
UCB-based Algorithms for Multinomial Logistic Regression Bandits

Sanae Amani
University of California, Los Angeles
samani@ucla.edu

Christos Thrampoulidis
University of British Columbia
cthrampo@ece.ubc.ca

Abstract

Out of the rich family of generalized linear bandits, perhaps the most well studied ones are logistic bandits that are used in problems with binary rewards: for instance, when the learner aims to maximize the profit over a user that can select one of two possible outcomes (e.g., ‘click’ vs ‘no-click’). Despite remarkable recent progress and improved algorithms for logistic bandits, existing works do not address practical situations where the number of outcomes that can be selected by the user is larger than two (e.g., ‘click’, ‘show me later’, ‘never show again’, ‘no click’). In this paper, we study such an extension. We use multinomial logit (MNL) to model the probability of each one of $K + 1 \geq 2$ possible outcomes (+1 stands for the ‘not click’ outcome): we assume that for a learner’s action \mathbf{x}_t , the user selects one of $K + 1 \geq 2$ outcomes, say outcome i , with a MNL probabilistic model with corresponding unknown parameter $\bar{\theta}_{*i}$. Each outcome i is also associated with a revenue parameter ρ_i and the goal is to maximize the expected revenue. For this problem, we present MNL-UCB, an upper confidence bound (UCB)-based algorithm, that achieves regret $\tilde{O}(dK\sqrt{T})$ with small dependency on problem-dependent constants that can otherwise be arbitrarily large and lead to loose regret bounds. We present numerical simulations that corroborate our theoretical results.

1 Introduction

Linear stochastic bandits provide simple, yet commonly encountered, models for a variety of sequential decision-making problems under uncertainty. Specifically, linear bandits generalize the classical multi-armed bandit (MAB) problem of K arms that each yields reward sampled independently from an underlying distribution with unknown parameters, to a setting where the expected reward of each arm is a linear function that depends on the same unknown parameter vector [Dani et al. \(2008\)](#); [Abbasi-Yadkori et al. \(2011\)](#); [Rusmevichientong and Tsitsiklis \(2010\)](#). Linear bandits have been successfully applied over the years in online advertising, recommendation services, resource allocation, etc. [Lattimore and Szepesvári \(2020\)](#). More recently, researchers have explored the potentials of such algorithms in more complex systems, such as in robotics, wireless networks, the power grid, medical trials, e.g., [Li et al. \(2013\)](#); [Avner and Mannor \(2019\)](#); [Berkenkamp et al. \(2016\)](#); [Sui et al. \(2018\)](#). However, linear bandits fail to model a host of other applications. This has called for extensions of linear bandits to a broader range of reward structures beyond linear models. One of the leading lines of work addressing these extensions relies on the Generalized Linear Model (GLM) framework of statistics. In GLMs the expected reward associated with an arm \mathbf{x} is given by $\mu(\bar{\theta}^T \mathbf{x})$, where $\bar{\theta} \in \mathbb{R}^d$ is the system unknown parameter and μ is a non-linear link function. Specifically, *logistic bandits*, that are appropriate for modeling binary reward structures, are a special case of generalized linear bandits (GLBs) with $\mu(x) = (1 + \exp(-x))^{-1}$. UCB-based algorithms for GLBs were first introduced in [Filippi et al. \(2010\)](#); [Li et al. \(2017\)](#); [Fauray et al. \(2020\)](#). The same problem, but with a Thompson Sampling- (TS) strategy was also studied in [Abeille et al. \(2017\)](#);

Russo and Van Roy (2013, 2014); Dong and Van Roy (2018). Beyond GLMs, an even more general framework for modeling reward is the semi-parametric index model (see for example Yang et al. (2017); Gamarnik and Gaudio (2020) for a list of applications in statistics). A semi-parametric index model relates the reward $y \in \mathbb{R}$ and the action/arm $\mathbf{x} \in \mathbb{R}^d$ as $y = \mu(\bar{\boldsymbol{\theta}}_1^T \mathbf{x}, \bar{\boldsymbol{\theta}}_2^T \mathbf{x}, \dots, \bar{\boldsymbol{\theta}}_K^T \mathbf{x}) + \epsilon$, where $\mu : \mathbb{R}^K \rightarrow \mathbb{R}$ and $\bar{\boldsymbol{\theta}}_1, \dots, \bar{\boldsymbol{\theta}}_K \in \mathbb{R}^d$ are K system's unknown parameters. GLBs are special cases of this for $K = 1$, also known as single-index models (SIM) in statistics. In this paper, we formulate an extension of the problem of binary logistic bandits (i.e., a special case of SIM) to multinomial logit (MNL) bandits, a special case of multi-index models (MIM) to account for settings with more than two possible outcomes on the user choices ($K \geq 1$). For this model, we present an algorithm and corresponding regret bound. Our algorithmic and analytic contribution is in large inspired by recent exciting progress on binary logistic bandits by Faury et al. (2020).

To motivate MNL bandits, consider ad placement. When an ad is shown to a user, the user may have several options to react to the ad. For example, she can choose to 1) click on the ad; 2) click on “show me later”; 3) click on “never show me this ad”; 4) not click at all, etc. The user selects each of these options based on an unknown probability distribution that inherently involves linear combinations of the selected feature vector denoting the ad and unknown parameters denoting the user's preferences about the ad. In this setting, each option is associated with a specific notion of reward. The agent's goal is to determine ads with maximum expected rewards to increase the chance of a successful advertisement.

Outline. In Section 1.1, we formally define the problem. In Sections 2.1, 2.2 and 2.4, we elaborate on the challenges that the generalization of the Logistic-UCB-1 by Faury et al. (2020) to the settings with MIM rewards brings to our theoretical analysis. We then summarize our proposed MNL-UCB in Algorithm 1 and provide a regret bound for it in Section 2.5. In Section 3, we present a detailed discussion on the challenges and computation of necessary problem-dependent constants. Finally, we complement our theoretical results with numerical simulations in Section 4.

Notation. We use lower-case letters for scalars, lower/upper-case bold letters for vectors/matrices. $\|\mathbf{x}\|_2$ denotes the Euclidean norm and $\mathbf{x}^T \mathbf{y}$ inner product. δ_{ij} denotes the Kronecker delta, i.e., $\delta_{ij} = 1$ if $i \neq j$ and $\delta_{ij} = 0$ if $i = j$. $\mathbf{A} \otimes \mathbf{B}$ denotes the Kronecker product. For square matrices \mathbf{A} and \mathbf{B} , we use $\mathbf{A} \preceq \mathbf{B}$ to denote $\mathbf{B} - \mathbf{A}$ is positive semi-definite. We denote the minimum and maximum eigenvalues of \mathbf{A} by $\lambda_{\min}(\mathbf{A})$ and $\lambda_{\max}(\mathbf{A})$. For $\mathbf{A} \succeq 0$, the weighted 2-norm of $\boldsymbol{\nu}$ with respect to \mathbf{A} is defined by $\|\boldsymbol{\nu}\|_{\mathbf{A}} = \sqrt{\boldsymbol{\nu}^T \mathbf{A} \boldsymbol{\nu}}$. For positive integers n and $m \leq n$, $[n]$ and $[m : n]$ denote the sets $\{1, 2, \dots, n\}$ and $\{m, \dots, n\}$, respectively. For any $\boldsymbol{\nu} \in \mathbb{R}^{Kd}$, $\bar{\boldsymbol{\nu}}_i = \boldsymbol{\nu}_{[(i-1)d+1:d]} \in \mathbb{R}^d$ denotes the vector containing the i -th set of d entries of $\boldsymbol{\nu}$. We use $\mathbf{1}$ and \mathbf{e}_i to denote the vector of all 1's and the i -th standard basis vector, respectively. Finally, we use \tilde{O} for big-Oh notation that ignores logarithmic factors.

1.1 Problem formulation

Reward Model. The agent is given a decision set¹ $\mathcal{D} \subset \mathbb{R}^d$. At each round t , the agent chooses an action $\mathbf{x}_t \in \mathcal{D}$ and observes the user purchase decision $y_t \in [K] \cup \{0\}$. Here, $\{0\}$ denotes the “outside decision”, which means the user did not select any of the presented options. The agent's decision at round t is based on the information gathered until time t , which can be formally encoded in the filtration $\mathcal{F}_t := (\mathcal{F}_0, \sigma(\{\mathbf{x}_s, y_s\}_{s=1}^{t-1}))$, where \mathcal{F}_0 represents any prior knowledge. Let each option $i \in [K]$ be associated with an unknown vector $\bar{\boldsymbol{\theta}}_{*i} \in \mathbb{R}^d$ and let $\boldsymbol{\theta}_* = [\bar{\boldsymbol{\theta}}_{*1}^T, \bar{\boldsymbol{\theta}}_{*2}^T, \dots, \bar{\boldsymbol{\theta}}_{*K}^T]^T \in \mathbb{R}^{Kd}$. The user's choice of what to click on is given by a multinomial logit (MNL) choice model. Under this model, the probability distribution of the user purchase decision is given by

$$\mathbb{P}(y_t = i | \mathbf{x}_t, \mathcal{F}_t) := \begin{cases} \frac{1}{1 + \sum_{j=1}^K \exp(\bar{\boldsymbol{\theta}}_{*j}^T \mathbf{x}_t)}, & \text{if } i = 0, \\ \frac{\exp(\bar{\boldsymbol{\theta}}_{*i}^T \mathbf{x}_t)}{1 + \sum_{j=1}^K \exp(\bar{\boldsymbol{\theta}}_{*j}^T \mathbf{x}_t)}, & \text{if } i \in [K]. \end{cases} \quad (1)$$

When the user clicks on the i -th option, a corresponding reward ρ_i is revealed to the agent and we set $\rho_0 = 0$. Then, the expected reward observed by the agent when she plays action \mathbf{x}_t is

¹Our results extend easily to time-varying decision sets.

