
Preference-based Reinforcement Learning beyond Pairwise Comparisons: Benefits of Multiple Options

Joongkyu Lee

Seoul National University
jklee0717@snu.ac.kr

Seouh-won Yi

Seoul National University
uniquesteou@snu.ac.kr

Min-hwan Oh

Seoul National University
minoh@snu.ac.kr

Abstract

We study online preference-based reinforcement learning (PbRL) with the goal of improving sample efficiency. While a growing body of theoretical work has emerged—motivated by PbRL’s recent empirical success, particularly in aligning large language models (LLMs)—most existing studies focus only on pairwise comparisons. A few recent works [93, 49, 76] have explored using multiple comparisons and ranking feedback, but their performance guarantees fail to improve—and can even deteriorate—as the feedback length increases, despite the richer information available. To address this gap, we adopt the Plackett–Luce (PL) model for ranking feedback over action subsets and propose M-AUPO, an algorithm that selects multiple actions by maximizing the average uncertainty within the offered subset.

We prove that M-AUPO achieves a suboptimality gap of $\tilde{\mathcal{O}}\left(\frac{d}{T}\sqrt{\sum_{t=1}^T \frac{1}{|S_t|}}\right)$, where T is the total number of rounds, d is the feature dimension, and $|S_t|$ is the size of the subset at round t . This result shows that larger subsets directly lead to improved performance and, notably, the bound avoids the exponential dependence on the unknown parameter’s norm, which was a fundamental limitation in most previous works. Moreover, we establish a near-matching lower bound of $\Omega\left(\frac{d}{K\sqrt{T}}\right)$, where K is the maximum subset size. To the best of our knowledge, this is the first theoretical result in PbRL with ranking feedback that explicitly shows improved sample efficiency as a function of the subset size.

1 Introduction

The framework of *Preference-based Reinforcement Learning* (PbRL) [12, 80, 81, 70] was introduced to address the difficulty of designing effective reward functions, which often demands substantial and complex engineering effort [79, 81]. PbRL has been successfully applied in diverse domains, including robot training, stock prediction, recommender systems, and clinical trials [30, 65, 18, 37, 52]. Notably, PbRL also serves as a foundational framework for Reinforcement Learning from Human Feedback (RLHF) when feedback is provided in the form of preferences rather than explicit scalar rewards. This preference-based approach has proven highly effective in aligning Large Language Models (LLMs) with human values and preferences [18, 57, 62].

Given its practical success, the field has also seen significant theoretical advances [16, 47, 70, 93, 86, 91, 90, 83, 72, 51, 13, 64, 22, 19, 49, 74, 71, 76, 85, 14, 38]. However, despite this progress, most existing models remain limited to handling only *pairwise* comparison feedback. A few works [93, 49, 76] explore the more general setting of *multiple* comparisons, offering a strict extension beyond the pairwise case. Zhu et al. [93] study the offline setting, where a dataset of questions (or contexts) along with corresponding ranking feedback over K answers (or actions), labeled by human annotators, is available. Mukherjee et al. [49] investigate the online learning-to-rank problem [61], where a dataset of questions with K candidate answers is provided, but no feedback is initially available.

Table 1: Comparisons of settings and theoretical guarantees in related works on PbRL with ranking feedback. Here, T denotes the number of rounds (or the number of data points in the offline setting), K is the (maximum) size of the offered action set (i.e., *assortment*), and d is the feature dimension. ρ represents the unknown context distribution. Here, $\tilde{\mathcal{O}}$ hides logarithmic factors and polynomial dependencies on B . “Pred. Error” refers to the prediction error.

	Setting	Context	Assortment	Measure	Result
Zhu et al. [93]	Offline	Accessible \mathcal{X}	Given	Suboptimality	$\tilde{\mathcal{O}}\left(e^B K^2 \sqrt{\frac{d}{T}}\right)$
Mukherjee et al. [49]	Online	Accessible \mathcal{X}	Given	Pred. Error	$\tilde{\mathcal{O}}\left(e^B K^3 \frac{d}{\sqrt{T}}\right)$
Thekumparampil et al. [76]	Online	No context	Select K	Pred. Error	$\tilde{\mathcal{O}}\left(e^B K^3 \frac{d}{\sqrt{T}}\right)$
This work (Theorem 1, 2)	Online	Sampled $x \sim \rho$	Select $\leq K$	Suboptimality	$\tilde{\mathcal{O}}\left(\frac{d}{T} \sqrt{\sum_{t=1}^T \frac{1}{ S_t }}\right)$
This work (Theorem 3)	Lower Bound	Sampled $x \sim \rho$	Select $\leq K$	Suboptimality	$\Omega\left(\frac{d}{K\sqrt{T}}\right)$

Thekumparampil et al. [76] consider a context-free setting (i.e., a singleton context), and the goal is to learn the ranking of $N \geq K$ answers based on ranking feedback obtained from subsets of size K . However, all of their theoretical performance guarantees fail to show that using multiple comparisons provides any advantage over the pairwise setting (see Table 1). This is counterintuitive, as ranking feedback is inherently more informative than pairwise feedback. Specifically, since a ranking over K actions provides $\binom{K}{2}$ pairwise comparisons, it should, in principle, enable faster learning and lead to stronger performance guarantees. Thus, the following fundamental question remains open:

Can we design an algorithm that achieves a strictly better theoretical guarantee under multiple-option feedback compared to the pairwise comparisons in the online PbRL setting?

In this paper, we assume that the ranking feedback follows the Plackett-Luce (PL) model [60, 45], where, in each round, the learner receives ranking feedback over a subset of up to K actions (with $K \leq N$) selected from a universe of N actions. This problem setup is closely related to that of Thekumparampil et al. [76]; however, unlike their work, which focuses solely on a context-free setting (or equivalently, a fixed singleton context), we study a more general setting where contexts are diverse and drawn from an unknown distribution.

Under this problem setup, we provide an affirmative answer to the above question by introducing a novel algorithm, *Maximizing Average Uncertainty for Preference Optimization* (M-AUPO), which explicitly exploits the richer information available from ranking feedback under the Plackett-Luce (PL) model. M-AUPO selects action subsets by maximizing *average uncertainty* and achieves a suboptimality gap that strictly improves upon what is attainable with pairwise comparisons. In particular, we show that its suboptimality gap decreases with longer ranking feedback.

Furthermore, our suboptimality gap eliminates the *exponential* dependence on the parameter norm bound, $\mathcal{O}(e^B)$, in the leading term. This improvement stems from *analytically* dividing the total rounds into *warm-up* and *non-warm-up* phases (see the proof sketch in Section 5.1). This represents a significant improvement over most prior works, where performance guarantees depend on $\mathcal{O}(e^B)$ [66, 70, 93, 86, 91, 19, 85, 76, 38]. Very recently, a few works [14, 20] have successfully avoided the $\mathcal{O}(e^B)$ dependency by relying on auxiliary techniques or additional information—such as specialized sampling schemes [14] or prior knowledge of κ [20]—which, however, are often impractical. Moreover, their methods are limited to pairwise comparison settings. In contrast, our approach eliminates the $\mathcal{O}(e^B)$ dependency without using any auxiliary techniques and considers more general ranking feedback beyond pairwise comparisons. Our main contributions are summarized as follows:

- **Improved sample efficiency via larger subsets:** We propose M-AUPO, a novel algorithm for online PbRL (or RLHF) with PL ranking feedback, which achieves a suboptimality gap of $\tilde{\mathcal{O}}\left(\frac{d}{T} \sqrt{\sum_{t=1}^T \frac{1}{|S_t|}}\right)$, where $|S_t|$ is the size of the action subset offered at round t . This result provides the first rigorous theoretical guarantee that larger subsets directly improve sample efficiency. To the best of our knowledge, this is the first theoretical work in PbRL that explicitly demonstrates performance improvements as a function of the subset size $|S_t|$.

- **Improvement on $\mathcal{O}(e^B)$ dependency:** Our result eliminates the exponential dependence on the parameter norm bound, $\mathcal{O}(e^B)$, in the leading term—without relying on any auxiliary techniques—through a *fully* refined analysis. This shows that the $\mathcal{O}(e^B)$ dependence widely observed in PbRL (or RLHF) and dueling bandit analyses is *not inherent* but rather an artifact of loose analysis. Moreover, our key technique for removing the $\mathcal{O}(e^B)$ dependence—by dividing the rounds into warm-up and non-warm-up phases only in the analysis—can be seamlessly incorporated into existing PbRL or dueling bandit analyses, including regret-minimization frameworks without altering the underlying algorithms (see Appendix G.1). To the best of our knowledge, this is the first PbRL work with ranking feedback beyond pairwise comparisons that avoids the $\mathcal{O}(e^B)$ dependence.
- **Efficiency of rank-breaking (RB):** We present both naive PL loss-based and rank-breaking (RB) loss-based learning approaches, and show that the RB formulation achieves superior computational efficiency and empirical performance compared to the PL counterpart.
- **Lower bound:** We establish a near-matching lower bound of $\Omega\left(\frac{d}{K\sqrt{T}}\right)$ under the PL model with ranking feedback, matching our upper bound up to a K factor. This shows that leveraging richer ranking information (larger K) provably improves sample efficiency.
- **Experiment:** We empirically evaluate M-AUPO on both synthetic and real-world datasets, showing its improved performance for larger K and its superiority over existing baselines.

2 Related Works

Fueled by the remarkable success of LLMs [18, 57, 62], the theoretical study of PbRL has rapidly emerged as a central focus within the research community. Early work in this area traces back to the dueling bandits literature [88, 95, 68, 8].

Dueling bandits. The dueling bandit framework, introduced by Yue et al. [88], departs from the classical multi-armed bandit setting by requiring the learner to select two arms and observe only their pairwise preference. For general preferences, a single best arm that is globally dominant may not exist. To address this, various alternative winners have been proposed, including the Condorcet winner [94, 35], Copeland winner [95, 82, 36], Borda winner [31, 25, 28, 69, 84], and von Neumann winner [63, 24, 7], each with its own corresponding performance metric.

To address scalability and contextual information, Saha [66] proposed a structured contextual dueling bandit setting in which preferences are modeled using a Bradley–Terry–Luce (BTL) model [11] based on the unknown intrinsic rewards of each arm. In a similar setting, Bengs et al. [9] studied a contextual linear stochastic transitivity model, and Di et al. [21] proposed a layered algorithm that achieves variance-aware regret bounds. However, most prior dueling bandit works suffer from an $\mathcal{O}(e^B)$ dependence. Only a few recent studies [20, 14] have succeeded in removing this $\mathcal{O}(e^B)$ term, either by introducing additional complex subroutines [14] or by relying on prior knowledge of κ [20].

Preference-based reinforcement learning (PbRL). Building upon this line of work, subsequent research has extended the dueling bandit framework to the RL, considering both online [87, 52, 16, 70, 83] and offline settings [93, 91, 44]. More recently, under the active learning framework—where the full set of contexts \mathcal{X} is accessible—many studies aim to improve sample efficiency by selecting prompts either based on the differences in estimated rewards for their responses [50] or through D-optimal design methods [47, 71, 19, 49, 76, 38]. However, most of these works focus exclusively on pairwise preference feedback and cannot be extended to more general ranking feedback cases. Mukherjee et al. [49] study the online learning-to-rank problem when prompts are given along with K candidate answers, while Thekumpampil et al. [76] investigate learning to rank $N \geq K$ answers from partial rankings over K answers, but under a context-free setting. In this paper, we consider a stochastic contextual setting (more general than Thekumpampil et al. [76]), where contexts are sampled from an unknown but fixed distribution, and aim to minimize the suboptimality gap using ranking feedback of up to length K . For further related work, see Appendix A.

3 Problem Setting and Preliminaries

Notations. Given a set \mathcal{X} , we use $|\mathcal{X}|$ to denote its cardinality. For a positive integer n , we denote $[n] := \{1, 2, \dots, n\}$. For a real-valued matrix A , we let $\|A\|_2 := \sup_{x: \|x\|_2=1} \|Ax\|_2$ which is the

maximum singular value of A . We write $A \succeq A'$ if $A - A'$ is positive semidefinite. For a univariate function f , we denote \dot{f} as its derivative.

We have a set of contexts (or prompts), denoted by \mathcal{X} , and a set of possible actions (or answers), denoted by $\mathcal{A} := \{a_1, \dots, a_N\}$.¹. We consider preference feedback in the form of partial rankings over subsets of \mathcal{A} , and model this feedback using the Plackett-Luce (PL) distribution:

Definition 1 (PL model). *Let $\mathcal{S} := \{S \subseteq \mathcal{A} \mid 2 \leq |S| \leq K\}$ be the collection of all action subsets whose sizes range from 2 to K . For any $S \in \mathcal{S}$, let σ denote the labeler’s ranking feedback—that is, a permutation of the elements in S . We write σ_j for the j -th most preferred action under σ . We model the distribution of such rankings using the Plackett-Luce (PL) model [60, 45], defined as:*

$$\mathbb{P}(\sigma|x, S; \boldsymbol{\theta}^*) = \prod_{j=1}^{|S|} \frac{\exp(r_{\boldsymbol{\theta}^*}(x, \sigma_j))}{\sum_{k=j}^{|S|} \exp(r_{\boldsymbol{\theta}^*}(x, \sigma_k))}, \quad \text{where } (x, S) \in \mathcal{X} \times \mathcal{S}. \quad (1)$$

Here, $r_{\boldsymbol{\theta}^*}$ represents a reward model parameterized by the unknown parameter $\boldsymbol{\theta}^*$.

When $K=2$, this reduces to the pairwise comparison framework considered in the Bradley-Terry-Luce (BTL) model [11]. The probability that a is preferred to a' given x can be expressed as:

$$\mathbb{P}(a > a'|x; \boldsymbol{\theta}^*) = \frac{\exp(r_{\boldsymbol{\theta}^*}(x, a))}{\exp(r_{\boldsymbol{\theta}^*}(x, a)) + \exp(r_{\boldsymbol{\theta}^*}(x, a'))} = \mu(r_{\boldsymbol{\theta}^*}(x, a) - r_{\boldsymbol{\theta}^*}(x, a')), \quad (2)$$

where $\mu(w) = \frac{1}{1+e^{-w}}$ is the sigmoid function. In this work, we assume a linear reward model:

Assumption 1. *Let $\phi : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}^d$ be a known feature map satisfying $\max_{x, a} \|\phi(x, a)\|_2 \leq 1$, and let $\boldsymbol{\theta}^* \in \mathbb{R}^d$ denote the true but unknown parameter. The reward is assumed to follow a linear structure given by $r_{\boldsymbol{\theta}^*}(x, a) = \phi(x, a)^\top \boldsymbol{\theta}^*$. We further assume realizability, i.e., $\boldsymbol{\theta}^* \in \Theta := \{\boldsymbol{\theta} \in \mathbb{R}^d \mid \|\boldsymbol{\theta}\|_2 \leq B\}$. Without loss of generality, we assume $B \geq 1$.*

At each round $t \in [T]$, a context $x_t \in \mathcal{X}$ is drawn from a fixed but unknown distribution ρ . Given the context x_t , the learning agent selects a subset of actions $S_t \in \mathcal{S}$ —referred to as an *assortment* throughout the paper—and receives a ranking over S_t as feedback, generated according to the PL model. Let $\pi^*(x) = \operatorname{argmax}_a r_{\boldsymbol{\theta}^*}(x, a)$ be the optimal policy under the true reward $r_{\boldsymbol{\theta}^*}$. After T rounds of interaction with the labeler, the goal is to output a policy $\hat{\pi}_T : \mathcal{X} \rightarrow \mathcal{A}$ that minimizes the *suboptimality gap*, defined as:

$$\mathbf{SubOpt}(T) := \mathbb{E}_{x \sim \rho} [r_{\boldsymbol{\theta}^*}(x, \pi^*(x)) - r_{\boldsymbol{\theta}^*}(x, \hat{\pi}_T(x))].$$

3.1 Loss Functions and Rank-Breaking

In this paper, we consider two different losses for estimating the parameter: one directly induced by the PL model, and the other obtained by splitting the ranking feedback into pairwise comparisons.

Plackett-Luce (PL) loss. The PL loss function for round t is defined as follows:

$$\ell_t(\boldsymbol{\theta}) := \sum_{j=1}^{|S_t|} \ell_t^{(j)}(\boldsymbol{\theta}), \quad \text{where } \ell_t^{(j)}(\boldsymbol{\theta}) := -\log \left(\frac{\exp(\phi(x_t, \sigma_{tj})^\top \boldsymbol{\theta})}{\sum_{k=j}^{|S_t|} \exp(\phi(x_t, \sigma_{tk})^\top \boldsymbol{\theta})} \right). \quad (3)$$

Here, $\ell_t^{(j)}(\boldsymbol{\theta})$ denotes the negative log-likelihood loss under the Multinomial Logit (MNL) model [46], conditioned on the assortment being the remaining actions in S_t after removing the previously selected actions $\sigma_{t1}, \dots, \sigma_{t(j-1)}$ —that is, over the set $S_t \setminus \{\sigma_{t1}, \dots, \sigma_{t(j-1)}\}$.

Rank-Breaking (RB) loss. In addition to this standard approach, one can replace the full $|S_t|$ -action ranking with its $\binom{|S_t|}{2}$ pairwise comparisons. This technique, referred to as *rank-breaking* (RB), decomposes (partial) ranking data into individual pairwise comparisons, treating each comparison as independent [6, 34, 32, 67]. Thus, the RB loss is defined as:

$$\ell_t(\boldsymbol{\theta}) := \sum_{j=1}^{|S_t|-1} \sum_{k=j+1}^{|S_t|} \ell_t^{(j,k)}(\boldsymbol{\theta}), \quad \text{where } \ell_t^{(j,k)}(\boldsymbol{\theta}) := -\log \left(\frac{\exp(\phi(x_t, \sigma_{tj})^\top \boldsymbol{\theta})}{\sum_{m \in \{j, k\}} \exp(\phi(x_t, \sigma_{tm})^\top \boldsymbol{\theta})} \right). \quad (4)$$

This approach is applied in the current RLHF for LLM (e.g., Ouyang et al. [57]) and is also studied in the theoretical RLHF paper [93] under the offline setting.

¹For simplicity, we assume a stationary action space \mathcal{A} , though it may depend on the context $x \in \mathcal{X}$.

Procedure 1 OMD-PL, OMD for PL Loss

Input: $\hat{\theta}_t^{(1)}, S_t, H_t$
for $j = 1$ to $|S_t|$ **do**
 Update $\tilde{H}_t^{(j)}, \hat{\theta}_t^{(j+1)}$ via (5)
end for
return $\hat{\theta}_t^{(|S_t|+1)}$

Procedure 2 OMD-RB, OMD for RB Loss

Input: $\hat{\theta}_t^{(1,2)}, S_t, H_t$
for each (j, k) such that $j < k \leq |S_t|$ **do**
 Update $\tilde{H}_t^{(j,k)}, \hat{\theta}_t^{(j,k+1)}$ via (7)
end for
return $\hat{\theta}_t^{(|S_t|-1, |S_t|+1)}$

3.2 Online Parameter Estimation

Motivated by recent advances in Multinomial Logit (MNL) bandits [92, 39, 41], we adopt an online mirror descent (OMD) algorithm to estimate the underlying parameter θ^* , instead of relying on maximum likelihood estimation (MLE). This enables a constant per-round computational cost, in contrast to the MLE-based approach, whose cost grows linearly with the number of rounds t .

OMD update for PL loss. For the the PL loss (3), we estimate the true parameter θ^* as follows:

$$\hat{\theta}_t^{(j+1)} = \underset{\theta \in \Theta}{\operatorname{argmin}} \langle \nabla \ell_t^{(j)}(\hat{\theta}_t^{(j)}), \theta \rangle + \frac{1}{2\eta} \|\theta - \hat{\theta}_t^{(j)}\|_{\tilde{H}_t^{(j)}}^2, \quad j = 1, \dots, |S_t|, \quad (5)$$

where we write $\hat{\theta}_t^{(|S_t|+1)} = \hat{\theta}_{t+1}^{(1)}$, and η is the step-size parameter to be specified later. The matrix $\tilde{H}_t^{(j)}$ is given by $\tilde{H}_t^{(j)} := H_t + \eta \sum_{j'=1}^j \nabla^2 \ell_t^{(j')}(\hat{\theta}_t^{(j')})$, where

$$H_t := \sum_{s=1}^{t-1} \sum_{j=1}^{|S_s|} \nabla^2 \ell_s^{(j)}(\hat{\theta}_s^{(j+1)}) + \lambda \mathbf{I}_d, \quad \lambda > 0. \quad (6)$$

The optimization problem (5) can be solved using a single projected gradient step [55], which enjoys a computational cost of only $\mathcal{O}(Kd^3)$ —independent of t [48], unlike MLE—and requires only $\mathcal{O}(d^2)$ storage, thanks to the incremental updates of $\tilde{H}_t^{(j)}$ and H_t .

OMD update for RB loss. Similarly, for the RB loss (4), we estimate the underlying parameter as:

$$\hat{\theta}_t^{(j,k+1)} = \underset{\theta \in \Theta}{\operatorname{argmin}} \langle \nabla \ell_t^{(j,k)}(\hat{\theta}_t^{(j,k)}), \theta \rangle + \frac{1}{2\eta} \|\theta - \hat{\theta}_t^{(j,k)}\|_{\tilde{H}_t^{(j,k)}}^2, \quad 1 \leq j < k \leq |S_t|, \quad (7)$$

where we set $\hat{\theta}_t^{(j,|S_t|+1)} = \hat{\theta}_t^{(j+1,j+2)}$ for all $j < |S_t| - 1$ and for the final pair, let $\hat{\theta}_t^{(|S_t|-1, |S_t|+1)} = \hat{\theta}_{t+1}^{(1,2)}$. The matrix $\tilde{H}_t^{(j,k)}$ is defined as $\tilde{H}_t^{(j,k)} := H_t + \eta \sum_{(j',k') \leq (j,k)} \nabla^2 \ell_t^{(j',k')}(\hat{\theta}_t^{(j',k')})^2$, where

$$H_t := \sum_{s=1}^{t-1} \sum_{j=1}^{|S_s|-1} \sum_{k=j+1}^{|S_s|} \nabla^2 \ell_s^{(j,k)}(\hat{\theta}_s^{(j,k+1)}) + \lambda \mathbf{I}_d, \quad \lambda > 0. \quad (8)$$

Remark 1 (Computational cost of OMD). *The per-round computational cost of the PL parameter update is $\mathcal{O}(K^2 d^3)$, since the parameter is updated $|S_t| \leq K$ times per round. Similarly, the cost for the RB parameter update is $\mathcal{O}(K^3 d^3)$, as the parameter is updated $\binom{|S_t|}{2}$ times per round.*

4 M-AUPO: Maximizing Average Uncertainty

In this section, we propose a new algorithm, M-AUPO, which selects an assortment that maximizes the *average uncertainty* of S_t , thereby exploiting the potential benefits of a larger K . For clarity of presentation, we first define the MNL probability [46] for a given assortment S at round t as follows:

$$P_t(a|S; \theta) := \frac{\exp(\phi(x_t, a)^\top \theta)}{\sum_{a' \in S} \exp(\phi(x_t, a')^\top \theta)}, \quad \forall a \in S. \quad (9)$$

²We write $(j', k') \leq (j, k)$ to indicate lexicographic order, i.e., $j' < j$ or $j' = j$ and $k' \leq k$.

Algorithm 3 M-AUPO: Maximizing Average Uncertainty for Preference Optimization

- 1: **Inputs:** maximum assortment size K , regularization parameter λ , step size η
- 2: **Initialize:** $H_1 = \lambda \mathbf{I}_d$, $\hat{\boldsymbol{\theta}}_1 \in \Theta$
- 3: **for** round $t = 1$ to T **do**
- 4: Observe $x_t \sim \rho$ and select S_t via (11)
- 5: Observe ranking feedback σ_t for S_t
- 6: $\hat{\boldsymbol{\theta}}_{t+1} \leftarrow \text{OMD-PL}(\hat{\boldsymbol{\theta}}_t, S_t, H_t)$ (Proc. 1) ▷ or $\text{OMD-RB}(\hat{\boldsymbol{\theta}}_t, S_t, H_t)$ (Proc. 2) if RB loss
- 7: Update $H_{t+1} \leftarrow H_t + \sum_{j=1}^{|S_t|} \nabla^2 \ell_t^{(j)}(\hat{\boldsymbol{\theta}}_t^{(j+1)})$ via (6) ▷ or via (8) if RB loss
- 8: **end for**
- 9: **Return:** $\hat{\pi}_T(x) \leftarrow \text{argmax}_{a \in \mathcal{A}} \phi(x, a)^\top \hat{\boldsymbol{\theta}}_{T+1}$

Given an assortment S and ranking feedback σ , let $S_\sigma^{(j)} := \{\sigma_j, \dots, \sigma_{|S|}\}$ denote the remaining actions in S after removing the first $j - 1$ actions. Let $\mathbb{P}_t(\sigma|S; \boldsymbol{\theta}) = \mathbb{P}(x_t, \sigma|S; \boldsymbol{\theta})$, for simplicity. Then, the PL model in Equation (1) can be expressed as follows:

$$\mathbb{P}_t(\sigma|S; \boldsymbol{\theta}) = P_t(\sigma_1|S; \boldsymbol{\theta}) \cdot P_t(\sigma_2|S \setminus \{\sigma_1\}; \boldsymbol{\theta}) \cdot \dots \cdot P_t(\sigma_{|S|}|S_\sigma^{(j)}; \boldsymbol{\theta}) = \prod_{j=1}^{|S|} P_t(\sigma_j|S_\sigma^{(j)}; \boldsymbol{\theta}),$$

Greedy assortment selection for PL loss. Given H_t and $\hat{\boldsymbol{\theta}}_t$, we define the function $f_t(S)$ as follows:

$$f_t(S) := \frac{1}{|S|} \sum_{j=1}^{|S|} \mathbb{E}_{\substack{\sigma \sim \mathbb{P}_t(\cdot|S; \hat{\boldsymbol{\theta}}_t) \\ a \sim P_t(\cdot|S_\sigma^{(j)}; \hat{\boldsymbol{\theta}}_t)}} \left[\left\| \phi(x_t, a) - \mathbb{E}_{a' \sim P_t(\cdot|S_\sigma^{(j)}; \hat{\boldsymbol{\theta}}_t)} [\phi(x_t, a')] \right\|_{H_t^{-1}}^2 \right]. \quad (10)$$

Intuitively, finding S that exactly maximizes $f_t(S)$ amounts to maximizing the *average (mean-centered) uncertainty*—that is, the Mahalanobis dispersion of the feature vectors evaluated under H_t^{-1} . However, computing the exact maximizer of Equation (10) is generally NP-hard. In our setting, it is sufficient to add actions sequentially in a *greedy manner*. This is because our analysis centers on the suboptimality gap between two actions—one optimal and one chosen by the policy. As a result, the suboptimality gap is (approximately) upper bounded by a term involving $f_t(S)$ for pairwise action sets (i.e., $|S| = 2$). Therefore, greedily adding actions that increase $f_t(S)$ is sufficient.

We initialize S with a pair of actions ($|S| = 2$) that maximizes the average information gain, as defined in Equation (10). Then, iteratively add one action at a time by

$$a^* \in \text{argmax}_{a \in \mathcal{A} \setminus S} \Delta_t(a | S), \quad \text{where } \Delta_t(a | S) := f_t(S \cup \{a\}) - f_t(S), \quad (11)$$

and accept a^* if $\Delta_t(a | S) \geq 0$, until $|S| = K$ or no non-negative gain remains. This selection rule is central to our algorithm, as it facilitates a rapid decrease in the reward estimation error by favoring assortments that provide more informative feedback, especially when the assortment size $|S_t|$ is large. Notably, the greedy variant of the selection rule can be implemented efficiently with a computational cost of $\mathcal{O}(N^2 d^3 + L N K^2 d^2)$ ³, where L denotes the (approximate) expectation cost over the ranking σ . This implementation avoids enumerating all $\binom{N}{K}$ possible subsets.

Greedy assortment selection for RB loss. Analogous to Equation (10), we define the following function to select an assortment that maximizes the average uncertainty:

$$\begin{aligned} f_t(S) &:= \frac{1}{|S|} \sum_{a, a' \in S} \mathbb{E}_{\substack{a \sim P_t(\cdot|S; \hat{\boldsymbol{\theta}}_t) \\ a' \sim P_t(\cdot|S; \hat{\boldsymbol{\theta}}_t)}} \left[\left\| \phi(x_t, a) - \mathbb{E}_{\tilde{a} \sim P_t(\cdot|S; \hat{\boldsymbol{\theta}}_t)} [\phi(x_t, \tilde{a})] \right\|_{H_t^{-1}}^2 \right] \\ &= \frac{1}{2|S|} \sum_{a, a' \in S} \dot{\mu} \left((\phi(x_t, a) - \phi(x_t, a'))^\top \hat{\boldsymbol{\theta}}_t \right) \left\| \phi(x_t, a) - \phi(x_t, a') \right\|_{H_t^{-1}}^2. \end{aligned} \quad (12)$$

We then perform greedy assortment selection according to Equation (11). Note that, unlike in the PL loss case, assortment selection under the RB loss does not require taking expectations over rankings, as all pairs within S are directly compared. Therefore, the RB-based assortment selection is both *exact* and *computationally efficient*, without any additional expectation-approximation cost L , and incurs a total computational cost of $\mathcal{O}(N^2 d^3 + N K^3 d^2)$ ⁴.

