
Order-Optimal Instance-Dependent Bounds for Offline Reinforcement Learning with Preference Feedback

Zhirui Chen¹ Vincent Y. F. Tan¹

Abstract

We consider offline reinforcement learning (RL) with preference feedback in which the implicit reward is a linear function of an unknown parameter. Given an offline dataset, our objective consists in ascertaining the optimal action for each state, with the ultimate goal of minimizing the *simple regret*. We propose an algorithm, RL with Locally Optimal Weights or RL-LOW, which yields a simple regret of $\exp(-\Omega(n/H))$ where n is the number of data samples and H denotes an instance-dependent hardness quantity that depends explicitly on the suboptimality gap of each action. Furthermore, we derive a first-of-its-kind instance-dependent lower bound in offline RL with preference feedback. Interestingly, we observe that the lower and upper bounds on the simple regret match order-wise in the exponent, demonstrating order-wise optimality of RL-LOW. In view of privacy considerations in practical applications, we also extend RL-LOW to the setting of (ϵ, δ) -differential privacy and show, somewhat surprisingly, that the hardness parameter H is unchanged in the asymptotic regime as n tends to infinity; this underscores the inherent efficiency of RL-LOW in terms of preserving the privacy of the observed rewards. Given our focus on establishing instance-dependent bounds, our work stands in stark contrast to previous works that focus on establishing worst-case regrets for offline RL with preference feedback.

1. Introduction

Reinforcement Learning (RL) (Sutton and Barto, 2018) has been widely recognized for its capacity to facilitate agents in learning a sequence of optimal actions through iterative in-

teractions with their environments. However, RL encounters significant hurdles in environments that are characterized by uncertainty or lacking explicit reward signals. To address these shortcomings, the concept of RL with human feedback (or RLHF) has emerged as a prominent paradigm. Preference-based RL (PbRL) (Christiano et al., 2017; Chen et al., 2022; Ibarz et al., 2018; Palan et al., 2019) has stood out as one of the most widely used frameworks for RLHF. In this regard, preference-based RL has achieved remarkable performances in practical applications, with particular importance lying in its ability to align large language models (LLMs) with human intent, thereby mitigating the output of toxic and dishonest information (Ouyang et al., 2022; Ziegler et al., 2019; Glaese et al., 2022; Bai et al., 2022; Liu et al., 2023), and improving the quality of applying to the specific tasks (Stiennon et al., 2020; Wu et al., 2021; Nakano et al., 2021).

In this work, we tackle the problem of offline RL with preference feedback, wherein the learning mechanism operates solely on pre-existing (or offline) data without dynamically engaging with the environment. Given the high cost associated with human interaction, offline RL has assumed particular importance in the context of incorporating human feedback. The significance of this offline framework has been justified by many previous prominent works (Shin et al., 2023; Ouyang et al., 2022; Zhu et al., 2023; Kim et al., 2023). For instance, within the learning process of InstructGPT (Ouyang et al., 2022), a pivotal procedure involves the training of a reward model utilizing pre-trained LLM feature vectors, coupled with the utilization of pre-collected human preference feedback as the training dataset. Conceptually, this procedure can be construed as treating the current prompt context as a state within a certain Markov Decision Process (MDP), while the responses generated by the LLM serve as actions within this process. Empirical findings presented by Ouyang et al. (2022) demonstrate the efficacy of this offline framework in effectively aligning human intent with the outputs of LLMs.

However, the literature concerning theoretical analyses within the domain of offline PbRL remains rather scant. Previous theoretical analyses (Zhu et al., 2023; Zhan et al., 2024) of offline PbRL predominantly focused on the worst-case (or minimax) regret, often resulting in the derivation

¹National University of Singapore. Correspondence to: Zhirui Chen <zhiruichen@u.nus.edu>, Vincent Y. F. Tan <vtan@nus.edu.sg>.

ICML 2024 Workshop on Models of Human Feedback for AI Alignment, Vienna, Austria. Copyright 2024 by the author(s).

of regret upper bounds for their algorithms of the form $\tilde{O}(n^{-1/2})$, where n is the size of the offline dataset. In this work, we adopt a different approach that is centered on instance-dependent guarantees. In other words, we wish to derive performance guarantees that are functions of the specific problem instance, thus elucidating the role of fundamental hardness parameters. This yields complementary insights to the existing worst-case analyses. To this end, we design and analyze RL-LOW, a preference-based RL algorithm. Our analysis of the performance RL-LOW unveils an instance-dependent simple regret bound of $\exp(-\Omega(n/H))$, where H is a hardness parameter. This reveals that the simple regret decays exponentially fast in the size of the dataset n and the exponential rate of convergence has also been identified. Complementarily, by proving an instance-dependent lower bound, we show that any algorithm will suffer from a simple regret of at least $\exp(-O(n/H))$. Thus, the dependence of the problem on H is fundamental and cannot be improved upon, thereby demonstrating the efficacy of RL-LOW and the tightness of our analyses.

1.1. Related works

Preference-based RL: From the empirical viewpoint, Christiano et al. (2017) initially demonstrated that RL systems can effectively address complex tasks like Atari games and simulated robot locomotion by learning from human preferences between trajectory segments. Later, numerous researchers started to employ human preference feedback to enhance the performance of LLMs, e.g., aligning the LLMs’ behavior with human intent (Ouyang et al., 2022; Ziegler et al., 2019; Glaese et al., 2022; Bai et al., 2022; Liu et al., 2023), and enhancing the efficacy of application to specific tasks (Stiennon et al., 2020; Wu et al., 2021; Nakano et al., 2021).

From the theoretical perspective, the existing literature remains sparse in offline RL with preference feedback. Zhu et al. (2023) elucidated the failure of the maximum likelihood estimation (MLE) procedure in some scenarios. Motivated by this, they theoretically prove the (near) minimax optimality of the PESSIMISTIC MLE approach with a high probability guarantee. In addition, Zhan et al. (2024) introduced a novel paradigm for general reward functions, and they introduce ε -bracket approximations for reward models, accompanied by a rigorous theoretical analysis delineating sample complexity in terms of approximation error ε and the high-probability parameter δ .

We observe that the above theoretical investigations, while invaluable, are not instance-dependent. Typically, the above minimax or worst-case guarantees yield upper bounds in the form of $\tilde{O}(n^{-1/2})$ and do not depend on any problem-specific factors (such as suboptimality gaps). Our research stands out as a pioneering attempt in offering an instance-

dependent examination for offline RL with preference feedback, thereby bridging a critical gap in the existing literature.

Label-Differential Privacy: In our study, we also consider the notion of *label privacy*, acknowledging that the labels in our offline dataset originate from users, thus highlighting the imperative to protect user privacy. Chaudhuri and Hsu (2011) were among the pioneers in exploring the concept of *label privacy* within the context of supervised learning for binary classification. Their foundational work posits that the sensitive information primarily resides in the labels, while considering the unlabeled attributes as non-sensitive. Later, the concept of label privacy has been investigated across various machine learning paradigms, including but not limited to PAC learning (Beimel et al., 2013) and deep learning frameworks (Ghazi et al., 2021). This broadened examination underscores the significance and relevance of label privacy considerations across diverse areas of machine learning research and applications.

More recently, Chowdhury et al. (2024) investigated the use of label differential privacy to protect the privacy of human labelers in the process of estimating rewards from preference-based feedback. Chowdhury et al. (2024) derive an upper bound for their proposed algorithm on the estimation error. They show that it also decays as $O(n^{-1/2})$ and the implied constant here depends on (ε, δ) , the parameters that define differential privacy. This bound only applies in the scenario of estimating the reward value and is not applicable if we want to understand how it depends on the simple regret of a specific instance. In our work, we consider the effect of (ε, δ) -DP on the simple regret.

1.2. Our contributions

We summarize our main contributions as follows:

1. We establish the first-of-its-kind instance-dependent lower bound characterized by suboptimality gaps for a given problem instance. Our analysis reveals that this lower bound takes the form $\exp(-O(n/H))$, where H is a hardness parameter that is an explicit function of the suboptimality gaps. This finding furnishes a novel, and possibly generalizable, analytical approach for assessing algorithmic performance within the realm of preference-based RL.
2. We design a simple algorithm RL-LOW based on the novel concept of *locally optimal weights*. Our analysis demonstrates that its expected simple regret matches the aforementioned instance-dependent lower bound (in the exponential decay rate of the simple regret), thus revealing our algorithm’s achievement of instance-dependent optimality.
3. We adapt RL-LOW to the (ε, δ) -differential privacy paradigm by combining the Gaussian mechanism with the aforementioned locally optimal weights. Our analy-

sis demonstrates that, for large datasets, this combination enables our algorithm to achieve differential privacy without weakening the bound on the simple regret, underscoring the superiority of the design and analysis of RL-LOW.

4. As a by-product of our analyses, we show that RL-LOW achieves a worst-case bound of the form $O(n^{-1/2})$. If we translate the high-probability upper bound in Zhu et al. (Zhu et al., 2023) to the same worst-case setting, we obtain a bound of the form $O(\sqrt{n^{-1} \log n})$. Thus, our work provides a noticeable (albeit small) improvement over the state-of-the-art theoretical result in (Zhu et al., 2023).

2. Preliminaries and Problem Setup

Let $\mathcal{S} = \{1, \dots, S\}$ denote the state space, and $\mathcal{A} = \{1, \dots, A\}$ denote the action set. We assume that there is an unknown non-degenerate distribution $\rho = (\rho_1, \dots, \rho_S)$ over the states, i.e., $\rho_k > 0$ for all $k \in \mathcal{S}$. The i -th action of state k is associated with the feature vector $\phi(k, i) \in \mathbb{R}^d$, and its associated (unknown) reward is

$$r_{k,i} = \langle \phi(k, i), \theta \rangle, \quad (1)$$

where $\theta \in \mathbb{R}^d$ is an unknown parameter vector. The collection of all feature vectors is denoted as $\phi = \{\phi(k, i)\}_{k \in \mathcal{S}, i \in \mathcal{A}}$. For all $k \in \mathcal{S}$, we denote the suboptimality gap of action $i \in \mathcal{A}$ as $\Delta_{k,i} = \max_{j \in \mathcal{A}} r_{k,j} - r_{k,i}$. Let $(a^{(0)}, a^{(1)}) \in \mathcal{A}^2$ be a pair of comparisons and let $s \in \mathcal{S}$ be a state. Then, we define a stochastic label σ , following the Bradley–Terry–Luce (BTL) model as

$$\mathbb{P}(\sigma = 1 \mid a^{(0)}, a^{(1)}, s) = \frac{\exp(r_{s,a^{(1)}})}{\exp(r_{s,a^{(0)}}) + \exp(r_{s,a^{(1)}})}. \quad (2)$$

Given this model, we assume throughout that we have access to an *offline dataset*, which we denote as $\mathcal{D} = \{(s_i, a_i^{(0)}, a_i^{(1)}, \sigma_i)\}_{i=1}^n$. Note that this dataset consists of n tuples of states, pairs of actions for comparison, and stochastic labels. Without loss of generality, we assume that the comparisons are arranged such that $a_i^{(0)} < a_i^{(1)}$ for all $i = 1, \dots, n$, and $a_i^{(0)} < a_j^{(0)}$ (or $a_i^{(1)} \leq a_j^{(1)}$ if $a_i^{(0)} = a_j^{(0)}$) for all $i < j$. For simplicity, we assume that the feature vectors satisfy $\phi(k, i) \neq \phi(k, j)$ for all states $k \in \mathcal{S}$ and all actions $i \neq j$. In addition, we assume that for each state k , the best action $i_k^* = \arg \max_{j \in \mathcal{A}} r_{k,j}$ is unique. Broadly speaking, our objective is to use the offline dataset \mathcal{D} to estimate the best action i_k^* for each state $k \in \mathcal{S}$. Following Zhu et al. (2023), we aim to design a (possibly randomised) policy π that uses the dataset \mathcal{D} to output a set of actions $\{\hat{i}_k\}_{k \in \mathcal{S}}$ that minimizes the *simple regret*¹

¹The term ‘‘simple regret’’ is referred to as ‘‘performance gap’’ in some existing works (e.g., Zhu et al. (2023)).

(averaged over all S states), defined as

$$R_n = \sum_{k \in \mathcal{S}} \rho_k (r_{k, \hat{i}_k} - r_{k, i_k^*}). \quad (3)$$

Let $N \in \mathbb{R}^{S \times A \times A}$ be a tensor that collects the proportions of each comparison in the dataset \mathcal{D} . In particular, $N_{k,i,j} := \frac{1}{n} \sum_{\ell=1}^n \mathbb{1}\{s_\ell = k, a_\ell^{(1)} = i, a_\ell^{(2)} = j\}$ is the proportion of the number of times actions i and j have been compared under state k . A problem instance, denoted as v , is completely characterised by the tuple $(\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta)$. In the following, we index instance-specific parameters with the instance v to indicate their dependence on v ; this will be omitted when the instance is clear from the context. In addition, we write \mathbb{P}_v^π (resp. \mathbb{E}_v^π) to denote the probability measure (reps. the expectation) induced under policy π and under the instance v . Other assumptions are as follows.

Assumption 2.1. (Bounded reward) There exists a finite and known constant L such that for any $k \in \mathcal{S}$ and $i \in \mathcal{A}$, it holds that $|\langle \phi(k, i), \theta \rangle| \leq L$.

In previous works (Zhu et al., 2023), the authors assume that the norms of the feature vectors $\phi(k, i)$ and parameter vector θ are separately bounded. This clearly implies that Assumption 2.1 is satisfied, but Assumption 2.1 is weaker as it is a bound on the rewards.

Definition 2.2. (Consistent instance) A problem instance $v = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta)$ is *consistent* if for all $(k, i, j) \in \mathcal{S} \times \mathcal{A}^2$, it holds that $\phi(k, i) - \phi(k, j) \in \text{Span}\{\phi(k', i') - \phi(k', j') : (k', i', j') \in \mathcal{S} \times \mathcal{A}^2 \text{ and } N_{k', i', j'} > 0\}$.

We say an instance v is *inconsistent* if it is not consistent. In the following, we will be only concerned with those instances that are consistent as the following result shows that it is impossible to design a policy that achieves vanishing simple regret for inconsistent instances.

Proposition 2.3. (Impossibility result) *For any inconsistent instance $v = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta)$, there exists an instance $v' = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta')$ such that for all policies π*

$$\liminf_{n \rightarrow \infty} \{\mathbb{E}_v^\pi[R_n] + \mathbb{E}_{v'}^\pi[R_n]\} > 0. \quad (4)$$

3. The Proposed Algorithm: RL-LOW

In this section, we describe our computationally and statistically efficient algorithm for offline RL with preference feedback based on the novel idea of *locally optimal weights* for estimating the relative reward of each pair of actions. This algorithm, called RL-LOW, is simple and is presented formally in Algorithm 1. Before we describe its components, we introduce some notations.

First, we denote $B_{k,i,j}$ as the empirical success rate with the comparison of actions i and j , i.e., for $k \in \mathcal{S}$ and $i, j \in \mathcal{A}$

Algorithm 1 Reinforcement Learning with Locally Optimal Weights (RL-LOW)

Input: Dataset $\mathcal{D} = \{(s_i, a_i^{(1)}, a_i^{(2)}, \sigma_i)\}_{i=1}^n$ and feature maps $\phi = \{\phi(k, i)\}_{k \in \mathcal{S}, i \in \mathcal{A}}$.

