

Instance-Sensitive Algorithms for Pure Exploration in Multinomial Logit Bandit*

Nikolai Karpov,¹ Qin Zhang¹

¹Indiana University Bloomington
Luddy Hall, RM 3044
700 North Woodlawn Avenue
Bloomington, IN 47408-3901, USA
nkarpov@iu.edu, qzhangcs@indiana.edu

Abstract

Motivated by real-world applications such as fast fashion retailing and online advertising, the Multinomial Logit Bandit (MNL-bandit) is a popular model in online learning and operations research, and has attracted much attention in the past decade. In this paper, we give efficient algorithms for *pure exploration* in MNL-bandit. Our algorithms achieve *instance-sensitive* pull complexities. We also complement the upper bounds by an almost matching lower bound.

1 Introduction

We study a model in online learning called *multinomial logit bandit* (MNL-bandit for short), where we have N substitutable items $\{1, 2, \dots, N\}$, each of which is associated with a known reward $r_i \in (0, 1]$ and an *unknown* preference parameter $v_i \in (0, 1]$. We further introduce a null item 0 with reward $r_0 = 0$, which stands for the case of “no-purchase”. We set $v_0 = 1$, that is, we assume that the no-purchase decision is the most frequent case, which is a convention in the MNL-bandit literature and can be justified by many real-world applications to be mentioned shortly.

Denote $[n] \triangleq \{1, 2, \dots, n\}$. Given a subset (called an *assortment*) $S \subseteq [N]$, the probability that one chooses $i \in S \cup \{0\}$ is given by

$$p_i(S) = \frac{v_i}{v_0 + \sum_{j \in S} v_j} = \frac{v_i}{1 + \sum_{j \in S} v_j}.$$

Intuitively, the probability of choosing the item i in S is proportional to its preference v_i . This choice model is called the *MNL choice model*, introduced independently by Luce [24] and Plackett [26]. We are interested in finding an assortment $S \subseteq [N]$ such that the following expected reward is maximized.

Definition 1 (expected reward). *Given an assortment $S \subseteq [N]$ and a vector of item preferences $\mathbf{v} = (v_1, \dots, v_N)$, the expected reward of S with respect to \mathbf{v} is defined to be*

$$R(S, \mathbf{v}) = \sum_{i \in S} r_i p_i(S) = \sum_{i \in S} \frac{r_i v_i}{1 + \sum_{j \in S} v_j}.$$

*Nikolai Karpov and Qin Zhang are supported in part by NSF CCF-1844234 and CCF-2006591.

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

The MNL-bandit problem was initially motivated by fast fashion retailing and online advertising, and finds many applications in online learning, recommendation systems, and operations research (see [6] for an overview). For instance, in fast fashion retailing, each item corresponds to a product and its reward is simply the revenue generated by selling the product. The assumption that $v_0 \geq \max\{v_1, \dots, v_N\}$ can be justified by the fact that most customers do not buy anything in a shop visit. A similar phenomenon is also observed in online advertising where it is most likely that a user does not click any of the ads on a webpage when browsing. We naturally want to select a set of products/ads $S \subseteq [N]$ to display in the shop/webpage so that $R(S, \mathbf{v})$, which corresponds to revenue generated by customer/user per visit, is maximized.

We further pose a capacity constraint K on the cardinality of S , since in most applications the size of the assortment cannot exceed a certain size. For example, the number of products presented at a retail shop is capped due to shelf space constraints, and the number of ads placed on a webpage cannot exceed a certain threshold.

In the MNL-bandit model, we need to *simultaneously* learn the item preference vector \mathbf{v} and find the assortment with the maximum expected reward under \mathbf{v} . We approach this by repeatedly selecting an assortment to present to the user, observing the user’s choice, and then trying to update the assortment selection policy. We call each observation of the user choice given an assortment a *pull*. We are interested in minimizing the number of pulls, which is the most expensive part of the learning process.

In bandit theory we are interested in two objectives. The first is called *regret minimization*: given a pull budget T , try to minimize the accumulated difference (called *regret*) between the sum of expected rewards of the optimal strategy in the T pulls and that of the proposed learning algorithm; in the optimal strategy we always present the best assortment (i.e., the assortment with the maximum expected reward) to the user at each pull. The second is called *pure exploration*, where the goal is simply to identify the best assortment.

Regret minimization in MNL-bandit has been studied extensively in the literature [27, 30, 12, 2, 3, 11]. The algorithms proposed in [27, 30] for the regret minimization problem make use of an “exploration then exploitation” strategy, that is, they first try to find the best assortment and then stick to it.

However, they need the prior knowledge of the *gap* between the expected reward of the optimal assortment and that of the second-best assortment, which, in our opinion, is *unrealistic* in practice since the preference vector \mathbf{v} is unknown at the beginning. We will give a more detailed discussion on these works in Section 1.1.

In this paper we focus on pure exploration. Pure exploration is useful in many applications. For example, the retailer may want to perform a set of customer preference tests (e.g., crowdsourcing) to select a good assortment before the actual store deployment. We propose algorithms for pure exploration in MNL-bandit *without* any prior knowledge of preference vector. Our algorithms achieve *instance-sensitive* pull complexities which we elaborate next.

Instance Complexity. Before presenting our results, we give a few definitions and introduce *instance complexities* for pure exploration in MNL-bandit.

Definition 2 (best assortment $S_{\mathbf{v}}$ and optimal expected reward $\theta_{\mathbf{v}}$). *Given a capacity parameter K and a vector of item preferences \mathbf{v} , let*

$$S_{\mathbf{v}} \triangleq \arg \max_{S \subseteq [N]: |S| \leq K} R(S, \mathbf{v})$$

denote the best assortment with respect to \mathbf{v} . If the solution is not unique then we choose the one with the smallest cardinality which is unique (see the discussion after Lemma 2). Let $\theta_{\mathbf{v}} \triangleq R(S_{\mathbf{v}}, \mathbf{v})$ be the optimal expected reward.

Denote $\eta_i \triangleq (r_i - \theta_{\mathbf{v}})v_i$; we call η_i the *advantage* of item i . Suppose we have sorted the N items according to η_i , let $\eta^{(j)}$ be the j -th largest value in the sorted list.

Definition 3 (reward gap Δ_i). *For any item $i \in [N] \setminus S_{\mathbf{v}}$, we define its reward gap to be*

$$\Delta_i \triangleq \begin{cases} \eta^{(K)} - \eta_i, & \text{if } |S_{\mathbf{v}}| = K, \\ -\eta_i, & \text{if } |S_{\mathbf{v}}| < K. \end{cases}$$

and for any item $i \in S_{\mathbf{v}}$, we define

$$\Delta_i \triangleq \bar{\Delta} = \min \left\{ \left(\eta^{(K)} - \eta^{(K+1)} \right), \min_{j \in S_{\mathbf{v}}} \{r_j - \theta_{\mathbf{v}}\} \right\}. \quad (1)$$

Definition 3 may look a bit cumbersome. The extra term $\min_{j \in S_{\mathbf{v}}} \{r_j - \theta_{\mathbf{v}}\}$ in (1) is added for a technical reason when handling the case that $|S_{\mathbf{v}}| < K$; we will discuss this in more detail in Remark 10. If $|S_{\mathbf{v}}| = K$, then the task of finding the best assortment is equivalent to the task of identifying the K items with the largest advantage values η_i , and the reward gap in Definition 3 can be simplified as

$$\Delta_i = \begin{cases} \eta^{(K)} - \eta_i, & \forall i \in [N] \setminus S_{\mathbf{v}}, \\ \eta^{(K)} - \eta^{(K+1)}, & \forall i \in S_{\mathbf{v}}. \end{cases}$$

We now give two instance complexities for pure exploration in MNL-bandit. The second can be viewed as a refinement of the first.

Definition 4 (instance complexity H_1). *We define the first instance complexity for pure exploration in MNL-bandit to be*

$$H_1 \triangleq \sum_{i \in [N]} \frac{1}{\Delta_i^2}.$$

In this paper we assume that $\forall i \in [N], \Delta_i \neq 0$, since otherwise the complexity H_1 will be infinity. This assumption implies that the best assortment is unique, which is also an essential assumption for works of literature whose pull complexities are based on ‘‘assortment-level’’ gaps, as we will discuss in Section 1.1.

Definition 4 bears some similarity to the instance complexity defined for pure exploration in the *multi-armed bandits* (MAB) model, where we have N items each of which is associated with an unknown distribution, and the goal is to identify the item whose distribution has the largest mean. In MAB the instance complexity is defined to be $H_{\text{MAB}} = \sum_{i=2}^N 1/\Delta_i^2$, where $\Delta_i = \mu^{(1)} - \mu^{(i)}$ where $\mu^{(1)}$ is the largest mean of the N items and $\mu^{(i)}$ is the i -th largest mean of the N items [5]. Our definition of Δ_i is more involved due to the more complicated combinatorial structure of the MNL-bandit model.

Definition 5 (instance complexity H_2).

$$H_2 \triangleq \sum_{i \in [N]} \frac{v_i + 1/K}{\Delta_i^2} + \max_{i \in [N]} \frac{1}{\Delta_i^2}.$$

It is easy to see that $H_2 = O(H_1)$ (more precisely, $\frac{H_1}{K} \leq H_2 \leq 3H_1$). We comment that the $\max_{i \in [N]} \frac{1}{\Delta_i^2}$ term is needed only when $|S_{\mathbf{v}}| < K$.

Our Results. We propose two fixed-confidence algorithms for pure exploration in MNL-bandit. The first one (Algorithm 3 in Section 3) gives a pull complexity of $O(K^2 H_1 \ln(\frac{N}{\delta} \ln(KH_1)))$ where δ is the confidence parameter. We then modify the algorithm using a more efficient preference exploration procedure at each pull, and improve the asymptotic pull complexity to $O(K^2 H_2 \ln(\frac{N}{\delta} \ln(KH_2)))$. The second algorithm is presented in Algorithm 5 in Section 4.

Both algorithms can be implemented efficiently: the time complexity of Algorithm 3 is bounded by $\tilde{O}(T + N^2)$ where T is the pull complexity and ‘ $\tilde{O}()$ ’ hides some logarithmic factors. That of Algorithm 5 is bounded by $\hat{O}(TN + N^2)$.¹

As we shall discuss in Remark 12, though having a larger pull complexity, Algorithm 3 still has the advantage that it better fits the *batched* model where we try to minimize the number of changes of the learning policy.

To complement our upper bounds, we prove that $\Omega(H_2/K^2)$ pulls is needed in order to identify the best assortment with probability at least 0.6. This is presented in Section 5. Note that when K is a constant, our upper and lower bounds match up to a logarithmic factor.

