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

Nikolai Karpov, Qin Zhang<br>Indiana University Bloomington<br>Luddy Hall, RM 3044<br>700 North Woodlawn Avenue<br>Bloomington, IN 47408-3901, USA<br>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 instancesensitive pull complexities. We also complement the upper bounds by an almost matching lower bound.


## Introduction

We study a model in online learning called multinomial logit bandit (MNL-bandit for short), where we have $N$ substitutable items $\{1,2, \ldots, N\}$, 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 realworld applications to be mentioned shortly.

Denote $[n] \triangleq\{1,2, \ldots, n\}$. 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 (1959) and Plackett (1975). 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 $\mathbf{v}=\left(v_{1}, \ldots, v_{N}\right)$, the expected reward of $S$ with respect to $\mathbf{v}$ is defined to be

$$
R(S, \mathbf{v})=\sum_{i \in S} r_{i} p_{i}(S)=\sum_{i \in S} \frac{r_{i} v_{i}}{1+\sum_{j \in S} v_{j}}
$$

[^0]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 Avadhanula (2019) 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 $v_{0} \geq \max \left\{v_{1}, \ldots, v_{N}\right\}$ 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 (Rusmevichientong, Shen, and Shmoys 2010; Sauré and Zeevi 2013; Davis, Gallego, and Topaloglu 2013; Agrawal et al. 2016, 2017; Chen and Wang 2018). The algorithms proposed by Rusmevichientong, Shen,
and Shmoys (2010); Sauré and Zeevi (2013) 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.

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 $\theta_{\mathbf{v}} \triangleq R\left(S_{\mathbf{v}}, \mathbf{v}\right)$ be the optimal expected reward.

Denote $\eta_{i} \triangleq\left(r_{i}-\theta_{\mathbf{v}}\right) v_{i}$; we call $\eta_{i}$ the advantage of item $i$. Suppose we have sorted the $N$ items according to $\eta_{i}$, let $\eta^{(j)}$ be the $j$-th largest value in the sorted list.
Definition 3 (reward gap $\Delta_{i}$ ). For any item $i \in[N] \backslash S_{\mathbf{v}}$, we define its reward gap to be

$$
\Delta_{i} \triangleq \begin{cases}\eta^{(K)}-\eta_{i}, & \text { if }\left|S_{\mathbf{v}}\right|=K \\ -\eta_{i}, & \text { if }\left|S_{\mathbf{v}}\right|<K\end{cases}
$$

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

$$
\begin{equation*}
\Delta_{i} \triangleq \bar{\Delta}=\min \left\{\left(\eta^{(K)}-\eta^{(K+1)}\right), \min _{j \in S_{\mathbf{v}}}\left\{r_{j}-\theta_{\mathbf{v}}\right\}\right\} \tag{1}
\end{equation*}
$$

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

$$
\Delta_{i}= \begin{cases}\eta^{(K)}-\eta_{i}, & \forall i \in[N] \backslash S_{\mathbf{v}} \\ \eta^{(K)}-\eta^{(K+1)}, & \forall i \in S_{\mathbf{v}}\end{cases}
$$

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.

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 from work of Audibert, Bubeck, and Munos (2010). Our definition of $\Delta_{i}$ is more involved due to the more complicated combinatorial structure of the MNL-bandit model.
Definition 5 (instance complexity $\mathrm{H}_{2}$ ).

$$
H_{2} \triangleq \sum_{i \in[N]} \frac{v_{i}+1 / K}{\Delta_{i}^{2}}+\max _{i \in[N]} \frac{1}{\Delta_{i}^{2}}
$$

It is easy to see that $H_{2}=O\left(H_{1}\right)$ (more precisely, $\frac{H_{1}}{K} \leq$ $H_{2} \leq 3 H_{1}$ ). We comment that the $\max _{i \in[N]} \frac{1}{\Delta_{i}^{2}}$ term is needed only when $\left|S_{\mathbf{v}}\right|<K$.
Our Results. We propose two fixed-confidence algorithms for pure exploration in MNL-bandit. Algorithm 3 gives a pull complexity of $O\left(K^{2} H_{1} \ln \left(\frac{N}{\delta} \ln \left(K H_{1}\right)\right)\right)$ where $\delta$ 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 $O\left(K^{2} H_{2} \ln \left(\frac{N}{\delta} \ln \left(K H_{2}\right)\right)\right)$. The second algorithm is presented in Algorithm 5.

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

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 $\Omega\left(H_{2} / K^{2}\right)$ pulls is needed in order to identify the best assortment with probability at least 0.6 . Note that when $K$ is a constant, our upper and lower bounds match up to a logarithmic factor.

