

MAKING RL WITH PREFERENCE-BASED FEEDBACK EFFICIENT VIA RANDOMIZATION

Runzhe Wu

Department of Computer Science
Cornell University
rw646@cornell.edu

Wen Sun

Department of Computer Science
Cornell University
ws455@cornell.edu

ABSTRACT

Reinforcement Learning algorithms that learn from human feedback (RLHF) need to be efficient in terms of *statistical complexity*, *computational complexity*, and *query complexity*. In this work, we consider the RLHF setting where the feedback is given in the format of preferences over pairs of trajectories. In the linear MDP model, using randomization in algorithm design, we present an algorithm that is sample efficient (i.e., has near-optimal worst-case regret bounds) and has polynomial running time (i.e., computational complexity is polynomial with respect to relevant parameters). Our algorithm further minimizes the query complexity through a novel randomized active learning procedure. In particular, our algorithm demonstrates a near-optimal tradeoff between the regret bound and the query complexity. To extend the results to more general nonlinear function approximation, we design a model-based randomized algorithm inspired by the idea of Thompson sampling. Our algorithm minimizes Bayesian regret bound and query complexity, again achieving a near-optimal tradeoff between these two quantities. Computation-wise, similar to the prior Thompson sampling algorithms under the regular RL setting, the main computation primitives of our algorithm are Bayesian supervised learning oracles which have been heavily investigated on the empirical side when applying Thompson sampling algorithms to RL benchmark problems.

1 INTRODUCTION

Reinforcement learning from human feedback (RLHF) has been widely used across various domains, including robotics (Jain et al., 2013; 2015) and natural language processing (Stiennon et al., 2020; Ouyang et al., 2022). Unlike standard RL, RLHF requires the agent to learn from feedback in the format of preferences between pairs of trajectories instead of per-step reward since assigning a dense reward function for each state is challenging in many tasks. For instance, in natural language generation, rating each generated token individually is challenging. Hence, it is more realistic to ask humans to compare two pieces of text and indicate their preference. Recent works have shown that, by integrating preference-based feedback into the training process, we can align models with human intention and enable high-quality human-machine interaction.

Despite the existing empirical applications of RLHF, its theoretical foundation remains far from satisfactory. Empirically, researchers first learn reward models from preference-based feedback and then optimize the reward models via policy gradient-based algorithms such as PPO (Schulman et al., 2017). Questions such as whether or not the learned reward model is accurate, whether PPO is sufficient for deep exploration, and how to strategically collect more feedback on the fly are often ignored. Theoretically, prior works study the regret bound for RL with preference-based feedback (Saha et al., 2023; Chen et al., 2022). Despite achieving sublinear worst-case regret, these algorithms are computationally intractable even for simplified models such as tabular Markov Decision Processes (MDPs). This means that we cannot easily leverage the algorithmic ideas in prior work to guide or improve how we perform RLHF in practice.

In addition to maximizing reward, another important metric in RLHF is the query complexity since human feedback is expensive to collect. To illustrate, we note that InstructGPT’s training data comprises a mere 30K instances of human feedback (Ouyang et al., 2022), which is significantly fewer

than the internet-scale dataset for training the GPT-3 base model. This underscores the challenge of scaling up the size of human feedback datasets. Ross et al. (2013); Laskey et al. (2016) also pointed out that extensively querying for feedback puts too much burden on human experts. Empirically, Lightman et al. (2023) observes that active learning reduces query complexity and improves the learned reward model. In theory, query complexity is mostly studied in the settings of active learning, online learning, and bandits (Cesa-Bianchi et al., 2005; Dekel et al., 2012; Agarwal, 2013; Hanneke & Yang, 2021; Zhu & Nowak, 2022; Sekhari et al., 2023a;b), but overlooked in RL.

In this work, we aim to design new RL algorithms that can learn from preference-based feedback and can be efficient in statistical complexity (i.e., regret), computational complexity, and query complexity. In particular, we strike a near-optimal balance between regret minimization and query complexity minimization. To achieve this goal, our key idea is to use *randomization* in algorithm design. We summarize our new algorithmic ideas and key contributions as follows.

1. For MDPs with linear structure (i.e., linear MDP (Jin et al., 2020)), we propose the first RL algorithm that achieves sublinear worst-case regret and computational efficiency simultaneously with preference-based feedback. Even when reduced to tabular MDPs, it is still the first to achieve a no-regret guarantee and computational efficiency. Moreover, it has an active learning procedure and attains a near-optimal tradeoff between the regret and the query complexity. Our algorithm adds *random Gaussian noises* to the learned state-action-wise reward model and the least-squares value iteration (LSVI) procedure. Using random noise instead of the UCB-style technique (Azar et al., 2017) preserves the Markovian property in the reward model and allows one to use dynamic programming to achieve computation efficiency.
2. For function approximation beyond linear, we present a model-based Thompson-sampling (TS) algorithm that forms posterior distributions over the transitions and reward models. Assuming the transition and the reward model class both have small ℓ_1 -norm eluder dimension – a structural condition introduced in Liu et al. (2022a) that is more general than the common ℓ_2 -norm eluder dimension (Russo & Van Roy, 2013), we show that our algorithm again achieves a near-optimal tradeoff between the Bayesian regret and the Bayesian query complexity. Computation-wise, similar to previous TS algorithms for regular RL (e.g., Osband et al. (2013)), the primary computation primitives are Bayesian supervised learning oracles for transition and reward learning.
3. Our query conditions for both algorithms are based on variance-style uncertainty quantification of the preference induced by the randomness of the reward model. We query for preference feedback only when the uncertainty of the preference on a pair of trajectories is large. Approximately computing the uncertainty can be easily done using i.i.d. random reward models drawn from the reward model distribution, which makes the active query procedure computationally tractable.

Overall, while our main contribution is on the theoretical side, our theoretical investigation provides several new practical insights. For instance, for regret minimization, our algorithms propose to draw a pair of trajectories with one from the latest policy and the other from an older policy instead of drawing two trajectories from the same policy (e.g., Christiano et al. (2017)), avoiding the situation of drawing two similar trajectories when the policy becomes more and more deterministic. Our theory shows that drawing two trajectories from a combination of new and older policies balances exploration and exploitation better. Another practical insight is the variance-style uncertainty measure for designing the query condition. Compared to more standard active learning procedure that relies on constructing version space and confidence intervals (Dekel et al., 2012; Puchkin & Zhivotovskiy, 2021; Zhu & Nowak, 2022; Sekhari et al., 2023a;b), our new approach comes with strong theoretical guarantees and is more computationally tractable. It is also amenable to existing implementations of Thompson sampling RL algorithms (e.g., using bootstrapping to approximate the posterior sampling (Osband et al., 2016a; 2023)).

2 COMPARISON TO PRIOR WORK

RL with preference-based feedback. Many recent works have obtained statistically efficient algorithms but are computationally inefficient even for tabular MDPs due to intractable policy search and version space construction (Chen et al., 2022; Zhan et al., 2023a;b; Saha et al., 2023). For example, Zhan et al. (2023b); Saha et al. (2023) use the idea from optimal design and rely on the computation oracle: $\arg \max_{\pi, \pi' \in \Pi} \|\mathbb{E}_{s, a \sim \pi} \phi(s, a) - \mathbb{E}_{s, a \sim \pi'} \phi(s, a)\|_A$ with some positive definite matrix A .

Here $\|x\|_A^2 := x^\top Ax$, and ϕ is some state-action feature.¹ It is unclear how to implement this oracle since standard planning approaches based on dynamic programming cannot be applied here. In addition, these methods also actively maintain a policy space by eliminating potentially sub-optimal policies. The policy class can be exponentially large even in tabular settings, so how to maintain it computationally efficiently is unclear. We provide a more detailed discussion on the challenges in achieving computational efficiency in RLHF in Appendix A.

While the work mentioned above is intractable even for tabular MDPs, there are some other works that could be computationally efficient but have weaker statistical results. For instance, very recently, Wang et al. (2023) proposed a reduction framework that can be computationally efficient (depending on the base algorithm used in the reduction). However, their algorithms have PAC bounds while we focus on regret minimization. Moreover, we achieve a near-optimal balance between regret and query complexity. Novoseller et al. (2020) proposed a posterior sampling algorithm for tabular MDP but their analysis is asymptotic (i.e., they do not address exploration, exploitation, and query complexity tradeoff). Xu et al. (2020) proposed efficient algorithms that do reward-free exploration. However, it is limited to tabular MDPs and PAC bounds.

In contrast to the above works, our algorithms aim to achieve efficiency in statistical, computational, and query complexities simultaneously. Our algorithms leverage *randomization* to balance exploration, exploitation, and feedback query. Randomization allows us to avoid non-standard computational oracles and only use standard Dynamic Programming (DP) based oracles (e.g., value iteration), which makes our algorithm computationally more tractable. Prior works that simultaneously achieve efficiency in all three aspects are often restricted in the bandit and imitation learning settings where the exploration problem is much easier (Sekhari et al., 2023a).

RL via randomization. There are two lines of work that study RL via randomization. The first injects random noise into the learning object to encourage exploration. A typical example is the randomized least-squares value iteration (RLSVI) (Osband et al., 2016b), which adds Gaussian noise into the least-squares estimation and achieves near-optimal worst-case regret (Zanette et al., 2020; Agrawal et al., 2021) for linear MDPs. The other line of work is Bayesian RL and uses Thompson sampling (TS) (Osband et al., 2013; Osband & Van Roy, 2014b;a; Gopalan & Mannor, 2015; Agrawal & Jia, 2017; Efroni et al., 2021; Zhong et al., 2022; Agarwal & Zhang, 2022). They achieve provable Bayesian regret upper bound by maintaining posterior distributions over models.

Active learning. Numerous studies have studied active learning across various settings (Cesa-Bianchi et al., 2005; Dekel et al., 2012; Agarwal, 2013; Hanneke & Yang, 2015; 2021; Zhu & Nowak, 2022; Sekhari et al., 2023b;a). However, most of them focus on the bandits and online learning settings, and their active learning procedures are usually computationally intractable due to computing version spaces or upper and lower confidence bounds. In contrast, we design a variance-style uncertainty quantification for our query condition, which can be easily estimated by random samples of reward model. This makes our active learning procedure more computationally tractable.

3 PRELIMINARY

Notations. For two real numbers a and b , we denote $[a, b] := \{x : a \leq x \leq b\}$. For an integer N , we denote $[N] := \{1, 2, \dots, N\}$. For a set \mathcal{S} , we denote $\Delta(\mathcal{S})$ as the set of distributions over \mathcal{S} . Let $d_{\text{TV}}(\cdot, \cdot)$ denote the total variation distance.

We consider a finite-horizon Markov decision process (MDP), which is a tuple $M(\mathcal{S}, \mathcal{A}, r^*, P^*, H)$ where \mathcal{S} is the state space, \mathcal{A} is the action space, $P^* : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ is the transition kernel, $r^* : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is the reward function, and H is the length of the episode. The interaction proceeds for T rounds. At each round $t \in [T]$, we need to select two policies π_t^0 and π_t^1 and execute them separately, which generates two trajectories τ_t^0 and τ_t^1 where $\tau_t^i = (s_{t,1}^i, a_{t,1}^i, \dots, s_{t,H}^i, a_{t,H}^i)$ for $i \in \{0, 1\}$. For the ease of notation, we assume a fixed initial state s_1 . Then, we need to decide whether to make a query for the preference between τ_t^0 and τ_t^1 . If making a query, we obtain a preference feedback $o_t \in \{0, 1\}$ that is sampled from the Bernoulli distribution:

$$\Pr(o_t = 1 \mid \tau_t^1, \tau_t^0, r^*) = \Pr(\tau_t^1 \text{ is preferred to } \tau_t^0 \mid r^*) = \Phi(r^*(\tau_t^1) - r^*(\tau_t^0))$$

¹These works typically assume trajectory-wise feature $\phi(\tau)$ for a trajectory τ . However, even when specified to state-action-wise features, these algorithms are still computationally intractable, even in tabular MDPs.

where $r^*(\tau_t^i) = \sum_{h=1}^H r^*(s_{t,h}^i, a_{t,h}^i)$ for $i \in \{0, 1\}$ is the trajectory reward, and $\Phi : \mathbb{R} \rightarrow [0, 1]$ is a monotonically increasing link function. We note that, by symmetry, we have $\Phi(r^*(\tau_t^0) - r^*(\tau_t^1)) + \Phi(r^*(\tau_t^1) - r^*(\tau_t^0)) = 1$. If not making a query, we receive no feedback.

This feedback model is weaker than the standard RL where the per-step reward signal is revealed. We impose the following assumption on the link function Φ , which has appeared in many existing works of RLHF (Saha et al., 2023; Zhu et al., 2023; Zhan et al., 2023a).

Assumption 3.1. *We assume Φ is differentiable and there exists constants $\kappa, \bar{\kappa} > 0$ such that $\kappa^{-1} \leq \Phi'(x) \leq \bar{\kappa}^{-1}$ for any $x \in [-H, H]$.*

The constants κ and $\bar{\kappa}$ characterize the non-linearity of Φ and determine the difficulty of estimating the reward from preference feedback. It is noteworthy that, in the theoretical results of our algorithms, the bounds depend polynomially on κ but logarithmically on $\bar{\kappa}$. Some typical examples of the link functions are provided below.

Example 3.2 (Link functions). *It is common to have $\Phi(x) = 1/(1 + \exp(-x))$, which recovers the Bradley-Terry-Luce (BTL) model (Bradley & Terry, 1952), and we have $\kappa = 2 + \exp(-H) + \exp(H)$ and $\bar{\kappa} = 4$. Additionally, if the trajectory-wise reward is scaled within the interval of $[0, 1]$, then the difference in reward will be within the range of $[-1, 1]$. In this case, another common choice of the link function is $\Phi(x) = (x + 1)/2$, which results in $\kappa = \bar{\kappa} = 2$.*

The goal is to minimize the worst-case regret and the query complexity simultaneously:

$$\text{Regret}_T := \sum_{t=1}^T \left(2V^*(s_1) - V^{\pi_t^0}(s_1) - V^{\pi_t^1}(s_1) \right), \quad \text{Queries}_T := \sum_{t=1}^T Z_t.$$

Here $V^\pi(s_1) := \mathbb{E}_\pi[\sum_{h=1}^H r^*(s_h, a_h)]$ denotes the state-value function of policy π , and we define $V^*(s_1) := V^{\pi^*}(s_1)$ where π^* is the optimal policy that maximizes the state-value function. The variable $Z_t \in \{0, 1\}$ indicates whether a query is made at round t . Note that the regret looks at the sum of the performance gaps between two pairs of policies: (π^*, π_t^0) and (π^*, π_t^1) . This is standard in dueling bandits (Yue & Joachims, 2011; Yue et al., 2012; Dudík et al., 2015; Bengs et al., 2022; Wu et al., 2023b) and RL with preference-based feedback (Saha et al., 2023; Chen et al., 2022).

Bayesian RL. We also consider Bayesian RL in this work when learning with general function approximation. In the Bayesian setting, P^* and r^* are sampled from some known prior distributions ρ_P and ρ_r . The goal is to minimize the Bayesian regret and the Bayesian query complexity:

$$\text{BayesRegret}_T := \mathbb{E} \left[\sum_{t=1}^T \left(2V^*(s_1) - V^{\pi_t^0}(s_1) - V^{\pi_t^1}(s_1) \right) \right], \quad \text{BayesQueries}_T := \mathbb{E} \left[\sum_{t=1}^T Z_t \right].$$

Here the expectation is taken with respect to the prior distribution over P^* and r^* . We will use Bayesian supervised learning oracles to compute posteriors over the transition and reward model.

4 A MODEL-FREE RANDOMIZED ALGORITHM FOR LINEAR MDPs

In this section, we present a model-free algorithm for linear MDPs which is defined as follows.

Assumption 4.1 (Linear MDP (Jin et al., 2020)). *We assume a known feature map $\phi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$, an unknown (signed) measure $\mu : \mathcal{S} \rightarrow \mathbb{R}^d$, and an unknown vector θ_r^* such that for any $(s, a) \in \mathcal{S} \times \mathcal{A}$, we have $P^*(s' | s, a) = \phi^\top(s, a) \cdot \mu(s')$ and $r^*(s, a) = \phi^\top(s, a) \cdot \theta_r^*$. We assume $\|\phi(s, a)\|_2 \leq 1$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$, $\int_{\mathcal{S}} \|\mu(s)\|_2 ds \leq \sqrt{d}$, and $\|\theta_r^*\|_2 \leq B$ for some $B > 0$. For a trajectory $\tau = (s_1, a_1, \dots, s_H, a_H)$, we define $\phi(\tau) = \sum_{h=1}^H \phi(s_h, a_h)$ and assume $\|\phi(\tau)\|_2 \leq 1$.*

Linear MDPs can capture tabular MDPs by setting $d = |\mathcal{S}| |\mathcal{A}|$ and $\phi(s, a)$ to be the one-hot encoding of (s, a) . In this case, we have $\|\phi(\tau)\|_2 \leq H$. However, we can scale it down to get $\|\phi(\tau)\|_2 \leq 1$ at the expense of scaling B up by H . We define $\Theta_B = \{\theta \in \mathbb{R}^d : \|\theta\|_2 \leq B\}$, which contains θ_r^* .

4.1 ALGORITHM

The algorithm, called PR-LSVI, is presented in Algorithm 1. At the beginning of episode k , it first computes the maximum likelihood estimate $\hat{\theta}_{r,t}$ (Line 3). Computation-wise, while the likelihood

objective is not guaranteed to be concave due to the generality of Φ , efficient algorithms exist in certain common scenarios. For example, if $\Phi(x) = 1/(1 + \exp(-x))$, it recovers the BTL model (Example 3.2). In this case, the MLE objective is concave in θ and thus can be solved in polynomial running time. Moreover, we emphasize that the reward is learned under trajectory-wise features, which is different from the standard RL setting where it is learned under state-action features.

Given the MLE $\hat{\theta}_{r,t}$, it next samples $\bar{\theta}_{r,t}$ from a Gaussian distribution centered at $\hat{\theta}_{r,t}$ (Line 4). Note that the covariance matrix Σ_{t-1}^{-1} uses trajectory-wise features (Line 16) which allows the randomized Gaussian vector to capture trajectory-wise uncertainty of the learned reward. The noise aims to encourage exploration. Then, it computes the least-squares estimate of the state-action value function $\hat{\theta}_{P,t,h}$ for each $h \in [H]$ and samples $\bar{\theta}_{P,t,h}$ from a Gaussian distribution centered at $\hat{\theta}_{P,t,h}$ (Lines 7-8). Similar to the reward model, the noise is added to the state-value function to encourage exploration. We then define the value function $\bar{Q}_{t,h}$ and $\bar{V}_{t,h}$ as

$$\bar{Q}_{t,h}(s, a) := \phi(s, a)^\top \bar{\theta}_{r,t} + \omega_{t,h}(s, a), \quad \bar{V}_{t,h}(s) := \max_a \bar{Q}_{t,h}(s, a) \quad (1)$$

and the function $\omega : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is defined as

$$\omega_{t,h}(s, a) = \begin{cases} \phi(s, a)^\top \bar{\theta}_{P,t,h} & \text{if } \|\phi(s, a)\|_{\Sigma_{t-1,h}^{-1}} \leq \alpha_L \\ \rho(s, a) \left(\phi(s, a)^\top \bar{\theta}_{P,t,h} \right) + (1 - \rho(s, a))(H - h) & \text{if } \alpha_L < \|\phi(s, a)\|_{\Sigma_{t-1,h}^{-1}} \leq \alpha_U \\ H - h & \text{if } \|\phi(s, a)\|_{\Sigma_{t-1,h}^{-1}} > \alpha_U \end{cases}$$

where $\rho(s, a) = (\alpha_U - \|\phi(s, a)\|_{\Sigma_{t-1,h}^{-1}}) / (\alpha_U - \alpha_L)$ interpolates between the two regimes to ensure continuity. This truncation trick is from Zanette et al. (2020) and is crucial. It controls the abnormally high value estimates. Specifically, when $\|\phi(s, a)\|_{\Sigma_{t-1,h}^{-1}}$ is large, the uncertainty in the direction of $\phi(s, a)$ is large, which makes the estimate $\phi(s, a)^\top \bar{\theta}_{P,t,h}$ abnormally large. In this case, we have to truncate it to $H - h$. Moreover, we note that the usual ‘‘value clipping’’ trick (i.e., simply constraining the value function within the range of $[0, H - h + 1]$ by clipping) cannot easily work here since it introduces bias to the random walk analysis, also pointed out by Zanette et al. (2020).

Then, the algorithm computes the greedy policy π_t^0 with respect to $\bar{Q}_{t,h}$. The comparator policy π_t^1 is simply set to the greedy policy from the previous episode, π_{t-1}^0 . In other words, we are comparing the two most recent greedy policies. This is different from previous work, which compares the current greedy policy with a fixed comparator (Wang et al., 2023). Analytically, for our algorithm, the cumulative regret incurred by π_t^1 for all $t \in [T]$ is equivalent to that incurred by π_t^0 for all $t \in [T]$. Hence, it suffices to compute the regret for one of them and multiply it by two to get the total regret.

Given the trajectories τ_t^0 and τ_t^1 generated by π_t^0 and π_t^1 , we compute the *expected absolute reward difference* between the trajectories under the same noisy distribution of the reward parameter:

$$\mathbb{E}_{\theta_0, \theta_1 \sim \mathcal{N}(\bar{\theta}_{r,t}, \sigma_r^2 \Sigma_{t-1}^{-1})} \left[\left| (\phi(\tau_t^0) - \phi(\tau_t^1))^\top (\theta_0 - \theta_1) \right| \right]. \quad (2)$$

This represents the uncertainty of the preference between the two trajectories, and we make a query only when it is larger than a threshold ϵ (Line 13). Intuitively, we only make a query on two trajectories when we are uncertain about the preference (e.g., the expected disagreement between two randomly sampled reward models is large). Computationally, we can estimate this expectation by drawing polynomially many reward models from the distribution $\mathcal{N}(\bar{\theta}_{r,t}, \sigma_r^2 \Sigma_{t-1}^{-1})$ and computing the empirical average. The deviation of the empirical average to the true mean can be easily bounded by standard concentration inequalities. We simply use expectation here for analytical simplicity. If the query condition is triggered, we make a query for feedback on τ_t^0, τ_t^1 , and update the trajectory-wise feature covariance matrix accordingly.

4.2 ANALYSIS

The theoretical results of Algorithm 1 are stated in Theorem 4.2. The detailed assignment of hyperparameters can be found in Table 1, and the proof is provided in Appendix B.

Theorem 4.2. *Define $\gamma = \sqrt{\kappa + B^2}$, which characterizes the difficulty of estimating the reward model. Set $\sigma_r = \tilde{\Theta}(\gamma\sqrt{d})$, $\sigma_P = \tilde{\Theta}(H^{3/2}d^2\gamma)$, $\alpha_U = (d^{5/2}H^{3/2}\gamma)^{-1}$, $\alpha_L = \alpha_U/2$, and $\lambda = 1$.*

Algorithm 1 Preference-based and Randomized Least-Squares Value Iteration (PR-LSVI)

Require: STD σ_r, σ_p , threshold ϵ , value cutoff parameters α_L, α_U , and regularization parameter λ .

- 1: Let π_0^0 be an arbitrary policy, $\Sigma_0 \leftarrow \lambda I, \Sigma_{0,h} \leftarrow \lambda I (\forall h \in [H])$.
- 2: **for** $t = 1, \dots, T$ **do**
- 3: $\hat{\theta}_{r,t} \leftarrow \arg \max_{\theta \in \Theta_B} \sum_{s=1}^{t-1} Z_s \ln(o_s \Phi((\phi(\tau_s^1) - \phi(\tau_s^0))^\top \theta) + (1 - o_s) \Phi((\phi(\tau_s^0) - \phi(\tau_s^1))^\top \theta))$
- 4: $\bar{\theta}_{r,t} \sim \mathcal{N}(\hat{\theta}_{r,t}, \sigma_r^2 \Sigma_{t-1}^{-1})$
- 5: $\hat{\theta}_{P,t,H} \leftarrow 0, \bar{\theta}_{P,t,H} \leftarrow 0$
- 6: **for** $h = H - 1, \dots, 1$ **do**
- 7: $\hat{\theta}_{P,t,h} \leftarrow \Sigma_{t-1,h}^{-1} (\sum_{i=1}^{t-1} \phi(s_{i,h}^0, a_{i,h}^0) \bar{V}_{t,h+1}(s_{i,h+1}^0))$
- 8: $\bar{\theta}_{P,t,h} \sim \mathcal{N}(\hat{\theta}_{P,t,h}, \sigma_p^2 \Sigma_{t-1,h}^{-1})$
- 9: Define $\bar{Q}_{t,h}$ and $\bar{V}_{t,h}$ as in (1).
- 10: **end for**
- 11: Set $\pi_t^0 \leftarrow \{\pi_{t,h}^0 : \pi_{t,h}^0(s) = \arg \max_a \bar{Q}_{t,h}(s, a), \forall s \in \mathcal{S}, h \in [H]\}$ and $\pi_t^1 \leftarrow \pi_{t-1}^0$.
- 12: Sample $\tau_t^0 \sim \pi_t^0$ and $\tau_t^1 \sim \pi_t^1$.
- 13: $Z_t \leftarrow \mathbb{1}\{\mathbb{E}_{\theta_0, \theta_1 \sim \mathcal{N}(\bar{\theta}_{r,t}, \sigma_r^2 \Sigma_{t-1}^{-1})} [|(\phi(\tau_t^0) - \phi(\tau_t^1))^\top (\theta_0 - \theta_1) |] > \epsilon\}$
- 14: **if** $Z_t = 1$ **then**
- 15: Query preference feedback o_t on $\{\tau_t^0, \tau_t^1\}$
- 16: $\Sigma_t \leftarrow \Sigma_{t-1} + (\phi(\tau_t^0) - \phi(\tau_t^1))(\phi(\tau_t^0) - \phi(\tau_t^1))^\top$
- 17: **else**
- 18: $\Sigma_t \leftarrow \Sigma_{t-1}$
- 19: **end if**
- 20: $\Sigma_{t,h} \leftarrow \Sigma_{t-1,h} + \phi(s_{t,h}^0, a_{t,h}^0) \phi^\top(s_{t,h}^0, a_{t,h}^0) (\forall h \in [H])$.
- 21: **end for**

Then, PR-LSVI (Algorithm 1) guarantees the following with probability at least $1 - \delta$:

$$\text{Regret}_T = \tilde{O}\left(\epsilon T d^{1/2} + \sqrt{T} \cdot d^3 H^{5/2} \gamma + d^{17/2} H^{11/2} \gamma^3\right), \quad \text{Queries}_T = \tilde{O}(d^4 \gamma^4 / \epsilon^2).$$

To further study the balance between the regret and the query complexity, we let $\epsilon = T^{-\beta}$ for some $\beta \leq 1/2$. Then, the upper bounds in Theorem 4.2 can be rewritten as

$$\text{Regret}_T = \tilde{O}(T^{1-\beta}), \quad \text{Queries}_T = \tilde{O}(T^{2\beta})$$

where we only focus on the dependence on T and omit any other factors for simplicity. We see that there is a tradeoff in T between the regret and the query complexity — the smaller regret we want, the more queries we need to make. For example, when $\beta = 0$, the regret is $\tilde{O}(T)$, and the query complexity is $\tilde{O}(1)$, meaning that we will incur linear regret if we don't make any query. If we increase β to $1/2$, the regret decreases to $\tilde{O}(\sqrt{T})$ while the query complexity increases to $\tilde{O}(T)$, meaning that the regret bound is optimal in T but we make queries every episode.

We emphasize that this tradeoff in T is *optimal*, as evidenced by a lower bound result established by Sekhari et al. (2023a). Their lower bound was originally proposed for contextual dueling bandits, which is a special case of our setting. Their results are stated below.

Theorem 4.3. (Sekhari et al., 2023a, Theorem 5) *The following two claims hold: (1) For any algorithm, there exists an instance that leads to $\text{Regret}_T = \Omega(\sqrt{T})$; (2) For any algorithm achieving an expected regret upper bound in the form of $\mathbb{E}[\text{Regret}_T] = O(T^{1-\beta})$ for some $\beta > 0$, there exists an instance that results in $\mathbb{E}[\text{Queries}_T] = \Omega(T^{2\beta})$.*

However, the dependence on other parameters (e.g., d and H) can be loose, and further improvement may be possible. We leave further investigation of these factors as future work.

Although injecting random noise is inspired by RLSVI (Zanette et al., 2020), we highlight five key differences between ours and theirs: (1) Since the feedback is trajectory-wise, we need to design random noise that preserves the state-action-wise format (so that it can be used in DP) but captures the trajectory-wise uncertainty. We do this by maintaining Σ_t , which uses trajectory-wise feature differences; (2) Since the preference feedback is generated from some probabilistic model, we learn the

reward model via MLE and use MLE generalization bound (Geer, 2000) to capture the uncertainty in learning. This allows us to use a more general link function Φ ; (3) We design a new regret decomposition technique to accommodate preference-based feedback. Particularly, we decompose regret to characterizes the *reward difference* between π_t^0 and π_t^1 : $\text{Regret}_T \lesssim \sum_{t=1}^T (\bar{V}_t - \tilde{V}_t) - (V^{\pi_t^0} - V^{\pi_t^1})$ where \bar{V}_t is an estimate of $V^{\pi_t^0}$, and $\tilde{V}_t := \mathbb{E}_{\tau \sim \pi_t^1} [\sum_{h=1}^H \phi(s_h, a_h)^\top \theta_{r,t}]$ is an estimate of $V^{\pi_t^1}$ under the real transition and the learned reward model. This is different from standard RL (Zanette et al., 2020), and is necessary since we cannot guarantee the learned reward model will be accurate in a state-action-wise manner under the preference-based feedback. (4) Our algorithms have a new randomized active learning procedure for reducing the number of queries, and our analysis achieves a near-optimal tradeoff between regret and query complexity; (5) In every round t , we propose to draw a pair of trajectories where one is from the current greedy policy π_t^0 and the other is from the greedy policy of the previous round, π_{t-1}^0 . This ensures π_t^1 is conditionally independent of the Gaussian noises at round t , which is the key to optimism (with a constant probability).

Running time. To assess the time complexity of Algorithm 1, assuming finite number of actions², all steps can be computed in polynomial running time (i.e., polynomial in d, H, A) except the MLE of the reward model (Line 3), which depends on the link function Φ . For the popular BTL model where $\Phi(x) = 1/(1 + \exp(-x))$, the MLE objective is concave with respect to θ and θ belongs to a convex set Θ_B . In this case, we can use any convex programming algorithms for the MLE procedure (e.g., projected gradient ascent).

5 A MODEL-BASED THOMPSON SAMPLING ALGORITHM

In this section, we aim to extend to nonlinear function approximation. We do so in a model-based framework with Thompson sampling (TS). The motivation is that TS is often considered a computationally more tractable alternative to UCB-style algorithms.

5.1 ALGORITHM

The algorithm, called PbTS, is presented in Algorithm 2. At the beginning of episode k , it computes the reward model posterior $\rho_{r,t}$ and the transition model posterior $\rho_{P,t}$ (Line 3). Then, it samples P_t and r_t from the posteriors and computes the optimal policy π_t^0 assuming the true reward function is r_t and the true model is P_t (Line 5). Here we denote $V_{r,P}^\pi$ as the state-value function of π under reward function r and model P . Note that this oracle is a standard planning oracle. The comparator policy π_t^1 is simply set to be the policy from the previous episode, π_{t-1}^0 , as we did in Algorithm 1. The two policies then generate respective trajectories τ_t^0 and τ_t^1 . To decide whether we should make a query, we compute the uncertainty quantity under the posterior distribution of the reward: $\mathbb{E}_{r,r' \sim \rho_{r,t}} [|r(\tau_t^0) - r(\tau_t^1) - (r'(\tau_t^0) - r'(\tau_t^1))|]$, which is analogous to (2) in Algorithm 1. We make a query only when it is larger than a threshold ϵ . Similar to Algorithm 1, we can approximate this expectation by sampling polynomial many pairs of r and r' and then compute the empirical average.

5.2 ANALYSIS

The theoretical results of Algorithm 2 should rely on the complexity of the reward and the transition model. In our analysis, we employ two complexity measures — eluder dimension and bracketing number. We start by introducing a generic notion of ℓ_p -eluder dimension (Russo & Van Roy, 2013).

Definition 5.1 (ℓ_p -norm ϵ -dependence). *Let $p > 0$. Let \mathcal{X} and \mathcal{Y} be two sets and $d(\cdot, \cdot)$ be a distance function on \mathcal{Y} . Let $\mathcal{F} \subseteq \mathcal{X} \rightarrow \mathcal{Y}$ be a function class. We say an element $x \in \mathcal{X}$ is ℓ_p -norm ϵ -dependent on $\{x_1, x_2, \dots, x_n\} \subseteq \mathcal{X}$ with respect to \mathcal{F} and d if any pair of functions $f, f' \in \mathcal{F}$ satisfying $\sum_{i=1}^n d^p(f(x_i), f'(x_i)) \leq \epsilon^p$ also satisfies $d(f(x), f'(x)) \leq \epsilon$. Otherwise, we say x is ℓ_p -norm ϵ -independent of $\{x_1, x_2, \dots, x_n\}$.*

Definition 5.2 (ℓ_p -norm eluder dimension). *The ℓ_p -norm ϵ -eluder dimension of function class $\mathcal{F} \subseteq \mathcal{X} \rightarrow \mathcal{Y}$, denoted by $\text{dim}_p(\mathcal{F}, \epsilon, d)$, is the length of the longest sequence of elements in \mathcal{X} satisfying that there exists $\epsilon' \geq \epsilon$ such that every element in the sequence is ℓ_p -norm ϵ' -independent of its predecessors.*

²This is to ensure that $\arg \max_a Q(s, a)$ can be computed efficiently.

Algorithm 2 Preference-based Thompson Sampling (PbTS)**Require:** priors ρ_P and ρ_r , threshold ϵ .

- 1: Let π_0^0 be an arbitrary policy.
- 2: **for** $t = 1, \dots, T$ **do**
- 3: Compute posteriors:

$$\rho_{P,t}(P) \propto \rho_P(P) \prod_{i=1}^{t-1} \prod_{h=1}^H P(s_{i,h+1}^0 | s_{i,h}^0, a_{i,h}^0),$$

$$\rho_{r,t}(r) \propto \rho_r(r) \prod_{i=1}^{t-1} \left(o_i \Phi(r(\tau_i^1) - r(\tau_i^0)) + (1 - o_i) \Phi(r(\tau_i^0) - r(\tau_i^1)) \right)^{Z_i}.$$

- 4: Sample $P_t \sim \rho_{P,t}$ and $r_t \sim \rho_{r,t}$.
- 5: Compute $\pi_t^0 \leftarrow \arg \max_{\pi} V_{r_t, P_t}^{\pi}(s_1)$ and $\pi_t^1 \leftarrow \pi_{t-1}^0$.
- 6: Sample $\tau_t^0 \sim \pi_t^0$ and $\tau_t^1 \sim \pi_t^1$.
- 7: $Z_t \leftarrow \mathbb{1}\{\mathbb{E}_{r, r' \sim \rho_{r,t}}[|r(\tau_t^0) - r(\tau_t^1) - (r'(\tau_t^0) - r'(\tau_t^1))|] > \epsilon\}$
- 8: **if** $Z_t = 1$ **then**
- 9: Query preference feedback o_t on $\{\tau_t^0, \tau_t^1\}$
- 10: **end if**
- 11: **end for**

The eluder dimension is non-decreasing in p , i.e., $\dim_p(\mathcal{F}, \epsilon, d) \leq \dim_q(\mathcal{F}, \epsilon, d)$ for any $p \leq q$. In the analysis, we will focus on ℓ_1 - and ℓ_2 -norm eluder dimension, which have been used in nonlinear bandits and RL extensively (Wen & Van Roy, 2013; Osband & Van Roy, 2014a; Jain et al., 2015; Wang et al., 2020; Ayoub et al., 2020; Foster et al., 2021; Ishfaq et al., 2021; Chen et al., 2022; Liu et al., 2022a; Sekhari et al., 2023a;b). Examples where eluder dimension is small include linear functions, generalized linear models, and functions in Reproducing Kernel Hilbert Space (RKHS).