$$\mathbb{E}[R_t | \mathbf{x}_t, \mathcal{F}_t] = \frac{\sum_{j=1}^K \rho_j \exp(\bar{\theta}_{*j}^T \mathbf{x}_t)}{1 + \sum_{j=1}^K \exp(\bar{\theta}_{*j}^T \mathbf{x}_t)} = \boldsymbol{\rho}^T \mathbf{z}(\mathbf{x}_t, \boldsymbol{\theta}_*), \text{ where } \boldsymbol{\rho} = [\rho_1, \rho_2, \dots, \rho_K]^T, \mathbf{z}(\mathbf{x}_t, \boldsymbol{\theta}_*) = [z_1(\mathbf{x}_t, \boldsymbol{\theta}_*), z_2(\mathbf{x}_t, \boldsymbol{\theta}_*), \dots, z_K(\mathbf{x}_t, \boldsymbol{\theta}_*)]^T, \text{ and}$$

$$z_i(\mathbf{x}_t, \boldsymbol{\theta}_*) = \mathbb{P}(y_t = i | \mathbf{x}_t, \mathcal{F}_t), \forall i \in [K]. \quad (2)$$

Note that $\mathbb{E}[R_t | \mathbf{x}_t, \mathcal{F}_t] = \mu(\bar{\theta}_{*1}^T \mathbf{x}, \bar{\theta}_{*2}^T \mathbf{x}, \dots, \bar{\theta}_{*K}^T \mathbf{x})$ is *not* directly a generalized linear model, i.e., a function $\mu(\bar{\boldsymbol{\theta}}_*^T \mathbf{x}_t)$, but rather it is a *multi-index model*, where $\mu : \mathbb{R}^K \rightarrow \mathbb{R}$.

Goal. Let T be the total number of rounds and \mathbf{x}_* be the optimal action that maximizes the reward in expectation, i.e., $\mathbf{x}_* \in \arg \max_{\mathbf{x} \in \mathcal{D}} \boldsymbol{\rho}^T \mathbf{z}(\mathbf{x}, \boldsymbol{\theta}_*)$. The agent’s goal is to minimize the *cumulative pseudo-regret* defined by $R_T = \sum_{t=1}^T \boldsymbol{\rho}^T \mathbf{z}(\mathbf{x}_*, \boldsymbol{\theta}_*) - \boldsymbol{\rho}^T \mathbf{z}(\mathbf{x}_t, \boldsymbol{\theta}_*)$.

1.2 Contributions

We study MNL logistic regression bandits, a generalization of binary logistic bandits, that address applications where the number of outcomes that can be selected by the user is larger than two. The probability of selecting any possible $K + 1 > 2$ options (+1 stands for the ‘not click’ outcome aka ‘outside decision’) is modeled using a multinomial logit (MNL) model. For this problem:

- We identify a critical parameter κ , which we interpret as the degree of (non)-smoothness (less smooth for larger values of κ) of the MNL model over the agent’s decision set. We prove that κ scales exponentially with the size of the agent’s decision set creating a challenge in the design of low-regret algorithms, similar to the special binary case previously studied in the literature.
- We develop a UCB-type algorithm for MNL logistic regression bandits. At every step, the algorithm decides on the inclusion of a K -tuple of parameter vectors $\bar{\boldsymbol{\theta}}_{*1}, \dots, \bar{\boldsymbol{\theta}}_{*K}$ in the confidence region in a way that captures the local smoothness of the MNL model around this K -tuple and past actions. We show that this is critical for the algorithm’s favorable regret performance in terms of κ .
- Specifically, we prove that the regret of our MNL-UCB scales as $\tilde{\mathcal{O}}(dK\sqrt{\kappa}\sqrt{T})$. Instead, we show that a confidence ellipsoid that fails to capture local dependencies described above results in regret that scales linearly with κ rather than with $\sqrt{\kappa}$. Moreover, our regret bound scales optimally in terms of the number of options K .
- We propose an improved algorithm that achieves a regret bound with problem-dependent constant κ being pushed into a second order term that vanishes quickly.
- We complement our theoretical results with numerical simulations and corresponding discussions on the performance of our algorithm.

1.3 Related works

Generalized Linear Bandits. GLBs were studied in [Filippi et al. \(2010\)](#); [Li et al. \(2017\)](#); [Abeille et al. \(2017\)](#); [Russo and Van Roy \(2013, 2014\)](#); [Dong and Van Roy \(2018\)](#) where the stochastic reward is modeled through an appropriate strictly increasing link function μ . All these works provide regret bounds $\tilde{\mathcal{O}}(\kappa\sqrt{T})$, where the multiplicative factor κ is a problem-dependent constant and characterizes the degree of non-linearity of the link function.

Logistic Bandits. In [Fauray et al. \(2020\)](#), the authors focused on the *logistic bandit* problem as a special case of GLBs. By introducing a novel Bernstein-like self-normalized martingale tail-inequality, they reduced the dependency of the existing GLB algorithms’ regret bounds on the constant κ by a factor of $\sqrt{\kappa}$ and obtained a $\mathcal{O}(d\sqrt{\kappa T})$ regret for the logistic bandit problem. They further discussed the crucial role of κ , which can be arbitrarily large as it scales exponentially with the size of the decision set, on the performance of existing algorithms. Motivated by such considerations, with careful algorithmic designs, they achieved to drop entirely the dependence on κ leading to a regret of $\tilde{\mathcal{O}}(d\sqrt{T})$.

Multinomial Logit Bandits. In a different line of work, [Agrawal et al. \(2017, 2019\)](#); [Wang et al. \(2018\)](#); [Oh and Iyengar \(2019\)](#); [Chen et al. \(2018\)](#); [Dong et al. \(2020\)](#); [Agrawal et al. \(2020\)](#) used the multinomial logit choice model to address the *dynamic assortment selection* problem, which is a combinatorial variant of the bandit problem. In this problem, the agent chooses a so-called assortment which is a subset of a set $\mathcal{S} = [N]$ of N items. At round t , feature vectors \mathbf{x}_{it} , for every item $i \in \mathcal{S}$, are revealed to the agent, and given this contextual information, the agent selects an assortment $S_t \subset \mathcal{S}$

and observes the user choice $y_t = i, i \in S_t \cup \{0\}$ where $\{0\}$ corresponds to the user not selecting any item in S_t . The user choice is given by a MNL model with an unknown parameter $\bar{\theta}_* \in \mathbb{R}^d$ such that the probability that the user selects item $i \in S_t$ is $\exp(\bar{\theta}_*^T \mathbf{x}_{it}) / (1 + \sum_{i \in S_t} \exp(\bar{\theta}_*^T \mathbf{x}_{it}))$. Furthermore, a revenue parameter denoted by r_{it} for each item i is also revealed at round t . The goal of the agent is to offer assortments with size at most K to maximize the expected cumulative revenue or to minimize the cumulative regret $\sum_{t=1}^T R_t(S_t^*, \bar{\theta}_*) - R_t(S_t, \bar{\theta}_*)$, where $R_t(S_t, \bar{\theta}_*) := \sum_{i \in S_t} r_{it} \exp(\bar{\theta}_*^T \mathbf{x}_{it}) / (1 + \sum_{i \in S_t} \exp(\bar{\theta}_*^T \mathbf{x}_{it}))$. Finally, the closely related paper [Cheung and Simchi-Levi \(2017\)](#) studies a problem where at each round, the agent observes a user-specific context based on which, it recommends a set of items to the user. The probability distribution of each one of the items in that set being selected by the user is given by an MNL model. This problem can be categorized as an online assortment optimization problem. Despite similarities in the use of an MNL model, there are certain differences between [Cheung and Simchi-Levi \(2017\)](#) and our paper in terms of problem formulation. In our setting, the user may have multiple reactions (one of $K + 1$ options), to a single selected item. In contrast, in [Cheung and Simchi-Levi \(2017\)](#), the agent must select a set of items to each of which the user reacts by either clicking or not clicking. Also, here the probability distribution of different user reactions remains the same at all rounds, while in [Cheung and Simchi-Levi \(2017\)](#) the response to an item from the recommended set depends on the other items in that set. We defer to future work studying implications of our techniques to the setting of [Cheung and Simchi-Levi \(2017\)](#).

2 Multinomial Logit UCB Algorithms

In this section, we introduce two key quantities: (i) θ_t , an estimate of θ_* ; (ii) $\epsilon_t(\mathbf{x})$, an exploration bonus for each $\mathbf{x} \in \mathcal{D}$ at each round $t \in [T]$. Based on these, we design a UCB-type algorithm, called MNL-UCB. At each round t , the algorithm computes an estimate θ_t of θ_* , that we present in [Section 2.4](#). For each $\mathbf{x} \in \mathcal{D}$ and $t \in [T]$, let $\epsilon_t(\mathbf{x})$ be such that the following holds with high probability:

$$\Delta(\mathbf{x}, \theta_t) := |\rho^T \mathbf{z}(\mathbf{x}, \theta_*) - \rho^T \mathbf{z}(\mathbf{x}, \theta_t)| \leq \epsilon_t(\mathbf{x}). \quad (3)$$

At round t , having knowledge of $\epsilon_t(\mathbf{x})$, the agent computes the following upper bound on the expected reward $\rho^T \mathbf{z}(\mathbf{x}, \theta_*)$ for all $\mathbf{x} \in \mathcal{D}$:

$$\rho^T \mathbf{z}(\mathbf{x}, \theta_*) \leq \rho^T \mathbf{z}(\mathbf{x}, \theta_t) + \epsilon_t(\mathbf{x}). \quad (4)$$

Then, the learner follows a UCB decision rule to select an action \mathbf{x}_t according to the following rule:

$$\mathbf{x}_t := \arg \max_{\mathbf{x} \in \mathcal{D}} \rho^T \mathbf{z}(\mathbf{x}, \theta_t) + \epsilon_t(\mathbf{x}). \quad (5)$$

To see how the UCB decision rule in [\(5\)](#) helps us control the cumulative regret, we show how it controls the instantaneous regret by the following standard argument [Abbasi-Yadkori et al. \(2011\)](#); [Fauray et al. \(2020\)](#):

$$\begin{aligned} r_t &= \rho^T \mathbf{z}(\mathbf{x}_*, \theta_*) - \rho^T \mathbf{z}(\mathbf{x}_t, \theta_*) \stackrel{(4)}{\leq} \rho^T \mathbf{z}(\mathbf{x}_*, \theta_t) + \epsilon_t(\mathbf{x}_*) - \rho^T \mathbf{z}(\mathbf{x}_t, \theta_*) \\ &\stackrel{(5)}{\leq} \rho^T \mathbf{z}(\mathbf{x}_t, \theta_t) - \rho^T \mathbf{z}(\mathbf{x}_t, \theta_*) + \epsilon_t(\mathbf{x}_t) \stackrel{(3)}{\leq} 2\epsilon_t(\mathbf{x}_t). \end{aligned} \quad (6)$$

In view of this, our goal is to design an algorithm that appropriately chooses the estimator θ_t and the exploration bonus $\epsilon_t(\mathbf{x})$ such that its regret is sub-linear.