1.1 Related Work.

Regret minimization in MNL-bandit was first studied by Rusmevichientong et al. [27] in the setting of *dynamic assortment selection* under the MNL choice model. Since then there have been a number of follow-ups that further improve the regret bound and/or remove some artificial assumptions.

¹When we talk about *time complexity*, we only count the running time of the algorithm itself, and do not include the time for obtaining the pull results which depends on users’ response time.

tions [30, 12, 2, 3, 11]. Agarwal et al. [1] studied *choice bandits* which can be seen as a generalization of MNL-bandit.

As mentioned previously, the algorithms in [27, 30] also have a component of identifying the best assortment. In [27, 30], the following ‘‘assortment-level’’ gap was introduced:

$$\Delta_{\text{asso}} = \theta_{\mathbf{v}} - \max_{S \subseteq [N], |S| \leq K, S \neq S_{\mathbf{v}}} R(S, \mathbf{v}),$$

that is, the difference between the reward of the best assortment and that of the second-best assortment. The pull complexity of the component in [30] for finding the best assortment can be written as $\tilde{O}(KN/\Delta_{\text{asso}}^2)$, where ‘ $\tilde{O}()$ ’ hides some logarithmic factors. This result is better than that in [27]. There are two critical differences between these results and our results: (1) More critically, in [27, 30] it is assumed that the ‘‘assortment-level’’ gap Δ_{asso} is known at the beginning, which is *not* practical since the fact that the preference vector is unknown at the beginning is a key feature of the MNL-bandit problem. (2) Our reward gaps Δ_i are defined at the ‘‘item-level’’; the instance complexity H_1 (or H_2) is defined as the sum of the inverse square of these item-level gaps and the total pull complexity is $\tilde{O}(K^2H_1)$ (or $\tilde{O}(K^2H_2)$). Though the two complexities are not directly comparable, the following example shows that for certain input instances, our pull complexity is significantly better.

Example 1. $K = 1, r_1 = \dots = r_N = 1, v_1 = 1, v_2 = 1 - 1/\sqrt{N}, v_3 = \dots = v_N = 1/\sqrt{N}$. We have $KN/\Delta_{\text{asso}} = \Omega(N^2)$, while $K^2H_1 = O(N)$. Thus, the pull complexity of the algorithm in [27] is quadratic of ours (up to logarithmic factors).

The best assortment identification problem has also been studied in the static setting (e.g., [31, 13]), where the user preference vector \mathbf{v} is known a priori and our task is to conduct an offline computation to find the assortment that maximizes the expected reward. We refer readers to [23] for an overview of this setting.

Chen et al. [10] studied the problem of top- k ranking under the MNL choice model (but without the ‘‘no purchase’’ option). Their problem is different from ours: They aimed to find the k items in $[N]$ with the largest preference v_i (instead of the advantage $\eta_i = (r_i - \theta_{\mathbf{v}})v_i$). In some sense, their problem can be thought of as a special case of ours, where $r_1 = r_2 = \dots = r_N$ (that is, the rewards of all items are the same); but in their model, there is *no* null item. It seems difficult to extend their approach to our setting. We would also like to mention the work on battling-bandits by Saha and Gopalan [29], who considered the problem of using the MNL choice model as one of the natural models to draw a winner from a set of items. But their problem settings and the notion of the optimal solution are again different from the problem we consider here.

Pure exploration has been studied extensively in the model of MAB [14, 25, 5, 16, 15, 21, 18, 22, 17, 28, 9]. MNL-bandit can be viewed as an MAB-type model with $\sum_{j \in [K]} \binom{N}{j}$ items, each corresponding to an assortment $S \subseteq [N]$ with $|S| \leq K$. However, these items may ‘‘intersect’’ with each other since assortments may contain the same items. Due to such dependencies, the algorithms designed for pure

exploration in the MAB model cannot be adopted to the MNL-bandit model. Audibert et al. [5] designed an instance-sensitive algorithm for the pure exploration problem in the MAB model. The result in [5] was later improved by Karnin et al. [21] and Chen et al. [9], and extended into the problem of identifying multiple items [7, 33, 8].

Finally, we note that recently, concurrent and independent of our work, Yang [32] has also studied pure exploration in MNL-bandit. But the definition of instance complexity in [32] is again at the ‘‘assortment-level’’ (and thus the results are not directly comparable), and the algorithmic approaches in [32] are also different from ours. The pull complexity of [32] can be written as $\tilde{O}(H_{\text{yang}})$ where $H_{\text{yang}} = \sum_{i \in [N]} \frac{1}{(\Delta'_i)^2}$, where Δ'_i is defined to be the difference between the best reward among assortments that include item i and that among assortments that exclude item i . The following example shows that for certain input instances, our pull complexity based on item-level gaps is better.

Example 2. $r_1 = \dots = r_N = 1, v_1 = \dots = v_K = 1, v_{K+1} = \dots = v_N = \epsilon$. For $\epsilon \in (0, 1/K)$ and $\omega(1) \leq K \leq o(N)$, we have $H_{\text{yang}} = \Theta(NK^4)$, while our $K^2H_2 = \Theta(K^5 + NK^3) = o(H_{\text{yang}})$.

2 Preliminaries

Before presenting our algorithms, we would like to introduce some tools in probability theory and give some basic properties of the MNL-bandit model. Due to space constraints, we leave the tools in probability theory (including Hoeffding’s inequality, concentration results for the sum of geometric random variables, etc.) to Appendix A.1.

The following (folklore) observation gives an effective way to check whether the expected reward of S with respect to \mathbf{v} is at least θ for a given value θ . The proof can be found in Appendix A.2.

Observation 1. For any $\theta \in [0, 1]$, $R(S, \mathbf{v}) \geq \theta$ if and only if $\sum_{i \in S} (r_i - \theta)v_i \geq \theta$.

With Observation 1, to check whether the maximum expected reward is at least θ for a given value θ , we only need to check whether the expected reward of the particular set $S \subseteq [N]$ containing the up to K items with the largest *positive* values $(r_i - \theta)v_i$ is at least θ .

To facilitate the future discussion we introduce the following definition.

Definition 6 ($\text{Top}(I, \mathbf{v}, \theta)$). Given a set of items I where the i -th item has reward r_i and preference v_i , and a value θ , let T be the set of $\min\{K, |I|\}$ items with the largest values $(r_i - \theta)v_i$. Define $\text{Top}(I, \mathbf{v}, \theta) \triangleq T \setminus \{i \in I \mid (r_i - \theta)v_i \leq 0\}$, where \mathbf{v} stands for $(v_1, \dots, v_{|I|})$.

The following lemma shows that $\text{Top}(I, \mathbf{v}, \theta_{\mathbf{v}})$ is exactly the best assortment. Its proof can be found in Appendix A.3.

Lemma 2. $\text{Top}(I, \mathbf{v}, \theta_{\mathbf{v}}) = S_{\mathbf{v}}$.

Note that the set $\text{Top}(I, \mathbf{v}, \theta_{\mathbf{v}})$ is unique by its definition. Therefore by Lemma 2 the set $S_{\mathbf{v}}$ is also uniquely defined.

We next show a monotonicity property of the expected reward function $R(\cdot, \cdot)$. Given two vectors \mathbf{v}, \mathbf{w} of the same

Algorithm 1: EXPLORE(i)

Input: Item i .**Output:** 0/1 (choose or not choose i).

- 1 Offer a singleton set $S_i \leftarrow \{i\}$ and observe a feedback a ;
 - 2 **if** $a = 0$ **then return** 1;
 - 3 **return** 0
-

Algorithm 2: PRUNE($I, K, \mathbf{a}, \mathbf{b}$)

Input: a set of items $I = \{1, \dots, N\}$, capacity parameter K , two vectors $\mathbf{a} = (a_1, \dots, a_N)$, $\mathbf{b} = (b_1, \dots, b_N)$ such that for any $i \in [N]$ it holds that $a_i \leq v_i \leq b_i$, where $\mathbf{v} = (v_1, \dots, v_N)$ is the (unknown) preference vector of the N items.**Output:** a set of candidate items for constructing the best assortment.

- 1 $\theta_{\mathbf{a}} \leftarrow \max_{S \subseteq I: |S| \leq K} R(S, \mathbf{a})$, $\theta_{\mathbf{b}} \leftarrow \max_{S \subseteq I: |S| \leq K} R(S, \mathbf{b})$;
 - 2 $C \leftarrow \emptyset$;
 - 3 **foreach** $i \in I$ **do**
 - 4 form a vector $\mathbf{g} = (g_1, \dots, g_N)$ s.t. $g_j = a_j$ for $j \neq i$, and $g_i = b_i$;
 - 5 **if** $\exists \theta \in [\theta_{\mathbf{a}}, \theta_{\mathbf{b}}]$ s.t. $i \in \text{Top}(I, \mathbf{g}, \theta)$ **then** add i to C ;
 - 6 **return** C
-

dimension, we write $\mathbf{v} \preceq \mathbf{w}$ if $\forall i, v_i \leq w_i$. We comment that similar properties appeared in [2, 3], but were formulated a bit differently from ours. The proof of Lemma 3 can be found in Appendix A.4.

Lemma 3. *If $\mathbf{v} \preceq \mathbf{w}$, then $(\theta_{\mathbf{v}} =) R(S_{\mathbf{v}}, \mathbf{v}) \leq (\theta_{\mathbf{w}} =) R(S_{\mathbf{w}}, \mathbf{w})$, and for any $S \subseteq I$ it holds that*

$$R(S, \mathbf{w}) - R(S, \mathbf{v}) \leq \sum_{i \in S} (w_i - v_i).$$

The following is an immediate corollary of Lemma 3.

Corollary 4. *If $\forall i : v_i \leq w_i \leq v_i + \frac{\epsilon}{K}$, then $\theta_{\mathbf{v}} \leq \theta_{\mathbf{w}} \leq \theta_{\mathbf{v}} + \epsilon$.*

3 The Basic Algorithm

In this section, we present our first algorithm for pure exploration in MNL-bandit. The main algorithm is described in Algorithm 3, which calls PRUNE (Algorithm 2) and EXPLORE (Algorithm 1) as subroutines. EXPLORE describes a pull of the assortment consisting of a single item.