³ $\mathcal{O}(N^2 d^3)$ to choose the first pair, followed by at most $K - 2$ additions, each of which requires $\mathcal{O}(L N K d^2)$.

⁴ $\mathcal{O}(N^2 d^3)$ to choose the first pair, followed by up to $K - 2$ additional steps, each costing $\mathcal{O}(N K^2 d^3)$.

Once the assortment S_t is selected, the algorithm receives the ranking feedback σ_t from the labeler (Line 5) and updates the parameter using Procedure 1 for the PL loss or Procedure 2 for the RB loss (Line 6). After T rounds, the algorithm returns the final policy $\hat{\pi}_T$, which selects actions by maximizing the estimated reward under the final parameter estimate $\hat{\theta}_{T+1}$ (Line 7).

5 Main Results

5.1 Suboptimality Gap of M-AUPO

We begin by presenting the online confidence bound for the PL loss, derived by extending the results of Lee and Oh [41], who analyzed the MNL model [46]. Since the PL model constructs ranking probabilities as a product of MNL probabilities, their confidence bound can be directly applied to our setting by replacing the round t with the cumulative number of updates $\sum_{s=1}^t |S_s|$.

Corollary 1 (Online confidence bound for PL loss). *Let $\delta \in (0, 1]$. We set $\eta = (1 + 3\sqrt{2}B)/2$ and $\lambda = \max\{12\sqrt{2}B\eta, 144\eta d, 2\}$. Then, under Assumption 1, with probability at least $1 - \delta$, we have*

$$\|\hat{\theta}_t^{(j)} - \theta^*\|_{H_t^{(j)}} \leq \beta_t(\delta) = \mathcal{O}\left(B\sqrt{d \log(tK/\delta)} + B\sqrt{\lambda}\right), \quad \forall t \geq 1, j \leq |S_t|,$$

$$\text{where } H_t^{(j)} := H_t + \sum_{j'=1}^{j-1} \nabla^2 \ell_s^{(j')}(\hat{\theta}_s^{(j'+1)}) + \lambda \mathbf{I}_d.$$

This confidence bound is free of any polynomial dependency on K , which is primarily made possible by the improved self-concordant-like properties proposed by Lee and Oh [41]. Moreover, for the RB loss, we can derive a confidence bound of the same order (see Proposition D.1). Based on this confidence bound, we derive the suboptimality gap for M-AUPO, with the proof deferred to Appendix C.

Theorem 1 (Suboptimality gap for PL loss). *Let $\delta \in (0, 1]$. Set $\lambda = \Omega(d \log(KT/\delta) + \eta(B + d))$ and $\eta = \frac{1}{2}(1 + 3\sqrt{2}B)$. Define $\kappa := e^{-6B}$. If Assumption 1 holds, then for any $T \geq e^{K/d}$, with probability at least $1 - \delta$, M-AUPO (Algorithm 3) achieves the following suboptimality gap:*

$$\mathbf{SubOpt}(T) = \tilde{\mathcal{O}}\left(\frac{d}{T} \sqrt{\sum_{t=1}^T \frac{1}{|S_t|}} + \frac{d^2 K^4}{\kappa T}\right).$$

Discussion of Theorem 1. For sufficiently large T , the second (non-leading) term becomes negligible, and Theorem 1 shows that the suboptimality gap of M-AUPO decreases as the assortment size $|S_t|$ increases. This establishes a strict advantage of receiving ranking feedback over larger assortments. Moreover, our result does not involve any $\mathcal{O}(e^B)$ dependency in the leading term, a harmful dependency that commonly appears in prior works [66, 70, 93, 86, 91, 19, 76, 38]. Although very recent studies [20, 14] also achieve $\mathcal{O}(e^B)$ -free performance in the leading term, they rely on auxiliary techniques and are restricted to pairwise preference feedback. To the best of our knowledge, this is the first theoretical study that simultaneously establishes (i) the performance benefits of utilizing richer ranking feedback over larger assortments, and (ii) the elimination of the $\mathcal{O}(e^B)$ dependence in the leading term of the PbRL framework when accommodating multiple (i.e., more than two) options.

Proof sketch of Theorem 1. We provide a proof sketch for Theorem 1. For simplicity, we define $\psi_{t,a,a'} := \phi(x_t, a) - \phi(x_t, a')$. We begin by defining the set of *warm-up rounds*, denoted by \mathcal{T}^w , which consists of rounds with large (non-centered) uncertainty:

$$\mathcal{T}^w := \left\{ t \in [T] : \max_{a,a' \in \mathcal{A}} \|\psi_{t,a,a'}\|_{H_t^{-1}} \geq 1/(3\sqrt{2}K\beta_{T+1}(\delta)) \right\}.$$

1) Regret decomposition and assortment selection. The proof begins by decomposing the suboptimality gap into two components: the *realized regrets* and a *martingale difference sequence (MDS)*. Since the MDS term can be readily bounded using the Azuma–Hoeffding inequality, the analysis focuses on bounding the realized regrets.

$$\begin{aligned} \mathbf{SubOpt}(T) &= \frac{1}{T} \sum_{t=1}^T \underbrace{\left(\psi_{t, \pi^*(x_t), \hat{\pi}_T(x_t)} \right)^\top \theta^*}_{\text{realized regret of } \hat{\pi}_T \text{ at round } t} + \underbrace{\frac{1}{T} \sum_{t=1}^T \text{MDS}_t}_{=\tilde{\mathcal{O}}(1/\sqrt{T})} \\ &\lesssim \frac{1}{T} \sum_{t \notin \mathcal{T}^w} \|\psi_{t, \pi^*(x_t), \hat{\pi}_T(x_t)}\|_{H_t^{-1}} \underbrace{\|\theta^* - \hat{\theta}_{T+1}\|_{H_T}}_{\leq \beta_t(\delta) = \tilde{\mathcal{O}}(\sqrt{d})} + \tilde{\mathcal{O}} \left(\frac{1}{\sqrt{T}} + \frac{d^2 K^4}{\kappa T} \right). \end{aligned}$$

In the inequality, we first use the fact that $-(\psi_{t,\pi^\star(x_t),\hat{\pi}_T(x_t)})^\top \theta_{T+1} \geq 0$, which follows from definition of $\hat{\pi}_T$. We then apply Hölder's inequality together with the inequality $H_{T+1} \geq H_t$. Moreover, we invoke Lemma C.4, which states that the regret incurred during the warm-up rounds \mathcal{T}^w is a non-leading term. Finally, by applying Corollary 1, we bound $\|\theta^\star - \hat{\theta}_{T+1}\|_{H_T}$ by $\tilde{\mathcal{O}}(\sqrt{d})$.

2) **Avoiding $\mathcal{O}(e^B)$ for non-warm-up rounds.** For $t \notin \mathcal{T}^w$, we have

$$(\psi_{t, \pi^\star(x_t), \hat{\pi}_T(x_t)})^\top \boldsymbol{\theta}^\star \leq \|\psi_{t, \pi^\star(x_t), \hat{\pi}_T(x_t)}\|_{H_t^{-1}} \|\boldsymbol{\theta}^\star - \hat{\boldsymbol{\theta}}_{T+1}\|_{H_T} \leq \frac{\beta_t(\delta)}{3\sqrt{2}K\beta_{T+1}(\delta)} \leq 1,$$

which implies that $1/\dot{\mu}(\psi_{t,\pi^*(x_t),\hat{\pi}_T(x_t)})^\top \theta^* \leq (1+e)^2$. Thus, we obtain

$$\begin{aligned} \frac{1}{T} \sum_{t \notin \mathcal{T}^w} \|\psi_{t, \pi^*(x_t), \hat{\pi}_T(x_t)}\|_{H_t^{-1}} &\leq \frac{(1+e)^2}{T} \sum_{t \notin \mathcal{T}^w} \dot{\mu}\left(\left(\psi_{t, \pi^*(x_t), \hat{\pi}_T(x_t)}\right)^\top \boldsymbol{\theta}^*\right) \|\psi_{t, \pi^*(x_t), \hat{\pi}_T(x_t)}\|_{H_t^{-1}} \\ &\leq \frac{e(1+e)^2}{T} \sum_{t \notin \mathcal{T}^w} \dot{\mu}\left(\left(\psi_{t, \pi^*(x_t), \hat{\pi}_T(x_t)}\right)^\top \hat{\boldsymbol{\theta}}_{t+1}\right) \|\psi_{t, \pi^*(x_t), \hat{\pi}_T(x_t)}\|_{H_t^{-1}}, \end{aligned}$$

where the last inequality follows from Lemma E.2, together with the fact that for any $a, a' \in \mathcal{A}$ and $t \notin \mathcal{T}^w$, $|\psi_{t,a,a'}^\top(\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{t+1})| \leq \|\psi_{t,a,a'}\|_{H_t^{-1}} \|\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{t+1}\|_{H_t} \leq \beta_t(\delta)/\beta_{T+1}(\delta) = 1$. Then, by the Cauchy–Schwarz inequality, and ignoring constants, we have

$$\begin{aligned} & \frac{1}{T} \sqrt{\sum_{t \notin \mathcal{T}^w} \frac{1}{|S_t|}} \sqrt{\sum_{t \notin \mathcal{T}^w} |S_t| \cdot \dot{\mu} \left(\left(\psi_{t, \pi^*}(x_t), \hat{\pi}_T(x_t) \right)^\top \hat{\theta}_{t+1} \right) \| \psi_{t, \pi^*}(x_t), \hat{\pi}_T(x_t) \|_{H_t^{-1}}^2} \\ & \lesssim \frac{1}{T} \sqrt{\sum_{t=1}^T \frac{1}{|S_t|}} \sqrt{\sum_{t \notin \mathcal{T}^w} |S_t| \cdot f_t(\{\pi^*(x_t), \hat{\pi}_T(x_t)\})} \leq \frac{1}{T} \sqrt{\sum_{t=1}^T \frac{1}{|S_t|}} \sqrt{\sum_{t \notin \mathcal{T}^w} |S_t| \cdot f_t(S_t)} \end{aligned} \quad (S_t \text{ selection})$$

where $M_t = \sum_{s \in [t-1] \setminus \mathcal{T}^w} \sum_{j=1}^{|S_s|} \mathbb{E}_{\sigma \sim \mathbb{P}_s(\cdot | S_s, \boldsymbol{\theta}^*)} \left[\nabla^2 \ell_{s, \sigma}^{(j)}(\hat{\boldsymbol{\theta}}_s) \right] + \lambda \mathbf{I}_d$ is the expected version of H_t . Finally, applying the elliptical potential lemma (Lemma C.3), we conclude the proof.

Remark 2 (Generality of our technique for avoiding $\mathcal{O}(e^B)$). *Our approach of dividing the total rounds into warm-up and non-warm-up phases to improve the $\mathcal{O}(e^B)$ dependency is broadly applicable. In particular, this technique can be readily incorporated into most existing PbRL and dueling bandit algorithms without modifying their original algorithms. See Appendix G.1 for more details.*

Furthermore, under the rank-breaking (RB) parameter update (7) and assortment selection rule (12), we establish a comparable suboptimality gap for the RB loss. The proof is provided in Appendix D.

Theorem 2 (Online confidence bound for RB loss). *Under the same setting as Theorem 1, if the parameter is updated and the assortment is selected according to the RB-loss setting and $K^2 \leq (2\sqrt{6} + 12\sqrt{3}B)d$, then for any $T \geq e^{K^2/d}$, with probability at least $1 - \delta$, M-AUPO (Algorithm 3) satisfies*

$$\mathbf{SubOpt}(T) = \tilde{\mathcal{O}}\left(\frac{d}{T}\sqrt{\sum_{t=1}^T \frac{1}{|S_t|}} + \frac{d^2}{\kappa T}\right).$$

Discussion of Theorem 2. For sufficiently large T , the suboptimality gap in Theorem 2 matches the leading-order term of Theorem 1, while its second (non-leading) term is tighter by a factor of $\mathcal{O}(K^4)$. In addition, the assortment selection cost for RB (12) is lower, since it avoids the expensive expectation over rankings required in the PL-loss case (Equation (10)). This yields an important insight: under ranking feedback, *breaking the ranking* and updating pairs independently can be more computationally efficient—and even more stable—than learning directly from the PL loss. Although the bound is stated for the small- K regime, the assumption appears mild in practice, since in LLM applications the dimension d is typically much larger than K . Moreover, this result provides a theoretical explanation for the empirical success of RLHF in LLMs (e.g., Ouyang et al. [57]), where ranking feedback is commonly decomposed into pairwise comparisons for parameter estimation.

5.2 Lower Bound

Theorem 3 (Lower bound). *Suppose $T \geq d^2/(8K^2)$. Define the feature space as $\Phi := \mathcal{S}^{d-1}$, the unit sphere in \mathbb{R}^d , and let the parameter space be $\Theta = \{-\mu, \mu\}^d$, where $\mu = \sqrt{d/(8K^2T)}$. Then, for any final policy $\hat{\pi}_T \in \Delta_\Phi$ returned after collecting T samples (using any sampling policy), the expected suboptimality gap is lower bounded as:*

$$\mathbf{SubOpt}(T) = \Omega\left(\frac{d}{K\sqrt{T}}\right).$$

Discussion of Theorem 3. The proof is deferred to Appendix F. Theorem 3 provides theoretical support for our upper bounds, especially in their dependence on K ; in particular, it matches the upper bound in Theorem 2 when $|S_t| = K$. To the best of our knowledge, this is the first lower bound on the suboptimality gap that incorporates PL ranking feedback in PbRL and formally shows that the suboptimality gap can diminish as K grows, highlighting the advantage of utilizing ranking feedback over simple pairwise comparisons.

6 Numerical Experiments

We conduct two sets of experiments to empirically validate our theoretical findings: (i) one using synthetic data (Subsection 6.1), and (ii) another using two real-world datasets (Subsection 6.2). We compare our proposed algorithm, M-AUPO, against three baselines: (i) `DopeWolfe` [76], which selects K actions in a non-contextual setting; (ii) `Uniform`, which uniformly samples assortments of size K at random; and (iii) `Best&Ref` constructs an action pair ($|S_t| = 2$) by combining the action that maximizes the current reward estimate with another sampled from a reference policy (e.g., uniform random or SFT), following the setup in Online GSHF [86] and XPO [85]. In our experiments, the reference policy for `Best&Ref` is set to the uniform random policy.

6.1 Synthetic Data

In the synthetic data experiment, for each instance, we sample the underlying parameter $\theta^* \sim \mathcal{N}(0, I_d)$ and normalize it to ensure that $\|\theta^*\|_2 \leq 1$. At every round t , a context $x \in \mathcal{X}$ is drawn uniformly at random, and its feature vector $\phi(x, \cdot)$ lies within the unit ball. We set $d = 5$, $|\mathcal{A}| = N = 100$, and $|\mathcal{X}| = 100$. We measure the suboptimality gap every 25 rounds and report the mean over 20 independent runs, together with one standard error.

The first two plots in Figure 1 show the suboptimality gap of M-AUPO under both the PL loss (3) and RB loss (4) as the maximum assortment size K varies. The results clearly show that performance improves as K increases, supporting our theoretical findings. In the third plot of Figure 1, we compare the performance of M-AUPO with three baseline methods at the final round for $K = 5$, showing that our algorithm significantly outperforms all baselines. While `DopeWolfe` also considers the selection of K actions from N actions, it treats each context x independently and is specifically designed for the context-free setting (i.e., a singleton context). As a result, `DopeWolfe` cannot leverage information sharing across varying contexts and performs poorly in our setting. Furthermore, M-AUPO outperforms naive assortment selection strategies such as `Uniform` and `Best&Ref`, as it explicitly chooses assortments that maximize the average uncertainty, thereby achieving more efficient exploration. Finally, the results also show that the RB version of M-AUPO empirically outperforms the PL-based approach, since RB does not incur approximation errors when computing expectations. See Appendix H.1 for additional experimental results.

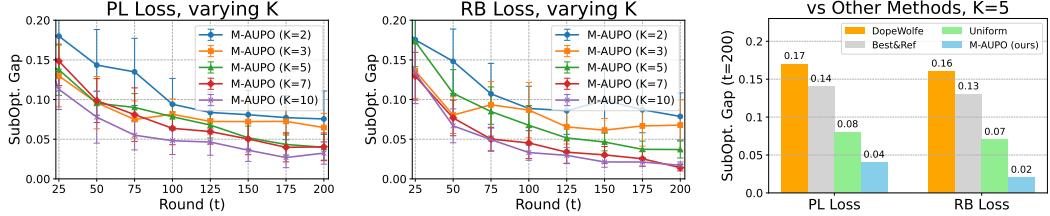


Figure 1: Synthetic data experiment: suboptimality gap of M-AUPO under varying K , evaluated with PL loss (left) and RB loss (middle), along with comparison against other baselines (right).

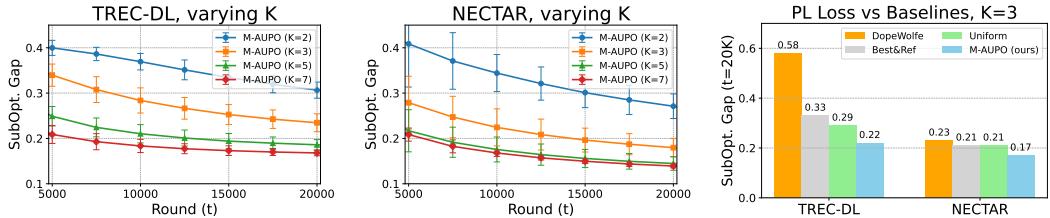


Figure 2: Real-world dataset experiment: suboptimality gap of M-AUPO under varying K on the TREC-DL dataset (left) and the NECTAR dataset (middle), along with comparison against other baselines (right). The results are rescaled to align the performances between the two datasets.

6.2 Real-World Dataset

We also conduct experiments using real-world datasets from TREC Deep Learning (TREC-DL)⁵ and NECTAR⁶. The TREC-DL dataset provides 100 candidate answers for each question, while the NECTAR dataset offers 7 candidate answers per question. We sample $|\mathcal{X}| = 5000$ prompts from each dataset, with the corresponding set of actions (100 or 7 actions, respectively).

We use the gemma-2b⁷ [75] LLM to construct the feature $\phi(x, a)$. Specifically, $\phi(x, a)$ is obtained by extracting the embedding of the concatenated prompt and response from the last hidden layer of the LLM, with size $d = 2048$. Additionally, we use the Mistral-7B [33] reward model⁸ as the true reward model r_{θ^*} to generate ranking feedback and compute the suboptimality gap accordingly. We measure the suboptimality gap every 2,500 rounds and report the average over 10 independent runs, along with the standard error. In these experiments, we report results only for the RB version of M-AUPO, as it is computationally more efficient and empirically performs better, as shown in Figure 1.

The first two plots in Figure 2 show that the performance improves as K increases on two real-world datasets. In the third plot of Figure 2, we compare the performance of M-AUPO with other baselines with $K = 3$ at the final round, showing that M-AUPO outperforms baselines by a large margin. See Appendix H.2 for additional experimental details and results.

7 Conclusion

To the best of our knowledge, this work provides the first theoretical result in online PbRL showing that the suboptimality gap decreases as more options are revealed to the labeler for ranking feedback. Our analysis also removes the $\mathcal{O}(e^B)$ dependency in the leading term without modifying the algorithm, implying that existing PbRL and dueling bandit methods can similarly avoid this dependence through our refined analysis. Together, these results advance the theoretical understanding of PbRL, revealing both the value of richer feedback and the opportunity for more refined and efficient analysis.

⁵<https://microsoft.github.io/msmarco/TREC-Deep-Learning>

⁶<https://huggingface.co/datasets/berkeley-nest/Nectar>

⁷<https://huggingface.co/google/gemma-2b-it>

⁸<https://huggingface.co/Ray2333/reward-model-Mistral-7B-instruct-Unified-Feedback>

Acknowledgements

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2022-NR071853, RS-2023-00222663, and RS-2025-25420849), by the Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. RS-2025-02263754 and RS-2025-25463302), by the AI-Bio Research Grant through Seoul National University, and by the 2025 Global Google PhD Fellowship funded with support from Google.org.

References

- [1] Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 24:2312–2320, 2011.
- [2] Marc Abeille, Louis Faury, and Clément Calauzènes. Instance-wise minimax-optimal algorithms for logistic bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 3691–3699. PMLR, 2021.
- [3] Priyank Agrawal, Theja Tulabandhula, and Vashist Avadhanula. A tractable online learning algorithm for the multinomial logit contextual bandit. *European Journal of Operational Research*, 310(2):737–750, 2023.
- [4] Shipra Agrawal, Vashist Avadhanula, Vineet Goyal, and Assaf Zeevi. Thompson sampling for the mnl-bandit. In *Conference on learning theory*, pages 76–78. PMLR, 2017.
- [5] Shipra Agrawal, Vashist Avadhanula, Vineet Goyal, and Assaf Zeevi. Mnl-bandit: A dynamic learning approach to assortment selection. *Operations Research*, 67(5):1453–1485, 2019.
- [6] Hossein Azari Soufiani, William Chen, David C Parkes, and Lirong Xia. Generalized method-of-moments for rank aggregation. *Advances in Neural Information Processing Systems*, 26, 2013.
- [7] Akshay Balsubramani, Zohar Karnin, Robert E Schapire, and Masrour Zoghi. Instance-dependent regret bounds for dueling bandits. In *Conference on Learning Theory*, pages 336–360. PMLR, 2016.
- [8] Viktor Bengs, Róbert Busa-Fekete, Adil El Mesaoudi-Paul, and Eyke Hüllermeier. Preference-based online learning with dueling bandits: A survey. *Journal of Machine Learning Research*, 22(7):1–108, 2021.
- [9] Viktor Bengs, Aadirupa Saha, and Eyke Hüllermeier. Stochastic contextual dueling bandits under linear stochastic transitivity models. In *International Conference on Machine Learning*, pages 1764–1786. PMLR, 2022.
- [10] Alina Beygelzimer, John Langford, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandit algorithms with supervised learning guarantees. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 19–26. JMLR Workshop and Conference Proceedings, 2011.
- [11] Ralph Allan Bradley and Milton E Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952.
- [12] Róbert Busa-Fekete, Balázs Szörényi, Paul Weng, Weiwei Cheng, and Eyke Hüllermeier. Preference-based reinforcement learning: evolutionary direct policy search using a preference-based racing algorithm. *Machine learning*, 97:327–351, 2014.
- [13] Shicong Cen, Jincheng Mei, Katayoon Goshvadi, Hanjun Dai, Tong Yang, Sherry Yang, Dale Schuurmans, Yuejie Chi, and Bo Dai. Value-incentivized preference optimization: A unified approach to online and offline RLHF. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=S0qnitDuow6>.

[14] Mingyu Chen, Yiding Chen, Wen Sun, and Xuezhou Zhang. Avoiding exp (r) scaling in rlhf through preference-based exploration. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.

[15] Xi Chen, Yining Wang, and Yuan Zhou. Dynamic assortment optimization with changing contextual information. *The Journal of Machine Learning Research*, 21(1):8918–8961, 2020.

[16] Xiaoyu Chen, Han Zhong, Zhuoran Yang, Zhaoran Wang, and Liwei Wang. Human-in-the-loop: Provably efficient preference-based reinforcement learning with general function approximation. In *International Conference on Machine Learning*, pages 3773–3793. PMLR, 2022.

[17] Wooseong Cho, Taehyun Hwang, Joongkyu Lee, and Min-hwan Oh. Randomized exploration for reinforcement learning with multinomial logistic function approximation. *Advances in Neural Information Processing Systems*, 37:76643–76720, 2024.

[18] Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30, 2017.

[19] Nirjhar Das, Souradip Chakraborty, Aldo Pacchiano, and Sayak Ray Chowdhury. Active preference optimization for sample efficient rlhf. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 96–112. Springer, 2025.

[20] Qiwei Di, Jiafan He, and Quanquan Gu. Nearly optimal algorithms for contextual dueling bandits from adversarial feedback. In *Forty-second International Conference on Machine Learning*, 2024.

[21] Qiwei Di, Tao Jin, Yue Wu, Heyang Zhao, Farzad Farnoud, and Quanquan Gu. Variance-aware regret bounds for stochastic contextual dueling bandits. In *The Twelfth International Conference on Learning Representations*, 2024.

[22] Hanze Dong, Wei Xiong, Bo Pang, Haoxiang Wang, Han Zhao, Yingbo Zhou, Nan Jiang, Doyen Sahoo, Caiming Xiong, and Tong Zhang. Rlhf workflow: From reward modeling to online rlhf. *Transactions on Machine Learning Research*, 2024.

[23] Shi Dong, Tengyu Ma, and Benjamin Van Roy. On the performance of thompson sampling on logistic bandits. In *Conference on Learning Theory*, pages 1158–1160. PMLR, 2019.

[24] Miroslav Dudík, Katja Hofmann, Robert E Schapire, Aleksandrs Slivkins, and Masrour Zoghi. Contextual dueling bandits. In *Conference on Learning Theory*, pages 563–587. PMLR, 2015.

[25] Moein Falahatgar, Yi Hao, Alon Orlitsky, Venkatadheeraj Pichapati, and Vaishakh Ravindrakumar. Maxing and ranking with few assumptions. *Advances in Neural Information Processing Systems*, 30, 2017.

[26] Louis Faury, Marc Abeille, Clément Calauzènes, and Olivier Fercoq. Improved optimistic algorithms for logistic bandits. In *International Conference on Machine Learning*, pages 3052–3060. PMLR, 2020.

[27] Louis Faury, Marc Abeille, Kwang-Sung Jun, and Clément Calauzènes. Jointly efficient and optimal algorithms for logistic bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 546–580. PMLR, 2022.

[28] Reinhard Heckel, Max Simchowitz, Kannan Ramchandran, and Martin Wainwright. Approximate ranking from pairwise comparisons. In *International Conference on Artificial Intelligence and Statistics*, pages 1057–1066. PMLR, 2018.

[29] Taehyun Hwang and Min-hwan Oh. Model-based reinforcement learning with multinomial logistic function approximation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pages 7971–7979, 2023.

[30] Ashesh Jain, Brian Wojcik, Thorsten Joachims, and Ashutosh Saxena. Learning trajectory preferences for manipulators via iterative improvement. *Advances in neural information processing systems*, 26, 2013.

[31] Kevin Jamieson, Sumeet Katariya, Atul Deshpande, and Robert Nowak. Sparse dueling bandits. In *Artificial Intelligence and Statistics*, pages 416–424. PMLR, 2015.

[32] Minje Jang, Sunghyun Kim, Changho Suh, and Sewoong Oh. Optimal sample complexity of m -wise data for top- k ranking. *Advances in Neural Information Processing Systems*, 30, 2017.

[33] Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. Mistral 7b, 2023. URL <https://arxiv.org/abs/2310.06825>.

[34] Ashish Khetan and Sewoong Oh. Data-driven rank breaking for efficient rank aggregation. *Journal of Machine Learning Research*, 17(193):1–54, 2016.

[35] Junpei Komiyama, Junya Honda, Hisashi Kashima, and Hiroshi Nakagawa. Regret lower bound and optimal algorithm in dueling bandit problem. In *Conference on learning theory*, pages 1141–1154. PMLR, 2015.

[36] Junpei Komiyama, Junya Honda, and Hiroshi Nakagawa. Copeland dueling bandit problem: Regret lower bound, optimal algorithm, and computationally efficient algorithm. In *International Conference on Machine Learning*, pages 1235–1244. PMLR, 2016.

[37] Andras Kupcsik, David Hsu, and Wee Sun Lee. Learning dynamic robot-to-human object handover from human feedback. In *Robotics Research: Volume 1*, pages 161–176. Springer, 2017.

[38] Branislav Kveton, Xintong Li, Julian McAuley, Ryan Rossi, Jingbo Shang, Junda Wu, and Tong Yu. Active learning for direct preference optimization. *arXiv preprint arXiv:2503.01076*, 2025.

[39] Joongkyu Lee and Min-hwan Oh. Nearly minimax optimal regret for multinomial logistic bandit. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.

[40] Joongkyu Lee and Min-hwan Oh. Combinatorial reinforcement learning with preference feedback. In *Forty-second International Conference on Machine Learning*, 2025.

[41] Joongkyu Lee and Min-hwan Oh. Improved online confidence bounds for multinomial logistic bandits. In *Forty-second International Conference on Machine Learning*, 2025.

[42] Junghyun Lee, Se-Young Yun, and Kwang-Sung Jun. Improved regret bounds of (multinomial) logistic bandits via regret-to-confidence-set conversion. In *International Conference on Artificial Intelligence and Statistics*, pages 4474–4482. PMLR, 2024.

[43] Junghyun Lee, Se-Young Yun, and Kwang-Sung Jun. A unified confidence sequence for generalized linear models, with applications to bandits. *Advances in Neural Information Processing Systems*, 37:124640–124685, 2024.

[44] Zhihan Liu, Miao Lu, Shenao Zhang, Boyi Liu, Hongyi Guo, Yingxiang Yang, Jose Blanchet, and Zhaoran Wang. Provably mitigating overoptimization in RLHF: Your SFT loss is implicitly an adversarial regularizer. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=2cQ31Phke0>.

[45] R Duncan Luce et al. *Individual choice behavior*, volume 4. Wiley New York, 1959.