## Related Work

Regret minimization in MNL-bandit was first studied by Rusmevichientong, Shen, and Shmoys (2010) in the setting of

[^1]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 (Sauré and Zeevi 2013; Davis, Gallego, and Topaloglu 2013; Agrawal et al. 2016, 2017; Chen and Wang 2018). Agarwal, Johnson, and Agarwal (2020) studied choice bandits which can be seen as a generalization of MNL-bandit.

As mentioned previously, the algorithms by Rusmevichientong, Shen, and Shmoys (2010); Sauré and Zeevi (2013) also have a component of identifying the best assortment. Rusmevichientong, Shen, and Shmoys (2010); Sauré and Zeevi (2013) introduced the following "assortment-level" gap:

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

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 Sauré and Zeevi (2013) for finding the best assortment can be written as $\tilde{O}\left(K N / \Delta_{\text {asso }}^{2}\right)$, where ' $\tilde{O}()$ ' hides some logarithmic factors. This result is better than that in Rusmevichientong, Shen, and Shmoys (2010). There are two critical differences between these results and our results: (1) More critically, in Rusmevichientong, Shen, and Shmoys (2010); Sauré and Zeevi (2013) it is assumed that the "assortment-level" gap $\Delta_{\text {asso }}$ 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 $\Delta_{i}$ 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 $\tilde{O}\left(K^{2} H_{1}\right)$ (or $\left.\tilde{O}\left(K^{2} H_{2}\right)\right)$. 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. $K=1, r_{1}=\ldots=r_{N}=1, v_{1}=1, v_{2}=1-$ $1 / \sqrt{N}, v_{3}=\ldots=v_{N}=1 / \sqrt{N}$. We have $K N / \Delta_{\text {asso }}=$ $\Omega\left(N^{2}\right)$, while $K^{2} H_{1}=O(N)$. Thus, the pull complexity of the algorithm in Rusmevichientong, Shen, and Shmoys (2010) is quadratic of ours (up to logarithmic factors).

The best assortment identification problem has also been studied in the static setting (Talluri and van Ryzin 2004; Désir, Goyal, and Zhang 2014), where the user preference vector $\mathbf{v}$ is known as a priori and our task is to conduct an offline computation to find the assortment that maximizes the expected reward. We refer readers to Kök and Fisher (2007) for an overview of this setting.

Chen, Li , and Mao (2018) 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 $\eta_{i}=\left(r_{i}-\theta_{\mathbf{v}}\right) v_{i}$ ). In some sense, their problem can be thought of as a special case of ours, where $r_{1}=r_{2}=\ldots=r_{N}$ (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 (2018), 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 (Even-Dar, Mannor, and Mansour 2002; Mannor and Tsitsiklis 2004; Audibert, Bubeck, and Munos 2010; Gabillon et al. 2011; Gabillon, Ghavamzadeh, and Lazaric 2012; Karnin, Koren, and Somekh 2013; Jamieson et al. 2014; Kaufmann, Cappé, and Garivier 2016; Garivier and Kaufmann 2016; Russo 2016; Chen, Li, and Qiao 2017). 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, Bubeck, and Munos (2010) designed an instance-sensitive algorithm for the pure exploration problem in the MAB model. The result of Audibert, Bubeck, and Munos (2010) was later improved by Karnin, Koren, and Somekh (2013) and Chen, Li, and Qiao (2017), and extended into the problem of identifying multiple items (Bubeck, Wang, and Viswanathan 2013; Zhou, Chen, and Li 2014; Chen et al. 2017).

Finally, we note that recently, concurrent and independent of our work, Yang (2021) has also studied pure exploration in MNL-bandit. But the definition of instance complexity in Yang (2021) is again at the "assortment-level" (and thus the results are not directly comparable), and the algorithmic approaches in Yang (2021) are also different from ours. The pull complexity of Yang (2021) can be written as $\tilde{O}\left(H_{\text {yang }}\right)$ where $H_{\text {yang }}=\sum_{i \in[N]} \frac{1}{\left(\Delta_{i}^{\prime}\right)^{2}}$, where $\Delta_{i}^{\prime}$ 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. $r_{1}=\ldots=r_{N}=1, v_{1}=\ldots=v_{K}=$ $1, v_{K+1}=\ldots=v_{N}=\epsilon$. For $\epsilon \in(0,1 / K)$ and $\omega(1) \leq K \leq o(N)$, we have $H_{\text {yang }}=\Theta\left(N K^{4}\right)$, while our $K^{2} H_{2}=\Theta\left(K^{5}+N K^{3}\right)=o\left(H_{y a n g}\right)$.

## 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 the full version of this paper.

The following (folklore) observation gives an effective way to check whether the expected reward of $S$ with respect to $\mathbf{v}$ is at least $\theta$ for a given value $\theta$. The proof can be found in the full version of this paper.
Observation 1. For any $\theta \in[0,1], R(S, \mathbf{v}) \geq \theta$ if and only if $\sum_{i \in S}\left(r_{i}-\theta\right) v_{i} \geq \theta$.

