
Combinatorial Bandits with Relative Feedback

Aadirupa Saha

Indian Institute of Science, Bangalore
aadirupa@iisc.ac.in

Aditya Gopalan

Indian Institute of Science, Bangalore
aditya@iisc.ac.in

Abstract

We consider combinatorial online learning with subset choices when only relative feedback information from subsets is available, instead of bandit or semi-bandit feedback which is absolute. Specifically, we study two regret minimisation problems over subsets of a finite ground set $[n]$, with subset-wise relative preference information feedback according to the Multinomial logit choice model. In the first setting, the learner can play subsets of size bounded by a maximum size and receives top- m rank-ordered feedback, while in the second setting the learner can play subsets of a fixed size k with a full subset ranking observed as feedback. For both settings, we devise instance-dependent and order-optimal regret algorithms with regret $O(\frac{n}{m} \ln T)$ and $O(\frac{n}{k} \ln T)$, respectively. We derive fundamental limits on the regret performance of online learning with subset-wise preferences, proving the tightness of our regret guarantees. Our results also show the value of eliciting more general top- m rank-ordered feedback over single winner feedback ($m = 1$). Our theoretical results are corroborated with empirical evaluations.

1 Introduction

Online learning over subsets with absolute or *cardinal* utility feedback is well-understood in terms of statistically efficient algorithms for bandits or semi-bandits with large, combinatorial subset action spaces [15, 31]. In such settings the learner aims to find the subset with highest value, and upon testing a subset observes either noisy rewards from its constituents or an aggregate reward. In many natural settings, however, information obtained about the utilities of alternatives chosen is inherently relative or *ordinal*, e.g., recommender systems [25, 34], crowdsourcing [16], multi-player game ranking [22], market research and social surveys [9, 6, 24], and in other systems where humans are often more inclined to express comparative preferences.

The framework of dueling bandits [43, 46] represents a promising attempt to model online optimisation with *pairwise* preference feedback. However, our understanding of the more general and realistic online learning setting of combinatorial subset choices and *subset-wise* feedback is relatively less developed than the case of observing absolute, subset-independent reward information.

In this work, we consider a generalisation of the dueling bandit problem where the learner, instead of choosing only two arms, selects a subset of (up to) $k \geq 2$ many arms in each round. The learner subsequently observes as feedback a rank-ordered list of $m \geq 1$ items from the subset, generated probabilistically according to an underlying subset-wise preference model – in this work the Plackett-Luce distribution on rankings based on the multinomial logit (MNL) choice model [7] – in which each arm has an unknown positive value. Simultaneously, the learner earns as reward the average value of the subset played in the round. The goal of the learner is to play subsets to minimise its cumulative regret with respect to the subset with highest value.

Achieving low regret with subset-wise preference feedback is relevant in settings where deviating from choosing an optimal subset of alternatives comes with a cost (driven by considerations like revenue) even during the learning phase, but where the feedback information provides purely relative feedback. For instance, consider a beverage company that experimentally develops several variants of a drink (arms or alternatives), a best-selling subset of which it wants to learn to put up in the open market by trial and error. Each time a subset of items is put up, in parallel the company elicits relative preference feedback about the subset from, say, a team of expert tasters or through crowdsourcing.

The value of a subset can be modelled as the average value of items in it, which is however not directly observable, it being function of the open market response to the offered subset. The challenge thus lies in optimizing the subset selection over time by observing only relative preferences (made precise by the notion of Top- k -regret, Section 2.2).

A challenging feature of this problem, with subset plays and relative feedback, is the combinatorially large action and feedback space, much like those in combinatorial bandits [14, 19]. The key question here is whether (and if so, how) structure in the subset choice model – defined compactly by only a few parameters (as many as the number of arms) – can be exploited to give algorithms whose regret does not explode combinatorially. The contributions of this paper are:

(1). We consider the problem of regret minimisation when subsets of items $\{1, \dots, n\}$ of size at most k can be played, top $m \leq k$ rank-ordered feedback is received according to the MNL model, and the value of a subset is the mean MNL-parameter value of the items in the subset. We propose an upper confidence bound (UCB)-based algorithm, with a new max-min subset-building rule and a lightweight space requirement of tracking $O(n^2)$ pairwise item estimates, showing that it enjoys instance-dependent regret in T rounds of $O(\frac{n}{m} \ln T)$. This is shown to be order-optimal by exhibiting a lower bound of $\Omega(\frac{n}{m} \ln T)$ on the regret for any *No-regret* algorithm. Our results imply that the optimal regret does not vary with the maximum subset size (k) that can be played, but improves multiplicatively with the length of top m -rank-ordered feedback received per round (Sec. 3).

(2). We consider a related regret minimisation setting in which subsets of size exactly k must be played, after which a ranking of the k items is received as feedback, and where the zero-regret subset consists of the k items with the highest MNL-parameter values. In this case, our analysis reveals a fundamental lower bound on regret of $\Omega(\frac{n-k}{k\Delta_{(k)}} \ln T)$, where the problem complexity now depends on the parameter difference between the k^{th} and $(k+1)^{\text{th}}$ best item of the MNL model. We follow this up with a subset-playing algorithm (Alg. 3) for this problem – a recursive variant of the earlier UCB-based algorithm – with a matching, optimal regret guarantee of $O(\frac{(n-k) \ln T}{k\Delta_{(k)}})$ (Sec. 4).

We also provide extensive numerical evaluations supporting our theoretical findings. Due to space constraints, a discussion on related work appears in the Appendix.

2 Preliminaries and Problem Statement

Notation. We denote by $[n]$ the set $\{1, 2, \dots, n\}$. For any subset $S \subseteq [n]$, we let $|S|$ denote the cardinality of S . When there is no confusion about the context, we often represent (an unordered) subset S as a vector (or ordered subset) S of size $|S|$ according to, say, a fixed global ordering of all the items $[n]$. In this case, $S(i)$ denotes the item (member) at the i th position in subset S . For any ordered set S , $S(i : j)$ denotes the set of items from position i to j , $i < j$, $\forall i, j \in [|S|]$. $\Sigma_S = \{\sigma \mid \sigma \text{ is a permutation over items of } S\}$, where for any permutation $\sigma \in \Sigma_S$, $\sigma(i)$ denotes the element at the i -th position in σ , $i \in [|S|]$. We also denote by Σ_S^m the set of permutations of any m -subset of S , for any $m \in [k]$, i.e. $\Sigma_S^m := \{\Sigma_{S'} \mid S' \subseteq S, |S'| = m\}$. $\mathbf{1}(\varphi)$ is generically used to denote an indicator variable that takes the value 1 if the predicate φ is true, and 0 otherwise. $Pr(A)$ is used to denote the probability of event A , in a probability space that is clear from the context.

Definition 1 (Multinomial logit probability model). A Multinomial logit (MNL) probability model $\text{MNL}(n, \theta)$, specified by positive parameters $(\theta_1, \dots, \theta_n)$, is a collection of probability distributions $\{Pr(\cdot|S) : S \subset [n], S \neq \emptyset\}$, where for each non-empty subset $S \subseteq [n]$, $Pr(i|S) = \frac{\theta_i \mathbf{1}(i \in S)}{\sum_{j \in S} \theta_j}$ $\forall 1 \leq i \leq n$. The indices $1, \dots, n$ are referred to as ‘items’ or ‘arms’.

(i). *Best-Item:* Given an $\text{MNL}(n, \theta)$ instance, we define the *Best-Item* $a^* \in [n]$, to be the item with highest MNL parameter if such a unique item exists, i.e. $a^* := \arg \max_{i \in [n]} \theta_i$.

(ii). *Top- k Best-Items:* Given any instance of $\text{MNL}(n, \theta)$ we define the *Top- k Best-Items* $S_{(k)} \subseteq [n]$, to be the set of k distinct items with highest MNL parameters if such a unique set exists, i.e. for any pair of items $i \in S_{(k)}$ and $j \in [n] \setminus S_{(k)}$, $\theta_i > \theta_j$, such that $|S_{(k)}| = k$. For this problem, we assume $\theta_1 \geq \theta_2 \geq \dots \theta_k > \theta_{k+1} \geq \dots \geq \theta_n$, implying $S_{(k)} = [k]$. We also denote $\Delta_{(k)} = \theta_k - \theta_{k+1}$.

2.1 Feedback models

An *online learning algorithm* interacts with a $\text{MNL}(n, \theta)$ probability model over n items as follows. At each round $t = 1, 2, \dots$, the algorithm plays a subset $S_t \subseteq [n]$ of (distinct) items, with $|S_t| \leq k$, upon which it receives stochastic feedback defined as:

1. **Winner Feedback:** In this case, the environment returns a single item J drawn independently from probability distribution $Pr(\cdot|S)$, i.e., $Pr(J = j|S) = \frac{\theta_j}{\sum_{\ell \in S} \theta_\ell} \forall j \in S$.

2. **Top- m -ranking Feedback** ($1 \leq m \leq k-1$): Here, the environment returns an ordered list of m items sampled without replacement from the $MNL(n, \theta)$ probability model on S . More formally, the environment returns a partial ranking $\sigma \in \Sigma_S^m$, drawn from the probability distribution $Pr(\sigma = \sigma|S) = \prod_{i=1}^m \frac{\theta_{\sigma^{-1}(i)}}{\sum_{j \in S \setminus \{\sigma^{-1}(1), \dots, \sigma^{-1}(i-1)\}} \theta_j}$, $\sigma \in \Sigma_S^m$. This can also be seen as picking an item $\sigma^{-1}(1) \in S$ according to *Winner Feedback* from S , then picking $\sigma^{-1}(2)$ from $S \setminus \{\sigma^{-1}(1)\}$, and so on, until all elements from S are exhausted. When $m = 1$, Top- m -ranking Feedback is the same as Winner Feedback. To incorporate sets with $|S| < m$, we set $m = \min(|S|, m)$. Clearly this model reduces to Winner Feedback for $m = 1$, and a full rank ordering of the set S when $m = |S| - 1$.

2.2 Decisions (Subsets) and Regret

We consider two regret minimisation problems in terms of their decision spaces and notions of regret:

(1). **Winner-regret:** This is motivated by learning to identify the *Best-Item* a^* . At any round t , the learner can play sets of size $1, \dots, k$, but is penalised for playing any item other than a^* . Formally, we define the learner's instantaneous regret at round t as $r_t^1 = \sum_{i \in S_t} \frac{(\theta_{a^*} - \theta_i)}{|S_t|}$, and its cumulative regret from T rounds as $R_T^1 = \sum_{t=1}^T r_t^1 = \sum_{t=1}^T \left(\sum_{i \in S_t} \frac{(\theta_{a^*} - \theta_i)}{|S_t|} \right)$,

The learner aims to play sets S_t to keep the regret as *low* as possible, i.e., to play only the singleton set $S_t = \{a^*\}$ over time, as that is the only set with 0 regret. The instantaneous Winner-regret can be interpreted as a shortfall in value of the played set S_t with respect to $\{a^*\}$, where the value of a set S is simply the mean parameter value $\frac{\sum_{i \in S} \theta_i}{|S|}$ of its items.

Remark 1. Assuming $\theta_{a^*} = 1$ (we can do this without loss of generality since the MNL model is positive scale invariant, see Defn. 1), it is easy to note that for any item $i \in [n] \setminus \{a^*\}$ $p_{a^*i} := Pr(a^*|\{a^*, i\}) = \frac{\theta_{a^*}}{\theta_{a^*} + \theta_i} \geq \frac{1}{2} + \frac{\theta_{a^*} - \theta_i}{4}$ (as $\theta_i < \theta_{a^*}$, $\forall i$). Consequently, the Winner-regret as defined above, can be further bounded above (up to constant factors) as $\tilde{R}_T^1 = \sum_{t=1}^T \sum_{i \in S_t} \frac{(p_{a^*i} - \frac{1}{2})}{|S_t|}$, which, for $k = 2$, is standard dueling bandit regret [45, 47, 42].

Remark 2. An alternative notion of instantaneous regret is the shortfall in the preference probability of the best item a^* in the selected set S_t , i.e., $\tilde{r}_t^1 = \sum_{i \in S_t} \left(Pr(a^*|S_t \cup \{a^*\}) - Pr(i|S_t \cup \{a^*\}) \right)$. However, if all the MNL parameters are bounded, i.e., $\theta_i \in [a, b]$, $\forall i \in [n]$, then $\frac{1}{b} \left(\sum_{i \in S_t} \frac{(\theta_{a^*} - \theta_i)}{|S_t|+1} \right) \leq \tilde{r}_t^1 \leq \frac{1}{a} \left(\sum_{i \in S_t} \frac{(\theta_{a^*} - \theta_i)}{|S_t|+1} \right)$, implying that these two notions of regret, r_t^1 and \tilde{r}_t^1 , are only constant factors apart.

(2). **Top- k -regret:** This setting is motivated by learning to identify the set of *Top- k Best-Items* $S_{(k)}$ of the $MNL(n, \theta)$ model. Correspondingly, we assume that the learner can play sets of k distinct items at each round $t \in [T]$. The instantaneous regret of the learner, in this case, in the t -th round is defined to be $r_t^k = \left(\frac{\theta_{S_{(k)}} - \sum_{i \in S_t} \theta_i}{k} \right)$, where $\theta_{S_{(k)}} = \sum_{i \in S_{(k)}} \theta_i$. Consequently, the cumulative

regret of the learner at the end of round T becomes $R_T^k = \sum_{t=1}^T r_t^k = \sum_{t=1}^T \left(\frac{\theta_{S_{(k)}} - \sum_{i \in S_t} \theta_i}{k} \right)$. As with the Winner-regret, the Top- k -regret also admits a natural interpretation as the shortfall in *value* of the set S_t with respect to the set $S_{(k)}$, with value of a set being the mean θ parameter of its arms.

3 Minimising Winner-regret

We first consider the problem of minimising Winner-regret. We start by analysing a regret lower bound for the problem, followed by designing an optimal algorithm with matching upper bound.

3.1 Fundamental lower bound on Winner-regret

Along the lines of [32], we define the following consistency property of any reasonable online learning algorithm in order to state a fundamental lower bound on regret performance.

Definition 2 (*No-regret algorithm*). An online learning algorithm \mathcal{A} is defined to be a *No-regret* algorithm for Winner-regret if for each problem instance $\text{MNL}(n, \theta)$, the expected number of times \mathcal{A} plays any suboptimal set $S \subseteq [n]$ is sublinear in T , i.e., $\forall S \neq \arg \max_i \theta_i : \mathbf{E}_\theta[N_S(T)] = o(T^\alpha)$, for some $\alpha \in [0, 1]$ (potentially depending on θ), where $N_S(T) := \sum_{t=1}^T \mathbf{1}(S_t = S)$ is the number plays of set S in T rounds. $\mathbf{E}_\theta[\cdot]$ denotes expectation under the algorithm and $\text{MNL}(n, \theta)$ model.

Theorem 3 (Winner-regret Lower Bound). *For any No-regret learning algorithm \mathcal{A} for Winner-regret that uses Winner Feedback, and for any problem instance $\text{MNL}(n, \theta)$ s.t. $a^* = \arg \max_{i \in [n]} \theta_i$,*

$$\text{the expected regret incurred by } \mathcal{A} \text{ satisfies } \liminf_{T \rightarrow \infty} \mathbf{E}_\theta \left[\frac{R_T^1(\mathcal{A})}{\ln T} \right] \geq \frac{\theta_{a^*}}{\left(\min_{i \in [n] \setminus \{a^*\}} \frac{\theta_{a^*}}{\theta_i} - 1 \right)} (n-1).$$

Note: This is a problem-dependent lower bound with $\theta_{a^*} \left(\min_{i \in [n] \setminus \{a^*\}} \frac{\theta_{a^*}}{\theta_i} - 1 \right)^{-1}$ denoting a complexity or hardness term ('gap') for regret performance under any 'reasonable learning algorithm'.

Remark 3. The result suggests the regret rate with only Winner Feedback cannot improve with k , uniformly across all problem instances. Rather strikingly, there is no reduction in hardness (measured in terms of regret rate) in learning the *Best-Item* using Winner Feedback from large (k -size) subsets as compared to using pairwise (dueling) feedback ($k = 2$). It could be tempting to expect an improved learning rate with subset-wise feedback as the number of items being tested per iteration is more ($k \geq 2$), so information-theoretically one may expect to 'learn more' about the underlying model per subset query. On the contrary, it turns out that it is intuitively 'harder' for a good (i.e., near-optimal) item to prove its competitiveness in just a single winner draw against a large population of its $k - 1$ other competitors, as compared to winning over just a single competitor for $k = 2$ case.

Proof sketch. The proof of the result is based on the change of measure technique for bandit regret lower bounds presented by, say, Garivier et al. [21], that uses the information divergence between two nearby instances $\text{MNL}(n, \theta)$ (the original instance) and $\text{MNL}(n, \theta')$ (an alternative instance) to quantify the hardness of learning the best arm in either environment. In our case, each bandit instance corresponds to an instance of the $\text{MNL}(n, \theta)$ problem with the arm set containing all subsets of $[n]$ of size upto k : $\mathcal{A} = \{S \subseteq [n] \mid |S| \in [k]\}$. The key of the proof relies on carefully crafting a true instance, with optimal arm $a^* = 1$, and a family of 'slightly perturbed' alternative instances $\{\nu^a : a \neq 1\}$, each with optimal arm $a \neq 1$, chosen as: **(1). True Instance:** $\text{MNL}(n, \theta^1) : \theta_1^1 > \theta_2^1 = \dots = \theta_n^1 = \theta$ (for some $\theta \in \mathbb{R}_+$), and for each suboptimal item $a \in [n] \setminus \{1\}$, the **(2). Altered instances:** $\text{MNL}(n, \theta^a) : \theta_a^a = \theta_1^1 + \epsilon = \theta + (\Lambda + \epsilon)$; $\theta_i^a = \theta_i^1, \forall i \in [n] \setminus \{a\}$ for some $\epsilon > 0$. The result of Thm. 3 now follows by applying Lemma 13 on pairs of problem instances $(\nu, \nu^{(a)})$ with suitable upper bounds on the divergences. (Complete proof given in Appendix C.3). \square

Note: We also show an alternate version of the regret lower bound of $\Omega\left(\frac{n}{\left(\min_{i \in [n] \setminus \{a^*\}} p_{a^*, i} - 0.5\right)} \ln T\right)$ in terms of pairwise preference-based instance complexities (details are moved to Appendix C.4).

Improved regret lower bound with Top- m -ranking Feedback. In contrast to the situation with only winner feedback, the following (more general) result shows a reduced lower bound when Top- m -ranking Feedback is available in each play of a subset, opening up the possibility of improved learning (regret) performance when ranked-order feedback is available.

Theorem 4 (Regret Lower Bound: Winner-regret with Top- m -ranking Feedback). *For any No-regret algorithm \mathcal{A} for the Winner-regret problem with Top- m -ranking Feedback, there exists a problem instance $\text{MNL}(n, \theta)$ such that the expected Winner-regret incurred by \mathcal{A} satisfies*

$$\liminf_{T \rightarrow \infty} \mathbf{E}_\theta \left[\frac{R_T^1(\mathcal{A})}{\ln T} \right] \geq \frac{\theta_{a^*}}{\left(\min_{i \in [n] \setminus \{a^*\}} \frac{\theta_{a^*}}{\theta_i} - 1 \right)} \frac{(n-1)}{m},$$
where as in Thm. 3, $\mathbf{E}_\theta[\cdot]$ denotes expectation under the algorithm and the MNL model $\text{MNL}(n, \theta)$, and recall $a^ := \arg \max_{i \in [n]} \theta_i$.*

Proof sketch. The main observation made here is that the KL divergences for Top- m -ranking Feedback are m times compared to the case of Winner Feedback, which we show using chain rule for KL divergences [20]: $KL(p_S^1, p_S^a) = KL(p_S^1(\sigma_1), p_S^a(\sigma_1)) + \sum_{i=2}^m KL(p_S^1(\sigma_i \mid \sigma(1:i-1)), p_S^a(\sigma_i \mid \sigma(1:i-1)))$, where $\sigma_i = \sigma(i)$ and $KL(P(Y \mid X), Q(Y \mid X)) := \sum_x Pr(X = x) [KL(P(Y \mid X = x), Q(Y \mid X = x))]$ denotes the conditional KL-divergence. Using this, along

with the upper bound on KL divergences for Winner Feedback (derived for Thm. 3), we show that $KL(p_S^1, p_S^a) \leq \frac{m\Delta_a'^2}{\theta_S^2(\theta_1^2 + \epsilon)}$, $\forall a \in [n] \setminus \{1\}$ (where $\theta_S = \sum_{i \in S} \theta_i$ and $\Delta'_a = O(\theta_{a^*} - \theta_a)$), which precisely gives the $\frac{1}{m}$ -factor reduction in the lower bound compared to Winner Feedback. The bound now can be derived following a similar technique used for Thm. 3 (details in Appendix C.5). \square .

