

ASYMPTOTIC INSTANCE-OPTIMAL ALGORITHMS FOR INTERACTIVE DECISION MAKING

Kefan Dong
Stanford University
kefandong@stanford.edu

Tengyu Ma
Stanford University
tengyuma@stanford.edu

ABSTRACT

Past research on interactive decision making problems (bandits, reinforcement learning, etc.) mostly focuses on the minimax regret that measures the algorithm’s performance on the hardest instance. However, an ideal algorithm should adapt to the complexity of a particular problem instance and incur smaller regrets on easy instances than worst-case instances. In this paper, we design the first asymptotic instance-optimal algorithm for general interactive decision making problems with finite number of decisions under mild conditions. On *every* instance f , our algorithm outperforms *all* consistent algorithms (those achieving non-trivial regrets on all instances), and has asymptotic regret $\mathcal{C}(f) \ln n$, where $\mathcal{C}(f)$ is an exact characterization of the complexity of f . The key step of the algorithm involves hypothesis testing with active data collection. It computes the most economical decisions with which the algorithm collects observations to test whether an estimated instance is indeed correct; thus, the complexity $\mathcal{C}(f)$ is the minimum cost to test the instance f against other instances. Our results, instantiated on concrete problems, recover the classical gap-dependent bounds for multi-armed bandits (Lai et al., 1985) and prior works on linear bandits (Lattimore & Szepesvari, 2017), and improve upon the previous best instance-dependent upper bound (Xu et al., 2021) for reinforcement learning.

1 INTRODUCTION

Bandit and reinforcement learning (RL) algorithms demonstrated a wide range of successful real-life applications (Silver et al., 2016; 2017; Mnih et al., 2013; Berner et al., 2019; Vinyals et al., 2019; Mnih et al., 2015; Degraeve et al., 2022). Past works have theoretically studied the regret or sample complexity of various interactive decision making problems, such as contextual bandits, reinforcement learning (RL), partially observable Markov decision process (see Azar et al. (2017); Jin et al. (2018); Dong et al. (2021); Li et al. (2019); Agarwal et al. (2014); Foster & Rakhlin (2020); Jin et al. (2020), and references therein). Recently, Foster et al. (2021) present a unified algorithmic principle for achieving the minimax regret—the optimal regret for the worst-case problem instances.

However, minimax regret bounds do not necessarily always present a full picture of the statistical complexity of the problem. They characterize the complexity of the most difficult instances, but potentially many other instances are much easier. An ideal algorithm should adapt to the complexity of a particular instance and incur smaller regrets on easy instances than the worst-case instances. Thus, an ideal regret bound should be instance-dependent, that is, depending on some properties of each instance. Prior works designed algorithms with instance-dependent regret bounds that are stronger than minimax regret bounds, but they are often not directly comparable because they depend on different properties of the instances, such as the gap conditions and the variance of the value function (Zanette & Brunskill, 2019; Xu et al., 2021; Foster et al., 2020; Tirinzoni et al., 2021).

A more ambitious and challenging goal is to design *instance-optimal* algorithms that outperform, on *every* instance, *all* consistent algorithms (those achieving non-trivial regrets on all instances). Past works designed instance-optimal algorithms for multi-armed bandit (Lai et al., 1985), linear bandits (Kirschner et al., 2021; Hao et al., 2020), Lipschitz bandits (Magureanu et al., 2014), and ergodic MDPs (Ok et al., 2018). However, instance-optimal regret bounds for tabular reinforcement learning remain an open question, despite recent progress (Tirinzoni et al., 2021; 2022). The challenge partly

stems from the fact that the *existence* of such an instance-optimal algorithm is even a priori unclear for general interactive decision making problems. Conceivably, each algorithm can have its own speciality on a subset of instances, and no algorithm can dominate all others on all instances.

Somewhat surprisingly, we prove that there exists a simple algorithm (T2C, stated in Alg. 1) that is asymptotic instance-optimal for general interactive decision making problems with finite number of decisions. We determine the *exact* leading term of the optimal asymptotic regret for instance f to be $\mathcal{C}(f) \ln n$. Here, n is the number of interactions and $\mathcal{C}(f)$ is a complexity measure for the instance f that intuitively captures the difficulty of distinguishing f from other instances (that have different optimal decisions) using observations. Concretely, under mild conditions on the interactive decision problem, our algorithm achieves an asymptotic regret $\mathcal{C}(f) \ln n$ (Theorem 5.2) for every instance f , while every consistent algorithm must have an asymptotic regret at least $\mathcal{C}(f) \ln n$ (Theorem 3.2).

Our algorithm consists of three simple steps. First, it explores uniformly for $o(1)$ -fraction of the steps and computes the MLE estimate of the instance with relatively low confidence. Then, it tests whether the estimate instance (or, precisely, its associated optimal decision) is indeed correct using the most economical set of queries/decisions. Concretely, it computes a set of decisions with minimal regret, such that, using a log-likelihood ratio test, it can either distinguish the estimated instance from all other instances (with different optimal decisions) with high confidence, or determine that our estimation was incorrect. Finally, with the high-confidence estimate, it commits to the optimal decision of the estimated instance in the rest of the steps. The algorithmic framework essentially reduces the problem to the key second step — optimal hypothesis testing with active data collection.

Our results recover the classical gap-dependent regret bounds for multi-armed bandits (Lai et al., 1985) and prior works on linear bandits (Lattimore & Szepesvari, 2017; Hao et al., 2020). As an instantiation of the general algorithm, we present the first asymptotic instance-optimal algorithm for tabular RL, improving upon prior instance-dependent algorithms (Xu et al., 2021; Simchowitz & Jamieson, 2019; Tirinzoni et al., 2021; 2022).

1.1 ADDITIONAL RELATED WORKS

Some algorithms are proved instance-optimal for specific interactive decision making problems. Variants of UCB algorithms are instance-optimal for bandits with various assumptions (Lattimore & Szepesvári, 2020; Gupta et al., 2021; Tirinzoni et al., 2020; Degenne et al., 2020; Magureanu et al., 2014), but are suboptimal for linear bandits (Lattimore & Szepesvari, 2017). These algorithms rely on the optimism-in-face-of-uncertainty principle to deal with exploration-exploitation tradeoff, whereas our algorithm explicitly computes the best tradeoff. Kirschner et al. (2021); Lattimore & Szepesvari (2017); Hao et al. (2020) design non-optimistic instance-optimal algorithms for linear bandits. There are also instance-optimal algorithms for ergodic MDPs where the regret is less sensitive to the exploration policy (Ok et al., 2018; Burnetas & Katehakis, 1997; Graves & Lai, 1997), interactive decision making with finite hypothesis class, finite state-action space, and known rewards (Rajeev et al., 1989), and interactive decision making with finite observations (Komiya et al., 2015).

The most related problem setup is structured bandits (Combes et al., 2017; Van Parys & Golrezaei, 2020; Jun & Zhang, 2020), where the instances also belong to an abstract and arbitrary family \mathcal{F} . The structured bandits problem is a very special case of general decision making problems and does not contain RL because the observation is a scalar. In contrast, the observation in general decision making problems could be high-dimensional (e.g., a trajectory with multiple states, actions, and rewards for episodic RL), which makes our results technically challenging.

Many algorithms' regret bounds depend on some properties of instances such as the gap condition. Foster et al. (2020) prove a gap-dependent regret bound for contextual bandits. For reinforcement learning, the regret bound may depend on the variance of the optimal value function (Zanette & Brunskill, 2019) or the gap of the Q -function (Xu et al., 2021; Simchowitz & Jamieson, 2019; Yang et al., 2021). Xu et al. (2021); Foster et al. (2020) also prove that the gap-dependent bounds cannot be improved on some instances. To some extent, these instance-dependent lower bounds can be viewed as minimax bounds for a fine-grained instance family (e.g., all instances with the same Q -function gap), and therefore are different from ours.

2 SETUP AND NOTATIONS

Interactive decision making problems. We focus on interactive decision making problem with structured observations (Foster et al., 2021), which includes bandits and reinforcement learning as special cases. An interactive decision making problem is defined by a space of decisions Π , observations O , a reward function $R : O \rightarrow \mathbb{R}$, and a function f (also called an instance) that maps a decision $\pi \in \Pi$ to a distribution over observations $f[\pi]$. We use $f[\pi](\cdot) : O \rightarrow \mathbb{R}_+$ to denote the density function of the distribution $f[\pi]$. We assume that the reward R is a deterministic and known.

An environment picks an ground-truth instance f^* from a instance family \mathcal{F} , and then an agent (which knows the instance family \mathcal{F} but not the ground-truth f^*) interacts with the environment for n total rounds. In round $t \leq n$,

- (a) the learner selects a decision π_t from the decision class Π , and
- (b) the environment generates an observation o_t following the ground-truth distribution $f^*[\pi_t]$ and reveals the observation. Then the agent receives a reward $R(o_t)$.

For example, multi-armed bandits, linear bandits, and *episodic* reinforcement learning all belong to interactive decision making problems. For bandits, a decision π corresponds to an action and an observation o corresponds to a reward. For reinforcement learning, a decision π is a deterministic policy that maps from states to actions, and an observation o is a trajectory (including the reward at each step). In other words, one round of interactions in the interactive decision making problem corresponds to one episode of the RL problem. We will formally discuss the RL setup in Section 5.

Let $R_f(\pi) = \mathbb{E}_{o \sim f[\pi]}[R(o)]$ be the expected reward for decision π under instance f , and $\pi^*(f) \triangleq \operatorname{argmax}_{\pi} R_f(\pi)$ the optimal decision of instance f . The expected regret measures how much worse the agent’s decision is than the optimal decision:

$$\operatorname{Reg}_{f,n} = \mathbb{E}_f \left[\sum_{t=1}^n \left(\max_{\pi \in \Pi} R_f(\pi) - R_f(\pi_t) \right) \right]. \quad (1)$$

We consider the case where the decision family Π is *finite*, and every instance $f \in \mathcal{F}$ has a unique optimal decision, denoted by $\pi^*(f)$. We also assume that for every $f \in \mathcal{F}$ and $\pi \in \Pi$, $0 \leq R(o) \leq R_{\max}$ almost surely when o is drawn from the distribution $f[\pi]$, and $\sup_o f[\pi](o) < \infty$.

Consistent algorithms and asymptotic instance-optimality. We first note that it’s unreasonable to ask for an algorithm that can outperform or match any *arbitrary* algorithm on *every* instance. This is because, for any instance $f \in \mathcal{F}$, a bespoke algorithm that always outputs $\pi^*(f)$ can achieve zero regret on instance f (though it has terrible regrets for other instances). Therefore, if an algorithm can outperform or match any algorithm on any instance, it must have zero regret on every instance, which is generally impossible. Instead, we are interested in finding an algorithm that is as good as any other reasonable algorithms that are not completely customized to a single instance.

We say an algorithm is consistent if its expected regret satisfies $\operatorname{Reg}_{f,n} = o(n^p)$ for *every* $p > 0$ and $f \in \mathcal{F}$ (Lai et al., 1985). Most of the reasonable algorithms are consistent, such as the UCB algorithm for multi-armed bandits (Lattimore & Szepesvári, 2020), UCBVI, and UCB-Q algorithm for tabular reinforcement learning (Simchowitz & Jamieson, 2019; Yang et al., 2021), because all of them achieve asymptotically $\mathcal{O}(\ln n)$ regrets on any instance, where \mathcal{O} hides constants that depend on the property of the particular instance.¹ However, consistent algorithms exclude the algorithm mentioned in the previous paragraph which is purely customized to a particular instance.

We say an algorithm is asymptotically *instance-optimal* if on *every* instance $f \in \mathcal{F}$, $\limsup_{n \rightarrow \infty} \operatorname{Reg}_{f,n} / \ln n$ is the smallest among all *consistent* algorithms (Lai et al., 1985). We note that even though an instance-optimal algorithm only needs to perform as well as every consistent algorithm, a priori, it’s still unclear if such an algorithm exists.

Some prior works have also used a slightly weaker definition in place of consistent algorithms, e.g., α -uniformly good algorithms in Tirinzoni et al. (2021), which allows a sublinear regret bound $\mathcal{O}(n^\alpha)$ for *some* constant $\alpha < 1$. The alternative definition, though apparently includes more algorithms, does not change the essence. Our algorithm is still instance-optimal *up to a constant factor*—simple modification of the lower bound part of the proof shows that its asymptotic regret is at most

¹Here the regret scales in $\log n$ because we are in the asymptotic setting where the instance is fixed and n tends to infinity. If the instance can depend on n (which is not the setting we are interested in), then the minimax regret typically scales in $O(\sqrt{n})$.

$(1 - \alpha)^{-1}$ factor bigger than any α -uniformly good algorithms on any instance. This paper thus primarily compares with consistent algorithms for simplicity.

3 MAIN RESULTS

In this section, we present an intrinsic complexity measure $\mathcal{C}(f)$ of an instance f (Section 3.1), and an instance-optimal algorithm that achieves an asymptotic regret $\mathcal{C}(f) \ln n$ (Section 3.2).

3.1 COMPLEXITY MEASURE AND REGRET LOWER BOUNDS

In this section, our goal is to understand the minimum regret of consistent algorithms, which is an intrinsic complexity measure of an instance. We define an instance-dependent complexity measure $\mathcal{C}(f)$ and prove that any consistent algorithm must have asymptotic regret at least $\mathcal{C}(f) \ln n$.

The key observation is that any consistent algorithm has to collect enough information from the observations to tell apart the ground-truth instance from other instances with different optimal decisions. Consider the situation when a sequence of n decisions is insufficient to distinguish two instances, denoted by f and g , with different optimal decisions. If an algorithm achieves a sublinear regret on f , the sequence must contain $\pi^*(f)$ (the optimal decision for f) at least $n - o(n)$ times. As a result, if the true instance were g , the algorithm suffers a linear regret (due to the linear number of $\pi^*(f)$), and therefore cannot be consistent.

An ideal algorithm should find out decisions that can *most efficiently* identify the ground-truth instance (or precisely the family of instances with the same optimal decision as the ground-truth). However, decisions collect information but also incur regrets. So the algorithm should pick a list of decisions with the best tradeoff between these two effects—minimizing the decisions’ regret and collecting sufficient information to identify the true instance. Concretely, suppose a sequence of decisions includes w_π occurrences of decision π . The regret of these decisions is $\sum_{\pi \in \Pi} w_\pi \Delta(f, \pi)$, where $\Delta(f, \pi) \triangleq R_f(\pi^*(f)) - R_f(\pi)$ is the sub-optimality gap of decision π . We use KL divergence between the observations of π under two instances f and g (denoted by $D_{\text{KL}}(f[\pi] \| g[\pi])$) to measure π ’s power to distinguish them. The following optimization problem defines the optimal mixture of the decisions (in terms of the regret) that has sufficient power to distinguish f from all other instances with different optimal decision.

$$\mathcal{C}(f, n) \triangleq \min_{w \in \mathbb{R}_+^{|\Pi|}} \sum_{\pi \in \Pi} w_\pi \Delta(f, \pi) \quad (2)$$

$$\text{s.t. } \sum_{\pi \in \Pi} w_\pi D_{\text{KL}}(f[\pi] \| g[\pi]) \geq 1, \forall g \in \mathcal{F}, \pi^*(g) \neq \pi^*(f), \quad (3)$$

$$\|w\|_\infty \leq n. \quad (4)$$

The last constraint makes sure that $w_\pi < \infty$ even if the decision has no regret (i.e., $\exists \pi, \Delta(f, \pi) = 0$). We only care about the case when n approaches infinity. The asymptotic complexity of f is

$$\mathcal{C}(f) \triangleq \lim_{n \rightarrow \infty} \mathcal{C}(f, n). \quad (5)$$

In Eq. (3), we only require separation between instances f, g when they have different optimal decisions. As there are only finite decisions, there are finite equivalent classes, so f and g are in principle separable with sufficient number of observations. Thus, the complexity measure $\mathcal{C}(f, n)$ is well defined as stated in the following lemma.

Lemma 3.1. *For any $f \in \mathcal{F}$, $\mathcal{C}(f, n)$ is non-increasing in n , and there exists $n_0 > 0$ such that for all $n > n_0$, $\mathcal{C}(f, n) < \infty$. As a corollary, $\mathcal{C}(f) < \infty$ and is well defined.*

Proof of Lemma 3.1 is deferred to Appendix B.1. In Section 3.1, we formalize the intuition above and prove that $\mathcal{C}(f)$ is a lower bound of the asymptotic regret, as stated in the following theorem.

Theorem 3.2. *For every instance $f \in \mathcal{F}$, the expected regret of any consistent algorithm satisfies $\limsup_{n \rightarrow \infty} \frac{\text{Reg}_{f,n}}{\ln n} \geq \mathcal{C}(f)$.*

We prove Theorem 3.2 in Appendix B.2. Theorem 3.2 is inspired by previous works (Lattimore & Szepesvari, 2017; Tirinzoni et al., 2021). When applied to tabular RL problems, our lower bound is very similar to that in Tirinzoni et al. (2021) with the only difference that their optimization problems omit the second constraint (Eq. (4)), which can make it slightly less tight (see Section B.5).

The complexity measure $\mathcal{C}(f)$ reduces to the well-known inverse action gap bound $\mathcal{O}(\sum_{a \in \mathcal{A}, \Delta(a) > 0} 1/\Delta(a))$ for multi-armed bandits (Proposition B.2), and recovers the result of Lattimore & Szepesvari (2017) for linear bandits (Proposition B.3). For reinforcement learning, Tirinzoni et al. (2021) prove that the instance-optimal bound can be smaller than the gap-dependent bound $\mathcal{O}(\sum_{s,a: \Delta(s,a) > 0} 1/\Delta(s,a))$ in Xu et al. (2021) and Simchowitz & Jamieson (2019).

3.2 INSTANCE-OPTIMAL ALGORITHMS

Algorithm 1 Test-to-Commit (T2C)

1: Parameters: the number of rounds of interactions n .

Step 1: Initialization.

- 2: Play each decision $\lceil \frac{\ln n}{\ln \ln n} \rceil$ times. We denote these decisions by $\{\hat{\pi}_i\}_{i=1}^{m_{\text{init}}}$, where $m_{\text{init}} = \lceil \Pi \lceil \frac{\ln n}{\ln \ln n} \rceil \rceil$, and the corresponding observations by $\{\hat{o}_i\}_{i=1}^{m_{\text{init}}}$.
- 3: Compute the max likelihood estimation (MLE) with arbitrary tie-breaking

$$\hat{f} = \operatorname{argmax}_{f \in \mathcal{F}} \sum_{i=1}^{m_{\text{init}}} \ln f[\hat{\pi}_i](\hat{o}_i). \quad (6)$$

Step 2: Identification.

- 4: Let $\eta = (\ln \ln n)^{1/4}$ for shorthand, and \hat{w} be the solution of the program defining $\mathcal{C}(\hat{f}, \eta)$. Compute $\bar{w}_\pi = (1 + 1/\eta)\hat{w}_\pi + 1/\eta$ for all $\pi \in \Pi$.
- 5: Play each decision π for $\lceil \bar{w}_\pi \ln n \rceil$ times. Denote these decisions by $w = \{\pi_i\}_{i=1}^m$ where $m = \sum_{\pi} \lceil \bar{w}_\pi \ln n \rceil$, and the corresponding observations by $\{o_i\}_{i=1}^m$.
- 6: Run the log-likelihood ratio test on instance \hat{f} , the sequence of decision $\{\pi_i\}_{i=1}^m$ and its corresponding observations $\{o_i\}_{i=1}^m$ (that is, compute the event $\mathcal{E}_{\text{acc}}^{\hat{f}}$ defined in Eq. (8)).

Step 3: Exploitation.

- 7: **if** $\mathcal{E}_{\text{acc}}^{\hat{f}} = \text{true}$ **then**
- 8: Commit to $\pi^*(\hat{f})$ (i.e., run $\pi^*(\hat{f})$ for the remaining steps).
- 9: **else**
- 10: Run UCB for the remaining steps.
-

We first present the regret bound for our algorithm of the T2C algorithm. For simplicity, we consider a finite hypothesis (i.e., $|\mathcal{F}| < \infty$) here, and extend to infinite hypothesis case in Section 5.

We start by stating a condition that excludes abnormal observation distributions $f[\pi]$. Recall that for any $\zeta \in (0, 1)$, the Rényi divergence of two distributions p, q is

$$D_\zeta(p||q) = \frac{1}{\zeta-1} \ln \int_x p(x)^\zeta q(x)^{1-\zeta} dx. \quad (7)$$

The Rényi divergence $D_\zeta(p||q)$ is non-decreasing in ζ , and $\lim_{\zeta \uparrow 1} D_\zeta(p||q) = D_{\text{KL}}(p||q)$ (Van Erven & Harremoës, 2014). We require the limits converge uniformly for all instances $g \in \mathcal{F}$ and decisions $\pi \in \Pi$ with a ζ bounded away from 1 (the choice of the constants in Condition 1 is not critical to our result), as stated below.

Condition 1. For any fixed $\alpha > 0, \epsilon > 0$, instance $f \in \mathcal{F}$, there exists $\lambda_0(\alpha, \epsilon, f) > 0$ such that for all $\lambda \leq \lambda_0(\alpha, \epsilon, f)$, $g \in \mathcal{F}$ and $\pi \in \Pi$, $D_{1-\lambda}(f[\pi]||g[\pi]) \geq \min\{D_{\text{KL}}(f[\pi]||g[\pi]) - \epsilon, \alpha\}$. Moreover, we require $\lambda_0(\alpha, \epsilon, f) \geq \epsilon^{c_1} \min\{1/\alpha, c_2\}^{c_3} \iota(f)$ for some universal constants $c_1, c_2, c_3 > 0$, where $\iota(f) > 0$ is a function that only depends on f .

Condition 1 holds for a wide range of distributions, such as Gaussian, Bernoulli, multinomial, Laplace with bounded mean, Log-normal (Gil et al., 2013), and tabular RL where $f[\pi]$ is a distribution over a trajectory consists of state, action and reward tuples (see Theorem 5.1 for the proof of tabular RL). A stronger but more interpretable variant of Condition 1 is that the log density ratio of $f[\pi]$ and $g[\pi]$ has finite fourth moments (see Proposition B.4), therefore Condition 1 can also be potentially verified for other distributions.

The main theorem analyzing Alg. 1 is shown below. The asymptotic regret of Alg. 1 matches the constants in the lower bound (Theorem 3.2), indicating its asymptotic instance-optimality.

Theorem 3.3. Suppose \mathcal{F} is a finite hypothesis class and satisfies Condition 1 The regret of Alg. 1 satisfies $\limsup_{n \rightarrow \infty} \frac{\text{Reg}_{f^*, n}}{\ln n} \leq \mathcal{C}(f^*)$.

Our algorithm is stated in Alg. 1 and consists of three steps: initialization, identification, and exploitation. In the initialization step, we explore uniformly for a short period and compute the MLE estimate \hat{f} of the true instance (Line 3 of Alg. 1), where we only requires \hat{f} to be accurate with moderate probability (i.e., $1 - 1/\ln n$). The estimation is used to compute the lowest-regret list of decisions that can distinguish the optimal decision of \hat{f} . Since we only collect $o(\ln n)$ samples, the regret of this step is negligible asymptotically.

In the identification step, we hold the belief that \hat{f} is the true instance and solve $\mathcal{C}(\hat{f}, (\ln \ln n)^{1/4})$ to get the best list of decisions to fortify this belief (see Line 4 of Alg. 1). Then we collect more samples using this list (with minor decorations) to boost the confidence of our initial estimation to $1 - 1/n$ (or reject it when the estimation is not accurate) by the following log-likelihood ratio test:

$$\mathcal{E}_{\text{acc}}^{\hat{f}} = \mathbb{I}[\forall g \in \mathcal{F} \text{ and } \pi^*(g) \neq \pi^*(\hat{f}), \sum_{i=1}^m \ln \frac{\hat{f}[\pi_i](o_i)}{g[\pi_i](o_i)} \geq \ln n]. \quad (8)$$

Intuitively, if \hat{f} is not the ground-truth f^* , $\mathcal{E}_{\text{acc}}^{\hat{f}}$ is unlikely to hold because the expected log-likelihood ratio is non-positive: $\mathbb{E}_{o \sim f^*[\pi]}[\ln(\hat{f}[\pi](o)/f^*[\pi](o))] = -D_{\text{KL}}(f^*[\pi] \parallel \hat{f}[\pi]) \leq 0$. Hence, with high probability a wrong guess cannot be accepted. On the other hand, the first constraint (Eq. (3)) in the definition of $\mathcal{C}(\hat{f})$ guarantees that when $\hat{f} = f^*$, the expected log-likelihood ratio is large for all $g \in \mathcal{F}$ and $\pi^*(g) \neq \pi^*(\hat{f})$, so an accurate guess will be accepted. In this step $m = \tilde{\Theta}(\ln n)$, so Step 2 dominates the regret of Alg. 1. In other words, the efficiency of the log-likelihood test is critical to our analysis.

Finally in the exploitation step, the algorithm commits to the optimal decision of \hat{f} if it believes this estimation with confidence $1 - 1/n$, or run a standard UCB algorithm as if on a multi-armed bandits problem (with a well-known $\mathcal{O}(\ln n)$ regret) when the initial estimation is not accurate (happens with probability at most $1/\ln n$). Hence, the expected regret here is $\mathcal{O}(\ln n)/\ln n + \mathcal{O}(n)/n = \mathcal{O}(1)$.

Our three-stage algorithm is inspired by Lattimore & Szepesvari (2017), where their test in Step 2 only uses the reward information. However, the observation for general interactive decision making problem contains much more information such as the transition in tabular RL problems. In contrast, the T2C algorithm relies on a general hypothesis testing method (i.e., log-likelihood ratio test) and achieves optimal regret for general interactive decision making problems.

4 PROOF SKETCHES OF THE MAIN RESULTS

In this section, we discuss the proof sketch of our main results. Section 4.1 discusses the proof sketch of the lower bound, and Section 4.2 shows the main lemmas for each step in Alg. 1. In Section 4.3, we discuss the log-likelihood ratio test in detail.

4.1 PROOF SKETCH OF THEOREM 3.2

Recall that $\mathcal{C}(f)$ is the minimum regret of a list of decisions that can distinguish the instance f with all other instances g (with a different optimal decision). Hence, to prove the lower bound, we show that the sequence of decisions played by any consistent algorithm must also distinguish f , and thus, $\mathcal{C}(f)$ lower bounds the regret of any consistent algorithm.

For any consistent algorithm, number of interactions $n > 0$, and two instances $f, g \in \mathcal{F}$ with different optimal decisions, let $P_{f,n}$ and $P_{g,n}$ denote the probability space generated by running the algorithm on instances f and g respectively for n rounds. Since f, g have different optimal decisions, $P_{f,n}$ and $P_{g,n}$ should be very different — following the same proof strategy in (Lattimore & Szepesvari, 2017), we can show that

$$D_{\text{KL}}(P_{f,n} \parallel P_{g,n}) \geq (1 + o(1)) \ln n.$$

Let the random variable π_i be the decision of the algorithm in round i . The chain rule of KL divergence shows

$$D_{\text{KL}}(P_{f,n} \parallel P_{g,n}) = \mathbb{E}_{f,n} \left[\sum_{i=1}^n D_{\text{KL}}(f[\pi_i] \parallel g[\pi_i]) \right] = \sum_{\pi} \mathbb{E}_{f,n} [N_{\pi}] D_{\text{KL}}(f[\pi] \parallel g[\pi]).$$

Now consider the vector $w \in \mathbb{R}_+^{|\Pi|}$ where $w_{\pi} = \mathbb{E}_f [N_{\pi}] / ((1 + o(1)) \ln n)$. Based on the derivation above, we can verify that w is a valid solution to $\mathcal{C}(f, n)$, and therefore

$$\text{Reg}_{f,n} = \sum_{\pi} \mathbb{E}_f [N_{\pi}] \Delta(f, \pi) = \sum_{\pi} w_{\pi} \Delta(f, \pi) (1 + o(1)) \ln n \geq \mathcal{C}(f, n) (1 + o(1)) \ln n.$$

Then the final result is proved by the fact that $\mathcal{C}(f) = \lim_{n \rightarrow \infty} \mathcal{C}(f, n)$.

