. This result leads to the Bellman optimality equation:

$$V_h^*(s) = \max_{a \in \mathcal{A}} Q_h^*(s, a), \quad Q_h^*(s, a) = r_h(s, a) + \sum_{s' \in \mathcal{S}} p(s'|s, a) V_h^*(s').$$

The performance of an algorithm is often evaluated by its regret against the optimal policy $\mathcal{R}^T = \sum_{t=1}^T [V_1^*(s_1) - V_1^{\pi^t}(s_1)]$. Note that $V_1^{\pi^t}(s_1) \leq V_1^*(s_1)$ for any $t \in [T]$.

Upper confidence bound momentum Q-learning (UCBMQ) UCBMQ is designed for training an EMDP model [24]. Its objective is to approximate the optimal policy by controlling the regret. It optimises four functions: estimated Q-value function $Q_h^t(s, a)$, bias-value function $V_{h,s,a}^t(s')$, upper bound on the optimal Q-value function $\bar{Q}_h^t(s, a)$, and upper bound on the optimal value function $\bar{V}_h^t(s)$. In iteration t , the agent interacts with the environment to collect episode t , updating the visitation counts $n_h^t(s, a)$ for each state-action pair (s, a) . Next, we update the count-dependent parameters, including the bonus term $\beta_h^t(s, a)$. Finally, we update the four value functions. The core method of updating $Q_h^t(s, a)$ retains the standard Q-learning update, which is inspired by the Bellman optimality equation, and adds a momentum correction to remove the bias. $V_{h,s,a}^t(s')$ is updated as a convex combination of the previous bias-value and the optimistic next-step value, $\bar{Q}_h^t(s, a)$ is updated by adding $\beta_h^t(s, a)$ to $Q_h^t(s, a)$ and $\bar{V}_h^t(s)$ is updated according to Bellman optimality equation. The procedure of UCBMQ is described in Appendix B. UCBMQ can guarantee a regret bound of $\tilde{O}(\sqrt{H^3 SAT} + H^4 SA)$, where H , S , A and T are the horizon, the size of the state space, the size of the action space and the number of learning episodes, respectively. This regret bound matches the lower bound $\Omega(\sqrt{SAH^3 T})$ [10] for large enough T .

4 Growing Awareness in Extremely Under-explored, Uncertain Environment

The awareness of an agent can be considerably restricted and problematic in an extremely under-explored and uncertain environment. We are aiming to grow the awareness along the agent’s course of exploring the environment.

4.1 Mathematically Grounding Awareness in Reinforcement Learning

An example of a spacecraft Consider a spacecraft controlled by a reinforcement learning agent. She has explored her homeland (or home planet), which is not necessarily fully completed. The agent’s value functions are naturally defined on her homeland, which is her current aware domain. After taking an action calculated from the value functions defined on the aware domain, the spacecraft may go out of her home planet. The spacecraft then starts to explore the previously unaware domain. The exploration requires an initialisation for the value functions to kick off. This calls for an expansion strategy for the value functions to newly discovered states. The value functions are optimised afterwards. After sufficient exploration, a good spacecraft is expected to develop an (near-)optimal policy – this is exactly the goal of this paper. Figure 2 illustrates this example.

Mathematically, we translate awareness to the domain of value functions, as the definitions below.

Definition 1 (aware domain). *The value functions are defined over the awareness domain. In this domain, the agent is aware of each state’s existence but needs to fully explore it to ascertain its value.*

Definition 2 (unaware domain). *The unaware domain comprises the states whose existence and values are both unknown to the agent.*

Remark 1. *In our framework, once a state in the unaware domain is explored, it moves to the aware domain and never reverts to the unaware domain. States in the awareness domain remain there permanently. Hence, the agent’s awareness monotonically increases without forgetting previously discovered states.*

4.2 Episodic Markov Decision Processes with Growing Awareness

We now define our *episodic Markov decision processes with growing awareness* (EMDP-GA) model as a tuple $M_{GA} = (\mathcal{S}, \mathcal{S}_0, \mathcal{A}, P, r, H, T)$, where $\mathcal{S}_0 \subseteq \mathcal{S}$ is the initial aware domain. Note that the initial state of each learning episode is fixed as $s_1 \in \mathcal{S}_0$. If the agent reaches an unknown state s after taking an action, she will be aware of the existence of s , so that the aware domain grows by the state s . Let \mathcal{S}_t denote the aware domain after exploring t episodes. At episode $t + 1$, the agent’s trajectory is $(s_1^{t+1}, \dots, s_H^{t+1})$, and the aware domain is updated accordingly as $\mathcal{S}_{t+1} = \mathcal{S}_t \cup \{s_1^{t+1}, \dots, s_H^{t+1}\}$. Note that $s_1 \in \mathcal{S}_0 \subseteq \dots \subseteq \mathcal{S}_T \subseteq \mathcal{S}$.

After t episodes, we denote by U_t the unknown unknown states. The extended aware domain is thus $\tilde{\mathcal{S}}_t = \mathcal{S}_t \cup U_t$. EMDP usually defines state space and action space independently. Inheriting this




Figure 2: Illustration of a spacecraft

convention, in this paper, the agent selects actions from a fixed action space \mathcal{A} , despite the fact that the aware domain is growing, even when the agent is placed at new states. Note that the policy π^t , which is used to get episode t , is defined on the aware domain after exploring $t - 1$ episodes. Thus, the policy $\pi^t = (\pi_1^t, \dots, \pi_H^t)$ is defined on $\tilde{\mathcal{S}}_{t-1}$; i.e., $\pi_h^t : \tilde{\mathcal{S}}_{t-1} \times \mathcal{A} \rightarrow [0, 1]$ for each h . Here, $\pi_h^t(\cdot|s)$ is a distribution over \mathcal{A} . Consequently, π^t effectively constitutes a policy over the entire state space \mathcal{S} .

Homeland condition Exploring in an extremely unexplored, uncertain world requires a ‘homeland’. Space exploration needs an earth. We assume that with high probability, unknown unknowns are no more valuable than the average on known states; a very high-value target still exists but with a high risk. Mathematically, with probability at least $1 - \frac{\delta}{3}$, the average value of the known states exceeds that of any unknown unknown state:

$$\forall s \in \mathcal{S} - \mathcal{S}_t, Q_h^*(s, a) \leq \frac{1}{|\mathcal{S}_t|} \sum_{\hat{s} \in \mathcal{S}_t} Q_h^*(\hat{s}, a), \quad V_h^*(s) \leq \frac{1}{|\mathcal{S}_t|} \sum_{\hat{s} \in \mathcal{S}_t} V_h^*(\hat{s}).$$

It coincides with the fact that exploring in an extremely under-explored, uncertain environment is a high-risk, high-gain mission.

4.3 Noninformative Prior for Domain Expansion to Newly Aware States

When an agent encounters an unknown unknown state that lies outside the currently aware domain, the reinforcement learning algorithm becomes invalid. The unknown unknown status implies that it exhibits properties that significantly differ from any known state, and prior information is completely absent. Out of respect for this extreme lack of knowledge, we design a *noninformative value expansion* (NIVE) approach to expand value functions to the newly discovered states – the value functions are initialised using the averages over the current aware domain. We introduce NIVE with more details below.

Algorithm of NIVE NIVE updates three value functions: estimated Q-value function $Q_h^t(s, a)$, bias-value function $V_{h,s,a}^t(s')$, and upper bound on the optimal value function $\bar{V}_h^t(s)$. When the aware domain grows from \mathcal{S}_{t-1} to \mathcal{S}_t , we would like to expand the domains of three value functions: $Q_h^{t-1}(s, a)$, $\bar{V}_h^{t-1}(s)$ and $V_{h,s,a}^{t-1}(s')$, accordingly. NIVE calculates the average values over the aware domain \mathcal{S}_{t-1} as estimators for states in $\mathcal{S}_t - \mathcal{S}_{t-1}$:

$$Q_{h,avg}^t(a) = \frac{1}{|\mathcal{S}_t|} \sum_{s \in \mathcal{S}_t} Q_h^t(s, a), \quad \bar{V}_{h,avg}^t = \frac{1}{|\mathcal{S}_t|} \sum_{s \in \mathcal{S}_t} \bar{V}_h^t(s), \quad V_{h,s,a,avg}^t = \frac{1}{|\mathcal{S}_t|} \sum_{s' \in \mathcal{S}_t} V_{h,s,a}^t(s'),$$

$$V_{h,a,avg}^t(s') = \frac{1}{|\mathcal{S}_t|} \sum_{s \in \mathcal{S}_t} V_{h,s,a}^t(s'), \quad \text{and} \quad V_{h,a,avg}^t = \frac{1}{|\mathcal{S}_t|^2} \sum_{s,s' \in \mathcal{S}_t} V_{h,s,a}^t(s').$$

In many reinforcement learning algorithms, such as UCBMQ [24] used later, upper bounds $\bar{Q}_h^t(s, a)$ and $\bar{V}_h^t(s)$ are used to calculate actions. After the surprising encounter, the NIVE expands the estimators of value functions to newly discovered states by assigning them the average values

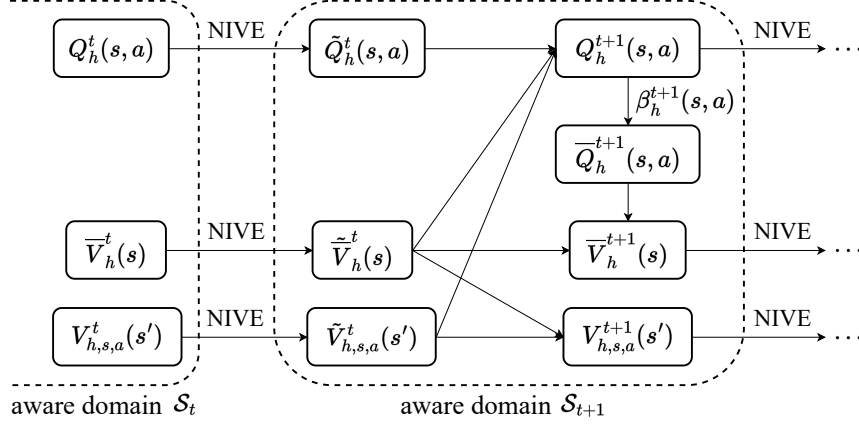


Figure 3: Illustration of updating functions in UCBMQ-GA

computed over known states. Crucially, the expanded functions $\tilde{Q}_h^t(s, a)$ and $\tilde{V}_h^t(s)$ remain valid upper bounds for $Q_h^*(s, a)$ and $V_h^*(s)$, respectively, as the following lemma. The proof is given in Appendix I (implied in Lemma 9).

Lemma 1. *For any time step $t \in [T]$ in episode $h \in [H]$, the following inequality holds with probability at least $1 - \frac{\delta}{2}$,*

$$\tilde{Q}_h^t(s, a) \geq Q_h^*(s, a), \quad \tilde{V}_h^t(s) \geq V_h^*(s).$$

Further, we may prove the following corollary. The proof is given in Appendix E.

Corollary 1. *If one were to use any scalar multiple $d < 1$ of the average values to expand the functions, the upper-bound property could become invalid. In some cases, assigning $d\bar{V}_{h,avg}^t$, ($d \neq 1$) to a newly discovered state would produce an expanded function that no longer upper bounds $V_h^*(s)$.*

This corollary confirms the efficiency of NIVE – multiplying any constant $d < 1$ to NIVE will make it invalid.

4.4 Learning an Agent with Growing Awareness to Unknown Unknowns

We now design an optimisation method for training our EMDP-GA model for obtaining an agent with growing awareness to handle the unknown unknowns. We adapt upper confidence bound momentum Q-learning (UCBMQ) [24] to the growing awareness, termed as UCBMQ-GA.

Algorithm of UCBMQ-GA UCBMQ-GA optimises four functions: estimated Q-value function $Q_h^t(s, a)$, bias-value function $V_{h,s,a}^t(s')$, upper bound on the optimal Q-value function $\tilde{Q}_h^t(s, a)$, and upper bound on the optimal value function $\tilde{V}_h^t(s)$.

Let $\chi_h^t(s, a) = \mathbb{I}(s_h^t = s, a_h^t = a)$ denote the indicator function for the occurrence of (s, a) at time step h in episode t . Thus, $n_h^t(s, a) = \sum_{l=1}^t \chi_h^l(s, a)$ as the visitation count for the pair (s, a) appearing at time step h over t episodes. Define $\chi_h^t(s, a) = n_h^t(s, a) = 0$ if state $s \notin S_t$. Following the convention, and to the brevity, we define $0 \times \infty = 0$ and $\frac{1}{0} = \infty$.

For any transition function p and any function $f : S \rightarrow \mathbb{R}$, we define $pf(s, a) = \mathbb{E}_{s' \sim p(\cdot|s,a)}[f(s')]$. We denote by p_h^t the Dirac distribution concentrated at (s_{h+1}^t) , and then, $p_h^t f(s, a) = f(s_{h+1}^t)$. The procedure of UCBMQ-GA for training EMDA-GA is also explained below. More details are shown in Appendix C.

(0) Initialisation For state $s \in S_0 \subseteq S$, we set $V_{h,s,a}^0 = \bar{V}_h^0 = H$, and $Q_h^0 = 0$. Then, go to (1).

(1) Decision making In episode t , if the current state is aware (i.e., $s \in \mathcal{S}_{t-1}$), the agent takes an action that maximises $\bar{Q}_h^{t-1}(s, a)$. Conversely, if the agent encounters an unknown unknown state (i.e., $s \notin \mathcal{S}_{t-1}$), it chooses an action that maximises $Q_{h,avg}^{t-1}(a)$.

The algorithm terminates here if the termination conditions have been met. Otherwise, if the agent go beyond the aware domain, go to (2); otherwise, go to (3).

(2) Embedded NIVE for growing awareness After completing an episode, the awareness set is updated from \mathcal{S}_{t-1} to \mathcal{S}_t . UCBMQ-GA then employs NIVE to expand $\tilde{Q}_h^{t-1}(s, a)$, $\tilde{V}_h^{t-1}(s)$ and $\tilde{V}_{h,s,a}^{t-1}(s')$ to the newly discovered area, as follows,

$$\tilde{Q}_h^{t-1}(s, a) = \begin{cases} Q_h^{t-1}(s, a), & s \in \mathcal{S}_{t-1}, \\ Q_{h,avg}^{t-1}(a), & s \in \mathcal{S}_t - \mathcal{S}_{t-1}, \end{cases} \quad \tilde{V}_h^{t-1}(s) = \begin{cases} \bar{V}_h^{t-1}(s), & s \in \mathcal{S}_{t-1}, \\ \bar{V}_{h,avg}^{t-1}, & s \in \mathcal{S}_t - \mathcal{S}_{t-1}, \end{cases} \quad \text{and}$$

$$\tilde{V}_{h,s,a}^{t-1}(s') = \begin{cases} V_{h,s,a}^{t-1}(s'), & s, s' \in \mathcal{S}_{t-1}, \\ V_{h,s,a,avg}^{t-1}, & s \in \mathcal{S}_{t-1}, s' \in \mathcal{S}_t - \mathcal{S}_{t-1}, \\ V_{h,a,avg}^{t-1}(s'), & s \in \mathcal{S}_t - \mathcal{S}_{t-1}, s' \in \mathcal{S}_{t-1}, \\ V_{h,a,avg}^{t-1}, & s, s' \in \mathcal{S}_t - \mathcal{S}_{t-1}. \end{cases}$$

The exploration bonus is then calculated with the expanded functions: if $n_h^t(s, a) = 0$ or $s \notin \mathcal{S}_t$, $\beta_h^t(s, a) = H$; otherwise,

$$\beta_h^t(s, a) = 2\sqrt{\frac{\zeta W_h^t(s, a)}{n_h^t(s, a)}} + 53H^3 \frac{\zeta \log(T)}{n_h^t(s, a)} + \sum_{k=1}^t \frac{\chi_h^k(s, a) \hat{\gamma}_h^k(s, a)}{H \log(T) n_h^t(s, a)} p_h^k \left(\tilde{V}_{h,s,a}^{k-1} - \tilde{V}_{h+1}^{k-1} \right) (s, a),$$

$$\text{where } W_h^t(s, a) = \sum_{k=1}^t \frac{\chi_h^k(s, a)}{n_h^k(s, a)} p_h^k \left(\tilde{V}_{h+1}^{k-1} - \sum_{l=1}^t \frac{\chi_h^l(s, a)}{n_h^l(s, a)} p_h^l \tilde{V}_{h+1}^{l-1} \right)^2 (s, a).$$

Then, go to (3).

(3) Optimising value functions The value functions $Q_h^t(s, a)$, $V_{h,s,a}^t(s')$, $\bar{Q}_h^t(s, a)$ and $\bar{V}_h^t(s)$ are then optimised as illustrated in Figure 3. The learning rate and momentum term are chosen as

$$\alpha_h^t(s, a) = \frac{\chi_h^t(s, a)}{n_h^t(s, a)}, \quad \gamma_h^t(s, a) = \chi_h^t(s, a) \frac{H}{H + n_h^t(s, a)} \frac{n_h^t(s, a) - 1}{n_h^t(s, a)}, \quad \text{and}$$

$$\eta_h^t(s, a) = \alpha_h^t(s, a) + \gamma_h^t(s, a).$$

ζ is the exploration threshold which helps in regulating the exploration bonus.