With Observation 1, to check whether the maximum expected reward is at least $\theta$ for a given value $\theta$, we only need

```
Algorithm 1: EXPLORE \((i)\)
    Input: Item \(i\).
    Output: \(0 / 1\) (choose or not choose \(i\) ).
    Offer a singleton set \(S_{i} \leftarrow\{i\}\) and observe a feedback
    \(a\);
    if \(a=0\) then return 1 ;
    return 0
```

to check whether the expected reward of the particular set $S \subseteq[N]$ containing the up to $K$ items with the largest positive values $\left(r_{i}-\theta\right) v_{i}$ is at least $\theta$.

To facilitate the future discussion we introduce the following definition.
Definition $6(\operatorname{Top}(I, \mathbf{v}, \theta))$. Given a set of items $I$ where the $i$-th item has reward $r_{i}$ and preference $v_{i}$, and a value $\theta$, let $T$ be the set of $\min \{K,|I|\}$ items with the largest values $\left(r_{i}-\theta\right) v_{i}$. Define $\operatorname{Top}(I, \mathbf{v}, \theta) \triangleq T \backslash\left\{i \in I \mid\left(r_{i}-\theta\right) \leq 0\right\}$, where $\mathbf{v}$ stands for $\left(v_{1}, \ldots, v_{|I|}\right)$.

The following lemma shows that $\operatorname{Top}\left(I, \mathbf{v}, \theta_{\mathbf{v}}\right)$ is exactly the best assortment. Its proof can be found in the full version of this paper.
Lemma 2. $\operatorname{Top}\left(I, \mathbf{v}, \theta_{\mathbf{v}}\right)=S_{\mathbf{v}}$.
Note that the set $\operatorname{Top}\left(I, \mathbf{v}, \theta_{\mathbf{v}}\right)$ is unique by its definition. Therefore by Lemma 2 the set $S_{\mathrm{v}}$ 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 dimension, we write $\mathbf{v} \preceq \mathbf{w}$ if $\forall i, v_{i} \leq w_{i}$. We comment that similar properties appeared in Agrawal et al. $(2016,2017)$, but were formulated a bit differently from ours. The proof of Lemma 3 can be found in the full version of this paper.
Lemma 3. If $\mathbf{v} \preceq \mathbf{w}$, then $\left(\theta_{\mathbf{v}}=\right) R\left(S_{\mathbf{v}}, \mathbf{v}\right) \leq\left(\theta_{\mathbf{w}}=\right.$ ) $R\left(S_{\mathbf{w}}, \mathbf{w}\right)$, and for any $S \subseteq I$ it holds that

$$
R(S, \mathbf{w})-R(S, \mathbf{v}) \leq \sum_{i \in S}\left(w_{i}-v_{i}\right)
$$

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 $\theta_{\mathbf{v}} \leq \theta_{\mathbf{w}} \leq$ $\theta_{\mathbf{v}}+\epsilon$.

## 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 $\tau$, each "surviving" item in the set $I_{\tau}$ has been pulled by $T_{\tau}$ times in total. We try to construct two vectors $\mathbf{a}$ and $\mathbf{b}$ based on the empirical means of the items in $I_{\tau}$ such that the (unknown) true preference vector $\mathbf{v}$ of $I_{\tau}$ is tightly sandwiched by a and $\mathbf{b}$ (Line 7-8). We then feed $I_{\tau}$, $\mathbf{a}$, and $\mathbf{b}$ to the Prune subroutine which reduces the size of $I_{\tau}$ by removing items that have no chance to be included in the best assortment (Line 9). Finally, we test

```
Algorithm 2: \(\operatorname{PrUNE}(I, K, \mathbf{a}, \mathbf{b})\)
    Input: a set of items \(I=\{1, \ldots, N\}\), capacity
                parameter \(K\), two vectors
                \(\mathbf{a}=\left(a_{1}, \ldots, a_{N}\right), \mathbf{b}=\left(b_{1}, \ldots, b_{N}\right)\) such that
                for any \(i \in[N]\) it holds that \(a_{i} \leq v_{i} \leq b_{i}\),
                where \(\mathbf{v}=\left(v_{1}, \ldots, v_{N}\right)\) is the (unknown)
                preference vector of the \(N\) items.
    Output: a set of candidate items for constructing the
                best assortment.
    \(\theta_{\mathbf{a}} \leftarrow \max _{S \subseteq I:|S| \leq K} R(S, \mathbf{a}), \theta_{\mathbf{b}} \leftarrow \max _{S \subseteq I:|S| \leq K} R(S, \mathbf{b}) ;\)
    \(C \leftarrow \emptyset ;\)
    foreach \(i \in I\) do
        form a vector \(\mathbf{g}=\left(g_{1}, \ldots, g_{N}\right)\) s.t. \(g_{j}=a_{j}\) for
        \(j \neq i\), and \(g_{i}=b_{i}\);
        if \(\exists \theta \in\left[\theta_{\mathbf{a}}, \theta_{\mathbf{b}}\right]\) s.t. \(i \in \operatorname{Top}(I, \mathbf{g}, \theta)\) then add \(i\) to
        \(C\);
    return \(C\)
```