4.2 PROOF SKETCH OF THEOREM 3.3

In the following, we discuss main lemmas and their proof sketches for the three steps of Alg. 1.

Step 1: Initialization. In this step we show that the max likelihood estimation can find the exact instance (i.e., $\hat{f} = f^*$) with probability at least $1 - 1/\ln n$ and negligible regret. Note that the failure probability is not small enough to directly commit to the optimal decision of \hat{f} (i.e., play $\pi^*(\hat{f})$ forever), which would incur linear regret when $\hat{f} \neq f^*$. Formally speaking, the main lemma for Step 1 is stated as follows, whose proof is deferred to Appendix A.1.

Lemma 4.1. *Under Condition 1, with probability at least $1 - 1/\ln n$ we get $\hat{f} = f^*$. In addition, the regret of Step 1 is upper bounded by $\mathcal{O}(\frac{\ln n}{\ln \ln n})$.*

Lemma 4.1 is not surprising since the MLE only requires $\mathcal{O}(\log(1/\delta))$ samples to reach a failure probability δ in general (Van de Geer, 2000, Section 7). Here we require $1/\ln n$ failure probability but allow $\ln n/(\ln \ln n) = \Omega(\ln \ln n)$ samples, hence the result is expected for large enough n .

Step 2: Identification. In the identification step, we boost the failure probability to $1/n$ using the log-likelihood test. To this end, we compute the optimal list of decisions $w = \{\pi_1, \dots, \pi_m\}$ that distinguishes \hat{f} by solving $\mathcal{C}(\hat{f}, (\ln \ln n)^{1/4})$ (Line 4 of Alg. 1). The choice of $(\ln \ln n)^{1/4}$ is not critical, and could be replaced by any smaller quantity that approaches infinity as $n \rightarrow \infty$. Then we run the log-likelihood ratio test using the observations collected by executing the list of decision w , and achieve the following:

- (a) when the true instance f^* and the estimation \hat{f} have different optimal decisions, accept \hat{f} with probability at most $1/n$;
- (b) when $\hat{f} = f^*$, accept \hat{f} with probability at least $1 - 1/\ln n$.

In Section 4.3, we will discuss the log-likelihood ratio test in detail. Note that the regret after we falsely reject the true instance is $\mathcal{O}(\ln n)$ (by the regret bound of UCB algorithm), so we only require a $1/\ln n$ failure probability for (b) because then it leads to a $\mathcal{O}(\ln n)/\ln n = \mathcal{O}(1)$ expected regret. The main lemma for Step 2 is stated as follows, and its proof is deferred to Appendix A.2

Lemma 4.2. *Under Condition 1, for any finite hypothesis \mathcal{F} , for large enough n the following holds:*

- (a) *conditioned on the event $\pi^*(\hat{f}) \neq \pi^*(f^*)$, $\mathcal{E}_{\text{acc}}^{\hat{f}}$ is true with probability at most $1/n$;*
- (b) *conditioned on the event $\hat{f} = f^*$, $\mathcal{E}_{\text{acc}}^{\hat{f}}$ is true with probability at least $1 - 1/\ln n$;*
- (c) *conditioned on the event $\hat{f} = f^*$, the expected regret of Step 2 is upper bounded by $(\mathcal{C}(f^*, (\ln \ln n)^{1/4}) + o(1)) \ln n$. Otherwise, the expected regret of Step 2 is upper bounded by $\mathcal{O}(\ln n \ln \ln n)$.*

Step 3: Exploitation. Finally, in Step 3 we either commits to the optimal decision $\pi^*(\hat{f})$ when the estimation \hat{f} is accepted, or run a standard UCB algorithm with $\mathcal{O}(\ln n)$ regret (Lattimore & Szepesvári, 2020). Combining Lemma 4.1 and Lemma 4.2 we can prove that

- (a) the probability of Step 3 incurring a $\mathcal{O}(n)$ regret is at most $1/n$, and
- (b) the probability of Step 3 incurring a $\mathcal{O}(\ln n)$ regret is at most $1/\ln n$.

As a result, the expected regret in Step 3 is $\mathcal{O}(1)$ and negligible. Finally Theorem 3.3 is proved by stitching the three steps together, and deferred to Appendix A.4.

4.3 THE LOG-LIKELIHOOD RATIO TEST

The goal of the log-likelihood ratio test is to boost the confidence of our initial estimation \hat{f} to $1 - 1/n$, so that the algorithm can safely commit to its optimal decision in Step 3. In other words, the test should reject a wrong estimation but also accept a correct one.

Formally speaking, in the test we observe a list of observations $\{o_1, \dots, o_m\}$ collected by the list of decision $\{\pi_1, \dots, \pi_m\}$ on the *true instance* f^* , and achieve the following two goals simultaneously,

- (a) when \hat{f} and f^* are sufficiently different (i.e., their KL divergence is large in the sense that $\sum_{i=1}^m D_{\text{KL}}(\hat{f}[\pi_i] \| f^*[\pi_i]) \geq (1 + o(1)) \ln n$), accept \hat{f} with probability at most $1/n$, and
- (b) when $\hat{f} = f^*$, accept \hat{f} with probability at least $1 - 1/\ln n$.

We prove that the log-likelihood ratio test with proper parameters achieves (a) and (b) under Condition 1 in the next lemma, whose proof is deferred to Appendix A.3.

Lemma 4.3. *Given two sequences of distributions $P = \{P_i\}_{i=1}^m$ and $Q = \{Q_i\}_{i=1}^m$, and a sequence of independent random variables $o = \{o_i\}_{i=1}^m$. For any fixed $\lambda > 0, c > 0$, and $\beta = \frac{1}{m} \sum_{i=1}^m D_{1-\lambda}(P_i \| Q_i)$, the test event $\mathcal{E}_{\text{acc}} = \mathbb{I}[\sum_{i=1}^m \ln \frac{P_i(o_i)}{Q_i(o_i)} \geq c]$ satisfies*

$$\Pr_{o \sim Q}(\mathcal{E}_{\text{acc}}) \leq \exp(-c), \quad (9)$$

$$\Pr_{o \sim P}(\mathcal{E}_{\text{acc}}) \geq 1 - \exp(-\lambda(m\beta - c)). \quad (10)$$

To use Lemma 4.3 in the proof of Lemma C.2, we set $c = \ln n$ and $\lambda = \text{poly}(\ln \ln n)$. Note that the choice of w in Line 4 of Alg. 1 and Condition 1 implies $m\beta = \sum_{i=1}^m D_{1-\lambda}(P_i \| Q_i) \gtrsim \sum_{i=1}^m D_{\text{KL}}(\hat{f}[\pi_i] \| f^*[\pi_i]) = (1 + o(1)) \ln n$ when \hat{f}, f^* have different optimal policies. So the conclusion of this lemma matches the first two items in Lemma 4.2.

Lemma 4.3 is closely related to the Chernoff-Stein lemma (see Chernoff (1959) and Mao (2021, Theorem 4.11)), with the difference that the failure probability of Eq. (10) in classical Chernoff-Stein lemma is a constant, while here it decreases with m . The proof of Eq. (9) is the same as in Chernoff-Stein lemma, and proof of Eq. (10) only requires an one-sided concentration of the empirical log-likelihood ratio. Indeed, the expectation of empirical log-likelihood ratio can be lower bounded by $\mathbb{E}_P[\sum_{i=1}^m \ln \frac{P_i(o_i)}{Q_i(o_i)}] = \sum_{i=1}^m D_{\text{KL}}(P_i \| Q_i) \geq \sum_{i=1}^m D_{1-\lambda}(P_i \| Q_i) = m\beta$. Hence, $\Pr_P(\mathcal{E}_{\text{acc}})$ is the probability that the empirical log-likelihood ratio is (approximately) larger than its expectation. We also note that the concentration is non-trivial because we do not make assumptions on the boundedness on the tail of Q_i .

5 EXTENSIONS TO INFINITE HYPOTHESIS CLASS

Now we extend our analysis to infinite hypothesis settings, and instantiate our results on tabular RL. Our results in the infinite hypothesis case need two additional conditions. The first condition requires an upper bound on covering number of the hypothesis, and is used to prove a infinite-hypothesis version of Lemma 4.2. Formally speaking, for any $f, g \in \mathcal{F}$, define their distance as

$$d(f, g) = \sup_{\pi \in \Pi, o} |f[\pi](o) - g[\pi](o)|. \quad (11)$$

An ϵ (proper) covering of \mathcal{F} is a subset $\mathcal{C} \subseteq \mathcal{F}$, such that for any $g \in \mathcal{F}$, there exists $g' \in \mathcal{C}$ with $d(g, g') \leq \epsilon$. The covering number $\mathcal{N}(\mathcal{F}, \epsilon)$ is the size of the minimum ϵ covering of \mathcal{F} .

Condition 2. *There exists constant c that depends on \mathcal{F} such that $\ln \mathcal{N}(\mathcal{F}, \epsilon) \leq \mathcal{O}(c \ln(1/\epsilon))$ for every $\epsilon > 0$. In addition, the base measure of the probability space has a finite volume $\text{Vol} < \infty$.*

The second condition upper bounds the distance of two instances by a polynomial of their KL divergence, and is used to prove a stronger version of Lemma 4.1.

Condition 3. *There exists a constant $c_{\min} > 0$ (which may depend on f^*) such that for all $\pi \in \Pi$ and $o \in \text{supp}(f^*[\pi])$ $f^*[\pi](o) > c_{\min}$. In addition, there exists a constant $\iota(f^*) > 0$ that only depends on f^* and $c_5 > 0$, such that for all $f \in \mathcal{F}, \pi \in \Pi$, when $D_{\text{KL}}(f^*[\pi] \| f[\pi]) \leq 1$*

$$\|f^*[\pi] - f[\pi]\|_{\infty} \leq \iota(f^*) D_{\text{KL}}(f^*[\pi] \| f[\pi])^{c_5}.$$

A lot of interactive decision making problems satisfy Conditions 2, 3, such as multi-armed bandits and linear bandits with Bernoulli reward or truncated Gaussian reward. For bandit problems, Conditions 2 and 3 only depend on the noise of the reward function and can be verified easily.

For tabular RL, an observation o denotes a trajectory $(s_1, a_1, r_1, \dots, s_H, a_H, r_H)$ where the state-action pairs (s_h, a_h) are discrete random variables but the rewards r_h are continuous. A decision

$\pi : \mathcal{S} \rightarrow \mathcal{A}$ is a mapping from the state space to action space, so $|\Pi| = |\mathcal{A}|^{|\mathcal{S}|} < \infty$. In Appendix D, we formally define the tabular RL problems where the reward r_h given (s_h, a_h) follows a truncated Gaussian distribution, and prove that Conditions 1-3 are satisfied as stated in the following theorem.

Theorem 5.1. *Let \mathcal{F} be the family of tabular RL with truncated Gaussian reward and unique optimal policies. Then Conditions 1, 2, 3 holds simultaneously.*

For tabular RL with truncated Gaussian reward, the observation o is a mixture of discrete random variables (i.e., the states and actions s_h, a_h) and continuous random variables (i.e., the rewards r_h). To prove Conditions 1-3, we deal with the discrete and continuous part of the observation separately. We also have flexibility in the reward distribution, e.g., our proof technique can deal with other distributions such as truncated Laplace distribution. We present these unified conditions for bounded random variables, but Conditions 2 and 3 do not hold for unbounded random variables because the base measure has an infinite volume (which contradicts to Condition 2), and the density of the observation cannot be lower bounded (which contradicts to Condition 3). However, we can still prove the same result using a slightly different approach (see Appendix E).

With Conditions 1-3, we can prove our main results for infinite hypothesis class. The asymptotic regret of Alg. 1 in this case still matches the lower bound exactly. The proof of Theorem 5.2 is deferred to Appendix C.1.

Theorem 5.2. *Suppose \mathcal{F} is an infinite hypothesis class that satisfies Conditions 1-3, the regret of Alg. 1 satisfies $\limsup_{n \rightarrow \infty} \frac{\text{Reg}_{f^*, n}}{\ln n} \leq \mathcal{C}(f^*)$.*

As a corollary of Theorem 5.1 and Theorem 5.2, our algorithm strictly improves upon previous best gap-dependent bounds $\mathcal{O}(\sum_{s,a:\Delta(s,a)>0} 1/\Delta(s,a))$ in Xu et al. (2021) and Simchowitz & Jamieson (2019) because the gap-dependent bounds are not tight for many instances (Tirinzi et al., 2021).

In the following we discuss the challenges when extending to infinite hypothesis settings.

Covering Number. With infinite hypothesis, we need to accept an accurate estimation even when there are infinitely many other choices. Recall that the accept event is

$$\mathcal{E}_{\text{acc}}^{\hat{f}} = \mathbb{I} \left[\forall g \in \mathcal{F} \text{ and } \pi^*(g) \neq \pi^*(\hat{f}), \sum_{i=1}^m \ln \frac{\hat{f}[\pi_i](o_i)}{g[\pi_i](o_i)} \geq \ln n \right]. \quad (12)$$

So informally speaking, we need to show that with high probability,

$$\forall g \in \mathcal{F} \text{ and } \pi^*(g) \neq \pi^*(\hat{f}), \quad \sum_{i=1}^m \ln \frac{\hat{f}[\pi_i](o_i)}{g[\pi_i](o_i)} \gtrsim (D_{1-\lambda}^w(\hat{f} \| g) - \epsilon)m \geq \ln n, \quad (13)$$

which is essentially a uniform concentration inequality as discussed in Section 4.3. So we resort to a covering number argument. The standard covering argument is not directly suitable in this case — even if $d(g, g')$ is small, it's still possible that $\ln(f[\pi](o)/g[\pi](o))$ is very different from $\ln(f'[\pi](o)/g'[\pi](o))$ (especially when $g[\pi](o)$ is very close to 0). Instead, we consider the distribution with density $\hat{g}[\pi](o) \triangleq (g'[\pi](o) + \epsilon)/Z$ where Z is the normalization factor, and only prove a one-sided covering (that is, $\ln(f[\pi](o)/g[\pi](o)) \geq \ln(f[\pi](o)/\hat{g}[\pi](o)) - \mathcal{O}(\epsilon)$). We state and prove the uniform concentration in Appendix F.

Initialization. With infinite hypothesis, we cannot hope to recover the true instance f^* exactly — some instances can be arbitrarily close to f^* and thus indistinguishable. Instead, we prove that the estimation in Step 1 satisfies $\sup_{\pi \in \Pi} D_{\text{KL}}(f^*[\pi] \| \hat{f}[\pi]) \leq \text{poly}(\frac{\ln \ln n}{\ln n})$. The main lemma of Step 1 in the infinite hypothesis class case is stated in Lemma C.1.

6 CONCLUSION

In this paper, we design instance-optimal algorithms for general interactive decision making problems with finite decisions. As an instantiation, our algorithm is the first instance-optimal algorithm for tabular MDPs. For future works, we raise the following open questions.

- To implement Alg. 1, we need to solve $\mathcal{C}(f, (\ln \ln n)^{1/4}/2)$, which is a linear programming with $|\Pi|$ variables and infinitely many constraints. However, $|\Pi|$ is exponentially large for tabular MDPs. Can we compute this optimization problem efficiently for tabular MDPs?
- Although our algorithm is asymptotically optimal, the lower order terms may dominate the regret unless n is very large. Can we design non-asymptotic instance optimal algorithms?

ACKNOWLEDGMENT

The authors would like to thank Yuanhao Wang, Jason D. Lee for helpful discussions, and the support from JD.com.

REFERENCES

- Alekh Agarwal, Daniel Hsu, Satyen Kale, John Langford, Lihong Li, and Robert Schapire. Taming the monster: A fast and simple algorithm for contextual bandits. In *International Conference on Machine Learning*, pp. 1638–1646, 2014.
- Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning*, pp. 263–272, 2017.
- Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemyslaw Debiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, et al. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*, 2019.
- Apostolos N Burnetas and Michael N Katehakis. Optimal adaptive policies for Markov decision processes. *Mathematics of Operations Research*, 22(1):222–255, 1997.
- Herman Chernoff. Sequential design of experiments. *The Annals of Mathematical Statistics*, 30(3): 755–770, 1959.
- Richard Combes, Stefan Magureanu, and Alexandre Proutiere. Minimal exploration in structured stochastic bandits. *Advances in Neural Information Processing Systems*, 30, 2017.
- Thomas M Cover. *Elements of information theory*. John Wiley & Sons, 1999.
- Rémy Degenne, Han Shao, and Wouter Koolen. Structure adaptive algorithms for stochastic bandits. In *International Conference on Machine Learning*, pp. 2443–2452. PMLR, 2020.
- Jonas Degraeve, Federico Felici, Jonas Buchli, Michael Neunert, Brendan Tracey, Francesco Carpanese, Timo Ewalds, Roland Hafner, Abbas Abdolmaleki, Diego de Las Casas, et al. Magnetic control of tokamak plasmas through deep reinforcement learning. *Nature*, 602(7897):414–419, 2022.
- Kefan Dong, Jiaqi Yang, and Tengyu Ma. Provable model-based nonlinear bandit and reinforcement learning: Shelve optimism, embrace virtual curvature. *arXiv preprint arXiv:2102.04168*, 2021.
- Dylan Foster and Alexander Rakhlin. Beyond UCB: Optimal and efficient contextual bandits with regression oracles. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 3199–3210. PMLR, 13–18 Jul 2020.
- Dylan J Foster, Alexander Rakhlin, David Simchi-Levi, and Yunzong Xu. Instance-dependent complexity of contextual bandits and reinforcement learning: A disagreement-based perspective. *arXiv preprint arXiv:2010.03104*, 2020.
- Dylan J Foster, Sham M Kakade, Jian Qian, and Alexander Rakhlin. The statistical complexity of interactive decision making. *arXiv preprint arXiv:2112.13487*, 2021.
- Manuel Gil, Fady Alajaji, and Tamas Linder. Rényi divergence measures for commonly used univariate continuous distributions. *Information Sciences*, 249:124–131, 2013.
- Todd L Graves and Tze Leung Lai. Asymptotically efficient adaptive choice of control laws in controlled markov chains. *SIAM journal on control and optimization*, 35(3):715–743, 1997.
- Samarth Gupta, Shreyas Chaudhari, Gauri Joshi, and Osman Yağan. Multi-armed bandits with correlated arms. *IEEE Transactions on Information Theory*, 67(10):6711–6732, 2021.
- Botao Hao, Tor Lattimore, and Csaba Szepesvari. Adaptive exploration in linear contextual bandit. In *International Conference on Artificial Intelligence and Statistics*, pp. 3536–3545. PMLR, 2020.

- Chi Jin, Zeyuan Allen-Zhu, Sebastien Bubeck, and Michael I Jordan. Is q-learning provably efficient? In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pp. 4868–4878, 2018.
- 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.
- Kwang-Sung Jun and Chicheng Zhang. Crush optimism with pessimism: Structured bandits beyond asymptotic optimality. *Advances in Neural Information Processing Systems*, 33:6366–6376, 2020.
- Johannes Kirschner, Tor Lattimore, Claire Vernade, and Csaba Szepesvári. Asymptotically optimal information-directed sampling. In *Conference on Learning Theory*, pp. 2777–2821. PMLR, 2021.
- Junpei Komiyama, Junya Honda, and Hiroshi Nakagawa. Regret lower bound and optimal algorithm in finite stochastic partial monitoring. *Advances in Neural Information Processing Systems*, 28, 2015.
- Tze Leung Lai, Herbert Robbins, et al. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- Tor Lattimore and Csaba Szepesvari. The end of optimism? an asymptotic analysis of finite-armed linear bandits. In *Artificial Intelligence and Statistics*, pp. 728–737. PMLR, 2017.
- Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- Yingkai Li, Yining Wang, and Yuan Zhou. Nearly minimax-optimal regret for linearly parameterized bandits. In *Conference on Learning Theory*, pp. 2173–2174. PMLR, 2019.
- Stefan Magureanu, Richard Combes, and Alexandre Proutiere. Lipschitz bandits: Regret lower bound and optimal algorithms. In *Conference on Learning Theory*, pp. 975–999. PMLR, 2014.
- Cheng Mao. Hypothesis testing: Lecture notes for math 6263 at georgia tech. 2021.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.
- Jungseul Ok, Alexandre Proutiere, and Damianos Tranos. Exploration in structured reinforcement learning. In *Advances in Neural Information Processing Systems*, pp. 8874–8882, 2018.
- A Rajeev, Demosthenis Teneketzis, and Venkatachalam Anantharam. Asymptotically efficient adaptive allocation schemes for controlled iid processes: Finite parameter space. *IEEE Transactions on Automatic Control*, 34(3), 1989.
- David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.
- David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of Go without human knowledge. *Nature*, 550(7676):354, 2017.
- Max Simchowitz and Kevin G Jamieson. Non-asymptotic gap-dependent regret bounds for tabular MDPs. In *Advances in Neural Information Processing Systems*, pp. 1153–1162, 2019.
- Andrea Tirinzoni, Alessandro Lazaric, and Marcello Restelli. A novel confidence-based algorithm for structured bandits. In *International Conference on Artificial Intelligence and Statistics*, pp. 3175–3185. PMLR, 2020.

- Andrea Tirinzoni, Matteo Pirota, and Alessandro Lazaric. A fully problem-dependent regret lower bound for finite-horizon mdps. *arXiv preprint arXiv:2106.13013*, 2021.
- Andrea Tirinzoni, Aymen Al-Marjani, and Emilie Kaufmann. Near instance-optimal pac reinforcement learning for deterministic mdps. *arXiv preprint arXiv:2203.09251*, 2022.
- Sara Van de Geer. *Empirical Processes in M-estimation*. Cambridge University Press, 2000.
- Tim Van Erven and Peter Harremo. Rényi divergence and kullback-leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820, 2014.
- Bart PG Van Parys and Negin Golrezaei. Optimal learning for structured bandits. *arXiv preprint arXiv:2007.07302*, 2020.
- Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.
- Haike Xu, Tengyu Ma, and Simon S Du. Fine-grained gap-dependent bounds for tabular mdps via adaptive multi-step bootstrap. *arXiv preprint arXiv:2102.04692*, 2021.
- Kunhe Yang, Lin Yang, and Simon Du. Q-learning with logarithmic regret. In *International Conference on Artificial Intelligence and Statistics*, pp. 1576–1584. PMLR, 2021.
- Andrea Zanette and Emma Brunskill. Tighter problem-dependent regret bounds in reinforcement learning without domain knowledge using value function bounds. In *International Conference on Machine Learning*, pp. 7304–7312, 2019.

Appendix

Table of Contents

A Proofs for Finite Hypothesis Class	13
A.1 Proof of Lemma 4.1	13
A.2 Proof of Lemma 4.2	14
A.3 Proof of Lemma 4.3	17
A.4 Proof of Theorem 3.3	18
B Missing Proofs in Section 3.1	19
B.1 Proof of Lemma 3.1	19
B.2 Proof of Theorem 3.2	20
B.3 Instantiation of the Complexity Measure	21
B.4 A Sufficient Condition for Condition 1	23
B.5 Comparison with Existing Lower Bounds	23
C Extension to Infinite Hypothesis Class	24
C.1 Proof of Theorem 5.2	24
C.2 Proof of Lemma C.1	24
C.3 Proof of Lemma C.2	26
D Proof of Conditions for Tabular Reinforcement Learning	28
D.1 Proof of Theorem 5.1	29
D.2 Proof of Condition 1	30
D.3 Proof of Condition 2	32
D.4 Proof of Condition 3	33
E RL with General Reward Distribution	34
E.1 Lemmas for RL with General Reward Distribution	34
E.2 Proof of Condition 4 for Gaussian Distribution	39
F Uniform Concentration	39
F.1 Uniform Concentration with ℓ_∞ Covering	39
F.2 Uniform Concentration for RL with General Reward Distribution	41
G Helper Lemmas	43

A PROOFS FOR FINITE HYPOTHESIS CLASS

In this section we show the missing proofs in Section 4.

A.1 PROOF OF LEMMA 4.1

In this section, we prove Lemma 4.1.

Proof of Lemma 4.1. Recall that $m_{\text{init}} = \lceil |\Pi| \lceil \frac{\ln n}{\ln |\Pi|} \rceil \rceil$. Let $w = \{\hat{\pi}_i\}_{i=1}^{m_{\text{init}}}$ be the sequence of decisions in the initialization step of Alg. 1, and $\{\hat{o}_i\}_{i=1}^{m_{\text{init}}}$ the corresponding observations. Define

$$\rho = \min_{g \in \mathcal{F}, g \neq f^*} \max_{\pi \in \Pi} D_{1/2}(f^*[\pi] \| g[\pi]). \quad (14)$$

For finite hypothesis we have $\rho > 0$. Then Lemma A.2, for any $\epsilon > 0$ we get

$$\Pr \left(\frac{1}{m_{\text{init}}} \sum_{i=1}^{m_{\text{init}}} \ln \frac{f^*[\hat{\pi}_i](\hat{o}_i)}{g[\hat{\pi}_i](\hat{o}_i)} \geq \frac{1}{m_{\text{init}}} \sum_{i=1}^{m_{\text{init}}} D_{1/2}(f^*[\hat{\pi}_i] \| g[\hat{\pi}_i]) - \epsilon \right) \geq 1 - \exp(-m_{\text{init}}\epsilon/2). \quad (15)$$

Let $\epsilon = \frac{\rho}{2|\Pi|}$. By definition of ρ we get, for all $g \in \mathcal{F} \setminus \{f^*\}$

$$\Pr \left(\frac{1}{m_{\text{init}}} \sum_{i=1}^{m_{\text{init}}} \ln \frac{f^*[\hat{\pi}_i](\hat{o}_i)}{g[\hat{\pi}_i](\hat{o}_i)} \geq \frac{\rho}{2|\Pi|} \right) \quad (16)$$

$$\geq \Pr \left(\frac{1}{m_{\text{init}}} \sum_{i=1}^{m_{\text{init}}} \ln \frac{f^*[\hat{\pi}_i](\hat{o}_i)}{g[\hat{\pi}_i](\hat{o}_i)} \geq \frac{1}{m_{\text{init}}} \sum_{i=1}^{m_{\text{init}}} D_{1/2}(f^*[\hat{\pi}_i] \| g[\hat{\pi}_i]) - \frac{\rho}{2|\Pi|} \right) \quad (17)$$

$$\geq 1 - \exp \left(-\frac{\rho \ln n}{4 \ln \ln n} \right). \quad (18)$$

By union bound, with probability at least $1 - |\mathcal{F}| \exp \left(-\frac{\rho \ln n}{4 \ln \ln n} \right)$ we have

$$\forall g \in \mathcal{F} \setminus \{f^*\}, \quad \sum_{i=1}^{m_{\text{init}}} \ln \frac{f^*[\hat{\pi}_i](\hat{o}_i)}{g[\hat{\pi}_i](\hat{o}_i)} > 0 = \sum_{i=1}^{m_{\text{init}}} \ln \frac{f^*[\hat{\pi}_i](\hat{o}_i)}{f^*[\hat{\pi}_i](\hat{o}_i)}, \quad (19)$$

which implies that

$$\forall g \in \mathcal{F} \setminus \{f^*\}, \quad \sum_{i=1}^{m_{\text{init}}} \ln g[\hat{\pi}_i](\hat{o}_i) < \sum_{i=1}^{m_{\text{init}}} \ln f^*[\hat{\pi}_i](\hat{o}_i). \quad (20)$$

Recalling that $\hat{f} = \operatorname{argmax}_{f \in \mathcal{F}} \sum_{i=1}^{m_{\text{init}}} \ln f[\hat{\pi}_i](\hat{o}_i)$, Eq. (20) implies $\hat{f} = f^*$.