2.1 Maximum likelihood estimate

The problem of estimating θ_* at round t given \mathcal{F}_t is identical to a multi-class linear classification problem, where $\bar{\theta}_{*i}$ is the ‘‘classifier’’ for class i . A natural way to compute the estimator of the unknown parameter θ_* of the MNL model given \mathcal{F}_t is to use the maximum likelihood principle. At round t , the regularized log-likelihood (aka negative cross-entropy loss) with regularizer $\lambda > 0$ writes

$$\mathcal{L}_t^\lambda(\theta) := \sum_{s=1}^{t-1} \sum_{i=0}^K \mathbb{1}\{y_s = i\} \log(z_i(\mathbf{x}_s, \theta)) - \frac{\lambda}{2} \|\theta\|_2^2. \quad (7)$$

Then, the maximum likelihood estimate of θ_* is defined as $\hat{\theta}_t := \arg \max_{\theta} \mathcal{L}_t^\lambda(\theta)$. Taking the gradient of (7) with respect to θ we obtain

$$\nabla_{\theta} \mathcal{L}_t^\lambda(\theta) := \sum_{s=1}^{t-1} [\mathbf{m}_s - \mathbf{z}(\mathbf{x}_s, \theta)] \otimes \mathbf{x}_s - \lambda \theta, \quad (8)$$

where \mathbf{m}_s is the ‘one-hot encoding’ vector of the user’s selection at round s , i.e., $\mathbf{m}_s := [\mathbb{1}\{y_s = 1\}, \dots, \mathbb{1}\{y_s = K\}]^T$. It will also be convenient to define the Hessian of $-\mathcal{L}_t^\lambda(\theta)$:

$$\mathbf{H}_t(\theta) := \lambda I_{Kd} + \sum_{s=1}^{t-1} \mathbf{A}(\mathbf{x}_s, \theta) \otimes \mathbf{x}_s \mathbf{x}_s^T, \quad (9)$$

where $\mathbf{A}(\mathbf{x}, \theta)_{ij} := z_i(\mathbf{x}, \theta) (\delta_{ij} - z_j(\mathbf{x}, \theta))$ for all $i, j \in [K]$. Equivalently, in matrix form

$$\mathbf{A}(\mathbf{x}, \theta) := \text{diag}(\mathbf{z}(\mathbf{x}, \theta)) - \mathbf{z}(\mathbf{x}, \theta) \mathbf{z}^T(\mathbf{x}, \theta). \quad (10)$$

Note here that $\mathbf{A}(\mathbf{x}, \theta)$ is a matrix function that depends on θ and \mathbf{x} via the inner products $\bar{\theta}_1^T \mathbf{x}, \bar{\theta}_2^T \mathbf{x}, \dots, \bar{\theta}_K^T \mathbf{x}$. Also, the matrix $\mathbf{A}(\mathbf{x}, \theta)$ has nice algebraic properties (discussed more in Section 3) that turn out to be critical in the execution and analysis of our algorithm.

Now, we introduce assumptions on the problem structure, under which our theoretical results hold.

Assumption 1 (Boundedness). *Without loss of generality, $\|\mathbf{x}\|_2 \leq 1$ for all $\mathbf{x} \in \mathcal{D}$. Also, $\theta_* \in \Theta := \{\theta \in \mathbb{R}^{Kd} : \|\theta\|_2 \leq S\}$ and $\|\rho\|_2 \leq R$. Both upper bounds S and R are known to the agent.*

The assumption that S is known is standard in the literature of GLBs. Knowledge of R is also reasonable to assume because ρ_i ’s represent the revenue parameters that are typically known or set by the system operator.

Assumption 2 (Problem-dependent constants). *There exist strictly positive constants $0 < L < \infty$ and $0 < \kappa < \infty$ such that $\sup_{\mathbf{x} \in \mathcal{D}, \theta \in \Theta} \lambda_{\max}(\mathbf{A}(\mathbf{x}, \theta)) := L$ and $\inf_{\mathbf{x} \in \mathcal{D}, \theta \in \Theta} \lambda_{\min}(\mathbf{A}(\mathbf{x}, \theta)) := \frac{1}{\kappa}$.*

We comment further on the knowledge of κ and L in Section 3. Here, we note that κ is reminiscent of the corresponding quantity in binary logistic bandits which is defined accordingly as $\kappa := \sup_{\mathbf{x} \in \mathcal{D}, \|\theta\|_2 \leq S} 1/\dot{\mu}(\bar{\theta}^T \mathbf{x})$, where $\dot{\mu}$ is the first derivative of the logistic function $\mu(x) = 1/(1 + \exp(-x))$. As Filippi et al. (2010); Li et al. (2017); Fauray et al. (2020) have shown, this quantity plays a key role in characterizing the behavior of binary ($K = 1$) logit bandit algorithms. In this paper, we will show that the proper analogue of this quantity to multinomial ($K > 1$) logit bandit algorithms is the parameter κ defined in Assumption 2.

2.2 Confidence set around θ_*

We introduce a confidence set $\mathcal{C}_t(\delta)$ that will include θ_* with high probability thus allowing us to upper bound $\Delta(\mathbf{x}, \theta_t)$ in the following subsection. We start by defining the key quantity

$$\mathbf{g}_t(\theta) := \lambda \theta + \sum_{s=1}^{t-1} \mathbf{z}(\mathbf{x}_s, \theta) \otimes \mathbf{x}_s. \quad (11)$$

To see why this is useful, note that by the first-order optimality condition $\nabla_{\theta} \mathcal{L}_t^\lambda(\hat{\theta}_t) = \mathbf{0}$ we have θ_* satisfying $\mathbf{g}_t(\theta_*) - \mathbf{g}_t(\hat{\theta}_t) = \lambda \theta_* + \mathbf{s}_t$, with $\mathbf{s}_t := \sum_{s=1}^{t-1} (\mathbf{z}(\mathbf{x}_s, \theta_*) - \mathbf{m}_s) \otimes \mathbf{x}_s$. This in turn motivates us to define a confidence set \mathcal{C}_t at the beginning of each round $t \in [T]$ such that

$$\mathcal{C}_t(\delta) := \{\theta \in \Theta : \|\mathbf{g}_t(\theta) - \mathbf{g}_t(\hat{\theta}_t)\|_{\mathbf{H}_t^{-1}(\theta)} \leq \beta_t(\delta)\}, \quad (12)$$

where $\beta_t(\delta)$ is chosen as in Theorem 1 below to guarantee $\theta_* \in \mathcal{C}_t(\delta)$ with high probability $1 - \delta$.

Theorem 1 (Confidence set). *Let the Assumption 1 hold and for $\delta \in (0, 1)$, define $\beta_t(\delta) := \frac{K^{3/2}d}{\sqrt{\lambda}} \log(1 + \frac{t}{d\lambda}) + \frac{\sqrt{\lambda/K}}{2} + \frac{2K^{3/2}d}{\sqrt{\lambda}} \log(\frac{2}{\delta}) + \sqrt{\lambda}S$. Then with probability at least $1 - \delta$, for all $t \in [T]$ it holds that $\theta_* \in \mathcal{C}_t(\delta)$.*

Once we have properly identified the key quantities $\mathbf{g}_t(\boldsymbol{\theta}_*)$ and $\mathbf{H}_t(\boldsymbol{\theta})$, the proof of the theorem above rather naturally extends Lemma 1 in [Fauray et al. \(2020\)](#). Compared to the special case $K = 1$ studied in [Fauray et al. \(2020\)](#), extra care is needed here to properly track the scaling of $\beta_t(\delta)$ with respect to the new parameter K that is of interest for us. The details are deferred to [Appendix A](#). To a great extent, the similarities between the binary and multinomial cases end here. It will turn out that bounding $\Delta(\mathbf{x}, \boldsymbol{\theta}_t)$, for an appropriate choice of $\boldsymbol{\theta}_t$, is significantly more intricate here than in the binary case. Our main technical contribution towards this direction is given in the lemmas presented in [Section 2.3](#) that are used to prove the following key result.

Lemma 1. *Let Assumptions 1 and 2 hold. For all $\mathbf{x} \in \mathcal{D}$, $t \in [T]$ and $\boldsymbol{\theta} \in \mathcal{C}_t(\delta)$, with probability at least $1 - \delta$:*

$$\Delta(\mathbf{x}, \boldsymbol{\theta}) \leq 2RL\beta_t(\delta)\sqrt{\kappa(1+2S)}\|\mathbf{x}\|_{\mathbf{V}_t^{-1}}, \quad (13)$$

where $\mathbf{V}_t := \kappa\lambda I + \sum_{s=1}^{t-1} \mathbf{x}_s \mathbf{x}_s^T$.

The complete proof is in [Appendix B](#). In the following section, we give a proof sketch.