[46] Daniel McFadden. Modelling the choice of residential location. 1977.

[47] Viraj Mehta, Vikramjeet Das, Ojash Neopane, Yijia Dai, Ilija Bogunovic, Jeff Schneider, and Willie Neiswanger. Sample efficient reinforcement learning from human feedback via active exploration. *arXiv preprint arXiv:2312.00267*, 2023.

[48] Zakaria Mhammedi, Wouter M Koolen, and Tim Van Erven. Lipschitz adaptivity with multiple learning rates in online learning. In *Conference on Learning Theory*, pages 2490–2511. PMLR, 2019.

[49] Subhjoyoti Mukherjee, Anusha Lalitha, Kousha Kalantari, Aniket Anand Deshmukh, Ge Liu, Yifei Ma, and Branislav Kveton. Optimal design for human preference elicitation. *Advances in Neural Information Processing Systems*, 37:90132–90159, 2024.

[50] William Muldrew, Peter Hayes, Mingtian Zhang, and David Barber. Active preference learning for large language models. In *Forty-first International Conference on Machine Learning*, 2024.

[51] Rémi Munos, Michal Valko, Daniele Calandriello, Mohammad Gheshlaghi Azar, Mark Rowland, Zhaohan Daniel Guo, Yunhao Tang, Matthieu Geist, Thomas Mesnard, Côme Fiebel, et al. Nash learning from human feedback. In *Forty-first International Conference on Machine Learning*, 2023.

[52] Ellen Novoseller, Yibing Wei, Yanan Sui, Yisong Yue, and Joel Burdick. Dueling posterior sampling for preference-based reinforcement learning. In *Conference on Uncertainty in Artificial Intelligence*, pages 1029–1038. PMLR, 2020.

[53] Min-hwan Oh and Garud Iyengar. Thompson sampling for multinomial logit contextual bandits. *Advances in Neural Information Processing Systems*, 32, 2019.

[54] Min-hwan Oh and Garud Iyengar. Multinomial logit contextual bandits: Provable optimality and practicality. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 9205–9213, 2021.

[55] Francesco Orabona. A modern introduction to online learning. *arXiv preprint arXiv:1912.13213*, 2019.

[56] Mingdong Ou, Nan Li, Shenghuo Zhu, and Rong Jin. Multinomial logit bandit with linear utility functions. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 2602–2608. International Joint Conferences on Artificial Intelligence Organization, 2018.

[57] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022.

[58] Jaehyun Park, Junyeop Kwon, and Dabeen Lee. Infinite-horizon reinforcement learning with multinomial logit function approximation. In *International Conference on Artificial Intelligence and Statistics*, pages 361–369. PMLR, 2025.

[59] Noemie Perivier and Vineet Goyal. Dynamic pricing and assortment under a contextual mnl demand. *Advances in Neural Information Processing Systems*, 35:3461–3474, 2022.

[60] Robin L Plackett. The analysis of permutations. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 24(2):193–202, 1975.

[61] Filip Radlinski, Robert Kleinberg, and Thorsten Joachims. Learning diverse rankings with multi-armed bandits. In *Proceedings of the 25th international conference on Machine learning*, pages 784–791, 2008.

[62] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36:53728–53741, 2023.

[63] Siddartha Y Ramamohan, Arun Rajkumar, and Shivani Agarwal. Dueling bandits: Beyond condorcet winners to general tournament solutions. *Advances in Neural Information Processing Systems*, 29, 2016.

[64] Sayak Ray Chowdhury, Anush Kini, and Nagarajan Natarajan. Provably robust DPO: Aligning language models with noisy feedback. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 42258–42274. PMLR, 21–27 Jul 2024.

[65] Dorsa Sadigh, Anca Dragan, Shankar Sastry, and Sanjit Seshia. *Active preference-based learning of reward functions*. 2017.

[66] Aadirupa Saha. Optimal algorithms for stochastic contextual preference bandits. *Advances in Neural Information Processing Systems*, 34:30050–30062, 2021.

[67] Aadirupa Saha and Pierre Gaillard. Finally rank-breaking conquers mnl bandits: Optimal and efficient algorithms for mnl assortment. In *The Thirteenth International Conference on Learning Representations*, 2025.

[68] Aadirupa Saha and Aditya Gopalan. Battle of bandits. In *UAI*, pages 805–814, 2018.

[69] Aadirupa Saha, Tomer Koren, and Yishay Mansour. Adversarial dueling bandits. In *International Conference on Machine Learning*, pages 9235–9244. PMLR, 2021.

[70] Aadirupa Saha, Aldo Pacchiano, and Jonathan Lee. Dueling rl: Reinforcement learning with trajectory preferences. In *International conference on artificial intelligence and statistics*, pages 6263–6289. PMLR, 2023.

[71] Antoine Scheid, Etienne Boursier, Alain Durmus, Michael I Jordan, Pierre Ménard, Eric Moulines, and Michal Valko. Optimal design for reward modeling in rlhf. *arXiv preprint arXiv:2410.17055*, 2024.

[72] Ayush Sekhari, Karthik Sridharan, Wen Sun, and Runzhe Wu. Contextual bandits and imitation learning with preference-based active queries. *Advances in Neural Information Processing Systems*, 36:11261–11295, 2023.

[73] Ohad Shamir. On the complexity of bandit and derivative-free stochastic convex optimization. In *Conference on learning theory*, pages 3–24. PMLR, 2013.

[74] Gokul Swamy, Christoph Dann, Rahul Kidambi, Zhiwei Steven Wu, and Alekh Agarwal. A minimaximalist approach to reinforcement learning from human feedback. In *Proceedings of the 41st International Conference on Machine Learning*, pages 47345–47377, 2024.

[75] Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*, 2024.

[76] Kiran Koshy Thekumpampil, Gaurush Hiranandani, Kousha Kalantari, Shoham Sabach, and Branislav Kveton. Comparing few to rank many: Active human preference learning using randomized frank-wolfe method. In *Forty-second International Conference on Machine Learning*, 2024.

[77] Jean Ville. *Etude critique de la notion de collectif*, volume 3. Gauthier-Villars Paris, 1939.

[78] Andrew J Wagenmaker, Yifang Chen, Max Simchowitz, Simon Du, and Kevin Jamieson. Reward-free rl is no harder than reward-aware rl in linear markov decision processes. In *International Conference on Machine Learning*, pages 22430–22456. PMLR, 2022.

[79] Christian Wirth and Johannes Fürnkranz. Preference-based reinforcement learning: A preliminary survey. In *Proceedings of the ECML/PKDD-13 Workshop on Reinforcement Learning from Generalized Feedback: Beyond Numeric Rewards*, 2013.

[80] Christian Wirth, Johannes Fürnkranz, and Gerhard Neumann. Model-free preference-based reinforcement learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30, 2016.

[81] Christian Wirth, Riad Akroud, Gerhard Neumann, and Johannes Fürnkranz. A survey of preference-based reinforcement learning methods. *Journal of Machine Learning Research*, 18(136):1–46, 2017.

[82] Huasen Wu and Xin Liu. Double thompson sampling for dueling bandits. *Advances in neural information processing systems*, 29, 2016.

[83] Runzhe Wu and Wen Sun. Making rl with preference-based feedback efficient via randomization. In *The Twelfth International Conference on Learning Representations*, 2023.

[84] Yue Wu, Tao Jin, Qiwei Di, Hao Lou, Farzad Farnoud, and Quanquan Gu. Borda regret minimization for generalized linear dueling bandits. In *International Conference on Machine Learning*, pages 53571–53596. PMLR, 2024.

[85] Tengyang Xie, Dylan J Foster, Akshay Krishnamurthy, Corby Rosset, Ahmed Hassan Awadallah, and Alexander Rakhlin. Exploratory preference optimization: Harnessing implicit q^* -approximation for sample-efficient RLHF. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=QYigQ6gXNw>.

[86] Wei Xiong, Hanze Dong, Chenlu Ye, Ziqi Wang, Han Zhong, Heng Ji, Nan Jiang, and Tong Zhang. Iterative preference learning from human feedback: Bridging theory and practice for rlhf under kl-constraint. In *Forty-first International Conference on Machine Learning*, 2024.

[87] Yichong Xu, Ruosong Wang, Lin Yang, Aarti Singh, and Artur Dubrawski. Preference-based reinforcement learning with finite-time guarantees. *Advances in Neural Information Processing Systems*, 33:18784–18794, 2020.

[88] Yisong Yue, Josef Broder, Robert Kleinberg, and Thorsten Joachims. The k-armed dueling bandits problem. *Journal of Computer and System Sciences*, 78(5):1538–1556, 2012.

[89] Andrea Zanette, Ching-An Cheng, and Alekh Agarwal. Cautiously optimistic policy optimization and exploration with linear function approximation. In *Conference on Learning Theory*, pages 4473–4525. PMLR, 2021.

[90] Wenhao Zhan, Masatoshi Uehara, Nathan Kallus, Jason D Lee, and Wen Sun. Provable offline preference-based reinforcement learning. In *The Twelfth International Conference on Learning Representations*, 2024.

[91] Wenhao Zhan, Masatoshi Uehara, Wen Sun, and Jason D Lee. Provable reward-agnostic preference-based reinforcement learning. In *The Twelfth International Conference on Learning Representations*, 2024.

[92] Yu-Jie Zhang and Masashi Sugiyama. Online (multinomial) logistic bandit: Improved regret and constant computation cost. *Advances in Neural Information Processing Systems*, 36, 2024.

[93] Banghua Zhu, Michael Jordan, and Jiantao Jiao. Principled reinforcement learning with human feedback from pairwise or k-wise comparisons. In *International Conference on Machine Learning*, pages 43037–43067. PMLR, 2023.

[94] Masrour Zoghi, Shimon Whiteson, Remi Munos, and Maarten Rijke. Relative upper confidence bound for the k-armed dueling bandit problem. In *International conference on machine learning*, pages 10–18. PMLR, 2014.

[95] Masrour Zoghi, Zohar S Karnin, Shimon Whiteson, and Maarten De Rijke. Copeland dueling bandits. *Advances in neural information processing systems*, 28, 2015.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: Abstract and Section 1

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: Appendix I

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[Yes\]](#)

Justification: Assumption 1 and Appendix

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [\[Yes\]](#)

Justification: Section 6 and Appendix H

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: We have included the code in the supplementary material. After our paper is accepted, we will provide open access to the code.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Section 6 and Appendix H

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: Figure 1 and 2.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: Appendix H

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [\[Yes\]](#)

Justification: We comply with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [\[NA\]](#)

Justification: This paper focuses on theoretical results and therefore does not discuss societal impacts.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: No risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: Section 6

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: **[Yes]**

Justification: We will release the full code with documentation to support reproducibility.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: **[NA]**

Justification: This paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: **[NA]**

Justification: The study does not involve human participants or sensitive data.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorosity, or originality of the research, declaration is not required.

Answer: [NA]

Justification: We used LLMs only for writing and editing.

Guidelines: The study does not utilize LLMs in the core methodology or experimental pipeline.

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

Appendix

Table of Contents

A	Further Related Work	24
B	Notation	25
C	Proof of Theorem 1	26
C.1	Main Proof of Theorem 1	26
C.2	Proofs of Lemmas for Theorem 1	31
D	Proof of Theorem 2	36
D.1	Main Proof of Theorem 2	36
D.2	Proofs of Propositions and Lemmas for Theorem 2	38
E	Technical Lemmas	42
F	Proof of Theorem 3	44
F.1	Main Proof of Theorem 3	44
F.2	Proof of Lemmas for Theorem 3	47
G	Additional Discussions	48
G.1	Avoiding e^B Scaling in Regret Minimization	48
G.2	Extension to Active Learning Setting	50
H	Experimental Details and Additional Results	51
H.1	Synthetic Data	51
H.2	Real-World Dataset	54
I	Limitations	56

A Further Related Work

In this section, we provide additional related work that complements Section 2.

Logistic and MNL bandits. Our work is also closely related to logistic bandits and multinomial logit (MNL) bandits. The logistic bandit problem [23, 26, 2, 27, 42, 43] is a special case of the MNL bandit model in which the agent offers only a single item (i.e., $K = 1$) at each round and receives binary feedback indicating whether the item was selected (1) or not (0). Faury et al. [26] examined how the regret in logistic bandits depends on the non-linearity parameter κ of the logistic link function and proposed the first algorithm whose regret bound eliminates explicit dependence on $1/\kappa = \mathcal{O}(e^B)$. Abeille et al. [2] further improved the theoretical dependency on $1/\kappa$ and established a matching, problem-dependent lower bound. Building on this, Faury et al. [27] developed a computationally efficient algorithm whose regret still matches the lower bound established by Abeille et al. [2].

Multinomial logit (MNL) bandits tackle a more sophisticated problem than logistic bandits. Instead of offering a single item and observing binary feedback, the learner chooses a subset of items—underscoring the combinatorial nature of the task—and receives non-uniform rewards driven by an MNL choice model [5, 4, 56, 15, 53, 54, 59, 3, 39, 41]. A recent breakthrough by Lee and Oh [39] closed a long-standing gap by providing a computationally efficient algorithm that attains the minimax-optimal regret for this setting. Building on this result, Lee and Oh [41] further reduced the regret bound by a factor polynomial in B and logarithmic in K , and established the first variance-dependent regret bounds for MNL bandits.

Our work extends the online confidence bound analysis of Lee and Oh [41] to the Plackett–Luce (PL) model. This extension is natural because the PL probability distribution decomposes into a sequence of MNL probabilities over successive choices. Crucially, we leverage their key insight—that the MNL loss exhibits an ℓ_∞ -self-concordant property—to eliminate the harmful $\mathcal{O}(e^B)$ dependence. This is one of the main contributions of our work (see Lemma C.2).

RL with MNL models. Recent work has extended the Multinomial Logit (MNL) framework beyond bandit formulations to reinforcement learning. Lee and Oh [40] introduced *combinatorial RL with preference feedback*, a framework in which an agent learns to select subsets of items so as to maximize long-term cumulative rewards.

Another line of research incorporates MNL models directly into the transition dynamics. Hwang and Oh [29] proposed MNL-MDPs, a class of Markov decision processes whose transition probabilities follow an MNL parameterization. Building upon this formulation, Cho et al. [17] improved the regret bounds by improving the exponential dependence on B , and Park et al. [58] extended the analysis to the infinite-horizon setting.

B Notation

Let T denote the total number of rounds, with $t \in [T]$ representing the current round. We use N for the total number of items, K for the maximum assortment size, d for the feature vector dimension, and B as an upper bound on the norm of the unknown parameter. For notational convenience, we provide Table B.1.

For clarity, we derive the first- and second-order derivatives (i.e., gradients and Hessians) of the loss functions. For the PL loss at round t for the j 'th ranking, let $y_{ti}^{(j)} = 1$ if $i = j$, and $y_{ti}^{(j)} = 0$ for otherwise. Then, we have

$$\begin{aligned} \ell_t^{(j)}(\boldsymbol{\theta}) &= -\log \left(\frac{\exp(\phi(x_t, \sigma_{tj})^\top \boldsymbol{\theta})}{\sum_{k=j}^{|S_t|} \exp(\phi(x_t, \sigma_{tk})^\top \boldsymbol{\theta}))} \right) = -\sum_{i=j}^{|S_t|} y_{ti}^{(j)} \log \left(\underbrace{\frac{\exp(\phi(x_t, \sigma_{ti})^\top \boldsymbol{\theta})}{\sum_{k=j}^{|S_t|} \exp(\phi(x_t, \sigma_{tk})^\top \boldsymbol{\theta}))}}_{=: P_{t,\boldsymbol{\theta}}^{(j)}(\sigma_{ti})} \right) \\ &= -\sum_{i=j}^{|S_t|} y_{ti}^{(j)} \log P_{t,\boldsymbol{\theta}}^{(j)}(\sigma_{ti}), \\ \nabla \ell_t^{(j)}(\boldsymbol{\theta}) &= \sum_{i=j}^{|S_t|} \left(P_{t,\boldsymbol{\theta}}^{(j)}(\sigma_{ti}) - y_{ti}^{(j)} \right) \phi(x_t, \sigma_{ti}), \\ \nabla^2 \ell_t^{(j)}(\boldsymbol{\theta}) &= \sum_{i=j}^{|S_t|} P_{t,\boldsymbol{\theta}}^{(j)}(\sigma_{ti}) \phi(x_t, \sigma_{ti}) \phi(x_t, \sigma_{ti})^\top - \sum_{i=j}^{|S_t|} \sum_{k=j}^{|S_t|} P_{t,\boldsymbol{\theta}}^{(j)}(\sigma_{ti}) P_{t,\boldsymbol{\theta}}^{(j)}(\sigma_{tk}) \phi(x_t, \sigma_{ti}) \phi(x_t, \sigma_{tk})^\top \\ &= \frac{1}{2} \sum_{i=j}^{|S_t|} \sum_{k=j}^{|S_t|} P_{t,\boldsymbol{\theta}}^{(j)}(\sigma_{ti}) P_{t,\boldsymbol{\theta}}^{(j)}(\sigma_{tk}) (\phi(x_t, \sigma_{ti}) - \phi(x_t, \sigma_{tk})) (\phi(x_t, \sigma_{ti}) - \phi(x_t, \sigma_{tk}))^\top. \end{aligned}$$

For the RB loss at round t for the pairwise comparison between σ_{tj} and σ_{tk} , let $y_{ti}^{(j,k)} = 1$ if $i = j$, and $y_{ti}^{(j,k)} = 0$ for otherwise (i.e., when $i = k$). Then, we have

$$\begin{aligned} \ell_t^{(j,k)}(\boldsymbol{\theta}) &= -\log \left(\frac{\exp(\phi(x_t, \sigma_{tj})^\top \boldsymbol{\theta})}{\exp(\phi(x_t, \sigma_{tj})^\top \boldsymbol{\theta}) + \exp(\phi(x_t, \sigma_{tk})^\top \boldsymbol{\theta})} \right) \\ &= -\log \mu \left((\phi(x_t, \sigma_{tj}) - \phi(x_t, \sigma_{tk}))^\top \boldsymbol{\theta} \right), \quad \text{where } \mu(w) = \frac{1}{1 + e^{-w}}, \\ \nabla \ell_t^{(j,k)}(\boldsymbol{\theta}) &= \left(\mu \left((\phi(x_t, \sigma_{tj}) - \phi(x_t, \sigma_{tk}))^\top \boldsymbol{\theta} \right) - 1 \right) (\phi(x_t, \sigma_{tj}) - \phi(x_t, \sigma_{tk})), \\ \nabla^2 \ell_t^{(j,k)}(\boldsymbol{\theta}) &= \dot{\mu} \left((\phi(x_t, \sigma_{tj}) - \phi(x_t, \sigma_{tk}))^\top \boldsymbol{\theta} \right) (\phi(x_t, \sigma_{tj}) - \phi(x_t, \sigma_{tk})) (\phi(x_t, \sigma_{tj}) - \phi(x_t, \sigma_{tk}))^\top. \end{aligned}$$

Table B.1: Symbols

$\mathcal{X}, \mathcal{A}, \mathcal{S}$	context (prompt) space, action (answer) space, assortment space
$\phi(x, a) \in \mathbb{R}^d$	feature representation of context-action pair (x, a)
z_{tjk}	$:= \phi(x_t, \sigma_{tj}) - \phi(x_t, \sigma_{tk})$, feature difference between σ_{tj} and σ_{tk} under context x_t
S_t	assortment chosen by an algorithm at round t
$\ell_t^{(j)}(\boldsymbol{\theta})$	$:= -\log \left(\frac{\exp(\phi(x_t, \sigma_{tj})^\top \boldsymbol{\theta})}{\sum_{k=j}^{ S_t } \exp(\phi(x_t, \sigma_{tk})^\top \boldsymbol{\theta})} \right)$, PL loss at round t for j 'th ranking
$\ell_t^{(j,k)}(\boldsymbol{\theta})$	$:= -\log \left(\frac{\exp(\phi(x_t, \sigma_{tj})^\top \boldsymbol{\theta})}{\sum_{m \in \{j, k\}} \exp(\phi(x_t, \sigma_{tm})^\top \boldsymbol{\theta})} \right)$, RB loss at round t for comparison σ_{tj} vs σ_{tk}
$\nabla^2 \ell_t^{(j)}(\boldsymbol{\theta})$	$= \sum_{k=j}^{ S_t } \sum_{k'=j}^{ S_t } \frac{\exp((\phi(x_t, \sigma_{tk}) + \phi(x_t, \sigma_{tk'}))^\top \boldsymbol{\theta})}{2 \left(\sum_{k'=j}^{ S_t } \exp(\phi(x_t, \sigma_{tk'})^\top \boldsymbol{\theta}) \right)^2} \cdot z_{tjk} z_{tjk'}^\top$
$\nabla^2 \ell_t^{(j,k)}(\boldsymbol{\theta})$	$= \dot{\mu}(z_{tjk}^\top \boldsymbol{\theta}) z_{tjk} z_{tjk}^\top$, where $\mu(w) = \frac{1}{1+e^{-w}}$ is sigmoid function
$\hat{\boldsymbol{\theta}}_t^{(j+1)}$	online parameter estimate using PL loss at round t , after j 'th update
$\hat{\boldsymbol{\theta}}_t^{(j,k+1)}$	online parameter estimate using RB loss at round t , after (j, k) 'th comparison update
η	$:= \frac{1}{2}(1 + 3\sqrt{2}B)$, step-size parameter
λ	$:= \Omega(d \log(KT/\delta) + \eta(B+d))$, regularization parameter
H_t	$:= \sum_{s=1}^{t-1} \sum_{j=1}^{ S_s } \nabla^2 \ell_s^{(j)}(\hat{\boldsymbol{\theta}}_s^{(j+1)}) + \lambda \mathbf{I}_d$ (or $\sum_{s=1}^{t-1} \sum_{j=1}^{ S_s -1} \sum_{k=j+1}^{ S_s } \nabla^2 \ell_s^{(j,k)}(\hat{\boldsymbol{\theta}}_s^{(j,k+1)}) + \lambda \mathbf{I}_d$)
$\tilde{H}_t^{(j)}$	$:= H_t + \eta \sum_{j'=1}^j \nabla^2 \ell_t^{(j')}(\hat{\boldsymbol{\theta}}_t^{(j')})$ (for PL loss)
$\tilde{H}_t^{(j,k)}$	$:= H_t + \eta \sum_{(j',k') \leq (j,k)} \nabla^2 \ell_t^{(j',k')}(\hat{\boldsymbol{\theta}}_t^{(j',k')})$ (for RB loss)
$\beta_t(\delta)$	$:= \mathcal{O}(B\sqrt{d \log(tK/\delta)} + B\sqrt{\lambda})$, confidence radius for $\boldsymbol{\theta}_t$ at round t
\mathcal{T}^w	$:= \left\{ t \in [T] : \max_{a, a' \in \mathcal{A}} \ \phi(x_t, a) - \phi(x_t, a')\ _{H_t^{-1}} \geq \frac{1}{3\sqrt{2}K\beta_{T+1}(\delta)} \right\}$, warm-up rounds
M_t	$:= \sum_{s \in [t-1] \setminus \mathcal{T}^w} \sum_{j=1}^{ S_s } \mathbb{E}_{\sigma \sim \mathbb{P}_s(\cdot S_s, \boldsymbol{\theta}^*)} \left[\nabla^2 \ell_s^{(j)}(\hat{\boldsymbol{\theta}}_s) \right] + \lambda \mathbf{I}_d$

C Proof of Theorem 1

In this section, we present the proof of Theorem 1.

C.1 Main Proof of Theorem 1

PL loss and OMD. We begin by recalling the loss function and the parameter update rule. Specifically, we use the PL loss defined in Equation (3) and update the parameter according to Equation (5).

$$\ell_t(\boldsymbol{\theta}) := \sum_{j=1}^{|S_t|} \underbrace{-\log \left(\frac{\exp(\phi(x_t, \sigma_{tj})^\top \boldsymbol{\theta})}{\sum_{k=j}^{|S_t|} \exp(\phi(x_t, \sigma_{tk})^\top \boldsymbol{\theta})} \right)}_{=: \ell_t^{(j)}(\boldsymbol{\theta})} = \sum_{j=1}^{|S_t|} \ell_t^{(j)}(\boldsymbol{\theta}),$$

and

$$\hat{\boldsymbol{\theta}}_t^{(j+1)} = \underset{\boldsymbol{\theta} \in \Theta}{\operatorname{argmin}} \langle \nabla \ell_t^{(j)}(\hat{\boldsymbol{\theta}}_t^{(j)}), \boldsymbol{\theta} \rangle + \frac{1}{2\eta} \|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_t^{(j)}\|_{\tilde{H}_t^{(j)}}^2, \quad j = 1, \dots, |S_t|,$$

where $\hat{\boldsymbol{\theta}}_t^{(|S_t|+1)} = \hat{\boldsymbol{\theta}}_{t+1}^{(1)}$, and $\eta := \frac{1}{2}(1 + 3\sqrt{2}B)$ is the step-size parameter. The matrix $\tilde{H}_t^{(j)}$ is given by $\tilde{H}_t^{(j)} := H_t + \eta \sum_{j'=1}^j \nabla^2 \ell_t^{(j')}(\hat{\boldsymbol{\theta}}_t^{(j')})$, where

$$H_t := \sum_{s=1}^{t-1} \sum_{j=1}^{|S_s|} \nabla^2 \ell_s^{(j)}(\hat{\boldsymbol{\theta}}_s^{(j+1)}) + \lambda \mathbf{I}_d, \quad \lambda > 0.$$

Online confidence bound for PL loss. Now, we present the confidence bound for online parameter estimation in MNL models, as recently proposed by Lee and Oh [41].

Lemma C.1 (Online confidence bound, Theorem 4.2 of Lee and Oh 41). *Let $\delta \in (0, 1]$. We set $\eta = (1 + 3\sqrt{2}B)/2$ and $\lambda = \max\{12\sqrt{2}B\eta, 144\eta d, 2\}$. Then, under Assumption 1, with probability*

at least $1 - \delta$, we have

$$\|\hat{\boldsymbol{\theta}}_t - \boldsymbol{\theta}^*\|_{H_t} \leq \beta_t(\delta) = \mathcal{O}\left(B\sqrt{d\log(t/\delta)} + B\sqrt{\lambda}\right), \quad \forall t \geq 1.$$

We now extend this result to our setting. Since the total number of updates up to round t is $\sum_{s=1}^t |S_s|$, the corresponding confidence bound can be expressed as follows:

Corollary C.1 (Restatement of Corollary 1, Online confidence bound for PL loss). *Let $\delta \in (0, 1]$. We set $\eta = (1 + 3\sqrt{2}B)/2$ and $\lambda = \max\{12\sqrt{2}B\eta, 144\eta d, 2\}$. Then, under Assumption 1, with probability at least $1 - \delta$, we have*

$$\|\hat{\boldsymbol{\theta}}_t^{(j)} - \boldsymbol{\theta}^*\|_{H_t^{(j)}} \leq \beta_t(\delta) = \mathcal{O}\left(B\sqrt{d\log(tK/\delta)} + B\sqrt{\lambda}\right), \quad \forall t \geq 1, j \leq |S_t|,$$

where $H_t^{(j)} := H_t + \sum_{j'=1}^{j-1} \nabla^2 \ell_s^{(j')}(\hat{\boldsymbol{\theta}}_s^{(j'+1)}) + \lambda \mathbf{I}_d$ and $\hat{\boldsymbol{\theta}}_t^{(1)} = \hat{\boldsymbol{\theta}}_t$.

Useful definitions. For clarity, we write $\ell_{t,\sigma}(\boldsymbol{\theta})$ to make explicit that the loss function is random with respect to the ranking σ . When the ranking is realized as σ_t , we simply denote it by $\ell_t(\boldsymbol{\theta})$. We define the set of *warm-up rounds*, denoted by \mathcal{T}^w , which consists of rounds with large uncertainty, as:

$$\mathcal{T}^w := \left\{ t \in [T] : \max_{a, a' \in \mathcal{A}} \|\phi(x_t, a) - \phi(x_t, a')\|_{H_t^{-1}} \geq \frac{1}{3\sqrt{2}K\beta_{T+1}(\delta)} \right\}, \quad (\text{C.1})$$

where $\beta_{T+1}(\delta)$ denotes the confidence radius as defined in Corollary C.1.

Furthermore, we define the expected version of the Hessian matrix (taken with respect to the randomness of the ranking feedback) as follows:

$$M_t = \sum_{s \in [t-1] \setminus \mathcal{T}^w} \sum_{j=1}^{|S_s|} \mathbb{E}_{\sigma \sim \mathbb{P}_s(\cdot | S_s, \boldsymbol{\theta}^*)} \left[\nabla^2 \ell_{s,\sigma}^{(j)}(\hat{\boldsymbol{\theta}}_s) \right] + \lambda \mathbf{I}_d \quad (\text{C.2})$$

Note that $\nabla^2 \ell_{s,\sigma}^{(j)}(\hat{\boldsymbol{\theta}}_s)$ is a random matrix conditional on the assortment S_s , where the randomness arises from the ranking feedback σ .

Key lemmas. We now present key lemmas needed to prove Theorem 1. The following lemma, one of our main contributions, is crucial for avoiding the $1/\kappa = \mathcal{O}(e^B)$ dependency in the leading term.