Let us describe the Algorithm 2 and 3 in more detail. Algorithm 3 proceeds in rounds. In round τ , each “surviving” item in the set I_{τ} has been pulled by T_{τ} times in total. We try to construct two vectors \mathbf{a} and \mathbf{b} based on the empirical means of the items in I_{τ} such that the (unknown) true preference vector \mathbf{v} of I_{τ} is tightly sandwiched by \mathbf{a} and \mathbf{b} (Line 7-8). We then feed I_{τ} , \mathbf{a} , and \mathbf{b} to the PRUNE subroutine which reduces the size of I_{τ} by removing items that have no chance

Algorithm 3: The Fixed Confidence Algorithm for MNL-Bandit

Input: a set of items $I = \{1, \dots, N\}$, a capacity parameter K , a confidence parameter δ .**Output:** the best assortment.

- 1 $I_0 \leftarrow I$, set $\epsilon_{\tau} = 2^{-\tau-3}$ for $\tau \geq 0$;
 - 2 set $T_{-1} \leftarrow 0$ and $T_{\tau} \leftarrow \left\lceil \frac{32}{\epsilon_{\tau}^2} \ln \frac{16N(\tau+1)^2}{\delta} \right\rceil$ for $\tau \geq 0$;
 - 3 **for** $\tau = 0, 1, \dots$ **do**
 - 4 **foreach** $i \in I_{\tau}$ **do** call EXPLORE(i) for $(T_{\tau} - T_{\tau-1})$ times;
 - 5 let $x_i^{(\tau)}$ be the mean of the outputs of the T_{τ} calls of EXPLORE(i);
 - 6 **foreach** $i \in I_{\tau}$ **do**
 - 7 set $v_i^{(\tau)} \leftarrow \min\{\frac{1}{x_i^{(\tau)}} - 1, 1\}$,
 - 7 $a_i^{(\tau)} \leftarrow \max\{v_i^{(\tau)} - \epsilon_{\tau}, 0\}$, and
 - 7 $b_i^{(\tau)} \leftarrow \min\{v_i^{(\tau)} + \epsilon_{\tau}, 1\}$;
 - 8 let $\mathbf{a}^{(\tau)}$ be the vector containing the $|I_{\tau}|$ estimated preferences $a_i^{(\tau)}$, and $\mathbf{b}^{(\tau)}$ be the vector containing the $|I_{\tau}|$ estimated preferences $b_i^{(\tau)}$;
 - 9 $C \leftarrow \text{PRUNE}(I_{\tau}, \mathbf{a}^{(\tau)}, \mathbf{b}^{(\tau)})$;
 - 10 **if** $(|C| \leq K) \wedge \left(\bigwedge_{i \in C} (r_i > R(C, \mathbf{b}^{(\tau)})) \right)$ **then**
 - 11 **return** C ;
 - 12 $I_{\tau+1} \leftarrow C$;
-

to be included in the best assortment (Line 9). Finally, we test whether the output of PRUNE is indeed the best assortment (Line 10). If not we proceed to the next round, otherwise we return the solution.

Now we turn to the PRUNE subroutine (Algorithm 2), which is the most interesting part of the algorithm. Recall that the two vectors \mathbf{a} and \mathbf{b} are constructed such that $\mathbf{a} \preceq \mathbf{v} \preceq \mathbf{b}$. We try to prune items in I by the following test: For each $i \in I$, we form another vector \mathbf{g} such that $\mathbf{g} = \mathbf{a}$ in all coordinates except the i -th coordinate where $g_i = b_i$ (Line 4). We then check whether there exists a value $\theta \in [\theta_{\mathbf{a}}, \theta_{\mathbf{b}}]$ such that $i \in \text{Top}(I, \mathbf{g}, \theta)$, where $\theta_{\mathbf{a}}, \theta_{\mathbf{b}}$ are the maximum expected rewards with \mathbf{a} and \mathbf{b} as the item preference vectors respectively; if the answer is Yes then item i survives, otherwise it is pruned (Line 5). Note that our test is fairly conservative: we try to put item i in a more favorable position by using the upper bound b_i as its preference, while for other items we use the lower bounds a_j as their preferences. Such a conservative pruning step makes sure that the output C of the PRUNE subroutine is always a superset of the best assortment $S_{\mathbf{v}}$.

Theorem 5. *For any confidence parameter $\delta > 0$, Algorithm 3 returns the best assortment with probability $(1 - \delta)$ using at most $\Gamma = O\left(K^2 H_1 \ln\left(\frac{N}{\delta} \ln(KH_1)\right)\right)$ pulls. The running time of Algorithm 3 is bounded by $O\left(N\Gamma + N^2 \ln N \ln\left(\frac{K}{\min_{i \in I} \Delta_i}\right)\right)$.*

In the rest of this section we prove Theorem 5.

Correctness. We start by introducing the following event which we will condition on in the rest of the proof. The event states that in any round τ , the estimated preference $v_i^{(\tau)}$ for each item i (computed at Line 7 of Algorithm 3) is at most $\epsilon_\tau = 2^{-\tau-3}$ away from the true preference v_i .

$$\mathcal{E}_1 \triangleq \{\forall \tau \geq 0, \forall i \in I_\tau : |v_i^{(\tau)} - v_i| < \epsilon_\tau\}.$$

The proof of the following lemma can be found in Appendix B. This lemma states that event \mathcal{E}_1 holds with high probability.

Lemma 6. $\Pr[\mathcal{E}_1] \geq 1 - \delta$.

It is easy to see from Line 7 of Algorithm 3 that conditioned on \mathcal{E}_1 , we have

$$\forall \tau \geq 0 : \mathbf{a}^{(\tau)} \preceq \mathbf{v}^{(\tau)} \preceq \mathbf{b}^{(\tau)}, \quad (2)$$

where $\mathbf{v}^{(\tau)}$ is the preference vector of items in I_τ .

The following lemma shows that if (2) holds, then the PRUNE subroutine (Algorithm 2) always produces a set of candidate items C which is a superset of the best assortment.

Lemma 7. *If the preference vector \mathbf{v} of I satisfies $\mathbf{a} \preceq \mathbf{v} \preceq \mathbf{b}$, then $\text{PRUNE}(I, K, \mathbf{a}, \mathbf{b})$ (Algorithm 2) returns a set C such that $S_{\mathbf{v}} \subseteq C$.*

Proof. First, if $\mathbf{a} \preceq \mathbf{v} \preceq \mathbf{b}$, then by Lemma 3 we have $\theta_{\mathbf{v}} \in [\theta_{\mathbf{a}}, \theta_{\mathbf{b}}]$.

Consider any item $i \in S_{\mathbf{v}}$, by the construction of \mathbf{g} (Line 4 of Algorithm 2) we have for every $j \in I$:

- if $j \neq i$, then $(r_j - \theta_{\mathbf{v}})g_j \leq \max\{(r_j - \theta_{\mathbf{v}})v_j, 0\}$;
- if $j = i$, then $(r_j - \theta_{\mathbf{v}})g_j \geq (r_j - \theta_{\mathbf{v}})v_j$.

By these two facts and the definition of $\text{Top}(I, \mathbf{v}, \theta_{\mathbf{v}})$, we know that if $i \in \text{Top}(I, \mathbf{v}, \theta_{\mathbf{v}})$, then $i \in \text{Top}(I, \mathbf{g}, \theta_{\mathbf{v}})$. Therefore for the particular value $\theta = \theta_{\mathbf{v}} \in [\theta_{\mathbf{a}}, \theta_{\mathbf{b}}]$ we have $i \in \text{Top}(I, \mathbf{g}, \theta)$, and consequently i will be added to the candidate set C at Line 5, implying that $S_{\mathbf{v}} \subseteq C$. \square

Now suppose Algorithm 3 stops after round τ and outputs a set $C \supseteq S_{\mathbf{v}}$ of size at most K (Line 10-11), then for any $i \in C$, we have $r_i > \theta_{\mathbf{b}}$. By Lemma 3 we also have $\theta_{\mathbf{b}} \geq \theta_{\mathbf{v}}$ (since $\mathbf{v} \preceq \mathbf{b}$). We thus have $r_i > \theta_{\mathbf{v}}$. Consequently, it holds that for every $i \in C$, $(r_i - \theta_{\mathbf{v}}) > 0$. We thus have $C = S_{\mathbf{v}}$.

Up to this point we have shown that conditioned on \mathcal{E}_1 , if Algorithm 3 stops, then it outputs the best assortment $S_{\mathbf{v}}$. We next bound the number of pulls the algorithm uses.

Pull Complexity. We again conditioned on event \mathcal{E}_1 . The next lemma essentially states that an item $i \in I \setminus S_{\mathbf{v}}$ will be pruned if its reward gap Δ_i is much larger than K times its preference estimation error $\max\{b_i - v_i, v_i - a_i\}$.

Lemma 8. *In $\text{PRUNE}(I, K, \mathbf{a}, \mathbf{b})$ (Algorithm 2), if $\mathbf{a} \preceq \mathbf{v} \preceq \mathbf{b}$, and $\forall i \in I : \max\{b_i - v_i, v_i - a_i\} \leq \epsilon/K$ for any $\epsilon \in (0, 1)$, then any item $i \in I \setminus S_{\mathbf{v}}$ satisfying $\Delta_i > 8\epsilon$ will not be added to set C .*

Proof. By Corollary 4, if $\mathbf{a} \preceq \mathbf{v} \preceq \mathbf{b}$, and $\forall i \in I : \max\{b_i - v_i, v_i - a_i\} \leq \epsilon/K$, then we have

$$\theta_{\mathbf{v}} - \epsilon \leq \theta_{\mathbf{a}} \leq \theta_{\mathbf{v}} \leq \theta_{\mathbf{b}} \leq \theta_{\mathbf{v}} + \epsilon. \quad (3)$$

Consider any item $i \in I \setminus S_{\mathbf{v}}$ with $\Delta_i > 8\epsilon$. We analyze in two cases.

Case 1: $\theta_{\mathbf{v}} - r_i > 8\epsilon$. By (3) we have $\theta_{\mathbf{a}} - r_i > 7\epsilon$. Therefore, for any $\theta \in [\theta_{\mathbf{a}}, \theta_{\mathbf{b}}]$ we have $r_i < \theta_{\mathbf{a}} \leq \theta$, and consequently $i \notin \text{Top}(I, \mathbf{g}, \theta)$ for any $\theta \in [\theta_{\mathbf{a}}, \theta_{\mathbf{b}}]$ by the definition of $\text{Top}()$.

Case 2: $\theta_{\mathbf{v}} - r_i \leq 8\epsilon$. First, note that if $|S_{\mathbf{v}}| < K$, then we have

$$\Delta_i = -(r_i - \theta_{\mathbf{v}})v_i = (\theta_{\mathbf{v}} - r_i)v_i \leq \theta_{\mathbf{v}} - r_i \leq 8\epsilon,$$

contradicting our assumption that $\Delta_i > 8\epsilon$. We thus focus on the case that $|S_{\mathbf{v}}| = K$. We analyze two subcases.