2.3 Proof sketch of Lemma 1

To prove [Lemma 1](#), we will use the high probability confidence set $\mathcal{C}_t(\delta)$ in [\(12\)](#) paired with the problem setting's properties encapsulated in [Assumptions 1 and 2](#). To see the key challenges in establishing [\(13\)](#), consider the following. By definition of $\Delta(\mathbf{x}, \boldsymbol{\theta})$ in [\(3\)](#), Cauchy-Schwarz inequality, and [Assumption 1](#), for any $\mathbf{x} \in \mathcal{D}$, $t \in [T]$, and $\boldsymbol{\theta} \in \mathcal{C}_t(\delta)$:

$$\Delta(\mathbf{x}, \boldsymbol{\theta}) \leq R\|\mathbf{z}(\mathbf{x}, \boldsymbol{\theta}_*) - \mathbf{z}(\mathbf{x}, \boldsymbol{\theta})\|_2. \quad (14)$$

Thus, our goal becomes relating $\|\mathbf{z}(\mathbf{x}, \boldsymbol{\theta}_*) - \mathbf{z}(\mathbf{x}, \boldsymbol{\theta})\|_2$ to $\left\| \mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t) \right\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta}_*)}$ and/or $\left\| \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t) - \mathbf{g}_t(\boldsymbol{\theta}) \right\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta})}$. The reason is that the two latter quantities are both known to be bounded by $\beta_t(\delta)$ with high probability since $\boldsymbol{\theta}_* \in \mathcal{C}_t(\delta)$ with probability at least $1 - \delta$ (cf. [Theorem 1](#)). We accomplish our goal in three steps. First, in [Lemma 2](#), we connect $\mathbf{z}(\mathbf{x}, \boldsymbol{\theta}_1) - \mathbf{z}(\mathbf{x}, \boldsymbol{\theta}_2)$ to $\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2$ for any $\mathbf{x} \in \mathbb{R}^d$ and $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \mathbb{R}^{Kd}$.

Lemma 2. *For any $\mathbf{x} \in \mathbb{R}^d$, $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \mathbb{R}^{Kd}$, recall the definition of $\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})$ in [\(10\)](#) and define*

$$\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2) := \int_0^1 \mathbf{A}(\mathbf{x}, v\boldsymbol{\theta}_1 + (1-v)\boldsymbol{\theta}_2) dv. \quad (15)$$

Then, we have $\mathbf{z}(\mathbf{x}, \boldsymbol{\theta}_1) - \mathbf{z}(\mathbf{x}, \boldsymbol{\theta}_2) = [\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \otimes \mathbf{x}^T] (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2)$.

The proof of the lemma above in [Appendix B.1](#) relies on a proper application of the mean-value Theorem. Next, in [Lemma 3](#) below, we relate $\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2$ to $\mathbf{g}_t(\boldsymbol{\theta}_1) - \mathbf{g}_t(\boldsymbol{\theta}_2)$.

Lemma 3. *Let*

$$\mathbf{G}_t(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2) := \lambda I_{Kd} + \sum_{s=1}^{t-1} \mathbf{B}(\mathbf{x}_s, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \otimes \mathbf{x}_s \mathbf{x}_s^T. \quad (16)$$

Then, for any $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \mathbb{R}^{Kd}$, we have

$$\mathbf{g}_t(\boldsymbol{\theta}_1) - \mathbf{g}_t(\boldsymbol{\theta}_2) = \mathbf{G}_t(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)(\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2). \quad (17)$$

The proof in [Appendix B.1](#) uses the definition of $\mathbf{g}_t(\boldsymbol{\theta})$, [Lemma 2](#) and a proper application of the mixed-product property of the Kronecker product. Note that $\mathbf{G}_t(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \succeq 0$; thus, it is invertible.

Now, combining [Lemmas 2 and 3](#), our new goal becomes bounding $\left\| [\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_*, \boldsymbol{\theta}) \otimes \mathbf{x}^T] \mathbf{G}_t^{-1}(\boldsymbol{\theta}_*, \boldsymbol{\theta}) (\mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\boldsymbol{\theta})) \right\|_2$. Our key technical contribution is establishing good bounds for the spectral norm $\left\| [\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_*, \boldsymbol{\theta}) \otimes \mathbf{x}^T] \mathbf{G}_t^{-1/2}(\boldsymbol{\theta}_*, \boldsymbol{\theta}) \right\|_2$ and the weighted Euclidean norm $\left\| \mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\boldsymbol{\theta}) \right\|_{\mathbf{G}_t^{-1}(\boldsymbol{\theta}_*, \boldsymbol{\theta})}$. We start by briefly explaining how we bound the spectral norm above; see [Appendix B.2](#) for details. By using the cyclic property of

the maximum eigenvalue and the mixed-product property of the Kronecker product, it suffices to bound $\lambda_{\max} \left(\mathbf{G}_t^{-1/2} \left((\mathbf{B}^T \mathbf{B}) \otimes (\mathbf{x} \mathbf{x}^T) \right) \mathbf{G}_t^{-1/2} \right)$, where we denote $\mathbf{B} = \mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_*, \boldsymbol{\theta})$ and $\mathbf{G}_t = \mathbf{G}_t(\boldsymbol{\theta}_*, \boldsymbol{\theta})$ for simplicity. There are two essential ideas to do so. First, thanks to our Assumption 2, which upper bounds the eigenvalues of $\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})$ by L , and by recalling the definition of $\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$, we manage to show that $(\mathbf{B} \mathbf{B}^T) \otimes (\mathbf{x} \mathbf{x}^T) \preceq L^2 (I_K \otimes \mathbf{x}) (I_K \otimes \mathbf{x}^T)$. Our second idea is to relate the matrix \mathbf{G}_t to the Gram matrix of actions

$$\mathbf{V}_t := \kappa \lambda I_d + \sum_{s=1}^{t-1} \mathbf{x}_s \mathbf{x}_s^T. \quad (18)$$

Specifically, using our Assumption 2 that the minimum eigenvalue of $\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})$ is lower bounded by $1/\kappa$ and standard spectral properties of the Kronecker product, we prove in Lemma 12 in the appendix that $\mathbf{G}_t \succeq \frac{1}{\kappa} I_K \otimes \mathbf{V}_t$ or $\mathbf{G}_t^{-1} \preceq \kappa I_K \otimes \mathbf{V}_t^{-1}$. By properly combining the above, we achieve the following convenient upper bound:

$$\left\| [\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_*, \boldsymbol{\theta}) \otimes \mathbf{x}^T] \mathbf{G}_t^{-1/2}(\boldsymbol{\theta}_*, \boldsymbol{\theta}) \right\|_2 \leq L \sqrt{\kappa} \|\mathbf{x}\|_{\mathbf{V}_t^{-1}}. \quad (19)$$

To see why this is useful, note that compared to $\mathbf{H}_t(\boldsymbol{\theta})$ in (9), which is a ‘matrix-weighted’ version of the Gram matrix, the definition of \mathbf{V}_t in (18) is the same as the definition of the Gram matrix in linear bandits. Thus, we are now able to use standard machinery in Abbasi-Yadkori et al. (2011) to bound $\sum_{t=1}^T \|\mathbf{x}_t\|_{\mathbf{V}_t^{-1}}$.

Finally, we discuss how to control the remaining term $\|\mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\boldsymbol{\theta})\|_{\mathbf{G}_t(\boldsymbol{\theta}_*, \boldsymbol{\theta})^{-1}}$. By adding and subtracting $\mathbf{g}_t(\hat{\boldsymbol{\theta}}_t)$ to and from the argument inside the norm, it suffices to bound

$$\left\| \mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t) \right\|_{\mathbf{G}_t(\boldsymbol{\theta}_*, \boldsymbol{\theta})^{-1}} + \left\| \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t) - \mathbf{g}_t(\boldsymbol{\theta}) \right\|_{\mathbf{G}_t(\boldsymbol{\theta}_*, \boldsymbol{\theta})^{-1}}. \quad (20)$$

To do this, we exploit the definition of $\mathcal{C}_t(\delta)$ in (12) by first relating the $\mathbf{G}_t^{-1}(\boldsymbol{\theta}_*, \boldsymbol{\theta})$ -norms with those in terms of $\mathbf{H}_t^{-1}(\boldsymbol{\theta}_*)$ and $\mathbf{H}_t^{-1}(\boldsymbol{\theta})$. To do this, we rely on the *generalized self-concordance* property of the strictly convex *log-sum-exp* (lse : $\mathbb{R}^K \rightarrow \mathbb{R} \cup \infty$) function Tran-Dinh et al. (2015) $\text{lse}(\mathbf{s}) := \log(1 + \sum_{i=1}^K \exp(s_i))$. Notice that our $\mathbf{z}(\mathbf{x}, \boldsymbol{\theta})$ is the gradient of the lse function at point $[\bar{\boldsymbol{\theta}}_1^T \mathbf{x}, \bar{\boldsymbol{\theta}}_2^T \mathbf{x}, \dots, \bar{\boldsymbol{\theta}}_K^T \mathbf{x}]^T$, that is $\mathbf{z}(\mathbf{x}, \boldsymbol{\theta}) = \nabla \text{lse}([\bar{\boldsymbol{\theta}}_1^T \mathbf{x}, \bar{\boldsymbol{\theta}}_2^T \mathbf{x}, \dots, \bar{\boldsymbol{\theta}}_K^T \mathbf{x}]^T)$. Thanks to the generalized self-concordance property of the lse, upper bounds on its Hessian matrix (essentially the matrix $\mathbf{H}_t(\boldsymbol{\theta})$) have been developed in Tran-Dinh et al. (2015); Sun and Tran-Dinh (2019). Proper use of such bounds leads to lower bounds on $\mathbf{G}_t(\boldsymbol{\theta}_*, \boldsymbol{\theta})$ as follows.

Lemma 4 (Generalized self-concordance). *For any $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \Theta$, we have $(1 + 2S)^{-1} \mathbf{H}_t(\boldsymbol{\theta}_1) \preceq \mathbf{G}_t(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$ and $(1 + 2S)^{-1} \mathbf{H}_t(\boldsymbol{\theta}_2) \preceq \mathbf{G}_t(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$.*

The proof is given in Appendix B.1. Finally, plugging the above lower bounds on matrix \mathbf{G}_t into (20) for $\boldsymbol{\theta}_*$ and $\boldsymbol{\theta}$ gives the final bound on $\Delta(\mathbf{x}, \boldsymbol{\theta})$ in Lemma 1.