Output: The estimated best action $\hat{i}_k \in \mathcal{A}$ of each state $k \in \mathcal{S}$.

- 1: Compute the sample proportions $N_{k,i,j} \leftarrow \frac{1}{n} \sum_{l=1}^n \mathbb{1}\{s_l = k, a_l^{(1)} = i, a_l^{(2)} = j\}$.
- 2: For $k \in \mathcal{S}$ and $i, j \in \mathcal{A}$ such that $N_{k,i,j} > 0$, compute the success rate of each pair of comparisons using Eqn. (5).
- 3: Clip the success rate using the knowledge of L and using Eqn. (6).
- 4: For each state $k \in \mathcal{S}$ and distinct actions $i, j \in \mathcal{A}$ with $i < j$, compute the locally optimal weights $(w_{k',i',j'}^{(k,i,j)})_{k' \in \mathcal{S}, i', j' \in \mathcal{A}}$ using Eqn. (7).
- 5: Compute the empirical relative reward $\hat{r}_{k,i,j}$ for each $k \in \mathcal{S}, i, j \in \mathcal{A}$ using Eqn. (8).
- 6: **return** For each $k \in \mathcal{S}$, let $\hat{i}_k \in \{i \in \mathcal{A} : \hat{r}_{k,i,j} \geq 0, \forall j \neq i\}$; resolve ties uniformly.

with $N_{k,i,j} > 0$,

$$B_{k,i,j} := \frac{1}{nN_{k,i,j}} \sum_{l=1}^n \sigma_l \mathbb{1}\{s_l = k, a_l^{(1)} = i, a_l^{(2)} = j\}, \quad (5)$$

and $B_{k,j,i} := 1 - B_{k,i,j}$. If $N_{k,i,j} = N_{k,j,i} = 0$, we define $B_{k,i,j} = B_{k,j,i} = 0$. Subsequently, certain empirical success rates may exhibit magnitudes that are either excessively large or small. We clip them by means of the following operation: $B_{k,i,j}^{\text{CLP}} = \text{CLIP}_L(B_{k,i,j})$, where

$$\text{CLIP}_L(a) = \begin{cases} \frac{\exp(2L)}{1+\exp(2L)} & a > \frac{\exp(2L)}{1+\exp(2L)} \\ \frac{1}{1+\exp(2L)} & a < \frac{1}{1+\exp(2L)} \\ a & \text{otherwise} \end{cases} \quad (6)$$

In accordance with Assumption 2.1, the implicit rewards are bounded by L . Consequently, within our BTL model framework, the success rate of each comparison necessarily falls within the interval $[\frac{1}{1+\exp(2L)}, \frac{\exp(2L)}{1+\exp(2L)}]$. We exploit this in (6) to ensure that the implementation of our clip operation is consistent with the model’s dynamics. We are now ready to introduce the notion of *locally optimal* weights, which plays a central role in the estimation of the rewards.

Definition 3.1. (Locally Optimal Weight) For an consistent instance v , let $\mathcal{U}_{k,i,j} = \{u \in \mathbb{R}^{\mathcal{S} \times \mathcal{A} \times \mathcal{A}} : \phi(k, i) - \phi(k, j) = \sum_{k' \in \mathcal{S}, i', j' \in \mathcal{A}} u_{k',i',j'} (\phi(k', i') - \phi(k', j')) \text{ and } u_{k',i',j'} = 0 \text{ if } N_{k',i',j'} = 0\}$. We say that $w^{(k,i,j)} = (w_{k',i',j'}^{(k,i,j)})_{k' \in \mathcal{S}, i', j' \in \mathcal{A}}$ is a set of *locally optimal*

weights for $(k, i, j) \in \mathcal{S} \times \mathcal{A}^2$ with $i \neq j$ if

$$w^{(k,i,j)} \in \operatorname{argmin}_{u \in \mathcal{U}_{k,i,j}} \left\{ \sum_{k' \in \mathcal{S}, i', j' \in \mathcal{A}: N_{k',i',j'} > 0} \frac{(u_{k',i',j'})^2}{N_{k',i',j'}} \right\}. \quad (7)$$

The weights in Eqn. (7) are described as “locally optimal” because they are customized to each (k, i, j) tuple. Hence, $w^{(k,i,j)}$ is *local* to (k, i, j) . This is a novelty in the design of our algorithm.

By the definition of the consistency of an instance (cf. Definition 2.2), there exists a subset $\beta \subset \mathcal{S} \times \mathcal{A}^2$ such that $\phi(k, i) - \phi(k, j) \in \operatorname{Span}\{\phi(k', i') - \phi(k', j') : (k', i', j') \in \beta\}$. Hence, there exists a locally optimal weight for every pair of actions given a consistent instance. In addition, $w^{(k,i,j)}$ can be calculated efficiently by its analytic form (see details in Appendix C).

Equipped with the definition of locally-optimal weights, we now provide an estimate of the relative reward for state $k \in \mathcal{S}$ and pair of action $(i, j) \in \mathcal{A}^2$ with $i \neq j$ as follows:

$$\hat{r}_{k,i,j} = \sum_{k' \in \mathcal{S}, i', j' \in \mathcal{A}} w_{k',i',j'}^{(k,i,j)} \log \left(\frac{B_{k',i',j'}^{\text{CLP}}}{1 - B_{k',i',j'}^{\text{CLP}}} \right). \quad (8)$$

The term $\frac{(u_{k',i',j'})^2}{N_{k',i',j'}}$ in Eqn. (7) is a proxy for the variance introduced by the pair of actions (i', j') in state k' when associated with the coefficient $u_{k',i',j'}$ in the linear combination of the definition of $\mathcal{U}_{k,i,j}$. Our objective is to minimize the cumulative variance proxy for (k, i, j) , thus enhancing the precision of the estimate of the relative reward for (k, i, j) for the purposes of establishing the tightest possible concentration result for subGaussian random variables (see Appendix E).

Finally, for any $k \in \mathcal{S}$, let $\hat{i}_k \in \hat{\mathcal{I}}_k := \{i \in \mathcal{A} : \hat{r}_{k,i,j} \geq 0, \forall j \neq i\}$ be any estimate of the best action under state k . It is natural to wonder whether \hat{i}_k exists, i.e., whether the set $\hat{\mathcal{I}}_k$ is empty. The following proposition answers this in the affirmative.

Proposition 3.2. For any consistent instance v and using estimate of the best action under each state k as prescribed by RL-LOW, we have $|\hat{\mathcal{I}}_k| \geq 1$, i.e., \hat{i}_k exists.

3.1. Upper Bound of RL-LOW

In this section, we provide an instance-dependent upper bound of the simple regret for RL-LOW. In addition, we also provide a worst-case upper bound as a by-product. First, we define an instance-dependent hardness parameter $H(v)$. Let

$$H(v) := \max_{k \in \mathcal{S}, i \in \mathcal{A}: i \neq i_k^*} \frac{\gamma_{k,i}}{\Delta_{k,i}^2}, \quad (9)$$

where

$$\gamma_{k,i} := \sum_{k' \in \mathcal{S}, i', j' \in \mathcal{A}: N_{k', i', j'} > 0} \frac{(w_{k', i', j'}^{(k, i, i_k^*)})^2}{N_{k', i', j'}}.$$

The parameter $\gamma_{k,i}$ exhibits a positive correlation with the variance proxy of the relative empirical reward \hat{r}_{k, i, i_k^*} . Consequently, the ratio $\frac{\gamma_{k,i}}{\Delta_{k,i}^2}$ in the definition of $H(v)$ serves as a quantitative measure of the difficulty that the empirical reward of a suboptimal action i surpasses that of i_k^* in state k .

Theorem 3.3. (*Instance-Dependent Upper Bound*) *For any consistent instance v , under RL-LOW, we have for all sufficiently large n ,*

$$\mathbb{E}_v^{\text{RL-LOW}} [R_n] \leq \exp\left(-\frac{n}{C_{\text{up}} \cdot H(v)}\right), \quad (10)$$

where C_{up} is a universal constant.²

From Theorem 3.3, it is evident that the upper bound decays exponentially fast and the exponent is a function of an instance-dependent hardness term $H(v)$. This is the first instance-dependent analysis in offline reinforcement learning with preference feedback.

It is natural to wonder why we do not devise an instance-dependent analysis of or modification to PESSIMISTIC MLE which was developed by Zhu et al. (2023). Note that PESSIMISTIC MLE is designed to perform well *with high probability* and not necessarily *in expectation*. In particular, the regret bound of PESSIMISTIC MLE holds with probability at least $1 - \delta$. Hence, to ensure the regret is less than $\exp(-\Omega(n/H(v)))$, one should set the failure probability δ to be $\exp(-\Theta(n/H(v)))$, which is not possible as $H(v)$ is unknown to the algorithm (since θ is also unknown).

We further provide a worst-case upper bound for RL-LOW as follows.

Proposition 3.4. (*Worst-Case Upper Bound*) *For any consistent instance v and for all $n \geq 1$,*

$$\mathbb{E}_v^{\text{RL-LOW}} [R_n] \leq \frac{\sum_{k, i: i \neq i_k^*} \rho_k(\sqrt{\gamma_{k,i}} + \tilde{\gamma}_{k,i})}{C_{\text{wup}} \sqrt{n}} \quad (11)$$

where

$$\tilde{\gamma}_{k,i} = \sum_{k', i', j': N_{k', i', j'} > 0} \frac{|w_{k', i', j'}^{(k, i, i_k^*)}|}{\sqrt{N_{k', i', j'}}} \quad (12)$$

and $C_{\text{wup}} > 0$ is a universal constant.

²In this paper, our universal constants depend on L , which is known and fixed throughout.

We note that in Zhu et al. (2023), the high probability upper bound is of the form $O(\sqrt{n^{-1} \log(1/\delta)})$. Hence, if we desire a bound in expectation, we obtain, through the law of total probability, a bound of the form $\mathbb{E}[R_n] = O(\sqrt{n^{-1} \log(1/\delta)} + \delta)$. Minimizing this bound over δ yields $\mathbb{E}[R_n] = O(\sqrt{n^{-1} \log n})$, which is slightly inferior to our upper bound $O(n^{-1/2})$.

4. Instance-Dependent Lower Bound

In this section, we derive the first-of-its-kind instance-dependent lower bound on offline RL with preference feedback. Before we present our bound, we present some auxiliary lemmas that are potentially instrumental in deriving lower bounds on other preference-based RL problems.

Given any instance v , we let $P_v^{(n)}$ denote the joint distribution of the associated labels $\{\sigma_i\}_{i=1}^n$. The following lemma provides an estimate of the Kullback–Leibler (KL) divergence between instances v and v' that share the same parameters except for the latent vector θ that defines the reward in (1).

Lemma 4.1. *For any instance $v = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta)$ and $v' = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta')$, it holds that*

$$2n \exp(-4R_{\text{max}}) \leq \frac{D_{\text{KL}}(P_v^{(n)} \| P_{v'}^{(n)})}{\tilde{D}(v, v')} \leq 2n \exp(2R_{\text{max}}) \quad (13)$$

where

$$\tilde{D}(v, v') = \sum_{k \in \mathcal{S}, i, j \in \mathcal{A}} N_{k, i, j} (\langle \phi(k, i) - \phi(k, j), \theta - \theta' \rangle)^2, \quad (14)$$

and $R_{\text{max}} = \max_{k \in \mathcal{S}, i \in \mathcal{A}} \max\{|\langle \phi(k, i), \theta \rangle|, |\langle \phi(k, i), \theta' \rangle|\}$ is the maximum absolute reward in these two instances.

This lemma demonstrates that when the rewards are bounded, the weighted sum of squared differences of the relative rewards can be used to approximate the KL divergence between the distributions of two instances. The approximation is precisely $\tilde{D}(v, v')$ defined in (14). Furthermore, for any $\mathbf{z} \in \mathbb{R}^d$, $\eta \in \mathbb{R}$ and consistent instance $v = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta)$, we let $\text{Alt}(v, \mathbf{z}, \eta)$ be the set of instances that share the same instance parameters except for θ and satisfies $\langle \mathbf{z}, \theta' - \theta \rangle = \eta$ for all $v' = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta') \in \text{Alt}(v, \mathbf{z}, \eta)$. The following lemma states a useful property that relates the Alt set to the “approximate KL divergence” \tilde{D} .

Lemma 4.2. *Let \mathcal{G} be an arbitrary orthonormal basis of $\text{Span}\{\phi(k', i') - \phi(k', j') : (k', i', j') \in \mathcal{S} \times \mathcal{A}^2 \text{ and } N_{k', i', j'} > 0\}$. Also let $[\mathbf{w}]_{\mathcal{G}}$ denote the column vector that represents \mathbf{w} under basis \mathcal{G} (Meyer, 2000, Chap-*

ter 4). Define the matrix

$$V := \sum_{k \in \mathcal{S}, i, j \in \mathcal{A}} N_{k,i,j} [\phi(k,i) - \phi(k,j)]_{\mathcal{G}} [\phi(k,i) - \phi(k,j)]_{\mathcal{G}}^{\top}. \quad (15)$$

Then for any consistent instance $v = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta)$, $\eta \in \mathbb{R}$, and $\mathbf{z} \in \text{Span}\{\phi(k', i') - \phi(k', j') : (k', i', j') \in \mathcal{S} \times \mathcal{A}^2 \text{ and } N_{k', i', j'} > 0\}$,

$$\min_{v' \in \text{Alt}(v, \mathbf{z}, \eta)} \tilde{D}(v, v') = \frac{\eta^2}{\|\mathbf{z}\|_{\mathcal{G}}^2 \|V^{-1}\|}. \quad (16)$$

Lemma 4.2 provides an estimate of the KL divergence between instance v and $v' \in \text{Alt}(v, \mathbf{z}, \eta)$. This, in turn, provides a convenient means to apply the ubiquitous *change of measure* technique to derive the lower bound.

In addition, let (\bar{i}, \bar{k}) be the state-action pair that attains maximum in the definition of hardness in Eqn. (9). Define the subset of instances

$$\mathcal{Q} = \left\{ v \text{ consistent} : \frac{\gamma_{\bar{k}, \bar{i}}}{\Delta_{\bar{k}, \bar{i}}^2} \geq 4 \frac{\gamma_{k,i}}{\Delta_{k,i}^2} \right. \\ \left. \forall (k, i) \neq (\bar{k}, \bar{i}) \text{ and } i \neq i_k^* \right\}. \quad (17)$$

We are now ready to state our lower bound.

Theorem 4.3. (Instance-Dependent Lower Bound) For any instance $v = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta) \in \mathcal{Q}$, there exists another instance $v' = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta')$ with $H(v) \leq H(v') \leq 8H(v)$ such that for all sufficiently large n ,

$$\inf_{\pi} \{ \mathbb{E}_{\pi}^v [R_n] + \mathbb{E}_{\pi}^{v'} [R_n] \} \geq \exp \left(- \frac{n}{C_{\text{lo}} \cdot H(v)} \right),$$

where $C_{\text{lo}} > 0$ is a universal constant.