The other complexity measure we use is the bracketing number (Van de Geer, 2000).

Definition 5.3 (Bracketing number). *Consider a function class $\mathcal{F} \subseteq \mathcal{X} \rightarrow \mathbb{R}$. Given two functions $l, u : \mathcal{X} \rightarrow \mathbb{R}$, the bracket $[l, u]$ is defined as the set of functions $f \in \mathcal{F}$ with $l(x) \leq f(x) \leq u(x)$ for all $x \in \mathcal{X}$. It is called an ω -bracket if $\|l - u\| \leq \omega$. The bracketing number of \mathcal{F} w.r.t. the metric $\|\cdot\|$, denoted by $N_{[]}(\omega, \mathcal{F}, \|\cdot\|)$, is the minimum number of ω -brackets needed to cover \mathcal{F} .*

The logarithm of the bracketing number is small in many common scenarios, which has been extensively examined by previous studies (e.g., Van de Geer (2000)) for deriving MLE generalization bound (Agarwal et al., 2020; Uehara & Sun, 2021; Liu et al., 2022b; 2023). For example, when \mathcal{F} is finite, the bracketing number is bounded by its size. When \mathcal{F} is a d -dimensional linear function class, the logarithm of the bracketing number is upper bounded by d up to logarithmic factors.

It is worth noting that while we will employ both measures to the model class \mathcal{P} , we can not similarly apply them to the reward class \mathcal{R} . Instead, we have to rely on the complexity of the following function class, which comprises functions mapping pairs of trajectories to reward differences:

$$\tilde{\mathcal{R}} := \left\{ \tilde{r} : \tilde{r}(\tau^0, \tau^1) = \sum_{h=1}^H r(s_h^0, a_h^0) - r(s_h^1, a_h^1), \forall \tau^i = \{s_h^i, a_h^i\}_h, i \in \{0, 1\}, r \in \mathcal{R} \right\}. \quad (3)$$

We have to use $\tilde{\mathcal{R}}$ instead of \mathcal{R} because we only receive preference feedback, and thus we cannot guarantee that the learned reward model is accurate state-action-wise. Now we are ready to state our main results. The proofs are provided in Appendix C.

Theorem 5.4. *PbTS (Algorithm 2) guarantees that*

$$\text{BayesRegret}_T = \tilde{O} \left(T\epsilon + H^2 \cdot \dim_1(\mathcal{P}, 1/T) \cdot \sqrt{T \cdot \nu_{\mathcal{P}}} + \kappa \cdot \dim_1(\tilde{\mathcal{R}}, 1/T) \cdot \sqrt{T \cdot \nu_{\mathcal{R}}} \right),$$

$$\text{BayesQueries}_T = \tilde{O} \left(\min \left\{ \frac{\kappa \sqrt{T \cdot \nu_{\mathcal{R}}}}{\epsilon} \cdot \dim_1(\tilde{\mathcal{R}}, \epsilon/2), \frac{\kappa^2 \cdot \nu_{\mathcal{R}}}{\epsilon^2} \cdot \dim_2(\tilde{\mathcal{R}}, \epsilon/2) \right\} \right)$$

where we denote $\nu_{\mathcal{P}} := \log(N_{[]}((HT|\mathcal{S})^{-1}, \mathcal{P}, \|\cdot\|_{\infty}))$ and $\nu_{\mathcal{R}} := \log(N_{[]}(\kappa(2T)^{-1}, \tilde{\mathcal{R}}, \|\cdot\|_{\infty}))$.

Similar to the analysis of Algorithm 1, we study the balance between the Bayesian regret and the query complexity by setting $\epsilon = T^{-\beta}$ for some $\beta \leq 1/2$. Then, we can simplify the bounds into $\text{BayesRegret}_T = \tilde{O}(T^{1-\beta})$ and

$$\text{BayesQueries}_T = \tilde{O}\left(\min\left\{\underbrace{T^{\beta+\frac{1}{2}} \cdot \dim_1(\tilde{\mathcal{R}}, \epsilon/2)}_{(i)}, \underbrace{T^{2\beta} \cdot \dim_2(\tilde{\mathcal{R}}, \epsilon/2)}_{(ii)}\right\}\right)$$

where we have hidden factors except T and the eluder dimension for brevity. We see that there is again a tradeoff in T between the Bayesian regret and the query complexity, similar to the one in Theorem 4.2. Term (ii) demonstrates that the tradeoff in T is again *optimal*, evidenced by the lower bound (Theorem 4.3). Moreover, term (i) further improves the dependence on the eluder dimension (recalling that ℓ_1 -norm version is smaller than the ℓ_2 -norm version). However, the T -dependence is worse. It is desired to derive a query complexity upper bound that scales as $\tilde{O}(T^{2\beta} \cdot \dim_1(\tilde{\mathcal{R}}, \epsilon/2))$, attaining the favorable dependence on both T and the eluder dimension. We leave it as future work.

We emphasize that the Bayesian regret analysis in Theorem 5.4 is not a simple extension of previous TS works. We highlight four main differences: (1) The feedback is preference-based, which necessitates a new Bayesian regret decomposition:

$$\text{BayesRegret}_T = \underbrace{\sum_{t=0}^T \mathbb{E} \left[V_{r_t, P_t}^{\pi_t^0} - V_{r_t, P^*}^{\pi_t^0} \right]}_{\text{T}_{\text{model}}} + \underbrace{\sum_{t=0}^T \mathbb{E} \left[\left(V_{r_t, P^*}^{\pi_t^0} - V_{r_t, P^*}^{\pi_t^1} \right) - \left(V_{r^*, P^*}^{\pi_t^0} - V_{r^*, P^*}^{\pi_t^1} \right) \right]}_{\text{T}_{\text{reward}}}.$$

Here T_{model} and T_{reward} are the respective regret incurred due to model and reward misspecification. We highlight that T_{reward} characterizes the misspecification in terms of the *reward difference* between π_t^0 and π_t^1 , which is different from the standard Bayesian RL. (2) Unlike prior works (Russo & Van Roy, 2014), we do not rely on upper confidence bounds (UCB) or optimism. Instead, we construct version spaces by classic MLE generalization bound. Taking the reward learning as an example, given the preference data $\{\tau_i^0, \tau_i^1, o_i\}_{i=1}^{t-1}$, we construct the version space at round t as

$$\mathcal{V}_t = \left\{ r \in \mathcal{R} : \sum_{i=1}^{t-1} d_{\text{TV}}^2(\Pr(\cdot | \tau_i^1, \tau_i^0, \hat{r}_t), \Pr(\cdot | \tau_i^1, \tau_i^0, r)) \leq \beta \right\}$$

where $\hat{r}_t := \arg \max_r \log \sum_{i=1}^{t-1} \Pr(o_i | \tau_i^1, \tau_i^0, r)$ is the MLE from the preference data and β is tuned appropriately to ensure $r^* \in \mathcal{V}_t$ with high probability. We then show the posterior probability of r_t and r^* not belonging to \mathcal{V}_t is small. (3) Our analysis uses the tighter ℓ_1 -norm eluder dimension, which is strictly better than the ℓ_2 -norm eluder dimension used in prior work. (4) We also equipped it with a randomized active learning procedure for query complexity minimization.

Computation. The computational bottleneck of Algorithm 2 lies in the computation of the posterior distribution (Line 3). Prior TS works have used Bootstrapping to approximate posterior sampling (Osband et al., 2016a; 2023) and achieved competitive performance in common RL benchmarks.

Non-Markovian reward. Algorithm 2 can also be applied to non-Markovian reward (i.e., reward model is trajectory-wise) without any change. Here we consider Markovian reward for the consistency with Algorithm 1 and for the purpose of using a standard planning oracle for computing an optimal policy from a reward and transition model. While non-Markovian reward is more general, it is unclear how to solve the planning problem efficiently even in tabular MDPs. This computational intractability makes non-Markovian rewards not easily applicable in practice.

Extension to SEC. In Appendix C.4, we extend the eluder dimension in Theorem 5.4 to the Sequential Extrapolation Coefficient (SEC) (Xie et al., 2022), which is more general.

6 CONCLUSION

We use randomization to design algorithms for RL with preference-based feedback. Randomization allows us to minimize regret and query complexity while at the same time maintaining computation efficiency. For linear models, our algorithms achieve a near-optimal balance between the worst-case reward regret and query complexity with computational efficiency. For models beyond linear, using eluder dimension, we present a TS-inspired algorithm that balances Bayesian regret and Bayesian query complexity nearly optimally.

ACKNOWLEDGEMENTS

Both RW and WS acknowledge support from NSF IIS-2154711, NSF CAREER 2339395, and Cornell Infosys Collaboration.

REFERENCES

- Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 24, 2011.
- Marc Abeille and Alessandro Lazaric. Linear thompson sampling revisited. In *Artificial Intelligence and Statistics*, pp. 176–184. PMLR, 2017.
- Alekh Agarwal. Selective sampling algorithms for cost-sensitive multiclass prediction. In *International Conference on Machine Learning*, pp. 1220–1228. PMLR, 2013.
- Alekh Agarwal and Tong Zhang. Model-based rl with optimistic posterior sampling: Structural conditions and sample complexity. *Advances in Neural Information Processing Systems*, 35: 35284–35297, 2022.
- Alekh Agarwal, Sham Kakade, Akshay Krishnamurthy, and Wen Sun. Flambe: Structural complexity and representation learning of low rank mdps. *Advances in neural information processing systems*, 33:20095–20107, 2020.
- Priyank Agrawal, Jinglin Chen, and Nan Jiang. Improved worst-case regret bounds for randomized least-squares value iteration. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 6566–6573, 2021.
- Shipra Agrawal and Randy Jia. Optimistic posterior sampling for reinforcement learning: worst-case regret bounds. *Advances in Neural Information Processing Systems*, 30, 2017.
- Alex Ayoub, Zeyu Jia, Csaba Szepesvari, Mengdi Wang, and Lin Yang. Model-based reinforcement learning with value-targeted regression. In *International Conference on Machine Learning*, pp. 463–474. PMLR, 2020.
- Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforcement learning. In *International Conference on Machine Learning*, pp. 263–272. PMLR, 2017.
- Viktor Bengs, Aadirupa Saha, and Eyke Hüllermeier. Stochastic contextual dueling bandits under linear stochastic transitivity models. In *International Conference on Machine Learning*, pp. 1764–1786. PMLR, 2022.
- Ralph Allan Bradley and Milton E Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952.
- Nicolo Cesa-Bianchi, Gábor Lugosi, and Gilles Stoltz. Minimizing regret with label efficient prediction. *IEEE Transactions on Information Theory*, 51(6):2152–2162, 2005.
- Niladri Chatterji, Aldo Pacchiano, Peter Bartlett, and Michael Jordan. On the theory of reinforcement learning with once-per-episode feedback. *Advances in Neural Information Processing Systems*, 34:3401–3412, 2021.
- Xiaoyu Chen, Han Zhong, Zhuoran Yang, Zhaoran Wang, and Liwei Wang. Human-in-the-loop: Provably efficient preference-based reinforcement learning with general function approximation. In *International Conference on Machine Learning*, pp. 3773–3793. PMLR, 2022.
- Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30, 2017.
- Ofer Dekel, Claudio Gentile, and Karthik Sridharan. Selective sampling and active learning from single and multiple experts. *The Journal of Machine Learning Research*, 13(1):2655–2697, 2012.

- Simon Du, Sham Kakade, Jason Lee, Shachar Lovett, Gaurav Mahajan, Wen Sun, and Ruosong Wang. Bilinear classes: A structural framework for provable generalization in rl. In *International Conference on Machine Learning*, pp. 2826–2836. PMLR, 2021.
- Miroslav Dudík, Katja Hofmann, Robert E Schapire, Aleksandrs Slivkins, and Masrour Zoghi. Contextual dueling bandits. In *Conference on Learning Theory*, pp. 563–587. PMLR, 2015.
- Yonathan Efroni, Nadav Merlis, and Shie Mannor. Reinforcement learning with trajectory feedback. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pp. 7288–7295, 2021.
- Dylan Foster, Alexander Rakhlin, David Simchi-Levi, and Yunzong Xu. Instance-dependent complexity of contextual bandits and reinforcement learning: A disagreement-based perspective. In *Conference on Learning Theory*, pp. 2059–2059. PMLR, 2021.
- Sara A Geer. *Empirical Processes in M-estimation*, volume 6. Cambridge university press, 2000.
- Aditya Gopalan and Shie Mannor. Thompson sampling for learning parameterized markov decision processes. In *Conference on Learning Theory*, pp. 861–898. PMLR, 2015.
- Steve Hanneke and Liu Yang. Minimax analysis of active learning. *J. Mach. Learn. Res.*, 16(1): 3487–3602, 2015.
- Steve Hanneke and Liu Yang. Toward a general theory of online selective sampling: Trading off mistakes and queries. In *International Conference on Artificial Intelligence and Statistics*, pp. 3997–4005. PMLR, 2021.
- Haque Ishfaq, Qiwen Cui, Viet Nguyen, Alex Ayoub, Zhuoran Yang, Zhaoran Wang, Doina Precup, and Lin Yang. Randomized exploration in reinforcement learning with general value function approximation. In *International Conference on Machine Learning*, pp. 4607–4616. PMLR, 2021.
- Ashesh Jain, Brian Wojcik, Thorsten Joachims, and Ashutosh Saxena. Learning trajectory preferences for manipulators via iterative improvement. *Advances in neural information processing systems*, 26, 2013.
- Ashesh Jain, Shikhar Sharma, Thorsten Joachims, and Ashutosh Saxena. Learning preferences for manipulation tasks from online coactive feedback. *The International Journal of Robotics Research*, 34(10):1296–1313, 2015.
- Chi Jin, Zhuoran Yang, Zhaoran Wang, and Michael I Jordan. Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*, pp. 2137–2143. PMLR, 2020.
- Chi Jin, Qinghua Liu, and Sobhan Miryoosefi. Bellman eluder dimension: New rich classes of rl problems, and sample-efficient algorithms. *Advances in neural information processing systems*, 34:13406–13418, 2021.
- Michael Laskey, Sam Staszak, Wesley Yu-Shu Hsieh, Jeffrey Mahler, Florian T Pokorny, Anca D Dragan, and Ken Goldberg. Shiv: Reducing supervisor burden in dagger using support vectors for efficient learning from demonstrations in high dimensional state spaces. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 462–469. IEEE, 2016.
- Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step. In *The Twelfth International Conference on Learning Representations*, 2023.
- Qinghua Liu, Alan Chung, Csaba Szepesvári, and Chi Jin. When is partially observable reinforcement learning not scary? In *Conference on Learning Theory*, pp. 5175–5220. PMLR, 2022a.
- Qinghua Liu, Csaba Szepesvári, and Chi Jin. Sample-efficient reinforcement learning of partially observable markov games. *Advances in Neural Information Processing Systems*, 35:18296–18308, 2022b.
- Qinghua Liu, Praneeth Netrapalli, Csaba Szepesvari, and Chi Jin. Optimistic mle: A generic model-based algorithm for partially observable sequential decision making. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pp. 363–376, 2023.

- Ellen Novoseller, Yibing Wei, Yanan Sui, Yisong Yue, and Joel Burdick. Dueling posterior sampling for preference-based reinforcement learning. In *Conference on Uncertainty in Artificial Intelligence*, pp. 1029–1038. PMLR, 2020.
- Ian Osband and Benjamin Van Roy. Model-based reinforcement learning and the eluder dimension. *Advances in Neural Information Processing Systems*, 27, 2014a.
- Ian Osband and Benjamin Van Roy. Near-optimal reinforcement learning in factored mdps. *Advances in Neural Information Processing Systems*, 27, 2014b.
- Ian Osband, Daniel Russo, and Benjamin Van Roy. (more) efficient reinforcement learning via posterior sampling. *Advances in Neural Information Processing Systems*, 26, 2013.
- Ian Osband, Charles Blundell, Alexander Pritzel, and Benjamin Van Roy. Deep exploration via bootstrapped dqn. *Advances in neural information processing systems*, 29, 2016a.
- Ian Osband, Benjamin Van Roy, and Zheng Wen. Generalization and exploration via randomized value functions. In *International Conference on Machine Learning*, pp. 2377–2386. PMLR, 2016b.
- Ian Osband, Zheng Wen, Seyed Mohammad Asghari, Vikranth Dwaracherla, Morteza Ibrahimi, Xiuyuan Lu, and Benjamin Van Roy. Approximate thompson sampling via epistemic neural networks. In *Uncertainty in Artificial Intelligence*, pp. 1586–1595. PMLR, 2023.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35: 27730–27744, 2022.
- David Pollard. Empirical processes: theory and applications. Ims, 1990.
- Nikita Puchkin and Nikita Zhivotovskiy. Exponential savings in agnostic active learning through abstention. In *Conference on Learning Theory*, pp. 3806–3832. PMLR, 2021.
- Stéphane Ross, Narek Melik-Barkhudarov, Kumar Shaurya Shankar, Andreas Wendel, Debadeepta Dey, J Andrew Bagnell, and Martial Hebert. Learning monocular reactive uav control in cluttered natural environments. In *2013 IEEE international conference on robotics and automation*, pp. 1765–1772. IEEE, 2013.
- Daniel Russo and Benjamin Van Roy. Eluder dimension and the sample complexity of optimistic exploration. *Advances in Neural Information Processing Systems*, 26, 2013.
- Daniel Russo and Benjamin Van Roy. Learning to optimize via posterior sampling. *Mathematics of Operations Research*, 39(4):1221–1243, 2014.
- Aadirupa Saha, Aldo Pacchiano, and Jonathan Lee. Dueling rl: Reinforcement learning with trajectory preferences. In *International Conference on Artificial Intelligence and Statistics*, pp. 6263–6289. PMLR, 2023.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Ayush Sekhari, Karthik Sridharan, Wen Sun, and Runzhe Wu. Contextual bandits and imitation learning with preference-based active queries. *Advances in Neural Information Processing Systems*, 36, 2023a.
- Ayush Sekhari, Karthik Sridharan, Wen Sun, and Runzhe Wu. Selective sampling and imitation learning via online regression. *Advances in Neural Information Processing Systems*, 36, 2023b.
- Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. Learning to summarize with human feedback. *Advances in Neural Information Processing Systems*, 33:3008–3021, 2020.
- Masatoshi Uehara and Wen Sun. Pessimistic model-based offline reinforcement learning under partial coverage. In *International Conference on Learning Representations*, 2021.

- Sara Van de Geer. *Empirical Processes in M-estimation*, volume 6. Cambridge university press, 2000.
- Ruosong Wang, Russ R Salakhutdinov, and Lin Yang. Reinforcement learning with general value function approximation: Provably efficient approach via bounded eluder dimension. *Advances in Neural Information Processing Systems*, 33:6123–6135, 2020.
- Yuanhao Wang, Qinghua Liu, and Chi Jin. Is rlhf more difficult than standard rl? *arXiv preprint arXiv:2306.14111*, 2023.
- Zheng Wen and Benjamin Van Roy. Efficient exploration and value function generalization in deterministic systems. *Advances in Neural Information Processing Systems*, 26, 2013.
- Runzhe Wu, Masatoshi Uehara, and Wen Sun. Distributional offline policy evaluation with predictive error guarantees. In *International Conference on Machine Learning*, pp. 37685–37712. PMLR, 2023a.
- Yue Wu, Tao Jin, Hao Lou, Farzad Farnoud, and Quanquan Gu. Borda regret minimization for generalized linear dueling bandits. *arXiv preprint arXiv:2303.08816*, 2023b.
- Tengyang Xie, Dylan J Foster, Yu Bai, Nan Jiang, and Sham M Kakade. The role of coverage in online reinforcement learning. In *The Eleventh International Conference on Learning Representations*, 2022.
- Yichong Xu, Ruosong Wang, Lin Yang, Aarti Singh, and Artur Dubrawski. Preference-based reinforcement learning with finite-time guarantees. *Advances in Neural Information Processing Systems*, 33:18784–18794, 2020.
- Yisong Yue and Thorsten Joachims. Beat the mean bandit. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pp. 241–248. Citeseer, 2011.
- Yisong Yue, Josef Broder, Robert Kleinberg, and Thorsten Joachims. The k-armed dueling bandits problem. *Journal of Computer and System Sciences*, 78(5):1538–1556, 2012.
- Andrea Zanette, David Brandfonbrener, Emma Brunskill, Matteo Pirotta, and Alessandro Lazaric. Frequentist regret bounds for randomized least-squares value iteration. In *International Conference on Artificial Intelligence and Statistics*, pp. 1954–1964. PMLR, 2020.
- Wenhao Zhan, Masatoshi Uehara, Nathan Kallus, Jason D Lee, and Wen Sun. Provable offline reinforcement learning with human feedback. *arXiv preprint arXiv:2305.14816*, 2023a.
- Wenhao Zhan, Masatoshi Uehara, Wen Sun, and Jason D Lee. How to query human feedback efficiently in rl? *arXiv preprint arXiv:2305.18505*, 2023b.
- Han Zhong, Wei Xiong, Sirui Zheng, Liwei Wang, Zhaoran Wang, Zhuoran Yang, and Tong Zhang. Gec: A unified framework for interactive decision making in mdp, pomdp, and beyond. *CoRR*, 2022.
- Banghua Zhu, Michael Jordan, and Jiantao Jiao. Principled reinforcement learning with human feedback from pairwise or k-wise comparisons. In *International Conference on Machine Learning*, pp. 43037–43067. PMLR, 2023.
- Yinglun Zhu and Robert Nowak. Efficient active learning with abstention. *Advances in Neural Information Processing Systems*, 35:35379–35391, 2022.

A COMPUTATIONAL CHALLENGES IN RL WITH PREFERENCE-BASED FEEDBACK

For RL with preference-based feedback, there is currently no algorithm that achieves sublinear worst-case regret and computational efficiency simultaneously, even for tabular MDPs. In this section, we discuss the challenges that hindered previous works from being computationally efficient. Specifically, the reasons are twofold:

(1) *Trajectory-wise information.* The feedback is trajectory-wise, meaning that we only receive information about cumulative rewards instead of per-step rewards. In this case, there is no longer a ground truth per-step reward signal. To see this, consider an MDP with two steps and a trajectory with a cumulative reward of 1. Then, we cannot decide the respective per-step reward for the two steps — it could be that the first step has reward 1 and the second has reward 0, or both have reward 0.5. The preference-based feedback is strictly harder than trajectory-wise reward feedback, and thus the same issue persists. This invalidates all algorithms relying on per-step reward information (e.g., UCBCVI (Azar et al., 2017)) and necessitates leveraging feedback signal at a trajectory level. However, trajectory-level approaches typically entail maintaining a version space via trajectory constraints (Saha et al., 2023; Chen et al., 2022; Zhan et al., 2023a), and the computational complexity of searching within this version space is at least exponential in the length of the episode. Some works circumvent this computational obstacle by additional assumptions (e.g., explorability (Chatterji et al., 2021)), which are restrictive and do not generally hold even in tabular MDPs.

(2) *Preference-based information.* The feedback relies on a pair of policies. Standard algorithms based on a single policy become computationally intractable when adapting to this setting since optimizing over a pair of policies simultaneously is qualitatively different. For example, Zhan et al. (2023b); Saha et al. (2023) use the idea from optimal design and need the computation oracle: $\arg \max_{\pi, \pi' \in \Pi} \|\mathbb{E}_{s, a \sim \pi} \phi(s, a) - \mathbb{E}_{s, a \sim \pi'} \phi(s, a)\|_A$ for some positive definite matrix A . Here $\|x\|_A^2 := x^\top A x$, and ϕ is some state-action wise feature.³ It is unclear how to implement this oracle since standard planning approaches relying on dynamic programming cannot be applied here. Additionally, these methods also require actively maintaining a policy space Π by eliminating potentially sub-optimal policies. The policy class can be exponentially large even in tabular settings, so it is unclear how to maintain a valid policy space in a computationally tractable manner.

These challenges motivate us to devise a novel algorithm using a technique distinct from previous approaches. Our solution centers around the concept of *randomization*, which allows us to balance exploration and exploitation and thus enable standard efficient computation oracles (e.g., DP-style planning oracle like value iteration).

B PROOF OF THEOREM 4.2

B.1 NOTATIONS

We define some symbols and their values in Table 1. We have categorized them into four classes for the ease of reference.

The concept of covering number is defined below, which will be used to bound the statistical error of our algorithm.

Definition B.1 (Covering number). *Consider a function class $\mathcal{F} \subseteq \mathcal{X} \rightarrow \mathbb{R}$. The ω -cover of a function $\hat{f} \in \mathcal{F}$ is defined as the set of functions $f \in \mathcal{F}$ for which $\|f - \hat{f}\| \leq \omega$. The covering number of \mathcal{F} w.r.t. the metric $\|\cdot\|$ denoted by $N(\omega, \mathcal{F}, \|\cdot\|)$ is the minimum number of ω -covers needed to cover \mathcal{F} .*

³These prior work typically assume trajectory wise feature $\phi(\tau)$ for a state-action wise trajectory τ . However, even when specializing to state-action-wise features, these algorithms are still not computationally tractable even in tabular MDPs.

Table 1: Symbols and their respective values.

Symbol	Value
<i>(1) Error components</i>	
$\xi_{r,t}$	$\bar{\theta}_{r,t} - \hat{\theta}_{r,t}$
$\xi_{P,t,h}$	$\bar{\theta}_{P,t,h} - \hat{\theta}_{P,t,h}$
$\lambda_{t,h}$	$\lambda \Sigma_{t-1,h}^{-1} \int_{s'} \mu_h^*(s') \bar{V}_{t,h+1}(s') ds'$
$\eta_{r,t}$	$\hat{\theta}_{r,t} - \theta_r^*$
$\eta_{P,t,h}$	$\Sigma_{t-1,h}^{-1} \left(\sum_{i=1}^{t-1} \phi(s_{i,h}, a_{i,h}) (\bar{V}_{t,h+1}(s_{i,h+1}) - \mathbb{E}_{s_{i,h+1}} [\bar{V}_{t,h+1}(s_{i,h+1}) s_{i,h}, a_{i,h}]) \right)$
$\theta_{P,t,h}^*$	$\int_{s'} \mu_h^*(s') [\bar{V}_{t,h+1}(s')] ds'$
<i>(2) Statistical upper bounds</i>	
$\epsilon_{r,\xi}$	$\sigma_r \sqrt{2d \log(2dT/\delta)}$
$\epsilon_{r,\eta}$	$\sqrt{80\kappa d \log(24BT^2/(\bar{\kappa}\delta)) + 168B^2 d \log(6BT^2/\delta) + 4\lambda B^2}$
V_{\max}	$H(2 + (\epsilon_{r,\xi} + \epsilon_{r,\eta})/\sqrt{\lambda})$
ϵ_λ	$V_{\max} \sqrt{\lambda d}$
$\epsilon_{P,\xi}$	$\sigma_P \sqrt{2d \log(2dHT/\delta)}$
ι_ϵ	$\log \left(12HT^2(T + \lambda) V_{\max} \left(B + \frac{2V_{\max} \sqrt{dt} + \epsilon_{P,\xi} + \epsilon_{r,\xi}}{\sqrt{\lambda}} \right) \right)$
χ	(defined in Lemma B.3)
$\epsilon'_{P,\eta}$	$\frac{6}{\sqrt{\lambda}} + 16dV_{\max} \sqrt{\iota_\epsilon - \log((\alpha_U - \alpha_L)\delta\lambda)}$
$\epsilon_{P,\eta}$	$\chi \cdot \left(\frac{6}{\sqrt{\lambda}} + 16dV_{\max} \sqrt{\iota_\epsilon - \log(\delta\lambda)} \right)$
<i>(3) Value cutoff</i>	
$L_{t,h}(s)$	$\mathbb{1}\{\ \phi(s, \pi_{t,h}(s))\ _{\Sigma_{t-1,h}^{-1}} \leq \alpha_L\}$
$L_{t,h}^c(s)$	$\mathbb{1}\{\ \phi(s, \pi_{t,h}(s))\ _{\Sigma_{t-1,h}^{-1}} > \alpha_L\}$
L_{\max}	$\frac{2dH}{\alpha_L^2} \cdot \log\left(\frac{\lambda+T}{\lambda}\right)$
<i>(4) Hyperparameters</i>	
σ_r	$\epsilon_{r,\eta}$
σ_P	$(\epsilon_{P,\eta} + \epsilon_\lambda) \sqrt{H}$
α_U	$(\epsilon_{P,\xi} + \epsilon_{P,\eta} + \epsilon_\lambda)^{-1}$
α_L	$(\epsilon_{P,\xi} + \epsilon_{P,\eta} + \epsilon_\lambda)^{-1/2}$
λ	1

B.2 SUPPORTING LEMMAS

Lemma B.2 (Covering number of Euclidean balls). (*Pollard, 1990*) Let $\Theta_B := \{\theta \in \mathbb{R}^d : \|\theta\|_2 \leq B\}$. Then we have

$$N(\omega, \Theta_B, \|\cdot\|_2) \leq (3B/\omega)^d.$$

Lemma B.3. There exists $\chi = O(\log(\epsilon_{P,\xi} + \epsilon_{P,\eta} + \epsilon_\lambda))$ such that $\epsilon'_{P,\eta} \leq \epsilon_{P,\eta}$ (recalling that χ is in the definition of $\epsilon_{P,\eta}$).

Proof of Lemma B.3. By definition, we have

$$\begin{aligned} \epsilon'_{P,\eta} &= \frac{6}{\sqrt{\lambda}} + 16dV_{\max} \sqrt{\iota_\epsilon - \log((\alpha_U - \alpha_L)\delta\lambda)} \\ &= \frac{6}{\sqrt{\lambda}} + 16dV_{\max} \sqrt{\iota_\epsilon - \log(\delta\lambda) - \log(\alpha_U - \alpha_L)} \\ &= \frac{6}{\sqrt{\lambda}} + 16dV_{\max} \sqrt{\iota_\epsilon - \log(\delta\lambda) + \log 2(\epsilon_{P,\xi} + \epsilon_{P,\eta} + \epsilon_\lambda)} \\ &\leq \chi \left(\frac{6}{\sqrt{\lambda}} + 16dV_{\max} \sqrt{\iota_\epsilon - \log(\delta\lambda)} \right) \\ &= \epsilon_{P,\eta}. \end{aligned}$$

Here the third equality is by the definition of α_L and α_U . The inequality holds by setting $\chi = O(\log(\epsilon_{P,\xi} + \epsilon_{P,\eta} + \epsilon_\lambda))$ to be large enough. \square

Lemma B.4. It holds that $\eta_{P,t,h} + \lambda_{t,h} = \widehat{\theta}_{P,t,h} - \theta_{P,t,h}^*$ for any $t \in [T]$ and $h \in [H]$.

Proof of Lemma B.4. By definition, we have

$$\begin{aligned} &\eta_{P,t,h} + \lambda_{t,h} \\ &= \Sigma_{t-1,h}^{-1} \left(\sum_{i=1}^{t-1} \phi(s_{i,h}, a_{i,h}) \left(\bar{V}_{t,h+1}(s_{i,h+1}) - \mathbb{E}_{s_{h+1}} [\bar{V}_{t,h+1}(s_{i,h+1}) | s_{i,h}, a_{i,h}] \right) \right) \\ &\quad + \lambda \Sigma_{t-1,h}^{-1} \int_{s'} \mu_h^*(s') \bar{V}_{t,h+1}(s') ds' \\ &= \widehat{\theta}_{P,t,h} - \Sigma_{t-1,h}^{-1} \left(\sum_{i=1}^{t-1} \phi(s_{i,h}, a_{i,h}) \phi(s_{i,h}, a_{i,h}) + \lambda I \right)^\top \int_{s'} \mu_h^*(s') [\bar{V}_{t,h+1}(s')] ds' \\ &= \widehat{\theta}_{P,t,h} - \int_{s'} \mu_h^*(s') [\bar{V}_{t,h+1}(s')] ds' \\ &= \widehat{\theta}_{P,t,h} - \theta_{P,t,h}^*. \end{aligned}$$

\square

The following lemma is adapted from [Zanette et al. \(2020, Lemma 1\)](#) to our setting.