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 $\theta \in\left[\theta_{\mathbf{a}}, \theta_{\mathbf{b}}\right]$ such that $i \in \operatorname{Top}(I, \mathbf{g}, \theta)$, 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 $\Gamma=O\left(K^{2} H_{1} \ln \left(\frac{N}{\delta} \ln \left(K H_{1}\right)\right)\right)$ 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 $\tau$, 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 $v_{i}$.

$$
\mathcal{E}_{1} \triangleq\left\{\forall \tau \geq 0, \forall i \in I_{\tau}:\left|v_{i}^{(\tau)}-v_{i}\right|<\epsilon_{\tau}\right\}
$$

The proof of the following lemma can be found in full version of this paper. This lemma states that event $\mathcal{E}_{1}$ holds with high probability.

```
Algorithm 3: The Fixed Confidence Algorithm for
MNL-Bandit
    Input: a set of items \(I=\{1, \ldots, N\}\), a capacity
            parameter \(K\), a confidence parameter \(\delta\).
    Output: the best assortment.
    \(I_{0} \leftarrow I\), set \(\epsilon_{\tau}=2^{-\tau-3}\) for \(\tau \geq 0\);
    set \(T_{-1} \leftarrow 0\) and \(T_{\tau} \leftarrow\left\lceil\frac{32}{\epsilon_{\tau}^{2}} \ln \frac{16 N(\tau+1)^{2}}{\delta}\right\rceil\) for \(\tau \geq 0\);
    for \(\tau=0,1, \ldots\) do
        foreach \(i \in I_{\tau}\) do call \(\operatorname{ExPlORE}(i)\) for
        \(\left(T_{\tau}-T_{\tau-1}\right)\) times;
        let \(x_{i}^{(\tau)}\) be the mean of the outputs of the \(T_{\tau}\) calls
        of EXPLORE \((i)\);
        foreach \(i \in I_{\tau}\) do
            set \(v_{i}^{(\tau)} \leftarrow \min \left\{\frac{1}{x_{i}^{(\tau)}}-1,1\right\}\),
                \(\begin{aligned} a_{i}^{(\tau)} & \leftarrow \max \left\{v_{i}^{(\tau)}-\epsilon_{\tau}, 0\right\}, \text { and } \\ b_{i}^{(\tau)} & \leftarrow \min \left\{v_{i}^{(\tau)}+\epsilon_{\tau}, 1\right\} ;\end{aligned}\)
        let \(\mathbf{a}^{(\tau)}\) be the vector containing the \(\left|I_{\tau}\right|\) estimated
        preferences \(a_{i}^{(\tau)}\), and \(\mathbf{b}^{(\tau)}\) be the vector
        containing the \(\left|I_{\tau}\right|\) estimated preferences \(b_{i}^{(\tau)}\);
        \(C \leftarrow \operatorname{Prune}\left(I_{\tau}, \mathbf{a}^{(\tau)}, \mathbf{b}^{(\tau)}\right) ;\)
        if \((|C| \leq K) \wedge\left(\bigwedge_{i \in C}\left(r_{i}>R\left(C, \mathbf{b}^{(\tau)}\right)\right)\right)\) then
            return \(C\);
        \(I_{\tau+1} \leftarrow C ;\)
```

Lemma 6. $\operatorname{Pr}\left[\mathcal{E}_{1}\right] \geq 1-\delta$.

It is easy to see from Line 7 of Algorithm 3 that conditioned on $\mathcal{E}_{1}$, we have

$$
\begin{equation*}
\forall \tau \geq 0: \mathbf{a}^{(\tau)} \preceq \mathbf{v}^{(\tau)} \preceq \mathbf{b}^{(\tau)} \tag{2}
\end{equation*}
$$

where $\mathbf{v}^{(\tau)}$ is the preference vector of items in $I_{\tau}$.
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 $\operatorname{Prune}(I, K, \mathbf{a}, \mathbf{b})$ (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 $\theta_{\mathbf{v}} \in\left[\theta_{\mathbf{a}}, \theta_{\mathbf{b}}\right]$.