The alternative instance v' that appears in Theorem 4.3 is judiciously chosen to be $v' \in \text{Alt}(v, \phi(\bar{k}(v), \bar{i}(v)) - \phi(\bar{k}(v), i_k^*(v)), 2\Delta_{\bar{k}(v), \bar{i}(v)}(v))$. In particular, it is designed so that the optimal action i_k^* under state k of instance v will become suboptimal under instance v' , and its suboptimality gap is at least $\Delta_{\bar{k}(v), \bar{i}(v)}(v)$ under v' .

Theorem 4.3 is an instance-dependent lower bound for all instances in the set \mathcal{Q} . The condition that defines \mathcal{Q} in Eqn. (17) ensures that the hardness quantities $H(v)$ and $H(v')$ have the same order. Since instances in \mathcal{Q} cover all possible hardness values $H(v)$ (i.e., for every hardness values $h > 0$, there exists an instance in \mathcal{Q} of hardness h), we conclude that for any consistent instance v and (small) $\epsilon \in (0, 1)$, there does not exist any policy π that achieves

$$\mathbb{E}_{\pi}^v [R_n] = \exp \left(- \Omega \left(\frac{n}{H(v)^{1-\epsilon}} \right) \right). \quad (18)$$

In this sense, the exponential decay rate of the simple regret of RL-LOW presented in Theorem 3.3 is asymptotically tight (or optimal) and the exponential dependence on the hardness parameter $H(v)$ is necessary, fundamental, and cannot be improved upon.

5. Extension to (ϵ, δ) -Differential Privacy (DP)

In this section, we extend our algorithm RL-LOW to be amenable to (ϵ, δ) -differential privacy. To formalize our results, we provide the definition of (ϵ, δ) -DP, following Dwork et al. (2014). We say that two sets of preference labels, $\sigma := \{\sigma_i\}_{i=1}^n$ and $\sigma' := \{\sigma'_i\}_{i=1}^n$ are *neighboring* if there exists $s \in [n]$ such that $\sigma_s \neq \sigma'_s$ and $\sigma_j = \sigma'_j$ for all $j \neq s$.

Definition 5.1. (Differential Privacy) Fix any label-free dataset $\{(s_i, a_i^{(1)}, a_i^{(2)})\}_{i=1}^n$. A (randomized) policy $\mathcal{M} : \{0, 1\}^n \rightarrow \mathcal{A}^S$ (that takes as inputs a set of labels and outputs a set of actions, one for each state) satisfies (ϵ, δ) -DP if for all neighboring labels $\sigma := \{\sigma_i\}_{i=1}^n$ and $\sigma' := \{\sigma'_i\}_{i=1}^n$,

$$\mathbb{P}(\mathcal{M}(\sigma) \in \mathcal{Z}) \leq e^{\epsilon} \mathbb{P}(\mathcal{M}(\sigma') \in \mathcal{Z}) + \delta \quad \forall \mathcal{Z} \subset \mathcal{A}^S. \quad (19)$$

Note that Definition 5.1 primarily concerns protecting the privacy of users' *labels*. In particular, the DP mechanism guarantees that any alteration in a user's label must not substantially affect the output of our algorithm. Otherwise, there exists a risk that a user's label might be inferred through the algorithm's output. Our definition of differential privacy (DP) aligns with that of (Chowdhury et al., 2024).

We now adapt our RL-LOW to (ϵ, δ) -DP by using the Gaussian mechanism (Dwork et al., 2014). Firstly, we introduce the private version of the empirical success rate (analogous to $B_{k,i,j}$ in (5)) which we denote as $\forall k \in \mathcal{S}, (i, j) \in \mathcal{A}^2$ with $N_{k,i,j} > 0$,

$$\tilde{B}_{k,i,j} := \frac{1}{nN_{k,i,j}} \sum_{\iota=1}^n \sigma_{\iota} \mathbb{1}\{s_{\iota} = k, a_{\iota}^{(1)} = i, a_{\iota}^{(2)} = j\} + \tilde{\xi}_{k,i,j}$$

where $\tilde{\xi}_{k,i,j}$ is an independent (across k, i , and j) Gaussian noise with zero mean and variance $\frac{2 \log(1.25/\delta)}{(\epsilon n N_{k,i,j})^2}$, and we let

$\tilde{B}_{k,j,i} := 1 - \tilde{B}_{k,i,j}$. If $N_{k,i,j} = N_{k,j,i} = 0$, we define $\tilde{B}_{k,i,j} = \tilde{B}_{k,j,i} = 0$. Again, analogously to the operation in (6), we clip $\tilde{B}_{k,i,j}$ to form

$$\tilde{B}_{k,i,j}^{\text{CLP}} = \text{CLIP}_L(\tilde{B}_{k,i,j}) \quad (20)$$

Similarly to Eqn. (8), the perturbed estimated relative rewards are given as follows

$$\tilde{r}_{k,i,j} = \sum_{k' \in \mathcal{S}, (i', j') \in \mathcal{A}^2} w_{k', i', j'}^{(k, i, j)} \log \left(\frac{\tilde{B}_{k', i', j'}^{\text{CLP}}}{1 - \tilde{B}_{k', i', j'}^{\text{CLP}}} \right), \quad (21)$$




Figure 1: Comparison of RL-LOW and DP-RL-LOW to PESSIMISTIC MLE (Zhu et al., 2023) on average simple regret and standard deviation (shaded area). In the left figure, we set $\delta = 0.2$ and $\varepsilon = 2.0$ for DP-RL-LOW. In the right figure, we set $n = 300$ for all policies, and ε does not apply to RL-LOW.

where $w^{(k,i,j)}$ is defined in Definition 3.1. Finally, the empirical best action is $\hat{i}_k \in \tilde{\mathcal{I}}_k := \{i \in \mathcal{A} : \tilde{r}_{k,i,j} \geq 0, \forall j \neq i\}$. A similar argument as Proposition 3.2 shows that \hat{i}_k exists; see details in Appendix C for the details. The algorithm described above is a differentially private version of the RL-LOW policy and hence, it is named DP-RL-LOW.

DP-RL-LOW with the carefully chosen variance of $\xi_{k,i,j}$ fulfils the requirement of (ε, δ) -DP.

Proposition 5.2. *Given privacy parameters $\varepsilon, \delta > 0$, DP-RL-LOW satisfies (ε, δ) -DP.*

The proof of Proposition 5.2 follows exactly along the lines of the proof of Dwork et al. (2014, Theorem A.1) and is omitted. We next upper bound the expected simple regret of DP-RL-LOW.

Theorem 5.3. *(Instance-Dependent Upper Bound for DP-RL-LOW) Given any consistent instance v , for all sufficiently large n ,*

$$\begin{aligned} & \mathbb{E}_v^{\text{DP-RL-LOW}} [R_n] \\ & \leq \exp \left(-C_{\text{DP}} \cdot \left(\frac{n}{H(v)} \wedge \left(\frac{n}{H_{\text{DP}}^{(\varepsilon, \delta)}(v)} \right)^2 \right) \right), \end{aligned}$$

where $C_{\text{DP}} > 0$ is a universal constant, and

$$H_{\text{DP}}^{(\varepsilon, \delta)}(v) = \max_{k \in \mathcal{S}, i \in \mathcal{A}: i \neq i_k^*} \frac{\sqrt{\log(\frac{1.25}{\delta})} \gamma_{k,i}^{\text{DP}}}{\sqrt{\varepsilon} \Delta_{k,i}} \quad (22)$$

and

$$\gamma_{k,i}^{\text{DP}} = \sum_{k', i', j' \in \mathcal{A}: N_{k', i', j'} > 0} \left(\frac{w_{k', i', j'}^{(k, i, i_k^*)}}{N_{k', i', j'}} \right)^2.$$

Consequently,

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{E}_v^{\text{DP-RL-LOW}} [R_n] \leq -\frac{C_{\text{DP}}}{H(v)}. \quad (23)$$

The limiting statement in (23) implies that DP-RL-LOW has the same order of the exponential decay rate as its non-differentially privacy counterpart RL-LOW when n is sufficiently large; in particular, $n > (H_{\text{DP}}^{(\varepsilon, \delta)}(v))^2 / H(v)$ suffices to nullify the effect of the privacy requirement. In other words, in the sense of the exponent, privacy comes “for free” for sufficiently large offline datasets.

6. Numerical Simulations

In this section, we present some numerical simulations of our algorithm RL-LOW and its differentially private counterpart DP-RL-LOW. We compare them to the state-of-the-art (non-private algorithm) PESSIMISTIC MLE developed by Zhu et al. (2023). We conduct the experiments on a synthetic dataset. Specifically, we set the number of states $S = 2$, the number of actions $A = 10$, the dimensionality of the data $d = 5$, the unknown parameter vector $\theta = [1, 1, 1, 1, 1]^\top$, and the state distribution $\rho = [0.4, 0.6]$. The feature vector of each action is generated as follows: For the i -th action of state $k \in \{1, 2\}$, we first uniformly generate a d -dimensional vector $\phi'(k, i)$ with all non-negative elements and $\|\phi'(k, i)\|_1 = 1$. Then, for each state $k \in \{1, 2\}$, we set the feature vector of i -th action as $\phi(k, i) = \phi'(k, i) - 0.01(i-1)\theta$. That is, in both state 1 and 2, the best action is the first action, and the suboptimality gap of the i -th action is $0.05i$. In addition, for both states $k \in \{1, 2\}$ and $i < j$, we set $N_{k,i,j} = \frac{1}{A(A-1)}$, i.e., the proportions of comparisons for this instance are uniform. In

the simulation, we use $\lceil N_{k,i,j}n \rceil$ as the number of samples involved in the comparison between actions i and j under state k .

As for the hyperparameters of PESSIMISTIC MLE, we follow the default setting of Zhu et al. (2023, Section 3). In addition, PESSIMISTIC MLE only works under the assumption that $\langle 1, \theta \rangle = 0$ (Zhu et al., 2023, Assumption 2.1). Therefore, when running the experiments of PESSIMISTIC MLE, we further set $d = 6$, $\theta = [1, 1, 1, 1, 1, -5]$ and the 6-th element of each feature vector is set to 0. Then, this new instance is mathematically equivalent to the original instance and additionally satisfies (Zhu et al., 2023, Assumption 2.1) which is needed for PESSIMISTIC MLE.

The simulation results are shown in Figure 1. We run each experiment 200 times, and report the average and standard deviation. From Figure 1 (left), we observe that RL-LOW is inferior to PESSIMISTIC MLE for small n . However, since RL-LOW is instance-dependent optimal in the exponential decay rate and in its dependence on the hardness parameter $H(v)$, the experimental findings depicted in Figure 1 (left) corroborate the empirical superiority of our proposed RL-LOW algorithm over PESSIMISTIC MLE for n sufficiently large ($n > 150$ suffices). This observation underscores the efficacy of our novel algorithmic design based on locally optimal weights. Furthermore, from Figure 1 (left), we also observe that as the sample size n increases, the performance of DP-RL-LOW converges to that of RL-LOW, consistent with our theoretical findings in Theorem 5.3.

Lastly, as shown in Figure 1 (right), it is evident that achieving comparable performance between RL-LOW and DP-RL-LOW may necessitate substantially larger sample sizes n when considering small privacy parameters of ϵ and δ . This observation is again consistent with our theoretical findings in Theorem 5.3.

7. Concluding Remarks

This paper addresses the problem of offline RL with preference feedback, aiming to determine the optimal action for each state to minimize the simple regret. We introduced a novel algorithm, RL-LOW, which achieves a simple regret of $\exp(-\Omega(n/H(v)))$, where n represents the number of data samples and $H(v)$ characterizes an instance-dependent hardness parameter related to the suboptimality gaps of each action. Additionally, we established a first-of-its-kind instance-dependent lower bound for offline RL with preference feedback, demonstrating the order-wise optimality of RL-LOW (in the exponential decay rate) through the matching of lower and upper bounds on the simple regret. To address privacy concerns, we extended RL-LOW to be amenable to (ϵ, δ) -differential privacy, revealing that the hardness parameter $H(v)$ remains unchanged in the

asymptotic regime as n tends to infinity. This underscores RL-LOW’s effectiveness and robustness in preserving the privacy of observed rewards. Our focus on establishing instance-dependent bounds sets this work apart from previous research that focuses primarily on worst-case regret analyses in offline RL with preference feedback.

Some interesting directions for future research include extending our work to incorporate general reward functions (Zhan et al., 2024). In particular, a natural question concerns whether or not there exist an algorithm that is instance-dependent and order-optimal for general reward functions?

Acknowledgements

The authors are supported by a Singapore Ministry of Education Tier 2 grant (Grant Number: A-8000423-00-00).

References

- Bai, Y., Jones, A., Ndousse, K., Askell, A., Chen, A., Das-Sarma, N., Drain, D., Fort, S., Ganguli, D., Henighan, T., et al. (2022). Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- Beimel, A., Nissim, K., and Stemmer, U. (2013). Private learning and sanitization: Pure vs. approximate differential privacy. In *International Workshop on Approximation Algorithms for Combinatorial Optimization*, pages 363–378. Springer.
- Chaudhuri, K. and Hsu, D. (2011). Sample complexity bounds for differentially private learning. In *Proceedings of the 24th Annual Conference on Learning Theory*, pages 155–186. JMLR Workshop and Conference Proceedings.
- Chen, X., Zhong, H., Yang, Z., Wang, Z., and Wang, L. (2022). Human-in-the-loop: Provably efficient preference-based reinforcement learning with general function approximation. In *International Conference on Machine Learning*, pages 3773–3793. PMLR.
- Chowdhury, S. R., Zhou, X., and Natarajan, N. (2024). Differentially private reward estimation with preference feedback. In *International Conference on Artificial Intelligence and Statistics*, pages 4843–4851. PMLR.
- Christiano, P. F., Leike, J., Brown, T., Martic, M., Legg, S., and Amodei, D. (2017). Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30.
- Dwork, C., Roth, A., et al. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407.