By algebraic manipulation, for large enough n we have

$$\exp \left(-\frac{\rho \ln n}{4 \ln \ln n} \right) \leq \exp(-\ln |\mathcal{F}| - \ln \ln n) = \frac{1}{|\mathcal{F}| \ln n}. \quad (21)$$

As a result, for large enough n we get $\Pr(\hat{f} = f^*) \geq 1 - 1/\ln n$. In addition, the regret of Step 1 is upper bounded by $\mathcal{O}(m_{\text{init}}) = \mathcal{O}(\frac{\ln n}{\ln \ln n})$. \square

A.2 PROOF OF LEMMA 4.2

In this section, we prove Lemma 4.2. To this end, we need the following lemma.

Lemma A.1. *Consider an instance $f \in \mathcal{F}$ and $n > 0$. Let $\delta = (\ln \ln n)^{-1/4}$ and $\epsilon = (\ln \ln n)^{-1}$. Define*

$$\lambda = \lambda_0 \left(4(\ln \ln n)^{3/4}, \frac{1}{\ln \ln n}, f \right) \quad (22)$$

as the value that Condition 1 holds with corresponding parameters.

Consider any $\hat{w} \in \mathbb{R}_+^{|\Pi|}$ such that $\|\hat{w}\|_\infty \leq (\ln \ln n)^{1/4}$. Let $w = \{\pi_i\}_{i=1}^m$ be a list of decisions where a decision π occurs $\lceil ((1 + \delta)\hat{w}_\pi + \delta) \ln n \rceil$ times for every $\pi \in \Pi$, and $m = \sum_{\pi} \lceil ((1 + \delta)\hat{w}_\pi + \delta) \ln n \rceil$.

Define the set $\mathcal{F}(\hat{w}, f) = \{g \in \mathcal{F} : \sum_{\pi \in \Pi} \hat{w}_\pi D_{\text{KL}}(f[\pi] \| g[\pi]) \geq 1\}$. For any constant $c > 0$, there exists $n_0 > 0$ such that for all $n > n_0$,

$$D_{1-\lambda}^w(f \| g) \geq \frac{\ln n}{m} + c\epsilon, \quad \forall g \in \mathcal{F}(\hat{w}, f). \quad (23)$$

In addition, we have $\lambda^{-1} = O(\text{poly}(\ln \ln n))$ and $m \geq \frac{|\Pi| \ln n}{(\ln \ln n)^{1/4}}$.

Proof. To prove this lemma, we invoke Condition 1 with proper parameters.

First we bound the value of m . By the assumption of this lemma, we have $\|\hat{w}\|_\infty \leq (\ln \ln n)^{1/4}$. Consequently, for large enough n we get

$$m = \sum_{\pi} \lceil ((1 + \delta)\hat{w}_\pi + \delta) \ln n \rceil \leq 2|\Pi| \ln n (\ln \ln n)^{1/4}. \quad (24)$$

On the other hand,

$$m = \sum_{\pi} \lceil ((1 + \delta)\hat{w}_\pi + \delta) \ln n \rceil \geq |\Pi| \delta \ln n = \frac{|\Pi| \ln n}{(\ln \ln n)^{1/4}}. \quad (25)$$

Define $\alpha = \frac{\ln n}{m} + c\epsilon$. Then for large enough n we have $\alpha + \epsilon \leq \frac{2}{|\Pi|} (\ln \ln n)^{1/4}$. We can also lower bound $\gamma = \frac{1}{m} \min_{\pi} \lceil ((1 + \delta)\hat{w}_\pi + \delta) \ln n \rceil$. By the upper bound of m and the fact that $\lceil ((1 + \delta)\hat{w}_\pi + \delta) \ln n \rceil \geq \delta \ln n = \frac{\ln n}{(\ln \ln n)^{1/4}}$, we get $\gamma \geq \frac{1}{2|\Pi|(\ln \ln n)^{1/2}}$.

We invoke Condition 1 with parameters $((\alpha + \epsilon)/\gamma, \epsilon, f)$. Note that λ defined in Eq. (22) satisfies $\lambda \leq \lambda_0((\alpha + \epsilon)/\gamma, \epsilon, f)$ because $(\alpha + \epsilon)/\gamma \leq 4(\ln \ln n)^{3/4}$. Therefore we get

$$D_{1-\lambda}(f[\pi] \| g[\pi]) \geq \min\{D_{\text{KL}}(f[\pi] \| g[\pi]) - \epsilon, (\alpha + \epsilon)/\gamma\}, \quad \forall g \in \mathcal{F}, \pi \in \Pi. \quad (26)$$

In the following, we prove that $D_{1-\lambda}^w(f \| g) \geq \alpha$ for all $g \in \mathcal{F}(\hat{w}, f)$. First of all, for large enough n , for all $g \in \mathcal{F}(\hat{w}, f)$ we get

$$\begin{aligned} D_{\text{KL}}^w(f \| g) &= \frac{1}{m} \sum_{i=1}^m D_{\text{KL}}(f[\pi_i] \| g[\pi_i]) = \frac{1}{m} \sum_{\pi \in \Pi} \lceil ((1 + \delta)\hat{w}_\pi + \delta) \ln n \rceil D_{\text{KL}}(f[\pi] \| g[\pi]) \quad (27) \\ &\geq \frac{1}{m} \sum_{\pi \in \Pi} (1 + \delta)(\ln n)\hat{w}_\pi D_{\text{KL}}(f[\pi] \| g[\pi]) \geq \frac{1}{m} (1 + \delta) \ln n \geq \frac{1}{m} (\ln n + (c + 1)m\epsilon), \quad (28) \end{aligned}$$

where the last inequality comes from the fact that $\delta \ln n = \frac{\ln n}{(\ln \ln n)^{1/4}} \gtrsim 2(c + 1)|\Pi| \frac{\ln n}{(\ln \ln n)^{3/4}} \geq (c + 1)m\epsilon$.

Now for a fixed $g \in \mathcal{F}(\hat{w}, f)$, consider the following two cases.

Case 1: $\exists \bar{\pi} \in \Pi, D_{\text{KL}}(f[\bar{\pi}] \| g[\bar{\pi}]) \geq (\alpha + \epsilon)/\gamma$. In this case we have

$$D_{1-\lambda}^w(f \| g) = \frac{1}{m} \sum_{i=1}^m D_{1-\lambda}(f[\pi_i] \| g[\pi_i]) \quad (29)$$

$$\geq \frac{\lceil ((1 + \delta)\hat{w}_{\bar{\pi}} + \delta) \ln n \rceil}{m} D_{1-\lambda}(f[\bar{\pi}] \| g[\bar{\pi}]) \quad (30)$$

$$\geq \gamma D_{1-\lambda}(f[\bar{\pi}] \| g[\bar{\pi}]) \geq \alpha, \quad (31)$$

where the last inequality comes from Eq. (26).

Case 2: $\forall \pi \in \Pi, D_{\text{KL}}(f[\pi] \| g[\pi]) \leq (\alpha + \epsilon)/\gamma$. In this case, Eq. (26) implies that

$$D_{1-\lambda}(f[\pi] \| g[\pi]) \geq D_{\text{KL}}(f[\pi] \| g[\pi]) - \epsilon, \quad \forall \pi \in \Pi. \quad (32)$$

Consequently,

$$D_{1-\lambda}^w(f \| g) = \frac{1}{m} \sum_{i=1}^m D_{1-\lambda}(f[\pi_i] \| g[\pi_i]) \geq \frac{1}{m} \sum_{i=1}^m (D_{\text{KL}}(f[\pi_i] \| g[\pi_i]) - \epsilon) = D_{\text{KL}}^w(f \| g) - \epsilon.$$

By Eq. (28) we get $D_{\text{KL}}^w(f \| g) \geq \alpha + \epsilon$. Therefore $D_{1-\lambda}^w(f \| g) \geq \alpha$.

Combining the two cases together we get the desired result. The lower bound of λ^{-1} follows directly from Condition 1. \square

Now we are ready to prove Lemma 4.2.

Proof of Lemma 4.2. We prove the four items in Lemma 4.2 separately. We will invoke Lemma A.1 and Lemma 4.3 in the proof. Following Lemma A.1, let $\delta = (\ln \ln n)^{1/4}$, $\epsilon = (\ln \ln n)^{-1}$ and

$$\lambda = \lambda_0 \left(4(\ln \ln n)^{3/4}, \frac{1}{\ln \ln n}, f \right). \quad (33)$$

Recall that $\Lambda(f) = \{g \in \mathcal{F} : \pi^*(g) \neq \pi^*(f)\}$.

Proof of item (a). In this case we have $f^* \in \Lambda(\hat{f})$. By Lemma 4.3 we get

$$\Pr_{f^*} \left(\mathcal{E}_{\text{acc}}^{\hat{f}} \right) = \Pr_{f^*} \left(\forall g \in \Lambda(\hat{f}), \sum_{i=1}^m \ln \frac{\hat{f}[\pi_i](o_i)}{g[\pi_i](o_i)} \geq \ln n \right) \quad (34)$$

$$\leq \Pr_{f^*} \left(\sum_{i=1}^m \ln \frac{\hat{f}[\pi_i](o_i)}{f^*[\pi_i](o_i)} \geq \ln n \right) \leq \exp(-\ln n) = 1/n. \quad (35)$$

Proof of item (b). Recall that in this case we have $\hat{f} = f^*$. Since \hat{w} is the solution to $\mathcal{C}(\hat{f}, (\ln \ln n)^{1/4})$ (Line 4 of Alg. 1), we have $\sum_{\pi \in \Pi} \hat{w}_\pi D_{\text{KL}}(\hat{f}[\pi] \| g[\pi]) \geq 1$ for all $g \in \Lambda(\hat{f})$. Recall that w is the list of decisions computed by Line 5 of Alg. 1. By Lemma A.1 we get

$$D_{1-\lambda}^w(f \| g) \geq \frac{\ln n}{m} + \epsilon, \quad \forall g \in \Lambda(\hat{f}). \quad (36)$$

Let $\beta = \frac{\ln n}{m} + \epsilon$. By Lemma 4.3, for every $g \in \Lambda(\hat{f})$ we have

$$\Pr_{f^*} \left(\sum_{i=1}^m \ln \frac{f^*[\pi_i](o_i)}{g[\pi_i](o_i)} \geq \ln n \right) \geq 1 - \exp(-m\lambda\epsilon). \quad (37)$$

Lemma A.1 also yields $\lambda^{-1} \leq \text{poly}(\ln \ln n)$ and $m \geq \frac{|\Pi| \ln n}{(\ln \ln n)^{1/4}}$. Therefore $m\lambda\epsilon \geq \frac{\ln n}{\text{poly}(\ln \ln n)}$. Consequently, for large enough n we have

$$\exp(-m\lambda\epsilon) \leq \frac{\ln n}{|\mathcal{F}|}. \quad (38)$$

Applying union bound, under the event $\mathbb{I}[\hat{f} = f^*]$ we get

$$\Pr_{f^*} \left(\mathcal{E}_{\text{acc}}^{\hat{f}} \right) = \Pr_{f^*} \left(\forall g \in \Lambda(\hat{f}), \sum_{i=1}^m \ln \frac{\hat{f}[\pi_i](o_i)}{g[\pi_i](o_i)} \geq \ln n \right) \geq 1 - \ln n. \quad (39)$$

Proof of item (c). Recall that $\Delta(f, \pi)$ is the sub-optimality gap of decision π on instance f , and $\Delta_{\max}(f) = \max_{\pi} \Delta(f, \pi)$ is the maximum decision gap of instance f . Since \hat{w} is the solution to $\mathcal{C}(\hat{f}, (\ln \ln n)^{1/4})$, we have $\|\hat{w}\|_\infty \leq (\ln \ln n)^{1/4}$. As a result, the regret of Step 2 is upper bounded by

$$\sum_{\pi} [((1 + \delta)\hat{w}_\pi + \delta) \ln n] \Delta_{\max}(f^*) \lesssim |\Pi| \ln n (\ln \ln n)^{1/4}, \quad (40)$$

which proves the second part of (c). For the first part, when $\hat{f} = f^*$ we have

$$\sum_{\pi} [((1 + \delta)\hat{w}_\pi + \delta) \ln n] \Delta(f^*, \pi) \quad (41)$$

$$\leq \sum_{\pi} (1 + \delta)\hat{w}_\pi (\ln n) \Delta(f^*, \pi) + |\Pi| \Delta_{\max}(f^*) (1 + \delta \ln n) \quad (42)$$

$$= ((1 + \delta)\mathcal{C}(f^*, (\ln \ln n)^{1/4}) + o(1)) \ln n. \quad (43)$$

By Lemma B.1, $\mathcal{C}(f^*, (\ln \ln n)^{1/4}) = \mathcal{O}(1)$. As a result,

$$((1 + \delta)\mathcal{C}(f^*, (\ln \ln n)^{1/4}) + o(1)) \ln n \leq (\mathcal{C}(f^*, (\ln \ln n)^{1/4}) + o(1)) \ln n. \quad (44)$$

□

A.3 PROOF OF LEMMA 4.3

In this section, we prove Lemma 4.3. To this end, we need the following concentration inequality.

Lemma A.2. *Consider two sequences of distributions $P = \{P_i\}_{i=1}^m$ and $Q = \{Q_i\}_{i=1}^m$, and a sequence of random variables $o = \{o_i\}_{i=1}^m$ independently drawn from P . For any $\lambda \in (0, 1)$ and $\epsilon > 0$ we have*

$$\Pr_{o \sim P} \left(\frac{1}{m} \sum_{i=1}^m \ln \frac{P_i(o_i)}{Q_i(o_i)} \geq \frac{1}{m} \sum_{i=1}^m D_{1-\lambda}(P_i \| Q_i) - \epsilon \right) \geq 1 - \exp(-m\lambda\epsilon). \quad (45)$$

Proof of Lemma A.2. We prove Lemma A.2 by moment method. Consider any $\lambda \in (0, 1)$, $\epsilon > 0$ and let $\beta = \frac{1}{m} \sum_{i=1}^m D_{1-\lambda}(P_i \| Q_i)$. Then we have

$$\Pr_{o \sim P} \left(\sum_{i=1}^m \ln \frac{P_i(o_i)}{Q_i(o_i)} \leq m(\beta - \epsilon) \right) \quad (46)$$

$$= \Pr_{o \sim P} \left(\sum_{i=1}^m \ln \frac{Q_i(o_i)}{P_i(o_i)} \geq -m(\beta - \epsilon) \right) \quad (47)$$

$$= \Pr_{o \sim P} \left(\exp \left(\lambda \sum_{i=1}^m \ln \frac{Q_i(o_i)}{P_i(o_i)} \right) \geq \exp(-\lambda m(\beta - \epsilon)) \right) \quad (48)$$

$$\leq \exp(\lambda m(\beta - \epsilon)) \mathbb{E}_{o \sim P} \left[\exp \left(\lambda \sum_{i=1}^m \ln \frac{Q_i(o_i)}{P_i(o_i)} \right) \right] \quad (49)$$

$$\leq \exp(\lambda m(\beta - \epsilon)) \prod_{i=1}^m \mathbb{E}_{o_i \sim P_i} [(Q_i(o_i)^\lambda P_i(o_i)^{-\lambda})] \quad (50)$$

$$\leq \exp(\lambda m(\beta - \epsilon)) \prod_{i=1}^m \left(\int_o (Q_i(o)^\lambda P_i(o)^{1-\lambda}) do \right) \quad (51)$$

$$\leq \exp(\lambda m(\beta - \epsilon)) \prod_{i=1}^m \exp((\lambda - 1)D_\lambda(Q_i \| P_i)). \quad (52)$$

where Eq. (49) follows from Markov inequality. By Van Erven & Harremoës (2014, Proposition 2) we get $D_\lambda(Q \| P) = \frac{\lambda}{1-\lambda} D_{1-\lambda}(P \| Q)$ for any distributions P, Q . As a result,

$$\exp(\lambda m(\beta - \epsilon)) \prod_{i=1}^m \exp((\lambda - 1)D_\lambda(Q_i \| P_i)) \quad (53)$$

$$= \exp(\lambda m(\beta - \epsilon)) \prod_{i=1}^m \exp(-\lambda D_{1-\lambda}(P_i \| Q_i)) \quad (54)$$

$$= \exp(\lambda m(\beta - \epsilon)) \exp(-\lambda m\beta) \quad (55)$$

$$= \exp(-m\lambda\epsilon). \quad (56)$$

□

Now we are ready to prove Lemma 4.3.

Proof of Lemma 4.3. First we prove Eq. (9). By the moment method, for any $c > 0$

$$\Pr_Q(\mathcal{E}_{\text{acc}}) = \Pr_Q \left(\sum_{i=1}^m \ln \frac{P_i(o_i)}{Q_i(o_i)} \geq c \right) \quad (57)$$

$$= \Pr_Q \left(\exp \left(\sum_{i=1}^m \ln \frac{P_i(o_i)}{Q_i(o_i)} \right) \geq \exp(c) \right) \quad (58)$$

$$\begin{aligned}
&\leq \exp(-c) \mathbb{E}_Q \left[\prod_{i=1}^m \frac{P_i(o_i)}{Q_i(o_i)} \right] && \text{(Markov inequality)} \\
&\leq \exp(-c) \prod_{i=1}^m \mathbb{E}_{o_i \sim Q_i} \left[\frac{P_i(o_i)}{Q_i(o_i)} \right] && \text{(independence of } o_i \text{'s)} \\
&= \exp(-c), && (59)
\end{aligned}$$

where the last line comes from the fact that both P_i, Q_i are valid distributions. For Eq. (10), we can directly invoke Lemma A.2 with $\epsilon = \beta - c/m$. \square

A.4 PROOF OF THEOREM 3.3

We prove Theorem 3.3 by stitching Lemma 4.1 and Lemma 4.2 together.

Proof of Theorem 3.3. We upper bound the regret of Alg. 1 by discussing the regret under following events separately. Let $\text{Reg}_{\text{Step1}}, \text{Reg}_{\text{Step2}},$ and $\text{Reg}_{\text{Step3}}$ be the regret incurred in Step 1, 2, and 3 respectively. Let $\mathcal{E}_{\text{init}} = \mathbb{I}[\hat{f} = f^*]$ be the event that the initial estimation is accurate.

Regret of Step 1: By Lemma 4.1 we have,

$$\limsup_{n \rightarrow \infty} \frac{1}{\ln n} \mathbb{E}[\text{Reg}_{\text{Step1}}] \leq 0.$$

Regret of Step 2. By item (c) of Lemma 4.2 and Lemma 4.1,

$$\begin{aligned}
\mathbb{E}[\text{Reg}_{\text{Step2}}] &= \mathbb{E}[\text{Reg}_{\text{Step2}} \mid \mathcal{E}_{\text{init}}] \Pr(\mathcal{E}_{\text{init}}) + \mathbb{E}[\text{Reg}_{\text{Step2}} \mid \mathcal{E}_{\text{init}}^c] \Pr(\mathcal{E}_{\text{init}}^c) && (60) \\
&\leq (\mathcal{C}(f^*, (\ln \ln n)^{1/4}) + o(1)) \ln(n) + \mathcal{O}(\ln n \ln \ln n) / \ln n. && (61)
\end{aligned}$$

As a result,

$$\limsup_{n \rightarrow \infty} \frac{1}{\ln n} \mathbb{E}[\text{Reg}_{\text{Step2}}] \leq \limsup_{n \rightarrow \infty} \mathcal{C}(f^*, (\ln \ln n)^{1/4}) = \mathcal{C}(f^*). \quad (62)$$

Regret of Step 3. First focus on the event $\mathcal{E}_{\text{init}}$. Under event $\mathcal{E}_{\text{acc}}^{\hat{f}} \cap \mathcal{E}_{\text{init}}$, there is no regret in Step 3. On the other hand, by item (b) of Lemma 4.2 we have $\Pr(\mathcal{E}_{\text{init}} \cap (\mathcal{E}_{\text{acc}}^{\hat{f}})^c) \leq 1/\ln n$. Since UCB gives logarithmic regret, we have

$$\limsup_{n \rightarrow \infty} \frac{1}{\ln n} \mathbb{E}[\mathbb{I}[\mathcal{E}_{\text{init}}, (\mathcal{E}_{\text{acc}}^{\hat{f}})^c] \text{Reg}_{\text{Step3}}] \leq \limsup_{n \rightarrow \infty} \frac{\mathcal{O}(\ln n)}{\ln n} \Pr(\mathcal{E}_{\text{init}} \cap (\mathcal{E}_{\text{acc}}^{\hat{f}})^c) \leq 0. \quad (63)$$

As a result,

$$\limsup_{n \rightarrow \infty} \frac{1}{\ln n} \mathbb{E}[\mathbb{I}[\mathcal{E}_{\text{init}}] \text{Reg}_{\text{Step3}}] \leq 0. \quad (64)$$

Now we focus on the event $\mathcal{E}_{\text{init}}^c$. Let $\mathcal{E}_1 = \{\pi^*(\hat{f}) = \pi^*(f^*)\}$. Under event $\mathcal{E}_{\text{acc}}^{\hat{f}} \cap \mathcal{E}_1$ the algorithm has no regret in Step 3:

$$\mathbb{E}[\mathbb{I}[\mathcal{E}_{\text{init}}^c, \mathcal{E}_1, \mathcal{E}_{\text{acc}}^{\hat{f}}] \text{Reg}_{\text{Step3}}] = 0. \quad (65)$$

On the other hand, consider the event $\mathcal{E}_{\text{acc}}^{\hat{f}} \cap \mathcal{E}_1^c$. By item (a) of Lemma 4.2 we have

$$\mathbb{E}[\mathbb{I}[\mathcal{E}_{\text{init}}^c, \mathcal{E}_1^c, \mathcal{E}_{\text{acc}}^{\hat{f}}] \text{Reg}_{\text{Step3}}] \leq n \Delta_{\max} \Pr(\mathbb{I}[\pi^*(f^*) \neq \pi^*(\hat{f}), \mathcal{E}_{\text{acc}}^{\hat{f}}]) \leq \frac{n \Delta_{\max}}{n} \leq \mathcal{O}(1). \quad (66)$$

Under the event $(\mathcal{E}_{\text{acc}}^{\hat{f}})^c$, Step 3 incurs logarithmic regret. As a result,

$$\mathbb{E}[\mathbb{I}[\mathcal{E}_{\text{init}}^c, (\mathcal{E}_{\text{acc}}^{\hat{f}})^c] \text{Reg}_{\text{Step3}}] \leq \mathcal{O}(\ln n) \Pr(\mathcal{E}_{\text{init}}^c) \leq \mathcal{O}(1). \quad (67)$$

Combining Eqs. (65), (66) and (67) we get

$$\limsup_{n \rightarrow \infty} \frac{1}{\ln n} \mathbb{E} [\mathbb{I} [\mathcal{E}_{\text{init}}^c] \text{Reg}_{\text{Step3}}] = 0. \quad (68)$$

Therefore, combining Eqs. (64) and (68),

$$\limsup_{n \rightarrow \infty} \frac{1}{\ln n} \mathbb{E} [\text{Reg}_{\text{Step3}}] = 0. \quad (69)$$

Stitching the regrets from steps 1-3 together we prove Theorem 3.3. \square

B MISSING PROOFS IN SECTION 3.1

In this section, we present the missing proofs in section 3.1.

B.1 PROOF OF LEMMA 3.1

To prove Lemma 3.1, we need the following lemma.

Lemma B.1. *For any $f \in \mathcal{F}$, let $n_0 = 5(\Delta_{\min}(f)/R_{\max})^{-2}$. Consider $w \in \mathbb{R}_+^{|\Pi|}$ such that $w_\pi = n_0$ for all $\pi \in \Pi$. Then w is a valid solution to $\mathcal{C}(f, n)$ for all $n > n_0$. As a corollary, $\mathcal{C}(f, n) \leq \Delta_{\max}(f)|\Pi|n_0$ for all $n > n_0$.*

Proof of Lemma B.1. Recall that $\Delta(f, \pi) = \max_{\pi' \in \Pi} R_f(\pi') - R_f(\pi)$ is the reward gap of decision π under instance f , and $\Delta_{\min}(f) = \min_{\pi: \Delta(f, \pi) > 0} \Delta(f, \pi)$ is the minimum decision gap of the instance f .

Let $\Lambda(f) = \{g \in \mathcal{F}, \pi^*(g) \neq \pi^*(f)\}$. For any instance $g \in \Lambda(f)$, consider the following two decisions: $\pi_1 = \pi^*(f)$, $\pi_2 = \pi^*(g)$. We claim that

$$\max_{\pi \in \{\pi_1, \pi_2\}} |R_f(\pi) - R_g(\pi)| \geq \frac{\Delta_{\min}(f)}{3}. \quad (70)$$

This claim can be proved by contradiction. Suppose on the contrary that

$$\max_{\pi \in \{\pi_1, \pi_2\}} |R_f(\pi) - R_g(\pi)| < \Delta_{\min}(f)/3.$$

Then we have

$$R_g(\pi_1) \geq R_f(\pi_1) - \frac{\Delta_{\min}(f)}{3} \geq R_f(\pi_2) + \frac{2\Delta_{\min}(f)}{3} \geq R_g(\pi_2) + \frac{\Delta_{\min}(f)}{3} > R_g(\pi_2),$$

which contradicts with the fact that $\pi_2 = \pi^*(g)$.

Recall that $R_f(\pi) = \mathbb{E}_{o \sim f[\pi]} [R(o)]$. As a result, $|R_f(\pi) - R_g(\pi)| \leq R_{\max} D_{\text{TV}}(f[\pi] \| g[\pi])$. Now, by Pinsker's inequality, for any $\pi \in \Pi$ we have

$$D_{\text{KL}}(f[\pi] \| g[\pi]) \geq 2D_{\text{TV}}(f[\pi] \| g[\pi])^2 \geq 2(|R_f(\pi) - R_g(\pi)|/R_{\max})^2. \quad (71)$$

Combining with Eq. (70), for any $g \in \Lambda(f)$ we get

$$\sum_{\pi \in \Pi} n_0 D_{\text{KL}}(f[\pi] \| g[\pi]) \geq 1. \quad (72)$$

Since this inequality holds for every $g \in \Lambda(f)$, we prove the desired result. \square

Now we present the proof of Lemma 3.1.

Proof of Lemma 3.1. By Lemma B.1, for $n > 5(\Delta_{\min}(f)/R_{\max})^{-2}$, we get $\mathcal{C}(f, n) < \infty$. \square

B.2 PROOF OF THEOREM 3.2

Now we show the proof of Theorem 3.2.

Proof of Theorem 3.2. In the following we fix a consistent algorithm \mathcal{A} . Consider any instance $g \in \mathcal{F}$ such that $\pi^*(g) \neq \pi^*(f)$.

Recall that $\Delta(f, \pi) = \max_{\pi' \in \Pi} R_f(\pi') - R_f(\pi)$ is the reward gap of decision π under instance f and $\Delta_{\min}(f) = \min_{\pi: \Delta(f, \pi) > 0} \Delta(f, \pi)$ is the minimum decision gap of the instance f . Let $\epsilon = \min\{\Delta_{\min}(g), \Delta_{\min}(f)\}/2$. Consider two distributions (over observations and decisions) $P_{f,n}, P_{g,n}$ induced by running n steps of algorithm \mathcal{A} on instance f, g respectively. In addition, let N_π be the random variable indicates the number of times decision π is executed, and π_i the random variable indicating the decision executed at step i .