Remark 4. Thm. 4 shows a $\Omega\left(\frac{n \ln T}{m}\right)$ lower bound on regret, containing the instance-dependent constant term $\frac{\theta_{a^*}}{\left(\min_{i \in [n] \setminus \{a^*\}} \frac{\theta_{a^*}}{\theta_i} - 1\right)}$ which exposes the hardness of the regret minimisation problem in terms

of the ‘gap’ between the best a^* and the second best item $\min_{i \in [n] \setminus \{a^*\}} \theta_i$: $(\theta_{a^*} - \max_{i \in [n] \setminus \{a^*\}} \theta_i)$. The $\frac{1}{m}$ factor improvement in learning rate with Top- m -ranking Feedback can be intuitively interpreted as follows: revealing an m -ranking of a k -set is worth about $\ln\left(\binom{k}{m} m!\right) = O(m \ln k)$ bits of information, which is about m times as large compared to revealing a single winner.

3.2 An order-optimal algorithm for Winner-regret

We here show that above fundamental lower bounds on Winner-regret are, in fact, achievable with carefully designed online learning algorithms. We design an upper-confidence bound (UCB)-based algorithm for Winner-regret with Top- m -ranking Feedback model based on the following ideas:

(1). **Playing sets of only $(m+1)$ sizes:** It is enough for the algorithm to play subsets of size either $(m+1)$ (to fully exploit the Top- m -ranking Feedback) or 1 (singleton sets), and not play a singleton unless there is a high degree of confidence about the single item being the best item.

(2). **Parameter estimation from pairwise preferences:** It is possible to play the subset-wise game just by maintaining pairwise preference estimates of all n items of the $\text{MNL}(n, \theta)$ model using the idea of *Rank-Breaking*—the idea of extracting pairwise comparisons from (partial) rankings and applying estimators on the obtained pairs treating each comparison independently over the received subset-wise feedback—this is possible owing to the independence of irrelevant attributes (IIA) property of the MNL model (Defn. 10).

(3). **A new UCB-based max-min set building rule for playing large sets (*build_S*):** Main novelty of *MaxMin-UCB* lies in its underlying set building subroutine (Alg. 2, Appendix C.1), that constructs S_t by applying a recursive max-min strategy on the UCB estimates of empirical pairwise preferences.

Algorithm description. *MaxMin-UCB* maintains an pairwise preference matrix $\hat{\mathbf{P}} \in [0, 1]^{n \times n}$, whose (i, j) -th entry \hat{p}_{ij} records the empirical probability of i having beaten j in a pairwise duel, and a corresponding upper confidence bound u_{ij} for each pair (i, j) . At any round t , it plays a subset $S_t \subseteq [n]$, $|S_t| \in [k]$ using the *Max-Min* set building rule *build_S* (see Alg. 2), receives Top- m -ranking Feedback $\sigma_t \in \Sigma_{S_t}^m$ from S_t , and updates the \hat{p}_{ij} entries of pairs in S_t by applying *Rank-Breaking* (Line 10). The set building rule *build_S* is at the heart of *MaxMin-UCB* which builds the subset S_t from a set of potential Condorcet winners (\mathcal{C}_t) of round t : By recursively picking the strongest opponents of the already selected items using a max-min selection strategy on u_{ij} . The complete algorithm is presented in Alg. 1, Appendix C.1.

The following result establishes that *MaxMin-UCB* enjoys $O\left(\frac{n}{m} \ln T\right)$ regret with high probability.

Theorem 5 (*MaxMin-UCB: High Probability Regret bound*). *Fix a time horizon T and $\delta \in (0, 1)$, $\alpha > \frac{1}{2}$. With probability at least $(1 - \delta)$, the regret of *MaxMin-UCB* for Winner-regret with Top- m -*

ranking Feedback satisfies $R_T^1 \leq \left(2 \left[\frac{2\alpha n^2}{(2\alpha-1)\delta}\right]^{\frac{1}{2\alpha-1}} + 2D \ln 2D\right) \hat{\Delta}_{\max} + \frac{\ln T}{m+1} \sum_{i=2}^n (D_{\max} \hat{\Delta}_i)$,

where $\forall i \in [n] \setminus \{a^\}$, $\hat{\Delta}_i = (\theta_{a^*} - \theta_i)$, $\Delta_i = \frac{\theta_{a^*} - \theta_i}{2(\theta_{a^*} + \theta_i)}$, $\hat{\Delta}_{\max} = \max_{i \in [n] \setminus \{a^*\}} \hat{\Delta}_i$, $D_{1i} = \frac{4\alpha}{\Delta_i^2}$, $D := \sum_{i < j} D_{ij}$, $D_{\max} = \max_{i \in [n] \setminus \{a^*\}} D_{1i}$.*

Proof sketch. The proof hinges on analysing the entire run of *MaxMin-UCB* by breaking it up into 3 phases: (1). **Random-Exploration** (2). **Progress**, and (3). **Saturation**.

(1). **Random-Exploration:** This phase runs from time 1 to $f(\delta) = \left[\frac{2\alpha n^2}{(2\alpha-1)\delta}\right]^{\frac{1}{2\alpha-1}}$, for any $\delta \in (0, 1)$, such that for any $t > f(\delta)$, the upper confidence bounds u_{ij} are guaranteed to be correct for the true values p_{ij} for all pairs $(i, j) \in [n] \times [n]$ (i.e. $p_{ij} \leq u_{ij}$), with high probability $(1 - \delta)$.

(2). **Progress:** After $t > f(\delta)$, the algorithm can be viewed as starting to explore the ‘confusing items’, appearing in \mathcal{C}_t , as potential candidates for the *Best-Item* a^* , and trying to capture a^* in a holding set \mathcal{B}_t . At any time, the set \mathcal{B}_t is either empty or a singleton by construction, and once

$a^* \in \mathcal{B}_t$ it stays there forever (with high probability). The **Progress** phase ensures that the algorithm explores fast enough so that within a constant number of rounds, \mathcal{B}_t captures $\{a^*\}$.

(3). Saturation: This is the last phase from time $T_0(\delta) + 1$ to T . As the name suggests, *MaxMin-UCB* shows relatively stable behavior here, mostly playing $S_t = \{a^*\}$ and incurring almost no regret.

Although Thm. 5 shows a $(1 - \delta)$ -high probability regret bound for *MaxMin-UCB* it is important to note that the algorithm itself does not require to take the probability of failure (δ) as input. As a consequence, by simply integrating the bound obtained in Thm. 5 over the entire range of $\delta \in [0, 1]$, we get an expected regret bound of *MaxMin-UCB* for Winner-regret with Top- m -ranking Feedback:

Theorem 6. *The expected regret of MaxMin-UCB for Winner-regret with Top- m -ranking Feedback*

is: $\mathbf{E}[R_T^1] \leq \left(2 \left[\frac{2\alpha n^2}{(2\alpha-1)} \right]^{\frac{1}{2\alpha-1}} \frac{2\alpha-1}{\alpha-1} + 2D \ln 2D \right) \hat{\Delta}_{\max} + \frac{\ln T}{m+1} \sum_{i=2}^n (D_{\max} \hat{\Delta}_i)$, in T rounds.

Remark 5. This is an upper bound on expected regret of the same order as that in the lower bound of Thm. 3, which shows that the algorithm is essentially regret-optimal. From Thm. 6, note that the

first two terms $\left(2 \left[\frac{2\alpha n^2}{(2\alpha-1)} \right]^{\frac{1}{2\alpha-1}} \frac{2\alpha-1}{\alpha-1} + 2D \ln 2D \right) \hat{\Delta}_{\max}$ of $\mathbf{E}[R_T^1]$ are essentially instance specific

constants, its only the third term which makes expected regret $O\left(\frac{n \ln T}{m}\right)$ which is in fact optimal in

terms of its dependencies on n and T (since it matches the $\Omega\left(\frac{n \ln T}{m}\right)$ lower bound of Thm. 4). More-

over the problem dependent complexity terms $(D_{\max} \hat{\Delta}_i) = \frac{16\alpha(\theta_{a^*} - \theta_i)(\theta_{a^*} + \max_{j \in [n] \setminus \{a^*\}} \theta_j)^2}{(\theta_{a^*} - \max_{j \in [n] \setminus \{a^*\}} \theta_j)^2} \leq \frac{64\alpha(\theta_{a^*} - \theta_i)(\theta_{a^*})}{(\theta_{a^*} - \max_{j \in [n] \setminus \{a^*\}} \theta_j)^2} = O\left(\frac{\theta_{a^*}}{(\theta_{a^*} - \max_{j \in [n] \setminus \{a^*\}} \theta_j)}\right)$, also brings out the inverse dependency on the ‘gap-term’ $(\theta_{a^*} - \max_{j \in [n] \setminus \{a^*\}} \theta_j)$ as discussed in Rem. 4.

4 Minimising Top- k -regret

In this section, we study the problem of minimising Top- k -regret with Top- k -ranking Feedback. As before, we first derive a regret lower bound, for this learning setting, of the form $\Omega\left(\frac{n-k}{k\Delta_{(k)}} \ln T\right)$ (recall $\Delta_{(k)}$ from Sec. 2). We next propose an UCB based algorithm (Alg. 3) for the same, along with a matching upper bound regret analysis (Thm. 8,9) showing optimality of our proposed algorithm.

4.1 Regret lower bound for Top- k -regret with Top- k -ranking Feedback

Theorem 7 (Regret Lower Bound: Top- k -regret with Top- m -ranking Feedback). *For any No-regret learning algorithm \mathcal{A} for Top- k -regret that uses Top- k -ranking Feedback, and for any problem instance $\text{MNL}(n, \theta)$, the expected regret incurred by \mathcal{A} when run on it satisfies $\liminf_{T \rightarrow \infty} \mathbf{E}_{\theta} \left[\frac{R_T^k(\mathcal{A})}{\ln T} \right] \geq \frac{\theta_1 \theta_{k+1}}{\Delta_{(k)}} \frac{(n-k)}{k}$, where $\mathbf{E}_{\theta}[\cdot]$ denotes expectation under the algorithm and $\text{MNL}(n, \theta)$ model.*

Proof sketch. Similar to 4, the proof again relies on carefully constructing a true instance, with optimal set of Top- k Best-Items $S_{(k)} = [k]$, and a family of slightly perturbed alternative instances $\{\nu^a : a \in [n] \setminus S_{(k)}\}$, for each suboptimal arm $a \in [n] \setminus S_{(k)}\}$, which we design as:

(1). True Instance: $\text{MNL}(n, \theta^1) : \theta_1^1 = \theta_2^1 = \dots = \theta_{k-1}^1 = \theta + 2\epsilon; \theta_n^1 = \theta + \epsilon; \theta_{k+1}^1 = \theta_{k+2}^1 = \dots = \theta_{n-1}^1 = \theta$, for some $\theta \in \mathbb{R}_+$ and $\epsilon > 0$. Clearly Top- k Best-Items of $\text{MNL}(n, \theta^1)$ is $S_{(k)}[1] = [k-1] \cup \{n\}$. **(2). Altered Instances:** For every $n-k$ suboptimal items $a \notin S_{(k)}[1]$, now consider an altered instance **Instance a**, denoted by $\text{MNL}(n, \theta^a)$, such that $\theta_a^a = \theta + 2\epsilon; \theta_i^a = \theta_i^1, \forall i \in [n] \setminus \{a\}$. The result of Thm. 7 now can be obtained by following an exactly same procedure as described for the proof of Thm. 4. The complete details is given in Appendix D.1. \square

Remark 6. The regret lower bound of Thm. 7 is $\Omega\left(\frac{(n-k) \ln T}{k}\right)$, with an instance-dependent term $\frac{\theta_1 \theta_{k+1}}{(\theta_k - \theta_{k+1})}$ which shows for recovering the Top- k Best-Items, the problem complexity is governed by the ‘gap’ between the k^{th} and $(k+1)^{\text{th}}$ best item $\Delta_{(k)} = (\theta_k - \theta_{k+1})$, as consistent with intuition.

4.2 An order-optimal algorithm with low Top- k -regret with Top- k -ranking Feedback

Main idea: A recursive set-building rule: As with the *MaxMin-UCB* algorithm (Alg. 1), we maintain pairwise UCB estimates (u_{ij}) of empirical pairwise preferences \hat{p}_{ij} via *Rank-Breaking*. However the main difference here lies in the set building rule, as here it is required to play sets of

size exactly k . The core idea here is to recursively try to capture the set of *Top-k Best-Items* in an ordered set \mathcal{B}_t , and, once the set is assumed to be found with confidence (formally $|\mathcal{B}_t| = k$), to keep playing \mathcal{B}_t unless some other potential good item emerges, which is then played replacing the weakest element ($\mathcal{B}_t(k)$) of \mathcal{B}_t . The algorithm is described in Alg. 3, Appendix D.2.

Theorem 8 (*Rec-MaxMin-UCB: High Probability Regret bound*). *Given a fixed time horizon T and $\delta \in (0, 1)$, with high probability $(1 - \delta)$, the regret incurred by Rec-MaxMin-UCB for Top-k-regret admits the bound $R_T^k \leq \left(2 \left[\frac{2\alpha n^2}{(2\alpha-1)\delta}\right]^{\frac{1}{2\alpha-1}} + 2\bar{D}^{(k)} \ln(2\bar{D}^{(k)})\right) \Delta'_{\max} + \frac{4\alpha \ln T}{k} \left(\sum_{b=k+1}^n \frac{(\theta_k - \theta_b)}{\bar{D}^2}\right)$, where $D^{(k)}$ is an instance dependent constant (see Lem. 26, Appendix), $\Delta'_{\max} = \frac{(\sum_{i=1}^k \theta_i - \sum_{i=n-k+1}^n \theta_i)}{k}$, and $\hat{D} = \min_{g \in [k-1]} (p_{kg} - p_{bg})$.*

Proof sketch. Similar to Thm. 5, we prove the above bound dividing the entire run of algorithm *Rec-MaxMin-UCB* into three phases and applying an recursive argument:

(1). **Random-Exploration:** Same as Thm. 5, in this case also this phase runs from time 1 to $f(\delta) = \left\lceil \frac{2\alpha n^2}{(2\alpha-1)\delta} \right\rceil^{\frac{1}{2\alpha-1}}$, for any $\delta \in (0, 1)$, after which, for any $t > f(\delta)$, one can guarantee $p_{ij} \leq u_{ij}$ for all pairs $(i, j) \in [n] \times [n]$, with high probability at least $(1 - \delta)$. (Lem. 15)

(2). **Progress:** The analysis of this phase is quite different from that of Thm. 5: After $t > f(\delta)$, the algorithm starts exploring the items in the set of *Top-k Best-Items* in a *recursive* manner—It first tries to capture (one of) the *Best-Items* in $\mathcal{B}_t(1)$. Once that slot is secured, it goes on for searching the second *Best-Item* from remaining pool of items and try capturing it in $\mathcal{B}_t(2)$ and so on upto $\mathcal{B}_t(k)$. By definition, the phase ends at, say $t = T_0(\delta)$, when $\mathcal{B}_t = S_{(k)}$. Moreover the update rule of *Rec-MaxMin-UCB* (along with Lem. 15) ensures that $\mathcal{B}_t = S_{(k)} \forall t > T_0(\delta)$. The novelty of our analysis lies in showing that $T_0(\delta)$ is bounded by just a instance dependent complexity term which does not scale with t (Lem. 26), and hence the regret incurred in this phase is also constant.

(3). **Saturation:** In the last phase from time $T_0(\delta) + 1$ to T *Rec-MaxMin-UCB* has already captured $S_{(k)}$ in \mathcal{B}_t , and $\mathcal{B}_t = S_{(k)}$ henceforth. Hence the algorithm mostly plays $S_t = S_{(k)}$ without incurring any regret. Only if any item outside \mathcal{B}_t enters into the list of potential *Top-k Best-Items*, it takes a very conservative approach of replacing the ‘weakest of \mathcal{B}_t by that element to make sure whether it indeed lies in or outside $S_{(k)}$. However we are able to show that any such suboptimal item $i \notin S_{(k)}$ can not occur for more than $O(\frac{\ln T}{\bar{D}^2})$ times (Lem. 27), combining which over all $[n] \setminus [k]$ suboptimal items finally leads to the desired regret. *The complete details are moved to Appendix D.3.* \square

From Theorem 8, we can also derive an expected regret bound for *Rec-MaxMin-UCB* in T rounds is:

Theorem 9. *The expected regret incurred by MaxMin-UCB for Top-k-regret is:*

$$\mathbf{E}[R_T^1] \leq \left(2 \left[\frac{2\alpha n^2}{(2\alpha-1)}\right]^{\frac{1}{2\alpha-1}} \frac{2\alpha-1}{\alpha-1} + 2\bar{D}^{(k)} \ln(2\bar{D}^{(k)})\right) \Delta'_{\max} + \frac{4\alpha \ln T}{k} \left(\sum_{b=k+1}^n \frac{(\theta_k - \theta_b)}{\bar{D}^2}\right).$$

Remark 7. In Thm. 9, the first two terms $\left(2 \left[\frac{2\alpha n^2}{(2\alpha-1)\delta}\right]^{\frac{1}{2\alpha-1}} + 2\bar{D}^{(k)} \ln(2\bar{D}^{(k)})\right) \Delta'_{\max}$ of $\mathbf{E}[R_T^k]$ are just some $\text{MNL}(n, \theta)$ model dependent constants which do not contribute to the learning rate of *Rec-MaxMin-UCB*, and the third term is $O\left(\frac{(n-k) \ln T}{k}\right)$ which varies optimally in terms of on n, k, T matching the $\Omega\left(\frac{(n-k) \ln T}{k}\right)$ lower bound of Thm. 7). Also Rem. 6 indicates an inverse dependency on the ‘gap-complexity’ $(\theta_k - \theta_{k+1})$, which also shows up in above bound through the component $\frac{(\theta_k - \theta_b)}{\bar{D}^2}$: Let $g^* \in [k-1]$ is the minimizer of \hat{D} , then $\frac{(\theta_k - \theta_b)}{\bar{D}^2} = \frac{(\theta_{g^*} + \theta_k)(\theta_b + \theta_{g^*})}{\theta_{g^*}^2 (\theta_k - \theta_b)} \leq \frac{4}{\theta_{g^*} (\theta_k - \theta_{k+1})}$, where the upper bounding follows as $\theta_{g^*} \geq \theta_k > \theta_b$ for any $b \in [n] \setminus [k]$, and $\theta_b \leq \theta_{k+1}$ for any b .

5 Experiments

In this section we present the empirical evaluations of our proposed algorithm *MaxMin-UCB* (abbreviated as **MM**) on different synthetic datasets, and also compare them with different algorithms. All results are reported as average across 50 runs along with the standard deviations. For this we use 7 different $\text{MNL}(n, \theta)$ environments as described below:

MNL(n, θ) Environments. 1. *g1*, 2. *g4*, 3. *arith*, 4. *geo*, 5. *har* all with $n = 16$, and two larger models 6. *arithb*, and 7. *geob* with $n = 50$ items in both. Details are moved to Appendix E.

We compare our proposed methods with the following two baselines which closely applies to our problem setup. Note, as discussed in Sec. 1, none of the existing work exactly addresses our problem.

Algorithms. 1. **BD:** The *Battling-Duel* algorithm of [36] with *RUCB* algorithm [47] as the dueling bandit blackbox, and 2. **Sp-TS:** The *Self-Sparring* algorithm of [39] with Thompson Sampling [1], and 3. **MM:** Our proposed method *MaxMin-UCB* for Winner-regret (Alg. 1).

Comparing Winner-regret with Top- m -ranking Feedback (Fig. 1): We first compare the regret performances for $k = 10$ and $m = 5$. From Fig. 1, it clearly follows that in all cases *MaxMin-UCB* uniformly outperforms the other two algorithms taking the advantage of Top- m -ranking Feedback which the other two fail to make use of as they both allow repetitions in the played subsets which can not exploit the rank-ordered feedback to the full extent. Furthermore, the thompson sampling based *Sp-TS* in general exhibits a much higher variance compared to the rest due to its bayesian nature. Also as expected, *g1* and *g4* being comparatively easier instances, i.e. with larger ‘gap’ $\hat{\Delta}_{\max}$ (see Thm. 3, 4, 5, 6 etc. for a formal justification), our algorithm converges much faster on these models.