- Ghazi, B., Golowich, N., Kumar, R., Manurangsi, P., and Zhang, C. (2021). Deep learning with label differential privacy. *Advances in neural information processing systems*, 34:27131–27145.
- Glaese, A., McAleese, N., Trębacz, M., Aslanides, J., Firoiu, V., Ewalds, T., Rauh, M., Weidinger, L., Chadwick, M., Thacker, P., et al. (2022). Improving alignment of dialogue agents via targeted human judgements. *arXiv preprint arXiv:2209.14375*.
- Ibarz, B., Leike, J., Pohlen, T., Irving, G., Legg, S., and Amodei, D. (2018). Reward learning from human preferences and demonstrations in atari. *Advances in neural information processing systems*, 31.
- Kim, C., Park, J., Shin, J., Lee, H., Abbeel, P., and Lee, K. (2023). Preference transformer: Modeling human preferences using transformers for RL. In *The Eleventh International Conference on Learning Representations*.
- Liu, H., Sferrazza, C., and Abbeel, P. (2023). Chain of hindsight aligns language models with feedback. *arXiv preprint arXiv:2302.02676*.
- Meyer, C. D. (2000). *Matrix analysis and applied linear algebra*, volume 71. SIAM.
- Nakano, R., Hilton, J., Balaji, S., Wu, J., Ouyang, L., Kim, C., Hesse, C., Jain, S., Kosaraju, V., Saunders, W., et al. (2021). Webgpt: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332*.
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al. (2022). Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Palan, M., Landolfi, N. C., Shevchuk, G., and Sadigh, D. (2019). Learning reward functions by integrating human demonstrations and preferences. *arXiv preprint arXiv:1906.08928*.
- Shin, D., Dragan, A., and Brown, D. S. (2023). Benchmarks and algorithms for offline preference-based reward learning. *Transactions on Machine Learning Research*.
- Stiennon, N., Ouyang, L., Wu, J., Ziegler, D., Lowe, R., Voss, C., Radford, A., Amodei, D., and Christiano, P. F. (2020). Learning to summarize with human feedback. *Advances in Neural Information Processing Systems*, 33:3008–3021.
- Sutton, R. S. and Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.
- Tsybakov, A. (2009). *Introduction to Nonparametric Estimation*. Springer Series in Statistics. Springer.
- Vershynin, R. (2018). *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press.
- Wu, J., Ouyang, L., Ziegler, D. M., Stiennon, N., Lowe, R., Leike, J., and Christiano, P. (2021). Recursively summarizing books with human feedback. *arXiv preprint arXiv:2109.10862*.
- Zhan, W., Uehara, M., Kallus, N., Lee, J. D., and Sun, W. (2024). Provable offline preference-based reinforcement learning. In *The Twelfth International Conference on Learning Representations*.
- Zhu, B., Jordan, M., and Jiao, J. (2023). Principled reinforcement learning with human feedback from pairwise or k-wise comparisons. In *International Conference on Machine Learning*, pages 43037–43067. PMLR.
- Ziegler, D. M., Stiennon, N., Wu, J., Brown, T. B., Radford, A., Amodei, D., Christiano, P., and Irving, G. (2019). Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593*.

Supplementary Material for “Order-Optimal Instance-Dependent Bounds for Offline Reinforcement Learning with Preference Feedback”

A. Useful facts

In this section, we collate some useful facts that will be used in the subsequent proofs.

Definition A.1 (SubGaussian norm). A random variable X is *subGaussian* if it has a finite subGaussian norm denoted as $\|X\|_{\psi_2}$ and defined as

$$\|X\|_{\psi_2} = \inf \left\{ c > 0 : \mathbb{E} \left[\exp \left(\frac{X^2}{c^2} \right) \right] \leq 2 \right\} < +\infty.$$

Definition A.2 (Variance proxy). The variance proxy of a subGaussian random variable X is denoted as $\|X\|_{\text{vp}}^2$ and defined as

$$\|X\|_{\text{vp}}^2 := \inf \left\{ s^2 > 0 : \mathbb{E} [\exp((X - \mathbb{E}[X])t)] \leq e^{\frac{s^2 t^2}{2}}, \forall t > 0 \right\}.$$

Lemma A.3 (Linear combination of subGaussian random variables). *Let X_1, \dots, X_n be independent subGaussian random variables, where the variance proxy of X_i is σ_i^2 . Then, for any $a_1, \dots, a_n \in \mathbb{R}$, the random variable $\sum_{i=1}^n a_i X_i$ is a subGaussian random variable with variance proxy $\sigma^2 = \sum_{i=1}^n a_i^2 \sigma_i^2$.*

Lemma A.4. (Tail bound of subGaussian random variables) *Suppose X is subGaussian with variance proxy σ^2 . Then, for any $\epsilon > 0$, we have*

$$\Pr(X - \mathbb{E}[X] \geq \epsilon) \leq \exp(-\epsilon^2 / (2\sigma^2)),$$

and

$$\Pr(X - \mathbb{E}[X] \leq -\epsilon) \leq \exp(-\epsilon^2 / (2\sigma^2)),$$

Lemma A.5. (Adapted from Vershynin (2018, Proposition 2.5.2)) *For any subGaussian random variable X ,*

$$\|X\|_{\text{vp}} \leq C \|X\|_{\psi_2},$$

where $C \leq 6\sqrt{2e} \cdot (3\sqrt{\log 2} + 1)$. *If $\mathbb{E}[X] = 0$, then we have*

$$\|X\|_{\psi_2} \leq \sqrt{6} \|X\|_{\text{vp}}.$$

B. Proof of Proposition 2.3

Lemma B.1. *For any inconsistent instance $v = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta)$, there exists $(k, i) \in \mathcal{S} \times \mathcal{A}$ with $i \neq i_k^*$ such that*

$$\phi(k, i) - \phi(k, i_k^*) \notin \text{Span}\{\phi(k', i') - \phi(k', j') \mid N_{k', i', j'} > 0, (k', i', j') \in \mathcal{S} \times \mathcal{A}^2\}.$$

Proof. We prove this result by contradiction. Fix any inconsistent instance $v = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta)$. Assume that for all $(k, i) \in \mathcal{S} \times \mathcal{A}$ with $i \neq i_k^*$, it holds

$$\phi(k, i) - \phi(k, i_k^*) \in \text{Span}\{\phi(k', i') - \phi(k', j') \mid N_{k', i', j'} > 0, (k', i', j') \in \mathcal{S} \times \mathcal{A}^2\}. \quad (24)$$

By the fact that $\phi(k, i) - \phi(k, j) = (\phi(k, i) - \phi(k, i_k^*)) - (\phi(k, j) - \phi(k, i_k^*))$, Eqn. (24) implies that for any $(k, i, j) \in \mathcal{S} \times \mathcal{A} \times \mathcal{A}$ with $i \neq j$, it holds

$$\phi(k, i) - \phi(k, j) \in \text{Span}\{\phi(k', i') - \phi(k', j') \mid N_{k', i', j'} > 0, (k', i', j') \in \mathcal{S} \times \mathcal{A}^2\}, \quad (25)$$

which is a contradiction to that v is inconsistent. This completes the proof of Lemma B.1. \square

Proof of Proposition 2.3. Fix any inconsistent instance $v = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta)$. By Lemma B.1, there exists $(\underline{k}, \underline{i}) \in \mathcal{S} \times \mathcal{A}$ with $\underline{i} \neq i_{\underline{k}}^*$ such that

$$\phi(\underline{k}, \underline{i}) - \phi(\underline{k}, i_{\underline{k}}^*) \notin \text{Span}\{\phi(k', i') - \phi(k', j') \mid N_{k', i', j'} > 0, (k', i', j') \in \mathcal{S} \times \mathcal{A}^2\}.$$

That is, there exists a vector $\mathbf{z} \in \mathbb{R}^d$ such that

$$\langle \mathbf{z}, \phi(\underline{k}, \underline{i}) - \phi(\underline{k}, \underline{i}_k^*) \rangle = -2\langle \theta, \phi(\underline{k}, \underline{i}) - \phi(\underline{k}, \underline{i}_k^*) \rangle$$

and for all $(k', i', j') \in \mathcal{S} \times \mathcal{A}^2$ with $N_{k', i', j'} > 0$,

$$\langle \mathbf{z}, \phi(k', i') - \phi(k', j') \rangle = 0.$$

Finally, we let $\theta' = \theta + \mathbf{z}$, and instance $v' = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta')$. By the fact that $\langle \theta, \phi(k', i') - \phi(k', j') \rangle = \langle \theta', \phi(k', i') - \phi(k', j') \rangle$ for all $(k', i', j') \in \mathcal{S} \times \mathcal{A}^2$ with $N_{k', i', j'} > 0$ we get that for all $n \geq 1$

$$D_{\text{KL}}(P_v^{(n)}, P_{v'}^{(n)}) = 0.$$

Therefore, we get that $P_v^{(n)}$ is equal to $P_{v'}^{(n)}$. In addition, by definition of R_n , we get that

$$\begin{aligned} & \mathbb{E}_v^\pi[R_n] + \mathbb{E}_{v'}^\pi[R_n] \\ &= \mathbb{E}_v^\pi \left[\sum_{k \in \mathcal{S}} \rho_k \left(\max_{j \in \mathcal{A}} \langle \phi(k, j) - \phi(k, \hat{i}_k), \theta \rangle \right) \right] + \mathbb{E}_{v'}^\pi \left[\sum_{k \in \mathcal{S}} \rho_k \left(\max_{j \in \mathcal{A}} \langle \phi(k, j) - \phi(k, \hat{i}_k), \theta' \rangle \right) \right] \\ &\stackrel{(a)}{=} \mathbb{E}_v^\pi \left[\sum_{k \in \mathcal{S}} \rho_k \left(\max_{j \in \mathcal{A}} \langle \phi(k, j) - \phi(k, \hat{i}_k), \theta \rangle + \max_{j \in \mathcal{A}} \langle \phi(k, j) - \phi(k, \hat{i}_k), \theta' \rangle \right) \right] \\ &\geq \rho_{\underline{k}} \min_{i \in \mathcal{A}} \left[\max_{j \in \mathcal{A}} \langle \phi(\underline{k}, j) - \phi(\underline{k}, i), \theta \rangle + \max_{j \in \mathcal{A}} \langle \phi(\underline{k}, j) - \phi(\underline{k}, i), \theta' \rangle \right], \end{aligned} \quad (26)$$

where (a) follows from the fact that $P_v^{(n)}$ is equivalent with $P_{v'}^{(n)}$.

Further, by the definition of v and v' , we get that

$$\min_{i \in \mathcal{A}} \left[\max_{j \in \mathcal{A}} \langle \phi(\underline{k}, j) - \phi(\underline{k}, i), \theta \rangle + \max_{j \in \mathcal{A}} \langle \phi(\underline{k}, j) - \phi(\underline{k}, i), \theta' \rangle \right] > 0,$$

and recall that $\rho_{\underline{k}} > 0$. This completes the proof of Proposition 2.3 □

C. The existence of \hat{i}_k in RL-LOW and DP-RL-LOW

We first provide analytical solutions of $w^{(k, i, j)}$ and γ_{k, i, i_k^*} .

Lemma C.1. Fix any consistent instance $v = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta)$. Recall the definitions of $w^{(k, i, j)}$ and γ_{k, i, i_k^*} in Eqn. (7) and Eqn. (9), respectively. Then, for any $(k, i) \in \mathcal{S} \times \mathcal{A}$ with $i \neq i_k^*$,

$$\gamma_{k, i, i_k^*} = \|\phi(k, i_k^* - \phi(k, i))\|_{\mathcal{G}}^2_{V^{-1}} \quad (27)$$

and for any $(k, i, j) \in \mathcal{S} \times \mathcal{A}^2$ with $i \neq j$,

$$w_{k', i', j'}^{(k, i, j)} = N_{k', i', j'} [\phi(k', i') - \phi(k', j')]_{\mathcal{G}}^\top V^{-1} [\phi(k, i) - \phi(k, j)]_{\mathcal{G}} \quad (28)$$

where V and \mathcal{G} are as defined in Lemma 4.2.

Proof of Lemma C.1. Fix any $(k, i, j) \in \mathcal{S} \times \mathcal{A}^2$ with $i \neq j$. By definition, the optimization problem of Eqn. (7) is equivalent to

$$\min_{u \in \mathbb{R}^{\mathcal{S} \times \mathcal{A} \times \mathcal{A}}} \sum_{k' \in \mathcal{S}, i', j' \in \mathcal{A}: N_{k', i', j'} > 0} \frac{(u_{k', i', j'})^2}{N_{k', i', j'}} \quad (29)$$

subject to

$$[\phi(k, i) - \phi(k, j)]_{\mathcal{G}} = \sum_{k' \in \mathcal{S}, i', j' \in \mathcal{A}: N_{k', i', j'} > 0} u_{k', i', j'} [\phi(k', i') - \phi(k', j')]_{\mathcal{G}}.$$

The Lagrangian of the above constrained optimization problem is

$$L(u, \lambda) = \sum_{k' \in \mathcal{S}, i', j' \in \mathcal{A}: N_{k', i', j'} > 0} \frac{(u_{k', i', j'})^2}{N_{k', i', j'}} + \lambda^\top \left([\phi(k, i) - \phi(k, j)]_{\mathcal{G}} - \sum_{k' \in \mathcal{S}, i', j' \in \mathcal{A}: N_{k', i', j'} > 0} u_{k', i', j'} [\phi(k', i') - \phi(k', j')]_{\mathcal{G}} \right),$$

for $u \in \mathbb{R}^{\mathcal{S} \times \mathcal{A} \times \mathcal{A}}$ and $\lambda \in \mathbb{R}^{|\mathcal{G}|}$. Then, by solving $\frac{dL}{du} = 0$ and $\frac{dL}{d\lambda} = 0$ for all $(k', i', j') \in \mathcal{S} \times \mathcal{A}^2$ with $N_{k', i', j'} > 0$, we obtain that the minimum of (29) is

$$\lambda = -2V^{-1}[\phi(k, i) - \phi(k, j)]_{\mathcal{G}}$$

and

$$u_{k', i', j'} = N_{k', i', j'} [\phi(k', i') - \phi(k', j')]_{\mathcal{G}}^\top V^{-1} [\phi(k, i) - \phi(k, j)]_{\mathcal{G}} \quad (30)$$

for all $(k', i', j') \in \mathcal{S} \times \mathcal{A}^2$ with $N_{k', i', j'} > 0$. That is, for any $(k, i) \in \mathcal{S} \times \mathcal{A}$ with $i \neq i_k^*$,

$$\gamma_{k, i, i_k^*} = \|\phi(k, i_k^*) - \phi(k, i)\|_{V^{-1}}^2 \quad (31)$$

This completes the desired proof. \square

Then, we are ready to prove Proposition 3.2

Proof of Propostion 3.2. By Lemma C.1, we get that under RL-LOW, for any $(k, i, j, j_2) \in \mathcal{S} \times \mathcal{A}^3$,

$$\begin{aligned} & \hat{r}_{k, i, j} + \hat{r}_{k, j, j_2} \\ &= \sum_{(k', i', j') \in \mathcal{S} \times \mathcal{A}^2} N_{k', i', j'} [\phi(k', i') - \phi(k', j')]_{\mathcal{G}}^\top V^{-1} [\phi(k, i) - \phi(k, j)]_{\mathcal{G}} \log \left(\frac{B_{k', i', j'}^{\text{CLP}}}{1 - B_{k', i', j'}^{\text{CLP}}} \right) \\ & \quad + \sum_{(k', i', j') \in \mathcal{S} \times \mathcal{A}^2} N_{k', i', j'} [\phi(k', i') - \phi(k', j')]_{\mathcal{G}}^\top V^{-1} [\phi(k, j) - \phi(k, j_2)]_{\mathcal{G}} \log \left(\frac{B_{k', i', j'}^{\text{CLP}}}{1 - B_{k', i', j'}^{\text{CLP}}} \right) \\ &= \sum_{(k', i', j') \in \mathcal{S} \times \mathcal{A}^2} N_{k', i', j'} [\phi(k', i') - \phi(k', j')]_{\mathcal{G}}^\top V^{-1} [\phi(k, i) - \phi(k, j_2)]_{\mathcal{G}} \log \left(\frac{B_{k', i', j'}^{\text{CLP}}}{1 - B_{k', i', j'}^{\text{CLP}}} \right) \\ &= \hat{r}_{k, i, j_2}, \end{aligned} \quad (32)$$

which implies that $|\hat{\mathcal{I}}_k| \geq 1$ under RL-LOW. This completes the proof of Proposition 3.2. \square

Following the same lines as the proof of Propostion 3.2, we get the corollary below.

