Tractable Multinomial Logit Contextual Bandits with Non-Linear Utilities

Taehyun Hwang Seoul National University th.hwang@snu.ac.kr **Dahngoon Kim** Seoul National University dahngoon.k@snu.ac.kr

Min-hwan Oh Seoul National University minoh@snu.ac.kr

Abstract

We study the multinomial logit (MNL) contextual bandit problem for sequential assortment selection. Although most existing research assumes utility functions to be linear in item features, this linearity assumption restricts the modeling of intricate interactions between items and user preferences. A recent work [41] has investigated general utility function classes, yet its method faces fundamental tradeoffs between computational tractability and statistical efficiency. To address this limitation, we propose a computationally efficient algorithm for MNL contextual bandits leveraging the upper confidence bound principle, specifically designed for non-linear parametric utility functions, including those modeled by neural networks. Under a realizability assumption and a mild geometric condition on the utility function class, our algorithm achieves a regret bound of $\mathcal{O}(\sqrt{T})$, where T denotes the total number of rounds. Our result establishes that sharp $\mathcal{O}(\sqrt{T})$ -regret is attainable even with neural network-based utilities, without relying on strong assumptions such as neural tangent kernel approximations. To the best of our knowledge, our proposed method is the first computationally tractable algorithm for MNL contextual bandits with non-linear utilities that provably attains $\mathcal{O}(\sqrt{T})$ regret. Comprehensive numerical experiments validate the effectiveness of our approach, showing robust performance not only in realizable settings but also in scenarios with model misspecification.

1 Introduction

The *multinomial logit* (MNL) contextual bandit [5, 30, 28, 4] is a model for sequential assortment selection under uncertainty prevalent in applications such as online recommendation systems, personalized marketing, and online retailing. In this setting, an agent repeatedly selects subsets of items to present to users, where user choice behavior is modeled using the MNL model [27]. Specifically, the probability that a user selects an item is determined by its latent utility, which typically depends on contextual features associated with items and possibly user characteristics. The primary goal is to minimize expected regret, defined as the cumulative difference in rewards obtained by the algorithm compared to an oracle that knows the true utility parameters in advance.

Almost all prior studies on MNL contextual bandits assume the item utility functions to be linear in contextual features [5, 30, 28, 4, 29, 32, 2, 44, 20, 21]. Although the linearity assumption simplifies computations and theoretical analyses, it significantly limits the expressive power required to capture complex and nonlinear user behaviors typically observed in real-world scenarios. Recent work by Zhang and Luo [41] addresses this limitation by considering more general utility functions. However, they encounter critical trade-offs: algorithms with computational tractability yield suboptimal regret (e.g., $\widetilde{\mathcal{O}}(T^{2/3})$) in stochastic and $\widetilde{\mathcal{O}}(T^{5/6})$ in adversarial settings), whereas their statistically efficient approaches become computationally infeasible due to non-convex optimization challenges.

Consequently, designing computationally efficient algorithms that achieve optimal statistical guarantees under general, non-linear utility functions remains an open and significant challenge.

Motivated by these fundamental limitations, we investigate the MNL contextual bandit problem with utilities represented by general, non-linear parametric models, such as neural networks. This setting introduces several fundamental challenges compared to the linear case. First, the negative log-likelihood function used for parameter estimation lacks the self-concordant-like properties crucial for theoretical analysis in linear utility functions [36]. Second, gradient-based analyses that rely on parameter-independent gradients in linear models do not extend naturally to the non-linear setting, as gradients become parameter-dependent and more intricate. Lastly, non-linear utility functions typically induce a non-convex optimization landscape characterized by multiple global minima due to symmetries (e.g., neuron permutations in neural networks), complicating both the optimization process and subsequent regret analysis.

Despite these challenges, this paper presents a computationally tractable algorithm for MNL contextual bandits that can handle non-linear utility functions and achieve a sharp regret bound. Our key contributions are as follows:

- We propose an *upper confidence bound* (UCB)-based algorithm designed explicitly for MNL contextual bandits with non-linear utilities. Under a realizability assumption and a newly proposed generalized geometric condition, our algorithm achieves a regret bound of $\widetilde{\mathcal{O}}(\sqrt{T})$. To our knowledge, this constitutes the first algorithm with such near-optimal regret guarantees independent of the total number of items N in this non-linear utility setting.
- Unlike the existing approach [41], which either compromise statistical efficiency or computational feasibility, our algorithm simultaneously attains computational tractability and the sharp $\widetilde{\mathcal{O}}(\sqrt{T})$ regret bound under non-linear utility settings.
- To overcome technical challenges arising in parameter learning for non-linear utility based MNL models, we introduce a novel generalized geometric condition on the squared error of the utility function (Assumption 4). This condition significantly generalizes and weakens assumptions used in previous studies [25, 22, 34] and accommodates broad classes of non-linear functions, including neural networks. Moreover, we derive a new concentration inequality for parameter estimation in MNL models, independent of the regularization parameter up to logarithmic factors. These technical advancements allow us to construct computationally tractable, optimistic utility estimators suitable for efficient exploration. Crucially, our approach circumvents the commonly used *neural tangent kernel* (NTK) assumptions [18], eliminating the impractical requirement of extensive over-parameterization prevalent in NTK-based analyses.
- We demonstrate the superior performance of our algorithm through extensive numerical
 experiments. While benchmark methods suffer significant performance degradation in nonlinear utility scenarios, our proposed method maintains robust performance in both realizable
 and misspecified contexts, underscoring its practical effectiveness and adaptability.

2 Preliminary

Notations. For a positive integer $N \in \mathbb{N}$, we denote the set $\{1,\ldots,N\}$ by [N]. For a vector $\mathbf{x} \in \mathbb{R}^d$, its ℓ_p -norm is defined as $\|\mathbf{x}\|_p = (\sum_{j=1}^d |x_j|^p)^{1/p}$ for $1 \leq p < \infty$, and its ℓ_∞ -norm is defined as $\|\mathbf{x}\|_\infty = \max_{j \in [d]} |x_j|$. For a matrix \mathbf{A} , its operator norm is denoted by $\|\mathbf{A}\|_{\mathrm{op}}$. For a vector $\mathbf{x} \in \mathbb{R}^d$ and a symmetric positive semi-definite matrix $\mathbf{A} \in \mathbb{R}^{d \times d}$, we define $\|\mathbf{x}\|_{\mathbf{A}} := \sqrt{\mathbf{x}^\top \mathbf{A} \mathbf{x}}$. We denote a real-valued function defined on \mathcal{X} and parameterized by \mathbf{w} as $f_{\mathbf{w}} : \mathcal{X} \to \mathbb{R}$, and denote its gradient and Hessian with respect to \mathbf{w} by $\nabla f_{\mathbf{w}}$ and $\nabla^2 f_{\mathbf{w}}$, respectively. Throughout, we use standard big- \mathcal{O} notation to hide universal constants, and $\widetilde{\mathcal{O}}$ notation to additionally hide logarithmic and poly-logarithmic factors.

2.1 Problem Setting

MNL Contextual Bandits. We consider a sequential assortment selection problem in which, at each round t, the agent observes a set of feature vectors for all items, denoted by $X_t := \{\mathbf{x}_{t1}, \dots, \mathbf{x}_{tN}\} \subset$

Table 1: Comparison of results and main assumptions for prior MNL contextual bandits with non-linear utilities. Zhang and Luo [41] consider a bounded and Lipschitz-continuous utility class, whereas our work assumes a bounded, Lipschitz, and smooth utility class. The geometric condition (Assumption 4) is imposed in Phase I to establish the convergence rate of the pilot estimator. In contrast, the analysis in [41] relies on the generalization error of the (offline or online) regression oracle, measured in terms of log-loss, rather than the distance between the estimator and the true parameter. The dimensionality d_w appearing in the regret bounds of [41] corresponds to the dimension of the Lipschitz function parameter, which may differ from the context dimension d.

Algorithm	Utility	Assumptions	Tractability [†]	Regret
ε -greedy [41]	Lipschitzness	Offline reg. oracle	✓	$\widetilde{\mathcal{O}}((d_w NK)^{1/3} T^{2/3})$
log-barrier regularizer [41]	Lipschitzness	Offline reg. oracle	X	$\widetilde{\mathcal{O}}(K^2\sqrt{d_wNT})$
ε -greedy [41]	Lipschitzness	Online reg. oracle	✓	$\widetilde{\mathcal{O}}((NK)^{1/3}T^{5/6})$
Feel-Good TS[41]	Lipschitzness	Online reg. oracle	X	$\widetilde{\mathcal{O}}(K^2\sqrt{d_wNT})$
ONL-MNL (this work)	Smoothness	Geometric condition	✓	$\widetilde{\mathcal{O}}(\kappa^{-2}\mu^{-1}d_w\sqrt{T})$

[†]Tractability of each method refers to the computational feasibility of the exploration strategy, i.e., the process of selecting an assortment given a current utility estimator. The computation required to obtain our pilot estimator is comparable to what is already used in [41]. Following [41], any tractable empirical risk minimization method (e.g., SGD) is sufficient for computing our pilot estimator in practice.

 \mathbb{R}^d , where each feature vector belongs to a general context space $\mathcal{X} \subset \mathbb{R}^d$. Based on this information, the agent selects an assortment $S_t = \{i_1, \dots, i_l\} \in \mathcal{S} := \{S \subset [N] : |S| \leq K\}$, where $l \leq K$, and observes a user choice $i_t \in S_t \cup \{0\}$, where 0 denotes the "outside option", i.e., the user chooses none of the offered items. The user's selection $i_t \in S_t \cup \{0\}$ is modeled by the MNL choice model:

$$\mathbb{P}(i_t = i \mid X_t, S_t) =: p(i \mid X_t, S_t, \mathbf{w}^*) := \frac{\exp(f_{\mathbf{w}^*}(\mathbf{x}_{ti}))}{1 + \sum_{j \in S_t} \exp(f_{\mathbf{w}^*}(\mathbf{x}_{tj}))},$$

$$\mathbb{P}(i_t = 0 \mid X_t, S_t) =: p(0 \mid X_t, S_t, \mathbf{w}^*) := \frac{1}{1 + \sum_{j \in S_t} \exp(f_{\mathbf{w}^*}(\mathbf{x}_{tj}))},$$

where $f_{\mathbf{w}^*}(\cdot): \mathcal{X} \to \mathbb{R}$ is an *unknown* utility function, specifying the user's value for each item under the given context. At each round t, the agent receives feedback in the form of a choice response indicating which item from the offered set $S_t \cup \{0\}$ was selected. The choice outcome is represented as a one-hot vector $\mathbf{y}_t := (y_{t0}, y_{t1}, \dots, y_{tl})$, where $y_{ti} = 1$ if item i was chosen and 0 otherwise. By definition of the model, this choice is drawn according to the following MNL distribution:

$$\mathbf{y}_t \sim \text{Multinomial} \{1, [p(0 \mid X_t, S_t, \mathbf{w}^*), p(i_1 \mid X_t, S_t, \mathbf{w}^*), \dots, p(i_l \mid X_t, S_t, \mathbf{w}^*)]\}$$

where 1 denotes \mathbf{y}_t is a single-trial sample, i.e., $y_{t0} + \sum_{k=1}^{l} y_{tk} = 1$.

For a given revenue parameter vector $\mathbf{r}_t = [r_{t0}, r_{t1}, \dots, r_{tN}]$, where $r_{t0} = 0$ denotes the revenue associated with the outside option and r_{ti} is the revenue for item $i \in [N]$, the expected reward of offering an assortment $S \in \mathcal{S}$ under the context feature set X_t is defined as

$$R_{X_t,\mathbf{r}_t}(S,\mathbf{w}^*) =: R_t(S,\mathbf{w}^*) = \sum_{i \in S} p(i \mid X_t, S, \mathbf{w}^*) r_{ti} = \sum_{i \in S} \frac{\exp(f_{\mathbf{w}^*}(\mathbf{x}_{ti})) r_{ti}}{1 + \sum_{j \in S} \exp(f_{\mathbf{w}^*}(\mathbf{x}_{tj}))}.$$

The goal of the agent is to minimize the cumulative regret over T rounds, defined as the total difference in expected reward between the offline optimal assortment $S_t^* := \operatorname{argmax}_{S \in \mathcal{S}} R_t(S, \mathbf{w}^*)$

and the assortment chosen by the agent,
$$\operatorname{Regret}_T := \sum_{t=1}^T \left[R_t(S_t^*, \mathbf{w}^*) - R_t(S_t, \mathbf{w}^*) \right].$$

Non-Linear Parametric Functions. We consider a parametric function class $\mathcal{F}:=\{f_{\mathbf{w}}:\mathcal{X}\to\mathbb{R}\mid \mathbf{w}\in\mathcal{W}\}$, where $\mathcal{W}\subset\mathbb{R}^{d_w}$ is a set of d_w -dimensional parameter vectors. Each function $f_{\mathbf{w}}\in\mathcal{F}$ may be highly non-convex and is not necessarily differentiable with respect to the context vector $\mathbf{x}\in\mathcal{X}$, but we assume it is differentiable with respect to the parameter \mathbf{w} . We further assume that the function class \mathcal{F} is expressive enough to contain the true (unknown) utility function $f_{\mathbf{w}^*}$.

Assumption 1 (Realizability). We assume $\mathbf{w}^* \in \mathcal{W}$ and \mathcal{F} is given to the agent.

Assumption 1 is standard in the contextual bandit literature using general function approximation [13, 12, 35, 25, 41]. While this assumption is necessary for our theoretical guarantees, we will also demonstrate empirically that our proposed algorithm performs well even under model misspecification, particularly when the function class under consideration has high representational capacity, such as neural networks.

We also define the equivalence set $\mathcal{W}^* \subset \mathcal{W}$ as the set of parameter vectors that yield the same function values as $f_{\mathbf{w}^*}$ for all $\mathbf{x} \in \mathcal{X}$, i.e., $\mathcal{W}^* := \{\mathbf{w} \in \mathcal{W} : f_{\mathbf{w}}(\mathbf{x}) = f_{\mathbf{w}^*}(\mathbf{x}) \text{ for all } \mathbf{x} \in \mathcal{X} \}$. This equivalence set arises naturally in expressive function classes such as neural networks. For instance, consider a two-layer neural network defined as $f_{\mathbf{w}^*}(\mathbf{x}) = \sum_{k=1}^m \sigma(\mathbf{x}^\top \mathbf{w}_k^*)$, where $\mathbf{w}^* = [\mathbf{w}_1^*, \dots, \mathbf{w}_m^*] \in \mathbb{R}^{d \times m}$ is the weight matrix, and σ denotes the activation function. Any column permutation of \mathbf{w}^* yields the same function output as $f_{\mathbf{w}^*}$. Therefore, rather than recovering \mathbf{w}^* exactly, it suffices to identify any parameter within the equivalence set \mathcal{W}^* .

Following prior work on contextual bandits [1, 28, 9, 20, 25, 22, 34], we make the following boundedness assumption.

Assumption 2 (Boundedness). For all $t \ge 1$ and $i \in [N]$, we assume that $\|\mathbf{x}_{ti}\|_2 \le 1$ and $r_{ti} \in [0,1]$, and $\|\mathbf{w}\|_2 \le 1$ for all $\mathbf{w} \in \mathcal{W}$. Also, we assume that there exist constants $C_f, C_g, C_h > 0$ such that for all $\mathbf{x} \in \mathcal{X}$ and $\mathbf{w} \in \mathcal{W}$, it holds that

$$|f_{\mathbf{w}}(\mathbf{x})| \le C_f$$
, $\|\nabla f_{\mathbf{w}}(\mathbf{x})\|_2 \le C_q$, $\|\nabla^2 f_{\mathbf{w}}(\mathbf{x})\|_{op} \le C_h$,

where $\nabla f_{\mathbf{w}}(\mathbf{x})$ and $\nabla^2 f_{\mathbf{w}}(\mathbf{x})$ denote the gradient and Hessian of $f_{\mathbf{w}}$ with respect to \mathbf{w} .

2.2 Challenges of Learning in Non-Linear Utility

Learning utility parameters in non-linear MNL bandit models presents significant challenges compared to linear models. In linear utility settings, the negative log-likelihood function exhibits a self-concordant-like property [36], simplifying theoretical analysis and enabling tight regret bounds [32, 20]. Non-linear utility models, however, typically lack this property, complicating analysis significantly. Another key difference lies in the gradient structure. In linear models, gradients are simply feature vectors, independent of parameters, allowing straightforward concentration inequalities. In contrast, non-linear gradients depend explicitly on unknown parameters, complicating the application of these analytical methods. Moreover, the optimization landscape for non-linear utilities is notably more complex. Linear models yield convex loss functions with unique minima, whereas non-linear models (e.g., neural networks) produce highly non-convex loss surfaces. Additionally, symmetries like neuron permutations cause multiple global minima, further complicating optimization and theoretical analysis in non-linear settings.

3 Main Results

In this section, we introduce an algorithm for the MNL bandit problem with a non-linear parametric utility function that is both computationally tractable and statistically efficient. We begin by explaining the key design principles of the algorithm, and then present its regret bound along with the main technical components of the analysis.

3.1 Algorithm: ONL-MNL

We propose ONL-MNL, an algorithm that leverages optimistic estimates derived from non-linear utility functions to guide exploration. The complete procedure is presented in Algorithm 1, which consists of two main stages.

Uniform exploration phase. The first stage (Phase I) is a uniform exploration phase that lasts for n rounds. During this phase, the agent selects an assortment of size at most K, uniformly at random. The goal of this exploration phase is to obtain a "pilot" estimate of the unknown parameter \mathbf{w}^* or any parameter in the equivalence set \mathcal{W}^* —that is, a good initial estimator. At the end of Phase I, the agent invokes a regression oracle to estimate the initial parameter $\hat{\mathbf{w}}_0$ based on the data collected

Algorithm 1 ONL-MNL (Optimistic Non-Linear Utility for Contextual MNL Bandit)

```
1: Inputs: Regularization parameter \lambda, confidence radius \{\beta_t\}_{t\geq 1}, exploration length n, number of
       rounds T
 2: Initialization: V_1 = \lambda I_{d_w}
 3: for m = 1, 2, \dots, n do
                                                                                                                                ⊳ Phase I: Uniform exploration
               Observe X_m = \{\mathbf{x}_{m1}, \dots, \mathbf{x}_{mN}\}, offer S_m \sim \text{Unif}(\mathcal{S}) and observe i_m
 6: Compute \hat{\mathbf{w}}_0 = \operatorname{argmin}_{\mathbf{w} \in \mathcal{W}} \mathcal{L}_n(\mathbf{w})
 7: for t = 1, 2, ..., T do
                                                                                                                           ▶ Phase II: Optimistic exploration
              Observe X_t = \{\mathbf{x}_{t1}, \dots, \mathbf{x}_{tN}\}
Compute \hat{\mathbf{w}}_t = \operatorname{argmin}_{\mathbf{w} \in \mathcal{W}} \ell_t(\mathbf{w})
 9:
              Compute z_{ti} = f_{\hat{\mathbf{w}}_t}(\mathbf{x}_{ti}) + \sqrt{\beta_t} \|\nabla f_{\hat{\mathbf{w}}_t}(\mathbf{x}_{ti})\|_{\mathbf{V}_{\bullet}^{-1}} + \frac{\beta_t C_h}{\lambda} for all i \in [N]
10:
              Offer S_t \in \operatorname{argmax}_{S \in \mathcal{S}} \widetilde{R}_t(S) and observe i_t
11:
              Update \mathbf{V}_{t+1} = \mathbf{V}_t + \sum_{i \in S_t} \nabla f_{\hat{\mathbf{w}}_t}(\mathbf{x}_{ti}) \nabla f_{\hat{\mathbf{w}}_t}(\mathbf{x}_{ti})^{\top}
12:
13: end for
```

during the exploration rounds. This is done by minimizing the negative log-likelihood of the observed item choices under the MNL model defined as follows:

$$\hat{\mathbf{w}}_0 = \operatorname*{argmin}_{\mathbf{w} \in \mathcal{W}} \mathcal{L}_n(\mathbf{w}) := -\sum_{m=1}^n \log p(i_m \mid X_m, S_m, \mathbf{w}).$$

Here, $\mathcal{L}_n(\mathbf{w})$ denotes the empirical loss over the first n rounds, where i_m is the item chosen from assortment S_m with the context set X_m . Using standard results from empirical risk minimization over Lipschitz function classes, together with the reverse Lipschitz property of the MNL model, we show that the minimum squared error of the pilot estimator $\hat{\mathbf{w}}_0$ with respect to the equivalence set \mathcal{W}^* , namely $\min_{\tilde{\mathbf{w}} \in \mathcal{W}^*} \|\hat{\mathbf{w}}_0 - \tilde{\mathbf{w}}\|_2^2$, converges at a rate of $\mathcal{O}(1/n)$ (Lemma 1).

