

# FINALLY RANK-BREAKING CONQUERS MNL BANDITS: OPTIMAL AND EFFICIENT ALGORITHMS FOR MNL ASSORTMENT

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

We address the problem of active online assortment optimization problem with preference feedback, which is a framework for modeling user choices and subsetwise utility maximization. The framework is useful in various real-world applications including ad placement, online retail, recommender systems, and fine-tuning language models, amongst many others. The problem, although has been studied in the past, lacks an intuitive and practical solution approach with simultaneously efficient algorithm and optimal regret guarantee. E.g., popularly used assortment selection algorithms often require the presence of a ‘strong reference’ which is always included in the choice sets, further they are also designed to offer the same assortments repeatedly until the reference item gets selected—all such requirements are quite unrealistic for practical applications. In this paper, we designed efficient algorithms for the problem of regret minimization in assortment selection with *Multinomial Logit* (MNL) based user choices. We designed a novel concentration guarantee for estimating the score parameters of the PL model using ‘*Pairwise Rank-Breaking*’, which builds the foundation of our proposed algorithms. Moreover, our methods are practical, provably optimal, and devoid of the aforementioned limitations. Empirical evaluations corroborate our findings and outperform the existing baselines.

## 1 INTRODUCTION

Studies have shown that it is often easier, faster and less expensive to collect feedback on a relative scale rather than asking ratings on an absolute scale. E.g., to understand the liking for a given pair of items, say (A,B), it is easier for the users to answer preference-based queries like: “Do you prefer Item A over B?”, rather than their absolute counterparts: “How much do you score items A and B in a scale of [0-10]?” (Musallam et al., 2004; Kahneman & Tversky, 1982). Due to the widespread applicability and ease of data collection with relative feedback, learning from preferences has gained much popularity in the machine-learning community, especially the active learning literature which has applications in Medical surveys, AI tutoring systems, Multi-player sports/games, or any real-world systems that have ways to collect feedback in terms of preferences. The problem is famously studied as the *Dueling-Bandit* (DB) problem in the active learning community Yue et al. (2012); Ailon et al. (2014); Zoghi et al. (2014a;b; 2015), which is an online learning framework for identifying a set of ‘good’ items from a fixed decision-space (set of items) by querying preference feedback of actively chosen item-pairs. Consequently, the generalization of Dueling-Bandits, with *subset-wise* preferences has also been developed into an active field of research. For instance, applications like Web search, language models, online shopping, recommender systems (e.g. Youtube, Netflix, Google News/Maps, Spotify) typically involve users expressing preferences by choosing one result (or a handful of results) from a subset of offered items and often the objective of the system is to identify the ‘most-profitable’ subset to offer to their users. The problem, popularly termed as ‘Assortment Optimization’ is studied in many interdisciplinary literature, e.g. Online learning and bandits Bengs et al. (2021a), Operations research Talluri & Van Ryzin (2004); Agrawal et al. (2019), Game theory Chatterji et al. (2021), RLHF Christiano et al. (2017); Ouyang et al. (2022), to name a few.

**Problem (Informal): Active Optimal Assortment (AOA)** Active Assortment Optimization (a.k.a. Utility Maximization with Subset Choices) Berbeglia & Joret (2016); Agrawal et al. (2019); Désir et al. (2016b;a) is an active learning framework for find-

ing the ‘optimal’ profit-maximizing subset. Formally, assume we have a decision set of  $[K] := \{1, 2, \dots, K\}$  of  $K$  items, with each item being associated with the score (or utility) parameters  $\theta := (\theta_1, \theta_2, \dots, \theta_K)$  (without loss of generality assume  $\theta_1 \geq \theta_2 \geq \dots \geq \theta_K \geq 0$ ). At each round  $t = 1, 2, \dots$ , the learner or the algorithm gets to query an assortment (typically subsets containing up to  $m$ -items)  $S_t \subseteq [K]$ , upon which it gets to see some (noisy) relative preferences across the items in  $S_t$ , typically generated according to an underlying Multinomial Logit (MNL) model with parameters  $\theta$  (1). Further, to allow the event where no items are selected, we also model a No-Choice (NC) item, indexed by item-0, with PL parameter  $\theta_0 \in \mathbb{R}_+$ .

**(Objective 1.) Top- $m$ :** identify the top- $m$  item-set:  $\{\theta_1, \dots, \theta_m\}$ , for some  $m \in [1, K]$ .

**(Objective 2.) Wtd-Top- $m$ :** A more general objective could also consider a weight (or price)  $r_i \in \mathbb{R}_+$  associated with the item  $i \in [K]$ , and the goal could be to identify the assortment (subset) with maximum weighted utility <sup>1</sup>, as detailed in Sec. 2.

**Related Works and Limitations:** As stated above, the problem of AOA is fundamental in many practical scenarios, and thus widely studied in multiple research areas, including Online ML/learning theory and operations research.

- In the Online ML literature, the problem is well-studied as *Multi-Dueling Bandits* Sui et al. (2017); Brost et al. (2016), or *Battling Bandits* Saha & Gopalan (2019a; 2018); Bengs et al. (2021b), which is an extension of the famous *Dueling Bandit* problem Zoghi et al. (2014b;a). The main limitation of this line of work is the lack of practical objectives, which either aim to identify the ‘best-item’  $1 (= \arg \max_{i \in [K]} \theta_i)$  within a PAC (probably approximately correct) framework Saha & Gopalan (2019b); Chen et al. (2017; 2018); Ren et al. (2018) or quantifying regret against the best items Saha & Gopalan (2019a); Bengs et al. (2022). Note the latter actually leads to the optimal subset choice of repeatedly selecting the optimal item,  $\arg \max_i \theta_i$ ,  $m$  times, i.e.  $(1, 1, \dots, 1)$ , which is unrealistic from the viewpoint of real-world system design. Selecting an assortment of distinct top- $m$  items (Top- $m$ -AOA) or maximum expected utility (Wtd-Top- $m$ -AOA) makes more sense.

- On the other hand, a similar line of the problem has been studied in operations research and dynamic assortment selection literature, where the goal is to offer a subset of items to the customers in order to maximize expected revenue. The problem has been studied under different user choice models, e.g. PL or Multinomial-Logit models (Agrawal et al., 2019), Mallows and mixture of Mallows (Désir et al., 2016a), Markov chain-based choice models (Désir et al., 2016b), single transition model (Nip et al., 2017) etc. While these works indeed consider a more practical objective of finding the best assortment (subset) with the highest expected utility for a regret minimization objective, (1) a major drawback in their approach lies in the algorithm design which *requires to keep on querying the same set multiple times*, e.g. Agrawal et al. (2019); Ou et al. (2018); Chen et al. (2021); Agrawal et al. (2017). Such design techniques could be impractical to be deployed in real systems where users could easily get annoyed if the same items are shown again and again. For example, in ad-placement, music/movies/news/tweets/reels recommendations, offering the same assortment could increase user dissatisfaction and disengagement.

(2) The second major drawback of this line of work lies in the *structural assumption of their underlying choice models which requires the existence of a reference/default item, that needs to be part of every assortment  $S_t$* . This leads to assuming a No-Choice item, typically denoted as item-0, which is a default choice of any assortment  $S_t$ . Further a stronger and more unrealistic assumption lies in the fact that they require to assume that the above pivot is stronger than the rest of the  $K$  items, i.e.  $\theta_0 \geq \max_{i \in [K]} \theta_i$ , i.e. the No-Choice (NC) action is the most likely outcome of any assortment  $S_t$ . This is often unrealistic, e.g., during user interactions with language models, or online shopping, or maps recommendation, users typically make choices as the user needs to commute or book a flight and a NC action is highly improbable, e.g. in recommender systems like YouTube, Spotify, Netflix or even Yahoo News, users typically make choices as they actually wanted to consume a video, new article or song, etc. Similarly, in shopping recommendations like flights (Expedia or Google

<sup>1</sup>This is equivalent to finding the set with maximum expected revenue when  $r_i$ s represents the price of item  $i$  Agrawal et al. (2019)

flights), hotel (Booking.com), restaurants (Grubhub), or Google Maps recommendations, NC is unlikely. In fact, in some applications, NC might not even be an available (feasible) option, e.g., while interacting with ChatGPT/ Gemini, the language model often requests the users to definitively select one outcome in order to proceed with the thread; Similarly, in robotics applications, training of autonomous vehicles, or more generally in any preference-based RL (a.k.a. PbRL) applications, the teacher/ demonstrator/ human feedback provider must choose an option out of the multiple options (or RL trajectories) towards training the RL policy. Consequently, such assumption limits the use in real-systems. In the existing literature Agrawal et al. (2019); Oh & Iyengar (2019); Agrawal et al. (2017); Grant & Leslie (2023), such assumptions are primarily adapted solely for theoretical needs, precisely for maintaining concentration bounds of the PL parameters  $\theta$ , and hence not well justified from a practical viewpoint.