Corollary C.2. *For any consistent instance v and using the estimate of the best action under each state k as prescribed by RL-LOW-DP, we have \hat{i}_k exists.*

D. Proof of Lower Bound

Before we prove the lower bound, we first give a useful corollary. The following corollary is a direct result of Pinsker's inequality.

Corollary D.1. *Fix any $C \in (0, \frac{1}{2})$. For any $p, q \in (0, 1)$ with $\min(p, 1-p) \geq C$ and $\min(q, 1-q) \geq C$, we have*

$$2(p-q)^2 \leq d_{\text{KL}}(p, q) \leq \frac{2}{C}(p-q)^2$$

where $d_{\text{KL}}(p, q)$ denotes the KL divergence between the Bernoulli distributions with parameters of p and q .

Then, we prove Lemma 4.1 that reveals the KL divergence between instances. Recall that given any instance v , we let $P_v^{(n)}$ denote the distribution of $(\sigma_i)_{i=1}^n$. The following lemma gives an estimation of the KL divergence between instance v and v' that share the same parameters but θ .

Lemma 4.1. For any instance $v = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta)$ and $v' = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta')$, it holds that

$$2n \exp(-4R_{\max}) \leq \frac{D_{\text{KL}}(P_v^{(n)} \| P_{v'}^{(n)})}{\tilde{D}(v, v')} \leq 2n \exp(2R_{\max}) \quad (33)$$

where

$$\tilde{D}(v, v') = \sum_{k \in \mathcal{S}, i, j \in \mathcal{A}} N_{k, i, j} (\langle \phi(k, i) - \phi(k, j), \theta - \theta' \rangle)^2, \quad (34)$$

and where $R_{\max} = \max_{k \in \mathcal{S}, i, j \in \mathcal{A}} \max\{|\langle \phi(k, i), \theta \rangle|, |\langle \phi(k, i), \theta' \rangle|\}$ is the maximum absolute reward in these two instances.

Proof. Fix any consistent instance $v = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta)$. By the chain rule of the KL divergence, we have

$$\begin{aligned} & D_{\text{KL}}(P_v^{(n)} \| P_{v'}^{(n)}) \\ &= n \sum_{k \in \mathcal{S}, i, j \in \mathcal{A}} N_{k, i, j} d_{\text{KL}} \left(\frac{\exp(\langle \phi(k, i), \theta \rangle)}{\exp(\langle \phi(k, i), \theta \rangle) + \exp(\langle \phi(k, j), \theta \rangle)}, \right. \\ & \quad \left. \frac{\exp(\langle \phi(k, i), \theta' \rangle)}{\exp(\langle \phi(k, i), \theta' \rangle) + \exp(\langle \phi(k, j), \theta' \rangle)} \right) \\ &= n \sum_{k \in \mathcal{S}, i, j \in \mathcal{A}} N_{k, i, j} d_{\text{KL}}(\text{Sig}(\langle \phi(k, i) - \phi(k, j), \theta \rangle), \text{Sig}(\langle \phi(k, i) - \phi(k, j), \theta' \rangle)) \end{aligned} \quad (35)$$

where $\text{Sig}(\cdot)$ represents the Sigmoid function. By the fact that $-2R_{\max} \leq \langle \phi(k, i) - \phi(k, j), \theta \rangle \leq 2R_{\max}$ and $-2R_{\max} \leq \langle \phi(k, i) - \phi(k, j), \theta' \rangle \leq 2R_{\max}$ and that $\frac{d\text{Sig}(x)}{dx} = \frac{\exp(-x)}{(\exp(-x)+1)^2}$, we further get that

$$\begin{aligned} & |\text{Sig}(\langle \phi(k, i) - \phi(k, j), \theta \rangle) - \text{Sig}(\langle \phi(k, i) - \phi(k, j), \theta' \rangle)| \\ & \leq |\langle \phi(k, i) - \phi(k, j), \theta - \theta' \rangle| \end{aligned} \quad (36)$$

and

$$\begin{aligned} & |\text{Sig}(\langle \phi(k, i) - \phi(k, j), \theta \rangle) - \text{Sig}(\langle \phi(k, i) - \phi(k, j), \theta' \rangle)| \\ & \geq \exp(-2R_{\max}) |\langle \phi(k, i) - \phi(k, j), \theta - \theta' \rangle|. \end{aligned} \quad (37)$$

Then, by Corollary D.1, we get that Eqn. (36) and Eqn. (37) imply that

$$\begin{aligned} & d_{\text{KL}}(\text{Sig}(\langle \phi(k, i) - \phi(k, j), \theta \rangle), \text{Sig}(\langle \phi(k, i) - \phi(k, j), \theta' \rangle)) \\ & \leq 2 \exp(2R_{\max}) |\langle \phi(k, i) - \phi(k, j), \theta - \theta' \rangle|^2 \end{aligned} \quad (38)$$

and

$$\begin{aligned} & d_{\text{KL}}(\text{Sig}(\langle \phi(k, i) - \phi(k, j), \theta \rangle), \text{Sig}(\langle \phi(k, i) - \phi(k, j), \theta' \rangle)) \\ & \geq 2 \exp(-4R_{\max}) |\langle \phi(k, i) - \phi(k, j), \theta - \theta' \rangle|^2. \end{aligned} \quad (39)$$

Finally, combining Eqn. (35), Eqn. (38) and Eqn. (39), we complete the proof of Lemma 4.1. \square

Then, recall that we denote $\tilde{D}(\cdot, \cdot)$ as the approximation of KL divergence between instance v and v' that share the same parameter except θ , i.e., for any $v = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta)$ and $v' = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta')$,

$$\tilde{D}(v, v') := \sum_{k \in \mathcal{S}, i, j \in \mathcal{A}} N_{k, i, j} (\langle \phi(k, i) - \phi(k, j), \theta - \theta' \rangle)^2.$$

In addition, recall that for any $\mathbf{z} \in \mathbb{R}^d$, $\eta \in \mathbb{R}$ and consistent instance $v = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta)$, we denote $\text{Alt}(v, \mathbf{z}, \eta)$ as the set of instances that share the same parameter except θ and satisfy

$$\langle \mathbf{z}, \theta' - \theta \rangle = \eta$$

for all $v' = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta') \in \text{Alt}(v, \mathbf{z}, \eta)$.

We are ready to prove Lemma 4.2 that reveals a useful property for $\text{Alt}(\cdot)$ and $\tilde{D}(\cdot)$

Lemma 4.2. Let \mathcal{G} be an arbitrary orthonormal basis of $\text{Span}\{\phi(k', i') - \phi(k', j') : (k', i', j') \in \mathcal{S} \times \mathcal{A}^2 \text{ and } N_{k', i', j'} > 0\}$. Also let $[\mathbf{w}]_{\mathcal{G}}$ denote the column vector that represents \mathbf{w} under basis \mathcal{G} (Meyer, 2000, Chapter 4). Define the matrix

$$V := \sum_{k \in \mathcal{S}, i, j \in \mathcal{A}} N_{k, i, j} [\phi(k, i) - \phi(k, j)]_{\mathcal{G}} [\phi(k, i) - \phi(k, j)]_{\mathcal{G}}^{\top}.$$

Then for any consistent instance $v = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta)$, $\eta \in \mathbb{R}$, and $\mathbf{z} \in \text{Span}\{\phi(k', i') - \phi(k', j') : (k', i', j') \in \mathcal{S} \times \mathcal{A}^2 \text{ and } N_{k', i', j'} > 0\}$,

$$\min_{v' \in \text{Alt}(v, \mathbf{z}, \eta)} \tilde{D}(v, v') = \frac{\eta^2}{\|[\mathbf{z}]_{\mathcal{G}}\|_{V^{-1}}^2}. \quad (40)$$

and the minimum of (40) is attained in $v' = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta')$ with

$$\theta' = \theta - \frac{\eta}{\|[\mathbf{z}]_{\mathcal{G}}\|_{V^{-1}}^2} V^{-1} \mathbf{z}.$$

Proof. By definition, we equivalently write down the optimization problem of Eqn. (40) as follows.

$$\min_{\mathbf{x} \in \mathbb{R}^d} \sum_{k \in \mathcal{S}, i, j \in \mathcal{A}} N_{k, i, j} (\langle [\phi(k, i) - \phi(k, j)]_{\mathcal{G}}, [\mathbf{x}]_{\mathcal{G}} \rangle)^2 \quad (41)$$

subject to

$$\langle \mathbf{x}, \mathbf{z} \rangle = \eta.$$

The Lagrangian of the above constrained optimization problem is,

$$L(\mathbf{x}, \lambda) = \sum_{k \in \mathcal{S}, i, j \in \mathcal{A}} N_{k, i, j} (\langle [\phi(k, i) - \phi(k, j)]_{\mathcal{G}}, [\mathbf{x}]_{\mathcal{G}} \rangle)^2 + \lambda (\langle \mathbf{x}, \mathbf{z} \rangle - \eta).$$

By solving $\frac{dL}{d\lambda} = 0$ and $\frac{dL}{d\mathbf{x}_i} = 0$ for all $i \in [d]$, we attain the minimum of Eqn. (41) at

$$\begin{cases} \lambda = \frac{2\eta}{\|[\mathbf{z}]_{\mathcal{G}}\|_{V^{-1}}^2} \\ \mathbf{x} = -\frac{\eta}{\|[\mathbf{z}]_{\mathcal{G}}\|_{V^{-1}}^2} V^{-1} \mathbf{z} + \mathbf{g}, \end{cases}$$

where \mathbf{g} is any vector that is orthogonal with vector space $\text{Span}\{\phi(k', i') - \phi(k', j') : (k', i', j') \in \mathcal{S} \times \mathcal{A}^2 \text{ and } N_{k', i', j'} > 0\}$, which implies

$$L(\mathbf{z}, \lambda) = \frac{\eta^2}{\|[\mathbf{z}]_{\mathcal{G}}\|_{V^{-1}}^2}.$$

This completes the desired proof. \square

Lemma D.2. Fix any $v = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta) \in \mathcal{Q}$. Let

$$v' = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta') \in \underset{u \in \text{Alt}(v, \phi(\bar{k}(v), \bar{i}(v)) - \phi(\bar{k}(v), \hat{i}_k^*(v)), 2\Delta_{\bar{k}(v), \bar{i}(v)}(v))}{\text{argmin}} \tilde{D}(v, u).$$

Then, it holds

$$\begin{cases} \langle \phi(k, \hat{i}_k^*(v)) - \phi(k, i), \theta - \theta' \rangle = 2\Delta_{k, i}(v) & \text{for } (k, i) = (\bar{k}(v), \bar{i}(v)) \\ |\langle \phi(k, \hat{i}_k^*(v)) - \phi(k, i), \theta - \theta' \rangle| \leq \frac{1}{2}\Delta_{k, i}(v) & \forall i \neq \hat{i}_k^*(v) \text{ and } (k, i) \neq (\bar{k}(v), \bar{i}(v)) \end{cases} \quad (42)$$

Consequently,

$$\begin{cases} i_k^*(v) = i_k^*(v') & \forall k \neq \bar{k}(v) \\ \bar{k}(v) = \bar{k}(v') \\ i_{\bar{k}(v')}^*(v') = \bar{i}(v) \end{cases} \quad (43)$$

Proof. Fix any v and v' as in Lemma D.2. By the definition of v' , we can immediately obtain the first inequality of Eqn. (42), i.e.,

$$\langle \phi(k, i_k^*(v)) - \phi(k, i), \theta - \theta' \rangle = 2\Delta_{k,i}(v) \text{ for } (k, i) = (\bar{k}(v), \bar{i}(v)). \quad (44)$$

Then, we will use proof by contradiction to prove the second inequality of Eqn. (42).

Assume that there exists $(\tilde{k}, \tilde{i}) \in \mathcal{S} \times \mathcal{A}$ with for all $\tilde{i} \neq i_{\tilde{k}}^*$ and $(\tilde{k}, \tilde{i}) \neq (\bar{k}(v), \bar{i}(v))$ such that

$$|\langle \phi(\tilde{k}, i_{\tilde{k}}^*) - \phi(\tilde{k}, \tilde{i}), \theta - \theta' \rangle| > \frac{\Delta_{\tilde{k}, \tilde{i}}(v)}{2} \quad (45)$$

Then, we let $v'' = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta'')$, where

$$\theta'' = \theta + \frac{(\theta - \theta')2\Delta_{\tilde{k}, \tilde{i}}(v)}{\langle \phi(\tilde{k}, i_{\tilde{k}}^*) - \phi(\tilde{k}, \tilde{i}), \theta - \theta' \rangle},$$

which implies that $v'' \in \text{Alt}(v, \phi(\tilde{k}, \tilde{i}) - \phi(\tilde{k}, i_{\tilde{k}}^*), 2\Delta_{\tilde{k}, \tilde{i}}(v))$. Note that from Lemma 4.2, we have

$$\tilde{D}(v, v') = \frac{4\Delta_{\bar{k}(v), \bar{i}(v)}(v)^2}{\|[\phi(\bar{k}(v), \bar{i}(v)) - \phi(\bar{k}(v), i_{\bar{k}(v)}^*)]_{\mathcal{G}}\|_{V^{-1}}^2}, \quad (46)$$

where V and \mathcal{G} are as defined in Lemma 4.2. In addition, by definition, we also can get that

$$\begin{aligned} \tilde{D}(v, v'') &= \left(\frac{2\Delta_{\tilde{k}, \tilde{i}}(v)}{\langle \phi(\tilde{k}, i_{\tilde{k}}^*) - \phi(\tilde{k}, \tilde{i}), \theta - \theta' \rangle} \right)^2 \cdot \tilde{D}(v, v') \\ &\stackrel{(a)}{<} 16\tilde{D}(v, v') \\ &\stackrel{(b)}{<} \frac{16\Delta_{\tilde{k}(v), \tilde{i}(v)}(v)^2}{\|[\phi(\bar{k}(v), \bar{i}(v)) - \phi(\bar{k}(v), i_{\bar{k}(v)}^*)]_{\mathcal{G}}\|_{V^{-1}}^2} \\ &\stackrel{(c)}{<} \frac{64}{\gamma_{\bar{k}(v), \bar{i}(v)}(v)} \end{aligned} \quad (47)$$

where (a) follows from Eqn. (45) and (b) follows from Eqn. (46), and (c) follows from Lemma C.1.

Similarly, from Lemma 4.2, we get that

$$\begin{aligned} \min_{u \in \text{Alt}(v, \phi(\tilde{k}, \tilde{i}) - \phi(\tilde{k}, i_{\tilde{k}}^*), 2\Delta_{\tilde{k}, \tilde{i}}(v))} \tilde{D}(v, u) &= \frac{4\Delta_{\tilde{k}, \tilde{i}}(v)^2}{\|[\phi(\tilde{k}, \tilde{i}) - \phi(\tilde{k}, i_{\tilde{k}}^*)]_{\mathcal{G}}\|_{V^{-1}}^2} \\ &= \frac{4}{\gamma_{\tilde{k}, \tilde{i}}(v)} \end{aligned} \quad (48)$$

By the fact that $v'' \in \text{Alt}(v, \phi(\tilde{k}, \tilde{i}) - \phi(\tilde{k}, i_{\tilde{k}}^*))$ as well as Eqn. (24) and Eqn. (25), we further get that

$$\frac{4}{\gamma_{\tilde{k}, \tilde{i}}(v)} < \frac{64}{\gamma_{\bar{k}(v), \bar{i}(v)}(v)}.$$

That is, $16\gamma_{\tilde{k}, \tilde{i}}(v) > \gamma_{\bar{k}(v), \bar{i}(v)}(v)$, which contradicts the fact that $v \in \mathcal{Q}$.