Figure 1: Comparative performances on Winner-regret for $k = 10$, $m = 5$

Comparing Top- k -regret performances for Top- k -ranking Feedback (Fig. 2): We are not aware of any existing algorithm for Top- k -regret objective with Top- k -ranking Feedback. We thus use a modified version of *Sp-TS* algorithm [39] described above for the purpose—it simply draws k -items without repetition and uses *Rank-Breaking* updates to maintain the Beta posteriors. Here again, we see that our method *Rec-MaxMin-UCB* (*Rec-MM*) uniformly outperforms *Sp-TS* in all cases, and as before *Sp-TS* shows a higher variability as well. Interestingly, our algorithm converges the fastest on *g4*, it being the easiest model with largest ‘gap’ $\Delta_{(k)}$ between the k^{th} and $(k+1)^{th}$ best item (see Thm. 7, 8, 9 etc.), and takes longest time for *har* since it has the smallest $\Delta_{(k)}$.




Figure 2: Comparative performances on Top- k -regret for $k = 10$

Effect of varying m with fixed k (Fig. 3): We also studied our algorithm *MaxMin-UCB*, with varying size rank-ordered feedback (m), keeping the subset size (k) fixed, both for Winner-regret and Top- k -regret objective, on the larger models *arithb* and *geob* which has $n = 50$ items. As expected, in both cases, regret scales down with increasing m (justifying the bounds in Thm. 5, 6, 8, 9).




Figure 3: Regret with varying m with fixed $k = 40$ (on our proposed algorithm *MaxMin-UCB*)

6 Conclusion and Future Work

Although we have analysed low-regret algorithms for learning with subset-wise preferences, there are several avenues for investigation that open up with these results. The case of learning with contextual subset-wise models is an important and practically relevant problem, as is the problem of considering mixed cardinal and ordinal feedback structures in online learning. Other directions of interest could be studying the budgeted version where there are costs associated with the amount of preference information that may be elicited in each round, or analysing the current problem on a variety of subset choice models, e.g. multinomial probit, Mallows, or even adversarial preference models etc.

Acknowledgements

The authors are grateful to the anonymous reviewers for valuable feedback. This work is supported by a Qualcomm Innovation Fellowship 2019, and the Indigenous 5G Test Bed project grant from the Dept. of Telecommunications, Government of India. Aadirupa Saha thanks Arun Rajkumar for the valuable discussions, and the Tata Trusts and ACM-India/IARCS Travel Grants for travel support.

References

- [1] Shipra Agrawal and Navin Goyal. Analysis of Thompson sampling for the multi-armed bandit problem. In *Conference on Learning Theory*, pages 39–1, 2012.
- [2] Shipra Agrawal, Vashist Avadhanula, Vineet Goyal, and Assaf Zeevi. A near-optimal exploration-exploitation approach for assortment selection. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 599–600. ACM, 2016.
- [3] Shipra Agrawal, Vashist Avadhanula, Vineet Goyal, and Assaf Zeevi. Thompson sampling for the mnl-bandit. *Machine Learning Research*, 65:1–3, 2017.
- [4] 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.
- [5] Nir Ailon, Zohar Shay Karnin, and Thorsten Joachims. Reducing dueling bandits to cardinal bandits. In *ICML*, volume 32, pages 856–864, 2014.
- [6] Duane F Alwin and Jon A Krosnick. The measurement of values in surveys: A comparison of ratings and rankings. *Public Opinion Quarterly*, 49(4):535–552, 1985.
- [7] Hossein Azari, David Parkes, and Lirong Xia. Random utility theory for social choice. In *Advances in Neural Information Processing Systems*, pages 126–134, 2012.
- [8] Gábor Bartók, Dávid Pál, and Csaba Szepesvári. Minimax regret of finite partial-monitoring games in stochastic environments. In *Proceedings of the 24th Annual Conference on Learning Theory*, pages 133–154, 2011.
- [9] Moshe Ben-Akiva, Mark Bradley, Takayuki Morikawa, Julian Benjamin, Thomas Novak, Harmen Oppewal, and Vithala Rao. Combining revealed and stated preferences data. *Marketing Letters*, 5(4):335–349, 1994.
- [10] Austin R Benson, Ravi Kumar, and Andrew Tomkins. On the relevance of irrelevant alternatives. In *Proceedings of the 25th International Conference on World Wide Web*, pages 963–973. International World Wide Web Conferences Steering Committee, 2016.
- [11] Brian Brost, Yevgeny Seldin, Ingemar J. Cox, and Christina Lioma. Multi-dueling bandits and their application to online ranker evaluation. *CoRR*, abs/1608.06253, 2016.
- [12] Róbert Busa-Fekete and Eyke Hüllermeier. A survey of preference-based online learning with bandit algorithms. In *International Conference on Algorithmic Learning Theory*, pages 18–39. Springer, 2014.
- [13] Róbert Busa-Fekete, Eyke Hüllermeier, and Balázs Szörényi. Preference-based rank elicitation using statistical models: The case of mallows. In *Proceedings of The 31st International Conference on Machine Learning*, volume 32, 2014.
- [14] Nicolo Cesa-Bianchi and Gábor Lugosi. Combinatorial bandits. *Journal of Computer and System Sciences*, 78(5):1404–1422, 2012.
- [15] Wei Chen, Yajun Wang, and Yang Yuan. Combinatorial multi-armed bandit: General framework and applications. In *International Conference on Machine Learning*, pages 151–159, 2013.
- [16] Xi Chen, Paul N Bennett, Kevyn Collins-Thompson, and Eric Horvitz. Pairwise ranking aggregation in a crowdsourced setting. In *Proceedings of the sixth ACM international conference on Web search and data mining*, pages 193–202. ACM, 2013.

- [17] Xi Chen, Yuanzhi Li, and Jieming Mao. A nearly instance optimal algorithm for top-k ranking under the multinomial logit model. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2504–2522. SIAM, 2018.
- [18] Yuxin Chen and Changho Suh. Spectral mle: Top-k rank aggregation from pairwise comparisons. In *International Conference on Machine Learning*, pages 371–380, 2015.
- [19] Richard Combes, Mohammad Sadegh Talebi Mazraeh Shahi, Alexandre Proutiere, et al. Combinatorial bandits revisited. In *Advances in Neural Information Processing Systems*, pages 2116–2124, 2015.
- [20] Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- [21] Aurélien Garivier, Pierre Ménard, and Gilles Stoltz. Explore first, exploit next: The true shape of regret in bandit problems. *Mathematics of Operations Research*, 44(2):377–399, 2018.
- [22] Thore Graepel and Ralf Herbrich. Ranking and matchmaking. *Game Developer Magazine*, 25: 34, 2006.
- [23] Bruce Hajek, Sewoong Oh, and Jiaming Xu. Minimax-optimal inference from partial rankings. In *Advances in Neural Information Processing Systems*, pages 1475–1483, 2014.
- [24] David A Hensher. Stated preference analysis of travel choices: the state of practice. *Transportation*, 21(2):107–133, 1994.
- [25] Katja Hofmann. Fast and reliable online learning to rank for information retrieval. In *SIGIR Forum*, volume 47, page 140, 2013.
- [26] Minje Jang, Sunghyun Kim, Changho Suh, and Sewoong Oh. Optimal sample complexity of m-wise data for top-k ranking. In *Advances in Neural Information Processing Systems*, pages 1685–1695, 2017.
- [27] Sumeet Katariya, Branislav Kveton, Csaba Szepesvari, and Zheng Wen. Dcm bandits: Learning to rank with multiple clicks. In *International Conference on Machine Learning*, pages 1215–1224, 2016.
- [28] Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best-arm identification in multi-armed bandit models. *The Journal of Machine Learning Research*, 17(1): 1–42, 2016.
- [29] Ashish Khetan and Sewoong Oh. Data-driven rank breaking for efficient rank aggregation. *Journal of Machine Learning Research*, 17(193):1–54, 2016.
- [30] Junpei Komiyama, Junya Honda, Hisashi Kashima, and Hiroshi Nakagawa. Regret lower bound and optimal algorithm in dueling bandit problem. In *COLT*, pages 1141–1154, 2015.
- [31] Branislav Kveton, Zheng Wen, Azin Ashkan, and Csaba Szepesvari. Tight regret bounds for stochastic combinatorial semi-bandits. In *Artificial Intelligence and Statistics*, pages 535–543, 2015.
- [32] Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- [33] Pantelimon G Popescu, Silvestru Dragomir, Emil I Slusanschi, and Octavian N Stanasila. Bounds for Kullback-Leibler divergence. *Electronic Journal of Differential Equations*, 2016, 2016.
- [34] Filip Radlinski, Madhu Kurup, and Thorsten Joachims. How does clickthrough data reflect retrieval quality? In *Proceedings of the 17th ACM conference on Information and knowledge management*, pages 43–52. ACM, 2008.
- [35] Wenbo Ren, Jia Liu, and Ness B Shroff. PAC ranking from pairwise and listwise queries: Lower bounds and upper bounds. *arXiv preprint arXiv:1806.02970*, 2018.

- [36] Aadirupa Saha and Aditya Gopalan. Battle of bandits. In *Uncertainty in Artificial Intelligence*, 2018.
- [37] Aadirupa Saha and Aditya Gopalan. PAC Battling Bandits in the Plackett-Luce Model. In *Algorithmic Learning Theory*, pages 700–737, 2019.
- [38] Hossein Azari Soufiani, David C Parkes, and Lirong Xia. Computing parametric ranking models via rank-breaking. In *ICML*, pages 360–368, 2014.
- [39] Yanan Sui, Vincent Zhuang, Joel Burdick, and Yisong Yue. Multi-dueling bandits with dependent arms. In *Conference on Uncertainty in Artificial Intelligence*, UAI’17, 2017.
- [40] Balázs Szörényi, Róbert Busa-Fekete, Adil Paul, and Eyke Hüllermeier. Online rank elicitation for plackett-luce: A dueling bandits approach. In *Advances in Neural Information Processing Systems*, pages 604–612, 2015.
- [41] Tanguy Urvoy, Fabrice Clerot, Raphael Féraud, and Sami Naamane. Generic exploration and k -armed voting bandits. In *International Conference on Machine Learning*, pages 91–99, 2013.
- [42] Huasen Wu and Xin Liu. Double Thompson sampling for dueling bandits. In *Advances in Neural Information Processing Systems*, pages 649–657, 2016.
- [43] Yisong Yue and Thorsten Joachims. Interactively optimizing information retrieval systems as a dueling bandits problem. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 1201–1208. ACM, 2009.
- [44] Yisong Yue and Thorsten Joachims. Beat the mean bandit. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 241–248, 2011.
- [45] 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.
- [46] Masrour Zoghi, Shimon Whiteson, Remi Munos, and Maarten de Rijke. Relative upper confidence bound for the k -armed dueling bandit problem. *arXiv preprint arXiv:1312.3393*, 2013.
- [47] Masrour Zoghi, Shimon Whiteson, Remi Munos, Maarten de Rijke, et al. Relative upper confidence bound for the k -armed dueling bandit problem. In *JMLR Workshop and Conference Proceedings*, number 32, pages 10–18. JMLR, 2014.

Supplementary: Combinatorial Bandits with Relative Feedback

A Related Works

Over the last decade, online learning from pairwise preferences has seen a widespread resurgence in the form of the Dueling Bandit problem, from the points of view of both pure-exploration (PAC) settings [44, 40, 13, 12], and regret minimisation [45, 41, 47, 5, 30, 42]. In contrast, bandit learning with combinatorial, subset-wise preferences, though a natural and practical generalisation, has not received a commensurate treatment.

There have been a few attempts in the batch (i.e., non-adaptive) setting for parameter estimation in utility-based subset choice models, e.g. Plackett-Luce or Thurstonian models [23, 18, 29, 26]. In the online setup, a recent work by Brost et al. [11] considers an extension of the dueling bandits framework where multiple arms are chosen in each round, but they receive comparisons for each pair, and there are no regret guarantees stated for their algorithm. Another similar work is DCM-bandits [27], where a list of k distinct items are offered at each round and the users choose one or more from it scanning the list from top to bottom. However due to this cascading nature of their feedback model, this is also not strictly a relative subset-wise preference model unlike ours, since the utility or attraction weight of an item is assumed to be independently drawn, and so their learning objective differs substantially.

A related body of literature lies in dynamic assortment selection, where the goal is to offer a subset of items to customers in order to maximise expected revenue. A specific, bandit (online) counterpart of this problem has been studied in the recent work of Agrawal et al. [2, 3], although it takes items' prices into account due to which their notion of the 'best subset' is rather different from our 'benchmark subset', and the two settings are incomparable in general. More specifically, in this setting,

1. Their assumption of a *no-purchase* option, say item-0, necessarily present in every set and having the *known* and *highest* MNL parameter value $\theta_0 = 1$, is crucial for their algorithm design as well as the regret analysis — more specifically this helps them to estimate the MNL model parameters easily. We however do not make this assumption, due to which it is more challenging to estimate the MNL model parameters in our case. This is also precisely the reason why the algorithm of Agrawal et al. [2] cannot be directly applied for solving our problem.
2. The regret objective boils down to the top- k best arm identification problem when all item prices are same, say $r_i = 1, \forall i \in [n]$. So in a sense we actually solve a special case of the assortment selection objective – the top k item(s) – but without assumptions on the no-purchase item with known highest parameter value.
3. Agrawal et al. [2] show gap independent $\tilde{O}(\sqrt{NT})$ regret for their algorithm and this is later improved to gap-dependent $O(n^2 \ln T)$ regret [4]; however, the latter guarantee is suboptimal by a factor of n , whereas we show tightness of the regret performance of our proposed algorithms by proving matching lower bound guarantees.

Some recent work addresses the probably approximately correct (PAC) version of the best arm(s) identification problem from subsetwise preferences [17, 35], which is qualitatively different than the optimisation objective considered here. The work which is perhaps closest in spirit to ours is that of Saha and Gopalan [36], but they consider a much more elementary subset choice model based on pairwise preferences, unlike the standard MNL model rooted in choice theory. Sui et al. [39] also address a similar problem; however, a key difference lies in the feedback which consists of outcomes of one or more pairs from the played subset, as opposed to our winner or Top- m -ranking Feedback which is often practical.

Lastly, like the dueling bandit, our more general MNL regret problem can be viewed as a stochastic partial monitoring problem [8], in which the reward or loss of a subset play is not directly observed; instead, only stochastic feedback depending on the subset's parameters is observed. Moreover, under one of the regret structures we consider (Winner-regret, Sec. 3.2), playing the optimal subset (the single item with the highest value) yields no useful information.

B Properties of MNL model

Definition 10 (Independence of Irrelevant Alternatives (IIA) property). A choice model is said to possess the *Independence of Irrelevant Attributes (IIA)* property if the ratio of probabilities of choosing any two items, say i_1 and i_2 from within any choice set $S \ni i_1, i_2$ is independent of a third alternative j present in S [10]. More specifically, $\frac{Pr(i_1|S_1)}{Pr(i_2|S_1)} = \frac{Pr(i_1|S_2)}{Pr(i_2|S_2)}$ for any two distinct subsets $S_1, S_2 \subseteq [n]$ that contain i_1 and i_2 . One such example is the MNL choice model as follows from Defn. 1.

IIA turns out to be very valuable in estimating the parameters of a PL model, with high confidence, via *Rank-Breaking* – the idea of extracting pairwise comparisons from (partial) rankings and applying estimators on the obtained pairs, treating each comparison independently, as described below.

Definition 11 (Rank-Breaking [38, 29]). This is a procedure of deriving pairwise comparisons from multiwise (subsetwise) preference information. Formally, given any set $S \subseteq [n]$, $m \leq |S| < n$, if $\sigma \in \Sigma_S^m$ denotes a possible Top- m -ranking Feedback of S , *Rank-Breaking* considers each item in S to be beaten by its preceding items in σ in a pairwise sense and extracts out total $\sum_{i=1}^m (k-i) = \frac{m(2k-m-1)}{2}$ such pairwise comparisons. For instance, given a full ranking of a set of 4 elements $S = \{a, b, c, d\}$, say $b \succ a \succ c \succ d$, Rank-Breaking generates the set of 6 pairwise comparisons: $\{(b \succ a), (b \succ c), (b \succ d), (a \succ c), (a \succ d), (c \succ d)\}$. Similarly, given the ranking of only 2 most preferred items say $b \succ a$, it yields the 5 pairwise comparisons $(b, a \succ c), (b, a \succ d)$ and $(b \succ a)$ etc. See Line 10 of Algorithm 1 for example.

Owing to the IIA property of $MNL(n, \theta)$ model, one can show the following guarantee on the empirical pairwise estimates $\hat{p}_{ij}(T) = \frac{n_i(T)}{n_{ij}(T)}$ obtained via *Rank-Breaking* on MNL based subsetwise preferences:

Lemma 12 ([37]). *Consider a $MNL(n, \theta)$ model, and fix two distinct items $i, j \in [n]$. Let S_1, \dots, S_T be a sequence of (possibly random) subsets of $[n]$ of size at least 2, where T is a positive integer, and i_1, \dots, i_T a sequence of random items with each $i_t \in S_t$, $1 \leq t \leq T$, such that for each $1 \leq t \leq T$, (a) S_t depends only on S_1, \dots, S_{t-1} , and (b) i_t is distributed as the Plackett-Luce winner of the subset S_t , given $S_1, i_1, \dots, S_{t-1}, i_{t-1}$ and S_t , and (c) $\forall t : \{i, j\} \subseteq S_t$ with probability 1. Let $n_i(T) = \sum_{t=1}^T \mathbf{1}(i_t = i)$ and $n_{ij}(T) = \sum_{t=1}^T \mathbf{1}(\{i_t \in \{i, j\}\})$. Then, for any positive integer v , and $\eta \in (0, 1)$,*

$$Pr\left(\frac{n_i(T)}{n_{ij}(T)} - \frac{\theta_i}{\theta_i + \theta_j} \geq \eta, n_{ij}(T) \geq v\right) \vee Pr\left(\frac{n_i(T)}{n_{ij}(T)} - \frac{\theta_i}{\theta_i + \theta_j} \leq -\eta, n_{ij}(T) \geq v\right) \leq e^{-2v\eta^2}.$$

Remark 8. Above lemma is crucially used in proving the regret bounds of our proposed algorithms (Alg. 1 and 3), in particular see the derivation of Lem. 15.