Agrawal et al. (2019) is the classical MNL-Assortment work with MNL-UCB: The idea is to estimate the true PL parameter  $\theta = (\theta_1, \dots, \theta_K)$ s by repeatedly querying the same set (i.e. assortment  $S_t$ ) multiple times and keeping a count of the average number of times an item  $i \in [K]$  is selected until no items (NC) are selected. They further maintain a UCB of the estimated PL parameters,  $(\hat{\theta}_1, \dots, \hat{\theta}_K)$ , and the assortment of the next phase optimistically based on the UCB estimates. The process is repeated until time step  $T$ . Agrawal et al. (2017) is a follow-up work of Agrawal et al. (2019) from the same group of authors and hence their algorithm MNL-TS is almost the same as MNL-UCB above, with the exception of using Thompson Sampling (TS) with Beta posteriors, instead of the UCB estimates. The regret guarantee and set of restrictive assumptions imposed on the MNL model is also identical to that of Agrawal et al. (2019). The objective of Grant & Leslie (2023) is slightly different from that of above two as their objective is ‘learning to rank’ (LTR), i.e. to find the best ordered subset based on some position bias  $\lambda_i > 0$  for position  $i \in [m]$ . More generally, their preference model is different which assumes the probability of the item played at position- $k$  getting selected (or clicked) for any  $m$ -length ordered set  $S = (S(1), \dots, S(m))$  is given by  $P(k | S) := \frac{\lambda_k \theta_{S(k)}}{\theta_0 + \sum_{j \in [m]} \lambda_j \theta_{S(j)}}$ , where as usual  $\theta_0$  is the score parameter of the no-choice item.

We summarize these existing works in Table 1.

Some recent developments also generalized the AOA problem to linear MNL scores to incorporate large actions embedded in  $d$ -dimension Zhang & Ji (2019); Zhang & Sugiyama (2024); Oh & Iyengar (2019), however, their approaches are either limited to the above restrictions or suffer sub-optimal regret guarantees without those assumptions (e.g. the regret bound of Oh & Iyengar (2019) is  $O(d^{3/2}\sqrt{T})$  which is suboptimal by a  $d$ -factor). Considering the above limitations of the AOA literature, we set to answer two questions:

- (1) Can we consider a general AOA model where the default item, like the NC item defined above, is not necessarily the strongest one, i.e.  $\theta_0 \geq \max_{i \in [K]} \theta_i$ ?
- (2) Can we design a practical and regret optimal algorithm for the AOA framework, without needing to play the same repetitive actions and yet converge to the optimal assortment?

Work	Framework	Assume $\theta_0 = \theta_{\max} = 1$	Regret
Our (Alg. 1)	MNL model (Obj. 2)	No	$\sqrt{\min\{\theta_{\max}, K\}KT \log T}$
Agrawal et al. (2019) (Thm 1)	MNL model (Obj. 2)	Yes	$\sqrt{KT \log T}$
Agrawal et al. (2019) (Thm 4)	MNL model (Obj. 2)	No	$\sqrt{\theta_{\max}KT \log T}$
Agrawal et al. (2017)	MNL model (Obj. 2)	Yes	$\sqrt{KT \log(mT)}$
Grant & Leslie (2023)	MNL model with constraints (Obj. 2)	No	$\sqrt{\frac{KT}{\min_i r_i} \log T}$

Table 1: Our Contribution vs the Existing Results in the  $K$ -armed MNL-Assortment literature. The regret statements include only the leading asymptotic term, ignoring constant factors.

**Contributions** We answer these questions in the affirmative and present best of all scenarios. We design practical algorithms on practical AOA framework with practical objectives—Unlike the existing approaches of the AOA, literature Agrawal et al. (2019); Chen et al. (2021), we do not have to keep playing the same assortment multiple times,

neither require a strongest default item (like NC satisfying  $\theta_0 \geq \max_{i \in [K]} \theta_i$ ). Moreover, our objectives do not require us to converge to a multiset of replicated arms like  $(1, 1, \dots, 1)$ , but converge to the utility-maximizing set of distinct items. We list our contributions below:

**1. A General AOA Setup:** We work with a general problem of AOA for PL model, which requires no additional structural assumption of the  $\theta$  parameters such as  $\theta_0 \geq \max_i \theta_i$ , unlike the existing works. We designed algorithms for two separate objectives Top- $m$  and Wtd-Top- $m$  as discussed above (Sec. 2).

**2. Efficient and Optimal Algorithm using Rank-Breaking MNL-Parameter Estimation:** In Sec. 3, we give a practical, efficient and optimal algorithm for MNL Assortment (up to log factors and the magnitude of  $\theta_{\max}$ ). The regret bound of our algorithm AOA-RB<sub>PL</sub> (Alg. 1) yields  $\tilde{O}(\sqrt{KT})$  regret for both Top- $m$  and Wtd-Top- $m$  objective. Our algorithms use a novel parameter estimation technique for discrete choice models based on the concept of *Rank-Breaking* (RB) which is one of our key contributions towards designing the efficient and optimal algorithm. This enables our algorithm to perform optimally without requiring the No-Choice item to be the strongest. Appendix A details the key concept of our parameter estimation technique exploiting the concept of RB. Our resulting algorithm plays optimistically based on the UCB estimates of PL parameters and does not require repeating the same subset multiple times, justifying our title.

**3. Improvement with Adaptive Pivots:** In Sec. 4, we refine the performance of our algorithm by employing the novel idea of ‘adaptive pivots’ (a reference item) and proposed AOA-RB<sub>PL</sub>-Adaptive. Performance-wise this removes the asymptotic dependence on  $\theta_{\max} = \max_i \theta_i / \theta_0$  in the regret analysis. This enables the algorithm to work effectively in scenarios where the No-Choice item is less likely to be selected, i.e.,  $\theta_{\max} \gg 1$ . This leads to a huge improvement in our experiments, especially in the range of low  $\theta_0$ , where AOA-RB<sub>PL</sub>-Adaptive drastically outperforms over the existing baseline. Comparison of our regret bound with existing work is detailed in Table 1.

**4. Empirical Analysis.** Finally, we corroborate our theoretical results with empirical evaluations (Sec. 5), which certify our superior performance in the general AOA setups.

It is also worth mentioning that our proposed algorithm and their respective regret analysis could be extended to any general random utility (RUM) based preference models Soufiani et al. (2014); Saha & Gopalan (2020), as explained in Rem. 2, the techniques. However, to keep the focus on the AOA problem and ease the presentation, we stick to the special case of MNL choice model based preferences.

## 2 PROBLEM SETUP

We write  $[n] = \{1, 2, \dots, n\}$  and  $\mathbf{1}\{\cdot\}$  denotes the indicator function. The symbol  $\lesssim$ , employed in the proof sketches, represents a coarse inequality.