Then, go to (1).

5 Theoretical Analysis

This section presents theoretical analysis of our algorithm.

5.1 Asymptotical Consistency of Learning Unknown Unknowns

Oracle agent The theoretical analysis aims to compare our approach against the (unknown) optimal policy $\pi^* = (\pi_1^*, \dots, \pi_H^*)$, $\pi_h^* : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$. This is characterised by the regret of UCBMQ-GA against this optimal policy, defined as $\mathcal{R}^T = \sum_{t=1}^T [V^{\pi^*}(s_1) - V^{\pi^t}(s_1)]$. To characterise this optimal policy, we assume the existence of an oracle agent that is aware of the entire state space \mathcal{S} and can maximise the value function accordingly.

We make the following assumption about the transition function.

Table 1: Comparisons with the State of the Art

| Algorithm | Regret Bound | Comput. Complex. | Spac. Complex. |
|------------------------|--|--------------------------|-----------------------|
| UCBMQ [24] | $\tilde{\mathcal{O}}(\sqrt{H^3 SAT} + H^4 SA)$ | $\mathcal{O}(H(S + A)T)$ | $\mathcal{O}(HS^2 A)$ |
| MVP [38] | $\tilde{\mathcal{O}}(\sqrt{H^3 SAT})$ | $\mathcal{O}(HSAT)$ | $\mathcal{O}(HS^2 A)$ |
| UCB-Adv [40] | $\tilde{\mathcal{O}}(\sqrt{H^3 SAT} + H^{\frac{33}{4}} S^2 A^{\frac{3}{2}} T^{\frac{1}{4}})$ | $\mathcal{O}(HAT)$ | $\mathcal{O}(HSA)$ |
| Bayes-UCBVI [31] | $\tilde{\mathcal{O}}(\sqrt{H^3 SAT} + H^3 S^2 A)$ | $\mathcal{O}(BHS^2 AT)$ | $\mathcal{O}(HS^2 A)$ |
| UCBMQ-GA (ours) | $\tilde{\mathcal{O}}(\sqrt{H^3 SAT} + H^4 SA + H^2 S\sqrt{T})$ | $\mathcal{O}(H(S + A)T)$ | $\mathcal{O}(HS^2 A)$ |

Assumption 1 (non-zero probability of transition). *For the transition function, We assume that $1 - (\frac{\delta}{3})^{\frac{1}{H[\sqrt{T}]}} \leq \frac{1}{S}$, and the following holds*

$$\forall(s, a, s') \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}, P(s'|s, a) \geq 1 - \left(\frac{\delta}{3}\right)^{\frac{1}{H[\sqrt{T}]}}.$$

Remark 2. When T is large enough, $1 - (\frac{\delta}{3})^{\frac{1}{H[\sqrt{T}]}} \leq \frac{1}{S}$ holds and guarantees that $P(\cdot|s, a)$ can be a distribution over \mathcal{S} . Assumption 1 means that the agent can probably reach any state from the current state when taking any action.

We then define an *aware moment* of a state as follows.

Definition 3 (aware moment). *For any state $s \in \mathcal{S}$, define its aware moment as $t(s) = \min_{t \in [T]} \{t : s \in \mathcal{S}_t\}$; i.e., the first episode after which the agent is aware of s .*

Naturally, $t(s) \geq 0$. We then have an upper bound on the aware moment. The proof is detailed in Appendix F.

Theorem 1. *With probability at least $1 - \frac{\delta}{3}$, the aware moment is upper-bounded: $t(s) \leq \sqrt{T}$, where T is the number of learning episodes.*

Remark 3. *This theorem implies that, with high probability, the agent becomes aware of all states within \sqrt{T} episodes.*

This finding is instrumental in deriving regret upper bounds. Then, we obtain the regret bound as follows. A detailed proof is given in Appendix J.

Theorem 2. *We assume that $T > 3$ and $\zeta = \log\left(\frac{96eHSA(2T+1)}{\delta}\right)$. When the UCNMQ-GA is used to train EMDP-GA, with probability at least $1 - \delta$, the regret satisfies*

$$\mathcal{R}^T \leq \tilde{\mathcal{O}}(\sqrt{H^3 SAT} + H^4 SA + H^2 S\sqrt{T}),$$

where H is the horizon, S is the size of the state space, A is the size of the action space and T is the number of learning episodes.

Remark 4. *This regret bound is sublinear with respect to the number of episodes T . It matches the SOTA, despite that they are not exposed to unknown unknowns, but we are.*

Remark 5. *We have now proved that unknown unknowns could be surprising, but can be asymptotically consistently discovered.*

Comparison with existing methods, but without exposure to unknown unknowns. Compared with SOTA, for $H \leq \frac{A}{S}$, the regret bound of our UCBMQ-GA matches UCBMQ ($\mathcal{O}(\sqrt{SAH^3 T} + SAH^4)$) [24]. For $T \geq H^5 SA$ and $H \leq \frac{A}{S}$, the regret bound of UCBMQ-GA matches monotonic value propagation (MVP, $\mathcal{O}(\sqrt{H^3 SAT})$) [38]. Other methods, upper Bayes-confidence bound value iteration (Bayes-UCBVI) [31] and upper confidence bound-advantage(UCB-Adv) [40], are worse than the three methods in either case. Table 1 shows the comparisons.

Proof sketch The proof is in five steps. $\overset{\circ}{Q}_h^t(s, a)$ and $\overset{\circ}{V}_h^t(s)$ are defined in Appendix A.

Step 1: Upper bound $(\bar{Q}_h^t - Q_h^{\pi^{t+1}})(s, a)$ Combining the estimation of $\bar{Q}_h^t(s, a)$ in Lemma 10 and the upper-bound on $\beta_h^t(s, a)$ in Lemma 11, we obtain the upper-bound of $(\bar{Q}_h^t - Q_h^{\pi^{t+1}})(s, a)$.

Step 2: Upper bound the local optimistic regret $\hat{R}_h^T(s, a)$ We define the local optimistic regret as

$$\hat{R}_h^T(s, a) \triangleq \sum_{t=0}^{T-1} \chi_h^{t+1}(s, a)(\bar{Q}_h^t - Q_h^{\pi^{t+1}})(s, a) = \sum_{t=t(s)-1}^{T-1} \chi_h^{t+1}(s, a)(\bar{Q}_h^t - Q_h^{\pi^{t+1}})(s, a).$$

With the result in **Step 1**, we can decompose $\hat{R}_h^T(s, a)$ and upper bound each term.

Step 3: Replace χ_h^t with \bar{p}_h^t in the upper-bound on $\hat{R}_h^T(s, a)$ Following **Step 2**, we modify the upper bound on the local optimistic regret using $\bar{p}_h^t(s, a)$, the probability to reach (s, a) at time step h in the episode t as below,

$$\begin{aligned} \hat{R}_h^T(s, a) \leq & 63 \log(T) \sqrt{\zeta \sum_{t=t(s)-1}^{T-1} \bar{p}_h^{t+1}(s, a) \text{Var}_{p_h}(V_{h+1}^{\pi^{t+1}})(s, a) + 1754 H^3 \log(T)^2 \zeta} \\ & + (1 + \frac{83}{H}) \sum_{t=t(s)-1}^{T-1} \bar{p}_h^{t+1}(s, a) p_h(\bar{V}_{h+1}^t - V_{h+1}^{\pi^{t+1}})(s, a). \end{aligned}$$

Step 4: Upper bound the regret \hat{R}_h^T at step h We define the step h regret as $\hat{R}_h^T \triangleq \sum_{s \in S} \sum_{t=0}^{T-1} \bar{p}_h^{t+1}(s) (\bar{V}_h^t - V_h^{\pi^{t+1}})(s)$. For each s , the composition of the step h regret follows by $t(s)$. Employing the Cauchy-Schwarz inequality, we get the upper bound.

Step 5: Upper bound the regret \mathcal{R}^T With all the results in the previous steps, we can eventually prove that the regret \mathcal{R}^T can be upper-bounded by \hat{R}_h^T .

5.2 Computational Complexity and Space Complexity

We then prove the computational complexity of employing UCBMQ-GA to train EMDP-GA as below.

Theorem 3. *The computational complexity of employing UCBMQ-GA to train EMDP-GA is of order $\mathcal{O}(H(S + A)T)$, where H is the horizon, S is the size of the state space, A is the size of the action space and T is the number of learning episodes.*

The proof is given in Appendix G. Here, we present an intuitive sketch. The updates of \bar{Q}_h^t , \bar{V}_h^t and $V_{h,s,a}^t$ are carried out in an online manner. At time step h of the episode t , if $(s_h^t, a_h^t) \neq (s, a)$, both learning rate $\alpha_H^t(s, a)$ and momentum term $\gamma_h^t(s, a)$ become zero. Thus, we do not need to update the functions on such (s, a) pairs. Furthermore, the expansions of Q_h^{t-1} , \bar{V}_h^{t-1} , and $V_{h,s,a}^{t-1}$ in each episode depend on their corresponding averaged values on \mathcal{S}_{t-1} .

In parallel, we have the following result for space complexity.

Theorem 4. *The space complexity of employing UCBMQ-GA to train EMDP-GA is $\mathcal{O}(HS^2A)$.*

A detailed proof is presented in Appendix G. Intuitively, this lemma holds because we need to store all the bias-value function $V_{h,s,a}^t(s')$.

Comparison with existing results As shown in Table 1, the computational complexity of UCBMQ-GA is of the same order as our baseline UCBMQ, smaller than MVP ($\mathcal{O}(HSAT)$) and Bayes-UCBVI ($\mathcal{O}(BHS^2AT)$), and a bit larger than UCB-Adv ($\mathcal{O}(HAT)$). The space complexity of UCBMQ-GA is of the same order as UCBMQ, MVP, and Bayes-UCBVI, and still a bit larger than that of UCB-Adv. However, the complexity advances of UCB-Adv are in cost of regret, especially when state space size S and action space size A are large compared with training time T , which is the case in an under-explored environment.

6 Conclusion

We mathematically ground the concept of unknown unknowns in reinforcement learning through our proposed episodic Markov decision process with growing awareness (EMDP-GA). By expanding value functions Q and V to newly encountered states with the noninformative value expansion (NIVE) method, our approach effectively addresses the challenge of unknown unknowns. We adapt upper confidence bound momentum Q-learning (UCBMQ) to train the EMDP-GA model at an affordable cost, achieving a regret bound competitive with existing methods not exposed to unknown unknowns.

Applicability This paper focuses on problems with finite state space and action space. Future works include extending it to continuous state and action spaces.

References

- [1] Priyank Agrawal and Shipra Agrawal. Optimistic Q-learning for average reward and episodic reinforcement learning. *arXiv preprint arXiv:2407.13743*, 2024.
- [2] Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforcement learning. In *International Conference on Machine Learning*, pages 263–272. PMLR, 2017.
- [3] Christoph K Becker, Tigran Melkonyan, Eugenio Proto, Andis Sofianos, and Stefan T Trautmann. Reverse Bayesianism: Revising beliefs in light of unforeseen events. 2020.
- [4] Surajeet Chakravarty, David Kelsey, and Joshua C Teitelbaum. Reverse Bayesianism and act independence. *Journal of Economic Theory*, 203:105495, 2022.
- [5] Michael Chang, Sid Kaushik, S Matthew Weinberg, Tom Griffiths, and Sergey Levine. Decentralized reinforcement learning: Global decision-making via local economic transactions. In *International Conference on Machine Learning*, pages 1437–1447. PMLR, 2020.
- [6] Christoph Dann and Emma Brunskill. Sample complexity of episodic fixed-horizon reinforcement learning. *Annual Conference on Neural Information Processing Systems*, 28, 2015.
- [7] Christoph Dann, Lihong Li, Wei Wei, and Emma Brunskill. Policy certificates: Towards accountable reinforcement learning. In *International Conference on Machine Learning*, pages 1507–1516. PMLR, 2019.
- [8] Christoph Dann, Mehryar Mohri, Tong Zhang, and Julian Zimmert. A provably efficient model-free posterior sampling method for episodic reinforcement learning. *Annual Conference on Neural Information Processing Systems*, 34:12040–12051, 2021.
- [9] Omar D Domingues, Pierre Ménard, Matteo Pirodda, Emilie Kaufmann, and Michal Valko. Regret bounds for kernel-based reinforcement learning. In *International Conference on Machine Learning*. PMLR, 2020.
- [10] Omar Darwiche Domingues, Pierre Ménard, Emilie Kaufmann, and Michal Valko. Episodic reinforcement learning in finite mdps: Minimax lower bounds revisited. In *International Conference on Algorithmic Learning Theory*, pages 578–598. PMLR, 2021.
- [11] Yonathan Efroni, Nadav Merlis, Mohammad Ghavamzadeh, and Shie Mannor. Tight regret bounds for model-based reinforcement learning with greedy policies. *Annual Conference on Neural Information Processing Systems*, 32, 2019.
- [12] Claude-Nicolas Fiechter. Efficient reinforcement learning. In *Annual Conference on Computational Learning Theory*, pages 88–97, 1994.
- [13] Simon Grant and John Quiggin. A preference model for choice subject to surprise. *Theory and Decision*, 79:167–180, 2015.
- [14] Edi Karni, Quitzé Valenzuela-Stookey, and Marie-Louise Vierø. Reverse bayesianism: A generalization. *The B.E. Journal of Theoretical Economics*, 21(2):557–569, 2021.

- [15] Edi Karni and Marie-Louise Vierø. “Reverse Bayesianism”: A choice-based theory of growing awareness. *American Economic Review*, 103(7):2790–2810, 2013.
- [16] Edi Karni and Marie-Louise Vierø. Probabilistic sophistication and reverse bayesianism. *Journal of Risk and Uncertainty*, 50:189–208, 2015.
- [17] Edi Karni and Marie-Louise Vierø. Awareness of unawareness: A theory of decision making in the face of ignorance. *Journal of Economic Theory*, 168:301–328, 2017.
- [18] Emilie Kaufmann, Pierre Ménard, Omar Darwiche Domingues, Anders Jonsson, Edouard Leurent, and Michal Valko. Adaptive reward-free exploration. In *International Conference on Algorithmic Learning Theory*, pages 865–891. PMLR, 2021.
- [19] Asen Kochov. A behavioral definition of unforeseen contingencies. *Journal of Economic Theory*, 175:265–290, 2018.
- [20] Gen Li, Laixi Shi, Yuxin Chen, and Yuejie Chi. Breaking the sample complexity barrier to regret-optimal model-free reinforcement learning. *Information and Inference-A Journal of the IMA*, 12(2):969–1043, 2023.
- [21] Gen Li, Yuting Wei, Yuejie Chi, and Yuxin Chen. Breaking the sample size barrier in model-based reinforcement learning with a generative model. *Operations Research*, 72(1):203–221, 2024.
- [22] Xiao-Yang Liu, Hongyang Yang, Jiechao Gao, and Christina Dan Wang. FinRL: Deep reinforcement learning framework to automate trading in quantitative finance. In *International Conference on AI in Finance*, pages 1–9. ACM, 2021.
- [23] Weichao Mao, Kaiqing Zhang, Ruihao Zhu, David Simchi-Levi, and Tamer Basar. Near-optimal model-free reinforcement learning in non-stationary episodic mdps. In *International Conference on Machine Learning*, pages 7447–7458. PMLR, 2021.
- [24] Pierre Ménard, Omar Darwiche Domingues, Xuedong Shang, and Michal Valko. UCB Momentum Q-learning: Correcting the bias without forgetting. In *International Conference on Machine Learning*, pages 7609–7618. PMLR, 2021.
- [25] Thomas M Moerland, Joost Broekens, Aske Plaat, Catholijn M Jonker, et al. Model-based reinforcement learning: A survey. *Foundations and Trends® in Machine Learning*, 16(1):1–118, 2023.
- [26] Gergely Neu, Andras Gyorgy, and Csaba Szepesvári. The adversarial stochastic shortest path problem with unknown transition probabilities. In *International Conference on Artificial Intelligence and Statistics*, pages 805–813. PMLR, 2012.
- [27] United States Department of Defense. Defense.gov Transcript: DoD News Briefing - Secretary Rumsfeld and Gen. Myers, 2002.
- [28] Evan Piermont. Introspective unawareness and observable choice. *Games and Economic Behavior*, 106:134–152, 2017.
- [29] Steven E Shreve and Dimitri P Bertsekas. Alternative theoretical frameworks for finite horizon discrete-time stochastic optimal control. *SIAM Journal on Control and Optimization*, 16(6):953–978, 1978.
- [30] Richard S Sutton, Andrew G Barto, et al. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge, 1998.
- [31] Daniil Tiapkin, Denis Belomestny, Eric Moulines, Alexey Naumov, Sergey Samsonov, Yunhao Tang, Michal Valko, and Pierre Ménard. From dirichlet to rubin: Optimistic exploration in rl without bonuses. In *International Conference on Machine Learning*, pages 21380–21431. PMLR, 2022.
- [32] Robert Tibshirani. Noninformative priors for one parameter of many. *Biometrika*, 76(3):604–608, 1989.