C Supplementary for Sec. 3

C.1 Algorithm Pseudocode for Winner-regret

Algorithm 1 MaxMin-UCB

```

1: init:  $\alpha > 0.5$ ,  $\mathbf{W} \leftarrow [0]_{n \times n}$ ,  $\mathcal{B}_0 \leftarrow \emptyset$ 
2: for  $t = 1, 2, 3, \dots, T$  do
3:   Set  $\mathbf{N} = \mathbf{W} + \mathbf{W}^\top$ , and  $\hat{\mathbf{P}} = \frac{\mathbf{W}}{\mathbf{N}}$ . Denote  $N = [n_{ij}]_{n \times n}$  and  $\hat{P} = [\hat{p}_{ij}]_{n \times n}$ .
4:   Define  $u_{ij} = \hat{p}_{ij} + \sqrt{\frac{\alpha \ln t}{n_{ij}}}$ ,  $\forall i, j \in [n], i \neq j$ ,  $u_{ii} = \frac{1}{2}$ ,  $\forall i \in [n]$ .  $\mathbf{U} = [u_{ij}]_{n \times n}$ 
5:    $\mathcal{C}_t \leftarrow \{i \in [n] \mid u_{ij} > \frac{1}{2}, \forall j \in [n] \setminus \{i\}\}$ ;  $\mathcal{B}_t \leftarrow \mathcal{C}_t \cap \mathcal{B}_{t-1}$ 
6:   if  $|\mathcal{C}_t| = 1$ , then set  $\mathcal{B}_t \leftarrow \mathcal{C}_t$ ,  $S_t \leftarrow \mathcal{C}_t$ , and go to Line 9
7:   if  $\mathcal{B}_t \neq \emptyset$  then set  $S_t \leftarrow \mathcal{B}_t$ , else select any item  $a \in \mathcal{C}_t$ , and set  $S_t \leftarrow \{a\}$ 
8:    $S_t \leftarrow S_t \cup \text{build\_S}(\mathbf{U}, S_t, [n] \setminus S_t, m)$ 
9:   Play  $S_t$ , and receive:  $\sigma_t \in \Sigma_{S_t}^m$ 
10:   $W(\sigma_t(k'), i) \leftarrow W(\sigma_t(k'), i) + 1 \ \forall i \in S_t \setminus \sigma_t(1 : k')$  for all  $k' = 1, 2, \dots, \min(|S_t| - 1, m)$ 
11: end for

```

Algorithm 2 $\text{build_}S(\mathbf{U}, S, I, \ell)$

```
1: input:  $\mathbf{U}$ : UCB matrix of  $\hat{\mathbf{P}}$ ,  $S$ : Set to build,  $I$ : pool of items  $I$ ,  $\ell > 0$ : Number of items to draw
2:  $\mathcal{C} \leftarrow \{i \in I \mid u_{ij} > \frac{1}{2}, \forall j \in I \setminus \{i\}\}$ 
3: while  $|\mathcal{C}| < \ell$  do
4:    $S \leftarrow S \cup \mathcal{C}$ ;  $I \leftarrow I \setminus \mathcal{C}$ ;  $\mathcal{C} \leftarrow \{i \in I \mid u_{ij} > \frac{1}{2}, \forall j \in I \setminus \{i\}\}$ ;  $\ell \leftarrow \ell - |\mathcal{C}|$ 
5: end while
6: for  $k' = 2, 3, \dots, \ell$  do
7:    $a \leftarrow \arg \max_{c \in I \setminus S} \left[ \min_{i \in S} u_{ci} \right]$ ;  $S \leftarrow S \cup \{a\}$ 
8: end for
9: return:  $S$ 
```

C.2 Restating the change of measure Lemma 1 of Kaufmann et al. [28]

Lemma 13 ([21]). *Given any bandit instance (A, μ) , with A being the arm set of MAB, and $\mu = \{\mu_i, \forall i \in A\}$ being the set of reward distributions associated to A with arm 1 having the highest expected reward, for any suboptimal arm $a \in A \setminus \{1\}$, consider an altered bandit instance μ^a with a being the (unique) optimal arm (the one with highest expected reward) for μ^a , and let μ and μ^a be mutually absolutely continuous for all $a \in A \setminus \{1\}$. At any round t , let A_t and Z_t denote the arm played and the observation (reward) received, respectively. Let $\mathcal{F}_t = \sigma(A_1, Z_1, \dots, A_t, Z_t)$ be the sigma algebra generated by the trajectory of a sequential bandit algorithm upto round t . Then, for any \mathcal{F}_T -measurable random variable Z with values in $[0, 1]$ it satisfies:*

$\sum_{i \in A} \mathbf{E}_\mu[N_i(T)] KL(\mu_i, \mu_i^a) \geq kl(\mathbf{E}_\mu[Z], \mathbf{E}_{\mu^a}[Z])$, where $N_i(T)$ denotes the number of pulls of arm $i \in [n]$ in T rounds, KL is the Kullback-Leibler divergence between distributions, and $kl(p, q)$ is the Kullback-Leibler divergence between Bernoulli distributions with parameters p and q .

C.3 Proof of Thm. 3

Theorem 3 (Winner-regret Lower Bound). *For any No-regret learning algorithm \mathcal{A} for Winner-regret that uses Winner Feedback, and for any problem instance $\text{MNL}(n, \theta)$ s.t. $a^* = \arg \max_{i \in [n]} \theta_i$,*

the expected regret incurred by \mathcal{A} satisfies $\liminf_{T \rightarrow \infty} \mathbf{E}_\theta \left[\frac{R_T(\mathcal{A})}{\ln T} \right] \geq \frac{\theta_{a^}}{\left(\min_{i \in [n] \setminus \{a^*\}} \frac{\theta_{a^*}}{\theta_i} - 1 \right)} (n - 1)$.*

Proof. The foundation of the current lower bound analysis stands on the ground on constructing $\text{MNL}(n, \theta)$ instances, and slightly modified versions of it such that no algorithm can achieve *No-regret* property on these instances without incurring $\Omega(n \ln T)$ regret. We describe the our constructed problem instances below:

Consider an $\text{MNL}(n, \theta)$ instance with the arm (item) set A containing all subsets of sizes $1, 2, \dots$ upto k of $[n]$: $A = \{S = (S(1), \dots, S(k')) \subseteq [n] \mid k' \in [k]\}$. Let $\text{MNL}(n, \theta^1)$ be the true distribution associated to the bandit arms $[n]$, given by the MNL parameters $\theta^1 = (\theta_1^1, \dots, \theta_n^1)$, such that $\theta_1^1 > \theta_i^1, \forall i \in [n] \setminus \{1\}$ such that,

True Instance: $\text{MNL}(n, \theta^1) : \theta_1^1 > \theta_2^1 = \dots = \theta_n^1 = \theta$ (say).

for some $\theta \in \mathbb{R}_+$. We moreover denote $\Lambda = (\theta_1^1 - \theta)$. Clearly, the *Best-Item* of $\text{MNL}(n, \theta^1)$ is $a^* = 1$. Now for every suboptimal item $a \in [n] \setminus \{1\}$, consider the altered problem instance $\text{MNL}(n, \theta^a)$ such that:

Instance a: $\text{MNL}(n, \theta^a) : \theta_a^a = \theta_1^1 + \epsilon = \theta + (\Lambda + \epsilon); \theta_i^a = \theta_i^1, \forall i \in [n] \setminus \{a\}$

for some $\epsilon > 0$. Clearly, the *Best-Item* of $\text{MNL}(n, \theta^a)$ is $a^* = a$. Note that, for problem instance $\text{MNL}(n, \theta^a)$ $a \in [n]$, the probability distribution associated to arm $S \in A$ is given by

$$p_S^a \sim \text{Categorical}(p_1, p_2, \dots, p_k), \text{ where } p_i = \Pr(i|S) = \frac{\theta_i^a}{\sum_{j \in S} \theta_j^a}, \forall i \in [k], \forall S \in A, \forall a \in [n],$$

since recall that $Pr(i|S)$ is as defined in Defn. 1. Now applying Lem. 13 we get,

$$\sum_{\{S \in \mathcal{A} \setminus \{a\} \mid a \in S\}} \mathbf{E}_{\theta^1}[N_S(T)] KL(p_S^1, p_S^a) \geq kl(\mathbf{E}_{\theta^1}[Z], \mathbf{E}_{\theta^a}[Z]). \quad (1)$$

The above result holds from the straightforward observation that for any arm $S \in \mathcal{A}$ with $a \notin S$, p_S^1 is same as p_S^a , hence $KL(p_S^1, p_S^a) = 0$, $\forall S \in \mathcal{A}$, $a \notin S$ or if $S = \{a\}$.

For the notational convenience we will henceforth denote $S^a = \{S \in \mathcal{A} \setminus \{a\} \mid a \in S\}$. Now let us analyse the right hand side of (1), for any set $S \in S^a$. We further denote $\Lambda' = \Lambda + \epsilon = (\theta_1^1 - \theta) + \epsilon$, $k' = |S| \in [k]$, $r = \mathbf{1}(1 \in S)$, $q = (k' - r)$, and $\theta_S^a = \sum_{i \in S} \theta_i^a$ for any $a \in [n]$.

Note that by construction of above problem instances we can further derive that for any $i \in S$:

$$p_S^1(i) = \begin{cases} \frac{r\theta_1^1}{\theta_S^1} = \frac{\theta + \Lambda}{\theta|S| + r\Lambda}, & \text{such that } i = 1, \\ \frac{\theta}{\theta_S^1} = \frac{\theta}{\theta|S| + r\Lambda}, & \text{otherwise.} \end{cases}$$

On the other hand, for problem **Instance-a**, we have that:

$$p_S^a(i) = \begin{cases} \frac{r\theta_1^1}{\theta_S^1 + \Lambda'} = \frac{\theta + \Lambda}{\theta|S| + \Lambda(1+r) + \epsilon}, & \text{such that } i = 1, \\ \frac{\theta_1^1 + \epsilon}{\theta_S^1 + \Lambda'} = \frac{\theta + \Lambda + \epsilon}{\theta|S| + \Lambda(1+r) + \epsilon}, & \text{such that } i = a, \\ \frac{\theta}{\theta_S^1 + \Lambda'} = \frac{\theta}{\theta|S| + \Lambda(1+r) + \epsilon}, & \text{otherwise.} \end{cases}$$

Now using the following upper bound on $KL(\mathbf{p}, \mathbf{q}) \leq \sum_{x \in \mathcal{X}} \frac{p^2(x)}{q(x)} - 1$, \mathbf{p} and \mathbf{q} be two probability mass functions on the discrete random variable \mathcal{X} [33], we get:

$$\begin{aligned} KL(p_S^1, p_S^a) &\leq \sum_{i \in S \setminus \{a\}} \left(\frac{\theta_i^1}{\theta_S^1} \right)^2 \left(\frac{\theta_S^a}{\theta_i^a} \right) + \left(\frac{\theta_a^1}{\theta_S^1} \right)^2 \left(\frac{\theta_S^a}{\theta_a^a} \right) - 1 \\ &= \sum_{i \in S \setminus \{a\}} \left(\frac{\theta_i^1}{\theta_S^1} \right)^2 \left(\frac{\theta_S^1 + \Lambda'}{\theta_i^1} \right) + \left(\frac{\theta_a^1}{\theta_S^1} \right)^2 \left(\frac{\theta_S^1 + \Lambda'}{\theta_a^1} \right) - 1 \\ &= \left(\frac{\theta_S^1 + \Lambda'}{(\theta_S^1)^2} \right) \left(\sum_{i \in [n] \setminus \{a\}} \theta_i^1 + \frac{(\theta_a^1)^2}{\theta_a^1 + \Lambda'} \right) - 1 \\ &= \left(\frac{\theta_S^1 + \Lambda'}{(\theta_S^1)^2} \right) \left(\frac{\theta_a^1 \theta_S^1 + \Lambda'(\theta_S^1 - \theta_a^1)}{\theta_a^1 + \Lambda'} \right) - 1 \quad \left[\text{replacing } \sum_{i \in [n] \setminus \{a\}} \theta_i^1 = (\theta_S^1 - \theta_a^1) \right] \\ &= \frac{\Lambda'^2 (\theta_S^1 - \theta_a^1)}{(\theta_S^1)^2 (\theta_a^1 + \Lambda')} \leq \frac{\Lambda'^2}{\theta_S^1 (\theta_a^1 + \Lambda')} = \frac{\Lambda'^2}{\theta_S^1 (\theta_1^1 + \epsilon)} \\ &= \frac{(\Lambda + \epsilon)^2}{(\theta|S| + r\Lambda)(\theta_1^1 + \epsilon)} \leq \frac{(\Lambda + \epsilon)^2}{\theta|S|(\theta_1^1 + \epsilon)} \end{aligned} \quad (2)$$

Let us now analyze the left hand side of (1), with $Z = \frac{N_a(T)}{T}$, where $N_a(T)$ simply denotes the number of times the singleton set containing item $\{a\}$ is played by \mathcal{A} , for any suboptimal item $a \in [n] \setminus \{1\}$. Thus we get,

$$kl(\mathbf{E}_{\theta^1}[Z], \mathbf{E}_{\theta^a}[Z]) \geq \left(1 - \frac{\mathbf{E}_{\theta^1}[N_1(T)]}{T} \right) \ln \frac{T}{T - \mathbf{E}_{\theta^a}[N_1(T)]} - \ln 2, \quad (3)$$

where the inequality follows from the fact that for all $(p, q) \in [0, 1]^2$, $kl(p, q) = p \ln \frac{1}{q} + (1 - p) \ln \frac{1}{1-q} + (p \ln p + (1 - p) \ln(1 - p))$, and $p \ln \frac{1}{q} \geq 0$, $(p \ln p + (1 - p) \ln(1 - p)) \geq -\ln 2$.

But now owing to the *No-regret* property (see Defn. 2) of Algorithm \mathcal{A} , we have $\mathbf{E}_{\theta^1}[N_a(T)] = o(T^\alpha)$ and $T - \mathbf{E}_{\theta^a}[N_a(T)] = \mathbf{E}_{\theta^a}[\sum_{S \in A, S \neq \{a\}} N_S(T)] = o(T^\alpha)$, $0 < \alpha \leq 1$. Thus from (3), we get

$$\begin{aligned} \lim_{T \rightarrow \infty} \frac{kl(\mathbf{E}_{\theta^1}[Z], \mathbf{E}_{\theta^a}[Z])}{\ln T} &\geq \lim_{T \rightarrow \infty} \frac{1}{\ln T} \left[\left(1 - \frac{\mathbf{E}_{\theta^1}[N_a(T)]}{T} \right) \ln \frac{T}{T - \mathbf{E}_{\theta^a}[N_a(T)]} - \ln 2 \right] \\ &= \lim_{T \rightarrow \infty} \frac{1}{\ln T} \left[\left(1 - \frac{o(T^\alpha)}{T} \right) \ln \frac{T}{T^\alpha} - \ln 2 \right] = (1 - \alpha). \end{aligned}$$

Combining above with (2) we get:

$$\begin{aligned} \lim_{T \rightarrow \infty} \frac{1}{\ln T} \sum_{\{S \in S^a\}} \mathbf{E}_{\theta^1}[N_S(T)] KL(p_S^1, p_S^a) &\geq (1 - \alpha) \\ \implies \lim_{T \rightarrow \infty} \frac{1}{\ln T} \sum_{\{S \in S^a\}} \mathbf{E}_{\theta^1}[N_S(T)] \frac{\Lambda'^2}{\theta|S|(\theta_1^1 + \epsilon)} &\geq (1 - \alpha) \\ \implies \lim_{T \rightarrow \infty} \frac{1}{\ln T} \sum_{\{S \in S^a\}} \mathbf{E}_{\theta^1}[N_S(T)] \frac{\Lambda'}{|S|} &\geq (1 - \alpha) \frac{\theta(\theta_1^1 + \epsilon)}{\Lambda'} \end{aligned} \quad (4)$$

Now applying (4) for each modified bandit **Instance- θ^a** , and summing over $(n - 1)$ suboptimal items $a \in [n] \setminus \{1\}$ we get,

$$\lim_{T \rightarrow \infty} \frac{1}{\ln T} \sum_{a=2}^n \sum_{\{S \in S^a\}} \mathbf{E}_{\theta^1}[N_S(T)] \frac{\Lambda'}{|S|} \geq (1 - \alpha) \theta(\theta_1^1 + \epsilon) \frac{(n - 1)}{\Lambda'} \quad (5)$$

Now recall that regret of \mathcal{A} on the true instance $\text{MNL}(n, \theta^1)$, is given by: $R_T^1(\mathcal{A}) = \sum_{t=1}^T \left(\sum_{i \in S_t} \frac{(\theta_1^1 - \theta_i^1)}{|S_t|} \right)$. But this can be equivalently written as:

$$\begin{aligned} \mathbf{E}_{\theta^1}[R_T^1(\mathcal{A})] &= \mathbf{E}_{\theta^1} \left[\sum_{t=1}^T \sum_{i \in S_t} \frac{(\theta_1^1 - \theta_i^1)}{|S_t|} \right] \\ &= \mathbf{E}_{\theta^1} \left[\sum_{t=1}^T \sum_{S \in A} \mathbf{1}(S_t = S) \sum_{a=2}^n \mathbf{1}(a \in S) \frac{(\theta_1^1 - \theta_a^1)}{|S_t|} \right] \\ &= \mathbf{E}_{\theta^1} \left[\sum_{a=2}^n \sum_{t=1}^T \sum_{S \in A} \mathbf{1}(S_t = S) \mathbf{1}(a \in S) \frac{(\theta_1^1 - \theta_a^1)}{|S|} \right] \\ &= \sum_{a=2}^n \sum_{t=1}^T \mathbf{E}_{\theta^1} \left[\sum_{S \in A} \mathbf{1}(S_t = S) \mathbf{1}(a \in S) \frac{(\theta_1^1 - \theta_a^1)}{|S|} \right] \\ &= \sum_{a=2}^n \sum_{S \in A} \mathbf{E}_{\theta^1} \left[\sum_{t=1}^T \mathbf{1}(S_t = S) \mathbf{1}(a \in S) \frac{(\theta_1^1 - \theta_a^1)}{|S|} \right] \\ &= \sum_{a=2}^n \sum_{S \in A} \left[\mathbf{E}_{\theta^1}[N_S(T)] \mathbf{1}(a \in S) \frac{(\theta_1^1 - \theta_a^1)}{|S|} \right] \\ &= \sum_{a=2}^n \sum_{\{S \in A | a \in S\}} \mathbf{E}_{\theta^1}[N_S(T)] \frac{\Lambda}{|S|} \end{aligned} \quad (6)$$

Then combining (6) with (5) we get and taking $\epsilon \rightarrow 0$:

$$\begin{aligned} \lim_{T \rightarrow \infty} \frac{1}{\ln T} \mathbf{E}_{\theta^1} [R_T^1(\mathcal{A})] &\geq \lim_{T \rightarrow \infty} \frac{1}{\ln T} \sum_{a=2}^n \sum_{\{S \in S^a\}} \mathbf{E}_{\theta^1} [N_S(T)] \frac{\Lambda}{|S|} \\ &\geq (1 - \alpha) \theta(\theta_1^1) \frac{(n-1)}{\Lambda} = (1 - \alpha) \theta_1^1 \frac{(n-1)}{(\frac{\theta_1^1}{\theta} - 1)}. \end{aligned}$$

Finally, since α is a fixed constant in $(0, 1]$, above construction shows the existence of a $\text{MNL}(n, \theta)$ problem instance, precisely $\text{MNL}(n, \theta^1)$, such that for large T , $\mathbf{E}_{\theta^1} [R_T^1] = \Omega\left(\frac{\theta_1^1}{(\frac{\theta_1^1}{\theta} - 1)} (n-1) \ln T\right)$, which concludes the proof. \square

C.4 An alternate version of the regret lower bound (Thm. 4) with pairwise preference-based instance complexities

Theorem 14 (Alternate version of Thm. 4 with pairwise preference based instance complexities). *For any No-regret algorithm \mathcal{A} for Winner-regret with Winner Feedback, there exists a problem instance of $\text{MNL}(n, \theta)$ model, such that the expected regret incurred by \mathcal{A} on it satisfies $\liminf_{T \rightarrow \infty} \mathbf{E}_{\theta} \left[\frac{R_T^1(\mathcal{A})}{\ln T} \right] \geq \frac{\theta_{a^*}}{4 \left(\min_{i \in [n] \setminus \{a^*\}} p_{a^*, i} - 0.5 \right)} (n-1)$, where $p_{ij} := \Pr(i | \{i, j\}) = \frac{\theta_i}{\theta_i + \theta_j} \forall i, j \in [n]$, and $\mathbf{E}_{\theta}[\cdot]$, a^* are same as that of Thm. 3. Thus the only difference lies in terms of the instance dependent complexity term ('gap') which is now expressed in terms of pairwise preference of the best item a^* over the second best item: $\left(\min_{i \in [n] \setminus \{a^*\}} p_{a^*, i} - 0.5 \right)$.*

Proof. Firstly, is easy to note that $\arg \min_{i \in [n] \setminus \{a^*\}} (p_{a^*, i} - 0.5) = \arg \max_{i \in [n] \setminus \{a^*\}} \theta_i =: b$ (say). The proof now follows from the fact that

$$p_{a^*b} - 0.5 = \frac{\theta_{a^*} - \theta_b}{2(\theta_{a^*} + \theta_b)} \leq \frac{\theta_{a^*} - \theta_b}{4\theta_b} \quad (\text{since } \theta_b \leq \theta_{a^*})$$

Thus using the lower bound from Thm. 4, one can further derive

$$\liminf_{T \rightarrow \infty} \frac{1}{\ln T} \mathbf{E}_{\theta} [R_T^1(\mathcal{A})] \geq \frac{4\theta_{a^*}\theta_b}{4(\theta_{a^*} - \theta_b)} (n-1) \geq \frac{\theta_{a^*}}{4 \left(\min_{i \in [n] \setminus \{a^*\}} p_{a^*, i} - 0.5 \right)} (n-1),$$

which proves the claim. \square

C.5 Proof of Thm. 4

Theorem 4 (Regret Lower Bound: Winner-regret with Top- m -ranking Feedback). *For any No-regret algorithm \mathcal{A} for the Winner-regret problem with Top- m -ranking Feedback, there exists a problem instance $\text{MNL}(n, \theta)$ such that the expected Winner-regret incurred by \mathcal{A} satisfies $\liminf_{T \rightarrow \infty} \mathbf{E}_{\theta} \left[\frac{R_T^1(\mathcal{A})}{\ln T} \right] \geq \frac{\theta_{a^*}}{\left(\min_{i \in [n] \setminus \{a^*\}} \frac{\theta_{a^*}}{\theta_i} - 1 \right)} \frac{(n-1)}{m}$, where as in Thm. 3, $\mathbf{E}_{\theta}[\cdot]$ denotes expectation under the algorithm and the MNL model $\text{MNL}(n, \theta)$, and recall $a^* := \arg \max_{i \in [n]} \theta_i$.*

Proof. The proof proceeds almost same as the proof of Thm. 3, the only difference lies in the analysis of the KL-divergence terms with Top- m -ranking Feedback.