2.4 Error bound on $\Delta(\mathbf{x}, \boldsymbol{\theta}_t)$

In this section, we specify $\boldsymbol{\theta}_t$ and $\epsilon_t(\mathbf{x})$ for all $t \in [T]$ and $\mathbf{x} \in \mathcal{D}$. In view of Lemma 1, the exploration bonus $\epsilon_t(\mathbf{x})$ can be set equal to the RHS of (13), only if $\boldsymbol{\theta}_t \in \mathcal{C}_t(\delta)$. Recall from the definition of $\mathcal{C}_t(\delta)$ in (12) that for $\boldsymbol{\theta}_t \in \mathcal{C}_t(\delta)$, it must be that $\boldsymbol{\theta}_t \in \Theta$. Since the ML estimator $\hat{\boldsymbol{\theta}}_t$ does not necessarily satisfy $\hat{\boldsymbol{\theta}}_t \in \Theta$, we introduce the following ‘feasible estimator’:

$$\boldsymbol{\theta}_t := \arg \min_{\boldsymbol{\theta} \in \Theta} \left\| \mathbf{g}_t(\boldsymbol{\theta}) - \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t) \right\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta})}, \quad (21)$$

which is guaranteed to be in the confidence set $\mathcal{C}_t(\delta)$ for all $t \in [T]$ since $\|\mathbf{g}_t(\boldsymbol{\theta}_t) - \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t)\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta})} \leq \|\mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t)\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta})} \leq \beta_t(\delta)$ and $\boldsymbol{\theta}_t \in \Theta$. Thus, we have proved the following.

Corollary 1 (Exploration bonus). *For all $\mathbf{x} \in \mathcal{D}$ and $t \in [T]$, with probability at least $1 - \delta$, we have*

$$\Delta(\mathbf{x}, \boldsymbol{\theta}_t) \leq \epsilon_t(\mathbf{x}) := 2RL\beta_t(\delta) \sqrt{\kappa(1 + 2S)} \|\mathbf{x}\|_{\mathbf{V}_t^{-1}}. \quad (22)$$

Algorithm 1: MNL-UCB

```

1 for  $t = 1, \dots, T$  do
2   Compute  $\boldsymbol{\theta}_t$  as in (21).
3   Compute  $\mathbf{x}_t := \arg \max_{\mathbf{x} \in \mathcal{D}} \boldsymbol{\rho}^T \mathbf{z}(\mathbf{x}, \boldsymbol{\theta}_t) + \epsilon_t(\mathbf{x})$  with  $\epsilon_t(\mathbf{x})$  defined in (22).
4   Play  $\mathbf{x}_t$  and observe  $y_t$ .
```

With these, we are now ready to summarize MNL-UCB in Algorithm 1.

Remark 1. The projection step in (21) is similar to the ones used in Filippi et al. (2010); Faury et al. (2020) for binary logistic bandits. In particular, this step in Filippi et al. (2010) involves norms with respect to \mathbf{V}_t instead of $\mathbf{H}_t(\boldsymbol{\theta})$ (for $K = 1$). All of these involve non-convex optimization problems. Empirically, we observe that it occurs frequently that $\hat{\boldsymbol{\theta}}_t \in \Theta$. Then, $\boldsymbol{\theta}_t = \hat{\boldsymbol{\theta}}_t$ and these complicated projection steps do not need to be implemented.

2.5 Regret bound of MNL-UCB

In the following theorem, we state the regret bound of MNL-UCB as our main result.

Theorem 2 (Regret of MNL-UCB). *Fix $\delta \in (0, 1)$. Let Assumptions 1 and 2 hold. Then, with probability at least $1 - \delta$, it holds that $R_T \leq 4RL\beta_T(\delta)\sqrt{2 \max(1, \frac{1}{\lambda\kappa})\kappa(1 + 2S)} dT \log(1 + \frac{T}{\kappa\lambda d})$. In particular, choosing $\lambda = Kd \log(T)$ yields $R_T \leq \mathcal{O}\left(RLKd \log(T)\sqrt{\kappa T}\right)$.*

Now we comment on the regret order with respect to the key problem parameters T, d, K and $\sqrt{\kappa}$. With respect to these the theorem shows a scaling of the regret of Algorithm 1 as $\mathcal{O}(Kd \log(T)\sqrt{\kappa T})$. Specifically, for $K = 1$ we retrieve the exact same scaling as in Theorem 2 of Faury et al. (2020) for the binary case in terms of T, d and κ . In particular, the bound is optimal with respect to the action-space dimension d as $\mathcal{O}(d)$ is the optimal order in the simpler setting of linear bandits Dani et al. (2008). Of course, compared to Faury et al. (2020) our result applies for general $K \geq 1$ and implies a linear scaling with the number K of possible outcomes that can be selected by the user. In fact, our bound suggests our algorithm on a K -multinomial problem has performance of same order as the performance of the algorithm of Faury et al. (2020) for a binary problem with dimension Kd instead of d . On the one hand, this is intuitive since the MNL reward model indeed involves Kd unknown parameters. Thus, we cannot expect regret better than $\mathcal{O}(Kd)$ for Kd unknown parameters. On the other hand, the MNL is a special case of multi-index models rather than a GLM. Thus, it is a-priori unclear whether it differs from a binary logistic model with Kd parameters in terms of regret performance. In fact, our proof does *not* treat the MNL reward model as a GLM. Despite that, it results to the linear order $\mathcal{O}(K)$ same as the optimal order of a binary logistic setting with Kd unknown parameters. Finally, as previously mentioned, the parameter κ defined in Assumption 2 generalizes the corresponding parameter for binary bandits. As in Theorem 2 of Faury et al. (2020) our bound for the multinomial case scales with $\sqrt{\kappa}$. Next, we show how this scaling is non-trivial and improves upon standard approaches.

Remark 2. Using same tools as for GLM-UCB in Filippi et al. (2010) for single-index reward models ($K = 1$), we can define the following alternative confidence set around the parameter $\boldsymbol{\theta}_*$ of the MNL model:

$$\mathcal{E}_t(\delta) := \{\boldsymbol{\theta} \in \Theta : \|\boldsymbol{\theta} - \tilde{\boldsymbol{\theta}}_t\|_{\tilde{\mathbf{V}}_t^{-1}} \leq \kappa\gamma_t(\delta)\}, \quad (23)$$

where $\gamma_t(\delta)$ is a slowly increasing function of t with similar order as $\beta_t(\delta)$ (see Lemma 14 in appendix), $\tilde{\mathbf{V}}_t = I_K \otimes \mathbf{V}_t$, and $\tilde{\boldsymbol{\theta}}_t := \arg \min_{\boldsymbol{\theta} \in \Theta} \|\mathbf{g}_t(\boldsymbol{\theta}) - \mathbf{g}_t(\tilde{\boldsymbol{\theta}}_t)\|_{\tilde{\mathbf{V}}_t^{-1}}$. Due to the appearance of an extra κ factor above compared to (12), relying on $\mathcal{E}_t(\delta)$ (23) in our analysis would lead to the following error bound (see Appendix C.2):

$$\Delta(\mathbf{x}, \tilde{\boldsymbol{\theta}}_t) \leq \tilde{\epsilon}_t(\mathbf{x}) := 2RL\kappa\gamma_t(\delta)\|\mathbf{x}\|_{\mathbf{V}_t^{-1}}. \quad (24)$$

This bound is significantly looser compared to our bound in (22) since the parameter κ can become arbitrarily large depending on the size of set $\mathcal{D} \times \Theta$.

2.6 Improved MNL-UCB

The regret of MNL-UCB scales as $\mathcal{O}(\sqrt{\kappa})$ with respect to κ . A more careful treatment of $\mathcal{C}_t(\delta)$ in (12) leads to Improved MNL-UCB, a modification of MNL-UCB, with a newly defined estimator and improved exploration bonus. Compared to MNL-UCB, the new algorithm is computationally intractable. Yet, it is interesting from a theory perspective as we show a regret bound where the κ -dependence is pushed into a quickly-vanishing second order term, trading-off a slight increase in the regret's dependence on K . In particular, we prove improved regret of order $\mathcal{O}(K^{1.5}d\sqrt{T}\log(T) + \kappa K^2 d \log(T))$. Due to space limitations, we defer a detailed description (see Algorithm 2) and analysis (see Theorem 3) of Improved MNL-UCB in Appendix D.

3 Discussion on κ and L

Knowledge of the problem-dependent constants κ and L is required to implement MNL-UCB as they appear in the definition of $\epsilon_t(\mathbf{x})$ in (22). While specifying their true values (defined in Assumption 2) requires solving non-convex optimization problems in general, here we present computationally efficient ways to obtain simple upper bounds. Also, we show that κ can indeed scale poorly with the size of $\mathcal{D} \times \Theta$ by deriving an appropriate lower bound for it. The upper bound on L is rather straightforward and it can be easily checked that $L \leq \max_{\mathbf{x} \in \mathcal{D}} \frac{e^{S\|\mathbf{x}\|_2}}{1 + e^{S\|\mathbf{x}\|_2} + (K-1)e^{-S\|\mathbf{x}\|_2}}$. Next, we focus on upper/lower bounding κ , which is more interesting.

In (26), we will show that κ scales unfavorably with the size of the set $\mathcal{D} \times \Theta$. In our regret analysis in the previous sections, it was useful to assume in Assumption 1 that $\|\mathbf{x}\|_2 \leq 1$. This assumption was without loss of generality because of the following. Suppose a general setting with $X = \max_{\mathbf{x} \in \mathcal{D}} \|\mathbf{x}\|_2 > 1$. We can then define an equivalent MNL model with actions $\tilde{\mathbf{x}} = \mathbf{x}/X$ and new parameter vector $\tilde{\boldsymbol{\theta}} = X\boldsymbol{\theta}$. The new problem satisfies the unit norm constraint on the radius of \mathcal{D} and has a new parameter norm bound $\tilde{S} = SX$. For clarity with regards to the goal of this section, it is convenient to keep track of the radius of \mathcal{D} (rather than push it in S). Thus, we let $X := \max_{\mathbf{x} \in \mathcal{D}} \|\mathbf{x}\|_2$ (possibly large). We will prove that κ grows at least exponentially in SX , thus it can be very large if the action decision set is large.