Optimistic exploration phase. In the second stage (Phase II), we adopt a more sophisticated strategy to balance exploration and exploitation. In particular, we construct a confidence region for the parameter in the equivalence set \mathcal{W}^* that is closest to the pilot estimator $\hat{\mathbf{w}}_0$. We then estimate the utility parameters within \mathcal{W}^* by minimizing the following regularized negative log-likelihood:

$$\hat{\mathbf{w}}_t = \operatorname*{argmin}_{\mathbf{w} \in \mathcal{W}} \ell_t(\mathbf{w}) := -\sum_{s=1}^{t-1} \sum_{i \in S_s} y_{si} \log \hat{p}(i \mid X_s, S_s, \mathbf{w}) + \frac{\lambda}{2} \|\mathbf{w} - \hat{\mathbf{w}}_0\|_2^2, \tag{1}$$

where $\hat{p}(i \mid X_s, S_s, \mathbf{w})$ is a *linearized* MNL model given by

$$\hat{p}(i \mid X_s, S_s, \mathbf{w}) := \frac{\exp(f_{\hat{\mathbf{w}}_s}(\mathbf{x}_{si}) + \nabla f_{\hat{\mathbf{w}}_s}(\mathbf{x}_{si})^{\top}(\mathbf{w} - \hat{\mathbf{w}}_s))}{1 + \sum_{j \in S_s} \exp(f_{\hat{\mathbf{w}}_s}(\mathbf{x}_{sj}) + \nabla f_{\hat{\mathbf{w}}_s}(\mathbf{x}_{sj})^{\top}(\mathbf{w} - \hat{\mathbf{w}}_s))}.$$

The key intuition is that we apply a first-order Taylor approximation of the non-linear utility function $f_{\mathbf{w}}(\cdot)$ around the estimate $\hat{\mathbf{w}}_s$, since we are dealing with a general non-linear model. Moreover, instead of regularizing around the origin $\mathbf{0}_{d_w}$, we regularize around the pilot estimator $\hat{\mathbf{w}}_0$, as it provides a meaningful reference point that is already close to the parameters in \mathcal{W}^* . Based on the estimated parameter $\hat{\mathbf{w}}_t$, we define the confidence set \mathcal{C}_t as follows:

$$C_t := \left\{ \mathbf{w} \in \mathcal{W} : \|\mathbf{w} - \hat{\mathbf{w}}_t\|_{\mathbf{V}_t}^2 \le \beta_t \right\},\,$$

where β_t is a pre-defined monotonically increasing sequence specified later, and $\mathbf{V}_t := \lambda \mathbf{I}_{d_w} + \sum_{s=1}^{t-1} \sum_{i \in S_s} \nabla f_{\hat{\mathbf{w}}_s}(\mathbf{x}_{si}) \nabla f_{\hat{\mathbf{w}}_s}(\mathbf{x}_{si})^{\top}$ is the Gram matrix. With suitably set β_t , we will show that with high probability, $\underset{\mathbf{w} \in \mathcal{W}^*}{\operatorname{argmin}}_{\tilde{\mathbf{w}} \in \mathcal{W}^*} \|\hat{\mathbf{w}}_0 - \tilde{\mathbf{w}}\|_2 \in \mathcal{C}_t$, i.e., the equivalent utility parameter in \mathcal{W}^* that is closest to the pilot estimator $\hat{\mathbf{w}}_0$ lies within the constructed confidence set \mathcal{C}_t . Based on this observation, we compute the optimistic utility estimate z_{ti} as follows:

$$z_{ti} = f_{\hat{\mathbf{w}}_t}(\mathbf{x}_{ti}) + \sqrt{\beta_t} \|\nabla f_{\hat{\mathbf{w}}_t}(\mathbf{x}_{ti})\|_{\mathbf{V}_t^{-1}} + \frac{\beta_t C_h}{\lambda}, \quad \forall i \in [N].$$

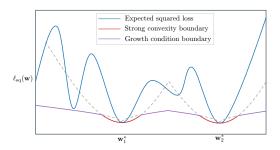


Figure 1: Example of a highly non-convex loss $\ell_{\mathrm{sq}}(\mathbf{w})$ that satisfies Assumption 4. The solid line shows the lower bound given by the pointwise minimum of the strong convexity and growth conditions over the equivalence set \mathcal{W}^* . For neural networks, the squared loss may have multiple global minima due to parameter symmetries or over-parameterization. This highlights that Assumption 4 is strictly weaker than prior geometric conditions, enabling it to cover a wider range of utility functions.

The optimistic utility z_{ti} is composed of the parts: the mean utility estimate $f_{\hat{\mathbf{w}}_t}(\mathbf{x}_{ti})$ and the uncertainty estimate $\sqrt{\beta_t} \|\nabla f_{\hat{\mathbf{w}}_t}(\mathbf{x}_{ti})\|_{\mathbf{V}_t^{-1}} + \frac{\beta_t C_h}{\lambda}$. We will show that z_{ti} is an optimistic estimate of the true utility $f_{\mathbf{w}^*}(\mathbf{x}_{ti})$, assuming $\mathop{\rm argmin}_{\widetilde{\mathbf{w}} \in \mathcal{W}^*} \|\hat{\mathbf{w}}_0 - \widetilde{\mathbf{w}}\|_2 \in \mathcal{C}_t$. Based on z_{ti} , we construct the optimistic expected reward for the assortment, defined as $\widetilde{R}_t(S) := \sum_{i \in S} \frac{\exp(z_{ti})r_{ti}}{1 + \sum_{j \in S} \exp(z_{tj})}$. Then, we offer the set S_t that maximizes the optimistic expected reward, i.e., $S_t = \mathop{\rm argmax} \widetilde{R}_t(S)$.

3.2 Regret Bound

In this section, we present the cumulative regret upper bound for the proposed algorithm. First, we introduce technical assumptions used to derive the regret bound.

Assumption 3 (Stochastic context in Phase I). We assume that during Phase I, the context feature vectors $X_m := \{\mathbf{x}_{m1}, \dots, \mathbf{x}_{mN}\}$ are independently and identically distributed (i.i.d.) samples drawn from an unknown distribution \mathcal{D} supported on \mathcal{X} .

Assumption 4 (Generalized geometric condition on the squared loss). *Let*

$$\ell_{\mathrm{sq}}(\mathbf{w}) = \mathbb{E}_{X \sim \mathcal{D}, S \sim \mathrm{Unif}(\mathcal{S})} \left[\sum_{i \in S} (f_{\mathbf{w}}(\mathbf{x}_j) - f_{\mathbf{w}^*}(\mathbf{x}_j))^2 \right]$$

be the expected squared loss function over both the context feature distribution \mathcal{D} and a uniformly sampled assortment. For the equivalence set \mathcal{W}^* , we assume that $\ell_{sq}(\mathbf{w})$ satisfies either a (τ, γ) -growth condition or μ -local strong convexity with respect to \mathcal{W}^* , i.e., $\forall \mathbf{w} \in \mathcal{W} \setminus \mathcal{W}^*$,

$$\min_{\widetilde{\mathbf{w}} \in \mathcal{W}^*} \left\{ \frac{\mu}{2} \|\mathbf{w} - \widetilde{\mathbf{w}}\|_2^2 + \ell_{\mathrm{sq}}(\widetilde{\mathbf{w}}), \frac{\tau}{2} \|\mathbf{w} - \widetilde{\mathbf{w}}\|_2^{\gamma} + \ell_{\mathrm{sq}}(\widetilde{\mathbf{w}}) \right\} \leq \ell_{\mathrm{sq}}(\mathbf{w}),$$

for constants $\mu > 0$ and $0 < \gamma < 2$.

Discussion of assumptions. Assumption 3 is used to bound the difference between the expected negative log-likelihoods of the true utility parameter \mathbf{w}^* and the pilot estimator $\hat{\mathbf{w}}_0$, a type of assumption commonly used in the empirical risk minimization literature [37, 41]. Note that Assumption 3 is only required for Phase I; after Phase I, the context vectors may be chosen even adversarially.

Assumption 4 is on the expected squared loss between utility functions over the context distribution, rather than the unknown true utility function $f_{\mathbf{w}^*}$ itself. Assumption 4 is also required in Phase I to establish the convergence rate of the pilot estimator $\hat{\mathbf{w}}_0$. The growth condition assumes that any utility function $f_{\mathbf{w}}$, whose parameters lie far from the equivalence set \mathcal{W}^* , cannot approximate the true utility function well over the context distribution. The local strong convexity condition assumes that the expected loss exhibits quadratic growth within a local neighborhood around the equivalence set \mathcal{W}^* . Notably, this condition is strictly weaker than the global strong convexity assumptions commonly adopted in prior work on parametric bandits—such as linear and generalized linear bandits [11, 1]—since it does not require the squared loss to be convex outside the local region near the optimal parameters. Moreover, Assumption 4 is more general than the geometric conditions employed in prior works [25, 22, 34]. That is, any function satisfying the geometric conditions in [25, 22, 34] also satisfies ours, but the converse does not necessarily hold. Previous geometric assumptions typically require the existence of a unique global minimum, whereas we explicitly allow for multiple equivalent minima. As previously discussed, this is a more realistic assumption in settings such as neural networks, where the squared error loss can admit multiple

global minima—e.g., due to parameter symmetries or overparameterization. This flexibility allows our framework to accommodate highly non-convex and non-linear expected loss landscapes, which may contain arbitrarily many spurious local minima and multiple global minima (Figure 1).

Now we are ready to state our main result that upper bounds the cumulative regret of the Algorithm 1.

Theorem 1 (Regret Bound of ONL-MNL). Suppose Assumption 1, 2, 3 and 4 hold. For any $\delta \in (0,1)$, if we set the algorithmic parameters in Algorithm 1 as follows: $T \geq \widetilde{C} \kappa^{-1} C_f^2 \zeta^2 \left(\frac{\mu^{\gamma/(2-\gamma)}}{\tau^{2/(2-\gamma)}}\right)^2$, $n = \kappa^{-3/2} d_w \sqrt{T}$, $\lambda = \widetilde{\mathcal{O}}(\kappa^{-5/2} \mu^{-1} d_w \sqrt{T})$, $\beta_t = \widetilde{\mathcal{O}}\left(\mu^{-2} \kappa^{-4} d_w\right)$, where \widetilde{C} is a universal constant, $\kappa := \min_{\mathbf{w} \in \mathcal{W}, X \subset \mathcal{X}, S \in \mathcal{S}, i \in S} p(0 \mid X, S, \mathbf{w}) p(i \mid X, S, \mathbf{w})$, and, ζ is the logarithmic term depending on $K, n, C_g, 1/\delta$, then with probability at least $1-2\delta$, Algorithm 1 achieves the following cumulative regret:

$$\operatorname{Regret}_{n+T} = \widetilde{\mathcal{O}}\left(\kappa^{-3/2} d_w \sqrt{T} + \kappa^{-2} \mu^{-1} d_w \sqrt{T}\right).$$

Discussion of Theorem 1. In terms of the total time step T, Algorithm 1 achieves $\widetilde{\mathcal{O}}(\sqrt{T})$ cumulative regret. This implies that Algorithm 1 is a no-regret algorithm, as $\lim_{T\to\infty} \mathrm{Regret}_{n+T}/T = 0$.

When the utility function family consists of neural networks, our proposed method still achieves regret bound of $\widetilde{\mathcal{O}}(\sqrt{T})$ without relying on the NTK assumption, which is commonly used in the neural network-based bandit analysis [45, 43, 17], and thus eliminates the impractical requirement of extensive over-parameterization prevalent in NTK-based analyses. In particular, existing neural bandits [38, 15, 39] that do not rely on NTK assumptions still require restrictive conditions—such as specific network architectures (e.g., two-layered ReLU networks with fixed second-layer weights set to one [39], or networks with quadratic activation functions [38]) or specific context distributions (e.g., Gaussian [15], or uniform distribution over the unit sphere [39]). In contrast, our approach applies to a broader class of neural networks that satisfy Assumption 4, without imposing such architectural or distributional constraints.

Compared to the existing contextual MNL bandit algorithms with general utility functions [41], our algorithm is both computationally tractable and statistically efficient. While the uniform exploration method in Zhang and Luo [41] is computationally tractable, it results in sub-optimal regret—for example, $\widetilde{\mathcal{O}}(T^{2/3})$ in the stochastic setting and $\widetilde{\widetilde{\mathcal{O}}}(T^{5/6})$ in the adversarial setting. Their other methods—log-barrier regularization-based exploration and the Feel-Good Thompson Sampling [42] approach—are computationally intractable, meaning they cannot be solved in polynomial time. In contrast, our algorithm operates within polynomial time. The main computational cost lies in minimizing the regularized negative log-likelihood, which can be efficiently approximated using gradient-based methods such as gradient descent with a per-round complexity of $\mathcal{O}(t)$ [10]. The assortment selection step is also tractable: by combining the optimistic utility estimates z_{ti} with efficient assortment optimization techniques [33, 8], as done in Oh and Iyengar [28, 29], it can be solved with $\mathcal{O}(N)$ computational cost. In terms of regret bounds, although our regret depends on a problem-dependent instance factor κ , it does not scale directly with the total number of items N or directly with the assortment size K. In contrast, the computationally intractable algorithm with $\mathcal{O}(\sqrt{T})$ -regret by Zhang and Luo [41] has super-linear dependence on K—specifically $\mathcal{O}(K^2)$ —, and additionally depend polynomially on N, limiting their applicability in settings with a large set of items, hence restricting real-world applications. To the best of our knowledge, our method is the only known algorithm that is both computationally tractable and provably achieves a $\widetilde{\mathcal{O}}(\sqrt{T})$ regret in the contextual MNL bandit with non-linear utility functions.

Remark 1. By definition, κ^{-1} can scale as $\mathcal{O}(K^2)$ in the worst-case scenario. However, in practice, the behavior of κ is often much more favorable. For example, as shown in the experiments of Lee and Oh [20] (Figure 1 in [20]), the gap between the regret of κ^{-1} -dependent algorithms [28, 29] and that of κ^{-1} -improved algorithm [20] does not grow significantly with K, even as K increases. This observation suggests that equating κ^{-1} with K^2 or treating it as equivalent to K is overly pessimistic. Also, we note that the exponential dependence in κ can be adjusted by rescaling the utility values. For instance, Zhang and Luo [41] define the MNL probability model using bounded utilities without applying the exponential transformation.

3.3 Regret Analysis

In this section, we present key technical ingredients that enable the regret bound of Algorithm 1, as established in Theorem 1. All the detailed proofs are provided in the appendix.

Convergence of $\hat{\mathbf{w}}_0$. The purpose of Phase I in Algorithm 1 is to explore the unknown utility function $f_{\mathbf{w}^*}$ through n rounds of uniform sampling and to obtain a pilot estimate $\hat{\mathbf{w}}_0$ that is sufficiently accurate. We establish a convergence guarantee for $\hat{\mathbf{w}}_0$ relative to the equivalence set \mathcal{W}^* ; specifically, we bound the quantity $\min_{\widetilde{\mathbf{w}} \in \mathcal{W}^*} \|\hat{\mathbf{w}}_0 - \widetilde{\mathbf{w}}\|_2^2$.

Lemma 1 (Convergence rate of $\hat{\mathbf{w}}_0$). Suppose Assumption 1, 2, 3 and 4 hold. There exist an absolute constant C>0 such that after n rounds in Phase I of Algorithm 1 where n satisfies $n \geq 2^{-1}C\kappa^{-2}d_wC_f^2\zeta\frac{\mu^{\gamma/(2-\gamma)}}{\tau^{2/(2-\gamma)}}$, with probability at least $1-\delta$, we have

$$\min_{\widetilde{\mathbf{w}} \in \mathcal{W}^*} \|\hat{\mathbf{w}}_0 - \widetilde{\mathbf{w}}\|_2^2 \le C \frac{\kappa^{-2} d_w C_f^2 \zeta}{\mu n},$$

where $\kappa := \min_{\mathbf{w} \in \mathcal{W}, X \subset \mathcal{X}, S \in \mathcal{S}, i \in S} p(0 \mid X, S, \mathbf{w}) p(i \mid X, S, \mathbf{w})$ and ζ is the logarithmic term depending on $K, n, C_g, 1/\delta$.

Compared to the contextual multi-armed bandit setting with general reward function [12, 35], one of the key challenges in estimating the utility parameter in the MNL bandit is that the agent does not receive direct feedback on the utility of each item. Instead, it only observes the user's choice from the offered assortment. Despite this challenge, in Lemma 1, we first adapt existing generalization error bounds of the empirical risk minimizer under *i.i.d.* sampling, expressed in terms of the logloss [37, 41]. Then, leveraging the reverse Lipschitz property of the MNL model (Lemma 4), we lower bound the difference in log-likelihoods by the squared error between the utility values. Unlike Zhang and Luo [41], which derives a lower bound based on the assortment size K, our analysis incorporates a problem-dependent instance factor κ .

Confidence set & Optimism. Although the 1/n convergence rate of $\hat{\mathbf{w}}_0$ is quite tight, it is not sufficient on its own. For example, Xu et al. [38] show that in the contextual bandit setting, one can achieve a regret of $\widetilde{\mathcal{O}}(T^{2/3})$ by first learning the parameters of a quadratic neural network through uniform exploration (yielding a 1/n convergence rate), and then switching entirely to exploitation using the learned parameters. In contrast, we go beyond this by constructing a confidence set for parameters in equivalence set \mathcal{W}^* . We then use an optimistic estimate of the true utility within this set to guide assortment selection, ultimately achieving a regret of $\widetilde{\mathcal{O}}(\sqrt{T})$. The following lemma explains how this confidence set is constructed.

Lemma 2 (Confidence set). Suppose Assumption 1, 2, 3 and 4 hold. If we choose $n = d_w \kappa^{-3/2} \sqrt{T}$ and $\lambda = \widetilde{\mathcal{O}}(\kappa^{-3/2} \mu^{-1} d_w \sqrt{T})$, then $\forall t \geq 1$ in Phase II of Algorithm 1, with probability at least $1 - 2\delta$, the following holds:

$$\min_{\widetilde{\mathbf{w}} \in \mathcal{W}^*} \|\hat{\mathbf{w}}_t - \widetilde{\mathbf{w}}\|_{\mathbf{V}_t}^2 \le \beta_t := \widetilde{\mathcal{O}} \left(\mu^{-2} \kappa^{-4} d_w \right).$$

The construction of the confidence set in Lemma 2 relies on two key components. The first is the careful choice of the exploration phase length n and the regularization parameter λ , both of which play a critical role in ensuring the confidence radius β_t scales as $\mathcal{O}(\log T)$, which is essential for achieving a cumulative regret of $\widetilde{\mathcal{O}}(\sqrt{T})$.