Consider the exact same $\text{MNL}(n, \theta)$ instances, $\text{MNL}(n, \theta^a)$ we constructed for Thm. 3. It is now interesting to note that how Top- m -ranking Feedback affects the KL-divergence analysis, precisely the KL-divergence shoots up by a factor of m which in fact triggers an $\frac{1}{m}$ reduction in regret learning rate. Note that for Top- m -ranking Feedback for any problem instance $\text{MNL}(n, \theta^a)$, $a \in [n]$, each

k -set $S \subseteq [n]$ (such that $|S| = k$) is associated to $\binom{k}{m}(m!)$ number of possible outcomes, each representing one possible ranking of set of m items of S , say S_m . Also the probability of any permutation $\sigma \in \Sigma_{S_m}$ is given by $p_S^a(\sigma) = Pr_{\theta^a}(\sigma|S)$, where $Pr_{\theta^a}(\sigma|S)$ is as defined for Top- m -ranking Feedback (in Sec. 2.1). For ease of analysis let us first assume $1 \notin S$ and let $k' = |S|$ be the cardinality of S and $m' = \min(m, k)$. (Note if $m' \leq m + 1$ the corresponding Top- m -ranking Feedback becomes a full ranking feedback on the entire m' items). In this case we get

$$p_S^1(\sigma) = \prod_{i=1}^{m'} \frac{\theta_{\sigma(i)}^1}{\sum_{j=i}^{m'} \theta_{\sigma(j)}^1 + \sum_{j \in S \setminus \sigma(1:m')} \theta_{\sigma(j)}^1} = \frac{1}{k'(k'-1)(k'-2) \cdots (k'-m+1)}, \quad \forall \sigma \in \Sigma_S^{m'}.$$

On the other hand, for problem **Instance-a**, we have that:

$$\begin{aligned} p_S^a(\sigma) &= \prod_{i=1}^{m'} \frac{\theta_{\sigma(i)}^a}{\sum_{j=i}^{m'} \theta_{\sigma(j)}^a + \sum_{j \in S \setminus \sigma(1:m')} \theta_{\sigma(j)}^a} \\ &= \begin{cases} \frac{x}{(x+k'-1)(x+k'-2) \cdots (x+k'-i)(k'-i)(k'-i-1) \cdots (k'-m+1)}, & \text{such that } \sigma(i) = a, \\ \frac{1}{(x+k'-1)(x+k'-2) \cdots (x+k'-i)(k'-i)(k'-i-1) \cdots (k'-m+1)}, & \text{such that } a \notin \sigma(1:m'), \end{cases} \end{aligned}$$

where we denote by $x = 1 + \frac{\Lambda'}{\theta}$, where recall that we denote $\Lambda' = \Lambda + \epsilon$. Similarly we can derive the probability distribution associated to sets including item 1.

The important thing now to note is that $KL(p_S^1, p_S^a) = 0$ for any set $S \not\ni a$. Hence while comparing the KL-divergence of instances θ^1 vs θ^a , we need to focus only on sets containing a . Applying *Chain-Rule* of KL-divergence, we now get

$$\begin{aligned} KL(p_S^1, p_S^a) &= KL(p_S^1(\sigma_1), p_S^a(\sigma_1)) + KL(p_S^1(\sigma_2 | \sigma_1), p_S^a(\sigma_2 | \sigma_1)) + \cdots \\ &\quad + KL(p_S^1(\sigma_m | \sigma(1:m-1)), p_S^a(\sigma_m | \sigma(1:m-1))), \end{aligned} \tag{7}$$

where we abbreviate $\sigma(i)$ as σ_i and following the usual convention the notation $KL(P(Y | X), Q(Y | X)) := \sum_x Pr(X = x) [KL(P(Y | X = x), Q(Y | X = x))]$ denotes the conditional KL-divergence. Moreover it is easy to note that for any $\sigma \in \Sigma_S^m$ such that $\sigma(i) = a$, we have $KL(p_S^1(\sigma_{i+1} | \sigma(1:i)), p_S^a(\sigma_{i+1} | \sigma(1:i))) := 0$, for all $i \in [m]$.

Now as derived in (2) in the proof of Thm. 3, we have

$$KL(p_S^1(\sigma_1), p_S^a(\sigma_1)) \leq \frac{(\Lambda + \epsilon)^2}{\theta|S|(\theta_1^1 + \epsilon)}.$$

To bound the remaining terms of (7), note that for all $i \in [m-1]$

$$\begin{aligned} &KL(p_S^1(\sigma_{i+1} | \sigma(1:i)), p_S^a(\sigma_{i+1} | \sigma(1:i))) \\ &= \sum_{\sigma' \in \Sigma_S^i} Pr(\sigma') KL(p_S^1(\sigma_{i+1} | \sigma(1:i)) = \sigma', p_S^a(\sigma_{i+1} | \sigma(1:i)) = \sigma') \\ &= \sum_{\sigma' \in \Sigma_S^i | a \notin \sigma'} \left[\prod_{j=1}^i \left(\frac{\theta_{\sigma'_j}^1}{\theta_S^1 - \sum_{j'=1}^{j-1} \theta_{\sigma'_{j'}}} \right) \right] \frac{\Lambda'^2}{(|S| - i)\theta(\theta_1^1 + \epsilon)} \\ &= \prod_{j=1}^i (|S| - j) \frac{\theta^i}{\prod_{j=1}^i (\theta(|S| - i + 1) + \Lambda')} \frac{(\Lambda')^2}{(|S| - i)\theta(\theta_1^1 + \epsilon)} = \frac{\emptyset}{(\theta|S| + \Lambda')} \frac{\Lambda'^2}{\emptyset(\theta_1^1 + \epsilon)} \\ &= \frac{\Lambda'^2}{(\theta|S| + \Lambda')(\theta_1^1 + \epsilon)}, \end{aligned}$$

where for simplicity we assumed $1 \notin S$. It is easy to note that the similar analysis would lead to the same upper bound for sets S containing 1 as well. Thus applying above in (7) we get:

$$\begin{aligned} KL(p_S^1, p_S^a) &= KL(p_S^1(\sigma_1) + \dots + KL(p_S^1(\sigma_m \mid \sigma(1:m-1)), p_S^a(\sigma_m \mid \sigma(1:m-1))) \\ &\leq \frac{m\Lambda'^2}{|S|\theta(\theta_1^1 + \epsilon)}. \end{aligned} \quad (8)$$

Eqn. (8) gives the main result to derive Thm. 4 as it shows an m -factor blow up in the KL-divergence terms owing to Top- m -ranking Feedback. The rest of the proof follows exactly the same argument used in 3. We add the steps below for convenience. Firstly, considering $Z = \frac{N_a(T)}{T}$, in this case as well, one can show that:

$$\begin{aligned} \lim_{T \rightarrow \infty} \frac{kl(\mathbf{E}_{\theta^1}[Z], \mathbf{E}_{\theta^a}[Z])}{\ln T} &\geq \lim_{T \rightarrow \infty} \frac{1}{\ln T} \left[\left(1 - \frac{\mathbf{E}_{\theta^1}[N_a(T)]}{T} \right) \ln \frac{T}{T - \mathbf{E}_{\theta^a}[N_a(T)]} - \ln 2 \right] \\ &= \lim_{T \rightarrow \infty} \frac{1}{\ln T} \left[\left(1 - \frac{o(T^\alpha)}{T} \right) \ln \frac{T}{T^\alpha} - \ln 2 \right] = (1 - \alpha). \end{aligned}$$

Now combining above with (8) we get:

$$\begin{aligned} \lim_{T \rightarrow \infty} \frac{1}{\ln T} \sum_{\{S \in S^a\}} \mathbf{E}_{\theta^1}[N_S(T)] KL(p_S^1, p_S^a) &\geq (1 - \alpha) \\ \Rightarrow \lim_{T \rightarrow \infty} \frac{1}{\ln T} \sum_{\{S \in S^a\}} \mathbf{E}_{\theta^1}[N_S(T)] \frac{m\Lambda'^2}{\theta|S|(\theta_1^1 + \epsilon)} &\geq (1 - \alpha) \\ \Rightarrow \lim_{T \rightarrow \infty} \frac{1}{\ln T} \sum_{\{S \in S^a\}} \mathbf{E}_{\theta^1}[N_S(T)] \frac{\Lambda'}{|S|} &\geq (1 - \alpha) \frac{\theta(\theta_1^1 + \epsilon)}{m\Lambda'} \end{aligned} \quad (9)$$

Applying (9) for each modified bandit **Instance- θ^a** , and summing over $(n-1)$ suboptimal items $a \in [n] \setminus \{1\}$ we get,

$$\lim_{T \rightarrow \infty} \frac{1}{\ln T} \sum_{a=2}^n \sum_{\{S \in S^a\}} \mathbf{E}_{\theta^1}[N_S(T)] \frac{\Lambda'}{|S|} \geq (1 - \alpha) \theta(\theta_1^1 + \epsilon) \frac{(n-1)}{m\Lambda'} \quad (10)$$

Further recall that we derived earlier that $\mathbf{E}_{\theta^1}[R_T^1(\mathcal{A})] = \sum_{a=2}^n \sum_{\{S \in \mathcal{A} \mid a \in S\}} \mathbf{E}_{\theta^1}[N_S(T)] \frac{\Lambda}{|S|}$, using which combined with (10), and taking $\epsilon \rightarrow 0$ we get:

$$\begin{aligned} \lim_{T \rightarrow \infty} \frac{1}{\ln T} \mathbf{E}_{\theta^1}[R_T^1(\mathcal{A})] &\geq \lim_{T \rightarrow \infty} \frac{1}{\ln T} \sum_{a=2}^n \sum_{\{S \in S^a\}} \mathbf{E}_{\theta^1}[N_S(T)] \frac{\Lambda}{|S|} \\ &\geq (1 - \alpha) \theta(\theta_1^1) \frac{(n-1)}{m\Lambda} = (1 - \alpha) \theta_1^1 \frac{(n-1)}{m(\frac{\theta_1^1}{\theta} - 1)}. \end{aligned}$$

Now since α is a fixed constant in $(0, 1]$, we thus prove the existence of a $\text{MNL}(n, \theta)$ problem instance (precisely $\text{MNL}(n, \theta^1)$), such that for large T , $\mathbf{E}_{\theta^1}[R_T^1] = \Omega\left(\frac{\theta_1^1}{\left(\frac{\theta_1^1}{\theta} - 1\right)} \frac{(n-1)}{m} \ln T\right)$, which concludes the proof. \square

C.6 Proof of Thm. 5

Theorem 5 (MaxMin-UCB: High Probability Regret bound). Fix a time horizon T and $\delta \in (0, 1)$, $\alpha > \frac{1}{2}$. With probability at least $(1 - \delta)$, the regret of MaxMin-UCB for Winner-regret with Top- m -ranking Feedback satisfies $R_T^1 \leq \left(2 \left[\frac{2\alpha n^2}{(2\alpha-1)\delta} \right]^{\frac{1}{2\alpha-1}} + 2D \ln 2D \right) \hat{\Delta}_{\max} + \frac{\ln T}{m+1} \sum_{i=2}^n (D_{\max} \hat{\Delta}_i)$, where $\forall i \in [n] \setminus \{a^*\}$, $\hat{\Delta}_i = (\theta_{a^*} - \theta_i)$, $\Delta_i = \frac{\theta_{a^*} - \theta_i}{2(\theta_{a^*} + \theta_i)}$, $\hat{\Delta}_{\max} = \max_{i \in [n] \setminus \{a^*\}} \hat{\Delta}_i$, $D_{1i} = \frac{4\alpha}{\Delta_i^2}$, $D := \sum_{i < j} D_{ij}$, $D_{\max} = \max_{i \in [n] \setminus \{a^*\}} D_{1i}$.

Proof. For the notational convenience we will assume $\theta_1 > \theta_2 \dots \geq \theta_n$, so $a^* = 1$. We also use $\hat{p}_{ij}(t)$, $n_{ij}(t)$ and $u_{ij}(t)$ to denote the values of the respective quantities at time iteration t , for any $t \in [T]$, just to be precise

$$u_{ij}(t) = \hat{p}_{ij}(t) + \sqrt{\frac{\alpha \ln t}{n_{ij}(t)}}, \forall i, j \in [n], i \neq j,$$

and $u_{ii}(t) = \frac{1}{2}$ for all $i \in [n]$. We also find it convenient to denote

$$c_{ij}(t) = \sqrt{\frac{\alpha \ln t}{n_{ij}(t)}}, \text{ and } l_{ij}(t) = 1 - u_{ij}(t).$$

We also denote $T_0(\delta) = 2f(\delta) + 2D \ln 2D$, where $D := \sum_{i < j} D_{ij}$. We start with the following crucial lemma that analyzes the confidence bounds $[l_{ij}(t), u_{ij}(t)]$ on the pairwise probability estimates \hat{p}_{ij} for each pair (i, j) , $i \neq j$.

Lemma 15. Suppose $\mathbf{P} := [p_{ij}]$ be the pairwise probability matrix associated to the underlying MNL(n, θ) model, i.e. $p_{ij} = \Pr(i | \{i, j\}) = \frac{\theta_i}{\theta_i + \theta_j}$. Then for any $\alpha > \frac{1}{2}$, $\delta \in (0, 1)$,

$$\Pr(\forall t > f(\delta), \forall i, j, p_{ij} \in [l_{ij}(t), u_{ij}(t)]) > (1 - \delta)$$

$$\text{where } f(\delta) = \left\lceil \frac{2\alpha n^2}{(2\alpha-1)\delta} \right\rceil^{\frac{1}{2\alpha-1}}.$$

Proof. The proof of this lemma is adapted from a similar result (Lemma 1) of [47]. Suppose $\mathcal{G}_{ij}(t)$ denotes the event that at time t , $p_{ij} \in [l_{ij}(t), u_{ij}(t)]$, $\forall i, j \in [n]$. $\mathcal{G}_{ij}^c(t)$ denotes its complement.

Case 1: ($i = j$) Note that for any such that pair (i, i) , $\mathcal{G}_{ii}(t)$ always holds true for any $t \in [T]$ and $i \in [n]$, as $p_{ii} = u_{ii} = l_{ii} = \frac{1}{2}$.

Case 2: ($i \neq j$) Recall from the definition of $u_{ij}(t)$ that $\mathcal{G}_{ij}(t)$ equivalently implies at round t , $|\hat{p}_{ij}(t) - p_{ij}| \leq \sqrt{\frac{\alpha \ln(t)}{n_{ij}(t)}}$, $\forall i, j \in [n]$. Moreover, for any t and i, j , $\mathcal{G}_{ij}(t)$ holds if and only if $\mathcal{G}_{ij}(t)$ as $|\hat{p}_{ji}(t) - p_{ji}| = |(1 - \hat{p}_{ij}(t)) - (1 - p_{ij})| = |\hat{p}_{ij}(t) - p_{ij}|$. Thus we will restrict our focus only to pairs $i < j$ for the rest of the proof.

Let $\tau_{ij}(n)$ the time step $t \in [T]$ when the pair (i, j) was updated for the n^{th} time. Clearly for any $n \in \mathbb{N}$, $\tau_{ij}(n+1) \geq \tau_{ij}(n)$ and $\tau_{ij}(n+k) > \tau_{ij}(n)$. For convenience of notation we use $F = f(\delta)$. It is now straightforward to note that we want to find F such that:

$$\begin{aligned} \Pr(\forall t > F, \forall i, j \text{ such that } i < j, \mathcal{G}_{ij}(t)) &> (1 - \delta) \text{ or equivalently,} \\ \Pr(\exists t > F \text{ and atleast a pair } (i < j), \text{ with } \mathcal{G}_{ij}^c(t)) &< \delta. \end{aligned} \quad (11)$$

Further decomposing the right hand side of above we get:

$$\begin{aligned}
& Pr(\exists t > F, i < j, \text{ such that } \mathcal{G}_{ij}^c(t)) \\
& \leq \sum_{i < j} \left[Pr\left(\exists n \geq 0, \tau_{ij}(n) > F, |p_{ij} - \hat{p}_{ij}(\tau_{ij}(n))| > \sqrt{\frac{\alpha \ln(\tau_{ij}(n))}{n_{ij}(\tau_{ij}(n))}}\right) \right] \\
& \leq \sum_{i < j} \left[Pr\left(\exists n \leq F, \tau_{ij}(n) > F, |p_{ij} - \hat{p}_{ij}(n)| > \sqrt{\frac{\alpha \ln(\tau_{ij}(n))}{n}}\right) + Pr\left(\exists n > F, |p_{ij} - \hat{p}_{ij}(n)| > \sqrt{\frac{\alpha \ln(\tau_{ij}(n))}{n}}\right) \right],
\end{aligned}$$

where $\hat{p}(n) = \frac{w_{ij}(\tau_{ij}(n))}{w_{ij}(\tau_{ij}(n)) + w_{ij}(\tau_{ji}(n))}$ is the frequentist estimate of p_{ij} after n comparisons between arm i and j . Now the above inequality can be further upper bounded as:

$$\begin{aligned}
& Pr(\exists t > F, i < j, \text{ such that } \mathcal{G}_{ij}^c(t)) \\
& \leq \sum_{i < j} \left[Pr\left(\exists n \leq F, \tau_{ij}(n) > F, |p_{ij} - \hat{p}_{ij}(n)| > \sqrt{\frac{\alpha \ln(F)}{n}}\right) + Pr\left(\exists n > F, |p_{ij} - \hat{p}_{ij}(n)| > \sqrt{\frac{\alpha \ln(n)}{n}}\right) \right],
\end{aligned}$$

since in the second term $\tau_{ij}(n) > F$, and for the third term $n < \tau_{ij}(n)$ since at a particular time iteration, any pair (i, j) , can be updated at most once, implying $n \leq \tau_{ij}(n)$. Using Lem. 12 we now get:

$$\begin{aligned}
& Pr(\exists t > F, i < j, \text{ such that } \mathcal{G}_{ij}^c(t)) \\
& \leq \sum_{i < j} \left[\sum_{n=1}^F 2e^{-2n \frac{\alpha \ln F}{n}} + \sum_{n=F+1}^{\infty} 2e^{-2n \frac{\alpha \ln n}{n}} \right] \\
& = \frac{n(n-1)}{2} \left[2 \sum_{n=1}^F \frac{1}{F^{2\alpha}} + \sum_{n=F+1}^{\infty} \frac{2}{n^{2\alpha}} \right] \\
& \leq \frac{n^2}{F^{2\alpha-1}} + n^2 \int_F^{\infty} \frac{dx}{x^{2\alpha}} \leq \frac{n^2}{F^{2\alpha-1}} - \frac{n^2}{(1-2\alpha)F^{2\alpha-1}} = \frac{(2\alpha)n^2}{(2\alpha-1)F^{2\alpha-1}}.
\end{aligned}$$

Now from (11), we want to find F such that

$$\frac{2\alpha n^2}{(2\alpha-1)F^{2\alpha-1}} \leq \delta,$$

which suffices by setting $F = \left\lceil \frac{2\alpha n^2}{(2\alpha-1)\delta} \right\rceil^{\frac{1}{2\alpha-1}}$, and recall that we assumed $f(\delta) = F$, for any given $\delta \in [0, 1]$, which concludes the claim. \square

Lem. 15 ensures the termination of the **Random-Exploration** phase. We now proceed to analyse **Progress** phase which shows that the set \mathcal{B}_t captures the *Best-Item* $a^* = 1$ ‘soon after’ $f(\delta)$ within a constant number of rounds $T_0(\delta)$ which is independent of T (see Lem. 20). Once the 1 is captured in \mathcal{B}_t , the algorithm goes into **Saturation** phase where the suboptimal items can not stay too long in the set of potential *Best-Items* \mathcal{C}_t , and thus the regret bound follows (Lem. 22). More formally, the rest of the proof follows based on the following main observations:

In Progress:

- **Observation 1:** At any iteration, the set \mathcal{B}_t is either singleton or an empty set.
- **Observation 2:** For any $\delta \in (0, 1)$, suppose $T_0(\delta) := \min_{t > f(\delta)} \mathcal{B}_t = \{1\}$, then for any $t > T_0(\delta)$, $\mathcal{B}_t = \{1\}$.
- **Observation 3:** $T_0(\delta)$ is not far from $f(\delta)$ (see Lem. 20 which holds due to Lem. 18 and 19)

In Saturation:

- **Observation 4:** After $T_0(\delta)$, $\mathcal{B}_t = \{1\}$ thereafter, and thus it is always played in S_t , i.e. $1 \in S_t$ for all $t > T_0(\delta)$. Now the suboptimal items start getting frequently compared to item 1 every time they are played alongside with 1, and thus they can not stay too long in the set of ‘good’ items \mathcal{C}_t and eventually $\mathcal{C}_t = \{1\}$, when the algorithm *MaxMin-UCB* plays the optimal set $S_t = \{1\}$ only, and thus the regret bound follows. (see Lem. 21 and 22)

Observation 1 is straightforward to follow from Alg. 1. Observation 2 follows from Lem. 15, as for any $t > f(\delta)$, $1 \in \mathcal{C}_t$ always, since $u_{1i} \geq p_{1i} > \frac{1}{2}$, $\forall i \in [n] \setminus \{1\}$. We next recall the notations before proceeding to the next results: Let $\Delta_i = P(1 \succ i) - \frac{1}{2} = \frac{\theta_1 - \theta_i}{2(\theta_1 + \theta_i)}$, for all $i \in [n] \setminus \{1\}$. For any pair (i, j) such that $1 \notin \{i, j\}$, we define $D_{ij} = \frac{4\alpha}{\min\{\Delta_i^2, \Delta_j^2\}}$. For any $i \in [n] \setminus \{1\}$, $D_{1i} = \frac{4\alpha}{\Delta_i^2}$.