We consider the sequential decision-making problem of Active Optimal Assortment (AOA), with preference/choice feedback. Formally, the learner is given  $[K]$ , a finite set of  $K$  items ( $K > 2$ ). At each decision round  $t = 1, 2, \dots$ , the learner selects a subset  $S_t \subseteq [K]$  of up to  $m$  items, and receives some (stochastic) feedback about the item preferences of  $S_t$ , drawn according to some unknown underlying MNL choice model (1) with parameters  $\theta = (\theta_1, \theta_2, \dots, \theta_K) \in \mathbb{R}_+^K$ . We assume  $\theta_1 \geq \theta_2 \geq \dots \geq \theta_K$  without loss of generality. An interested reader may check App. A.1 for a detailed discussion on PL models. Given any assortment  $S_t$  we also consider the possibility of ‘no-selection’ of any items given an  $S_t$ . Following the literature of Agrawal et al. (2019), we model this mathematically as a No-Choice (NC) item, indexed by item-0, and its corresponding PL utility parameter  $\theta_0$ . Unlike most existing literature on assortment selection, we are not assuming  $\theta_0 \gtrsim \max_{i \in [K]} \theta_i$ . Further, since the PL model is scale independent, we set  $\theta_0 = 1$  and scale the rest of the PL parameters.

**Feedback model** The feedback model formulates the information received (from the ‘environment’) once the learner plays a subset  $S_t \subseteq [K]$  of at most  $m$  items. Given  $S_t$  we consider the algorithm receives a winner feedback (or index of an item)  $i_t \in S_t \cup \{0\}$ , drawn

according to the underlying PL choice model as:

$$\mathbb{P}(i_t = i | S_t) = \theta_i / (\theta_0 + \sum_{j \in S_t} \theta_j), \quad \forall i \in S_t. \quad (1)$$

We consider the following two objectives for the learner:

**1. Top- $m$ -Objective.** One simple objective could be to identify the top- $m$  item-set:  $\{\theta_1, \dots, \theta_m\}$ , for some  $m \in [1, K]$ . The performance of the learner can be captured by minimizing the following regret:

$$Reg_T^{\text{top}} := \sum_{t=1}^T \frac{\Theta_{S^*} - \Theta_{S_t}}{m}, \quad \text{where } S^* := \operatorname{argmax}_{S \subseteq [K]: |S|=m} \left\{ \Theta_S := \sum_{i \in S} \theta_i \right\}.$$

**2. Wtd-Top- $m$ -Objective.** Here, each item- $i$  is associated with a weight (for example price)  $r_i \in \mathbb{R}_+$ , and the goal is to identify the set of size at most  $m$  with maximum weighted utility. One could measure the regret of the learner as:

$$Reg_T^{\text{wtd}} := \sum_{t=1}^T (\mathcal{R}(S^*, \theta) - \mathcal{R}(S_t, \theta)), \quad \text{where } \mathcal{R}(S, \theta) := \sum_{i \in S} \frac{r_i \theta_i}{\theta_0 + \sum_{j \in S} \theta_j}, \quad \forall S \subseteq [K], \quad (2)$$

denotes  $S^* := \operatorname{argmax}_{S \subseteq [K]: |S| \leq m} \mathcal{R}(S, \theta)$  is the optimal utility-maximizing subset. This objective corresponds to the standard objective in the MNL literature Agrawal et al. (2019).

### 3 A PRACTICAL AND EFFICIENT ALGORITHM FOR AOA WITH PL

#### 3.1 ALGORITHM DESIGN

**Main Idea.** The crux of our novelty lies in our PL parameter estimation technique which maintains an estimate of pairwise scores of  $p_{ij} = \frac{\theta_i}{\theta_i + \theta_j}$  for each pair of item  $(i, j)$  using *Rank-Breaking* (RB) Khetan & Oh (2016)—a classical technique of extracting pairwise comparisons from choice (or partial ranking) feedback by breaking each win-loss pair independently in the choice data. A formal description is given in App. A.2. More precisely, using *rank-breaking* we estimate the relative (pairwise) strength  $\hat{p}_{ij,t}$  of each item pair  $(i, j)$  at round  $t$ , as explained in (3). Further, noting  $\theta_i = p_{i0}/(1 - p_{i0})$  (as  $\theta_0 = 1$ ), we use  $\hat{p}_{i0,t}$  to estimate the MNL score  $\hat{\theta}_{i,t}$  of the  $i$ -th item using the NC (0-th item) as the ‘pivot’ item to benchmark against. Next, we prove a crucial concentration result in Lemma 1 showing indeed  $\hat{\theta}_{i,t}$  is a ‘sharp’ estimate of  $\theta_i$ , which is then subsequently used to prove the final regret guarantees Theorem 3 and Theorem 4 respectively. Our proposed algorithm Alg. 1 is described below:

**Estimate upper-confidence-bounds  $\theta_t^{\text{ucb}}$  from Pairwise Estimates.** At each time  $t$ , our algorithm (Alg. 1) maintains a pairwise preference matrix  $\hat{\mathbf{P}}_t \in [0, 1]^{n \times n}$ , whose  $(i, j)$ -th entry  $\hat{p}_{ij,t}$  records the empirical probability of  $i$  having beaten  $j$  in a pairwise duel, and a corresponding upper confidence bound  $p_{ij,t}^{\text{ucb}}$ . Let  $[\hat{K}] := [K] \cup \{0\}$ . We define for each pair  $(i, j) \in [\hat{K}] \times [\hat{K}]$ ,

$$p_{ij,t}^{\text{ucb}} := \hat{p}_{ij,t} + \sqrt{\frac{2\hat{p}_{ij,t}(1 - \hat{p}_{ij,t})x}{n_{ij,t}}} + \frac{3x}{n_{ij,t}}, \quad \text{where } \hat{p}_{ij,t} := \frac{w_{ij,t}}{n_{ij,t}}, \quad (3)$$

where  $x > 0$  is an input of Alg. 1 and  $w_{ij,t} = \sum_{s=1}^{t-1} \mathbb{1}\{i_s = i, j \in S_s\}$  denotes the number of pairwise wins of item- $i$  over  $j$  after *rank-breaking* and  $n_{ij,t} = w_{ij,t} + w_{ji,t}$  being the total number of times  $(i, j)$  has been ‘rank-broken’ till time  $t$  (details in App. A.2). Noting that  $\theta_i = p_{i0}/(1 - p_{i0})$ , the above UCB estimates  $p_{ij,t}^{\text{ucb}}$  are further used to design UCB estimates of the PL parameters  $\theta_i$  as follows

$$\theta_{i,t}^{\text{ucb}} = p_{i0,t}^{\text{ucb}} / (1 - p_{i0,t}^{\text{ucb}})_+, \quad \text{where } (\cdot)_+ := \max\{\cdot, 0\}.$$

**Optimistic Assortment Selection** The estimates  $\theta_{i,t}^{\text{ucb}}$ s are then used to select the set  $S_t$ , that maximizes the underlying objective. This optimization problem transforms into a static assortment optimization problem with upper confidence bounds  $\theta_{i,t}^{\text{ucb}}$  as the parameters, and efficient solution methods for this case are available (see e.g., Avadhanula et al. (2016); Davis et al. (2013); Rusmevichientong et al. (2010)).