In order to bound κ , we identify and take advantage of the following key property of $\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})$ stated as Lemma 5. Recall that a matrix $\mathbf{A} \in \mathbb{R}^{K \times K}$ is *strictly diagonally dominant* if each of its diagonal entries is greater than the sum of absolute values of all other entries in the corresponding row/column. We also need the definition of an *M-matrix*: A matrix \mathbf{A} is an *M-matrix* if all its off-diagonal entries are non-positive and the real parts of its eigenvalues are non-negative Berman and Plemmons (1994).

Lemma 5. *For any $\mathbf{x} \in \mathbb{R}^d$ and $\boldsymbol{\theta} \in \mathbb{R}^{Kd}$, the matrix $\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})$ in (10) is a strictly diagonally dominant M-matrix.*

This key observation (see Appendix C for the proof) allows us to use Theorem 1.1 in Tian and Huang (2010) that provides upper and lower bounds on the minimum eigenvalue of a strictly diagonally dominant M-matrix. Specifically, we find that for all \mathbf{x} and $\boldsymbol{\theta}$:

$$\min_{i \in [K]} \sum_{j=1}^K \mathbf{A}(\mathbf{x}, \boldsymbol{\theta})_{ij} \leq \lambda_{\min}(\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})) \leq \max_{i \in [K]} \sum_{j=1}^K \mathbf{A}(\mathbf{x}, \boldsymbol{\theta})_{ij}. \quad (25)$$

Starting from this and setting $X := \max_{\mathbf{x} \in \mathcal{D}} \|\mathbf{x}\|_2$ we show the following bounds (Appendix C.4):

$$e^{\frac{SX}{\sqrt{K}}} (1 + Ke^{-\frac{SX}{\sqrt{K}}})^2 \leq \kappa \leq e^{SX} (1 + Ke^{SX})^2. \quad (26)$$

4 Experiments

We present numerical simulations to complement and confirm our theoretical findings. In all experiments, we used the upper bound on κ in (26) to compute the exploration bonus $\epsilon_t(\mathbf{x})$.

We evaluate the performance of MNL-UCB on synthetic data. All the results shown depict averages over 20 realizations, for which we have chosen $\delta = 0.01$, $d = 2$, and $T = 1000$. We considered time-independent decision sets \mathcal{D} of 20 arms in \mathbb{R}^2 and the reward vector $\boldsymbol{\rho} = [1, \dots, K]^T$. Moreover,

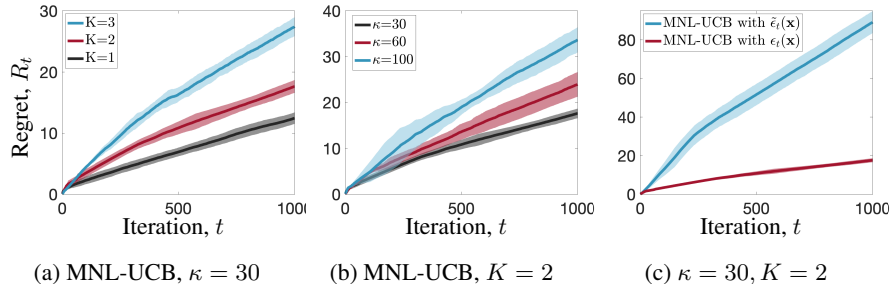


Figure 1: The shaded regions show standard deviation around the average over 20 problem realizations. See text for detailed description.

the arms and $\bar{\theta}_{*i}$ are drawn from $\mathcal{N}(0, I_d)$ and $\mathcal{N}(0, I_d/K)$, respectively. The normalization of the latter by K is so that it guarantees that the problem’s signal-to-noise ratio $\|\theta_*\|_2$ does not change with varying K . Figure 1a depicts the average regret of MNL-UCB for problem settings with different values of $K = 1, 2, 3$. The plot verifies that larger K leads to larger regret and seems to match the proved scaling of the regret bound of MNL-UCB as $\mathcal{O}(K)$ with respect to K . Figure 1b showcases the average regret of MNL-UCB for problem settings with fixed $K = 2$ and different values of $\kappa = 30, 60, 100$ (upper bounds computed using (26)). Observe that larger κ leads to larger regret. This is consistent with our theoretical findings on the impacts of κ on the algorithm’s performance. Finally, Figure 1c emphasizes the value of using the exploration bonus $\epsilon_t(\mathbf{x})$ in (22) compared to $\tilde{\epsilon}_t(\mathbf{x})$ introduced in (24) in the UCB decision making step. In this figure, we fixed $K = 2$ and the average regret curves are associated with a problem setting with $\kappa = 30$. A comparison between regret curves further confirms the worse regret performance of MNL-UCB when it exploits $\hat{\theta}_t$ and $\tilde{\epsilon}_t(\mathbf{x})$ rather than θ_t and $\epsilon_t(\mathbf{x})$ in the UCB decision rule at Line 3 of the Algorithm 1.

5 Conclusion

For the MNL regression bandit problem, we developed MNL-UCB and showed a regret $\tilde{\mathcal{O}}(Kd\sqrt{\kappa}\sqrt{T})$ that scales favorably with the critical problem-dependent parameter κ and optimally with respect to the number of options K . We further proposed Improved MNL-UCB that achieves a regret bound with problem-dependent constant κ being pushed into a second order logarithmic term, trading-off a slight increase in the regret’s dependence on K . After this work was completed, we became aware of Abeille et al. (2021) that improves Faury et al. (2020) in the binary case. Our work shows that extension of Faury et al. (2020) to MNL model is non-trivial and requires several careful analysis adjustments. It is an exciting future direction to investigate whether the algorithmic improvements of Abeille et al. (2021) can be properly adjusted in the multiclass case of interest here. In light of our discussion after Theorem 2 in Section 2.5, another promising future direction is to determine a lower bound that formally justifies the optimal dependence is $\mathcal{O}(K)$ or determine an improved algorithm/analysis that lowers the dependency on K . It is also interesting to study the efficacy of Thompson sampling-based algorithms for this new problem. Also, extending our results to other multi-index models is yet another important future direction.

Acknowledgments and Disclosure of Funding

This work is supported in part by the National Science Foundation under Grant Number (1934641). Part of this work was completed while the second author was affiliated with the University of California, Santa Barbara.

References

Abbasi-Yadkori, Y., Pál, D., and Szepesvári, C. (2011). Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*, pages 2312–2320.

- Abeille, M., Faury, L., and Calauzènes, C. (2021). Instance-wise minimax-optimal algorithms for logistic bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 3691–3699. PMLR.
- Abeille, M., Lazaric, A., et al. (2017). Linear thompson sampling revisited. *Electronic Journal of Statistics*, 11(2):5165–5197.
- Agrawal, P., Avadhanula, V., and Tulabandhula, T. (2020). A tractable online learning algorithm for the multinomial logit contextual bandit. *arXiv preprint arXiv:2011.14033*.
- Agrawal, S., Avadhanula, V., Goyal, V., and Zeevi, A. (2017). Thompson sampling for the mnl-bandit. In *Conference on Learning Theory*, pages 76–78.
- Agrawal, S., Avadhanula, V., Goyal, V., and Zeevi, A. (2019). Mnl-bandit: A dynamic learning approach to assortment selection. *Operations Research*, 67(5):1453–1485.
- Avner, O. and Mannor, S. (2019). Multi-user communication networks: A coordinated multi-armed bandit approach. *IEEE/ACM Transactions on Networking*, 27(6):2192–2207.
- Berkenkamp, F., Krause, A., and Schoellig, A. P. (2016). Bayesian optimization with safety constraints: safe and automatic parameter tuning in robotics. *arXiv preprint arXiv:1602.04450*.
- Berman, A. and Plemmons, R. J. (1994). *Nonnegative matrices in the mathematical sciences*. SIAM.
- Chen, X., Wang, Y., and Zhou, Y. (2018). Dynamic assortment optimization with changing contextual information. *arXiv preprint arXiv:1810.13069*.
- Cheung, W. C. and Simchi-Levi, D. (2017). Thompson sampling for online personalized assortment optimization problems with multinomial logit choice models. *Available at SSRN 3075658*.
- Dani, V., Hayes, T. P., and Kakade, S. M. (2008). Stochastic linear optimization under bandit feedback. In *Conference on Learning Theory*.
- Dong, K., Li, Y., Zhang, Q., and Zhou, Y. (2020). Multinomial logit bandit with low switching cost. *arXiv preprint arXiv:2007.04876*.
- Dong, S. and Van Roy, B. (2018). An information-theoretic analysis for thompson sampling with many actions. In *Advances in Neural Information Processing Systems*, pages 4157–4165.
- Faury, L., Abeille, M., Calauzènes, C., and Fercoq, O. (2020). Improved optimistic algorithms for logistic bandits. *arXiv preprint arXiv:2002.07530*.
- Filippi, S., Cappe, O., Garivier, A., and Szepesvári, C. (2010). Parametric bandits: The generalized linear case. In *Advances in Neural Information Processing Systems*, pages 586–594.
- Gamarnik, D. and Gaudio, J. (2020). Estimation of monotone multi-index models. *arXiv preprint arXiv:2006.02806*.
- Horn, R. A. and Johnson, C. R. (2012). *Matrix analysis*. Cambridge university press.
- Lattimore, T. and Szepesvári, C. (2020). *Bandit algorithms*. Cambridge University Press.
- Li, L., Lu, Y., and Zhou, D. (2017). Provably optimal algorithms for generalized linear contextual bandits. *arXiv preprint arXiv:1703.00048*.
- Li, S., Hao, F., Li, M., and Kim, H.-C. (2013). Medicine rating prediction and recommendation in mobile social networks. In *International conference on grid and pervasive computing*, pages 216–223. Springer.
- Oh, M.-h. and Iyengar, G. (2019). Thompson sampling for multinomial logit contextual bandits. In *Advances in Neural Information Processing Systems*, pages 3151–3161.
- Rusmevichientong, P. and Tsitsiklis, J. N. (2010). Linearly parameterized bandits. *Mathematics of Operations Research*, 35(2):395–411.