Lemma B.5 (One-step decomposition). For any $t \in [T]$, $h \in [H]$, $s \in \mathcal{S}$, $a \in \mathcal{A}$, and policy π , we have

$$\begin{aligned} &\phi(s, a)^\top (\bar{\theta}_{r,t} + \bar{\theta}_{P,t,h}) - Q_h^\pi(s, a) \\ &= \phi(s, a)^\top (\xi_{r,t} + \xi_{P,t,h} + \eta_{r,t} + \eta_{P,t,h} - \lambda_{t,h}) + \mathbb{E}_{s'} [\bar{V}_{t,h+1}(s') - V_{h+1}^\pi(s') | s, a]. \end{aligned}$$

Proof. We have

$$\begin{aligned} &\phi(s, a)^\top (\bar{\theta}_{r,t} + \bar{\theta}_{P,t,h}) - Q_h^\pi(s, a) \\ &= \phi(s, a)^\top (\bar{\theta}_{r,t} + \bar{\theta}_{P,t,h}) - \phi(s, a)^\top \theta_r^* - \mathbb{E}_{s'} [V_{h+1}^\pi(s') | s, a] \\ &= \phi(s, a)^\top (\bar{\theta}_{r,t} - \widehat{\theta}_{r,t}) + \phi(s, a)^\top (\bar{\theta}_{P,t,h} - \widehat{\theta}_{P,t,h}) + \phi(s, a)^\top (\widehat{\theta}_{r,t} - \theta_r^*) \\ &\quad + \phi(s, a)^\top \widehat{\theta}_{P,t,h} - \mathbb{E}_{s'} [V_{h+1}^\pi(s') | s, a]. \end{aligned}$$

For the last but one term, we note that

$$\begin{aligned}
& \phi(s, a)^\top \widehat{\theta}_{\mathbb{P}, t, h} \\
&= \phi(s, a)^\top \Sigma_{t-1, h}^{-1} \left(\sum_{i=1}^{t-1} \phi(s_{i, h}, a_{i, h}) \bar{V}_{t, h+1}(s_{i, h+1}) \right) \\
&= \phi(s, a)^\top \Sigma_{t-1, h}^{-1} \left(\sum_{i=1}^{t-1} \phi(s_{i, h}, a_{i, h}) \mathbb{E}_{s_{h+1}} [\bar{V}_{t, h+1}(s_{i, h+1}) \mid s_{i, h}, a_{i, h}] \right) + \phi(s, a)^\top \eta_{\mathbb{P}, t, h} \\
&= \phi(s, a)^\top \Sigma_{t-1, h}^{-1} \left(\sum_{i=1}^{t-1} \phi(s_{i, h}, a_{i, h}) \phi(s_{i, h}, a_{i, h})^\top \int_{s'} \mu_h^*(s') \bar{V}_{t, h+1}(s') \, ds' \right) + \phi(s, a)^\top \eta_{\mathbb{P}, t, h} \\
&= \phi(s, a)^\top \int_{s'} \mu_h^*(s') \bar{V}_{t, h+1}(s') \, ds' - \lambda \phi(s, a)^\top \Sigma_{t-1, h}^{-1} \int_{s'} \mu_h^*(s') \bar{V}_{t, h+1}(s') \, ds' + \phi(s, a)^\top \eta_{\mathbb{P}, t, h} \\
&= \mathbb{E}_{s'} [\bar{V}_{t, h+1}(s') \mid s, a] - \phi(s, a)^\top \lambda_{t, h} + \phi(s, a)^\top \eta_{\mathbb{P}, t, h}.
\end{aligned}$$

Plugging this back, we get

$$\begin{aligned}
& \phi(s, a)^\top (\bar{\theta}_{r, t} + \bar{\theta}_{\mathbb{P}, t, h}) - Q_h^\pi(s, a) \\
&= \phi(s, a)^\top (\bar{\theta}_{r, t} - \widehat{\theta}_{r, t}) + \phi(s, a)^\top (\bar{\theta}_{\mathbb{P}, t, h} - \widehat{\theta}_{\mathbb{P}, t, h}) + \phi(s, a)^\top (\widehat{\theta}_{r, t} - \theta_r^*) \\
&\quad + \mathbb{E}_{s'} [\bar{V}_{t, h+1}(s') - V_{h+1}^\pi(s') \mid s, a] - \phi(s, a)^\top \lambda_{t, h} + \phi(s, a)^\top \eta_{\mathbb{P}, t, h} \\
&= \phi(s, a)^\top (\xi_{r, t} + \xi_{\mathbb{P}, t, h} + \eta_{r, t} + \eta_{\mathbb{P}, t, h} - \lambda_{t, h}) + \mathbb{E}_{s'} [\bar{V}_{t, h+1}(s') - V_{h+1}^\pi(s') \mid s, a].
\end{aligned}$$

This completes the proof. \square

Lemma B.6. For any $t \in [T]$, if $Z_t = 0$, then we have

$$\|\phi(\tau_t^0) - \phi(\tau_t^1)\|_{\Sigma_{t-1}^{-1}} \leq \epsilon \sqrt{\pi} / (2\sigma_r).$$

Proof. By the definition of Z_t , we have

$$\begin{aligned}
\epsilon &\geq \mathbb{E}_{\theta_0, \theta_1 \sim \mathcal{N}(\bar{\theta}_{r, t}, \sigma_r^2 \Sigma_{t-1}^{-1})} |(\phi(\tau_t^0) - \phi(\tau_t^1))^\top (\theta_0 - \theta_1)| \\
&= \mathbb{E}_{u_0, u_1 \sim \mathcal{N}(0, I_d)} |(\phi(\tau_t^0) - \phi(\tau_t^1))^\top \sigma_r \Sigma_{t-1}^{-1/2} (u_0 - u_1)| \\
&= \mathbb{E}_{u \sim \mathcal{N}(0, I_d)} |(\phi(\tau_t^0) - \phi(\tau_t^1))^\top \sqrt{2} \sigma_r \Sigma_{t-1}^{-1/2} u| \\
&= \sigma_r \sqrt{2} \mathbb{E}_{u \sim \mathcal{N}(0, I_d)} \sqrt{u^\top \Sigma_{t-1}^{-1/2} (\phi(\tau_t^0) - \phi(\tau_t^1)) (\phi(\tau_t^0) - \phi(\tau_t^1))^\top \Sigma_{t-1}^{-1/2} u} \\
&= \sigma_r \sqrt{2} \mathbb{E}_{u \sim \mathcal{N}(0, I_d)} \|u\|_{\Sigma_{t-1}^{-1/2} (\phi(\tau_t^0) - \phi(\tau_t^1)) (\phi(\tau_t^0) - \phi(\tau_t^1))^\top \Sigma_{t-1}^{-1/2}}.
\end{aligned}$$

We then apply Lemma D.9 and obtain

$$\begin{aligned}
\epsilon &\geq \sigma_r \sqrt{2} \cdot \sqrt{\frac{2 \operatorname{tr} \left(\Sigma_{t-1}^{-1/2} (\phi(\tau_t^0) - \phi(\tau_t^1)) (\phi(\tau_t^0) - \phi(\tau_t^1))^\top \Sigma_{t-1}^{-1/2} \right)}{\pi}} \\
&= 2\sigma_r \cdot \sqrt{\frac{(\phi(\tau_t^0) - \phi(\tau_t^1))^\top \Sigma_{t-1}^{-1/2} \Sigma_{t-1}^{-1/2} (\phi(\tau_t^0) - \phi(\tau_t^1))}{\pi}} \\
&= \frac{2\sigma_r}{\sqrt{\pi}} \cdot \|\phi(\tau_t^0) - \phi(\tau_t^1)\|_{\Sigma_{t-1}^{-1}}.
\end{aligned}$$

Hence, we complete the proof. \square

B.3 STATISTICAL UPPER BOUNDS

Lemma B.7 (Gaussian noise). *Each of the following holds with probability at least $1 - \delta$:*

1. $\|\xi_{r,t}\|_{\Sigma_{t-1}} \leq \epsilon_{r,\xi}$ for all $t \in [T]$,
2. $\|\xi_{P,t,h}\|_{\Sigma_{t-1,h}} \leq \epsilon_{P,\xi}$ for all $t \in [T]$ and $h \in [H]$.

Proof of Lemma B.7. It simply follows from Lemma D.3 and the union bound. \square

Lemma B.8. *Fix $t \in [T]$. It holds with probability at least $1 - \delta$ that*

$$\begin{aligned} & \sum_{s=1}^{t-1} \left| (\phi(\tau_s^0) - \phi(\tau_s^1))^\top (\theta_r^* - \hat{\theta}_{r,t}) \right|^2 \\ & \leq 8 \sum_{s=1}^{t-1} \mathbb{E}_{\tau_s^0, \tau_s^1} \left| (\phi(\tau_s^0) - \phi(\tau_s^1))^\top (\theta_r^* - \hat{\theta}_{r,t}) \right|^2 + 168B^2 d \log(3TB/\delta). \end{aligned}$$

Here the expectation is taken over the randomness of sampling τ_s^0 and τ_s^1 from π_s^0 and π_s^1 , respectively.

Proof of Lemma B.8. Let $\widetilde{\Theta}_B$ denote an ω -covering of Θ_B with respect to $\|\cdot\|_2$ for some $\omega > 0$. By Lemma B.2, we have $|\widetilde{\Theta}_B| \leq (3B/\omega)^d$. Moreover, for any $s \in [t-1]$, by Cauchy-Schwarz inequality, it holds that

$$\left| (\phi(\tau_s^0) - \phi(\tau_s^1))^\top (\theta_r^* - \tilde{\theta}) \right|^2 \leq \|\phi(\tau_s^0) - \phi(\tau_s^1)\|_2^2 \cdot \|\theta_r^* - \tilde{\theta}\|_2^2 \leq 16B^2.$$

Then, by Lemma D.4 and union bound over $\widetilde{\Theta}_B$, it holds that

$$\begin{aligned} & \sum_{s=1}^{t-1} \left| (\phi(\tau_s^0) - \phi(\tau_s^1))^\top (\theta_r^* - \tilde{\theta}) \right|^2 \\ & \leq 2 \sum_{s=1}^{t-1} \mathbb{E}_{\tau_s^0, \tau_s^1} \left| (\phi(\tau_s^0) - \phi(\tau_s^1))^\top (\theta_r^* - \tilde{\theta}) \right|^2 + 64B^2 d \log(3B/(\omega\delta)) \end{aligned} \quad (4)$$

for all $\tilde{\theta} \in \widetilde{\Theta}_B$ with probability at least $1 - \delta$. Now we consider an arbitrary $\theta \in \Theta_B$ and let $\tilde{\theta} \in \widetilde{\Theta}_B$ be the closest to θ . Then, by triangle inequality and Cauchy-Schwarz inequality, we have the following two inequalities:

$$\begin{aligned} & \left| (\phi(\tau_s^0) - \phi(\tau_s^1))^\top (\theta_r^* - \theta) \right|^2 \\ & \leq 2 \left| (\phi(\tau_s^0) - \phi(\tau_s^1))^\top (\theta_r^* - \tilde{\theta}) \right|^2 + 2 \|\phi(\tau_s^0) - \phi(\tau_s^1)\|_2^2 \|\theta - \tilde{\theta}\|_2^2 \\ & \leq 2 \left| (\phi(\tau_s^0) - \phi(\tau_s^1))^\top (\theta_r^* - \tilde{\theta}) \right|^2 + 8\omega^2, \end{aligned}$$

and

$$\begin{aligned} & \left| (\phi(\tau_s^0) - \phi(\tau_s^1))^\top (\theta_r^* - \tilde{\theta}) \right|^2 \\ & \leq 2 \left| (\phi(\tau_s^0) - \phi(\tau_s^1))^\top (\theta_r^* - \theta) \right|^2 + 2 \|\phi(\tau_s^0) - \phi(\tau_s^1)\|_2^2 \|\theta - \tilde{\theta}\|_2^2 \\ & \leq 2 \left| (\phi(\tau_s^0) - \phi(\tau_s^1))^\top (\theta_r^* - \theta) \right|^2 + 8\omega^2 \end{aligned}$$

Applying these inequalities and (4), we get

$$\begin{aligned}
& \sum_{s=1}^{t-1} \left| (\phi(\tau_s^0) - \phi(\tau_s^1))^\top (\theta_r^* - \theta) \right|^2 \\
& \leq 2 \sum_{s=1}^{t-1} \left| (\phi(\tau_s^0) - \phi(\tau_s^1))^\top (\theta_r^* - \tilde{\theta}) \right|^2 + 8\omega^2 t \\
& \leq 4 \sum_{s=1}^{t-1} \mathbb{E}_{\tau_s^0, \tau_s^1} \left| (\phi(\tau_s^0) - \phi(\tau_s^1))^\top (\theta_r^* - \tilde{\theta}) \right|^2 + 128B^2 d \log(3B/(\omega\delta)) + 8\omega^2 t \\
& \leq 8 \sum_{s=1}^{t-1} \mathbb{E}_{\tau_s^0, \tau_s^1} \left| (\phi(\tau_s^0) - \phi(\tau_s^1))^\top (\theta_r^* - \theta) \right|^2 + 128B^2 d \log(3B/(\omega\delta)) + 40\omega^2 t \\
& \leq 8 \sum_{s=1}^{t-1} \mathbb{E}_{\tau_s^0, \tau_s^1} \left| (\phi(\tau_s^0) - \phi(\tau_s^1))^\top (\theta_r^* - \theta) \right|^2 + 168B^2 d \log(3TB/\delta)
\end{aligned}$$

where the last step holds by setting $\omega = 1/t$. \square

Lemma B.9 (Reward estimation error). *It holds with probability at least $1 - \delta$ that $\|\eta_{r,t}\|_{\Sigma_{t-1}} \leq \epsilon_{r,\eta}$ for all $t \in [T]$.*

Proof of Lemma B.9. Since $\hat{\theta}_{r,t}$ is the MLE, by Lemma D.10, we have

$$\sum_{s=1}^{t-1} \mathbb{E}_{\tau_s^0, \tau_s^1} d_{\text{TV}}^2 \left(\Pr(\cdot \mid \tau_s^0, \tau_s^1, \theta_r^*), \Pr(\cdot \mid \tau_s^0, \tau_s^1, \hat{\theta}_{r,t}) \right) \leq 10 \log \left(N_{\square}((2T)^{-1}, \overline{\mathcal{R}}, \|\cdot\|_{\infty}) / \delta \right) \quad (5)$$

with probability at least $1 - \delta$, where $\overline{\mathcal{R}}$ denotes the set of probability distributions of preference feedback and $\Pr(\cdot)$ denotes the preference feedback generating probability. Now we upper bound the covering number of $\overline{\mathcal{R}}$ by the covering number of Θ_B . We notice the following:

$$\begin{aligned}
& \sup_{o, \tau^0, \tau^1} \left| \Pr(o \mid \tau^0, \tau^1, \theta_r^*) - \Pr(o \mid \tau^0, \tau^1, \hat{\theta}_{r,t}) \right| \\
& = \sup_{\tau^0, \tau^1} \left| \Phi \left((\phi(\tau^0) - \phi(\tau^1))^\top \theta_r^* \right) - \Phi \left((\phi(\tau^0) - \phi(\tau^1))^\top \hat{\theta}_{r,t} \right) \right| \\
& \leq \bar{\kappa}^{-1} \sup_{\tau^0, \tau^1} \left| (\phi(\tau^0) - \phi(\tau^1))^\top (\theta_r^* - \hat{\theta}_{r,t}) \right| \\
& \leq \bar{\kappa}^{-1} \sup_{\tau^0, \tau^1} \|\phi(\tau^0) - \phi(\tau^1)\|_2 \|\theta_r^* - \hat{\theta}_{r,t}\|_2 \\
& \leq 2\bar{\kappa}^{-1} \|\theta_r^* - \hat{\theta}_{r,t}\|_2
\end{aligned}$$

where the equality holds since the feedback o is binary, the first inequality is Lemma D.1, and the last inequality is by the condition that $\|\phi(\tau)\|_2 \leq 1$ for any trajectory τ . This means that an ω -covering of the space of parameter θ implies a $2\omega\bar{\kappa}^{-1}$ -covering of the space of the probability distribution of the preference feedbacks. Hence, (5) can be further upper bounded by

$$\begin{aligned}
& \sum_{s=1}^{t-1} \mathbb{E}_{\tau_s^0, \tau_s^1} d_{\text{TV}}^2 \left(\Pr(\cdot \mid \tau_s^0, \tau_s^1, \theta_r^*), \Pr(\cdot \mid \tau_s^0, \tau_s^1, \hat{\theta}_{r,t}) \right) \\
& \leq 10 \log \left(N_{\square}((2T)^{-1}, \overline{\mathcal{R}}, \|\cdot\|_{\infty}) / \delta \right) \\
& \leq 10 \log \left(N(\bar{\kappa}/(4T), \Theta_B, \|\cdot\|_2) / \delta \right) \\
& \leq 10d \log \left(12BT / (\bar{\kappa}\delta) \right) \quad (6)
\end{aligned}$$

where the last inequality is Lemma B.2. For the left side, we note that

$$\begin{aligned}
& \sum_{s=1}^{t-1} \mathbb{E}_{\tau_s^0, \tau_s^1} d_{\text{TV}}^2 \left(\Pr(\cdot \mid \tau_s^0, \tau_s^1, \theta_r^*), \Pr(\cdot \mid \tau_s^0, \tau_s^1, \hat{\theta}_{r,t}) \right) \\
&= \sum_{s=1}^{t-1} \mathbb{E}_{\tau_s^0, \tau_s^1} \left| \Phi \left((\phi(\tau_s^0) - \phi(\tau_s^1))^\top \theta_r^* \right) - \Phi \left((\phi(\tau_s^0) - \phi(\tau_s^1))^\top \hat{\theta}_{r,t} \right) \right|^2 \\
&\geq \kappa^{-1} \sum_{s=1}^{t-1} \mathbb{E}_{\tau_s^0, \tau_s^1} \left| (\phi(\tau_s^0) - \phi(\tau_s^1))^\top (\theta_r^* - \hat{\theta}_{r,t}) \right|^2 \\
&\geq \frac{1}{8} \kappa^{-1} \sum_{s=1}^{t-1} \left| (\phi(\tau_s^0) - \phi(\tau_s^1))^\top (\theta_r^* - \hat{\theta}_{r,t}) \right|^2 - 21\kappa^{-1} B^2 d \log(3TB/\delta) \tag{7}
\end{aligned}$$

with probability at least $1 - \delta$, where the first inequality is by Lemma D.1 and the last inequality is by Lemma B.8. Furthermore, we have

$$\begin{aligned}
& \sum_{s=1}^{t-1} \left| (\phi(\tau_s^0) - \phi(\tau_s^1))^\top (\theta_r^* - \hat{\theta}_{r,t}) \right|^2 \\
&= (\theta_r^* - \hat{\theta}_{r,t})^\top \sum_{s=1}^{t-1} (\phi(\tau_s^0) - \phi(\tau_s^1)) (\phi(\tau_s^0) - \phi(\tau_s^1))^\top (\theta_r^* - \hat{\theta}_{r,t}) \\
&= (\theta_r^* - \hat{\theta}_{r,t})^\top \Sigma_{t-1} (\theta_r^* - \hat{\theta}_{r,t}) - (\theta_r^* - \hat{\theta}_{r,t})^\top \lambda I (\theta_r^* - \hat{\theta}_{r,t}) \geq \|\eta_{r,t}\|_{\Sigma_{t-1}}^2 - 4\lambda B^2 \tag{8}
\end{aligned}$$

Putting (6), (7), and (8) together, we get

$$\|\eta_{r,t}\|_{\Sigma_{t-1}}^2 \leq 80\kappa d \log(12BT/(\bar{\kappa}\delta)) + 168B^2 d \log(3TB/\delta) + 4\lambda B^2$$

with probability at least $1 - 2\delta$. Adjusting δ to $\delta/2$ and taking the union bound over $t \in [T]$ yields the desired result. \square

Lemma B.10 (One-step transition estimation error). *Assume the events defined in Lemma B.7 hold. Then, the following claim holds for all $t \in [T]$ and $h \in [H]$ with probability at least $1 - \delta$: if $\max_s |\bar{V}_{t,h+1}(s)| \leq V_{\max}$, it holds that $\|\eta_{P,t,h}\|_{\Sigma_{t-1,h}} \leq \epsilon_{P,\eta}$.*

Proof of Lemma B.10. By definition, we have

$$\|\eta_{P,t,h}\|_{\Sigma_{t-1,h}} = \left\| \sum_{i=1}^{t-1} \phi(s_{i,h}, a_{i,h}) \left(\bar{V}_{t,h+1}(s_{i,h+1}) - \mathbb{E}_{s_{i,h+1}} [\bar{V}_{t,h+1}(s_{i,h+1}) \mid s_{i,h}, a_{i,h}] \right) \right\|_{\Sigma_{t-1,h}^{-1}}.$$

Towards an upper bound of the above, we will conduct some covering arguments. To that end, we first derive the upper bounds of norms of $\bar{\theta}_{r,t}$ and $\bar{\theta}_{P,t,h}$.

Bounding the ℓ_2 -norms of $\bar{\theta}_{r,t}$ and $\bar{\theta}_{P,t,h}$. For $\bar{\theta}_{r,t}$, applying triangle inequality, we have $\|\bar{\theta}_{r,t}\|_2 \leq \|\xi_{r,t}\|_2 + \|\hat{\theta}_{r,t}\|_2 \leq \|\xi_{r,t}\|_2 + B$. For $\|\xi_{r,t}\|_2$, we have

$$\|\xi_{r,t}\|_2 = \sqrt{\xi_{r,t}^\top \Sigma_{t-1}^{-1/2} \Sigma_{t-1}^{1/2} \xi_{r,t}} \leq \sqrt{\|\xi_{r,t}\|_{\Sigma_{t-1}} \|\xi_{r,t}\|_{\Sigma_{t-1}^{-1}}} \leq \sqrt{\epsilon_{r,\xi} \cdot \|\xi_{r,t}\|_2 / \sqrt{\lambda}}$$

where in the last step we applied Lemma B.7 and the fact that $\|\xi_{r,t}\|_{\Sigma_{t-1}^{-1}} \leq \|\xi_{r,t}\|_2 \|\Sigma_{t-1}^{-1/2}\|_2 \leq \|\xi_{r,t}\|_2 / \sqrt{\lambda}$. This implies that $\|\xi_{r,t}\|_2 \leq \epsilon_{r,\xi} / \sqrt{\lambda}$, and thus, $\|\bar{\theta}_{r,t}\|_2 \leq \epsilon_{r,\xi} / \sqrt{\lambda} + B$.

For $\bar{\theta}_{P,t,h}$, we have $\|\bar{\theta}_{P,t,h}\|_2 \leq \|\xi_{P,t,h}\|_2 + \|\hat{\theta}_{P,t,h}\|_2$. For $\|\xi_{P,t,h}\|_2$, following a similar argument as above, we have

$$\|\xi_{P,t,h}\|_2 \leq \sqrt{\|\xi_{P,t,h}\|_{\Sigma_{t-1,h}} \|\xi_{P,t,h}\|_{\Sigma_{t-1,h}^{-1}}} \leq \sqrt{\epsilon_{P,\xi} \cdot \|\xi_{P,t,h}\|_2 / \sqrt{\lambda}},$$

which implies that $\|\xi_{P,t,h}\|_2 \leq \epsilon_{P,\xi}/\sqrt{\lambda}$. For $\|\widehat{\theta}_{P,t,h}\|_2$, we have

$$\begin{aligned} \|\widehat{\theta}_{P,t,h}\|_2 &= \left\| \Sigma_{t-1,h}^{-1/2} \Sigma_{t-1,h}^{-1/2} \left(\sum_{i=1}^{t-1} \phi(s_{i,h}, a_{i,h}) \bar{V}_{t,h+1}(s_{i,h+1}) \right) \right\|_2 \\ &\leq \|\Sigma_{t-1,h}^{-1/2}\|_2 \cdot \left\| \sum_{i=1}^{t-1} \phi(s_{i,h}, a_{i,h}) \bar{V}_{t,h+1}(s_{i,h+1}) \right\|_{\Sigma_{t-1,h}^{-1}} \\ &\leq \sqrt{\lambda^{-1}} \cdot \sqrt{t} \cdot \sqrt{\sum_{i=1}^{t-1} \|\phi(s_{i,h}, a_{i,h}) \bar{V}_{t,h+1}(s_{i,h+1})\|_{\Sigma_{t-1,h}^{-1}}^2} \\ &\leq \sqrt{\lambda^{-1}} \cdot \sqrt{t} \cdot V_{\max} \cdot \sqrt{\sum_{i=1}^{t-1} \|\phi(s_{i,h}, a_{i,h})\|_{\Sigma_{t-1,h}^{-1}}^2} \\ &\leq \frac{V_{\max} \sqrt{dt}}{\sqrt{\lambda}} \end{aligned}$$

where the second inequality is by the Jensen's inequality, the third inequality is by the condition that $\max_s \|\bar{V}_{t,h}(s)\| \leq V_{\max}$, and the last inequality is by Lemma D.7. Putting these together, we get $\|\widehat{\theta}_{P,t,h}\|_2 \leq \epsilon_{P,\xi}/\sqrt{\lambda} + V_{\max} \sqrt{dt}/\sqrt{\lambda}$.

Covering construction. With these bounds in hand, we can now proceed with a covering construction. First define

$$\begin{aligned} Q_{\theta_r, \theta_P, \Sigma}(s, a) &:= \phi(s, a)^\top \theta_r + \begin{cases} \phi(s, a)^\top \theta_P & \text{if } \|\phi(s, a)\|_\Sigma \leq \alpha_L \\ \rho(s, a) \left(\phi(s, a)^\top \theta_P \right) + (1 - \rho(s, a))(H - h) & \text{if } \alpha_L < \|\phi(s, a)\|_\Sigma \leq \alpha_U, \\ H - h & \text{if } \|\phi(s, a)\|_\Sigma > \alpha_U \end{cases} \\ V_{\theta_r, \theta_P, \Sigma}(s) &:= \max_a Q_{\theta_r, \theta_P, \Sigma}(s, a). \end{aligned} \quad (9)$$

where $\rho(s, a) := \frac{\alpha_U - \|\phi(s, a)\|_\Sigma}{\alpha_U - \alpha_L}$. This definition aims to mimic the behavior of Algorithm 1. Then, we define the space of parameters as follows

$$\mathcal{U} := \left\{ (\theta_r, \theta_P, \Sigma) : \|\theta_r\|_2 \leq B + \epsilon_{r,\xi}/\sqrt{\lambda}, \|\theta_P\|_2 \leq (V_{\max} \sqrt{dt} + \epsilon_{P,\xi})/\sqrt{\lambda}, \|\Sigma\|_F \leq \sqrt{d}/\lambda, \Sigma = \Sigma^\top, \Sigma \succeq 0, \|V_{\theta_r, \theta_P, \Sigma}\|_\infty \leq V_{\max} \right\}.$$

Clearly, $\bar{\theta}_{r,t}$ and $\bar{\theta}_{P,t,h}$ satisfy the constraints of θ_r and θ_P . Moreover, $\|\Sigma_{t-1}^{-1}\|_F \leq \sqrt{d} \|\Sigma_{t-1}^{-1}\|_2 \leq \sqrt{d}/\lambda$, and thus Σ_{t-1}^{-1} satisfies the constraint of Σ . Hence, we have $(\bar{\theta}_{r,t}, \bar{\theta}_{P,t,h}, \Sigma_{t-1}) \in \mathcal{U}$.

By Lemma B.2, the size of an ω -covering of \mathcal{U} by treating it as a subset of \mathbb{R}^{2d+d^2} can be bounded by

$$N(\omega, \mathcal{U}, \|\cdot\|_2) \leq \left(\frac{3 \left((B + \epsilon_{r,\xi}/\sqrt{\lambda}) + (V_{\max} \sqrt{dt} + \epsilon_{P,\xi})/\sqrt{\lambda} + \sqrt{d}/\lambda \right)}{\omega} \right)^{2d+d^2}$$

We denote the covering set by $\tilde{\mathcal{U}}$. We will write $N := N(\omega, \mathcal{U}, \|\cdot\|_2)$ for simplicity when there is no ambiguity. We note that it is possible that not all points in $\tilde{\mathcal{U}}$ belong to \mathcal{U} . However, we can project all points to \mathcal{U} so it becomes a 2ω -covering of \mathcal{U} . We will omit this subtlety.

Covering argument. For any $(\theta_r, \theta_P, \Sigma) \in \mathcal{U}$, we define

$$x_{i,\theta_r, \theta_P, \Sigma} := V_{\theta_r, \theta_P, \Sigma}(s_{i,h+1}) - \mathbb{E}_{s'} [V_{\theta_r, \theta_P, \Sigma}(s') \mid s_{i,h}, a_{i,h}].$$

By construction, we have $|x_{i,\theta_r,\theta_P,\Sigma}| \leq 2V_{\max}$, so it is $4V_{\max}$ -subgaussian conditioning on $(s_{i,h}, a_{i,h})$. Hence, for all $(\theta_r, \theta_P, \Sigma) \in \tilde{\mathcal{U}}$, we have

$$\left\| \sum_{i=1}^{t-1} \phi(s_{i,h}, a_{i,h}) x_{i,\theta_r,\theta_P,\Sigma} \right\|_{\Sigma_{t-1,h}^{-1}} \leq \sqrt{32V_{\max}^2 \left(d \log \left(1 + \frac{t}{\lambda} \right) + \log(N/\delta) \right)} \quad (10)$$

with probability at least $1 - \delta$ by Lemma D.5 and the union bound over $\tilde{\mathcal{U}}$. Now consider an arbitrary tuple $(\theta_r, \theta_P, \Sigma) \in \mathcal{U}$ and the nearest point $(\tilde{\theta}_r, \tilde{\theta}_P, \tilde{\Sigma}) \in \tilde{\mathcal{U}}$. Then, we have

$$\begin{aligned} & \left\| \sum_{i=1}^{t-1} \phi(s_{i,h}, a_{i,h}) x_{i,\theta_r,\theta_P,\Sigma} \right\|_{\Sigma_{t-1,h}^{-1}} \\ & \leq \left\| \sum_{i=1}^{t-1} \phi(s_{i,h}, a_{i,h}) x_{i,\tilde{\theta}_r,\tilde{\theta}_P,\tilde{\Sigma}} \right\|_{\Sigma_{t-1,h}^{-1}} + \left\| \sum_{i=1}^{t-1} \phi(s_{i,h}, a_{i,h}) (x_{i,\theta_r,\theta_P,\Sigma} - x_{i,\tilde{\theta}_r,\tilde{\theta}_P,\tilde{\Sigma}}) \right\|_{\Sigma_{t-1,h}^{-1}} \\ & \leq \left\| \sum_{i=1}^{t-1} \phi(s_{i,h}, a_{i,h}) x_{i,\tilde{\theta}_r,\tilde{\theta}_P,\tilde{\Sigma}} \right\|_{\Sigma_{t-1,h}^{-1}} + \sum_{i=1}^{t-1} \|\phi(s_{i,h}, a_{i,h})\|_{\Sigma_{t-1,h}^{-1}} \cdot \max_i |x_{i,\theta_r,\theta_P,\Sigma} - x_{i,\tilde{\theta}_r,\tilde{\theta}_P,\tilde{\Sigma}}| \\ & \leq \left\| \sum_{i=1}^{t-1} \phi(s_{i,h}, a_{i,h}) x_{i,\tilde{\theta}_r,\tilde{\theta}_P,\tilde{\Sigma}} \right\|_{\Sigma_{t-1,h}^{-1}} + \sqrt{\lambda^{-1}} \cdot t \cdot \max_i |x_{i,\theta_r,\theta_P,\Sigma} - x_{i,\tilde{\theta}_r,\tilde{\theta}_P,\tilde{\Sigma}}| \end{aligned}$$

where the last step is by the fact that $\|\phi(s_{i,h}, a_{i,h})\|_{\Sigma_{t-1,h}^{-1}} \leq \sqrt{\lambda^{-1}}$. Observing the derived upper bound above, the first term is already bounded by (10) since $(\tilde{\theta}_r, \tilde{\theta}_P, \tilde{\Sigma}) \in \tilde{\mathcal{U}}$. To bound the second term, we note that, by the definition of $x_{i,\theta_r,\theta_P,\Sigma}$ and $x_{i,\tilde{\theta}_r,\tilde{\theta}_P,\tilde{\Sigma}}$,

$$\begin{aligned} |x_{i,\theta_r,\theta_P,\Sigma} - x_{i,\tilde{\theta}_r,\tilde{\theta}_P,\tilde{\Sigma}}| & \leq 2 \max_s \left| V_{\theta_r,\theta_P,\Sigma}(s) - V_{\tilde{\theta}_r,\tilde{\theta}_P,\tilde{\Sigma}}(s) \right| \\ & = 2 \max_s \left| \max_a Q_{\theta_r,\theta_P,\Sigma}(s, a) - \max_a Q_{\tilde{\theta}_r,\tilde{\theta}_P,\tilde{\Sigma}}(s, a) \right| \\ & \leq 2 \max_{s,a} \left| Q_{\theta_r,\theta_P,\Sigma}(s, a) - Q_{\tilde{\theta}_r,\tilde{\theta}_P,\tilde{\Sigma}}(s, a) \right| \end{aligned}$$

We assume $\omega \leq (\alpha_U - \alpha_L)^2$ (and this will be satisfied later when we specify the value of ω). Recall that we have three cases in the definition of $Q_{\theta_r,\theta_P,\Sigma}(s, a)$ (see (9)), and we will refer to them as Case L, Case M, and Case U. There are in total 6 possible combinations of cases for the pair $Q_{\theta_r,\theta_P,\Sigma}(s, a)$, $Q_{\tilde{\theta}_r,\tilde{\theta}_P,\tilde{\Sigma}}(s, a)$, and we will discuss them one by one below. For the ease of notation, we denote $Q(s, a)$ and $\tilde{Q}(s, a)$ as shorthand for $Q_{\theta_r,\theta_P,\Sigma}(s, a)$ and $Q_{\tilde{\theta}_r,\tilde{\theta}_P,\tilde{\Sigma}}(s, a)$.

(1) Case L + Case U. This is impossible, because, in this case, we have $|\|\phi\|_{\Sigma} - \|\phi\|_{\tilde{\Sigma}}| > \alpha_U - \alpha_L$, but we also have $|\|\phi\|_{\Sigma} - \|\phi\|_{\tilde{\Sigma}}| = |\sqrt{\phi^\top \Sigma \phi} - \sqrt{\phi^\top \tilde{\Sigma} \phi}| \leq \sqrt{|\phi^\top (\Sigma - \tilde{\Sigma}) \phi|} \leq \sqrt{\|\phi\|_2 \|\Sigma - \tilde{\Sigma}\|_F \|\phi\|_2} \leq \sqrt{\omega} < \alpha_U - \alpha_L$.

(2) Both are Case L. Then we immediately have

$$|Q(s, a) - \tilde{Q}(s, a)| = \left| \phi(s, a)^\top (\theta_r + \theta_P - \tilde{\theta}_r - \tilde{\theta}_P) \right| \leq \|\phi(s, a)\|_2 \left(\|\theta_r - \tilde{\theta}_r\|_2 + \|\theta_P - \tilde{\theta}_P\|_2 \right) \leq 2\omega.$$

(3) Both are Case U. Then we immediately have

$$|Q(s, a) - \tilde{Q}(s, a)| = \left| \phi(s, a)^\top (\theta_r - \tilde{\theta}_r) \right| \leq \|\phi(s, a)\|_2 \cdot \|\theta_r - \tilde{\theta}_r\|_2 \leq \omega.$$

(4) Case L + Case M. Without loss of generality, we assume $Q(s, a)$ is in Case M and $\tilde{Q}(s, a)$ is in Case L. Then we have

$$\begin{aligned} |Q(s, a) - \tilde{Q}(s, a)| &\leq \rho(s, a) \left| \phi(s, a)^\top (\theta_r + \theta_P - \tilde{\theta}_r - \tilde{\theta}_P) \right| \\ &\quad + (1 - \rho(s, a)) \left| \phi(s, a)^\top (\tilde{\theta}_r + \tilde{\theta}_P) - (H - h) \right| \\ &\leq \rho(s, a) \cdot 2\omega + (1 - \rho(s, a)) \cdot 2V_{\max} \end{aligned}$$

Recall that $\rho(s, a) := \frac{\alpha_U - \|\phi(s, a)\|_\Sigma}{\alpha_U - \alpha_L}$, so $1 - \rho(s, a) := \frac{\|\phi(s, a)\|_\Sigma - \alpha_L}{\alpha_U - \alpha_L}$. Furthermore,

$$\begin{aligned} 1 - \rho(s, a) &= \frac{\|\phi(s, a)\|_\Sigma - \alpha_L}{\alpha_U - \alpha_L} = \frac{\|\phi(s, a)\|_\Sigma - \|\phi(s, a)\|_{\tilde{\Sigma}} + \|\phi(s, a)\|_{\tilde{\Sigma}} - \alpha_L}{\alpha_U - \alpha_L} \\ &\leq \frac{\|\phi(s, a)\|_\Sigma - \|\phi(s, a)\|_{\tilde{\Sigma}}}{\alpha_U - \alpha_L} = \frac{\sqrt{\phi^\top \Sigma \phi} - \sqrt{\phi^\top \tilde{\Sigma} \phi}}{\alpha_U - \alpha_L} \\ &\leq \frac{\sqrt{|\phi^\top (\Sigma - \tilde{\Sigma}) \phi|}}{\alpha_U - \alpha_L} \leq \frac{\sqrt{\|\phi\|_2 \|\Sigma - \tilde{\Sigma}\|_2 \|\phi\|_2}}{\alpha_U - \alpha_L} \leq \frac{\sqrt{\omega}}{\alpha_U - \alpha_L} \end{aligned}$$

where the first inequality is by the condition that \tilde{Q} is in Case L. Inserting this back, we get $|Q(s, a) - \tilde{Q}(s, a)| \leq 2\omega + \frac{2V_{\max}\sqrt{\omega}}{\alpha_U - \alpha_L}$.