- [33] Jessica Van Brummelen, Marie O’Brien, Dominique Gruyer, and Homayoun Najjaran. Autonomous vehicle perception: The technology of today and tomorrow. *Transportation Research Part C: Emerging Technologies*, 89:384–406, 2018.
- [34] Marie-Louise Vierø. An intertemporal model of growing awareness. *Journal of Economic Theory*, 197:105351, 2021.
- [35] Chelsea C White III and Douglas J White. Markov decision processes. *European Journal of Operational Research*, 39(1):1–16, 1989.
- [36] 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.
- [37] Carlos Zednik and Frank Jäkel. Bayesian reverse-engineering considered as a research strategy for cognitive science. *Synthese*, 193:3951–3985, 2016.
- [38] Zihan Zhang, Yuxin Chen, Jason D Lee, and Simon S Du. Settling the sample complexity of online reinforcement learning. In *Annual Conference on Learning Theory*, pages 5213–5219. PMLR, 2024.
- [39] Zihan Zhang, Xiangyang Ji, and Simon Du. Is reinforcement learning more difficult than bandits? a near-optimal algorithm escaping the curse of horizon. In *Annual Conference on Learning Theory*, pages 4528–4531. PMLR, 2021.
- [40] Zihan Zhang, Yuan Zhou, and Xiangyang Ji. Almost optimal model-free reinforcement learning via reference-advantage decomposition. *Annual Conference on Neural Information Processing Systems*, 33:15198–15207, 2020.
- [41] Alexander Zimin and Gergely Neu. Online learning in episodic markovian decision processes by relative entropy policy search. *Annual Conference on Neural Information Processing Systems*, 26, 2013.

A Additional Notations and Preliminaries

Table 2: List of Notation

| Description | Notation |
|---|--|
| Size of the state space | $S = \mathcal{S} $ |
| Size of the action space | $A = \mathcal{A} $ |
| Number of learning episodes | T |
| Horizon (length of an episode) | H |
| Aware domain | $\mathcal{S}_t, t \in [T]$ |
| immediate reward | r_h |
| Exploration threshold | $\zeta = \log(\frac{96eHSA(2T+1)}{\delta})$ |
| Indicator function for the occurrence of (s, a) | $\chi_h^t(s, a)$ |
| Visitation count for (s, a) | $n_h^t(s, a)$ |
| First aware domain including s | $t(s)$ |
| Policy at episode t | π^t |
| Optimal policy | π^* |
| Optimal Q-value function | $Q_h^*(s, a), s \in \mathcal{S}, a \in \mathcal{A}$ |
| Optimal value function | $V_h^*(s), s \in \mathcal{S}$ |
| Dirac distribution concentrated at (s_{h+1}^t) | $p_h^t(s' s, a), s, s' \in \mathcal{S}_t$ $\bar{p}_h^t(s' s, a), s, s' \in \mathcal{S}$ |
| Estimated Q-value function | $Q_h^t(s, a), s \in \mathcal{S}_t, a \in \mathcal{A}$ $\bar{Q}_h^t(s, a), s \in \mathcal{S}_{t+1}, a \in \mathcal{A}$ $\overset{\circ}{Q}_h^t(s, a), s \in \mathcal{S}, a \in \mathcal{A}$ |
| Bias-value function | $V_{h,s,a}^t(s'), s, s' \in \mathcal{S}_t$ $\bar{V}_{h,s,a}^t(s'), s, s' \in \mathcal{S}_{t+1}$ $\overset{\circ}{V}_{h,s,a}^t(s'), s, s' \in \mathcal{S}$ |
| Upper bound on Q-value function | $\bar{Q}_h^t(s, a), s \in \mathcal{S}_t, a \in \mathcal{A}$ $\tilde{\bar{Q}}_h^t(s, a), s \in \mathcal{S}_{t+1}, a \in \mathcal{A}$ $\overset{\circ}{\bar{Q}}_h^t(s, a), s \in \mathcal{S}, a \in \mathcal{A}$ |
| Upper bound on value function | $\bar{V}_h^t(s), s \in \mathcal{S}_t$ $\tilde{\bar{V}}_h^t(s), s \in \mathcal{S}_{t+1}$ $\overset{\circ}{\bar{V}}_h^t(s), s \in \mathcal{S}$ |

For $t \in [T]$, we expand $p_h^t(s'|s, a)$, $Q_h^t(s, a)$, $\bar{V}_h^t(s)$, $V_{h,s,a}^t(s)$, $\bar{Q}_h^t(s, a)$ to the whole state space \mathcal{S} .

$$\bar{p}_h^t(s'|s, a) = \begin{cases} p_h^t(s'|s, a), & s, s' \in \mathcal{S}_t, \\ 0, & s \in \mathcal{S}_t, s' \in \mathcal{S} - \mathcal{S}_t, \\ p_h(s'|s, a), & s \in \mathcal{S} - \mathcal{S}_t, \end{cases}$$

$$\overset{\circ}{Q}_h^t(s, a) = \begin{cases} Q_h^t(s, a), & s \in \mathcal{S}_t, \\ Q_{h,avg}^t(a), & s \in \mathcal{S} - \mathcal{S}_t, \end{cases}$$

$$\overset{\circ}{\bar{V}}_h^t(s) = \begin{cases} \bar{V}_h^t(s), & s \in \mathcal{S}_t, \\ \bar{V}_{h,avg}^t, & s \in \mathcal{S} - \mathcal{S}_t, \end{cases}$$

$$\mathring{V}_{h,s,a}^t(s') = \begin{cases} V_{h,s,a}^t(s'), & s, s' \in \mathcal{S}_t, \\ V_{h,s,a,avg}^t, & s \in \mathcal{S}_t, s' \in \mathcal{S} - \mathcal{S}_t, \\ V_{h,a,avg}^t(s'), & s \in \mathcal{S} - \mathcal{S}_t, s' \in \mathcal{S}_t, \\ V_{h,a,avg}^t, & s, s' \in \mathcal{S} - \mathcal{S}_t, \end{cases}$$

$$\mathring{Q}_h^t(s, a) = \begin{cases} Q_h^t(s, a) + \beta_h^t(s, a), & n_h^t(s, a) > 0 \text{ (which implies } s \in \mathcal{S}_t), \\ \mathring{Q}_h^t(s, a) + H, & n_h^t(s, a) = 0 \text{ or } s \in \mathcal{S} - \mathcal{S}_t. \end{cases}$$

From the definitions, it is easy to see that for any s, s' in the aware domain \mathcal{S}_{t+1} , we have

$$\mathring{Q}_h^t(s, a) = \tilde{Q}_h^t(s, a), \quad \mathring{V}_h^t(s) = \tilde{V}_h^t(s), \quad \mathring{V}_{h,s,a}^t(s') = \tilde{V}_{h,s,a}^t(s'),$$

and for any state $s \in \mathcal{S}_t$,

$$\mathring{Q}_h^t(s, a) = \overline{Q}_h^t(s, a).$$

When the count of a state-action pair $n_h^t(s, a) > 0$, we can obtain explicit formulas for the estimate of the Q-function:

$$\begin{aligned} n_h^t(s, a)Q_h^t(s, a) &= \chi_h^t(s, a)[r_h(s, a) + p_h^t \tilde{V}_{h+1}^{t-1}(s, a) + \mathring{\gamma}_h^t(s, a)p_h^t(\tilde{V}_{h+1}^{t-1} - \tilde{V}_{h,s,a}^{t-1})(s, a)] \\ &\quad + (n_h^t(s, a) - \chi_h^t(s, a))\tilde{Q}_h^{t-1}(s, a), \end{aligned}$$

where $\mathring{\gamma}_h^t(s, a) = H \frac{n_h^t(s, a) - 1}{n_h^t(s, a) + H}$.

For $s \in \mathcal{S}_t - \mathcal{S}_{t-1}$, which means $t(s) = t$, we have $n_h^t(s, a) = \chi_h^t(s, a) = 1$, hence $\mathring{\gamma}_h^t(s, a) = 0$. Therefore, the following holds

$$Q_h^t(s, a) = r_h(s, a) + p_h^t \tilde{V}_{h+1}^{t-1}(s, a).$$

For $s \in \mathcal{S}_{t-1}$, we have

$$\begin{aligned} n_h^t(s, a)Q_h^t(s, a) &= \chi_h^t(s, a)[r_h(s, a) + p_h^t \tilde{V}_{h+1}^{t-1}(s, a) + \mathring{\gamma}_h^t(s, a)p_h^t(\tilde{V}_{h+1}^{t-1} - \tilde{V}_{h,s,a}^{t-1})(s, a)] \\ &\quad + n_h^{t-1}(s, a)Q_h^{t-1}(s, a), \end{aligned}$$

$$Q_h^t(s, a) = r_h(s, a) + \frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a)[p_h^k \tilde{V}_{h+1}^{k-1}(s, a) + \mathring{\gamma}_h^k(s, a)p_h^k(\tilde{V}_{h+1}^{k-1} - \tilde{V}_{h,s,a}^{k-1})(s, a)].$$

From all above, we can obtain explicit formulas for Q_h^t ,

$$Q_h^t(s, a) = r_h(s, a) + \frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a)[p_h^k \mathring{V}_{h+1}^{k-1}(s, a) + \mathring{\gamma}_h^k(s, a)p_h^k(\mathring{V}_{h+1}^{k-1} - \mathring{V}_{h,s,a}^{k-1})(s, a)].$$

We can do the same with the bias value function when $n_h^t(s, a) > 0$. For $s, s' \in \mathcal{S}_t$,

$$\begin{aligned} \mathring{V}_{h,s,a}^t(s') &= V_{h,s,a}^t(s') \\ &= \eta_h^t(s, a)\tilde{V}_{h+1}^{t-1}(s') + (1 - \eta_h^t(s, a))\tilde{V}_{h,s,a}^{t-1}(s') \\ &= \eta_h^t(s, a)\mathring{V}_{h+1}^{t-1}(s') + (1 - \eta_h^t(s, a))\mathring{V}_{h,s,a}^{t-1}(s'). \end{aligned}$$

For $s \in \mathcal{S}, s' \in \mathcal{S} - \mathcal{S}_t$,

$$\begin{aligned} \mathring{V}_{h,s,a}^t(s') &= V_{h,s,a,avg}^t \\ &= \frac{1}{|\mathcal{S}_t|} \sum_{s' \in \mathcal{S}_t} V_{h,s,a}^t(s') \\ &= \frac{1}{|\mathcal{S}_t|} \sum_{s' \in \mathcal{S}_t} \eta_h^t(s, a)\tilde{V}_{h+1}^{t-1}(s') + (1 - \eta_h^t(s, a))\tilde{V}_{h,s,a}^{t-1}(s') \\ &= \eta_h^t(s, a)\mathring{V}_{h+1}^{t-1}(s') + (1 - \eta_h^t(s, a))\mathring{V}_{h,s,a}^{t-1}(s'). \end{aligned}$$

For $s \in \mathcal{S} - \mathcal{S}_t$, we get $n_h^t(s, a) = 0$.

We can conclude

$$\begin{aligned}\hat{V}_{h,s,a}^t(s') &= \eta_h^t(s, a) \bar{V}_{h+1}^{\circ t-1}(s') + (1 - \eta_h^t(s, a)) \hat{V}_{h,s,a}^{t-1}(s') \\ &= \frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) [\bar{V}_{h+1}^{\circ k-1}(s') + \gamma_h^k(s, a) (\bar{V}_{h+1}^{\circ k-1} - \hat{V}_{h,s,a}^{\circ k-1})(s')] \quad (1)\end{aligned}$$

$$= \sum_{k=t(s)}^t \tilde{\eta}_h^{t,k}(s, a) \bar{V}_{h+1}^{\circ k-1}(s'), \quad (2)$$

where $\tilde{\eta}_h^{t,k}(s, a) = \eta_h^k(s, a) \prod_{l=k+1}^t (1 - \eta_h^l(s, a))$.

We prove the following lemma to have a basic knowledge about $\bar{V}_h^t(s)$ and $V_{h,s,a}^t(s')$.

Lemma 2. For $s, s' \in \mathcal{S}_t$, the following holds almost surely:

- the sequence $(\bar{V}_h^t(s))_{t \geq t(s)}$ is non-increasing,
- $0 \leq \bar{V}_h^t(s) \leq H$,
- $\bar{V}_{h+1}^t(s') \leq V_{h,s,a}^t(s') \leq H$.

When $t \geq t(s)$, we have $\tilde{V}_h^t = \bar{V}_h^t$.

We can get from the construction of $\bar{V}_h^t(s)$:

$$\bar{V}_h^t(s) = \text{clip}(\max_{a \in \mathcal{A}} \bar{Q}_h^t(s, a), 0, \tilde{V}_h^{t-1}(s))$$

that

$$\tilde{V}_h^{t(s)-1}(s) \geq \bar{V}_h^{t(s)}(s) \geq \bar{V}_h^{t(s)+1}(s) \geq \dots \geq 0.$$

To prove that $\bar{V}_h^t(s) \leq H$, $s \in \mathcal{S}_t$, we proceed by induction. We have $\bar{V}_h^0 = H$, hence this claim holds when $t = 0$. Assuming that $\bar{V}_h^{t-1}(s) \leq H$, $s \in \mathcal{S}_{t-1}$. For $s \in \mathcal{S}_{t-1} \subset \mathcal{S}_t$, we have

$$\bar{V}_h^t(s) \leq \bar{V}_h^{t-1}(s) \leq H.$$

For $s \in \mathcal{S}_t - \mathcal{S}_{t-1}$, we have

$$\bar{V}_h^t(s) \leq \tilde{V}_h^{t-1}(s) = \frac{1}{|\mathcal{S}_{t-1}|} \sum_{\hat{s} \in \mathcal{S}_{t-1}} \bar{V}_h^{t-1}(\hat{s}) \leq H.$$

To prove the third claim, we proceed by induction. When $t = 0$, the claims holds. Assuming $\bar{V}_{h+1}^{t-1}(s') \leq V_{h,s,a}^{t-1}(s') \leq H$ for $s, s' \in \mathcal{S}_{t-1}$. We consider two different cases of $s \in \mathcal{S}_t$.

Case 1: $s \in \mathcal{S}_t - \mathcal{S}_{t-1}$, which means $t(s) = t$ and $\chi_h^t(s, a) = n_h^t(s, a) \leq 1$. Since

$$\eta_h^t(s, a) = \chi_h^t(s, a) \frac{H}{H + n_h^t(s, a)} \frac{n_h^t(s, a) - 1}{n_h^t(s, a)} + \chi_h^t(s, a) \frac{1}{n_h^t(s, a)} = 0,$$

the algorithm updates the bias-value function $V_{h,s,a}^t(s') = \tilde{V}_{h,s,a}^{t-1}(s')$.