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 $\left(r_{j}-\theta_{\mathbf{v}}\right) g_{j} \leq \max \left\{\left(r_{j}-\theta_{\mathbf{v}}\right) v_{j}, 0\right\}$;
- if $j=i$, then $\left(r_{j}-\theta_{\mathbf{v}}\right) g_{j} \geq\left(r_{j}-\theta_{\mathbf{v}}\right) v_{j}$.

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

Now suppose Algorithm 3 stops after round $\tau$ 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 $i \in C,\left(r_{i}-\theta_{\mathbf{v}}\right)>0$. 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 $i \in I \backslash S_{\mathbf{v}}$ will be pruned if its reward gap $\Delta_{i}$ is much larger than $K$ times its preference estimation error $\max \left\{b_{i}-v_{i}, v_{i}-a_{i}\right\}$.
Lemma 8. In $\operatorname{PRUNE}(I, K, \mathbf{a}, \mathbf{b})$ (Algorithm 2), if $\mathbf{a} \preceq \mathbf{v} \preceq$ b, and $\forall i \in I: \max \left\{b_{i}-v_{i}, v_{i}-a_{i}\right\} \leq \epsilon / K$ for any $\epsilon \in(0,1)$, then any item $i \in I \backslash S_{\mathbf{v}}$ 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 $\forall i \in I: \max \left\{b_{i}-\right.$ $\left.v_{i}, v_{i}-a_{i}\right\} \leq \epsilon / K$, then we have

$$
\begin{equation*}
\theta_{\mathbf{v}}-\epsilon \leq \theta_{\mathbf{a}} \leq \theta_{\mathbf{v}} \leq \theta_{\mathbf{b}} \leq \theta_{\mathbf{v}}+\epsilon \tag{3}
\end{equation*}
$$

Consider any item $i \in I \backslash S_{\mathbf{v}}$ 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 $\theta \in\left[\theta_{\mathbf{a}}, \theta_{\mathbf{b}}\right]$ we have $r_{i}<\theta_{\mathbf{a}} \leq \theta$, and consequently $i \notin \operatorname{Top}(I, \mathbf{g}, \theta)$ for any $\theta \in\left[\theta_{\mathbf{a}}, \theta_{\mathbf{b}}\right]$ by the definition of $\operatorname{Top}()$.

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

$$
\Delta_{i}=-\left(r_{i}-\theta_{\mathbf{v}}\right) v_{i}=\left(\theta_{\mathbf{v}}-r_{i}\right) v_{i} \leq \theta_{\mathbf{v}}-r_{i} \leq 8 \epsilon
$$

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

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

$$
\begin{aligned}
& \left(r_{i}-\theta\right) g_{i}-\left(r_{j}-\theta\right) g_{j} \\
= & \left(r_{i}-\theta\right) b_{i}-\left(r_{j}-\theta\right) a_{j} \\
\leq & \left(r_{i}-\theta\right)\left(v_{i}+\epsilon\right)-\left(r_{j}-\theta\right) a_{j} \\
\leq & \left(r_{i}-\theta\right) v_{i}-\left(r_{j}-\theta\right) a_{j}+\epsilon \\
\leq & \left(r_{i}-\theta_{\mathbf{v}}\right) v_{i}-\left(r_{j}-\theta_{\mathbf{v}}\right) a_{j}+\left(1+a_{j}+v_{i}\right) \epsilon \\
\leq & \left(r_{i}-\theta_{\mathbf{v}}\right) v_{i}-\left(r_{j}-\theta_{\mathbf{v}}\right)\left(v_{j}-\epsilon\right)+3 \epsilon \\
\leq & \left(r_{i}-\theta_{\mathbf{v}}\right) v_{i}-\left(r_{j}-\theta_{\mathbf{v}}\right) v_{j}+4 \epsilon \\
\leq & -\Delta_{i}+4 \epsilon \\
< & \left.-4 \epsilon . \quad \text { by the assumption } \Delta_{i}>8 \epsilon\right)
\end{aligned}
$$

We thus have that for any $\theta \in\left[\theta_{\mathbf{a}}, \theta_{\mathbf{b}}\right] \cap\left[0, r_{i}\right],\left(r_{i}-\theta\right) g_{i}<$ $\left(r_{j}-\theta\right) g_{j}$ for any $j \in S_{\mathbf{v}}$, therefore $i \notin \operatorname{Top}(I, \mathbf{g}, \theta)$ for any $\theta \in\left[\theta_{\mathbf{a}}, \theta_{\mathbf{b}}\right]$, and consequently $i \notin C$.