Lemma C.2 (Empirical-to-expected Hessian lower bound). *Let M_t be defined as in Equation (C.2). Set $\lambda = \Omega(d\log(KT/\delta))$. Then, for all $t \in [T]$, with probability at least $1 - \delta$, we have*

$$H_t \geq \frac{1}{3e^2} M_t.$$

The proof is deferred to Appendix C.2.1.

The following lemma is the elliptical potential lemma adapted to our setting.

Lemma C.3 (Elliptical potential for expected mean-centered uncertainty). *Let M_t be defined as in Equation (C.2). Let $P_t(a | S_t^{(j)}; \boldsymbol{\theta})$ denote the MNL probability for the subset of remaining actions in S_t after removing the first $j - 1$ actions according to the ranking σ . Then, for any $\lambda > 0$, we have*

$$\begin{aligned} & \sum_{t \notin \mathcal{T}^w} \min \left\{ 1, \sum_{j=1}^{|S_t|} \mathbb{E}_{\sigma \sim \mathbb{P}_t(\cdot | S_t; \boldsymbol{\theta}^*)} \left[\left\| \phi(x_t, a) - \mathbb{E}_{a \sim P_t(\cdot | S_t^{(j)}; \hat{\boldsymbol{\theta}}_t)} [\phi(x_t, a')] \right\|_{M_t^{-1}}^2 \right] \right\} \\ & \leq 2d \log \left(1 + \frac{KT}{d\lambda} \right). \end{aligned}$$

The proof is deferred to Appendix C.2.2.

The size of the set \mathcal{T}^w is bounded as described in the following lemma.

Lemma C.4. *Let $\mathcal{T}^w = \{t \in [T] : \max_{a, a' \in \mathcal{A}} \|\phi(x_t, a) - \phi(x_t, a')\|_{H_t^{-1}} \geq \frac{1}{3\sqrt{2}K\beta_{T+1}(\delta)}\}$. Define $\kappa := e^{-6B}$. Set $\lambda \geq 1$. Then, the size of the set \mathcal{T}^w is bounded as follows:*

$$|\mathcal{T}^w| \leq \frac{288K^4}{\kappa} \beta_{T+1}(\delta)^2 d \log \left(1 + \frac{2KT}{d\lambda} \right).$$

The proof is deferred to Appendix C.2.3.

We are now ready to provide the proof of Theorem 1.

Proof of Theorem 1. To begin, we define a martingale difference sequence (MDS) ζ_t as follows:

$$\zeta_t := \mathbb{E}_{x \sim \rho} \left[(\phi(x, \pi^*(x)) - \phi(x, \hat{\pi}_T(x)))^\top \boldsymbol{\theta}^* \right] - (\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t)))^\top \boldsymbol{\theta}^*,$$

which satisfies $|\zeta_t| \leq 2B$. Then, by the definition of the suboptimality gap, we have

$$\begin{aligned} \mathbf{SubOpt}(T) &= \mathbb{E}_{x \sim \rho} \left[(\phi(x, \pi^*(x)) - \phi(x, \hat{\pi}_T(x)))^\top \boldsymbol{\theta}^* \right] \\ &= \frac{1}{T} \sum_{t=1}^T (\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t)))^\top \boldsymbol{\theta}^* + \frac{1}{T} \sum_{t=1}^T \zeta_t && \text{(Def. of } \zeta_t\text{)} \\ &\leq \frac{1}{T} \sum_{t=1}^T (\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t)))^\top (\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{T+1}) + \frac{1}{T} \sum_{t=1}^T \zeta_t \\ &\quad (\hat{\pi}_T(x_t) = \operatorname{argmax}_{a \in \mathcal{A}} \phi(x_t, a)^\top \hat{\boldsymbol{\theta}}_{T+1}) \\ &\leq \frac{1}{T} \sum_{t=1}^T (\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t)))^\top (\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{T+1}) + \tilde{\mathcal{O}}\left(\frac{1}{\sqrt{T}}\right), && \text{(C.3)} \end{aligned}$$

where the last inequality follows from the Azuma–Hoeffding inequality. Specifically, for any $T \geq 1$, with probability at least $1 - \delta$, we have

$$\frac{1}{T} \sum_{t=1}^T \zeta_t \leq \frac{1}{T} \sqrt{8B^2 T \log(1/\delta)} = \tilde{\mathcal{O}}\left(\frac{1}{\sqrt{T}}\right).$$

To complete the proof, it remains to bound the first term in Equation (C.3).

$$\begin{aligned} &\frac{1}{T} \sum_{t=1}^T (\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t)))^\top (\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{T+1}) \\ &= \frac{1}{T} \sum_{t \in \mathcal{T}^w} (\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t)))^\top (\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{T+1}) \\ &\quad + \frac{1}{T} \sum_{t \notin \mathcal{T}^w} (\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t)))^\top (\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{T+1}) \\ &\leq \frac{4B}{T} |\mathcal{T}^w| + \frac{1}{T} \sum_{t \notin \mathcal{T}^w} (\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t)))^\top (\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{T+1}) && \text{(Assumption 1)} \\ &\leq \mathcal{O}\left(\frac{BK^4}{\kappa T} \beta_{T+1}(\delta)^2 d \log\left(1 + \frac{T}{d\lambda}\right)\right) && \text{(Lemma C.4)} \\ &\quad + \frac{1}{T} \sum_{t \notin \mathcal{T}^w} (\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t)))^\top (\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{T+1}). && \text{(C.4)} \end{aligned}$$

To further bound the last term of Equation (C.4), we get

$$\begin{aligned} &\frac{1}{T} \sum_{t \notin \mathcal{T}^w} (\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t)))^\top (\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{T+1}) \\ &\leq \frac{1}{T} \sum_{t \notin \mathcal{T}^w} \|\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t))\|_{H_{T+1}^{-1}} \|\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{T+1}\|_{H_{T+1}} && \text{(Hölder's ineq.)} \\ &\leq \frac{1}{T} \sum_{t \notin \mathcal{T}^w} \|\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t))\|_{H_t^{-1}} \|\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{T+1}\|_{H_{T+1}} && (H_{T+1} \geq H_t) \\ &\leq \frac{\beta_{T+1}(\delta)}{T} \sum_{t \notin \mathcal{T}^w} \|\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t))\|_{H_t^{-1}}. \\ &\quad \text{(Corollary C.1, with prob. } 1 - \delta\text{)} \end{aligned}$$

Define the sigmoid function $\mu(w) = \frac{1}{1+e^{-w}}$. Then, we have

$$\begin{aligned}
& \sum_{t \notin \mathcal{T}^w} \|\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t))\|_{H_t^{-1}} \\
&= \sum_{t \notin \mathcal{T}^w} \frac{\dot{\mu}\left(\left(\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t))\right)^\top \boldsymbol{\theta}^*\right)}{\dot{\mu}\left(\left(\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t))\right)^\top \boldsymbol{\theta}^*\right)} \|\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t))\|_{H_t^{-1}} \\
&= (1+e)^2 \sum_{t \notin \mathcal{T}^w} \dot{\mu}\left(\left(\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t))\right)^\top \boldsymbol{\theta}^*\right) \|\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t))\|_{H_t^{-1}},
\end{aligned}$$

where the last equality holds due to the fact that

$$\begin{aligned}
& \frac{1}{\dot{\mu}\left(\left(\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t))\right)^\top \boldsymbol{\theta}^*\right)} \\
&= \frac{\left(1 + e^{\left(\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t))\right)^\top \boldsymbol{\theta}^*}\right)^2}{e^{\left(\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t))\right)^\top \boldsymbol{\theta}^*}} \\
&\leq \left(1 + e^{\left(\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t))\right)^\top \boldsymbol{\theta}^*}\right)^2 \quad ((\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t)))^\top \boldsymbol{\theta}^* \geq 0) \\
&\leq \left(1 + e^{\left(\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t))\right)^\top (\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{T+1})}\right)^2 \quad (\hat{\pi}_T(x_t) = \operatorname{argmax}_{a \in \mathcal{A}} \phi(x_t, a)^\top \hat{\boldsymbol{\theta}}_{T+1}) \\
&\leq \left(1 + e^{\|\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t))\|_{H_t^{-1}} \|\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{T+1}\|_{H_{T+1}}}\right)^2 \quad (\text{Hölder's ineq. and } H_{T+1} \geq H_t) \\
&\leq \left(1 + e^{\frac{\beta_{T+1}(\delta)}{3\sqrt{2K}\beta_{T+1}(\delta)}}\right)^2 \quad (t \notin \mathcal{T}^w \text{ and Corollary C.1}) \\
&= (1+e^{1/K})^2 \leq (1+e)^2.
\end{aligned}$$

Recall that $S_\sigma^{(j)} := \{\sigma_j, \dots, \sigma_{|S|}\}$ denotes the subset of remaining actions in S after removing the first $j-1$ actions, given S and σ . We denote the Multinomial Logit (MNL) model at round t by $P_t(a | S; \boldsymbol{\theta})$ (see Equation (9)). Then, by our assortment selection rule in Equation (11), we obtain

$$\begin{aligned}
& \dot{\mu}\left(\left(\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t))\right)^\top \boldsymbol{\theta}^*\right) \|\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t))\|_{H_t^{-1}} \\
&\leq \dot{\mu}\left(\left(\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t))\right)^\top \hat{\boldsymbol{\theta}}_t\right) e^{|\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t))^\top (\hat{\boldsymbol{\theta}}_t - \boldsymbol{\theta}^*)|} \\
&\quad \cdot \|\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t))\|_{H_t^{-1}} \quad (\text{Lemma E.2}) \\
&\leq e \cdot \dot{\mu}\left(\left(\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t))\right)^\top \hat{\boldsymbol{\theta}}_t\right) \|\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t))\|_{H_t^{-1}} \\
&\quad (\text{Eqn. C.5}) \\
&= \frac{e}{2} \cdot \mathbb{E}_{a \sim P_t(\cdot | S_t^*; \hat{\boldsymbol{\theta}}_t)} \left[\left\| \phi(x_t, a) - \mathbb{E}_{a' \sim P_t(\cdot | S_t^*; \hat{\boldsymbol{\theta}}_t)} [\phi(x_t, a')] \right\|_{H_t^{-1}} \right] \\
&\quad (\dot{\mu}(z) = \mu(z)(1 - \mu(z)) \text{ and let } S_t^* := \{\pi^*(x_t), \hat{\pi}_T(x_t)\}) \\
&= \frac{e}{|S_t^*|} \mathbb{E}_{\sigma \sim \mathbb{P}_t(\cdot | S_t^*; \hat{\boldsymbol{\theta}}_t)} \sum_{j=1}^{|S_t^*|} \mathbb{E}_{a \sim P_t(\cdot | S_t^{*,(j)}; \hat{\boldsymbol{\theta}}_t)} \left[\left\| \phi(x_t, a) - \mathbb{E}_{a' \sim P_t(\cdot | S_t^{*,(j)}; \hat{\boldsymbol{\theta}}_t)} [\phi(x_t, a')] \right\|_{H_t^{-1}} \right], \\
&\quad (|S_t^*| = 2)
\end{aligned}$$

where the second inequality holds since, for any $t \notin \mathcal{T}^w$ and $a, a' \in \mathcal{A}$, the following property holds:

$$\begin{aligned}
& \left| (\phi(x_t, a) - \phi(x_t, a'))^\top (\hat{\boldsymbol{\theta}}_t - \boldsymbol{\theta}^*) \right| \\
& \leq \|\phi(x_t, a) - \phi(x_t, a')\|_{H_t^{-1}} \|\hat{\boldsymbol{\theta}}_t - \boldsymbol{\theta}^*\|_{H_t} && \text{(Hölder's inequality)} \\
& \leq \frac{1}{3\sqrt{2}K\beta_{T+1}(\delta)} \|\hat{\boldsymbol{\theta}}_t - \boldsymbol{\theta}^*\|_{H_t} && (t \neq \mathcal{T}^w) \\
& \leq \frac{\beta_t(\delta)}{K\beta_{T+1}(\delta)} && \text{(Corollary C.1)} \\
& \leq \frac{1}{K}. && \text{(C.5)}
\end{aligned}$$

Combining the above results, we get

$$\begin{aligned}
& \sum_{t \notin \mathcal{T}^w} \|\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t))\|_{H_t^{-1}} \\
& \leq e(1+e)^2 \sum_{t \notin \mathcal{T}^w} \frac{1}{|S_t^*|} \sum_{j=1}^{|S_t^*|} \mathbb{E}_{\mathbb{P}_t(\cdot|S_t^*; \hat{\boldsymbol{\theta}}_t) / P_t(\cdot|S_t^{*,(j)}; \hat{\boldsymbol{\theta}}_t)} \left[\|\phi(x_t, a) - \mathbb{E}_{P_t(\cdot|S_t^{*,(j)}; \hat{\boldsymbol{\theta}}_t)} [\phi(x_t, a')]\|_{H_t^{-1}} \right] && \text{(C.6)} \\
& \leq e(1+e)^2 \sqrt{\sum_{t=1}^T \frac{1}{|S_t|}} \sqrt{\sum_{t \notin \mathcal{T}^w} \frac{|S_t|}{|S_t^*|} \sum_{j=1}^{|S_t^*|} \mathbb{E}_{\mathbb{P}_t(\cdot|S_t^*; \hat{\boldsymbol{\theta}}_t) / P_t(\cdot|S_t^{*,(j)}; \hat{\boldsymbol{\theta}}_t)} \left[\|\phi(x_t, a) - \mathbb{E}_{P_t(\cdot|S_t^{*,(j)}; \hat{\boldsymbol{\theta}}_t)} [\phi(x_t, a')]\|_{H_t^{-1}}^2 \right]} && \text{(Cauchy-Schwartz ineq., and } \mathcal{T}^w \subseteq [T]) \\
& \leq e(1+e)^2 \sqrt{\sum_{t=1}^T \frac{1}{|S_t|}} \sqrt{\sum_{t \notin \mathcal{T}^w} \frac{|S_t|}{|\mathcal{S}_t|} \sum_{j=1}^{|\mathcal{S}_t|} \mathbb{E}_{\mathbb{P}_t(\cdot|\mathcal{S}_t; \hat{\boldsymbol{\theta}}_t) / P_t(\cdot|\mathcal{S}_t^{(j)}; \hat{\boldsymbol{\theta}}_t)} \left[\|\phi(x_t, a) - \mathbb{E}_{P_t(\cdot|\mathcal{S}_t^{(j)}; \hat{\boldsymbol{\theta}}_t)} [\phi(x_t, a')]\|_{H_t^{-1}}^2 \right]} && \text{(Assortment selection rule, Equation (11))} \\
& \leq e^2(1+e)^2 \sqrt{\sum_{t=1}^T \frac{1}{|S_t|}} \sqrt{\sum_{t \notin \mathcal{T}^w} \sum_{j=1}^{|S_t|} \mathbb{E}_{\mathbb{P}_t(\cdot|S_t; \boldsymbol{\theta}^*) / P_t(\cdot|S_t^{(j)}; \hat{\boldsymbol{\theta}}_t)} \left[\|\phi(x_t, a) - \mathbb{E}_{P_t(\cdot|S_t^{(j)}; \hat{\boldsymbol{\theta}}_t)} [\phi(x_t, a')]\|_{H_t^{-1}}^2 \right]} && \text{(Eqn. (C.7))} \\
& \leq \sqrt{3}e^3(1+e)^2 \sqrt{\sum_{t=1}^T \frac{1}{|S_t|}} \sqrt{\sum_{t \notin \mathcal{T}^w} \sum_{j=1}^{|S_t|} \mathbb{E}_{\mathbb{P}_t(\cdot|S_t; \boldsymbol{\theta}^*) / P_t(\cdot|S_t^{(j)}; \hat{\boldsymbol{\theta}}_t)} \left[\|\phi(x_t, a) - \mathbb{E}_{P_t(\cdot|S_t^{(j)}; \hat{\boldsymbol{\theta}}_t)} [\phi(x_t, a')]\|_{\mathcal{M}_t^{-1}}^2 \right]}, && \text{(Lemma C.2, with prob. } 1 - \delta)
\end{aligned}$$

where the second-to-last inequality holds because, for any $t \notin \mathcal{T}^w$, $S \in \mathcal{S}$, and $a \in S$, we have

$$\begin{aligned}
P_t(a|S; \hat{\boldsymbol{\theta}}_t) &= \frac{\exp(\phi(x_t, a)^\top \hat{\boldsymbol{\theta}}_t)}{\sum_{a' \in S} \exp(\phi(x_t, a')^\top \hat{\boldsymbol{\theta}}_t)} \\
&= \frac{\exp((\phi(x_t, a) - \phi(x_t, \bar{a}))^\top \hat{\boldsymbol{\theta}}_t)}{1 + \sum_{a' \in S \setminus \{\bar{a}\}} \exp((\phi(x_t, a') - \phi(x_t, \bar{a}))^\top \hat{\boldsymbol{\theta}}_t)} && \text{(any } \bar{a} \in S) \\
&\leq \frac{\exp((\phi(x_t, a) - \phi(x_t, \bar{a}))^\top \boldsymbol{\theta}^* + 1/K)}{1 + \sum_{a' \in S \setminus \{\bar{a}\}} \exp((\phi(x_t, a') - \phi(x_t, \bar{a}))^\top \boldsymbol{\theta}^* - 1/K)} && \text{(Eqn. (C.5))} \\
&= e^{2/K} P_t(a|S; \boldsymbol{\theta}^*),
\end{aligned}$$

which implies that, for the ranking $\sigma = (\sigma_1, \dots, \sigma_{|S|})$,

$$\begin{aligned}\mathbb{P}_t(\sigma|S; \hat{\theta}_t) &= \prod_{j=1}^{|S|} P_t(\sigma_j|S_\sigma^{(j)}; \hat{\theta}_t) \leq e^{2|S|/K} \prod_{j=1}^{|S|} P_t(\sigma_j|S_\sigma^{(j)}; \theta^\star) \leq e^2 \prod_{j=1}^{|S|} P_t(\sigma_j|S_\sigma^{(j)}; \theta^\star) \\ &= e^2 \mathbb{P}_t(\sigma|S; \theta^\star).\end{aligned}\quad (\text{C.7})$$

Moreover, by setting $\lambda = d \log T$ and $T \geq e^{K/d}$, for any t , we have

$$\sum_{j=1}^{|S_t|} \mathbb{E}_{\mathbb{P}_t(\cdot|S_t; \theta^\star)} \left[\left\| \phi(x_t, a) - \mathbb{E}_{P_t(\cdot|S_t^{(j)}; \hat{\theta}_t)} [\phi(x_t, a')] \right\|_{M_t^{-1}}^2 \right] \leq \frac{2K}{\lambda} \leq 2. \quad (\text{C.8})$$

Hence, by the elliptical potential for expected mean-centered uncertainty (Lemma C.3), we obtain

$$\begin{aligned}&\frac{1}{T} \sum_{t \notin \mathcal{T}^w} (\phi(x_t, \pi^\star(x_t)) - \phi(x_t, \hat{\pi}_T(x_t)))^\top (\theta^\star - \hat{\theta}_{T+1}) \\ &\leq \frac{\beta_{T+1}(\delta)}{T} \sum_{t \notin \mathcal{T}^w} \|\phi(x_t, \pi^\star(x_t)) - \phi(x_t, \hat{\pi}_T(x_t))\|_{H_t^{-1}} \\ &\leq \mathcal{O}\left(\frac{\beta_{T+1}(\delta)}{T} \sqrt{\sum_{t=1}^T \frac{1}{|S_t|}} \sqrt{\sum_{t \notin \mathcal{T}^w} \sum_{j=1}^{|S_t|} \mathbb{E}_{\mathbb{P}_t(\cdot|S_t; \theta^\star)} \left[\left\| \phi(x_t, a) - \mathbb{E}_{P_t(\cdot|S_t^{(j)}; \hat{\theta}_t)} [\phi(x_t, a')] \right\|_{M_t^{-1}}^2 \right]} \right) \\ &= \mathcal{O}\left(\frac{\beta_{T+1}(\delta)}{T} \sqrt{\sum_{t=1}^T \frac{1}{|S_t|}} \sqrt{d \log\left(1 + \frac{KT}{d\lambda}\right)}\right) \quad (\text{Lemma C.3 with Eqn. (C.8)}) \\ &= \mathcal{O}\left(\frac{\beta_{T+1}(\delta)}{T} \sqrt{\sum_{t=1}^T \frac{1}{|S_t|}} \cdot \sqrt{d \log(KT)}\right).\end{aligned}\quad (\text{C.9})$$

By combining Equations (C.3), (C.4), and (C.9), and setting $\beta_{T+1}(\delta) = \mathcal{O}(B\sqrt{d \log(KT)} + B\sqrt{\lambda})$, we derive that, with probability at least $1 - 3\delta$ (omitting logarithmic terms and polynomial dependencies on B for brevity),

$$\mathbf{SubOpt}(T) = \tilde{\mathcal{O}}\left(\frac{d}{T} \sqrt{\sum_{t=1}^T \frac{1}{|S_t|}} + \frac{d^2 K^4}{\kappa T}\right).$$

Substituting $\delta \leftarrow \frac{\delta}{3}$, we conclude the proof of Theorem 1. \square

C.2 Proofs of Lemmas for Theorem 1

C.2.1 Proof of Lemma C.2

Proof of Lemma C.2. Recall the definition of H_t .

$$H_t = \sum_{s=1}^{t-1} \sum_{j=1}^{|S_s|} \nabla^2 \ell_s^{(j)}(\hat{\theta}_s^{(j+1)}) + \lambda \mathbf{I}_d \geq \sum_{s \in [t-1] \setminus \mathcal{T}^w} \sum_{j=1}^{|S_s|} \nabla^2 \ell_s^{(j)}(\hat{\theta}_s^{(j+1)}) + \lambda \mathbf{I}_d$$

Here, we can equivalently express the MNL loss at step j and round s , denoted by $\nabla^2 \ell_s^{(j)}(\hat{\theta}_s^{(j+1)})$, as follows:

$$\begin{aligned}\ell_s^{(j)}(\hat{\theta}_s^{(j+1)}) &= -\log \left(\frac{\exp(\phi(x_s, \sigma_{sj})^\top \hat{\theta}_s^{(j+1)})}{\sum_{k=j}^{|S_s|} \exp(\phi(x_s, \sigma_{sk})^\top \hat{\theta}_s^{(j+1)})} \right) = -\log \left(\frac{\exp(a_{sj})}{\sum_{k=j}^{|S_s|} \exp(a_{sk})} \right) \\ &=: \tilde{\ell}_s^{(j)}(\mathbf{a}_s^{(j)}),\end{aligned}\quad (\text{C.10})$$

where $a_{sj} = \phi(x_s, \sigma_{sj})^\top \hat{\boldsymbol{\theta}}_s^{(j+1)}$, $\mathbf{a}_s^{(j)} = (a_{sk})_{k=j}^{|S_s|} \in \mathbb{R}^{|S_s|-j+1}$. Define the matrix

$$\boldsymbol{\Phi}_s^{(j)} = \begin{pmatrix} \phi(x_s, \sigma_{sj})^\top \\ \vdots \\ \phi(x_s, \sigma_{s|S_s|})^\top \end{pmatrix} \in \mathbb{R}^{(|S_s|-j+1) \times d},$$

where each row corresponds to the feature vector of an action ranked from position j to $|S_s|$ in the ranking σ_s . Moreover, we define $\hat{a}_{sj} = \phi(x_s, \sigma_{sj})^\top \hat{\boldsymbol{\theta}}_s$ and $\hat{\mathbf{a}}_s^{(j)} = (\hat{a}_{sk})_{k=j}^{|S_s|} \in \mathbb{R}^{|S_s|-j+1}$

Then, using the ℓ_∞ -norm self-concordant property of the MNL loss [41], for any $s \in [t-1] \setminus \mathcal{T}^w$, we obtain

$$\begin{aligned} \nabla^2 \ell_s^{(j)}(\hat{\boldsymbol{\theta}}_s^{(j+1)}) &= \left(\boldsymbol{\Phi}_s^{(j)} \right)^\top \nabla_{\mathbf{a}}^2 \bar{\ell}_s^{(j)}(\mathbf{a}_s^{(j)}) \boldsymbol{\Phi}_s^{(j)} && \text{(Eqn. (C.10))} \\ &\geq e^{-3\sqrt{2}\|\mathbf{a}_s^{(j)} - \hat{\mathbf{a}}_s^{(j)}\|_\infty} \left(\boldsymbol{\Phi}_s^{(j)} \right)^\top \nabla_{\mathbf{a}}^2 \bar{\ell}_s^{(j)}(\hat{\mathbf{a}}_s^{(j)}) \boldsymbol{\Phi}_s^{(j)} && \text{(Lemma E.1)} \\ &\geq \frac{1}{e^2} \left(\boldsymbol{\Phi}_s^{(j)} \right)^\top \nabla_{\mathbf{a}}^2 \bar{\ell}_s^{(j)}(\hat{\mathbf{a}}_s^{(j)}) \boldsymbol{\Phi}_s^{(j)} && (\|\mathbf{a}_s^{(j)} - \hat{\mathbf{a}}_s^{(j)}\|_\infty \leq \frac{2}{3\sqrt{2}}) \\ &= \frac{1}{e^2} \nabla^2 \ell_s^{(j)}(\hat{\boldsymbol{\theta}}_s), && \text{(Eqn. (C.10))} \end{aligned}$$

where the last inequality holds because, for any $s \in [t-1] \setminus \mathcal{T}^w$ and $j \leq |S_s|$, the following holds:

$$\begin{aligned} \|\mathbf{a}_s^{(j)} - \hat{\mathbf{a}}_s^{(j)}\|_\infty &\leq \max_{k=j, \dots, |S_s|} \left| \phi(x_k, \sigma_{sk})^\top (\hat{\boldsymbol{\theta}}_s^{(k+1)} - \hat{\boldsymbol{\theta}}_s) \right| \\ &\leq \max_{k=j, \dots, |S_s|} \|\phi(x_k, \sigma_{sk})\|_{H_s^{-1}} \left(\left\| \hat{\boldsymbol{\theta}}_s^{(k+1)} - \boldsymbol{\theta}^* \right\|_{H_s} + \left\| \hat{\boldsymbol{\theta}}_s - \boldsymbol{\theta}^* \right\|_{H_s} \right) && \text{(Hölder's inequality)} \\ &\leq \frac{1}{3\sqrt{2}\beta_{T+1}(\delta)} \max_{k=j, \dots, |S_s|} \left(\left\| \hat{\boldsymbol{\theta}}_s^{(k+1)} - \boldsymbol{\theta}^* \right\|_{H_s^{(k+1)}} + \left\| \hat{\boldsymbol{\theta}}_s - \boldsymbol{\theta}^* \right\|_{H_s} \right) && (s \notin \mathcal{T}^w, H_s \leq H_s^{(k+1)}) \\ &\leq \frac{2\beta_{T+1}(\delta)}{3\sqrt{2}\beta_{T+1}(\delta)} && \text{(Corollary C.1, } \beta_t(\delta) \text{ is non-decreasing)} \\ &= \frac{2}{3\sqrt{2}}. \end{aligned}$$

Therefore, we get

$$H_t \geq \sum_{s \in [t-1] \setminus \mathcal{T}^w} \sum_{j=1}^{|S_s|} \nabla^2 \ell_s^{(j)}(\hat{\boldsymbol{\theta}}_s^{(j+1)}) + \lambda \mathbf{I}_d \geq \frac{1}{e^2} \sum_{s \in [t-1] \setminus \mathcal{T}^w} \sum_{j=1}^{|S_s|} \nabla^2 \ell_s^{(j)}(\hat{\boldsymbol{\theta}}_s) + \lambda \mathbf{I}_d.$$

Recall the definition of the MNL choice probability (9):

$$P_s(a|S; \boldsymbol{\theta}) := \frac{\exp(\phi(x_s, a)^\top \boldsymbol{\theta})}{\sum_{a' \in S} \exp(\phi(x_s, a')^\top \boldsymbol{\theta})}, \quad \forall a \in S.$$

Given $S_s = \{\sigma_{s1}, \dots, \sigma_{s|S_s|}\}$ and the ranking feedback σ_s , we define $S_{\sigma_s}^{(j)} := \{\sigma_{sj}, \dots, \sigma_{s|S_s|}\}$ as the subset of remaining actions in S_s after removing the first $j-1$ actions. For simplicity, let