1. $\theta \in (r_i, 1]$. In this case, by the definition of $\text{Top}()$ and the fact that $r_i - \theta < 0$, we have $i \notin \text{Top}(I, \mathbf{g}, \theta)$.
2. $\theta \in [\theta_{\mathbf{a}}, \theta_{\mathbf{b}}] \cap [0, r_i]$. For any $j \in S_{\mathbf{v}}$, we have

$$\begin{aligned} & (r_i - \theta)g_i - (r_j - \theta)g_j \\ &= (r_i - \theta)b_i - (r_j - \theta)a_j \\ &\leq (r_i - \theta)(v_i + \epsilon) - (r_j - \theta)a_j \quad (\text{since } r_i \geq \theta) \\ &\leq (r_i - \theta)v_i - (r_j - \theta)a_j + \epsilon \\ &\leq (r_i - \theta_{\mathbf{v}})v_i - (r_j - \theta_{\mathbf{v}})a_j + (1 + a_j + v_i)\epsilon \quad (\text{by (3)}) \\ &\leq (r_i - \theta_{\mathbf{v}})v_i - (r_j - \theta_{\mathbf{v}})(v_j - \epsilon) + 3\epsilon \quad (\text{since } r_j > \theta_{\mathbf{v}}) \\ &\leq (r_i - \theta_{\mathbf{v}})v_i - (r_j - \theta_{\mathbf{v}})v_j + 4\epsilon \\ &\leq -\Delta_i + 4\epsilon \\ &< -4\epsilon. \quad (\text{by the assumption } \Delta_i > 8\epsilon) \end{aligned}$$

We thus have that for any $\theta \in [\theta_{\mathbf{a}}, \theta_{\mathbf{b}}] \cap [0, r_i]$, $(r_i - \theta)g_i < (r_j - \theta)g_j$ for any $j \in S_{\mathbf{v}}$, therefore $i \notin \text{Top}(I, \mathbf{g}, \theta)$ for any $\theta \in [\theta_{\mathbf{a}}, \theta_{\mathbf{b}}]$, and consequently $i \notin C$. \square

For any $i \in I$, we define

$$\tau(i) \triangleq \min \left\{ \tau \geq 0 : \epsilon_\tau \leq \frac{\Delta_i}{32K} \right\}. \quad (4)$$

The next lemma shows that item i will *not* appear in any set I_τ with $\tau > \tau(i)$, and thus will not be pulled further after round $\tau(i)$.

Lemma 9. *In Algorithm 3, for any item $i \in I$, we have $i \notin I_\tau$ for any $\tau > \tau(i)$.*

Proof. For any $i \in I \setminus S_{\mathbf{v}}$, setting $\epsilon = \Delta_i/16$. By (4) we have that for any $j \in I_{\tau(i)}$ it holds that

$$\max \left\{ v_j - a_j^{(\tau(i))}, b_j^{(\tau(i))} - v_j \right\} \leq \frac{\Delta_i}{16K} = \frac{\epsilon}{K}. \quad (5)$$

Moreover, we have,

$$\Delta_i = 16\epsilon > 8\epsilon. \quad (6)$$

By (5), (6) and Lemma 8, we have $i \notin I_{\tau(i)+1}$.

We next consider items in $S_{\mathbf{v}}$. Note that by Definition 3, all $i \in S_{\mathbf{v}}$ have the same reward gap:

$$\Delta_i = \bar{\Delta} \triangleq \min \left\{ \min_{j \in I \setminus S_{\mathbf{v}}} \{\Delta_j\}, \min_{j \in S_{\mathbf{v}}} \{r_j - \theta_{\mathbf{v}}\} \right\} \leq \min_{j \in I \setminus S_{\mathbf{v}}} \{\Delta_j\}.$$

Let

$$\bar{\tau} \triangleq \min \left\{ \tau \geq 0 : \epsilon_\tau \leq \frac{\bar{\Delta}}{32K} \right\}. \quad (7)$$

We thus have $\bar{\tau} = \tau(i)$ for all $i \in S_v$, and $\bar{\tau} \geq \tau(j)$ for any $j \in I \setminus S_v$. Therefore, at the end of round $\bar{\tau}$, all items in $I \setminus S_v$ have already been pruned, and consequently,

$$|C| \leq K. \quad (8)$$

By (4) and Corollary 4 we have $\theta_{\mathbf{b}(\bar{\tau})} \leq \theta_v + \bar{\Delta}/16$. Consequently we have

$$\begin{aligned} r_i - R(C, \mathbf{b}(\bar{\tau})) &= r_i - \theta_{\mathbf{b}(\bar{\tau})} = (r_i - \theta_v) - (\theta_{\mathbf{b}(\bar{\tau})} - \theta_v) \\ &\geq \bar{\Delta} - \frac{\bar{\Delta}}{16} > 0. \end{aligned} \quad (9)$$

By (8) and (9), we know that Algorithm 3 will stop after round $\bar{\tau}$ and return $C = S_v$. \square

With Lemma 9 we can easily bound the total number of pulls made by Algorithm 3. By (4) we have $\tau(i) = O\left(\ln\left(\frac{K}{\Delta_i}\right)\right)$. By the definition of T_τ (Line 2 of Algorithm 3), the total number of pulls is at most

$$\begin{aligned} \sum_{i \in I} T_{\tau(i)} &\leq O\left(\sum_{i \in I} \frac{K^2}{\Delta_i^2} \ln \frac{N\tau^2(i)}{\delta}\right) \\ &= O\left(K^2 H_1 \ln\left(\frac{N}{\delta} \ln(KH_1)\right)\right). \end{aligned}$$

Remark 10. *The reason that we introduce an extra term $\min_{j \in S_v} \{r_j - \theta_v\}$ in the definition of reward gap Δ_i for all $i \in S_v$ (Definition 3) is for handling the case when $|S_v| < K$. More precisely, in the case $|S_v| < K$ we have to make sure that for all items $i \in I$ that we are going to add into the best assortment S_v , it holds that $r_i > \theta_v$. In our proof this is guaranteed by (9). On the other hand, if we are given the promise that $|S_v| = K$ (or $|S_v| = K'$ for a fixed value $K' \leq K$), then we do not need this extra term: we know when to stop simply by monitoring the size of I_τ , since at the end all items $i \in I/S_v$ will be pruned.*

Running Time. Finally, we analyze the time complexity of Algorithm 3. Although the time complexity of the algorithm is not the first consideration in the MNL-bandit model, we believe it is important for the algorithm to finish in a reasonable amount of time for real-time decision making. Observe that the running time of Algorithm 3 is dominated by the sum of the total number of pulls and the running time of the PRUNE subroutine, which is the main object that we shall bound next.

Let us analyze the running time of PRUNE. Let $n \triangleq |I|$. First, θ_a and θ_b can be computed in $O(n^2)$ time by an algorithm proposed by Rusmevichientong et al. [27]. We next show that Line 5 of Algorithm 2 can be implemented in $O(n \ln n)$ time, with which the total running time of PRUNE is bounded by $O(n^2 \ln n)$.

Consider any item $i \in I$. We can restrict our search of possible θ in the range of $\Theta_i = [\theta_a, \theta_b] \cap [0, r_i]$, since if

$i \in \text{Top}(I, \mathbf{g}, \theta)$, then by the definition of $\text{Top}()$ we have $\theta < r_i$. For each $j \neq i, j \in I$, define

$$\Theta_j = \{\theta \in \Theta_i \mid (r_j - \theta)g_j > (r_i - \theta)g_i\}.$$

Intuitively speaking, Θ_j contains all θ values for which item j is ‘‘preferred to’’ item i for $\text{Top}(I, \mathbf{g}, \theta)$. Consequently, for any $\theta \in \Theta_i$, if the number of Θ_j that contain θ is at least K , then we have $i \notin \text{Top}(I, \mathbf{g}, \theta)$; otherwise if the number of such Θ_j is less than K , then we have $i \in \text{Top}(I, \mathbf{g}, \theta)$. Note that each set Θ_j can be computed in $O(1)$ time.

Now think each set Θ_j as an interval. The problem of testing whether there exists a $\theta \in [\theta_a, \theta_b] \cap [0, r_i]$ such that $i \in \text{Top}(I, \mathbf{g}, \theta)$ can be reduced to the problem of checking whether there is a $\theta \in [\theta_a, \theta_b] \cap [0, r_i]$ such that θ is contained in fewer than K intervals Θ_j ($j \neq i$). The later problem can be solved by the standard sweep line algorithm in $O(n \ln n)$ time.

Recall that the total number of rounds can be bounded by $\tau_{\max} = \max_{i \in I} \tau(i) = O\left(\ln\left(\frac{K}{\min_{i \in I} \Delta_i}\right)\right)$. Therefore the total running time of Algorithm 3 can be bounded by

$$O\left(\Gamma + \sum_{\tau=0}^{\tau_{\max}} |I_\tau|^2 \ln |I_\tau|\right) = O\left(\Gamma + N^2 \ln N \ln\left(\frac{K}{\min_{i \in I} \Delta_i}\right)\right),$$

where $\Gamma = O\left(K^2 H_1 \ln\left(\frac{N}{\delta} \ln(KH_1)\right)\right)$ is the total number of pulls made by the algorithm.

4 The Improved Algorithm

In this section we try to improve our basic algorithm presented in Section 3. We design an algorithm whose pull complexity depends on H_2 which is asymptotically at most H_1 . The improved algorithm is described in Algorithm 5.

The structure of Algorithm 5 is very similar to that of Algorithm 3. The main difference is that instead of using EXPLORE to pull a singleton assortment at each time, we use a new procedure EXPLORESET (Algorithm 4) which pulls an assortment of size up to K (Line 6 of Algorithm 5). We construct the assortments by partitioning the whole set of items I_τ into subsets of size up to K (Line 4-5). In the EXPLORESET procedure, we keep pulling the assortment S until the output is 0 (i.e., a no-purchase decision is made). We then estimate the preference of item i using the average number of times that item i is chosen in those EXPLORESET calls that involve item i (Line 8).

Intuitively, EXPLORESET has the advantage over EXPLORE in that at each pull, the probability for EXPLORESET to return an item instead of a no-purchase decision is higher, and consequently EXPLORESET extracts more information about the item preferences. We note that the EXPLORESET procedure was first introduced in [4] in the setting of regret minimization.