For $s' \in \mathcal{S}_{t-1}$, we have

$$V_{h,s,a}^t(s') = \tilde{V}_{h,s,a}^{t-1}(s') = \frac{1}{|\mathcal{S}_{t-1}|} \sum_{\hat{s} \in \mathcal{S}_{t-1}} V_{h,\hat{s},a}^{t-1}(s') \geq \bar{V}_{h+1}^{t-1}(s') \geq \bar{V}_{h+1}^t(s').$$

For $s' \in \mathcal{S}_t - \mathcal{S}_{t-1}$, which means $t = t(s')$, we have

$$\begin{aligned}
V_{h,s,a}^t(s') &= \tilde{V}_{h,s,a}^{t-1}(s') \\
&= \frac{1}{|\mathcal{S}_{t-1}|^2} \sum_{\hat{s}, \hat{s}' \in \mathcal{S}_{t-1}} V_{h,\hat{s},a}^{t-1}(\hat{s}') \\
&\geq \frac{1}{|\mathcal{S}_{t-1}|} \sum_{\hat{s} \in \mathcal{S}_{t-1}} \bar{V}_{h+1}^{t-1}(\hat{s}) \\
&= \tilde{\bar{V}}_{h+1}^t(s') \geq \bar{V}_{h+1}^t(s').
\end{aligned}$$

Case 2: $s \in \mathcal{S}_{t-1}$, the algorithm updates the bias-value function

$$V_{h,s,a}^t(s') = \eta_h^t(s, a) \tilde{\bar{V}}_{h+1}^{t-1}(s') + (1 - \eta_h^t(s, a)) \tilde{V}_{h,s,a}^{t-1}(s').$$

Then we consider the cases of $s' \in \mathcal{S}_{t-1}$ and $s' \in \mathcal{S}_t - \mathcal{S}_{t-1}$. For $s' \in \mathcal{S}_{t-1}$, we have

$$\begin{aligned}
V_{h,s,a}^t(s') &= \eta_h^t(s, a) \tilde{\bar{V}}_{h+1}^{t-1}(s') + (1 - \eta_h^t(s, a)) \tilde{V}_{h,s,a}^{t-1}(s') \\
&= \eta_h^t(s, a) \bar{V}_{h+1}^{t-1}(s') + (1 - \eta_h^t(s, a)) V_{h,s,a}^{t-1}(s') \\
&\geq \eta_h^t(s, a) \bar{V}_{h+1}^{t-1}(s') + (1 - \eta_h^t(s, a)) \bar{V}_{h+1}^{t-1}(s') \\
&= \bar{V}_{h+1}^{t-1}(s') \geq \bar{V}_{h+1}^t(s').
\end{aligned}$$

For $s' \in \mathcal{S}_t - \mathcal{S}_{t-1}$, we have

$$\begin{aligned}
V_{h,s,a}^t(s') &= \eta_h^t(s, a) \tilde{\bar{V}}_{h+1}^{t-1}(s') + (1 - \eta_h^t(s, a)) \tilde{V}_{h,s,a}^{t-1}(s') \\
&= \eta_h^t(s, a) \frac{1}{|\mathcal{S}_{t-1}|} \sum_{\hat{s} \in \mathcal{S}_{t-1}} \bar{V}_{h+1}^{t-1}(\hat{s}) + (1 - \eta_h^t(s, a)) \frac{1}{|\mathcal{S}_{t-1}|} \sum_{\hat{s} \in \mathcal{S}_{t-1}} V_{h,s,a}^{t-1}(\hat{s}) \\
&\geq \eta_h^t(s, a) \frac{1}{|\mathcal{S}_{t-1}|} \sum_{\hat{s} \in \mathcal{S}_{t-1}} \bar{V}_{h+1}^{t-1}(\hat{s}) + (1 - \eta_h^t(s, a)) \frac{1}{|\mathcal{S}_{t-1}|} \sum_{\hat{s} \in \mathcal{S}_{t-1}} \bar{V}_{h+1}^{t-1}(\hat{s}) \\
&= \tilde{\bar{V}}_{h+1}^t(s') \geq \bar{V}_{h+1}^t(s').
\end{aligned}$$

B A Brief Introduction to Upper Confidence Bound Momentum Q-learning

This Appendix gives an overview of the procedure of UCBMQ. The notations $\chi_h^t(s, a)$, $n_h^t(s, a)$, $\alpha_h^t(s, a)$ and $\gamma_h^t(s, a)$ are the same in UCBMQ and UCBMQ-GA. For all $(s, a, h, s') \in \mathcal{S} \times \mathcal{A} \times [H] \times \mathcal{S}$, UCBMQ is initialized as follows:

$$\bar{V}_h^0(s) = H, V_{h,s,a}^0(s') = H, \bar{V}_{H+1}^t(s) = 0, Q_h^0(s, a) = 0.$$

In episode t , the agent follows a greedy policy based on $\bar{Q}_h^t(s, a)$ to obtain an episode of length H . Next, $Q_h^t(s, a)$, $V_{h,s,a}^t(s')$, $\bar{Q}_h^t(s, a)$ and $\bar{V}_h^t(s)$ are updated using the data collected in this episode. The algorithm terminates after T episodes.

The updates in UCBMQ described as follows. $Q_h^t(s, a)$ is updated using the momentum term $\gamma_h^t(s, a)$:

$$\begin{aligned}
Q_h^t(s, a) &= \alpha_h^t(s, a) [r_h(s, a) + p_h^t \bar{V}_{h+1}^{t-1}(s, a)] \\
&\quad + \gamma_h^t(s, a) p_h^t (\bar{V}_{h+1}^{t-1} - V_{h,s,a}^{t-1})(s, a) + (1 - \alpha_h^t(s, a)) Q_h^{t-1}(s, a),
\end{aligned}$$

where p_h^t denotes the Dirac distribution concentrated at (s_{h+1}^t) and, for any transition function p and any function $f : \mathcal{S} \rightarrow \mathbb{R}$, we define

$$pf(s, a) = \mathbb{E}_{s' \sim p(\cdot | s, a)} [f(s')].$$

$V_{h,s,a}^t(s')$ is updated according to

$$V_{h,s,a}^t(s') = \eta_h^t(s, a) \bar{V}_{h+1}^{t-1}(s') + (1 - \eta_h^t(s, a)) V_{h,s,a}^{t-1}(s'),$$

where

$$\eta_h^t(s, a) = \alpha_h^t(s, a) + \gamma_h^t(s, a).$$

$\bar{Q}_h^t(s, a)$ is used to upper-bound the optimal Q-value function, which is defined as

$$\bar{Q}_h^t(s, a) = Q_h^t(s, a) + \beta_h^t(s, a),$$

where $\beta_h^t(s, a)$ denotes the bonus term. In particular, if $n_h^t(s, a) = 0$ then $\beta_h^t(s, a) = H$; otherwise

$$\beta_h^t(s, a) = 2\sqrt{\frac{\zeta W_h^t(s, a)}{n_h^t(s, a)}} + 53H^3 \frac{\zeta \log(T)}{n_h^t(s, a)} + \sum_{k=1}^t \frac{\chi_h^k(s, a) \hat{\gamma}_h^k(s, a)}{H \log(T) n_h^t(s, a)} p_h^k(V_{h,s,a}^{k-1} - \bar{V}_{h+1}^{k-1})(s, a).$$

Here,

$$\hat{\gamma}_h^t(s, a) = H \frac{n_h^t(s, a) - 1}{n_h^t(s, a) + H}.$$

ζ is an exploration threshold and $W_h^t(s, a)$ is a proxy for the variance term defined as

$$W_h^t(s, a) = \sum_{k=1}^t \frac{\chi_h^k(s, a)}{n_h^k(s, a)} p_h^k(\bar{V}_{h+1}^{k-1} - \sum_{l=1}^t \frac{\chi_h^l(s, a)}{n_h^l(s, a)} p_h^l \bar{V}_{h+1}^{l-1})^2(s, a).$$

$\bar{V}_h^t(s)$ is used to upper-bound the optimal value function, which is updated as

$$\bar{V}_h^t(s) = \text{clip}(\max_{a \in \mathcal{A}} \bar{Q}_h^t(s, a), 0, \bar{V}_h^{t-1}(s)).$$

C Algorithm Chart of Upper Confidence Bound Momentum Q-learning with Growing Awareness

Algorithm 1 UCBMQ-GA for Training EMDA-GA.

```

1: Initialise: For all  $(s, a, h)$ , where  $s \in \mathcal{S}_0 \subset \mathcal{S}$ ,  $V_{h,s,a}^0 = \bar{V}_h^0 = H$ ,  $Q_h^0 = 0$ 
2: for  $t \in [T]$  do
3:   for  $h \in [H]$  do
4:     if  $s_h^t \in \mathcal{S}_{t-1}$  then
5:       Play  $a_h^t \in \arg \max \bar{Q}_h^{t-1}(s_h^t, a)$ 
6:     else
7:       Play  $a_h^t \in \arg \max Q_{h,avg}^{t-1}(a)$ 
8:     end if
9:     Observe  $s_{h+1}^t$ 
10:  end for
11:   $\mathcal{S}_t = \mathcal{S}_{t-1} \cup \{s_1^t, \dots, s_H^t\}$ 
12:  Expand  $Q_h^{t-1}$ ,  $\bar{V}_h^{t-1}$  and  $V_{h,s,a}^{t-1}$ :

```

$$\tilde{Q}_h^{t-1}(s, a) = \begin{cases} Q_h^{t-1}(s, a), & s \in \mathcal{S}_{t-1}, \\ Q_{h,avg}^{t-1}(a), & s \in \mathcal{S}_t - \mathcal{S}_{t-1}, \end{cases}$$

$$\tilde{\bar{V}}_h^{t-1}(s) = \begin{cases} \bar{V}_h^{t-1}(s), & s \in \mathcal{S}_{t-1}, \\ \bar{V}_{h,avg}^{t-1}, & s \in \mathcal{S}_t - \mathcal{S}_{t-1}, \end{cases}$$

$$\tilde{V}_{h,s,a}^{t-1}(s') = \begin{cases} V_{h,s,a}^{t-1}(s'), & s, s' \in \mathcal{S}_{t-1}, \\ V_{h,s,a,avg}^{t-1}, & s \in \mathcal{S}_{t-1}, s' \in \mathcal{S}_t - \mathcal{S}_{t-1}, \\ V_{h,a,avg}^{t-1}(s'), & s \in \mathcal{S}_t - \mathcal{S}_{t-1}, s' \in \mathcal{S}_{t-1}, \\ V_{h,a,avg}^{t-1}, & s, s' \in \mathcal{S}_t - \mathcal{S}_{t-1}. \end{cases}$$

```

13:  for all  $s \in \mathcal{S}_t, a, h$  do
14:     $Q_h^t(s, a) = \alpha_h^t(s, a)(r_h(s, a) + p_h^t \tilde{\bar{V}}_{h+1}^{t-1}(s, a)) + \gamma_h^t(s, a)p_h^t(\tilde{\bar{V}}_{h+1}^{t-1} - \tilde{V}_{h,s,a}^{t-1})(s, a) + (1 - \alpha_h^t(s, a))\tilde{Q}_h^{t-1}(s, a)$ 
15:     $V_{h,s,a}^t(s') = \eta_h^t(s, a)\tilde{\bar{V}}_{h+1}^{t-1}(s') + (1 - \eta_h^t(s, a))\tilde{V}_{h,s,a}^{t-1}(s')$ 
16:     $\bar{Q}_h^t(s, a) = Q_h^t(s, a) + \beta_h^t(s, a)$ 
17:     $\bar{V}_h^t(s) = \text{clip}(\max_{a \in \mathcal{A}} \bar{Q}_h^t(s, a), 0, \tilde{\bar{V}}_h^{t-1}(s))$ 
18:  end for
19: end for

```

D When a New Observation Hurts Awareness Confidence – A Reverse Bayesian Perspective

Encountering an unknown unknown clearly tells an agent her currently understanding to the environment is problematic. If the agent were a human, this surprise should hurt severely her confidence. We want to model this scenario in agent, and below we show our approach exactly has this property.

We first define a *awareness confidence* below, as a quantitative measure of a reinforcement learning agent's confidence to her understanding to the environment, which underpins her confidence to her decisions.

Definition 4 (awareness confidence of an agent). *Suppose an agent has an estimator to the hidden optimal value function V^* over a subset of the state space \mathcal{S} , defined as function $F : \mathcal{S}' \rightarrow \mathbb{R}$, where*

$S' \subset S$. The awareness confidence of this agent is defined as the average absolute difference between F and V^* over S' :

$$AC(F) = -\frac{1}{|S'|} \sum_{s \in S'} |F(s) - V^*(s)|.$$

Remark 6. The awareness confidence quantifies the potential for improving the estimator. By definition, low awareness confidence indicates that the estimator deviates substantially from the optimal value function, implying that the agent should remain not confident of its estimators. Conversely, high awareness confidence indicates that the estimator closely approximates the optimal value function, thereby justifying the agent's confidence.

We now prove that the awareness confidence drastically increases when an unknown unknown is encountered.

Theorem 5. For $t \in [T]$ and $h \in [H]$, where T is the number of learning episodes and H is the horizon, with probability at least $1 - \frac{\delta}{2}$, the awareness confidence associated with \bar{V}_h^t is lower than that associated with \tilde{V}_h^t :

$$AC(\bar{V}_h^t) < AC(\tilde{V}_h^t).$$

Proof. From Lemma 9, we have

$$\begin{aligned} AC(\tilde{V}_h^t) &= \frac{1}{|\mathcal{S}_{t+1}|} \sum_{s \in \mathcal{S}_{t+1}} |\tilde{V}_h^t(s) - V^*(s)| = \frac{1}{|\mathcal{S}_{t+1}|} \sum_{s \in \mathcal{S}_{t+1}} \tilde{V}_h^t(s) - V^*(s), \\ AC(\bar{V}_h^t) &= \frac{1}{|\mathcal{S}_t|} \sum_{s \in \mathcal{S}_t} |\bar{V}_h^t(s) - V^*(s)| = \frac{1}{|\mathcal{S}_t|} \sum_{s \in \mathcal{S}_t} \bar{V}_h^t(s) - V^*(s). \end{aligned}$$

With Homeland condition in Section 4.2, we have

$$\forall s \in \mathcal{S}_{t+1} - \mathcal{S}_t, V_h^*(s) \leq \frac{1}{|\mathcal{S}_t|} \sum_{\hat{s} \in \mathcal{S}_t} V_h^*(\hat{s}) \leq \frac{1}{|\mathcal{S}_t|} \sum_{\hat{s} \in \mathcal{S}_t} \bar{V}_h^t(\hat{s}) = \tilde{V}_h^t(s).$$

Then we get

$$\begin{aligned} \frac{1}{|\mathcal{S}_{t+1} - \mathcal{S}_t|} \sum_{\hat{s} \in \mathcal{S}_{t+1} - \mathcal{S}_t} \tilde{V}_h^t(\hat{s}) - V_h^*(\hat{s}) &\geq \frac{1}{|\mathcal{S}_{t+1} - \mathcal{S}_t|} \sum_{\hat{s} \in \mathcal{S}_{t+1} - \mathcal{S}_t} \left[\frac{1}{|\mathcal{S}_t|} \sum_{s' \in \mathcal{S}_t} \bar{V}_h^t(s') - V_h^*(s') \right] \\ &= \frac{1}{|\mathcal{S}_t|} \sum_{s \in \mathcal{S}_t} \bar{V}_h^t(s) - V^*(s). \end{aligned}$$

Thus, we can conclude $AC(\tilde{V}_h^t) \geq AC(\bar{V}_h^t)$. \square

Remark 7. This theorem provides a mathematical characterisation to our intuition mentioned in the beginning of this section. In the view of the agent, the environment becomes larger but the agent's knowledge about the environment is not able to increase at the same time. The knowledge acquired is thus diluted, leading to a degradation in average estimation accuracy, reflecting in the awareness confidence.

E Proof of Corollary 1

Proof of Corollary 1. Here we give a special case where $d\bar{V}_{h,avg}^t$ ($d < 1$) no longer upper-bounds the optimal function on a newly discovered state. Consider an EMDP with a fixed immediate reward function: $\forall (s, a) \in \mathcal{S} \times \mathcal{A}, r_h(s, a) = 1$. It is easy to get that for every state $s \in \mathcal{S}$, the optimal value function is $V_h^*(s) = H$. According to Lemma 2, we have $\bar{V}_{h,avg}^t \leq H$ for all $h \in [H]$. Therefore, for any $d < 1$, we get: $\forall s \in \mathcal{S}, d\bar{V}_{h,avg}^t < H = V_h^*(s)$. \square

F Proof of Theorem 1

Proof of Theorem 1. We prove the lemma under Assumption 1. For $s' \in \mathcal{S}_0$, we have $t(s') = 0 \leq \sqrt{T}$. For $s' \in \mathcal{S} - \mathcal{S}_0$, we have $\forall (s, a) \in \mathcal{S} \times \mathcal{A}$, $P(s'|s, a) \geq 1 - (\frac{\delta}{3})^{\frac{1}{H[\sqrt{T}]}} \triangleq \theta$. Thus, the following holds $\mathbb{P}(s' \notin \mathcal{O}_{[\sqrt{T}]} \leq (1 - \theta)^{H[\sqrt{T}]}$. From the relationship between \mathcal{O}_t and \mathcal{S}_t , we get

$$\mathbb{P}(s' \notin \mathcal{S}_{[\sqrt{T}]} \leq (1 - \theta)^{H[\sqrt{T}]} = \frac{\delta}{3}.$$

Finally, we get $t(s') \leq \lceil \sqrt{T} \rceil \leq \sqrt{T}$ with probability at least $1 - \frac{\delta}{3}$. \square

G Proof of Lemma 3 and Lemma 4

Proof of Lemma 3. We prove the computational complexity of UCBMQ-GA by analysing the cost of each time-step. The computational complexity of UCBMQ-GA per time step comprises function updates and average-value maintenance, which can be performed concurrently. At time-step h of an episode t , we only need to update the functions concerning (s_h^t, a_h^t) since for other (s, a) , $\alpha_h^t(s, a) = 0$ and $\gamma_h^t(s, a) = 0$. Thus, updating Q_h^t and \bar{Q}_h^t takes $\mathcal{O}(1)$ time. Updating the bias-value function $V_{h,s,a}^t$ takes $\mathcal{O}(1)$ time because the update is performing linear interpolation over a vector of length S . Updating \bar{V}_h^t on s_h^t involves taking the maximum over all A , which takes $\mathcal{O}(A)$ time. We can calculate average values every time the functions are updated, and the calculation takes $\mathcal{O}(1)$ time. Hence, each step of updating functions costs $\mathcal{O}(S + A)$, each step of calculating average values costs $\mathcal{O}(1)$ and the total cost is $\mathcal{O}HT(S + A)$. \square

Proof of Lemma 4. UCBMQ-GA needs space to store the functions. Q_h^t and \bar{Q}_h^t both require $\mathcal{O}(SAH)$ space. \bar{V}_h^t needs $\mathcal{O}(HS)$ space, while $V_{h,s,a}^t$ needs $\mathcal{O}(HS^2A)$ space. So the space complexity is $\mathcal{O}(HS^2A)$. \square

H Concentration Events

We firstly introduce the Bernstein-type concentration inequality.