Definition 16 (Unsaturated Pairs). At any time $t \in [T]$, and any pair of two distinct items $i, j \in [n]$, we term the pair (i, j) to be *unsaturated* at time t if $n_{ij}(t) \leq D_{ij} \ln t$. Otherwise, we call the pair *saturated* at t .

Lemma 17. For any set $S \subseteq [n]$ such that $|S| \geq m + 1$, and given a *Top- m -ranking Feedback* $\sigma \in \Sigma_S^m$ (for any $m \in [k - 1]$), applying pairwise Rank-Breaking on S according to σ , updates each element $i \in S$ for atleast m distinct pairs.

Proof. For any item $i \in S$, one can make the following two case analyses:

Case 1: ($i \in \sigma(1 : m)$). If the item i occurs in one of the top- m position, it is clearly compared with rest of the $|S| - 1 \geq m$ elements of S , as it is beaten by the preceding items in σ and wins over the rest.

Case 2: ($i \notin \sigma(1 : m)$). In this case i gets updated for m many pairs since it is considered to be beaten by all items in $\sigma(1 : m)$ in a pairwise duel.

The claim follows combining Case 1 and 2 above. \square

Lemma 18. Assuming $\forall t > f(\delta)$, and $\forall i, j \in [n]$ $p_{ij} \in [l_{ij}(t), u_{ij}(t)]$, for some $\delta \in (0, 1)$: At any iteration $t > f(\delta)$, if \exists a suboptimal item $i \in [n] \setminus \{1\}$, such that $i \in \mathcal{C}_t$, then the pair $(1, i)$ is unsaturated at t .

Proof. Firstly note that for any $t > f(\delta)$, $1 \in \mathcal{C}_t$ always, since $u_{1i} \geq p_{1i} > \frac{1}{2}$, $\forall i \in [n] \setminus \{1\}$.

Now suppose $(1, i)$ is indeed saturated at time t , i.e. $n_{1i}(t) > D_{1i} \ln t$, then this implies:

$$u_{i1}(t) = \hat{p}_{i1}(t) + c_{i1}(t) \leq p_{i1}(t) + 2c_{i1}(t) = p_{1i}(t) + \Delta_i = \frac{1}{2},$$

which implies $i \notin \mathcal{C}_t$, at t . Thus $(1, i)$ must be unsaturated at t . \square

Lemma 19. Assuming $\forall t > f(\delta)$, and $\forall i, j \in [n]$ $p_{ij} \in [l_{ij}(t), u_{ij}(t)]$, for some $\delta \in (0, 1)$: At any iteration $t > f(\delta)$, for any set $S \not\ni 1$ if \exists a suboptimal item $a \in [n] \setminus \{1\}$, such that $a = \arg \max_{c \in I \setminus S} \left[\min_{i \in S} u_{ci}(t) \right]$, then \exists atleast one suboptimal item $i \in S$ such that the pair (i, a) is unsaturated at t .

Proof. We start by noting that for any $i \in \mathcal{C}_t \setminus \{1\}$ if $u_{ji}(t) > u_{1i}(t)$, then $n_{ij} \leq D_{ij} \ln t$, i.e. the pair (i, j) must be unsaturated at round t . Suppose not and $n_{ij}(t) > D_{ij} \ln t$. Then we have that

$$u_{ij}(t) - l_{ij}(t) = 2c_{ij}(t) \leq \sqrt{\min\{\Delta_i^2, \Delta_j^2\}} = \min\{\Delta_i, \Delta_j\}.$$

But on the other hand, since $u_{ji}(t) > u_{1i}(t)$, this implies:

$$u_{ij}(t) - l_{ij}(t) = u_{ij}(t) + u_{ji}(t) - 1 > \frac{1}{2} + u_{1i}(t) - 1 > \frac{1}{2} + p_{1i}(t) - 1 = \Delta_i \geq \min\{\Delta_i, \Delta_j\},$$

where the first inequality is because $i \in \mathcal{C}_t$, hence $u_{ij}(t) > \frac{1}{2}$ and $u_{ji}(t) > u_{1i}(t)$. This leads to a contradiction implying that (i, j) has to be unsaturated at t .

The proof now follows noting that, by definition of a , $\min_{i \in S} u_{ai}(t) > \min_{i \in S} u_{1i}(t) \implies \exists$ atleast one item $i \in S$ such that $u_{ai}(t) > u_{1i}$. But following above chain of argument that leads to a contradiction unless the pair (i, a) is unsaturated at round t . \square

Combining Lem. 18 and 19 we can conclude that it does not take too long to reach to a time $T_0(\delta) > f(\delta)$, such that $\mathcal{C}_{T_0(\delta)} = \{1\}$ and thus $\mathcal{B}_t = \{1\}$ for all $t > T_0(\delta)$.

Lemma 20. Assume $\forall t > f(\delta)$, and $\forall i, j \in [n]$ $p_{ij} \in [l_{ij}(t), u_{ij}(t)]$, for some $\delta \in (0, 1)$. Then if we define $T_0(\delta)$ such that: $T_0(\delta) = \min\{t > f(\delta) \mid \mathcal{C}_t = \{1\}\}$, it can be upper bounded as $T_0(\delta) \leq 2f(\delta) + 2D \ln 2D$, where $D := \sum_{i < j} D_{ij}$.

Proof. The first observation for this is to note that: For any $t > f(\delta)$, $1 \in \mathcal{C}_t$ since $u_{1i} \geq p_{1i} > \frac{1}{2}$, $\forall i \in [n] \setminus \{1\}$. So, until $T_0(\delta)$, for all $t \in \{f(\delta), f(\delta) + 1, \dots, T_0(\delta) - 1\}$, $|\mathcal{C}_t| \geq 2$.

Secondly, for any $t \in \{f(\delta), f(\delta) + 1, \dots, T_0(\delta) - 1\}$, there exists atleast $\min(m, |\mathcal{C}_t| - 1)$ unsaturated pairs in S_t which gets updated. This holds from the following two case analyses:

Case 1: ($1 \in S_t$). This is the easy case since for any item $i \in \mathcal{C}_t \setminus \{1\}$, we know that $(1, i)$ is unsaturated from Lem. 18, and item 1 has to be updated for atleast $\min(m, |\mathcal{C}_t| - 1)$ many unsaturated pairs as follows from Lem. 17.

Case 2: ($1 \notin S_t$). From Lem. 19 we know that for any item $i \in S_t \cap \mathcal{C}_t \setminus \{1\}$ has to be unsaturated with atleast another item $j \in S_t$. Since *MaxMin-UCB* makes sure $|S_t| \geq m + 1$, again owing to Lem. 17, any item $i \in S_t \cap \mathcal{C}_t \setminus \{1\}$ gets compared for atleast m pairs out of which atleast one pair has to be unsaturated which proves the claim.

Moreover, as argued above, at any round $t \in \{f(\delta), f(\delta) + 1, \dots, T_0(\delta) - 1\}$, since $|\mathcal{C}_t| \geq 2$, any such round t updates atleast $\min(m, |\mathcal{C}_t| - 1) \geq 1$ unsaturated pair.

Thirdly, at any time t , if all pairs (i, j) , $i \neq j$, $i, j \in [n]$ are saturated, then $\mathcal{C}_t = \{1\}$.

So to bound $T_0(\delta)$, all we need to figure out is the worst possible number of iterations *MaxMin-UCB* would take to saturate all possible unsaturated pairs, precisely $\sum_{i < j} D_{ij} \ln t$ many pairwise updates. But as we argued before, since any round $t > f(\delta)$ updates atleast one unsaturated pair, we find that

$$T_0(\delta) = \min\{t > f(\delta) \mid t > f(\delta) + \sum_{i < j} D_{ij} \ln t\}$$

Now it is easy to see that the above inequality ($t > f(\delta) + \sum_{i < j} D_{ij} \ln t$) certainly satisfies for $t = 2f(\delta) + 2D \ln 2D$, where $D := \sum_{i < j} D_{ij}$ as:

$$\begin{aligned} f(\delta) + D \ln t &= f(\delta) + D \ln(2f(\delta) + 2D \ln 2D) \\ &\leq f(\delta) + D \ln(2D \ln 2D) + D \frac{2f(\delta)}{2D \ln 2D} \\ &\leq f(\delta) + D \ln(2D)^2 + f(\delta), \quad [\text{since, } \ln 2D > 1] \\ &= 2f(\delta) + 2D \ln 2D = t. \end{aligned}$$

Since $T_0(\delta)$ is the minimum time index at which $t > f(\delta) + D \ln t$ is satisfied, clearly $T_0(\delta) \leq 2f(\delta) + 2D \ln 2D$. \square

Finally we are ready to prove Thm. 5 based on the the following two claims:

Lemma 21. Assume $\forall t > f(\delta)$, and $\forall i, j \in [n]$ $p_{ij} \in [l_{ij}(t), u_{ij}(t)]$, for some $\delta \in (0, 1)$. For any time step $t > T_0(\delta)$, $1 \in S_t$ always. Moreover for any $|S_t| > 1$, item 1 gets compared with atleast m suboptimal items $a \in [n] \setminus \{1\}$.

Proof. For any $t > f(\delta)$, $1 \in \mathcal{C}_t$ since $u_{1i} \geq p_{1i} > \frac{1}{2}$, $\forall i \in [n] \setminus \{1\}$. Moreover as $T_0(\delta)$ ensures $\mathcal{C}_{T_0(\delta)} = \{1\}$, at this round, the algorithm set $\mathcal{B}_{T_0(\delta)} = \{1\}$. For the subsequent rounds $t > T_0(\delta)$, thus the algorithm continues setting $\mathcal{B}_t = \mathcal{B}_{t-1} \cap \mathcal{C}_t = \{1\}$.

Moreover note that for any $t > T_0(\delta)$, unless $|\mathcal{C}_t| = 1$, the algorithm always plays a set S_t such that $|S_t| = m + 1$, and in which item 1 always resides. Then by Lem. 17 we can conclude that item 1 is compared with atleast m distinct items at any round after pairwise *Rank-Breaking* update. \square

Lemma 22. *For any $\delta \in (0, 1)$, with probability atleast $(1 - \delta)$, the total cumulative regret of *MaxMin-UCB* is upper bounded as:*

$$R_T^1 \leq \left(2 \left[\frac{2\alpha n^2}{(2\alpha - 1)\delta} \right]^{\frac{1}{2\alpha - 1}} + 2D \ln 2D \right) \hat{\Delta}_{\max} + \frac{\ln T}{m + 1} \sum_{i=2}^n \hat{\Delta}_i (\mathbf{1}(m = 1)D_{1i} + \mathbf{1}(m > 1)D_{\max})$$

where recall that $\forall i \in [n] \setminus \{a^*\}$, $\hat{\Delta}_i = (\theta_1 - \theta_i)$, $\Delta_i = p_{1i} = \frac{\theta_1 - \theta_i}{2(\theta_1 + \theta_i)}$, $\hat{\Delta}_{\max} = \max_{i \in [n] \setminus \{1\}} \hat{\Delta}_i$, $D_{1i} = \frac{4\alpha}{\Delta_i^2}$, $D := \sum_{i < j} D_{ij}$, $D_{\max} = \max_{i \in [n] \setminus \{1\}} D_{1i}$.

Proof. Given Lem. 21 in place, the crucial observation now is to note that for any $t > T_0(\delta)$, *MaxMin-UCB*, always explores as long as there exists any suboptimal item $i \in [n] \setminus \{1\}$ such that the pair $(1, i)$ is unsaturated and thus $i \in \mathcal{C}_t$. In other words, our set building rule (*build_S*) always picks items from \mathcal{C}_t first before picking anything from $[n] \setminus \mathcal{C}_t$. However, any suboptimal item $i \in [n] \setminus \{1\}$ can belong to \mathcal{C}_t only if the pair $(1, i)$ is unsaturated, as follows from Lem. 18.

Thus for any time t , if the pair $(1, i)$ is already saturated (i.e. $n_{1i}(t) > D_{1i}(t) \ln t$), then $i \notin S_t$ unless item 1 is saturated with every suboptimal item in $[n] \setminus \{1\}$. But then $\mathcal{C}_t = \{1\}$ by Lem. 18 and the algorithm would go on playing $S_t = \{1\}$ until some pair $(1, i)$ gets unsaturated again. This argument holds true even for $t = T$.

Now lets try to analyse what is the maximum number of time an item $i \in [n] \setminus \{1\}$ can show up at any round post **Saturation** (i.e. for any $t > T_0(\delta)$). But since post **Saturation**, for any $t > T_0(\delta)$, $1 \in S_t$ always, the quantity $n_{1i}(T) - n_{1i}(T_0(\delta))$ is same as above. We hence analyse $n_{1i}(T) - n_{1i}(T_0(\delta))$ with the following two cases:

Case-1 ($m = 1$): This case is easy to analyse since at any round t , $|S_t| = 2$ and since $1 \in S_t$, so 1 gets compared with exactly one other suboptimal element $i \in [n] \setminus \{1\}$ at any t such that $i \in \mathcal{C}_t$. So clearly $n_{1i}(T) - n_{1i}(T_0(\delta)) \leq D_{1i} \ln T$ as by Lem. 18 after 1 is compared to i for $D_{1i} \ln T$ times $i \notin \mathcal{C}_t$ henceforth.

Case-2 ($m > 1$): In this case there are two possible ways i can show up in S_t : (i). If its unsaturated with 1 for which it can show up for at most $D_{1i} \ln T$ times as argued in Case-1, and (ii). When $i \notin \mathcal{C}_t$ but it shows up as a place holder for one of the $m + 1$ slot of S_t as long as some other element $j \in [n] \setminus \{1\}$, $i \neq j$ is unsaturated with 1 and $j \in \mathcal{C}_t$. But in the worst case once all item $i \in [n] \setminus \{1\}$ has appeared in S_t for $D_{\max} \ln T (> D_{1i} \ln T)$ times by Lem. 1 we have $u_{i1} < \frac{1}{2} \forall i \in [n] \setminus \{1\}$ and then \mathcal{C}_t has to be the singleton $\{1\}$ thereafter. So it has to be that $n_{1i}(T) - n_{1i}(T_0(\delta)) \leq D_{\max} \ln T$.

Finally note that all our above results holds good under the assumption that $\forall t > f(\delta)$, and $\forall i, j \in [n]$ $p_{ij} \in [l_{ij}(t), u_{ij}(t)]$, for some $\delta \in (0, 1)$, which itself holds good with probability atleast $(1 - \delta)$. Thus we have the maximum regret incurred by *MaxMin-UCB* in T rounds is

$$\begin{aligned} R_T &\leq T_0(\delta) \hat{\Delta}_{\max} + \frac{1}{m + 1} \sum_{i=2}^n (n_{1i}(T) - n_{1i}(T_0(\delta))) (\mathbf{1}(m = 1)D_{1i} + \mathbf{1}(m > 1)D_{\max}) \\ &= \left(2 \left[\frac{2\alpha n^2}{(2\alpha - 1)\delta} \right]^{\frac{1}{2\alpha - 1}} + 2D \ln 2D \right) \hat{\Delta}_{\max} + \frac{\ln T}{m + 1} \sum_{i=2}^n \hat{\Delta}_i (\mathbf{1}(m = 1)D_{1i} + \mathbf{1}(m > 1)D_{\max}), \end{aligned}$$

with probability atleast $(1 - \delta)$, were first term in the right hand side of the inequality holds since the maximum possible per trial regret that could be incurred by *MaxMin-UCB* in initial $T_0(\delta)$ rounds is $\hat{\Delta}_{\max}$. The proof now follows further upper bounding T_0 using Lem. 20. \square

This also concludes the proof of Thm. 5 using the exact value of $f(\delta)$ as derived in Lem. 15. \square

C.7 Proof of Theorem 6

Theorem 6. *The expected regret of MaxMin-UCB for Winner-regret with Top-m-ranking Feedback is: $\mathbf{E}[R_T^1] \leq \left(2 \left[\frac{2\alpha n^2}{(2\alpha-1)}\right]^{\frac{1}{2\alpha-1}} \frac{2\alpha-1}{\alpha-1} + 2D \ln 2D\right) \hat{\Delta}_{\max} + \frac{\ln T}{m+1} \sum_{i=2}^n (D_{\max} \hat{\Delta}_i)$, in T rounds.*

Proof. Recall from the statement of Thm. 5 that the only term in R_T^1 that depends on δ is $2f(\delta)$, where recall that $T_0(\delta) = 2f(\delta) + 2D \ln 2D$. Then by integrating $f(\delta)$ for δ from 0 to 1 as follows:

$$\int_0^1 f(\delta) d\delta = \int_0^1 \left[\frac{2\alpha n^2}{(2\alpha-1)\delta}\right]^{\frac{1}{2\alpha-1}} d\delta = \left[\frac{2\alpha n^2}{(2\alpha-1)}\right]^{\frac{1}{2\alpha-1}} \int_0^1 \left(\frac{1}{\delta}\right)^{\frac{1}{2\alpha-1}} d\delta = \left[\frac{2\alpha n^2}{(2\alpha-1)}\right]^{\frac{1}{2\alpha-1}} \frac{2\alpha-1}{2\alpha-2}$$

Thus expected regret $\mathbf{E}_\delta[R_T]$ can be upper bounded as:

$$\mathbf{E}_\delta[R_T^1] \leq \left(2 \left[\frac{2\alpha n^2}{(2\alpha-1)}\right]^{\frac{1}{2\alpha-1}} \frac{2\alpha-1}{\alpha-1} + 2D \ln 2D\right) \hat{\Delta}_{\max} + \frac{\ln T}{m+1} \sum_{i=2}^n \hat{\Delta}_i (\mathbf{1}(m=1)D_{1i} + \mathbf{1}(m>1)D_{\max}).$$

□

D Supplementary for Section 4

D.1 Proof of Thm. 7

Theorem 7 (Regret Lower Bound: Top- k -regret with Top- m -ranking Feedback). *For any No-regret learning algorithm \mathcal{A} for Top- k -regret that uses Top- k -ranking Feedback, and for any problem instance $\text{MNL}(n, \theta)$, the expected regret incurred by \mathcal{A} when run on it satisfies $\liminf_{T \rightarrow \infty} \mathbf{E}_\theta \left[\frac{R_T^k(\mathcal{A})}{\ln T} \right] \geq \frac{\theta_1 \theta_{k+1}}{\Delta_{(k)}} \frac{(n-k)}{k}$, where $\mathbf{E}_\theta[\cdot]$ denotes expectation under the algorithm and $\text{MNL}(n, \theta)$ model.*

Proof. The main idea lies in constructing ‘hard enough’ problem instances for which any No-regret algorithm has to incur $\Omega\left(\frac{n}{k\Delta_{(k)}} \ln T\right)$ regret.