The second challenge is that for $\lambda = \mathcal{O}(\sqrt{T})$, existing concentration results for utility parameter estimation in MNL or GLM bandits [9, 32] are no longer applicable. This is because these results rely on Bernstein-type inequalities, which yield a confidence radius of order $\mathcal{O}\left(\sqrt{\lambda} + \frac{d\log T}{\sqrt{\lambda}}\right)$ where d is the dimension of the context feature vector. To ensure the confidence radius is $\mathcal{O}(\log T)$, these bounds require $\lambda = \mathcal{O}(d\log T)$, which limits their applicability when $\lambda = \mathcal{O}(\sqrt{T})$. To overcome this, we first establish a new concentration inequality for the utility parameter in the MNL choice model that is independent of λ (up to logarithmic factors), presented in Lemma 5. Combined with carefully chosen values of λ and n, we then use a mathematical induction argument to show that for all rounds t in Phase II, the estimation error satisfies $\min_{\widetilde{\mathbf{w}} \in \mathcal{W}^*} \|\hat{\mathbf{w}}_t - \widetilde{\mathbf{w}}\|_2^2 = \widetilde{\mathcal{O}}(1/n)$. Finally, we prove that this implies $\min_{\widetilde{\mathbf{w}} \in \mathcal{W}^*} \|\hat{\mathbf{w}}_t - \widetilde{\mathbf{w}}\|_{\mathbf{V}_t}^2 \leq \beta_t$.

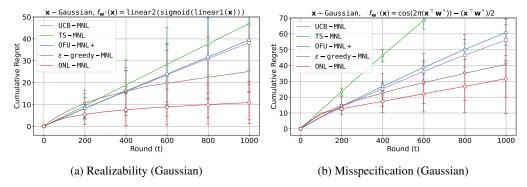


Figure 2: Cumulative regret comparison between ONL-MNL (ours) and baselines under Gaussian contexts. The results for the uniform context distribution are provided in Figure 4.

Next, based on the confidence radius β_t established in Lemma 2, we show that the optimistic utility estimate z_{ti} is greater than or equal to the true utility $f_{\mathbf{w}^*}(\mathbf{x}_{ti})$ as follows:

Lemma 3 (Optimistic utility). Suppose Assumption 1, 2, 3 and 4 hold and $\min_{\tilde{\mathbf{w}} \in \mathcal{W}^*} \|\hat{\mathbf{w}}_t - \tilde{\mathbf{w}}\|_{\mathbf{V}t}^2 \le \beta_t$ for all $t \in [T]$. Then for all $t \in [T]$ and $i \in S_t$, we have

$$0 \le z_{ti} - f_{\mathbf{w}^*}(\mathbf{x}_{ti}) \le 2\sqrt{\beta_t} \|\nabla f_{\hat{\mathbf{w}}_t}(\mathbf{x}_{ti})\|_{\mathbf{V}_t^{-1}} + \frac{2\beta_t C_h}{\lambda}.$$

Unlike the linear utility case [29], our optimistic utility estimate for $f_{\mathbf{w}^*}$ introduces an item-independent extra term of the form $2\beta_t C_h/\lambda$. This term arises as a residual from the first-order approximation of the unknown utility function $f_{\mathbf{w}^*}$, and unlike standard deviation terms $\|\nabla f_{\hat{\mathbf{w}}_t}(\mathbf{x}_{ti})\|_{\mathbf{V}_t^{-1}}$, it does not vanish over time. As a result, it poses an additional challenge in achieving a cumulative regret of $\widetilde{\mathcal{O}}(\sqrt{T})$. Nevertheless, we show that with a regularization parameter $\lambda = \widetilde{\mathcal{O}}(\sqrt{T})$, it is still possible to achieve both optimism and cumulative regret of $\mathcal{O}(\sqrt{T})$. We then show that the expected reward of the assortment S_t , chosen based on the optimistic utility estimates z_{ti} , exceeds the true expected reward of the optimal assortment S_t^* with high probability.

4 Numerical Experiments

MNL Choice model with non-linear utilities. In this section, we present numerical experiments to evaluate the performance of the proposed algorithm ONL-MNL. As baselines, we compare with algorithms for contextual MNL bandits with linear utilities—UCB-MNL [29], TS-MNL [28], and OFU-MNL+ [20]—as well as ε -greedy-MNL [41], a regression-oracle-based method with uniform exploration designed for general utilities. Each experiment was conducted independently over 30 random seeds, and we report the mean and standard deviation of the cumulative regret. We set N=100, K=5, d=3, and $r_{ti}=1$. Experiments were performed under two different context distributions: $\mathbf{x} \sim \mathcal{N}(\mathbf{0}_d, \mathbf{I}_d)$ and $x_i \sim \mathrm{Unif}[-3,3], \forall i \in [d]$. We use a two-layered neural network with sigmoid activation as the estimator \hat{f} for our proposed algorithm, i.e., $\hat{f}(\mathbf{x})=1\mathrm{inear2}(\mathrm{sigmoid}(1\mathrm{inear1}(\mathbf{x})))$.

For the realizable setting, the true utility function $f_{\mathbf{w}^*}$ is defined as a two-layer neural network with the same architecture as \hat{f} , and we set the number of neurons in the hidden layer to 3. Each entry of the true utility parameter \mathbf{w}^* is independently sampled from the uniform distribution over the interval [-1,1]. We set the length of Phase I to n=50 and ran experiments for T=1000.

For the misspecified setting, the true utility function is defined as $f_{\mathbf{w}^*}(\mathbf{x}) = \cos(2\pi (\mathbf{x}^\top \mathbf{w}^*)) - \frac{1}{2}(\mathbf{x}^\top \mathbf{w}^*)$, which differs in form from the estimator architecture. To account for this increased complexity, we set the number of neurons in the hidden layer of the estimator to 15. Accordingly, we also increase the length of Phase I to n = 100.

The experimental results in Figure 2 clearly demonstrate the effectiveness of our proposed algorithm, ONL-MNL, in both realizable and misspecified settings. Existing linear utility-based methods exhibit steadily increasing regret under non-linear utilities, as they fail to capture the complex structure of the

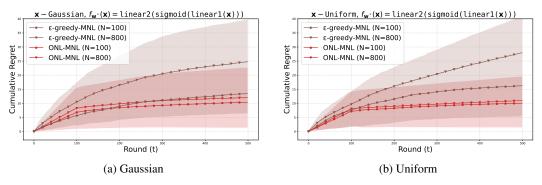


Figure 3: Cumulative regret comparison between ONL-MNL (ours) and ε -greedy-MNL [41] under varying number of items N.

true utility function. Compared to the ε -greedy-MNL algorithm designed for general utilities [41], our approach is not only statistically efficient—achieving low cumulative regret—but also exhibits robust and consistently strong performance across different scenarios in both realizable and misspecified settings. See Appendix G for additional experimental results and details.

Effect of the number of item on regret performance. To further highlight the distinct advantages of our proposed algorithm, ONL-MNL, over the ε -greedy-MNL [41], we compare the performances of those algorithms as the total number of items N increases. Beyond being computationally tractable, a key advantage of our algorithm is that its regret bound is independent of N, whereas the regret bound for ε -greedy-MNL grows polynomially with N. To verify this, we conduct experiments under a realizable setting where the true utility function is a two-layer neural network with sigmoid activation and 10 hidden neurons. We evaluate the cumulative regret of both algorithms at T=500 for different values of $N \in \{100,800\}$. Following [41], where the optimal exploration rate satisfies $\varepsilon = \mathcal{O}(N^{1/3})$ (Theorem 3.4 in [41]), we set $\varepsilon=0.1$ and 0.2 for N=100 and 800, respectively.

Figure 3-(a) and (b) show the results under Gaussian and uniform context distributions, respectively. As expected, the cumulative regret of ε -greedy-MNL increases as N grows, which aligns with its theoretical regret bound that scales polynomially with N. While our algorithm also exhibits a slight increase in cumulative regret as N grows, this is primarily due to the increased regret incurred during the uniform exploration phase (Phase I). As the total number of items increases, the size of the assortment set $\mathcal S$ also grows, leading to a larger gap between the expected reward of a uniformly sampled assortment and that of the optimal assortment. Nevertheless, in Phase II, our algorithm consistently converges toward the optimal policy, regardless of the value of N. In other words, the slope of the regret curve during Phase II remains relatively stable across different values of N, indicating that the learning efficiency of our method is unaffected by the size of the item set. These results demonstrate that our method not only offers provable computational and statistical efficiency but also remains robust and effective even in settings with extremely large item pools.

5 Conclusion

We studied contextual MNL bandits with non-linear parametric utilities and proposed a UCB-based algorithm that is both computationally tractable and statistically efficient. Our algorithm attains $\widetilde{\mathcal{O}}(\sqrt{T})$ regret independent of the number of items, to the best of our knowledge the first tractable method with this guarantee for non-linear utilities. Relative to prior approaches for general utility models, our result achieves the same $\widetilde{\mathcal{O}}(\sqrt{T})$ rate while remaining implementable under non-linearity.

Looking ahead, the current regret bound depends on the problem-dependent instance factor κ ; reducing or removing this dependence in the non-linear setting likely requires new analytical techniques and remains an important direction. Another promising avenue is to extend our framework to MNL MDPs [16], where the transition kernel is specified by an MNL model. Prior work analyzes only the case in which the logit is a linear function of state—action features [16, 6, 24, 31]; generalizing to *non-linear* parameterizations while preserving computational tractability and performance guarantees would broaden the scope and impact of this line of research.

Acknowledgements

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2022-NR071853 and RS-2023-00222663), by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. RS-2025-02263754), and by AI-Bio Research Grant and Ascending SNU Future Leader Fellowship through Seoul National University.

References

- [1] Y. Abbasi-Yadkori, D. Pál, and C. Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 24:2312–2320, 2011.
- [2] P. Agrawal, T. Tulabandhula, and V. Avadhanula. A tractable online learning algorithm for the multinomial logit contextual bandit. *European Journal of Operational Research*, 2023.
- [3] S. Agrawal, V. Avadhanula, V. Goyal, and A. Zeevi. Mnl-bandit: A dynamic learning approach to assortment selection. *Operations Research*, 67(5):1453–1485, 2019.
- [4] X. Chen, Y. Wang, and Y. Zhou. Dynamic assortment optimization with changing contextual information. *Journal of machine learning research*, 2020.
- [5] W. C. Cheung and D. Simchi-Levi. Thompson sampling for online personalized assortment optimization problems with multinomial logit choice models. *Available at SSRN 3075658*, 2017.
- [6] W. Cho, T. Hwang, J. Lee, and M.-h. Oh. Randomized exploration for reinforcement learning with multinomial logistic function approximation. *Advances in Neural Information Processing Systems*, 37:76643–76720, 2024.
- [7] V. Dani, T. P. Hayes, and S. M. Kakade. Stochastic linear optimization under bandit feedback. pages 355–366, 2008. URL http://colt2008.cs.helsinki.fi/papers/80-Dani.pdf.
- [8] J. M. Davis, G. Gallego, and H. Topaloglu. Assortment optimization under variants of the nested logit model. *Operations Research*, 62(2):250–273, 2014.
- [9] L. Faury, M. Abeille, C. Calauzènes, and O. Fercoq. Improved optimistic algorithms for logistic bandits. In *International Conference on Machine Learning*, pages 3052–3060. PMLR, 2020.
- [10] L. Faury, M. Abeille, K.-S. Jun, and C. Calauzènes. Jointly efficient and optimal algorithms for logistic bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 546–580. PMLR, 2022.
- [11] S. Filippi, O. Cappé, A. Garivier, and C. Szepesvári. Parametric bandits: The generalized linear case. In *Proceedings of the 23rd International Conference on Neural Information Processing Systems Volume 1*, NIPS'10, page 586–594, Red Hook, NY, USA, 2010. Curran Associates Inc.
- [12] D. Foster and A. Rakhlin. Beyond ucb: Optimal and efficient contextual bandits with regression oracles. In *International conference on machine learning*, pages 3199–3210. PMLR, 2020.
- [13] D. Foster, A. Agarwal, M. Dudik, H. Luo, and R. Schapire. Practical contextual bandits with regression oracles. In *International Conference on Machine Learning*, pages 1539–1548. PMLR, 2018.
- [14] D. A. Freedman. On tail probabilities for martingales. the Annals of Probability, pages 100–118, 1975.
- [15] B. Huang, K. Huang, S. Kakade, J. D. Lee, Q. Lei, R. Wang, and J. Yang. Going beyond linear rl: Sample efficient neural function approximation. *Advances in Neural Information Processing Systems*, 34:8968–8983, 2021.

- [16] T. Hwang and M.-h. Oh. Model-based reinforcement learning with multinomial logistic function approximation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pages 7971–7979, 2023.
- [17] T. Hwang, K. Chai, and M.-h. Oh. Combinatorial neural bandits. In *Proceedings of the 40th International Conference on Machine Learning*. PMLR, 2023.
- [18] A. Jacot, F. Gabriel, and C. Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in Neural Information Processing Systems*, volume 31, pages 8571–8580, 2018.
- [19] P. Kassraie and A. Krause. Neural contextual bandits without regret. In *International Conference on Artificial Intelligence and Statistics*, pages 240–278. PMLR, 2022.
- [20] J. Lee and M.-h. Oh. Nearly minimax optimal regret for multinomial logistic bandit. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- [21] J. Lee and M.-h. Oh. Improved online confidence bounds for multinomial logistic bandits. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267, pages 33576–33615. PMLR, 13–19 Jul 2025.
- [22] C. Li, C. Liu, and Y.-X. Wang. Communication-efficient federated non-linear bandit optimization. In *The Twelfth International Conference on Learning Representations*, 2024.
- [23] L. Li, Y. Lu, and D. Zhou. Provably optimal algorithms for generalized linear contextual bandits. In *International Conference on Machine Learning*, pages 2071–2080. PMLR, 2017.
- [24] L.-F. Li, Y.-J. Zhang, P. Zhao, and Z.-H. Zhou. Provably efficient reinforcement learning with multinomial logit function approximation. *Advances in Neural Information Processing Systems*, 37:58539–58573, 2024.
- [25] C. Liu and Y.-X. Wang. Global optimization with parametric function approximation. In *International Conference on Machine Learning*, pages 22113–22136. PMLR, 2023.
- [26] A. Maas, R. E. Daly, P. T. Pham, D. Huang, A. Y. Ng, and C. Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies*, pages 142–150, 2011.
- [27] D. McFadden et al. Modelling the choice of residential location. 1978.
- [28] M.-h. Oh and G. Iyengar. Thompson sampling for multinomial logit contextual bandits. *Advances in Neural Information Processing Systems*, 32:3151–3161, 2019.
- [29] M.-h. Oh and G. Iyengar. Multinomial logit contextual bandits: Provable optimality and practicality. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 9205–9213, 2021.
- [30] M. Ou, N. Li, S. Zhu, and R. Jin. Multinomial logit bandit with linear utility functions. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 2602–2608, 2018.
- [31] J. Park, J. Kwon, and D. Lee. Infinite-horizon reinforcement learning with multinomial logit function approximation. In *International Conference on Artificial Intelligence and Statistics*, pages 361–369. PMLR, 2025.
- [32] N. Perivier and V. Goyal. Dynamic pricing and assortment under a contextual mnl demand. *Advances in Neural Information Processing Systems*, 35:3461–3474, 2022.
- [33] P. Rusmevichientong, Z.-J. M. Shen, and D. B. Shmoys. Dynamic assortment optimization with a multinomial logit choice model and capacity constraint. *Operations research*, 58(6): 1666–1680, 2010.
- [34] Z. S. Siam, C. Guan, and C. Liu. Quantum non-linear bandit optimization. arXiv preprint arXiv:2503.03023, 2025.

- [35] D. Simchi-Levi and Y. Xu. Bypassing the monster: A faster and simpler optimal algorithm for contextual bandits under realizability. *Mathematics of Operations Research*, 47(3):1904–1931, 2022.
- [36] Q. Tran-Dinh, Y.-H. Li, and V. Cevher. Composite convex minimization involving self-concordant-like cost functions. *arXiv preprint arXiv:1502.01068*, pages 155–168, 2015.
- [37] T. Van Erven, P. D. Grünwald, N. A. Mehta, M. D. Reid, and R. C. Williamson. Fast rates in statistical and online learning. *The Journal of Machine Learning Research*, 16(1):1793–1861, 2015.
- [38] K. Xu, H. Bastani, and O. Bastani. Robust generalization of quadratic neural networks via function identification. *arXiv preprint arXiv:2109.10935*, 2021.
- [39] K. Xu, H. Bastani, S. Goel, and O. Bastani. Stochastic bandits with relu neural networks. In *Forty-first International Conference on Machine Learning*, 2024.
- [40] P. Xu, Z. Wen, H. Zhao, and Q. Gu. Neural contextual bandits with deep representation and shallow exploration. In *International Conference on Learning Representations*, 2021.
- [41] M. Zhang and H. Luo. Contextual multinomial logit bandits with general value functions. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- [42] T. Zhang. Feel-good thompson sampling for contextual bandits and reinforcement learning. *SIAM Journal on Mathematics of Data Science*, 4(2):834–857, 2022.
- [43] W. Zhang, D. Zhou, L. Li, and Q. Gu. Neural thompson sampling. In *International Conference on Learning Representation (ICLR)*, 2021.
- [44] Y.-J. Zhang and M. Sugiyama. Online (multinomial) logistic bandit: Improved regret and constant computation cost. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [45] D. Zhou, L. Li, and Q. Gu. Neural contextual bandits with ucb-based exploration. In *International Conference on Machine Learning*, pages 11492–11502. PMLR, 2020.

NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [NA].
- [NA] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

The checklist answers are an integral part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a proper justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "[No]" or "[NA]" is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- Delete this instruction block, but keep the section heading "NeurIPS Paper Checklist",
- Keep the checklist subsection headings, questions/answers and guidelines below.
- Do not modify the questions and only use the provided macros for your answers.

1. Claims

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

Answer: [Yes]

Justification: The main claims made in the abstract is propose a new computationally tractable MNL bandit algorithm based on the upper confidence bound principle, which is designed to handle general non-linear parametric utility functions, including neural networks. Also, in **Introduction**, we summarize the main contributions of this paper.

Guidelines:

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

2. Limitations

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

Answer: [Yes]

Y

Justification: We discuss the limitation along with the future direction of this work in **Conclusion**.

Guidelines:

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

3. Theory assumptions and proofs

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

Answer: [Yes]

Justification: We provide the full set of assumptions in Section 2 & 3. Also we provide all proof of main results in the appendix.

Guidelines:

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

4. Experimental result reproducibility

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

Answer: [Yes]

Justification: We provide numerical experiments in Section 4 and the detailed information of experiments in the appendix.

Guidelines:

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

5. Open access to data and code

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

Answer: [Yes]

Justification: We have attached the data and code with sufficient instructions to reproduce the main experimental results in the supplementary material.

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how
 to access the raw data, preprocessed data, intermediate data, and generated data, etc.

- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

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

Answer: [Yes]

Justification: We provide the detailed explanations for the experimental setting in the appendix.

Guidelines:

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

7. Experiment statistical significance

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

Answer: [Yes]

Justification: We report error bars in our experiments shown in Section 4.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error
 of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

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

Answer: [Yes]

Justification: We provide sufficient information on the computer resources needed to reproduce the experiments in the appendix.