For any $i \in I$, we define

$$
\begin{equation*}
\tau(i) \triangleq \min \left\{\tau \geq 0: \epsilon_{\tau} \leq \frac{\Delta_{i}}{32 K}\right\} \tag{4}
\end{equation*}
$$

The next lemma shows that item $i$ will not appear in any set $I_{\tau}$ 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 $i \in I \backslash S_{\mathbf{v}}$, setting $\epsilon=\Delta_{i} / 16$. By (4) we have that for any $j \in I_{\tau(i)}$ it holds that

$$
\begin{equation*}
\max \left\{v_{j}-a_{j}^{(\tau(i))}, b_{j}^{(\tau(i))}-v_{j}\right\} \leq \frac{\Delta_{i}}{16 K}=\frac{\epsilon}{K} \tag{5}
\end{equation*}
$$

Moreover, we have,

$$
\begin{equation*}
\Delta_{i}=16 \epsilon>8 \epsilon \tag{6}
\end{equation*}
$$

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:

$$
\Delta_{i}=\min \left\{\min _{j \in I \backslash S_{\mathbf{v}}}\left\{\Delta_{j}\right\}, \min _{j \in S_{\mathbf{v}}}\left\{r_{j}-\theta_{\mathbf{v}}\right\}\right\} \leq \min _{j \in I \backslash S_{\mathbf{v}}}\left\{\Delta_{j}\right\} .
$$

Let

$$
\begin{equation*}
\bar{\tau} \triangleq \min \left\{\tau \geq 0: \epsilon_{\tau} \leq \frac{\bar{\Delta}}{32 K}\right\} \tag{7}
\end{equation*}
$$

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

$$
\begin{equation*}
|C| \leq K \tag{8}
\end{equation*}
$$

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

$$
\begin{align*}
r_{i}-R\left(C, \mathbf{b}^{(\bar{\tau})}\right) & =\left(r_{i}-\theta_{\mathbf{v}}\right)-\left(\theta_{\mathbf{b}}(\bar{\tau})\right. \\
& \left.\geq \theta_{\mathbf{v}}\right)  \tag{9}\\
& \geq \bar{\Delta}-\frac{\bar{\Delta}}{16}>0
\end{align*}
$$

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

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 $T_{\tau}$ (Line 2 of Algorithm 3), the total number of pulls is at most

$$
\begin{aligned}
\sum_{i \in I} T_{\tau(i)} & \leq O\left(\sum_{i \in I} \frac{K^{2}}{\Delta_{i}^{2}} \ln \frac{N \tau^{2}(i)}{\delta}\right) \\
& =O\left(K^{2} H_{1} \ln \left(\frac{N}{\delta} \ln \left(K H_{1}\right)\right)\right)
\end{aligned}
$$

Remark 10. The reason that we introduce an extra term $\min _{j \in S_{\mathbf{v}}}\left\{r_{j}-\theta_{\mathbf{v}}\right\}$ in the definition of reward gap $\Delta_{i}$ for all $i \in S_{\mathbf{v}}$ (Definition 3) is for handling the case when $\left|S_{\mathbf{v}}\right|<K$. More precisely, in the case $\left|S_{\mathbf{v}}\right|<K$ we have to make sure that for all items $i \in I$ that we are going to add into the best assortment $S_{\mathbf{v}}$, it holds that $r_{i}>\theta_{\mathbf{v}}$. In our proof this is guaranteed by (9). On the other hand, if we are given the promise that $\left|S_{\mathbf{v}}\right|=K$ (or $\left|S_{\mathbf{v}}\right|=K^{\prime}$ for a fixed value $K^{\prime} \leq K$ ), then we do not need this extra term: we know when to stop simply by monitoring the size of $I_{\tau}$, since at the end all items $i \in I / S_{\mathbf{v}}$ 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, $\theta_{\mathbf{a}}$ and $\theta_{\mathbf{b}}$ can be computed in $O\left(n^{2}\right)$ time by an algorithm proposed by Rusmevichientong, Shen, and Shmoys (2010). 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 $O\left(n^{2} \ln n\right)$.

Consider any item $i \in I$. We can restrict our search of possible $\theta$ in the range of $\Theta_{i}=\left[\theta_{\mathbf{a}}, \theta_{\mathbf{b}}\right] \cap\left[0, r_{i}\right)$, since if $i \in \operatorname{Top}(I, \mathbf{g}, \theta)$, then by the definition of $\operatorname{Top}()$ we have $\theta<r_{i}$. For each $j \neq i, j \in I$, define

$$
\Theta_{j}=\left\{\theta \in \Theta_{i} \mid\left(r_{j}-\theta\right) g_{j}>\left(r_{i}-\theta\right) g_{i}\right\}
$$