$\hat{P}_s^{(j)} = P_s(\cdot | S_{\sigma_s}^{(j)}; \hat{\theta}_t)$. Then, we can express the Hessian of the loss as follows:

$$\begin{aligned}
\sum_{j=1}^{|S_s|} \nabla^2 \ell_s^{(j)}(\hat{\theta}_s) &= \sum_{j=1}^{|S_s|} \sum_{k=j}^{|S_s|} \sum_{k'=j}^{|S_s|} \frac{\exp\left((\phi(x_s, \sigma_{sk}) + \phi(x_s, \sigma_{sk'}))^{\top} \hat{\theta}_s\right)}{2 \left(\sum_{k'=j}^{|S_s|} \exp\left(\phi(x_s, \sigma_{sk'})^{\top} \hat{\theta}_s\right)\right)^2} \cdot z_{skk'} z_{skk'}^{\top} \\
&= \frac{1}{2} \sum_{j=1}^{|S_s|} \sum_{k=j}^{|S_s|} \sum_{k'=j}^{|S_s|} P_s(\sigma_{sk} | S_s^{(j)}; \hat{\theta}_s) P_s(\sigma_{sk'} | S_s^{(j)}; \hat{\theta}_s) z_{skk'} z_{skk'}^{\top} \\
&= \frac{1}{2} \sum_{j=1}^{|S_s|} \sum_{k=j}^{|S_s|} \sum_{k'=j}^{|S_s|} \mathbb{E}_{(a, a') \sim \hat{P}_s^{(j)} \times \hat{P}_s^{(j)}} \left[(\phi(x_s, a) - \phi(x_s, a')) (\phi(x_s, a) - \phi(x_s, a'))^{\top} \right] \\
&= \sum_{j=1}^{|S_s|} \mathbb{E}_{\hat{P}_s^{(j)}} \left[(\phi(x_s, a) - \mathbb{E}_{\hat{P}_s^{(j)}} [\phi(x_s, a')]) (\phi(x_s, a) - \mathbb{E}_{\hat{P}_s^{(j)}} [\phi(x_s, a')])^{\top} \right] \geq 0,
\end{aligned} \tag{C.11}$$

where $z_{skk'} = \phi(x_s, \sigma_{sk}) - \phi(x_s, \sigma_{sk'})$. Therefore, $\sum_{j=1}^{|S_s|} \nabla^2 \ell_s^{(j)}(\hat{\theta}_s)$ is a random PSD matrix conditional on the assortment S_s , where the randomness arises from the realized ranking σ_s . Note that the true PL distribution $\mathbb{P}_s(\cdot | S_s; \theta^*)$ is measurable with respect to the filtration $\mathcal{F}_{s-1} = \sigma(x_1, S_1, \sigma_1, \dots, x_s, S_s)$. Moreover, the maximum eigenvalue of the Hessian satisfies $\lambda_{\max} \left(\sum_{j=1}^{|S_s|} \nabla^2 \ell_s^{(j)}(\hat{\theta}_s) \right) \leq K$. Then, by the PSD matrices concentration lemma (Lemma E.3), with probability $1 - \delta$, we have

$$\begin{aligned}
H_t &\geq \frac{1}{e^2} \sum_{s \in [t-1] \setminus \mathcal{T}^w} \sum_{j=1}^{|S_s|} \nabla^2 \ell_s^{(j)}(\hat{\theta}_s) + \lambda \mathbf{I}_d \\
&= \frac{K}{e^2} \sum_{s \in [t-1] \setminus \mathcal{T}^w} \frac{1}{K} \sum_{j=1}^{|S_s|} \nabla^2 \ell_s^{(j)}(\hat{\theta}_s) + \lambda \mathbf{I}_d \\
&\geq \frac{K}{3e^2} \left(\sum_{s \in [t-1] \setminus \mathcal{T}^w} \mathbb{E}_{\sigma \sim \mathbb{P}_s(\cdot | S_s, \theta^*)} \left[\frac{1}{K} \sum_{j=1}^{|S_s|} \nabla^2 \ell_{s,\sigma}^{(j)}(\hat{\theta}_s) \right] + \lambda \mathbf{I}_d \right) \\
&\quad \text{(Lemma E.3, with prob. } 1 - \delta \text{)} \\
&= \frac{1}{3e^2} \left(\sum_{s \in [t-1] \setminus \mathcal{T}^w} \sum_{j=1}^{|S_s|} \mathbb{E}_{\sigma \sim \mathbb{P}_s(\cdot | S_s, \theta^*)} \left[\nabla^2 \ell_{s,\sigma}^{(j)}(\hat{\theta}_s) \right] + \lambda \mathbf{I}_d \right) = \frac{1}{3e^2} M_t.
\end{aligned}$$

This concludes the proof of Lemma C.2. \square

C.2.2 Proof of Lemma C.3

Proof of Lemma C.3. By the definition of M_t (C.2) and Equation (C.11), for $t \notin \mathcal{T}^w$, we have

$$\begin{aligned}
\det(M_{t+1}) &= \det \left(M_t + \sum_{j=1}^{|S_t|} \mathbb{E}_{\sigma \sim \mathbb{P}_t(\cdot | S_t, \theta^*)} \left[\nabla^2 \ell_{t,\sigma}^{(j)}(\hat{\theta}_t) \right] \right) \\
&= \det \left(M_t + \sum_{j=1}^{|S_t|} \mathbb{E}_{\sigma \sim \mathbb{P}_t(\cdot | S_t, \theta^*)} \left[\bar{\phi}_t^{(j)}(a) \bar{\phi}_t^{(j)}(a)^{\top} \right] \right) \\
&= \det(M_t) \left(1 + \sum_{j=1}^{|S_t|} \mathbb{E}_{\sigma \sim \mathbb{P}_t(\cdot | S_t, \theta^*)} \left[\|\bar{\phi}_t^{(j)}(a)\|_{M_t^{-1}}^2 \right] \right) \\
&\geq \det(\lambda \mathbf{I}_d) \prod_{s \in [t] \setminus \mathcal{T}^w} \left(1 + \min \left\{ 1, \sum_{j=1}^{|S_s|} \mathbb{E}_{\sigma \sim \mathbb{P}_s(\cdot | S_s, \theta^*)} \left[\|\bar{\phi}_s^{(j)}(a)\|_{M_s^{-1}}^2 \right] \right\} \right),
\end{aligned} \tag{Eqn. (C.11)}$$

where, for simplicity, we denote $\bar{\phi}_t^{(j)}(a) = \phi(x_t, a) - \mathbb{E}_{a' \sim P_t(\cdot | S_{t,\sigma}^{(j)}; \hat{\theta}_t)}[\phi(x_t, a')]$.

For $t \in \mathcal{T}^w$, it is clear that $\det(M_{t+1}) = \det(M_t)$.

Hence, using the fact that $a \leq 2 \log(1 + a)$ for any $a \in [0, 1]$, we get

$$\begin{aligned} & \sum_{t \notin \mathcal{T}^w} \min \left\{ 1, \sum_{j=1}^{|S_t|} \mathbb{E}_{\sigma \sim \mathbb{P}_t(\cdot | S_t, \theta^*)} \left[\|\bar{\phi}_t^{(j)}(a)\|_{M_t^{-1}}^2 \right] \right\} \\ & \leq 2 \sum_{t \notin \mathcal{T}^w} \log \left(1 + \min \left\{ 1, \sum_{j=1}^{|S_t|} \mathbb{E}_{\sigma \sim \mathbb{P}_t(\cdot | S_t, \theta^*)} \left[\|\bar{\phi}_t^{(j)}(a)\|_{M_t^{-1}}^2 \right] \right\} \right) \\ & \leq 2 \log \left(\frac{\det(M_{T+1})}{\det(\lambda \mathbf{I}_d)} \right) \\ & \leq 2d \log \left(1 + \frac{4KT}{d\lambda} \right), \end{aligned}$$

where the last inequality holds because

$$\begin{aligned} \det(M_{T+1}) & \leq \left(\frac{\lambda_1 + \dots + \lambda_d}{d} \right)^d \quad (\lambda_1, \dots, \lambda_d \text{ are eigenvalues of } M_{T+1}, \text{ AM-GM ineq.}) \\ & \leq \left(\frac{\text{trace}(M_{T+1})}{d} \right)^d \\ & = \left(\frac{\lambda d + \sum_{t=1}^T \sum_{j=1}^{|S_t|} \mathbb{E} \|\bar{\phi}_t^{(j)}(a)\|_2^2}{d} \right)^d \leq \left(\lambda + \frac{KT}{d} \right)^d. \quad (\mathbb{E} \|\bar{\phi}_t^{(j)}(a)\|_2^2 \leq 1) \end{aligned}$$

This concludes the proof of Lemma C.3. \square

C.2.3 Proof of Lemma C.4

Proof of Lemma C.4. Let

$$(\tilde{a}_t, \bar{a}_t) = \underset{a, a' \in \mathcal{A}}{\text{argmax}} \dot{\mu} \left((\phi(x_t, a) - \phi(x_t, a'))^\top \hat{\theta}_t \right) \|\phi(x_t, a) - \phi(x_t, a')\|_{H_t^{-1}}^2,$$

where $\mu(w) = \frac{1}{1+e^{-w}}$ denote the sigmoid function. Note that, by our assortment selection rule in Equation (11), the two actions $a, a' \in \mathcal{A}$ that maximize $\dot{\mu} \left((\phi(x_t, a) - \phi(x_t, a'))^\top \hat{\theta}_t \right) \|\phi(x_t, a) - \phi(x_t, a')\|_{H_t^{-1}}^2$ are always included in S_t . It is because

$$\begin{aligned} & \dot{\mu} \left((\phi(x_t, a) - \phi(x_t, a'))^\top \hat{\theta}_t \right) \|\phi(x_t, a) - \phi(x_t, a')\|_{H_t^{-1}}^2 \\ & = 2\mathbb{E}_{a \sim P_t(\cdot | \{a, a'\}; \hat{\theta}_t)} \left[\left\| \phi(x_t, a) - \mathbb{E}_{a' \sim P_t(\cdot | \{a, a'\}; \hat{\theta}_t)} [\phi(x_t, a')] \right\|_{H_t^{-1}}^2 \right] \\ & = 2\mathbb{E}_{\substack{\sigma \sim \mathbb{P}_t(\cdot | \{a, a'\}; \hat{\theta}_t) \\ a \sim P_t(\cdot | \{a, a'\}; \hat{\theta}_t)}} \left[\left\| \phi(x_t, a) - \mathbb{E}_{a' \sim P_t(\cdot | \{a, a'\}; \hat{\theta}_t)} [\phi(x_t, a')] \right\|_{H_t^{-1}}^2 \right] = 2f_t(\{a, a'\}), \end{aligned}$$

which directly implies that $\tilde{a}_t, \bar{a}_t \in S_t$.

For simplicity, we denote $z_{skk'} = \phi(x_s, \sigma_{sk}) - \phi(x_s, \sigma_{sk'})$. Then, by the definition of H_t , we have

$$\begin{aligned}
H_t &= \sum_{s=1}^{t-1} \sum_{j=1}^{|S_s|} \nabla^2 \ell_s^{(j)}(\hat{\boldsymbol{\theta}}_s^{(j+1)}) + \lambda \mathbf{I}_d \\
&= \sum_{s=1}^{t-1} \sum_{j=1}^{|S_s|} \sum_{k=j}^{|S_s|} \sum_{k'=j}^{|S_s|} \frac{\exp\left((\phi(x_s, \sigma_{sk}) + \phi(x_s, \sigma_{sk'}))^{\top} \hat{\boldsymbol{\theta}}_s^{(j+1)}\right)}{2 \left(\sum_{k'=j}^{|S_s|} \exp\left(\phi(x_s, \sigma_{sk'})^{\top} \hat{\boldsymbol{\theta}}_s^{(j+1)}\right)\right)^2} \cdot z_{skk'} z_{skk'}^{\top} + \lambda \mathbf{I}_d \\
&\geq \frac{e^{-4B}}{2K^2} \sum_{s=1}^{t-1} \sum_{j=1}^{|S_s|} \sum_{k=j}^{|S_s|} \sum_{k'=j}^{|S_s|} z_{skk'} z_{skk'}^{\top} + \lambda \mathbf{I}_d \\
&\geq \frac{e^{-4B}}{2K^2} \sum_{s=1}^{t-1} (\phi(x_s, \tilde{a}_s) - \phi(x_s, \bar{a}_s)) (\phi(x_s, \tilde{a}_s) - \phi(x_s, \bar{a}_s))^{\top} + \lambda \mathbf{I}_d \quad (\tilde{a}_s, \bar{a}_s \in S_s) \\
&\geq \frac{e^{-4B}}{2K^2} \left(\sum_{s \in \mathcal{T}_t^w} \dot{\mu} \left((\phi(x_s, \tilde{a}_s) - \phi(x_s, \bar{a}_s))^{\top} \hat{\boldsymbol{\theta}}_s \right) (\phi(x_s, \tilde{a}_s) - \phi(x_s, \bar{a}_s)) (\phi(x_s, \tilde{a}_s) - \phi(x_s, \bar{a}_s))^{\top} \right. \\
&\quad \left. + \lambda \mathbf{I}_d \right), \tag{C.12}
\end{aligned}$$

where the last inequality holds since $\dot{\mu} \left((\phi(x_s, a) - \phi(x_s, a'))^{\top} \hat{\boldsymbol{\theta}}_s \right) \leq \frac{1}{4}$. In addition, we denote $\mathcal{T}_t^w = \mathcal{T}^w \setminus \{t, \dots, T\}$ in the last inequality. To better presentation, we define

$$\begin{aligned}
\Lambda_t^w &:= \sum_{s \in \mathcal{T}_t^w} \dot{\mu} \left((\phi(x_s, \tilde{a}_s) - \phi(x_s, \bar{a}_s))^{\top} \hat{\boldsymbol{\theta}}_s \right) (\phi(x_s, \tilde{a}_s) - \phi(x_s, \bar{a}_s)) (\phi(x_s, \tilde{a}_s) - \phi(x_s, \bar{a}_s))^{\top} \\
&\quad + \lambda \mathbf{I}_d.
\end{aligned}$$

On the other hand, by setting $\lambda \geq 1$, for any t , we can ensure that

$$\begin{aligned}
\dot{\mu} \left((\phi(x_t, \tilde{a}_t) - \phi(x_t, \bar{a}_t))^{\top} \hat{\boldsymbol{\theta}}_t \right) \|\phi(x_t, \tilde{a}_t) - \phi(x_t, \bar{a}_t)\|_{(\Lambda_t^w)^{-1}}^2 &\leq \frac{1}{4} \cdot \frac{2}{\lambda} \quad (\dot{\mu}(\cdot) \leq \frac{1}{4}) \\
&\leq 1. \tag{C.13}
\end{aligned}$$

Then, we can derive that

$$\begin{aligned}
&\sum_{t \in \mathcal{T}^w} \max_{a, a' \in \mathcal{A}} \|\phi(x_t, a) - \phi(x_t, a')\|_{H_t^{-1}}^2 \\
&= \sum_{t \in \mathcal{T}^w} \max_{a, a' \in \mathcal{A}} \frac{\dot{\mu} \left((\phi(x_t, a) - \phi(x_t, a'))^{\top} \hat{\boldsymbol{\theta}}_t \right)}{\dot{\mu} \left((\phi(x_t, a) - \phi(x_t, a'))^{\top} \hat{\boldsymbol{\theta}}_t \right)} \|\phi(x_t, a) - \phi(x_t, a')\|_{H_t^{-1}}^2 \\
&\leq 4e^{2B} \sum_{t \in \mathcal{T}^w} \max_{a, a' \in \mathcal{A}} \dot{\mu} \left((\phi(x_t, a) - \phi(x_t, a'))^{\top} \hat{\boldsymbol{\theta}}_t \right) \|\phi(x_t, a) - \phi(x_t, a')\|_{H_t^{-1}}^2 \quad (\text{Assumption 1}) \\
&= 4e^{2B} \sum_{t \in \mathcal{T}^w} \dot{\mu} \left((\phi(x_t, \tilde{a}_t) - \phi(x_t, \bar{a}_t))^{\top} \hat{\boldsymbol{\theta}}_t \right) \|\phi(x_t, \tilde{a}_t) - \phi(x_t, \bar{a}_t)\|_{H_t^{-1}}^2 \\
&\leq 8K^2 e^{6B} \sum_{t \in \mathcal{T}^w} \dot{\mu} \left((\phi(x_t, \tilde{a}_t) - \phi(x_t, \bar{a}_t))^{\top} \hat{\boldsymbol{\theta}}_t \right) \|\phi(x_t, \tilde{a}_t) - \phi(x_t, \bar{a}_t)\|_{(\Lambda_t^w)^{-1}}^2 \quad (\text{Eqn. (C.12)}) \\
&= 8K^2 e^{6B} \sum_{t \in \mathcal{T}^w} \min \left\{ 1, \dot{\mu} \left((\phi(x_t, \tilde{a}_t) - \phi(x_t, \bar{a}_t))^{\top} \hat{\boldsymbol{\theta}}_t \right) \|\phi(x_t, \tilde{a}_t) - \phi(x_t, \bar{a}_t)\|_{(\Lambda_t^w)^{-1}}^2 \right\} \quad (\text{Eqn. (C.13)}) \\
&\leq 16K^2 e^{6B} d \log \left(1 + \frac{2T}{d\lambda} \right). \quad (\text{Lemma E.6})
\end{aligned}$$

On the other hand, for $t \in \mathcal{T}^w$, we know that

$$\sum_{t \in \mathcal{T}^w} \max_{a, a' \in \mathcal{A}} \|\phi(x_t, a) - \phi(x_t, a')\|_{H_t^{-1}}^2 \geq \frac{|\mathcal{T}^w|}{18K^2\beta_{T+1}(\delta)^2}.$$

By combining the two results above, and setting $\kappa = e^{-6}$, we obtain

$$|\mathcal{T}^w| \leq \frac{288K^4}{\kappa} \beta_{T+1}(\delta)^2 d \log \left(1 + \frac{2KT}{d\lambda} \right),$$

which concludes the proof. \square

D Proof of Theorem 2

D.1 Main Proof of Theorem 2

In this section, we present the proof of Theorem 2, which is obtained by using the RB loss (4) instead of the PL loss (3). Note that this approach is based on the concept of *rank-breaking* (RB), which decomposes (partial) ranking data into individual pairwise comparisons, treats each comparison as independent, and has been extensively studied in previous works [6, 34, 32, 67]. Moreover, this RB approach is applied in the current RLHF for LLM (e.g., Ouyang et al. [57]) and is also studied theoretically in Zhu et al. [93] under the offline setting.

RB loss and OMD. We begin by recalling the loss function and the parameter update rule. Specifically, we use the PL loss defined in Equation (4) and update the parameter according to Equation (7).

$$\ell_t(\boldsymbol{\theta}) := \sum_{j=1}^{|S_t|-1} \underbrace{\sum_{k=j+1}^{|S_t|} -\log \left(\frac{\exp(\phi(x_t, \sigma_{tj})^\top \boldsymbol{\theta})}{\exp(\phi(x_t, \sigma_{tj})^\top \boldsymbol{\theta}) + \exp(\phi(x_t, \sigma_{tk})^\top \boldsymbol{\theta})} \right)}_{=: \ell_t^{(j,k)}(\boldsymbol{\theta})} = \sum_{j=1}^{|S_t|-1} \sum_{k=j+1}^{|S_t|} \ell_t^{(j,k)}(\boldsymbol{\theta}).$$

and

$$\hat{\boldsymbol{\theta}}_t^{(j,k+1)} = \operatorname{argmin}_{\boldsymbol{\theta} \in \Theta} \langle \nabla \ell_t^{(j,k)}(\hat{\boldsymbol{\theta}}_t^{(j,k)}), \boldsymbol{\theta} \rangle + \frac{1}{2\eta} \|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_t^{(j,k)}\|_{\tilde{H}_t^{(j,k)}}^2, \quad 1 \leq j < k \leq |S_t|,$$

where if $k = |S_t|$, we set $\hat{\boldsymbol{\theta}}_t^{(j,k+1)} = \hat{\boldsymbol{\theta}}_t^{(j+1,j+2)}$, and for the final pair, let $\hat{\boldsymbol{\theta}}_t^{(|S_t|-1, |S_t|+1)} = \hat{\boldsymbol{\theta}}_{t+1}^{(1,2)}$. Also, the matrix $\tilde{H}_t^{(j,k)}$ is defined as $\tilde{H}_t^{(j,k)} := H_t + \eta \sum_{(j',k') \leq (j,k)} \nabla^2 \ell_t^{(j',k')}(\hat{\boldsymbol{\theta}}_t^{(j',k')})$ ⁹, where

$$\begin{aligned} H_t &:= \sum_{s=1}^{t-1} \sum_{j=1}^{|S_s|-1} \sum_{k=j+1}^{|S_s|} \nabla^2 \ell_s^{(j,k)}(\hat{\boldsymbol{\theta}}_s^{(j,k+1)}) + \lambda \mathbf{I}_d \\ &= \sum_{s=1}^{t-1} \sum_{j=1}^{|S_s|-1} \sum_{k=j+1}^{|S_s|} \dot{\mu} \left((\phi(x_t, \sigma_j) - \phi(x_t, \sigma_k))^\top \hat{\boldsymbol{\theta}}_s^{(j,k+1)} \right) \\ &\quad \cdot (\phi(x_t, \sigma_j) - \phi(x_t, \sigma_k)) (\phi(x_t, \sigma_j) - \phi(x_t, \sigma_k))^\top + \lambda \mathbf{I}_d. \end{aligned}$$

Online confidence bound for RB loss. Now, we introduce the online confidence bound for RB loss. Since the total number of updates up to round t is $\sum_{s=1}^t \binom{|S_s|}{2}$, a modification of Lemma C.1 yields the following result:

Proposition D.1 (Online confidence bound for RB loss). *Let $\delta \in (0, 1]$. We set $\eta = (1 + 3\sqrt{2}B)/2$ and $\lambda = \max\{12\sqrt{2}B\eta, 144\eta d, 2\}$. If $K^2 \leq (2\sqrt{6} + 12\sqrt{3}B)d$, then under Assumption 1, with probability at least $1 - \delta$, we have*

$$\|\hat{\boldsymbol{\theta}}_t^{(j,k)} - \boldsymbol{\theta}^*\|_{H_t^{(j,k)}} \leq \beta_t(\delta) = \mathcal{O} \left(B \sqrt{d \log(tK/\delta)} + B\sqrt{\lambda} \right), \quad \forall t \geq 1, 1 \leq j < k \leq |S_t|,$$

where $H_t^{(j,k)} := H_t + \sum_{(j',k') \leq (j,k)} \nabla^2 \ell_t^{(j',k')}(\hat{\boldsymbol{\theta}}_t^{(j',k'+1)}) + \lambda \mathbf{I}_d$ and $\hat{\boldsymbol{\theta}}_t^{(1,2)} = \hat{\boldsymbol{\theta}}_t$.

⁹We write $(j', k') \leq (j, k)$ to indicate lexicographic order, i.e., $j' < j$ or $j' = j$ and $k' \leq k$.

The proof is deferred to Appendix D.2.1.

Useful definitions. We define the set of *warm-up rounds* \mathcal{T}^w in a similar manner to the analysis of Theorem 1:

$$\mathcal{T}^w := \left\{ t \in [T] : \max_{a, a' \in \mathcal{A}} \|\phi(x_t, a) - \phi(x_t, a')\|_{H_t^{-1}} \geq \frac{1}{\beta_{T+1}(\delta)} \right\}. \quad (\text{warm-up rounds})$$

Moreover, we define the matrix

$$\begin{aligned} \Lambda_t = & \sum_{s \in [t-1] \setminus \mathcal{T}^w} \sum_{a, a' \in S_s} \mu \left((\phi(x_t, a) - \phi(x_t, a'))^\top \hat{\theta}_s \right) (\phi(x_t, a) - \phi(x_t, a')) (\phi(x_t, a) - \phi(x_t, a'))^\top \\ & + \lambda \mathbf{I}_d \end{aligned} \quad (\text{D.1})$$

Key lemmas. The relationship between H_t and Λ_t is as follows:

Lemma D.1. *Let Λ_t be defined as in Equation (D.1). Then, for all $t \in [T]$, we have*

$$H_t \geq \frac{1}{2e^2} \Lambda_t.$$

The proof is deferred to Appendix D.2.2.

The size of the set \mathcal{T}^w is bounded as described in the following lemma:

Lemma D.2. *Let $\mathcal{T}^w = \{t \in [T] : \max_{a, a' \in \mathcal{A}} \|\phi(x_t, a) - \phi(x_t, a')\|_{H_t^{-1}} \geq \frac{1}{\beta_{T+1}(\delta)}\}$. Define $\kappa := e^{-6B}$. Set $\lambda \geq 1$. Then, the size of the set \mathcal{T}^w is bounded as follows:*

$$|\mathcal{T}^w| \leq \frac{32}{\kappa} \beta_{T+1}(\delta)^2 d \log \left(1 + \frac{2T}{d\lambda} \right).$$

The proof is deferred to Appendix D.2.3. Note that the bound on the cardinality of \mathcal{T}^w is tighter than in the PL loss case (Lemma C.4), by up to a factor of K^4 (and up to logarithmic factors in K).

We are now ready to provide the proof of Theorem 2.

Proof of Theorem 2. The overall proof structure is similar to that of Theorem 1. We start from Equation (C.4), but apply Lemma D.2 instead of Lemma C.4. Then, with probability at least $1 - \delta$, we have

$$\begin{aligned} \mathbf{SubOpt}(T) &= \mathbb{E}_{x \sim \rho} \left[(\phi(x, \pi^*(x)) - \phi(x, \hat{\pi}_T(x)))^\top \theta^* \right] \\ &\leq \tilde{\mathcal{O}} \left(\frac{1}{\sqrt{T}} \right) + \mathcal{O} \left(\frac{B}{\kappa T} \beta_{T+1}(\delta)^2 d \log \left(1 + \frac{T}{d\lambda} \right) \right) \quad (\text{Lemma D.2}) \\ &+ \frac{1}{T} \sum_{t \notin \mathcal{T}^w} (\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t)))^\top (\theta^* - \hat{\theta}_{T+1}). \end{aligned} \quad (\text{D.2})$$

To further bound the last term of Equation (D.2), by following the similar reasoning from Equation (C.4) to Equation (C.6), with probability at least $1 - \delta$, we obtain

$$\begin{aligned}
& \sum_{t \notin \mathcal{T}^w} \|\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t))\|_{H_t^{-1}} \\
& \leq e(1+e)^2 \sum_{t \notin \mathcal{T}^w} \frac{1}{|S_t^*|} \sum_{a, a' \in S_t^*} \dot{\mu}(\psi_t(a, a')^\top \hat{\theta}_t) \|\psi_t(a, a')\|_{H_t^{-1}} \\
& \quad (S_t^* := \{\pi^*(x_t), \hat{\pi}_T(x_t)\}, \text{ and let } \psi_t(a, a') = \phi(x_t, a) - \phi(x_t, a')) \\
& \leq e(1+e)^2 \sqrt{\sum_{t=1}^T \frac{1}{|S_t|}} \sqrt{\sum_{t \notin \mathcal{T}^w} \frac{|S_t|}{|S_t^*|} \sum_{a, a' \in S_t^*} \dot{\mu}(\psi_t(a, a')^\top \hat{\theta}_t) \|\psi_t(a, a')\|_{H_t^{-1}}^2} \\
& \quad (\text{Cauchy-Schwartz ineq., and } \mathcal{T}^w \subseteq [T]) \\
& \leq e(1+e)^2 \sqrt{\sum_{t=1}^T \frac{1}{|S_t|}} \sqrt{\sum_{t \notin \mathcal{T}^w} \frac{|S_t|}{|S_t^*|} \sum_{a, a' \in S_t^*} \dot{\mu}(\psi_t(a, a')^\top \hat{\theta}_t) \|\psi_t(a, a')\|_{H_t^{-1}}^2} \\
& \leq \sqrt{2}e^2(1+e)^2 \sqrt{\sum_{t=1}^T \frac{1}{|S_t|}} \sqrt{\sum_{t \notin \mathcal{T}^w} \sum_{a, a' \in S_t} \dot{\mu}(\psi_t(a, a')^\top \hat{\theta}_t) \|\psi_t(a, a')\|_{\Lambda_t^{-1}}^2}. \quad (\text{Lemma D.1})
\end{aligned}$$

Moreover, by setting $\lambda = d \log T$ and $T \geq e^{K^2/d}$, for any t , we have

$$\sum_{a, a' \in S_t} \dot{\mu}(\psi_t(a, a')^\top \hat{\theta}_t) \|\psi_t(a, a')\|_{\Lambda_t^{-1}}^2 \leq \frac{2K^2}{\lambda} \leq 2. \quad (\text{D.3})$$

Hence, by the elliptical potential lemma (Lemma E.6), we derive that

$$\begin{aligned}
& \frac{1}{T} \sum_{t \notin \mathcal{T}^w} (\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t)))^\top (\theta^* - \hat{\theta}_{T+1}) \\
& \leq \frac{\beta_{T+1}(\delta)}{T} \sum_{t \notin \mathcal{T}^w} \|\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t))\|_{H_t^{-1}} \\
& = \mathcal{O}\left(\frac{\beta_{T+1}(\delta)}{T} \sqrt{\sum_{t=1}^T \frac{1}{|S_t|}} \cdot \sqrt{d \log(KT)}\right). \quad (\text{D.4})
\end{aligned}$$

By plugging Equation (D.4) into Equation (D.2) and setting $\beta_{T+1}(\delta) = \mathcal{O}(B\sqrt{d \log(KT)} + B\sqrt{\lambda})$, then, with probability at least $1 - 2\delta$, we derive that

$$\mathbf{SubOpt}(T) = \tilde{\mathcal{O}}\left(\frac{d}{T} \sqrt{\sum_{t=1}^T \frac{1}{|S_t|}} + \frac{d^2}{\kappa T}\right).$$

Substituting $\delta \leftarrow \frac{\delta}{2}$, we conclude the proof of Theorem 2. \square

D.2 Proofs of Propositions and Lemmas for Theorem 2

D.2.1 Proof of Proposition D.1

Proof of Proposition D.1. Since the rank-breaking pairs generated from the same set S_t are not mutually independent, Lemma C.1 cannot be applied directly. Therefore, instead of Lemma C.2 of Lee and Oh [41], we use the following lemma.