Let $\text{Reg}_{f,n}$ and $\text{Reg}_{g,n}$ be the regret of running algorithm \mathcal{A} on instances f and g respectively. By definition we have

$$\frac{\text{Reg}_{f,n} + \text{Reg}_{g,n}}{\epsilon n} \geq \Pr_f \left(N_{\pi^*(f)} \leq \frac{n}{2} \right) + \Pr_g \left(N_{\pi^*(f)} > \frac{n}{2} \right) \quad (73)$$

By basic inequality of KL divergence (Lattimore & Szepesvari, 2017, Lemma 5) we get,

$$\Pr_f \left(N_{\pi^*(f)} \leq \frac{n}{2} \right) + \Pr_g \left(N_{\pi^*(f)} > \frac{n}{2} \right) \geq \frac{1}{2} \exp(-D_{\text{KL}}(P_{f,n} \| P_{g,n})). \quad (74)$$

Combining Eq. (73) and Eq. (74) we get,

$$D_{\text{KL}}(P_{f,n} \| P_{g,n}) \geq \ln \left(\frac{\epsilon n}{2 (\text{Reg}_{f,n} + \text{Reg}_{g,n})} \right). \quad (75)$$

Now applying the chain rule for KL divergence (see, e.g., (Cover, 1999, Theorem 2.5.3)), we get

$$D_{\text{KL}}(P_{f,n} \| P_{g,n}) = \mathbb{E}_f \left[\sum_{i=1}^n D_{\text{KL}}(f[\pi_i] \| g[\pi_i]) \right] = \sum_{\pi \in \Pi} \mathbb{E}_f[N_\pi] D_{\text{KL}}(f[\pi] \| g[\pi]). \quad (76)$$

Therefore we have

$$\sum_{\pi \in \Pi} \frac{\mathbb{E}_f[N_\pi]}{\ln(\epsilon n) - \ln(2 (\text{Reg}_{f,n} + \text{Reg}_{g,n}))} D_{\text{KL}}(f[\pi] \| g[\pi]) \geq 1. \quad (77)$$

Consider the mixture of decisions

$$w_\pi = \frac{\mathbb{E}_f[N_\pi]}{\ln(\epsilon n) - \ln(2 (\text{Reg}_{f,n} + \text{Reg}_{g,n}))}.$$

Since $\text{Reg}_{f,n} + \text{Reg}_{g,n} \leq n\Delta_{\max}$ and $\mathbb{E}_f[N_\pi] \leq n$, w_π satisfies the constraint of $\mathcal{C}(f, n)$ for large enough n . In addition, the expected regret of algorithm \mathcal{A} is

$$\text{Reg}_{f,n} = \mathbb{E}_f \left[\sum_{i=1}^n \Delta_{\pi_i} \right] = \sum_{\pi} \mathbb{E}_f[N_\pi] \Delta_\pi \quad (78)$$

$$\geq \sum_{\pi} \Delta_\pi w_\pi \ln \left(\frac{\epsilon n}{2 (\text{Reg}_{f,n} + \text{Reg}_{g,n})} \right) \quad (79)$$

$$\geq \mathcal{C}(f, n) \ln \left(\frac{\epsilon n}{2 (\text{Reg}_{f,n} + \text{Reg}_{g,n})} \right). \quad (80)$$

Since \mathcal{A} is consistent, for any $p > 0$ we have $\text{Reg}_{f,n} + \text{Reg}_{g,n} = O(n^p)$. Consequently, for any $p > 0$,

$$\limsup_{n \rightarrow \infty} \frac{\text{Reg}_{f,n}}{\ln(n)} \quad (81)$$

$$\geq \limsup_{n \rightarrow \infty} \mathcal{C}(f, n) \frac{\ln(n) - \ln(2(\text{Reg}_{f,n} + \text{Reg}_{g,n})/\epsilon)}{\ln(n)} \quad (82)$$

$$\geq \mathcal{C}(f)(1-p). \quad (83)$$

Since $p > 0$ is an arbitrary constant, we get

$$\limsup_{n \rightarrow \infty} \frac{\text{Reg}_{f,n}}{\ln(n)} \geq \mathcal{C}(f). \quad (84)$$

□

B.3 INSTANTIATION OF THE COMPLEXITY MEASURE

In this section, we instantiate our complexity measure $\mathcal{C}(f)$ on concrete settings.

First we consider multi-armed bandits with unit Gaussian reward. The family of decisions is $\Pi = \{1, 2, \dots, A\}$. Let $(\mu_1, \mu_2, \dots, \mu_A)$ be the mean rewards of each decision. An instance f in this case is characterized by the mean rewards, and $f[i] = \mathcal{N}(\mu_i, 1)$.

Proposition B.2. *For an multi-armed bandit instance f with unique optimal decision and unit Gaussian noise, let $(\mu_1, \mu_2, \dots, \mu_A)$ be the mean reward of each action. Then*

$$\mathcal{C}(f) \leq \sum_{i \in [1, A] \text{ and } \Delta_i > 0} \frac{2}{\Delta_i}$$

where $\Delta_i = \max_{i'} \mu_{i'} - \mu_i$.

Proof. We prove this proposition by constructing a solution $w \in \mathbb{R}_+^A$ to the optimization problem $\mathcal{C}(f, n)$ for large enough n .

Without loss of generality, we assume $\mu_1 > \mu_i$ for all $i \geq 2$. Consider the action frequency

$$w_i = \begin{cases} \frac{2n}{n\Delta_i^2 - 3}, & i \geq 2, \\ n, & i = 1. \end{cases} \quad (85)$$

Then when n is large enough we have $\|w\|_\infty \leq n$. On the other hand, consider any $g \in \mathcal{F}$ such that $\pi^*(g) \neq \pi^*(f)$. Suppose $\pi^*(g) = i$. In the following we show that

$$\sum_{i=1}^n w_i D_{\text{KL}}(f[i] \| g[i]) \geq 1. \quad (86)$$

Let $(\mu'_1, \mu'_2, \dots, \mu'_A)$ be the mean reward of instance g . For multi-armed bandits with unit Gaussian noise, we have

$$\sum_{i=1}^n w_i D_{\text{KL}}(f[i] \| g[i]) = \frac{1}{2} \sum_{i=1}^n w_i (\mu_i - \mu'_i)^2 \geq \frac{1}{2} (w_1 (\mu_1 - \mu'_1)^2 + w_i (\mu_i - \mu'_i)^2). \quad (87)$$

By the condition that $\pi^*(g) = i$, we get $\mu'_i \geq \mu'_1$. Combining with the fact that $\mu_1 > \mu_i$ we get

$$\min_{\mu'_1, \mu'_i; \mu'_1 \leq \mu'_i} (w_1 (\mu_1 - \mu'_1)^2 + w_i (\mu_i - \mu'_i)^2) \quad (88)$$

$$= \min_{\mu'_1 \in [\mu_i, \mu_1]} (w_1 (\mu_1 - \mu'_1)^2 + w_i (\mu_i - \mu'_1)^2) \quad (89)$$

$$= \frac{w_1 w_i}{w_1 + w_i} (\mu_1 - \mu_i)^2. \quad (90)$$

By the definition of w_i , we have

$$\frac{w_1 w_i}{w_1 + w_i} (\mu_1 - \mu_i)^2 = \frac{\frac{2n^2}{n\Delta_i^2 - 3}}{n + \frac{2n}{n\Delta_i^2 - 3}} \Delta_i^2 = \frac{2\Delta_i^2 n}{n\Delta_i^2 - 1} \geq 2. \quad (91)$$

As a result,

$$\mathcal{C}(f, n) \leq \sum_{i=2}^A \frac{2n}{n\Delta_i^2 - 3} \Delta_i. \quad (92)$$

It follows that

$$\mathcal{C}(f) = \lim_{n \rightarrow \infty} \mathcal{C}(f, n) \leq \lim_{n \rightarrow \infty} \sum_{i=2}^A \frac{2n}{n\Delta_i^2 - 3} \Delta_i = \sum_{i=2}^A \frac{2}{\Delta_i}. \quad (93)$$

□

In the following, we focus on a linear bandit instance f where the action set is $\mathcal{A} \subset \mathbb{R}^d$. The mean of an action $x \in \mathcal{A}$ is given by $\mu_x = \langle x, \theta \rangle$ for some $\theta \in \mathbb{R}^d$. Let $x^* = \operatorname{argmax}_x \mu_x$ be the optimal action, and $\Delta_x \triangleq \mu_{x^*} - \mu_x$ be the sub-optimality gap of action x . Define $\mathcal{A}^- = \mathcal{A} \setminus \{x^*\}$. A linear bandit instance is characterized by the vector θ , and $f[x] = \mathcal{N}(\langle x, \theta \rangle, 1)$.

We assume that \mathcal{A} is discrete, full rank, and $\|x\|_2 \leq 1$ for all $x \in \mathcal{A}$. Then we have the following proposition.

Proposition B.3. *For an linear bandit instance f with unique optimal decision and unit Gaussian noise, our complexity measure $\mathcal{C}(f)$ recovers that in Lattimore & Szepesvari (2017). That is,*

$$\mathcal{C}(f) \leq \inf_{w \in \mathbb{R}_+^A} \sum_{x \in \mathcal{A}^-} w_x \Delta_x, \quad (94)$$

$$\text{s.t. } \|x\|_{H(w)^{-1}}^2 \leq \frac{\Delta_x^2}{2}, \quad \forall x \in \mathcal{A}^-, \quad (95)$$

where $H(w) = \sum_x w_x x x^\top$.

Proof. Let \hat{w} be the solution to the RHS of Eq. (94). In the following we construct a solution to our optimization problem $\mathcal{C}(f, n)$ from \hat{w} . Recall that

$$\mathcal{C}(f, n) \triangleq \min_{w \in \mathbb{R}_+^A} \sum_{x \in \mathcal{A}} w_x \Delta(f, x) \quad (96)$$

$$\text{s.t. } \sum_{x \in \mathcal{A}} w_x D_{\text{KL}}(f[x] \| g[x]) \geq 1, \quad \forall g \in \mathcal{F}, \pi^*(g) \neq \pi^*(f), \quad (97)$$

$$\|w\|_\infty \leq n. \quad (98)$$

For any $w \in \mathbb{R}_+^A$ and any instance $g \in \mathcal{F}$ associates with parameter θ' , we have

$$\sum_{x \in \mathcal{A}} w_x D_{\text{KL}}(f[x] \| g[x]) = \sum_{x \in \mathcal{A}} w_x D_{\text{KL}}(\mathcal{N}(\langle x, \theta \rangle, 1) \| \mathcal{N}(\langle x, \theta' \rangle, 1)) = \frac{\|\theta - \theta'\|_{H(w)}^2}{2}. \quad (99)$$

Consider any g such that $\pi^*(g) \neq \pi^*(f)$. Suppose $\pi^*(g) = x \neq x^*$. It follows that $\langle x - x^*, \theta' \rangle > 0$. Recall that $\Delta_x = \langle x^* - x, \theta \rangle$. By algebraic manipulation we have,

$$\min_{\theta': \langle x^* - x, \theta - \theta' \rangle > \Delta_x} \frac{1}{2} \|\theta - \theta'\|_{H(w)}^2 = \frac{1}{2} \frac{\Delta_x^2}{\|x^* - x\|_{H(w)^{-1}}^2} \quad (100)$$

Therefore to satisfy Eq. (97), it's enough to construct a solution w such that

$$\frac{1}{2} \frac{\Delta_x^2}{\|x^* - x\|_{H(w)^{-1}}^2} \geq 1, \quad \forall x \in \mathcal{A}^-. \quad (101)$$

Define $A = \sum_{x \in \mathcal{A}^-} \hat{w}_x x x^\top + (x^*)(x^*)^\top$. When the action set \mathcal{A} is full rank, A is positive definite (see Lattimore & Szepesvari (2017, Appendix C)). We use $\sigma_{\max}(A), \sigma_{\min}(A)$ to denote the

maximum/minimum singular value of a matrix A respectively. Then for any $n > 0$, consider the following solution

$$w_x = \begin{cases} \hat{w}_x \left(1 - \frac{8}{\Delta_{\min}^2} \frac{\sigma_{\max}(A^{-1})}{1 + (n-1)\sigma_{\min}(A^{-1})}\right)^{-1}, & \text{when } x \neq x^*, \\ n, & \text{when } x = x^*. \end{cases} \quad (102)$$

For large enough n we get $\|w\|_\infty \leq n$. In the following, we prove that w satisfies Eq. (100). Let

$$c_n = \left(1 - \frac{8}{\Delta_{\min}^2} \frac{\sigma_{\max}(A^{-1})}{1 + (n-1)\sigma_{\min}(A^{-1})}\right)^{-1}$$

for shorthand. Since $\Delta_{x^*} = 0$, we have $\hat{w}_{x^*} = \infty$. Therefore $\|x^*\|_{H(\hat{w})^{-1}} = 0$. Then for any $x \in \mathcal{A}^-$, by Eq. (95) we have

$$(x^* - x)^\top H(c_n \hat{w})^{-1} (x^* - x) = c_n^{-1} (x^* - x)^\top H(\hat{w})^{-1} (x^* - x) \quad (103)$$

$$= c_n^{-1} x^\top H(\hat{w})^{-1} x \leq \frac{\Delta_x^2}{2} - \frac{4\sigma_{\max}(A^{-1})}{1 + (n-1)\sigma_{\min}(A^{-1})}. \quad (104)$$

Invoking Lemma G.7 we get

$$(x^* - x)^\top H(w)^{-1} (x^* - x) \leq (x^* - x)^\top H(c_n \hat{w})^{-1} (x^* - x) + \frac{4\sigma_{\max}(A^{-1})}{1 + (n-1)\sigma_{\min}(A^{-1})} = \frac{\Delta_x^2}{2},$$

which implies Eq. (101). As a result, w is a valid solution of $\mathcal{C}(f, n)$. Consequently,

$$\mathcal{C}(f, n) \leq \sum_x \Delta_x w_x \leq c_n \sum_x \hat{w}_x \Delta_x. \quad (105)$$

By definition we have $\lim_{n \rightarrow \infty} c_n = 1$. Consequently, $\mathcal{C}(f) = \lim_{n \rightarrow \infty} \mathcal{C}(f, n) \leq \sum_x \hat{w}_x \Delta_x$. \square

B.4 A SUFFICIENT CONDITION FOR CONDITION 1

In this section, we discuss a sufficient but more interpretable condition to Condition 1.

Proposition B.4. *Suppose there exists a constant c_M such that for every $f, g \in \mathcal{F}$ and $\pi \in \Pi$,*

$$\mathbb{E}_{o \sim f[\pi]} \left[\left(\ln \frac{f[\pi](o)}{g[\pi](o)} \right)^4 \right] \leq c_M^4.$$

Then Condition 1 holds with $\lambda_0(\alpha, \epsilon, f) = \min\{2\epsilon/c_M^2, 1/2\}$.

Proof. For any fixed $f, g \in \mathcal{F}$, $\pi \in \Pi$ and $\lambda < 1/2$, by Lemma G.8 we get

$$D_{\text{KL}}(f[\pi] \| g[\pi]) - D_{1-\lambda}(f[\pi] \| g[\pi]) \leq \frac{\lambda}{2} \mathbb{E}_{o \sim f[\pi]} \left[\left(\ln \frac{f[\pi](o)}{g[\pi](o)} \right)^4 \right]^{1/2}. \quad (106)$$

Therefore we have

$$D_{\text{KL}}(f[\pi] \| g[\pi]) - D_{1-\lambda}(f[\pi] \| g[\pi]) \leq \frac{\lambda}{2} c_M^2. \quad (107)$$

Since $D_{1-\lambda}(f[\pi] \| g[\pi])$ is monotonically decreasing with λ , we prove the desired result. \square

B.5 COMPARISON WITH EXISTING LOWER BOUNDS

Recall that when applied to tabular RL problems, our lower bound is very similar to that in Tirinzoni et al. (2021), and the only difference is that their optimization problems omit the second constraint (Eq. (4)). On a high level, the constraint $\|w_\pi\|_\infty \leq n$ is necessary to prevent degenerate cases and guarantee mathematical rigor.

The value $\mathcal{C}(f, n)$ can be different from the solution without this constraint for some artificial hypothesis class \mathcal{F} . For example, we can construct a multi-armed bandit hypothesis class

$\mathcal{F} = \{\mu \in \mathbb{R}^A : \mu(1) \neq 0.5\} \cup \{[0.5, 0.1, \dots, 0.1]\}$ (where $\mu \in \mathbb{R}^A$ represents the mean reward of each arm). Then for $f = [0.5, 0.1, \dots, 0.1]$, $\mathcal{C}(f, n) > 0$ for every $n > 0$ (because there exists other instances in \mathcal{F} whose mean reward of action 1 is arbitrarily close to 0.5). As a result, $\mathcal{C}(f) = \lim_{n \rightarrow \infty} \mathcal{C}(f, n) > 0$ by the definition of limits and in this case $\lim_{n \rightarrow \infty} \mathcal{C}(f, n) \neq \mathcal{C}(f, \infty)$. However, without the constraint $\|w_\pi\|_\infty \leq n$, the solution will be 0, achieved by letting $w_1 = \infty$ and $w_i = 0, \forall i \neq 1$.

For other hypothesis classes (such as the standard MAB and linear bandits discussed in Proposition B.2 & B.3, and tabular RL), however, this constraint does not change the value of $\mathcal{C}(f, n)$.

C EXTENSION TO INFINITE HYPOTHESIS CLASS

In this section, we extend our results to the infinite hypothesis case.

C.1 PROOF OF THEOREM 5.2

To prove Theorem 5.2, we require the following lemmas for steps 1 and 2 by analogy with the finite hypothesis case.

Lemma C.1 (Main lemma for Initialization). *Let $\mathcal{E}_{\text{init}}$ be the event that, there exists a universal constant $c_4 > 0$ such that*

- (a) $\max_{\pi} D_{\text{KL}}(f^*[\pi] \|\hat{f}[\pi]) \leq \left(\frac{\ln \ln n}{\ln n}\right)^{c_4}$,
- (b) $\left| R_{\hat{f}}(\pi) - R_{f^*}(\pi) \right| \leq R_{\max} \left(\frac{\ln \ln n}{\ln n}\right)^{c_4}$, for all $\pi \in \Pi$,
- (c) $\pi^*(\hat{f}) = \pi^*(f^*)$.

Under Conditions 1 and 2, there exists $n_0 > 0$ such that when $n > n_0$, $\Pr(\mathcal{E}_{\text{init}}) \geq 1 - 1/\ln n$. In addition, the regret of Step 1 is upper bounded by $\mathcal{O}\left(\frac{\ln n}{\ln \ln n}\right)$.

Proof of Lemma C.1 is deferred to Appendix C.2.

Lemma C.2 (Main lemma for Identification). *Under Conditions 1, 2, and 3, there exists $n_0 > 0$ such that when $n > n_0$, the following holds.*

- (a) *Conditioned on the event $\pi^*(\hat{f}) \neq \pi^*(f^*)$, $\Pr(\mathcal{E}_{\text{acc}}) \leq 1/n$.*
- (b) *Conditioned on the event $\mathcal{E}_{\text{init}}$, $\Pr(\mathcal{E}_{\text{acc}}) \geq 1 - 1/\ln n$.*
- (c) *The expected regret of Step 2 is always upper bounded by $\mathcal{O}(\ln n \ln \ln n)$.*
- (d) *Conditioned on the event $\mathcal{E}_{\text{init}}$, the expected regret of Step 2 is upper bounded by*

$$\left(\mathcal{C}(f^*, (\ln \ln n)^{1/4}/2) + o(1)\right) \ln n.$$

Proof of Lemma C.2 is deferred to Appendix C.3.

Then, proof of Theorem 5.2 is the same as that of Theorem 3.3 by plugging in Lemma C.1 and Lemma C.2,

C.2 PROOF OF LEMMA C.1

Proof of Lemma C.1. We prove the two items in this lemma separately.

Proof of item (a): Let $c_1 > 0$ be the constant from Condition 1. Set $c_6 = \frac{1}{2c_1+3}$, $\alpha = \left(\frac{\ln \ln n}{\ln n}\right)^{c_6}$, and $\epsilon = \frac{\alpha}{5}$. Let $w = (\pi_1, \dots, \pi_{m_{\text{init}}})$ be the list of decisions run by Step 1, and $o_1, \dots, o_{m_{\text{init}}}$ the corresponding observations. Recall that by definition,

$$\hat{f} = \operatorname{argmax}_{g \in \mathcal{F}} \sum_{i=1}^{m_{\text{init}}} \ln g[\pi_i](o_i). \quad (108)$$

Combining with the fact that $f^* \in \mathcal{F}$, we have

$$\sum_{i=1}^{m_{\text{init}}} \ln \frac{f^*[\pi_i](o_i)}{\hat{f}[\pi_i](o_i)} \leq \sum_{i=1}^{m_{\text{init}}} \ln \frac{f^*[\pi_i](o_i)}{f^*[\pi_i](o_i)} \leq 0. \quad (109)$$

Let $\mathcal{G}(\alpha) = \{g \in \mathcal{F} : \exists \pi : D_{\text{KL}}(f^*[\pi] \| g[\pi]) \geq \alpha\}$. We will prove that for all $g \in \mathcal{G}(\alpha)$, we have $\sum_{i=1}^{m_{\text{init}}} \ln \frac{f^*[\pi_i](o_i)}{g[\pi_i](o_i)} > 0$. Combining with Eq. (109) we get $D_{\text{KL}}(f^*[\pi] \| \hat{f}[\pi]) \leq \alpha, \forall \pi$.

To this end, we apply Lemma F.1 with parameters $(\alpha/|\Pi|, \epsilon/|\Pi|, w)$. Following Lemma F.1, define

$$\gamma = \frac{1}{m_{\text{init}}} \min_{\pi} \sum_{i=1}^{m_{\text{init}}} \mathbb{I}[\pi_i = \pi].$$

Then we have $\gamma = 1/|\Pi|$. Let $\lambda = \lambda_0(\alpha, \epsilon/|\Pi|, f)$ be the value that satisfies Condition 1, and

$$\epsilon_0 = \frac{\exp(-\alpha)(\epsilon\lambda/|\Pi|)^{1/\lambda}}{3\text{Vol}}.$$

Recall that the condition of Lemma F.1 states

$$m_{\text{init}} \geq \frac{1}{\lambda\epsilon} (\ln \mathcal{N}(\mathcal{F}, \epsilon_0) + \ln(1/\delta)). \quad (110)$$

The failure probability in this case is $\delta = 1/\ln n$. By Condition 1, for large enough n we get $1/\lambda \lesssim \epsilon^{-c_1}$. By Condition 2 we get

$$\ln \mathcal{N}(\mathcal{F}, \epsilon_0) \lesssim \ln(1/\epsilon_0) \lesssim \mathcal{O}(1) + \frac{1}{\lambda} \ln(1/(\epsilon\lambda)) + \alpha.$$

As a result, when n is large enough

$$\frac{1}{\lambda\epsilon} (\ln \mathcal{N}(\mathcal{F}, \epsilon_0) + \ln(1/\delta)) \lesssim \frac{1}{\epsilon^{c_1+1}} \left(\mathcal{O}(1) + \ln \ln n + \frac{1}{\epsilon^{c_1}} \ln \frac{1}{\epsilon^{c_1}} \right) \quad (111)$$

$$\lesssim \epsilon^{-(2c_1+2)} \lesssim \left(\frac{\ln n}{\ln \ln n} \right)^{1-\frac{1}{2c_1+3}}, \quad (112)$$

where the last inequality comes from the definition of ϵ , i.e., $\epsilon = \mathcal{O}\left(\left(\frac{\ln \ln n}{\ln n}\right)^{\frac{1}{2c_1+3}}\right)$. Recall that $m_{\text{init}} \geq |\Pi| \frac{\ln n}{\ln \ln n}$. When n is large enough, the condition of Lemma F.1 (i.e., Eq. (110)) is satisfied.

Because every policy appears in w exactly the same number of times, we have $D_{\text{KL}}^w(f^* \| g) \geq \alpha/|\Pi|$ for all $g \in \mathcal{G}(\alpha)$. Therefore, by Lemma F.1 with parameters $(\alpha/|\Pi|, \epsilon/|\Pi|, w)$,

$$\forall g \in \mathcal{G}(\alpha), \sum_{i=1}^{m_{\text{init}}} \ln \frac{f^*[\pi_i](o_i)}{g[\pi_i](o_i)} \geq \left(\frac{\alpha}{|\Pi|} - 4 \frac{\epsilon}{|\Pi|} \right) m > 0. \quad (113)$$

Combining with Eq. (109), we get $\hat{f} \notin \mathcal{G}(\alpha)$. As a result

$$\forall \pi \in \Pi, \quad D_{\text{KL}}(f^*[\pi] \| \hat{f}[\pi]) \leq \alpha = \left(\frac{\ln \ln n}{\ln n} \right)^{c_6}. \quad (114)$$

Proof of item (b): Now we focus on item (b). For any fixed $\pi \in \Pi$, by Pinsker's inequality and Eq. (114) we get

$$D_{\text{TV}}(f^*[\pi] \| \hat{f}[\pi]) \leq \sqrt{\frac{1}{2} D_{\text{KL}}(f^*[\pi] \| \hat{f}[\pi])} \leq \left(\frac{\ln \ln n}{\ln n} \right)^{c_6/2}. \quad (115)$$

By assumption we have $0 \leq R(o) \leq R_{\text{max}}$ almost surely for both $f^*[\pi]$ and $\hat{f}[\pi]$. It follows that

$$\left| R_{\hat{f}}(\pi) - R_{f^*}(\pi) \right| \leq R_{\text{max}} D_{\text{TV}}(f^*[\pi] \| \hat{f}[\pi]). \quad (116)$$

Then we prove item (b) with $c_4 = c_6/2$ and $\iota(f^*) = R_{\text{max}}$.

Proof of item (c): Since $\Delta_{\text{min}}(f^*) > 0$, (c) follows from (b) directly when n is large enough.

Proof of regret: The number of samples collected in Step 1 is upper bounded by $m_{\text{init}} = |\Pi| \lceil \frac{\ln n}{\ln \ln n} \rceil$. As a result, the regret is upper bounded by

$$\mathcal{O}(\Delta_{\text{max}} m_{\text{init}}) = \mathcal{O}\left(\frac{\ln n}{\ln \ln n}\right). \quad (117)$$

□

C.3 PROOF OF LEMMA C.2

Proof of Lemma C.2. We prove the four items in Lemma C.2 separately.

Item (a): First we prove item (a) of Lemma C.2. By Markov inequality, for any $c > 0$, we have

$$\Pr_{f^*} \left(\sum_{i=1}^m \ln \frac{\hat{f}[\pi_i](o_i)}{f^*[\pi_i](o_i)} \geq \ln n \right) = \Pr_{f^*} \left(\exp \left(\sum_{i=1}^m \ln \frac{\hat{f}[\pi_i](o_i)}{f^*[\pi_i](o_i)} \right) \geq \exp(\ln n) \right) \quad (118)$$

$$\leq \exp(-\ln n) \mathbb{E}_{f^*} \left[\exp \left(\sum_{i=1}^m \ln \frac{\hat{f}[\pi_i](o_i)}{f^*[\pi_i](o_i)} \right) \right] = \exp(-\ln n) \prod_{i=1}^m \mathbb{E}_{f^*} \left[\frac{\hat{f}[\pi_i](o_i)}{f^*[\pi_i](o_i)} \right] \quad (119)$$

$$= 1/n. \quad (120)$$

The last equality follows from the fact that $\hat{f}[\pi]$ and $f^*[\pi]$ are both probability distributions given any decision $\pi \in \Pi$.