Guidelines:

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

9. Code of ethics

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

Answer: [Yes]

Justification: The research conducted in this paper adheres to the NeurIPS Code of Ethics in all aspects.

Guidelines:

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

10. Broader impacts

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

Answer: [NA]

Justification: here is no negative societal impacts of the work performed because this research focuses on theoretical aspects.

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

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

Answer: [NA]

Justification: The research conducted in this paper does not pose any such risks.

Guidelines:

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

12. Licenses for existing assets

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

Answer: [NA]

Justification: This paper does not use any external assets such as code, data, or models.

Guidelines:

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

13. New assets

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

Answer: [NA]

Justification: This paper does not release new assets.

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.

 At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

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

Answer: [NA]

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

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

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

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

Answer: [NA]

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

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

16. Declaration of LLM usage

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

Answer: [NA]

Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

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

Contents of Appendix

A	Related Works				
В	Proof of Lemma 1				
	B.1 Proofs of Technical Lemmas	24			
	B.1.1 Proof of Lemma 4	24			
C	Proof of Lemma 2	24			
	C.1 Proofs of Technical Lemmas	29			
	C.1.1 Proof of Lemma 5	29			
	C.1.2 Proof of Lemma 6	31			
	C.1.3 Proof of Lemma 7	32			
	C.1.4 Proof of Lemma 8	33			
D	Proof of Lemma 3	33			
E	Proof of Theorem 1				
F	Auxiliary Lemmas				
G	G Details on Experiments & Additional Results				

A Related Works

Contextual MNL bandits. For contextual MNL bandits under linear utility assumptions, Ou et al. [30] consider a setting in which the item-context features remain fixed over time and achieve $\widetilde{\mathcal{O}}(dK\sqrt{T})$ regret bound. Chen et al. [4] extended the fixed context setting to stochastic and time-varying contexts, deriving a regret bound of $\widetilde{\mathcal{O}}(d\sqrt{T}+d^2K^2)$; however, their method is computationally intractable due to the need to enumerate all possible assortments. Oh and Iyengar [29] introduced a polynomial-time algorithm that maintains confidence bounds in the parameter space and computes UCB scores for each item, achieving regret $\widetilde{\mathcal{O}}(dK\sqrt{T/\kappa})$, where κ is the problem-dependent instance factor. Perivier and Goyal [32] further refined this under adversarial contexts and non-uniform rewards, improving the dependency on κ and achieving $\widetilde{\mathcal{O}}(dK\sqrt{\kappa^*T}+d^2K^4/\kappa)$, but their method is not computationally tractable. Zhang and Sugiyama [44] proposed a computationally efficient algorithm using online parameter updates for a multi-parameter setting, where a shared contextual feature vector is used across items, but each item is associated with its own utility parameter. Lee and Oh [20] proposed a computationally efficient algorithm that achieves minimax optimal regret (up to logarithmic factors) in the single-parameter setting.

Extending beyond the linearity assumption, Zhang and Luo [41] is, to the best of our knowledge, the only work that addresses contextual MNL bandits with a general function class. For the stochastic context setting, they proposed two offline regression oracle-based algorithms. Assuming the function class is Lipschitz, the uniform exploration method achieves a regret bound of $\widetilde{\mathcal{O}}((dKN)^{1/3}T^{2/3})$, while the log-barrier regularization method achieves a tighter bound of $\widetilde{\mathcal{O}}(K^2\sqrt{dNT})$. For the adversarial context setting, they proposed two online regression oracle-based algorithms. Under the same Lipschitz assumption, the uniform exploration algorithm attains a regret of $\widetilde{\mathcal{O}}((NK)^{1/3}T^{5/6})$, and the log-barrier method again achieves $\widetilde{\mathcal{O}}(K^2\sqrt{dN}T^{3/4})$. Additionally, they showed that a variant of Feel-Good Thompson Sampling [42] achieves a regret of $\widetilde{\mathcal{O}}(K^2\sqrt{dNT})$. Even though these results offer significant theoretical insights for general utility function classes, the log-barrier regularized exploration and Thompson sampling-based algorithms are not computationally tractable—that is,

they cannot be solved in polynomial time. Furthermore, all of their algorithms exhibit polynomial dependence on the total number of items N, limiting their applicability to environments with large or infinite item sets. In contrast, our proposed algorithm is not only computationally tractable but also achieves a regret bound of $\widetilde{\mathcal{O}}(\sqrt{T})$, which remains independent of N, even in the presence of non-linear utilities.

Beyond linear models in contextual bandits. For generalized linear reward models, a growing body of work has provided algorithms with provable guarantees [11, 23, 9, 10]. Most of these methods achieve a regret bound of $\widetilde{\mathcal{O}}(d\sqrt{T})$, with subsequent improvements focusing on instance-dependent analysis and computational efficiency. For more general parametric reward models, Liu and Wang [25] consider the infinite-armed bandit setting and establish a regret bound of $\widetilde{\mathcal{O}}(d_w^2\sqrt{T})$, where d_w denotes the dimension of the parameterized reward model. Their analysis also relies on a geometric condition on the squared loss of the parametric model, but crucially assumes a unique global minimum—a stronger condition than our Assumption 4.

Beyond parametric reward models, there has been a surge of interest in bandit algorithms utilizing neural networks [45, 40, 43, 19]. However, these works typically rely on the assumption that the reward function lies in the reproducing kernel Hilbert space and heavily depend on NTK theory. Moreover, regret bounds derived under NTK assumptions are known to suffer from unfavorable worst-case information gain of order $\widetilde{\mathcal{O}}(T^{\frac{d-1}{d}})$, leading to a regret bound of $\widetilde{\mathcal{O}}(T^{\frac{2d-1}{2d}})$, which is significantly worse than the standard sublinear regret of $\widetilde{\mathcal{O}}(\sqrt{T})$ [19]. Without relying on NTK theory, Xu et al. [38] analyze a quadratic neural network with a single hidden layer and establish a regret bound of $\widetilde{\mathcal{O}}(T^{2/3})$. Huang et al. [15] consider a two-layer ReLU neural network with fewer hidden units than input dimensions under the assumption of Gaussian input distributions. Xu et al. [39] study a specific two-layer ReLU architecture in which all second-layer weights are fixed to one, assuming a unit sphere context distribution. While these approaches achieve provably efficient algorithms without relying on NTK assumptions, their dependence on specific network architectures and restrictive context distributions limits their applicability to more general settings.

B Proof of Lemma 1

Proof of Lemma 1. For notational simplicity, we abbreviate the following process:

$$X \sim \mathcal{D}, S \sim \text{Unif}(\mathcal{S}), i \sim \text{Multinomial}\left(1, [p(0 \mid X, S, \mathbf{w}^*), \dots, p(i_{|S|} \mid X, S, \mathbf{w}^*)]\right)$$
 (2)

as $(X,S,i) \sim \mathcal{H}$. Under Assumption 2, each function $f_{\mathbf{w}} \in \mathcal{F}$ is Lipschitz continuous with constant C_f . Then, by the regression oracle guarantee for Lipschitz function classes (Lemma 9), there exists a constant $C_1 > 0$ such that, with probability at least $1 - \delta$, the following bound holds:

$$\mathbb{E}_{(X,S,i)\sim\mathcal{H}}\left[\log p(i\mid X,S,\mathbf{w}^*) - \log p(i\mid X,S,\hat{\mathbf{w}}_0)\right] \le C_1 \frac{d_w C_f^2 \zeta}{n},$$

where $\zeta = \log K \log(nC_q) \log(1/\delta)$ is the logarithmic term.

For any $X = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}, S \in \mathcal{S}$, and $\mathbf{w} \in \mathcal{W}$, we define

$$\mathbf{p}(X, S, \mathbf{w}) := [p(i_1 \mid X, S, \mathbf{w}), \dots, p(i_{|S|} \mid X, S, \mathbf{w})],$$

$$\mathbf{p}_0(X, S, \mathbf{w}) := [p(0 \mid X, S, \mathbf{w}), p(i_1 \mid X, S, \mathbf{w}), \dots, p(i_{|S|} \mid X, S, \mathbf{w})].$$

Then, we have

$$C_{1} \frac{d_{w}C_{f}^{2}\zeta}{n} \geq \mathbb{E}_{(X,S,i)\sim\mathcal{H}}\left[\log p(i \mid X, S, \mathbf{w}^{*}) - \log p(i \mid X, S, \hat{\mathbf{w}}_{0})\right]$$

$$= \mathbb{E}_{X,S}\left[\mathbb{E}_{i\sim\mathbf{p}_{0}(X,S,\mathbf{w}^{*})}\left[\log p(i \mid X, S, \mathbf{w}^{*}) - \log p(i \mid X, S, \hat{\mathbf{w}}_{0}) \mid X, S\right]\right]$$

$$= \mathbb{E}_{X,S}\left[D_{KL}(\mathbf{p}_{0}(X, S, \mathbf{w}^{*}), \mathbf{p}_{0}(X, S, \hat{\mathbf{w}}_{0})) \mid X, S\right]$$

$$\geq \mathbb{E}_{X,S}\left[\frac{1}{2}\|\mathbf{p}_{0}(X, S, \mathbf{w}^{*}) - \mathbf{p}_{0}(X, S, \hat{\mathbf{w}}_{0})\|_{2}^{2} \mid X, S\right]$$

$$\geq \mathbb{E}_{X,S}\left[\frac{1}{2}\|\mathbf{p}_{0}(X, S, \mathbf{w}^{*}) - \mathbf{p}_{0}(X, S, \hat{\mathbf{w}}_{0})\|_{2}^{2} \mid X, S\right], \tag{3}$$

where the second inequality follows from Pinsker's inequality, and the final inequality follows from the definitions of \mathbf{p} and \mathbf{p}_0 . For the next step, we introduce the following lemma showing reverse Lipschitzness of the MNL model. The proof of Lemma 4 is provided in Section B.1.

Lemma 4 (Reverse Lipschitzness of the MNL model). For $\mathbf{a} \in \mathcal{A} := {\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\|_2 \le C}$, let $\mathbf{h}(\mathbf{a}) = [h_1(\mathbf{a}), \dots, h_d(\mathbf{a})]$ be defined as

$$h_i(\mathbf{a}) = \frac{\exp(a_i)}{1 + \sum_{j=1}^d \exp(a_j)}.$$

Then, for any $\mathbf{a}, \mathbf{b} \in \mathcal{A}$ *,*

$$\|\mathbf{h}(\mathbf{a}) - \mathbf{h}(\mathbf{b})\|_2 \ge \kappa_0 \|\mathbf{a} - \mathbf{b}\|_2$$

where $\kappa_0 := \min_{\mathbf{a} \in \mathcal{A}} \min_{i \in [d]} h_i(\mathbf{a}) (1 - \sum_{j=1}^d h_j(\mathbf{a})).$

Note that if we denote $\mathbf{f}(X,S,\mathbf{w}):=[f_{\mathbf{w}}(\mathbf{x}_j)]_{j\in S}\in\mathbb{R}^{|S|}$, then we have $\mathbf{h}(\mathbf{f}(X,S,\mathbf{w}))=\mathbf{p}(X,S,\mathbf{w})$. Therefore by applying Lemma 4, Eq. (3) can be further bounded by

$$(3) \ge \frac{\kappa^2}{2} \mathbb{E}_{X,S} \left[\sum_{j \in S} \left(f_{\hat{\mathbf{w}}_0}(\mathbf{x}_j) - f_{\mathbf{w}^*}(\mathbf{x}_j) \right)^2 \mid X, S \right] = \frac{\kappa^2}{2} \ell_{\mathrm{sq}}(\hat{\mathbf{w}}_0).$$

Therefore, we have

$$\ell_{\rm sq}(\hat{\mathbf{w}}_0) \le 2C_1 \frac{\kappa^{-2} d_w C_f^2 \zeta}{n} \,.$$

On the other hand, let $\widetilde{\mathbf{w}}^* = \operatorname{argmin}_{\widetilde{\mathbf{w}} \in \mathcal{W}^*} \|\hat{\mathbf{w}}_0 - \widetilde{\mathbf{w}}\|_2$. Then, by the condition on n,

$$n \ge 2^{-1} C \kappa^{-2} d_w C_f^2 \zeta \frac{\mu^{\gamma/(2-\gamma)}}{\tau^{2/(2-\gamma)}},$$
 (4)

where $C = 4C_1$, we can obtain

$$\|\hat{\mathbf{w}}_0 - \widetilde{\mathbf{w}}^*\|_2 < (\tau/\mu)^{1/(2-\gamma)}$$
.

We can prove it by contradiction. Suppose that

$$\|\hat{\mathbf{w}}_0 - \widetilde{\mathbf{w}}^*\|_2 > (\tau/\mu)^{1/(2-\gamma)}$$
.

Then, we have

$$\frac{\mu}{2} \|\hat{\mathbf{w}}_0 - \widetilde{\mathbf{w}}^*\|_2^2 > \frac{\tau}{2} \|\hat{\mathbf{w}}_0 - \widetilde{\mathbf{w}}^*\|_2^{\gamma}, \tag{5}$$

which implies that the (τ, γ) -growth condition holds by Assumption 4. On the other hand, by the (τ, γ) -growth condition in Assumption 4, Eq. (5) implies

$$\frac{\tau}{2} \|\hat{\mathbf{w}}_0 - \widetilde{\mathbf{w}}^*\|_2^{\gamma} \le \ell_{\mathrm{sq}}(\hat{\mathbf{w}}_0) - \ell_{\mathrm{sq}}(\widetilde{\mathbf{w}}^*) = \ell_{\mathrm{sq}}(\hat{\mathbf{w}}_0) \le 2C_1 \frac{\kappa^{-2} d_w C_f^2 \zeta}{n}.$$

By the condition on n in Eq. (4), this implies

$$\|\hat{\mathbf{w}}_0 - \widetilde{\mathbf{w}}^*\|_2 \le (\tau/\mu)^{1/(2-\gamma)}$$

which leads to a contradiction. This shows that when n satisfies Eq. (4), then $\hat{\mathbf{w}}_0$ lies in the region where the local strong convexity of $\ell_{\mathrm{sq}}(\cdot)$ at $\widetilde{\mathbf{w}}^*$ holds. Therefore, we obtain

$$\frac{\mu}{2} \|\hat{\mathbf{w}}_0 - \widetilde{\mathbf{w}}^*\|_2^2 \le \ell_{\mathrm{sq}}(\hat{\mathbf{w}}_0) - \ell_{\mathrm{sq}}(\widetilde{\mathbf{w}}^*) = \ell_{\mathrm{sq}}(\hat{\mathbf{w}}_0) - \ell_{\mathrm{sq}}(\mathbf{w}^*) \le 2C_1 \frac{\kappa^{-2} d_w C_f^2 \zeta}{n},$$

which results in

$$\|\hat{\mathbf{w}}_0 - \widetilde{\mathbf{w}}^*\|_2^2 \le C \frac{\kappa^{-2} d_w C_f^2 \zeta}{\mu n},$$

where we again denote $C = 4C_1$.

B.1 Proofs of Technical Lemmas

B.1.1 Proof of Lemma 4

Proof of Lemma 4. For any $\mathbf{a}, \mathbf{b} \in \mathcal{A}$, by the mean value theorem there exists $\mathbf{c} \in \operatorname{conv}(\mathbf{a}, \mathbf{b})$ such that

$$\mathbf{h}(\mathbf{a}) - \mathbf{h}(\mathbf{b}) = \nabla \mathbf{h}(\mathbf{c})^{\top} (\mathbf{a} - \mathbf{b}).$$

Now, it suffices to show that $\lambda_{\min}(\nabla(\mathbf{h}(\mathbf{c})^{\top}) \geq \kappa_0$. Since $\nabla \mathbf{h}(\mathbf{c}) = \operatorname{diag}(\mathbf{h}(\mathbf{c})) - \mathbf{h}(\mathbf{c})\mathbf{h}(\mathbf{c})^{\top}$, for any $\mathbf{x} \in \mathbb{R}^d \setminus \{\mathbf{0}_d\}$,

$$\mathbf{x}^{\top} \nabla \mathbf{h}(\mathbf{c}) \mathbf{x} = \mathbf{x}^{\top} \operatorname{diag}(\mathbf{h}(\mathbf{c})) \mathbf{x} - \left(\mathbf{x}^{\top} \mathbf{h}(\mathbf{c})\right)^{2}$$

$$= \sum_{i=1}^{d} h_{i}(\mathbf{c}) x_{i}^{2} - \left(\sum_{i=1}^{d} h_{i}(\mathbf{c}) x_{i}\right)^{2}$$

$$\geq \sum_{i=1}^{d} h_{i}(\mathbf{c}) x_{i}^{2} - \left(\sum_{i=1}^{d} h_{i}(\mathbf{c})\right) \left(\sum_{i=1}^{d} h_{i}(\mathbf{c}) x_{i}^{2}\right)$$

$$= \left(\sum_{i=1}^{d} h_{i}(\mathbf{c}) x_{i}^{2}\right) \left(1 - \sum_{i=1}^{d} h_{i}(\mathbf{c})\right)$$

$$= \sum_{i=1}^{d} h_{i}(\mathbf{c}) \left(1 - \sum_{i=1}^{d} h_{i}(\mathbf{c})\right) x_{i}^{2}$$

$$\geq \sum_{i=1}^{d} \kappa_{0} x_{i}^{2},$$

where the first inequality uses Cauchy-Schwarz inequality, and the last inequality follows by the definition of κ_0 . This concludes the proof.