(5) Case M + Case U. Without loss of generality, we assume $Q(s, a)$ is in Case M and $\tilde{Q}(s, a)$ is in Case U. Then we have

$$\begin{aligned} |Q(s, a) - \tilde{Q}(s, a)| &\leq \rho(s, a) \left| \phi(s, a)^\top (\theta_r + \theta_P - (H - h)) \right| + (1 - \rho(s, a)) |(H - h) - (H - h)| \\ &= \rho(s, a) \left| \phi(s, a)^\top (\theta_r + \theta_P - (H - h)) \right| \leq \rho(s, a) \cdot 2V_{\max} \leq \frac{2V_{\max}\sqrt{\omega}}{\alpha_U - \alpha_L} \end{aligned}$$

where the last inequality is from

$$\begin{aligned} \rho(s, a) &= \frac{\alpha_U - \|\phi(s, a)\|_\Sigma}{\alpha_U - \alpha_L} = \frac{\alpha_U - \|\phi(s, a)\|_\Sigma - \|\phi(s, a)\|_{\tilde{\Sigma}} + \|\phi(s, a)\|_{\tilde{\Sigma}}}{\alpha_U - \alpha_L} \\ &\leq \frac{\|\phi(s, a)\|_{\tilde{\Sigma}} - \|\phi(s, a)\|_\Sigma}{\alpha_U - \alpha_L} = \frac{\sqrt{\phi^\top \tilde{\Sigma} \phi} - \sqrt{\phi^\top \Sigma \phi}}{\alpha_U - \alpha_L} \\ &\leq \frac{\sqrt{|\phi^\top (\Sigma - \tilde{\Sigma}) \phi|}}{\alpha_U - \alpha_L} \leq \frac{\sqrt{\|\phi\|_2 \|\Sigma - \tilde{\Sigma}\|_2 \|\phi\|_2}}{\alpha_U - \alpha_L} \leq \frac{\sqrt{\omega}}{\alpha_U - \alpha_L}. \end{aligned}$$

where the first inequality is by the condition that \tilde{Q} is in Case U.

(6) Both are Case M. Denote $\rho(s, a) = \frac{\alpha_U - \|\phi(s, a)\|_\Sigma}{\alpha_U - \alpha_L}$ and $\tilde{\rho}(s, a) = \frac{\alpha_U - \|\phi(s, a)\|_{\tilde{\Sigma}}}{\alpha_U - \alpha_L}$. Then, we have

$$|\rho(s, a) - \tilde{\rho}(s, a)| = \frac{\left| \|\phi(s, a)\|_{\tilde{\Sigma}} - \|\phi(s, a)\|_\Sigma \right|}{\alpha_U - \alpha_L} \leq \frac{\sqrt{\|\phi\|_2 \|\tilde{\Sigma} - \Sigma\|_F \|\phi\|_2}}{\alpha_U - \alpha_L} \leq \frac{\sqrt{\omega}}{\alpha_U - \alpha_L},$$

which also implies $|(1 - \rho(s, a)) - (1 - \tilde{\rho}(s, a))| \leq \frac{\sqrt{\omega}}{\alpha_U - \alpha_L}$. Hence, we have

$$\begin{aligned} |Q(s, a) - \tilde{Q}(s, a)| &\leq \left| \rho(s, a) \phi(s, a)^\top (\theta_r + \theta_P) - \tilde{\rho}(s, a) \phi(s, a)^\top (\tilde{\theta}_r + \tilde{\theta}_P) \right| \\ &\quad + |(1 - \rho(s, a))(H - h) - (1 - \tilde{\rho}(s, a))(H - h)| \\ &\leq \left| \rho(s, a) \phi(s, a)^\top (\theta_r + \theta_P) - \rho(s, a) \phi(s, a)^\top (\tilde{\theta}_r + \tilde{\theta}_P) \right| \\ &\quad + \left| \rho(s, a) \phi(s, a)^\top (\tilde{\theta}_r + \tilde{\theta}_P) - \tilde{\rho}(s, a) \phi(s, a)^\top (\tilde{\theta}_r + \tilde{\theta}_P) \right| \\ &\quad + |(1 - \rho(s, a))(H - h) - (1 - \tilde{\rho}(s, a))(H - h)| \\ &\leq 2\omega + \frac{V_{\max}\sqrt{\omega}}{\alpha_U - \alpha_L} + \frac{H\sqrt{\omega}}{\alpha_U - \alpha_L} \end{aligned}$$

Putting Everything Together. Taking the maximum of the six cases, we conclude that

$$|Q(s, a) - \tilde{Q}(s, a)| \leq 2\omega + \frac{2V_{\max}\sqrt{\omega}}{\alpha_U - \alpha_L}$$

We set $\omega = \left(\frac{\alpha_U - \alpha_L}{2tV_{\max}}\right)^2$ and obtain

$$|Q(s, a) - \tilde{Q}(s, a)| \leq 2 \left(\frac{\alpha_U - \alpha_L}{2tV_{\max}}\right)^2 + \frac{1}{t} \leq \frac{3}{t}$$

where we leverage the condition that $\alpha_L, \alpha_U \leq 1$. Thus, we have

$$|x_{i, \theta_r, \theta_P, \Sigma} - x_{i, \tilde{\theta}_r, \tilde{\theta}_P, \tilde{\Sigma}}| \leq 2 \max_{s, a} |Q(s, a) - \tilde{Q}(s, a)| \leq \frac{6}{t}.$$

Final Bound. The covering argument is now complete, and we can use it to derive an upper bound for $\|\eta_{P, t, h}\|_{\Sigma_{t-1, h}}$. To that end, we note that, for $\eta_{P, t, h}$, there must exist $(\theta_r, \theta_P, \Sigma) \in \mathcal{U}$ and its closest element $(\tilde{\theta}_r, \tilde{\theta}_P, \tilde{\Sigma}) \in \tilde{\mathcal{U}}$ such that

$$\begin{aligned} & \|\eta_{P, t, h}\|_{\Sigma_{t-1, h}} \\ &= \left\| \sum_{i=1}^{t-1} \phi(s_{i, h}, a_{i, h}) \left(\bar{V}_{t, h+1}(s_{i, h+1}) - \mathbb{E}_{s_{i, h+1}} [\bar{V}_{t, h+1}(s_{i, h+1}) | s_{i, h}, a_{i, h}] \right) \right\|_{\Sigma_{t-1, h}^{-1}} \\ &= \left\| \sum_{i=1}^{t-1} \phi(s_{i, h}, a_{i, h}) x_{i, \theta_r, \theta_P, \Sigma} \right\|_{\Sigma_{t-1, h}^{-1}} \\ &= \left\| \sum_{i=1}^{t-1} \phi(s_{i, h}, a_{i, h}) x_{i, \tilde{\theta}_r, \tilde{\theta}_P, \tilde{\Sigma}} \right\|_{\Sigma_{t-1, h}^{-1}} + \left\| \sum_{i=1}^{t-1} \phi(s_{i, h}, a_{i, h}) (x_{i, \theta_r, \theta_P, \Sigma} - x_{i, \tilde{\theta}_r, \tilde{\theta}_P, \tilde{\Sigma}}) \right\|_{\Sigma_{t-1, h}^{-1}} \end{aligned}$$

where the second term is bounded by

$$\begin{aligned} & \left\| \sum_{i=1}^{t-1} \phi(s_{i, h}, a_{i, h}) (x_{i, \theta_r, \theta_P, \Sigma} - x_{i, \tilde{\theta}_r, \tilde{\theta}_P, \tilde{\Sigma}}) \right\|_{\Sigma_{t-1, h}^{-1}} \\ & \leq \sum_{i=1}^{t-1} \left\| \phi(s_{i, h}, a_{i, h}) (x_{i, \theta_r, \theta_P, \Sigma} - x_{i, \tilde{\theta}_r, \tilde{\theta}_P, \tilde{\Sigma}}) \right\|_{\Sigma_{t-1, h}^{-1}} \\ & \leq \sqrt{\lambda^{-1}} \cdot t \cdot \max_i |x_{i, \theta_r, \theta_P, \Sigma} - x_{i, \tilde{\theta}_r, \tilde{\theta}_P, \tilde{\Sigma}}| \\ & \leq 6/\sqrt{\lambda}, \end{aligned}$$

and, applying (10), the first term is bounded by

$$\begin{aligned} & \left\| \sum_{i=1}^{t-1} \phi(s_{i, h}, a_{i, h}) x_{i, \tilde{\theta}_r, \tilde{\theta}_P, \tilde{\Sigma}} \right\|_{\Sigma_{t-1, h}^{-1}} \\ & \leq \sqrt{32V_{\max}^2 \left(d \log \left(1 + \frac{t}{\lambda} \right) + \log(N/\delta) \right)} \\ & \leq \sqrt{32V_{\max}^2 \left(d \log \left(1 + \frac{t}{\lambda} \right) + (2d + d^2) \log \left(\frac{3 \left((B + \epsilon_{r, \xi}/\sqrt{\lambda}) + (V_{\max}\sqrt{dt} + \epsilon_{P, \xi})/\sqrt{\lambda} + \sqrt{d}/\lambda \right)}{\omega \delta} \right) \right)} \\ & = V_{\max} \sqrt{32d \log \left(1 + \frac{t}{\lambda} \right) + 32(2d + d^2) \log \left(\frac{12V_{\max}^2 t^2 \left((B + \epsilon_{r, \xi}/\sqrt{\lambda}) + (V_{\max}\sqrt{dt} + \epsilon_{P, \xi})/\sqrt{\lambda} + \sqrt{d}/\lambda \right)}{(\alpha_U - \alpha_L)^2 \delta} \right)} \end{aligned}$$

where the second inequality is by inserting the upper bound of the covering number N , and the equality is by inserting the value of ω . The last line looks complicated and we can further simplify it and get

$$\begin{aligned} & \left\| \sum_{i=1}^{t-1} \phi(s_{i,h}, a_{i,h}) x_{i, \bar{\theta}_t, \bar{\theta}_P, \Sigma} \right\|_{\Sigma_{t-1,h}^{-1}} \\ & \leq 16dV_{\max} \sqrt{\log \left(\frac{12V_{\max}t(t+\lambda) \left((B + \epsilon_{r,\xi}/\sqrt{\lambda}) + (V_{\max}\sqrt{dt} + \epsilon_{P,\xi})/\sqrt{\lambda} + \sqrt{d}/\lambda \right)}{(\alpha_U - \alpha_L)\delta\lambda} \right)}. \end{aligned}$$

Plugging these upper bounds back, we obtain

$$\begin{aligned} & \|\eta_{P,t,h}\|_{\Sigma_{t-1,h}} \\ & \leq \frac{6}{\sqrt{\lambda}} + 16dV_{\max} \sqrt{\log \left(\frac{12V_{\max}t(t+\lambda) \left((B + \epsilon_{r,\xi}/\sqrt{\lambda}) + (V_{\max}\sqrt{dt} + \epsilon_{P,\xi})/\sqrt{\lambda} + \sqrt{d}/\lambda \right)}{(\alpha_U - \alpha_L)\delta\lambda} \right)}. \end{aligned}$$

Applying union bound over all $t \in [T]$ and $h \in [H]$, the upper bound exactly becomes $\epsilon'_{P,\eta}$. Further invoking Lemma B.3 finishes the proof. \square

Lemma B.11 (Boundness of value functions and transition estimation Error). *Assume the events define in Lemmas B.7 and B.9 hold. Then, the following holds with probability at least $1 - \delta$:*

1. $\max_{s,a} |\bar{Q}_{t,h}(s,a)| \leq V_{\max}$ for all $t \in [T]$ and $h \in [H]$.
2. $\|\eta_{P,t,h}\|_{\Sigma_{t-1,h}} \leq \epsilon_{P,\eta}$ for all $t \in [T]$ and $h \in [H]$.
3. $\|\lambda_{t,h}\|_{\Sigma_{t-1,h}} \leq \epsilon_\lambda$ for all $t \in [T]$ and $h \in [H]$.

The first statement implies $\max_s |\bar{V}_{t,h}(s)| \leq V_{\max}$ for all $t \in [T]$ and $h \in [H]$.

Proof. We prove the three statements together by induction.

For the first statement, we define $V_{\max,h} := (H - h + 1)(1 + (\epsilon_{r,\xi} + \epsilon_{r,\eta})/\sqrt{\lambda})$, and we will actually show that $\max_{s,a} |\bar{Q}_{t,h}(s,a) - Q_h^*(s,a)| \leq V_{\max,h}$ in the induction, which immediately leads to $\max_{s,a} |\bar{Q}_{t,h}(s,a)| \leq V_{\max,h} + (H - h + 1) \leq V_{\max}$.

Base. For $h = H$, we can directly apply Lemma B.10 without the transition argument, which leads to the desired upper bound of $\|\eta_{P,t,H}\|_{\Sigma_{t-1,H}}$. Moreover, the upper bound of $\|\lambda_{t,h}\|_{\Sigma_{t-1,h}}$ also trivially holds. For an upper bound on the value function, we immediately have

$$|\bar{Q}_{t,H}(s,a) - Q_H^*(s,a)| = |\phi(s,a)^\top (\bar{\theta}_{r,t} - \theta_r^*)| \leq \|\phi(s,a)\|_{\Sigma_{t-1}^{-1}} \|\bar{\theta}_{r,t} - \theta_r^*\|_{\Sigma_{t-1}}$$

by Cauchy-Schwarz inequality. We note that $\|\phi(s,a)\|_{\Sigma_{t-1}^{-1}} \leq \sqrt{\lambda^{-1}}$ and $\|\bar{\theta}_{r,t} - \theta_r^*\|_{\Sigma_{t-1}} \leq \|\xi_{r,t} + \eta_{r,t}\|_{\Sigma_{t-1}} \leq \|\xi_{r,t}\|_{\Sigma_{t-1}} + \|\eta_{r,t}\|_{\Sigma_{t-1}} \leq \epsilon_{r,\xi} + \epsilon_{r,\eta}$. Hence, we have

$$|\bar{Q}_{t,H}(s,a) - Q_H^*(s,a)| \leq (\epsilon_{r,\xi} + \epsilon_{r,\eta})/\sqrt{\lambda} \leq V_{\max,H}.$$

Inductive Steps. Now consider $h < H$. By induction hypothesis, we have $\max_{t,s} |\bar{V}_{t,h+1}(s)| \leq V_{\max}$. Thus, we can apply Lemma B.10 and get $\|\eta_{P,t,h}\|_{\Sigma_{t-1,h}} \leq \epsilon_{P,\eta}$. For the upper bound of $\|\lambda_{t,h}\|_{\Sigma_{t-1,h}}$, by definition, we have

$$\begin{aligned} \|\lambda_{t,h}\|_{\Sigma_{t-1,h}} &= \left\| \lambda \int_{s'} \mu_h^*(s') \bar{V}_{t,h+1}(s') ds' \right\|_{\Sigma_{t-1,h}^{-1}} \\ &\leq \sqrt{\lambda} \left\| \int_{s'} \mu_h^*(s') \bar{V}_{t,h+1}(s') ds' \right\|_2 \\ &\leq V_{\max} \sqrt{\lambda d} = \epsilon_\lambda \end{aligned}$$

where the second inequality is by $\|\Sigma_{t-1,h}^{-1/2}\|_2 \leq \sqrt{\lambda}$, and the last inequality is by Assumption 4.1.

For the upper bound on the value function, consider the following three cases.

Case 1: $\|\phi(s, a)\|_{\Sigma_{t-1,h}^{-1}} \leq \alpha_L$. We apply Lemma B.5 and get

$$\begin{aligned}
& |\phi(s, a)^\top (\bar{\theta}_{r,t} + \bar{\theta}_{P,t,h}) - Q_h^*(s, a)| \\
&= \left| \phi(s, a)^\top (\xi_{r,t} + \xi_{P,t,h} + \eta_{r,t} + \eta_{P,t,h} - \lambda_{t,h}) + \mathbb{E}_{s'} [\bar{V}_{t,h+1}(s') - V_{h+1}^*(s') \mid s, a] \right| \\
&\leq \left| \phi(s, a)^\top (\xi_{r,t} + \eta_{r,t}) \right| + \left| \phi(s, a)^\top (\xi_{P,t,h} + \eta_{P,t,h} - \lambda_{t,h}) \right| + \left| \mathbb{E}_{s'} [\bar{V}_{t,h+1}(s') - V_{h+1}^*(s') \mid s, a] \right| \\
&\leq (\epsilon_{r,\xi} + \epsilon_{r,\eta})/\sqrt{\lambda} + \|\phi(s, a)\|_{\Sigma_{t-1,h}^{-1}} (\epsilon_{P,\xi} + \epsilon_{P,\eta} + \epsilon_\lambda) + \left((H-h)(1 + (\epsilon_{r,\xi} + \epsilon_{r,\eta})/\sqrt{\lambda}) \right) \\
&\leq (\epsilon_{r,\xi} + \epsilon_{r,\eta})/\sqrt{\lambda} + 1 + \left((H-h)(1 + (\epsilon_{r,\xi} + \epsilon_{r,\eta})/\sqrt{\lambda}) \right) \\
&= (H-h+1)(1 + (\epsilon_{r,\xi} + \epsilon_{r,\eta})/\sqrt{\lambda}) \\
&= V_{\max,h}
\end{aligned}$$

where the second inequality is by Cauchy-Schwarz inequality, Lemmas B.7 and B.9, and the induction hypothesis. The third inequality is by the condition that $\|\phi(s, a)\|_{\Sigma_{t-1,h}^{-1}} \leq \alpha_L$ and the definition of α_L .

Case 2: $\|\phi(s, a)\|_{\Sigma_{t-1,h}^{-1}} > \alpha_U$. We have

$$\begin{aligned}
|\bar{Q}_{t,h}(s, a) - Q_h^*(s, a)| &\leq |\phi(s, a)^\top \bar{\theta}_{r,t} + H - h - Q_h^*(s, a)| \\
&= |\phi(s, a)^\top (\xi_{r,t} + \eta_{r,t}) + H - h + \phi(s, a)^\top \theta_r^* - Q_h^*(s, a)| \\
&\leq (\epsilon_{r,\xi} + \epsilon_{r,\eta})/\sqrt{\lambda} + (H - h + 1) \\
&\leq V_{\max,h}
\end{aligned}$$

where we used the fact that $0 \leq \phi(s, a)^\top \theta_r^* \leq 1$ and $0 \leq Q_h^*(s, a) \leq H - h + 1$, which leads to $-(H - h + 1) \leq \phi(s, a)^\top \theta_r^* - Q_h^*(s, a) \leq 1$.

Case 3: $\alpha_L < \|\phi(s, a)\|_{\Sigma_{t-1,h}^{-1}} \leq \alpha_U$. Denoting $\rho := \rho(s, a)$ for simplicity, we have

$$\begin{aligned}
& |\bar{Q}_{t,h}(s, a) - Q_h^*(s, a)| \\
&\leq \rho \left| \phi(s, a)^\top (\bar{\theta}_{r,t} + \bar{\theta}_{P,t,h}) - Q_h^*(s, a) \right| + (1 - \rho) \left| \phi(s, a)^\top \bar{\theta}_{r,t} + H - h - Q_h^*(s, a) \right| \\
&\leq \rho(H - h + 1) \left(1 + (\epsilon_{r,\xi} + \epsilon_{r,\eta})/\sqrt{\lambda} \right) + (1 - \rho) \left| \phi(s, a)^\top (\xi_{r,t} + \eta_{r,t}) + H - h + \phi(s, a)^\top \theta_r^* - Q_h^*(s, a) \right| \\
&\leq \rho(H - h + 1) \left(1 + (\epsilon_{r,\xi} + \epsilon_{r,\eta})/\sqrt{\lambda} \right) + (1 - \rho) \left((\epsilon_{r,\xi} + \epsilon_{r,\eta})/\sqrt{\lambda} + (H - h + 1) \right) \\
&\leq V_{\max,h}
\end{aligned}$$

where we have used the similar arguments from Case 1 and 2.

Taking the maximum over the three cases, we conclude that

$$|\bar{Q}_{t,h}(s, a) - Q_h^*(s, a)| \leq V_{\max,h}$$

which implies $\max_{s,a} |\bar{Q}_{t,h}(s, a)| \leq V_{\max}$. \square

Lemma B.12 (High probability bounds – summary). *All of the following events hold simultaneously with probability at least $1 - \delta$:*

1. $\|\xi_{r,t}\|_{\Sigma_{t-1}} \leq \epsilon_{r,\xi}$ for all $t \in [T]$
2. $\|\xi_{P,t,h}\|_{\Sigma_{t-1,h}} \leq \epsilon_{P,\xi}$ for all $t \in [T]$ and $h \in [H]$.

3. $\|\eta_{r,t}\|_{\Sigma_{t-1}} \leq \epsilon_{r,\eta}$ for all $t \in [T]$.
4. $\|\eta_{P,t,h}\|_{\Sigma_{t-1,h}} \leq \epsilon_{P,\eta}$ for all $t \in [T]$ and $h \in [H]$.
5. $\|\lambda_{t,h}\|_{\Sigma_{t-1,h}} \leq \epsilon_\lambda$ for all $t \in [T]$ and $h \in [H]$.
6. $|\bar{Q}_{t,h}(s,a)| \leq V_{\max}$ and $|\bar{V}_{t,h}(s)| \leq V_{\max}$ for all $(s,a) \in \mathcal{S} \times \mathcal{A}$, $t \in [T]$, and $h \in [H]$.

Proof of Lemma B.12. The first and second statements are by Lemma B.7. The third is by Lemma B.9. The last three are by Lemma B.11. Then we take a union bound over all of them. \square

B.4 BOUNDING REGRET

We define

$$\tilde{V}_t = \mathbb{E}_{\tau \sim \pi_t^1} \left[\sum_{h=1}^H \phi(s_h, a_h)^\top \bar{\theta}_{r,t} \right].$$

Lemma B.13. *Assume all events listed in Lemma B.12 hold. Fix $t \in [T]$. Then, we have $(V^* - \bar{V}_t) + (\tilde{V} - V^{\pi_t^1}) \leq 0$ with probability at least $F^2(-1)$ where F denotes the CDF of a standard normal distribution (i.e., $\mathcal{N}(0, 1)$).*

Proof. Define the indicators

$$\begin{aligned} \mathbf{L}_{t,h}^*(s) &:= \mathbb{1} \left\{ \|\phi(s, \pi^*(s))\|_{\Sigma_{t-1,h}^{-1}} \leq \alpha_L \right\}, \\ \mathbf{M}_{t,h}^*(s) &:= \mathbb{1} \left\{ \alpha_L < \|\phi(s, \pi^*(s))\|_{\Sigma_{t-1,h}^{-1}} \leq \alpha_U \right\}, \\ \mathbf{U}_{t,h}^*(s) &:= \mathbb{1} \left\{ \|\phi(s, \pi^*(s))\|_{\Sigma_{t-1,h}^{-1}} > \alpha_U \right\}. \end{aligned}$$

We note that they are independent of the random noises at episode t and only depends on the information up to episode $t-1$. Then, we can decompose $V^* - \bar{V}_t$:

$$\begin{aligned} V^* - \bar{V}_t &= \mathbb{E}_{s_1} [V_1^*(s_1) - \bar{V}_{t,1}(s_1)] = \mathbb{E}_{s_1} [Q_1^*(s_1, \pi^*(s_1)) - \bar{Q}_{t,1}(s_1, \pi_{t,1}^0(s_1))] \\ &\leq \mathbb{E}_{s_1} [Q_1^*(s_1, \pi^*(s_1)) - \bar{Q}_{t,1}(s_1, \pi^*(s_1))] \\ &= \mathbb{E}_{s_1} [\mathbf{L}_{t,1}^*(s_1) (Q_1^*(s_1, \pi^*(s_1)) - \bar{Q}_{t,1}(s_1, \pi^*(s_1)))] \\ &\quad + \mathbb{E}_{s_1} [\mathbf{M}_{t,1}^*(s_1) (Q_1^*(s_1, \pi^*(s_1)) - \bar{Q}_{t,1}(s_1, \pi^*(s_1)))] \\ &\quad + \mathbb{E}_{s_1} [\mathbf{U}_{t,1}^*(s_1) (Q_1^*(s_1, \pi^*(s_1)) - \bar{Q}_{t,1}(s_1, \pi^*(s_1)))] \\ &=: \mathbf{T}_L + \mathbf{T}_M + \mathbf{T}_U. \end{aligned}$$

Now we establish upper bounds for each term. For \mathbf{T}_U , we have

$$\begin{aligned} &\mathbb{E}_{s_1} [\mathbf{U}_{t,1}^*(s_1) (Q_1^*(s_1, \pi^*(s_1)) - \bar{Q}_{t,1}(s_1, \pi^*(s_1)))] \\ &= \mathbb{E}_{s_1} [\mathbf{U}_{t,1}^*(s_1) (Q_1^*(s_1, \pi^*(s_1)) - \phi(s_1, \pi^*(s_1))^\top \bar{\theta}_{r,t} - (H-1))] \\ &= \mathbb{E}_{s_1} [\mathbf{U}_{t,1}^*(s_1) (Q_1^*(s_1, \pi^*(s_1)) - \phi(s_1, \pi^*(s_1))^\top \theta_r^* - \phi(s_1, \pi^*(s_1))^\top (\eta_{r,t} + \xi_{r,t}) - (H-1))] \\ &\leq \mathbb{E}_{s_1} [(-\mathbf{U}_{t,1}^*(s_1) \phi(s_1, \pi^*(s_1)))^\top (\eta_{r,t} + \xi_{r,t})] \end{aligned}$$

where the inequality holds since

$$\begin{aligned} &Q_1^*(s_1, \pi^*(s_1)) - \phi(s_1, \pi^*(s_1))^\top \theta_r^* - (H-1) \\ &= \phi(s_1, \pi^*(s_1))^\top \theta_r^* + \mathbb{E}_{s_2} [V_2^*(s_2) | s_1, \pi^*(s_1)] - \phi(s_1, \pi^*(s_1))^\top \theta_r^* - (H-1) \leq 0. \end{aligned}$$

For T_L , we have

$$\begin{aligned}
& \mathbb{E}_{s_1} [\mathbf{L}_{t,1}^*(s_1) (Q_1^*(s_1, \pi^*(s_1)) - \bar{Q}_{t,1}(s_1, \pi^*(s_1)))] \\
&= \mathbb{E}_{s_1} [\mathbf{L}_{t,1}^*(s_1) (Q_1^*(s_1, \pi^*(s_1)) - \phi(s_1, \pi^*(s_1)))^\top (\bar{\theta}_{r,t} + \bar{\theta}_{P,t,h})] \\
&= \mathbb{E}_{s_1} \left[\mathbf{L}_{t,1}^*(s_1) \left(-\phi(s_1, \pi^*(s_1))^\top (\xi_{r,t} + \xi_{P,t,1} + \eta_{r,t} + \eta_{P,t,1} - \lambda_{t,1}) + \mathbb{E}_{s_2} [V_2^*(s_2) - \bar{V}_{t,2}(s') \mid s_1, \pi^*(s_1)] \right) \right]
\end{aligned}$$

where the last equality is by Lemma B.5.

For T_M , we have

$$\begin{aligned}
& \mathbb{E}_{s_1} [\mathbf{M}_{t,1}^*(s_1) (Q_1^*(s_1, \pi^*(s_1)) - \bar{Q}_{t,1}(s_1, \pi^*(s_1)))] \\
&= \mathbb{E}_{s_1} \left[\mathbf{M}_{t,1}^*(s_1) \left(\rho(s_1, \pi^*(s_1)) (Q_1^*(s_1, \pi^*(s_1)) - \phi(s_1, \pi^*(s_1)))^\top (\bar{\theta}_{r,t} + \bar{\theta}_{P,t,1}) \right. \right. \\
&\quad \left. \left. + (1 - \rho(s_1, \pi^*(s_1))) (Q_1^*(s_1, \pi^*(s_1)) - \phi(s_1, \pi^*(s_1)))^\top \bar{\theta}_{r,t} - (H - h) \right) \right] \\
&\leq \mathbb{E}_{s_1} \left[\mathbf{M}_{t,1}^*(s_1) \left(\rho(s_1, \pi^*(s_1)) \left(-\phi(s_1, \pi^*(s_1))^\top (\xi_{r,t} + \xi_{P,t,1} + \eta_{r,t} + \eta_{P,t,1} - \lambda_{t,1}) \right. \right. \right. \\
&\quad \left. \left. + \mathbb{E}_{s_2} [V_2^*(s_2) - \bar{V}_{t,2}(s') \mid s_1, \pi^*(s_1)] \right) + (1 - \rho(s_1, \pi^*(s_1))) (-\phi(s_1, \pi^*(s_1)))^\top (\eta_{r,t} + \xi_{r,t}) \right) \right]
\end{aligned}$$

where the inequality is by similar arguments as in the case of T_U and T_L .

So putting T_L, T_M, T_U together, we have

$$\begin{aligned}
& V^* - \bar{V}_t \\
&\leq \mathbb{E}_{s_1} \left[\underbrace{\left(\mathbf{U}_{t,1}^*(s_1) + \mathbf{M}_{t,1}^*(s_1) (1 - \rho(s_1, \pi^*(s_1))) \right)}_{=: \mathbf{UM}_{t,1}^*(s_1)} (-\phi(s_1, \pi^*(s_1)))^\top (\eta_{r,t} + \xi_{r,t}) \right. \\
&\quad \left. + \underbrace{\left(\mathbf{L}_{t,1}^*(s_1) + \mathbf{M}_{t,1}^*(s_1) \rho(s_1, \pi^*(s_1)) \right)}_{=: \mathbf{LM}_{t,1}^*(s_1)} \left(-\phi(s_1, \pi^*(s_1))^\top (\xi_{r,t} + \xi_{P,t,1} + \eta_{r,t} + \eta_{P,t,1} - \lambda_{t,1}) \right. \right. \\
&\quad \left. \left. + \mathbb{E}_{s_2} [V_2^*(s_2) - \bar{V}_{t,2}(s') \mid s_1, \pi^*(s_1)] \right) \right]
\end{aligned}$$

Keeping expanding the last term, we arrive at

$$\begin{aligned}
V^* - \bar{V}_t &\leq \mathbb{E}_{\tau \sim \pi^*} \left[\sum_{h=1}^H \left(\mathbf{UM}_{t,h}^*(s_h) \prod_{i=1}^{h-1} \mathbf{LM}_{t,i}^*(s_i) \right) (-\phi(s_h, a_h))^\top (\xi_{r,t} + \eta_{r,t}) \right. \\
&\quad \left. + \sum_{h=1}^H \left(\prod_{i=1}^h \mathbf{LM}_{t,i}^*(s_i) \right) (-\phi(s_h, a_h))^\top (\xi_{r,t} + \xi_{P,t,h} + \eta_{r,t} + \eta_{P,t,h} - \lambda_{t,h}) \right] \\
&= \mathbb{E}_{\tau \sim \pi^*} \left[\sum_{h=1}^H \left(\mathbf{UM}_{t,h}^*(s_h) \prod_{i=1}^{h-1} \mathbf{LM}_{t,i}^*(s_i) + \prod_{i=1}^h \mathbf{LM}_{t,i}^*(s_i) \right) (-\phi(s_h, a_h))^\top (\xi_{r,t} + \eta_{r,t}) \right. \\
&\quad \left. + \sum_{h=1}^H \left(\prod_{i=1}^h \mathbf{LM}_{t,i}^*(s_i) \right) (-\phi(s_h, a_h))^\top (\xi_{P,t,h} + \eta_{P,t,h} - \lambda_{t,h}) \right] \\
&=: (-\phi_r^*)^\top (\xi_{r,t} + \eta_{r,t}) + \sum_{h=1}^H (-\phi_{P,h}^*)^\top (\xi_{P,t,h} + \eta_{P,t,h} - \lambda_{t,h})
\end{aligned}$$

where in the last step we defined ϕ_r^* and $\phi_{P,h}^*$ as

$$\begin{aligned}\phi_r^* &:= \mathbb{E}_{\tau \sim \pi^*} \left[\sum_{h=1}^H \left(U M_{t,h}^*(s_h) \prod_{i=1}^{h-1} L M_{t,i}^*(s_i) + \prod_{i=1}^h L M_{t,i}^*(s_i) \right) \phi(s_h, a_h) \right], \\ \phi_{P,h}^* &:= \mathbb{E}_{\tau \sim \pi^*} \left[\left(\prod_{i=1}^h L M_{t,i}^*(s_i) \right) \phi(s_h, a_h) \right].\end{aligned}$$

Now we also decompose $\tilde{V}_t - V^{\pi_t^1}$ and get

$$\begin{aligned}\tilde{V}_t - V^{\pi_t^1} &= \mathbb{E}_{\tau \sim \pi_t^1} \left[\sum_{h=1}^H \phi(s_h, a_h) \right]^\top (\bar{\theta}_{r,t} - \theta_r^*) \\ &= \mathbb{E}_{\tau \sim \pi_t^1} \left[\sum_{h=1}^H \phi(s_h, a_h) \right]^\top (\xi_{r,t} + \eta_{r,t}) \\ &=: (\tilde{\phi}_t)^\top (\xi_{r,t} + \eta_{r,t}).\end{aligned}$$

Combining the decomposition together, we obtain

$$\begin{aligned}& (V^* - \bar{V}_t) + (\tilde{V} - V^{\pi_t^1}) \\ & \leq \left((-\phi_r^*)^\top (\xi_{r,t} + \eta_{r,t}) + \sum_{h=1}^H (-\phi_{P,h}^*)^\top (\xi_{P,t,h} + \eta_{P,t,h} - \lambda_{t,h}) \right) + \tilde{\phi}_t^\top (\xi_{r,t} + \eta_{r,t}) \\ & = (\tilde{\phi}_t - \phi^*)^\top (\xi_{r,t} + \eta_{r,t}) + \sum_{h=1}^H (-\phi_{P,h}^*)^\top (\xi_{P,t,h} + \eta_{P,t,h} - \lambda_{t,h}) \\ & \leq \|\tilde{\phi}_t - \phi^*\|_{\Sigma_{t-1}^{-1}} \|\eta_{r,t}\|_{\Sigma_{t-1}} + (\tilde{\phi}_t - \phi^*)^\top \xi_{r,t} \\ & \quad + \sum_{h=1}^H \|\phi_{P,h}^*\|_{\Sigma_{t-1,h}^{-1}} (\|\eta_{P,t,h}\|_{\Sigma_{t-1,h}} + \|\lambda_{t,h}\|_{\Sigma_{t-1,h}}) - \sum_{h=1}^H \phi_{P,h}^*{}^\top \xi_{P,t,h} \\ & \leq \|\tilde{\phi}_t - \phi^*\|_{\Sigma_{t-1}^{-1}} \epsilon_{r,\eta} - (\tilde{\phi}_t - \phi^*)^\top \xi_{r,t} + \sqrt{\sum_{h=1}^H \|\phi_{P,h}^*\|_{\Sigma_{t-1,h}^{-1}}^2} \cdot \sqrt{H(\epsilon_{P,\eta} + \epsilon_\lambda)^2} - \sum_{h=1}^H \phi_{P,h}^*{}^\top \xi_{P,t,h} \\ & = \underbrace{\|\tilde{\phi}_t - \phi^*\|_{\Sigma_{t-1}^{-1}} \epsilon_{r,\eta} - (\tilde{\phi}_t - \phi^*)^\top \xi_{r,t}}_{(a)} \\ & \quad + \underbrace{\sqrt{\sum_{h=1}^H \min \left\{ 1, \|\phi_{P,h}^*\|_{\Sigma_{t-1,h}^{-1}}^2 \right\}} \cdot \sqrt{H(\epsilon_{P,\eta} + \epsilon_\lambda)^2} - \sum_{h=1}^H \phi_{P,h}^*{}^\top \xi_{P,t,h}}_{(b)}\end{aligned}$$

where the second inequality is by Cauchy-Schwarz inequality, the third inequality is Lemma B.12 and Cauchy-Schwarz inequality again, and the last step is by the fact that $\|\phi_{P,h}^*\|_{\Sigma_{t-1,h}^{-1}}^2 \leq 1/\lambda \leq 1$.