Hence, for all $(k, i) \in \mathcal{S} \times \mathcal{A}$ with $\forall \bar{i} \neq i_k^*$ and $(k, i) \neq (\bar{k}(v), \bar{i}(v))$, we have

$$|\langle \phi(k, i_k^*(v)) - \phi(\bar{k}(v), \bar{i}(v)), \theta - \theta' \rangle| \geq 2\Delta_{k,i}(v)$$

Consequently, we have

$$\begin{cases} i_k^*(v) = i_k^*(v') & \forall k \neq \bar{k}(v) \\ \bar{k}(v) = \bar{k}(v') \\ i_{\bar{k}(v')}^*(v') = \bar{i}(v) \end{cases}$$

This completes the proof of Lemma D.2. \square

Lemma D.3. Fix any $v = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta) \in \mathcal{Q}$. Let

$$v' = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta') \in \underset{u \in \text{Alt}(v, \phi(\bar{k}(v), \bar{i}(v)) - \phi(\bar{k}(v), i_{\bar{k}(v)}^*), 2\Delta_{\bar{k}(v), \bar{i}(v)}(v))}{\text{argmin}} \tilde{D}(v, u).$$

Then, it holds

$$H(v) \leq H(v') \leq 8H(v).$$

Proof. Fix any v and v' as in Lemma D.2. Note that from Lemma D.2, we get that instance v and instance v' share the same optimal action for the state $k \neq \bar{k}$ and the optimal action of state $\bar{k}(v)$ is $\bar{i}(v)$ under the instance v' , i.e., $i_{\bar{k}(v')}^*(v') = \bar{i}(v)$ and $i_k^*(v) = i_k^*(v')$ for all $k \neq \bar{k}(v)$. Then, by definition, the hardness of instance v' is

$$\begin{aligned} H(v') &= \max_{k \in \mathcal{S}, i \in \mathcal{A}: i \neq i_k^*(v')} \frac{\gamma_{k,i}(v')}{\Delta_{k,i}^2(v')} \\ &\stackrel{(a)}{=} \max \left\{ \max_{k \in \mathcal{S}, i \in \mathcal{A}: i \neq i_k^*(v'), k \neq \bar{k}(v)} \frac{\gamma_{k,i}(v)}{\Delta_{k,i}^2(v')}, \max_{i \in \mathcal{A}: i \neq \bar{i}(v), k = \bar{k}(v)} \frac{\gamma_{k,i}(v')}{\Delta_{k,i}^2(v')} \right\} \end{aligned}$$

where (a) follows the fact that $i_k^*(v) = i_k^*(v')$ for all $k \neq \bar{k}(v)$ and both instances v and v' share the same parameters of ϕ and N , which implies $\gamma_{k,i}(v) = \gamma_{k,i}(v')$ for all $k \neq \bar{k}(v)$.

Then, by the fact that $|\langle \phi(k, i_k^*(v)) - \phi(k, i), \theta - \theta' \rangle| \leq \frac{1}{2}\Delta_{k,i}(v)$ for all $i \neq i_k^*(v)$ and $(k, i) \neq (\bar{k}(v), \bar{i}(v))$ from Lemma D.2, we get that

$$\begin{aligned} \max_{k \in \mathcal{S}, i \in \mathcal{A}: i \neq i_k^*(v'), k \neq \bar{k}(v)} \frac{\gamma_{k,i}(v)}{\Delta_{k,i}^2(v')} &\leq \max_{k \in \mathcal{S}, i \in \mathcal{A}: i \neq i_k^*(v'), k \neq \bar{k}(v)} \frac{4\gamma_{k,i}(v)}{\Delta_{k,i}^2(v)} \\ &\leq 4H(v) \end{aligned} \tag{49}$$

Further, by the fact that $\langle \phi(k, i_k^*(v)) - \phi(k, i), \theta - \theta' \rangle = 2\Delta_{k,i}(v)$ for $(k, i) = (\bar{k}(v), \bar{i}(v))$ Lemma D.2, we get that

$$\begin{aligned} &\max_{i \in \mathcal{A}: i \neq \bar{i}(v), k = \bar{k}(v)} \frac{\gamma_{k,i}(v')}{\Delta_{k,i}^2(v')} \\ &= \max \left\{ \frac{\gamma_{\bar{k}(v), \bar{i}(v)}(v)}{\Delta_{\bar{k}(v), \bar{i}(v)}^2(v)}, \max_{i \in \mathcal{A}: i \neq \bar{i}(v), i \neq i_{\bar{k}(v)}^*(v)} \frac{\gamma_{\bar{k}(v), i}(v')}{\Delta_{\bar{k}(v), i}^2(v')} \right\} \\ &= \max \left\{ H(v), \max_{i \in \mathcal{A}: i \neq \bar{i}(v), i \neq i_{\bar{k}(v)}^*(v)} \frac{\gamma_{\bar{k}(v), i}(v')}{\Delta_{\bar{k}(v), i}^2(v')} \right\} \\ &\stackrel{(a)}{=} \max \left\{ H(v), \max_{i \in \mathcal{A}: i \neq \bar{i}(v), i \neq i_{\bar{k}(v)}^*(v)} \frac{\|[\phi(\bar{k}(v), \bar{i}(v)) - \phi(\bar{k}(v), i)]_{\mathcal{G}}\|_{V^{-1}}^2}{\Delta_{\bar{k}(v), i}^2(v')} \right\} \end{aligned} \tag{50}$$

where (a) follows from Lemma C.1, V and \mathcal{G} are as defined in Lemma 4.2. Similarly, by the fact that $|\langle \phi(k, i_k^*(v)) - \phi(k, i), \theta - \theta' \rangle| \leq \frac{1}{2}\Delta_{k,i}(v) \forall i \neq i_k^*(v)$ and $(k, i) \neq (\bar{k}(v), \bar{i}(v))$, we obtain for $i \in \mathcal{A}$ with $i \neq \bar{i}(v), i \neq i_{\bar{k}(v)}^*(v)$

$$\frac{\|[\phi(\bar{k}(v), \bar{i}(v)) - \phi(\bar{k}(v), i)]_{\mathcal{G}}\|_{V^{-1}}^2}{\Delta_{\bar{k}(v), i}^2(v')} \leq \frac{4\|[\phi(\bar{k}(v), \bar{i}(v)) - \phi(\bar{k}(v), i)]_{\mathcal{G}}\|_{V^{-1}}^2}{(\Delta_{\bar{k}(v), i}(v) \vee \Delta_{\bar{k}(v), \bar{i}(v)}(v))^2}$$

$$\begin{aligned}
 &= \frac{4\|\phi(\bar{k}(v), \bar{i}(v)) - \phi(\bar{k}(v), i_{\bar{k}(v)}^*(v)) + \phi(\bar{k}(v), i_{\bar{k}(v)}^*(v)) - \phi(\bar{k}(v), i)\|_{\mathcal{G}}\|_{V^{-1}}^2}{(\Delta_{\bar{k}(v), i}(v) \vee \Delta_{\bar{k}(v), \bar{i}(v)}(v))^2} \\
 &\leq \frac{4(\|\phi(\bar{k}(v), \bar{i}(v)) - \phi(\bar{k}(v), i_{\bar{k}(v)}^*(v))\|_{\mathcal{G}}\|_{V^{-1}}^2 + \|\phi(\bar{k}(v), i_{\bar{k}(v)}^*(v)) - \phi(\bar{k}(v), i)\|_{\mathcal{G}}\|_{V^{-1}}^2)}{(\Delta_{\bar{k}(v), i}(v) \vee \Delta_{\bar{k}(v), \bar{i}(v)}(v))^2} \\
 &\leq 8H(v).
 \end{aligned} \tag{51}$$

Finally, combining Eqn. (49), Eqn. (50) and Eqn. (51), we complete the proof of Lemma D.3. \square

With the ingredients of the above lemmas, we are ready to prove Theorem 4.3.

Proof of Theorem 4.3. Fix any $v = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta) \in \mathcal{Q}$ and

$$v' = (\rho, \mathcal{S}, \mathcal{A}, \phi, N, \theta') \in \underset{u \in \text{Alt}(v, \phi(\bar{k}(v), \bar{i}(v)) - \phi(\bar{k}(v), i_{\bar{k}(v)}^*(v)), 2\Delta_{\bar{k}(v), \bar{i}(v)}(v))}{\text{argmin}} \tilde{D}(v, u).$$

By Lemma D.3, we get that $H(v') \leq H(v) \leq 8H(v')$. By Lemma 4.2, we get that

$$\begin{aligned}
 \tilde{D}(v, v') &= \frac{\|\phi(\bar{k}(v), i_{\bar{k}(v)}^*(v)) - \phi(\bar{k}(v), \bar{i}(v))\|_{\mathcal{G}}\|_{V^{-1}}^2}{4\Delta_{\bar{k}(v), \bar{i}(v)}^2(v)} \\
 &= \frac{1}{4H(v)},
 \end{aligned}$$

where \mathcal{G} and V are as defined in Lemma 4.2. Further, by Lemma 4.1 and Lemma D.2, we get that for any $n > 0$

$$\begin{aligned}
 D_{\text{KL}}(P_v^{(n)} \| P_{v'}^{(n)}) &\leq 2 \exp(2L) \cdot n \tilde{D}(v, v') \\
 &= \frac{n}{2 \exp(-2L) H(v)}.
 \end{aligned} \tag{52}$$

Then, we let $\bar{\Delta}$ be the minimum suboptimality of state $\bar{k}(v)$ in both instances v and v' , i.e.,

$$\bar{\Delta} = \min_{i \in \mathcal{A}} \left[\max_{j \in \mathcal{A}} \langle \phi(\bar{k}(v), j) - \phi(\bar{k}(v), i), \theta \rangle + \max_{j \in \mathcal{A}} \langle \phi(\bar{k}(v), j) - \phi(\bar{k}(v), i), \theta' \rangle \right]$$

By the definitions of v and v' , we can obtain that $\bar{\Delta} > 0$. Then, we get that for any policy π

$$\begin{aligned}
 \mathbb{E}_v^\pi[R_n] + \mathbb{E}_{v'}^\pi[R_n] &\geq \rho_{\bar{k}(v)} \bar{\Delta} (1 - D_{\text{TV}}(P_v^{(n)}, P_{v'}^{(n)})) \\
 &\stackrel{(a)}{\geq} \frac{1}{2} \rho_{\bar{k}(v)} \bar{\Delta} \exp(-D_{\text{KL}}(P_v^{(n)} \| P_{v'}^{(n)})) \\
 &\stackrel{(b)}{\geq} \frac{1}{2} \rho_{\bar{k}(v)} \bar{\Delta} \exp\left(-\frac{n}{2 \exp(-2L) H(v)}\right),
 \end{aligned} \tag{53}$$

where $D_{\text{TV}}(\cdot, \cdot)$ denotes the total variance distance, and (a) follows from Bretagnolle–Huber inequality (Tsybakov, 2009, Lemma 2.6) and (c) follows from Eqn. (52). Finally, for all sufficiently large n , we have

$$\mathbb{E}_v^\pi[R_n] + \mathbb{E}_{v'}^\pi[R_n] \geq \exp\left(-\frac{n}{C_{\text{lo}} \cdot H(v)}\right),$$

which completes the proof of Theorem 4.3. \square

E. Proof of Upper bound

Lemma E.1. Let Y_n be a random variable sampled from the Binomial distribution with n trials and probability of success $p \in [1 - \beta, \beta]$ for $\beta \in (\frac{1}{2}, 1)$. Then,

$$|\mathbb{E}(f(X_n)) - f(\mathbb{E}(X_n))| \leq \frac{3}{(1 - \beta)^4 \sqrt{n}}$$

where $X_n = \frac{Y_n}{n}$ and

$$f(x) = \begin{cases} \log(\frac{x}{1-x}) & \text{if } x \in (1 - \beta, \beta) \\ \log(\frac{\beta}{1-\beta}) & \text{if } x \geq \beta \\ \log(\frac{1-\beta}{\beta}) & \text{if } x \leq 1 - \beta. \end{cases} \quad (54)$$

Proof of Lemma E.1. For simplicity of notation, we let $x_0 := \mathbb{E}(X_n)$, which implies $x_0 = p \in [1 - \beta, \beta]$. Then, by the smoothness of $f(\cdot)$ on $[1 - \beta, \beta]$, we obtain the equivalent expression of $f(\cdot)$ from the Taylor expansion of $f(\cdot)$ on $[1 - \beta, \beta]$,

$$f(x) = \begin{cases} f(x_0) + f'(x_0) \cdot (x - x_0) + \frac{1}{2} f''(\xi_x) \cdot (x - x_0)^2 & \text{if } x \in [1 - \beta, \beta] \\ f(x_0) + f'(x_0) \cdot (\beta - x_0) + \frac{1}{2} f''(\xi_\beta) \cdot (\beta - x_0)^2 & \text{if } x > \beta \\ f(x_0) + f'(x_0) \cdot (1 - \beta - x_0) + \frac{1}{2} f''(\xi_\beta) \cdot (1 - \beta - x_0)^2 & \text{if } x < 1 - \beta, \end{cases} \quad (55)$$

where $\xi_x \in (\min(x, x_0), \max(x, x_0))$ only depends on x and x_0 in the Tolor expansion.

By using the fact that $f'(x_0) \cdot (\beta - x_0) = f'(x_0) \cdot (\beta - x) + f'(x_0) \cdot (x - x_0)$ and $f'(x_0) \cdot (1 - \beta - x_0) = f'(x_0) \cdot (1 - \beta - x) + f'(x_0) \cdot (x - x_0)$, we obtain from Eqn. (55) that

$$\begin{aligned} & |\mathbb{E}(f(X_n)) - f(\mathbb{E}(X_n))| \\ & \leq \sup_{x \in (1 - \beta, \beta)} |f''(x)| \cdot \text{Var}(X_n) + |f'(x_0)| \cdot \mathbb{E}([X_n - \beta]_+) + |f'(x_0)| \cdot \mathbb{E}([1 - \beta - X_n]_+) \end{aligned} \quad (56)$$

where $\text{Var}(\cdot)$ represents the variance of the random variable \cdot and $[x]_+ = \max\{x, 0\}$.