Lemma D.3. *Let $\delta \in (0, 1]$. We define a pseudo-inverse function of $\mu_t^{(j,k)}(\cdot)$ as $\mu_t^{(j,k),+} : \mathbb{R} \rightarrow \mathbb{R}$, where $[\mu_t^{(j,k),+}(\mathbf{q})]_i = \log(q_i / (1 - \|\mathbf{q}\|_1))$ for any $\mathbf{q} \in \{\mathbf{p} \in [0, 1]^d \mid \|\mathbf{p}\|_1 < 1\}$. Let $\psi_t(a, a') = \phi(x_t, a) - \phi(x_t, a')$, and the intermediary parameter denote $\tilde{z}_t^{(j,k)} :=$*

$\mu_t^{(j,k),+} \left(\mathbb{E}_{\boldsymbol{\theta} \sim P_t^{(j,k)}} \left[\mu_t^{(j,k)} (\psi_t(\sigma_j, \sigma_k)^\top \boldsymbol{\theta}) \right] \right)$, where $P_t^{(j,k)} := \mathcal{N} \left(\hat{\boldsymbol{\theta}}_t^{(j,k)}, cH_t^{-1} \right)$ is a multivariate normal distribution with mean $\hat{\boldsymbol{\theta}}_t^{(j,k)}$ and covariance cH_t^{-1} ($c > 0$). Then, for any $t \in [T]$, with probability at least $1 - \delta$, we have

$$\sum_{s=1}^t \sum_{j < k \leq |S_s|} \left(\ell_s^{(j,k)}(\boldsymbol{\theta}^*) - \ell_s^{(j,k)}(\tilde{z}_s^{(j,k)}) \right) \leq \frac{K(K-1)}{2} \log \frac{1}{\delta},$$

where $\bar{\ell}_t(z) = -\log \left(\frac{\exp(z)}{1+\exp(z)} \right)$.

Proof of Lemma D.3. The proof of Lemma D.3 mostly follows the analysis of Lemma C.2 in Lee and Oh [41]. However, unlike their setting, the rank-breaking pairs generated from the same assortment S_t are generally dependent. Therefore, instead of applying Ville's inequality [77] to individual pairs, we apply it at the *block level*, where each block corresponds to one round t , by aggregating pairwise likelihood ratios through a mixture construction. This preserves the supermartingale property and allows us to obtain the desired bound.

For simplicity, we write

$$\mu_t^{(j,k)} := \mu_t^{(j,k)} (\psi_t(\sigma_j, \sigma_k)^\top \boldsymbol{\theta}^*), \quad \text{and} \quad \tilde{p}_t^{(j,k)} := \mathbb{E}_{\boldsymbol{\theta} \sim P_t^{(j,k)}} \left[\mu_t^{(j,k)} (\psi_t(\sigma_j, \sigma_k)^\top \boldsymbol{\theta}) \right].$$

We define

$$M_t := \sum_{j < k \leq |S_t|} \frac{1}{\binom{|S_t|}{2}} \exp \left(\ell_t^{(j,k)}(\boldsymbol{\theta}^*) - \ell_t^{(j,k)}(\tilde{z}_s^{(j,k)}) \right) = \sum_{j < k \leq |S_t|} \frac{1}{\binom{|S_t|}{2}} \cdot \frac{\tilde{p}_t^{(j,k)}(Y_t^{(j,k)})}{p_t^{(j,k)}(Y_t^{(j,k)})},$$

where $p_t^{(j,k)}(Y_t^{(j,k)}) = (p_t^{(j,k)})^{Y_t^{(j,k)}} (1 - p_t^{(j,k)})^{1 - Y_t^{(j,k)}}$. Then, we have

$$\mathbb{E}[M_t | \mathcal{F}_{t-1}] = \frac{1}{\binom{|S_t|}{2}} \sum_{j < k \leq |S_t|} \mathbb{E} \left[\frac{\tilde{p}_t^{(j,k)}(Y_t^{(j,k)})}{p_t^{(j,k)}(Y_t^{(j,k)})} \mid \mathcal{F}_{t-1} \right] = \frac{1}{\binom{|S_t|}{2}} \sum_{j < k \leq |S_t|} \sum_{t \in \{0,1\}} \tilde{p}_t^{(j,k)}(y) = 1.$$

Let $A_t := \prod_{s=1}^t M_s$ and $A_0 := 1$. Then, it follows that $(A_t)_{t \geq 1}$ is a nonnegative supermartingale with respect to the filtration $\mathcal{F}_{t-1} = \sigma(x_1, S_1, \sigma_1, \dots, x_t, S_t)$.

Hence, following the analysis of Lemma C.2 in Lee and Oh [41] and applying Ville's inequality, we obtain that

$$\mathbb{P} \left(\sup_t A_t \geq \frac{1}{\delta} \right) \leq \mathbb{E}[A_0] \delta = \delta, \quad (\text{Ville's inequality})$$

which is equivalent to

$$\mathbb{P} \left(\exists t : \sum_{s=1}^t \log M_s \geq \log \frac{1}{\delta} \right) \leq \delta.$$

Therefore, with probability at least $1 - \delta$, we have

$$\begin{aligned} \sum_{s=1}^t \sum_{j < k \leq |S_s|} \left(\ell_s^{(j,k)}(\boldsymbol{\theta}^*) - \bar{\ell}_s^{(j,k)}(\tilde{z}_s^{(j,k)}) \right) &\leq \frac{K(K-1)}{2} \sum_{s=1}^t \frac{1}{\binom{|S_s|}{2}} \sum_{j < k \leq |S_s|} \left(\ell_s^{(j,k)}(\boldsymbol{\theta}^*) - \bar{\ell}_s^{(j,k)}(\tilde{z}_s^{(j,k)}) \right) \\ &\leq \frac{K(K-1)}{2} \sum_{s=1}^t \log M_s \quad (\text{Jensen's inequality}) \\ &\leq \frac{K(K-1)}{2} \log \frac{1}{\delta}. \end{aligned}$$

This concludes the proof of Lemma D.3. \square

Now we follow the proof of Theorem 4.2 in Lee and Oh [41], with substituting Lemma C.2 of theirs with our Lemma D.3, then, we obtain that

$$\begin{aligned} \|\hat{\boldsymbol{\theta}}_t^{(j,k)} - \boldsymbol{\theta}^*\|_{H_t^{(j,k)}}^2 &\leq \frac{1+3\sqrt{2}B}{2} K(K-1) \log \frac{1}{\delta} + 4\sqrt{6} \left(\frac{1+3\sqrt{2}B}{2} \right)^2 d \log(tK^2 + 2) + 4B^2 \lambda \\ &= \mathcal{O}(B^2 d \log(tK/\delta) + B^2 \lambda). \quad (\text{Let } K^2 \leq (2\sqrt{6} + 12\sqrt{3}B)d) \end{aligned}$$

This completes the proof of Proposition D.1. \square

D.2.2 Proof of Lemma D.1

Proof of Lemma D.1. Recall that, under the Bradley–Terry–Luce (BTL) model defined in Equation (2), the probability that action a is preferred over action a' is given by:

$$\mathbb{P}(a > a'|x_t, ; \boldsymbol{\theta}) = \frac{\exp(\phi(x_t, a)^\top \boldsymbol{\theta})}{\exp(\phi(x_t, a)^\top \boldsymbol{\theta}) + \exp(\phi(x_t, a')^\top \boldsymbol{\theta})} = \mu((\phi(x_t, a) - \phi(x_t, a'))^\top \boldsymbol{\theta}).$$

Then, we can derive a lower bound on the matrix H_t as follows:

$$\begin{aligned} H_t &= \sum_{s=1}^{t-1} \sum_{j=1}^{|S_s|-1} \sum_{k=j+1}^{|S_s|} \nabla^2 \ell_s^{(j,k)}(\widehat{\boldsymbol{\theta}}_s^{(j,k+1)}) + \lambda \mathbf{I}_d \\ &\geq \sum_{s \in [t-1] \setminus \mathcal{T}^w} \sum_{j=1}^{|S_s|-1} \sum_{k=j+1}^{|S_s|} \nabla^2 \ell_s^{(j,k)}(\widehat{\boldsymbol{\theta}}_s^{(j,k+1)}) + \lambda \mathbf{I}_d \\ &= \sum_{s \in [t-1] \setminus \mathcal{T}^w} \sum_{j=1}^{|S_s|-1} \sum_{k=j+1}^{|S_s|} \dot{\mu}\left(z_{sjk}^\top \widehat{\boldsymbol{\theta}}_s^{(j,k+1)}\right) z_{sjk} z_{sjk}^\top + \lambda \mathbf{I}_d \\ &\geq \sum_{s \in [t-1] \setminus \mathcal{T}^w} \sum_{j=1}^{|S_s|-1} \sum_{k=j+1}^{|S_s|} \dot{\mu}\left(z_{sjk}^\top \widehat{\boldsymbol{\theta}}_s\right) e^{-|z_{sjk}^\top (\widehat{\boldsymbol{\theta}}_s^{(j,k+1)} - \widehat{\boldsymbol{\theta}}_s)|} z_{sjk} z_{sjk}^\top + \lambda \mathbf{I}_d \quad (\text{Lemma E.2}) \\ &\geq e^{-2} \sum_{s \in [t-1] \setminus \mathcal{T}^w} \sum_{j=1}^{|S_s|-1} \sum_{k=j+1}^{|S_s|} \dot{\mu}\left(z_{sjk}^\top \widehat{\boldsymbol{\theta}}_s\right) z_{sjk} z_{sjk}^\top + \lambda \mathbf{I}_d \\ &\geq \frac{e^{-2}}{2} \Lambda_t, \end{aligned}$$

where the second-to-last inequality holds because, for any $s \notin \mathcal{T}^w$, the following property is satisfied:

$$\begin{aligned} |z_{sjk}^\top (\widehat{\boldsymbol{\theta}}_s^{(j,k+1)} - \widehat{\boldsymbol{\theta}}_s)| &= |(\phi(x_s, \sigma_{sj}) - \phi(x_s, \sigma_{sk}))^\top (\widehat{\boldsymbol{\theta}}_s^{(j,k+1)} - \widehat{\boldsymbol{\theta}}_s)| \\ &\leq \|\phi(x_s, \sigma_{sj}) - \phi(x_s, \sigma_{sk})\|_{H_s^{-1}} \left(\|\widehat{\boldsymbol{\theta}}_s^{(j,k+1)} - \boldsymbol{\theta}^*\|_{H_s} + \|\widehat{\boldsymbol{\theta}}_s - \boldsymbol{\theta}^*\|_{H_s} \right) \\ &\quad (\text{Hölder's inequality}) \\ &\leq \frac{1}{\beta_{T+1}(\delta)} \left(\|\widehat{\boldsymbol{\theta}}_s^{(j,k+1)} - \boldsymbol{\theta}^*\|_{H_s^{(j,k+1)}} + \|\widehat{\boldsymbol{\theta}}_s - \boldsymbol{\theta}^*\|_{H_s} \right) \\ &\quad (s \neq \mathcal{T}^w, H_s \leq H_s^{(j,k+1)}) \\ &\leq \frac{2\beta_{T+1}(\delta)}{\beta_{T+1}(\delta)} \quad (\text{Proposition D.1 and } \beta_t(\delta) \text{ is non-decreasing}) \\ &\leq 2. \end{aligned}$$

This concludes the proof of Lemma D.1. \square

D.2.3 Proof of Lemma D.2

Proof of Lemma D.2. The proof follows the same structure as the analysis of Lemma C.4, with the main difference that we use the Hessian corresponding to the RB loss. Let

$$(\tilde{a}_t, \bar{a}_t) = \underset{a, a' \in \mathcal{A}}{\text{argmax}} \dot{\mu}\left((\phi(x_t, a) - \phi(x_t, a'))^\top \widehat{\boldsymbol{\theta}}_t\right) \|\phi(x_t, a) - \phi(x_t, a')\|_{H_t^{-1}}^2.$$

By the assortment selection rule in (11), the two actions \tilde{a}_t, \bar{a}_t are always included in S_t . It is because, for any pair $\{a, a'\}$, we have

$$\dot{\mu}\left((\phi(x_t, a) - \phi(x_t, a'))^\top \widehat{\boldsymbol{\theta}}_t\right) \|\phi(x_t, a) - \phi(x_t, a')\|_{H_t^{-1}}^2 = 2f_t(\{a, a'\}),$$

which directly implies that $\tilde{a}_t, \bar{a}_t \in S_t$.

For simplicity, we denote $z_{skk'} = \phi(x_s, \sigma_{sk}) - \phi(x_s, \sigma_{sk'})$. Recall that by the definition of H_t ,

$$\begin{aligned}
H_t &= \sum_{s=1}^{t-1} \sum_{j=1}^{|S_s|-1} \sum_{k=j+1}^{|S_s|} \nabla^2 \ell_s^{(j,k)}(\hat{\boldsymbol{\theta}}_s^{(j,k+1)}) + \lambda \mathbf{I}_d \\
&\geq \frac{e^{-4B}}{4} \sum_{s=1}^{t-1} \sum_{j=1}^{|S_s|-1} \sum_{k=j+1}^{|S_s|} z_{sjk} z_{sjk}^\top + \lambda \mathbf{I}_d \\
&\geq \frac{e^{-4B}}{4} \sum_{s=1}^{t-1} (\phi(x_s, \tilde{a}_s) - \phi(x_s, \bar{a}_s)) (\phi(x_s, \tilde{a}_s) - \phi(x_s, \bar{a}_s))^\top + \lambda \mathbf{I}_d \quad (\tilde{a}_s, \bar{a}_s \in S_s) \\
&\geq \frac{e^{-4B}}{4} \left(\sum_{s \in \mathcal{T}_t^w} \dot{\mu} \left((\phi(x_s, \tilde{a}_s) - \phi(x_s, \bar{a}_s))^\top \hat{\boldsymbol{\theta}}_s \right) (\phi(x_s, \tilde{a}_s) - \phi(x_s, \bar{a}_s)) (\phi(x_s, \tilde{a}_s) - \phi(x_s, \bar{a}_s))^\top \right. \\
&\quad \left. + \lambda \mathbf{I}_d \right), \tag{D.5}
\end{aligned}$$

where in the last inequality, we use the fact that $\dot{\mu}(\cdot) \leq \frac{1}{4}$ and let $\mathcal{T}_t^w = \mathcal{T}^w \setminus \{t, \dots, T\}$. Define

$$\begin{aligned}
\Lambda_t^w &:= \sum_{s \in \mathcal{T}_t^w} \dot{\mu} \left((\phi(x_s, \tilde{a}_s) - \phi(x_s, \bar{a}_s))^\top \hat{\boldsymbol{\theta}}_s \right) (\phi(x_s, \tilde{a}_s) - \phi(x_s, \bar{a}_s)) (\phi(x_s, \tilde{a}_s) - \phi(x_s, \bar{a}_s))^\top \\
&\quad + \lambda \mathbf{I}_d.
\end{aligned}$$

Furthermore, by setting $\lambda \geq 1$, for any t , the following holds

$$\begin{aligned}
\dot{\mu} \left((\phi(x_t, \tilde{a}_t) - \phi(x_t, \bar{a}_t))^\top \hat{\boldsymbol{\theta}}_t \right) \|\phi(x_t, \tilde{a}_t) - \phi(x_t, \bar{a}_t)\|_{(\Lambda_t^w)^{-1}}^2 &\leq \frac{1}{4} \cdot \frac{2}{\lambda} \quad (\dot{\mu}(\cdot) \leq \frac{1}{4}) \\
&\leq 1. \tag{D.6}
\end{aligned}$$

Then, we obtain

$$\begin{aligned}
&\sum_{t \in \mathcal{T}^w} \max_{a, a' \in \mathcal{A}} \|\phi(x_t, a) - \phi(x_t, a')\|_{H_t^{-1}}^2 \\
&= \sum_{t \in \mathcal{T}^w} \max_{a, a' \in \mathcal{A}} \frac{\dot{\mu} \left((\phi(x_t, a) - \phi(x_t, a'))^\top \hat{\boldsymbol{\theta}}_t \right)}{\dot{\mu} \left((\phi(x_t, a) - \phi(x_t, a'))^\top \hat{\boldsymbol{\theta}}_t \right)} \|\phi(x_t, a) - \phi(x_t, a')\|_{H_t^{-1}}^2 \\
&\leq 4e^{2B} \sum_{t \in \mathcal{T}^w} \max_{a, a' \in \mathcal{A}} \dot{\mu} \left((\phi(x_t, a) - \phi(x_t, a'))^\top \hat{\boldsymbol{\theta}}_t \right) \|\phi(x_t, a) - \phi(x_t, a')\|_{H_t^{-1}}^2 \\
&\quad \text{(Assumption 1)} \\
&= 4e^{2B} \sum_{t \in \mathcal{T}^w} \dot{\mu} \left((\phi(x_t, \tilde{a}_t) - \phi(x_t, \bar{a}_t))^\top \hat{\boldsymbol{\theta}}_t \right) \|\phi(x_t, \tilde{a}_t) - \phi(x_t, \bar{a}_t)\|_{H_t^{-1}}^2 \\
&\leq 16e^{6B} \sum_{t \in \mathcal{T}^w} \dot{\mu} \left((\phi(x_t, \tilde{a}_t) - \phi(x_t, \bar{a}_t))^\top \hat{\boldsymbol{\theta}}_t \right) \|\phi(x_t, \tilde{a}_t) - \phi(x_t, \bar{a}_t)\|_{(\Lambda_t^w)^{-1}}^2 \quad \text{(Eqn. (D.5))} \\
&= 16e^{6B} \sum_{t \in \mathcal{T}^w} \min \left\{ 1, \dot{\mu} \left((\phi(x_t, \tilde{a}_t) - \phi(x_t, \bar{a}_t))^\top \hat{\boldsymbol{\theta}}_t \right) \|\phi(x_t, \tilde{a}_t) - \phi(x_t, \bar{a}_t)\|_{(\Lambda_t^w)^{-1}}^2 \right\} \\
&\quad \text{(Eqn. (D.6))} \\
&\leq 32e^{6B} d \log \left(1 + \frac{2T}{d\lambda} \right). \tag{Lemma E.6}
\end{aligned}$$

On the other hand, for $t \in \mathcal{T}^w$, we know that

$$\sum_{t \in \mathcal{T}^w} \max_{a, a' \in \mathcal{A}} \|\phi(x_t, a) - \phi(x_t, a')\|_{H_t^{-1}}^2 \geq \frac{|\mathcal{T}^w|}{\beta_{T+1}(\delta)^2}.$$

By combining the two results above, and setting $\kappa = e^{-6}$, we derive that

$$|\mathcal{T}^w| \leq \frac{32}{\kappa} \beta_{T+1}(\delta)^2 d \log \left(1 + \frac{2T}{d\lambda} \right),$$

which concludes the proof. \square

E Technical Lemmas

Lemma E.1 (Proposition B.5 of Lee and Oh 41). *The Hessian of the multinomial logistic loss $\bar{\ell} : \mathbb{R}^M \rightarrow \mathbb{R}$ satisfies that, for any $\mathbf{a}_1, \mathbf{a}_2 \in \mathbb{R}^M$, we have:*

$$e^{-3\sqrt{2}\|\mathbf{a}_1 - \mathbf{a}_2\|_\infty} \nabla^2 \bar{\ell}(\mathbf{a}_1) \preceq \nabla^2 \bar{\ell}(\mathbf{a}_2) \preceq e^{3\sqrt{2}\|\mathbf{a}_1 - \mathbf{a}_2\|_\infty} \nabla^2 \bar{\ell}(\mathbf{a}_1).$$

Lemma E.2 (Lemma 9 of Abeille et al. 2). *Let f be a strictly increasing function such that $|\dot{f}| \leq \ddot{f}$, and let \mathcal{Z} be any bounded interval of \mathbb{R} . Then, for all $z_1, z_2 \in \mathcal{Z}$, we have*

$$\dot{f}(z_2) \exp(-|z_2 - z_1|) \leq \dot{f}(z_1) \leq \dot{f}(z_2) \exp(|z_2 - z_1|).$$

We also provide a concentration lemma for positive semi-definite (PSD) random matrices.

Lemma E.3 (Concentration of PSD matrices). *Let μ_i denote the conditional distribution of a positive semi-definite $M \in \mathbb{R}^{d \times d}$ conditioned on the filtration \mathcal{F}_{i-1} . Assume $\lambda_{\max}(M) \leq 1$. Define $\bar{M} := \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{M \sim \mu_i} M$. If $\lambda = \Omega(d \log(n/\delta))$, then with probability at least $1 - \delta$, for any $n \geq 1$,*

$$\frac{1}{3} (n\bar{M} + \lambda \mathbf{I}_d) \leq \sum_{i=1}^n M_i + \lambda \mathbf{I}_d \leq \frac{5}{3} (n\bar{M} + \lambda \mathbf{I}_d).$$

Proof of Lemma E.3. The overall structure of the proof closely follows that of Lemma 39 in Zanette et al. [89]. For completeness, we provide the full proof below.

Fix $x \in \mathbb{R}^d$ such that $\|x\|_2 = 1$. Let $\bar{M}_i = \mathbb{E}_{M \sim \mu_i} M$ and $\bar{M} = \frac{1}{n} \sum_{i=1}^n \bar{M}_i$. Then, we have

$$\mathbb{E}_{M \sim \mu_i} x^\top M x = x^\top \mathbb{E}_{M \sim \mu_i} M x = x^\top \bar{M}_i x.$$

Since M is a positive semi-definite matrix, the random variable $x^\top M x$ is non-negative, and it satisfies $x^\top M x \leq \lambda_{\max}(M) \|x\|_2^2 \leq 1$. Thus, the conditional variance is at most $x^\top \bar{M}_i x$ because

$$\text{Var}_{M \sim \mu_i} (x^\top M x) \leq \mathbb{E}_{M \sim \mu_i} (x^\top M x)^2 \leq \mathbb{E}_{M \sim \mu_i} x^\top M x = x^\top \bar{M}_i x.$$

Applying Lemma E.4 with the filtration \mathcal{F}_i , we obtain that, with probability at least $1 - \delta$, there exists a universal constant c such that

$$\left| \frac{1}{n} \sum_{i=1}^n (x^\top M_i x - x^\top \bar{M}_i x) \right| = \left| \frac{1}{n} \sum_{i=1}^n x^\top M_i x - x^\top \bar{M} x \right| \leq c \left(\sqrt{\frac{2x^\top \bar{M} x}{n} \log(2/\delta)} + \frac{\log(2/\delta)}{3n} \right).$$

Now, we will show that if $\lambda = \Omega(\log(1/\delta))$, we can derive

$$c \left(\sqrt{\frac{2x^\top \bar{M} x}{n} \log(2/\delta)} + \frac{\log(2/\delta)}{3n} \right) \leq \frac{1}{2} \left(x^\top \bar{M} x + \frac{\lambda}{n} \right). \quad (\text{E.1})$$

Case 1. $x^\top \bar{M} x \leq \frac{\lambda}{n}$.

In this case, it is sufficient to satisfy for some constants c', c''

$$\begin{aligned} \sqrt{\frac{2 \log(2/\delta)}{n}} &\leq c' \sqrt{\frac{\lambda}{n}} \quad \longleftrightarrow \quad \Omega(\log(1/\delta)) \leq \lambda \\ \frac{\log(2/\delta)}{3n} &\leq c'' \left(\frac{\lambda}{n} \right) \quad \longleftrightarrow \quad \Omega(\log(1/\delta)) \leq \lambda. \end{aligned}$$

Case 2. $x^\top \bar{M} x > \frac{\lambda}{n}$.

In this case, it is sufficient to satisfy for some constants c''', c''''

$$\begin{aligned} \sqrt{\frac{2x^\top \bar{M} x}{n} \log(2/\delta)} &\leq c''' \left(\frac{\lambda}{n} \right) \quad \longleftrightarrow \quad \Omega(\log(1/\delta)) \leq \lambda \\ \frac{\log(2/\delta)}{3n} &\leq c'''' \left(\frac{\lambda}{n} \right) \quad \longleftrightarrow \quad \Omega(\log(1/\delta)) \leq \lambda. \end{aligned}$$

Therefore, Equation (E.1) is satisfied. Since $\|x\|_2 \leq 1$, this implies

$$\left| x^\top \left(\frac{1}{n} \sum_{i=1}^n M_i - \bar{M} \right) x \right| \leq \frac{1}{2} x^\top \left(\bar{M} + \frac{\lambda}{n} \mathbf{I}_d \right) x. \quad (\text{E.2})$$

We denote the boundary of the unit ball by $\partial\mathcal{B} = \{\|x\|_2 = 1\}$. Then, for any $x \in \partial\mathcal{B}$, we know there exists a x' in the ϵ -covering such that $\|x - x'\|_2 \leq \epsilon$. Let \mathcal{N}_ϵ be the ϵ -covering number of $\partial\mathcal{B}$. Then, by the covering number of Euclidean ball lemma (Lemma E.5), we get

$$\mathcal{N}_\epsilon \leq \left(\frac{3}{\epsilon} \right)^d. \quad (\text{E.3})$$

Taking a union bound over x' and the number of samples n , with probability at least $1 - n\mathcal{N}_\epsilon\delta$, we obtain

$$\begin{aligned} \left| x^\top \left(\frac{1}{n} \sum_{i=1}^n M_i - \bar{M} \right) x \right| &\leq \left| (x')^\top \left(\frac{1}{n} \sum_{i=1}^n M_i - \bar{M} \right) x' \right| + \left| (x - x')^\top \left(\frac{1}{n} \sum_{i=1}^n M_i - \bar{M} \right) x' \right| \\ &\quad + \left| (x')^\top \left(\frac{1}{n} \sum_{i=1}^n M_i - \bar{M} \right) (x - x') \right| \\ &\leq \left| (x')^\top \left(\frac{1}{n} \sum_{i=1}^n M_i - \bar{M} \right) x' \right| + 4\epsilon. \\ &\quad (\|x - x'\|_2 \leq \epsilon \text{ and } M_i, \|\bar{M}\|_2 \leq 1) \\ &\leq \frac{1}{2} (x')^\top \left(\bar{M} + \frac{\lambda}{n} \mathbf{I}_d \right) x' + 4\epsilon \quad (\text{Eqn. (E.2)}) \\ &\leq \frac{1}{2} x^\top \left(\bar{M} + \frac{\lambda}{n} \mathbf{I}_d \right) x + \frac{9}{2}\epsilon \quad (\|x - x'\|_2 \leq \epsilon \text{ and } \|\bar{M}\|_2 \leq 1) \\ &\leq \frac{2}{3} x^\top \left(\bar{M} + \frac{\lambda}{n} \mathbf{I}_d \right) x, \quad (\text{set } \epsilon = \mathcal{O}(\frac{1}{n})) \end{aligned}$$

where $\lambda = \Omega(\log(\frac{2n\mathcal{N}_\epsilon}{\delta}))$. By substituting $\delta \leftarrow \delta/(n\mathcal{N}_\epsilon + 1)$ and combining this with Equation (E.3), we obtain:

$$\frac{1}{3} \left(\bar{M} + \frac{\lambda}{n} \mathbf{I}_d \right) \leq \frac{1}{n} \sum_{i=1}^n M_i + \frac{\lambda}{n} \mathbf{I}_d \leq \frac{5}{3} \left(\bar{M} + \frac{\lambda}{n} \mathbf{I}_d \right),$$

which concludes the proof. \square

Lemma E.4 (Bernstein for martingales, Theorem 1 of Beygelzimer et al. 10 and Lemma 45 of Zanette et al. 89). *Consider the stochastic process $\{X_n\}$ adapted to the filtration $\{\mathcal{F}_n\}$. Assume $\mathbb{E}X_n = 0$ and $cX_n \leq 1$ for every n ; then for every constant $z \neq 0$ it holds that*

$$\Pr \left(\sum_{n=1}^N X_n \leq z \sum_{n=1}^N \mathbb{E}(X_n^2 | \mathcal{F}_n) + \frac{1}{z} \log \frac{1}{\delta} \right) \geq 1 - \delta.$$

By optimizing the bound as a function of z , we also have

$$\Pr \left(\sum_{n=1}^N X_n \leq c \sqrt{\sum_{n=1}^N \mathbb{E}(X_n^2 | \mathcal{F}_n) \log \frac{1}{\delta}} + \log \frac{1}{\delta} \right) \geq 1 - \delta.$$

Lemma E.5 (Covering number of Euclidean ball). *For any $\epsilon > 0$, the ϵ -covering number of the Euclidean ball in \mathbb{R}^d with radius $R > 0$ is upper bounded by $(1 + 2R/\epsilon)^d$.*

Lemma E.6 (Elliptical potential lemma Abbasi-Yadkori et al. 1). *Let $\{z_t\}_{t \geq 1}$ be a bounded sequence in \mathbb{R}^d satisfying $\max_{t \geq 1} \|z_t\|_2 \leq X$. For any $t \geq 1$, we define $\Lambda_t := \sum_{s=1}^{t-1} z_s z_s^\top + \lambda \mathbf{I}_d$ with $\lambda > 0$. Then, we have*

$$\sum_{t=1}^T \min \left\{ 1, \|z_t\|_{\Lambda_t^{-1}}^2 \right\} \leq 2d \log \left(1 + \frac{X^2 T}{d\lambda} \right).$$

F Proof of Theorem 3

F.1 Main Proof of Theorem 3

Throughout the proof, we consider the setting where the context space is a singleton, i.e., $\mathcal{X} = \{x\}$. As a result, the problem reduces to a context-free setting, and we focus solely on the action space \mathcal{A} . Note that this is equivalent to assuming that ρ is a Dirac distribution.