We choose our *true problem instance* with MNL parameters $\theta^1 = (\theta_1^1, \dots, \theta_n^1)$, such that:

$$\begin{aligned} \text{True Instance: } \text{MNL}(n, \theta^1) : \theta_1^1 &= \theta_2^1 = \dots = \theta_{k-1}^1 = \theta + 2\epsilon; \\ \theta_n^1 &= \theta + \epsilon; \theta_{k+1}^1 = \theta_{k+2}^1 = \dots = \theta_{n-1}^1 = \theta. \end{aligned}$$

for some $\theta \in \mathbb{R}_+$ and $\epsilon > 0$. Clearly, the *Top- k Best-Items* (recall the definition from Def. 1, Sec. 2) of $\text{MNL}(n, \theta^1)$ is $S_{(k)}[1] = [k-1] \cup \{n\}$. Now for every $n-k$ suboptimal items $a \notin S_{(k)}[1]$, consider the altered problem instance $\text{MNL}(n, \theta^a)$ such that:

$$\text{Instance a: } \text{MNL}(n, \theta^a) : \theta_a^a = \theta + 2\epsilon; \theta_i^a = \theta_i^1, \forall i \in [n] \setminus \{a\}$$

And now the *Top- k Best-Items* of $\text{MNL}(n, \theta^a)$ is $S_{(k)}[a] = [k-1] \cup \{a\}$. Same as the case for proof of Thm. 3 or Thm. 4, we now again use the results of [21] (Lem. 13) for proving the lower bound. Precisely, the main trick lies in analyzing the KL-divergence terms for the above problem instances. For ease of analysis we first assume analyse the case with just the Winner Feedback. Borrowing same notations from Thm. 4, and denoting $x = |S \cap S_{(k)}[1]| - r$, $r = \mathbf{1}(n \in S)$, $y = k - (x + r)$, for any set $S \in S^a$, we now get that for any $i \in S$:

$$p_S^1(i) = \begin{cases} \frac{\theta+2\epsilon}{\theta_S^1} = \frac{\theta+2\epsilon}{x(\theta+2\epsilon)+r(\theta+\epsilon)+y\theta} = \frac{\theta+2\epsilon}{k\theta+\epsilon(2x+r)}, & \text{such that } i \in S_{(k)}[1] \cap S, \\ \frac{\theta+\epsilon}{\theta_S^1} = \frac{\theta+\epsilon}{x(\theta+2\epsilon)+r(\theta+\epsilon)+y\theta} = \frac{\theta+\epsilon}{k\theta+\epsilon(2x+r)}, & \text{such that } i = n, \\ \frac{\theta}{\theta_S^1} = \frac{\theta}{x(\theta+2\epsilon)+r(\theta+\epsilon)+y\theta} = \frac{\theta+\epsilon}{k\theta+\epsilon(2x+r)}, & \text{otherwise.} \end{cases}$$

On the other hand, for problem **Instance-a**, we have that:

$$p_S^a(i) = \begin{cases} \frac{\theta+2\epsilon}{\theta_S^a} = \frac{\theta+2\epsilon}{(x+1)(\theta+2\epsilon)+r(\theta+\epsilon)+(y-1)\theta} = \frac{\theta+2\epsilon}{k\theta+\epsilon(2(x+1)+r)}, & \text{such that } i \in (S \cap S_{(k)}[1]) \cup \{a\}, \\ \frac{\theta+\epsilon}{\theta_S^a} = \frac{\theta+\epsilon}{(x+1)(\theta+2\epsilon)+r(\theta+\epsilon)+(y-1)\theta} = \frac{\theta+\epsilon}{k\theta+\epsilon(2(x+1)+r)}, & \text{such that } i = n, \\ \frac{\theta}{\theta_S^a} = \frac{\theta}{(x+1)(\theta+2\epsilon)+r(\theta+\epsilon)+(y-1)\theta} = \frac{\theta+\epsilon}{k\theta+\epsilon(2(x+1)+r)}, & \text{otherwise.} \end{cases}$$

For ease of notation we denote $\theta_S = \theta_S^1$. Now using the following upper bound on $KL(\mathbf{p}, \mathbf{q}) \leq \sum_{x \in \mathcal{X}} \frac{p^2(x)}{q(x)} - 1$, \mathbf{p} and \mathbf{q} be two probability mass functions on the discrete random variable \mathcal{X} [33], we get for any $S \in S^a$:

$$\begin{aligned} KL(p_S^1, p_S^a) &\leq \sum_{i \in S} \left(\frac{\theta_i^1}{\theta_S^1} \right)^2 \left(\frac{\theta_S^a}{\theta_i^a} \right) - 1 \\ &= \sum_{i \in S \cap S_{(k)}[1]} \left(\frac{\theta+2\epsilon}{\theta_S} \right)^2 \left(\frac{\theta_S+2\epsilon}{\theta+2\epsilon} \right) \\ &\quad + \sum_{i \in S \cap (\{a, n\} \cup S_{(k)}[1])^c} \left(\frac{\theta}{\theta_S} \right)^2 \left(\frac{\theta_S+2\epsilon}{\theta} \right) + \left(\frac{\theta+\epsilon}{\theta_S} \right)^2 \left(\frac{\theta_S+2\epsilon}{\theta+\epsilon} \right) + \left(\frac{\theta}{\theta_S} \right)^2 \left(\frac{\theta_S+2\epsilon}{\theta+2\epsilon} \right) - 1 \\ &= \frac{\theta+2\epsilon}{\theta_S^2} \left[\theta_S + \frac{\theta^2}{\theta+2\epsilon} - \theta \right] = \frac{2\epsilon}{\theta_S} \left[1 - \frac{\theta}{\theta+2\epsilon} \right] - \frac{4\epsilon^2\theta}{\theta_S^2(\theta+2\epsilon)} \\ &\leq \frac{(2\epsilon)^2[\theta_S - \theta]}{\theta_S^2(\theta+2\epsilon)} \leq \frac{(2\epsilon)^2}{\theta_S(\theta+2\epsilon)} = \frac{4\epsilon^2}{[k\theta + (2x+r)\epsilon](\theta+2\epsilon)} \leq \frac{4\epsilon^2}{k\theta(\theta+2\epsilon)} \end{aligned} \quad (12)$$

Now coming back to the Top- k -ranking Feedback applying chain rule of KL-divergence (similar to the analysis of Eqn. (7)), we can write

$$KL(p_S^1, p_S^a) = KL(p_S^1(\sigma_1) + \dots + KL(p_S^1(\sigma_k \mid \sigma(1:k-1)), p_S^a(\sigma_k \mid \sigma(1:k-1))).$$

for any ranking $\sigma \in \Sigma_S^k$. And following the same argument that of (8), we further get

$$KL(p_S^1, p_S^a) \leq \frac{4k\epsilon^2}{k\theta(\theta+2\epsilon)}. \quad (13)$$

The rest of the proof follows exactly the same argument used in 4. We add the steps below for convenience. As before, considering $Z = \frac{N_{S_{(k)}[1]}(T)}{T}$, for large T , in this case we get:

$$\lim_{T \rightarrow \infty} \frac{kl(\mathbf{E}_{\theta^1}[Z], \mathbf{E}_{\theta^a}[Z])}{\ln T} \geq (1 - \alpha), \quad (14)$$

which follows from an exact similar analysis shown in the proof of Thm. 3 along with the facts that:

$$\begin{aligned} \mathbf{E}_{\theta^1}[N_{S_{(k)}[1]}(T)] &= 1 - o(T^\alpha) \quad \text{since } \mathcal{A} \text{ is assumed to be } No\text{-}regret \text{ and} \\ \mathbf{E}_{\theta^a}[N_{S_{(k)}[1]}(T)] &= o(T^\alpha). \end{aligned}$$

Then using the results of Eqn. (13) and (14) in Lem. 13, we further get:

$$\begin{aligned}
& \lim_{T \rightarrow \infty} \frac{1}{\ln T} \sum_{\{S \in S^a\}} \mathbf{E}_{\theta^1}[N_S(T)] KL(p_S^1, p_S^a) \geq (1 - \alpha) \\
& \implies \lim_{T \rightarrow \infty} \frac{1}{\ln T} \sum_{\{S \in S^a\}} \mathbf{E}_{\theta^1}[N_S(T)] \frac{4k\epsilon^2}{k\theta(\theta + 2\epsilon)} \geq (1 - \alpha) \\
& \implies \lim_{T \rightarrow \infty} \frac{1}{\ln T} \sum_{\{S \in S^a\}} \mathbf{E}_{\theta^1}[N_S(T)] \frac{\epsilon}{k} \geq (1 - \alpha) \frac{\theta(\theta + 2\epsilon)}{4k\epsilon}
\end{aligned} \tag{15}$$

Now applying (15) for each $n - k$ modified bandit **Instance- θ^a** (i.e. for each $a \in [n] \setminus S_{(k)}[1]$), we get:

$$\lim_{T \rightarrow \infty} \frac{1}{\ln T} \sum_{a=k+1}^n \sum_{\{S \in S^a\}} \mathbf{E}_{\theta^1}[N_S(T)] \frac{\epsilon}{k} \geq (1 - \alpha) \theta(\theta + 2\epsilon) \frac{(n - k)}{4k\epsilon} \tag{16}$$

Further recall from Eqn. (2.2) the expected regret of \mathcal{A} on problem instance $\text{MNL}(n, \theta^1)$ is given by:

$\mathbf{E}_{\theta^1}[R_T^k(\mathcal{A})] = \sum_{t=1}^T r_t^k = \sum_{t=1}^T \left(\frac{\sum_{i \in [k]} \theta_i^1 - \sum_{i \in S_t} \theta_i^1}{k} \right)$ which can be rewritten as:

$$\begin{aligned}
\mathbf{E}_{\theta^1}[R_T^k(\mathcal{A})] &= \mathbf{E}_{\theta^1} \left[\sum_{t=1}^T \left(\frac{\sum_{i \in [k]} \theta_i^1 - \sum_{i \in S_t} \theta_i^1}{k} \right) \right] \\
&\geq \mathbf{E}_{\theta^1} \left[\sum_{t=1}^T \left(\frac{\sum_{i \in [k]} \theta_k^1 - \sum_{i \in S_t} \theta_i^1}{k} \right) \right] \\
&= \mathbf{E}_{\theta^1} \left[\sum_{t=1}^T \left(\sum_{i \in [S_t]} \frac{\theta_k^1 - \theta_i^1}{k} \right) \right] \\
&= \mathbf{E}_{\theta^1} \left[\sum_{t=1}^T \sum_{S \in A} \mathbf{1}(S_t = S) \sum_{a=k+1}^n \mathbf{1}(a \in S) \frac{(\theta_k^1 - \theta_a^1)}{|S_t|} \right] \\
&= \mathbf{E}_{\theta^1} \left[\sum_{a=k+1}^n \sum_{t=1}^T \sum_{S \in A} \mathbf{1}(S_t = S) \mathbf{1}(a \in S) \frac{((\theta + \epsilon) - \theta)}{k} \right] \quad (\text{since } \theta_k^1 = \theta_n^1 = \theta + \epsilon) \\
&= \sum_{a=k+1}^n \sum_{t=1}^T \mathbf{E}_{\theta^1} \left[\sum_{S \in A} \mathbf{1}(S_t = S) \mathbf{1}(a \in S) \frac{\epsilon}{k} \right] \\
&= \sum_{a=k+1}^n \sum_{S \in A} \mathbf{E}_{\theta^1} \left[\sum_{t=1}^T \mathbf{1}(S_t = S) \mathbf{1}(a \in S) \frac{\epsilon}{k} \right] \\
&= \sum_{a=k+1}^n \sum_{S \in A} \left[\mathbf{E}_{\theta^1}[N_S(T)] \mathbf{1}(a \in S) \frac{\epsilon}{k} \right] \\
&= \sum_{a=k+1}^n \sum_{\{S \in A | a \in S\}} \mathbf{E}_{\theta^1}[N_S(T)] \frac{\epsilon}{k}
\end{aligned} \tag{17}$$

Using above combined with (16) we get:

$$\begin{aligned} \lim_{T \rightarrow \infty} \frac{1}{\ln T} \mathbf{E}_{\theta^1} [R_T^k(\mathcal{A})] &\geq \lim_{T \rightarrow \infty} \frac{1}{\ln T} \sum_{a=k+1}^n \sum_{\{S \in \mathcal{A} \mid a \in S\}} \mathbf{E}_{\theta^1} [N_S(T)] \frac{\epsilon}{k} \\ &\geq \lim_{T \rightarrow \infty} \frac{1}{\ln T} \sum_{a=k+1}^n \sum_{\{S \in \mathcal{S}^a\}} \mathbf{E}_{\theta^1} [N_S(T)] \frac{\epsilon}{k} \geq (1 - \alpha) \theta (\theta + 2\epsilon) \frac{(n - k)}{4k\epsilon}. \end{aligned}$$

Now since α is a fixed constant in $(0, 1]$, we thus prove the existence of a $\text{MNL}(n, \theta)$ problem instance (precisely $\text{MNL}(n, \theta^1)$), such that for large T , $\mathbf{E}_{\theta^1} [R_T^1] = \Omega \left(\frac{\theta_1 \theta_{k+1}}{\Delta_{(k)}} \frac{(n-k)}{k} \ln T \right)$ (noting that for instance $\text{MNL}(n, \theta^1)$, $\Delta_{(k)} = \epsilon$), which concludes the proof \square

D.2 Algorithm pseudocode: Top- k -regret

Algorithm 3 Rec-MaxMin-UCB

```

1: init:  $\alpha > 0.5$ ,  $\mathbf{W} \leftarrow [0]_{n \times n}$ ,  $\mathcal{B}_0 \leftarrow [\emptyset]_k$ 
2: for  $t = 1, 2, 3, \dots, T$  do
3:   Set  $I \leftarrow [n]$ ,  $\mathbf{N} = \mathbf{W} + \mathbf{W}^\top$ , and  $\hat{\mathbf{P}} = \frac{\mathbf{W}}{\mathbf{N}}$ .  $N = [n_{ij}]_{n \times n}$  and  $\hat{P} = [\hat{p}_{ij}]_{n \times n}$ .
4:   Define  $u_{ij} = \hat{p}_{ij} + \sqrt{\frac{\alpha \ln t}{n_{ij}}}$ ,  $\forall i, j \in [n], i \neq j$ ,  $u_{ii} = \frac{1}{2}$ ,  $\forall i \in [n]$ .  $\mathbf{U} = [u_{ij}]_{n \times n}$ 
5:   for  $h = 1, 2, \dots, k - 1$  do
6:      $\mathcal{C}_t^h \leftarrow \{i \in I \mid u_{ij} > \frac{1}{2}, \forall j \in I \setminus \{i\}\}$ ;  $\mathcal{B}_t(h) \leftarrow \mathcal{C}_t^h \cap \mathcal{B}_{t-1}(h)$ 
7:     if  $\mathcal{B}_t(h) \neq \emptyset$ , then set  $I \leftarrow I \setminus \mathcal{B}_t(h)$  and  $S_t \leftarrow S_t \cup \mathcal{B}_t(h)$ 
8:     if  $\mathcal{B}_t(h) = \emptyset$  then
9:       if  $\mathcal{C}_t^h = \emptyset$  then  $\mathcal{C}_t^h \leftarrow I$ ; elseif  $|\mathcal{C}_t^h| > 0$  set  $\mathcal{B}_t(h) \leftarrow \text{build\_S}(\mathbf{U}, S_t, I, 1)$ 
10:       $S_t \leftarrow S_t \cup \text{build\_S}(\mathbf{U}, S_t, I, k - |\mathcal{B}_t(h)|)$ . Exit and Goto Line 15)
11:   end if
12:   end for
13:    $\mathcal{C}_t^k \leftarrow \{i \in I \mid u_{ij} > \frac{1}{2}, \forall j \in I \setminus \{i\}\}$ ;  $\mathcal{B}_t(k) \leftarrow \mathcal{C}_t^k \cap \mathcal{B}_{t-1}(k)$ 
14:   if  $|\mathcal{C}_t^k| = 1$ , then  $\mathcal{B}_t(k) \leftarrow \mathcal{C}_t^k$ , and set  $S_t \leftarrow S_t \cup \mathcal{C}_t^k$ ; else  $S_t \leftarrow S_t \cup \text{build\_S}(\mathbf{U}, S_t, I, 1)$ 
15:   Play  $S_t$ , and receive:  $\sigma_t \in \Sigma_{S_t}^k$ 
16:    $W(\sigma_t(k'), i) \leftarrow W(\sigma_t(k'), i) + 1 \quad \forall i \in S_t \setminus \sigma_t(1 : k')$  for all  $k' = 1, 2, \dots, k$ 
17: end for

```

D.3 Proof of Thm. 8

Theorem 8 (Rec-MaxMin-UCB: High Probability Regret bound). *Given a fixed time horizon T and $\delta \in (0, 1)$, with high probability $(1 - \delta)$, the regret incurred by Rec-MaxMin-UCB for Top- k -regret admits the bound $R_T^k \leq \left(2 \left[\frac{2\alpha n^2}{(2\alpha - 1)\delta} \right]^{\frac{1}{2\alpha - 1}} + 2\bar{D}^{(k)} \ln(2\bar{D}^{(k)}) \right) \Delta'_{\max} + \frac{4\alpha \ln T}{k} \left(\sum_{b=k+1}^n \frac{(\theta_k - \theta_b)}{\bar{D}^2} \right)$, where $D^{(k)}$ is an instance dependent constant (see Lem. 26, Appendix), $\Delta'_{\max} = \frac{(\sum_{i=1}^k \theta_i - \sum_{i=n-k+1}^n \theta_i)}{k}$, and $\hat{D} = \min_{g \in [k-1]} (p_{kg} - p_{bg})$.*

Proof. For ease of analysis we assume $\theta_1 \geq \theta_2 \geq \dots \theta_k > \theta_{k+1} \geq \dots \geq \theta_n$ and hence $S_{(k)} = [k]$.

We use the same notations as introduced in the proof of Thm. 5. Note that Lem. 15 holds in this case as well. So that concludes the **Random-Exploration** phase.

Analysis of the Progress phase: We next proceed to analyse the **Progress** phase from round $f(\delta) + 1$ to $T_0(\delta)$, where $T_0(\delta)$ is defined to be such that

$$T_0(\delta) = \arg \min_{t > f(\delta)} \mathcal{B}_t = [k]$$

The goal of this phase is to show that the length of interval $[f(\delta) + 1, T_0(\delta)]$ is ‘small’, precisely $T_0(\delta) \leq 2f(\delta) + 2\bar{D}^{(k)} \ln(2\bar{D}^{(k)})$, where $\bar{D}^{(k)} := \sum_{r=1}^k D^{(r)}$, and $D^{(r)} := \sum_{i \in [n]} \sum_{j \in W(r)} D_{ij}^r$ (see Lem. 26). Note here $\bar{D}^{(k)}$ is a problem dependent constant, independent of T .

Notations: We first define few notations for ease of analysis: Let us first define the set of items $W(i) = \{j \in [k] \mid \theta_j > \theta_i\}$ be the set of items in the *Top-k Best-Items* strictly better than i , and $Z(i) = \{j \in [k] \mid \theta_j < \theta_i\}$ be the set of items in $[k]$ worse than item i , for some $i \in [k]$.

For any $g \in [k]$,

$$D_{ij}^g = \begin{cases} \frac{4\alpha}{\min((p_{gi}-\frac{1}{2})^2, (p_{gj}-\frac{1}{2})^2)} & \text{if } i, j \in Z(g) \\ \frac{4\alpha}{(p_{gi}-p_{ji})^2} & \text{if } i \in W(g) \text{ and } j \in Z(g) \\ D_{gj}^g & \text{if } \theta_i = \theta_g \end{cases}, D_{ij}^g = D_{ji}^g$$

for any pair $(i, j) \in [n] \times [n]$, $i \neq g, j \neq g$. and $D_{gi}^g = \frac{4\alpha}{(p_{gi}-\frac{1}{2})^2}$, where $p_{ij} = \frac{1}{2} + \frac{\theta_i - \theta_j}{2(\theta_i + \theta_j)}$ for all $i, j \in [n]$.