- Russo, D. and Van Roy, B. (2013). Eluder dimension and the sample complexity of optimistic exploration. In *Advances in Neural Information Processing Systems*, pages 2256–2264.
- Russo, D. and Van Roy, B. (2014). Learning to optimize via posterior sampling. *Mathematics of Operations Research*, 39(4):1221–1243.
- Sui, Y., Zhuang, V., Burdick, J. W., and Yue, Y. (2018). Stagewise safe bayesian optimization with gaussian processes. *arXiv preprint arXiv:1806.07555*.
- Sun, T. and Tran-Dinh, Q. (2019). Generalized self-concordant functions: a recipe for newton-type methods. *Mathematical Programming*, 178(1-2):145–213.
- Tian, G.-X. and Huang, T.-Z. (2010). Inequalities for the minimum eigenvalue of m-matrices. *The Electronic Journal of Linear Algebra*, 20.
- Tran-Dinh, Q., Li, Y.-H., and Cevher, V. (2015). Composite convex minimization involving self-concordant-like cost functions. In *Modelling, Computation and Optimization in Information Systems and Management Sciences*, pages 155–168. Springer.
- Wang, Y., Chen, X., and Zhou, Y. (2018). Near-optimal policies for dynamic multinomial logit assortment selection models. In *Advances in Neural Information Processing Systems*, pages 3101–3110.
- Yang, Z., Balasubramanian, K., Wang, Z., and Liu, H. (2017). Learning non-gaussian multi-index model via second-order stein’s method. *Advances in Neural Information Processing Systems*, 30:6097–6106.

Notation

Before presenting the proofs of our results, we recall some key notation for the reader's convenience.

First, recall that $\boldsymbol{\theta} = \begin{bmatrix} \bar{\boldsymbol{\theta}}_1 \\ \bar{\boldsymbol{\theta}}_2 \\ \vdots \\ \bar{\boldsymbol{\theta}}_K \end{bmatrix} \in \mathbb{R}^{Kd}$, where $\bar{\boldsymbol{\theta}}_i \in \mathbb{R}^d$.

MNL model.

$$\text{lse}(\mathbf{s}) := \log\left(1 + \sum_{i=1}^K \exp(\mathbf{s}_i)\right).$$

$$\mathbf{z}(\mathbf{x}, \boldsymbol{\theta}) = \nabla \text{lse} \left([\bar{\boldsymbol{\theta}}_1^T \mathbf{x}, \bar{\boldsymbol{\theta}}_2^T \mathbf{x}, \dots, \bar{\boldsymbol{\theta}}_K^T \mathbf{x}]^T \right) = \frac{1}{1 + \sum_{j=1}^K \exp(\bar{\boldsymbol{\theta}}_j^T \mathbf{x})} \begin{bmatrix} \exp(\bar{\boldsymbol{\theta}}_1^T \mathbf{x}) \\ \exp(\bar{\boldsymbol{\theta}}_2^T \mathbf{x}) \\ \vdots \\ \exp(\bar{\boldsymbol{\theta}}_K^T \mathbf{x}) \end{bmatrix}.$$

$$\mathbf{A}(\mathbf{x}, \boldsymbol{\theta}) = \nabla^2 \text{lse} \left([\bar{\boldsymbol{\theta}}_1^T \mathbf{x}, \bar{\boldsymbol{\theta}}_2^T \mathbf{x}, \dots, \bar{\boldsymbol{\theta}}_K^T \mathbf{x}]^T \right) = \text{diag}(\mathbf{z}(\mathbf{x}, \boldsymbol{\theta})) - \mathbf{z}(\mathbf{x}, \boldsymbol{\theta})\mathbf{z}(\mathbf{x}, \boldsymbol{\theta})^T.$$

Confidence set.

$$\mathbf{H}_t(\boldsymbol{\theta}) := \lambda I_{Kd} + \sum_{s=1}^{t-1} \mathbf{A}(\mathbf{x}_s, \boldsymbol{\theta}) \otimes \mathbf{x}_s \mathbf{x}_s^T$$

$$\mathbf{g}_t(\boldsymbol{\theta}) := \lambda \boldsymbol{\theta} + \sum_{s=1}^{t-1} \mathbf{z}(\mathbf{x}_s, \boldsymbol{\theta}) \otimes \mathbf{x}_s.$$

Recall the definition of $\nabla_{\boldsymbol{\theta}} \mathcal{L}_t^\lambda(\boldsymbol{\theta})$ in (8), which we repeat here for the reader's convenience:

$$\nabla_{\boldsymbol{\theta}} \mathcal{L}_t^\lambda(\boldsymbol{\theta}) := \sum_{s=1}^{t-1} [\mathbf{m}_s - \mathbf{z}(\mathbf{x}_s, \boldsymbol{\theta})] \otimes \mathbf{x}_s - \lambda \boldsymbol{\theta}.$$

By definition, $\hat{\boldsymbol{\theta}}_t$ satisfies $\nabla_{\boldsymbol{\theta}} \mathcal{L}_t^\lambda(\hat{\boldsymbol{\theta}}_t) = \mathbf{0}$.

$$\mathcal{C}_t(\delta) := \left\{ \boldsymbol{\theta} \in \Theta : \left\| \mathbf{g}_t(\boldsymbol{\theta}) - \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t) \right\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta})} \leq \beta_t(\delta) \right\}.$$

A Proof of Theorem 1

In this section, we present the proof of Theorem 1.

To see what is necessary to prove Theorem 1, we start with the following standard argument to analyze $\left\| \mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t) \right\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta}_*)}$. By definition, $\hat{\boldsymbol{\theta}}_t$ satisfies $\nabla_{\boldsymbol{\theta}} \mathcal{L}_t^\lambda(\hat{\boldsymbol{\theta}}_t) = \mathbf{0}$. Consequently, it holds that

$$\mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t) = \lambda \boldsymbol{\theta}_* + \mathbf{s}_t, \tag{27}$$

where we defined

$$\boldsymbol{\epsilon}_s := \mathbf{z}(\mathbf{x}_s, \boldsymbol{\theta}_*) - \mathbf{m}_s \quad \text{and} \quad \mathbf{s}_t := \sum_{s=1}^{t-1} \boldsymbol{\epsilon}_s \otimes \mathbf{x}_s. \tag{28}$$

Assumption 1 and simple linear algebra then imply that

$$\left\| \mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t) \right\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta}^*)} \leq \|\mathbf{s}_t\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta}^*)} + \sqrt{\lambda}S. \quad (29)$$

Therefore, establishing an appropriate confidence set $\mathcal{C}_t(\delta)$ requires us controlling $\|\mathbf{s}_t\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta}^*)}$. We do this in several steps organized in a series of key lemmas presented next. 1.

A.1 Necessary lemmas

The following is a matrix version of Lemma 7 in Faury et al. (2020) that uses a rather standard argument to bound the moment generating function of a random vector with bounded components in terms of its covariance matrix.

Lemma 6. *Let $\boldsymbol{\epsilon} \in \mathbb{R}^K$ be a zero-mean random vector with covariance matrix \mathbf{C} . Then, for any vector $\boldsymbol{\eta} \in \mathbb{R}^K$ such that $|\boldsymbol{\epsilon}^T \boldsymbol{\eta}| \leq 1$, we have*

$$\mathbb{E} \left[\exp(\boldsymbol{\epsilon}^T \boldsymbol{\eta}) \right] \leq \exp \left(\boldsymbol{\eta}^T \mathbf{C} \boldsymbol{\eta} \right). \quad (30)$$

Proof. Starting with Taylor expansion of the exponential, we have the following chain of inequalities

$$\begin{aligned} \mathbb{E} \left[\exp(\boldsymbol{\epsilon}^T \boldsymbol{\eta}) \right] &= 1 + \mathbb{E} \left[\boldsymbol{\epsilon}^T \boldsymbol{\eta} \right] + \sum_{k=2}^{\infty} \frac{\mathbb{E} \left[(\boldsymbol{\epsilon}^T \boldsymbol{\eta})^k \right]}{k!} \\ &= 1 + \sum_{k=2}^{\infty} \frac{\mathbb{E} \left[(\boldsymbol{\epsilon}^T \boldsymbol{\eta})^k \right]}{k!} && (\mathbb{E}[\boldsymbol{\epsilon}] = \mathbf{0}) \\ &\leq 1 + \sum_{k=2}^{\infty} \frac{\mathbb{E} \left[(\boldsymbol{\epsilon}^T \boldsymbol{\eta})^2 \right]}{k!} && (|\boldsymbol{\epsilon}^T \boldsymbol{\eta}| \leq 1) \\ &= 1 + \left(\boldsymbol{\eta}^T \mathbf{C} \boldsymbol{\eta} \right) \sum_{k=2}^{\infty} \frac{1}{k!} \leq 1 + \boldsymbol{\eta}^T \mathbf{C} \boldsymbol{\eta}, \\ &\leq \exp \left(\boldsymbol{\eta}^T \mathbf{C} \boldsymbol{\eta} \right). && (1 + x \leq \exp(x), \forall x \in \mathbb{R}) \end{aligned}$$

This completes the proof. \square

Hereafter, we will denote the p -dimensional unit ball by $\mathcal{B}_2(p) := \{\boldsymbol{\nu} \in \mathbb{R}^p : \|\boldsymbol{\nu}\|_2 \leq 1\}$.