**Algorithm 1 AOA for PL model with RB (AOA-RB<sub>PL</sub>)**


---

```

270 1: input:  $x > 0$ 
271
272 2: init:  $\tilde{K} \leftarrow K + 1$ ,  $[\tilde{K}] = [K] \cup \{0\}$ ,  $\mathbf{W}_1 \leftarrow [0]_{\tilde{K} \times \tilde{K}}$ 
273
274 3: for  $t = 1, 2, 3, \dots, T$  do
275
276 4:   Set  $\mathbf{N}_t = \mathbf{W}_t + \mathbf{W}_t^\top$ , and  $\hat{\mathbf{P}}_t = \frac{\mathbf{W}_t}{\mathbf{N}_t}$ . Denote  $\mathbf{N}_t = [n_{ij,t}]_{\tilde{K} \times \tilde{K}}$  and  $\hat{\mathbf{P}}_t = [\hat{p}_{ij,t}]_{\tilde{K} \times \tilde{K}}$ .
277
278 5:   Define for all  $i$ ,  $p_{ii,t}^{\text{ucb}} = \frac{1}{2}$  and for all  $i, j \in [\tilde{K}], i \neq j$ 
279
280     
$$p_{ij,t}^{\text{ucb}} = \hat{p}_{ij,t} + \left( \frac{2\hat{p}_{ij,t}(1-\hat{p}_{ij,t})x}{n_{ij,t}} \right)^{1/2} + \frac{3x}{n_{ij,t}}$$

281
282 6:    $\theta_{i,t}^{\text{ucb}} := p_{i0,t}^{\text{ucb}} / (1 - p_{i0,t}^{\text{ucb}})_+$ 
283
284 7:    $S_t \leftarrow \begin{cases} \text{Top-}m \text{ items from } \text{argsort}(\{\theta_{1,t}^{\text{ucb}}, \dots, \theta_{K,t}^{\text{ucb}}\}), \\ \text{for Top-}m \text{ objective} \\ \text{argmax}_{S \subseteq [K] \mid |S| \leq m} \mathcal{R}(S, \theta_t^{\text{ucb}}), \\ \text{for Wtd-Top-}m \text{ objective} \end{cases}$ 
285
286 8:   Play  $S_t$ 
287
288 9:   Receive the winner  $i_t \in [\tilde{K}]$  (drawn as per (1))
289
290 10:  Update:  $\mathbf{W}_{t+1} = [w_{ij,t+1}]_{\tilde{K} \times \tilde{K}}$  s.t.  $w_{i_t j, t+1} \leftarrow w_{i_t j, t} + 1 \quad \forall j \in S_t \cup \{0\}$ 
291
292 11: end for

```

---

## 3.2 ANALYSIS: CONCENTRATION LEMMAS

We start the analysis by providing two technical lemmas, whose proofs are deferred to the appendix and that provide confidence bounds for the  $\theta_i$ .

**Lemma 1.** *Let  $T \geq 1$  and  $x > 0$ . Then, with probability at least  $1 - 3KTe^{-x}$ , for all  $t \in [T]$  and  $i \in [K]$ :  $\theta_i \leq \theta_{i,t}^{\text{ucb}}$  atleast one of the following two inequalities is satisfied*

$$n_{i0,t} < 69x(\theta_0 + \theta_i) \quad \text{or} \quad \theta_{i,t}^{\text{ucb}} \leq \theta_i + 4(\theta_0 + \theta_i) \sqrt{\frac{2\theta_0\theta_i x}{n_{i0,t}}} + \frac{22x(\theta_0 + \theta_i)^2}{n_{i0,t}}.$$

The above lemma depends on  $n_{i0,t}$  the number of times items  $i$  have been compared with item 0 up to round  $t$ . The latter is controlled using the following lemma:

**Lemma 2.** *Let  $T \geq 1$  and  $x > 0$ . Then, with probability at least  $1 - KTe^{-x}$ : simultaneously for all  $t \in [T]$  and  $i \in [K]$*

$$\tau_{i,t} < 2x(\theta_0 + \Theta_{S^*})^2 \quad \text{or} \quad n_{i0,t} \geq \frac{(\theta_0 + \theta_i)\tau_{i,t}}{2(\theta_0 + \Theta_{S^*})}, \quad (4)$$

where  $\tau_{i,t} = \sum_{s=1}^{t-1} \mathbb{1}\{i \in S_s\}$  denotes the number of rounds item  $i$  got selected before round  $t$ .

3.3 ANALYSIS: TOP- $m$  OBJECTIVE:

We are now ready to provide the regret upper bound for Algorithm 1 with Top- $m$  objective.

**Theorem 3** (Top- $m$  Objective). *Let  $\theta_{\max} \geq 1$ . Consider any instance of PL model on  $K$  items with parameters  $\theta \in [0, \theta_{\max}]^K$ ,  $\theta_0 = 1$ . The regret of Alg. 1 with parameter  $x = 2 \log T$  is bounded as*

$$\text{Reg}_T^{\text{top}} = O(\theta_{\max}^{3/2} \sqrt{KT \log T}) \quad \text{when } T \rightarrow \infty.$$

The above rate of  $\tilde{O}(\sqrt{KT})$  is optimal (up to log-factors), as a lower bound can be derived from standard multi-armed bandits Auer (2000); Auer et al. (2002). We only state here a sketch of the proof of Theorem 3. The detailed proof is deferred to the App. B.

*Proof Sketch of Theorem 3.* Let us define for any  $S \subseteq [K]$ ,

$$\Theta_S = \sum_{i \in S} \theta_i, \quad \text{and} \quad \Theta_S^{\text{ucb}} := \sum_{i \in S} \theta_i^{\text{ucb}}.$$

Let  $\mathcal{E}$  be the high-probability event such that both Lemma 1 and 2 holds true. Then,  $\mathbb{P}(\mathcal{E}) \geq 1 - 4TKe^{-x}$ . Let us first assume that  $\mathcal{E}$  holds true. Then, by Lemma 1,  $\Theta_{S^*} \leq \Theta_{S^*}^{\text{ucb}} \leq \Theta_{S_t}^{\text{ucb}}$ , which yields

$$\text{Reg}_T^{\text{top}} = \frac{1}{m} \sum_{t=1}^T \Theta_{S^*} - \Theta_{S_t} \leq \frac{1}{m} \sum_{t=1}^T \Theta_{S_t}^{\text{ucb}} - \Theta_{S_t} \lesssim \tau_0 + \frac{1}{m} \sum_{t=1}^T \sum_{i \in S_t} (\theta_{i,t}^{\text{ucb}} - \theta_i) \mathbb{1}\{\tau_{i,t} \geq \tau_0\},$$

where  $\tau_0 = 138x(m+1)^2\theta_{\max}^2$  corresponds to an exploration phase needed for the confidence upper bounds of Lem 1 and 2 to be satisfied. Then, noting that if  $\mathcal{E}$  holds true, we can show by Lemma 2, that  $\mathbb{1}\{\tau_{i,t} \geq \tau_0\} \leq \mathbb{1}\{n_{i0,t} \geq 69x(\theta_0 + \theta_i)\}$ . Therefore, we can apply Lemma 1 that entails,

$$\begin{aligned} \frac{1}{m} \sum_{t=1}^T \sum_{i \in S_t} (\theta_{i,t}^{\text{ucb}} - \theta_i) \mathbb{1}\{\tau_{i,t} \geq \bar{n}_{i0}\} &\lesssim \frac{1}{m} \sum_{t=1}^T \sum_{i \in S_t} \left( (\theta_0 + \theta_i) \sqrt{\frac{\theta_0 \theta_i x}{n_{i0,t}}} \mathbb{1}\{\tau_{i,t} \geq \tau_0\} \right) \\ &\stackrel{\text{Lem. 2}}{\lesssim} \frac{1}{m} \sum_{t=1}^T \sum_{i \in S_t} \theta_{\max}^{3/2} \sqrt{\frac{mx}{\tau_{i,t}}} \lesssim \frac{1}{m} \sum_{i=1}^K \theta_{\max}^{3/2} \sqrt{mx\tau_{i,t}} \lesssim \theta_{\max}^{3/2} \sqrt{xKT}. \end{aligned}$$

where we used  $\sum_{i=1}^n 1/\sqrt{i} \leq 2\sqrt{n}$  and  $\sum_i \tau_{i,t} = mT$  together with Jensen's inequality in the last inequality. We thus have under the event  $\mathcal{E}$  that  $\text{Reg}_T^{\text{top}} \leq O(\theta_{\max}^{3/2} \sqrt{xKT})$  and the proof is concluded by taking the expectation with  $x = 2 \log T$  to control  $\mathbb{P}(\mathcal{E}^c)$ .  $\square$

### 3.4 ANALYSIS: WTD-TOP- $m$ OBJECTIVE

In this section we analyze the regret guarantee of Alg. 1 for Wtd-Top- $m$  objective (2).