Theorem 6 (Bernstein-type concentration inequality [9]). *Let $(Y_t)_{t \in \mathbb{N}^*}, (\omega_t)_{t \in \mathbb{N}^*}$ be two sequences of random variables adapted to a filtration $(\mathcal{F}_t)_{t \in \mathbb{N}^*}$. We assume that the weights are in the unit interval $\omega_t \in [0, 1]$ and predictable, i.e. \mathcal{F}_{t-1} measurable. We also assume that the random variables Y_t are bounded $|Y_t| \leq b$ and centered $\mathbb{E}[Y_t | \mathcal{F}_{t-1}] = 0$. Consider the following quantities*

$$S_t \triangleq \sum_{s=1}^t \omega_s Y_s, \quad V_t \triangleq \sum_{s=1}^t \omega_s^2 \mathbb{E}[Y_s^2 | \mathcal{F}_{s-1}], \quad W_t \triangleq \sum_{s=1}^t \omega_s,$$

and let $h(s) \triangleq (x+1)\log(x+1) - x$ be the Cr  mer transform of a Poisson distribution of parameter 1. For all $\delta > 0$,

$$\mathbb{P}[\exists t \geq 1, (\frac{V_t}{b^2})h(\frac{b|S_t|}{V_t + b^2}) \geq \log(\frac{1}{\delta}) + \log(4e(2t+1))] < \delta.$$

The previous inequality can be weakened to obtain a more explicit bound: if $b \geq 1$ with probability at least $1 - \delta$, for all $t \geq 1$,

$$|S_t| \leq \sqrt{2V_t \log(\frac{4e(2t+1)}{\delta})} + 3b \log(\frac{4e(2t+1)}{\delta}).$$

We define the favorable events as follows:

$$\mathcal{E}^{v_1} \triangleq \{\forall t \in [T], \forall h \in [H], \forall (s, a) \in \mathcal{S}_t \times \mathcal{A} :$$

$$|\sum_{k=t(s)}^t \chi_h^k(s, a)(p_h^k - p_h) \bar{V}_{h+1}^{\circ k-1}(s, a)| \leq \sqrt{2\zeta \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(\bar{V}_h^{\circ k-1})(s, a) + 6H\zeta}\},$$

$$\begin{aligned}\mathcal{E}^{v_2} &\triangleq \{\forall t \in [T], \forall h \in [H], \forall (s, a) \in \mathcal{S}_t \times \mathcal{A} : | \sum_{k=t(s)}^t \chi_h^k(s, a) (\hat{p}_h^k - p_h) (\bar{V}_{h+1}^{\circ k-1})^2(s, a) | \\ &\leq \sqrt{8H^2\zeta \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(\bar{V}_{h+1}^{\circ k-1})(s, a) + 12H^2\zeta}\},\end{aligned}$$

$$\begin{aligned}\mathcal{E}^{m_\omega} &\triangleq \{\forall t \in [T], \forall h \in [H], \forall (s, a) \in \mathcal{S}_t \times \mathcal{A} : \\ &| \sum_{k=t(s)}^t \chi_h^k(s, a) \hat{\gamma}_h^k(s, a) (\hat{p}_h^k - p_h) (\bar{V}_{h+1}^{\circ k-1} - \hat{V}_{h,s,a}^{\circ k-1})(s, a) | \\ &\leq \sqrt{2\zeta \sum_{k=t(s)}^t \chi_h^k(s, a) \hat{\gamma}_h^k(s, a)^2 \text{Var}_{p_h}(\bar{V}_{h+1}^{\circ k-1} - \hat{V}_{h,s,a}^{\circ k-1})(s, a) + 6H^2\zeta}\},\end{aligned}$$

$$\begin{aligned}\mathcal{E}^m &\triangleq \{\forall t \in [T], \forall h \in [H], \forall (s, a) \in \mathcal{S}_t \times \mathcal{A} : | \sum_{k=t(s)}^t \chi_h^k(s, a) (\hat{p}_h^k - p_h) (\bar{V}_{h+1}^{\circ k-1} - \hat{V}_{h,s,a}^{\circ k-1})(s, a) | \\ &\leq \sqrt{2\zeta \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(\bar{V}_{h+1}^{\circ k-1} - \hat{V}_{h,s,a}^{\circ k-1})(s, a) + 6H\zeta}\}.\end{aligned}$$

We define $\mathcal{E} = \mathcal{E}^{v_1} \cap \mathcal{E}^{v_2} \cap \mathcal{E}^{m_\omega} \cap \mathcal{E}^m$ the intersection of all the events above. \mathcal{E} holds with high probability.

Lemma 3. *For the choice*

$$\zeta = \log\left(\frac{96e(2T+1)}{\delta}\right),$$

the following holds $\mathbb{P}(\mathcal{E}) \geq 1 - \frac{\delta}{6}$.

Proof. From Lemma 2 we can get

$$\forall s, s' \in \mathcal{S}, \forall t \in [T], 0 \leq \bar{V}_{h+1}^{\circ t}(s) \leq H, 0 \leq \hat{V}_{h,s,a}^{\circ t}(s') - \bar{V}_{h+1}^{\circ t}(s') \leq H.$$

Therefore, for $s \in \mathcal{S}_k$

$$\begin{aligned} |(\hat{p}_h^k - p_h) \bar{V}_{h+1}^{\circ k-1}(s, a)| &\leq 2H, \\ |(\hat{p}_h^k - p_h) (\bar{V}_{h+1}^{\circ k-1})^2(s, a)| &\leq 4H^2, \\ |(\hat{p}_h^k - p_h) (\bar{V}_{h+1}^{\circ k-1} - \hat{V}_{h,s,a}^{\circ k-1})(s, a)| &\leq 2H. \end{aligned}$$

For \mathcal{E}^{v_2} , based on Lemma 16, we have $\text{Var}_{p_h}(\bar{V}_{h+1}^{\circ k-1})^2(s, a) \leq 2H^2 \text{Var}_{p_h}(\bar{V}_{h+1}^{\circ k-1})(s, a)$. Based on $\zeta = \log(\frac{32e(2T+1)}{\delta})$ and Theorem 6, we get

$$\mathbb{P}((\mathcal{E}^{v_1})^c) \leq \frac{\delta}{24}, \quad \mathbb{P}((\mathcal{E}^{v_2})^c) \leq \frac{\delta}{24}, \quad \mathbb{P}((\mathcal{E}^{m_\omega})^c) \leq \frac{\delta}{24}, \quad \mathbb{P}((\mathcal{E}^m)^c) \leq \frac{\delta}{24}.$$

Combine all above, we can conclude $\mathbb{P}(\mathcal{E}) \geq 1 - \frac{\delta}{6}$. \square

Lemma 4. *On the event \mathcal{E} , $\forall t \in [T], \forall h \in [H], \forall (s, a) \in \mathcal{S}_t \times \mathcal{A}$, the following holds*

$$\begin{aligned} &| \sum_{k=t(s)}^t \chi_h^k(s, a) \hat{\gamma}_h^k(s, a) (\hat{p}_h^k - p_h) (\bar{V}_{h+1}^{\circ k-1} - \hat{V}_{h,s,a}^{\circ k-1})(s, a) | \\ &\leq \frac{1}{4H \log(T)} \sum_{k=t(s)}^t \chi_h^k(s, a) \hat{\gamma}_h^k(s, a) p_h (\hat{V}_{h,s,a}^{\circ k-1} - \bar{V}_{h+1}^{\circ k-1})(s, a) + 14H^3 \log(T) \zeta, \end{aligned}$$

$$| \sum_{k=t(s)}^t \chi_h^k(s, a) (\hat{p}_h^k - p_h) (\bar{V}_{h+1}^{\circ k-1} - \hat{V}_{h,s,a}^{\circ k-1})(s, a) | \leq \frac{1}{4} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h (\hat{V}_{h,s,a}^{\circ k-1} - \bar{V}_{h+1}^{\circ k-1})(s, a) + 14H\zeta,$$

and

$$| \sum_{k=t(s)}^t \chi_h^k(s, a) (\hat{p}_h^k - p_h) (\bar{V}_{h+1}^{\circ k-1})^2(s, a) | \leq \frac{1}{4} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(\bar{V}_{h+1}^{\circ k-1})(s, a) + 44H^2\zeta.$$

Proof. For the first claim, we use the fact that $\mathcal{E}^{m\omega}$ holds and for all $s \in \mathcal{S}_t, s' \in \mathcal{S}$ we have

$$\hat{\gamma}_h^t(s, a) \leq H, 0 \leq \hat{V}_{h,s,a}^{t-1}(s') - \hat{V}_{h+1}^{t-1}(s') \leq H, \sqrt{xy} \leq x + y,$$

and get

$$\begin{aligned} & | \sum_{k=t(s)}^t \chi_h^k(s, a) \hat{\gamma}_h^k(s, a) (\hat{p}_h^k - p_h) (\bar{V}_{h+1}^{\circ k-1} - \hat{V}_{h,s,a}^{\circ k-1})(s, a) | \\ & \leq \sqrt{2\zeta H \sum_{k=t(s)}^t \chi_h^k(s, a) \hat{\gamma}_h^k(s, a) \text{Var}_{p_h}(\bar{V}_{h+1}^{\circ k-1} - \hat{V}_{h,s,a}^{\circ k-1})(s, a) + 6H^2\zeta} \\ & \leq \sqrt{2\zeta H^2 \sum_{k=t(s)}^t \chi_h^k(s, a) \hat{\gamma}_h^k(s, a) p_h (\hat{V}_{h,s,a}^{\circ k-1} - \bar{V}_{h+1}^{\circ k-1})(s, a) + 6H^2\zeta} \\ & \leq \frac{1}{4H \log(T)} \sum_{k=t(s)}^t \chi_h^k(s, a) \hat{\gamma}_h^k(s, a) p_h (\hat{V}_{h,s,a}^{\circ k-1} - \bar{V}_{h+1}^{\circ k-1})(s, a) + 14H^3 \log(T)\zeta. \end{aligned}$$

For the second claim, on the event \mathcal{E}^m , we have

$$\begin{aligned} & | \sum_{k=t(s)}^t \chi_h^k(s, a) (\hat{p}_h^k - p_h) (\bar{V}_{h+1}^{\circ k-1} - \hat{V}_{h,s,a}^{\circ k-1})(s, a) | \\ & \leq \sqrt{2\zeta \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(\bar{V}_{h+1}^{\circ k-1} - \hat{V}_{h,s,a}^{\circ k-1})(s, a) + 6H\zeta} \\ & \leq \sqrt{2\zeta H \sum_{k=t(s)}^t \chi_h^k(s, a) p_h (\bar{V}_{h+1}^{\circ k-1} - \hat{V}_{h,s,a}^{\circ k-1})(s, a) + 6H\zeta} \\ & \leq \frac{1}{4} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h (\hat{V}_{h,s,a}^{\circ k-1} - \bar{V}_{h+1}^{\circ k-1})(s, a) + 14H\zeta. \end{aligned}$$

For the third claim, we use event \mathcal{E}^{v2} :

$$\begin{aligned} & | \sum_{k=t(s)}^t \chi_h^k(s, a) (\hat{p}_h^k - p_h) (\bar{V}_{h+1}^{\circ k-1})^2(s, a) | \\ & \leq \sqrt{8H^2\zeta \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(\bar{V}_{h+1}^{\circ k-1})(s, a) + 12H^2\zeta} \\ & \leq \frac{1}{4} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(\bar{V}_{h+1}^{\circ k-1})(s, a) + 44H^2\zeta. \end{aligned}$$

This concludes the proof. \square

For $s \in \mathcal{S}$, Let $\bar{p}_h^t(s, a)$ and $\bar{p}_h^t(s)$ denote the probabilities to reach (s, a) and s , respectively, at the time-step h under the policy π^t in the algorithm. We define the following events:

$$\mathcal{G}^{Var} \triangleq \{\forall t \in [T], \forall h \in [H], \forall (s, a) \in \mathcal{S}_t \times \mathcal{A} : \\ | \sum_{k=t(s)}^t (\chi_h^k - \bar{p}_h^k)(s, a) Var_{p_h}(V_{h+1}^{\pi^t})(s, a) | \leq \sum_{k=t(s)}^t \bar{p}_h^k(s, a) Var_{p_h}(V_{h+1}^{\pi^t})(s, a) + 8H^2\zeta\},$$

$$\mathcal{G}^{v_1} \triangleq \{\forall t \in [T], \forall h \in [H], \forall (s, a) \in \mathcal{S}_t \times \mathcal{A} : | \sum_{k=t(s)}^t (\chi_h^k - \bar{p}_h^k)(s, a) p_h(\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^{\pi^t})(s, a) | \\ \leq \frac{1}{4H} \sum_{k=t(s)}^t \bar{p}_n^k(s, a) p_h |\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^{\pi^t}|(s, a) + 14H^2\zeta\},$$

$$\mathcal{G}^{v_2} \triangleq \{\forall t \in [T], \forall h \in [H], \forall s \in \mathcal{S}_t : \\ | \sum_{k=t(s)}^t (\chi_h^k - \bar{p}_h^k)(s) (\bar{V}_h^{\circ k-1} - V_h^{\pi^t})(s) | \leq \frac{1}{4H} \sum_{k=t(s)}^t \bar{p}_n^k(s) |\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^{\pi^t}|(s) + 14H^2\zeta\},$$

We define $\mathcal{G} = \mathcal{G}^{Var} \cap \mathcal{G}^{v_1} \cap \mathcal{G}^{v_2}$ the intersection of the events. This event holds with high probability.

Lemma 5. *For the choice*

$$\zeta = \log\left(\frac{96eHSA(2T+1)}{\delta}\right),$$

the following holds $\mathbb{P}(\mathcal{G}) \geq 1 - \frac{\delta}{6}$.

Proof. Based on Theorem 6, with probability at least $1 - \frac{\delta}{24}$, $\forall t \in [T], \forall h \in [H], \forall (s, a) \in \mathcal{S}_t \times \mathcal{A}$, using the facts that for $X \sim \text{Ber}(q)$ the following holds $Var(X) = q(1-q) \leq q$ and $\sqrt{xy} \leq x + y$,

$$| \sum_{k=t(s)}^t (\chi_h^k - \bar{p}_h^k)(s, a) Var_{p_h}(V_{h+1}^{\pi^t})(s, a) | \leq \sqrt{2\zeta \sum_{k=t(s)}^t \bar{p}_h^k(s, a) Var_{p_h}(V_{h+1}^{\pi^t})(s, a)^2 + 6H^2\zeta} \\ \leq \sum_{k=t(s)}^t \bar{p}_h^k(s, a) Var_{p_h}(V_{h+1}^{\pi^t})(s, a) + 8H^2\zeta.$$

So, we have

$$| \sum_{k=t(s)}^t (\chi_h^k - \bar{p}_h^k)(s, a) p_h(\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^{\pi^t})(s, a) | \\ \leq \sqrt{2\zeta \sum_{k=t(s)}^t \bar{p}_n^k(s, a) p_h |\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^{\pi^t}|(s, a)^2 + 6\zeta H^2} \\ \leq \frac{1}{4H} \sum_{k=t(s)}^t \bar{p}_n^k(s, a) p_h |\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^{\pi^t}|(s, a) + 14H^2\zeta\}.$$

Further,

$$| \sum_{k=t(s)}^t (\chi_h^k - \bar{p}_h^k)(s) (\bar{V}_h^{\circ k-1} - V_h^{\pi^t})(s) | \leq \sqrt{2\zeta \sum_{k=t(s)}^t \bar{p}_n^k(s) |\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^{\pi^t}|(s)^2 + 6\zeta H^2} \\ \leq \frac{1}{4H} \sum_{k=t(s)}^t \bar{p}_n^k(s) |\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^{\pi^t}|(s) + 14H^2\zeta.$$

Therefore, we get

$$\mathbb{P}((\mathcal{G}^{Var})^c) \leq \frac{\delta}{24}, \quad \mathbb{P}((\mathcal{G}^{V_1})^c) \leq \frac{\delta}{24}, \quad \mathbb{P}((\mathcal{G}^{V_2})^c) \leq \frac{\delta}{24}.$$

We can conclude that $\mathbb{P}(\mathcal{G}) \geq 1 - \frac{\delta}{6}$. \square

We define the event $\mathcal{D} = \mathcal{E} \cap \mathcal{G}$. Combine Lemma 3 and Lemma 5, we get:

Lemma 6. *For the choice*

$$\zeta = \log\left(\frac{96eHSA(2T+1)}{\delta}\right),$$

the following holds $\mathbb{P}(\mathcal{D}) \geq 1 - \frac{\delta}{3}$.