Recall that in this case $\pi^*(\hat{f}) \neq \pi^*(f^*)$. Therefore,

$$\Pr(\mathcal{E}_{\text{acc}}^{\hat{f}}) = \Pr \left(\forall g \in \mathcal{F} \text{ and } \pi^*(g) \neq \pi^*(\hat{f}), \sum_{i=1}^m \ln \frac{\hat{f}[\pi_i](o_i)}{g[\pi_i](o_i)} \geq \ln n \right) \quad (121)$$

$$\leq \Pr \left(\sum_{i=1}^m \ln \frac{\hat{f}[\pi_i](o_i)}{f^*[\pi_i](o_i)} \geq \ln n \right) \leq 1/n. \quad (122)$$

Item (b): Let $\epsilon = 1/\ln \ln n$ and $\alpha = \frac{\ln n}{m}$. We prove this statement by invoking Lemma F.1 with parameters $(\alpha + 5\epsilon, \epsilon, w)$. Following Lemma F.1, let $\gamma = \frac{1}{m} \min_{\pi} \sum_{i=1}^m \mathbb{I}[\pi_i = \pi]$ and $\lambda = \lambda_0((\alpha + 5\epsilon)/\gamma, \epsilon, f^*)$ be the value that satisfies Condition 1. Let

$$\epsilon_0 = \frac{\exp(-(\alpha + 5\epsilon)/\gamma)(\epsilon\lambda)^{1/\lambda}}{3\text{Vol}}.$$

First of all, we prove that the condition for Lemma F.1 holds. That is,

$$m \geq \frac{1}{\lambda\epsilon} (\ln \mathcal{N}(\mathcal{F}, \epsilon_0) + \ln \ln n). \quad (123)$$

Recall that in Alg. 1, \hat{w} is the solution of $\mathcal{C}(\hat{f}, (\ln \ln n)^{1/4})$, $\bar{w}_\pi = \left(1 + \frac{1}{(\ln \ln n)^{1/4}}\right) \hat{w}_\pi + \frac{1}{(\ln \ln n)^{1/4}}$, and $m = \sum_x \lceil \bar{w}_\pi \ln n \rceil$. As a result, $m \geq \frac{|\Pi| \ln n}{(\ln \ln n)^{1/4}}$. Now consider the RHS of Eq. (123). By the definition of \bar{w}_π we get $m \leq 2|\Pi| \ln n (\ln \ln n)^{1/4}$, so $\alpha \leq \frac{1}{|\Pi|} (\ln \ln n)^{1/4}$ and $\gamma^{-1} \leq 2|\Pi| (\ln \ln n)^{1/2}$. It follows from Condition 1 that $\lambda \geq \text{poly}(1/\ln \ln n)$. By the definition of ϵ_0 and Condition 2 we get

$$\ln \mathcal{N}(\mathcal{F}, \epsilon_0) \lesssim \ln(1/\epsilon_0) \lesssim \text{poly}(\ln \ln n). \quad (124)$$

Consequently, when n is sufficiently large, Eq. (123) holds. By Lemma F.1 we get, with probability at least $1 - 1/\ln n$,

$$\sum_{i=1}^m \ln \frac{f^*[\pi_i](o_i)}{g[\pi_i](o_i)} \geq (\alpha + \epsilon)m, \quad \forall g \in \mathcal{F}(w, f^*, \alpha + 5\epsilon), \quad (125)$$

where $\mathcal{F}(w, f^*, \alpha + 5\epsilon) = \{g \in \mathcal{F} : D_{\text{KL}}^w(f^* \| g) \geq \alpha + 5\epsilon\}$. In the following, we prove that Eq. (125) implies $\mathcal{E}_{\text{acc}}^{\hat{f}}$. Recall that $\mathcal{E}_{\text{acc}}^{\hat{f}}$ is the event defined as follows:

$$\mathcal{E}_{\text{acc}}^{\hat{f}} = \mathbb{I} \left[\forall g \in \Lambda(\hat{f}), \sum_{i=1}^m \ln \frac{\hat{f}[\pi_i](o_i)}{g[\pi_i](o_i)} \geq \ln n \right]. \quad (126)$$

Recall that $\Lambda(\hat{f}) = \{g \in \mathcal{F}, \pi^*(g) \neq \pi^*(\hat{f})\}$. Next, we apply Lemma G.4 to show that $\Lambda(\hat{f}) \subseteq \mathcal{F}(w, f^*, \alpha + 5\epsilon)$. To verify the condition of Lemma G.4, we have $D_{\text{TV}}(f^*[\pi] \| \hat{f}[\pi]) \lesssim \left(\frac{\ln \ln n}{\ln n}\right)^{c_4}$ for all $\pi \in \Pi$ by item (a) of Lemma C.1. By the definition of \hat{w} we get

$$\forall g \in \Lambda(\hat{f}), \quad \sum_{\pi \in \Pi} \hat{w}_\pi D_{\text{KL}}(\hat{f}[\pi] \| g[\pi]) \geq 1. \quad (127)$$

Therefore when n is large enough,

$$\forall g \in \Lambda(\hat{f}), D_{\text{KL}}^w(f^* \| g) \geq \frac{\ln n}{m} + 5\epsilon = \alpha + 5\epsilon. \quad (128)$$

Then $\Lambda(\hat{f}) \subseteq \mathcal{F}(w, f^*, \alpha + 5\epsilon)$. It follows from Eq. (125) that

$$\sum_{i=1}^m \ln \frac{f^*[\pi_i](o_i)}{g[\pi_i](o_i)} \geq (\alpha + \epsilon)m, \quad \forall g \in \Lambda(\hat{f}). \quad (129)$$

Finally, by Condition 3 we get $f^*[\pi](o) > c_{\min}$ for any $\pi \in \Pi$ and $o \in \text{supp}(f^*[\pi])$. As a result

$$\left| \ln \frac{\hat{f}[\pi](o)}{f^*[\pi](o)} \right| \leq \left| \ln \left(1 + \frac{\hat{f}[\pi](o) - f^*[\pi](o)}{f^*[\pi](o)} \right) \right| \leq \left| \ln \left(1 + \frac{\hat{f}[\pi](o) - f^*[\pi](o)}{c_{\min}} \right) \right|. \quad (130)$$

When $\|\hat{f} - f^*\|_{\infty} \leq c_{\min}/2$, applying the basic inequality $|\ln(1+x)| \leq 2|x|, \forall |x| \leq 1/2$ we get

$$\left| \ln \left(1 + \frac{\hat{f}[\pi](o) - f^*[\pi](o)}{c_{\min}} \right) \right| \leq \frac{2}{c_{\min}} \|\hat{f} - f^*\|_{\infty} \lesssim \frac{2}{c_{\min}} \left(\frac{\ln \ln n}{\ln n} \right)^{c_4 c_5}, \quad (131)$$

where the last inequality comes from item (a) of Lemma C.1 and Condition 3. Therefore, for large enough n we get $\left| \ln \frac{\hat{f}[\pi](o)}{f^*[\pi](o)} \right| \leq \epsilon$. As a result,

$$\sum_{i=1}^m \ln \frac{\hat{f}[\pi_i](o_i)}{g[\pi_i](o_i)} = \sum_{i=1}^m \ln \frac{f^*[\pi_i](o_i)}{g[\pi_i](o_i)} + \sum_{i=1}^m \ln \frac{\hat{f}[\pi_i](o_i)}{f^*[\pi_i](o_i)} \geq \alpha m = \ln n. \quad (132)$$

Since $g \in \Lambda(\hat{f})$ is arbitrary, have

$$\Pr(\mathcal{E}_{\text{acc}}^{\hat{f}}) = \Pr \left(\forall g \in \Lambda(\hat{f}), \sum_{i=1}^m \ln \frac{\hat{f}[\pi_i](o_i)}{g[\pi_i](o_i)} \geq \ln n \right) \quad (133)$$

$$\geq \Pr \left(\forall g \in \mathcal{F}(w, f^*, \alpha + 5\epsilon), \sum_{i=1}^m \ln \frac{f^*[\pi_i](o_i)}{g[\pi_i](o_i)} \geq (\alpha + \epsilon)m \right) \geq 1 - 1/\ln n. \quad (134)$$

Item (c): By the definition of $\mathcal{C}(\hat{f}, (\ln \ln n)^{1/4})$, we have $\hat{w}_{\pi} \leq (\ln \ln n)^{1/4}$ for every $\pi \in \Pi$. As a result, $m \leq 2A \ln n (\ln \ln n)^{1/4}$. Therefore, the expect regret of Step 2 is upper bounded by

$$\Delta_{\max} 2A \ln n (\ln \ln n)^{1/4} = O(\ln n \ln \ln n). \quad (135)$$

Item (d): Recall that \hat{w} is the solution of $\mathcal{C}(\hat{f}, (\ln \ln n)^{1/4})$. As a result, the regret of Step 2 is upper bounded by

$$\left(\sum_{\pi \in \Pi} \hat{w}_{\pi} \Delta(f^*, \pi) + o(1) \right) \ln n \quad (136)$$

where $\Delta(f^*, \pi)$ is the sub-optimality gap of decision π under instance f^* . In the following, we prove that

$$\sum_{\pi \in \Pi} \hat{w}_{\pi} \Delta(f^*, \pi) \leq \mathcal{C}(f^*, (\ln \ln n)^{1/4}/2). \quad (137)$$

Let \hat{w}^* be the solution to $\mathcal{C}(f^*, (\ln \ln n)^{1/4}/2)$. Define $\delta = \frac{1}{(\ln \ln n)^{1/4}}$. Let $\bar{w}^* \triangleq \{(1 + \delta)\hat{w}_{\pi}^* + \delta\}_{\pi}$ and $m^* = \sum_{\pi \in \Pi} \lceil \bar{w}_{\pi}^* \ln n \rceil$. We will show that \bar{w}^* is also (approximately) a solution of $\mathcal{C}(\hat{f}, (\ln \ln n)^{1/4}/2)$

By the definition of $\mathcal{C}(f^*, (\ln \ln n)^{1/4}/2)$, for any $g \in \Lambda(f^*)$, we have

$$\sum_{\pi \in \Pi} (\hat{w}_{\pi}^*) D_{\text{KL}}(f^*[\pi] \| g[\pi]) \geq 1. \quad (138)$$

Let w^* be the list of decisions that π appears $\lceil \bar{w}^* \ln n \rceil$ times for every $\pi \in \Pi$. Define $\alpha^* = \frac{\ln n}{m^*}$ and $\epsilon = \frac{1}{\ln \ln n}$. Next we apply Lemma G.4 with parameters (ϵ, w^*) . To verify its condition, item (a) of Lemma C.1 gives

$$D_{\text{TV}}(f^*[\pi] \|\hat{f}[\pi]) \leq D_{\text{KL}}(f^*[\pi] \|\hat{f}[\pi])^{1/2} \lesssim \left(\frac{\ln \ln n}{\ln n} \right)^{c_4/2}, \quad \forall \pi \in \Pi, \quad (139)$$

which satisfies the condition of Lemma G.4. Consequently, we get $D_{\text{KL}}^{w^*}(\hat{f}[\pi] \| g[\pi]) \geq \frac{\ln n}{m^*}$ for every $g \in \Lambda(f^*)$. Therefore,

$$\forall g \in \Lambda(f^*), \quad \sum_{\pi \in \Pi} \frac{\lceil \bar{w}_\pi^* \ln n \rceil}{\ln n} D_{\text{KL}}(\hat{f}[\pi] \| g[\pi]) \geq 1. \quad (140)$$

By item (c) of Lemma C.1, $\Lambda(\hat{f}) = \Lambda(f^*)$. When n is large enough, we have $\frac{\lceil \bar{w}_\pi^* \ln n \rceil}{\ln n} \leq (\ln \ln n)^{1/4}$. Therefore, $\left\{ \frac{\lceil \bar{w}_\pi^* \ln n \rceil}{\ln n} \right\}_{\pi \in \Pi}$ satisfies all the constraints of $\mathcal{C}(\hat{f}, (\ln \ln n)^{1/4})$. Recall that \hat{w} is the solution to $\mathcal{C}(\hat{f}, (\ln \ln n)^{1/4})$. By the optimality of \hat{w} we have

$$\sum_{\pi} \hat{w}_\pi \Delta(\hat{f}, \pi) \leq \sum_{\pi} \frac{\lceil \bar{w}_\pi^* \ln n \rceil}{\ln n} \Delta(\hat{f}, \pi) \leq \sum_{\pi} \bar{w}_\pi^* \Delta(\hat{f}, \pi) + o(1). \quad (141)$$

By item (b) of Lemma C.1, $|\Delta(\hat{f}, \pi) - \Delta(f^*, \pi)| \leq \iota(f^*) \left(\frac{\ln \ln n}{\ln n} \right)^{c_4}$. As a result,

$$\sum_{\pi} \bar{w}_\pi^* \Delta(\hat{f}, \pi) \leq \sum_{\pi} \bar{w}_\pi^* \Delta(f^*, \pi) + o(1) \quad (142)$$

$$\leq \sum_{\pi} \hat{w}_\pi^* \Delta(f^*, \pi) + o(1) = \mathcal{C}(f^*, (\ln \ln n)^{1/4}/2) + o(1). \quad (143)$$

In addition,

$$\sum_{\pi \in \Pi} \hat{w}_\pi \Delta(f^*, \pi) \leq \sum_{\pi} \hat{w}_\pi \Delta(\hat{f}, \pi) + o(1). \quad (144)$$

Stitching the inequalities above we have

$$\sum_{\pi} \hat{w}_\pi \Delta(f^*, \pi) \leq \mathcal{C}(f^*, (\ln \ln n)^{1/4}/2) + o(1). \quad (145)$$

As a result, the regret in Step 2 is bounded by

$$\left(\mathcal{C}(f^*, (\ln \ln n)^{1/4}/2) + o(1) \right) \ln n.$$

□

D PROOF OF CONDITIONS FOR TABULAR REINFORCEMENT LEARNING

In this section, we first prove that Conditions 1- 3 holds for tabular RL with truncated Gaussian reward. Consequently, Theorem 5.2 gives the first asymptotic instance-optimal regret bound in this setting. Later in Appendix E, we extend our analysis to Gaussian reward.

First of all, we define some additional notations. A finite horizon Markov Decision Process (MDP) is denoted by a tuple $(\mathcal{S}, \mathcal{A}, p, r, H)$. \mathcal{S} and \mathcal{A} are the state and action spaces respectively. $H > 0$ denotes the horizon. The transition function p maps a state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$ to a distribution over \mathcal{S} , and the reward function r maps state-action pair (s, a) to a truncated Gaussian distribution. In particular, for a fixed (s, a) , the mean reward is $\mu(s, a) \in [-1, 1]$, and the density of the reward $r(s, a)$ is

$$r[s, a](x) = \mathbb{I}[x \in [-2, 2]] \frac{1}{Z} \exp\left(-\frac{(x - \mu(s, a))^2}{2}\right), \quad (146)$$

where $Z = \int_{x \in [-2, 2]} \exp(-(x - \mu(s, a))^2 / 2)$ is the normalization factor. We use a fixed truncation for every state-action pair regardless of $\mu(s, a)$ because otherwise the KL divergence of two instances will easily be infinity. Without loss of generality, we assume \mathcal{S} is partitioned into $\mathcal{S}_1, \dots, \mathcal{S}_H$ where for any state $s \in \mathcal{S}_h$ and action $a \in \mathcal{A}$, $\text{supp}(p(s, a)) \subseteq \mathcal{S}_{h+1}$. We also assume that the initial state s_1 is fixed.

A decision (a.k.a. policy) $\pi : \mathcal{S} \rightarrow \mathcal{A}$ is a deterministic mapping from a state to an action, and $\Pi = \mathcal{A}^{\mathcal{S}}$ is the set of possible decisions with $|\Pi| = |\mathcal{A}|^{|\mathcal{S}|} < \infty$. An observation o is a trajectory $(s_1, a_1, r_1, \dots, s_H, a_H, r_H)$. The reward of an observation is $R(o) = \sum_{h=1}^H r_h$. We emphasize that s_h, a_h are discrete random variables and $r_h \in \mathbb{R}$ is a continuous random variable. In the following, we use $o^{\text{sa}} = (s_1, a_1, \dots, s_H, a_H)$ to denote the set of state-action pairs in o and $o^r = (r_1, \dots, r_H)$ the set of rewards. As a basic property of MDP, the rewards are independent conditioned on o^{sa} , and $r_h \sim r(s_h, a_h)$.

An instance f is represented by the transition function p and reward function r (which is uniquely determined by $\mu : \mathcal{S} \times \mathcal{A} \rightarrow [-1, 1]$). The family of instances \mathcal{F} is defined as the set of all possible instances.

For an instance f and a decision π , the density of an observation $o = (s_1, a_1, r_1, \dots, s_H, a_H, r_H)$ is

$$f[\pi](o) = \prod_{h=1}^H p[s_h, a_h](s_{h+1}) r[s_h, a_h](r_h). \quad (147)$$

For simplicity, we also use $f_p[s, a](s')$ to denote the distribution $p[s, a](s')$, and $f_r[s, a](r)$ to denote $r[s, a](r)$.

By definition, for all s, a, r, s' we have $p[s, a](s') r[s, a](r) < 1$. We also define

$$f[\pi](o^{\text{sa}}) = \prod_{h=1}^H p[s_h, a_h](s_{h+1}) \quad (148)$$

to be the marginal density of o^{sa} , and

$$f[\pi](o^r | o^{\text{sa}}) = \prod_{h=1}^H r[s_h, a_h](r_h)$$

the conditional distribution of o^r .

For an instance $f \in \mathcal{F}$, we define $\mu_{\min}(f) = \min_{o^{\text{sa}}, \pi: f[\pi](o^{\text{sa}}) > 0} f[\pi](o^{\text{sa}})$. Since o^{sa}, π are finite, we have $\mu_{\min}(f) > 0$.

D.1 PROOF OF THEOREM 5.1

In this section, we prove Theorem 5.1.

Proof. In the following, we prove the three conditions separately. **Proof of Condition 1.** We prove this condition by invoking Lemma D.1. For instances $f, g \in \mathcal{F}$ and state-action pair $s, a \in \mathcal{S} \times \mathcal{A}$, let $f_r[s, a], g_r[s, a]$ be their reward distributions respectively. Since $f_r[s, a], g_r[s, a]$ are truncated Gaussian distributions, we get $\exp(-3) \leq g_r[s, a](x) \leq 1$ for every $x \in \text{supp}(f_r[s, a])$ and $g \in \mathcal{F}$. Therefore $\sup_x \left| \ln \frac{f_r[s, a](x)}{g_r[s, a](x)} \right| \leq 3$ for all $s \in \mathcal{S}, a \in \mathcal{A}$. By Lemma D.1 with $c_M = 3$, we get Condition 1.

Proof of Condition 2. The first part of this condition is proved by Lemma D.2. On the other hand, we have

$$\int 1 d\mathbf{o} = (2|\mathcal{S}||\mathcal{A}|)^H < \infty. \quad (149)$$

Proof of Condition 3. To prove the first part of this condition, recall that

$$f^*[\pi](o) = \prod_{h=1}^H p[s_h, a_h](s_{h+1}) r[s_h, a_h](r_h), \quad (150)$$

where p is the transition function and r is the reward distribution. Let

$$c_{\min} = \left(\min_{s,a \in \mathcal{S} \times \mathcal{A}, s' \in \text{supp } p[s,a](\cdot)} p[s,a](s') \right)^H \exp(-4H),$$

As a result, $f^*[\pi](o) > c_{\min}$ for all $o \in \text{supp}(f^*[\pi])$. We prove the second part of this condition in Lemma D.3. \square

D.2 PROOF OF CONDITION 1

In this section, we present a lemma that establishes Condition 1 for tabular RL.

Lemma D.1. *Consider any fixed RL instance f with discrete state, action and general reward distribution. Suppose there exists a constant $c_M > 0$ such that for any $g \in \mathcal{F}$, $s \in \mathcal{S}$, $a \in \mathcal{A}$, the reward distributions of instance f and g at state s and action a (denoted by $f_r[s,a], g_r[s,a]$ respectively) satisfy*

$$\mathbb{E}_{x \sim f_r[s,a]} \left[\left(\ln \frac{f_r[s,a](x)}{g_r[s,a](x)} \right)^4 \right] \leq c_M^4. \quad (151)$$

Then for every $\alpha > 0, \epsilon \in (0, 1)$, Condition 1 holds with

$$\lambda_0(\alpha, \epsilon, f) = \frac{\epsilon}{32H^2} \min \left\{ \frac{\mu_{\min}(f)}{4\alpha}, \frac{\mu_{\min}(f)}{10}, \frac{1}{c_M} \right\}^2. \quad (152)$$

Proof of Lemma D.1. Let

$$\kappa = \frac{\mu_{\min}(f)}{e^2} \exp \left(-\frac{2\alpha}{\mu_{\min}(f)} \right). \quad (153)$$

Recall that for reinforcement learning, an observation $o = (s_1, a_1, r_1, \dots, s_H, a_H, r_H)$ represents a trajectory, and $o^{\text{sa}} = (s_1, a_1, s_2, a_2, \dots, s_H, a_H)$ denotes the state-action pairs in the trajectory. Consider any fixed decision π , for any $\lambda < \lambda_0$, we prove the following two cases separately.

Case 1: $\min_{o^{\text{sa}}: f[\pi](o^{\text{sa}}) > 0} g[\pi](o^{\text{sa}}) < \kappa$. In this case, we prove that $D_{1-\lambda}(f[\pi] \| g[\pi]) \geq \alpha$.

Let $o^{\hat{\text{sa}}} = \text{argmin}_{o^{\text{sa}}: f[\pi](o^{\text{sa}}) > 0} g[\pi](o^{\text{sa}})$. By the condition of this case we have $g[\pi](o^{\hat{\text{sa}}}) < \kappa$. Applying Lemma G.1 we get

$$D_{1-\lambda}(f[\pi](o) \| g[\pi](o)) \geq D_{1-\lambda}(f[\pi](o^{\text{sa}}) \| g[\pi](o^{\text{sa}})). \quad (154)$$

In the following we prove the RHS of Eq. (154) is larger than α . We start with Hölder's inequality and the basic inequality that $(1-x)^t \leq 1-tx$ for any $t \in (0, 1), x \in (0, 1)$:

$$\sum_o^{\text{sa}} f[\pi](o^{\text{sa}})^{1-\lambda} g[\pi](o^{\text{sa}})^\lambda \quad (155)$$

$$= \sum_{o^{\text{sa}} \neq o^{\hat{\text{sa}}}} f[\pi](o^{\text{sa}})^{1-\lambda} g[\pi](o^{\text{sa}})^\lambda + f[\pi](o^{\hat{\text{sa}}})^{1-\lambda} g[\pi](o^{\hat{\text{sa}}})^\lambda \quad (156)$$

$$\leq \left(\sum_{o^{\text{sa}} \neq o^{\hat{\text{sa}}}} f[\pi](o^{\text{sa}}) \right)^{1-\lambda} \left(\sum_{o^{\text{sa}} \neq o^{\hat{\text{sa}}}} g[\pi](o^{\text{sa}}) \right)^\lambda + f[\pi](o^{\hat{\text{sa}}})^{1-\lambda} g[\pi](o^{\hat{\text{sa}}})^\lambda \quad (157)$$

$$\leq (1 - f[\pi](o^{\hat{\text{sa}}}))^{1-\lambda} (1 - g[\pi](o^{\hat{\text{sa}}}))^\lambda + f[\pi](o^{\hat{\text{sa}}})^{1-\lambda} g[\pi](o^{\hat{\text{sa}}})^\lambda \quad (158)$$

$$\leq (1 - f[\pi](o^{\hat{\text{sa}}}))^{1-\lambda} + f[\pi](o^{\hat{\text{sa}}})^{1-\lambda} g[\pi](o^{\hat{\text{sa}}})^\lambda \quad (159)$$

$$\leq 1 - f[\pi](o^{\hat{\text{sa}}})(1-\lambda) + f[\pi](o^{\hat{\text{sa}}})^{1-\lambda} g[\pi](o^{\hat{\text{sa}}})^\lambda. \quad (160)$$

Recall the basic inequality $\ln(1+x) \leq x$ for all $x > -1$. Therefore,

$$\frac{1}{\lambda} \ln \left(\sum_o^{\text{sa}} f[\pi](o^{\text{sa}})^{1-\lambda} g[\pi](o^{\text{sa}})^\lambda \right) \quad (161)$$

$$\leq \frac{1}{\lambda} (-f[\pi](o^{\hat{s}a})(1-\lambda) + f[\pi](o^{\hat{s}a})^{1-\lambda} g[\pi](o^{\hat{s}a})^\lambda) \quad (162)$$

$$\leq \frac{1}{\lambda} \left(f[\pi](o^{\hat{s}a}) \left(\left(\frac{g(o^{\hat{s}a})}{f(o^{\hat{s}a})} \right)^\lambda - 1 \right) + \lambda f[\pi](o^{\hat{s}a}) \right) \quad (163)$$

$$\leq \frac{1}{\lambda} \left(f[\pi](o^{\hat{s}a}) \left(\left(\frac{\kappa}{f(o^{\hat{s}a})} \right)^\lambda - 1 \right) + \lambda f[\pi](o^{\hat{s}a}) \right). \quad (164)$$

Recall the basic inequality that $\exp(x) \leq 1 + x/2$ for all $-1 \leq x \leq 0$. Since we have

$$\left(\frac{\kappa}{f[\pi](o^{\hat{s}a})} \right)^\lambda = \exp \left(\lambda \ln \left(\frac{\kappa}{f[\pi](o^{\hat{s}a})} \right) \right),$$

when $\lambda \leq (\ln(f[\pi](o^{\hat{s}a})/\kappa))^{-1}$ we get

$$\frac{1}{\lambda} \left(f[\pi](o^{\hat{s}a}) \left(\left(\frac{\kappa}{f[\pi](o^{\hat{s}a})} \right)^\lambda - 1 \right) + \lambda f[\pi](o^{\hat{s}a}) \right) \quad (165)$$

$$\leq \frac{1}{2} f[\pi](o^{\hat{s}a}) \ln(\kappa/f[\pi](o^{\hat{s}a})) + f[\pi](o^{\hat{s}a}) = \frac{1}{2} f[\pi](o^{\hat{s}a}) \ln(e^2 \kappa / f[\pi](o^{\hat{s}a})). \quad (166)$$

By the definition of κ we get

$$\frac{1}{2} f[\pi](o^{\hat{s}a}) \ln(e^2 \kappa / f[\pi](o^{\hat{s}a})) \leq -\alpha, \quad (167)$$

which leads to $D_{1-\lambda}(f[\pi](o^{\text{sa}}) \| g[\pi](o^{\text{sa}})) \geq \alpha$.