For (a), we note that $(\tilde{\phi}_t - \phi^*)^\top \xi_{r,t} \sim \mathcal{N}(0, \sigma_r^2 \|\tilde{\phi}_t - \phi^*\|_{\Sigma_{t-1}^{-1}}^2)$. So by setting $\sigma_r \geq \epsilon_{r,\eta}$, we have (a) ≤ 0 with probability at least $F(-1)$.

For (b), similarly, we note that $\sum_{h=1}^H \phi_{P,h}^*{}^\top \xi_{P,t,h} \sim \mathcal{N}(0, \sigma_P^2 \sum_{h=1}^H \|\phi_{P,h}^*\|_{\Sigma_{t-1,h}^{-1}}^2)$. So by setting $\sigma_P \geq \sqrt{H(\epsilon_{P,\eta} + \epsilon_\lambda)^2}$, we have (b) ≤ 0 with probability at least $F(-1)$.

Conclusion. Finally, we note that (a) and (b) are independent, so the probability that both (a) and (b) hold is at least $F^2(-1)$. \square

Lemma B.14. *It holds that*

$$\sum_{t=1}^T \sum_{h=1}^H \mathbf{L}_{t,h}^{\mathbf{G}}(s_{t,h}) \leq L_{\max}.$$

Proof. Denote $a_{t,h} = \pi_t(s_{t,h})$. Then, we have

$$\begin{aligned} \sum_{t=1}^T \sum_{h=1}^H \mathbf{L}_{t,h}^{\mathbf{G}}(s_{t,h}) &= \sum_{t=1}^T \sum_{h=1}^H \mathbb{1} \left\{ \|\phi(s_{t,h}, a_{t,h})\|_{\Sigma_{t-1,h}^{-1}}^2 > \alpha_L^2 \right\} \\ &\leq \sum_{t=1}^T \sum_{h=1}^H \min \left\{ \frac{\|\phi(s_{t,h}, a_{t,h})\|_{\Sigma_{t-1,h}^{-1}}^2}{\alpha_L^2}, 1 \right\} \\ &\leq \frac{1}{\alpha_L^2} \sum_{h=1}^H \sum_{t=1}^T \min \left\{ \|\phi(s_{t,h}, a_{t,h})\|_{\Sigma_{t-1,h}^{-1}}^2, 1 \right\} \\ &\leq \frac{H}{\alpha_L^2} \cdot 2d \log \left(\frac{\lambda + T}{\lambda} \right) = L_{\max} \end{aligned}$$

where the second inequality uses the fact that $\alpha_L < 1$, and the last inequality is by Lemma D.6. \square

Lemma B.15. *Assume all events listed in Lemma B.12 hold. Then, the following holds*

$$\begin{aligned} \left| \sum_t \left((\bar{V}_t - V^{\pi_t^0}) - (\tilde{V}_t - V^{\pi_t^1}) \right) \right| &= \tilde{O} \left(V_{\max} L_{\max} + V_{\max} \sqrt{HT} \right. \\ &\quad \left. + (\epsilon_{r,\xi} + \epsilon_{r,\eta}) \left(\sqrt{dT} + L_{\max} + \sqrt{HT} + \frac{\epsilon T \sqrt{\pi}}{2\sigma_r} \right) \right. \\ &\quad \left. + (\epsilon_{P,\xi} + \epsilon_{P,\eta} + \epsilon_\lambda) \left(H\sqrt{dT} + L_{\max} + \sqrt{HT} \right) \right) \end{aligned}$$

Proof. By definition, we have

$$\begin{aligned} &\left| \sum_{t=1}^T \left((\bar{V}_t - V^{\pi_t^0}) - (\tilde{V}_t - V^{\pi_t^1}) \right) \right| \\ &= \left| \sum_{t=1}^T \mathbb{E}_{s_1, a_1 \sim \pi_t^0} \left[\bar{Q}_{t,1}(s_1, a_1) - Q_1^{\pi_t^0}(s_1, a_1) \right] - \sum_{t=1}^T \sum_{h=1}^H \mathbb{E}_{s_h, a_h \sim \pi_t^1} \phi(s_h, a_h) (\bar{\theta}_{r,t} - \theta_r^*) \right| \end{aligned}$$

By triangle inequality, we have

$$\begin{aligned} &\leq \left| \sum_{t=1}^T \mathbb{E}_{s_1, a_1 \sim \pi_t^0} \left[\mathbf{L}_{t,1}(s_{t,1}) \left(\bar{Q}_{t,1}(s_1, a_1) - Q_1^{\pi_t^0}(s_1, a_1) \right) \right] - \sum_{t=1}^T \sum_{h=1}^H \mathbb{E}_{s_h, a_h \sim \pi_t^1} \phi(s_h, a_h) (\xi_{r,t} + \eta_{r,t}) \right| \\ &\quad + 2V_{\max} \sum_{t=1}^T \mathbb{E}_{s_1, a_1 \sim \pi_t^0} \left[\mathbf{L}_{t,1}^{\mathbf{G}}(s_{t,1}) \right] \end{aligned}$$

For the first term, conditioning on $\mathbf{L}_{t,1}(s_{t,1})$, we have

$$\begin{aligned} &= \left| \sum_{t=1}^T \mathbb{E}_{s_1, a_1 \sim \pi_t^0} \left[\mathbf{L}_{t,1}(s_{t,1}) \left(\phi(s_1, a_1)^\top (\bar{\theta}_{r,t} + \bar{\theta}_{P,t,1}) - Q_1^{\pi_t^0}(s_1, a_1) \right) \right] \right. \\ &\quad \left. - \sum_{t=1}^T \sum_{h=1}^H \mathbb{E}_{s_h, a_h \sim \pi_t^1} \phi(s_h, a_h) (\xi_{r,t} + \eta_{r,t}) \right| + 2V_{\max} \sum_{t=1}^T \mathbb{E}_{s_1, a_1 \sim \pi_t^0} \left[\mathbf{L}_{t,1}^{\mathbf{G}}(s_{t,1}) \right] \end{aligned}$$

Applying triangle inequality again, we get

$$\begin{aligned} &\leq \left| \sum_{t=1}^T \mathbb{E}_{s_1, a_1 \sim \pi_t^0} \left[\left(\phi(s_1, a_1)^\top (\bar{\theta}_{r,t} + \bar{\theta}_{p,t,1}) - Q_1^{\pi_t^0}(s_1, a_1) \right) \right] - \sum_{t=1}^T \sum_{h=1}^H \mathbb{E}_{s_h, a_h \sim \pi_t^1} \phi(s_h, a_h) (\xi_{r,t} + \eta_{r,t}) \right| \\ &\quad + 2V_{\max} \sum_{t=1}^T \mathbb{E}_{s_1, a_1 \sim \pi_t^0} \left[\mathbf{L}_{t,1}^{\mathcal{G}}(s_{t,1}) \right] \\ &\quad + \left| \sum_{t=1}^T \mathbb{E}_{s_1, a_1 \sim \pi_t^0} \left[\mathbf{L}_{t,1}^{\mathcal{G}}(s_{t,1}) \left(\phi(s_1, a_1)^\top (\bar{\theta}_{r,t} + \bar{\theta}_{p,t,1}) - Q_1^{\pi_t^0}(s_1, a_1) \right) \right] \right| \end{aligned}$$

Applying Lemma B.5 and the triangle inequality, we get

$$\begin{aligned} &\leq \left| \sum_{t=1}^T \mathbb{E}_{s_2 \sim \pi_t^0} \left[\bar{V}_{t,2}(s_2) - V_2^{\pi_t^0}(s_2) \right] - \sum_{t=1}^T \sum_{h=2}^H \mathbb{E}_{s_h, a_h \sim \pi_t^1} \phi(s_h, a_h) (\xi_{r,t} + \eta_{r,t}) \right. \\ &\quad \left. + \sum_{t=1}^T \left(\mathbb{E}_{s_1, a_1 \sim \pi_t^0} \phi^\top(s_1, a_1) (\xi_{r,t} + \xi_{p,t,1} + \eta_{r,t} + \eta_{p,t,1} - \lambda_{t,1}) - \mathbb{E}_{s_1, a_1 \sim \pi_t^1} \phi^\top(s_1, a_1) (\xi_{r,t} + \eta_{r,t}) \right) \right| \\ &\quad + 2V_{\max} \sum_{t=1}^T \mathbb{E}_{s_1, a_1 \sim \pi_t^0} \left[\mathbf{L}_{t,1}^{\mathcal{G}}(s_{t,1}) \right] \\ &\quad + \left| \sum_{t=1}^T \mathbb{E}_{s_1, a_1 \sim \pi_t^0} \left[\mathbf{L}_{t,1}^{\mathcal{G}}(s_{t,1}) \left(\phi(s_1, a_1)^\top (\bar{\theta}_{r,t} + \bar{\theta}_{p,t,1}) - Q_1^{\pi_t^0}(s_1, a_1) \right) \right] \right| \end{aligned}$$

Keep expanding the first term, we get

$$\begin{aligned} &\left| \sum_{t=1}^T \left((\bar{V}_t - V^{\pi_t^0}) - (\tilde{V}_t - V^{\pi_t^1}) \right) \right| \\ &\leq 2V_{\max} \sum_{h=1}^H \sum_{t=1}^T \mathbb{E}_{s_h, a_h \sim \pi_t^0} \left[\mathbf{L}_{t,h}^{\mathcal{G}}(s_{t,h}) \right] \\ &\quad + \sum_{h=1}^H \left| \sum_{t=1}^T \mathbb{E}_{s_h, a_h \sim \pi_t^0} \left[\mathbf{L}_{t,h}^{\mathcal{G}}(s_{t,h}) \left(\phi(s_h, a_h)^\top (\bar{\theta}_{r,t} + \bar{\theta}_{p,t,h}) - Q_h^{\pi_t^0}(s_h, a_h) \right) \right] \right| \\ &\quad + \left| \sum_{t=1}^T \sum_{h=1}^H \left(\mathbb{E}_{s_h, a_h \sim \pi_t^0} \mathbb{E}_{\tilde{s}_h, \tilde{a}_h \sim \pi_t^1} \left(\phi(s_h, a_h) - \phi(\tilde{s}_h, \tilde{a}_h) \right)^\top (\xi_{r,t} + \eta_{r,t}) \right) \right| \\ &\quad + \left| \sum_{t=1}^T \sum_{h=1}^H \mathbb{E}_{s_h, a_h \sim \pi_t^0} \phi^\top(s_h, a_h) (\xi_{p,t,h} + \eta_{p,t,h} - \lambda_{t,h}) \right| \\ &=: \mathbf{T}_1 + \mathbf{T}_2 + \mathbf{T}_3 + \mathbf{T}_4. \end{aligned}$$

We bound each term separately.

Bounding \mathbf{T}_1 . By Hoeffding's inequality and Lemma B.14, we have

$$\begin{aligned} \mathbf{T}_1 &\leq 2V_{\max} \sum_{h=1}^H \sum_{t=1}^T \left[\mathbf{L}_{t,h}^{\mathcal{G}}(s_{t,h}) \right] + 2V_{\max} \sqrt{\frac{HT}{2} \log(1/\delta)} \\ &\leq 2V_{\max} \left(L_{\max} + \sqrt{\frac{HT}{2} \log(1/\delta)} \right). \end{aligned}$$

with probability at least $1 - \delta$.

Bounding T_2 . By definition and Lemma B.4, we have

$$\begin{aligned}
T_2 &= \sum_{h=1}^H \left| \sum_{t=1}^T \mathbb{E}_{s_h, a_h \sim \pi_t^0} \left[\mathbf{L}_{t,h}^{\mathbf{C}}(s_{t,h}) \left(\phi(s_h, a_h)^\top (\xi_{r,t} + \eta_{r,t} + \theta_r^* + \xi_{P,t,h} + \eta_{P,t,h} + \lambda_{t,h} + \theta_{P,t,h}^*) \right) \right. \right. \\
&\quad \left. \left. - Q_h^{\pi_t^0}(s_h, a_h) \right) \right] \Big| \\
&\leq \sum_{h=1}^H \left| \sum_{t=1}^T \mathbb{E}_{s_h, a_h \sim \pi_t^0} \left[\mathbf{L}_{t,h}^{\mathbf{C}}(s_{t,h}) \left(\phi(s_h, a_h)^\top (\theta_r^* + \theta_{P,t,h}^*) - Q_h^{\pi_t^0}(s_h, a_h) \right) \right. \right. \\
&\quad \left. \left. + \phi^\top(s_h, a_h) (\xi_{r,t} + \eta_{r,t}) + \phi^\top(s_h, a_h) (\xi_{P,t,h} + \eta_{P,t,h} + \lambda_{t,h}) \right) \right] \Big| \\
&\leq \sum_{h=1}^H \left| \sum_{t=1}^T \mathbb{E}_{s_h, a_h \sim \pi_t^0} \left[\mathbf{L}_{t,h}^{\mathbf{C}}(s_{t,h}) \left(\phi(s_h, a_h)^\top (\theta_r^* + \theta_{P,t,h}^*) - Q_h^{\pi_t^0}(s_h, a_h) \right) \right. \right. \\
&\quad \left. \left. + \|\phi(s_h, a_h)\|_{\Sigma_{t-1}^{-1}} (\|\xi_{r,t}\|_{\Sigma_{t-1}} + \|\eta_{r,t}\|_{\Sigma_{t-1}}) \right. \right. \\
&\quad \left. \left. + \|\phi(s_h, a_h)\|_{\Sigma_{t-1,h}^{-1}} (\|\xi_{P,t,h}\|_{\Sigma_{t-1,h}} + \|\eta_{P,t,h}\|_{\Sigma_{t-1,h}} + \|\lambda_{t,h}\|_{\Sigma_{t-1,h}}) \right) \right] \Big|.
\end{aligned}$$

We note that $|\phi(s_h, a_h)^\top (\theta_r^* + \theta_{P,t,h}^*)| \leq V_{\max}$ and $|Q_h^{\pi_t^0}(s_h, a_h)| \leq V_{\max}$. Moreover, we have $\|\phi(s_h, a_h)\|_{\Sigma_{t-1}^{-1}} \leq 1/\sqrt{\lambda}$ and $\|\phi(s_h, a_h)\|_{\Sigma_{t-1,h}^{-1}} \leq 1/\sqrt{\lambda}$. Apply triangle inequality and inserting these upper bounds back, we obtain

$$\begin{aligned}
T_2 &\leq \sum_{h=1}^H \left| \sum_{t=1}^T \mathbb{E}_{s_h, a_h \sim \pi_t^0} \left[\mathbf{L}_{t,h}^{\mathbf{C}}(s_{t,h}) \left(\phi(s_h, a_h)^\top (\theta_r^* + \theta_{P,t,h}^*) - Q_h^{\pi_t^0}(s_h, a_h) \right) \right] \right| \\
&\quad + \sum_{h=1}^H \sum_{t=1}^T \mathbb{E}_{s_h, a_h \sim \pi_t^0} \left[\mathbf{L}_{t,h}^{\mathbf{C}}(s_{t,h}) \|\phi(s_h, a_h)\|_{\Sigma_{t-1}^{-1}} (\|\xi_{r,t}\|_{\Sigma_{t-1}} + \|\eta_{r,t}\|_{\Sigma_{t-1}}) \right] \\
&\quad + \sum_{h=1}^H \sum_{t=1}^T \mathbb{E}_{s_h, a_h \sim \pi_t^0} \left[\mathbf{L}_{t,h}^{\mathbf{C}}(s_{t,h}) \|\phi(s_h, a_h)\|_{\Sigma_{t-1,h}^{-1}} (\|\xi_{P,t,h}\|_{\Sigma_{t-1,h}} + \|\eta_{P,t,h}\|_{\Sigma_{t-1,h}} + \|\lambda_{t,h}\|_{\Sigma_{t-1,h}}) \right] \\
&\leq 2V_{\max} \sum_{h=1}^H \sum_{t=1}^T \mathbb{E}_{s_h, a_h \sim \pi_t^0} \left[\mathbf{L}_{t,h}^{\mathbf{C}}(s_{t,h}) \right] \\
&\quad + \sqrt{\lambda^{-1}} \sum_{h=1}^H \sum_{t=1}^T \mathbb{E}_{s_h, a_h \sim \pi_t^0} \left[\mathbf{L}_{t,h}^{\mathbf{C}}(s_{t,h}) (\epsilon_{r,\xi} + \epsilon_{r,\eta} + \epsilon_{P,\xi} + \epsilon_{P,\eta} + \epsilon_\lambda) \right] \\
&= \underbrace{\sum_{h=1}^H \sum_{t=1}^T \mathbb{E}_{s_h, a_h \sim \pi_t^0} \left[\mathbf{L}_{t,h}^{\mathbf{C}}(s_{t,h}) \right]}_{(*)} \left(2V_{\max} + \frac{\epsilon_{r,\xi} + \epsilon_{r,\eta} + \epsilon_{P,\xi} + \epsilon_{P,\eta} + \epsilon_\lambda}{\sqrt{\lambda}} \right).
\end{aligned}$$

We notice that $(*)$ can be bounded by Hoeffding's inequality and Lemma B.14, as we did similarly for T_1 :

$$(*) \leq L_{\max} + \sqrt{\frac{HT}{2} \log(1/\delta)}$$

Hence, we have

$$T_2 \leq \left(L_{\max} + \sqrt{\frac{HT}{2} \log(1/\delta)} \right) \left(2V_{\max} + \frac{\epsilon_{r,\xi} + \epsilon_{r,\eta} + \epsilon_{P,\xi} + \epsilon_{P,\eta} + \epsilon_\lambda}{\sqrt{\lambda}} \right)$$

with probability at least $1 - \delta$.

Bounding T_3 . By definition, we have

$$T_3 = \left| \sum_{t=1}^T \mathbb{E}_{\tau \sim \pi_t^0} \mathbb{E}_{\tilde{\tau} \sim \pi_t^1} (\phi(\tau) - \phi(\tilde{\tau})) (\xi_{r,t} + \eta_{r,t}) \right|$$

By Cauchy-Schwarz inequality, we have

$$\begin{aligned} &\leq \sum_{t=1}^T \mathbb{E}_{\tau \sim \pi_t^0} \mathbb{E}_{\tilde{\tau} \sim \pi_t^1} \left\| \phi(\tau) - \phi(\tilde{\tau}) \right\|_{\Sigma_{t-1}^{-1}} \left(\|\xi_{r,t}\|_{\Sigma_{t-1}} + \|\eta_{r,t}\|_{\Sigma_{t-1}} \right) \\ &\leq (\epsilon_{r,\xi} + \epsilon_{r,\eta}) \sum_{t=1}^T \mathbb{E}_{\tau \sim \pi_t^0} \mathbb{E}_{\tilde{\tau} \sim \pi_t^1} \left\| \phi(\tau) - \phi(\tilde{\tau}) \right\|_{\Sigma_{t-1}^{-1}} \end{aligned}$$

Since $\|\phi(\tau) - \phi(\tilde{\tau})\|_{\Sigma_{t-1}^{-1}} \leq 2/\sqrt{\lambda}$, by Hoeffding's inequality, we have

$$\begin{aligned} &\leq (\epsilon_{r,\xi} + \epsilon_{r,\eta}) \left(\sum_{t=1}^T \left\| \phi(\tau) - \phi(\tilde{\tau}) \right\|_{\Sigma_{t-1}^{-1}} + \sqrt{\frac{2T \log(1/\delta)}{\lambda}} \right) \\ &= (\epsilon_{r,\xi} + \epsilon_{r,\eta}) \left(\underbrace{\sum_{t=1}^T Z_t \left\| \phi(\tau) - \phi(\tilde{\tau}) \right\|_{\Sigma_{t-1}^{-1}}}_{(i)} + \underbrace{\sum_{t=1}^T (1 - Z_t) \left\| \phi(\tau) - \phi(\tilde{\tau}) \right\|_{\Sigma_{t-1}^{-1}}}_{(ii)} + \sqrt{\frac{2T \log(1/\delta)}{\lambda}} \right). \end{aligned}$$

To bound (i), we have

$$\begin{aligned} (i) &= \sum_{t=1}^T Z_t \min \left\{ 2/\sqrt{\lambda}, \left\| \phi(\tau) - \phi(\tilde{\tau}) \right\|_{\Sigma_{t-1}^{-1}} \right\} \\ &\leq \sqrt{T \sum_{t=1}^T Z_t \min \left\{ 4/\lambda, \left\| \phi(\tau) - \phi(\tilde{\tau}) \right\|_{\Sigma_{t-1}^{-1}}^2 \right\}} \\ &\leq 2 \sqrt{T \sum_{t=1}^T Z_t \min \left\{ 1, \left\| \phi(\tau) - \phi(\tilde{\tau}) \right\|_{\Sigma_{t-1}^{-1}}^2 \right\}} \\ &\leq 2 \sqrt{T \cdot 2d \log \left(\frac{\lambda + 4T}{\lambda} \right)} \end{aligned}$$

where the last inequality is Lemma D.6. To bound (ii), we have

$$\begin{aligned} (ii) &= \sum_{t=1}^T (1 - Z_t) \left\| \phi(\tau) - \phi(\tilde{\tau}) \right\|_{\Sigma_{t-1}^{-1}} \\ &\leq \sum_{t=1}^T (1 - Z_t) \cdot \epsilon \sqrt{\pi} / (2\sigma_r) \\ &\leq \epsilon T \sqrt{\pi} / (2\sigma_r) \end{aligned}$$

where the first inequality is by Lemma B.6. Putting the two upper bounds together, we have

$$T_3 \leq (\epsilon_{r,\xi} + \epsilon_{r,\eta}) \left(2 \sqrt{T \cdot 2d \log \left(\frac{\lambda + 4T}{\lambda} \right)} + \epsilon T \sqrt{\pi} / (2\sigma_r) + \sqrt{\frac{2T \log(1/\delta)}{\lambda}} \right).$$

Bounding T_4 . By Cauchy-Schwarz inequality, we have

$$\begin{aligned} T_4 &\leq \left| \sum_{t=1}^T \sum_{h=1}^H \mathbb{E}_{s_h, a_h \sim \pi_t^0} \left\| \phi(s_h, a_h) \right\|_{\Sigma_{t-1,h}^{-1}} \left(\|\xi_{P,t,h}\|_{\Sigma_{t-1,h}} + \|\eta_{P,t,h}\|_{\Sigma_{t-1,h}} + \|\lambda_{t,h}\|_{\Sigma_{t-1,h}} \right) \right| \\ &\leq (\epsilon_{P,\xi} + \epsilon_{P,\eta} + \epsilon_\lambda) \sum_{h=1}^H \sum_{t=1}^T \mathbb{E}_{s_h, a_h \sim \pi_t^0} \left\| \phi(s_h, a_h) \right\|_{\Sigma_{t-1,h}^{-1}} \end{aligned}$$

Since $\|\phi(s_h, a_h)\|_{\Sigma_{t-1,h}^{-1}} \leq 1/\sqrt{\lambda}$, by Hoeffding's inequality, we have

$$\begin{aligned} &\leq (\epsilon_{P,\xi} + \epsilon_{P,\eta} + \epsilon_\lambda) \left(\sum_{h=1}^H \sum_{t=1}^T \|\phi(s_h, a_h)\|_{\Sigma_{t-1,h}^{-1}} + \sqrt{\frac{TH \log(1/\delta)}{2\lambda}} \right) \\ &= (\epsilon_{P,\xi} + \epsilon_{P,\eta} + \epsilon_\lambda) \left(\sum_{h=1}^H \sum_{t=1}^T \min \left\{ 1/\sqrt{\lambda}, \|\phi(s_h, a_h)\|_{\Sigma_{t-1,h}^{-1}} \right\} + \sqrt{\frac{TH \log(1/\delta)}{2\lambda}} \right) \\ &\leq (\epsilon_{P,\xi} + \epsilon_{P,\eta} + \epsilon_\lambda) \left(\sum_{h=1}^H \sqrt{T \sum_{t=1}^T \min \left\{ 1/\lambda, \|\phi(s_h, a_h)\|_{\Sigma_{t-1,h}^{-1}}^2 \right\}} + \sqrt{\frac{TH \log(1/\delta)}{2\lambda}} \right) \end{aligned}$$

Since $\lambda \geq 1$, we have

$$\begin{aligned} &\leq (\epsilon_{P,\xi} + \epsilon_{P,\eta} + \epsilon_\lambda) \left(\sum_{h=1}^H \sqrt{T \sum_{t=1}^T \min \left\{ 1, \|\phi(s_h, a_h)\|_{\Sigma_{t-1,h}^{-1}}^2 \right\}} + \sqrt{\frac{TH \log(1/\delta)}{2\lambda}} \right) \\ &\leq (\epsilon_{P,\xi} + \epsilon_{P,\eta} + \epsilon_\lambda) \left(H \sqrt{T \cdot 2d \log \left(\frac{\lambda + T}{\lambda} \right)} + \sqrt{\frac{TH \log(1/\delta)}{2\lambda}} \right) \end{aligned}$$

where the last inequality is Lemma D.6.

Conclusion. Combining the above bounds, we have

$$\begin{aligned} &\left| \sum_{t=1}^T \left((\bar{V}_t - V^{\pi_t^0}) - (\tilde{V}_t - V^{\pi_t^1}) \right) \right| \\ &\leq 2V_{\max} \left(L_{\max} + \sqrt{\frac{HT}{2} \log(1/\delta)} \right) \\ &\quad + \left(L_{\max} + \sqrt{\frac{HT}{2} \log(1/\delta)} \right) \left(2V_{\max} + \frac{\epsilon_{r,\xi} + \epsilon_{r,\eta} + \epsilon_{P,\xi} + \epsilon_{P,\eta} + \epsilon_\lambda}{\sqrt{\lambda}} \right) \\ &\quad + (\epsilon_{r,\xi} + \epsilon_{r,\eta}) \left(2\sqrt{T \cdot 2d \log \left(\frac{\lambda + 4T}{\lambda} \right)} + \frac{\epsilon T \sqrt{\pi}}{2\sigma_r} + \sqrt{\frac{2T \log(1/\delta)}{\lambda}} \right) \\ &\quad + (\epsilon_{P,\xi} + \epsilon_{P,\eta} + \epsilon_\lambda) \left(H \sqrt{T \cdot 2d \log \left(\frac{\lambda + T}{\lambda} \right)} + \sqrt{\frac{TH \log(1/\delta)}{2\lambda}} \right) \\ &\leq 4V_{\max} \left(L_{\max} + \sqrt{\frac{HT}{2} \log(1/\delta)} \right) \\ &\quad + (\epsilon_{r,\xi} + \epsilon_{r,\eta}) \left(2\sqrt{T \cdot 2d \log \left(\frac{\lambda + 4T}{\lambda} \right)} + \frac{\epsilon T \sqrt{\pi}}{2\sigma_r} + \frac{L_{\max} + 2\sqrt{TH \log(1/\delta)/2}}{\sqrt{\lambda}} \right) \\ &\quad + (\epsilon_{P,\xi} + \epsilon_{P,\eta} + \epsilon_\lambda) \left(H \sqrt{T \cdot 2d \log \left(\frac{\lambda + T}{\lambda} \right)} + \frac{L_{\max} + 2\sqrt{TH \log(1/\delta)/2}}{\sqrt{\lambda}} \right) \\ &= \tilde{O} \left(V_{\max} L_{\max} + V_{\max} \sqrt{HT} \right. \\ &\quad + (\epsilon_{r,\xi} + \epsilon_{r,\eta}) \left(\sqrt{dT} + L_{\max} + \sqrt{HT} + \frac{\epsilon T \sqrt{\pi}}{2\sigma_r} \right) \\ &\quad \left. + (\epsilon_{P,\xi} + \epsilon_{P,\eta} + \epsilon_\lambda) \left(H \sqrt{dT} + L_{\max} + \sqrt{HT} \right) \right). \end{aligned}$$

□

Lemma B.16 (Regret decomposition). *Assume all events listed in Lemma B.12 hold. Then, we have*

$$\begin{aligned} & \sum_{t=1}^T (V^* - \bar{V}_t) + (\tilde{V}_t - V^{\pi_t^1}) \\ & \leq \frac{1}{F^2(-1)} \sum_{t=1}^T \left((\bar{V}_t - V^{\pi_t^0}) - (\tilde{V}_t - V^{\pi_t^1}) + (V^{\pi_t^0} - \bar{V}_t^-) - (V^{\pi_t^1} - \tilde{V}_t^-) \right) \\ & \quad + \frac{2V_{\max}}{F^2(-1)} \left(\sqrt{\frac{T \log(1/\delta)}{2}} + 4\delta \right) \end{aligned}$$

with probability at least $1 - \delta$.

Proof. We note that, at any round $t \in [T]$, conditioning on all information collected up to round $t - 1$, the randomness of \bar{V}_t and \tilde{V}_t only comes from the randomness of Gaussian noise variables $\xi_{r,t}, \xi_{P,t,1}, \dots, \xi_{P,t,H}$. In other words, the values of \bar{V}_t and \tilde{V}_t are determined once given these Gaussian noise variables. In light of this, we write out the dependence on the Gaussian noise variables explicitly: we treat \bar{V}_t and \tilde{V}_t as functions of $\xi_{r,t}, \xi_{P,t,1}, \dots, \xi_{P,t,H}$ and define $\bar{V}_t[\xi_{r,t}, \xi_{P,t,1}, \dots, \xi_{P,t,H}]$ and $\tilde{V}_t[\xi_{r,t}, \xi_{P,t,1}, \dots, \xi_{P,t,H}]$ as the values of \bar{V}_t and \tilde{V}_t obtained at round t with the Gaussian noise variables $\xi_{r,t}, \xi_{P,t,1}, \dots, \xi_{P,t,H}$. Then, we define a notion of “worst-case” Gaussian noise variables as follows:

$$\begin{aligned} \xi_{r,t}^-, \xi_{P,t,1}^-, \dots, \xi_{P,t,H}^- & := \arg \min_{\xi_{r,t}^-, \xi_{P,t,1}^-, \dots, \xi_{P,t,H}^-} \bar{V}_t[\xi_{r,t}^-, \xi_{P,t,1}^-, \dots, \xi_{P,t,H}^-] - \tilde{V}_t[\xi_{r,t}^-, \xi_{P,t,1}^-, \dots, \xi_{P,t,H}^-] \\ & \text{s.t. } \|\xi_{r,t}^-\|_{\Sigma_{t-1}} \leq \epsilon_{r,\xi} \quad \text{and} \quad \forall h \in [H] : \|\xi_{P,t,h}^-\|_{\Sigma_{t-1,h}} \leq \epsilon_{P,\xi}. \end{aligned}$$

And we denote \bar{V}_t^- and \tilde{V}_t^- as the value functions specified by $\xi_{r,t}^-, \xi_{P,t,1}^-, \dots, \xi_{P,t,H}^-$:

$$\bar{V}_t^- := \bar{V}_t[\xi_{r,t}^-, \xi_{P,t,1}^-, \dots, \xi_{P,t,H}^-], \quad \tilde{V}_t^- := \tilde{V}_t[\xi_{r,t}^-, \xi_{P,t,1}^-, \dots, \xi_{P,t,H}^-].$$

In other words, \bar{V}_t^- and \tilde{V}_t^- are counterparts of \bar{V}_t and \tilde{V}_t that attain the smallest difference, $\bar{V}_t - \tilde{V}_t$, while the noise variables still satisfy the high probability bounds. By Lemma B.12, we immediately have

$$\Pr(\mathcal{E}_{\text{low}}) := \Pr(\bar{V}_t^- - \tilde{V}_t^- \leq \bar{V}_t - \tilde{V}_t) \geq 1 - \delta/T. \quad (11)$$

We note that here we have δ/T instead of δ because we are considering a fixed t and the results of Lemma B.12 are derived by the union bound over all $t \in [T]$ — thus, the probability of the event \mathcal{E}_{low} is at least $1 - \delta/T$ for a single t .

Then, we denote \mathcal{E}_{opt} as the event that $\bar{V}_t - \tilde{V}_t \geq V^* - V^{\pi_t^1}$. Thus, by Lemma B.13, we have $\Pr(\mathcal{E}_{\text{opt}}) \geq F^2(-1)$. Moreover, we denote p_{alg} as the randomness of the algorithm. Then, we define the joint distribution p_{opt} of \bar{V}_t and \tilde{V}_t by restricting p_{alg} to the event \mathcal{E}_{opt} . Specifically, it is defined as

$$p_{\text{opt}}(\bar{V}_t, \tilde{V}_t) := \begin{cases} p_{\text{alg}}(\bar{V}_t, \tilde{V}_t) / \Pr(\mathcal{E}_{\text{opt}}) & \text{if } \mathcal{E}_{\text{opt}} \\ 0 & \text{otherwise} \end{cases}$$

Let $z := \bar{V}_t - \tilde{V}_t$. Then, we have

$$\begin{aligned} (V^* - \bar{V}_t) + (\tilde{V}_t - V^{\pi_t^1}) & \leq V^* - \bar{V}_t^- + \tilde{V}_t^- - V^{\pi_t^1} \\ & = \mathbb{E}_{z \sim p_{\text{opt}}} \left[V^* - \bar{V}_t^- - z + z + \tilde{V}_t^- - V^{\pi_t^1} \right] \\ & \leq \mathbb{E}_{z \sim p_{\text{opt}}} \left[z - \bar{V}_t^- + \tilde{V}_t^- \right] \\ & = \int_z \frac{\mathbf{1}\{\mathcal{E}_{\text{opt}}\} \cdot p_{\text{alg}}(z)}{\Pr(\mathcal{E}_{\text{opt}})} \cdot (z - \bar{V}_t^- + \tilde{V}_t^-) dz \\ & = \mathbb{E}_{z \sim p_{\text{alg}}} \left[\mathbf{1}\{\mathcal{E}_{\text{opt}}\} (z - \bar{V}_t^- + \tilde{V}_t^-) \right] / \Pr(\mathcal{E}_{\text{opt}}) \end{aligned}$$

where the second inequality is by the definition of p_{opt} , which rules out the case that \mathcal{E}_{opt} does not hold. Considering the event \mathcal{E}_{low} , we have

$$\begin{aligned}
& (V^* - \bar{V}_t) + (\tilde{V}_t - V^{\pi_t^1}) \\
& \leq \left(\mathbb{E}_{z \sim p_{\text{alg}}} \left[\mathbf{1}\{\mathcal{E}_{\text{opt}}\} \mathbf{1}\{\mathcal{E}_{\text{low}}\} \left(z - \bar{V}_t^- + \tilde{V}_t^- \right) \right] + \mathbb{E}_{z \sim p_{\text{alg}}} \left[\mathbf{1}\{\mathcal{E}_{\text{opt}}\} \mathbf{1}\{\mathcal{E}_{\text{low}}^c\} \left(z - \bar{V}_t^- + \tilde{V}_t^- \right) \right] \right) / \Pr(\mathcal{E}_{\text{opt}}) \\
& \leq \left(\mathbb{E}_{z \sim p_{\text{alg}}} \left[\mathbf{1}\{\mathcal{E}_{\text{opt}}\} \mathbf{1}\{\mathcal{E}_{\text{low}}\} \left(z - \bar{V}_t^- + \tilde{V}_t^- \right) \right] + \mathbb{E}_{z \sim p_{\text{alg}}} \left[\mathbf{1}\{\mathcal{E}_{\text{opt}}\} \mathbf{1}\{\mathcal{E}_{\text{low}}^c\} 4V_{\text{max}} \right] \right) / \Pr(\mathcal{E}_{\text{opt}}) \\
& \leq \left(\mathbb{E}_{z \sim p_{\text{alg}}} \left[\mathbf{1}\{\mathcal{E}_{\text{opt}}\} \mathbf{1}\{\mathcal{E}_{\text{low}}\} \left(z - \bar{V}_t^- + \tilde{V}_t^- \right) \right] + 4V_{\text{max}}\delta/T \right) / \Pr(\mathcal{E}_{\text{opt}})
\end{aligned}$$

where the last inequality is by (11). For the first term, we have

$$\begin{aligned}
& \mathbb{E}_{z \sim p_{\text{alg}}} \left[\mathbf{1}\{\mathcal{E}_{\text{opt}}\} \mathbf{1}\{\mathcal{E}_{\text{low}}\} \left(z - \bar{V}_t^- + \tilde{V}_t^- \right) \right] \\
& \leq \mathbb{E}_{z \sim p_{\text{alg}}} \left[\mathbf{1}\{\mathcal{E}_{\text{low}}\} \left(z - \bar{V}_t^- + \tilde{V}_t^- \right) \right] \\
& = \mathbb{E}_{z \sim p_{\text{alg}}} \left[\left(z - \bar{V}_t^- + \tilde{V}_t^- \right) \right] + \mathbb{E}_{z \sim p_{\text{alg}}} \left[\mathbf{1}\{\mathcal{E}_{\text{low}}^c\} \left(z - \bar{V}_t^- + \tilde{V}_t^- \right) \right] \\
& \leq \mathbb{E}_{z \sim p_{\text{alg}}} \left[\left(z - \bar{V}_t^- + \tilde{V}_t^- \right) \right] + 4V_{\text{max}}\delta/T \\
& = \mathbb{E} \left[\bar{V}_t - \tilde{V}_t - \bar{V}_t^- + \tilde{V}_t^- \right] + 4V_{\text{max}}\delta/T \\
& = \mathbb{E} \left[(\bar{V}_t - V^{\pi_t^0}) - (\tilde{V}_t - V^{\pi_t^1}) \right] + \mathbb{E} \left[(V^{\pi_t^0} - \bar{V}_t^-) - (V^{\pi_t^1} - \tilde{V}_t^-) \right] + 4V_{\text{max}}\delta/T.
\end{aligned}$$

where the first inequality holds since $z - \bar{V}_t^- + \tilde{V}_t^- \geq 0$ conditioning on \mathcal{E}_{low} , and the second inequality is by (11) again.