**Theorem 4** (Wtd-Top- $m$  Objective). *Let  $\theta_{\max} \geq 1$ . Then, for any  $\theta \in [0, \theta_{\max}]^K$  and weights  $\mathbf{r} \in [0, 1]^K$ , the weighted regret of AOA-RB<sub>PL</sub> (Alg. 1) with  $x = 2 \log T$*

$$\text{Reg}_T^{\text{wtd}} = O(\sqrt{\theta_{\max}KT \log T}) \quad \text{when } T \rightarrow \infty.$$

The complete proof is postponed to App. B. The rate  $\Omega(\sqrt{KT})$  is optimal as proved by the lower bound in Chen & Wang (2017) for MNL bandit problems for  $\theta_{\max} = 1$ . Our result recovers (up to a factor  $\sqrt{\log T}$ ) the one of Agrawal et al. (2019) when  $\theta_{\max} = 1$ . However, their algorithm relies on more sophisticated estimators that necessitate epochs repeating the same assortment until the No-Choice item is selected. Note for our problem setting, where it is possible to have  $\theta_{\max} \gg \theta_0 = 1$ , the length of these epochs could be of  $O(K\theta_{\max})$ , which could be potentially very large when  $\theta_{\max} \gg 1$ . This reduces the number of effective epochs, leading to poor estimation of the PL parameters. We see this tradeoff in our experiments (Sec. 5) where the MNL-UCB algorithm of Agrawal et al. (2019) yields linear  $O(T)$  regret for such choice of the problem parameters.

*Proof sketch of Thm. 4.* Let  $\mathcal{E}$  be the high-probability event such that both Lemma 1 and 2 are satisfied. Then,

$$\begin{aligned} \text{Reg}_T^{\text{wtd}} &= \sum_{t=1}^T \mathbb{E}[\mathcal{R}(S^*, \theta) - \mathcal{R}(S_t, \theta)] \lesssim \sum_{t=1}^T \mathbb{E}[(\mathcal{R}(S^*, \theta) - \mathcal{R}(S_t, \theta)) \mathbb{1}\{\mathcal{E}\}] + T\mathbb{P}(\mathcal{E}^c) \\ &\lesssim \sum_{t=1}^T \mathbb{E}[(\mathcal{R}(S_t, \theta_t^{\text{ucb}}) - \mathcal{R}(S_t, \theta)) \mathbb{1}\{\mathcal{E}\}] + T\mathbb{P}(\mathcal{E}^c) \end{aligned} \quad (5)$$

because  $\mathcal{R}(S_t, \theta_t^{\text{ucb}}) \geq \mathcal{R}(S^*, \theta_t^{\text{ucb}}) \geq \mathcal{R}(S^*, \theta)$  under the event  $\mathcal{E}$  by Lemma 8. We now upper-bound the first term of the right-hand-side

$$\begin{aligned} \sum_{t=1}^T \mathbb{E} \left[ \left( (\mathcal{R}(S_t, \theta_t^{\text{ucb}}) - \mathcal{R}(S_t, \theta)) \right) \mathbb{1}\{\mathcal{E}\} \right] &= \sum_{t=1}^T \mathbb{E} \left[ \left( \sum_{i \in S_t} \frac{r_i \theta_{i,t}^{\text{ucb}}}{\theta_0 + \Theta_{S_t, t}^{\text{ucb}}} - \frac{r_i \theta_i}{\theta_0 + \Theta_{S_t}} \right) \mathbb{1}\{\mathcal{E}\} \right] \\ &\leq \sum_{t=1}^T \mathbb{E} \left[ \left( \sum_{i \in S_t} \frac{r_i (\theta_{i,t}^{\text{ucb}} - \theta_i)}{\theta_0 + \Theta_{S_t}} \right) \mathbb{1}\{\mathcal{E}\} \right] \end{aligned}$$

Because  $\Theta_{S_t, t}^{\text{ucb}} \geq \Theta_{S_t}$  under the event  $\mathcal{E}$  by Lemma 1. Then, using  $r_i \leq 1$ , we further upper-bound using an exploration parameter  $\tau_0 = O(\log(T))$  so that the upper-confidence-bounds in Lemmas 1 and 2 are satisfied

$$\begin{aligned}
& \sum_{t=1}^T \mathbb{E} \left[ \left( \mathcal{R}(S_t, \theta_t^{\text{ucb}}) - \mathcal{R}(S_t, \theta) \right) \mathbb{1}\{\mathcal{E}\} \right] \leq \sum_{i=1}^K \mathbb{E} \left[ \sum_{t=1}^T \left( \frac{|\theta_{i,t}^{\text{ucb}} - \theta_i|}{\theta_0 + \Theta_{S_t}} \right) \mathbb{1}\{i \in S_t, \mathcal{E}\} \right] \\
& \lesssim O(\tau_0) + \sum_{i=1}^K \mathbb{E} \left[ \sum_{t=1}^T \frac{|\theta_{i,t}^{\text{ucb}} - \theta_i|}{\theta_0 + \Theta_{S_t}} \mathbb{1}\{i \in S_t, \tau_{i,t} \geq \tau_0, \mathcal{E}\} \right] \\
& \lesssim O(\tau_0) + \sum_{i=1}^K \sqrt{\sum_{t=1}^T \mathbb{E} \left[ \frac{\theta_i \mathbb{1}\{i \in S_t\}}{\theta_0 + \Theta_{S_t}} \right]} \times \underbrace{\sqrt{\sum_{t=1}^T \mathbb{E} \left[ \left( \frac{\theta_{i,t}^{\text{ucb}} - \theta_i}{\theta_0 + \Theta_{S_t}} \right)^2 \theta_0 + \Theta_{S_t} \frac{1}{\theta_i} \mathbb{1}\{i \in S_t, \tau_{i,t} \geq \tau_0, \mathcal{E}\} \right]}}_{=: A_T(i)}
\end{aligned} \tag{6}$$

where the last inequality is by Cauchy-Schwarz inequality. Now, the term  $A_T(i)$  above may be upper-bounded using Lemmas 1 and 2,

$$\begin{aligned}
A_T(i) &= \mathbb{E} \left[ \frac{(\theta_{i,t}^{\text{ucb}} - \theta_i)^2}{\theta_i(\theta_0 + \Theta_{S_t})} \mathbb{1}\{i \in S_t, \tau_{i,t} \geq \tau_0, \mathcal{E}\} \right] \lesssim \sum_{t=1}^T \mathbb{E} \left[ \frac{(\theta_0 + \theta_i)^2 x}{n_{i0,t}(\theta_0 + \Theta_{S_t})} \mathbb{1}\{i \in S_t\} \right] \\
&\lesssim \theta_{\max} x \sum_{t=1}^T \mathbb{E} \left[ \frac{(\theta_0 + \theta_i) \mathbb{1}\{i \in S_t\}}{(\theta_0 + \Theta_{S_t}) n_{i0,t}} \right] = \theta_{\max} x \mathbb{E} \left[ \sum_{t=1}^T \frac{\mathbb{1}\{i_t \in \{i, 0\}, i \in S_t\}}{n_{i0,t}} \right] \lesssim \theta_{\max} x \log T
\end{aligned}$$

where in the last inequality we used that  $\sum_{n=1}^T n^{-1} \leq 1 + \log T$ . Substituting into (6), Jensen's inequality entails,

$$\sum_{t=1}^T \mathbb{E} \left[ \left( \mathcal{R}(S_t, \theta_t^{\text{ucb}}) - \mathcal{R}(S_t, \theta) \right) \mathbb{1}\{\mathcal{E}\} \right] \lesssim O(\tau_0) + \mathbb{E} \left[ \sqrt{\theta_{\max} x \log T} \sum_{i=1}^K \sqrt{\sum_{t=1}^T \frac{\theta_i \mathbb{1}\{i \in S_t\}}{\theta_0 + \Theta_{S_t}}} \right]. \tag{7}$$