Theorem 11. *For any confidence parameter $\delta > 0$, Algorithm 5 returns the best assortment with probability $(1 - \delta)$ using at most $\Gamma = O\left(K^2 H_2 \ln\left(\frac{N}{\delta} \ln(KH_2)\right)\right)$ pulls. The running time of Algorithm 5 is bounded by $O\left(N\Gamma + N^2 \ln N \ln\left(\frac{K}{\min_{i \in I} \Delta_i}\right)\right)$.*

Algorithm 4: EXPLORESET(S)

Input: a set of items S of size at most K .

Output: a set of empirical preferences $\{f_i\}_{i \in S}$.

- 1 Initialize $f_i \leftarrow 0$ for $i \in S$;
 - 2 **repeat**
 - 3 offer assortment S and observe a feedback a ;
 - 4 **if** $a \in S$ **then** $f_a \leftarrow f_a + 1$;
 - 5 **until** $a = 0$;
 - 6 **return** $\{f_i\}_{i \in S}$
-

Compared with Theorem 5, the only difference in the pull complexity of Theorem 11 is that we have used H_2 instead of H_1 . Since $H_2 = O(H_1)$, the asymptotic pull complexity of Algorithm 5 is at least as good as that of Algorithm 3.

Remark 12. *Though having a higher pull complexity, Algorithm 3 still has an advantage against Algorithm 5 in that Algorithm 3 can be implemented in the batched setting with $\max_{i \in I} \tau(i) = O\left(\ln \frac{K}{\min_{i \in I} \Delta_i}\right)$ policy changes, which cannot be achieved by Algorithm 5 since the subroutine EXPLORESET is inherently sequential.*

Compared with the proof for Theorem 5, the challenge for proving Theorem 11 is that the number of pulls in each EXPLORESET is a random variable. We thus need slightly more sophisticated mathematical tools to bound the sum of these random variables. Due to the space constraints, we leave the technical proof of Theorem 11 to Appendix C.

5 Lower Bound

We manage to show the following lower bound to complement our algorithmic results.

Theorem 13. *For any algorithm \mathcal{A} for pure exploration in multinomial logit bandit, there exists an input instance such that \mathcal{A} needs $\Omega(H_2/K^2)$ pulls to identify the best assortment with probability at least 0.6.*

Note that Algorithm 5 identifies the best assortment with probability 0.99 using at most $\tilde{O}(K^2 H_2)$ pulls (setting $\delta = 0.01$). Therefore our upper and lower bounds match up to a logarithmic factor if $K = O(1)$.

The proof of Theorem 13 bears some similarity with the lower bound proof of the paper by Chen et al. [10], but there are some notable differences. As mentioned in the introduction, Chen et al. [10] considered the problem of top- k ranking under the MNL choice model, which differs from the best assortment searching problem in the following aspects:

1. The top- k ranking problem can be thought as a special case of the best assortment searching problem where the rewards of all items are equal to 1. While to prove Theorem 13 we need to choose hard instances in which items have *different* rewards.
2. There is *no* null item (i.e., the option of “no purchase”) in the top- k ranking problem. Note that we cannot treat the null item as the $(N + 1)$ -th item with reward 0 since the null item will appear implicitly in every selected assortment.

Algorithm 5: Improved Fixed Confidence Algorithm for MNL-bandit

Input: a set of items $I = \{1, \dots, N\}$, a capacity parameter K , and a confidence parameter δ .

Output: the best assortment.

- 1 set $I_0 \leftarrow I$, and $\epsilon_\tau = 2^{-\tau-3}$ for $\tau \geq 0$;
 - 2 set $T_{-1} \leftarrow 0$, and $T_\tau \leftarrow \left\lceil \frac{8}{\epsilon_\tau^2} \ln \frac{16N(\tau+1)^2}{\delta} \right\rceil$ for $\tau \geq 0$;
 - 3 **for** $\tau = 0, 1, \dots$ **do**
 - 4 $m_\tau \leftarrow \lceil |I_\tau| / K \rceil$;
 - 5 let $S_1^\tau \uplus \dots \uplus S_{m_\tau}^\tau$ be an arbitrary partition of I_τ into subsets of size at most K ;
 - 6 **foreach** $j \in [m_\tau]$ **do** call EXPLORESET(S_j^τ) for $(T_\tau - T_{\tau-1})$ times ;
 - 7 **foreach** $i \in I_\tau$ **do**
 - 8 let $v_i^{(\tau)}$ be the average of f_i ’s returned by the multiset of calls $\{\text{EXPLORESET}(S_j^\rho) \mid \rho \leq \tau, j \in [m_\rho], i \in S_j^\rho\}$;
 - 9 **foreach** $i \in I_\tau$ **do** set $a_i^{(\tau)} \leftarrow \max\{0, v_i^{(\tau)} - \epsilon_\tau\}$ and $b_i^{(\tau)} \leftarrow \min\{v_i^{(\tau)} + \epsilon_\tau, 1\}$;
 - 10 let $\mathbf{a}^{(\tau)}$ be the vector containing the $|I_\tau|$ estimated preferences $a_i^{(\tau)}$, and $\mathbf{b}^{(\tau)}$ be the vector containing the $|I_\tau|$ estimated preferences $b_i^{(\tau)}$;
 - 11 $C \leftarrow \text{PRUNE}(I_\tau, \mathbf{a}^{(\tau)}, \mathbf{b}^{(\tau)})$;
 - 12 **if** $(|C| \leq K) \wedge \left(\bigwedge_{i \in C} (r_i > R(C, \mathbf{b}^{(\tau)})) \right)$ **then**
 - 13 **return** C ;
 - 14 $I_{\tau+1} \leftarrow C$;
-

These two aspects prevent us to use the lower bound result in Chen et al. [10] as a blackbox. Due to the space constraints, we leave the technical proof to Appendix D.

6 Concluding Remarks

We would like to conclude the paper by making two remarks. First, our upper and lower bounds match only for constant K . For general K , there is a factor of K^4 discrepancy. Part of it is because when we estimate the preference of each arm with an accuracy ϵ , we can only estimate the expected reward of an assortment with an accuracy ϵK ; this increases the pull complexity by a factor of K^2 . Obtaining tight bounds for general K is an interesting open question.

Second, our algorithms for pure exploration fall into the category of *fixed-confidence* algorithms, that is, for a fixed confidence parameter δ , we want to identify the best assortment with probability at least $(1 - \delta)$ using the smallest number of pulls. Another variant of pure exploration is called *fixed-budget* algorithms, where given a fixed pull budget T , we try to identify the best assortment with the highest probability. We leave this variant as future work.

References

- [1] Agarwal, A.; Johnson, N.; and Agarwal, S. 2020. Choice Bandits. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *NeurIPS*.
- [2] Agrawal, S.; Avadhanula, V.; Goyal, V.; and Zeevi, A. 2016. A Near-Optimal Exploration-Exploitation Approach for Assortment Selection. In *EC*, 599–600.
- [3] Agrawal, S.; Avadhanula, V.; Goyal, V.; and Zeevi, A. 2017. Thompson Sampling for the MNL-Bandit. In *COLT*, 76–78.
- [4] Agrawal, S.; Avadhanula, V.; Goyal, V.; and Zeevi, A. 2019. MNL-bandit: A dynamic learning approach to assortment selection. *Operations Research*, 67(5): 1453–1485.
- [5] Audibert, J.; Bubeck, S.; and Munos, R. 2010. Best Arm Identification in Multi-Armed Bandits. In *COLT*, 41–53.
- [6] Avadhanula, V. 2019. *The MNL-Bandit Problem: Theory and Applications*. Ph.D. thesis, Columbia University.
- [7] Bubeck, S.; Wang, T.; and Viswanathan, N. 2013. Multiple Identifications in Multi-Armed Bandits. In *ICML*, 258–265.
- [8] Chen, J.; Chen, X.; Zhang, Q.; and Zhou, Y. 2017. Adaptive multiple-arm identification. In *ICML*, 722–730.
- [9] Chen, L.; Li, J.; and Qiao, M. 2017. Towards Instance Optimal Bounds for Best Arm Identification. In *COLT*, volume 65, 535–592.
- [10] Chen, X.; Li, Y.; and Mao, J. 2018. A Nearly Instance Optimal Algorithm for Top- k Ranking under the Multinomial Logit Model. In Czumaj, A., ed., *SODA*, 2504–2522.
- [11] Chen, X.; and Wang, Y. 2018. A note on a tight lower bound for capacitated MNL-bandit assortment selection models. *Oper. Res. Lett.*, 46(5): 534–537.
- [12] Davis, J.; Gallego, G.; and Topaloglu, H. 2013. Assortment planning under the multinomial logit model with totally unimodular constraint structures. *Technical Report*.
- [13] Désir, A.; Goyal, V.; and Zhang, J. 2014. Near-optimal algorithms for capacity constrained assortment optimization. Available at SSRN 2543309.
- [14] Even-Dar, E.; Mannor, S.; and Mansour, Y. 2002. PAC Bounds for Multi-armed Bandit and Markov Decision Processes. In *COLT*, 255–270.
- [15] Gabillon, V.; Ghavamzadeh, M.; and Lazaric, A. 2012. Best Arm Identification: A Unified Approach to Fixed Budget and Fixed Confidence. In *NIPS*, 3221–3229.
- [16] Gabillon, V.; Ghavamzadeh, M.; Lazaric, A.; and Bubeck, S. 2011. Multi-Bandit Best Arm Identification. In *NIPS*, 2222–2230.
- [17] Garivier, A.; and Kaufmann, E. 2016. Optimal Best Arm Identification with Fixed Confidence. In *COLT*, 998–1027.
- [18] Jamieson, K.; Malloy, M.; Nowak, R.; and Bubeck, S. 2014. lil'ucb : An optimal exploration algorithm for multi-armed bandits. In *COLT*, 423–439.
- [19] Janson, S. 2018. Tail bounds for sums of geometric and exponential variables. *Statistics & Probability Letters*, 135: 1–6.
- [20] Jin, Y.; Li, Y.; Wang, Y.; and Zhou, Y. 2019. On Asymptotically Tight Tail Bounds for Sums of Geometric and Exponential Random Variables. *CoRR*, abs/1902.02852.
- [21] Karnin, Z.; Koren, T.; and Somekh, O. 2013. Almost optimal exploration in multi-armed bandits. In *ICML*, 1238–1246.
- [22] Kaufmann, E.; Cappé, O.; and Garivier, A. 2016. On the Complexity of Best-Arm Identification in Multi-Armed Bandit Models. *J. Mach. Learn. Res.*, 17: 1:1–1:42.
- [23] Kök, A. G.; and Fisher, M. L. 2007. Demand Estimation and Assortment Optimization Under Substitution: Methodology and Application. *Operations Research*, 55(6): 1001–1021.
- [24] Luce, R. D. 1959. *Individual choice behavior: a theoretical analysis*. Wiley.
- [25] Mannor, S.; and Tsitsiklis, J. N. 2004. The Sample Complexity of Exploration in the Multi-Armed Bandit Problem. *J. Mach. Learn. Res.*, 5: 623–648.
- [26] Plackett, R. 1975. The analysis of permutations. *Applied Statistics*, 24: 193–302.
- [27] Rusmevichientong, P.; Shen, Z. M.; and Shmoys, D. B. 2010. Dynamic Assortment Optimization with a Multinomial Logit Choice Model and Capacity Constraint. *Operations Research*, 58(6): 1666–1680.
- [28] Russo, D. 2016. Simple Bayesian Algorithms for Best Arm Identification. In *COLT*, volume 49, 1417–1418. JMLR.org.
- [29] Saha, A.; and Gopalan, A. 2018. Battle of Bandits. In Globerson, A.; and Silva, R., eds., *UAI*, 805–814.
- [30] Sauré, D.; and Zeevi, A. 2013. Optimal Dynamic Assortment Planning with Demand Learning. *Manufacturing & Service Operations Management*, 15(3): 387–404.
- [31] Talluri, K. T.; and van Ryzin, G. J. 2004. Revenue Management Under a General Discrete Choice Model of Consumer Behavior. *Management Science*, 50(1): 15–33.
- [32] Yang, J. 2021. Fully Gap-Dependent Bounds for Multinomial Logit Bandit. In *AISTATS*, volume 130 of *Proceedings of Machine Learning Research*, 199–207.
- [33] Zhou, Y.; Chen, X.; and Li, J. 2014. Optimal PAC multiple arm identification with applications to crowd-sourcing. In *ICML*, 217–225.