Recall that $\Pr(\mathcal{E}_{\text{opt}}) \geq F^2(-1)$. Inserting all of these back, we obtain

$$\begin{aligned}
& \sum_{t=1}^T (V^* - \bar{V}_t) + (\tilde{V}_t - V^{\pi_t^1}) \\
& \leq F^{-2}(-1) \left(8V_{\text{max}}\delta + \sum_{t=1}^T \left(\mathbb{E} \left[(\bar{V}_t - V^{\pi_t^0}) - (\tilde{V}_t - V^{\pi_t^1}) \right] + \mathbb{E} \left[(V^{\pi_t^0} - \bar{V}_t^-) - (V^{\pi_t^1} - \tilde{V}_t^-) \right] \right) \right) \\
& \leq F^{-2}(-1) \left(8V_{\text{max}}\delta + \sum_{t=1}^T \left((\bar{V}_t - V^{\pi_t^0}) - (\tilde{V}_t - V^{\pi_t^1}) + (V^{\pi_t^0} - \bar{V}_t^-) - (V^{\pi_t^1} - \tilde{V}_t^-) \right) \right) \\
& \quad + 2V_{\text{max}}\sqrt{\frac{T \log(1/\delta)}{2}}.
\end{aligned}$$

where the last inequality is the Hoeffding's inequality. \square

Given all these lemmas, we are ready to establish an upper bound of regret. We first note that, since $\pi_t^1 = \pi_{t-1}^0$ for all t , the regret incurred by π_t^1 for all t is equivalent to that incurred by π_t^0 for all t . Hence, it suffices to compute the regret incurred by π_t^0 for $t \in [T]$ and multiply it by two to get the total regret.

We start with the following regret decomposition:

$$\text{Regret}_T \leq \sum_{t=1}^T (V^* - V^{\pi_t^0}) = \underbrace{\sum_{t=1}^T (V^* - \bar{V}_t + \tilde{V}_t - V^{\pi_t^1})}_{(*)} + \sum_{t=1}^T (\bar{V}_t - V^{\pi_t^0} + V^{\pi_t^1} - \tilde{V}_t).$$

By Lemma B.16, we can further decompose (*) and obtain

$$\begin{aligned} \text{Regret}_T &\leq \frac{1}{\mathbb{F}^2(-1)} \underbrace{\sum_{t=1}^T \left(V_t^{\pi_t^0} - \bar{V}_t^- + \tilde{V}_t^- - V_t^{\pi_t^1} \right)}_{(i)} + \left(1 + \frac{1}{\mathbb{F}^2(-1)} \right) \underbrace{\sum_{t=1}^T \left(\bar{V}_t - V_t^{\pi_t^0} + V_t^{\pi_t^1} - \tilde{V}_t \right)}_{(ii)} \\ &\quad + \frac{2V_{\max}}{\mathbb{F}^2(-1)} \left(\sqrt{\frac{T \log(1/\delta)}{2}} + 4\delta \right). \end{aligned}$$

We note that both (i) and (ii) can be bounded by Lemma B.15:

$$\begin{aligned} (i), (ii) &\leq \tilde{O} \left(V_{\max} L_{\max} + V_{\max} \sqrt{HT} \right. \\ &\quad \left. + (\epsilon_{r,\xi} + \epsilon_{r,\eta}) \left(\sqrt{dT} + L_{\max} + \sqrt{HT} + \frac{\epsilon T \sqrt{\pi}}{2\sigma_r} \right) \right. \\ &\quad \left. + (\epsilon_{P,\xi} + \epsilon_{P,\eta} + \epsilon_\lambda) \left(H\sqrt{dT} + L_{\max} + \sqrt{HT} \right) \right) \end{aligned}$$

Inserting this back, we obtain

$$\begin{aligned} \text{Regret}_T &= \tilde{O} \left(V_{\max} L_{\max} + V_{\max} \sqrt{HT} \right. \\ &\quad \left. + (\epsilon_{r,\xi} + \epsilon_{r,\eta}) \left(\sqrt{dT} + L_{\max} + \sqrt{HT} + \frac{\epsilon T \sqrt{\pi}}{2\sigma_r} \right) \right. \\ &\quad \left. + (\epsilon_{P,\xi} + \epsilon_{P,\eta} + \epsilon_\lambda) \left(H\sqrt{dT} + L_{\max} + \sqrt{HT} \right) \right. \\ &\quad \left. + V_{\max} \sqrt{T} \right) \end{aligned}$$

To compute this quantity, we note the following asymptotic rate:

- $\epsilon_{r,\xi} + \epsilon_{r,\eta} = \tilde{O} \left(d\sqrt{\kappa + B^2} \right)$
- $\epsilon_{P,\xi} + \epsilon_{P,\eta} + \epsilon_\lambda = \tilde{O} \left(d^{5/2} H^{3/2} \sqrt{\kappa + B^2} \right)$
- $L_{\max} = \tilde{O} \left(d^6 H^4 (\kappa + B^2) \right)$

Hence, we have

$$\text{Regret}_T = \tilde{O} \left(\epsilon T \sqrt{d} + \sqrt{T} \cdot d^3 H^{5/2} \sqrt{\kappa + B^2} + d^{17/2} H^{11/2} (\kappa + B^2)^{3/2} \right).$$

B.5 BOUNDING NUMBER OF QUERIES

Recall that the number of queries are computed via

$$\text{Queries}_T = \sum_{t=1}^T Z_t = \sum_{t=1}^T Z_t \mathbb{1} \left\{ \mathbb{E}_{\theta_0, \theta_1 \sim \mathcal{N}(\hat{\theta}_{r,t}, \sigma_r^2 \Sigma_{t-1}^{-1})} \left[\left| (\phi(\tau_t^0) - \phi(\tau_t^1))^\top (\theta_0 - \theta_1) \right| \right] > \epsilon \right\}.$$

Notice that Cauchy-Schwarz inequality implies

$$\begin{aligned} \mathbb{E}_{\theta_0, \theta_1} \left| (\phi(\tau_t^0) - \phi(\tau_t^1))^\top (\theta_0 - \theta_1) \right| &\leq \|\phi(\tau_t^0) - \phi(\tau_t^1)\|_{\Sigma_{t-1}^{-1}} \mathbb{E}_{\theta_0, \theta_1} \|\theta_0 - \theta_1\|_{\Sigma_{t-1}} \\ &\leq \frac{1}{\sqrt{\lambda}} \|\phi(\tau_t^0) - \phi(\tau_t^1)\|_2 \mathbb{E}_{\theta_0, \theta_1} \|\theta_0 - \theta_1\|_{\Sigma_{t-1}} \leq \frac{2}{\sqrt{\lambda}} \mathbb{E}_{\theta_0, \theta_1} \|\theta_0 - \theta_1\|_{\Sigma_{t-1}} \leq \frac{2\sqrt{2\sigma_r^2 d}}{\sqrt{\lambda}} \end{aligned}$$

where the second inequality is by $\|\Sigma_{t-1}^{-1}\|_2 = 1/\lambda$, and the last step is by the fact that $\theta_0 - \theta_1$ follows $\mathcal{N}(0, 2\sigma_r^2 \Sigma_{t-1}^{-1})$ and Lemma D.3. We denote $\zeta := 2\sigma_r \sqrt{2d/\lambda}$ for the ease of notation. Then, we

have

$$\begin{aligned} \text{Queries}_T &= \sum_{t=1}^T Z_t \mathbb{1} \left\{ \min \left\{ \zeta, \mathbb{E}_{\theta_0, \theta_1} \left[\left| (\phi(\tau_t^0) - \phi(\tau_t^1))^\top (\theta_0 - \theta_1) \right| \right] \right\} > \epsilon \right\} \\ &\leq \sum_{t=1}^T Z_t \mathbb{1} \left\{ \min \left\{ 1, \mathbb{E}_{\theta_0, \theta_1} \left[\left| (\phi(\tau_t^0) - \phi(\tau_t^1))^\top (\theta_0 - \theta_1) \right| \right] \right\} > \epsilon/\zeta \right\} \end{aligned}$$

Applying Cauchy-Schwarz inequality, we have

$$\text{Queries}_T \leq \sum_{t=1}^T Z_t \mathbb{1} \left\{ \min \left\{ 1, \mathbb{E}_{\theta_0, \theta_1} \left[\left((\phi(\tau_t^0) - \phi(\tau_t^1))^\top (\theta_0 - \theta_1) \right)^2 \right] \right\} > \epsilon^2/\zeta^2 \right\}.$$

Let u_0, u_1 denotes two independent standard Gaussian variables with zero mean and identity covariance matrix. Then, $\theta_0 - \theta_1$ has the same joint distribution as $\sigma_r \Sigma_{t-1}^{-1/2} (u_0 - u_1)$. Hence, we can rewrite the expectation in the indicator as

$$\mathbb{E}_{\theta_0, \theta_1} \left[\left((\phi(\tau_t^0) - \phi(\tau_t^1))^\top (\theta_0 - \theta_1) \right)^2 \right] = \mathbb{E}_{u_0, u_1} \left[\left((\phi(\tau_t^0) - \phi(\tau_t^1))^\top \sigma_r \Sigma_{t-1}^{-1/2} (u_0 - u_1) \right)^2 \right]$$

Furthermore, we have

$$\begin{aligned} &\mathbb{E}_{u_0, u_1} \left[\left((\phi(\tau_t^0) - \phi(\tau_t^1))^\top \sigma_r \Sigma_{t-1}^{-1/2} (u_0 - u_1) \right)^2 \right] \\ &= \mathbb{E}_{u_0, u_1} \left[(\phi(\tau_t^0) - \phi(\tau_t^1))^\top \sigma_r \Sigma_{t-1}^{-1/2} (u_0 - u_1) (u_0 - u_1)^\top \sigma_r \Sigma_{t-1}^{-1/2} (\phi(\tau_t^0) - \phi(\tau_t^1)) \right] \\ &= (\phi(\tau_t^0) - \phi(\tau_t^1))^\top \sigma_r \Sigma_{t-1}^{-1/2} \mathbb{E}_{u_0, u_1} \left[(u_0 - u_1) (u_0 - u_1)^\top \right] \sigma_r \Sigma_{t-1}^{-1/2} (\phi(\tau_t^0) - \phi(\tau_t^1)). \end{aligned}$$

For the expectation in the middle, we have

$$\mathbb{E}_{u_0, u_1} \left[(u_0 - u_1) (u_0 - u_1)^\top \right] = \mathbb{E}[u_0 u_0^\top] + \mathbb{E}[u_1 u_1^\top] = 2I$$

where we have used the fact that $\mathbb{E}[u_0 u_1^\top] = 0$ by independence. Therefore, we have

$$\begin{aligned} &\mathbb{E}_{\theta_0, \theta_1} \left[\left((\phi(\tau_t^0) - \phi(\tau_t^1))^\top (\theta_0 - \theta_1) \right)^2 \right] \\ &= (\phi(\tau_t^0) - \phi(\tau_t^1))^\top \sigma_r \Sigma_{t-1}^{-1/2} (2I) \sigma_r \Sigma_{t-1}^{-1/2} (\phi(\tau_t^0) - \phi(\tau_t^1)) \\ &= 2\sigma_r^2 \|\phi(\tau_t^0) - \phi(\tau_t^1)\|_{\Sigma_{t-1}^{-1}}^2. \end{aligned}$$

Inserting this back, we obtain

$$\begin{aligned} \text{Queries}_T &\leq \sum_{t=1}^T Z_t \mathbb{1} \left\{ \min \left\{ 1, 2\sigma_r^2 \|\phi(\tau_t^0) - \phi(\tau_t^1)\|_{\Sigma_{t-1}^{-1}}^2 \right\} > \epsilon^2/\zeta^2 \right\} \\ &\leq \sum_{t=1}^T Z_t \mathbb{1} \left\{ \min \left\{ 1, \|\phi(\tau_t^0) - \phi(\tau_t^1)\|_{\Sigma_{t-1}^{-1}}^2 \right\} > \frac{\epsilon^2}{2\zeta^2\sigma_r^2} \right\} \\ &\leq \frac{2\zeta^2\sigma_r^2}{\epsilon^2} \sum_{t=1}^T Z_t \min \left\{ 1, \|\phi(\tau_t^0) - \phi(\tau_t^1)\|_{\Sigma_{t-1}^{-1}}^2 \right\} \\ &\leq \frac{2\zeta^2\sigma_r^2}{\epsilon^2} \cdot 2d \log \left(\frac{\lambda + 4T}{\lambda} \right) \end{aligned}$$

where the last step is Lemma D.6. Plugging the value of ζ and all other variables, we obtain

$$\text{Queries}_T \leq \frac{32\sigma_r^4 d^2}{\lambda \epsilon^2} \log \left(\frac{\lambda + 4T}{\lambda} \right) = \tilde{O} \left(\frac{d^4 (\kappa + B^2)^2}{\epsilon^2} \right).$$

C PROOF OF THEOREM 5.4

We first present some supporting results in Appendix C.1. Then, we prove the upper bound of Bayesian regret in Appendix C.2 and the number of queries in Appendix C.3.

C.1 SUPPORTING LEMMAS

In Appendix C.1.1, we establish some supporting results for a probability estimation problem in the frequentist setting. In Appendix C.1.2, we adapt these results to the Bayesian setting. The Bayesian results will be heavily used later in the proof of Theorem 5.4 in Appendices C.2 and C.3.

C.1.1 SUPPORTING RESULTS FROM FREQUENTIST SETTING

Consider a conditional probability estimation problem. Let \mathcal{X} and \mathcal{Y} be the instance space and the target space, respectively. Let $\mathcal{F} : (\mathcal{X} \times \mathcal{Y}) \rightarrow \mathbb{R}$ be a function class. We are given a dataset $D := \{(x_i, y_i)\}_{i=1}^n$ where $x_i \sim \mathcal{D}_i$ and $y_i \sim f^*(x_i, \cdot)$. We assume $f^* \in \mathcal{F}$. Regarding the data generation process, we assume the data distribution \mathcal{D}_i is history-dependent, i.e., x_i can depend on the previous samples: $x_1, y_1, \dots, x_{i-1}, y_{i-1}$ for any $i \in [n]$. Our goal is to estimate the true conditional probability f^* using the dataset D .

At a high level, this problem is designed to capture both the reward learning and the model learning problems in the RL setting. Specifically, in the reward learning problem, we will instantiate $\mathcal{X} = \mathcal{S} \times \mathcal{A}$ and $\mathcal{Y} = \{0, 1\}$, where we recall that \mathcal{S} and \mathcal{A} are the state space and the action space, respectively, and the preference feedback is binary. In the model learning problem, we can instantiate $\mathcal{X} = \mathcal{S} \times \mathcal{A}$ and $\mathcal{Y} = \mathcal{S}$. We abstract these two problems into this conditional probability estimation problem to make the analysis more concise. However, one caveat is that all we derived in this section are *frequentist* results, while we are considering *Bayesian* RL. Thus, the results are not directly applicable. In Appendix C.1.2, we will adapt these frequentist results to the Bayesian setting so that they can be applied.

Now we establish some important results for this problem. First, we have the following lemma, which is a consequence of Lemma D.10.

Lemma C.1 (MLE generalization bound). *Fix $\delta \in (0, 1)$. Follow the setting stated above. Let \hat{f} be the maximum likelihood estimator:*

$$\hat{f} = \arg \max_{f \in \mathcal{F}} \sum_{i=1}^n \log f(x_i, y_i).$$

Define the version space:

$$\mathcal{V}^{\mathcal{F}} = \left\{ f \in \mathcal{F} : \sum_{i=1}^n d_{\text{TV}}^2 \left(\hat{f}(x_i, \cdot), f(x_i, \cdot) \right) \leq \beta_{\mathcal{F}}(n) \right\}$$

where $\beta_{\mathcal{F}}(n) := 98 \log(2N_{[]}((n|\mathcal{Y}|)^{-1}, \mathcal{F}, \|\cdot\|_{\infty})/\delta)$. Then, the following holds

- (i) $f^* \in \mathcal{V}^{\mathcal{F}}$ with probability at least $1 - \delta$,
- (ii) for any $f, f' \in \mathcal{V}^{\mathcal{F}}$, it holds that

$$\sum_{i=1}^n d_{\text{TV}}^2 \left(f(x_i, \cdot), f'(x_i, \cdot) \right) \leq 4\beta_{\mathcal{F}}(n).$$

Proof of Lemma C.1. We first construct an auxiliary version space $\tilde{\mathcal{V}}$ as follows

$$\tilde{\mathcal{V}}^{\mathcal{F}} = \left\{ f \in \mathcal{F} : \sum_{i=1}^n \mathbb{E}_{x \sim \mathcal{D}_i} d_{\text{TV}}^2 \left(\hat{f}(x, \cdot), f(x, \cdot) \right) \leq 10 \log \left(N_{[]} \left((n|\mathcal{Y}|)^{-1}, \mathcal{F}, \|\cdot\|_{\infty} \right) / \delta \right) \right\}$$

By Lemma D.10, $f^* \in \tilde{\mathcal{V}}^{\mathcal{F}}$ with probability at least $1 - \delta$. To prove (i), we will show that whenever $f^* \in \tilde{\mathcal{V}}^{\mathcal{F}}$, we have $f^* \in \mathcal{V}^{\mathcal{F}}$ as well with high probability. Let $\mathcal{F}_{[]} denote an $(n|\mathcal{Y}|)^{-1}$ -bracket of $\mathcal{F}$$

with respect to $\|\cdot\|_\infty$. Then for all $f_{\square} \in \mathcal{F}_{\square}$, the following holds with probability at least $1 - \delta$ by Lemma D.4 and the union bound on \mathcal{F}_{\square} ,

$$\begin{aligned} & \sum_{i=1}^n d_{\text{TV}}^2(f_{\square}(x_i, \cdot), f^*(x_i, \cdot)) \\ & \leq 2 \sum_{i=1}^n \mathbb{E}_{x \sim \mathcal{D}_i} d_{\text{TV}}^2(f_{\square}(x, \cdot), f^*(x, \cdot)) + 4 \log \left(N_{\square} \left((n|\mathcal{Y}|)^{-1}, \mathcal{F}, \|\cdot\|_\infty \right) / \delta \right). \end{aligned} \quad (12)$$

Now for any $f \in \mathcal{F}$, there must exist $f_{\square} \in \mathcal{F}_{\square}$ such that $\|f - f_{\square}\|_\infty \leq (n|\mathcal{Y}|)^{-1}$, which implies the following

$$\sum_{i=1}^n d_{\text{TV}}^2(f(x_i, \cdot), f_{\square}(x_i, \cdot)) \leq \sum_{i=1}^n |\mathcal{Y}|^2 \|f - f_{\square}\|_\infty^2 \leq \sum_{i=1}^n n^{-2} = 1/n, \quad (13)$$

$$\sum_{i=1}^n \mathbb{E}_{x \sim \mathcal{D}_i} d_{\text{TV}}^2(f(x, \cdot), f_{\square}(x, \cdot)) \leq \sum_{i=1}^n |\mathcal{Y}|^2 \|f - f_{\square}\|_\infty^2 \leq \sum_{i=1}^n n^{-2} = 1/n \quad (14)$$

Hence, for all $f \in \mathcal{F}$, we have

$$\begin{aligned} & \sum_{i=1}^n d_{\text{TV}}^2(f(x_i, \cdot), f^*(x_i, \cdot)) \\ & \leq 2 \sum_{i=1}^n d_{\text{TV}}^2(f_{\square}(x_i, \cdot), f^*(x_i, \cdot)) + 2 \sum_{i=1}^n d_{\text{TV}}^2(f(x_i, \cdot), f_{\square}(x_i, \cdot)) \\ & \leq 2 \sum_{i=1}^n d_{\text{TV}}^2(f_{\square}(x_i, \cdot), f^*(x_i, \cdot)) + 2/n \\ & \leq 4 \sum_{i=1}^n \mathbb{E}_{x \sim \mathcal{D}_i} d_{\text{TV}}^2(f_{\square}(x, \cdot), f^*(x, \cdot)) + 8 \log \left(N_{\square} \left((n|\mathcal{Y}|)^{-1}, \mathcal{F}, \|\cdot\|_\infty \right) / \delta \right) + 2/n \\ & \leq 8 \sum_{i=1}^n \mathbb{E}_{x \sim \mathcal{D}_i} d_{\text{TV}}^2(f(x, \cdot), f^*(x, \cdot)) + 8 \log \left(N_{\square} \left((n|\mathcal{Y}|)^{-1}, \mathcal{F}, \|\cdot\|_\infty \right) / \delta \right) + 10/n. \end{aligned}$$

with probability at least $1 - \delta$. Here the second inequality uses (13), the third uses (12), and the last uses (14). Therefore, for any possible value of the estimator \hat{f} , conditioning on $f^* \in \tilde{\mathcal{V}}^{\mathcal{F}}$, we have

$$\begin{aligned} & \sum_{i=1}^n d_{\text{TV}}^2(\hat{f}(x_i, \cdot), f^*(x_i, \cdot)) \\ & \leq 8 \sum_{i=1}^n \mathbb{E}_{x \sim \mathcal{D}_i} d_{\text{TV}}^2(\hat{f}(x, \cdot), f^*(x, \cdot)) + 8 \log \left(N_{\square} \left((n|\mathcal{Y}|)^{-1}, \mathcal{F}, \|\cdot\|_\infty \right) / \delta \right) + 10/n \\ & \leq 88 \log \left(N_{\square} \left((n|\mathcal{Y}|)^{-1}, \mathcal{F}, \|\cdot\|_\infty \right) / \delta \right) + 10/n \\ & \leq 98 \log \left(N_{\square} \left((n|\mathcal{Y}|)^{-1}, \mathcal{F}, \|\cdot\|_\infty \right) / \delta \right) \end{aligned}$$

with probability at least $1 - \delta$. Here the second inequality is by the definition of $\tilde{\mathcal{V}}^{\mathcal{F}}$. The last inequality holds since it is reasonable to assume that the bracketing number is at least some constant so that $\log(N_{\square}((n|\mathcal{Y}|)^{-1}, \mathcal{F}, \|\cdot\|_\infty)/\delta) > 1 \geq 1/n$.

The above means whenever $f^* \in \tilde{\mathcal{V}}^{\mathcal{F}}$, we have $f^* \in \mathcal{V}^{\mathcal{F}}$ as well with probability at least $1 - \delta$. Since $f^* \in \tilde{\mathcal{V}}^{\mathcal{F}}$ with probability at least $1 - \delta$, we have $f^* \in \mathcal{V}^{\mathcal{F}}$ with probability at least $1 - 2\delta$. Adjusting δ completes the proof of (i).

For (ii), we have

$$\begin{aligned} \sum_{i=1}^n d_{\text{TV}}^2\left(f(x, \cdot), f'(x, \cdot)\right) &\leq \sum_{i=1}^n \left(d_{\text{TV}}\left(f(x, \cdot), \widehat{f}(x, \cdot)\right) + d_{\text{TV}}\left(\widehat{f}(x, \cdot), f'(x, \cdot)\right)\right)^2 \\ &\leq \sum_{i=1}^n 2d_{\text{TV}}^2\left(f(x, \cdot), \widehat{f}(x, \cdot)\right) + \sum_{i=1}^n 2d_{\text{TV}}^2\left(\widehat{f}(x, \cdot), f'(x, \cdot)\right) \\ &\leq 4\beta_{\mathcal{F}}(n) \end{aligned}$$

where the second inequality is by the fact that $(a + b)^2 \leq 2a^2 + 2b^2$ for any a and b , and the third inequality is by the definition of $\mathcal{V}^{\mathcal{F}}$. \square

Next, we further assume that the function class \mathcal{F} is parameterized by a function class \mathcal{G} via a link function Φ . The reason for this assumption is that we want to capture the structure of the reward learning problem, where the feedback generating distribution is parameterized by the reward function via a link function Φ .

Assumption C.2 (Binary label and function parameterization.). *Assume $\mathcal{Y} = \{0, 1\}$ is binary, and there is a function class $\mathcal{G} \subseteq \mathcal{X} \rightarrow [0, G]$ that parameterizes \mathcal{F} via a link function Φ . Specifically, we assume*

$$\mathcal{F} = \left\{ f(x, 0) = \Phi(g(x)), f(x, 1) = 1 - \Phi(g(x)) : g \in \mathcal{G} \right\},$$

where we further assume Φ satisfies Assumption 3.1. For any $f \in \mathcal{F}$, let g_f denote the function g that parameterizes f . We define $\widehat{g} := g_{\widehat{f}}$ and $g^* := g_{f^*}$.

As a preliminary note, the function class \mathcal{G} will actually correspond to the function class $\widetilde{\mathcal{R}}$ (3) later in the proof.

We should emphasize again that, although the function class \mathcal{G} aims to capture the structure of the reward learning problem, the function class \mathcal{F} will be used to capture both the reward learning and the model learning problems. When it is applied to the model learning problem, we can simply ignore the function class \mathcal{G} and any results that are related to \mathcal{G} . On the other hand, when it is applied to the reward learning problem, such a function class \mathcal{G} will be helpful to derive some results in the reward function itself.

Before diving into the analysis of the function class \mathcal{G} , we first introduce the following lemma which shows that the bracketing numbers of \mathcal{F} and \mathcal{G} are bounded by each other.

Lemma C.3. *Under Assumption C.2, for any $\omega > 0$, we have*

$$N_{[]}(\omega, \mathcal{F}, \|\cdot\|_{\infty}) \leq N_{[]}(\bar{\kappa}\omega, \mathcal{G}, \|\cdot\|_{\infty}) \quad \text{and} \quad N_{[]}(\kappa\omega, \mathcal{G}, \|\cdot\|_{\infty}) \leq N_{[]}(\omega, \mathcal{F}, \|\cdot\|_{\infty}).$$

Proof. For any $f, f' \in \mathcal{F}$, we assume they are parameterized by $g := g_f$ and $g' := g_{f'}$, respectively. Then, we have

$$\sup_{x,y} |f(x, y) - f'(x, y)| = \sup_x |\Phi(g(x)) - \Phi(g'(x))| \leq \bar{\kappa}^{-1} \sup_x |g(x) - g'(x)|$$

where the inequality is Lemma D.1. Hence, if we have a $\bar{\kappa}\omega$ -bracket of \mathcal{G} in the infinite norm, then we have an ω -bracket of \mathcal{F} in the infinite norm. This proves the first claim. For the second claim, by Lemma D.1, we have

$$\sup_x |g(x) - g'(x)| \leq \kappa \sup_x |\Phi(g(x)) - \Phi(g'(x))| = \kappa \sup_{x,y} |f(x, y) - f'(x, y)|.$$

This proves the second claim. \square

Then, Lemma C.1 leads to similar results with a version space constructed on \mathcal{G} .

Corollary C.4. *Under Assumption C.2, define*

$$\mathcal{V}^{\mathcal{G}} = \left\{ g \in \mathcal{G} : \sum_{i=1}^n \left(\widehat{g}(x_i) - g(x_i)\right)^2 \leq \beta_{\mathcal{G}}(n) \right\}$$

where $\beta_{\mathcal{G}}(n) := 98\kappa^2 \log(2N_{[]}(\bar{\kappa}(n|\mathcal{Y})^{-1}, \mathcal{G}, \|\cdot\|_{\infty})/\delta)$ and we denote $\widehat{g} := g_{\widehat{f}}$ as the function that parameterizes the maximum likelihood estimator \widehat{f} . Then, the following holds

(1) $g^* \in \mathcal{V}^{\mathcal{G}}$ with probability at least $1 - \delta$

(2) for any $g, g' \in \mathcal{V}^{\mathcal{G}}$, we have

$$\sum_{i=1}^n \left(g(x_i) - g'(x_i) \right)^2 \leq 4\beta_{\mathcal{G}}(n).$$

Proof of Corollary C.4. To prove (i), we claim that for any $f \in \mathcal{V}^{\mathcal{F}}$ (defined in Lemma C.1), we have $g_f \in \mathcal{V}^{\mathcal{G}}$ as well. To see this, we note that

$$\begin{aligned} \sum_{i=1}^n \mathbb{E}_{x \sim \mathcal{D}_i} \left(\widehat{g}(x) - g_f(x) \right)^2 &\leq \kappa^2 \sum_{i=1}^n \mathbb{E}_{x \sim \mathcal{D}_i} \left| \Phi(\widehat{g}(x)) - \Phi(g_f(x)) \right|^2 \\ &= \kappa^2 \sum_{i=1}^n \mathbb{E}_{x \sim \mathcal{D}_i} \left| \widehat{f}(x, 0) - f(x, 0) \right|^2 = \kappa^2 \sum_{i=1}^n \mathbb{E}_{x \sim \mathcal{D}_i} d_{\text{TV}}^2 \left(\widehat{f}(x, \cdot), f(x, \cdot) \right) \end{aligned}$$

where the inequality is Lemma D.1, and the last equality holds since we assume $\mathcal{Y} = \{0, 1\}$ is binary. Hence, for any $f \in \mathcal{V}^{\mathcal{F}}$, it holds that

$$\begin{aligned} \sum_{i=1}^n \mathbb{E}_{x \sim \mathcal{D}_i} \left(\widehat{g}(x) - g_f(x) \right)^2 &\leq \kappa^2 \sum_{i=1}^n \mathbb{E}_{x \sim \mathcal{D}_i} d_{\text{TV}}^2 \left(\widehat{f}(x, \cdot), f(x, \cdot) \right) \\ &\leq \kappa^2 \beta_{\mathcal{F}}(n) \\ &= \kappa^2 \cdot 98 \log(2N_{\square}((n|\mathcal{Y}|)^{-1}, \mathcal{F}, \|\cdot\|_{\infty})/\delta) \end{aligned}$$

By Lemma C.3, we have

$$\log(2N_{\square}((n|\mathcal{Y}|)^{-1}, \mathcal{F}, \|\cdot\|_{\infty})/\delta) \leq \log(2N_{\square}(\bar{\kappa}(n|\mathcal{Y}|)^{-1}, \mathcal{G}, \|\cdot\|_{\infty})/\delta).$$

Hence, we have

$$\sum_{i=1}^n \mathbb{E}_{x \sim \mathcal{D}_i} \left(\widehat{g}(x) - g_f(x) \right)^2 \leq \beta_{\mathcal{G}}(n)$$

which implies $g_f \in \mathcal{V}^{\mathcal{G}}$. Therefore, whenever $f^* \in \mathcal{V}^{\mathcal{F}}$, we have $g_{f^*} \in \mathcal{V}^{\mathcal{G}}$ as well. Since $f^* \in \mathcal{V}^{\mathcal{F}}$ with probability at least $1 - \delta$, we have $g_{f^*} \in \mathcal{V}^{\mathcal{G}}$ with probability at least $1 - \delta$. This completes the proof of (i).

Now we prove (2). For any $g, g' \in \mathcal{V}^{\mathcal{G}}$, we have

$$\sum_{i=1}^n \left(g(x_i) - g'(x_i) \right)^2 \leq 2 \sum_{i=1}^n \left(\widehat{g}(x_i) - g(x_i) \right)^2 + 2 \sum_{i=1}^n \left(\widehat{g}(x_i) - g'(x_i) \right)^2 \leq 4\beta_{\mathcal{G}}(n)$$

where the first inequality is by the fact that $(a + b)^2 \leq 2a^2 + 2b^2$ for any a and b . \square

C.1.2 ADAPTING RESULTS INTO BAYESIAN SETTING

Now we change the setting defined in Appendix C.1.1 into a Bayesian online setting. The reason for this adaptation is to make these results applicable to the Bayesian RL setting. We formally define the Bayesian online conditional probability estimation problem below.

Let \mathcal{X} and \mathcal{Y} be the instance space and the target space, respectively. Let $\mathcal{F} : (\mathcal{X} \times \mathcal{Y}) \rightarrow \mathbb{R}$ be a function class. The interaction proceeds for T rounds. At each round $t \in [T]$, we observe an instance $x_t \in \mathcal{X}$ and need to outputs a function $f_t \in \mathcal{F}$ as our prediction. Then, the label $y_t \in \mathcal{Y}$ is revealed. We assume that $x_i \sim \mathcal{D}_i$ for some distribution \mathcal{D}_i and $y_i \sim f^*(x, \cdot)$. We assume the true conditional distribution f^* is sampled from some known prior distribution $\rho \in \Delta(\mathcal{F})$ at the beginning. Regarding the data generation process, we assume the data distribution \mathcal{D}_i is history-dependent, i.e., x_i can depend on the previous samples: $x_1, y_1, \dots, x_{i-1}, y_{i-1}$ for any $i \in [n]$.

Denote \mathcal{H}_t as the history up to round t , i.e., $\mathcal{H}_t = \{x_1, f_1, y_1, x_2, f_2, y_2, \dots, x_t, f_t, y_t\}$. Define the maximum likelihood estimator of f^* on the dataset $\{(x_s, y_s)\}_{s=1}^{t-1}$ as \widehat{f}_t , i.e.,

$$\widehat{f}_t = \arg \max_{f \in \mathcal{F}} \sum_{s=1}^{t-1} \log f(x_s, y_s).$$

Similar to the previous section, we will also consider the presence of a function class \mathcal{G} (Assumption C.2) in the analysis presented in this section. It's worth noting that although Assumption C.2 is initially introduced in the frequentist setting, when we mention that Assumption C.2 holds in this section, we are actually considering the Bayesian setting.