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

Now think each set $\Theta_{j}$ as an interval. The problem of testing whether there exists a $\theta \in\left[\theta_{\mathbf{a}}, \theta_{\mathbf{b}}\right] \cap\left[0, r_{i}\right)$ such that $i \in \operatorname{Top}(I, \mathbf{g}, \theta)$ can be reduced to the problem of checking whether there is a $\theta \in\left[\theta_{\mathbf{a}}, \theta_{\mathbf{b}}\right] \cap\left[0, r_{i}\right)$ such that $\theta$ is contained in fewer than $K$ intervals $\Theta_{j}(j \neq i)$. 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

$$
O\left(\Gamma+N^{2} \ln N \ln \left(\frac{K}{\min _{i \in I} \Delta_{i}}\right)\right)
$$

where $\Gamma=O\left(K^{2} H_{1} \ln \left(\frac{N}{\delta} \ln \left(K H_{1}\right)\right)\right)$ is the total number of pulls made by the algorithm.

## 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 $I_{\tau}$ 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

```
Algorithm 4: ExploreSet \((S)\)
    Input: a set of items \(S\) of size at most \(K\).
    Output: a set of empirical preferences \(\left\{f_{i}\right\}_{i \in S}\).
    Initialize \(f_{i} \leftarrow 0\) for \(i \in S\);
    repeat
        offer assortment \(S\) and observe a feedback \(a\);
        if \(a \in S\) then \(f_{a} \leftarrow f_{a}+1\);
    until \(a=0\);
    return \(\left\{f_{i}\right\}_{i \in S}\)
```

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 Agrawal et al. (2019) 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 $\Gamma=O\left(K^{2} H_{2} \ln \left(\frac{N}{\delta} \ln \left(K H_{2}\right)\right)\right)$ pulls. The running time of Algorithm 5 is bounded by $O\left(N \Gamma+N^{2} \ln N \ln \left(\frac{K}{\min _{i \in I} \Delta_{i}}\right)\right)$.

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 $H_{2}=O\left(H_{1}\right)$, 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 the full version of this paper.

## 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 $\Omega\left(H_{2} / K^{2}\right)$ 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 $\tilde{O}\left(K^{2} H_{2}\right)$ pulls (setting $\delta=$ 0.01 ). Therefore our upper and lower bounds match up to a logarithmic factor if $K=O(1)$.

```
Algorithm 5: Improved Fixed Confidence Algorithm
for MNL-bandit
    Input: a set of items \(I=\{1, \ldots, N\}\), a capacity
            parameter \(K\), and a confidence parameter \(\delta\).
    Output: the best assortment.
    set \(I_{0} \leftarrow I\), and \(\epsilon_{\tau}=2^{-\tau-3}\) for \(\tau \geq 0\);
    set \(T_{-1} \leftarrow 0\), and \(T_{\tau} \leftarrow\left\lceil\frac{8}{\epsilon_{\tau}^{2}} \ln \frac{16 N(\tau+1)^{2}}{\delta}\right\rceil\) for
    \(\tau \geq 0 ;\)
    for \(\tau=0,1, \ldots\) do
        \(m_{\tau} \leftarrow\left\lceil\left|I_{\tau}\right| / K\right\rceil ;\)
        let \(S_{1}^{\tau} \uplus \ldots \uplus S_{m_{\tau}}^{\tau}\) be an arbitrary partition of \(I_{\tau}\)
        into subsets of size at most \(K\);
        foreach \(j \in\left[m_{\tau}\right]\) do call \(\operatorname{ExploreSET}\left(S_{j}^{\tau}\right)\) for
        \(\left(T_{\tau}-T_{\tau-1}\right)\) times ;
        foreach \(i \in I_{\tau}\) do
                let \(v_{i}^{(\tau)}\) be the average of \(f_{i}\) 's returned by the
                multiset of calls \(\left\{\operatorname{ExploreSET}\left(S_{j}^{\rho}\right) \mid \rho \leq\right.\)
                \(\left.\tau, j \in\left[m_{\rho}\right], i \in S_{j}^{\rho}\right\} ;\)
        foreach \(i \in I_{\tau}\) do set \(a_{i}^{(\tau)} \leftarrow \max \left\{0, v_{i}^{(\tau)}-\epsilon_{\tau}\right\}\)
        and \(b_{i}^{(\tau)} \leftarrow \min \left\{v_{i}^{(\tau)}+\epsilon_{\tau}, 1\right\}\);
        let \(\mathbf{a}^{(\tau)}\) be the vector containing the \(\left|I_{\tau}\right|\) estimated
        preferences \(a_{i}^{(\tau)}\), and \(\mathbf{b}^{(\tau)}\) be the vector
        containing the \(\left|I_{\tau}\right|\) estimated preferences \(b_{i}^{(\tau)}\);
        \(C \leftarrow \operatorname{Prune}\left(I_{\tau}, \mathbf{a}^{(\tau)}, \mathbf{b}^{(\tau)}\right) ;\)
        if \((|C| \leq K) \wedge\left(\bigwedge_{i \in C}\left(r_{i}>R\left(C, \mathbf{b}^{(\tau)}\right)\right)\right)\) then
            return \(C\);
        \(I_{\tau+1} \leftarrow C\);
```

The proof of Theorem 13 bears some similarity with the lower bound proof of Chen, Li, and Mao (2018), but there are some notable differences. As mentioned in the introduction, Chen, Li, and Mao (2018) 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.