Appendix for Instance-Sensitive Algorithms for Pure Exploration in Multinomial Logit Bandit

A More Preliminaries

A.1 Tools in Probability Theory

We make use of the following standard concentration inequalities.

Lemma 14 (Hoeffding's inequality). *Let $X_1, \dots, X_n \in [0, 1]$ be independent random variables and $X = \sum_{i=1}^n X_i$.*

Then

$$\Pr[X > \mathbb{E}[X] + t] \leq \exp(-2t^2/n)$$

and

$$\Pr[X < \mathbb{E}[X] - t] \leq \exp(-2t^2/n).$$

Lemma 15 (Azuma's inequality). *Let the sequence Z_0, \dots, Z_n be a submartingale and*

$$\forall t \in [n] : |Z_t - Z_{t-1}| \leq d.$$

Then

$$\Pr[Z_n - Z_0 \leq -\epsilon] \leq \exp\left(\frac{-\epsilon^2}{2d^2n}\right).$$

Definition 7 (geometric random variable; the failure model). *Let $p \in [0, 1]$. If a random variable X with support \mathbb{Z}^+ satisfies $\Pr[X = k] = (1-p)^k p$ for any integer $k \geq 0$, then we say X follows the geometrical distribution with parameter p , denoted by $X \sim \text{Geo}(p)$.*

The following lemma gives the concentration result for sum of geometric random variables with a *multiplicative* error term.

Lemma 16 ([19]). *Let $p \geq 0$, $\lambda \geq 1$, and X_1, \dots, X_n be i.i.d. random variables from distribution $\text{Geo}(1/(1+p))$. We have*

$$\Pr\left[\sum_{i=1}^n (X_i + 1) \geq \lambda n(1+p)\right] \leq \exp(-n(\lambda - 1 - \ln \lambda)). \quad (10)$$

In our analysis we need the following concentration result for sum of geometric random variables with an *additive* error term.

Lemma 17. *Let $p \in [0, 1]$, $t \in [0, 1]$, and X_1, \dots, X_n be i.i.d. random variables from distribution $\text{Geo}(1/(1+p))$. We have*

$$\Pr\left[\left|\frac{1}{n} \sum_{i=1}^n (X_i - p)\right| \geq t\right] \leq 2 \exp\left(-\frac{nt^2}{8}\right). \quad (11)$$

Proof. We use the following lemma to derive Lemma 17.

Lemma 18 ([20]). *Let $p > 0$ and X_1, \dots, X_n be i.i.d. random variables from $\text{Geo}(1/(1+p))$, then for $\lambda \in (0, 1]$ we have*

$$\Pr\left[\frac{1}{n} \sum_{i=1}^n X_i \leq \lambda p\right] \leq \exp\left(-n \cdot \frac{p(\lambda - 1)^2}{2(1+p)}\right) \quad (12)$$

for $\lambda \in [1, 2]$

$$\Pr\left[\frac{1}{n} \sum_{i=1}^n X_i \geq \lambda p\right] \leq \exp\left(-n \cdot \frac{p(\lambda - 1)^2}{4(1+p)}\right) \quad (13)$$

and for $\lambda \geq 2$

$$\Pr\left[\frac{1}{n} \sum_{i=1}^n X_i \geq \lambda p\right] \leq \exp\left(-n \cdot \frac{p(\lambda - 1)}{4(1+p)}\right). \quad (14)$$

Note that the lemma holds trivially for $p = 0$. We thus focus on the case $p > 0$. We first show

$$\Pr\left[\frac{1}{n} \sum_{i=1}^n (X_i - p) \leq -t\right] \leq \exp\left(-\frac{nt^2}{8}\right). \quad (15)$$

We analyze in two cases.

1. If $1 \geq t \geq p$, then we have

$$\begin{aligned} & \Pr\left[\frac{1}{n} \sum_{i=1}^n (X_i - p) \leq -t\right] \\ & \leq \Pr\left[\sum_{i=1}^n X_i = 0\right] = \prod_{i=1}^n \Pr[X_i = 0] \\ & = \left(\frac{p}{1+p}\right)^n \leq 2^{-n} \leq \exp\left(-\frac{nt^2}{8}\right). \end{aligned}$$

2. If $t < p \leq 1$, then by (12), setting $\lambda = 1 - \frac{t}{p}$, we have

$$\begin{aligned} \Pr\left[\frac{1}{n} \sum_{i=1}^n (X_i - p) \leq -t\right] & \leq \exp\left(-\frac{np \cdot (t/p)^2}{2(1+p)}\right) \\ & \leq \exp\left(-\frac{nt^2}{8}\right). \end{aligned}$$

We next show

$$\Pr\left[\frac{1}{n} \sum_{i=1}^n (X_i - p) \geq t\right] \leq \exp\left(-\frac{nt^2}{8}\right). \quad (16)$$

We analyze in two cases.

1. If $t \leq p$, then by (13), setting $\lambda = 1 + \frac{t}{p}$, we have

$$\begin{aligned} \Pr\left[\frac{1}{n} \sum_{i=1}^n (X_i - p) \geq t\right] & \leq \exp\left(-\frac{np \cdot (t/p)^2}{4(1+p)}\right) \\ & \leq \exp\left(-\frac{nt^2}{8}\right). \end{aligned}$$

2. If $p < t \leq 1$, then by (14), setting $\lambda = 1 + \frac{t}{p}$, we have

$$\begin{aligned} \Pr\left[\frac{1}{n} \sum_{i=1}^n (X_i - p) \geq t\right] & \leq \exp\left(-\frac{np \cdot (t/p)}{4(1+p)}\right) \\ & \leq \exp\left(-\frac{nt^2}{8}\right). \end{aligned}$$

□

A.2 Proof of Observation 1

Proof. The observation follows directly from the definition of expected reward (Definition 1). That is, $R(S, \mathbf{v}) \geq \theta$ means $\frac{\sum_{i \in S} r_i v_i}{1 + \sum_{i \in S} v_i} \geq \theta$, which implies $\sum_{i \in S} (r_i - \theta) v_i \geq \theta$. The other direction can be shown similarly. \square

A.3 Proof of Lemma 2

The following is an easy observation by the definition of $\text{Top}(I, \mathbf{v}, \theta)$.

Observation 19. For any $S \subseteq I$ of size at most K and any $\theta \in [0, 1]$, it holds that

$$\sum_{i \in S} (r_i - \theta) v_i \leq \sum_{i \in \text{Top}(I, \mathbf{v}, \theta)} (r_i - \theta) v_i.$$

The following claim gives a crucial property of $\text{Top}(I, \mathbf{v}, \theta)$. Lemma 2 follows immediately from this claim.

Claim 20. For any $\theta \in [0, 1]$, $\theta \leq \theta_{\mathbf{v}}$ if and only if $\sum_{i \in \text{Top}(I, \mathbf{v}, \theta)} (r_i - \theta) v_i \geq \theta$.

Proof. First consider the case $\theta \leq \theta_{\mathbf{v}}$. By Observation 1, $R(S_{\mathbf{v}}, \mathbf{v}) = \theta_{\mathbf{v}} \geq \theta$ implies $\sum_{i \in S_{\mathbf{v}}} (r_i - \theta) v_i \geq \theta$. Then by Observation 19 we have $\sum_{i \in \text{Top}(I, \mathbf{v}, \theta)} (r_i - \theta) v_i \geq \sum_{i \in S_{\mathbf{v}}} (r_i - \theta) v_i \geq \theta$.

Next consider the case $\theta > \theta_{\mathbf{v}}$. For any $S \subseteq I$ with $|S| \leq K$, by the definition of $\theta_{\mathbf{v}}$ we have $R(S, \mathbf{v}) \leq \theta_{\mathbf{v}} < \theta$ for any $S \subseteq I$ with $|S| \leq K$. Consequently, we have $\sum_{i \in \text{Top}(I, \mathbf{v}, \theta)} (r_i - \theta) v_i < \theta$. \square

A.4 Proof of Lemma 3

Proof. If $\mathbf{v} \preceq \mathbf{w}$, then by definition of $S_{\mathbf{w}}$ we have $R(S_{\mathbf{w}}, \mathbf{w}) \geq R(S_{\mathbf{v}}, \mathbf{w})$, and

$$\begin{aligned} \sum_{i \in S_{\mathbf{v}}} (r_i - \theta_{\mathbf{v}}) w_i &\geq \sum_{i \in S_{\mathbf{v}}} (r_i - \theta_{\mathbf{v}}) v_i \quad (\mathbf{v} \preceq \mathbf{w}) \\ &\geq \theta_{\mathbf{v}} (R(S_{\mathbf{v}}, \mathbf{v}) \geq \theta_{\mathbf{v}} \text{ and Observation 1}) \\ &= R(S_{\mathbf{v}}, \mathbf{v}). \end{aligned}$$

We thus have $R(S_{\mathbf{w}}, \mathbf{w}) \geq R(S_{\mathbf{v}}, \mathbf{w}) \geq R(S_{\mathbf{v}}, \mathbf{v})$.