We first present the following theorem, which serves as the foundation for our analysis.

Theorem F.1 (Lower bound on adaptive PL model parameter estimation). *Let $\Phi = \mathcal{S}^{d-1}$ be the unit sphere in \mathbb{R}^d , and let $\Theta = \{-\mu, \mu\}^d$ for some $\mu \in (0, 1/\sqrt{d}]$. We consider a query model where, at each round $t = 1, \dots, T$, the learner selects a subset $S_t \subseteq \Phi$ of feature vectors, with cardinality satisfying $2 \leq |S_t| \leq K$, and then receives a ranking feedback σ_t drawn from the Plackett–Luce (PL) model defined as:*

$$\mathbb{P}(\sigma_t | S_t; \theta) = \prod_{j=1}^{|S_t|} \frac{\exp(\phi_{\sigma_{tj}}^\top \theta)}{\sum_{k=j}^{|S_t|} \exp(\phi_{\sigma_{tk}}^\top \theta)},$$

where $\sigma_t = (\sigma_{t1}, \dots, \sigma_{t|S_t|})$ is a permutation of the actions in S_t , $\phi_a \in \Phi$ denotes the feature vector associated with action $a \in \mathcal{A}$ in the selected subset at round t , and $\theta \in \Theta$. Then, we have

$$\inf_{\hat{\theta}, \pi} \max_{\theta \in \Theta} \mathbb{E}_{\theta} \left[\|\theta - \hat{\theta}\|_2^2 \right] \geq \frac{d\mu^2}{2} \left(1 - \sqrt{\frac{2K^2T\mu^2}{d}} \right),$$

where the infimum is over all measurable estimators $\hat{\theta}$ and measurable (but possibly adaptive) query rules π , and $\mathbb{E}_{\theta}[\cdot]$ denotes the expectation over the randomness in the observations and decision rules if θ is the true instance. In particular, if $T \geq \frac{d^2}{8K^2}$, by choosing $\mu = \sqrt{d/(8K^2T)}$, we obtain

$$\inf_{\hat{\theta}, \pi} \max_{\theta \in \Theta} \mathbb{E}_{\theta} \left[\|\theta - \hat{\theta}\|_2^2 \right] \geq \frac{d^2}{32K^2T}.$$

Proof of Theorem F.1. The analysis of this result closely follows the proof of Theorem 3 in Shamir [73]. The key distinction lies in the input structure: our setting involves a set of feature vectors, while theirs is restricted to a single feature vector.

To begin with, since the worst-case expected regret with respect to θ can be lower bounded by the average regret under the uniform prior over Θ , we have:

$$\begin{aligned} \max_{\theta \in \Theta} \mathbb{E}_{\theta} \left[\|\theta - \hat{\theta}\|_2^2 \right] &\geq \mathbb{E}_{\theta \sim \text{Unif}(\Theta)} \mathbb{E}_{\theta} \left[\|\theta - \hat{\theta}\|_2^2 \right] \\ &= \mathbb{E}_{\theta \sim \text{Unif}(\Theta)} \mathbb{E}_{\theta} \left[\sum_{i=1}^d (\theta_i - \hat{\theta}_i)^2 \right] \\ &\geq \mu^2 \cdot \mathbb{E}_{\theta \sim \text{Unif}(\Theta)} \mathbb{E}_{\theta} \left[\sum_{i=1}^d \mathbb{I}\{\theta_i \hat{\theta}_i < 0\} \right]. \end{aligned} \tag{F.1}$$

As in Shamir [73], we assume that the query strategy is deterministic conditioned on the past: that is, S_t is a deterministic function of the previous queries and observations, i.e., $S_1, \sigma_1, \dots, S_{t-1}, \sigma_{t-1}$. This assumption is made without loss of generality, since any randomized querying strategy can be viewed as a distribution over deterministic strategies. Therefore, a lower bound that holds uniformly for all deterministic strategies also applies to any randomized strategy. Then, we use the following lemma.

Lemma F.1 (Lemma 4 of Shamir 73). *Let θ be a random vector, none of whose coordinates is supported on 0, and let y_1, y_2, \dots, y_T be a sequence of queries obtained by a deterministic strategy returning a point $\hat{\theta}$ (that is, ψ_t is a deterministic function of $\psi_1, y_1, \dots, \psi_{t-1}, y_{t-1}$, and $\hat{\theta}$ is a deterministic function of y_1, \dots, y_T). Then, we have*

$$\mathbb{E}_{\theta \sim \text{Unif}(\Theta)} \mathbb{E}_{\theta} \left[\sum_{i=1}^d \mathbb{I}\{\theta_i \hat{\theta}_i < 0\} \right] \geq \frac{d}{2} \left(1 - \sqrt{\frac{1}{d} \sum_{i=1}^d \sum_{t=1}^T U_{ti}} \right),$$

where

$$U_{ti} := \sup_{\boldsymbol{\theta}_j, j \neq i} D_{KL}(P(y_t | \boldsymbol{\theta}_i > 0, \{\boldsymbol{\theta}_j\}_{j \neq i}, \{y_s\}_{s=1}^{t-1}) \| P(y_t | \boldsymbol{\theta}_i < 0, \{\boldsymbol{\theta}_j\}_{j \neq i}, \{y_s\}_{s=1}^{t-1})) .$$

In our setting, we interpret $y_t = \sigma_t$, and $\psi_t = \{\phi_a\}_{a \in S_t} \subseteq \Phi$. Then, we can write U_{ti} as follows:

$$U_{ti} = \sup_{\boldsymbol{\theta}_j, j \neq i} D_{KL}(\mathbb{P}(\sigma_t | S_t; \boldsymbol{\theta}_i > 0, \{\boldsymbol{\theta}_j\}_{j \neq i},) \| \mathbb{P}(\sigma_t | S_t; \boldsymbol{\theta}_i < 0, \{\boldsymbol{\theta}_j\}_{j \neq i},)) .$$

For simplicity, let $\mathbb{P}_{\boldsymbol{\theta}}(\sigma | S) = \mathbb{P}(\sigma | S; \boldsymbol{\theta})$. Then, we can upper bound U_{ti} using the following lemma.

Lemma F.2. *For any $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \mathbb{R}^d$, let $\mathbb{P}_{\boldsymbol{\theta}}(\cdot | S)$ denote the PL distribution over rankings induced by the action set S and parameter vector $\boldsymbol{\theta}$. Then, we have*

$$D_{KL}(\mathbb{P}_{\boldsymbol{\theta}}(\cdot | S) \| \mathbb{P}_{\boldsymbol{\theta}'}(\cdot | S)) \leq \frac{K}{2} \sum_{a \in S} (\phi_a^\top (\boldsymbol{\theta}' - \boldsymbol{\theta}))^2 .$$

The proof is deferred to Appendix F.2.1.

By applying Lemma F.2, we have

$$\begin{aligned} \sum_{i=1}^d U_{ti} &\leq \frac{K}{2} \sum_{i=1}^d \sum_{a \in S_t} (2\mu \cdot [\phi_a]_i)^2 = 2K\mu^2 \sum_{a \in S_t} \underbrace{\sum_{i=1}^d ([\phi_a]_i)^2}_{=1} \\ &= 2K\mu^2 \cdot |S_t| \quad (\phi_a \in \mathcal{S}^{d-1}) \\ &\leq 2K^2\mu^2. \quad (|S_t| \leq K) \end{aligned}$$

Hence, by Lemma F.1, we get

$$\begin{aligned} \mathbb{E}_{\boldsymbol{\theta} \sim \text{Unif}(\Theta)} \mathbb{E}_{\boldsymbol{\theta}} \left[\sum_{i=1}^d \mathbb{I}\{\boldsymbol{\theta}_i \hat{\boldsymbol{\theta}}_i < 0\} \right] &\geq \frac{d}{2} \left(1 - \sqrt{\frac{1}{d} \sum_{i=1}^d \sum_{t=1}^T U_{ti}} \right) \\ &\geq \frac{d}{2} \left(1 - \sqrt{\frac{2K^2 T \mu^2}{d}} \right). \end{aligned} \quad (\text{F.2})$$

Combining Equation (F.1) and (F.2), we prove the first inequality of Theorem F.1. The second inequality directly follows by choosing $\mu = \sqrt{d/(8K^2 T)}$. \square

We are now ready to present the proof of Theorem 3.

Proof of Theorem 3. The structure of our proof is similar to that of Theorem 2 in Wagenmaker et al. [78]. However, while they consider the linear bandit setting, we focus on the Plackett–Luce (PL) bandit setting.

We adopt the same instance construction as in Theorem F.1, where $\Phi = \mathcal{S}^{d-1}$ and $\Theta = \{-\mu, \mu\}^d$. Define $\phi^*(\boldsymbol{\theta}) = \text{argmax}_{a \in \mathcal{A}} \phi_a^\top \boldsymbol{\theta}$. Then, since $\phi^*(\boldsymbol{\theta}) \in \Phi$ and $\boldsymbol{\theta} \in \Theta$, it is clear that

$$\phi^*(\boldsymbol{\theta}) = \boldsymbol{\theta} / \|\boldsymbol{\theta}\|_2 = \boldsymbol{\theta} / (\sqrt{d}\mu), \quad \text{and } \phi^*(\boldsymbol{\theta})^\top \boldsymbol{\theta} = \sqrt{d}\mu. \quad (\text{F.3})$$

Fix the suboptimality gap $\epsilon > 0$. By definition, a policy $\pi \in \Delta_\Phi$ is said to be ϵ -optimal if it satisfies

$$\mathbb{E}_{\phi \sim \pi} [\phi^\top \boldsymbol{\theta}] = \left(\underbrace{\mathbb{E}_{\phi \sim \pi} [\phi]}_{=: \phi_\pi} \right)^\top \boldsymbol{\theta} \geq \phi^*(\boldsymbol{\theta})^\top \boldsymbol{\theta} - \epsilon = \sqrt{d}\mu - \epsilon. \quad (\text{F.4})$$

Moreover, by Jensen's inequality, we have

$$\|\phi_\pi\|_2^2 \leq \mathbb{E}_{\phi \sim \pi} [\|\phi\|_2^2] = 1.$$

Let $\Delta = \phi_\pi - \phi^*(\theta)$. Then, we get

$$\begin{aligned} 1 &\geq \|\phi_\pi\|_2^2 = \|\phi^*(\theta) + \Delta\|_2^2 = 1 + \|\Delta\|_2^2 + 2\phi^*(\theta)^\top \Delta \\ \iff \phi^*(\theta)^\top \Delta &\leq -\frac{1}{2}\|\Delta\|_2^2 \\ \iff \theta^* \Delta &\leq -\frac{\sqrt{d}\mu}{2}\|\Delta\|_2^2. \end{aligned} \tag{Eqn. (F.3)}$$

Hence, if a policy π is ϵ -optimal for a parameter θ , then the following bound holds:

$$\begin{aligned} -\epsilon &\leq -\frac{\sqrt{d}\mu}{2}\|\Delta\|_2^2. \\ \iff \|\Delta\|_2^2 &\leq \frac{2\epsilon}{\sqrt{d}\mu}, \quad \text{where } \theta = \sqrt{d}\mu(\phi_\pi - \Delta). \end{aligned} \tag{Eqn. (F.4)}$$

We now assume that we are given an ϵ -optimal policy $\hat{\pi}$. Define $\hat{\phi} := \phi_{\hat{\pi}}$ and the following estimator

$$\hat{\theta} = \begin{cases} \theta' & \text{if } \exists \theta' \in \Theta \text{ with } \theta' = \sqrt{d}\mu(\hat{\phi} - \Delta') \text{ for some } \Delta' \in \mathbb{R}^d, \|\Delta'\|_2^2 \leq \frac{2\epsilon}{\sqrt{d}\mu}; \\ \text{any } \theta' \in \Theta & \text{otherwise.} \end{cases}$$

If $\hat{\pi}$ is indeed ϵ -optimal for some $\theta \in \Theta$, then the first condition is satisfied, and we have:

$$\|\hat{\theta} - \theta\|_2 = \|\sqrt{d}\mu(\hat{\phi} - \Delta') - \sqrt{d}\mu(\hat{\phi} - \Delta)\|_2 \leq 2\sqrt{d}\mu\sqrt{\frac{2\epsilon}{\sqrt{d}\mu}} = \sqrt{8\sqrt{d}\mu\epsilon}. \tag{F.5}$$

We denote \mathcal{E} as the event that $\hat{\pi}$ is ϵ -optimal for $\theta \in \Theta$. Then, we get

$$\begin{aligned} \mathbb{E}_\theta [\|\hat{\theta} - \theta\|_2^2] &= \mathbb{E}_\theta [\|\hat{\theta} - \theta\|_2^2 \cdot \mathbb{I}\{\mathcal{E}\} + \|\hat{\theta} - \theta\|_2^2 \cdot \mathbb{I}\{\mathcal{E}^c\}] \\ &\leq 8\sqrt{d}\mu\epsilon + \mathbb{E}_\theta [\|\hat{\theta} - \theta\|_2^2 \cdot \mathbb{I}\{\mathcal{E}^c\}] \\ &\leq 8\sqrt{d}\mu\epsilon + 2d\mu^2 \cdot P_\theta[\mathcal{E}^c]. \end{aligned} \tag{Eqn. (F.5)}$$

On the other hand, by Theorem F.1, there exists a parameter $\theta \in \Theta$ such that, if we collect T samples and set $\mu = \sqrt{d}/(8K^2T)$, then the following lower bound holds:

$$\mathbb{E}_\theta [\|\hat{\theta} - \theta\|_2^2] \geq \frac{d^2}{32K^2T}.$$

To satisfy both inequalities, we require:

$$\begin{aligned} \frac{2\sqrt{2}d\epsilon}{\sqrt{K^2T}} + \frac{d^2}{4K^2T} \cdot P_\theta[\mathcal{E}^c] &\geq \frac{d^2}{32K^2T} \\ \iff P_\theta[\mathcal{E}^c] &\geq \frac{1}{8} - \frac{4\sqrt{2}K\sqrt{T}\epsilon}{d}. \end{aligned}$$

It follows that if

$$\frac{1}{8} - \frac{4\sqrt{2}K\sqrt{T}\epsilon}{d} \geq 0.1 \iff \frac{0.025^2}{32} \cdot \frac{d^2}{K^2\epsilon^2} \geq T,$$

then we have that $P_\theta[\mathcal{E}^c] \geq 0.1$. In words, this means that with constant probability, any algorithm must either collect more than $c \cdot \frac{d^2}{K^2\epsilon^2}$ samples, or output a policy that is not ϵ -optimal. This implies that $T = \Omega(\frac{d^2}{K^2\epsilon^2})$ samples are necessary to guarantee an ϵ -optimal policy. Equivalently, after T rounds, the suboptimality gap ϵ is lower bounded as

$$\mathbf{SubOpt}(T) = \Omega\left(\frac{d}{K\sqrt{T}}\right).$$

This concludes the proof of Theorem 3. \square

F.2 Proof of Lemmas for Theorem 3

F.2.1 Proof of Lemma F.2

Proof of Lemma F.2. By the definition of KL divergence, we have

$$D_{\text{KL}}(\mathbb{P}_{\boldsymbol{\theta}}(\cdot|S) \|\mathbb{P}_{\boldsymbol{\theta}'}(\cdot|S)) = \mathbb{E}_{\sigma \sim \mathbb{P}_{\boldsymbol{\theta}}(\cdot|S)} \left[\sum_{j=1}^{|S|} \left(\phi_{\sigma_j}^{\top} (\boldsymbol{\theta} - \boldsymbol{\theta}') - \log \frac{\sum_{k=j}^{|S|} e^{\phi_{\sigma_k}^{\top} \boldsymbol{\theta}}}{\sum_{k=j}^{|S|} e^{\phi_{\sigma_k}^{\top} \boldsymbol{\theta}'}} \right) \right]. \quad (\text{F.6})$$

Fix a stage j and a ranking σ . We define

$$p_{k'}(\boldsymbol{\theta}) := \frac{\exp(\phi_{\sigma_{k'}}^{\top} \boldsymbol{\theta})}{\sum_{k=j}^{|S|} \exp(\phi_{\sigma_k}^{\top} \boldsymbol{\theta})}, \quad \text{where } k' \in \{j, \dots, |S|\},$$

which corresponds to the Multinomial Logit (MNL) probability of selecting action $\sigma_{k'}$ at position j , given the parameter $\boldsymbol{\theta}$ and the choice set S . Moreover, we define

$$f(\boldsymbol{\theta}) := \log \left(\sum_{k=j}^{|S|} e^{\phi_{\sigma_k}^{\top} \boldsymbol{\theta}} \right).$$

Then, by applying the mean value form of Taylor's theorem, there exists $\bar{\boldsymbol{\theta}} = (1-c)\boldsymbol{\theta} + c\boldsymbol{\theta}'$ for some $c \in (0, 1)$ such that

$$\begin{aligned} -\log \frac{\sum_{k=j}^{|S|} e^{\phi_{\sigma_k}^{\top} \boldsymbol{\theta}}}{\sum_{k=j}^{|S|} e^{\phi_{\sigma_k}^{\top} \boldsymbol{\theta}'}} &= f(\boldsymbol{\theta}') - f(\boldsymbol{\theta}) \\ &= \nabla_{\boldsymbol{\theta}} f(\boldsymbol{\theta})^{\top} (\boldsymbol{\theta}' - \boldsymbol{\theta}) + \frac{1}{2} (\boldsymbol{\theta}' - \boldsymbol{\theta})^{\top} \nabla_{\boldsymbol{\theta}}^2 f(\bar{\boldsymbol{\theta}}) (\boldsymbol{\theta}' - \boldsymbol{\theta}) \quad (\text{Taylor's theorem}) \\ &\leq \nabla_{\boldsymbol{\theta}} f(\boldsymbol{\theta})^{\top} (\boldsymbol{\theta}' - \boldsymbol{\theta}) + \frac{1}{2} \sum_{k=j}^{|S|} p_k(\bar{\boldsymbol{\theta}}) (\phi_{\sigma_k}^{\top} (\boldsymbol{\theta}' - \boldsymbol{\theta}))^2 \\ &\leq \sum_{k=j}^{|S|} p_k(\boldsymbol{\theta}) \phi_{\sigma_k}^{\top} (\boldsymbol{\theta}' - \boldsymbol{\theta}) + \frac{1}{2} \sum_{a \in S} (\phi_a^{\top} (\boldsymbol{\theta}' - \boldsymbol{\theta}))^2, \end{aligned} \quad (\text{F.7})$$

where the first inequality holds because

$$\nabla_{\boldsymbol{\theta}}^2 f(\bar{\boldsymbol{\theta}}) = \sum_{k=j}^{|S|} p_k(\bar{\boldsymbol{\theta}}) \phi_{\sigma_k} \phi_{\sigma_k}^{\top} - \left(\sum_{k=j}^{|S|} p_k(\bar{\boldsymbol{\theta}}) \phi_{\sigma_k} \right) \left(\sum_{k=j}^{|S|} p_k(\bar{\boldsymbol{\theta}}) \phi_{\sigma_k} \right)^{\top} \leq \sum_{k=j}^{|S|} p_k(\bar{\boldsymbol{\theta}}) \phi_{\sigma_k} \phi_{\sigma_k}^{\top}.$$

Plugging Equation (F.7) into Equation (F.6), we get

$$\begin{aligned} D_{\text{KL}}(\mathbb{P}_{\boldsymbol{\theta}}(\cdot|S) \|\mathbb{P}_{\boldsymbol{\theta}'}(\cdot|S)) &\leq \mathbb{E}_{\sigma \sim \mathbb{P}_{\boldsymbol{\theta}}(\cdot|S)} \left[\sum_{j=1}^{|S|} \left(\phi_{\sigma_j}^{\top} (\boldsymbol{\theta} - \boldsymbol{\theta}') - \sum_{k=j}^{|S|} p_k(\boldsymbol{\theta}) \phi_{\sigma_k}^{\top} (\boldsymbol{\theta} - \boldsymbol{\theta}') + \frac{1}{2} \sum_{a \in S} (\phi_a^{\top} (\boldsymbol{\theta}' - \boldsymbol{\theta}))^2 \right) \right] \\ &= \mathbb{E}_{\sigma \sim \mathbb{P}_{\boldsymbol{\theta}}(\cdot|S)} \left[\sum_{j=1}^{|S|} \underbrace{\mathbb{E}_{\sigma_j} \left[\phi_{\sigma_j}^{\top} (\boldsymbol{\theta} - \boldsymbol{\theta}') - \sum_{k=j}^{|S|} p_k(\boldsymbol{\theta}) \phi_{\sigma_k}^{\top} (\boldsymbol{\theta} - \boldsymbol{\theta}') \mid \sigma_1, \dots, \sigma_{j-1} \right]}_{=0} \right] \\ &\quad + \frac{|S|}{2} \sum_{a \in S} (\phi_a^{\top} (\boldsymbol{\theta}' - \boldsymbol{\theta}))^2 \\ &\leq \frac{K}{2} \sum_{a \in S} (\phi_a^{\top} (\boldsymbol{\theta}' - \boldsymbol{\theta}))^2, \quad (|S| \leq K) \end{aligned}$$

which concludes the proof. \square

G Additional Discussions

This section provides further discussion of our approach. In Subsection G.1, we show how it extends to the regret-minimization setting and improves the existing $\mathcal{O}(e^B)$ dependence. In Subsection G.2, we describe how the approach also applies to the active-learning setting considered in [19].

G.1 Avoiding e^B Scaling in Regret Minimization

In this subsection, we show that our technique for eliminating the $\mathcal{O}(e^B)$ dependence in the leading term—primarily by dividing the total rounds into warm-up and non-warm-up phases—can also be applied in the regret-minimization setting (e.g., Bengs et al. 9).

First, we formally define the cumulative regret as:

$$\mathbf{Reg}_T := \sum_{t=1}^T \left(r_{\theta^*}(x_t, \pi^*(x_t)) - \max_{a \in S_t} r_{\theta^*}(x_t, a) \right),$$

where the contexts x_t can be given arbitrarily (in contrast to the main paper, where they are drawn from a fixed distribution). We assume the linear rewards (Assumption 1) as in the main paper. Moreover, we use the RB update rule (Procedure 2) in this subsection.

Next, we define

$$g_t(S) := \frac{1}{|S|} \sum_{a \in S} \dot{\mu} \left((\phi(x_t, a) - \phi(x_t, \hat{a}_t))^{\top} \hat{\theta}_t \right) \|\phi(x_t, a) - \phi(x_t, \hat{a}_t)\|_{H_t^{-1}}^2,$$

where $\hat{a}_t := \operatorname{argmax}_{a \in \mathcal{A}} \phi(x_t, a)^{\top} \hat{\theta}_t$. The quantity measures the average uncertainty between the feature vectors of actions in S and that of the current best action \hat{a}_t .

We then select the assortment S_t in a greedy manner, following the procedure in Equation (11). We initialize $S = \{\hat{a}_t\}$, and subsequently add actions one by one. Specifically, at each step, we select

$$a^* \in \operatorname{argmax}_{a \in \mathcal{A} \setminus S} \Delta_t(a \mid S), \quad \text{where } \Delta_t(a \mid S) := g_t(S \cup \{a\}) - g_t(S), \quad (\text{G.1})$$

and include a^* in S if $\Delta_t(a \mid S) \geq 0$. The process continues until either $|S| = K$ or no action yields a non-negative gain. Since $g_t(\{\hat{a}_t\}) = 0$, this rule guarantees that the selected set satisfies $|S| \geq 2$.

Theorem G.1 (Regret upper bound). *In the same setting as Theorem 2, suppose $T \geq e^{K/d}$. Then, with probability at least $1 - \delta$, the assortment selection rule in Equation (G.1), combined with the RB loss update strategy (Procedure 2), guarantees*

$$\mathbf{Reg}_T = \tilde{\mathcal{O}} \left(d \sqrt{\sum_{t=1}^T \frac{1}{|S_t|}} + \frac{d^2}{\kappa} \right).$$

Proof of Theorem G.1. By the definition of the cumulative regret, we have

$$\begin{aligned} \mathbf{Reg}_T &:= \sum_{t=1}^T \left(\phi(x_t, \pi^*(x_t)) - \max_{a \in S_t} \phi(x_t, a) \right)^{\top} \theta^* \\ &\leq \sum_{t=1}^T \left(\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{a}_t) \right)^{\top} \theta^* \quad (\hat{a}_t \in S_t) \\ &\leq \sum_{t=1}^T \left(\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{a}_t) \right)^{\top} \left(\theta^* - \hat{\theta}_t \right). \quad (\hat{a}_t := \operatorname{argmax}_{a \in \mathcal{A}} \phi(x_t, a)^{\top} \hat{\theta}_t) \end{aligned}$$

Recall the definition of the set of *warm-up rounds*, i.e., $\mathcal{T}^w = \{t \in [T] : \max_{a, a' \in \mathcal{A}} \|\phi(x_t, a) - \phi(x_t, a')\|_{H_t^{-1}} \geq 1/\beta_T(\delta)\}$. Then, we get

$$\begin{aligned}
\mathbf{Reg}_T &\leq 4B|\mathcal{T}^w| + \sum_{t \notin \mathcal{T}^w} (\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{a}_t))^\top (\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_t) \\
&\leq \mathcal{O}\left(\frac{B}{\kappa}\beta_{T+1}(\delta)^2 d \log\left(1 + \frac{T}{d\lambda}\right)\right) + \sum_{t \notin \mathcal{T}^w} \|\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{a}_t)\|_{H_t^{-1}} \|\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_t\|_{H_t} \\
&\quad \text{(Lemma D.2 and Cauchy-Schwartz ineq.)} \\
&\leq \mathcal{O}\left(\frac{B}{\kappa}\beta_{T+1}(\delta)^2 d \log\left(1 + \frac{T}{d\lambda}\right)\right) + \beta_T(\delta) \sum_{t \notin \mathcal{T}^w} \|\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{a}_t)\|_{H_t^{-1}}. \\
&\quad \text{(Proposition D.1)}
\end{aligned}$$

For simplify the presentation, let $\psi_t(a, a') = \phi(x_t, a) - \phi(x_t, a')$. Then, for $t \notin \mathcal{T}^w$, we have

$$\begin{aligned}
\frac{1}{\dot{\mu}(\psi_t(\pi^*(x_t), \hat{a}_t)^\top \boldsymbol{\theta}^*)} &= \frac{\left(1 + e^{\psi_t(\pi^*(x_t), \hat{a}_t)^\top \boldsymbol{\theta}^*}\right)^2}{e^{\psi_t(\pi^*(x_t), \hat{a}_t)^\top \boldsymbol{\theta}^*}} \\
&\leq \left(1 + e^{\psi_t(\pi^*(x_t), \hat{a}_t)^\top \boldsymbol{\theta}^*}\right)^2 \quad (\psi_t(\pi^*(x_t), \hat{a}_t)^\top \boldsymbol{\theta}^* \geq 0) \\
&\leq \left(1 + e^{\psi_t(\pi^*(x_t), \hat{a}_t)^\top (\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_t)}\right)^2 \quad (\hat{a}_t := \operatorname{argmax}_{a \in \mathcal{A}} \phi(x_t, a)^\top \hat{\boldsymbol{\theta}}_t) \\
&\leq \left(1 + e^{\|\psi_t(\pi^*(x_t), \hat{a}_t)^\top\|_{H_t^{-1}} \|\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_t\|_{H_t}}\right)^2 \\
&\leq \left(1 + e^{\frac{\beta_t(\delta)}{\beta_T(\delta)}}\right)^2 \leq (1 + e)^2. \quad (t \notin \mathcal{T}^w \text{ and Proposition D.1})
\end{aligned}$$

Therefore, the same line of reasoning used in the proof of Theorem 2 applies here as well.

$$\begin{aligned}
\sum_{t \notin \mathcal{T}^w} \|\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{a}_t)\|_{H_t^{-1}} &\leq e(1 + e)^2 \sum_{t \notin \mathcal{T}^w} \frac{1}{|S_t^*|} \sum_{a \in S_t^*} \dot{\mu}(\psi_t(a, \hat{a}_t)^\top \hat{\boldsymbol{\theta}}_t) \|\psi_t(a, \hat{a}_t)\|_{H_t^{-1}} \quad (S_t^* := \{\pi^*(x_t), \hat{a}_t\}) \\
&\leq e(1 + e)^2 \sqrt{\sum_{t=1}^T \frac{1}{|S_t|}} \sqrt{\sum_{t \notin \mathcal{T}^w} \frac{|S_t|}{|S_t^*|} \sum_{a \in S_t^*} \dot{\mu}(\psi_t(a, \hat{a}_t)^\top \hat{\boldsymbol{\theta}}_t) \|\psi_t(a, \hat{a}_t)\|_{H_t^{-1}}^2} \\
&\quad \text{(Cauchy-Schwartz ineq., and } \mathcal{T}^w \subseteq [T]) \\
&\leq e(1 + e)^2 \sqrt{\sum_{t=1}^T \frac{1}{|S_t|}} \sqrt{\sum_{t \notin \mathcal{T}^w} \frac{|S_t|}{|\mathcal{S}_t|} \sum_{a \in \mathcal{S}_t} \dot{\mu}(\psi_t(a, \hat{a}_t)^\top \hat{\boldsymbol{\theta}}_t) \|\psi_t(a, \hat{a}_t)\|_{H_t^{-1}}^2} \\
&\leq \sqrt{2}e^2(1 + e)^2 \sqrt{\sum_{t=1}^T \frac{1}{|S_t|}} \sqrt{\sum_{t \notin \mathcal{T}^w} \sum_{a \in S_t} \dot{\mu}(\psi_t(a, \hat{a}_t)^\top \hat{\boldsymbol{\theta}}_t) \|\psi_t(a, \hat{a}_t)\|_{\Sigma_t^{-1}}^2} \quad \text{(Lemma D.1)} \\
&= \mathcal{O}\left(\sqrt{\sum_{t=1}^T \frac{1}{|S_t|}} \cdot \sqrt{d \log(KT)}\right), \quad \text{(Lemma E.6)}
\end{aligned}$$

where, in the last inequality, we introduce $\Sigma_t = \sum_{s \in [t-1] \setminus \mathcal{T}^w} \sum_{a \in S_s} \dot{\mu}(\psi_t(a, \hat{a}_t)^\top \hat{\boldsymbol{\theta}}_s) \psi_t(a, \hat{a}_t) \psi_t(a, \hat{a}_t)^\top + \lambda \mathbf{I}_d$, and use the relation $H_t \geq \frac{1}{2e^2} \Lambda_t \geq \frac{1}{2e^2} \Sigma_t$ (Λ_t is defined in Equation (D.1)). In the final equality, we apply the fact that if $\lambda = d \log T$ and $T \geq e^{K/d}$, then, for any t ,

$$\dot{\mu}(\psi_t(a, \hat{a}_t)^\top \hat{\boldsymbol{\theta}}_t) \|\psi_t(a, \hat{a}_t)\|_{\Sigma_t^{-1}}^2 \leq \frac{2K}{\lambda} \leq 2.$$

Therefore, using $\beta_T(\delta) = \mathcal{O}(B\sqrt{d \log(KT)} + B\sqrt{\lambda})$, we obtain, with probability at least $1 - \delta$,

$$\mathbf{Reg}_T = \tilde{\mathcal{O}} \left(d \sqrt{\sum_{t=1}^T \frac{1}{|S_t|}} + \frac{d^2}{\kappa} \right).$$

This completes the proof of Theorem G.1. \square

G.2 Extension to Active Learning Setting

In this subsection, we consider a different setting—referred to as the *active learning setting*—where the learner has access to the entire context set \mathcal{X} , and the objective is to minimize the following *worst-case suboptimality gap*, defined as:

$$\mathbf{WorstSubOpt}(T) := \max_{x \in \mathcal{X}} [r_{\theta^*}(x, \pi^*(x)) - r_{\theta^*}(x, \hat{\pi}(x))].$$

This setting has received increasing attention in recent work [50, 47, 71, 19, 49, 76, 38]. However, most existing approaches focus exclusively on pairwise preference feedback. Mukherjee et al. [49] study an online learning-to-rank problem where, for each context, a fixed set of K actions is provided, and the goal is to recover the true ranking based on feedback over these K actions. In contrast, we consider a more general setting in which, for each context, a set of N actions is available. The learner selects at most K actions from this set and receives ranking feedback over the selected subset. Thekumparampil et al. [76] investigate the problem of ranking $N \geq K$ items using partial rankings over K candidates, but under a context-free setting. In contrast, we study a stochastic contextual setting, where contexts are drawn from an unknown (and fixed) distribution.