In addition, we note that

$$\begin{aligned} \mathbb{E}([X_n - \beta]_+) & \leq \int_{\beta}^1 (x - \beta) \mathbb{P}(X_n \geq x) dx \\ & \stackrel{(a)}{\leq} \int_{\beta}^1 (x - \beta) \exp(-2n(x - x_0)^2) dx \\ & \stackrel{(b)}{\leq} \int_{\beta}^1 (x - \beta) \exp(-n(x - \beta)^2) dx \\ & \leq \int_{\beta}^1 \frac{1}{\sqrt{n}} dx \\ & = \frac{1 - \beta}{\sqrt{n}}, \end{aligned} \quad (57)$$

where (a) follows from the fact X_n is a subGaussian random variable with variance proxy $\frac{1}{4n}$, and (b) follows from the fact that $(x - \beta) \exp(-n(x - \beta)^2) \leq \frac{1}{\sqrt{n}}$. Similarly, we also can get that

$$\mathbb{E}([1 - \beta - X_n]_+) \leq \frac{1 - \beta}{\sqrt{n}}. \quad (58)$$

Finally,

$$|\mathbb{E}(f(X_n)) - f(\mathbb{E}(X_n))|$$

$$\begin{aligned}
 &\stackrel{(a)}{\leq} \sup_{x \in (1-\beta, \beta)} |f''(x)| \cdot \text{Var}(X_n) + |f'(x_0)| \cdot \mathbb{E}([X_n - \beta]_+) + |f'(x_0)| \cdot \mathbb{E}([1 - \beta - X_n]_+) \\
 &\stackrel{(b)}{\leq} \frac{1}{(1-\beta)^4} \cdot \text{Var}(X_n) + f'(x_0) \cdot \mathbb{E}([X_n - \beta]_+) + f'(x_0) \cdot \mathbb{E}([1 - \beta - X_n]_+) \\
 &\stackrel{(c)}{\leq} \frac{1}{(1-\beta)^4} \cdot \frac{1}{n} + f'(x_0) \cdot \mathbb{E}([X_n - \beta]_+) + f'(x_0) \cdot \mathbb{E}([1 - \beta - X_n]_+) \\
 &\stackrel{(d)}{\leq} \frac{1}{(1-\beta)^4} \cdot \frac{1}{n} + \frac{1}{(1-\beta)^2} \cdot \mathbb{E}([X_n - \beta]_+) + \frac{1}{(1-\beta)^2} \cdot \mathbb{E}([1 - \beta - X_n]_+) \\
 &\stackrel{(e)}{\leq} \frac{1}{(1-\beta)^4} \cdot \frac{1}{n} + \frac{1}{(1-\beta)^2} \cdot \frac{2}{\sqrt{n}} \\
 &\leq \frac{3}{(1-\beta)^4 \sqrt{n}}, \tag{59}
 \end{aligned}$$

where (a) follows from Eqn. (56), (b) follows from $\text{Var}(X_n) \leq \frac{1}{n}$, (c) follows from $f''(x) = \frac{2x-1}{(x-1)^2 x^2}$ for $x \in (\beta, 1-\beta)$, (d) follows from $f'(x) = \frac{1}{x-x^2}$ for $x \in (\beta, 1-\beta)$, and (e) follows from Eqn. (57) and Eqn. (58). \square

Lemma E.2. Let Y_n be a random variable sampled from the binomial distribution with parameters $n > 0$ and $p \in [1-\beta, \beta]$ for $\beta \in (\frac{1}{2}, 1)$. Let $X_n = \frac{Y_n}{n}$. Then, $f(X_n)$ is subGaussian with

$$\|f(X_n)\|_{\text{vp}}^2 \leq \frac{C}{n},$$

where $C \leq (6\sqrt{2e} \cdot (3\sqrt{\log 2} + 1))^2 \cdot \frac{3}{2(1-\beta)^4}$ and $f(\cdot)$ is defined in Eqn. (54).

Proof. By the definition of X_n , we get that

$$\|X_n - p\|_{\text{vp}}^2 \leq \frac{1}{4n}.$$

Then, by Lemma A.5, we get that

$$\|X_n - p\|_{\phi_2}^2 \leq \frac{3}{2n}.$$

By the fact that $|f'(x)| \leq \frac{1}{(1-\beta)^2}$ for $x \in (1-\beta, \beta)$ and $f'(x) = 0$ for $x \in (0, 1-\beta) \cup (\beta, 1)$, we get that

$$\|f(X_n) - f(p)\|_{\phi_2}^2 \leq \frac{3}{2n(1-\beta)^4}. \tag{60}$$

Similarly, from Lemma A.5, we further get that

$$\|f(X_n) - f(p)\|_{\text{vp}}^2 \leq (6\sqrt{2e} \cdot (3\sqrt{\log 2} + 1))^2 \cdot \frac{3}{2n(1-\beta)^4}. \tag{61}$$

Note that by definition, we have $\|f(X_n) - f(p)\|_{\text{vp}}^2 = \|f(X_n)\|_{\text{vp}}^2$, which completes the proof of Lemma E.2. \square

With the ingredients of the above lemmas, we are ready to prove the upper bound of our algorithms

Proof of Theorem 3.3. Fix any consistent instance v and dataset size n under RL-LOW. By Lemma E.2 and Lemma A.3, we get that $\hat{r}_{k, i_k^*, i}$ is subGaussian with variance proxy as

$$\|\hat{r}_{k, i_k^*, i}\|_{\text{vp}}^2 \leq \frac{C\gamma_{k, i}}{n}, \tag{62}$$

where $C \leq (6\sqrt{2e} \cdot (3\sqrt{\log 2} + 1))^2 \cdot \frac{3}{2(1-\beta)^4}$ and $\beta = \frac{\exp(2L)}{1+\exp(2L)}$.

In addition, by Lemma E.1, we get that for any $(k', i', j') \in \mathcal{S} \times \mathcal{A} \times \mathcal{A}$ with $N_{k', i', j'} > 0$,

$$|\mathbb{E}_v^{\text{RL-LOW}} [f(B_{k', i', j'})] - f(\mathbb{E}_v^{\text{RL-LOW}} [B_{k', i', j'}])| \leq \frac{3}{(1-\beta)^4 \sqrt{n \cdot N_{k', i', j'}}},$$

which implies for any $(k, i) \in \mathcal{S} \times \mathcal{A}$ with $i \neq i_k^*$,

$$\begin{aligned} & \left| \sum_{k' \in \mathcal{S}, i' \in \mathcal{A}, j' \in \mathcal{A}} w_{k', i', j'}^{(k, i_k^*, i)} \mathbb{E}_v^{\text{RL-LOW}} [f(B_{k', i', j'})] - f(\mathbb{E}_v^{\text{RL-LOW}} [\hat{r}_{k, i_k^*, i}]) \right| \\ & \leq \sum_{k' \in \mathcal{S}, i' \in \mathcal{A}, j' \in \mathcal{A}: N_{k', i', j'} > 0} \frac{3|w_{k', i', j'}^{(k, i_k^*, i)}|}{(1-\beta)^4 \sqrt{n \cdot N_{k', i', j'}}}. \end{aligned} \quad (63)$$

That is,

$$|\mathbb{E}_v^{\text{RL-LOW}} [\hat{r}_{k, i_k^*, i}] - r_{k, i_k^*, i}| \leq \sum_{k' \in \mathcal{S}, i' \in \mathcal{A}, j' \in \mathcal{A}: N_{k', i', j'} > 0} \frac{3|w_{k', i', j'}^{(k, i_k^*, i)}|}{(1-\beta)^4 \sqrt{n \cdot N_{k', i', j'}}} \quad (64)$$

$$= \frac{3\tilde{\gamma}_{k, i}}{(1-\beta)^4 \sqrt{n}}, \quad (65)$$

where we denote $r_{k, i, j} := r_{k, i} - r_{k, j}$.

Note that by the definition of R_n we have

$$R_n \leq \sum_{(k, i) \in \mathcal{S} \times \mathcal{A}: i \neq i_k^*} \mathbb{1}_{\{\hat{r}_{k, i_k^*, i} \leq 0\}} \cdot \rho_k \Delta_{k, i} \quad (66)$$

$$\begin{aligned} & \leq \sum_{(k, i) \in \mathcal{S} \times \mathcal{A}: i \neq i_k^*} \left[\left(\mathbb{1}_{\{\hat{r}_{k, i_k^*, i} \leq \mathbb{E}_v^{\text{RL-LOW}} [\hat{r}_{k, i_k^*, i}] - \Delta_{k, i}/2\}} \right. \right. \\ & \quad \left. \left. \vee \mathbb{1}_{\{\mathbb{E}_v^{\text{RL-LOW}} [\hat{r}_{k, i_k^*, i}] - \Delta_{k, i} \geq \Delta_{k, i}/2\}} \right) \cdot \rho_k \Delta_{k, i} \right] \end{aligned} \quad (67)$$

By Lemma A.4 and Eqn. (62), we get that for any $(k, i) \in \mathcal{S} \times \mathcal{A}$ with $i \neq i_k^*$,

$$\mathbb{P}_v^{\text{RL-LOW}} \left(\hat{r}_{k, i_k^*, i} \leq \mathbb{E}_v^{\text{RL-LOW}} [\hat{r}_{k, i_k^*, i}] - \Delta_{k, i}/2 \right) \leq \exp \left(-\frac{2n\Delta_{k, i}^2}{C\gamma_{k, i}} \right), \quad (68)$$

and by Eqn. (65) we get that

$$|\mathbb{E}_v^{\text{RL-LOW}} [\hat{r}_{k, i_k^*, i}] - r_{k, i_k^*, i}| \leq \Delta_{k, i}/2$$

for all $n > \frac{18\tilde{\gamma}_{k, i}^2}{(1-\beta)^8 \Delta_{k, i}}$. That is, for all sufficiently large n , we have

$$\mathbb{E}_v^{\text{RL-LOW}} [R_n] \leq \sum_{k \in \mathcal{S}, i \in \mathcal{A}: i \neq i_k^*} \rho_k \Delta_{k, i} \exp \left(-\frac{2n\Delta_{k, i}^2}{C\gamma_{k, i}} \right),$$

which further implies that for all sufficiently large n , we have

$$\mathbb{E}_v^{\text{RL-LOW}} [R_n] \leq \exp \left(-\frac{n}{C_{\text{up}} \cdot H(v)} \right).$$

This completes the proof of Theorem 3.3. \square

In addition, we present the proof of Proposition 3.4 as follows.

Proof of Proposition 3.4. Note that by the definition of R_n we have

$$R_n \leq \sum_{(k, i) \in \mathcal{S} \times \mathcal{A}: i \neq i_k^*} \mathbb{1}_{\{\hat{r}_{k, i_k^*, i} \leq 0\}} \cdot \rho_k \Delta_{k, i}$$

$$\begin{aligned}
 &\stackrel{(a)}{\leq} \sum_{(k,i) \in \mathcal{S} \times \mathcal{A}: i \neq i_k^*} \rho_k \left[\mathbb{1}_{\{\Delta_{k,i} < \frac{6\tilde{\gamma}_{k,i}}{(1-\beta)^4\sqrt{n}}\}} \cdot \frac{6\tilde{\gamma}_{k,i}}{(1-\beta)^4\sqrt{n}} \right. \\
 &\quad \left. + \mathbb{1}_{\{|\mathbb{E}_v^{\text{RL-LOW}}[\hat{r}_{k,i,i_k^*}] - r_{k,i,i_k^*}| \leq \Delta_{k,i}/2\}} \cdot \Delta_{k,i} \right]
 \end{aligned}$$

where $\beta = \frac{\exp(2L)}{1+\exp(2L)}$, and (a) follows from Eqn. (65).

Hence, by Eqn. (68), we have

$$\mathbb{E}_v^{\text{RL-LOW}}[R_n] \leq \sum_{(k,i) \in \mathcal{S} \times \mathcal{A}: i \neq i_k^*} \rho_k \left[\frac{6\tilde{\gamma}_{k,i}}{(1-\beta)^4\sqrt{n}} + \exp\left(-\frac{2n\Delta_{k,i}^2}{C\gamma_{k,i}}\right) \Delta_{k,i} \right] \quad (69)$$

$$\leq \sum_{(k,i) \in \mathcal{S} \times \mathcal{A}: i \neq i_k^*} \rho_k \left[\frac{6\tilde{\gamma}_{k,i}}{(1-\beta)^4\sqrt{n}} + \sqrt{\frac{C\gamma_{k,i}}{2n}} \right] \quad (70)$$

$$= \frac{1}{\sqrt{n}} \sum_{(k,i) \in \mathcal{S} \times \mathcal{A}: i \neq i_k^*} \rho_k \left[\frac{6\tilde{\gamma}_{k,i}}{(1-\beta)^4} + \sqrt{\frac{C\gamma_{k,i}}{2}} \right] \quad (71)$$

where $C \leq (6\sqrt{2e} \cdot (3\sqrt{\log 2} + 1))^2 \cdot \frac{3}{2(1-\beta)^4}$.

This completes the desired proof. \square

F. Theoretical analysis of DP-RL-LOW

Lemma F.1. Fix any $\varepsilon > 0$ and $\delta > 0$. Let Y_n be a random variable sampled from the binomial distribution with number of trials $n \in \mathbb{N}$ and probability of success $p \in [1 - \beta, \beta]$ for $\beta \in (\frac{1}{2}, 1)$. Then,

$$|\mathbb{E}(f(\tilde{X}_n)) - f(\mathbb{E}(\tilde{X}_n))| \leq \frac{3}{(1-\beta)^4\sqrt{2n}} + \frac{4\sqrt{\log(1.25/\delta)}}{(1-\beta)^4(\varepsilon n)},$$

where $\tilde{X}_n = \frac{Y_n}{n} + \tilde{\xi}_n$, and $\tilde{\xi}_n$ is an independent Gaussian noise with zero mean and variance of $\frac{2\log(1.25/\delta)}{(\varepsilon n)^2}$, and $f(\cdot)$ is defined in (54).

Proof of Lemma F.1. For simplicity of notation, we let $x_0 := \mathbb{E}(\tilde{X}_n)$, which implies $x_0 = p \in [1 - \beta, \beta]$. Again, by the smoothness of $f(\cdot)$ on $[1 - \beta, \beta]$, we get the Taylor expansion of $f(\cdot)$ on $[1 - \beta, \beta]$,

$$f(x) = \begin{cases} f(x_0) + f'(x_0) \cdot (x - x_0) + \frac{1}{2}f''(\xi_x) \cdot (x - x_0)^2 & \text{if } x \in [1 - \beta, \beta] \\ f(x_0) + f'(x_0) \cdot (\beta - x_0) + \frac{1}{2}f''(\xi_\beta) \cdot (\beta - x_0)^2 & \text{if } x > \beta \\ f(x_0) + f'(x_0) \cdot (1 - \beta - x_0) + \frac{1}{2}f''(\xi_\beta) \cdot (1 - \beta - x_0)^2 & \text{if } x < 1 - \beta, \end{cases} \quad (72)$$

where $\xi_x \in (\min(x, x_0), \max(x, x_0))$ that only depends on x and x_0 in the Talor expansion.

Similarly, by the fact that $f'(x_0) \cdot (\beta - x_0) = f'(x_0) \cdot (\beta - x) + f'(x_0) \cdot (x - x_0)$ and $f'(x_0) \cdot (1 - \beta - x_0) = f'(x_0) \cdot (1 - \beta - x) + f'(x_0) \cdot (x - x_0)$, we can get from (72),

$$\begin{aligned}
 &|\mathbb{E}(f(\tilde{X}_n)) - f(\mathbb{E}(\tilde{X}_n))| \\
 &\leq \sup_{x \in (1-\beta, \beta)} |f''(x)| \text{Var}(\tilde{X}_n) + |f'(x_0)| \cdot \mathbb{E}([\tilde{X}_n - \beta]_+) + |f'(x_0)| \cdot \mathbb{E}([1 - \beta - \tilde{X}_n]_+) \quad (73)
 \end{aligned}$$

where $\text{Var}(\cdot)$ represents the variance and $[x]_+ = \max\{x, 0\}$.