I Optimism

We use $r_h(s, a) + p_h \bar{V}_{h,s,a}^t(s, a)$ to estimate $Q_h^t(s, a)$. The gap between $r_h(s, a) + p_h \bar{V}_{h,s,a}^t(s, a)$ and $Q_h^t(s, a)$ proves to be controlled by $\beta_h^t(s, a)$. Finally, we can use $\bar{Q}_h^t(s, a)$ to upper-bound $Q_h^*(s, a)$ and $\bar{V}_h^t(s)$ to upper-bound $V_h^*(s, a)$.

Lemma 7. *On the event \mathcal{E} , $\forall t \in [T]$, $\forall h \in [H]$, $\forall (s, a) \in \mathcal{S} \times \mathcal{A}$, if $n_h^t(s, a) > 0$, the following holds,*

$$\begin{aligned} & |Q_h^t(s, a) - r_h(s, a) - p_h \bar{V}_{h,s,a}^t(s, a)| \\ & \leq \sqrt{\frac{2}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(\bar{V}_h^{\circ k-1})(s, a) \frac{\zeta}{n_h^t(s, a)}} + 20H^3 \frac{\zeta \log(T)}{n_h^t(s, a)} \\ & \quad + \frac{1}{4H \log(T) n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \hat{\gamma}_h^k(s, a) p_h (\bar{V}_{h,s,a}^{k-1} - \bar{V}_{h+1}^{\circ k-1})(s, a), \end{aligned}$$

where $\hat{\gamma}_h^t(s, a) = H \frac{n_h^t(s, a) - 1}{n_h^t(s, a) + H}$.

Proof. $n_h^t(s, a) > 0$ implies $s \in \mathcal{S}_t$. So, we have

$$\begin{aligned} & Q_h^t(s, a) - r_h(s, a) - p_h \bar{V}_{h,s,a}^t(s, a) \\ & = \frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) [\bar{p}_h^k \bar{V}_{h+1}^{\circ k-1}(s, a) + \hat{\gamma}_h^k(s, a) \bar{p}_h^k (\bar{V}_{h+1}^{\circ k-1} - \bar{V}_{h,s,a}^{k-1})(s, a)] - p_h \bar{V}_{h,s,a}^t(s, a) \\ & = \frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) [(\bar{p}_h^k - p_h) \bar{V}_{h+1}^{\circ k-1}(s, a) + \hat{\gamma}_h^k(s, a) (\bar{p}_h^k - p_h) (\bar{V}_{h+1}^{\circ k-1} - \bar{V}_{h,s,a}^{k-1})(s, a)] \\ & \leq \left| \frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) (\bar{p}_h^k - p_h) \bar{V}_{h+1}^{\circ k-1}(s, a) \right| \\ & \quad + \left| \frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \hat{\gamma}_h^k(s, a) (\bar{p}_h^k - p_h) (\bar{V}_{h+1}^{\circ k-1} - \bar{V}_{h,s,a}^{k-1})(s, a) \right|. \end{aligned}$$

We can upper-bound the first term of the right-hand by the definition of \mathcal{E} :

$$\begin{aligned} & \left| \frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) (\bar{p}_h^k - p_h) \bar{V}_{h+1}^{\circ k-1}(s, a) \right| \\ & \leq \sqrt{\frac{2}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(\bar{V}_h^{\circ k-1})(s, a) \frac{\zeta}{n_h^t(s, a)}} + 6H \frac{\zeta}{n_h^t(s, a)}. \end{aligned}$$

For the second term, we use Lemma 4, and get

$$\begin{aligned} & \left| \frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \hat{\gamma}_h^k(s, a) (\hat{p}_h^k - p_h) (\bar{V}_{h+1}^{\circ k-1} - \hat{V}_{h, s, a}^{k-1})(s, a) \right| \\ & \leq \frac{1}{4H \log(T) n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \hat{\gamma}_h^k(s, a) p_h (\hat{V}_{h, s, a}^{k-1} - \bar{V}_{h+1}^{\circ k-1})(s, a) + 14H^3 \frac{\zeta \log(T)}{n_h^t(s, a)}. \end{aligned}$$

This concludes the proof. \square

We design a bonus to compensate for the bias in the previous lemma.

Lemma 8. *On the event \mathcal{E} , $\forall t \in [T]$, $\forall h \in [H]$, $\forall (s, a) \in \mathcal{S} \times \mathcal{A}$, if $n_n^t(s, a) > 0$, the following holds*

$$\begin{aligned} \beta_h^t(s, a) & \geq \sqrt{\frac{2}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(\bar{V}_{h+1}^{\circ k-1})(s, a) \frac{\zeta}{n_h^t(s, a)}} + 20H^3 \frac{\zeta \log(T)}{n_h^t(s, a)} \\ & \quad + \frac{1}{4H \log(T) n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \hat{\gamma}_h^k(s, a) p_h (\hat{V}_{h, s, a}^{k-1} - \bar{V}_{h+1}^{\circ k-1})(s, a). \end{aligned}$$

Proof. The definition of bonus is: for $s \in \mathcal{S}_t$

$$\begin{aligned} \beta_h^t(s, a) & = 2\sqrt{W_h^t(s, a) \frac{\zeta}{n_h^t(s, a)}} + 53H^3 \frac{\zeta \log(T)}{n_h^t(s, a)} \\ & \quad + \frac{1}{H \log(T) n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \hat{\gamma}_h^k(s, a) \hat{p}_h^k (\hat{V}_{h, s, a}^{k-1} - \bar{V}_{h+1}^{\circ k-1})(s, a), \end{aligned}$$

where

$$W_h^t(s, a) = \frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \hat{p}_h^k (\bar{V}_{h+1}^{\circ k-1})^2(s, a) - \left(\frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \hat{p}_h^k \bar{V}_{h+1}^{\circ k-1}(s, a) \right)^2.$$

Since we do not know the true transitions p_h , the design of the bonus only needs the observed data.

Based on Lemma 4, we can control the correction term as follows,

$$\begin{aligned} & \left| \frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \hat{\gamma}_h^k(s, a) (\hat{p}_h^k - p_h) (\bar{V}_{h+1}^{\circ k-1} - \hat{V}_{h, s, a}^{k-1})(s, a) \right| \\ & \leq \frac{1}{4H \log(T) n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \hat{\gamma}_h^k(s, a) p_h (\hat{V}_{h, s, a}^{k-1} - \bar{V}_{h+1}^{\circ k-1})(s, a) + 14H^3 \frac{\log(T) \zeta}{n_h^t(s, a)}. \end{aligned} \tag{3}$$

As for the proxy of the variance term $W_h^t(s, a)$, applying Lemma 4 and the definition of \mathcal{E} and $\sqrt{xy} \leq x + y$, we get

$$\begin{aligned} & \frac{1}{n_h^t(s, a)} \left| \sum_{k=t(s)}^t \chi_h^k(s, a) (\hat{p}_h^k - p_h) (\bar{V}_{h+1}^{\circ k-1})^2(s, a) \right| \\ & \leq \frac{1}{4n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(\bar{V}_{h+1}^{\circ k-1})(s, a) + 44H^2 \frac{\zeta}{n_h^t(s, a)}, \end{aligned} \tag{4}$$

and

$$\begin{aligned}
& \left| \left(\frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \hat{p}_h^k \bar{V}_{h+1}^{\circ k-1}(s, a) \right)^2 - \left(\frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h \bar{V}_{h+1}^{\circ k-1}(s, a) \right)^2 \right| \\
& \leq \left| \frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) (\hat{p}_h^k + p_h) \bar{V}_{h+1}^{\circ k-1}(s, a) \right| \cdot \left| \frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) (\hat{p}_h^k - p_h) \bar{V}_{h+1}^{\circ k-1}(s, a) \right| \\
& \leq \frac{2H}{n_h^t(s, a)} \left| \sum_{k=t(s)}^t \chi_h^k(s, a) (\hat{p}_h^k - p_h) \bar{V}_{h+1}^{\circ k-1}(s, a) \right| \\
& \leq H \sqrt{8 \frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(\bar{V}_{h+1}^{\circ k-1})(s, a) \frac{\zeta}{n_h^t(s, a)} + 12H^2 \frac{\zeta}{n_h^t(s, a)}} \\
& \leq \frac{1}{4n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(\bar{V}_{h+1}^{\circ k-1})(s, a) + 44H^2 \frac{\zeta}{n_h^t(s, a)}. \tag{5}
\end{aligned}$$

Then, combining eq. (4), eq. (5) and Jensen's inequality, we get

$$\begin{aligned}
W_h^t(s, a) &= \frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \hat{p}_h^k (\bar{V}_{h+1}^{\circ k-1})^2(s, a) - \left(\frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \hat{p}_h^k \bar{V}_{h+1}^{\circ k-1}(s, a) \right)^2 \\
&\geq \frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h (\bar{V}_{h+1}^{\circ k-1})^2(s, a) - \left(\frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h \bar{V}_{h+1}^{\circ k-1}(s, a) \right)^2 \\
&\quad - \frac{1}{2n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(\bar{V}_{h+1}^{\circ k-1})(s, a) - 88H^2 \frac{\zeta}{n_h^t(s, a)} \\
&\geq \frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h (\bar{V}_{h+1}^{\circ k-1})^2(s, a) - \frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) (p_h \bar{V}_{h+1}^{\circ k-1}(s, a))^2 \\
&\quad - \frac{1}{2n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(\bar{V}_{h+1}^{\circ k-1})(s, a) - 88H^2 \frac{\zeta}{n_h^t(s, a)} \\
&= \frac{1}{2n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(\bar{V}_{h+1}^{\circ k-1})(s, a) - 88H^2 \frac{\zeta}{n_h^t(s, a)}. \tag{6}
\end{aligned}$$

Finally, with eq. (3) and eq. (6), we can conclude

$$\begin{aligned}
\beta_h^t(s, a) &\geq 2 \sqrt{(W_h^t(s, a) + 88H^2 \frac{\zeta}{n_h^t(s, a)}) \frac{\zeta}{n_h^t(s, a)}} + (53 - 2\sqrt{88} - 14)H^3 \frac{\zeta \log(T)}{n_h^t(s, a)} \\
&\quad + \frac{3}{4H \log(T) n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \hat{\gamma}_h^k(s, a) p_h (\hat{V}_{h, s, a}^{k-1} - \bar{V}_{h+1}^{\circ k-1})(s, a) \\
&\geq 2 \sqrt{\frac{2}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(\bar{V}_{h+1}^{\circ k-1})(s, a) \frac{\zeta}{n_h^t(s, a)}} + 20H^3 \frac{\zeta \log(T)}{n_h^t(s, a)} \\
&\quad + \frac{1}{4H \log(T) n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \hat{\gamma}_h^k(s, a) p_h (\hat{V}_{h, s, a}^{k-1} - \bar{V}_{h+1}^{\circ k-1})(s, a),
\end{aligned}$$

which also relies on the fact that $\sqrt{x+y} \leq \sqrt{x} + \sqrt{y}$.

This concludes the proof. \square

Now we are ready to prove Lemma 9 on the event \mathcal{E} and with the Homeland condition in Section 4.2.

Lemma 9. *The following holds with probability at least $1 - \delta$, $\forall t \in [T]$, $\forall (s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$ (also holds for $H + 1$ for the value function), we have*

$$\bar{Q}_h^{\circ t}(s, a) \geq Q_h^*(s, a), \quad \text{and} \quad \bar{V}_h^{\circ t}(s) \geq V_h^*(s).$$

Proof. We proceed by induction on t . The two claims hold trivially for $t = 0$ as a result of the initialisation of the algorithm. Assume the two claims hold for $k \leq t - 1$. We consider the case of t . For $h = H + 1$, we have

$$\bar{Q}_h^{\circ t}(s, a) = Q_h^*(s, a) = 0,$$

and

$$\bar{V}_h^{\circ t}(s) = V_h^*(s) = 0.$$

We assume

$$\bar{Q}_{h+1}^{\circ t}(s, a) \geq Q_{h+1}^*(s, a),$$

and

$$\bar{V}_{h+1}^{\circ t}(s) \geq V_{h+1}^*(s).$$

If $n_h^t(s, a) = 0$ and $s \in \mathcal{S}_t$, we have

$$\begin{aligned} \bar{Q}_h^{\circ t}(s, a) &= \frac{1}{|\mathcal{S}_{t(s)-1}|} \sum_{\hat{s} \in \mathcal{S}_{t(s)-1}} Q_h^{t(s)-1}(\hat{s}, a) + H \\ &\geq \frac{1}{|\mathcal{S}_{t(s)-1}|} \sum_{\hat{s} \in \mathcal{S}_{t(s)-1}} \bar{Q}_h^{t(s)-1}(\hat{s}, a) \\ &\geq \frac{1}{|\mathcal{S}_{t(s)-1}|} \sum_{\hat{s} \in \mathcal{S}_{t(s)-1}} Q_h^*(\hat{s}, a) \\ &\geq Q_h^*(s, a). \end{aligned}$$

If $s \in \mathcal{S} - \mathcal{S}_t$, we have

$$\begin{aligned} \bar{Q}_h^{\circ t}(s, a) &= \frac{1}{|\mathcal{S}_t|} \sum_{\hat{s} \in \mathcal{S}_t} Q_h^t(\hat{s}, a) + H \\ &\geq \frac{1}{|\mathcal{S}_t|} \sum_{\hat{s} \in \mathcal{S}_t} \bar{Q}_h^t(\hat{s}, a) \\ &\geq \frac{1}{|\mathcal{S}_t|} \sum_{\hat{s} \in \mathcal{S}_t} Q_h^*(\hat{s}, a) \\ &\geq Q_h^*(s, a). \end{aligned}$$

If $n_h^t(s, a) > 0$, applying Lemmas 7 and 8, and because the event \mathcal{E} holds, we get

$$\begin{aligned} \bar{Q}_h^{\circ t}(s, a) &= Q_h^t(s, a) + \beta_h^t(s, a) \\ &\geq r_h(s, a) + p_h \bar{V}_{h,s,a}^t(s, a) \\ &= r_h(s, a) + p_h \left(\sum_{k=t(s)}^t \tilde{\eta}_h^{t,k}(s, a) \bar{V}_{h+1}^{\circ k-1}(s, a) \right) \\ &\geq r_h(s, a) + p_h V_h^*(s, a) = Q_h^*(s, a). \end{aligned}$$

For $s \in \mathcal{S}_t$, note that

$$\bar{V}_h^{\circ t}(s) = \bar{V}_h^t(s) \geq \text{clip}(\max_{a \in \mathcal{A}} \bar{Q}_h^t(s, a), 0, H) \geq \max_{a \in \mathcal{A}} Q_h^*(s, a) = V_h^*(s).$$

For $s \in \mathcal{S} - \mathcal{S}_t$, we have

$$\begin{aligned}
\bar{V}_h^t(s) &= \frac{1}{|\mathcal{S}_t|} \sum_{\hat{s} \in \mathcal{S}_t} \bar{V}_h^t(\hat{s}) \\
&\geq \frac{1}{|\mathcal{S}_t|} \sum_{\hat{s} \in \mathcal{S}_t} \text{clip}(\max_{a \in \mathcal{A}} \bar{Q}_h^t(\hat{s}, a), 0, H) \\
&\geq \frac{1}{|\mathcal{S}_t|} \sum_{\hat{s} \in \mathcal{S}_t} \max_{a \in \mathcal{A}} Q_h^*(\hat{s}, a) \\
&\geq \max_{a \in \mathcal{A}} \frac{1}{|\mathcal{S}_t|} \sum_{\hat{s} \in \mathcal{S}_t} Q_h^*(\hat{s}, a) \\
&\geq \max_{a \in \mathcal{A}} Q_h^*(s, a) = V_h^*(s).
\end{aligned}$$

The proof is completed. \square

J Proof of Theorem 2

Let $\tilde{n}_h^t(s, a) = \max\{n_h^t(s, a), 1\}$. We first give two important lemmas.