Lemma C.5. *It holds that*

$$\sum_{t=1}^T \mathbb{E}_{\mathcal{H}_{t-1}} \left[\mathbb{E}_{f, f'} \left[d_{\text{TV}}(f(x_t), f'(x_t)) \mid \mathcal{H}_{t-1} \right] \right] \leq 6\sqrt{T\beta'_{\mathcal{F}}(T)} \cdot \dim_1(\mathcal{F}, 1/T) \cdot \log(T) \quad (15)$$

and

$$\sum_{t=1}^T \mathbb{E}_{\mathcal{H}_{t-1}} \left[\mathbb{E}_{f, f'} \left[d_{\text{TV}}^2(f(x_t), f'(x_t)) \mid \mathcal{H}_{t-1} \right] \right] \leq 4\beta'_{\mathcal{F}}(T) + 2 \quad (16)$$

where $\beta'_{\mathcal{F}}(t) := 98 \log(2TN_{\square}((t|\mathcal{Y}|)^{-1}, \mathcal{F}, \|\cdot\|_{\infty}))$, and f, f' are sampled from the posterior of f^* conditioning on \mathcal{H}_{t-1} in the inner conditional expectation. Moreover, if \mathcal{F} is parameterized by \mathcal{G} (i.e., if Assumption C.2 holds), we have

$$\sum_{t=1}^T \mathbb{E}_{\mathcal{H}_{t-1}} \left[\mathbb{E}_{g, g'} \left[|g(x_t) - g'(x_t)| \mid \mathcal{H}_{t-1} \right] \right] \leq 6\sqrt{T\beta'_{\mathcal{G}}(T)} \cdot \dim_1(\mathcal{G}, 1/T) \cdot \log(GT) \quad (17)$$

and

$$\sum_{t=1}^T \mathbb{E}_{\mathcal{H}_{t-1}} \left[\mathbb{E}_{g, g'} \left[(g(x_t) - g'(x_t))^2 \mid \mathcal{H}_{t-1} \right] \right] \leq 4\beta'_{\mathcal{G}}(T) + 2 \quad (18)$$

where $\beta'_{\mathcal{G}}(t) := 98\kappa^2 \log(2GTN_{\square}(\bar{\kappa}(t|\mathcal{Y}|)^{-1}, \mathcal{G}, \|\cdot\|_{\infty}))$, and g, g' are sampled from the posterior of g^* conditioning on \mathcal{H}_{t-1} in the inner conditional expectation.

Proof of Lemma C.5. We will only prove (15) and (16) since the proof of (17) and (18) is almost identical.

Proof of (15). We define the version space at round t for $t \in [T]$ as follows

$$\mathcal{V}_t^{\mathcal{F}} = \left\{ f \in \mathcal{F} : \sum_{s=1}^{t-1} d_{\text{TV}}^2(\hat{f}_t(x_s, \cdot), f(x_s, \cdot)) \leq \beta_{\mathcal{F}}(t-1) \right\}$$

where \hat{f}_t is the maximum likelihood estimator on the dataset $\{(x_s, y_s)\}_{s=1}^{t-1}$. This construction aims to mimic the version space $\mathcal{V}^{\mathcal{F}}$ defined in Lemma C.1, which enable us to apply Lemma C.1.

We first split the left side of (15) into two cases based on whether f and f' belong to $\mathcal{V}_t^{\mathcal{F}}$:

$$\begin{aligned} & \sum_{t=1}^T \mathbb{E}_{\mathcal{H}_{t-1}} \left[\mathbb{E}_{f, f'} \left[d_{\text{TV}}(f(x_t), f'(x_t)) \mid \mathcal{H}_{t-1} \right] \right] \\ &= \sum_{t=1}^T \mathbb{E}_{\mathcal{H}_{t-1}} \left[\mathbb{E}_{f, f'} \left[\mathbb{1}\{f \in \mathcal{V}_t^{\mathcal{F}} \text{ and } f' \in \mathcal{V}_t^{\mathcal{F}}\} d_{\text{TV}}(f(x_t), f'(x_t)) \mid \mathcal{H}_{t-1} \right] \right. \\ & \quad \left. + \mathbb{E}_{f, f'} \left[\mathbb{1}\{f \notin \mathcal{V}_t^{\mathcal{F}} \text{ or } f' \notin \mathcal{V}_t^{\mathcal{F}}\} d_{\text{TV}}(f(x_t), f'(x_t)) \mid \mathcal{H}_{t-1} \right] \right] \\ &=: \mathbf{T}_1 + \mathbf{T}_2. \end{aligned}$$

To bound T_2 , since the total variation distance is upper bounded by 1, we have

$$\begin{aligned}
T_2 &\leq \sum_{t=1}^T \mathbb{E}_{\mathcal{H}_{t-1}} \left[\mathbb{E}_{f, f'} \left[\mathbb{1}\{f \notin \mathcal{V}_t^{\mathcal{F}} \text{ or } f' \notin \mathcal{V}_t^{\mathcal{F}}\} \mid \mathcal{H}_{t-1} \right] \right] \\
&= \sum_{t=1}^T \mathbb{E}_{\mathcal{H}_{t-1}} \left[\mathbb{E}_{f, f'} \left[\mathbb{1}\{f \notin \mathcal{V}_t^{\mathcal{F}} \text{ or } f' \notin \mathcal{V}_t^{\mathcal{F}}\} \mid \mathcal{H}_{t-1} \right] \right] \\
&\leq \sum_{t=1}^T \mathbb{E}_{\mathcal{H}_{t-1}} \left[2 \mathbb{E}_{f^*} \left[\mathbb{1}\{f^* \notin \mathcal{V}_t^{\mathcal{F}}\} \mid \mathcal{H}_{t-1} \right] \right] \\
&= 2 \sum_{t=1}^T \mathbb{E}_{\mathcal{H}_{t-1}, f^*} \left[\mathbb{1}\{f^* \notin \mathcal{V}_t^{\mathcal{F}}\} \right] \\
&= 2 \sum_{t=1}^T \mathbb{E}_{f^*} \left[\mathbb{E}_{\mathcal{H}_{t-1}} \left[\mathbb{1}\{f^* \notin \mathcal{V}_t^{\mathcal{F}}\} \mid f^* \right] \right] \\
&\leq 2T\delta
\end{aligned}$$

where the second inequality holds by the union bound and the condition that f and f' have the same posterior distributions as f^* conditioning on \mathcal{H}_{t-1} . The last two equalities holds by the law of probability. The last inequality is Lemma C.1.

Now let's bound T_1 . We first notice that we can replace the expectation with the supremum:

$$T_1 \leq \sum_{t=1}^T \mathbb{E}_{\mathcal{H}_{t-1}} \left[\sup_{f, f' \in \mathcal{V}_t^{\mathcal{F}}} d_{\text{TV}}(f(x_t), f'(x_t)) \right].$$

Applying Lemma D.12, we have

$$T_1 \leq 4\sqrt{T\beta_{\mathcal{F}}(t-1)} \cdot \dim_1(\mathcal{F}, 1/T) \cdot \log T$$

Combining the upper bounds of T_1 and T_2 and setting $\delta = 1/T$, we complete the proof of (15).

Proof of (16). We split the left side of (16) into two cases based on whether f and f' belong to $\mathcal{V}_t^{\mathcal{F}}$:

$$\begin{aligned}
&\sum_{t=1}^T \mathbb{E}_{\mathcal{H}_{t-1}} \left[\mathbb{E}_{f, f'} \left[d_{\text{TV}}^2(f(x_t), f'(x_t)) \mid \mathcal{H}_{t-1} \right] \right] \\
&\leq \sum_{t=1}^T \mathbb{E}_{\mathcal{H}_{t-1}} \left[\mathbb{E}_{f, f'} \left[\mathbb{1}\{f \in \mathcal{V}_t^{\mathcal{F}} \text{ and } f' \in \mathcal{V}_t^{\mathcal{F}}\} d_{\text{TV}}^2(f(x_t), f'(x_t)) \mid \mathcal{H}_{t-1} \right] \right] \\
&\quad + \sum_{t=1}^T \mathbb{E}_{\mathcal{H}_{t-1}} \left[\mathbb{E}_{f, f'} \left[\mathbb{1}\{f \notin \mathcal{V}_t^{\mathcal{F}} \text{ or } f' \notin \mathcal{V}_t^{\mathcal{F}}\} d_{\text{TV}}^2(f(x_t), f'(x_t)) \mid \mathcal{H}_{t-1} \right] \right] \\
&=: T_3 + T_4.
\end{aligned}$$

Following a similar argument as in the proof of (15) above, we have $T_4 \leq 2T\delta$. For T_3 , by the definition of $\mathcal{V}_t^{\mathcal{F}}$, we directly have $T_3 \leq 4\beta_{\mathcal{F}}(T)$. Setting $\delta = 1/T$, we complete the proof of (16).

Proof of (17) and (18). An identical argument can be applied to prove (17) and (18) leveraging Corollary C.4. \square

Lemma C.6. Under Assumption C.2, it holds that

$$\begin{aligned}
&\sum_{t=1}^T \mathbb{E}_{\mathcal{H}_{t-1}} \left[\mathbb{1} \left\{ \mathbb{E}_{g, g'} \left[|g(x_t) - g'(x_t)| \mid \mathcal{H}_{t-1} \right] > \epsilon \right\} \right] \\
&\leq \min \left\{ \frac{9\sqrt{T\beta'_{\mathcal{G}}(T)}}{\epsilon} \cdot \dim_1(\mathcal{G}, \epsilon/2), \frac{21\beta'_{\mathcal{G}}(T)}{\epsilon^2} \cdot \dim_2(\mathcal{G}, \epsilon/2) \right\}.
\end{aligned}$$

where $\beta'_{\mathcal{G}}(t) := 98\kappa^2 \log(2GTN_{\square}(\bar{\kappa}(t)|\mathcal{Y})^{-1}, \mathcal{G}, \|\cdot\|_{\infty})$. In the inner expectation, g, g' are sampled from the posterior of g^* conditioning on \mathcal{H}_{t-1} .

Proof of Lemma C.6. We define the version space at round t for $t \in [T]$ as follows

$$\mathcal{V}_t^{\mathcal{G}} = \left\{ g \in \mathcal{G} : \sum_{s=1}^{t-1} \left(\widehat{g}_t(x_s) - g(x_s) \right)^2 \leq \beta_{\mathcal{G}}(t-1) \right\}$$

where recall that \widehat{g}_t is defined as the function that parameterizes the maximum likelihood estimator on the dataset $\{(x_s, y_s)\}_{s=1}^{t-1}$ (i.e., $\widehat{g}_t = g_{\widehat{f}_t}$). This construction aims to mimic the version space $\mathcal{V}^{\mathcal{G}}$ defined in Corollary C.4, which enable us to apply Corollary C.4.

We first split the left side into two cases based on whether both g and g' belong to $\mathcal{V}_t^{\mathcal{G}}$:

$$\begin{aligned} & \sum_{t=1}^T \mathbb{E}_{\mathcal{H}_{t-1}} \left[\mathbb{1} \left\{ \mathbb{E}_{g, g'} \left[|g(x_t) - g'(x_t)| \mid \mathcal{H}_{t-1} \right] > \epsilon \right\} \right] \\ &= \sum_{t=1}^T \mathbb{E}_{\mathcal{H}_{t-1}} \left[\mathbb{1} \left\{ \mathbb{E}_{g, g'} \left[\mathbb{1} \{g \in \mathcal{V}_t^{\mathcal{G}} \text{ and } g' \in \mathcal{V}_t^{\mathcal{G}}\} |g(x_t) - g'(x_t)| \mid \mathcal{H}_{t-1} \right] \right. \right. \\ & \quad \left. \left. + \mathbb{E}_{g, g'} \left[\mathbb{1} \{g \notin \mathcal{V}_t^{\mathcal{G}} \text{ or } g' \notin \mathcal{V}_t^{\mathcal{G}}\} |g(x_t) - g'(x_t)| \mid \mathcal{H}_{t-1} \right] > \epsilon \right\} \right] \\ &\leq \sum_{t=1}^T \mathbb{E}_{\mathcal{H}_{t-1}} \left[\mathbb{1} \left\{ \mathbb{E}_{g, g'} \left[\mathbb{1} \{g \in \mathcal{V}_t^{\mathcal{G}} \text{ and } g' \in \mathcal{V}_t^{\mathcal{G}}\} |g(x_t) - g'(x_t)| \mid \mathcal{H}_{t-1} \right] > \epsilon/2 \right\} \right] \\ & \quad + \sum_{t=1}^T \mathbb{E}_{\mathcal{H}_{t-1}} \left[\mathbb{1} \left\{ \mathbb{E}_{g, g'} \left[\mathbb{1} \{g \notin \mathcal{V}_t^{\mathcal{G}} \text{ or } g' \notin \mathcal{V}_t^{\mathcal{G}}\} |g(x_t) - g'(x_t)| \mid \mathcal{H}_{t-1} \right] > \epsilon/2 \right\} \right] \\ &=: \mathsf{T}_1 + \mathsf{T}_2. \end{aligned}$$

Here the inequality is due to the fact that $\mathbb{1}\{a+b > c\} \leq \mathbb{1}\{a > c/2\} + \mathbb{1}\{b > c/2\}$ for any a, b, c . To bound T_2 , recalling that $g(\cdot) \in [0, G]$, then we have

$$\begin{aligned} \mathsf{T}_2 &\leq \sum_{t=1}^T \mathbb{E}_{\mathcal{H}_{t-1}} \left[\mathbb{1} \left\{ \mathbb{E}_{g, g'} \left[\mathbb{1} \{g \notin \mathcal{V}_t^{\mathcal{G}} \text{ or } g' \notin \mathcal{V}_t^{\mathcal{G}}\} \cdot G \mid \mathcal{H}_{t-1} \right] > \epsilon/2 \right\} \right] \\ &\leq \frac{2G}{\epsilon} \sum_{t=1}^T \mathbb{E}_{\mathcal{H}_{t-1}} \left[\mathbb{E}_{g, g'} \left[\mathbb{1} \{g \notin \mathcal{V}_t^{\mathcal{G}} \text{ or } g' \notin \mathcal{V}_t^{\mathcal{G}}\} \mid \mathcal{H}_{t-1} \right] \right] \\ &\leq \frac{2G}{\epsilon} \sum_{t=1}^T \mathbb{E}_{\mathcal{H}_{t-1}} \left[2 \mathbb{E}_{g^*} \left[\mathbb{1} \{g^* \notin \mathcal{V}_t^{\mathcal{G}}\} \mid \mathcal{H}_{t-1} \right] \right] \\ &= \frac{4G}{\epsilon} \sum_{t=1}^T \mathbb{E}_{g^*, \mathcal{H}_{t-1}} \left[\mathbb{1} \{g^* \notin \mathcal{V}_t^{\mathcal{G}}\} \right] \\ &= \frac{4G}{\epsilon} \sum_{t=1}^T \mathbb{E}_{g^*} \left[\mathbb{E}_{\mathcal{H}_{t-1}} \left[\mathbb{1} \{g^* \notin \mathcal{V}_t^{\mathcal{G}}\} \mid g^* \right] \right] \\ &\leq \frac{4GT\delta}{\epsilon}. \end{aligned}$$

where the second inequality is by the fact that $\mathbb{1}\{a > b\} \leq a/b$ for any $a \geq 0$ and $b > 0$, the third inequality holds since g and g' are identically distributed as g^* conditioning on \mathcal{H}_{t-1} , the two equalities holds by the law of probability, and the last inequality is Corollary C.4.

Now let us bound T_1 . We first notice that we can replace the expectation with the supremum:

$$\mathsf{T}_1 \leq \sum_{t=1}^T \mathbb{E}_{\mathcal{H}_{t-1}} \left[\mathbb{1} \left\{ \sup_{g, g' \in \mathcal{V}_t^{\mathcal{G}}} |g(x_t) - g'(x_t)| > \epsilon/2 \right\} \right]$$

Recall that, for any $g \in \mathcal{V}_t^{\mathcal{G}}$, we have

$$\sum_{s=1}^{t-1} (g(x_s) - \widehat{g}(x_s))^2 \leq \beta_{\mathcal{G}}(t-1).$$

Hence, applying Lemma D.11, we get

$$\begin{aligned} T_1 &\leq \min \left\{ \left(\frac{2\sqrt{T\beta_{\mathcal{G}}(t-1)}}{\epsilon/2} + 1 \right) \cdot \dim_1(\mathcal{G}, \epsilon/2), \left(\frac{4\beta_{\mathcal{G}}(t-1)}{\epsilon^2/4} + 1 \right) \dim_2(\mathcal{G}, \epsilon/2) \right\} \\ &\leq \min \left\{ \frac{5\sqrt{T\beta_{\mathcal{G}}(t-1)}}{\epsilon} \cdot \dim_1(\mathcal{G}, \epsilon/2), \frac{17\beta_{\mathcal{G}}(t-1)}{\epsilon^2} \cdot \dim_2(\mathcal{G}, \epsilon/2) \right\} \end{aligned}$$

Combining the upper bounds of T_1 and T_2 , we obtain

$$\begin{aligned} &\sum_{t=1}^T \mathbb{E} \left[\mathbb{1} \left\{ \mathbb{E}_{g, g'} \left[|g(x_t) - g'(x_t)| \mid \mathcal{H}_{t-1} \right] > \epsilon \right\} \right] \\ &\leq \frac{4GT\delta}{\epsilon} + \min \left\{ \frac{5\sqrt{T\beta_{\mathcal{G}}(t-1)}}{\epsilon} \cdot \dim_1(\mathcal{G}, \epsilon/2), \frac{17\beta_{\mathcal{G}}(t-1)}{\epsilon^2} \cdot \dim_2(\mathcal{G}, \epsilon/2) \right\}. \end{aligned}$$

Setting $\delta = 1/(TG)$, the parameter $\beta_{\mathcal{G}}(t-1)$ becomes $\beta'_{\mathcal{G}}(t-1)$, which is upper bounded by $\beta'_{\mathcal{G}}(T)$. Then, we finish the proof. \square

C.2 BOUNDING BAYESIAN REGRET

Now we are ready to prove Theorem 5.4. We will prove the upper bound on the Bayesian regret in this section and prove the upper bound on the number of queries in the next section.

For the ease of notation, we will omit the dependence of the state-value function on the initial state s_1 and simply write $V := V_1(s_1)$ for any state-value function V throughout the proof.

We start by simplifying the Bayesian regret. Since $\pi_t^1 = \pi_{t-1}^0$, we have

$$\begin{aligned} \text{BayesRegret}_T &= \mathbb{E}_{r^*, P^*} \left[\sum_{t=1}^T \left(2V^* - V^{\pi_t^0} - V^{\pi_t^1} \right) \right] \\ &= \mathbb{E}_{r^*, P^*} \left[\sum_{t=1}^T \left(2V^* - V^{\pi_t^0} - V^{\pi_{t-1}^0} \right) \right] \\ &= \mathbb{E}_{r^*, P^*} \left[\sum_{t=1}^T \left(V^* - V^{\pi_t^0} \right) \right] + \mathbb{E}_{r^*, P^*} \left[\sum_{t=0}^{T-1} \left(V^* - V^{\pi_t^0} \right) \right] \\ &\leq 2 \mathbb{E}_{r^*, P^*} \left[\sum_{t=0}^T \left(V^* - V^{\pi_t^0} \right) \right]. \end{aligned}$$

Hence, we only need to consider the regret incurred by π_t^0 . We defined $V_{r, P}^{\pi}$ as the state-value function of policy π with reward function r and model P . Given that, we can express the Bayesian regret as

$$\text{BayesRegret}_T \leq 2 \mathbb{E}_{r^*, P^*} \left[\sum_{t=0}^T \left(V_{r^*, P^*}^{\pi_t^*} - V_{r^*, P^*}^{\pi_t^0} \right) \right]$$

We reformulate the expectation by first taking the expectation of the historical data up to round $t-1$ and then taking the conditional expectation of P^* and r^* . Concretely, we denote $\mathcal{H}_{t-1} = \{\tau_1^0, \tau_1^1, (o_1), \dots, \tau_{t-1}^0, \tau_{t-1}^1, (o_{t-1})\}$ as the history up to round $t-1$, and then we have

$$\begin{aligned} \text{BayesRegret}_T &\leq 2 \sum_{t=0}^T \mathbb{E}_{\mathcal{H}_{t-1}} \mathbb{E}_{r^*, P^*} \left[V_{r^*, P^*}^{\pi_t^*} - V_{r^*, P^*}^{\pi_t^0} \mid \mathcal{H}_{t-1} \right] \\ &= 2 \sum_{t=0}^T \mathbb{E}_{\mathcal{H}_{t-1}} \mathbb{E}_{r^*, P^*} \left[V_{r_t, P_t}^{\pi_t^0} - V_{r^*, P^*}^{\pi_t^0} \mid \mathcal{H}_{t-1} \right]. \end{aligned}$$

For the equality above, we note that (r_t, P_t) and (r^*, P^*) are identically distributed given \mathcal{H}_{t-1} and π_t^0 and π^* are the respective optimal policies of (r_t, P_t) and (r^*, P^*) . Hence, $V_{r_t, P_t}^{\pi_t^0}$ and $V_{r^*, P^*}^{\pi^*}$ are identically distributed given \mathcal{H}_{t-1} , which is the reason that we can do the replacement in the conditional expectation. Next, we proceed by decomposing the regret:

$$\begin{aligned}
\text{BayesRegret}_T &= 2 \sum_{t=0}^T \mathbb{E}_{\mathcal{H}_{t-1}} \mathbb{E}_{r^*, P^*} \left[V_{r_t, P_t}^{\pi_t^0} - V_{r_t, P^*}^{\pi_t^0} + V_{r_t, P^*}^{\pi_t^0} - V_{r^*, P^*}^{\pi_t^0} + V_{r^*, P^*}^{\pi_t^0} - V_{r^*, P^*}^{\pi^*} \mid \mathcal{H}_{t-1} \right] \\
&= 2 \sum_{t=0}^T \mathbb{E}_{\mathcal{H}_{t-1}} \mathbb{E}_{r^*, P^*} \left[V_{r_t, P_t}^{\pi_t^0} - V_{r_t, P^*}^{\pi_t^0} + V_{r_t, P^*}^{\pi_t^0} - V_{r_t, P^*}^{\pi_t^1} + V_{r^*, P^*}^{\pi_t^1} - V_{r^*, P^*}^{\pi^*} \mid \mathcal{H}_{t-1} \right] \\
&= 2 \sum_{t=0}^T \mathbb{E} \left[V_{r_t, P_t}^{\pi_t^0} - V_{r_t, P^*}^{\pi_t^0} + V_{r_t, P^*}^{\pi_t^0} - V_{r_t, P^*}^{\pi_t^1} + V_{r^*, P^*}^{\pi_t^1} - V_{r^*, P^*}^{\pi^*} \right] \\
&= 2 \left(\underbrace{\sum_{t=0}^T \mathbb{E} \left[V_{r_t, P_t}^{\pi_t^0} - V_{r_t, P^*}^{\pi_t^0} \right]}_{=: \mathbf{T}_{\text{model}}} + \underbrace{\sum_{t=0}^T \mathbb{E} \left[V_{r_t, P^*}^{\pi_t^0} - V_{r_t, P^*}^{\pi_t^1} + V_{r^*, P^*}^{\pi_t^1} - V_{r^*, P^*}^{\pi^*} \right]}_{=: \mathbf{T}_{\text{reward}}} \right)
\end{aligned} \tag{19}$$

Here we added and subtracted $V_{r_t, P^*}^{\pi_t^0}$ and $V_{r^*, P^*}^{\pi_t^1}$ in the first equality. For the second equality, we note that $\pi_t^1 = \pi_{t-1}^0$ is measurable with respect to \mathcal{H}_{t-1} (i.e., $\mathbb{E}[\pi_t^1 \mid \mathcal{H}_{t-1}] = \pi_t^1$).⁴ Hence, $V_{r_t, P^*}^{\pi_t^0}$ and $V_{r^*, P^*}^{\pi_t^1}$ are identically distributed conditioning on \mathcal{H}_{t-1} .

Next, we will show that the two terms, $\mathbf{T}_{\text{model}}$ and $\mathbf{T}_{\text{reward}}$, arise from the estimation errors in the model and reward, respectively.

Bounding $\mathbf{T}_{\text{model}}$. By simulation lemma (Lemma D.2) and the tower rule, we have

$$\begin{aligned}
\mathbf{T}_{\text{model}} &\leq H \mathbb{E} \left[\sum_{t=0}^T \sum_{h=1}^H \mathbb{E}_{(s_h, a_h) \sim d_h^{\pi_t^0}} d_{\text{TV}} \left(P_t(s_h, a_h), P^*(s_h, a_h) \right) \right] \\
&= H \mathbb{E} \left[\sum_{t=0}^T \sum_{h=1}^H d_{\text{TV}} \left(P_t(s_h, a_h), P^*(s_h, a_h) \right) \right] \\
&= H \mathbb{E} \left[\sum_{t=0}^T \sum_{h=1}^H \mathbb{E}_{P_t, P^*} \left[d_{\text{TV}} \left(P_t(s_h, a_h), P^*(s_h, a_h) \right) \mid \mathcal{H}_{t-1} \right] \right] \\
&= H \sum_{h=1}^H \mathbb{E} \left[\sum_{t=0}^T \mathbb{E}_{P_t, P^*} \left[d_{\text{TV}} \left(P_t(s_h, a_h), P^*(s_h, a_h) \right) \mid \mathcal{H}_{t-1} \right] \right] \\
&\leq O \left(H^2 \cdot \dim_1(\mathcal{P}, 1/T) \cdot \log(T) \cdot \sqrt{T \log(TN_{\square}((HT|\mathcal{S})^{-1}, \mathcal{P}, \|\cdot\|_{\infty}))} \right)
\end{aligned}$$

where the last step is Lemma C.5.

Bounding $\mathbf{T}_{\text{reward}}$. By the definition of state-value function, we have

$$\begin{aligned}
\mathbf{T}_{\text{reward}} &= \sum_{t=0}^T \mathbb{E} \left[r_t(\tau_t^0) - r_t(\tau_t^1) + r^*(\tau_t^1) - r^*(\tau_t^0) \right] \\
&= \sum_{t=0}^T \mathbb{E} \left[(1 - Z_t) \left(r_t(\tau_t^0) - r_t(\tau_t^1) + r^*(\tau_t^1) - r^*(\tau_t^0) \right) \right] \\
&\quad + \sum_{t=0}^T \mathbb{E} \left[Z_t \left(r_t(\tau_t^0) - r_t(\tau_t^1) + r^*(\tau_t^1) - r^*(\tau_t^0) \right) \right].
\end{aligned}$$

⁴Strictly speaking, π_t^1 is measurable with respect to the σ -algebra generated by \mathcal{H}_t .

For the first term, by the query condition, it holds that

$$\begin{aligned}
& \sum_{t=0}^T \mathbb{E} \left[(1 - Z_t) \left(r_t(\tau_t^0) - r_t(\tau_t^1) + r^*(\tau_t^1) - r^*(\tau_t^0) \right) \right] \\
&= \sum_{t=0}^T \mathbb{E} \left[\mathbb{1} \left\{ \mathbb{E} \left[|r(\tau_t^0) - r(\tau_t^1) - (r'(\tau_t^0) - r'(\tau_t^1))| \mid \mathcal{H}_{t-1}, \tau_t^0, \tau_t^1 \right] \leq \epsilon \right\} \right. \\
&\quad \left. \left(r_t(\tau_t^0) - r_t(\tau_t^1) + r^*(\tau_t^1) - r^*(\tau_t^0) \right) \right] \\
&= \sum_{t=0}^T \mathbb{E}_{\mathcal{H}_{t-1}, \tau_t^0, \tau_t^1} \left[\mathbb{E} \left[\mathbb{1} \left\{ \mathbb{E} \left[|r(\tau_t^0) - r(\tau_t^1) - (r'(\tau_t^0) - r'(\tau_t^1))| \mid \mathcal{H}_{t-1}, \tau_t^0, \tau_t^1 \right] \leq \epsilon \right\} \right. \right. \\
&\quad \left. \left. \left(r_t(\tau_t^0) - r_t(\tau_t^1) + r^*(\tau_t^1) - r^*(\tau_t^0) \right) \mid \mathcal{H}_{t-1}, \tau_t^0, \tau_t^1 \right] \right] \\
&= \sum_{t=0}^T \mathbb{E}_{\mathcal{H}_{t-1}, \tau_t^0, \tau_t^1} \left[\mathbb{1} \left\{ \mathbb{E} \left[|r(\tau_t^0) - r(\tau_t^1) - (r'(\tau_t^0) - r'(\tau_t^1))| \mid \mathcal{H}_{t-1}, \tau_t^0, \tau_t^1 \right] \leq \epsilon \right\} \right. \\
&\quad \left. \mathbb{E} \left[\left(r_t(\tau_t^0) - r_t(\tau_t^1) + r^*(\tau_t^1) - r^*(\tau_t^0) \right) \mid \mathcal{H}_{t-1}, \tau_t^0, \tau_t^1 \right] \right] \\
&\leq T\epsilon.
\end{aligned}$$

Here the first equality is by definition. The second equality is by the law of probability. The third equality holds since the indicator is measurable with respect to $\mathcal{H}_{t-1}, \tau_t^0$, and τ_t^1 . The last inequality is by the indicator and the fact that r_t and r^* have the same posterior conditioning on \mathcal{H}_{t-1} . Plugging this upper bound back, we have

$$\begin{aligned}
\mathsf{T}_{\text{reward}} &\leq T\epsilon + \sum_{t=0}^T \mathbb{E} \left[Z_t \left(r_t(\tau_t^0) - r_t(\tau_t^1) + r^*(\tau_t^1) - r^*(\tau_t^0) \right) \right] \\
&= T\epsilon + \sum_{t=0}^T \mathbb{E}_{\mathcal{H}_{t-1}} \left[Z_t \mathbb{E}_{r_t, r^*} \left[\left| r_t(\tau_t^0) - r_t(\tau_t^1) + r^*(\tau_t^1) - r^*(\tau_t^0) \right| \mid \mathcal{H}_{t-1} \right] \right].
\end{aligned}$$

In the inner expectation of the second term above, we notice that both r_t and r^* are sampled from the posterior of r^* conditioning on \mathcal{H}_{t-1} . Thus, we can invoke the second statement in Lemma C.5 where the function g corresponds to r_t and the function g' corresponds to r^* . This gives

$$\mathsf{T}_{\text{reward}} \leq O \left(T\epsilon + \sqrt{T\kappa^2 \log \left(HTN_{\square}(\bar{\kappa}(2T)^{-1}, \tilde{\mathcal{R}}, \|\cdot\|_{\infty}) \right)} \cdot \dim_1(\tilde{\mathcal{R}}, 1/T) \cdot \log(HT) \right).$$

Conclusion. Given the upper bounds on $\mathsf{T}_{\text{model}}$ and $\mathsf{T}_{\text{reward}}$, ignoring logarithmic factors on T and H , we conclude that

$$\begin{aligned}
\text{BayesRegret}_T &= \tilde{O} \left(H^2 \cdot \dim_1(\mathcal{P}, 1/T) \cdot \sqrt{T \log(N_{\square}((HT|\mathcal{S})^{-1}, \mathcal{P}, \|\cdot\|_{\infty}))} \right. \\
&\quad \left. + T\epsilon + \kappa \cdot \dim_1(\tilde{\mathcal{R}}, 1/T) \cdot \sqrt{T \log \left(N_{\square}(\bar{\kappa}(2T)^{-1}, \tilde{\mathcal{R}}, \|\cdot\|_{\infty}) \right)} \right).
\end{aligned}$$

C.3 BOUNDING NUMBER OF QUERIES

It holds that

$$\begin{aligned} \text{BayesQueries}_T &= \mathbb{E} \left[\sum_{t=0}^T Z_t \right] \\ &= \sum_{t=0}^T \mathbb{E}_{\mathcal{H}_{t-1}} \left[Z_t \mathbb{1} \left\{ \mathbb{E}_{r, r' \sim \rho_{r,t}} \left[((r(\tau_t^0) - r(\tau_t^1)) - (r'(\tau_t^0) - r'(\tau_t^1))) \right] > \epsilon \right\} \right]. \end{aligned}$$

We notice that, in the inner expectation, r and r' are sampled from the posterior of r^* conditioning on \mathcal{H}_{t-1} (recalling the definition of $\rho_{r,t}$). Thus, we can invoke Lemma C.6 where the function g corresponds to r and the function g' corresponds to r' . This gives

$$\text{BayesQueries}_T \leq \min \left\{ \frac{9\sqrt{T}\beta_{\mathcal{R}}}{\epsilon} \cdot \text{dim}_1(\tilde{\mathcal{R}}, \epsilon/2), \frac{21\beta_{\mathcal{R}}}{\epsilon^2} \cdot \text{dim}_2(\tilde{\mathcal{R}}, \epsilon/2) \right\}.$$

where the inequality is Lemma C.6 and $\beta_{\mathcal{R}} := 98\kappa^2 \log(2THN_{\square}(\bar{\kappa}(2T)^{-1}, \tilde{\mathcal{R}}, \|\cdot\|_{\infty}))$.

C.4 GENERALIZING THEOREM 5.4 THROUGH SEC

In this section, we show that the dependence on the eluder dimension in Theorem 5.4 can be generalized to the *Sequential Extrapolation Coefficient (SEC)* (Xie et al., 2022). It have been shown by Xie et al. (2022) that the SEC can be upper bounded by the eluder dimension, the Bellman-eluder dimension (Jin et al., 2021), and the bilinear rank (Du et al., 2021). Thus, SEC is a more general measure of complexity.

We start by introducing the Bayesian Sequential Extrapolation Coefficient (Bayesian SEC) in the preference-based feedback, which is a variant of the original SEC to accommodate the Bayesian and preference-based learning setting. We define the Bayesian SEC for a model class \mathcal{P} as

$$\text{BayesSEC}_{\mathcal{P}}(\mathcal{P}) := \mathbb{E}_{r^*, P^*} \left[\sup_{\substack{P_1, \dots, P_T \in \mathcal{P} \\ (\pi_1^0, \pi_1^1), \dots, (\pi_T^0, \pi_T^1)}} \sum_{t=1}^T \frac{\left(\sum_{h=1}^H \mathbb{E}_{s, a \sim d_h^{\pi_t}} [d_{\text{TV}}(P_t(\cdot | s, a), P^*(\cdot | s, a))] \right)^2}{1 \vee \sum_{i=1}^{t-1} \sum_{h=1}^H \mathbb{E}_{s, a \sim d_h^{\pi_i}} [d_{\text{TV}}^2(P_t(\cdot | s, a), P^*(\cdot | s, a))]} \right]$$

and the Bayesian SEC for a reward function class \mathcal{R} as

$$\text{BayesSEC}_{\mathcal{R}}(\mathcal{R}) := \mathbb{E}_{r^*, P^*} \left[\sup_{\substack{r_1, \dots, r_T \in \mathcal{R} \\ (\pi_1^0, \pi_1^1), \dots, (\pi_T^0, \pi_T^1)}} \sum_{t=1}^T \frac{\mathbb{E}_{r_t^0 \sim \pi_t^0, r_t^1 \sim \pi_t^1} \left[(r_t(\tau_t^0) - r_t(\tau_t^1) + r^*(\tau_t^1) - r^*(\tau_t^0))^2 \right]}{1 \vee \sum_{i=1}^{t-1} \mathbb{E}_{r_i^0 \sim \pi_i^0, r_i^1 \sim \pi_i^1} \left[(r_i(\tau_i^0) - r_i(\tau_i^1) + r^*(\tau_i^1) - r^*(\tau_i^0))^2 \right]} \right]$$

We note that the Bayesian SEC can be easily reduced to the frequentist SEC by specifying the prior on r^* and P^* to be the Dirac delta function. We note that our definition of the reward function SEC involves squaring within the expectation in the numerator rather than externally, potentially making it larger.