The second part of the lemma is due to the following simple calculation. Recall that $r_i \in (0, 1]$ for any $i \in S$.

$$R(S, \mathbf{w}) - R(S, \mathbf{v}) \leq \frac{\sum_{i \in S} r_i (w_i - v_i)}{1 + \sum_{i \in S} v_i} \leq \sum_{i \in S} (w_i - v_i). \quad \square$$

B Proof of Lemma 6

Proof. The output of $\text{EXPLORE}(i)$ is a Bernoulli random variable with mean $x_i = \frac{1}{1+v_i}$. By Hoeffding's inequality (Lemma 14) we have

$$\Pr \left[\left| x_i^{(\tau)} - x_i \right| \geq \frac{\epsilon_{\tau}}{8} \right] \leq 2 \exp \left(-\frac{\epsilon_{\tau}^2 T_{\tau}}{32} \right) \leq \frac{\delta}{8N(\tau+1)^2}.$$

By a union bound we have

$$\begin{aligned} &\Pr \left[\forall \tau \geq 0, \forall i \in I_{\tau} : \left| x_i^{(\tau)} - x_i \right| < \frac{\epsilon_{\tau}}{8} \right] \\ &\geq 1 - \sum_{\tau=0}^{\infty} \sum_{i \in I_{\tau}} \frac{\delta}{8N(\tau+1)^2} \geq 1 - \delta. \end{aligned} \quad (17)$$

Since at Line 7 of Algorithm 3 we have set $v_i^{(\tau)} = \frac{1}{x_i^{(\tau)}} - 1$, with probability $(1 - \delta)$ we have

$$\begin{aligned} \left| v_i^{(\tau)} - v_i \right| &= \left| \frac{1}{x_i^{(\tau)}} - \frac{1}{x_i} \right| = \left| \frac{x_i - x_i^{(\tau)}}{x_i^{(\tau)} x_i} \right| \\ &\leq \frac{\epsilon_{\tau}/8}{x_i^{(\tau)} x_i} \quad (\text{holds with prob. } (1 - \delta) \text{ by (17)}) \\ &\leq \frac{\epsilon_{\tau}/8}{1/2 \cdot 3/8} \\ &< \epsilon_{\tau}, \end{aligned}$$

where the second inequality holds since (i) $x_i = \frac{1}{1+v_i} \geq 1/2$ given $v_i \in [0, 1]$, and (ii) $x_i^{(\tau)} \geq 3/8$ given $\left| x_i^{(\tau)} - x_i \right| < \epsilon_{\tau}/8 < 1/8$. \square

C Proof of Theorem 11

First, we have the following two observations for the procedure EXPLORESET .

Observation 21 ([4]). For any $i \in S$, $f_i \sim \text{Geo}(1/(1+v_i))$.

Observation 22. The number of pulls made in $\text{EXPLORESET}(S)$ is $(X + 1)$ where $X \sim \text{Geo}(1/(1 + \sum_{i \in S} v_i))$.

Correctness. We define the following event which we will condition on in the rest of the proof.

$$\mathcal{E}_2 \triangleq \{ \forall \tau \geq 0, i \in I_{\tau} : \left| v_i^{(\tau)} - v_i \right| < \epsilon_{\tau} \} \quad (18)$$

We have the following lemma regarding \mathcal{E}_2 .

Lemma 23. $\Pr[\mathcal{E}_2] \geq 1 - \delta/2$.

Proof. By Observation 21 and Lemma 17, we have that for any $\tau \geq 0$ and $i \in I_{\tau}$, it holds that

$$\Pr \left[\left| v_i^{(\tau)} - v_i \right| \geq \epsilon_{\tau} \right] \leq 2 \exp \left(-\frac{\epsilon_{\tau}^2 T_{\tau}}{8} \right) \leq \frac{\delta}{8N(\tau+1)^2}. \quad (19)$$

By a union bound we have

$$\begin{aligned} \Pr[\bar{\mathcal{E}}_2] &\leq \sum_{\tau=0}^{\infty} \Pr \left[\left| v_i^{(\tau)} - v_i \right| \geq \epsilon_{\tau} \right] \\ &\leq \sum_{\tau=0}^{\infty} \sum_{i \in I_{\tau}} \frac{\delta}{8N(\tau+1)^2} \leq \frac{\delta}{2}. \end{aligned} \quad \square$$

By the same arguments as that for Theorem 5, we can show that Algorithm 5 returns the correct answer given that event \mathcal{E}_2 holds. Then by Lemma 23, Algorithm 5 succeeds with probability at least $1 - \delta/2$.

Pull Complexity. Now we turn to the number of pulls that Algorithm 5 makes. For any $i \in I$ we again define

$$\tau(i) \triangleq \min \left\{ \tau \geq 0 : \epsilon_\tau \leq \frac{\Delta_i}{32K} \right\}. \quad (20)$$

The following lemma is identical to Lemma 9 in the proof for Theorem 5.

Lemma 24. *In Algorithm 5, for any item $i \in I$, we have $i \notin I_\tau$ for any $\tau > \tau(i)$.*

We next show that Algorithm 5 will not make too many pulls in each round.

The following lemma is a direct consequence of Observation 22 and Lemma 16 (setting $\lambda = 5$).

Lemma 25. *For any $T > 0$, let random variables X_t ($t = 1, \dots, T$) be the number of pulls made at the t -th call EXPLORESET(S). We have*

$$\Pr \left[\sum_{t=1}^T X_t \geq 5 \left(1 + \sum_{i \in S} v_i \right) T \right] \leq \exp(-2T).$$

For each round τ , applying Lemma 25 with $T = T_\tau - T_{\tau-1}$ for each $S \in \{S_1^\tau, \dots, S_{m_\tau}^\tau\}$ we get

$$\begin{aligned} & \Pr \left[\sum_{t=1}^{T_\tau - T_{\tau-1}} X_t \geq 5 \left(1 + \sum_{i \in S} v_i \right) (T_\tau - T_{\tau-1}) \right] \\ & \leq \exp(-2(T_\tau - T_{\tau-1})) \\ & \leq \exp(-T_\tau) \leq \frac{\delta}{8N(\tau+1)^2}, \end{aligned}$$

where in the second inequality we have used the fact $T_\tau - T_{\tau-1} \geq T_\tau/2$ (by the definition of T_τ).

By a union bound over $S \in \{S_1^\tau, \dots, S_{m_\tau}^\tau\}$ and $\tau \geq 0$, with probability

$$1 - \sum_{\tau \geq 0} \left(m_\tau \cdot \frac{\delta}{8N(\tau+1)^2} \right) \geq 1 - \frac{\delta}{2}, \quad (21)$$

the total number of pulls made by Algorithm 5 is bounded by

$$\begin{aligned} & 5 \sum_{\tau \geq 0, I_\tau \neq \emptyset} \left(\left\lceil \frac{|I_\tau|}{K} \right\rceil + \sum_{i \in I_\tau} v_i \right) (T_\tau - T_{\tau-1}) \quad (22) \\ & \leq 5 \sum_{\tau \geq 0, I_\tau \neq \emptyset} \left(\frac{|I_\tau|}{K} + 1 + \sum_{i \in I_\tau} v_i \right) (T_\tau - T_{\tau-1}) \\ & = 5 \sum_{\tau \geq 0, I_\tau \neq \emptyset} (T_\tau - T_{\tau-1}) \quad (23) \\ & + 5 \sum_{\tau \geq 0, I_\tau \neq \emptyset} \left(\sum_{i \in I_\tau} \left(v_i + \frac{1}{K} \right) \right) (T_\tau - T_{\tau-1}) \quad (24) \end{aligned}$$

By Lemma 24 we know that for any $\tau > \bar{\tau} \triangleq \max_{i \in I} \{\tau(i)\}$, it holds that $I_\tau = \emptyset$. We thus have

$$\sum_{\tau \geq 0, I_\tau \neq \emptyset} (T_\tau - T_{\tau-1}) \leq T_{\bar{\tau}}. \quad (25)$$

Again by Lemma 24 we have

$$\sum_{\tau \geq 0} \left(\sum_{i \in I_\tau} \left(v_i + \frac{1}{K} \right) \right) (T_\tau - T_{\tau-1}) \leq \sum_{i \in I} \left(v_i + \frac{1}{K} \right) T_{\tau(i)}. \quad (26)$$

Combining (21), (24), (25), (26) and Lemma 23, we have that with probability $1 - (\delta/2 + \delta/2) = 1 - \delta$, the total number of pulls made by Algorithm 5 is bounded by

$$O \left(T_{\bar{\tau}} + \sum_{i \in I_\tau} \left(v_i + \frac{1}{K} \right) T_{\tau(i)} \right). \quad (27)$$

By the definitions of $\tau(i)$ and T_τ we have

$$T_{\tau(i)} = O \left(\frac{K^2}{\Delta_i^2} \cdot \ln \left(\frac{N}{\delta} \tau(i) \right) \right),$$

where $\tau(i) = O(\ln(K/\Delta_i)) = O(\ln(KH_2))$. Plugging these values to (27) we can bound the total number of pulls by $O(K^2 H_2 \ln(\frac{N}{\delta} \ln(KH_2)))$.

Running Time. The analysis of the running time of Algorithm 5 is very similar as that for Algorithm 3. The main difference is that the time complexity for each call of EXPLORESET is bounded $O(N\beta)$ (instead of $O(\beta)$ for EXPLORE) in the worst case, where β is the number of pulls in the call. This is why the first term in the time complexity in Theorem 11 is NT instead of T as that in Theorem 5. The second term concerning the PRUNE subroutine is the same as that in Theorem 5.

D Proof of Theorem 13 (The Lower Bound)

We consider the following two input instances. Let $\delta \in (0, \frac{1}{4K})$ be a parameter.

- *Instance I_1 .* I_1 contains $N = K$ items with rewards $r_1 = \dots = r_{K-1} = 1, r_K = \frac{1-\delta}{2-\delta}$, and preferences $v_1 = \dots = v_{K-1} = \frac{1}{K-1}, v_K = 1$.
- *Instance I_2 .* I_2 contains $N = K$ items with rewards $r_1 = \dots = r_{K-1} = 1, r_K = \frac{1-\delta}{2-\delta}$, and preferences $v_1 = \dots = \frac{1}{K-1} - 2\delta, v_2 = \dots = v_{K-1} = \frac{1}{K-1}, v_K = 1$.