Case 2: $\min_{o^{\text{sa}}: f[\pi](o^{\text{sa}}) > 0} g[\pi](o^{\text{sa}}) \geq \kappa$. By Lemma G.8, for any $\lambda \in (0, 1/2)$ we get

$$D_{\text{KL}}(f[\pi] \| g[\pi]) - D_{1-\lambda}(f[\pi] \| g[\pi]) \leq \frac{\lambda}{2} \mathbb{E}_{o \sim f[\pi]} \left[\left(\ln \frac{f[\pi](o)}{g[\pi](o)} \right)^4 \right]^{1/2}. \quad (168)$$

Let $f_r, g_r : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathbb{R})$ be the reward distributions of instance f and g respectively, and $f_r[s, a](\cdot), g_r[s, a](\cdot)$ the densities of the reward given state-action pair (s, a) . Recall that for a trajectory $o = (s_1, a_1, r_1, \dots, s_H, a_H, r_H)$ we have

$$f[\pi](o) = f[\pi](o^{\text{sa}}) \prod_{h=1}^H f_r[s_h, a_h](r_h). \quad (169)$$

By Hölder's inequality we get

$$\mathbb{E}_{o \sim f[\pi]} \left[\left(\ln \frac{f[\pi](o)}{g[\pi](o)} \right)^4 \right] \quad (170)$$

$$= \mathbb{E}_{o \sim f[\pi]} \left[\left(\ln \frac{f[\pi](o^{\text{sa}})}{g[\pi](o^{\text{sa}})} + \sum_{h=1}^H \ln \frac{f_r[s_h, a_h](r_h)}{g_r[s_h, a_h](r_h)} \right)^4 \right] \quad (171)$$

$$= \mathbb{E}_{o \sim f[\pi]} \left[(H+1)^3 \left(\left(\ln \frac{f[\pi](o^{\text{sa}})}{g[\pi](o^{\text{sa}})} \right)^4 + \sum_{h=1}^H \left(\ln \frac{f_r[s_h, a_h](r_h)}{g_r[s_h, a_h](r_h)} \right)^4 \right) \right] \quad (172)$$

$$\leq (H+1)^3 \ln(1/\kappa)^4 + (H+1)^4 \sup_{s \in \mathcal{S}, a \in \mathcal{A}} \mathbb{E}_{x \sim f_r[s, a]} \left[\left(\ln \frac{f_r[s, a](x)}{g_r[s, a](x)} \right)^4 \right] \quad (173)$$

$$\leq (H+1)^3 \ln(1/\kappa)^4 + (H+1)^4 c_M^4, \quad (174)$$

where the last inequality comes from item (c) of Condition 4. Therefore, when $\lambda \leq \epsilon(H+1)^{-2} \min\{\ln(1/\kappa)^{-2}, c_M^{-2}\}$ we get

$$D_{\text{KL}}(f[\pi] \| g[\pi]) - D_{1-\lambda}(f[\pi] \| g[\pi]) \leq \frac{\lambda}{2} \mathbb{E}_{o \sim f[\pi]} \left[\left(\ln \frac{f[\pi](o)}{g[\pi](o)} \right)^4 \right]^{1/2} \quad (175)$$

$$\leq \frac{\epsilon}{2} (H+1)^{-2} \min\{\ln(1/\kappa)^{-2}, c_M^{-2}\} \left((H+1)^3 \ln(1/\kappa)^4 + (H+1)^4 c_M^4 \right)^{1/2} \quad (176)$$

$$\leq \frac{\epsilon}{2} \min\{\ln(1/\kappa)^{-2}, c_M^{-2}\} \left(\ln(1/\kappa)^4 + c_M^4 \right)^{1/2} \quad (177)$$

$$\leq \frac{\epsilon}{2} \min\{\ln(1/\kappa)^{-2}, c_M^{-2}\} \left(\ln(1/\kappa)^2 + c_M^2 \right) \quad (178)$$

$$\leq \epsilon. \quad (179)$$

Recall that

$$\kappa = \frac{\mu_{\min}(f)}{e^2} \exp\left(-\frac{2\alpha}{\mu_{\min}(f)}\right). \quad (180)$$

By algebraic manipulation we get

$$\frac{\epsilon}{32H^2} \min\left\{\frac{\mu_{\min}(f)}{4\alpha}, \frac{\mu_{\min}(f)}{10}, \frac{1}{c_M}\right\}^2 \leq \epsilon (H+1)^{-2} \min\{\ln(1/\kappa)^{-2}, c_M^{-2}\}, \quad (181)$$

which proves the desired result. \square

D.3 PROOF OF CONDITION 2

Recall that for two instances $f, g \in \mathcal{F}$, their distance is $d(f, g) = \sup_{\pi, o} |f[\pi](o) - g[\pi](o)|$.

Lemma D.2. *Suppose \mathcal{F} represents tabular RL with truncated Gaussian reward with state space \mathcal{S} and action space \mathcal{A} , then we have $\ln \mathcal{N}(\mathcal{F}, \epsilon) \leq \mathcal{O}(|\mathcal{S}||\mathcal{A}| \ln(1/\epsilon))$ for every $\epsilon > 0$.*

Proof. For instances $f \in \mathcal{F}$ and state-action pair $s, a \in \mathcal{S} \times \mathcal{A}$, let $f_r[s, a]$ be its reward distribution and $f_p[s, a]$ its transition. Recall that for tabular RL we have

$$f[\pi](o) = \prod_{h=1}^H f_p[s_h, a_h](s_{h+1}) f_r[s_h, a_h](r_h), \quad (182)$$

where the observation is $o = (s_1, a_1, r_1, \dots, s_H, a_H, r_H)$. Consequently,

$$\begin{aligned} & |f[\pi](o) - g[\pi](o)| \\ & \leq \sum_{h=1}^H \left(|f_p[s_h, a_h](s_{h+1}) f_r[s_h, a_h](r_h) - g_p[s_h, a_h](s_{h+1}) g_r[s_h, a_h](r_h)| \right. \\ & \quad \left. \prod_{h'=1}^{h-1} f_p[s_{h'}, a_{h'}](s_{h'+1}) f_r[s_{h'}, a_{h'}](r_{h'}) \prod_{h'=h+1}^H g_p[s_{h'}, a_{h'}](s_{h'+1}) g_r[s_{h'}, a_{h'}](r_{h'}) \right) \\ & \leq \sum_{h=1}^H |f_p[s_h, a_h](s_{h+1}) f_r[s_h, a_h](r_h) - g_p[s_h, a_h](s_{h+1}) g_r[s_h, a_h](r_h)| \\ & \leq \sum_{h=1}^H (|f_p[s_h, a_h](s_{h+1}) - g_p[s_h, a_h](s_{h+1})| + |f_r[s_h, a_h](r_h) - g_r[s_h, a_h](r_h)|) \end{aligned}$$

Therefore, we construct a covering \mathcal{C} such that for every $f \in \mathcal{F}$, there exists $g \in \mathcal{C}$ with

$$|f_p[s, a](s') - g_p[s, a](s')| \leq \frac{\epsilon}{2H}, \quad \forall s, a, s', \quad (183)$$

$$|f_r[s, a](r) - g_r[s, a](r)| \leq \frac{\epsilon}{2H}, \quad \forall s, a, r. \quad (184)$$

Since $f_p[s, a]$ is a discrete distribution, the covering number for the transition function is upper bounded by $(4H/\epsilon)^{|\mathcal{S}||\mathcal{A}|}$. On the other hand, $f_r[s, a](r)$ is a truncated Gaussian distribution, so covering number for the reward is also upper bounded by $(4H/\epsilon)^{|\mathcal{S}||\mathcal{A}|}$. As a result,

$$\ln \mathcal{N}(\mathcal{F}, \epsilon) \leq \mathcal{O}(|\mathcal{S}||\mathcal{A}| \ln(1/\epsilon)). \quad (185)$$

\square

D.4 PROOF OF CONDITION 3

Lemma D.3. *Consider tabular reinforcement learning with truncated Gaussian reward. For a fixed instance f^* , for all $f \in \mathcal{F}, \pi \in \Pi$ we have*

$$\|f^*[\pi] - f[\pi]\|_\infty \leq \frac{37H}{\mu_{\min}(f^*)^{1/6}} D_{\text{KL}}(f^*[\pi] \| f[\pi])^{1/6}. \quad (186)$$

Proof. Recall that for tabular RL, an observation is $o = (s_1, a_1, r_1, \dots, s_H, a_H, r_H)$. We use $o^{\text{sa}} = (s_1, a_1, s_2, a_2, \dots, s_H, a_H)$ to denote the collection of states and actions in the trajectory o , and $o^r = (r_1, \dots, r_H)$ the collection of rewards. For instances $f \in \mathcal{F}$ and state-action pair $s, a \in \mathcal{S} \times \mathcal{A}$, let $f_r[s, a]$ be its reward distribution and $f_p[s, a]$ its transition.

Let $\epsilon_0 = \frac{H}{\mu_{\min}(f^*)^{1/6}}$. Consider the random variables o^{sa} and o^r . By the chain rules of KL divergence we have

$$D_{\text{KL}}(f^*[\pi] \| f[\pi]) = D_{\text{KL}}(f^*[\pi](o^{\text{sa}}) \| f[\pi](o^{\text{sa}})) + \mathbb{E}_{o^{\text{sa}} \sim f^*} [D_{\text{KL}}(f^*[\pi](o^r | o^{\text{sa}}) \| f[\pi](o^r | o^{\text{sa}}))]. \quad (187)$$

Since o^{sa} is a discrete random variable, we have

$$|f^*[\pi](o^{\text{sa}}) - f[\pi](o^{\text{sa}})| \leq D_{\text{TV}}(f^*[\pi] \| f[\pi]) \leq D_{\text{KL}}(f^*[\pi] \| f[\pi])^{1/2}. \quad (188)$$

Therefore, for any $o^{\text{sa}} \notin \text{supp}(f^*[\pi])$,

$$|f^*[\pi](o^{\text{sa}}, o^r) - f[\pi](o^{\text{sa}}, o^r)| \leq f[\pi](o^{\text{sa}}, o^r) \leq f[\pi](o^{\text{sa}}) f[\pi](o^r | o^{\text{sa}}) \leq D_{\text{KL}}(f^*[\pi] \| f[\pi])^{1/2}, \quad (189)$$

where the last inequality comes from the fact that $f[\pi](o^r | o^{\text{sa}}) = \prod_{(s_h, a_h, r_h) \in o} f_r[s_h, a_h](r_h) \leq 1$.

Now, for any $o^{\text{sa}} \in \text{supp}(f^*[\pi])$, by Eq. (187) we get

$$D_{\text{KL}}(f^*[\pi](o^r | o^{\text{sa}}) \| f[\pi](o^r | o^{\text{sa}})) \leq \frac{1}{\mu_{\min}(f^*)} D_{\text{KL}}(f^*[\pi] \| f[\pi]). \quad (190)$$

By the chain rule of KL divergence,

$$D_{\text{KL}}(f^*[\pi](o^r | o^{\text{sa}}) \| f[\pi](o^r | o^{\text{sa}})) = \sum_{h=1}^H D_{\text{KL}}(f^*_r[s_h, a_h] \| f_r[s_h, a_h]). \quad (191)$$

As a result, for every $(s, a) \in o^{\text{sa}}$ we get

$$D_{\text{KL}}(f^*_r[s, a] \| f_r[s, a]) \leq \frac{1}{\mu_{\min}(f^*)} D_{\text{KL}}(f^*[\pi] \| f[\pi]). \quad (192)$$

By Lemma G.6, $|f^*_r[s, a](r) - f_r[s, a](r)| \leq 36 \left(\frac{1}{\mu_{\min}(f^*)} D_{\text{KL}}(f^*[\pi] \| f[\pi]) \right)^{1/6}$. Therefore,

$$|f^*[\pi](o^r | o^{\text{sa}}) - f[\pi](o^r | o^{\text{sa}})| \quad (193)$$

$$= \left| \prod_{h=1}^H f^*_r[s_h, a_h](r_h) - \prod_{h=1}^H f_r[s_h, a_h](r_h) \right| \quad (194)$$

$$\leq \sum_{h=1}^H |f^*_r[s_h, a_h](r_h) - f_r[s_h, a_h](r_h)| \quad (195)$$

$$\leq 36H \left(\frac{1}{\mu_{\min}(f^*)} D_{\text{KL}}(f^*[\pi] \| f[\pi]) \right)^{1/6}. \quad (196)$$

It follows that,

$$|f^*[\pi](o^{\text{sa}}, o^r) - f[\pi](o^{\text{sa}}, o^r)| \quad (197)$$

$$\leq |f^*[\pi](o^{\text{sa}}) - f[\pi](o^{\text{sa}})| + |f^*[\pi](o^r | o^{\text{sa}}) - f[\pi](o^r | o^{\text{sa}})| \quad (198)$$

$$\leq D_{\text{KL}}(f^*[\pi] \| f[\pi])^{1/2} + 36H \left(\frac{1}{\mu_{\min}(f^*)} D_{\text{KL}}(f^*[\pi] \| f[\pi]) \right)^{1/6} \quad (199)$$

$$\leq 37H \left(\frac{1}{\mu_{\min}(f^*)} D_{\text{KL}}(f^*[\pi] \| f[\pi]) \right)^{1/6} \quad (200)$$

□

E RL WITH GENERAL REWARD DISTRIBUTION

In this section, we extend our analysis to RL with general reward distribution, where the support of the reward may have infinite volume (e.g., the real line \mathbb{R} when the reward distribution is Gaussian). We assume that for any state-action pair (s, a) , the reward distribution $r[s, a]$ comes from a distribution family \mathcal{R} (e.g., $\mathcal{R} = \{\mathcal{N}(\mu, 1) : \mu \in [-1, 1]\}$ when the reward is Gaussian). For any $g \in \mathcal{R}$, let $g(\cdot)$ be its density function and $\mu(g) \triangleq \mathbb{E}_{x \sim g}[x]$ its mean and we assume $\sup_{g \in \mathcal{R}} \mu(g) \leq R_{\max}$. We emphasize that for general reward distributions, we do not require Conditions 2 and 3 because they do not hold for Gaussian distribution. Instead, we require the following.

Condition 4. Let \mathcal{R} be the reward distribution family. Then

- (a) for all $f \in \mathcal{R}$, there exists a constant $c_7 \in (0, 1]$, $c_8 \in (0, 1]$, $c_M > 0$, such that for every $\delta > 0$,

$$\Pr_{x \sim f} \left(\forall g, g' \in \mathcal{R}, \left| \ln \frac{g(x)}{g'(x)} \right| > D_{\text{TV}}(g \| g')^{c_7} \text{polylog}(1/\delta) \right) \leq \delta; \quad (201)$$

- (b) for all $g, g' \in \mathcal{R}$, $|\mu(g) - \mu(g')| \lesssim D_{\text{TV}}(g \| g')^{c_8}$;

- (c) for all $g, g' \in \mathcal{R}$,

$$\mathbb{E}_{x \sim g} \left[\left(\ln \frac{g(x)}{g'(x)} \right)^4 \right] \leq c_M^4;$$

- (d) for any $\epsilon > 0$, there exists a covering $\mathcal{C}(\mathcal{R}, \epsilon) \subseteq \mathcal{R}$ such that

$$\forall g \in \mathcal{R}, \exists g' \in \mathcal{C}(\mathcal{R}, \epsilon), D_{\text{TV}}(g \| g') \leq \epsilon. \quad (202)$$

And $\log |\mathcal{C}(\mathcal{R}, \epsilon)| = \mathcal{O}(\log(1/\epsilon))$.

Condition 1 still holds in this case because of Lemma D.1 and item (c) of Condition 4.

For tabular reinforcement learning problems, we only require Condition 4 for the reward distribution of any fixed state-action pair, which is one-dimensional. Condition 4 holds for Gaussian distribution (see Proposition E.5), Laplace distribution, etc.

Our main result for tabular RL is stated as follows.

Theorem E.1. Suppose \mathcal{F} is the hypothesis class representing tabular reinforcement learning with a reward distribution that satisfies Conditions 4, then the regret of Alg. 1 satisfies

$$\limsup_{n \rightarrow \infty} \frac{\text{Reg}_{f^*, n}}{\ln n} \leq \mathcal{C}(f^*). \quad (203)$$

In section E.1 we present the main lemmas for Steps 1-2. Then, Theorem E.1 is proved using exactly the same approach as Theorem 5.2 (by replacing the main lemmas for Steps 1-2).

E.1 LEMMAS FOR RL WITH GENERAL REWARD DISTRIBUTION

In this section, we state the main lemmas for RL with general reward distribution (namely, Lemma E.2 for Step 1 and Lemma E.3 for Step 2).

Lemma E.2. Let $\mathcal{E}_{\text{init}}$ be the event that, there exists a universal constant $c_4 > 0$ and a function $\iota(f^*)$ that only depends on f^* such that

- (a) for all $\pi \in \Pi$, $D_{\text{KL}}(f^*[\pi] \parallel \hat{f}[\pi]) \leq \left(\frac{\ln \ln n}{\ln n}\right)^{c_4}$,
- (b) $\left| R_{\hat{f}}(\pi) - R_{f^*}(\pi) \right| \lesssim 2HR_{\max} \left(\frac{\ln \ln n}{\ln n}\right)^{c_4}$, for all $\pi \in \Pi$,
- (c) $\pi^*(\hat{f}) = \pi^*(f^*)$.

For tabular reinforcement learning with general reward distribution that satisfies Condition 4, there exists $n_0 > 0$ such that when $n > n_0$, $\Pr(\mathcal{E}_{\text{init}}) \geq 1 - 1/\ln n$. In addition, the regret of Step 1 is upper bounded by $\mathcal{O}\left(\frac{\ln n}{\ln \ln n}\right)$.

Proof. In the following, we prove the three items separately. Recall that Condition 1 holds by Lemma D.1 and item (c) of Condition 4.

Items (a) and (c): Proofs of (a) and (c) are the same as that of Lemma C.1 if we replace Lemma F.1 by Lemma F.2.

Item (b): Recall that an observation is a sequence of state-action-reward tuples: $o = (s_1, a_1, r_1, \dots, s_H, a_H, r_H)$. Therefore, for all $\pi \in \Pi$ we have

$$\left| R_{\hat{f}}(\pi) - R_{f^*}(\pi) \right| = \left| \mathbb{E}_{o \sim \hat{f}[\pi]} \left[\sum_{h=1}^H r_h \right] - \mathbb{E}_{o \sim f^*[\pi]} \left[\sum_{h=1}^H r_h \right] \right| \quad (204)$$

$$\leq \left| \mathbb{E}_{o^{\text{sa}} \sim \hat{f}[\pi]} \left[\sum_{h=1}^H \hat{\mu}(s_h, a_h) \right] - \mathbb{E}_{o^{\text{sa}} \sim f^*[\pi]} \left[\sum_{h=1}^H \mu^*(s_h, a_h) \right] \right|, \quad (205)$$

where $\hat{\mu}$ and μ^* are the mean reward of instance \hat{f}, f^* respectively. Since $\hat{\mu} \in [0, 1]$, we have

$$\left| \mathbb{E}_{o^{\text{sa}} \sim \hat{f}[\pi]} \left[\sum_{h=1}^H \hat{\mu}(s_h, a_h) \right] - \mathbb{E}_{o^{\text{sa}} \sim f^*[\pi]} \left[\sum_{h=1}^H \mu^*(s_h, a_h) \right] \right| \quad (206)$$

$$\leq \left| \mathbb{E}_{o^{\text{sa}} \sim \hat{f}[\pi]} \left[\sum_{h=1}^H \hat{\mu}(s_h, a_h) \right] - \mathbb{E}_{o^{\text{sa}} \sim f^*[\pi]} \left[\sum_{h=1}^H \hat{\mu}(s_h, a_h) \right] \right| \quad (207)$$

$$+ \left| \mathbb{E}_{o^{\text{sa}} \sim f^*[\pi]} \left[\sum_{h=1}^H \hat{\mu}(s_h, a_h) \right] - \mathbb{E}_{o^{\text{sa}} \sim f^*[\pi]} \left[\sum_{h=1}^H \mu^*(s_h, a_h) \right] \right| \quad (208)$$

$$\leq HD_{\text{TV}}(f^*[\pi] \parallel \hat{f}[\pi]) + \mathbb{E}_{o^{\text{sa}} \sim f^*[\pi]} \left[\sum_{h=1}^H |\hat{\mu}(s_h, a_h) - \mu^*(s_h, a_h)| \right]. \quad (209)$$

In the following we upper bound the second term of Eq. (209). Let $f_r^*, \hat{f}_r : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathbb{R})$ be the reward distributions of instance f^* and \hat{f} respectively, and $f_r^*[s, a](\cdot), \hat{f}_r[s, a](\cdot)$ the densities of the reward given state-action pair (s, a) . By item (b) of Condition 4 and Cauchy-Schwarz we get

$$\mathbb{E}_{o^{\text{sa}} \sim f^*[\pi]} \left[\sum_{h=1}^H |\hat{\mu}(s_h, a_h) - \mu^*(s_h, a_h)| \right] \quad (210)$$

$$\lesssim \mathbb{E}_{o^{\text{sa}} \sim f^*[\pi]} \left[\sum_{h=1}^H D_{\text{KL}}(f_r^*[s_h, a_h] \parallel \hat{f}_r[s_h, a_h])^{c_8} \right] \quad (211)$$

$$\leq \mathbb{E}_{o^{\text{sa}} \sim f^*[\pi]} \left[H^{1-c_8} \left(\sum_{h=1}^H D_{\text{KL}}(f_r^*[s_h, a_h] \parallel \hat{f}_r[s_h, a_h]) \right)^{c_8} \right] \quad (212)$$

$$\leq H^{1-c_8} \mathbb{E}_{o^{\text{sa}} \sim f^*[\pi]} \left[\sum_{h=1}^H D_{\text{KL}}(f_r^*[s_h, a_h] \parallel \hat{f}_r[s_h, a_h]) \right]^{c_8}. \quad (213)$$

Recall that f_p^*, \hat{f}_p denotes the transition function of instances f^*, \hat{f} respectively. By the chain rule of KL divergence,

$$D_{\text{KL}}(f^*[\pi] \parallel \hat{f}[\pi]) \quad (214)$$

$$= \mathbb{E}_{o^{\text{sa}} \sim f^*[\pi]} \left[\sum_{h=1}^H (D_{\text{KL}}(f_p^*[s_h, a_h] \|\hat{f}_p[s_h, a_h]) + D_{\text{KL}}(f_r^*[s_h, a_h] \|\hat{f}_r[s_h, a_h])) \right] \quad (215)$$

$$\geq \mathbb{E}_{o^{\text{sa}} \sim f^*[\pi]} \left[\sum_{h=1}^H D_{\text{KL}}(f_r^*[s_h, a_h] \|\hat{f}_r[s_h, a_h]) \right]. \quad (216)$$

By item (a) of this lemma, $D_{\text{KL}}(f^*[\pi] \|\hat{f}[\pi]) \rightarrow 0$ as $n \rightarrow \infty$. Therefore for large enough n we get

$$\left| R_{\hat{f}}(\pi) - R_{f^*}(\pi) \right| \lesssim HD_{\text{TV}}(f^*[\pi] \|\hat{f}[\pi]) + H^{1-c_s} D_{\text{KL}}(f^*[\pi] \|\hat{f}[\pi])^{c_s} \leq 2HD_{\text{KL}}(f^*[\pi] \|\hat{f}[\pi])^{c_s}. \quad (217)$$

Combining with item (a) of this lemma, we prove the desired result. \square

Lemma E.3. *For tabular reinforcement learning with general reward distribution that satisfies Condition 4, there exists $n_0 > 0$ such that when $n > n_0$, the following holds.*

- (a) *Conditioned on the event $\pi^*(\hat{f}) \neq \pi^*(f^*)$, $\Pr(\mathcal{E}_{\text{acc}}) \leq 1/n$.*
- (b) *Conditioned on the event $\mathcal{E}_{\text{init}}$, $\Pr(\mathcal{E}_{\text{acc}}) \geq 1 - 1/\ln n$.*
- (c) *The expected regret of Step 2 is always upper bounded by $\mathcal{O}(\ln n \ln \ln n)$.*
- (d) *Conditioned on the event $\mathcal{E}_{\text{init}}$, the expected regret of Step 2 is upper bounded by*

$$\left(\mathcal{C}(f^*, (\ln \ln n)^{1/4}/2) + o(1) \right) \ln n.$$

Proof. We prove the four items above separately.

Items (a), (c), and (d): Proofs of (a), (c), and (d) are the same as that of Lemma C.2.

Item (b): Let $\epsilon = 1/\ln \ln n$, $\alpha = \frac{\ln n}{m}$ and $\delta = 1/(2 \ln n)$. We prove this statement by invoking Lemma F.2 with parameters $(\alpha + 5\epsilon, \epsilon, w)$. Following Lemma F.2, let $\gamma = \frac{1}{m} \min_{\pi} \sum_{i=1}^m \mathbb{I}[\pi_i = \pi]$ and $\lambda = \lambda_0((\alpha + 5\epsilon)/\gamma, \epsilon, f^*)$ be the value that satisfies Condition 1. Let

$$\epsilon_0 = \exp(-(\alpha + 5\epsilon)/\gamma)(\epsilon\lambda)^{1/\lambda}.$$

First of all, we prove that the condition for Lemma F.2 holds. That is,

$$m \gtrsim \frac{\text{poly}(|\mathcal{S}||\mathcal{A}|H)}{\lambda\epsilon} ((\ln(m/\epsilon_0)) + \ln(1/\delta)). \quad (218)$$

Recall that $m = \sum_x [\bar{w}_\pi \ln n]$. As a result, $m \geq \frac{|\Pi| \ln n}{(\ln \ln n)^{1/4}}$. Now consider the RHS of Eq. (218).

By the definition of \bar{w}_π we get $\alpha \leq \frac{2}{|\Pi|} (\ln \ln n)^{1/4}$ and $\gamma^{-1} \leq 2|\Pi| (\ln \ln n)^{1/2}$. It follows from Condition 1 that $\lambda \geq \text{poly}(1/\ln \ln n)$. By the definition of ϵ_0 and Condition 2 we get

$$\ln(m/\epsilon_0) \lesssim \text{poly}(\ln \ln n). \quad (219)$$

Consequently, the RHS of Eq. (218) is at most $\text{poly}(\ln \ln n)$, and Eq. (218) holds when n is sufficiently large. By Lemma F.2 we get, with probability at least $1 - 1/(2 \ln n)$,

$$\sum_{i=1}^m \ln \frac{f^*[\pi_i](o_i)}{g[\pi_i](o_i)} \geq (\alpha + \epsilon)m, \quad \forall g \in \mathcal{F}(w, f^*, \alpha + 5\epsilon), \quad (220)$$

where $\mathcal{F}(w, f^*, \alpha + 5\epsilon) = \{g \in \mathcal{F} : D_{\text{KL}}^w(f^* \| g) \geq \alpha + 5\epsilon\}$.

In the following, we prove that Eq. (220) implies $\mathcal{E}_{\text{acc}}^{\hat{f}}$. Recall that $\mathcal{E}_{\text{acc}}^{\hat{f}}$ is the event defined as follows:

$$\mathcal{E}_{\text{acc}}^{\hat{f}} = \mathbb{I} \left[\forall g \in \Lambda(\hat{f}), \sum_{i=1}^m \ln \frac{\hat{f}[\pi_i](o_i)}{g[\pi_i](o_i)} \geq \ln n \right]. \quad (221)$$

Next, we apply Lemma G.4 to any $g \in \Lambda(\hat{f})$. To verify its condition, we have $D_{\text{TV}}(f^*[\pi] \|\hat{f}[\pi]) \leq D_{\text{KL}}(f^*[\pi] \|\hat{f}[\pi])^{1/2} \lesssim \left(\frac{\ln \ln n}{\ln n}\right)^{c_4/2}$ for all $\pi \in \Pi$ by item (a) of Lemma E.2. Therefore when n is large enough,

$$D_{\text{KL}}^w(f^* \| g) \geq \frac{\ln n}{m} + 5\epsilon = \alpha + 5\epsilon. \quad (222)$$

Then $g \in \mathcal{F}(w, f^*, \alpha + 5\epsilon)$. It follows from Eq. (220) that

$$\sum_{i=1}^m \ln \frac{f^*[\pi_i](o_i)}{g[\pi_i](o_i)} \geq (\alpha + \epsilon)m, \quad \forall g \in \Lambda(\hat{f}). \quad (223)$$

By Lemma E.4, with probability at least $1 - 1/(2 \ln n)$ we have

$$\sum_{i=1}^m \ln \frac{\hat{f}[\pi_i](o_i)}{f^*[\pi_i](o_i)} \geq -m \left(\frac{\ln \ln n}{\ln n}\right)^{c_4 c_6} \text{poly}(\ln \ln n) \iota(f^*). \quad (224)$$

For large enough n , $\left(\frac{\ln \ln n}{\ln n}\right)^{c_4 c_6} \text{poly}(\ln \ln n) \iota(f^*) \leq \frac{1}{\ln \ln n} = \epsilon$. As a result, combining Eq. (223) and Eq. (224), with probability $1 - 1/\ln n$ we have

$$\forall g \in \Lambda(\hat{f}), \quad \sum_{i=1}^m \ln \frac{\hat{f}[\pi_i](o_i)}{g[\pi_i](o_i)} \geq \alpha m = \ln n. \quad (225)$$

□

The following lemma is used in the proof of Lemma E.3.