In the active learning setting, the algorithm jointly selects the context x_t —which is no longer given but actively chosen—and the assortment S_t by maximizing the average uncertainty objective. For each candidate context x , it first constructs the assortment $S_t(x)$ by solving Equation (11). It then selects $x_t = \operatorname{argmax}_{x \in \mathcal{X}} f_t(S_t(x))$, and sets $S_t = S_t(x_t)$. The rest of the algorithm proceeds in the same manner as Algorithm 3.

For simplicity, we focus on the RB loss and its corresponding update rule (Procedure 2). Then, we can obtain the following bound on the worst-case suboptimality gap.

Theorem G.2. *Under the same setting as Theorem 2, with probability at least $1 - \delta$, the selection rule for (x_t, S_t) described above achieves:*

$$\mathbf{WorstSubOpt}(T) = \tilde{\mathcal{O}} \left(\frac{d}{T} \sqrt{\sum_{t=1}^T \frac{1}{|S_t|}} + \frac{d^2}{\kappa T} \right).$$

Proof of Theorem G.2. By the definition of the worst-case suboptimality gap, we have

$$\begin{aligned} \mathbf{WorstSubOpt}(T) &= \max_{x \in \mathcal{X}} [(\phi(x, \pi^*(x)) - \phi(x, \hat{\pi}_T(x)))^\top \theta^*] \\ &\leq \max_{x \in \mathcal{X}} [(\phi(x, \pi^*(x)) - \phi(x, \hat{\pi}_T(x)))^\top (\theta^* - \hat{\theta}_{T+1})] \\ &\quad (\hat{\pi}_T(x) = \operatorname{argmax}_{a \in \mathcal{A}} \phi(x, a)^\top \hat{\theta}_{T+1}) \\ &= \frac{1}{T} \sum_{t=1}^T \max_{x \in \mathcal{X}} [(\phi(x, \pi^*(x)) - \phi(x, \hat{\pi}_T(x)))^\top (\theta^* - \hat{\theta}_{T+1})]. \end{aligned}$$

We adopt the same definitions for $\mathcal{T}^w := \{t \in [T] : \max_{a, a' \in \mathcal{A}} \|\phi(x_t, a) - \phi(x_t, a')\|_{H_t^{-1}} \geq 1/\beta_{T+1}(\delta)\}$. Thus, we get

$$\begin{aligned} &\frac{1}{T} \sum_{t=1}^T \max_{x \in \mathcal{X}} [(\phi(x, \pi^*(x)) - \phi(x, \hat{\pi}_T(x)))^\top (\theta^* - \hat{\theta}_{T+1})] \\ &\leq \mathcal{O} \left(\frac{B}{\kappa T} \beta_{T+1}(\delta)^2 d \log \left(1 + \frac{T}{d\lambda} \right) \right) \quad (\text{Lemma D.2}) \\ &\quad + \frac{1}{T} \sum_{t \notin \mathcal{T}^w} \max_{x \in \mathcal{X}} [(\phi(x, \pi^*(x)) - \phi(x, \hat{\pi}_T(x)))^\top (\theta^* - \hat{\theta}_{T+1})]. \quad (\text{G.2}) \end{aligned}$$

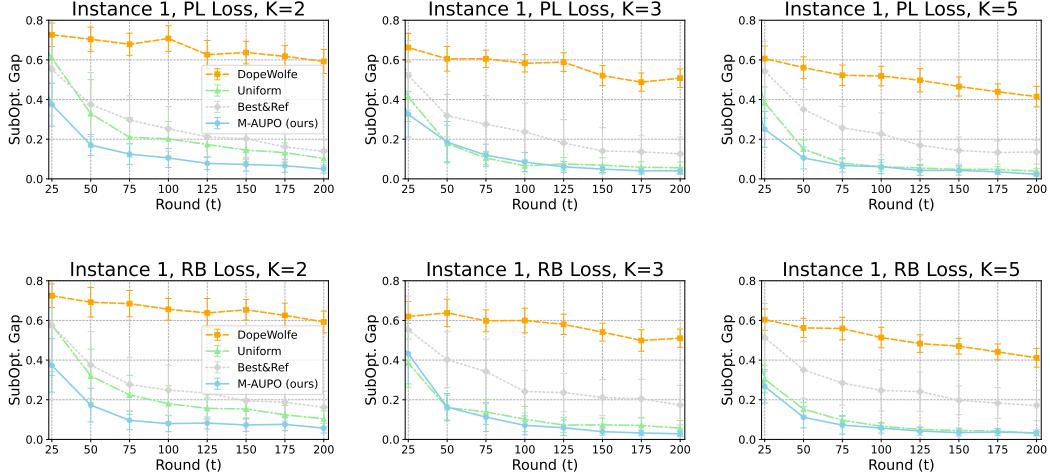


Figure G.1: Performance comparisons for Instance 1 (Stochastic contexts) with $K = 2, 3$, and 5 , evaluated under the PL loss (first row) and RB loss (second row).

To further bound the last term of Equation (G.2), we get

$$\begin{aligned}
& \frac{1}{T} \sum_{t \notin \mathcal{T}^w} \max_{x \in \mathcal{X}} \left[(\phi(x, \pi^*(x)) - \phi(x, \hat{\pi}_T(x)))^\top (\theta^* - \hat{\theta}_{T+1}) \right] \\
& \leq \frac{1}{T} \sum_{t \notin \mathcal{T}^w} \max_{x \in \mathcal{X}} \left[\|\phi(x, \pi^*(x)) - \phi(x, \hat{\pi}_T(x))\|_{H_{T+1}^{-1}} \right] \|\theta^* - \hat{\theta}_{T+1}\|_{H_{T+1}} \quad (\text{Hölder's ineq.}) \\
& \leq \frac{\beta_{T+1}(\delta)}{T} \sum_{t \notin \mathcal{T}^w} \max_{x \in \mathcal{X}} \left[\|\phi(x, \pi^*(x)) - \phi(x, \hat{\pi}_T(x))\|_{H_t^{-1}} \right].
\end{aligned}$$

$(H_{T+1} \succeq H_t \text{ and Corollary C.1, with prob. } 1 - \delta)$

For simplicity, let $x_t^* = \operatorname{argmax}_x \|\phi(x, \pi^*(x)) - \phi(x, \hat{\pi}_T(x))\|_{H_t^{-1}}$. It therefore suffices to bound $\sum_t \|\phi(x_t^*, \pi^*(x)) - \phi(x_t^*, \hat{\pi}_T(x))\|_{H_t^{-1}}$. From here, the same argument as in the proof of Theorem 2 applies, with the only modification that the procedure now selects both x_t and S_t . This completes the proof of Theorem G.2. \square

H Experimental Details and Additional Results

H.1 Synthetic Data

Setup. In the synthetic data experiment, we sample the true but unknown parameter $\theta^* \in \mathbb{R}^d$ from a d -dimensional standard normal distribution, i.e., $\theta^* \sim \mathcal{N}(0, I_d)$, and then normalize it to ensure $\|\theta^*\|_2 \leq 1$. We consider four different types of context sets \mathcal{X} :

1. **Instance 1 (Stochastic contexts):** For each $x \in \mathcal{X}$, the feature vectors $\phi(x, \cdot)$ are sampled from a standard normal distribution and then normalized to satisfy $\|\phi(x, \cdot)\|_2 \leq 1$. Here, $|\mathcal{X}| = 100$.
2. **Instance 2 (Non-contextual):** A single shared context is used for all rounds, i.e., $\mathcal{X} = \{x_1\}$ and $|\mathcal{X}| = 1$. The corresponding feature vectors $\phi(x_1, \cdot)$ are sampled from a standard normal distribution and then normalized to satisfy $\|\phi(x_1, \cdot)\|_2 \leq 1$.
3. **Instance 3 (Hard-to-learn contexts):** For each $x \in \mathcal{X}$, the feature vectors $\phi(x, \cdot)$ are constructed such that most of them are approximately orthogonal to the true parameter θ^* . Here, $|\mathcal{X}| = 100$.

4. Instance 4 (Skewed stochastic contexts): For each $x \in \mathcal{X}$, the feature vectors $\phi(x, \cdot)$ are sampled in a skewed or biased manner and then normalized to satisfy $\|\phi(x, \cdot)\|_2 \leq 1$. Here, $|\mathcal{X}| = 100$. This is our main experimental setup in Section 6.1.

Additionally, we set the feature dimension to $d = 5$ and the number of available actions to $|\mathcal{A}| = N = 100$. The suboptimality gap is measured every 25 rounds. All results are averaged over 20 independent runs with different random seeds, and standard errors are reported to indicate variability. The experiments are run on a Xeon(R) Gold 6226R CPU @ 2.90GHz (16 cores).

Baselines. We evaluate our proposed algorithm, M-AUPO, against three baselines: (i) **DopeWolfe** [76], a method designed for non-contextual K -subset selection; (ii) **Uniform**, which selects assortments of size K uniformly at random; and (iii) **Best&Ref**, which forms a pair of actions ($|S_t| = 2$) by combining one action from the current policy with another from a reference policy (e.g., uniform random or SFT), following the setup in Online GSHF [86] and XPO [85].

When using the PL loss-based update in our algorithm, M-AUPO, the exact expectation over rankings is computationally expensive. To mitigate this, we approximate the expectation via Monte Carlo sampling with 5 samples. Consequently, the computational cost of the PL-based approach is higher than that of the RB-based approach (see Table H.1).

Thekumparampil et al. [76] propose a D-optimal design approach for the Plackett-Luce objective to efficiently select informative subsets of items for comparison. Recognizing the computational complexity inherent in this method, they introduce a randomized Frank-Wolfe algorithm, named **DopeWolfe**, which approximates the optimal design by solving linear maximization sub-problems on randomly chosen variables. This approach reduces computational overhead while maintaining effective learning performance. However, their approach is specifically tailored to the single-context setting (e.g., **Instance 2**) and may not generalize well to the multiple-context scenarios (e.g., **Instances 1, 3, and 4**). While their original implementation updates the model parameters using a maximum likelihood estimation (MLE) procedure, we instead adopt an online update strategy (as described in Procedures 1 and 2) to ensure a fair comparison across all methods. For sampling size parameter R , we set $R = \min\{\binom{N}{K}, 100, 000\}$.

The uniform random selection strategy, **Uniform**, selects K actions uniformly at random from the available action set \mathcal{A} at each round, without utilizing any uncertainty or reward-based information.

Best&Ref constructs an action pair ($|S_t| = 2$) by combining two distinct sources of actions. The first action is chosen to maximize the current reward estimate, while the second is sampled from a reference policy—such as a uniform random policy or a supervised fine-tuned (SFT) model. This pairing mechanism follows the framework introduced in Online GSHF [86] and XPO [85]. In our experiments, we use the uniform random policy as the reference.

Performance measure. Since computing the exact suboptimality gap is challenging under a general distribution ρ , we instead evaluate the *realized regret*, which serves as a slightly relaxed proxy for the suboptimality gap.

$$\begin{aligned} \mathbf{SubOpt}(T) &\lesssim \frac{1}{T} \sum_{t=1}^T (\phi(x_t, \pi^*(x_t)) - \phi(x_t, \hat{\pi}_T(x_t)))^\top \theta^* + \underbrace{\tilde{\mathcal{O}}\left(\frac{1}{\sqrt{T}}\right)}_{\text{incurred by MDS terms}} \\ &\leq \underbrace{\frac{1}{T} \sum_{t=1}^T (\phi(x_t, \pi^*(x_t)) - \phi(x_t, \pi_t(x_t)))^\top \theta^*}_{=: \text{realized regret}} + \tilde{\mathcal{O}}\left(\frac{1}{\sqrt{T}}\right), \end{aligned}$$

where we define $\pi_t(x) := \operatorname{argmax}_a \phi(x, a)^\top \hat{\theta}_t$, and let $\hat{\pi}_T$ denote the best policy among $\{\pi_t\}_{t=1}^T$, possibly selected using a validation set.

Results. We present performance comparisons in Figures G.1 through H.3, corresponding to Instances 1 through 4, respectively. Overall, our algorithm, M-AUPO, consistently outperforms other baseline methods. The only exception is in Instance 2 (Figure H.1), a special case of the non-contextual setting, where M-AUPO performs slightly worse than DopeWolfe. This is an expected outcome, as DopeWolfe leverages a D-optimal design strategy, which is known to be highly effective in the single-context setting. However, it is important to note that DopeWolfe completely fails in more

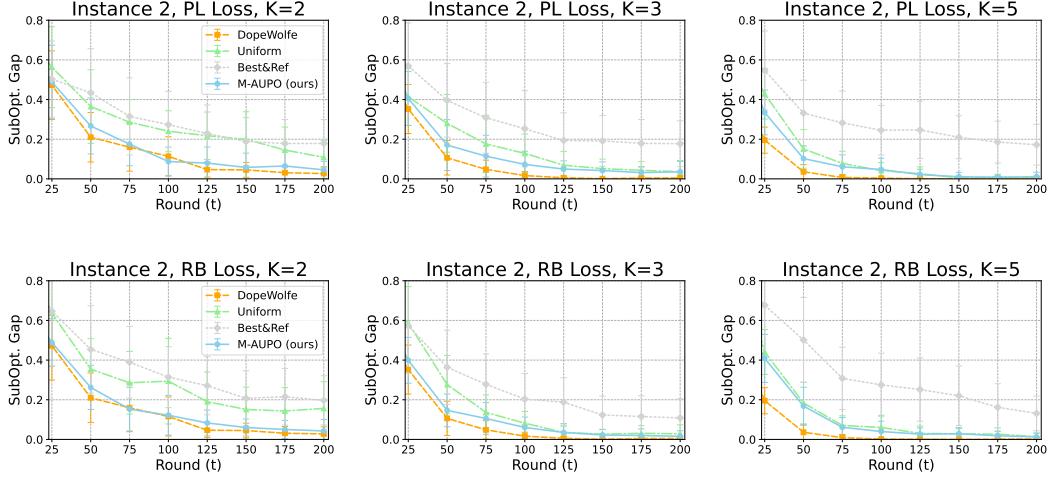


Figure H.1: Performance comparisons for Instance 2 (Non-contextual) with $K = 2, 3$, and 5 , evaluated under the PL loss (first row) and RB loss (second row).

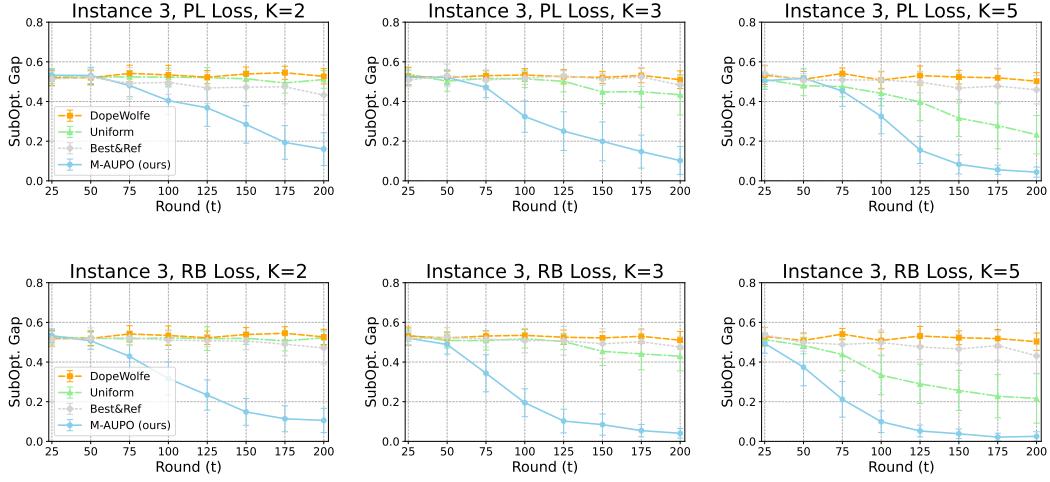


Figure H.2: Performance comparisons for Instance 3 (Hard-to-learn contexts) with $K = 2, 3$, and 5 , evaluated under the PL loss (first row) and RB loss (second row).

general contextual scenarios (Figures G.1, H.2, and H.3), and its computational cost is significantly higher than that of our approach (see Table H.1).

K	DopeWolfe	Uniform	Best&Ref	M-AUPO (PL)	M-AUPO (RB)
2	7.28 s	0.10 s	0.10s	1.09 s	1.04 s
3	99.6 s	0.18 s	0.10s	2.35 s	1.06 s
5	150.5 s	0.35 s	0.10s	2.71 s	1.10 s
7	218.8 s	0.58 s	0.10s	4.10 s	1.13 s
10	331.1 s	0.99 s	0.10s	6.89 s	1.31 s

Table H.1: Runtime comparison over 200 rounds (seconds)

The uniform random assortment selection strategy, Uniform, shows reasonable performance—though consistently inferior to M-AUPO—in Instances 1, 2, and 4, as presented in Figures G.1, H.1, and H.3,

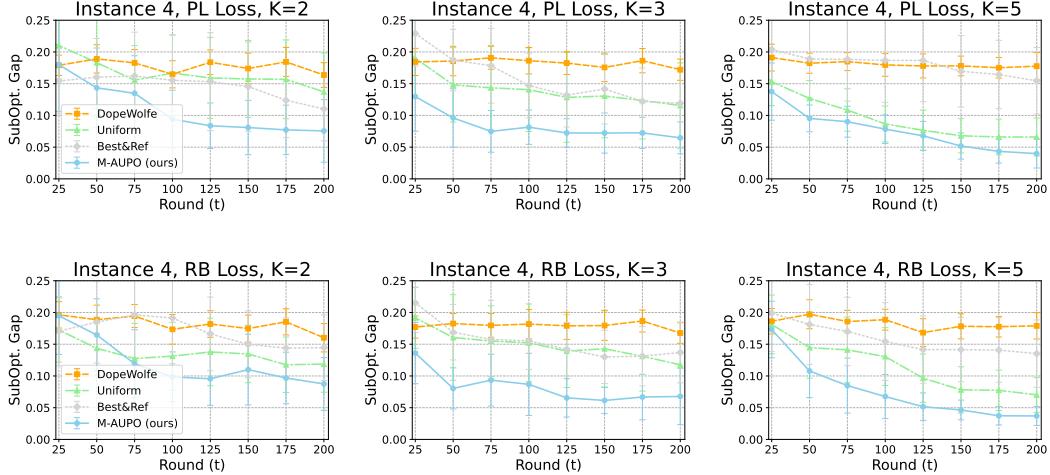


Figure H.3: Performance comparisons for Instance 4 (skewed stochastic contexts) with $K = 2, 3$, and 5 , evaluated under the PL loss (first row) and RB loss (second row).

respectively. In contrast, for Instance 3 (Figure H.2), where most features are uninformative (being nearly orthogonal to the true parameter) and thus careful action selection becomes crucial, Uniform performs substantially worse than M-AUPO.

The Best&Ref algorithm performs consistently worse than our algorithm and does not benefit from larger K , since it always selects only a pair of actions.

Moreover, the suboptimality gap consistently decreases with larger K across the three algorithms—M-AUPO, Uniform, and DopeWolfe—whereas Best&Ref shows no such improvement, as it always selects only two actions regardless of K . For M-AUPO, this trend is consistent with our theoretical results (Theorems 1 and 2). In contrast, the improvement observed for DopeWolfe suggests that its current theoretical guarantees may be loose, as their bound actually deteriorates with increasing K (recall that their theoretical guarantee worsens for larger K). This suggests that tighter bounds could be obtained by incorporating techniques similar to those introduced in our work.

Table H.2 presents the average assortment size $|S_t|$ of M-AUPO for various values of the maximum assortment size K . In most cases, the algorithm selects the full K actions, i.e., $|S_t| = K$. An exception occurs when K is large (e.g., 30 or more), which may be impractical in real-world applications due to the increased annotation burden on human labelers.

K	2	3	5	7	10	30	50
PL loss, $ S_t $	2.00	3.00	5.00	7.00	10.00	16.21	18.78
RB loss, $ S_t $	2.00	3.00	5.00	7.00	10.00	19.87	22.37

Table H.2: Assortment size $|S_t|$ of M-AUPO with varying K in the synthetic experiment

H.2 Real-World Dataset

Setup. In our real-world dataset experiments, we evaluate performance on two widely used benchmark datasets: TREC Deep Learning (TREC-DL) and NECTAR. The TREC-DL dataset provides 100 candidate answers for each query, offering a rich and diverse set of responses suitable for learning from listwise feedback. In contrast, the NECTAR dataset presents a more concise setup, with only 7 candidate answers per question. From each dataset, we randomly sample $|\mathcal{X}| = 5000$ prompts, each paired with its corresponding set of candidate actions—100 for TREC-DL and 7 for NECTAR.

We use the Gemma-2B language model [75] to construct the feature representation $\phi(x, a)$. To obtain $\phi(x, a)$, we first concatenate the input prompt x and the candidate response a into a single sequence,

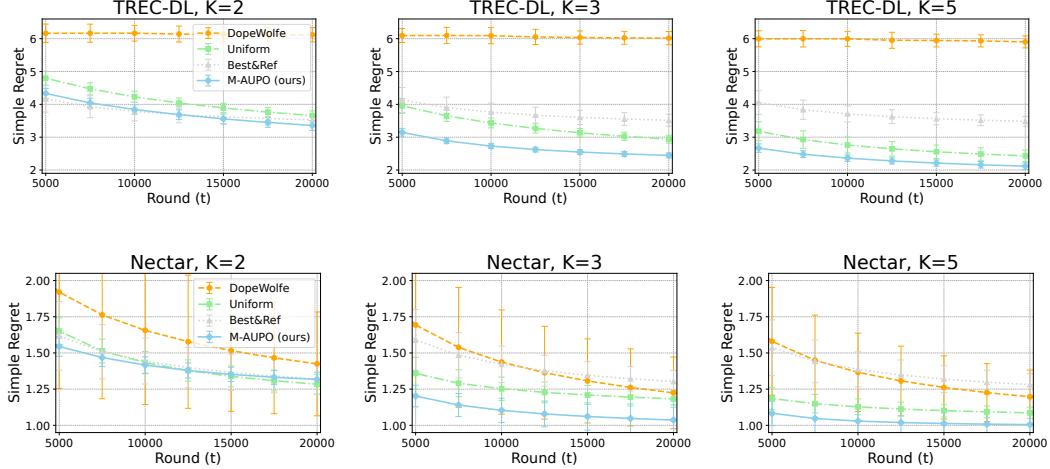


Figure H.4: Performance comparisons on the TREC-DL dataset (top row) and the NECTAR dataset (bottom row) for varying values of $K = 2, 3$, and 5 .

which is then fed into Gemma-2B. The resulting feature vector is extracted from the last hidden layer of the model and has a dimensionality of $d = 2048$. We then apply ℓ_1 normalization to enhance numerical stability and ensure consistent scaling. For each round t , we sample the context index from an exponential distribution with rate $\lambda = 0.1$, which assigns higher probability to smaller indices and thus biases the selection toward earlier contexts. To generate ranking feedback and evaluate the suboptimality gap, we use the Mistral-7B reward model [33] as the ground-truth reward function.

We measure the suboptimality gap every 2,500 rounds throughout the training process and report the average performance over 10 independent runs, each with a different random seed. Along with the average, we also include the standard error to indicate variability across runs. In these experiments, we report results under the RB loss only, due to its superior performance, as demonstrated in the synthetic data experiments. The experiments are conducted on a Xeon(R) Gold 6226R CPU @ 2.90GHz (16 cores) and a single GeForce RTX 3090 GPU.

Baselines. We use the same set of baselines as in the synthetic data experiments. For DopeWolfe [76], we set the sampling size parameter R as $R = \min\{\binom{N}{K}, 1000\}$. Although a small value of $R \leq 1000$ may introduce significant approximation error—since the theoretically minimal-error choice is $R = \mathcal{O}(\binom{N}{K})$ —we adopt this smaller value in our experiment to reduce computational overhead.

Performance measure. We measure the realized regret as in the synthetic experiment.

Results. We present performance comparisons in Figure H.4. Our algorithm, M-AUPO, consistently outperforms all baselines by a significant margin. As in the synthetic data experiments, the suboptimality gap for all methods decreases as K increases. Notably, DopeWolfe performs particularly poorly on the TREC-DL dataset. This may be attributed to the use of a small sampling size R , which is insufficient compared to the full subset space of size $\binom{N}{K} = \mathcal{O}(N^K) \gg 1000 \geq R$. This result highlights an important practical limitation of DopeWolfe: despite its use of approximate optimization to reduce runtime, the method still depends on combinatorial sampling to perform well, which becomes computationally infeasible in large-scale settings. In contrast, our algorithm, M-AUPO, achieves consistently strong performance while incurring only polynomial computational cost, demonstrating superior scalability and practicality for real-world deployment.

Table H.3 reports the actual assortment size $|S_t|$ selected by M-AUPO on both datasets. In the TREC-DL experiment, $|S_t|$ is nearly equal to K for all values of K , as the number of available actions is large ($N = 100$). In contrast, in the NECTAR experiment, where the number of available actions is much smaller ($N = 7$), the actual assortment size $|S_t|$ is often smaller than K , especially when $K = N$. This reduction occurs because the limited action space constrains the potential informativeness of larger assortments—for example, it becomes difficult to achieve high average uncertainty when there are too few actions to choose from.

K	2	3	5	7
TREC-DL dataset, $ S_t $	2.00	3.00	5.00	6.98
NECTAR dataset, $ S_t $	2.00	2.99	4.45	4.76

Table H.3: Assortment size $|S_t|$ of M-AUPO with varying K in the real-world dataset experiment

I Limitations

In this paper, we primarily focus on the online PbRL setting, where contexts are drawn stochastically from a fixed distribution. However, we do not explore the offline setting [93], which may involve a different set of challenges. As a result, it remains an open question whether similar improvements—such as better performance with larger K —can be achieved in the offline setting. We view this as a promising direction for future research and leave it as an open problem.