Lemma 10. *On the event \mathcal{E} , with Homeland condition in Section 4.2, $\forall t \in [T]$, $\forall h \in [H]$, $\forall (s, a) \in \mathcal{S} \times \mathcal{A}$, we have,*

$$\begin{aligned}
&|\bar{Q}_h^t(s, a) - r_h(s, a) - p_h \bar{V}_{h,s,a}^t(s, a)| \\
&\leq \sqrt{\frac{4}{\tilde{n}_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(V_{h+1}^{\pi_k})(s, a) \frac{\zeta}{\tilde{n}_h^t(s, a)}} + 24H^3 \frac{\zeta \log(T)}{\tilde{n}_h^t(s, a)} \\
&\quad + \frac{2}{H \log(T) \tilde{n}_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h (\bar{V}_{h,s,a}^{k-1} - V_{h+1}^{\pi_k})(s, a).
\end{aligned}$$

Proof. we first prove $|\bar{Q}_h^t(s, a)| \leq H^2$ by induction. When $t = 0$, the claim trivially holds. Assume that the claim holds for $k \leq t - 1$. For $s \in \mathcal{S}_t$, $n_h^t(s, a) = 0$, we have

$$|Q_h^t(s, a)| = |\bar{Q}_h^t(s, a)| \leq H^2.$$

For $s \in \mathcal{S}_t$, $n_h^t(s, a) > 0$, we have

$$\begin{aligned}
&|Q_h^t(s, a)| \\
&\leq |r_h(s, a)| + \left| \frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) [\hat{p}_h^k \bar{V}_{h+1}^{\circ k-1}(s, a) + \hat{\gamma}_h^k(s, a) \hat{p}_h^k (\bar{V}_{h+1}^{\circ k-1} - \bar{V}_{h,s,a}^{k-1})(s, a)] \right| \\
&\leq 1 + \left| \frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) [H + \frac{H^2}{2}] \right| \leq H^2.
\end{aligned}$$

For $s \in \mathcal{S} - \mathcal{S}_t$, we have

$$|\bar{Q}_h^t(s, a)| = \left| \frac{1}{|\mathcal{S}_t|} \sum_{\hat{s} \in \mathcal{S}_t} Q_h^t(\hat{s}, a) \right| \leq H^2.$$

Therefore, we can conclude $\forall (s, a) \in \mathcal{S} \times \mathcal{A}$: $|\bar{Q}_h^t(s, a)| \leq H^2$.

For $n_h^t(s, a) = 0$, the bound trivially holds because

$$|\bar{Q}_h^t(s, a) - r_h(s, a) - p_h \bar{V}_{h,s,a}^t(s, a)| \leq |\bar{Q}_h^t(s, a)| + |r_h(s, a)| + |p_h \bar{V}_{h,s,a}^t(s, a)| \leq H^2 + H + 1.$$

For $n_h^t(s, a) > 0$, we have on the event \mathcal{E} with Lemma 7

$$\begin{aligned}
& |\hat{Q}_h^t(s, a) - r_h(s, a) - p_h \hat{V}_{h,s,a}^t(s, a)| \\
& \leq \sqrt{\frac{2}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(\bar{V}_h^{\circ k-1})(s, a) \frac{\zeta}{n_h^t(s, a)}} + 20H^3 \frac{\zeta \log(T)}{n_h^t(s, a)} \\
& \quad + \frac{1}{4H \log(T) n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \gamma_h^k(s, a) p_h(\hat{V}_{h,s,a}^{\circ k-1} - \bar{V}_{h+1}^{\circ k-1})(s, a).
\end{aligned}$$

Applying $\hat{V}_{h,s,a}^t \geq V_{h+1}^* \geq V_{h+1}^{\pi^k}$ and eq. (1), we can get

$$\begin{aligned}
& \frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \gamma_h^k(s, a) p_h(\hat{V}_{h,s,a}^{\circ k-1} - \bar{V}_{h+1}^{\circ k-1})(s, a) \\
& \leq \frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h(\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^{\pi^k})(s, a).
\end{aligned} \tag{7}$$

Applying $H \geq \bar{V}_h^{\circ k} \geq V_h^* \geq V_h^{\pi^{k+1}}$ and Lemma 16, we have

$$\begin{aligned}
& \frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(\bar{V}_{h+1}^{\circ k-1})(s, a) \\
& \leq \frac{2}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(V_{h+1}^{\pi^k})(s, a) + \frac{2H}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h(\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^{\pi^k})(s, a).
\end{aligned}$$

Applying $\sqrt{x+y} \leq \sqrt{x} + \sqrt{y}$ and $\sqrt{xy} \leq x+y$, we get

$$\begin{aligned}
& \sqrt{\frac{2}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(\bar{V}_{h+1}^{\circ k-1})(s, a) \frac{\zeta}{n_h^t(s, a)}} \\
& \leq \sqrt{\frac{4}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(V_{h+1}^{\pi^k})(s, a) \frac{\zeta}{n_h^t(s, a)}} \\
& \quad + \sqrt{\frac{4H}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h(\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^{\pi^k})(s, a) \frac{\zeta}{n_h^t(s, a)}} \\
& \leq \sqrt{\frac{4}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(V_{h+1}^{\pi^k})(s, a) \frac{\zeta}{n_h^t(s, a)}} \\
& \quad + \frac{1}{H \log(T) n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h(\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^{\pi^k})(s, a) \frac{\zeta}{n_h^t(s, a)} + 4H^2 \frac{\log(T) \zeta}{n_h^t(s, a)}. \tag{8}
\end{aligned}$$

Combining eq. (7) and eq. (8), we can conclude

$$\begin{aligned}
& |\hat{Q}_h^t(s, a) - r_h(s, a) - p_h \hat{V}_{h,s,a}^t(s, a)| \\
& \leq \sqrt{\frac{4}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(V_{h+1}^{\pi^k})(s, a) \frac{\zeta}{n_h^t(s, a)}} + 24H^3 \frac{\zeta \log(T)}{n_h^t(s, a)} \\
& \quad + \frac{2}{H \log(T) n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h(\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^{\pi^k})(s, a).
\end{aligned}$$

□

Then, we give an upper bound on the bonus.

Lemma 11. *On the event \mathcal{E} , under Homeland condition in Section 4.2, $\forall t \in [T]$, $\forall h \in [H]$, $\forall (s, a) \in \mathcal{S} \times \mathcal{A}$, the following holds*

$$\begin{aligned} \beta_h^t(s, a) &\leq 2\sqrt{\frac{3}{\tilde{n}_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(V_{h+1}^{\pi^k})(s, a) \frac{\zeta}{\tilde{n}_h^t(s, a)} + 106H^3 \frac{\log(T)\zeta}{\tilde{n}_h^t(s, a)}} \\ &\quad + \frac{3}{H \log(T) \tilde{n}_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h(\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^{\pi^k})(s, a). \end{aligned}$$

Proof. If $n_h^t(s, a) = 0$, the bound is trivially true because $\beta_h^t(s, a) = H$. We now consider the case $n_h^t(s, a) > 0$. Using eq. (4), eq. (5) and $H \geq \bar{V}_h^{\circ k} \geq V_h^*$, we obtain

$$\begin{aligned} W_h^t(s, a) &\leq \frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h(\bar{V}_{h+1}^{\circ k-1})^2(s, a) - \left(\frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h \bar{V}_{h+1}^{\circ k-1}(s, a) \right)^2 \\ &\quad + \frac{1}{2n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(\bar{V}_{h+1}^{\circ k-1})(s, a) + 88H^2 \frac{\zeta}{n_h^t(s, a)} \\ &= \text{Var}_{p_h}(V_{h+1}^*)(s, a) + \frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h((\bar{V}_{h+1}^{\circ k-1})^2 - (V_{h+1}^*)^2)(s, a) \\ &\quad + (p_h V_{h+1}^*(s, a))^2 - \left(\frac{1}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h \bar{V}_{h+1}^{\circ k-1}(s, a) \right)^2 \\ &\quad + \frac{1}{2n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(\bar{V}_{h+1}^{\circ k-1})(s, a) + 88H^2 \frac{\zeta}{n_h^t(s, a)} \\ &\leq \text{Var}_{p_h}(V_{h+1}^*)(s, a) + \frac{1}{2n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(\bar{V}_{h+1}^{\circ k-1})(s, a) \\ &\quad + \frac{2H}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h(\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^*)(s, a) + 88H^2 \frac{\zeta}{n_h^t(s, a)}. \end{aligned}$$

Applying $H \geq \bar{V}_h^{\circ k} \geq V_h^* \geq V_h^{\pi^{k+1}}$ and applying Lemma 16 to the terms $Var_{p_h}(V_{h+1}^*)(s, a)$ and $Var_{p_h}(\bar{V}_{h+1}^{\circ k-1})(s, a)$, we obtain

$$\begin{aligned}
W_h^t(s, a) &\leq \frac{3}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) Var_{p_h}(V_{h+1}^{\pi^k})(s, a) \\
&\quad + \frac{2H}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h(V_{h+1}^* - V_{h+1}^{\pi^k})(s, a) \\
&\quad + \frac{H}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h(\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^{\pi^k})(s, a) \\
&\quad + \frac{2H}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h(\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^*)(s, a) + 88H^2 \frac{\zeta}{n_h^t(s, a)} \\
&\leq \frac{3}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) Var_{p_h}(V_{h+1}^{\pi^k})(s, a) \\
&\quad + \frac{5H}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h(\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^{\pi^k})(s, a) + 88H^2 \frac{\zeta}{n_h^t(s, a)}. \tag{9}
\end{aligned}$$

Combine eq. (9) with $\sqrt{x+y} \leq \sqrt{x} + \sqrt{y}$ and $\sqrt{xy} \leq x+y$ we can upper-bound the variance term of the bonus

$$\begin{aligned}
2\sqrt{W_h^t(s, a) \frac{\zeta}{n_h^t(s, a)}} &\leq 2\sqrt{\frac{3}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) Var_{p_h}(V_{h+1}^{\pi^k})(s, a) \frac{\zeta}{n_h^t(s, a)}} \\
&\quad + \sqrt{\frac{20H}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h(\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^{\pi^k})(s, a) \frac{\zeta}{n_h^t(s, a)}} \\
&\quad + 2\sqrt{88H} \frac{\zeta}{n_h^t(s, a)} \\
&\leq 2\sqrt{\frac{3}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) Var_{p_h}(V_{h+1}^{\pi^k})(s, a) \frac{\zeta}{n_h^t(s, a)}} \\
&\quad + \frac{1}{H \log(T) n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h(\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^{\pi^k})(s, a) + \frac{20H^2 \zeta}{n_h^t(s, a)} \\
&\quad + 19H \frac{\zeta}{n_h^t(s, a)} \\
&\leq 2\sqrt{\frac{3}{n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) Var_{p_h}(V_{h+1}^{\pi^k})(s, a) \frac{\zeta}{n_h^t(s, a)}} \\
&\quad + \frac{1}{H \log(T) n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h(\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^{\pi^k})(s, a) \\
&\quad + \frac{39H^2 \zeta}{n_h^t(s, a)}. \tag{10}
\end{aligned}$$

Applying Lemma 4 and eq. (7), we have

$$\begin{aligned}
& \frac{1}{H \log(T) n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \gamma_h^k(s, a) \hat{p}_h^k(\hat{V}_{h,s,a}^{k-1} - \hat{V}_{h+1}^{k-1})(s, a) \\
& \leq 14H^3 \frac{\log(T)\zeta}{n_h^t(s, a)} + \frac{5}{4H \log(T) n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \gamma_h^k(s, a) p_h(\hat{V}_{h,s,a}^{k-1} - \hat{V}_{h+1}^{k-1})(s, a) \\
& \leq 14H^3 \frac{\log(T)\zeta}{n_h^t(s, a)} + \frac{5}{4H \log(T) n_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h(\hat{V}_{h+1}^{k-1} - V_{h+1}^{\pi^k})(s, a). \tag{11}
\end{aligned}$$

Combining eq. (10) and eq. (11) we obtain the result of this lemma. \square

Now we are ready to prove Theorem 2.

Proof of Theorem 2. We assume the event \mathcal{D} , Homeland condition in Section 4.2, and Assumption 1 hold.

From Lemma 6, we have that \mathcal{D} , Homeland condition in Section 4.2 and Assumption 1 hold with probability at least $1 - \delta$. Fix $(s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$ and $t \in [T]$.

Step 1: Upper-bound $(\hat{Q}_h^t - Q_h^{\pi^{t+1}})(s, a)$ Applying Lemmas 10 and 11, we get

$$\begin{aligned}
(\hat{Q}_h^t - Q_h^{\pi^{t+1}})(s, a) & \leq p_h(\hat{V}_{h,s,a}^t - V_{h+1}^{\pi^{t+1}})(s, a) + 130H^3 \frac{\log(T)\zeta}{\tilde{n}_h^t(s, a)} \\
& \quad + \sqrt{\frac{30}{\tilde{n}_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(V_{h+1}^{\pi^k})(s, a) \frac{\zeta}{\tilde{n}_h^t(s, a)}} \\
& \quad + \frac{5}{H \log(T) \tilde{n}_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h(\hat{V}_{h+1}^{k-1} - V_{h+1}^{\pi^k})(s, a). \tag{12}
\end{aligned}$$

Step 2: Upper-bound the local optimistic regret $\hat{R}_h^T(s, a)$ We define

$$\hat{R}_h^T(s, a) \triangleq \sum_{t=0}^{T-1} \chi_h^{t+1}(s, a) (\hat{Q}_h^t - Q_h^{\pi^{t+1}})(s, a) = \sum_{t=t(s)-1}^{T-1} \chi_h^{t+1}(s, a) (\hat{Q}_h^t - Q_h^{\pi^{t+1}})(s, a).$$

Based on this, eq. (12) yields the follows,

$$\begin{aligned}
\hat{R}_h^T(s, a) & \leq \sum_{t=t(s)-1}^{T-1} \chi_h^{t+1}(s, a) p_h(\hat{V}_{h,s,a}^t - V_{h+1}^{\pi^{t+1}})(s, a) + 130H^3 \sum_{t=t(s)-1}^{T-1} \chi_h^{t+1}(s, a) \frac{\log(T)\zeta}{\tilde{n}_h^t(s, a)} \\
& \quad + \sum_{t=t(s)-1}^{T-1} \chi_h^{t+1}(s, a) \sqrt{\frac{30}{\tilde{n}_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(V_{h+1}^{\pi^k})(s, a) \frac{\zeta}{\tilde{n}_h^t(s, a)}} \\
& \quad + \sum_{t=t(s)-1}^{T-1} \chi_h^{t+1}(s, a) \frac{5}{H \log(T) \tilde{n}_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h(\hat{V}_{h+1}^{k-1} - V_{h+1}^{\pi^k})(s, a).
\end{aligned}$$

For the first term, we have the decomposition

$$p_h(\hat{V}_{h,s,a}^t - V_{h+1}^{\pi^{t+1}})(s, a) = p_h(\hat{V}_{h,s,a}^t - V_{h+1}^*(s, a)) + p_h(V_{h+1}^* - V_{h+1}^{\pi^{t+1}})(s, a). \tag{13}$$

Then, applying eq. (2) and Lemma 14, we obtain

$$\begin{aligned}
& \sum_{t=t(s)-1}^{T-1} \chi_h^{t+1}(s, a) p_h(\bar{V}_{h,s,a}^t - V_{h+1}^*)(s, a) \\
&= \sum_{t=t(s)-1}^{T-1} \chi_h^{t+1}(s, a) \mathbb{I}_{\{n_h^t(s,a)=0\}} p_h(\bar{V}_{h,s,a}^t - V_{h+1}^*)(s, a) \\
&\quad + \sum_{t=t(s)-1}^{T-1} \chi_h^{t+1}(s, a) \mathbb{I}_{\{n_h^t(s,a)>0\}} \sum_{k=t(s)}^t p_h(\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^*)(s, a) \\
&\leq H + \sum_{k=t(s)}^{T-1} \left(\sum_{t=k}^{T-1} \chi_h^{t+1}(s, a) \tilde{\eta}_h^{t,k}(s, a) \right) p_h(\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^*)(s, a) \\
&\leq H + \left(1 + \frac{1}{H}\right) \sum_{t=t(s)-1}^{T-1} \chi_h^{t+1}(s, a) p_h(\bar{V}_{h+1}^{\circ t-1} - V_{h+1}^*)(s, a). \tag{14}
\end{aligned}$$

Combining eq. (13), eq. (14), and $V_{h+1}^* \geq V_{h+1}^{\pi^{k+1}}$, we get

$$\begin{aligned}
& \sum_{t=0}^{T-1} \chi_h^{t+1}(s, a) p_h(\bar{V}_{h,s,a}^t - V_{h+1}^{\pi^{t+1}})(s, a) \\
&\leq \sum_{t=0}^{T-1} \chi_h^{t+1}(s, a) p_h(V_{h+1}^* - V_{h+1}^{\pi^{t+1}})(s, a) \\
&\quad + H + \left(1 + \frac{1}{H}\right) \sum_{t=0}^{T-1} \chi_h^{t+1}(s, a) p_h(\bar{V}_{h+1}^{\circ t-1} - V_{h+1}^*)(s, a) \\
&\leq H + \left(1 + \frac{1}{H}\right) \sum_{t=0}^{T-1} \chi_h^{t+1}(s, a) p_h(\bar{V}_{h+1}^{\circ t-1} - V_{h+1}^{\pi^{t+1}})(s, a). \tag{15}
\end{aligned}$$