C Proof of Lemma 2

Proof of Lemma 2. Let $\widetilde{\mathbf{w}}^* = \operatorname{argmin}_{\widetilde{\mathbf{w}} \in \mathcal{W}^*} \|\hat{\mathbf{w}}_0 - \widetilde{\mathbf{w}}\|_2$. And for notational simplicity, we will abbreviate $\hat{p}(i \mid X_t, S_t, \mathbf{w})$ as $\hat{p}_{ti}(\mathbf{w})$, and $p(i \mid X_t, S_t, \mathbf{w})$ as $p_{ti}(\mathbf{w})$ in this proof. Recall that $\ell_t(\mathbf{w})$ is the regularized negative log-likelihood defined as

$$\ell_t(\mathbf{w}) := -\sum_{s=1}^{t-1} \sum_{i \in S_s} y_{si} \log \hat{p}_{si}(\mathbf{w}) + \frac{\lambda}{2} \|\mathbf{w} - \hat{\mathbf{w}}_0\|_2^2.$$

Since $\hat{\mathbf{w}}_t = \operatorname{argmin}_{\mathbf{w} \in \mathcal{W}} \ell_t(\mathbf{w})$, we have $\nabla \ell_t(\mathbf{w}_t) = \mathbf{0}_{d_w}$, and this implies

$$\sum_{s=1}^{t-1} \sum_{i \in S_s} \hat{p}_{si}(\hat{\mathbf{w}}_t) \nabla f_{\hat{\mathbf{w}}_s}(\mathbf{x}_{si}) + \lambda (\hat{\mathbf{w}}_t - \hat{\mathbf{w}}_0) = \sum_{s=1}^{t-1} \sum_{i \in S_s} y_{si} \nabla f_{\hat{\mathbf{w}}_s}(\mathbf{x}_{si}).$$
 (6)

Let us define $\mathbf{g}_t(\mathbf{w}) = \sum_{s=1}^{t-1} \sum_{i \in S_s} \hat{p}_{si}(\mathbf{w}) \nabla f_{\hat{\mathbf{w}}_s}(\mathbf{x}_{si}) + \lambda \mathbf{w}$. Then, from Eq. (6) we have

$$\mathbf{g}_t(\hat{\mathbf{w}}_t) - \mathbf{g}_t(\widetilde{\mathbf{w}}^*) = \sum_{s=1}^{t-1} \sum_{i \in S_s} (y_{si} - \hat{p}_{si}(\widetilde{\mathbf{w}}^*)) \nabla f_{\hat{\mathbf{w}}_s}(\mathbf{x}_{si}) + \lambda (\hat{\mathbf{w}}_0 - \widetilde{\mathbf{w}}^*).$$

Note that since $\mathbb{E}[y_{si} \mid \mathcal{F}_s] \neq \hat{p}_{si}(\widetilde{\mathbf{w}}^*)$, we cannot apply conventional self-normalized martingale inequality used in previous MNL bandit [28, 29, 32, 44, 20]. Instead since $\mathbb{E}[y_{si} \mid \mathcal{F}_s] = p_{si}(\widetilde{\mathbf{w}}^*)$ we add and subtract $p_{si}(\widetilde{\mathbf{w}}^*)$ as follows:

$$\mathbf{g}_{t}(\hat{\mathbf{w}}_{t}) - \mathbf{g}_{t}(\widetilde{\mathbf{w}}^{*}) = \sum_{s=1}^{t-1} \sum_{i \in S_{s}} (y_{si} - p_{si}(\widetilde{\mathbf{w}}^{*})) \nabla f_{\hat{\mathbf{w}}_{s}}(\mathbf{x}_{si}) + \sum_{s=1}^{t-1} \sum_{i \in S_{s}} (p_{si}(\widetilde{\mathbf{w}}^{*}) - \hat{p}_{si}(\widetilde{\mathbf{w}}^{*})) \nabla f_{\hat{\mathbf{w}}_{s}}(\mathbf{x}_{si}) + \lambda (\hat{\mathbf{w}}_{0} - \widetilde{\mathbf{w}}^{*}).$$

By inequality $(a+b+c)^2 \le 4a^2+4b^2+4c^2$, we have

$$\|\mathbf{g}_{t}(\hat{\mathbf{w}}_{t}) - \mathbf{g}_{t}(\tilde{\mathbf{w}}^{*})\|_{\mathbf{V}_{t}^{-1}}^{2} \leq 4 \underbrace{\left\| \sum_{s=1}^{t-1} \sum_{i \in S_{s}} \left(y_{si} - p_{si}(\tilde{\mathbf{w}}^{*}) \right) \nabla f_{\hat{\mathbf{w}}_{s}}(\mathbf{x}_{si}) \right\|_{\mathbf{V}_{t}^{-1}}^{2}}_{I_{1}} + 4 \underbrace{\left\| \sum_{s=1}^{t-1} \sum_{i \in S_{s}} \left(p_{si}(\tilde{\mathbf{w}}^{*}) - \hat{p}_{si}(\tilde{\mathbf{w}}^{*}) \right) \nabla f_{\hat{\mathbf{w}}_{s}}(\mathbf{x}_{si}) \right\|_{\mathbf{V}_{t}^{-1}}^{2}}_{I_{2}} + \underbrace{4\lambda^{2} \|\hat{\mathbf{w}}_{0} - \tilde{\mathbf{w}}^{*}\|_{\mathbf{V}_{t}^{-1}}^{2}}_{I_{2}}$$

For the term I_1 , we introduce the following lemma.

Lemma 5. For all $t \ge 1$, with probability at least $1 - \delta$,

$$\left\| \sum_{s=1}^{t-1} \sum_{i \in S_s} (y_{si} - p_{si}(\widetilde{\mathbf{w}}^*)) \nabla f_{\widehat{\mathbf{w}}_s}(\mathbf{x}_{si}) \right\|_{\mathbf{V}_t^{-1}}^2 \le \alpha_t,$$

where $\alpha_t = \widetilde{C}_{\alpha} d_w \log(1 + T C_q^2/(d_w \lambda)) \log(t^2 \pi^2/\delta)$ for an absolute constant $\widetilde{C}_{\alpha} > 0$.

If we denote each $y_{si} - p_{si}(\mathbf{w}^*)$ as ε_{si} and note $\mathbb{E}[\varepsilon_{si}] = 0$, the quantity $\left\|\sum_{s=1}^{t-1}\sum_{i\in S_s}(y_{si}-p_{si}(\widetilde{\mathbf{w}}^*))\nabla f_{\widehat{\mathbf{w}}_s}(\mathbf{x}_{si})\right\|_{\mathbf{V}_t^{-1}}^2$ corresponds to the self-normalized vector-valued martingale term in Abbasi-Yadkori et al. [1]. However, due to correlations among items within each selected assortment $\{\varepsilon_{si}\}_{i\in S_s}$, standard self-normalized inequalities cannot be directly applied. To address this, Perivier and Goyal [32] introduced an analysis based on a global random vector $\mathbf{z}_s := \sum_{i\in S_s} \varepsilon_{si} \nabla f_{\widehat{\mathbf{w}}_s}(\mathbf{x}_{si})$, which, combined with a Bernstein-type tail inequality for the logistic bandit (Theorem 1 in Faury et al. [9]), yields a concentration bound of order $\widetilde{\mathcal{O}}\left(\sqrt{\lambda} + \frac{d_w \log T}{\sqrt{\lambda}}\right)$ (Theorem C.6 in Perivier and Goyal [32]). As a result, achieving a regret bound of $\widetilde{\mathcal{O}}(\sqrt{T})$ requires

Theorem C.6 in Perivier and Goyal [32]). As a result, achieving a regret bound of $\mathcal{O}(\sqrt{T})$ requires the regularization parameter as $\lambda = \mathcal{O}(d \log T)$. In contrast, by adapting techniques from Dani et al. [7], we derive a concentration result with a bound that is independent of λ (up to logarithmic factors). The proof of Lemma 5 is provided in Section C.1.

For the term I_2 , we introduce several notations to facilitate understanding. Let $K_s = |S_s|$ denote the size of the assortment selected at round s. Then we can write $S_s = \{i_1, \ldots, i_{K_s}\}$. Now we define $\mathbf{G}_s \in \mathbb{R}^{d \times K_s}$ as the matrix whose k-th column is the gradient of $f(\mathbf{x}_{si_k})$ evaluated at $\hat{\mathbf{w}}_s$, i.e.,

$$\mathbf{G}_s := \left[
abla f_{\hat{\mathbf{w}}_s}(\mathbf{x}_{si_1}), \dots,
abla f_{\hat{\mathbf{w}}_s}(\mathbf{x}_{si_{K_s}}) \right] \in \mathbb{R}^{d_w imes K_s}$$

For any vector $\mathbf{u} := [u_1, \dots, u_{K_s}]^{\top} \in \mathbb{R}^{K_s}$, we define $\mathbf{h}_s(\mathbf{u}) := [h_1(\mathbf{u}), \dots, h_{K_s}(\mathbf{u})]^{\top}$, where each h_k is given by:

$$h_k(\mathbf{u}) := \frac{\exp(u_k)}{1 + \sum_{j=1}^{K_s} \exp(u_j)}.$$

In addition, we define:

$$\begin{split} \hat{\mathbf{u}}_s &:= [\hat{u}_{si_1}, \dots, \hat{u}_{si_{K_s}}]^\top \in \mathbb{R}^{K_s} \quad \text{where} \quad \hat{u}_{si_k} = f_{\hat{\mathbf{w}}_s}(\mathbf{x}_{si_k}) + (\tilde{\mathbf{w}}^* - \hat{\mathbf{w}}_s)^\top \nabla f_{\hat{\mathbf{w}}_s}(\mathbf{x}_{si_k}) \,, \\ \mathbf{u}_s^* &:= [u_{si_1}^*, \dots, u_{si_{K_s}}^*]^\top \in \mathbb{R}^{K_s} \quad \text{where} \quad u_{si_k}^* = f_{\tilde{\mathbf{w}}^*}(\mathbf{x}_{si_k}) \,. \end{split}$$

Now we bound I_2 as follows:

$$I_{2} = 4 \left\| \sum_{s=1}^{t-1} \mathbf{G}_{s}(\mathbf{h}_{s}(\hat{\mathbf{u}}_{s}) - \mathbf{h}_{s}(\mathbf{u}_{s}^{*})) \right\|_{\mathbf{V}_{t}^{-1}}^{2}$$

$$= 4 \left\| \sum_{s=1}^{t-1} \mathbf{G}_{s} \nabla \mathbf{h}_{s}(\widetilde{\mathbf{u}}_{s})^{\top} (\hat{\mathbf{u}}_{s} - \mathbf{u}_{s}^{*}) \right\|_{\mathbf{V}_{t}^{-1}}^{2}$$

$$\leq 4 \left(\sum_{s=1}^{t-1} \left\| \mathbf{G}_{s} \nabla \mathbf{h}_{s}(\widetilde{\mathbf{u}}_{s})^{\top} (\hat{\mathbf{u}}_{s} - \mathbf{u}_{s}^{*}) \right\|_{\mathbf{V}_{t}^{-1}}^{2} \right)^{2}$$

$$\leq 4 \left(\sum_{s=1}^{t-1} \left\| \mathbf{G}_{s} \right\|_{\mathbf{V}_{t}^{-1}} \left\| \nabla \mathbf{h}_{s}(\widetilde{\mathbf{u}}_{s})^{\top} (\hat{\mathbf{u}}_{s} - \mathbf{u}_{s}^{*}) \right\|_{2}^{2} \right)^{2}$$

where the first inequality follows from the triangle inequality, and the second inequality uses the following bound: for any matrices $\mathbf{A}, \mathbf{B}, \mathbf{V} \in \mathbb{R}^{m \times m}$ and vector $\mathbf{x} \in \mathbb{R}^m$,

$$\|\mathbf{A}\mathbf{B}\mathbf{x}\|_{\mathbf{V}}^2 = (\mathbf{A}\mathbf{B}\mathbf{x})^{\top}\mathbf{V}(\mathbf{A}\mathbf{B}\mathbf{x}) = (\mathbf{B}\mathbf{x})^{\top}\mathbf{A}^{\top}\mathbf{V}\mathbf{A}(\mathbf{B}\mathbf{x}) \leq \lambda_{\max}(\mathbf{A}^{\top}\mathbf{V}\mathbf{A})\|\mathbf{B}\mathbf{x}\|_2^2 = \|\mathbf{A}\|_{\mathbf{V}}^2\|\mathbf{B}\mathbf{x}\|_2^2.$$

On the other hand, we can bound each $\|\nabla \mathbf{h}_s(\widetilde{\mathbf{u}}_s)^{\top}(\hat{\mathbf{u}}_s - \mathbf{u}_s^*)\|_2$ as follows:

$$\|\nabla \mathbf{h}_s(\widetilde{\mathbf{u}}_s)^{\top}(\hat{\mathbf{u}}_s - \mathbf{u}_s^*)\|_2 \leq \|\nabla \mathbf{h}_s(\widetilde{\mathbf{u}}_s)^{\top}\|_{\infty,2} \|\hat{\mathbf{u}}_s - \mathbf{u}_s^*\|_{\infty},$$

where the norm $\|\cdot\|_{\infty,2}$ is defined as follows:

$$\|\mathbf{A}\|_{\infty,2} = \sup\{\|\mathbf{A}\mathbf{x}\|_2 : \|\mathbf{x}\|_{\infty} \le 1\}.$$

Since $\nabla \mathbf{h}_s(\widetilde{\mathbf{u}}_s)^{\top} = \operatorname{diag}(\mathbf{h}_s(\widetilde{\mathbf{u}}_s)) - \mathbf{h}_s(\widetilde{\mathbf{u}}_s)\mathbf{h}_s(\widetilde{\mathbf{u}}_s)^{\top}$, for any $\mathbf{x} := [x_1, \dots, x_{K_s}]^{\top}$ with $\|\mathbf{x}\|_{\infty} \leq 1$, we have

$$\|\nabla \mathbf{h}_{s}(\widetilde{\mathbf{u}}_{s})^{\top} \mathbf{x}\|_{2}^{2} = \|\operatorname{diag}(\mathbf{h}_{s}(\widetilde{\mathbf{u}}_{s})) \mathbf{x} - (\mathbf{h}_{s}(\widetilde{\mathbf{u}}_{s})^{\top} \mathbf{x}) \mathbf{h}_{s}(\widetilde{\mathbf{u}}_{s})\|_{2}^{2}$$

$$= \sum_{k=1}^{K_{s}} \left(h_{k}(\widetilde{\mathbf{u}}_{s}) x_{k} - \left(\sum_{j=1}^{K_{s}} x_{j} h_{j}(\widetilde{\mathbf{u}}_{s}) \right) h_{k}(\widetilde{\mathbf{u}}_{s}) \right)^{2}$$

$$= \sum_{k=1}^{K_{s}} h_{k}^{2}(\widetilde{\mathbf{u}}_{s}) \left(x_{k} - \sum_{j=1}^{K_{s}} x_{j} h_{j}(\widetilde{\mathbf{u}}_{s}) \right)^{2}.$$

If we denote $\sum_{j=1}^{K_s} h_j(\widetilde{\mathbf{u}}_s) =: q$, since $\|\mathbf{x}\|_{\infty} \leq 1$, for all $1 \leq k \leq K_s$, we have

$$\left| x_k - \sum_{j=1}^{K_s} x_j h_j(\widetilde{\mathbf{u}}_s) \right| \le 1 + \left| \sum_{j=1}^{K_s} x_j h_j(\widetilde{\mathbf{u}}_s) \right| \le 1 + q.$$

Then, $\|\nabla \mathbf{h}_s(\widetilde{\mathbf{u}}_s)^{\top} \mathbf{x}\|_2^2$ can be further bounded as

$$\|\nabla \mathbf{h}_s(\widetilde{\mathbf{u}}_s)^{\top} \mathbf{x}\|_2^2 \leq \sum_{i=1}^{K_s} h_i^2(\widetilde{\mathbf{u}}_s)(1+q)^2 \leq \sum_{i=1}^{K_s} h_i(\widetilde{\mathbf{u}}_s)(1+q)^2 \leq q(1+q)^2 \leq \max_{p \in (0,1)} p(1+p)^2 \leq 4,$$

where the fourth inequality holds because 0 < q < 1. This implies $\|\nabla \mathbf{h}_s(\widetilde{\mathbf{u}}_s)^{\top}\|_{\infty,2} \le 2$ for all s, therefore we bound the term I_2 as follows:

$$I_{2} \leq 16 \left(\sum_{s=1}^{t-1} \|\mathbf{G}_{s}\|_{\mathbf{V}_{t}^{-1}} \|\hat{\mathbf{u}}_{s} - \mathbf{u}_{s}^{*}\|_{\infty} \right)^{2}$$

$$= 16 \left(\sum_{s=1}^{t-1} \|\mathbf{G}_{s}\|_{\mathbf{V}_{t}^{-1}} \max_{1 \leq k \leq K_{s}} |\hat{u}_{si_{k}} - u_{si_{k}}^{*}| \right)^{2}$$

$$= 16 \left(\sum_{s=1}^{t-1} \|\mathbf{G}_{s}\|_{\mathbf{V}_{t}^{-1}} \max_{i \in S_{s}} \frac{1}{2} \|\hat{\mathbf{w}}_{s} - \widetilde{\mathbf{w}}^{*}\|_{\nabla^{2} f_{\widetilde{\mathbf{w}}_{s}}(\mathbf{x}_{si})}^{2} \right)^{2}$$

$$\leq 4C_{h}^{2} \left(\sum_{s=1}^{t-1} \|\mathbf{G}_{s}\|_{\mathbf{V}_{t}^{-1}} \|\hat{\mathbf{w}}_{s} - \widetilde{\mathbf{w}}^{*}\|_{2}^{2} \right)^{2},$$

where the third equality follows from the second-order Taylor's theorem with $\widetilde{\mathbf{w}}_s$ lying between $\hat{\mathbf{w}}_s$ and $\widetilde{\mathbf{w}}^*$, and the final inequality follows from Assumption 2.

For I_3 , we have

$$4\lambda^2 \|\hat{\mathbf{w}}_0 - \widetilde{\mathbf{w}}^*\|_{\mathbf{V}_t^{-1}}^2 \le 4\lambda \|\hat{\mathbf{w}}_0 - \widetilde{\mathbf{w}}^*\|_2^2 \le \frac{4\lambda C\kappa^{-2} d_w C_f^2 \zeta}{\mu n}.$$

For next step, we introduce the following lemma, whose proof is provided in Section C.1.

Lemma 6. For any $t \in [T]$, it holds:

$$\|\hat{\mathbf{w}}_t - \widetilde{\mathbf{w}}^*\|_{\mathbf{V}_t}^2 \le (1 + 6\sqrt{2})^2 \kappa^{-1} \|\mathbf{g}_t(\hat{\mathbf{w}}_t) - \mathbf{g}_t(\widetilde{\mathbf{w}}^*)\|_{\mathbf{V}_t^{-1}}^2.$$

Combining Lemma 6 with the bounds of I_1, I_2, I_3 , we have

$$\|\hat{\mathbf{w}}_{t} - \widetilde{\mathbf{w}}^{*}\|_{2}^{2} \leq \frac{(1 + 6\sqrt{2})^{2}\kappa^{-1}}{\lambda} \|\mathbf{g}_{t}(\hat{\mathbf{w}}_{t}) - \mathbf{g}_{t}(\widetilde{\mathbf{w}}^{*})\|_{\mathbf{V}_{t}^{-1}}^{2}$$

$$\leq \frac{4(1 + 6\sqrt{2})^{2}\kappa^{-1}}{\lambda} \left(\alpha_{t} + C_{h}^{2} \left(\sum_{s=1}^{t-1} \|\mathbf{G}_{s}\|_{\mathbf{V}_{t}^{-1}} \|\hat{\mathbf{w}}_{s} - \widetilde{\mathbf{w}}^{*}\|_{2}^{2}\right)^{2} + \frac{\lambda\kappa^{-2}d_{w}C_{f}^{2}\zeta}{\mu n}\right)$$

$$= \frac{C_{0}\kappa^{-1}\alpha_{t}}{\lambda} + \frac{C_{0}C_{h}^{2}\kappa^{-1}}{\lambda} \left(\sum_{s=1}^{t-1} \|\mathbf{G}_{s}\|_{\mathbf{V}_{t}^{-1}} \|\hat{\mathbf{w}}_{s} - \widetilde{\mathbf{w}}^{*}\|_{2}^{2}\right)^{2} + \frac{C_{0}\kappa^{-3}d_{w}C_{f}^{2}\zeta}{\mu n}, \quad (7)$$

where we denote $C_0 := 4(1+6\sqrt{2})^2$. We now proceed by induction to establish the convergence rate of $\|\hat{\mathbf{w}}_t - \widetilde{\mathbf{w}}^*\|_2^2$ for every time step $t \in [T]$. Recall that by Lemma 1, with probability at least $1-\delta$, we have

$$\|\hat{\mathbf{w}}_0 - \widetilde{\mathbf{w}}^*\|_2^2 \le \frac{C\kappa^{-2} d_w C_f^2 \zeta}{\mu m}.$$

To formally set up the inductive argument, we assume that at round τ , there exists a universal constant \widetilde{C} such that, with probability at least $1-2\delta$, the following holds:

$$\|\hat{\mathbf{w}}_{\tau} - \widetilde{\mathbf{w}}^*\|_2^2 \le \frac{\widetilde{C}\kappa^{-2}d_w C_f^2 \zeta}{\mu n}.$$

Then, we want to show that at round $\tau + 1$ with probability at least $1 - 2\delta$,

$$\|\hat{\mathbf{w}}_{\tau+1} - \widetilde{\mathbf{w}}^*\|_2^2 \le \frac{\widetilde{C}\kappa^{-2}d_w C_f^2 \zeta}{\mu n}.$$