Towards analysing the **Progress** phase we first make the following key observations:

- **Observation 1:** At any round $t \in [T]$, $\mathcal{B}_t(i)$ is either singleton or an empty set, for all $i \in [k]$, which follows by the construction of $\mathcal{B}_t(i)$.
- **Observation 2:** For any item $i \in [k]$ in the *Top-k Best-Items*, at any round $t > f(\delta)$ if $j \in \mathcal{B}_t(|W(j)| + 1 : k - |Z(j)|)$ for all $j \in W(i)$, and $i \in \mathcal{B}_t$ such that $\mathcal{B}_t(x) = i$ for some $x \in \{|W(i)| + 1 \dots k - |Z(i)|\}$ and for any $x' \in \{|W(i)| + 1, \dots, (x-1)\}$, $\theta_{\mathcal{B}_t(x')} = \theta_i$, then $\mathcal{B}_{t'}(x) = \{i\} \forall t' > t$ —in other words i will continue to reside in slot $\mathcal{B}_{t'}(x)$ for any $t' > t$.

We next define another notation $T_0^i(\delta)$ for any $i \in [k]$ such that $T_0^i(\delta) = \arg \min_{t > f(\delta)} \{\forall j \in W(i), j \in \mathcal{B}_t(|W(j)| + 1 : k - |Z(j)|), i = \mathcal{B}_t(x), \text{ and } \forall x' \in \{|W(i)| + 1, \dots, (x-1)\}, \theta_{\mathcal{B}_t(x')} = \theta_i\}$. Clearly $\max_{i \in [k]} T_0^i(\delta) = T_0(\delta)$ as defined above.

- **Observation 3:** $T_0(\delta)$ is not far from $f(\delta)$ —we prove this in a stepwise manner, to explain it in an intuitive level assume $\theta_1 > \dots > \theta_k$. Then we first show that $T_0^1(\delta)$ is bounded. Once item 1 is secured in its slot $\mathcal{B}_{T_0^1(\delta)}(1)$, we proceed to bound $T_0^2(\delta)$, and so on till $T_0^k(\delta) = T_0(\delta)$ (see Lem. 26 for the formal details which holds due to Lem. 24 and 25).

We find it convenient to define one more definition before proving Lem. 26:

Definition 23 (*g-Unsaturated Pairs*). At any time $t \in [T]$, for any item $g \in [k]$ and any pair of two distinct pair of items $i, j \in [n]$, we call the pair (i, j) to be *g-unsaturated* at time t if $n_{ij}(t) \leq D_{ij}^g \ln t$. Otherwise, we call the pair *g-saturated* at t .

Lemma 24. Assuming $\forall t > f(\delta)$, and $\forall i, j \in [n]$ $p_{ij} \in [l_{ij}(t), u_{ij}(t)]$, for some $\delta \in (0, 1)$: At any iteration $t > f(\delta)$, for any $g \in [k]$, if \exists an item $i \in Z(g)$, i.e. $\theta_g > \theta_i$ and both $i, j \in \mathcal{C}_t^h$ for some $h \in [k]$, then the pair (g, i) is *g-unsaturated* at t .

Proof. By assumption $\forall i, j \in [n]$ $p_{ij} \leq u_{ij}(t)$. Now if $\exists h \in [k]$ such that a pair (g, i) such that both $i, g \in \mathcal{C}_t^h$, then it has to be the case that $u_{gi} > \frac{1}{2}$ and $u_{ig} > \frac{1}{2}$.

But then suppose (g, i) was indeed *g-saturated* at time t , i.e. $n_{1i}(t) > D_{gi}^g \ln t$, this implies:

$$u_{ig}(t) = \hat{p}_{ig}(t) + c_{ig}(t) \leq p_{ig}(t) + 2c_{ig}(t) = p_{ig}(t) + \Delta_i^g = \frac{1}{2},$$

which implies there cannot exist $i \notin \mathcal{C}_t^h$ if $g \in \mathcal{C}_t^h$ for any h , which leads to a contradiction. Hence the pair (g, i) must be *g-unsaturated* at t . \square

Lemma 25. Assuming $\forall t > f(\delta)$, and $\forall i, j \in [n] \ p_{ij} \in [l_{ij}(t), u_{ij}(t)]$, for some $\delta \in (0, 1)$. Consider any $g \in [k]$. At any iteration $t > f(\delta)$, for any set $S_t \not\supseteq g$, $0 \leq |S_t| < k$ if \exists an item $a \in Z(g)$, i.e. $\theta_g > \theta_a$, such that $a = \arg \max_{c \in I \setminus S_t} [\min_{i \in S_t} u_{ci}(t)]$, then \exists atleast one item $b \in S$ such that the pair (a, b) is unsaturated at t .

Proof. Firstly the important observation to make is at any round t , and in any of its sub-phase $h \in [k]$, our set building rule ensures that $u_{ji} > \frac{1}{2}$ for $j \in S_t$ and $i \notin S_t$.

Moreover since $a = \arg \max_{c \in I \setminus S_t} [\min_{i \in S_t} u_{ci}(t)]$ and $a \notin S_t$, there must exist an item b in S such that $u_{ab}(t) > u_{gb}(t)$ as otherwise g would have been picked instead of a . But following the argument above we also know that $u_{ba}(t) > \frac{1}{2}$. Now b can fall into the following three categories:

Case-1 $[b \in Z(g)]$: We first note that:

$$u_{ba}(t) - l_{ba}(t) = u_{ba}(t) + u_{ab}(t) - 1 > \frac{1}{2} + u_{gb}(t) - 1 > \frac{1}{2} + p_{gb}(t) - 1 = p_{gb} - \frac{1}{2},$$

but on the other hand if the pair (a, b) is indeed g -saturated at t , i.e. $n_{ba}(t) > D_{ba}^g \ln t$. Then we have that

$$u_{ba}(t) - l_{ba}(t) = 2c_{ba}(t) \leq \sqrt{\min \left(\left(p_{gb} - \frac{1}{2} \right)^2, \left(p_{ga} - \frac{1}{2} \right)^2 \right)} \leq \left(p_{gb} - \frac{1}{2} \right).$$

Case-2 $[b \in W(g)]$:

In this case suppose if the pair (a, b) is indeed g -saturated at t , i.e. $n_{ba}(t) > D_{ba}^g \ln t$ we have

$$2c_{ab}(t) \leq \sqrt{\left(p_{gb} - p_{ab} \right)^2} \leq \left(p_{gb} - p_{ab} \right).$$

It is important to note that the right hand side of the above inequality is positive since for this case $\theta_b > \theta_g > \theta_a$. But this implies

$$u_{ab}(t) \leq p_{ab}(t) + 2c_{ab}(t) = p_{gb} + 2c_{ab}(t) - (p_{gb} - p_{ab}) \leq p_{gb} < u_{gb}(t)$$

which leads to a contradiction again.

Case-3 $[b : \theta_b = \theta_g]$: The analysis in this case goes similar to **Case-2** above which finally leads to the contradiction that $u_{ab} < p_{gb} = \frac{1}{2} < u_{gb}(t)$.

Hence combining the above three cases, it follows that the unless the pair (b, a) is g -unsaturated at round t , $a \in Z(g)$ can not show up prior to g . \square

Assumption: Recall we assumed $\theta_1 \geq \theta_2 \geq \dots \geq \theta_k$. For ease of explanation (without loss of generality by relabelling the items) we also assume that at any time t , for any pair of items (i, j) such that $i, j \in [k]$, $\theta_i = \theta_j$, $i < j$, and if it happens to be the case that both $i, j \in \mathcal{B}_t$, with $\mathcal{B}_t(x) = i$, $\mathcal{B}_t(y) = j$ then $x < y$.

Lemma 26. Assume $\forall t > f(\delta)$, and $\forall i, j \in [n] \ p_{ij} \in [l_{ij}(t), u_{ij}(t)]$, for some $\delta \in (0, 1)$. Then if we define $T_0(\delta)$ such that: $T_0(\delta) = \min\{t > f(\delta) \mid \mathcal{B}_t = [k]\}$, it can be upper bounded as $T_0(\delta) \leq 2f(\delta) + 2\bar{D}^{(k)} \ln(2\bar{D}^{(k)})$, where $\bar{D}^{(k)} := \sum_{r=1}^k D^{(r)}$, and $D^{(r)} := \sum_{i \in [n]} \sum_{j \in W(r)} D_{ij}^r$ (recall the rest of the notations as defined above).

Proof. Combining Lem. 24 and 25 we first aim to bound the term $T_0^1(\delta)$ such that the first time after $T_0(\delta)$, when $\mathcal{B}_t(1) = 1$ and post which it follows that $\mathcal{B}_t = \{1\}$ for all $t > T_0(\delta)$.

Bounding $T_0^1(\delta)$: Note that for any $t > f(\delta)$, $Z(1) = n - 1$ always. So the only way 1 can miss the $\mathcal{B}_t(1)$ slot is if $\exists i \neq 1$ and $\theta_1 > \theta_i$ (due to the relabelling **Assumption** above) which occupies $\mathcal{B}_t(1)$. All we need to figure out is the worst possible number of rounds *Rec-MaxMin-UCB* would take to 1-saturate all the pairs, precisely $\sum_{i < j} D_{ij}^1 \ln t$ many pairwise updates should be done within t rounds. Now using a similar chain of argument given in Lem. 20 (along with Lem. 24 and 25), since any round $t > f(\delta)$ updates atleast one 1-unsaturated pair, we find that

$$T_0^1(\delta) = \min\{t > f(\delta) \mid t > f(\delta) + D^{(1)} \ln t\},$$

where $D^{(1)} := \sum_{i \in [n]} \sum_{j \in W(1)} D_{ij}^1$

Bounding $T_0^2(\delta)$: Note that once $1 \in \mathcal{B}_t(1)$, for any $t > T_0^1(\delta)$, $1 \in S_t$. And also either $\theta_2 = \theta_1$ or $\theta_2 < \theta_1$. But in either case $2 \in \mathcal{C}_t^2$, as $u_{2j} > \frac{1}{2}$ for all $j \in [n] \setminus \{1\}$. Then the only way the one can stop 2 occupying the slot $\mathcal{B}_t(2)$ is if there exists some other item $i \neq 2$ and $\theta_2 > \theta_i$ (due to the relabelling **Assumption** above) which occupies $\mathcal{B}_t(2)$. But then the algorithm picks $\mathcal{B}_t(2)$ in S_t , i.e. $i \in S_t$ alongside 1 and it get compared with 1 at each round it is picked. Moreover the last element of S_t is always picked by the *build_S* subroutine, so following the three case analyses of Lem. 25, the maximum number of rounds till which 2 can miss the slot $\mathcal{B}_t(2)$ is

$$T_0^2(\delta) = \min\{t \mid t > T_0^1(\delta) + D^{(2)} \ln t\},$$

where $D^{(2)} := \sum_{i \in [n]} \sum_{j \in W(2)} D_{ij}^2$.

Following the same argument we can state a general result that for any $r \in [k] \setminus \{1\}$:

Bounding $T_0^r(\delta)$ where $D^{(r)} := \sum_{i \in [n]} \sum_{j \in W(r)} D_{ij}^r$:

$$T_0^r(\delta) = \min\{t \mid t > T_0^{r-1}(\delta) + D^{(r)} \ln t\}$$

Then combining above for $r = 1, \dots, k$ we get $T_0(\delta) = T_0^k(\delta)$ should be such that

$$T_0^k(\delta) = \min\{t \mid t > f(\delta) + \sum_{r=1}^k D^{(r)} \ln t\}$$

And now following the exact same analysis of Lem. 20, it is easy to see that the above inequality satisfies for $t = 2f(\delta) + 2 \sum_{r=1}^k D^{(r)} \ln \left(\sum_{r=1}^k D^{(r)} \right)$. So that bounds $T_0(\delta) \leq 2f(\delta) + 2 \sum_{r=1}^k D^{(r)} \ln \left(\sum_{r=1}^k D^{(r)} \right)$. □

Analysis for Saturation phase ($t > T_0(\delta)$):

- **Observation 4:** After $T_0(\delta)$, now $\mathcal{B}_t = S_{(k)} = [k]$, with $\mathcal{B}_t(k) = \theta_k$, and thereafter i.e. $\mathcal{B}_t = [k]$ for all $t > T_0(\delta)$. Note that in this phase the algorithm always plays either $S_t = [k]$ or it plays $S_t = [k-1] \cup \{b\}$, $b \notin [k]$. So for any $t > T_0(\delta)$, $S_t \cap \mathcal{B}_t = [k-1]$. Then at any time t is a suboptimal item $b \in [n] \setminus [k]$ comes in S_t then it gets compared to all items in $[k-1]$ (owing to *Rank-Breaking*). But can not happen for too long and after a time $\mathcal{C}_t \cap \{b\} = \emptyset$, when the algorithm *Rec-MaxMin-UCB* will not play b any more. This holds true for any $b \in [n] \setminus [k]$, for which the algorithm will left with no other choice for S_t other than $S_t = [k]$ when it incurs no regret. See Lem. 27 and 28 for the formal claims.

Lemma 27. Assume $\mathcal{B}_t = S_{(k)} = [k]$ for all $t > T_0(\delta)$. Then the total cumulative regret of *Rec-MaxMin-UCB* post $T_0(\delta)$ is upper bounded by:

$$R_T^k(T_0(\delta) : T) := \sum_{t=T_0(\delta)}^T r_t^k \leq \frac{4\alpha \ln T}{k} \left(\sum_{b=k+1}^n \frac{(\theta_k - \theta_b)}{\hat{D}^2} \right),$$

where $\hat{D} = \min_{g \in [k-1]} (p_{kg} - p_{bg})$.

Proof. Note that when $\mathcal{B}_t = S_{(k)} = [k]$, at any such round *Rec-MaxMin-UCB* plays the set S_t such that the first $(k-1)$ items of \mathcal{B}_t are always included in S_t , i.e. $|S_t \cap \mathcal{B}_t(1 : k-1)| = k-1$. The k^{th} element of \mathcal{B}_t only gets replaced by a suboptimal element $b \in [n] \setminus [k]$ only if $\exists g \in \mathcal{B}_t(1 : k-1)$ such that the pair (g, b) is *unsaturated*, in a sense that $u_{bg} > u_{kg}$, and hence b got picked by the algorithm instead of k (in Line 14).

But is that possible for long? Precisely, we now show that any such suboptimal item $b \in [n] \setminus [k]$ can not get selected by the algorithm for more than $\frac{4\alpha \ln T}{\bar{D}^2}$ times. This is since for any $S_t = [k-1] \cup \{b\}$ (recall that $\mathcal{B}_t(1 : k-1) = [k-1]$), once played for $\frac{4\alpha \ln T}{\bar{D}^2}$ times (say this happens at time $t = \tau$), we know that $\forall g \in [k-1]$ such that number of times the pair (g, b) gets updated is exactly $\frac{4\alpha \ln T}{\bar{D}^2}$ too due to *Rank-Breaking* on Top- k -ranking Feedback. But this implies for any $g \in [k-1]$,

$$u_{bg}(\tau) \leq p_{bg} + 2c_{bg}(\tau) = p_{kg} + 2c_{bg}(\tau) - (p_{kg} - p_{bg}) \leq p_{kg} \leq u_{kg}(\tau) \quad (18)$$

where the first and last inequality follows by definition of u_{bg} and Lem. 15, the second last inequality follows due to the fact that since $n_{gb}(\tau) \geq \frac{4\alpha \ln T}{\bar{D}^2}$

$$2c_{gb}(\tau) \leq 2\sqrt{\frac{\alpha \ln T}{\frac{4\alpha \ln T}{\bar{D}^2}}} = \hat{D} \leq (p_{kg} - p_{bg}),$$

since by definition $\hat{D} = \min_{g \in [k-1]} (p_{kg} - p_{bg})$. Then Eqn. 18 leads to a contradiction showing $u_{bg}(\tau) < u_{kg}(\tau) \implies b$ can not replace k at any round $t > \tau$.

The rest of the analysis simply follows from the fact that since any $b \in [n] \setminus [k]$ can appear for only $\left(\frac{4\alpha \ln T}{\bar{D}^2}\right)$ times and it replaces the item $\mathcal{B}_t(k) = k$, hence the cost incurred for b is $\frac{(\theta_k - \theta_b)}{k}$ (by Eqn.

2.2). Thus the total regret incurred in saturation phase is $\frac{4\alpha \ln T}{k} \left(\sum_{b=k+1}^n \frac{(\theta_k - \theta_b)}{\bar{D}^2} \right)$. \square

Lemma 28. For any $\delta \in (0, 1)$, with probability at least $(1 - \delta)$, the total cumulative regret of *Rec-MaxMin-UCB* can be upper bounded by:

$$R_T^k \leq (2f(\delta) + 2\bar{D}^{(k)} \ln(2\bar{D}^{(k)}))\Delta'_{\max} + \frac{4\alpha \ln T}{k} \left(\sum_{b=k+1}^n \frac{(\theta_k - \theta_b)}{\hat{D}^2} \right)$$

where $\Delta'_{\max} = \frac{(\sum_{i=1}^k \theta_i - \sum_{i=n-k+1}^n \theta_i)}{k}$, and $f(\delta)$, $\bar{D}^{(k)}$ and \hat{D} is as defined in Lem. 15, 26, 27.

Proof. This can be proved just by combining the claims of Lem. 26 and 27. Note from Lem. 26 that till the **Progress** phase $T_0(\delta)$, the algorithm can play any arbitrary sets S_t for which the maximum regret incurred can be $\Delta'_{\max} = \frac{(\sum_{i=1}^k \theta_i - \sum_{i=n-k+1}^n \theta_i)}{k}$. Thereafter the algorithm enters into **Saturation** phase at which the maximum regret in can incur is $\frac{4\alpha \ln T}{k} \left(\sum_{b=k+1}^n \frac{(\theta_k - \theta_b)}{\bar{D}^2} \right)$ as follows from Lem. 28, which concludes the proof. \square

The entire analysis above thus concludes the proof of Thm. 8. \square

Proof of Thm. 9

Theorem 9. The expected regret incurred by *MaxMin-UCB* for Top- k -regret is:

$$\mathbf{E}[R_T^1] \leq \left(2 \left[\frac{2\alpha n^2}{(2\alpha - 1)} \right]^{\frac{1}{2\alpha - 1}} \frac{2\alpha - 1}{\alpha - 1} + 2\bar{D}^{(k)} \ln(2\bar{D}^{(k)}) \right) \Delta'_{\max} + \frac{4\alpha \ln T}{k} \left(\sum_{b=k+1}^n \frac{(\theta_k - \theta_b)}{\hat{D}^2} \right).$$

Proof. The proof essentially follows same as the proof of Thm. 6 by integrating the δ dependent term $\left[\frac{2\alpha n^2}{(2\alpha - 1)\delta} \right]^{\frac{1}{2\alpha - 1}}$ in R_T^k (see Thm. 8) from $\delta = 0$ to ∞ . \square

E Experiment Details

We report numerical results of the proposed algorithms run on the following $\text{MNL}(n, \theta)$ models:

MNL(n, θ) Environments. 1. *g1*, 2. *g4*, 3. *arith*, 4. *geo* all with $n = 16$ and two larger models 5. *arith-big*, and 6. *geo-big* each with $n = 50$ items. Their individual score parameters are as follows: **1. g1:** $\theta_1 = 0.8$, $\theta_i = 0.2$, $\forall i \in [16] \setminus \{1\}$ **2. g4:** $\theta_1 = 1$, $\theta_i = 0.7$, $\forall i \in \{2, \dots, 6\}$, $\theta_i = 0.5$, $\forall i \in \{7, \dots, 11\}$, and $\theta_i = 0.01$ otherwise. **3. arith:** $\theta_1 = 1$ and $\theta_i - \theta_{i+1} = 0.06$, $\forall i \in [15]$. **4. geo:** $\theta_1 = 1$, and $\frac{\theta_{i+1}}{\theta_i} = 0.8$, $\forall i \in [15]$. **5. har:** $\theta_1 = 1$ and $\theta_i = 1 - \frac{1}{i}$, $\forall i \in \{2, 3, \dots, 16\}$. **6. arithb:** $\theta_1 = 1$ and $\theta_i - \theta_{i+1} = 0.02$, $\forall i \in [49]$. **7. geob:** $\theta_1 = 1$, and $\frac{\theta_{i+1}}{\theta_i} = 0.9$, $\forall i \in [49]$.