Lemma E.4. *Let f^*, \hat{f} be any fixed tabular RL instances with reward distribution that satisfies Condition 4. For any sequence of policies $\{\pi_i\}_{i=1}^m$, let $o_i \sim f^*[\pi_i](\cdot), \forall i \in [m]$ be a sequence of random observations drawn from f^* . Then there exists constant $c_6 > 0$ and $\iota(f^*)$ that only depends on f^* such that, for any $\delta > 0$, with probability at least $1 - \delta$,*

$$\frac{1}{m} \sum_{i=1}^m \ln \frac{\hat{f}[\pi_i](o_i)}{f^*[\pi_i](o_i)} > - \left(\max_{\pi} D_{\text{KL}}(f^*[\pi] \|\hat{f}[\pi]) \right)^{c_6} \iota(f^*) \text{polylog}(mH/\delta) \quad (226)$$

assuming $\max_{\pi} D_{\text{KL}}(f^*[\pi] \|\hat{f}[\pi])^{1/2} \leq \mu_{\min}(f^*)/2$.

Proof. Recall that for reinforcement learning, an observation $o = (s_1, a_1, r_1, \dots, s_H, a_H, r_H)$ represents a trajectory, and $o^{\text{sa}} = (s_1, a_1, s_2, a_2, \dots, s_H, a_H)$ denotes the state-action pairs in the trajectory. By algebraic manipulation, we get

$$\sum_{i=1}^m \ln \frac{\hat{f}[\pi_i](o_i)}{f^*[\pi_i](o_i)} = \sum_{i=1}^m \ln \frac{\hat{f}[\pi_i](o_i^{\text{sa}})}{f^*[\pi_i](o_i^{\text{sa}})} + \sum_{i=1}^m \ln \frac{\hat{f}[\pi_i](o_i^r | o_i^{\text{sa}})}{f^*[\pi_i](o_i^r | o_i^{\text{sa}})}. \quad (227)$$

Recall that for the instance f^* , we have

$$\mu_{\min}(f^*) = \min_{\pi} \min_{o^{\text{sa}} \in \text{supp}(f^*[\pi](\cdot))} f^*[\pi](o^{\text{sa}}). \quad (228)$$

Since both o^{sa} and π are finite for tabular RL, we get $\mu_{\min}(f^*) > 0$. On the one hand, for any $o^{\text{sa}} \in \text{supp}(f^*[\pi])$, by Pinsker's inequality we get

$$\left| f^*[\pi](o^{\text{sa}}) - \hat{f}[\pi](o^{\text{sa}}) \right| \leq D_{\text{TV}}(f^*[\pi] \|\hat{f}[\pi]) \leq D_{\text{KL}}(f^*[\pi] \|\hat{f}[\pi])^{1/2}. \quad (229)$$

As a result

$$\left| \ln \frac{\hat{f}[\pi](o^{\text{sa}})}{f^*[\pi](o^{\text{sa}})} \right| \leq \left| \ln \left(1 + \frac{\hat{f}[\pi](o^{\text{sa}}) - f^*[\pi](o^{\text{sa}})}{f^*[\pi](o^{\text{sa}})} \right) \right| \leq \left| \ln \left(1 + \frac{\hat{f}[\pi](o^{\text{sa}}) - f^*[\pi](o^{\text{sa}})}{\mu_{\min}(f^*)} \right) \right|.$$

When $\max_{\pi} D_{\text{KL}}(f^*[\pi] \|\hat{f}[\pi])^{1/2} \leq \mu_{\min}(f^*)/2$, applying the basic inequality $|\ln(1+x)| \leq 2x, \forall |x| \leq 1/2$ we get for all $\pi \in \Pi$ and $o^{\text{sa}} \in \text{supp}(f^*[\pi](\cdot))$,

$$\left| \ln \left(1 + \frac{\hat{f}[\pi](o^{\text{sa}}) - f^*[\pi](o^{\text{sa}})}{\mu_{\min}(f^*)} \right) \right| \leq \frac{2}{\mu_{\min}(f^*)} \left| f^*[\pi](o^{\text{sa}}) - \hat{f}[\pi](o^{\text{sa}}) \right| \quad (230)$$

$$\leq \frac{2}{\mu_{\min}(f^*)} \max_{\pi} D_{\text{KL}}(f^*[\pi] \|\hat{f}[\pi])^{1/2}. \quad (231)$$

On the other hand, let $f^*_r[s, a]$ and $\hat{f}_r[s, a]$ be the reward distribution of instance f^* and \hat{f} given state-action pair (s, a) respectively. Then

$$\sum_{i=1}^m \ln \frac{\hat{f}[\pi_i](o_i^r | o_i^{\text{sa}})}{f^*[\pi_i](o_i^r | o_i^{\text{sa}})} = \sum_{i=1}^m \sum_{h=1}^H \ln \frac{\hat{f}_r[s_{i,h}, a_{i,h}](r_{i,h})}{f^*_r[s_{i,h}, a_{i,h}](r_{i,h})}. \quad (232)$$

For any $\pi \in \Pi$, by the chain rule of KL divergence we get

$$D_{\text{KL}}(f^*[\pi] \|\hat{f}[\pi]) \quad (233)$$

$$= \mathbb{E}_{o^{\text{sa}} \sim f^*[\pi]} \left[\sum_{h=1}^H (D_{\text{KL}}(f^*_p[s_h, a_h] \|\hat{f}_p[s_h, a_h]) + D_{\text{KL}}(f^*_r[s_h, a_h] \|\hat{f}_r[s_h, a_h])) \right] \quad (234)$$

$$\geq \mu_{\min}(f^*) D_{\text{KL}}(f^*_r[s, a] \|\hat{f}_r[s, a]) \mathbb{I}[(s, a) \in o^{\text{sa}} \text{ for some } o^{\text{sa}} \in \text{supp}(f^*[\pi])]. \quad (235)$$

Because $o_i \sim f^*[\pi_i]$, for any $i \in [m], h \in [H]$ we have $(s_{i,h}, a_{i,h}) \in o_i^{\text{sa}}$ and $o_i^{\text{sa}} \in \text{supp}(f^*[\pi_i])$. As a result, for all $i \in [m], h \in [H]$, item (a) of Condition 4 implies that with probability at least $1 - \delta/(mH)$

$$\left| \frac{\hat{f}_r[s_{i,h}, a_{i,h}](r_{i,h})}{f^*_r[s_{i,h}, a_{i,h}](r_{i,h})} \right| \leq D_{\text{TV}}(f^*_r[s_{i,h}, a_{i,h}](r_{i,h}) \|\hat{f}_r[s_{i,h}, a_{i,h}])^{c_7} \text{polylog}(mH/\delta) \quad (236)$$

$$\leq D_{\text{KL}}(f^*_r[s_{i,h}, a_{i,h}] \|\hat{f}_r[s_{i,h}, a_{i,h}])^{c_7/2} \text{polylog}(mH/\delta) \quad (237)$$

$$\leq \frac{1}{\mu_{\min}(f^*)^{c_7/2}} D_{\text{KL}}(f^*[\pi] \|\hat{f}[\pi])^{c_7/2} \text{polylog}(mH/\delta). \quad (238)$$

Let $\epsilon = \frac{1}{\mu_{\min}(f^*)^{c_7/2}} \left(\max_{\pi} D_{\text{KL}}(f^*[\pi] \|\hat{f}[\pi]) \right)^{c_7/2}$. Apply union bound and we get

$$\Pr_{o \sim f^*[\pi]} \left(\frac{1}{m} \sum_{i=1}^m \sum_{h=1}^H \left| \ln \frac{\hat{f}_r[s_{i,h}, a_{i,h}](r_{i,h})}{f^*_r[s_{i,h}, a_{i,h}](r_{i,h})} \right| \geq \epsilon H \text{polylog}(mH/\delta) \right) \leq 1 - \delta. \quad (239)$$

It follows that with probability at least $1 - \delta$,

$$\begin{aligned} & \frac{1}{m} \sum_{i=1}^m \ln \frac{\hat{f}[\pi_i](o_i)}{f^*[\pi_i](o_i)} \\ & \geq - \frac{2}{\mu_{\min}(f^*)} \max_{\pi} D_{\text{KL}}(f^*[\pi] \|\hat{f}[\pi])^{1/2} - H \text{polylog}(mH/\delta) \frac{1}{\mu_{\min}(f^*)^{c_7/2}} \left(\max_{\pi} D_{\text{KL}}(f^*[\pi] \|\hat{f}[\pi]) \right)^{c_7/2} \\ & \geq - \frac{2H}{\mu_{\min}(f^*)} \text{polylog}(mH/\delta) \left(\max_{\pi} D_{\text{KL}}(f^*[\pi] \|\hat{f}[\pi]) \right)^{c_7/2}. \end{aligned}$$

Therefore, Eq. (226) is satisfied by setting

$$\iota(f^*) = \frac{2H}{\mu_{\min}(f^*)}, \quad c_6 = c_7/2. \quad (240)$$

□

E.2 PROOF OF CONDITION 4 FOR GAUSSIAN DISTRIBUTION

In this section, we prove that the reward distribution family $\mathcal{R} = \{\mathcal{N}(\mu, 1), \mu \in [0, 1]\}$ satisfies Condition 4.

Proposition E.5. *Condition 4 holds for the reward distribution family $\mathcal{R} = \{\mathcal{N}(\mu, 1), \mu \in [0, 1]\}$*

Proof. In the following we prove the four items of Condition 4 respectively. **Item (a).** For any fixed $f \in \mathcal{R}$, let μ be the mean of f . In other words, $f = \mathcal{N}(\mu, 1)$. Then we have

$$\Pr_{x \sim f} \left(|x| > \mu + 2\sqrt{\log(2/\delta)} \right) \leq \delta. \quad (241)$$

By definition, for any $g = \mathcal{N}(\mu_g, 1), g' = \mathcal{N}(\mu'_g, 1) \in \mathcal{R}$, we have

$$\ln \frac{g(x)}{g'(x)} = \frac{1}{2} \left((x - \mu_g)^2 - (x - \mu'_g)^2 \right) = (\mu'_g - \mu_g)(x - (\mu'_g + \mu_g)/2).$$

Therefore when $\mu_g, \mu'_g \in [0, 1]$ we get

$$\left| \ln \frac{g(x)}{g'(x)} \right| \leq |\mu_g - \mu'_g| |x + 1|. \quad (242)$$

As a result, with probability at least $1 - \delta$ we have

$$\forall g, g' \in \mathcal{R}, \quad \left| \ln \frac{g(x)}{g'(x)} \right| \leq |\mu_g - \mu'_g| \left(1 + 2\sqrt{\log(2/\delta)} \right). \quad (243)$$

In addition, when $\mu_g, \mu'_g \in [0, 1]$ we have $D_{\text{TV}}(g||g') \geq \frac{1}{10}|\mu_g - \mu'_g|$. As a result, item (a) holds with $c_7 = 1$.

Item (b). Recall that for Gaussian distribution with unit variance, when $\mu_g, \mu'_g \in [0, 1]$ we have $D_{\text{TV}}(g||g') \geq \frac{1}{10}|\mu_g - \mu'_g|$. Therefore item (b) holds with $c_8 = 1$.

Item (c). Recall that for any $g = \mathcal{N}(\mu_g, 1), g' = \mathcal{N}(\mu'_g, 1) \in \mathcal{R}$, $\ln \frac{g(x)}{g'(x)} = \frac{1}{2} \left((x - \mu_g)^2 - (x - \mu'_g)^2 \right)$. Therefore when $\mu_g, \mu'_g \in [0, 1]$ we get

$$\mathbb{E}_{x \sim g} \left[\left(\ln \frac{g(x)}{g'(x)} \right)^4 \right] = \frac{1}{2} \mathbb{E}_{x \sim g} \left[\left((x - \mu_g)^2 - (x - \mu'_g)^2 \right)^4 \right] \quad (244)$$

$$\leq \frac{(\mu'_g - \mu_g)^4}{2} \mathbb{E}_{x \sim g} \left[(x - (\mu'_g + \mu_g)/2)^4 \right] \leq 3(\mu'_g - \mu_g)^4. \quad (245)$$

Therefore, item (c) holds with $c_M = 2$.

Item (d). For $g = \mathcal{N}(\mu_g, 1), g' = \mathcal{N}(\mu'_g, 1) \in \mathcal{R}$, we have

$$D_{\text{TV}}(g||g') \leq \sqrt{D_{\text{KL}}(g||g')/2} = \frac{|\mu_g - \mu'_g|}{2}. \quad (246)$$

Therefore, we can set $\mathcal{C}(\mathcal{R}, \epsilon) = \{\mathcal{N}(k\epsilon, 1) : k \in \{-[1/\epsilon], -[1/\epsilon] + 1, \dots, [1/\epsilon]\}\}$. Then $\log |\mathcal{C}(\mathcal{R}, \epsilon)| \leq \log \lceil 2/\epsilon \rceil = \mathcal{O}(\log(1/\epsilon))$. \square

F UNIFORM CONCENTRATION

In this section, we present the uniform concentration results.

F.1 UNIFORM CONCENTRATION WITH ℓ_∞ COVERING

In this section we prove uniform concentration results with ℓ_∞ covering. For two instances g, g' , define the ℓ_∞ distance as

$$\|g - g'\|_\infty \triangleq \sup_{\pi, o} |g[\pi](o) - g'[\pi](o)|. \quad (247)$$

Let $\mathcal{N}(\mathcal{F}, \epsilon)$ be the proper covering number of \mathcal{F} w.r.t. the distance ℓ_∞ . Then we have the following lemma.

Lemma F.1. Consider any fixed $\alpha > 0, 0 < \epsilon < \alpha/2$, list of decisions $w = (\pi_1, \dots, \pi_m)$, $f \in \mathcal{F}$. Let $\gamma = \frac{1}{m} \min_{\pi \in \Pi} \sum_{i=1}^m \mathbb{I}[\pi_i = \pi]$ and let $\lambda = \lambda_0(\alpha/\gamma, \epsilon, f)$ be the value that satisfies Condition 1. Let $\mathcal{F}(w, f, \alpha) = \{g \in \mathcal{F} : D_{\text{KL}}^w(f\|g) \geq \alpha\}$ and define

$$\epsilon_0 = \frac{\exp(-\alpha/\gamma)(\epsilon\lambda)^{1/\lambda}}{3\text{Vol}}.$$

Then for any $\delta > 0$, when $m \geq \frac{1}{\lambda\epsilon} (\ln \mathcal{N}(\mathcal{F}, \epsilon_0) + \ln(1/\delta))$ we have

$$\Pr_{o_i \sim f[\pi_i], \forall i} \left(\forall g \in \mathcal{F}(w, f, \alpha), \sum_{i=1}^m \ln \frac{f[\pi_i](o_i)}{g[\pi_i](o_i)} \geq (\alpha - 4\epsilon)m \right) \geq 1 - \delta. \quad (248)$$

Proof. For any $g \in \mathcal{F}$, we define a induced distribution \hat{g} for all $\pi \in \Pi$:

$$\hat{g}[\pi](o) = \frac{g[\pi](o) + \epsilon_0}{1 + \epsilon_0 \text{Vol}}. \quad (249)$$

Because $\hat{g}[\pi](o) \geq 0$ and

$$\|\hat{g}[\pi](o)\|_1 = \int \frac{g[\pi](o) + \epsilon_0}{1 + \epsilon_0 \text{Vol}} dx = \frac{1}{1 + \epsilon_0 \text{Vol}} \int g[\pi](o) + \epsilon_0 dx = 1,$$

the induced distribution $\hat{g}[\pi]$ is a valid distribution.

Properties of the induced covering. Let $Z = \epsilon_0 \text{Vol}$. Consider the minimum ϵ_0 covering set $\mathcal{C}(w, f, \alpha)$ of $\mathcal{F}(w, f, \alpha)$. For any $g \in \mathcal{F}(w, f, \alpha)$, let $g' \in \mathcal{C}(w, f, \alpha)$ be its cover and \hat{g} the induced distribution of g' . Now, we prove that \hat{g} satisfies the following two properties:

- (a) For any $\pi \in \Pi$ and $o \in \text{supp } f[\pi]$, $\ln \frac{f[\pi](o)}{g[\pi](o)} \geq \ln \frac{f[\pi](o)}{\hat{g}[\pi](o)} - \epsilon$, and
- (b) $D_{1-\lambda}^w(f\|\hat{g}) \geq \alpha - 2\epsilon$ for $g \in \mathcal{F}(w, f, \alpha)$.

First we prove item (a). Since g' is the ϵ_0 cover of g , we get $g'[\pi](o) + \epsilon_0 \geq g[\pi](o)$ for any $\pi \in \Pi$ and $x \in \text{supp } f[\pi]$. As a result,

$$\ln \frac{f[\pi](o)}{g[\pi](o)} \geq \ln \frac{f[\pi](o)}{g'[\pi](o) + \epsilon_0} = \ln \frac{f[\pi](o)}{\hat{g}[\pi](o)} - \ln(1 + Z). \quad (250)$$

By the basic inequality $\ln(1 + x) \leq x, \forall x \in (-1, \infty)$ we get

$$\ln(1 + Z) \leq Z = \epsilon_0 \text{Vol} \leq \epsilon. \quad (251)$$

As a result, item (a) follows directly. Now we prove item (b). By algebraic manipulation,

$$D_{\text{TV}}(g[\pi]\|\hat{g}[\pi]) \leq \int |\hat{g}[\pi](o) - g[\pi](o)| do = \int \left| \frac{g'[\pi](o) + \epsilon_0}{1 + \epsilon_0 \text{Vol}} - g[\pi](o) \right| do \quad (252)$$

$$\leq \int \left| \frac{g'[\pi](o) + \epsilon_0}{1 + \epsilon_0 \text{Vol}} - g'[\pi](o) - \epsilon_0 \right| do + \int |g'[\pi](o) + \epsilon_0 - g[\pi](o)| do \quad (253)$$

$$\leq \frac{1}{1 + \epsilon_0 \text{Vol}} \int |\epsilon_0 \text{Vol} (g'[\pi](o) + \epsilon_0)| do + 2\epsilon_0 \text{Vol} \quad (254)$$

$$\leq \epsilon_0 \text{Vol} + 2\epsilon_0 \text{Vol} = 3\epsilon_0 \text{Vol} = \exp(-\alpha/\gamma)(\epsilon\lambda)^{1/\lambda}. \quad (255)$$

Applying Lemma G.5, we get item (b).

Uniform concentration via covering. We define the induced covering set $\hat{\mathcal{C}}(w, f, \alpha) = \{\hat{g} : g' \in \mathcal{C}(w, f, \alpha)\}$. Then $|\hat{\mathcal{C}}(w, f, \alpha)| \leq |\mathcal{C}(w, f, \alpha)| \leq \mathcal{N}(\mathcal{F}, \epsilon_0)$. Applying Lemma A.2 and union bound we get, with probability at least $1 - \delta$,

$$\sum_{i=1}^m \ln \frac{f[\pi_i](o_i)}{\hat{g}[\pi_i](o_i)} \geq m(D_{1-\lambda}^w(f\|\hat{g}) - \epsilon), \quad \forall \hat{g} \in \hat{\mathcal{C}}(w, f, \alpha). \quad (256)$$

By item (a) of the covering property,

$$\sum_{i=1}^m \ln \frac{f[\pi_i](o_i)}{g[\pi_i](o_i)} \geq \sum_{i=1}^m \ln \frac{f[\pi_i](o_i)}{\hat{g}[\pi_i](o_i)} - m\epsilon. \quad (257)$$

By item (b) of the covering property,

$$D_{1-\lambda}^w(f \parallel \hat{g}) \geq \alpha - 2\epsilon \quad (258)$$

Combining Eqs. (256), (257), (258), with probability at least $1 - \delta$,

$$\sum_{i=1}^m \ln \frac{f[\pi_i](o_i)}{g[\pi_i](o_i)} \geq m(\alpha - 4\epsilon). \quad (259)$$

□

F.2 UNIFORM CONCENTRATION FOR RL WITH GENERAL REWARD DISTRIBUTION

In the following, we present a uniform concentration lemma for tabular RL with general reward.

Lemma F.2. *Let \mathcal{F} be the instance family representing tabular RL with general reward distribution that satisfies Condition 4. For any fixed $\alpha > 0$, $0 < \epsilon < \alpha/2$, list of decisions $w = (\pi_1, \dots, \pi_m)$, $f \in \mathcal{F}$, let $\gamma = \frac{1}{m} \min_{\pi: \pi \in w} \sum_{i=1}^m \mathbb{I}[\pi_i = \pi]$ and $\lambda = \lambda_0(\alpha/\gamma, \epsilon, f)$ be the value that satisfies Condition 1. Let $\mathcal{F}(w, f, \alpha) = \{g \in \mathcal{F} : D_{\text{KL}}^w(f \parallel g) \geq \alpha\}$ and define*

$$\epsilon_0 = \exp(-\alpha/\gamma)(\epsilon\lambda)^{1/\lambda}.$$

Then for any $\delta > 0$, when $m \gtrsim \frac{\text{poly}(|\mathcal{S}||\mathcal{A}|H)}{\lambda\epsilon} ((\ln(m/\epsilon_0)) + \ln(1/\delta))$ we have

$$\Pr_{o_i \sim f[\pi_i], \forall i} \left(\forall g \in \mathcal{F}(w, f, \alpha), \sum_{i=1}^m \ln \frac{f[\pi_i](o_i)}{g[\pi_i](o_i)} \geq (\alpha - 4\epsilon)m \right) \geq 1 - \delta. \quad (260)$$

Proof. Recall that a tabular RL instance $g \in \mathcal{F}$, we use $g_r[s, a]$ to denote its reward distribution given state-action pair (s, a) , and $g_p[s, a]$ its transition. We prove this lemma using the covering argument.

The covering set. Define

$$\epsilon_1 = \min \left\{ \frac{\epsilon}{2H|\mathcal{S}|}, \frac{\epsilon}{2H \text{polylog}(2mH/\delta)}, \frac{\exp(-\alpha/\gamma)(\epsilon\lambda)^{1/\lambda}}{4|\mathcal{S}|^2|\mathcal{A}|} \right\}.$$

Let $\mathcal{C} \subseteq \mathcal{F}$ be the minimum covering of \mathcal{F} such that for any $g \in \mathcal{F}$ (parameterized by p, μ), there exists $g' \in \mathcal{C}$ (parameterized by p', μ') such that

$$\sup_{s, a, s'} |g_p[s, a](s') - g'_p[s, a](s')| \leq \epsilon_1, \quad \sup_{s, a} D_{\text{TV}}(g_r[s, a] \parallel g'_r[s, a]) \leq \epsilon_1^{1/\epsilon_7}. \quad (261)$$

By item (d) of Condition 4 and a standard covering argument for discrete distributions, we have $\ln |\mathcal{C}| \lesssim |\mathcal{S}||\mathcal{A}| \ln(1/\epsilon_1)$. For any $g' \in \mathcal{C}$, we consider the induced instance \hat{g} defined by

$$\hat{g}_p[s, a](s') = \frac{g'_p[s, a](s') + \epsilon_1}{1 + \epsilon_1|\mathcal{S}|}, \quad \hat{g}_r[s, a](\cdot) = g'_r[s, a](\cdot). \quad (262)$$

In the following, we prove that

(a) with probability at least $1 - \delta/2$,

$$\Pr_{o_i \sim f[\pi_i], \forall i} \left(\forall g \in \mathcal{F}, \sum_{i=1}^m \ln \frac{\hat{g}[\pi_i](o_i)}{g[\pi_i](o_i)} \geq -m\epsilon \right) \geq 1 - \delta/2. \quad (263)$$

(b) $D_{1-\lambda}^w(f \parallel \hat{g}) \geq \alpha - 2\epsilon$ for $g \in \mathcal{F}(w, f, \alpha)$.

To prove (a), recall that $o_i = \{(s_{i,h}, a_{i,h}, r_{i,h})\}_{h=1}^H$ represents a trajectory. Consequently,

$$\sum_{i=1}^m \ln \frac{\hat{g}[\pi_i](o_i)}{g[\pi_i](o_i)} \quad (264)$$

$$= \sum_{i=1}^m \sum_{h=1}^H \ln \frac{\hat{g}_p[s_{i,h}, a_{i,h}](s_{i,h+1})}{g_p[s_{i,h}, a_{i,h}](s_{i,h+1})} + \sum_{i=1}^m \sum_{h=1}^H \ln \frac{\hat{g}_r[s_{i,h}, a_{i,h}](r_{i,h})}{g_r[s_{i,h}, a_{i,h}](r_{i,h})} \quad (265)$$

$$\geq - \sum_{i=1}^m \sum_{h=1}^H \ln(1 + |\mathcal{S}| \epsilon_1) + \sum_{i=1}^m \sum_{h=1}^H \ln \frac{\hat{g}_r[s_{i,h}, a_{i,h}](r_{i,h})}{g_r[s_{i,h}, a_{i,h}](r_{i,h})}. \quad (266)$$

By item (a) of Condition 4 and union bound we get with probability at least $1 - \delta/2$

$$\sum_{i=1}^m \sum_{h=1}^H \ln \frac{\hat{g}_r[s_{i,h}, a_{i,h}](r_{i,h})}{g_r[s_{i,h}, a_{i,h}](r_{i,h})} \geq -mH\epsilon_1 \text{polylog}(mH/\delta) \geq -m\epsilon/2. \quad (267)$$

Therefore,

$$\sum_{i=1}^m \ln \frac{\hat{g}[\pi_i](o_i)}{g[\pi_i](o_i)} \geq - \sum_{i=1}^m \sum_{h=1}^H |\mathcal{S}| \epsilon_1 - m\epsilon/2 \geq -m\epsilon, \quad (268)$$

which proves (a). For (b), by algebraic manipulation we have

$$D_{\text{TV}}(g[\pi] \|\hat{g}[\pi]) \leq \sum_{s,a} (D_{\text{TV}}(p[s,a] \|\hat{g}_p[s,a]) + D_{\text{TV}}(r[s,a] \|\hat{g}_r[s,a])) \quad (269)$$

$$\leq \sum_{s,a,s'} \left| \frac{p'[s,a](s') + \epsilon_1}{1 + \epsilon_1 |\mathcal{S}|} - p[s,a](s') \right| + |\mathcal{S}| |\mathcal{A}| \epsilon_1 \quad (270)$$

$$\leq \sum_{s,a,s'} \left| \frac{p'[s,a](s') + \epsilon_1}{1 + \epsilon_1 |\mathcal{S}|} - p'[s,a](s') - \epsilon_1 \right| + \sum_{s,a,s'} |p'[s,a](s') + \epsilon_1 - p[s,a](s')| + |\mathcal{S}| |\mathcal{A}| \epsilon_1 \quad (271)$$

$$\leq 4|\mathcal{S}|^2 |\mathcal{A}| \epsilon_1 \leq \exp(-\alpha/\gamma) (\epsilon \lambda)^{1/\lambda}. \quad (272)$$

Applying Lemma G.5, we get item (b).