Now, we are ready to state the generalization of Theorem 5.4.

Theorem C.7. *PbTS (Algorithm 2) guarantees that*

$$\begin{aligned} \text{BayesRegret}_T &= \tilde{O} \left(T\epsilon + H \sqrt{\text{BayesSEC}_{\mathcal{P}}(\mathcal{P}) \cdot TH \cdot \nu_{\mathcal{P}}} + \sqrt{\text{BayesSEC}_{\mathcal{R}}(\mathcal{R}) \cdot T \cdot \kappa^2 \cdot \nu_{\mathcal{R}}} \right), \\ \text{BayesQueries}_T &= \tilde{O} \left(\frac{\kappa^2 \cdot \text{BayesSEC}_{\mathcal{R}}(\mathcal{R}) \cdot \nu_{\mathcal{R}}}{\epsilon^2} \right). \end{aligned}$$

where we denote $\nu_{\mathcal{P}} := \log(N_{\square}((HT|\mathcal{S})^{-1}, \mathcal{P}, \|\cdot\|_{\infty}))$ and $\nu_{\mathcal{R}} := \log(N_{\square}(\bar{\kappa}(2T)^{-1}, \tilde{\mathcal{R}}, \|\cdot\|_{\infty}))$.

The proofs are provided in the following sections (Appendices C.4.1 and C.4.2).

C.4.1 BOUNDING BAYESIAN REGRET VIA SEC

We first prove the upper bound on the Bayesian regret. Following (19), it suffices to separately bound T_{model} and T_{reward} via SEC. We start with T_{model} .

Bounding T_{model} . By simulation lemma (Lemma D.2), we have

$$T_{\text{model}} \leq H \mathbb{E} \left[\sum_{t=0}^T \sum_{h=1}^H \mathbb{E}_{(s_h, a_h) \sim d_h^{\pi_t^0}} d_{\text{TV}} \left(P_t(s_h, a_h), P^*(s_h, a_h) \right) \right]$$

where $d_h^\pi(s, a)$ denotes the probability of π reaching (s, a) at time step h . By multiplying and dividing by the same term, we obtain the following.

$$\begin{aligned} T_{\text{model}} &\leq H \mathbb{E} \left[\sum_{t=0}^T \frac{\sum_{h=1}^H \mathbb{E}_{(s_h, a_h) \sim d_h^{\pi_t^0}} d_{\text{TV}} \left(P_t(s_h, a_h), P^*(s_h, a_h) \right)}{\sqrt{1 \vee \sum_{i=0}^{t-1} \sum_{h=1}^H \mathbb{E}_{(s_h, a_h) \sim d_h^{\pi_i^0}} d_{\text{TV}}^2 \left(P_t(s_h, a_h), P^*(s_h, a_h) \right)}} \right. \\ &\quad \left. \times \sqrt{1 \vee \sum_{i=0}^{t-1} \sum_{h=1}^H \mathbb{E}_{(s_h, a_h) \sim d_h^{\pi_i^0}} d_{\text{TV}}^2 \left(P_t(s_h, a_h), P^*(s_h, a_h) \right)} \right] \\ &\leq H \underbrace{\mathbb{E} \left[\sum_{t=0}^T \frac{\left(\sum_{h=1}^H \mathbb{E}_{(s_h, a_h) \sim d_h^{\pi_t^0}} d_{\text{TV}} \left(P_t(s_h, a_h), P^*(s_h, a_h) \right) \right)^2}{1 \vee \sum_{i=0}^{t-1} \sum_{h=1}^H \mathbb{E}_{(s_h, a_h) \sim d_h^{\pi_i^0}} d_{\text{TV}}^2 \left(P_t(s_h, a_h), P^*(s_h, a_h) \right)} \right]}_{\text{(i)}} \\ &\quad \times \underbrace{\sqrt{\sum_{t=0}^T \mathbb{E} \left[1 \vee \sum_{i=0}^{t-1} \sum_{h=1}^H \mathbb{E}_{(s_h, a_h) \sim d_h^{\pi_i^0}} d_{\text{TV}}^2 \left(P_t(s_h, a_h), P^*(s_h, a_h) \right) \right]}}_{\text{(ii)}} \end{aligned}$$

Here we note that $a \vee b := \max(a, b) \leq a + b$ for $a, b \geq 0$. The last step above is the Cauchy-Schwarz inequality. We observe that term (i) is exactly bounded by the Bayesian SEC of models. For term (ii), we can invoke (16) in Lemma C.5 since P_t has an identical posterior distribution as P^* given the historical data up to round $t - 1$. This gives (ii) $\leq \tilde{O}(\sqrt{TH \cdot \iota_{\mathcal{P}}})$. Plugging these back, we obtain

$$T_{\text{model}} \leq \tilde{O} \left(H \sqrt{\text{BayesSEC}_{\mathcal{P}}(\mathcal{P}) \cdot TH \cdot \iota_{\mathcal{P}}} \right)$$

Bounding T_{reward} . The reward part can be bounded similarly, except that we need to additionally consider the query conditions. To begin with, by the definition of state-value function, we have

$$\begin{aligned} T_{\text{reward}} &= \sum_{t=0}^T \mathbb{E} \left[r_t(\tau_t^0) - r_t(\tau_t^1) + r^*(\tau_t^1) - r^*(\tau_t^0) \right] \\ &= \underbrace{\sum_{t=0}^T \mathbb{E} \left[(1 - Z_t) \left(r_t(\tau_t^0) - r_t(\tau_t^1) + r^*(\tau_t^1) - r^*(\tau_t^0) \right) \right]}_{\text{(i)}} + \underbrace{\sum_{t=0}^T \mathbb{E} \left[Z_t \left(r_t(\tau_t^0) - r_t(\tau_t^1) + r^*(\tau_t^1) - r^*(\tau_t^0) \right) \right]}_{\text{(ii)}}. \end{aligned}$$

Following the same argument as in the proof of Appendix C.2, we can bound (i) as (i) $\leq T\epsilon$. Now we proceed to bound (ii), which is the regret incurred when making queries. By multiplying and

dividing by the same term, we obtain the following.

$$\begin{aligned}
\text{(ii)} &= \sum_{t=0}^T \mathbb{E}_{\pi_t^0, \pi_t^1, r_t, Z_t} \left[\frac{\mathbb{E} \left[Z_t \left(r_t(\tau_t^0) - r_t(\tau_t^1) + r^*(\tau_t^1) - r^*(\tau_t^0) \right) \mid \pi_t^0, \pi_t^1, r_t, Z_t \right]}{\sqrt{1 \vee \sum_{i=0}^{t-1} \mathbb{E} \left[Z_t \left(r_t(\tau_i^0) - r_t(\tau_i^1) + r^*(\tau_i^1) - r^*(\tau_i^0) \right)^2 \mid \pi_t^0, \pi_t^1, r_t, Z_t \right]}} \right. \\
&\quad \left. \times \sqrt{1 \vee \sum_{i=0}^{t-1} \mathbb{E} \left[Z_t \left(r_t(\tau_i^0) - r_t(\tau_i^1) + r^*(\tau_i^1) - r^*(\tau_i^0) \right)^2 \mid \pi_t^0, \pi_t^1, r_t, Z_t \right]} \right] \\
&\leq \underbrace{\sqrt{\sum_{t=0}^T \mathbb{E}_{\pi_t^0, \pi_t^1, r_t, Z_t} \left[\frac{\left(\mathbb{E} \left[Z_t \left(r_t(\tau_t^0) - r_t(\tau_t^1) + r^*(\tau_t^1) - r^*(\tau_t^0) \right) \mid \pi_t^0, \pi_t^1, r_t, Z_t \right] \right)^2}{1 \vee \sum_{i=0}^{t-1} \mathbb{E} \left[Z_t \left(r_t(\tau_i^0) - r_t(\tau_i^1) + r^*(\tau_i^1) - r^*(\tau_i^0) \right)^2 \mid \pi_t^0, \pi_t^1, r_t, Z_t \right]} \right]}}_{\text{(iii)}} \\
&\quad \times \underbrace{\sqrt{\sum_{t=0}^T \left(1 \vee \sum_{i=0}^{t-1} \mathbb{E} \left[Z_t \left(r_t(\tau_i^0) - r_t(\tau_i^1) + r^*(\tau_i^1) - r^*(\tau_i^0) \right)^2 \right] \right)}}_{\text{(iv)}}
\end{aligned}$$

where the last inequality is Cauchy-Schwarz inequality. We observe that term (iii) is exactly bounded by the Bayesian SEC of reward functions via Jensen's inequality. For term (iv), we invoke (18) in Lemma C.5 and get (iv) $\leq \tilde{O}(\sqrt{\kappa^2 T \cdot \nu_{\mathcal{R}}})$. Now plugging these upper bounds back, we get

$$T_{\text{reward}} \leq T\epsilon + \tilde{O}\left(\sqrt{\text{BayesSEC}_{\mathcal{R}}(\mathcal{R}) \cdot \kappa^2 T \cdot \nu_{\mathcal{R}}}\right).$$

Conclusion. Given the upper bounds on T_{model} and T_{reward} , we conclude that

$$\text{BayesRegret}_T \leq \tilde{O}\left(H\sqrt{\text{BayesSEC}_{\mathcal{P}}(\mathcal{P}) \cdot TH \cdot \nu_{\mathcal{P}}} + T\epsilon + \sqrt{\text{BayesSEC}_{\mathcal{R}}(\mathcal{R}) \cdot \kappa^2 T \cdot \nu_{\mathcal{R}}}\right).$$

C.4.2 BOUNDING NUMBER OF QUERIES VIA SEC

By the definition of Z_t , we have

$$\begin{aligned}
\text{BayesQueries}_T &= \mathbb{E} \left[\sum_{t=0}^T Z_t \right] = \mathbb{E} \left[\sum_{t=0}^T Z_t^2 \right] \\
&= \mathbb{E} \left[\sum_{t=0}^T Z_t \mathbb{1} \left\{ \mathbb{E}_{r, r' \sim \rho_{r,t}} \left[\left| (r(\tau_t^0) - r(\tau_t^1)) - (r'(\tau_t^0) - r'(\tau_t^1)) \right| \right] > \epsilon \right\} \right] \\
&= \mathbb{E} \left[\sum_{t=0}^T Z_t \mathbb{1} \left\{ \mathbb{E}_{r, r' \sim \rho_{r,t}} \left[\left| (r(\tau_t^0) - r(\tau_t^1)) - (r'(\tau_t^0) - r'(\tau_t^1)) \right|^2 \right] > \epsilon^2 \right\} \right] \\
&\leq \frac{1}{\epsilon^2} \mathbb{E} \left[\sum_{t=0}^T Z_t \mathbb{E}_{r, r' \sim \rho_{r,t}} \left[\left| (r(\tau_t^0) - r(\tau_t^1)) - (r'(\tau_t^0) - r'(\tau_t^1)) \right|^2 \right] \right] \\
&= \frac{1}{\epsilon^2} \mathbb{E} \left[\sum_{t=0}^T Z_t \mathbb{E}_{r, r'} \left[\left| (r_t(\tau_t^0) - r_t(\tau_t^1)) - (r^*(\tau_t^0) - r^*(\tau_t^1)) \right|^2 \mid \mathcal{H}_{t-1} \right] \right].
\end{aligned}$$

Here the last step is by the definition of $\rho_{r,t}$. Now we swap the order of expectation and get BayesQueries_T

$$\begin{aligned}
&\leq \frac{1}{\epsilon^2} \mathbb{E} \left[\sum_{t=0}^T \mathbb{E}_{\tau_t^0, \tau_t^1} \left[Z_t \left((r_t(\tau_t^0) - r_t(\tau_t^1)) - (r^*(\tau_t^0) - r^*(\tau_t^1)) \right)^2 \middle| \pi_t^0, \pi_t^1, r^*, r_t, Z_t \right] \right] \\
&= \frac{1}{\epsilon^2} \mathbb{E} \left[\sum_{t=0}^T \frac{\mathbb{E}_{\tau_t^0, \tau_t^1} \left[Z_t \left((r_t(\tau_t^0) - r_t(\tau_t^1)) - (r^*(\tau_t^0) - r^*(\tau_t^1)) \right)^2 \middle| \pi_t^0, \pi_t^1, r^*, r_t, Z_t \right]}{1 \vee \sum_{i=0}^{t-1} \mathbb{E}_{\tau_i^0, \tau_i^1} \left[Z_t \left((r_t(\tau_i^0) - r_t(\tau_i^1)) - (r^*(\tau_i^0) - r^*(\tau_i^1)) \right)^2 \middle| \pi_i^0, \pi_i^1, r^*, r_t, Z_t \right]} \right. \\
&\quad \left. \times \left(1 \vee \sum_{i=0}^{t-1} \mathbb{E}_{\tau_i^0, \tau_i^1} \left[Z_t \left((r_t(\tau_i^0) - r_t(\tau_i^1)) - (r^*(\tau_i^0) - r^*(\tau_i^1)) \right)^2 \middle| \pi_i^0, \pi_i^1, r^*, r_t, Z_t \right] \right) \right] \\
&\leq \frac{1}{\epsilon^2} \mathbb{E} \left[\sum_{t=0}^T \frac{\mathbb{E}_{\tau_t^0, \tau_t^1} \left[Z_t \left((r_t(\tau_t^0) - r_t(\tau_t^1)) - (r^*(\tau_t^0) - r^*(\tau_t^1)) \right)^2 \middle| \pi_t^0, \pi_t^1, r^*, r_t, Z_t \right]}{1 \vee \sum_{i=0}^{t-1} \mathbb{E}_{\tau_i^0, \tau_i^1} \left[Z_t \left((r_t(\tau_i^0) - r_t(\tau_i^1)) - (r^*(\tau_i^0) - r^*(\tau_i^1)) \right)^2 \middle| \pi_i^0, \pi_i^1, r^*, r_t, Z_t \right]} \right] \\
&\quad \times \tilde{O}(\kappa^2 \cdot \iota_{\mathcal{R}})
\end{aligned}$$

where in the last step we applied Jensen's inequality to the first multiplicative term and applied (18) in Lemma C.5 to the second term. We observe that the first term can be bounded by SEC, and thus we obtain

$$\text{BayesQueries}_T \leq \tilde{O} \left(\frac{\kappa^2 \cdot \text{BayesSEC}_{\mathcal{R}}(\mathcal{R}) \cdot \iota_{\mathcal{R}}}{\epsilon^2} \right).$$

D TECHNICAL LEMMAS

Lemma D.1. *Under Assumption 3.1, we have $\bar{\kappa}(\Phi(a) - \Phi(b)) \leq a - b \leq \kappa(\Phi(a) - \Phi(b))$ for any $a, b \in [0, H]$.*

Proof of Lemma D.1. We note that $\Phi(a) - \Phi(b) = \int_b^a \Phi'(x) dx$ and

$$\kappa^{-1}(a - b) \leq (a - b) \inf_{x \in [a, b]} \Phi'(x) \leq \int_b^a \Phi'(x) dx \leq (a - b) \sup_{x \in [a, b]} \Phi'(x) \leq \bar{\kappa}^{-1}(a - b).$$

□

Lemma D.2 (Simulation lemma). *For any models P, \hat{P} and any policy π , we have*

$$\begin{aligned}
\left| V_{\hat{P},1}^{\pi}(s_1) - V_{\hat{P},1}^{\pi}(s_1) \right| &\leq \sum_{h=1}^H \mathbb{E}_{(s_h, a_h) \sim d_{P,h}^{\pi}} \left| \mathbb{E}_{s' \sim P(s_h, a_h)} V_{\hat{P},h+1}^{\pi}(s') - \mathbb{E}_{s' \sim \hat{P}(s_h, a_h)} V_{\hat{P},h+1}^{\pi}(s') \right| \\
&\leq H \sum_{h=1}^H \mathbb{E}_{(s_h, a_h) \sim d_{P,h}^{\pi}} d_{\text{TV}} \left(P(\cdot | s_h, a_h), \hat{P}(\cdot | s_h, a_h) \right)
\end{aligned}$$

where $d_{P,h}^{\pi}(s, a)$ is the probability of π reaching (s, a) at time step h given model P .

Lemma D.3 (Gaussian concentration). *Let $\epsilon \sim \mathcal{N}(0, c\Sigma^{-1})$ for $c \in \mathbb{R}^+$ and Σ a positive definite matrix. Then, for any $\delta > 0$, we have*

1. $\Pr \left(\|\epsilon\|_{\Sigma} > \sqrt{2cd \log(2d/\delta)} \right) \leq \delta.$
2. $\mathbb{E} [\|\epsilon\|_{\Sigma}] \leq \sqrt{cd}.$

Proof of Lemma D.3. The proof of the first inequality is provided in Abeille & Lazaric (2017, Appendix A). We provide a proof here for completeness. Let $\eta \sim \mathcal{N}(0, I_d)$ where I_d denotes the d -dimensional identity matrix. Fix an arbitrary α . We have

$$\Pr \left(\|\eta\|_2 > \alpha\sqrt{d} \right) \leq \Pr \left(\exists i : |\eta_i| > \alpha \right) \leq d \Pr \left(|\eta_1| > \alpha \right) \leq d \cdot 2e^{-\alpha^2/2}$$

where the second inequality is the union bound, and the last inequality is the standard concentration inequality for one-dimensional Gaussian random variable. We choose $\alpha = \sqrt{2 \log(2d/\delta)}$ and get

$$\Pr \left(\|\eta\|_2 > \sqrt{2d \log(2d/\delta)} \right) \leq \delta.$$

Let B denote the square root of Σ , i.e., $BB = \Sigma$. Then, we have

$$\|\epsilon\|_\Sigma = \sqrt{\epsilon^\top \Sigma \epsilon} = \sqrt{(\sqrt{c}B^{-1}\eta)^\top BB(\sqrt{c}B^{-1}\eta)} = \sqrt{c}\|\eta\|_2 \quad (20)$$

Hence, it holds that

$$\Pr \left(\|\epsilon\|_\Sigma > \sqrt{2cd \log(2d/\delta)} \right) = \Pr \left(\|\eta\|_2 > \sqrt{2d \log(2d/\delta)} \right) \leq \delta.$$

This proves the first inequality. To prove the second inequality, we note that, by (20),

$$\mathbb{E} [\|\epsilon\|_\Sigma] = \sqrt{c} \mathbb{E} [\|\eta\|_2] \leq \sqrt{c} \sqrt{\mathbb{E} [\eta^\top \eta]} = \sqrt{cd}.$$

□

The following inequalities are well-known, and we use the version in [Zhu & Nowak \(2022\)](#).

Lemma D.4. ([Zhu & Nowak, 2022](#)) *Let $\{X_t\}_{t \leq T}$ be a sequence of positive valued random variables adapted to a filtration \mathfrak{F}_t , and let $\mathbb{E}_t[\cdot] := \mathbb{E}[\cdot | \mathfrak{F}_{t-1}]$. If $X_t \leq B$ almost surely, then with probability at least $1 - \delta$, the following holds:*

$$\begin{aligned} \sum_{t=1}^T X_t &\leq \frac{3}{2} \sum_{t=1}^T \mathbb{E}_t[X_t] + 4B \log(1/\delta), \\ \sum_{t=1}^T \mathbb{E}_t[X_t] &\leq 2 \sum_{t=1}^T X_t + 8B \log(1/\delta). \end{aligned}$$

Lemma D.5 (Self-normalized process). ([Abbasi-Yadkori et al., 2011](#)) *Let $\{x_i\}_{i=1}^\infty$ be a real valued stochastic process sequence over the filtration $\{\mathcal{F}_i\}_{i=1}^\infty$. Let x_i be conditionally B -subgaussian given \mathcal{F}_{i-1} . Let $\{\phi_i\}_{i=1}^\infty$ with $\phi_i \in \mathcal{F}_{i-1}$ be a stochastic process in \mathbb{R}^d with each $\|\phi_i\| \leq L_\phi$. Define $\Sigma_i = \lambda I + \sum_{j=1}^{i-1} \phi_j \phi_j^\top$. Then for any $\delta > 0$ and all $i \geq 0$, with probability at least $1 - \delta$*

$$\left\| \sum_{i=1}^{k-1} \phi_i x_i \right\|_{\Sigma_k^{-1}}^2 \leq 2B^2 \log \left(\frac{\det(\Sigma_i)^{1/2} \det(\lambda I)^{-1/2}}{\delta} \right) \leq 2B^2 \left(d \log \left(\frac{\lambda + kL_\phi^2}{\lambda} \right) + \log(1/\delta) \right)$$

Lemma D.6 (Elliptical potential lemma). ([Abbasi-Yadkori et al., 2011](#)) *Following the setting of Lemma D.5, we have*

$$\sum_{i=1}^k \min \left\{ 1, \|\phi_i\|_{\Sigma_i^{-1}}^2 \right\} \leq 2d \log \left(\frac{\lambda + kL_\phi^2}{\lambda} \right)$$

Lemma D.7 (Sum of features). ([Jin et al., 2020](#)) *Following the setting of Lemma D.5, we have*

$$\sum_{i=1}^{k-1} \|\phi_i\|_{\Sigma_k^{-1}}^2 \leq d$$

The following lemma converts the TV distance between Gaussian distributions into the distance between their means.

Lemma D.8. *Assume $p_1 = \mathcal{N}(\mu_1, \sigma^2 I)$ and $p_2 = \mathcal{N}(\mu_2, \sigma^2 I)$ are two Gaussian distributions over \mathbb{R}^d . Assume $\|\mu_1 - \mu_2\|_2 \leq m$. Then, we have*

$$\|\mu_1 - \mu_2\|_2 \cdot \underbrace{\min \left\{ \sqrt{\frac{1}{2e\pi\sigma^2}}, \frac{1}{2m} \right\}}_{=: C_1} \leq d_{\text{TV}}(p_1, p_2) \leq \|\mu_1 - \mu_2\|_2 \cdot \underbrace{\sqrt{\frac{1}{2\pi\sigma^2}}}_{=: C_2}$$

Proof. Without loss of generality, we can rotate and translate the \mathbb{R}^d space and assume that $p_1 = \mathcal{N}(0, \sigma^2 I)$ and $p_2 = \mathcal{N}(\alpha \cdot e_1, \sigma^2 I)$ where $\alpha = \|\mu_1 - \mu_2\|_2$ and $e_1 = [1, 0, \dots, 0] \in \mathbb{R}^d$. It is clear that the total variation distance is preserved. Then, by definition,

$$d_{\text{TV}}(p_1, p_2) = \frac{1}{2} \int_{x \in \mathbb{R}^d} |p_1(x) - p_2(x)| dx = \int_{x \cdot e_1 \leq \alpha/2} p_1(x) - p_2(x) dx$$

where the last step holds since $p_1(x) - p_2(x) \geq 0$ if and only if $x \cdot e_1 \leq \alpha/2$. Note that the above formulation is independent of any dimension except for first dimension, so we can marginize other dimensions and reduce into a one-dimensional problem:

$$\begin{aligned} d_{\text{TV}}(p_1, p_2) &= \int_{x \leq \alpha/2} \mathcal{N}(x | 0, \sigma^2) - \mathcal{N}(x | \alpha, \sigma^2) dx \\ &= \left(F(\alpha/(2\sigma)) - F(-\alpha/(2\sigma)) \right) \\ &= \int_{-\alpha/(2\sigma)}^{\alpha/(2\sigma)} \mathcal{N}(x | 0, 1) dx \end{aligned}$$

where we denote F as the CDF of standard Gaussian $\mathcal{N}(0, 1)$ in the middle step. Below we show that the above integral is upper bounded by $C_2\alpha$ and lower bounded by $C_1\alpha$.

Upper bound. For an upper bound, we note that for any $v > 0$, we have

$$\int_{-v}^v \mathcal{N}(x | 0, 1) dx \leq 2v \cdot \max_x \mathcal{N}(x | 0, 1) \leq v \sqrt{\frac{2}{\pi}}.$$

Inserting $v = \alpha/(2\sigma)$, we obtain the desired upper bound.

Lower bound. For a lower bound, we note that for any $0 < v \leq m/(2\sigma)$, if $v < 1$, we have

$$\int_{-v}^v \mathcal{N}(x | 0, 1) dx \geq 2v \max_{x \in [-v, v]} \mathcal{N}(x | 0, 1) \geq 2v \mathcal{N}(1 | 0, 1) = 2v \cdot \frac{e^{-1/2}}{\sqrt{2\pi}} = v \cdot \sqrt{\frac{2}{e\pi}}.$$

When $v \geq 1$, we have

$$\int_{-v}^v \mathcal{N}(x | 0, 1) dx \geq \int_{-\sigma}^{\sigma} \mathcal{N}(x | 0, 1) dx \approx 0.6827 \geq v \cdot \frac{1}{2v} \geq v \cdot \frac{\sigma}{m}.$$

Combining the two cases together, we have

$$\int_{-v}^v \mathcal{N}(x | 0, 1) dx \geq v \cdot \min \left\{ \sqrt{\frac{2}{e\pi}}, \frac{\sigma}{m} \right\}.$$

Inserting $v = \alpha/(2\sigma)$, we obtain the desired lower bound. \square

Lemma D.9. Let $x \sim \mathcal{N}(0, I_d)$ and A be a positive semi-definite matrix. Then, we have

$$\mathbb{E}_x \|x\|_A \geq \sqrt{\frac{2 \text{tr}(A)}{\pi}}.$$

Proof. Let P denote the orthogonal matrix diagonalizing A , i.e., $A = P\Lambda P^\top$ where $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_d)$. Then, we have $\mathbb{E}_x \|x\|_A = \mathbb{E}_x \sqrt{x^\top A x} = \mathbb{E}_y \sqrt{y^\top \Lambda y}$ where $y = P^\top x \sim \mathcal{N}(0, I_d)$. Now we consider the following optimization problem:

$$\min_{M \in \mathcal{M}} f(M) := \mathbb{E}_y \sqrt{y^\top M y} \quad \text{where } \mathcal{M} := \{M \succeq 0 : M \text{ is diagonal and } \text{tr}(M) = \text{tr}(\Lambda)\}.$$

Then, we have $\Lambda \in \mathcal{M}$ and $\mathbb{E}_x \|x\|_A = f(\Lambda)$. Thus, the solution to the above optimization problem is a lower bound of $\mathbb{E}_x \|x\|_A$. Since $f(M)$ is concave in M , it must attain the minimum on the boundary of \mathcal{M} . By symmetry of entries in the diagonal, we know that the minimizer must have $\text{tr}(\Lambda)$ at one entry and 0 at all other entries. Hence, the minimum value is

$$\mathbb{E}_y \sqrt{\text{tr}(\Lambda) y_1^2} = \sqrt{\text{tr}(\Lambda)} \mathbb{E}_y |y_1| = \sqrt{\frac{2 \text{tr}(\Lambda)}{\pi}}.$$

We finish the proof by noticing that $\text{tr}(\Lambda) = \text{tr}(A)$. \square

The following lemma was originally established by [Agarwal et al. \(2020\)](#) and later adapted by [Wu et al. \(2023a\)](#) to accommodate infinitely large function classes.

Lemma D.10 (Maximum likelihood estimation). ([Wu et al., 2023a, Lemma C.3](#)) Consider a sequential conditional probability estimation problem. Let \mathcal{X} and \mathcal{Y} be the instance space and the target space, respectively. Let $\mathcal{F} : (\mathcal{X} \times \mathcal{Y}) \rightarrow \mathbb{R}$ denote the function class. We are given a dataset $D := \{(x_i, y_i)\}_{i=1}^n$ where $x_i \sim \mathcal{D}_i$ and $y_i \sim f^*(x, \cdot)$. We assume $f^* \in \mathcal{F}$. For the data generating process, we assume the data distribution \mathcal{D}_i is history-dependent, i.e., x_i can depend on the previous samples: $x_1, y_1, \dots, x_{i-1}, y_{i-1}$ for any $i \in [n]$. We fix $\delta \in (0, 1)$. Let \hat{f} denote the maximum likelihood estimator,

$$\hat{f} = \arg \max_{f \in \mathcal{F}} \sum_{i=1}^n \log f(x_i, y_i).$$

Then, we have

$$\sum_{i=1}^n \mathbb{E}_{x \sim \mathcal{D}_i} d_{\text{TV}}^2(\hat{f}(x, \cdot), f^*(x, \cdot)) \leq 10 \log \left(N_{[]} \left((n|\mathcal{Y}|)^{-1}, \mathcal{F}, \|\cdot\|_{\infty} \right) / \delta \right) \quad (21)$$

with probability at least $1 - \delta$. Here $|\mathcal{Y}|$ denotes the total measure of space \mathcal{Y} , i.e., $|\mathcal{Y}| := \int_{\mathcal{Y}} 1 \, d\mathbf{y}$.

D.1 ELUDER DIMENSION

The following lemmas are adapted from [Russo & Van Roy \(2013\)](#) and [Liu et al. \(2022a\)](#).

Lemma D.11. Following the notation of [Definition 5.2](#), define

$$\mathcal{F}_t = \left\{ f \in \mathcal{F} : \sum_{s=1}^{t-1} d^2(f(x_s), \hat{f}_t(x_s)) \leq \beta \right\}$$

where $\hat{f}_t \in \mathcal{F}$ is an arbitrary function. Then, we have

$$\sum_{t=1}^T \mathbb{1} \left\{ \sup_{f, f' \in \mathcal{F}_t} d(f(x_t), f'(x_t)) \geq \omega \right\} \leq \left(\frac{2\sqrt{T\beta}}{\omega} + 1 \right) \dim_1(\mathcal{F}, \omega), \quad (22)$$

$$\sum_{t=1}^T \mathbb{1} \left\{ \sup_{f, f' \in \mathcal{F}_t} d(f(x_t), f'(x_t)) \geq \omega \right\} \leq \left(\frac{4\beta}{\omega^2} + 1 \right) \dim_2(\mathcal{F}, \omega) \quad (23)$$

for any constant $\omega > 0$.

Proof of Lemma D.11. We first prove (22) and then (23).

Proof of (22). We begin by showing that if

$$\sup_{f, f' \in \mathcal{F}_t} d(f(x_t), f'(x_t)) \geq \omega$$

for some $t \in [T]$, then x_t is ω -dependent on at most $2\beta/\omega$ disjoint subsequence of its predecessors. To see this, we note that, if x_t is ω -dependent on a subsequence $(x_{i_1}, x_{i_2}, \dots, x_{i_n})$ of its predecessors, we must have

$$\sum_{s=1}^n d(f(x_{i_s}) - f'(x_{i_s})) > \omega.$$

Hence, if x_t is ω -dependent on l disjoint subsequences, we have

$$\sum_{s=1}^{t-1} d(f(x_s), f'(x_s)) > l\omega. \quad (24)$$

For the left-hand side, we also have

$$\begin{aligned}
\sum_{s=1}^{t-1} d(f(x_s), f'(x_s)) &\leq \sum_{s=1}^{t-1} d(f(x_s), \widehat{f}_t(x_s)) + \sum_{s=1}^{t-1} d(\widehat{f}_t(x_s), f'(x_s)) \\
&\leq \sqrt{T} \sqrt{\sum_{s=1}^{t-1} d^2(f(x_s), \widehat{f}_t(x_s))} + \sqrt{T} \sqrt{\sum_{s=1}^{t-1} d^2(\widehat{f}_t(x_s), f'(x_s))} \\
&\leq 2\sqrt{T}\beta
\end{aligned} \tag{25}$$

where the first inequality is the triangle inequality, and the second inequality holds by the definition of \mathcal{F}_t . Combining (24) and (25), we get that $l \leq 2\sqrt{T}\beta/\omega$.

Next, we show that for any sequence (x'_1, \dots, x'_τ) , there is at least one element that is ω -dependent on at least $\tau/d - 1$ disjoint subsequence of its predecessors, where $d := \dim_1(\mathcal{F}, \omega)$. To show this, let m be the integer satisfying $md + 1 \leq \tau \leq md + d$. We will construct m disjoint subsequences, B_1, \dots, B_m . At the beginning, let $B_i = (x'_i)$ for $i \in [m]$. If x'_{m+1} is ω -dependent on each subsequence B_1, \dots, B_m , then we are done. Otherwise, we select a subsequence B_i which x'_{m+1} is ω -independent of and append x'_{m+1} to B_i . We repeat this process for all elements with indices $j > m + 1$ until either x'_j is ω -dependent on each subsequence or $j = \tau$. For the latter, we have $\sum_{i=1}^m |B_i| \geq md$, and since each element of a subsequence B_i is ω -independent of its predecessors, we must have $|B_i| = d$ for all i . Then, x_τ must be ω -dependent on each subsequence by the definition of eluder dimension.

Finally, let's set the sequence (x'_1, \dots, x'_τ) to be the subsequence of (x_1, \dots, x_T) consisting of elements x_t for which $d(f(x_t), f'(x_t)) > \omega$. As we have established, we have

1. each x'_i is ℓ_1 -norm ω -dependent on at most $2\sqrt{T}\beta/\omega$ disjoint subsequences, and
2. some x'_i is ℓ_1 -norm ω -dependent on at least $\tau/d - 1$ disjoint subsequences.

Therefore, we must have $\tau/d - 1 \leq 2\sqrt{T}\beta/\omega$, implying that $\tau \leq (2\sqrt{T}\beta/\omega + 1)d$.

Proof of (23). The proof is quite similar to the proof of (22). For ℓ_2 -norm, following the same argument, we can show the following:

1. each x'_i is ℓ_2 -norm ω -dependent on at most $4\beta/\omega^2$ disjoint subsequences, and
2. some x'_i is ℓ_2 -norm ω -dependent on at least $\tau/d - 1$ disjoint subsequences.

Therefore, we must have $\tau/d - 1 \leq 4\beta/\omega^2$, implying that $\tau \leq (4\beta/\omega^2 + 1)d$. \square

Lemma D.12. *Following the setting of Lemma D.11, assume $d(\cdot, \cdot)$ is upper bounded by $D \leq \beta$. Then, we have*

$$\sum_{t=1}^T \sup_{f, f' \in \mathcal{F}_t} d(f(x_t), f'(x_t)) \leq 4\sqrt{T}\beta \cdot \dim_1(\mathcal{F}, 1/T) \cdot \log(DT)$$

Proof of Lemma D.12. For notational simplicity, we denote

$$w_t := \sup_{f, f' \in \mathcal{F}_t} d(f(x_t), f'(x_t))$$

Now we fix a constant $\Delta > 0$, and then we have

$$\begin{aligned}
\sum_{t=1}^T w_t &= \sum_{t=1}^T \int_0^D \mathbb{1}\{w_t \geq \delta\} d\delta \\
&= \int_0^D \sum_{t=1}^T \mathbb{1}\{w_t \geq \delta\} d\delta \\
&\leq \Delta T + \int_{\Delta}^D \sum_{t=1}^T \mathbb{1}\{w_t \geq \delta\} d\delta \\
&\leq \Delta T + \int_{\Delta}^D \left(\frac{2\sqrt{T\beta}}{\delta} + 1 \right) \dim_1(\mathcal{F}, \delta) d\delta \\
&\leq \Delta T + \int_{\Delta}^D \frac{3\sqrt{T\beta}}{\delta} \cdot \dim_1(\mathcal{F}, \delta) d\delta \\
&\leq \Delta T + 3\sqrt{T\beta} \cdot \dim_1(\mathcal{F}, \Delta) \cdot \log(D/\Delta)
\end{aligned}$$

where the second inequality is Lemma D.11, and the last inequality uses the fact that $\dim_1(\mathcal{F}, \delta)$ is non-increasing in δ . We can conclude the proof by setting $\Delta = 1/T$. \square