The proof is finally concluded by applying Cauchy-Schwarz inequality which yields:

$$\sum_{i=1}^K \sqrt{\sum_{t=1}^T \frac{\theta_i \mathbb{1}\{i \in S_t\}}{\theta_0 + \Theta_{S_t}}} \leq \sqrt{K \sum_{t=1}^T \frac{\sum_{i=1}^K \theta_i \mathbb{1}\{i \in S_t\}}{\theta_0 + \Theta_{S_t}}} \leq \sqrt{KT}.$$

Finally, combining the above result with (5) and (7) concludes the proof

$$\text{Reg}_T^{\text{rtd}} \lesssim TP(\mathcal{E}^c) + O(\tau_0) + \sqrt{\theta_{\max} x KT \log T}.$$

Choosing  $x = 2 \log T$  ensures  $TP(\mathcal{E}^c) \leq O(1)$  and  $\tau_0 \leq O(\log T)$ .  $\square$

#### 4 IMPROVED DEPENDANCE ON $\theta_{\max}$ WITH ADAPTIVE PIVOT SELECTION

A problem with Algorithm 1 stems from estimating all  $\theta_i$  based on pairwise comparisons with item 0. When  $\theta_{\max} \gg \theta_0 = 1$ , item 0 may not be sampled enough as the winner, leading to poor estimators. This deficiency contributes to the suboptimal dependence on  $\theta_{\max}$  observed in Theorems 3 and 4 and in prior work, such as Agrawal et al. (2019). We propose the following fix to optimize the pivot. For all  $i, j \in [K] \cup \{0\}$  we define  $\gamma_{ij} = \frac{\theta_i}{\theta_j}$ ,

$$\gamma_{ij,t}^{\text{ucb}} = p_{ij,t}^{\text{ucb}} / (1 - p_{ij,t}^{\text{ucb}})_+ \quad \text{and} \quad \gamma_{ii,t}^{\text{ucb}} = 1,$$

where  $p_{ij,t}^{\text{ucb}}$  are defined in (3). For all rounds  $t$ , the algorithm AOA-RB<sub>PL</sub>-Adaptive selects

$$S_t = \underset{|S| \leq m}{\text{argmax}} \mathcal{R}(S, \hat{\theta}_t^{\text{ucb}}) \quad \text{where} \quad \hat{\theta}_{i,t}^{\text{ucb}} := \min_{j \in [K] \cup \{0\}} \gamma_{ij,t}^{\text{ucb}} \gamma_{j0,t}^{\text{ucb}}.$$



432 With the above definition of  $\widehat{\theta}_{i,t}^{\text{ucb}}$ , any item  $i$  is compared to the base item 0 through the  
 433 best possible item  $j$ . When  $j$  is a strong item that is often selected both  $\gamma_{ij,t}^{\text{ucb}}$  and  $\gamma_{j0,t}^{\text{ucb}}$   
 434 are sharp upper-bounds of  $\gamma_{ij}$  and  $\gamma_{j0}$ , making  $\widehat{\theta}_{i,t}^{\text{ucb}}$  itself a sharp upper-confidence bound for  $\theta_i$ .  
 435 This definition in turn also satisfies the condition  $\widehat{\theta}_{i,t}^{\text{ucb}} \geq \theta_i$  required by Lemma 8 and crucial  
 436 for our analysis. The condition would not hold if we used  $\gamma_{ij,t}^{\text{ucb}}$  directly without multiplying  
 437 it with  $\gamma_{j0,t}^{\text{ucb}}$ .

438 We offer below a regret bound that underscores the value of optimizing the pivot when  
 439  $\theta_{\max} \gg K$ . Note that while the algorithm and analysis are presented for the weighted  
 440 objective with winner feedback only, it can be adapted to other objectives by replacing  
 441  $\mathcal{R}(S, \theta)$  with the new objective in the analysis, as long as Lemma 8 remains valid.

442 **Theorem 5.** *Let  $\theta_{\max} \geq 1$ . For any  $\theta \in [0, \theta_{\max}]^K$  and weights  $\mathbf{r} \in [0, 1]^K$ , the weighted  
 443 regret of AOA-RB<sub>PL</sub>-Adaptive is upper-bounded as*

$$444 \text{Reg}_T^{\text{wt}d} = O(\sqrt{\min\{\theta_{\max}, K\}}KT \log T)$$

445 as  $T \rightarrow \infty$  for the choice  $x = 2 \log T$  (when defining  $p_{ij,t}^{\text{ucb}}$ ).

446 **Remark 1** (Drastic Improvement over Prior Works). *Asymptotically, when  $\theta_{\max}$  is constant,  
 448 the regret is  $O(K\sqrt{T} \log T)$ , eliminating any dependence on  $\theta_{\max}$ . This allows for handling  
 449 scenarios where the No-Choice item is highly unlikely, which is not achievable in previous  
 450 works such as Agrawal et al. (2019; 2017). Agrawal et al. (2019) did attempt in their Thm. 4 to  
 451 relax the assumption of  $\theta_{\max} = \theta_0$  and shows a bound of order  $O(\max\{\theta_{\max}/\theta_0, 1\}^{1/2}\sqrt{KT})$ ,  
 452 which unfortunately blows to  $\infty$  as  $\theta_0 \rightarrow 0$  or equivalently  $\theta_{\max} \rightarrow \infty$ , leading to a vacuous  
 453 bound. Here, lies the stark improvement and one of the key contributions, as also corroborated  
 454 in our experimental evaluation Sec. 5 (Fig. 2).*

455 **Remark 2** (Beyond MNL Assortment: Extending to any general RUM based Choice  
 456 Models). *Although, in this paper, we primarily focused on MNL based choice models, it is  
 457 worth mentioning that our proposed algorithms can be generalized to more general random  
 458 utility-based models (RUMs) Azari et al. (2012b); Saha & Ghoshal (2022) pursuing the ideas  
 459 from Saha & Gopalan (2020) that extends the RB based parameter estimation technique  
 460 to any RUM( $\theta$ ) choice models: Precisely, using the RB based RUM-parameter estimation  
 461 technique of Saha & Gopalan (2020), we can show a regret bound of  $\tilde{O}(\frac{\sqrt{\min(\theta_{\max}, K)}}{c_{\text{rum}}} \sqrt{KT})$  for  
 462 our proposed algorithm AOA-RB<sub>PL</sub>, where  $c_{\text{rum}}$  is the parameter associated to the minimum  
 463 advantage ratio (min-AR) of the underlying RUM( $\theta$ ) model, as defined in Thm6 of Saha &  
 464 Gopalan (2020). In particular,  $c_{\text{rum}}$  can shown to be a constant given a fixed RUM model,  
 465 e.g.  $c_{\text{rum}} = 1/4$  for Exp(1), Gamma(2, 1),  $c_{\text{rum}} = 1/(4\sigma)$  for Gumbel( $\mu, \sigma$ ),  $c_{\text{rum}} = \lambda/4$  for  
 466 Weibull( $\lambda, 1$ ),  $c_{\text{rum}} = 1/3$  for Gaussian(0, 1), etc (using Cor5 of Saha & Gopalan (2020)).*

467 Our algorithms and analyses thus apply to any general RUM( $\theta$ ) based choice models; we  
 468 stick to the special case of MNL models in this paper for brevity and keep the main focus on  
 469 the AOA problem and the related algorithmic novelties.