From Eq. (7), at round s + 1, we have

$$\|\hat{\mathbf{w}}_{\tau+1} - \widetilde{\mathbf{w}}^*\|_{2}^{2} \leq \frac{C_{0}\kappa^{-1}\alpha_{\tau+1}}{\lambda} + \frac{C_{0}C_{h}^{2}\kappa^{-1}}{\lambda} \left(\sum_{s=1}^{\tau} \|\mathbf{G}_{s}\|_{\mathbf{V}_{\tau+1}^{-1}} \|\hat{\mathbf{w}}_{s} - \widetilde{\mathbf{w}}^*\|_{2}^{2}\right)^{2} + \frac{C_{0}\kappa^{-3}d_{w}C_{f}^{2}\zeta}{\mu n}$$

$$\leq \frac{C_{0}\kappa^{-1}\alpha_{\tau+1}}{\lambda} + \frac{C_{0}C_{h}^{2}\widetilde{C}^{2}\kappa^{-5}d_{w}^{2}C_{f}^{4}\zeta^{2}}{\lambda\mu^{2}n^{2}} \left(\sum_{s=1}^{\tau} \|\mathbf{G}_{s}\|_{\mathbf{V}_{\tau+1}^{-1}}\right)^{2} + \frac{C_{0}\kappa^{-3}d_{w}C_{f}^{2}\zeta}{\mu n}$$

$$\leq \frac{C_{0}\kappa^{-1}\alpha_{\tau+1}}{\lambda} + \frac{C_{0}C_{h}^{2}\widetilde{C}^{2}\kappa^{-5}d_{w}^{2}C_{f}^{4}\zeta^{2}}{\lambda\mu^{2}n^{2}} \left(\sqrt{\sum_{s=1}^{\tau} \left\|\mathbf{G}_{s}\right\|_{\mathbf{V}_{\tau+1}^{-1}}^{2}}\right)^{2} + \frac{C_{0}\kappa^{-3}d_{w}C_{f}^{2}\zeta}{\mu n}$$

$$\leq \frac{C_{0}\kappa^{-1}\alpha_{\tau+1}}{\lambda} + \frac{C_{0}C_{h}^{2}\widetilde{C}^{2}\kappa^{-5}d_{w}^{2}C_{f}^{4}\zeta^{2}\tau}{\lambda\mu^{2}n^{2}} \left(\sum_{s=1}^{\tau} \|\mathbf{G}_{s}\|_{\mathbf{V}_{\tau+1}^{-1}}^{2}\right) + \frac{C_{0}\kappa^{-3}d_{w}C_{f}^{2}\zeta}{\mu n}$$

$$\leq \frac{C_{0}\kappa^{-1}\alpha_{\tau+1}}{\lambda} + \frac{C_{0}C_{h}^{2}\widetilde{C}^{2}\kappa^{-5}d_{w}^{3}C_{f}^{4}\zeta^{2}\tau}{\lambda\mu^{2}n^{2}} + \frac{C_{0}\kappa^{-3}d_{w}C_{f}^{2}\zeta}{\mu n}, \tag{8}$$

where the last inequality invokes the following lemma, with its proof provided in Section C.1.

Lemma 7. For all $t \in [T]$,

$$\sum_{s=1}^{t-1} \|\mathbf{G}_s\|_{\mathbf{V}_t^{-1}}^2 \le d_w.$$

Now, our goal is to show that there exist a universal constant \widetilde{C} such that the following bound holds:

$$\frac{C_0 \kappa^{-1} \alpha_{\tau+1}}{\lambda} + \frac{C_0 C_h^2 \widetilde{C}^2 \kappa^{-5} d_w^3 C_f^4 \zeta^2 \tau}{\lambda \mu^2 n^2} + \frac{C_0 \kappa^{-3} d_w C_f^2 \zeta}{\mu n} \le \frac{\widetilde{C} \kappa^{-2} d_w C_f^2 \zeta}{\mu n} \,. \tag{9}$$

For notational simplicity, let us denote $\alpha_t := \widetilde{C}_{\alpha} d_w \zeta_1$, where ζ_1 denotes a logarithmic term that depends on T, d_w, λ , and δ . Then, Eq. (9) can be rewritten as:

$$\frac{C_0\widetilde{C}_{\alpha}\kappa^{-1}d_w\zeta_1}{\lambda} + \frac{C_0C_h^2\widetilde{C}^2\kappa^{-5}d_w^3C_f^4\zeta^2\tau}{\lambda\mu^2n^2} + \frac{C_0\kappa^{-3}d_wC_f^2\zeta}{\mu n} \leq \frac{\widetilde{C}\kappa^{-2}d_wC_f^2\zeta}{\mu n} \,. \tag{10}$$

Note that the left-hand side of Eq. (10) is monotonically increasing in τ , and therefore the inequality must also hold at $\tau = T$; that is,

$$\begin{split} &\frac{C_0 \widetilde{C}_{\alpha} \kappa^{-1} d_w \zeta_1}{\lambda} + \frac{C_0 C_h^2 \widetilde{C}^2 \kappa^{-5} d_w^3 C_f^4 \zeta^2 T}{\lambda \mu^2 n^2} + \frac{C_0 \kappa^{-3} d_w C_f^2 \zeta}{\mu n} \leq \frac{\widetilde{C} \kappa^{-2} d_w C_f^2 \zeta}{\mu n} \\ &\iff \left(\frac{C_0 C_h^2 \kappa^{-5} d_w^3 C_f^4 \zeta^2 T}{\lambda \mu^2 n^2}\right) \widetilde{C}^2 - \left(\frac{\kappa^{-2} d_w C_f^2 \zeta}{\mu n}\right) \widetilde{C} + \left(\frac{C_0 \kappa^{-3} d_w C_f^2 \zeta}{\mu n} + \frac{C_0 \widetilde{C}_{\alpha} \kappa^{-1} d_w \zeta_1}{\lambda}\right) \leq 0 \,. \end{split}$$

A feasible choice of \widetilde{C} must satisfy the following condition:

$$\left(\frac{\kappa^{-2}d_w C_f^2 \zeta}{\mu n}\right)^2 - 4\left(\frac{C_0 C_h^2 \kappa^{-5} d_w^3 C_f^4 \zeta^2 T}{\lambda \mu^2 n^2}\right) \left(\frac{C_0 \kappa^{-3} d_w C_f^2 \zeta}{\mu n} + \frac{C_0 \widetilde{C}_\alpha \kappa^{-1} d_w \zeta_1}{\lambda}\right) \ge 0.$$

By substituting $\lambda = C_{\lambda} \sqrt{T}$ and $n = \kappa^{-3/2} d_w \sqrt{T}$, this condition becomes:

$$\begin{split} & \left(\frac{\kappa^{-1/2}C_f^2\zeta}{\mu\sqrt{T}}\right)^2 - 4\left(\frac{C_0C_h^2\kappa^{-2}d_wC_f^4\zeta^2}{C_\lambda\sqrt{T}\mu^2}\right) \left(\frac{C_0\kappa^{-3/2}C_f^2\zeta}{\mu\sqrt{T}} + \frac{C_0\widetilde{C}_\alpha\kappa^{-1}d_w\zeta_1}{C_\lambda\sqrt{T}}\right) \geq 0 \\ & \iff (\kappa^{-1}C_f^4\zeta^2\mu)C_\lambda^2 - (4C_0^2C_h^2\kappa^{-7/2}d_wC_f^6\zeta^3)C_\lambda - (4C_0^2C_h^2\widetilde{C}_\alpha\kappa^{-3}d_w^2C_f^4\zeta^2\zeta_1\mu) \geq 0 \,. \end{split}$$

Solving the above quadratic inequality for C_{λ} , it suffices to choose

$$C_{\lambda} \geq \frac{2C_{0}C_{h}d_{w}\kappa^{-5/2}\left(C_{0}C_{f}^{2}C_{h}\zeta + \sqrt{C_{0}^{2}C_{f}^{4}C_{h}^{2}\kappa^{4}\zeta^{2} + \widetilde{C}_{\alpha}\kappa^{7}\mu^{2}\zeta_{1}}\right)}{\mu} = \widetilde{\mathcal{O}}(\kappa^{-5/2}\mu^{-1}d_{w}).$$

Note that C_{λ} depends only logarithmically on T, and this choice guarantees that Eq.(10) holds, thereby confirming the existence of a universal constant \widetilde{C} . Hence, by induction, we conclude that for all $t \in [T]$, there exists a universal constant \widetilde{C} such that, with probability at least $1-2\delta$, the desired bound holds:

$$\|\hat{\mathbf{w}}_t - \widetilde{\mathbf{w}}^*\|_2^2 \le \frac{\widetilde{C}\kappa^{-2}d_w C_f^2 \zeta}{\mu n}.$$

Finally, combining again the results of Lemma 6 and similar argument used in Eq. (8), we have

$$\begin{split} &\|\hat{\mathbf{w}}_{t} - \widetilde{\mathbf{w}}^{*}\|_{\mathbf{V}_{t}}^{2} \\ &\leq (1 + 6\sqrt{2})^{2}\kappa^{-1}\|\mathbf{g}_{t}(\hat{\mathbf{w}}_{t}) - \mathbf{g}_{t}(\widetilde{\mathbf{w}}^{*})\|_{\mathbf{V}_{t}^{-1}}^{2} \\ &\leq 4(1 + 6\sqrt{2})^{2}\kappa^{-1}\left(\alpha_{t} + C_{h}^{2}\left(\sum_{s=1}^{t-1}\|\mathbf{G}_{s}\|_{\mathbf{V}_{t}^{-1}}\|\hat{\mathbf{w}}_{s} - \widetilde{\mathbf{w}}^{*}\|_{2}^{2}\right)^{2} + \frac{\lambda\kappa^{-2}d_{w}C_{f}^{2}\zeta}{\mu^{n}}\right) \\ &\leq C_{0}\kappa^{-1}\left(\alpha_{t} + \frac{C_{h}^{2}\widetilde{C}^{2}\kappa^{-4}d_{w}^{2}C_{f}^{4}\zeta^{2}}{\mu^{2}n^{2}}\left(\sum_{s=1}^{t-1}\|\mathbf{G}_{s}\|_{\mathbf{V}_{t}^{-1}}\right)^{2} + \frac{\lambda\kappa^{-2}d_{w}C_{f}^{2}\zeta}{\mu^{n}}\right) \\ &\leq C_{0}\kappa^{-1}\left(\alpha_{t} + \frac{C_{h}^{2}\widetilde{C}^{2}\kappa^{-4}d_{w}^{2}C_{f}^{4}\zeta^{2}}{\mu^{2}n^{2}}\left(\sqrt{\sum_{s=1}^{t-1}\|\mathbf{G}_{s}\|_{\mathbf{V}_{t}^{-1}}^{2}}\right)^{2} + \frac{\lambda\kappa^{-2}d_{w}C_{f}^{2}\zeta}{\mu^{n}}\right) \\ &\leq C_{0}\kappa^{-1}\left(\alpha_{t} + \frac{C_{h}^{2}\widetilde{C}^{2}\kappa^{-4}d_{w}^{2}C_{f}^{4}\zeta^{2}t}{\mu^{2}n^{2}}\left(\sum_{s=1}^{t-1}\|\mathbf{G}_{s}\|_{\mathbf{V}_{t}^{-1}}^{2}\right) + \frac{\lambda\kappa^{-2}d_{w}C_{f}^{2}\zeta}{\mu^{n}}\right) \\ &\leq C_{0}\kappa^{-1}\left(\alpha_{t} + \frac{C_{h}^{2}\widetilde{C}^{2}\kappa^{-4}d_{w}^{3}C_{f}^{4}\zeta^{2}t}{\mu^{2}n^{2}} + \frac{\lambda\kappa^{-2}d_{w}C_{f}^{2}\zeta}{\mu^{n}}\right) \\ &= \widetilde{\mathcal{O}}\left(\mu^{-2}\kappa^{-4}d_{w}\right)\,, \end{split}$$

where, in the last inequality, we substitute $\alpha_t = \widetilde{\mathcal{O}}(d_w)$, $\lambda = C_\lambda \sqrt{T} = \widetilde{\mathcal{O}}(\kappa^{-5/2}\mu^{-1}d_w\sqrt{T})$, $n = d_w \kappa^{-3/2} \sqrt{T}$.

C.1 Proofs of Technical Lemmas

C.1.1 Proof of Lemma 5

Proof of Lemma 5. If we denote $\mathbf{s}_t = \sum_{s=1}^{t-1} \sum_{i \in S_s} (y_{si} - p_{si}(\widetilde{\mathbf{w}}^*)) \nabla f_{\hat{\mathbf{w}}_s}(\mathbf{x}_{si})$, then,

$$\mathbf{s}_{t+1} = \mathbf{s}_t + \sum_{i \in S_t} (y_{ti} - p_{ti}(\widetilde{\mathbf{w}}^*)) \nabla f_{\widehat{\mathbf{w}}_t}(\mathbf{x}_{ti}) = \mathbf{s}_t + \phi_t,$$

where we denote $\phi_t := \sum_{i \in S_t} (y_{ti} - p_{ti}(\widetilde{\mathbf{w}}^*)) \nabla f_{\widehat{\mathbf{w}}_t}(\mathbf{x}_{ti})$. Then, we have

$$\begin{aligned} \|\mathbf{s}_{t+1}\|_{\mathbf{V}_{t+1}^{-1}}^{2} &= (\mathbf{s}_{t+1} + \boldsymbol{\phi}_{t})^{\top} \mathbf{V}_{t+1}^{-1} (\mathbf{s}_{t+1} + \boldsymbol{\phi}_{t}) \\ &= \mathbf{s}_{t+1}^{\top} \mathbf{V}_{t+1}^{-1} \mathbf{s}_{t+1} + 2 \mathbf{s}_{t+1}^{\top} \mathbf{V}_{t+1}^{-1} \boldsymbol{\phi}_{t} + \boldsymbol{\phi}_{t}^{\top} \mathbf{V}_{t+1}^{-1} \boldsymbol{\phi}_{t} \\ &\leq \mathbf{s}_{t+1}^{\top} \mathbf{V}_{t}^{-1} \mathbf{s}_{t+1} + 2 \mathbf{s}_{t+1}^{\top} \mathbf{V}_{t+1}^{-1} \boldsymbol{\phi}_{t} + \boldsymbol{\phi}_{t}^{\top} \mathbf{V}_{t}^{-1} \boldsymbol{\phi}_{t}, \end{aligned}$$

where the last inequality follows by $V_t \leq V_{t+1}$. Using this argument recursively, we have

$$\|\mathbf{s}\|_{\mathbf{V}_{t}^{-1}}^{2} \leq \sum_{s=1}^{t-1} \boldsymbol{\phi}_{s}^{\top} \mathbf{V}_{s}^{-1} \boldsymbol{\phi}_{s} + \sum_{s=1}^{t-1} 2 \boldsymbol{\phi}_{s} \mathbf{V}_{s+1}^{-1} \mathbf{s}_{s}.$$
 (11)

Now we will define a martingale difference sequence. We start to define an event $E_t = \mathbb{1}\left\{\|\mathbf{s}_s\|_{\mathbf{V}_s^{-1}}^2 \leq \beta_s, \forall s \leq t\right\}$, and define

$$J_t = 2E_t \boldsymbol{\phi}_t^{\top} \mathbf{V}_{t+1}^{-1} \mathbf{s}_t .$$

Then, we have $\mathbb{E}[J_t \mid \mathcal{F}_t] = 2\mathbb{E}[\phi_t^\top \mid \mathcal{F}_t]\mathbf{V}_{t+1}^{-1}\mathbf{s}_t = 0$. And since we have

$$|J_{t}| \leq 2E_{t}|\boldsymbol{\phi}_{t}^{\top}\mathbf{V}_{t+1}^{-1}\mathbf{s}_{t}|$$

$$\leq 2E_{t}||\boldsymbol{\phi}_{t}^{\top}\mathbf{V}_{t+1}^{-1/2}||_{2}||\mathbf{V}_{t+1}^{-1/2}\mathbf{s}_{t}||_{2}$$

$$= 2E_{t}\sqrt{\boldsymbol{\phi}_{t}^{\top}\mathbf{V}_{t+1}^{-1}\boldsymbol{\phi}_{t}}\sqrt{\mathbf{s}_{t}^{\top}\mathbf{V}_{t+1}^{-1}\mathbf{s}_{t}}$$

$$\leq 2E_{t}\sqrt{\boldsymbol{\phi}_{t}^{\top}\mathbf{V}_{t}^{-1}\boldsymbol{\phi}_{t}}\sqrt{\mathbf{s}_{t}^{\top}\mathbf{V}_{t}^{-1}\mathbf{s}_{t}}$$

$$\leq 2E_{t}\sqrt{\alpha_{t}}||\boldsymbol{\phi}_{t}||_{\mathbf{V}_{t}^{-1}}$$

where the last inequality follows trivially when $E_t = 0$, and by definition of E_t when $E_t = 1$. Additionally this gives us a family of uniform upper bound:

$$|J_s| \le 2\sqrt{\alpha_t}, \quad \forall s \le t.$$

Now we bound the conditional variance of J_t as follows:

$$V_{t} := \sum_{s=1}^{t-1} \operatorname{Var}(J_{s} \mid J_{0}, \dots, J_{s-1}) \leq \sum_{s=1}^{t-1} 4E_{s} \alpha_{s} \|\boldsymbol{\phi}_{s}\|_{\mathbf{V}_{s}^{-1}}^{2}$$

$$\leq 4(\max_{s \leq t} \alpha_{s}) \sum_{s=1}^{t-1} \|\boldsymbol{\phi}_{s}\|_{\mathbf{V}_{s}^{-1}}^{2}$$

$$\leq 16\alpha_{t} \sum_{s=1}^{t-1} \max_{i \in S_{s}} \|\nabla f_{\mathbf{w}_{s}}(\mathbf{x}_{si})\|_{\mathbf{V}_{s}^{-1}}^{2}$$

where the last inequality uses the following inequality

$$\begin{aligned} \|\boldsymbol{\phi}_{s}\|_{\mathbf{V}_{s}^{-1}} &\leq \sum_{i \in S_{s}} |\varepsilon_{si}| \|\nabla f_{\mathbf{w}_{s}}(\mathbf{x}_{si})\|_{\mathbf{V}_{s}^{-1}} \\ &\leq \left(\sum_{i \in S_{s}} |y_{si}| \|\nabla f_{\mathbf{w}_{s}}(\mathbf{x}_{si})\|_{\mathbf{V}_{s}^{-1}} + \sum_{i \in S_{s}} |p_{si}(\widetilde{\mathbf{w}}^{*})| \|\nabla f_{\mathbf{w}_{s}}(\mathbf{x}_{si})\|_{\mathbf{V}_{s}^{-1}} \right) \\ &\leq 2 \max_{i \in S_{s}} \|\nabla f_{\mathbf{w}_{s}}(\mathbf{x}_{si})\|_{\mathbf{V}_{s}^{-1}}. \end{aligned}$$

For the next step, we introduce the following lemma:

Lemma 8. For any $t \ge 1$, it holds that

$$\sum_{s=1}^{t-1} \min \left\{ 1, \max_{i \in S_s} \left\| \nabla f_{\hat{\mathbf{w}}_s}(\mathbf{x}_{si}) \right\|_{\mathbf{V}_s^{-1}}^2 \right\} \le 2d_w \log \left(1 + \frac{tC_g^2}{d_w \lambda} \right).$$