In addition, we note that

$$\mathbb{E}([\tilde{X}_n - \beta]_+) \leq \int_{\beta}^1 (x - \beta) \mathbb{P}(\tilde{X}_n \geq x) dx$$

$$\begin{aligned}
 &\stackrel{(a)}{\leq} \int_{\beta}^1 (x - \beta) \exp\left(-\frac{(x - x_0)^2}{2\left(\frac{1}{4n} + \frac{2\log(1.25/\delta)}{(\varepsilon n)^2}\right)}\right) dx \\
 &\leq \int_{\beta}^1 (x - \beta) \exp\left(-\left(2n(x - x_0)^2 \wedge \frac{(x - x_0)^2(\varepsilon n)^2}{4\log(1.25/\delta)}\right)\right) dx \\
 &= \int_{\beta}^1 (x - \beta) \exp\left(-\left(x - x_0\right)^2 \left(2n \wedge \frac{(\varepsilon n)^2}{4\log(1.25/\delta)}\right)\right) dx \\
 &\leq \int_{\beta}^1 (x - \beta) \exp\left(-\left(x - \beta\right)^2 \left(2n \wedge \frac{(\varepsilon n)^2}{4\log(1.25/\delta)}\right)\right) dx \\
 &\stackrel{(b)}{\leq} \int_{\beta}^1 \sqrt{\frac{1}{2n \wedge \frac{(\varepsilon n)^2}{4\log(1.25/\delta)}}} dx \\
 &= \frac{1 - \beta}{\sqrt{2n \wedge \frac{(\varepsilon n)^2}{4\log(1.25/\delta)}}}, \tag{74}
 \end{aligned}$$

where (a) follows from the fact that \tilde{X}_n is subGaussian with variance proxy $\frac{1}{4n} + \frac{2\log(1.25/\delta)}{(\varepsilon n)^2}$, and (b) follows from the fact that $x \exp(-x^2 y) < \sqrt{\frac{1}{y}}$ for any $y > 0$. Similarly, we also can get that

$$\mathbb{E}([1 - \beta - X_n]_+) \leq \frac{1 - \beta}{\sqrt{2n \wedge \frac{(\varepsilon n)^2}{4\log(1.25/\delta)}}}. \tag{75}$$

Finally,

$$\begin{aligned}
 &|\mathbb{E}(f(\tilde{X}_n)) - f(\mathbb{E}(\tilde{X}_n))| \\
 &\stackrel{(a)}{\leq} \sup_{x \in (1-\beta, \beta)} |f''(x)| \text{Var}(\tilde{X}_n) + |f'(x_0)| \cdot \mathbb{E}([\tilde{X}_n - \beta]_+) + |f'(x_0)| \cdot \mathbb{E}([1 - \beta - \tilde{X}_n]_+) \\
 &\stackrel{(b)}{\leq} \frac{1}{(1 - \beta)^4} \cdot \left(\frac{1}{n} + \frac{2\log(1.25/\delta)}{(\varepsilon n)^2}\right) + f'(x_0) \cdot \mathbb{E}([X_n - \beta]_+) + f'(x_0) \cdot \mathbb{E}([1 - \beta - X_n]_+) \\
 &\stackrel{(c)}{\leq} \frac{1}{(1 - \beta)^4} \cdot \left(\frac{1}{n} + \frac{2\log(1.25/\delta)}{(\varepsilon n)^2}\right) + f'(x_0) \cdot \frac{2}{\sqrt{2n \wedge \frac{(\varepsilon n)^2}{4\log(1.25/\delta)}}} \\
 &\stackrel{(d)}{\leq} \frac{1}{(1 - \beta)^4} \cdot \left(\frac{1}{n} + \frac{2\log(1.25/\delta)}{(\varepsilon n)^2}\right) + \frac{1}{(1 - \beta)^2} \cdot \frac{2}{\sqrt{2n \wedge \frac{(\varepsilon n)^2}{4\log(1.25/\delta)}}} \\
 &\leq \frac{3}{(1 - \beta)^4 \sqrt{2n}} + \frac{4\sqrt{\log(1.25/\delta)}}{(1 - \beta)^2(\varepsilon n)} + \frac{2\log(1.25/\delta)}{(1 - \beta)^4(\varepsilon n)^2} \\
 &\leq \frac{3}{(1 - \beta)^4 \sqrt{2n}} + \frac{4}{(1 - \beta)^4} \cdot \left(\frac{\sqrt{\log(1.25/\delta)}}{(\varepsilon n)} \vee \frac{\log(1.25/\delta)}{(\varepsilon n)^2}\right) \\
 &\stackrel{(e)}{\leq} \frac{3}{(1 - \beta)^4 \sqrt{2n}} + \frac{4\sqrt{\log(1.25/\delta)}}{(1 - \beta)^4(\varepsilon n)} \tag{76}
 \end{aligned}$$

where (a) follows from Eqn. (73), (b) follows from the fact that $\text{Var}(\tilde{X}_n) \leq \frac{1}{n} + \frac{2\log(1.25/\delta)}{(\varepsilon n)^2}$ and that $f''(x) = \frac{2x-1}{(x-1)^2 x^2}$, (c) follows from Eqn. (74) and Eqn. (75), (d) follows from the fact that $f'(x) = \frac{1}{x-x^2}$, and (e) follows the fact that $|\mathbb{E}(f(\tilde{X}_n)) - f(\mathbb{E}(\tilde{X}_n))| \leq \frac{4}{(1-\beta)^4}$ and that $\sqrt{x} \geq x$ for any $x \in (0, 1]$ \square

Lemma F.2. Fix any $\varepsilon > 0$ and $\delta > 0$. Let Y_n be a random variable sampled from the binomial distribution with number of trials $\in \mathbb{N}$ and probability of success $p \in [1 - \beta, \beta]$ for $\beta \in (\frac{1}{2}, 1)$. Let $\tilde{X}_n = \frac{Y_n}{n} + \tilde{\xi}_n$, and $\tilde{\xi}_n$ is an independent Gaussian noise with zero mean and variance of $\frac{2\log(1.25/\delta)}{(\varepsilon n)^2}$. Then, $f(\tilde{X}_n)$ is subGaussian with

$$\|f(\tilde{X}_n)\|_{\text{vp}}^2 \leq C \cdot \left(\frac{1}{n} + \frac{8 \log(1.25/\delta)}{(\varepsilon n)^2} \right),$$

where $C \leq (6\sqrt{2e} \cdot (3\sqrt{\log 2} + 1))^2 \cdot \frac{3}{2(1-\beta)^4}$ and $f(\cdot)$ is defined in (54).

Proof. By the definition of X_n , we get that

$$\|\tilde{X}_n - p\|_{\text{vp}}^2 \leq \frac{1}{4n} + \frac{2 \log(1.25/\delta)}{(\varepsilon n)^2}.$$

Then, by Lemma A.5, we get that

$$\|\tilde{X}_n - p\|_{\phi_2}^2 \leq \frac{3}{2n} + \frac{12 \log(1.25/\delta)}{(\varepsilon n)^2}.$$

By the fact that $|f'(x)| \leq \frac{1}{(1-\beta)^2}$, we get that

$$\|f(\tilde{X}_n) - f(p)\|_{\phi_2}^2 \leq \left(\frac{3}{2n} + \frac{12 \log(1.25/\delta)}{(\varepsilon n)^2} \right) \cdot \frac{1}{(1-\beta)^4}. \quad (77)$$

Similarly, from Lemma A.5, we further get that

$$\|f(\tilde{X}_n) - f(p)\|_{\text{vp}}^2 \leq (6\sqrt{2e} \cdot (3\sqrt{\log 2} + 1))^2 \cdot \frac{3}{2(1-\beta)^4} \left(\frac{1}{n} + \frac{8 \log(1.25/\delta)}{(\varepsilon n)^2} \right). \quad (78)$$

Note that by definition, we have $\|f(\tilde{X}_n) - f(p)\|_{\text{vp}}^2 = \|f(\tilde{X}_n)\|_{\text{vp}}^2$, which completes the proof of Lemma E.2. \square

Proof of Theorem 5.3. Fix any consistent instance v and $n > 0$ under DP-RL-LOW. By lemma E.2 and lemma A.3, we get that $\tilde{r}_{k, i_k^*, i}$ is subGaussian with variance proxy to be,

$$\|\tilde{r}_{k, i_k^*, i}\|_{\text{vp}}^2 \leq C \cdot \left(\frac{\gamma_{k, i}}{n} + \frac{8\gamma_{k, i}^{\text{DP}} \log(1.25/\delta)}{(\varepsilon n)^2} \right), \quad (79)$$

where $C \leq (6\sqrt{2e} \cdot (3\sqrt{\log 2} + 1))^2 \cdot \frac{3}{2(1-\beta)^4}$ and $\beta = \frac{\exp(2L)}{1+\exp(2L)}$.

In addition, by Lemma F.1, we get that for any $(k', i', j') \in \mathcal{S} \times \mathcal{A}^2$ with $N_{k', i', j'} > 0$,

$$\begin{aligned} & \left| \mathbb{E}_v^{\text{DP-RL-LOW}} \left[f(\tilde{B}_{k', i', j'}) \right] - f \left(\mathbb{E}_v^{\text{DP-RL-LOW}} \left[\tilde{B}_{k', i', j'} \right] \right) \right| \\ & \leq \frac{3}{(1-\beta)^4 \sqrt{2nN_{k', i', j'}}} + \frac{4\sqrt{\log(1.25/\delta)}}{(1-\beta)^4 (\varepsilon n N_{k', i', j'})}, \end{aligned} \quad (80)$$

which implies for any $(k, i) \in \mathcal{S} \times \mathcal{A}$ with $i \neq i_k^*$,

$$\begin{aligned} & \left| \sum_{k' \in \mathcal{S}, i' \in \mathcal{A}, j' \in \mathcal{A}} w_{k', i', j'}^{(k, i_k^*, i)} \mathbb{E}_v^{\text{DP-RL-LOW}} [f(B_{k', i', j'})] - f(\mathbb{E}_v^{\text{DP-RL-LOW}} [B_{k', i', j'}]) \right| \\ & \leq \sum_{k' \in \mathcal{S}, i' \in \mathcal{A}, j' \in \mathcal{A}: N_{k', i', j'} \neq 0} |w_{k', i', j'}^{(k, i_k^*, i)}| \left(\frac{3}{(1-\beta)^4 \sqrt{2nN_{k', i', j'}}} + \frac{4\sqrt{\log(1.25/\delta)}}{(1-\beta)^4 (\varepsilon n N_{k', i', j'})} \right). \end{aligned} \quad (81)$$

Recall that we denote $r_{k, i^* k, i} = r_{k, i_k^*, i} - r_{k, i}$. Then, from Eqn. (81), we get

$$\begin{aligned} & \left| \mathbb{E}_v^{\text{DP-RL-LOW}} [\tilde{r}_{k, i_k^*, i}] - r_{k, i_k^*, i} \right| \\ & \leq \sum_{k' \in \mathcal{S}, i' \in \mathcal{A}, j' \in \mathcal{A}: N_{k', i', j'} \neq 0} |w_{k', i', j'}^{(k, i, j)}| \left(\frac{3}{(1-\beta)^4 \sqrt{2nN_{k', i', j'}}} + \frac{4\sqrt{\log(1.25/\delta)}}{(1-\beta)^4 (\varepsilon n N_{k', i', j'})} \right) \end{aligned} \quad (82)$$

$$\leq \frac{3\tilde{\gamma}_{k,i}}{(1-\beta)^4\sqrt{n}} + \frac{4\tilde{\gamma}_{k,i}^{DP}\sqrt{\log(1.25/\delta)}}{(1-\beta)^4(\varepsilon n)}, \quad (83)$$

where $\tilde{\gamma}_{k,i}^{DP} := \sum_{k' \in \mathcal{S}, i' \in \mathcal{A}, j' \in \mathcal{A}: N_{k',i',j'} \neq 0} \frac{|w_{k',i',j'}^{(k,i^*,i)}|}{N_{k',i',j'}}$

Note that by the definition of R_n , under DP-RL-LOW we have

$$R_n \leq \sum_{(k,i) \in \mathcal{S} \times \mathcal{A}: i \neq i_k^*} \mathbb{1}_{\{\tilde{r}_{k,i^*,i} < 0\}} \cdot \rho_k \Delta_{k,i} \quad (84)$$

$$\leq \sum_{(k,i) \in \mathcal{S} \times \mathcal{A}: i \neq i_k^*} \left[\left(\mathbb{1}_{\{\tilde{r}_{k,i^*,i} \leq \mathbb{E}_v^{\text{DP-RL-LOW}}[\tilde{r}_{k,i^*,i}] - \Delta_{k,i}/2\}} \vee \mathbb{1}_{\{|\mathbb{E}_v^{\text{DP-RL-LOW}}[\tilde{r}_{k,i^*,i}] - r_{k,i^*,i}|\leq \Delta_{k,i}/2\}} \right) \cdot \rho_k \Delta_{k,i} \right] \quad (85)$$

By Lemma A.4 and (79), we get that for any $(k,i) \in \mathcal{S} \times \mathcal{A}$ with $i \neq i_k^*$,

$$\begin{aligned} & \mathbb{P}_v^{\text{DP-RL-LOW}}(\tilde{r}_{k,i^*,i} \leq \mathbb{E}_v^{\text{DP-RL-LOW}}[\tilde{r}_{k,i^*,i}] - \Delta_{k,i}/2) \\ & \leq \exp\left(-\frac{2\Delta_{k,i}^2}{C \cdot (\gamma_{k,i}/n + 8\gamma_{k,i}^{DP} \log(1.25/\delta)/(\varepsilon n^2))}\right), \end{aligned} \quad (86)$$

and by (83) we get that

$$|\mathbb{E}_v^{\text{DP-RL-LOW}}[\tilde{r}_{k,i^*,i}] - r_{k,i^*,i}| \leq \Delta_{k,i}/2$$

for all $n > \frac{12\tilde{\gamma}_{k,i}^2}{(1-\beta)^8\Delta_{k,i}^2} + \frac{8\tilde{\gamma}_{k,i}^{DP}\sqrt{\log(1.25/\delta)}}{(1-\beta)^4(\varepsilon\Delta_{k,i})}$. That is, for all sufficiently large n , we have

$$\mathbb{E}_v^{\text{RL-LOW}}(R_n) \leq \sum_{k \in \mathcal{S}, i \in \mathcal{A}: i \neq i_k^*} \rho_k \Delta_{k,i} \exp\left(-\frac{2\Delta_{k,i}^2}{C(\gamma_{k,i}/n + 8\gamma_{k,i}^{DP} \log(1.25/\delta)/(\varepsilon n^2))}\right),$$

which further implies that for all sufficiently large n , there exists a global constant C_{DP} , we have

$$\mathbb{E}_v^{\text{RL-LOW}}(R_n) \leq \exp\left(-C_{\text{DP}} \cdot \left(\frac{n}{H(v)} \wedge \left(\frac{n}{H_{\text{DP}}^{(\varepsilon,\delta)}(v)}\right)^2\right)\right).$$

This completes the proof of Theorem 5.3 □