470 The proof of Theorem 5 is deferred to the App. B, with a key step relying on selecting  
 471 the pivot  $j_t = \arg\max_{j \in S_t \cup \{0\}} \theta_j$ . The use of  $|\widehat{\theta}_{i,t}^{\text{ucb}} - \theta_i| \leq |\gamma_{ij_t,t}^{\text{ucb}} - \theta_i|$  provides confidence  
 472 upper-bounds with an improved dependence on  $\theta_{\max}$ , leveraging the fact that  $\theta_{j_t} \geq \theta_i$ . Due  
 473 to the varying pivot over time, a telescoping argument introduces an additive factor  $\sqrt{K}$ .  
 474

## 475 5 EXPERIMENTS

476 We run experiments to compare the performance of our method with the state-of-the-art  
 477 methods. All results are averaged across 100 runs. We evaluate the performance of our main  
 478 algorithm AOA-RB<sub>PL</sub>-Adaptive (Sec. 4), referred as “Our Alg-1 (Adaptive Pivot)”, with the  
 479 following algorithms: AOA-RB<sub>PL</sub> (Sec. 3) referred as “Our Alg-2 (No-Choice Pivot)”, and  
 480 MNL-UCB, the state-of-the-art algorithm for AOA (Agrawal et al. (2019), Alg. 1).  
 481

482 **Different PL ( $\theta$ ) Environments.** We report our experiment results on two datasets with  
 483  $K = 50$  items: (1) Arith50 with PL parameters  $\theta_i = 1 - (i - 1)0.02$ ,  $\forall i \in [50]$ . (2) Bad50  
 484 with PL parameters  $\theta_i = 0.6$ ,  $\forall i \in [50] \setminus \{25\}$  and  $\theta_{25} = 0.8$ . For simplicity of computing  
 485 the assortment choices  $S_t$ , we assume  $r_i = 1$ ,  $\forall i \in [K]$ .

(1). **Averaged Regret with weak NC** ( $\theta_{\max}/\theta_0 \gg 1$ ) (**Fig. 1**): In our first experiment, we set  $m = 5$  and  $\theta_0/\theta_{\max} = 0.01$  and report the average regret of the above three algorithms for our two objectives.




Figure 1: Averaged Regret for  $m = 5$ ,  $\theta_0 = 0.01$

Fig. 1 shows that our algorithm AOA-RB<sub>PL</sub>-Adaptive (with adaptive pivot) significantly outperforms the other two algorithms, while our algorithm AOA-RB<sub>PL</sub> with no-choice (NC) pivot still outperforms MNL-UCB.

(2). **Averaged Regret vs No-Choice PL Parameter** ( $\theta_{\max}/\theta_0$ ) (**Fig. 2**): In this experiment, we evaluate the regret performance of our algorithm AOA-RB<sub>PL</sub>-Adaptive. We report the experiment on Artith50 PL dataset and set the subsetsize  $m = 5$ ,  $\theta_{\max}/\theta_0 = \{1, 0.5, 0.1, 0.05, 0.01, 0.005, 0.001\}$ . Fig. 2 shows the increase in the performance gap between our algorithm AOA-RB<sub>PL</sub>-Adaptive (with adaptive pivot) with decreasing  $\theta_0/\theta_{\max}$ .




Figure 2: Comparative performance for varying  $\theta_0/\theta_{\max}$ ,  $m = 5$




Figure 3: Tradoff: Averaged Regret vs length of the  $k$  rank-ordered feedback

(3). **Averaged Regret vs Length of the rank-ordered feedback** ( $k$ ) (**Fig. 3**): We also run a thought experiment to understand the tradeoff between learning rate with  $k$ -length rank-ordered feedback, where given any assortment  $S_t \subseteq [K]$  of size  $m$ , the learner gets to see the top- $k$  draws ( $k \leq m$ ) from the PL model without replacement. This is a stronger feedback than the winner (i.e. top-1 for  $k = 1$ ) feedback and, as expected, we see in Fig. 3 an improved regret (for both notions) when increasing  $k$ . The experiment are run on the Arith50 dataset with  $m = 30$  and  $k \in \{1, 2, 4, 8\}$ .

## 6 CONCLUSION

We address the Active Optimal Assortment Selection problem with PL choice models, introducing a versatile framework (AOA) that eliminates the need for a strong default item, typically assumed as the No-Choice (NC) item in the existing literature. Our proposed algorithms employ a novel ‘Rank-Breaking’ technique to establish tight concentration guarantees for estimating the parameters of the PL model. Our approach stands out for its practicality and avoids the suboptimal practice of repeatedly selecting the same set of items until the default item prevails. This is beneficial when the default item’s quality ( $\theta_0$ ) is significantly lower than the quality of the best item ( $\theta_{\max}$ ). Our algorithms are computationally efficient, optimal (up to log factors), and free from restrictive assumptions on the default item.

**Future Works.** Among many interesting questions to address in the future, it will be interesting to understand the role of the No-Choice (NC) item in the algorithm design, precisely, can we design efficient algorithms without the existence of NC items with a regret rate still linear in  $\theta_{\max}$ ? Further, it will be interesting to extend our results to more general choice models beyond the PL model Chen et al. (2021); Désir et al. (2016a;b). What is the tradeoff between the subsetsize  $m$  and the regret for such general choice models? Extending our results to large (potentially infinite) decision spaces and contextual settings would also be a very useful and practical contribution to the literature of assortment optimization.

## REFERENCES

- 540  
541  
542 Shipra Agrawal, Vashist Avadhanula, Vineet Goyal, and Assaf Zeevi. Thompson sampling  
543 for the mnl-bandit. In *Conference on learning theory*, pp. 76–78. PMLR, 2017.
- 544  
545 Shipra Agrawal, Vashist Avadhanula, Vineet Goyal, and Assaf Zeevi. Mnl-bandit: A dynamic  
546 learning approach to assortment selection. *Operations Research*, 67(5):1453–1485, 2019.
- 547  
548 Nir Ailon, Zohar Karnin, and Thorsten Joachims. Reducing dueling bandits to cardinal  
549 bandits. In *International Conference on Machine Learning*, pp. 856–864. PMLR, 2014.
- 550  
551 Jean-Yves Audibert, Rémi Munos, and Csaba Szepesvári. Exploration–exploitation tradeoff  
552 using variance estimates in multi-armed bandits. *Theoretical Computer Science*, 410(19):  
1876–1902, 2009.
- 553  
554 Peter Auer. Using upper confidence bounds for online learning. In *Foundations of Computer  
555 Science, 2000. Proceedings. 41st Annual Symposium on*, pp. 270–279. IEEE, 2000.
- 556  
557 Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed  
558 bandit problem. *Machine learning*, 47(2-3):235–256, 2002.
- 559  
560 Vashist Avadhanula, Jalaj Bhandari, Vineet Goyal, and Assaf Zeevi. On the tightness of an  
561 lp relaxation for rational optimization and its applications. *Operations Research Letters*,  
44(5):612–617, 2016.
- 562  
563 Hossein Azari, David Parkes, and Lirong Xia. Random utility theory for social choice. In  
564 *Advances in Neural Information Processing Systems*, pp. 126–134, 2012a.
- 565  
566 Hossein Azari, David Parks, and Lirong Xia. Random utility theory for social choice.  
567 *Advances in Neural Information Processing Systems*, 25, 2012b.
- 568  
569 Viktor Bengs, Róbert Busa-Fekete, Adil El Mesaoudi-Paul, and Eyke Hüllermeier. Preference-  
570 based online learning with dueling bandits: A survey. *Journal of Machine Learning  
571 Research*, 2021a.
- 572  
573 Viktor Bengs, Róbert Busa-Fekete, Adil El Mesaoudi-Paul, and Eyke Hüllermeier. Preference-  
574 based online learning with dueling bandits: A survey. *J. Mach. Learn. Res.*, 22:7–1, 2021b.
- 575  
576 Viktor Bengs, Aadirupa Saha, and Eyke Hüllermeier. Stochastic contextual dueling bandits  
577 under linear stochastic transitivity models. In *International Conference on Machine  
578 Learning*, pp. 1764–1786. PMLR, 2022.
- 579  
580 Gerardo Berbeglia and Gwenaél Joret. Assortment optimisation under a general discrete  
581 choice model: A tight analysis of revenue-ordered assortments. *arXiv preprint  
582 arXiv:1606.01371*, 2016.
- 583  
584 Brian Brost, Yevgeny Seldin, Ingemar J. Cox, and Christina Lioma. Multi-dueling bandits  
585 and their application to online ranker evaluation. *CoRR*, abs/1608.06253, 2016.
- 586  
587 Niladri S Chatterji, Aldo Pacchiano, Peter L Bartlett, and Michael I Jordan. On the theory  
588 of reinforcement learning with once-per-episode feedback. *arXiv preprint arXiv:2105.14363*,  
589 2021.
- 590  
591 Xi Chen and Yining Wang. A note on a tight lower bound for mnl-bandit assortment selection  
592 models. *arXiv preprint arXiv:1709.06109*, 2017.
- 593  
594 Xi Chen, Sivakanth Gopi, Jieming Mao, and Jon Schneider. Competitive analysis of the top-k  
595 ranking problem. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on  
596 Discrete Algorithms*, pp. 1245–1264. SIAM, 2017.
- 597  
598 Xi Chen, Yuanzhi Li, and Jieming Mao. A nearly instance optimal algorithm for top-k  
599 ranking under the multinomial logit model. In *Proceedings of the Twenty-Ninth Annual  
600 ACM-SIAM Symposium on Discrete Algorithms*, pp. 2504–2522. SIAM, 2018.