Before proving Theorem 13, we first bound the instance complexities of I_1 and I_2 .

Instance complexity of I_1 . The optimal expected reward of I_1 is $1/2$, achieved on the set $[K-1]$. Indeed, all items from $[K-1]$ should be included in the best assortment since their rewards are all 1, and this already gives an expected reward of

$$\sum_{i \in [K-1]} \frac{1 \cdot v_i}{1 + \sum_{j \in [K-1]} v_j} = \frac{1}{2}.$$

While the reward of Item K is $\frac{1-\delta}{2-\delta} < \frac{1}{2}$, and thus Item K should be excluded in the best assortment.

By Definition 3, we have

$$\Delta_K = \frac{1}{2} - \frac{1-\delta}{2-\delta} = \frac{2-\delta-2+2\delta}{2(2-\delta)} \geq \frac{\delta}{4}.$$

For every $i \in [K - 1]$, we have

$$\Delta_i = \min \left\{ 1 - \frac{1}{2}, \Delta_K \right\} = \Delta_K.$$

We can thus bound

$$H_2(I_1) = \sum_{i \in [K]} \frac{v_i + 1/K}{\Delta_i^2} + \max_{i \in [K-1]} \left\{ \frac{1}{\Delta_i^2} \right\} \leq \frac{4}{\Delta_K^2} \leq \frac{64}{\delta^2}. \quad (28)$$

Instance complexity of I_2 . The optimal expected reward of I_2 is at least that of the assortment $[K]$, which can be bounded as

$$\left(\sum_{i \in [K-1]} \frac{1 \cdot v_i}{1 + \sum_{j \in [K]} v_j} \right) + \frac{\frac{1-\delta}{2-\delta} \cdot v_K}{1 + \sum_{j \in [K]} v_j} \geq \frac{1-2\delta}{2-2\delta}.$$

Thus, for every $i \in [K]$, we have

$$\Delta_K \geq \frac{1-\delta}{2-\delta} - \frac{1-2\delta}{2-2\delta} \geq \frac{\delta}{4}.$$

We can again bound

$$H_2(I_2) = \sum_{i \in [K]} \frac{v_i + 1/K}{\Delta_i^2} + \max_{i \in [K]} \left\{ \frac{1}{\Delta_i^2} \right\} \leq \frac{4}{\Delta_K^2} \leq \frac{64}{\delta^2}. \quad (29)$$

By (28) and (29), to prove Theorem 13 it suffices to show the following.

Lemma 26. *Any algorithm that uses less than $\frac{c}{4\delta^2 K^2}$ pulls for $c < 10^{-4}$ outputs the wrong answer on at least one instance among I_1 and I_2 with the probability at least 0.4.*

In the rest of this section we prove Lemma 26. We can focus on deterministic algorithms, since for any randomized algorithm we can always fix its randomness and obtain the deterministic algorithm with the smallest error on the input.

Let $\mathcal{T}_t = (U_1, o_1), \dots, (U_t, o_t)$ be the transcript of algorithm up to the t -th pull. We use $g_1(\mathcal{T}_t)$ and $g_2(\mathcal{T}_t)$ to denote the probabilities of observing the transcript \mathcal{T}_t on instances I_1 and I_2 respectively. The following lemma is the key for proving Lemma 26.

Lemma 27. *Let $c > 0$ and $T = \frac{c}{4\delta^2 K^2}$. For all $\epsilon > 0$, we have*

$$\Pr_{\mathcal{T}_T \sim g_1} \left[\ln \frac{g_2(\mathcal{T}_T)}{g_1(\mathcal{T}_T)} \leq -(\epsilon + c) \right] \leq \exp \left(\frac{-\epsilon^2}{9c} \right).$$

To see Lemma 27 implies Lemma 26, we set $\epsilon = \frac{1}{5}$, $c = \frac{1}{2250}$, and define event \mathcal{Q} as

$$\mathcal{Q} \triangleq \left\{ \ln \frac{g_2(\mathcal{T}_T)}{g_1(\mathcal{T}_T)} > -(\epsilon + c) \right\}. \quad (30)$$

By Lemma 27, it holds that $\Pr_{\mathcal{T}_T \sim g_1}[\bar{\mathcal{Q}}] \leq e^{-10}$. Let \mathcal{B} be

the event that algorithm \mathcal{A} outputs the set $[K - 1]$. We have

$$\begin{aligned} \Pr_{\mathcal{T}_T \sim g_1}[\mathcal{B}] &= \Pr_{\mathcal{T}_T \sim g_1}[\mathcal{B} \wedge \bar{\mathcal{Q}}] + \Pr_{\mathcal{T}_T \sim g_1}[\mathcal{B} \wedge \mathcal{Q}] \\ &\leq \Pr_{\mathcal{T}_T \sim g_1}[\bar{\mathcal{Q}}] + \Pr_{\mathcal{T}_T \sim g_1}[\mathcal{B} \wedge \mathcal{Q}] \\ &\leq e^{-10} + \Pr_{\mathcal{T}_T \sim g_1}[\mathcal{B} \wedge \mathcal{Q}] \\ &= e^{-10} + \sum_{\mathcal{T}_T: \mathcal{B} \wedge \mathcal{Q}} g_1(\mathcal{T}_T) \\ &\stackrel{(30)}{\leq} e^{-10} + e^{\epsilon+c} \sum_{\mathcal{T}_T: \mathcal{B} \wedge \mathcal{Q}} g_2(\mathcal{T}_T) \\ &\leq e^{-10} + e^{\epsilon+c} \Pr_{\mathcal{T}_T \sim g_2}[\mathcal{B}] \\ &= e^{-10} + e^{\epsilon+c} - e^{\epsilon+c} \Pr_{\mathcal{T}_T \sim g_2}[\bar{\mathcal{B}}]. \end{aligned}$$

Therefore, we have

$$\Pr_{\mathcal{T}_T \sim g_1}[\mathcal{B}] + e^{\epsilon+c} \Pr_{\mathcal{T}_T \sim g_2}[\bar{\mathcal{B}}] \leq e^{-10} + e^{\epsilon+c},$$

and consequently,

$$\min \left\{ \Pr_{\mathcal{T}_T \sim g_1}[\mathcal{B}], \Pr_{\mathcal{T}_T \sim g_2}[\bar{\mathcal{B}}] \right\} \leq \frac{e^{-10} + e^{\epsilon+c}}{1 + e^{\epsilon+c}} \leq 0.6. \quad (31)$$

(31) indicates that one of the followings hold: (1) Event \mathcal{B} holds with probability at most 0.6 when $\mathcal{T}_T \sim g_1$, and (2) Event $\bar{\mathcal{B}}$ holds with probability at most 0.6 when $\mathcal{T}_T \sim g_2$. In the first case, it indicates that algorithm \mathcal{A} errors on input instance I_1 with probability at least 0.4. In the second case, it indicates that algorithm \mathcal{A} errors on input instance I_2 with probability at least 0.4.

We now prove Lemma 27.

Proof. (of Lemma 27) We define a sequence of random variables Z_0, Z_1, \dots, Z_T when the transcript \mathcal{T}_t ($0 \leq t \leq T$) is produced by applying algorithm \mathcal{A} on the input instance I_1 :

$$Z_t = \ln \frac{g_2(\mathcal{T}_t)}{g_1(\mathcal{T}_t)}.$$

Let $V_t = \sum_{i \in U_t} v_i$. Z_i has the following properties.

- If $1 \notin U_t$, then $Z_t - Z_{t-1} = 0$, and $\mathbb{E}[Z_t - Z_{t-1} | Z_{t-1}] = 0$.
- If $1 \in U_t$, then with probability $\frac{1+V_t-v_1}{1+V_t}$,

$$Z_t - Z_{t-1} = -\ln \left(1 - \frac{2\delta}{1+V_t} \right),$$

and with probability $\frac{v_1}{1+V_t}$,

$$Z_t - Z_{t-1} = -\ln \left(1 - \frac{2\delta}{1+V_t} \right) + \ln \left(1 - \frac{2\delta}{v_1} \right).$$

We thus have

$$\begin{aligned} \mathbb{E}[Z_t - Z_{t-1} | Z_{t-1}] &= -\ln \left(1 - \frac{2\delta}{1+V_t} \right) \\ &\quad + \frac{v_1}{1+V_t} \ln \left(1 - \frac{2\delta}{v_1} \right) \end{aligned} \quad (32)$$

Using inequalities $\ln(1+x) \leq x$ and $\ln(1-x) \geq -x - x^2$ for $x \in [0, 0.5]$, and noting that $2\delta/v_1 = 2\delta(K-1) \leq 0.5$, we have

$$(32) \geq \frac{2\delta}{1+V_t} - \frac{2\delta}{1+V_t} - \frac{4\delta^2}{(1+V_t)v_1} \geq -\frac{4\delta^2(K-1)}{1+V_t} \geq -4\delta^2 K^2. \quad (33)$$

Note that in the case that $1 \notin U_t$, the inequality $\mathbb{E}[Z_t - Z_{t-1} \mid Z_{t-1}] = 0 \geq -4\delta^2 K^2$ holds trivially.

We can also bound the difference of two adjacent variables in the sequence $\{Z_0, Z_1, \dots, Z_T\}$.

$$|Z_t - Z_{t-1}| \leq \left| \ln \left(1 - \frac{2\delta}{1+V_t} \right) \right| + \left| \ln \left(1 - \frac{2\delta}{v_1} \right) \right| \leq 2\delta K. \quad (34)$$

Define $Z'_t \triangleq Z_t + 4\delta^2 K^2 t$. By (33) it follows that Z'_t is a submartingale and satisfies

$$\mathbb{E}[Z'_{t+1} \mid Z'_t] \geq Z'_t. \quad (35)$$

By (34) and the fact that $\delta < \frac{1}{4K}$, we have

$$|Z'_t - Z'_{t-1}| \leq 4\delta^2 K^2 + 2\delta K \leq 3\delta K. \quad (36)$$

By (36) and Azuma's inequality (Lemma 15), for $T = \frac{c}{4\delta^2 K^2}$, we get

$$\begin{aligned} \Pr_{\mathcal{T}_T \sim g_1} [Z_T \leq -(\epsilon + c)] &= \Pr_{\mathcal{T}_T \sim g_1} [Z'_T \leq -\epsilon] \\ &< \exp \left(\frac{-\epsilon^2}{18T\delta^2 K^2} \right) \leq \exp \left(\frac{-2\epsilon^2}{9c} \right). \end{aligned}$$

□

The lemma follows from (15) and (16).