Lemma 8 is conceptually similar to the elliptical potential lemma commonly used in the linear utility setting [28, 29, 20]. As a result, the proof of Lemma 8 can be readily extended from its linear counterpart. For completeness, we provide the proof of Lemma 8 at the end of this section. Suppose that $\lambda \geq C_g^2$, which guarantees that $\|\nabla f_{\hat{\mathbf{w}}_t}(\mathbf{x}_{ti})\|_{\mathbf{V}_t^{-1}}^2 \leq 1$ for all $t \in [T]$ and $i \in [N]$. By using Lemma 8, we bound V_t as follows:

$$V_t \le 32\alpha_t d_w \log(1 + TC_g^2/(d_w \lambda)) =: v_t.$$

Now we have established the sum of conditional variance of martingale difference sequence, denoted by V_t , is bounded by $v_t = \widetilde{\mathcal{O}}(\alpha_t d_w)$. Then, by Freedman's inequality (Lemma 10) with parameters

 $a = \alpha_t/2, b = 2\sqrt{\alpha_t}$, and $v = v_t$, we have

$$\begin{split} \mathbb{P}\left(\sum_{s=1}^{t-1} J_s \geq \alpha_t/2\right) &= \mathbb{P}\left(\sum_{s=1}^{t-1} J_s \geq \alpha_t/2, V_t \leq v_t\right) \\ &\leq \exp\left(\frac{-(\alpha_t/2)^2}{2v_t + 2(\alpha_t/2)(2\sqrt{\alpha_t})/3}\right) \\ &\leq \exp\left(\max\left\{-\frac{\alpha_t^2}{16v_t}\right\}, \left\{-\frac{3\sqrt{\alpha_t}}{16}\right\}\right) \\ &\leq \frac{\delta}{t^2\pi^2} \end{split}$$

where the last inequality holds if we set

$$\alpha_t \ge C_\beta d_w \log(1 + TC_q^2/(d_w \lambda)) \log(t^2 \pi^2/\delta) \tag{12}$$

for some absolute constant $C_{\beta} > 0$ satisfying that

$$\alpha_t^2 \ge 16v_t \log(t^2 \pi^2 / \delta) + (16/3)^2 \log^2(t^2 \pi^2 / \delta).$$

By taking union bound,

$$\mathbb{P}\left(\exists t, \quad \sum_{s=1}^{t-1} J_s \ge \frac{\alpha_t}{2}\right) \le \sum_{t=1}^{\infty} \mathbb{P}\left(\sum_{s=1}^{t-1} J_s \ge \frac{\alpha_t}{2}\right) \\
\le \sum_{t=1}^{\infty} \mathbb{P}\left(\sum_{s=1}^{t-1} J_s \ge \frac{\alpha_t}{2}\right) \\
\le \sum_{t=1}^{\infty} \frac{\delta}{t^2 \pi^2} = \frac{\pi^2}{\delta} \cdot \frac{\pi^2}{6} = \delta.$$

Until now, we have established the event $\{\forall t \geq 1, \sum_{s=1}^{t-1} J_s \leq \alpha_t/2\}$ holds with probability at least $1-\delta$. Now we will show that the event $E_t=1$ under the event $\{\forall t \geq 1, \sum_{s=1}^{t-1} J_s \leq \alpha_t/2\}$. We prove this by induction on t. Recall that $E_t=1$ $\{\|\mathbf{s}_s\|_{\mathbf{V}_s^{-1}}^2 \leq \alpha_s, \forall s \leq t\}$. For the base case (t=1), by definition of $\mathbf{s}_1=\mathbf{0}_{d_w}$, it holds that $\|\mathbf{s}_1\|_{\mathbf{V}_1^{-1}}^2 \leq \alpha_1$. For the induction step, suppose that $E_s=1$ for s < t. By Eq. (11),

$$\|\mathbf{s}\|_{\mathbf{V}_{t}^{-1}}^{2} \leq \sum_{s=1}^{t-1} \boldsymbol{\phi}_{s}^{\top} \mathbf{V}_{s}^{-1} \boldsymbol{\phi}_{s} + \sum_{s=1}^{t-1} 2 \boldsymbol{\phi}_{s} \mathbf{V}_{s+1}^{-1} \mathbf{s}_{s}$$

$$\leq \sum_{s=1}^{t-1} \min\{1, \|\boldsymbol{\phi}_{s}\|_{\mathbf{V}_{s}^{-1}}^{2}\} + \frac{\alpha_{t}}{2}$$

$$\leq 2d_{w} \log(1 + TC_{g}^{2}/(d_{w}\lambda)) + \frac{\alpha_{t}}{2},$$

where we invoke Lemma 8 again for the second inequality. Then, with Eq. (12), if we have

$$\alpha_t = \widetilde{C}_{\alpha} d_w \log(1 + TC_a^2/(d_w \lambda)) \log(t^2 \pi^2/\delta)$$

for sufficiently large constant $\widetilde{C}_{\alpha} > 0$, then we have $\|\mathbf{s}_t\|_{\mathbf{V}_{\tau}^{-1}}^2 \leq \alpha_t$. This completes the proof. \square

C.1.2 Proof of Lemma 6

Proof of Lemma 6. For any $\mathbf{w}_1, \mathbf{w}_2 \in \mathcal{W}$, by the mean value theorem, we have

$$\mathbf{g}_t(\mathbf{w}_1) - \mathbf{g}_t(\mathbf{w}_2) = \nabla \ell_t(\mathbf{w}_2) - \nabla \ell_t(\mathbf{w}_1) = \left(\int_0^1 \nabla^2 \ell_t(\mathbf{w}_1 + z(\mathbf{w}_2 - \mathbf{w}_1)) dx\right) (\mathbf{w}_2 - \mathbf{w}_1).$$

This implies

$$\|\mathbf{g}_{t}(\mathbf{w}_{1}) - \mathbf{g}_{t}(\mathbf{w}_{2})\|_{\mathbf{G}_{t}^{-1}(\mathbf{w}_{1}, \mathbf{w}_{2})} = \|\mathbf{w}_{1} - \mathbf{w}_{2}\|_{\mathbf{G}_{t}(\mathbf{w}_{1}, \mathbf{w}_{2})},$$
 (13)

where we denote $G_t(\mathbf{w}_1, \mathbf{w}_2) := \int_0^1 \nabla^2 \ell_t(\mathbf{w}_1 + z(\mathbf{w}_2 - \mathbf{w}_1)) dz$. On the other hand, Lee and Oh [20] show that the multinomial logistic loss with linear utility is $3\sqrt{2}$ -self-concordant-like (Lemma 11), then by the property of self-concordant-like function (Lemma 12), we have

$$\nabla^2 \ell_t(\mathbf{w}_1 - z(\mathbf{w}_2 - \mathbf{w}_1)) \succeq \exp\left(-3\sqrt{2}\|z(\mathbf{w}_2 - \mathbf{w}_1)\|_2\right) \nabla^2 \ell_t(\mathbf{w}_1).$$

This implies

$$\mathbf{G}_{t}(\mathbf{w}_{1}, \mathbf{w}_{2}) \succeq \int_{0}^{1} \exp\left(-3\sqrt{2}\|z(\mathbf{w}_{2} - \mathbf{w}_{1})\|_{2}\right) dz \nabla^{2} \ell_{t}(\mathbf{w}_{1})$$

$$\succeq \frac{1 - \exp\left(\|\mathbf{w}_{2} - \mathbf{w}_{1}\|_{2}\right)}{3\sqrt{2}\|\mathbf{w}_{2} - \mathbf{w}_{1}\|_{2}} \nabla^{2} \ell_{t}(\mathbf{w}_{1})$$

$$\succeq \frac{1}{1 + 3\sqrt{2}\|\mathbf{w}_{2} - \mathbf{w}_{1}\|_{2}} \nabla^{2} \ell_{t}(\mathbf{w}_{1})$$

$$\succeq \frac{1}{1 + 6\sqrt{2}} \nabla^{2} \ell_{t}(\mathbf{w}_{1}),$$

where the third inequality uses the inequality $\frac{1-\exp(-x)}{x} \ge \frac{1}{1+x}$ for $x \ge 0$. By applying this to Eq. (13), we obtain

$$\|\mathbf{w}_{1} - \mathbf{w}_{2}\|_{\nabla^{2}\ell_{t}(\mathbf{w}_{1})}^{2} \leq (1 + 6\sqrt{2})\|\mathbf{w}_{1} - \mathbf{w}_{2}\|_{\mathbf{G}_{t}(\mathbf{w}_{1}, \mathbf{w}_{2})}^{2}$$

$$\leq (1 + 6\sqrt{2})\|\mathbf{g}_{t}(\mathbf{w}_{1}) - \mathbf{g}_{t}(\mathbf{w}_{2})\|_{\mathbf{G}_{t}^{-1}(\mathbf{w}_{1}, \mathbf{w}_{2})}^{2}$$

$$\leq (1 + 6\sqrt{2})^{2}\|\mathbf{g}_{t}(\mathbf{w}_{1}) - \mathbf{g}_{t}(\mathbf{w}_{2})\|_{(\nabla^{2}\ell_{t}(\mathbf{w}_{1}))^{-1}}^{2}.$$
(14)

On the other hand, since

$$\nabla^{2} \ell_{t}(\mathbf{w}) := \lambda \mathbf{I}_{d_{w}} + \sum_{s=1}^{t-1} \left[\sum_{i \in S_{s}} \hat{p}_{si}(\mathbf{w}) \nabla f_{si} \nabla f_{si}^{\top} - \left(\sum_{i \in S_{s}} \hat{p}_{si}(\mathbf{w}) \nabla f_{si} \right) \left(\sum_{i \in S_{s}} \hat{p}_{si}(\mathbf{w}) \nabla f_{si} \right)^{\top} \right]$$

$$\succeq \lambda \mathbf{I}_{d_{w}} + \kappa \sum_{s=1}^{t-1} \sum_{i \in S_{s}} \nabla f_{si} \nabla f_{si}^{\top} \succeq \kappa \mathbf{V}_{t},$$

where we abbreviate $\nabla f_{\hat{\mathbf{w}}_s}(\mathbf{x}_{si}) =: \nabla f_{si}$. Then, by applying this to Eq. (14) we have

$$\|\mathbf{w}_1 - \mathbf{w}_2\|_{\mathbf{V}_t}^2 \le (1 + 6\sqrt{2})^2 \kappa^{-1} \|\mathbf{g}_t(\mathbf{w}_1) - \mathbf{g}_t(\mathbf{w}_2)\|_{\mathbf{V}_t^{-1}}^2$$
.

Substituting $\mathbf{w}_1 = \hat{\mathbf{w}}_t$ and $\mathbf{w}_2 = \widetilde{\mathbf{w}}^*$ completes the proof.

C.1.3 Proof of Lemma 7

Proof of Lemma 7. Note that

$$\sum_{s=1}^{t-1} \|\mathbf{G}_s\|_{\mathbf{V}_t^{-1}}^2 = \sum_{s=1}^{t-1} \lambda_{\max}(\mathbf{G}_s^{\top} \mathbf{V}_t^{-1} \mathbf{G}_s)$$

$$\leq \sum_{s=1}^{t-1} \operatorname{tr}(\mathbf{G}_s^{\top} \mathbf{V}_t^{-1} \mathbf{G}_s)$$

$$= \sum_{s=1}^{t-1} \sum_{i \in S_s} \nabla f_{\hat{\mathbf{w}}_s}(\mathbf{x}_{si})^{\top} \mathbf{V}_t^{-1} \nabla f_{\hat{\mathbf{w}}_s}(\mathbf{x}_{si})$$

$$= \sum_{s=1}^{t-1} \sum_{i \in S_s} \operatorname{tr} \left(\nabla f_{\hat{\mathbf{w}}_s}(\mathbf{x}_{si})^{\top} \mathbf{V}_t^{-1} \nabla f_{\hat{\mathbf{w}}_s}(\mathbf{x}_{si}) \right)$$

$$= \operatorname{tr} \left(\mathbf{V}_t^{-1} \sum_{s=1}^{t-1} \sum_{i \in S_s} \nabla f_{\hat{\mathbf{w}}_s}(\mathbf{x}_{si}) \nabla f_{\hat{\mathbf{w}}_s}(\mathbf{x}_{si})^{\top} \right).$$

If we denote $\lambda_1, \dots, \lambda_{d_w}$ the eigenvalues of $\sum_{s=1}^{t-1} \sum_{i \in S_s} \nabla f_{\hat{\mathbf{w}}_s}(\mathbf{x}_{si}) \nabla f_{\hat{\mathbf{w}}_s}(\mathbf{x}_{si})^{\top}$, then we have

$$\operatorname{tr}\left(\mathbf{V}_{t}^{-1} \sum_{s=1}^{t-1} \sum_{i \in S_{s}} \nabla f_{\hat{\mathbf{w}}_{s}}(\mathbf{x}_{si}) \nabla f_{\hat{\mathbf{w}}_{s}}(\mathbf{x}_{si})^{\top}\right) = \sum_{j=1}^{d_{w}} \lambda_{j} / (\lambda_{j} + \lambda) \leq d_{w}.$$

C.1.4 Proof of Lemma 8

Proof of Lemma 8. For notational simplicity, we abbreviate $\nabla f_{\hat{\mathbf{w}}_t}(\mathbf{x}_{ti})$ as ∇f_{ti} . Note that

$$\mathbf{V}_{t+1} = \mathbf{V}_t + \sum_{i \in S_t} \nabla f_{ti} \nabla f_{ti}^{\top} ,$$

which results in

$$\det \mathbf{V}_{t+1} = \det \mathbf{V}_{t} \left(1 + \sum_{i \in S_{t}} \|\nabla f_{ti}\|_{\mathbf{V}_{t}^{-1}}^{2} \right) \ge \det \mathbf{V}_{t} \left(1 + \max_{i \in S_{t}} \|\nabla f_{ti}\|_{\mathbf{V}_{t}^{-1}}^{2} \right)$$

$$\ge \det(\lambda \mathbf{I}_{d_{w}}) \prod_{i = 1}^{t} \left(1 + \max_{i \in S_{s}} \|\nabla f_{si}\|_{\mathbf{V}_{s}^{-1}}^{2} \right).$$

Then, we have

$$\sum_{s=1}^{t} \left(1 + \max_{i \in S_s} \|\nabla f_{si}\|_{\mathbf{V}_s^{-1}}^2 \right) \le \log \frac{\det \mathbf{V}_{t+1}}{\det(\lambda \mathbf{I}_{d_w})}.$$

Using the inequality $z \le 2\log(1+z)$ for $z \in [0,1]$, we have

$$\sum_{s=1}^{t} \min \left\{ 1, \max_{i \in S_s} \|\nabla f_{si}\|_{\mathbf{V}_s^{-1}}^2 \right\} \leq \sum_{s=1}^{t} 2 \log \left(1 + \max_{i \in S_s} \|\nabla f_{si}\|_{\mathbf{V}_s^{-1}}^2 \right)
\leq 2 \log \frac{\det \mathbf{V}_{t+1}}{\det(\lambda \mathbf{I}_{d_w})}
\leq 2 d_w \log \left(1 + \frac{(t+1)C_g^2}{d_w \lambda} \right),$$

where the last inequality follows by the determinant-trace inequality (Lemma 13).

D Proof of Lemma 3

Proof of Lemma 3. Note that

$$z_{ti} - f_{\mathbf{w}^*}(\mathbf{x}_{ti}) = f_{\hat{\mathbf{w}}_t}(\mathbf{x}_{ti}) - f_{\mathbf{w}^*}(\mathbf{x}_{ti}) + \sqrt{\beta_t} \|\nabla f_{\hat{\mathbf{w}}_t}(\mathbf{x}_{ti})\|_{\mathbf{V}_t^{-1}} + \frac{\beta_t C_h}{\lambda}$$

$$= (\hat{\mathbf{w}}_t - \mathbf{w}^*)^\top \nabla f_{\hat{\mathbf{w}}}(\mathbf{x}_{ti}) + \sqrt{\beta_t} \|\nabla f_{\hat{\mathbf{w}}_t}(\mathbf{x}_{ti})\|_{\mathbf{V}_t^{-1}} + \frac{\beta_t C_h}{\lambda}$$

$$= (\hat{\mathbf{w}}_t - \mathbf{w}^*)^\top \nabla f_{\hat{\mathbf{w}}_t}(\mathbf{x}_{ti}) + (\hat{\mathbf{w}}_t - \mathbf{w}^*)^\top (\nabla f_{\hat{\mathbf{w}}}(\mathbf{x}_{ti}) - \nabla f_{\hat{\mathbf{w}}_t}(\mathbf{x}_{ti}))$$

$$+ \sqrt{\beta_t} \|\nabla f_{\hat{\mathbf{w}}_t}(\mathbf{x}_{ti})\|_{\mathbf{V}_t^{-1}} + \frac{\beta_t C_h}{\lambda}$$

$$\geq (\hat{\mathbf{w}}_t - \mathbf{w}^*)^\top (\nabla f_{\hat{\mathbf{w}}}(\mathbf{x}_{ti}) - \nabla f_{\hat{\mathbf{w}}_t}(\mathbf{x}_{ti})) + \frac{\beta_t C_h}{\lambda}$$

$$= (\hat{\mathbf{w}}_t - \mathbf{w}^*)^\top \nabla^2 f_{\hat{\mathbf{w}}}(\mathbf{x}_{ti}) (\dot{\mathbf{w}} - \hat{\mathbf{w}}_t) + \frac{\beta_t C_h}{\lambda}$$

$$\geq -\|\hat{\mathbf{w}}_t - \mathbf{w}^*\|_{\mathbf{V}_t} \|\mathbf{V}_t^{-1/2} \nabla^2 f_{\hat{\mathbf{w}}}(\mathbf{x}_{ti}) \mathbf{V}_t^{-1/2} \|\|\dot{\mathbf{w}} - \hat{\mathbf{w}}_t\|_{\mathbf{V}_t} + \frac{\beta_t C_h}{\lambda}$$

$$\geq 0,$$

$$(15)$$

where Eq. (15) applies the mean value theorem with $\dot{\mathbf{w}}$ lying between $\hat{\mathbf{w}}_t$ and \mathbf{w}^* , Eq. (16) uses it again for $\ddot{\mathbf{w}}$ lying between $\dot{\mathbf{w}}$ and $\hat{\mathbf{w}}_t$, and the last inequality follows from the bound $\|\dot{\mathbf{w}} - \hat{\mathbf{w}}_t\|_{\mathbf{V}_t} \leq \|\hat{\mathbf{w}}_t - \mathbf{w}^*\|_{\mathbf{V}_t} \leq \sqrt{\beta_t}$. The upper bound can be derived by the similar argument.