- 594 Xi Chen, Chao Shi, Yining Wang, and Yuan Zhou. Dynamic assortment planning under  
595 nested logit models. *Production and Operations Management*, 30(1):85–102, 2021.
- 596
- 597 Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei.  
598 Deep reinforcement learning from human preferences. *Advances in neural information*  
599 *processing systems*, 30, 2017.
- 600 James Davis, Guillermo Gallego, and Huseyin Topaloglu. Assortment planning under the  
601 multinomial logit model with totally unimodular constraint structures. *Work in Progress*,  
602 2013.
- 603
- 604 Antoine Désir, Vineet Goyal, Srikanth Jagabathula, and Danny Segev. Assortment optimiza-  
605 tion under the mallows model. In *Advances in Neural Information Processing Systems*, pp.  
606 4700–4708, 2016a.
- 607 Antoine Désir, Vineet Goyal, Danny Segev, and Chun Ye. Capacity constrained assortment  
608 optimization under the markov chain based choice model. *Operations Research*, 2016b.
- 609
- 610 James A Grant and David S Leslie. Learning to rank under multinomial logit choice. *Journal*  
611 *of Machine Learning Research*, 24(260):1–49, 2023.
- 612 Minje Jang, Sunghyun Kim, Changho Suh, and Sewoong Oh. Optimal sample complexity of  
613 m-wise data for top-k ranking. In *Advances in Neural Information Processing Systems*, pp.  
614 1685–1695, 2017.
- 615 Daniel Kahneman and Amos Tversky. The psychology of preferences. *Scientific American*,  
616 246(1):160–173, 1982.
- 617
- 618 Ashish Khetan and Sewoong Oh. Data-driven rank breaking for efficient rank aggregation.  
619 *Journal of Machine Learning Research*, 17(193):1–54, 2016.
- 620
- 621 Sam Musallam, BD Corneil, Bradley Greger, Hans Scherberger, and Richard A Andersen.  
622 Cognitive control signals for neural prosthetics. *Science*, 305(5681):258–262, 2004.
- 623 Kameng Nip, Zhenbo Wang, and Zizhuo Wang. Assortment optimization under a single  
624 transition model. 2017.
- 625
- 626 Min-hwan Oh and Garud Iyengar. Thompson sampling for multinomial logit contextual  
627 bandits. *Advances in Neural Information Processing Systems*, 32, 2019.
- 628 Mingdong Ou, Nan Li, Shenghuo Zhu, and Rong Jin. Multinomial logit bandit with linear  
629 utility functions. *arXiv preprint arXiv:1805.02971*, 2018.
- 630
- 631 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin,  
632 Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language  
633 models to follow instructions with human feedback. *Advances in Neural Information*  
634 *Processing Systems*, 35:27730–27744, 2022.
- 635 Wenbo Ren, Jia Liu, and Ness B Shroff. PAC ranking from pairwise and listwise queries:  
636 Lower bounds and upper bounds. *arXiv preprint arXiv:1806.02970*, 2018.
- 637
- 638 Paat Rusmevichientong, Zuo-Jun Max Shen, and David B Shmoys. Dynamic assortment  
639 optimization with a multinomial logit choice model and capacity constraint. *Operations*  
640 *research*, 58(6):1666–1680, 2010.
- 641 Aadirupa Saha and Suprovat Ghoshal. Exploiting correlation to achieve faster learning rates  
642 in low-rank preference bandits. In *International Conference on Artificial Intelligence and*  
643 *Statistics*, pp. 456–482. PMLR, 2022.
- 644 Aadirupa Saha and Aditya Gopalan. Active ranking with subset-wise preferences. *Interna-*  
645 *tional Conference on Artificial Intelligence and Statistics (AISTATS)*, 2018.
- 646
- 647 Aadirupa Saha and Aditya Gopalan. Combinatorial bandits with relative feedback. In  
*Advances in Neural Information Processing Systems*, 2019a.

- 648 Aadirupa Saha and Aditya Gopalan. PAC Battling Bandits in the Plackett-Luce Model. In  
649 *Algorithmic Learning Theory*, pp. 700–737, 2019b.
- 650
- 651 Aadirupa Saha and Aditya Gopalan. Best-item learning in random utility models with  
652 subset choices. In *International Conference on Artificial Intelligence and Statistics*, pp.  
653 4281–4291. PMLR, 2020.
- 654 Hossein Azari Soufiani, David C Parkes, and Lirong Xia. Computing parametric ranking  
655 models via rank-breaking. In *ICML*, pp. 360–368, 2014.
- 656
- 657 Yanan Sui, Vincent Zhuang, Joel Burdick, and Yisong Yue. Multi-dueling bandits with  
658 dependent arms. In *Conference on Uncertainty in Artificial Intelligence*, UAI’17, 2017.
- 659 Kalyan Talluri and Garrett Van Ryzin. Revenue management under a general discrete choice  
660 model of consumer behavior. *Management Science*, 50(1):15–33, 2004.
- 661
- 662 Yisong Yue, Josef Broder, Robert Kleinberg, and Thorsten Joachims. The  $k$ -armed dueling  
663 bandits problem. *Journal of Computer and System Sciences*, 78(5):1538–1556, 2012.
- 664 Yu-Jie Zhang and Masashi Sugiyama. Online (multinomial) logistic bandit: Improved regret  
665 and constant computation cost. *Advances in Neural Information Processing Systems*, 36,  
666 2024.
- 667
- 668 Zihan Zhang and Xiangyang Ji. Regret minimization for reinforcement learning by evaluating  
669 the optimal bias function. In *Advances in Neural Information Processing Systems*, pp.  
670 2827–2836, 2019.
- 671 Masrour Zoghi, Shimon Whiteson, Remi Munos, Maarten de Rijke, et al. Relative upper  
672 confidence bound for the  $k$ -armed dueling bandit problem. In *JMLR Workshop and  
673 Conference Proceedings*, number 32, pp. 10–18. JMLR, 2014a.
- 674
- 675 Masrour Zoghi, Shimon A Whiteson, Maarten De Rijke, and Remi Munos. Relative confidence  
676 sampling for efficient on-line ranker evaluation. In *Proceedings of the 7th ACM international  
677 conference on Web search and data mining*, pp. 73–82. ACM, 2014b.
- 678 Masrour Zoghi, Zohar S Karnin, Shimon Whiteson, and Maarten De Rijke. Copeland dueling  
679 bandits. In *Advances in Neural Information Processing Systems*, pp. 307–315, 2015.
- 680
- 681
- 682
- 683
- 684
- 685
- 686
- 687
- 688
- 689
- 690
- 691
- 692
- 693
- 694
- 695
- 696
- 697
- 698
- 699
- 700
- 701