For the fourth term, we can proceed in a similar way but using Lemma 15 to get

$$\begin{aligned}
& \sum_{t=0}^{T-1} \frac{\chi_h^{t+1}(s, a)}{\tilde{n}_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h(\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^{\pi^k})(s, a) \\
&\leq \sum_{k=t(s)}^{T-1} \left(\sum_{t=k}^{T-1} \frac{\chi_h^{t+1}(s, a)}{\tilde{n}_h^t(s, a)} \right) \chi_h^k(s, a) p_h(\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^{\pi^k})(s, a) \\
&\leq 8 \log(T) \sum_{k=1}^{T-1} \chi_h^k(s, a) p_h(\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^{\pi^k})(s, a).
\end{aligned}$$

Thus, we have

$$\begin{aligned}
& \sum_{t=t(s)-1}^{T-1} \frac{5\chi_h^{t+1}(s, a)}{H \log(T) \tilde{n}_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) p_h(\bar{V}_{h+1}^{\circ k-1} - V_{h+1}^{\pi^k})(s, a) \\
&\leq \frac{40}{H} \sum_{t=t(s)-1}^{T-1} \chi_h^{t+1}(s, a) p_h(\bar{V}_{h+1}^{\circ t} - V_{h+1}^{\pi^{t+1}})(s, a). \tag{16}
\end{aligned}$$

For the third term, combining Cauchy-Schwarz inequality, Lemmas 15 and 13, we get

$$\begin{aligned}
& \sum_{t=t(s)-1}^{T-1} \chi_h^{t+1}(s, a) \sqrt{\frac{30}{\tilde{n}_h^t(s, a)} \sum_{k=t(s)}^t \chi_h^k(s, a) \text{Var}_{p_h}(V_{h+1}^{\pi^k})(s, a) \frac{\zeta}{\tilde{n}_h^t(s, a)}} \\
& \leq \sqrt{30 \sum_{t=t(s)-1}^{T-1} \frac{\chi_h^{t+1}(s, a)}{\tilde{n}_h^t(s, a)} \sum_{k=t(s)}^t \text{Var}_{p_h}(V_{h+1}^{\pi^k})(s, a)} \sqrt{\sum_{t=t(s)-1}^{T-1} \chi_h^{t+1}(s, a) \frac{\zeta}{\tilde{n}_h^t(s, a)}} \\
& \leq 44 \log(T) \sqrt{\zeta \sum_{t=t(s)-1}^{T-1} \chi_h^{t+1}(s, a) \text{Var}_{p_h}(V_{h+1}^{\pi^{t+1}})(s, a)}. \tag{17}
\end{aligned}$$

For the second term, we use Lemma 13 to get

$$130H^3 \sum_{t=t(s)-1}^{T-1} \chi_h^{t+1}(s, a) \frac{\log(T)\zeta}{\tilde{n}_h^t(s, a)} \leq 1040H^3 \log(T)^2 \zeta. \tag{18}$$

Finally, we combine eq. (15), eq. (16), eq. (17) and eq. (18) to conclude

$$\begin{aligned}
\hat{R}_h^T(s, a) & \leq 44 \log(T) \sqrt{\zeta \sum_{t=t(s)-1}^{T-1} \chi_h^{t+1}(s, a) \text{Var}_{p_h}(V_{h+1}^{\pi^{t+1}})(s, a) + 1041H^3 \log(T)^2 \zeta} \\
& \quad + (1 + \frac{41}{H}) \sum_{t=t(s)-1}^{T-1} \chi_h^{t+1}(s, a) p_h(\bar{V}_{h+1}^{\circ t} - V_{h+1}^{\pi^{t+1}})(s, a). \tag{19}
\end{aligned}$$

Step 3: Replace χ_h^t with \bar{p}_h^t in the upper-bound on $\hat{R}_h^T(s, a)$ Since \mathcal{G} holds, we have

$$\sqrt{\sum_{t=t(s)-1}^{T-1} \chi_h^{t+1}(s, a) \text{Var}_{p_h}(V_{h+1}^{\pi^{t+1}})(s, a)} \leq \sqrt{2 \sum_{t=t(s)-1}^{T-1} \bar{p}_h^{t+1}(s, a) \text{Var}_{p_h}(V_{h+1}^{\pi^{t+1}})(s, a)} + \sqrt{8\zeta}H$$

and

$$\begin{aligned}
& \sum_{t=t(s)-1}^{T-1} \chi_h^{t+1}(s, a) p_h(\bar{V}_{h+1}^{\circ t} - V_{h+1}^{\pi^{t+1}})(s, a) \\
& \leq (1 + \frac{1}{H}) \sum_{t=t(s)-1}^{T-1} \bar{p}_h^{t+1}(s, a) p_h(\bar{V}_{h+1}^{\circ t} - V_{h+1}^{\pi^{t+1}})(s, a) + 14H^2 \zeta.
\end{aligned}$$

Plugging the two inequalities in eq. (19) yields

$$\begin{aligned}
\hat{R}_h^T(s, a) & \leq 63 \log(T) \sqrt{\zeta \sum_{t=t(s)-1}^{T-1} \bar{p}_h^{t+1}(s, a) \text{Var}_{p_h}(V_{h+1}^{\pi^{t+1}})(s, a) + 1754H^3 \log(T)^2 \zeta} \\
& \quad + (1 + \frac{83}{H}) \sum_{t=t(s)-1}^{T-1} \bar{p}_h^{t+1}(s, a) p_h(\bar{V}_{h+1}^{\circ t} - V_{h+1}^{\pi^{t+1}})(s, a). \tag{20}
\end{aligned}$$

Step 4: Upper-bound the step h regret \hat{R}_h^T We define

$$\hat{R}_h^T \triangleq \sum_{s \in S} \sum_{t=0}^{T-1} \bar{p}_h^{t+1}(s) (\bar{V}_h^{\circ t} - V_h^{\pi^{t+1}})(s).$$

We use the definition of $\bar{V}_h^{\circ k}(s)$, Assumption 1 and the facts that \mathcal{D} holds, $\bar{Q}_h^{\circ k} \geq 0$ on \mathcal{D} and that for all $x > 1$ the following holds $\frac{1}{1-\frac{1}{4x}} \leq 1 + \frac{1}{x}$ to get

$$\begin{aligned}
& \sum_{t=0}^{T-1} \bar{p}_h^{t+1}(s)(\bar{V}_h^{\circ t} - V_h^{\pi^{t+1}})(s) \\
&= \sum_{t=0}^{(t(s)-2) \vee 0} \bar{p}_h^{t+1}(s)(\bar{V}_h^{\circ t} - V_h^{\pi^{t+1}})(s) + \sum_{t=t(s)-1}^{T-1} \bar{p}_h^{t+1}(s)(\bar{V}_h^{\circ t} - V_h^{\pi^{t+1}})(s) \\
&\leq H(t(s)-1) + \frac{1}{1-\frac{1}{4H}} \sum_{t=t(s)-1}^{T-1} \chi_h^{t+1}(s)(\bar{V}_h^{\circ t} - V_h^{\pi^{t+1}})(s) + \frac{56}{3} H^2 \zeta \\
&\leq H\sqrt{T} + (1 + \frac{1}{H}) \sum_{t=t(s)-1}^{T-1} \chi_h^{t+1}(s)(\bar{V}_h^{\circ t} - V_h^{\pi^{t+1}})(s) + 19H^2 \zeta \\
&\leq H\sqrt{T} + (1 + \frac{1}{H}) \sum_{t=t(s)-1}^{T-1} \chi_h^{t+1}(s) \pi_h^{t+1}(\bar{Q}_h^{\circ t} - Q_h^{\pi^{t+1}})(s) + 19H^2 \zeta.
\end{aligned}$$

Using Cauchy-Schwarz inequality, eq. (20) and the fact that the policy π^t is deterministic, we obtain

$$\begin{aligned}
\hat{R}_h^T &\leq HS\sqrt{T} + (1 + \frac{1}{H}) \sum_{s \in S} \sum_{t=t(s)-1}^{T-1} \chi_h^{t+1}(s) \pi_h^{t+1}(\bar{Q}_h^{\circ t} - Q_h^{\pi^{t+1}})(s) + 19H^2 S \zeta \\
&= HS\sqrt{T} + (1 + \frac{1}{H}) \sum_{(s,a) \in S \times A} \sum_{t=t(s)-1}^{T-1} \chi_h^{t+1}(s,a) (\bar{Q}_h^{\circ t} - Q_h^{\pi^{t+1}})(s,a) + 19H^2 S \zeta \\
&= HS\sqrt{T} + (1 + \frac{1}{H}) \sum_{(s,a) \in S \times A} \hat{R}_h^T(s,a) + 19H^2 S \zeta \\
&\leq HS\sqrt{T} + 126 \log(T) \sqrt{\zeta} \sum_{(s,a) \in S \times A} \sqrt{\sum_{t=t(s)-1}^{t-1} \bar{p}_h^{t+1}(s,a) \text{Var}_{p_h}(V_{h+1}^{\pi^{t+1}})(s,a)} \\
&\quad + (1 + \frac{167}{H}) \sum_{(s,a) \in S \times A} \sum_{t=t(s)-1}^{T-1} \bar{p}_h^{t+1}(s,a) p_h(\bar{V}_{h+1}^{\circ t} - V_{h+1}^{\pi^{t+1}})(s,a) \\
&\quad + 3527SAH^3 \log(T)^2 \zeta \\
&\leq HS\sqrt{T} + 3527SAH^3 \log(T)^2 \zeta + (1 + \frac{167}{H}) \hat{R}_{h+1}^T \\
&\quad + 126 \log(T) \sqrt{\zeta SA \sum_{(s,a) \in S \times A} \sum_{t=t(s)-1}^{t-1} \bar{p}_h^{t+1}(s,a) \text{Var}_{p_h}(V_{h+1}^{\pi^{t+1}})(s,a)}. \tag{21}
\end{aligned}$$

For the last inequality, we have

$$\begin{aligned}
& \sum_{(s,a) \in S \times A} \sum_{t=t(s)-1}^{T-1} \bar{p}_h^{t+1}(s,a) p_h(\bar{V}_{h+1}^{\circ t} - V_{h+1}^{\pi^{t+1}})(s,a) \\
&= \sum_{(s,a,s') \in S \times A \times S} \sum_{t=t(s)-1}^{T-1} \bar{p}_h^{t+1}(s,a) p_h(s'|s,a) (\bar{V}_{h+1}^{\circ t} - V_{h+1}^{\pi^{t+1}})(s') \\
&\leq \sum_{s' \in S} \sum_{(s,a) \in S \times A} \sum_{t=0}^{T-1} \bar{p}_h^{t+1}(s,a) p_h(s'|s,a) (\bar{V}_{h+1}^{\circ t} - V_{h+1}^{\pi^{t+1}})(s') \\
&= \sum_{s' \in S} \sum_{t=0}^{T-1} \bar{p}_{h+1}^{t+1}(s') (\bar{V}_{h+1}^{\circ t} - V_{h+1}^{\pi^{t+1}})(s') = \hat{R}_{h+1}^T.
\end{aligned}$$

Step 5: Upper-bound the regret \mathcal{R}^T Unfolding eq. (21), using Cauchy-Schwarz inequality, Lemma 12 and the fact that $\hat{R}_{H+1}^T = 0$, we get

$$\begin{aligned}
\hat{R}_1^T &\leq \sum_{h=1}^H (1 + \frac{167}{H})^{H-h} [126 \log(T) \sqrt{\zeta S A \sum_{(s,a) \in S \times A} \sum_{t=t(s)-1}^{t-1} \bar{p}_h^{t+1}(s,a) \text{Var}_{p_h}(V_{h+1}^{\pi^{t+1}})(s,a)} \\
&\quad + HS\sqrt{T} + 3527SAH^3 \log(T)^2 \zeta] \\
&\leq e^{167} 126 \log(T) \sqrt{\zeta S A H \sum_{(s,a,h) \in S \times A \times [H]} \sum_{t=0}^{t-1} \bar{p}_h^{t+1}(s,a) \text{Var}_{p_h}(V_{h+1}^{\pi^{t+1}})(s,a)} \\
&\quad + e^{167} H^2 S \sqrt{T} + 3527 e^{167} S A H^4 \log(T)^2 \zeta \\
&= e^{167} 126 \log(T) \sqrt{\zeta S A H \sum_{t=0}^{T-1} \mathbb{E}_{\pi^{t+1}}[(\sum_{h=1}^H r_h(s_h, a_h) - V_1^{\pi^{t+1}}(s_1))^2]} \\
&\quad + e^{167} H^2 S \sqrt{T} + 3527 e^{167} S A H^4 \log(T)^2 \zeta \\
&\leq \tilde{\mathcal{O}}(\sqrt{\zeta H^3 S A T}) + \tilde{\mathcal{O}}(\zeta H^4 S A) + \tilde{\mathcal{O}}(H^2 S \sqrt{T}).
\end{aligned}$$

Applying Lemma 9, we have

$$V_1^*(s_1) - V_h^{\pi^{t+1}}(s_1) \leq \bar{V}_1^{\circ t}(s_1) - V_h^{\pi^{t+1}}(s_1).$$

This allow us to conclude

$$\mathcal{R}^T \leq \hat{R}_1^T \leq \tilde{\mathcal{O}}(\sqrt{\zeta H^3 S A T}) + \tilde{\mathcal{O}}(\zeta H^4 S A) + \tilde{\mathcal{O}}(H^2 S \sqrt{T}).$$

□

K Preparation Lemmas

The proof of the following results can be seen in Ménard et al.[24]. For a deterministic policy π , we define Bellman-type equations for the variances as follows

$$\sigma Q_h^\pi(s,a) \triangleq \text{Var}_{p_h} V_{h+1}^\pi(s,a) + p_h \sigma V_{h+1}^\pi(s,a)$$

$$\sigma V_h^\pi(s) \triangleq \sigma Q_h^\pi(s, \pi(s))$$

$$\sigma V_{H+1}^\pi(s) \triangleq 0,$$

where $\text{Var}_{p_h}(f)(s,a) \triangleq \mathbb{E}_{s' \sim p_h(\cdot|s,a)}[(f(s') - p_h f(s,a))^2]$.

The definition above indicates that

$$\sigma V_1^\pi(s_1) = \sum_{h=1}^H \sum_{s,a} p_h^\pi(s,a) \text{Var}_{p_h}(V_{h+1}^\pi)(s,a).$$

We have the following lemmas.

Lemma 12. *For any deterministic policy π and for all $h \in [H]$,*

$$\mathbb{E}_\pi[(\sum_{h'=h}^H r_{h'}(s_{h'}, a_{h'}) - Q_h(s_h, a_h))^2 | (s_h, a_h) = (s, a)] = \sigma Q_h^\pi(s, a).$$

Particularly,

$$\mathbb{E}_\pi[(\sum_{h=1}^H r_h(s_h, a_h) - V_1^\pi(s_1))^2] = \sigma V_1^\pi(s_1) = \sum_{h=1}^H \sum_{s,a} p_h^\pi(s,a) \text{Var}_{p_h}(V_{h+1}^\pi)(s,a).$$

Lemma 13. *For $T \in \mathbb{N}^*$ and $(u_t)_{t \in \mathbb{N}^*}$, for a sequence where $u_t \in [0, 1]$ and $U_t \triangleq \sum_{l=1}^t u_l$, we get*

$$\sum_{t=1}^T \frac{u_{t+1}}{U_t \vee 1} \leq 4 \log(U_{T+1} + 1).$$

In particular if $T + 1 \geq 2$,

$$\sum_{t=1}^T \frac{u_{t+1}}{U_t \vee 1} \leq 8 \log(T + 1).$$

Lemma 14. *For all $(s, a) \in S \times A$ the following holds*

$$\sum_{k=l}^t \chi_h^{k+1}(s, a) \tilde{\eta}_h^{k,l}(s, a) \leq (1 + \frac{1}{H}) \chi_h^l(s, a),$$

$$\sum_{k=t(s)}^t \tilde{\eta}_h^{t,k}(s, a) = 1 \quad \text{if } n_h^t(s, a) > 0.$$

Lemma 15. *For all $(s, a) \in S \times A$ and $t \leq T - 1$ (with $T \geq 2$), the following holds*

$$\chi_h^l(s, a) \sum_{k=l}^t \frac{\chi_h^{k+1}(s, a)}{\tilde{n}_h^k(s, a)} \leq \sum_{k=0}^{T-1} \frac{\chi_h^{k+1}(s, a)}{\tilde{n}_h^k(s, a)} = \sum_{k=t(s)-1}^{T-1} \frac{\chi_h^{k+1}(s, a)}{\tilde{n}_h^k(s, a)} \leq 8 \log(T).$$

Lemma 16. *For $p, q \in \sum_S$, for $f, g : S \mapsto [0, b]$ two functions defined on S , we have that*

$$\text{Var}_p(f) \leq 2 \text{Var}_p(g) + 2bp|f - g|, \quad \text{Var}_p(f^2) \leq 2b \text{Var}_p(f),$$

where we denote the absolute operator by $|f|(s) = |f(s)|$ for all $s \in S$.