E Proof of Theorem 1

Proof of Theorem 1. Suppose that Lemma 2 holds. Then, by Lemma 3, for all $t \in [T]$ and $i \in S_t$, we have $f_{\mathbf{w}^*}(\mathbf{x}_{ti}) \leq z_{ti}$. By Lemma 14, the expected reward corresponding to the optimal assortment S_t^* is monotone in the MNL utilities, for all $t \in [T]$, we have

$$R_t(S_t^*, \mathbf{w}^*) \le \widetilde{R}_t(S_t^*) \le \widetilde{R}_t(S_t). \tag{17}$$

If we denote $z'_{ti} = f_{\mathbf{w}^*}(\mathbf{x}_{ti}) + 2\sqrt{\beta_t}\|\nabla f_{\hat{\mathbf{w}}_t}(\mathbf{x}_{ti})\|_{\mathbf{V}_t^{-1}} + 2\beta_t C_h/\lambda$, by Lemma 15, which states that the optimistic expected reward $\widetilde{R}_t(S)$ is monotonically increasing with respect to the utilities of the items in S, and by Lemma 3 we have $z_{ti} \leq z'_{ti}$, which implies

$$\widetilde{R}_t(S_t) = \sum_{i \in S_t} \frac{\exp(z_{ti}) r_{ti}}{1 + \sum_{j \in S_t} \exp(z_{tj})} \le \sum_{i \in S_t} \frac{\exp(z'_{ti}) r_{ti}}{1 + \sum_{j \in S_t} \exp(z'_{tj})}.$$
(18)

Let $\widetilde{\mathbf{w}}^* := \operatorname{argmin}_{\widetilde{\mathbf{w}} \in \mathcal{W}^*} \|\hat{\mathbf{w}}_0 - \widetilde{\mathbf{w}}\|_2$ denote the parameter in the equivalence set \mathcal{W}^* that is closest to $\hat{\mathbf{w}}_0$. By combining Eq. 17 and Eq. 18, we bound the cumulative regret in Phase II as follows:

$$\sum_{t=1}^{T} R_{t}(S_{t}, \mathbf{w}^{*}) - R_{t}(S_{t}, \mathbf{w}^{*}) = \sum_{t=1}^{T} R_{t}(S_{t}, \widetilde{\mathbf{w}}^{*}) - R_{t}(S_{t}, \widetilde{\mathbf{w}}^{*})$$

$$\leq \sum_{t=1}^{T} \widetilde{R}_{t}(S_{t}) - R_{t}(S_{t}, \widetilde{\mathbf{w}}^{*})$$

$$\leq \sum_{t=1}^{T} \left(\frac{\sum_{i \in S_{t}} \exp(z'_{ti}) r_{ti}}{1 + \sum_{j \in S_{t}} \exp(z'_{tj})} - \frac{\sum_{i \in S_{t}} \exp(f_{\widetilde{\mathbf{w}}^{*}}(\mathbf{x}_{ti})) r_{ti}}{1 + \sum_{j \in S_{t}} \exp(f_{\widetilde{\mathbf{w}}^{*}}(\mathbf{x}_{tj}))} \right)$$

$$\leq \sum_{t=1}^{T} \max_{i \in S_{t}} |z'_{ti} - f_{\widetilde{\mathbf{w}}^{*}}(\mathbf{x}_{ti})|$$

$$= \sum_{t=1}^{T} 2\sqrt{\beta_{t}} \|\nabla f_{\widehat{\mathbf{w}}_{t}}(\mathbf{x}_{ti})\|_{\mathbf{V}_{t}^{-1}} + \frac{2\beta_{t}C_{h}}{\lambda},$$
(19)

where Eq. (19) follows by Lemma 16. By applying Cauchy-Schwarz inequality, we have

$$\sum_{t=1}^{T} \left(2\sqrt{\beta_{t}} \|\nabla f_{\hat{\mathbf{w}}_{t}}(\mathbf{x}_{ti})\|_{\mathbf{V}_{t}^{-1}} + \frac{2\beta_{t}C_{h}}{\lambda} \right) \leq \sqrt{T \sum_{t=1}^{T} \left(2\sqrt{\beta_{t}} \|\nabla f_{\hat{\mathbf{w}}_{t}}(\mathbf{x}_{ti})\|_{\mathbf{V}_{t}^{-1}} + \frac{2\beta_{t}C_{h}}{\lambda} \right)^{2}} \\
\leq \sqrt{T \sum_{t=1}^{T} \left(8\beta_{t} \|\nabla f_{\hat{\mathbf{w}}_{t}}(\mathbf{x}_{ti})\|_{\mathbf{V}_{t}^{-1}}^{2} + \frac{8\beta_{t}^{2}C_{h}^{2}}{\lambda^{2}} \right)} \\
\leq \sqrt{8\beta_{T}T \sum_{t=1}^{T} \|\nabla f_{\hat{\mathbf{w}}_{t}}(\mathbf{x}_{ti})\|_{\mathbf{V}_{t}^{-1}}^{2} + \frac{8\beta_{t}^{2}C_{h}^{2}T^{2}}{\lambda^{2}}} \\
\leq \sqrt{16T\beta_{T}d_{w} \log(1 + (T+1)C_{g}^{2}/(d_{w}\lambda)) + \frac{8\beta_{t}^{2}C_{h}^{2}T^{2}}{\lambda^{2}}} \\
= \widetilde{\mathcal{O}}\left(\kappa^{-2}\mu^{-1}d_{w}\sqrt{T}\right),$$

where the second inequality uses the inequality $(a+b)^2 \leq 2a^2 + 2b^2$, the third inequality uses the monotonicity of β_t , and the last inequality follows by Lemma 8, and the final bound is obtained by setting $\beta_T = \widetilde{\mathcal{O}}(\kappa^{-4}d_w\mu^{-2}), \lambda = \widetilde{\mathcal{O}}(\kappa^{-5/2}\mu^{-1}d_w\sqrt{T})$. Then, the cumulative regret incurred over Phases I and II is bounded by

$$\operatorname{Regret}_{n+T} = \widetilde{\mathcal{O}}\left(\kappa^{-3/2} d_w \sqrt{T} + \kappa^{-2} \mu^{-1} d_w \sqrt{T}\right).$$

F Auxiliary Lemmas

Lemma 9 (Lemma D.1 in Zhang and Luo [41]). Let \mathcal{X} be an instance space, \mathcal{Y} a target space, and $p(y \mid x) = f^*(x,y)$ be the true conditional density for $(x,y) \in \mathcal{X} \times \mathcal{Y}$, where $f^* \in \mathcal{F}$. Assume $\mathcal{F} = \{f_{\boldsymbol{\theta}}: \mathcal{X} \times \mathcal{Y} \to [\beta,1] \mid \boldsymbol{\theta} \in [0,1]^d, \beta > 0\}$ is a 1-Lipschitz function class in the parameter $\boldsymbol{\theta} \in [0,1]^d$, i.e., $\|f_{\boldsymbol{\theta}_1}(x,\cdot) - f_{\boldsymbol{\theta}_2}(x,\cdot)\|_{\infty} \leq \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_{\infty}$ for all $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in [0,1]^d$ and $x \in \mathcal{X}$. Let $D = \{(x_i,y_i)\}_{i=1}^n$ be i.i.d. samples with $x_i \sim \mathcal{H}$ (unknown) and $y_i \sim p(\cdot \mid x_i)$, and define the empirical risk minimizer $\hat{f} = \operatorname{argmin}_{f \in \mathcal{F}} \sum_{i=1}^n -\log f(x_i,y_i)$. Then for any $\delta \in (0,1)$, with probability at least $1 - \delta$,

$$\mathbb{E}_{x \sim \mathcal{H}, y \sim p(\cdot|x)}[\log \hat{f}(x, y) - \log f^*(x, y)] \leq \mathcal{O}\left(\frac{d \log \frac{|\mathcal{Y}|}{\beta} \log \frac{n}{\beta} \log \frac{1}{\delta}}{n}\right).$$

Lemma 10 (Freedman's inequality [14]). Suppose X_1, \ldots, X_T is a martingale difference sequence with $|X_t| \leq b$ for all $t \in [T]$. Let V denote the sum of conditional variances,

$$V = \sum_{t=1}^{T} \text{Var}(X_i \mid X_1, \dots, X_{i-1}).$$

Then, for every a, v > 0,

$$\mathbb{P}\left(\sum_{i=1}^{T} X_i \ge a, V \le v\right) \le \exp\left(\frac{-a^2}{2v + 2ab/3}\right).$$

Lemma 11 (Proposition C.1 in Lee and Oh [20]). Let $f_{\mathbf{w}}(\mathbf{x}) = \mathbf{x}^{\top}\mathbf{w}$. Then the loss $\ell_t(\mathbf{w})$ in Eq. (1) is $3\sqrt{2}$ -self-concordant-like.

Lemma 12 (Theorem 4 in Tran-Dinh et al. [36]). Let $f : \mathbb{R}^n \to \mathbb{R}$ be a M_f -self-concordant-like function and let $x, y \in \text{dom}(f)$, then it holds:

$$\exp(-M_f ||y - x||_2) \nabla^2 f(x) \leq \nabla^2 f(y).$$

Lemma 13 (Lemma 10 in Abbasi-Yadkori et al. [1]). Let $\mathbf{x}_1, \dots, \mathbf{x}_t \in \mathbb{R}^d$ with $\|\mathbf{x}_s\|_2 \leq L$ for any $1 \leq s \leq t$. Let $\bar{\mathbf{V}}_t = \lambda \mathbf{I}_d + \sum_{s=1}^t \mathbf{x}_s \mathbf{x}_s^{\top}$ for some $\lambda > 0$. Then,

$$\det \bar{\mathbf{V}}_t \le (\lambda + tL^2/d)^d.$$

Lemma 14 (Lemma A.3 in Agrawal et al. [3]). Let $\mathbf{v} = [v_1, \dots, v_N]^{\top} \in \mathbb{R}_+^N$ be a utility vector and $\mathbf{r} = [r_1, \dots, r_N]^{\top} \in \mathbb{R}_+^N$ a revenue vector. The expected revenue of an assortment $S \subset [N]$ under the utility vector \mathbf{v} is defined as

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

Let $\mathbf{u} = [u_1, \dots, u_N]^{\top} \in \mathbb{R}_+^N$ be another utility vector such that $0 \le v_i \le u_i$ for all $i \in [N]$. Then the following holds:

$$R(S_{\mathbf{v}}, \mathbf{v}) \le R(S_{\mathbf{v}}, \mathbf{u}) \le R(S_{\mathbf{u}}, \mathbf{u}),$$

where $S_{\mathbf{v}} := \operatorname{argmax}_{S \subset [N]} R(S, \mathbf{v})$ and $S_{\mathbf{u}} := \operatorname{argmax}_{S \subset [N]} R(S, \mathbf{u})$.

Lemma 15 (Lemma H.2 in Lee and Oh [20]). Let $\mathbf{r} = [r_1, \dots, r_N] \in \mathbb{R}^N$ with $0 \le r_i \le 1, \forall i \in [N]$ be given. For $\mathbf{v} = [v_1, \dots, v_N] \in \mathbb{R}^N$ and $S \subset [N]$, let $R(S, \mathbf{v}) = \frac{\sum_{i \in S} \exp(v_i) r_i}{1 + \sum_{j \in S} \exp(v_j)}$. Then for any $\mathbf{u} = [u_1, \dots, u_N]^\top \in \mathbb{R}^N$ with $u_i \ge v_i$ for all $i \in [N]$, the following holds:

$$R(S_{\mathbf{v}}, \mathbf{v}) \leq R(S_{\mathbf{v}}, \mathbf{u})$$
,

where $S_{\mathbf{v}} = \operatorname{argmax}_{S \subset [N]} R(S, \mathbf{v})$.

Lemma 16 (Lemma 3 in Oh and Iyengar [28]). Let $\mathbf{u} = [u_1, \dots, u_d], \mathbf{v} = [v_1, \dots, v_d] \in \mathbb{R}^d$. Then for any $\mathbf{r} = [r_1, \dots, r_d] \in \mathbb{R}^d$ with $|r_i| \leq 1$,

$$\frac{\sum_{i=1}^d r_i \exp(u_i)}{1 + \sum_{j=1}^d \exp(u_j)} - \frac{\sum_{i=1}^d r_i \exp(v_i)}{1 + \sum_{j=1}^d \exp(v_j)} \le \max_{i \in [d]} |u_i - v_i|.$$

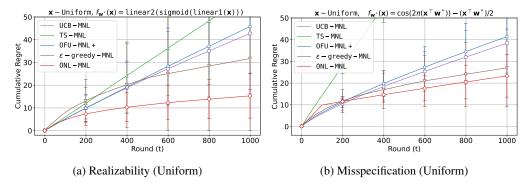


Figure 4: Cumulative regret comparison between ONL-MNL (ours) and baselines under realizable and misspecified settings with uniform context distributions.

G Details on Experiments & Additional Results

In this section, we provide additional details on the experimental settings and results discussed in the main text. All experiments are run on a computing cluster with Intel[®] Xeon[®] Gold 6526R (16-core, 2.8 GHz, 37.5 MB cache, 3 UPI, 195 W).

Additional experiments under varying parameters. We conduct experiments over a wide range of values for K and d. The experimental setup is identical to that described in Section 4. Figures 6 and 7 present additional evaluations of our proposed algorithm and baseline methods under Gaussian and uniform context distributions, respectively. The results provide strong evidence that our proposed algorithm consistently outperforms existing MNL contextual bandit algorithms.

Semi-synthetic experiment with real-world dataset. In this section, we provide an additional semi-synthetic experiment leveraging a real-world dataset. We used the IMDB Large Movie Review dataset [26], which consists of 50,000 movie reviews in text form, each labeled as either positive or negative. To evaluate online assortment selection algorithms, one needs access to ground-truth choice probabilities for given assortments, which are typically unobservable in real-world datasets. To address this, we first transformed the review texts into vector representations using TF-IDF (Term Frequency–Inverse Document Frequency), implemented via the 'TfidfVectorizer' from scikit-learn. We then applied truncated SVD to reduce the dimensionality of the vectors to d=30. This process yielded a dataset consisting of 50,000 context feature vectors with corresponding binary labels.

We then used 40,000 of these samples as training data to fit a binary classification model. Specifically, we trained a two-layered neural network with 32 hidden nodes to classify the binary labels. The learned model was subsequently used to define the true utility for each movie based on its extracted context feature, allowing us to formulate an online assortment selection task. We then evaluated both the MNL bandit baselines and our proposed algorithm under this setting.

At each round of the online experiment, we randomly sampled N=100 movies from the remaining 10,000 held-out samples, and asked the algorithm to choose an assortment of size K=5. As in our main experiments, we used a uniform revenue of 1 for each item in the assortment (since revenue information is not given in the dataset). The experimental results (cumulative regret over time step T=1000) are presented in Figure 5.

As shown in Figure 5, our proposed algorithm significantly outperforms the baseline methods. These results demonstrate that our proposed algorithm not only achieves superior performance on synthetic benchmarks but also generalizes effectively to semi-synthetic settings constructed from real-world data. The consistent performance advantage over existing baselines highlights the superior ability of our approach to learn complex, non-linear utility structures, demonstrating strong robustness and adaptability to realistic, high-dimensional feature representations. This experiment thus provides empirical evidence that our method can handle complex contextual structures beyond linear utility assumptions, supporting its potential applicability in real-world recommendation and decision-making environments.

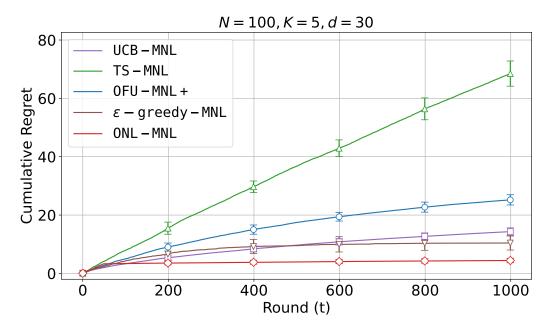


Figure 5: Cumulative regret comparison between ONL-MNL (ours) and baseline methods in the semi-synthetic experimental setting based on the IMDB Large Movie Review dataset [26].

Implementation of ε -greedy-MNL. We implement ε -greedy-MNL [41] based on the ε -greedy-style method described in Eq.(4) of Zhang and Luo [41]. Since the official code for ε -greedy-MNL is not publicly available, we tailor the implementation to our problem setting. To isolate and assess the effectiveness of the exploration strategy, we set the utility function class of ε -greedy-MNL to match the true utility function—a two-layer neural network with sigmoid activation. We adopt an epoch-based update schedule, where the first epoch starts with length 1, and each subsequent epoch doubles in length. That is, after time step $t=2^k-1$ ($k=1,2,3,\ldots$), we update the utility parameter for the next epoch using observations collected in the current epoch. We approximate the offline regression oracle using the Adam optimizer with a learning rate of 10^{-4} over 2000 iterations. The ε parameter starts at 0.1 and decays multiplicatively by a factor of 0.995 after each step, with a minimum threshold of 0.001.

Implementation of ONL-MNL. After Phase I, the pilot estimator $\hat{\mathbf{w}}_0$ is approximated by minimizing the negative log-likelihood using the Adam optimizer with a learning rate of 10^{-4} for 2000 iterations. The regularization parameter λ is set as $\lambda = c_\lambda \cdot \kappa^{-5/2} d_w \sqrt{T}$, and the confidence radius β_t is set as $\beta_t = c_\beta \cdot \kappa^{-4} d_w \frac{t}{T}$. Both scaling constants, c_λ and c_β , are selected via grid search.

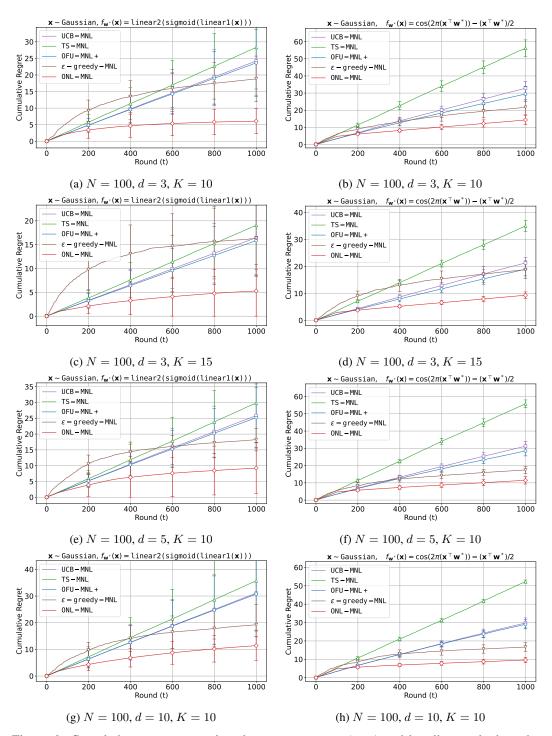


Figure 6: Cumulative regret comparison between ONL-MNL (ours) and baseline methods under Gaussian contexts with varying context feature dimensions $d \in \{3,5,10\}$ and assortment sizes $K \in \{10,15\}$. The left column shows results under the realizable setting, while the right column corresponds to the misspecified setting. Each curve represents the average cumulative regret over 10 independent runs.

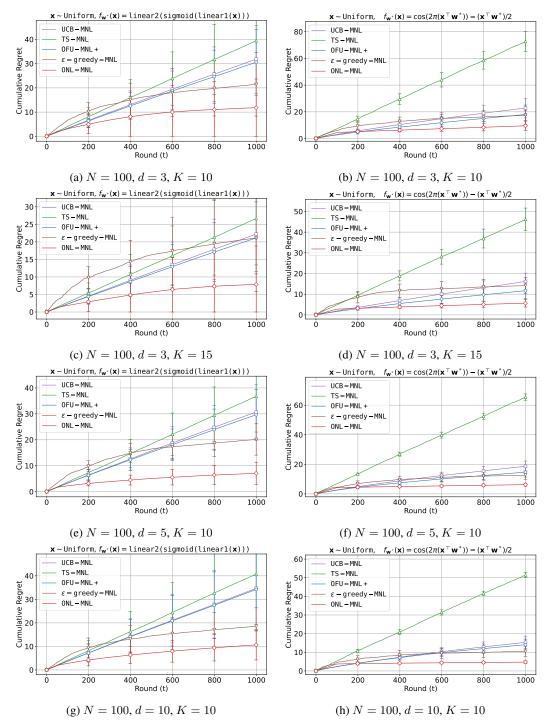


Figure 7: Cumulative regret comparison between ONL-MNL (ours) and baseline methods under uniform contexts with varying context feature dimensions $d \in \{3,5,10\}$ and assortment sizes $K \in \{10,15\}$. The left column shows results under the realizable setting, while the right column corresponds to the misspecified setting. Each curve represents the average cumulative regret over 10 independent runs.