Uniform concentration via covering. Now we apply Lemma A.2 and union bound. When

$$m \gtrsim \frac{|\mathcal{S}| |\mathcal{A}| \ln(1/\epsilon_1)}{\lambda \epsilon} \gtrsim \frac{\ln(|\mathcal{C}|/\delta)}{\lambda \epsilon},$$

with probability at least $1 - \delta/2$ we get

$$\sum_{i=1}^m \ln \frac{f[\pi_i](o_i)}{\hat{g}[\pi_i](o_i)} \geq m(D_{1-\lambda}^w(f \|\hat{g}) - \epsilon), \quad \forall \hat{g} \in \mathcal{C}. \quad (273)$$

By item (a) of the covering property, with probability at least $1 - \delta/2$,

$$\sum_{i=1}^m \ln \frac{f[\pi_i](o_i)}{g[\pi_i](o_i)} \geq \sum_{i=1}^m \ln \frac{f[\pi_i](o_i)}{\hat{g}[\pi_i](o_i)} - m\epsilon. \quad (274)$$

By item (b) of the covering property,

$$D_{1-\lambda}^w(f \|\hat{g}) \geq \alpha - 2\epsilon \quad (275)$$

Combining Eqs. (273), (274), (275), with probability at least $1 - \delta$,

$$\sum_{i=1}^m \ln \frac{f[\pi_i](o_i)}{g[\pi_i](o_i)} \geq m(\alpha - 4\epsilon). \quad (276)$$

□

G HELPER LEMMAS

The following lemma shows that the Rényi divergence of a marginal distribution is smaller than that of the original distribution.

Lemma G.1 (Theorem 1 of Van Erven & Harremoës (2014)). *Consider any distributions f, g over variables x, y . Let $f(x), g(x)$ denote their marginal distribution over x respectively. For any $\alpha \in (0, 1)$, we have*

$$D_\alpha(f(x)\|g(x)) \leq D_\alpha(f(x, y)\|g(x, y)). \quad (277)$$

Proof. When $\alpha \in (0, 1)$, we have $\frac{1}{\alpha-1} < 0$. As a result, we only need to prove that

$$\int f(x, y)^\alpha g(x, y)^{1-\alpha} dx dy \leq \int f(x)^\alpha g(x)^{1-\alpha} dx. \quad (278)$$

We prove this by Hölder’s inequality. In particular,

$$\int f(x, y)^\alpha g(x, y)^{1-\alpha} dx dy \quad (279)$$

$$= \int f(x)^\alpha g(x)^{1-\alpha} \int f(y|x)^\alpha g(y|x)^{1-\alpha} dy dx \quad (280)$$

$$\leq \int f(x)^\alpha g(x)^{1-\alpha} \left(\int f(y|x) dy \right)^\alpha \left(\int g(y|x) dy \right)^{1-\alpha} dx \quad (281)$$

$$\leq \int f(x)^\alpha g(x)^{1-\alpha} dx. \quad (282)$$

□

Intuitively, the following two lemmas upper bound the difference of Rényi divergence by the TV distance.

Lemma G.2. *For and fixed $\lambda \in (0, 1), \alpha > 0, \epsilon > 0$ and distribution f , consider two distributions g, \hat{g} such that $D_{TV}(g\|\hat{g}) \leq \exp(-\alpha)(\lambda\epsilon)^{1/\lambda}$. Then we have*

$$D_{1-\lambda}(f\|\hat{g}) \geq \min\{\alpha, D_{1-\lambda}(f\|g)\} - \epsilon. \quad (283)$$

Proof. Let $\kappa = \min\{\alpha, D_{1-\lambda}(f\|g)\}$. We start by proving

$$\left| \int f(x)^{1-\lambda} g(x)^\lambda dx - \int f(x)^{1-\lambda} \hat{g}(x)^\lambda dx \right| \leq (\exp(\lambda\epsilon) - 1) \exp(-\lambda\kappa). \quad (284)$$

By Hölder’s inequality we get

$$\left| \int f(x)^{1-\lambda} g(x)^\lambda dx - \int f(x)^{1-\lambda} \hat{g}(x)^\lambda dx \right| \quad (285)$$

$$\leq \left(\int |g(x) - \hat{g}(x)| dx \right)^\lambda \left(\int f(x) dx \right)^{1-\lambda} \quad (286)$$

$$\leq \exp(-\lambda\alpha) \lambda\epsilon \quad (287)$$

$$\leq \exp(-\lambda\alpha) (\exp(\lambda\epsilon) - 1) \quad (288)$$

$$\leq \exp(-\lambda\kappa) (\exp(\lambda\epsilon) - 1), \quad (289)$$

where the Eq. (288) follows from the basic inequality $1 + x \leq \exp(x)$ for $x > 0$.

By the definition of Rényi divergence,

$$\int f(x)^{1-\lambda} g(x)^\lambda dx = \exp(-\lambda D_{1-\lambda}(f\|g)) \leq \exp(-\lambda\kappa). \quad (290)$$

Combining Eq. (289) and Eq. (290) we get,

$$\int f(x)^{1-\lambda} \hat{g}(x)^\lambda dx \leq \exp(-\lambda(\kappa - \epsilon)). \quad (291)$$

It follows that

$$D_{1-\lambda}(f\|\hat{g}) = -\frac{1}{\lambda} \ln \int f(x)^{1-\lambda} \hat{g}(x)^\lambda dx \geq \kappa - \epsilon. \quad (292)$$

□

Lemma G.3. For and fixed $\lambda \in (0, 1/2)$, $\alpha > 0$, $\epsilon > 0$ and distribution g , consider two distributions f, \hat{f} such that $D_{\text{TV}}(f\|\hat{f}) \leq \exp\left(-\frac{\lambda}{1-\lambda}\alpha\right) (\lambda\epsilon)^{1/(1-\lambda)}$. Then we have

$$D_{1-\lambda}(\hat{f}\|g) \geq \min\{\alpha, D_{1-\lambda}(f\|g)\} - \epsilon. \quad (293)$$

Proof. We use a similar proof as Lemma G.2. Let $\kappa = \min\{\alpha, D_{1-\lambda}(f\|g)\}$. Then we have

$$\left| \int f(x)^{1-\lambda} g(x)^\lambda dx - \int \hat{f}(x)^{1-\lambda} g(x)^\lambda dx \right| \quad (294)$$

$$\leq \left(\int |f(x) - \hat{f}(x)| dx \right)^{1-\lambda} \left(\int g(x) dx \right)^\lambda \quad (295)$$

$$\leq \exp(-\lambda\alpha) \lambda\epsilon \quad (296)$$

$$\leq \exp(-\lambda\kappa) (\exp(\lambda\epsilon) - 1), \quad (297)$$

By the definition of Rényi divergence,

$$\int f(x)^{1-\lambda} g(x)^\lambda dx \leq \exp(-\lambda D_{1-\lambda}(f\|g)) \leq \exp(-\lambda\kappa). \quad (298)$$

Combining Eq. (297) and Eq. (298) we get,

$$\int \hat{f}(x)^{1-\lambda} g(x)^\lambda dx \leq \exp(-\lambda(\kappa - \epsilon)). \quad (299)$$

It follows that

$$D_{1-\lambda}(\hat{f}\|g) = -\frac{1}{\lambda} \ln \int \hat{f}(x)^{1-\lambda} g(x)^\lambda dx \geq \kappa - \epsilon. \quad (300)$$

□

The following lemma is used to prove that when \hat{f} is close to f^* (measured by TV distance), we can use \hat{f} to approximately solve $\mathcal{C}(f^*)$ (see Lemma C.2).

Lemma G.4. For an instance $f \in \mathcal{F}$ and $n > 0$. Let $\delta = (\ln \ln n)^{-1/4}$ and $\epsilon = (\ln \ln n)^{-1}$. For any $\hat{w} \in \mathbb{R}_+^{|\Pi|}$ such that $\|\hat{w}\|_\infty \leq (\ln \ln n)^{1/4}$, define $w = \{\pi_i\}_{i=1}^m$ be a list of decisions where a decision π occurs $\lceil ((1+\delta)\hat{w}_\pi + \delta) \ln n \rceil$ times for every $\pi \in \Pi$, and $m = \sum_{\pi} \lceil ((1+\delta)\hat{w}_\pi + \delta) \ln n \rceil$. Consider two instances $f, \hat{f} \in \mathcal{F}$ such that there exists constant $c_6 > 0$ with

$$D_{\text{TV}}(f[\pi]\|\hat{f}[\pi]) \lesssim \left(\frac{\ln \ln n}{\ln n} \right)^{c_6}, \forall \pi \in \Pi.$$

Define the set $\mathcal{F}(\hat{w}, f) = \{g \in \mathcal{F} : \sum_{\pi \in \Pi} \hat{w}_\pi D_{\text{KL}}(f[\pi]\|g[\pi]) \geq 1\}$. Then for any constant $c > 0$, there exists $n_0 > 0$ such that for all $n > n_0$,

$$D_{\text{KL}}^w(\hat{f}\|g) \geq \frac{\ln n}{m} + c\epsilon, \quad \forall g \in \mathcal{F}(\hat{w}, f). \quad (301)$$

Proof of Lemma G.4. First we invoke Condition 1 with proper parameters. Define

$$\lambda = \lambda_0 \left(4(\ln \ln n)^{3/4}, \frac{1}{\ln \ln n}, f \right) \quad (302)$$

By the definition of m and the fact that $\|\hat{w}\|_\infty \leq (\ln \ln n)^{1/4}$ we get

$$\frac{|\Pi| \ln n}{(\ln \ln n)^{1/4}} \leq m \leq 2|\Pi| \ln n (\ln \ln n)^{1/4}. \quad (303)$$

Let $\alpha \triangleq \frac{\ln n}{m} + (c+2)\epsilon \lesssim \frac{2}{|\Pi|} (\ln \ln n)^{1/4}$. Consider $\gamma = \min_\pi \frac{1}{m} \lceil ((1+\delta)\hat{w}_\pi + \delta) \ln n \rceil$. By the upper bound of m we get $\gamma \geq \frac{1}{2|\Pi|(\ln \ln n)^{1/2}}$.

We invoke Condition 1 with parameters $(\alpha/\gamma, \epsilon, f)$. For large enough n we have $\alpha/\gamma \leq 4(\ln \ln n)^{3/4}$, which implies that $\lambda < \lambda_0(\alpha/\gamma, \epsilon, f)$. Then for any $\pi \in \Pi$ we get

$$D_{1-\lambda}(f[\pi] \| g[\pi]) \geq \min\{\alpha/\gamma, D_{\text{KL}}(f[\pi] \| g[\pi]) - \epsilon\}. \quad (304)$$

We claim that $D_{\text{KL}}^w(f \| g) \geq \frac{\ln n}{m} + (c+2)\epsilon, \forall g \in \mathcal{F}(\hat{w}, f)$ for large enough n . Indeed, by the definition of w we get

$$m D_{\text{KL}}^w(f \| g) \geq \sum_\pi ((1+\delta)\hat{w}_\pi + \delta) (\ln n) D_{\text{KL}}(f[\pi] \| g[\pi]) \geq (1+\delta) \ln n \geq \ln n + (c+2)m\epsilon. \quad (305)$$

Let

$$\epsilon_1 = \exp(-\alpha/\gamma)(\lambda\epsilon)^2 = \Omega\left(\exp(-(\ln \ln n)^{3/4}) \text{poly}(\ln \ln n)\right).$$

By Lemma G.3 with parameters $(\lambda, \alpha/\gamma, \epsilon)$ and the assumption that $D_{\text{TV}}(f[\pi] \| \hat{f}[\pi]) \lesssim \left(\frac{\ln \ln n}{\ln n}\right)^{c_6} = o(\epsilon_1)$ for all $\pi \in \Pi$, we get

$$D_{1-\lambda}(\hat{f}[\pi] \| g[\pi]) \geq \min\{\alpha/\gamma, D_{1-\lambda}(f[\pi] \| g[\pi])\} - \epsilon, \forall \pi \in \Pi. \quad (306)$$

Now we consider the following two cases.

Case 1: There exists $\pi \in w$ such that $D_{\text{KL}}(f[\pi] \| g[\pi]) \geq \alpha/\gamma$. In this case Eq. (304) implies that $D_{1-\lambda}(f[\pi] \| g[\pi]) \geq \alpha/\gamma - \epsilon$. Combining with Eq. (306) we have

$$D_{1-\lambda}(\hat{f}[\pi] \| g[\pi]) \geq \alpha/\gamma - 2\epsilon. \quad (307)$$

As a result,

$$D_{1-\lambda}^w(\hat{f} \| g) \geq \frac{1}{m} \sum_{i=1}^m \mathbb{I}[\pi_i = \pi] (\alpha/\gamma - 2\epsilon) \geq \alpha - 2\epsilon = \frac{\ln n}{m} + c\epsilon. \quad (308)$$

Case 2: For all $\pi \in w$, $D_{\text{KL}}(f[\pi] \| g[\pi]) \leq \alpha/\gamma$. In this case we also have $D_{1-\lambda}(f[\pi] \| g[\pi]) \leq \alpha/\gamma, \forall \pi \in w$. Therefore Eq. (306) and Eq. (304) implies that

$$D_{1-\lambda}(\hat{f}[\pi] \| g[\pi]) \geq D_{1-\lambda}(f[\pi] \| g[\pi]) - \epsilon, \quad (309)$$

$$D_{1-\lambda}(f[\pi] \| g[\pi]) \geq D_{\text{KL}}(f[\pi] \| g[\pi]) - \epsilon. \quad (310)$$

As a result,

$$D_{1-\lambda}^w(\hat{f} \| g) \geq \frac{1}{m} \sum_{i=1}^m (D_{1-\lambda}(f[\pi_i] \| g[\pi_i]) - \epsilon) \geq \frac{1}{m} \sum_{i=1}^m (D_{\text{KL}}(f[\pi_i] \| g[\pi_i]) - 2\epsilon) \quad (311)$$

$$= D_{\text{KL}}^w(f \| g) - 2\epsilon \geq \frac{\ln n}{m} + c\epsilon. \quad (312)$$

Combining the two cases together, we get $D_{1-\lambda}^w(\hat{f} \| g) \geq \frac{\ln n}{m} + c\epsilon$. \square

The following lemma is used to prove a nice property of the covering (see Lemma F.1).

Lemma G.5. Consider any $\epsilon > 0, \alpha > 0$, sequence of decisions $w = \{\pi_i\}_{i=1}^m$. Let $\gamma = \frac{1}{m} \min_{\pi: \pi \in w} \sum_{i=1}^m \mathbb{I}[\pi_i = \pi]$ and $\lambda = \lambda_0(\alpha/\gamma, \epsilon, f)$ be the value that satisfies Condition 1. For two distributions $f, \hat{f} \in \mathcal{F}$ such that

$$D_{\text{TV}}(g[\pi] \| \hat{g}[\pi]) \leq \exp(-\alpha/\gamma)(\lambda\epsilon)^{1/\lambda}, \forall \pi \in \Pi.$$

we have

$$D_{1-\lambda}^w(f \| \hat{g}) \geq \min\{D_{\text{KL}}^w(f \| g), \alpha\} - 2\epsilon. \quad (313)$$

Proof of Lemma G.5. Let $\epsilon_1 = \exp(-\alpha/\gamma)(\lambda\epsilon)^{1/\lambda}$ and $\kappa = \min\{D_{\text{KL}}^w(f\|g), \alpha\}$.

By Lemma G.2 and the fact that $D_{\text{TV}}(g[\pi]\|\hat{g}[\pi]) \leq \epsilon_1, \forall \pi \in \Pi$, for any $\pi \in \Pi$ we have

$$D_{1-\lambda}(f[\pi]\|\hat{g}[\pi]) \geq \min\{\alpha/\gamma, D_{1-\lambda}(f[\pi]\|g[\pi])\} - \epsilon. \quad (314)$$

Applying Condition 1, for any $\pi \in \Pi$ we have

$$D_{1-\lambda}(f[\pi]\|g[\pi]) \geq \min\{\alpha/\gamma, D_{\text{KL}}(f[\pi]\|g[\pi]) - \epsilon\}. \quad (315)$$

Now we consider the following two cases.

Case 1: There exists $\pi \in w$ such that $D_{\text{KL}}(f[\pi]\|g[\pi]) \geq \kappa/\gamma$. By the definition of κ we have $\kappa \leq \alpha$. In this case Eq. (315) implies that $D_{1-\lambda}(f[\pi]\|g[\pi]) \geq \kappa/\gamma - \epsilon$. Combining with Eq. (314) we have

$$D_{1-\lambda}(f[\pi]\|\hat{g}[\pi]) \geq \kappa/\gamma - 2\epsilon. \quad (316)$$

As a result,

$$D_{1-\lambda}^w(f\|\hat{g}) \geq \frac{1}{m} \sum_{i=1}^m \mathbb{I}[\pi_i = \pi] (\kappa/\gamma - 2\epsilon) \geq \kappa - 2\epsilon. \quad (317)$$

Case 2: For all $\pi \in w$, $D_{\text{KL}}(f[\pi]\|g[\pi]) \leq \kappa/\gamma$. In this case we also have $D_{1-\lambda}(f[\pi]\|g[\pi]) \leq \kappa/\gamma, \forall \pi \in w$. Therefore Eq. (314) and Eq. (315) implies that

$$D_{1-\lambda}(f[\pi]\|\hat{g}[\pi]) \geq D_{1-\lambda}(f[\pi]\|g[\pi]) - \epsilon, \quad (318)$$

$$D_{1-\lambda}(f[\pi]\|g[\pi]) \geq D_{\text{KL}}(f[\pi]\|g[\pi]) - \epsilon. \quad (319)$$

As a result,

$$D_{1-\lambda}^w(f\|\hat{g}) \geq \frac{1}{m} \sum_{i=1}^m (D_{1-\lambda}(f[\pi_i]\|g[\pi_i]) - \epsilon) \geq \frac{1}{m} \sum_{i=1}^m (D_{\text{KL}}(f[\pi_i]\|g[\pi_i]) - 2\epsilon) \quad (320)$$

$$= D_{\text{KL}}^w(f\|g) - 2\epsilon \geq \kappa - 2\epsilon. \quad (321)$$

Combining the two cases together, we get $D_{1-\lambda}^w(f\|\hat{g}) \geq \kappa - 2\epsilon$. \square

The following lemma shows that for truncated Gaussian distributions, the difference in their density function can be upper bounded by their KL divergence. We use this lemma to prove Condition 3 for tabular RL with truncated Gaussian reward.

Lemma G.6. Consider two truncated Gaussian distributions p_1, p_2 with density

$$p_i(x) = \mathbb{I}[x \in [-2, 2]] \frac{1}{Z_i} \exp\left(-\frac{(x - \mu_i)^2}{2}\right), \quad (322)$$

where Z_i is the normalization factor. Assuming $\mu_1, \mu_2 \in [-1, 1]$, we have

$$\sup_{x \in [-2, 2]} |p_1(x) - p_2(x)| \leq 36 D_{\text{KL}}(p_1\|p_2)^{1/6}. \quad (323)$$

Proof. We first prove that $|\mu_1 - \mu_2| \lesssim D_{\text{KL}}(p_1\|p_2)^{1/6}$. Then we show that $\sup_{x \in [-2, 2]} |p_1(x) - p_2(x)| \lesssim |\mu_1 - \mu_2|$.

By Pinsker's inequality,

$$D_{\text{TV}}(p_1\|p_2) \lesssim D_{\text{KL}}(p_1\|p_2)^{1/2}. \quad (324)$$

Now we prove that $|\mu_1 - \mu_2|^3 \lesssim D_{\text{TV}}(p_1\|p_2)$. W.l.o.g., we assume $Z_1 \geq Z_2 > 1/\sqrt{2\pi}$ and $\mu_1 \leq \mu_2$. Then we have, for any $x \in [\mu_1, \mu_1 + \frac{1}{4}(\mu_2 - \mu_1)]$,

$$p_1(x) - p_2(x) \geq \frac{1}{\sqrt{2\pi}} \left[\exp\left(-\frac{(\mu_2 - \mu_1)^2}{32}\right) - \exp\left(-\frac{9(\mu_2 - \mu_1)^2}{32}\right) \right] \quad (325)$$

$$\geq \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{9(\mu_2 - \mu_1)^2}{32}\right) \left(\exp\left(\frac{1}{4}(\mu_2 - \mu_1)^2\right) - 1\right) \quad (326)$$

$$\geq \frac{1}{4e^2\sqrt{2\pi}}(\mu_2 - \mu_1)^2. \quad (327)$$

As a result,

$$D_{\text{TV}}(p_1 \| p_2) \geq \frac{1}{4}(\mu_2 - \mu_1)(p_1(x) - p_2(x)) \gtrsim |\mu_2 - \mu_1|^3. \quad (328)$$

Now we prove that $\sup_{x \in [-2, 2]} |p_1(x) - p_2(x)| \lesssim |\mu_1 - \mu_2|$. By definition, for any $x \in [-2, 2]$ we have

$$|p_1(x) - p_2(x)| = \left| \frac{1}{Z_1} \exp\left(-\frac{(x - \mu_1)^2}{2}\right) - \frac{1}{Z_2} \exp\left(-\frac{(x - \mu_2)^2}{2}\right) \right| \quad (329)$$

$$\leq \left| \frac{1}{Z_1} - \frac{1}{Z_2} \right| + \frac{1}{Z_1} \left| \exp\left(-\frac{(x - \mu_1)^2}{2}\right) - \exp\left(-\frac{(x - \mu_2)^2}{2}\right) \right|, \quad (330)$$

$$\leq 4|Z_1 - Z_2| + 2 \left| \exp\left(-\frac{(x - \mu_1)^2}{2}\right) - \exp\left(-\frac{(x - \mu_2)^2}{2}\right) \right|, \quad (331)$$

where the last inequality comes from the fact that $Z_i \geq 1/2$ when $|\mu_i| \leq 1$. For the second term,

$$\left| \exp\left(-\frac{(x - \mu_1)^2}{2}\right) - \exp\left(-\frac{(x - \mu_2)^2}{2}\right) \right| \quad (332)$$

$$= \exp\left(-\frac{(x - \mu_2)^2}{2}\right) \left| \exp\left(\frac{(\mu_2 - \mu_1)(\mu_2 + \mu_1 - x)}{2}\right) - 1 \right| \quad (333)$$

$$\leq 2|\mu_2 - \mu_1|. \quad (334)$$

For the first term,

$$|Z_1 - Z_2| \leq \int_{x=-2}^2 |p_1(x) - p_2(x)| dx \leq 8|\mu_2 - \mu_1|. \quad (335)$$

As a result, we get

$$|p_1(x) - p_2(x)| \leq 36|\mu_2 - \mu_1|. \quad (336)$$

Combining Eq. (328) and Eq. (336) we prove the this lemma. \square

The following lemma can be viewed as an perturbation analysis to the covariance matrix in linear bandits setting. We use the lemma to prove that our complexity measure recovers that in (Lattimore & Szepesvari, 2017).

Lemma G.7. *For a fixed positive definite matrix $A \in \mathbb{R}^{d \times d}$ and an unit vector $x \in \mathbb{R}^d$, let $G_1 = (A + nx x^\top)^{-1}$ and $G_2 = \lim_{\rho \rightarrow \infty} (A + \rho x x^\top)^{-1}$. Then*

$$\|G_1 - G_2\|_2 \leq \frac{\sigma_{\max}(A^{-1})}{1 + n\sigma_{\min}(A^{-1})}. \quad (337)$$

Proof. By Sherman–Morrison formula we get

$$G_1 = A^{-1} - \frac{nA^{-1}xx^\top A^{-1}}{1 + nx^\top A^{-1}x}, \quad (338)$$

$$G_2 = \lim_{\rho \rightarrow \infty} A^{-1} - \frac{\rho A^{-1}xx^\top A^{-1}}{1 + \rho x^\top A^{-1}x}. \quad (339)$$

Then for any $v \in \mathbb{R}^d$ such that $\|v\|_2 = 1$, we get

$$v^\top (G_1 - G_2)v = \lim_{\rho \rightarrow \infty} \frac{\rho v^\top A^{-1}xx^\top A^{-1}v}{1 + \rho x^\top A^{-1}x} - \frac{nv^\top A^{-1}xx^\top A^{-1}v}{1 + nx^\top A^{-1}x} \quad (340)$$

$$= \lim_{\rho \rightarrow \infty} \frac{(\rho - n)v^\top A^{-1}xx^\top A^{-1}v}{(1 + nx^\top A^{-1}x)(1 + \rho x^\top A^{-1}x)} \quad (341)$$

$$= \frac{v^\top A^{-1}xx^\top A^{-1}v}{x^\top A^{-1}x(1 + nx^\top A^{-1}x)} \quad (342)$$

$$\leq \frac{v^\top A^{-1}v}{1 + nx^\top A^{-1}x} \quad (343)$$

$$\leq \frac{\sigma_{\max}(A^{-1})}{1 + n\sigma_{\min}(A^{-1})}. \quad (344)$$

□

The following lemma upper bounds the difference between KL divergence and Rényi divergence, and is used to prove Condition 1.

Lemma G.8. For any two distribution f, g and constant $\lambda \in (0, 1/2)$, we have

$$D_{\text{KL}}(f\|g) - D_{1-\lambda}(f\|g) \leq \frac{\lambda}{2} \mathbb{E}_{o \sim f} \left[\left(\ln \frac{f(o)}{g(o)} \right)^4 \right]^{1/2}. \quad (345)$$

Proof. Recall that

$$D_{1-\lambda}(f\|g) = -\frac{1}{\lambda} \ln \int f(o)^{1-\lambda} g(o)^\lambda do.$$

Define the function $h(\lambda) \triangleq \int f(o)^{1-\lambda} g(o)^\lambda do$. By basic algebra we get

$$h'(\lambda) = \int f(o)^{1-\lambda} g(o)^\lambda \ln \frac{g(o)}{f(o)} do \quad (346)$$

$$h''(\lambda) = \int f(o)^{1-\lambda} g(o)^\lambda \ln^2 \frac{g(o)}{f(o)} do. \quad (347)$$

By Taylor expansion, there exists $\xi \in (0, \lambda)$ such that

$$h(\lambda) = h(0) + \lambda h'(0) + \frac{\lambda^2}{2} h''(\xi). \quad (348)$$

By definition we have $h(0) = 1$ and $h'(0) = -D_{\text{KL}}(f\|g)$. As a result, we get

$$D_{1-\lambda}(f\|g) = -\frac{1}{\lambda} \ln h(\lambda) \quad (349)$$

$$= -\frac{1}{\lambda} \ln \left(1 - \lambda D_{\text{KL}}(f\|g) + \frac{\lambda^2}{2} \int f(o)^{1-\xi} g(o)^\xi \ln^2 \frac{g(o)}{f(o)} do \right) \quad (350)$$

$$\geq D_{\text{KL}}(f\|g) - \frac{\lambda}{2} \int f(o)^{1-\xi} g(o)^\xi \ln^2 \frac{g(o)}{f(o)} do. \quad (351)$$

By Hölder's inequality, when $\xi < \lambda < 1/2$ we get

$$\int f(o)^{1-\xi} g(o)^\xi \ln^2 \frac{g(o)}{f(o)} do \quad (352)$$

$$= \mathbb{E}_{o \sim f} \left[\left(\frac{g(o)}{f(o)} \right)^\xi \ln^2 \frac{g(o)}{f(o)} \right] \quad (353)$$

$$\leq \mathbb{E}_{o \sim f} \left[\frac{g(o)}{f(o)} \right]^\xi \mathbb{E}_{o \sim f} \left[\left(\ln \frac{g(o)}{f(o)} \right)^{\frac{2}{1-\xi}} \right]^{1-\xi} \quad (354)$$

$$\leq \mathbb{E}_{o \sim f} \left[\left(\ln \frac{g(o)}{f(o)} \right)^4 \right]^{1/2}. \quad (355)$$

Combining the inequalities above we get the desired result. □