These two aspects prevent us to use the lower bound result in Chen, Li, and Mao (2018) as a blackbox. Due to the space constraints, we leave the technical proof to the full version of this paper.

## Acknowledgments

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

## References

Agarwal, A.; Johnson, N.; and Agarwal, S. 2020. Choice Bandits. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., NeurIPS.

Agrawal, S.; Avadhanula, V.; Goyal, V.; and Zeevi, A. 2016. A Near-Optimal Exploration-Exploitation Approach for Assortment Selection. In EC, 599-600.
Agrawal, S.; Avadhanula, V.; Goyal, V.; and Zeevi, A. 2017. Thompson Sampling for the MNL-Bandit. In COLT, 76-78.
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.
Audibert, J.; Bubeck, S.; and Munos, R. 2010. Best Arm Identification in Multi-Armed Bandits. In COLT, 41-53.
Avadhanula, V. 2019. The MNL-Bandit Problem: Theory and Applications. Ph.D. thesis, Columbia University.
Bubeck, S.; Wang, T.; and Viswanathan, N. 2013. Multiple Identifications in Multi-Armed Bandits. In ICML, 258-265.
Chen, J.; Chen, X.; Zhang, Q.; and Zhou, Y. 2017. Adaptive multiple-arm identification. In ICML, 722-730.
Chen, L.; Li, J.; and Qiao, M. 2017. Towards Instance Optimal Bounds for Best Arm Identification. In COLT, volume 65, 535-592.
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.
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.
Davis, J.; Gallego, G.; and Topaloglu, H. 2013. Assortment planning under the multinomial logit model with totally unimodular constraint structures. Technical Report.
Désir, A.; Goyal, V.; and Zhang, J. 2014. Near-optimal algorithms for capacity constrained assortment optimization. Available at SSRN 2543309.
Even-Dar, E.; Mannor, S.; and Mansour, Y. 2002. PAC Bounds for Multi-armed Bandit and Markov Decision Processes. In COLT, 255-270.
Gabillon, V.; Ghavamzadeh, M.; and Lazaric, A. 2012. Best Arm Identification: A Unified Approach to Fixed Budget and Fixed Confidence. In NIPS, 3221-3229.
Gabillon, V.; Ghavamzadeh, M.; Lazaric, A.; and Bubeck, S. 2011. Multi-Bandit Best Arm Identification. In NIPS, 2222-2230.
Garivier, A.; and Kaufmann, E. 2016. Optimal Best Arm Identification with Fixed Confidence. In COLT, 998-1027.
Jamieson, K.; Malloy, M.; Nowak, R.; and Bubeck, S. 2014. lil'ucb: An optimal exploration algorithm for multi-armed bandits. In COLT, 423-439.

Janson, S. 2018. Tail bounds for sums of geometric and exponential variables. Statistics \& Probability Letters, 135: 1-6.
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.
Karnin, Z.; Koren, T.; and Somekh, O. 2013. Almost optimal exploration in multi-armed bandits. In ICML, 1238-1246.
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.
Kök, A. G.; and Fisher, M. L. 2007. Demand Estimation and Assortment Optimization Under Substitution: Methodology and Application. Operations Research, 55(6): 1001-1021.
Luce, R. D. 1959. Individual choice behavior: a theoretical analysis. Wiley.
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.
Plackett, R. 1975. The analysis of permutations. Applied Statistics, 24: 193-302.
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.
Russo, D. 2016. Simple Bayesian Algorithms for Best Arm Identification. In COLT, volume 49, 1417-1418. JMLR.org. Saha, A.; and Gopalan, A. 2018. Battle of Bandits. In Globerson, A.; and Silva, R., eds., UAI, 805-814.
Sauré, D.; and Zeevi, A. 2013. Optimal Dynamic Assortment Planning with Demand Learning. Manufacturing \& Service Operations Management, 15(3): 387-404.
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.
Yang, J. 2021. Fully Gap-Dependent Bounds for Multinomial Logit Bandit. In AISTATS, volume 130 of Proceedings of Machine Learning Research, 199-207.
Zhou, Y.; Chen, X.; and Li, J. 2014. Optimal PAC multiple arm identification with applications to crowdsourcing. In ICML, 217-225.


[^0]:    Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

[^1]:    ${ }^{1}$ When 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.