Lemma 7. *For all $\boldsymbol{\xi} \in \frac{1}{\sqrt{K}}\mathcal{B}_2(Kd)$, let $M_0(\boldsymbol{\xi}) = 1$ and for $t > 1$*

$$M_t(\boldsymbol{\xi}) = \exp(\boldsymbol{\xi}^T \mathbf{s}_t - \|\boldsymbol{\xi}\|_{\bar{\mathbf{H}}_t(\boldsymbol{\theta}_*)}^2), \quad (31)$$

where $\bar{\mathbf{H}}_t(\boldsymbol{\theta}_*) := \sum_{s=1}^{t-1} \mathbf{A}(\mathbf{x}_s, \boldsymbol{\theta}_*) \otimes \mathbf{x}_s \mathbf{x}_s^T$. Then $\{M_t(\boldsymbol{\xi})\}_{t=1}^{\infty}$ is a non-negative super-martingale.

Proof. For $t > 1$, let $\boldsymbol{\eta}_t(\boldsymbol{\xi}) = [\bar{\boldsymbol{\xi}}_1^T \mathbf{x}_t, \dots, \bar{\boldsymbol{\xi}}_K^T \mathbf{x}_t]$. For all $\boldsymbol{\xi} \in \frac{1}{\sqrt{K}}\mathcal{B}_2(Kd)$ and $t > 1$, we have

$$\begin{aligned} \mathbb{E} \left[M_t(\boldsymbol{\xi}) | \mathcal{F}_{t-1} \right] &= \mathbb{E} \left[\exp(\boldsymbol{\xi}^T \mathbf{s}_t) | \mathcal{F}_{t-1} \right] \exp \left(-\|\boldsymbol{\xi}\|_{\bar{\mathbf{H}}_t(\boldsymbol{\theta}_*)}^2 \right) \\ &= \mathbb{E} \left[\exp(\boldsymbol{\xi}^T [\boldsymbol{\epsilon}_{t-1} \otimes \mathbf{x}_{t-1}]) | \mathcal{F}_{t-1} \right] \exp \left(\boldsymbol{\xi}^T \mathbf{s}_{t-1} - \|\boldsymbol{\xi}\|_{\bar{\mathbf{H}}_t(\boldsymbol{\theta}_*)}^2 \right) \\ &= \mathbb{E} \left[\exp \left(\boldsymbol{\epsilon}_{t-1}^T \boldsymbol{\eta}_{t-1}(\boldsymbol{\xi}) \right) | \mathcal{F}_{t-1} \right] \exp \left(\boldsymbol{\xi}^T \mathbf{s}_{t-1} - \|\boldsymbol{\xi}\|_{\bar{\mathbf{H}}_t(\boldsymbol{\theta}_*)}^2 \right). \end{aligned} \quad (32)$$

Note that $\mathbb{E}[\epsilon_t | \mathcal{F}_{t-1}] = \mathbf{0}$ and $\mathbb{E}[\epsilon_t \epsilon_t^T | \mathcal{F}_{t-1}] = \mathbf{A}(\mathbf{x}_t, \boldsymbol{\theta}_*)$ for all $t > 1$. Furthermore, $|\epsilon_{t-1}^T \boldsymbol{\eta}_{t-1}(\boldsymbol{\xi})| \leq 1$ because

$$\begin{aligned}
|\epsilon_{t-1}^T \boldsymbol{\eta}_{t-1}(\boldsymbol{\xi})| &\leq \|\epsilon_{t-1}\|_2 \|\boldsymbol{\eta}_{t-1}\|_2 && \text{(Cauchy-Schwarz)} \\
&\leq \sqrt{K} \sqrt{\sum_{i=1}^K (\bar{\boldsymbol{\xi}}_i^T \mathbf{x}_t)^2} && (|\epsilon_{t-1}|_i \leq 1) \\
&= \sqrt{K} \sqrt{\mathbf{x}_t^T \left(\sum_{i=1}^K \bar{\boldsymbol{\xi}}_i \bar{\boldsymbol{\xi}}_i^T \right) \mathbf{x}_t} \\
&\leq \|\mathbf{x}_t\|_2 \sqrt{K \lambda_{\max} \left(\sum_{i=1}^K \bar{\boldsymbol{\xi}}_i \bar{\boldsymbol{\xi}}_i^T \right)} \\
&\leq \sqrt{K \text{Tr} \left(\sum_{i=1}^K \bar{\boldsymbol{\xi}}_i \bar{\boldsymbol{\xi}}_i^T \right)} && \text{(Assumption 1)} \\
&= \sqrt{K \sum_{i=1}^K \|\bar{\boldsymbol{\xi}}_i\|_2^2} \\
&= \sqrt{K \|\boldsymbol{\xi}\|_2^2} \\
&\leq 1. && (\|\boldsymbol{\xi}\|_2 \leq \frac{1}{\sqrt{K}})
\end{aligned}$$

Thus, using Lemma 6, we conclude that

$$\begin{aligned}
\mathbb{E} \left[\exp \left(\epsilon_{t-1}^T \boldsymbol{\eta}_{t-1}(\boldsymbol{\xi}) \right) | \mathcal{F}_{t-1} \right] &\leq 1 + \boldsymbol{\eta}_{t-1}^T(\boldsymbol{\xi}) \mathbf{A}(\mathbf{x}_{t-1}, \boldsymbol{\theta}_*) \boldsymbol{\eta}_{t-1}(\boldsymbol{\xi}) \\
&\leq \exp \left(\boldsymbol{\xi}^T \left(\mathbf{A}(\mathbf{x}_{t-1}, \boldsymbol{\theta}_*) \otimes [\mathbf{x}_{t-1} \mathbf{x}_{t-1}^T] \right) \boldsymbol{\xi} \right), \quad (33)
\end{aligned}$$

Combining (32) and (33), we have

$$\mathbb{E} [M_t(\boldsymbol{\xi}) | \mathcal{F}_{t-1}] \leq \exp \left(\boldsymbol{\xi}^T \mathbf{s}_{t-1} - \|\boldsymbol{\xi}\|_{\mathbf{H}_t(\boldsymbol{\theta}_*)}^2 + \boldsymbol{\xi}^T \left(\mathbf{A}(\mathbf{x}_{t-1}, \boldsymbol{\theta}_*) \otimes [\mathbf{x}_{t-1} \mathbf{x}_{t-1}^T] \right) \boldsymbol{\xi} \right) = M_{t-1}(\boldsymbol{\xi}), \quad (34)$$

as desired. \square

Now, let $h(\boldsymbol{\xi})$ be a probability measure with support on $\frac{1}{\sqrt{K}} \mathcal{B}_2(Kd)$ and define

$$\bar{M}_t = \int_{\boldsymbol{\xi} \in \frac{1}{\sqrt{K}} \mathcal{B}_2(Kd)} M_t(\boldsymbol{\xi}) dh(\boldsymbol{\xi}) = \int_{\boldsymbol{\xi} \in \frac{1}{\sqrt{K}} \mathcal{B}_2(Kd)} \exp(\boldsymbol{\xi}^T \mathbf{s}_t - \|\boldsymbol{\xi}\|_{\mathbf{H}_t(\boldsymbol{\theta}_*)}^2) dh(\boldsymbol{\xi}). \quad (35)$$

By Lemma 20.3 in Lattimore and Szepesvári (2020) \bar{M}_t is also a non-negative super-martingale and $\mathbb{E}[\bar{M}_0] = 1$. Let τ be the stopping point with respect to the filtration $\{\mathcal{F}_t\}_{t=0}^\infty$. By Lemma 8 Abbasi-Yadkori et al. (2011), \bar{M}_τ is well-defined and $\mathbb{E}[\bar{M}_\tau] \leq 1$. Thanks to Markov inequality, for any $\delta \in (0, 1)$, we have

$$\mathbb{P}(\bar{M}_\tau \geq \frac{1}{\delta}) \leq \delta \mathbb{E}[\bar{M}_\tau] \leq \delta. \quad (36)$$

Lemma 8. Let $h(\boldsymbol{\xi})$ be the density of an isotropic normal distribution with precision matrix (i.e. inverse covariance matrix) $2\lambda I_{Kd}$ that is truncated on $\frac{1}{\sqrt{K}} \mathcal{B}_2(Kd)$ and $g(\boldsymbol{\xi})$ be the density of the

normal distribution with precision matrix $2\mathbf{H}_t(\boldsymbol{\theta}_*)$ that is truncated on the ball $\frac{1}{2\sqrt{K}}\mathcal{B}_2(Kd)$. Denote the normalization constant of h and g by $Z(h)$ and $Z(g)$, respectively. Then

$$\mathbb{P}\left(\left\|\mathbf{s}_t\right\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta}_*)} \geq \frac{\sqrt{\lambda/K}}{2} + 2\sqrt{K/\lambda} \log\left(\frac{Z(h)}{\delta Z(g)}\right)\right) \leq \delta. \quad (37)$$

Proof. We prove this lemma using similar techniques employed in [Abbasi-Yadkori et al. \(2011\)](#); [Faury et al. \(2020\)](#). We first specify $Z(h)$ and $Z(g)$ for completeness as follows:

$$\begin{aligned} Z(h) &:= \int_{\{\boldsymbol{\xi} \in \mathbb{R}^{Kd} : \|\boldsymbol{\xi}\|_2 \leq \frac{1}{\sqrt{K}}\}} \exp\left(-\lambda\|\boldsymbol{\xi}\|_2^2\right) d\boldsymbol{\xi}, \\ Z(g) &:= \int_{\{\boldsymbol{\xi} \in \mathbb{R}^{Kd} : \|\boldsymbol{\xi}\|_2 \leq \frac{1}{2\sqrt{K}}\}} \exp\left(-\frac{1}{2}\|\boldsymbol{\xi}\|_{2\mathbf{H}_t(\boldsymbol{\theta}_*)}^2\right) d\boldsymbol{\xi}. \end{aligned}$$

By the definition of \bar{M}_t in (35) and recalling $\mathbf{H}_t(\boldsymbol{\theta}_*) = \lambda\mathbf{I}_{Kd} + \bar{\mathbf{H}}_t(\boldsymbol{\theta}_*)$, we have

$$\bar{M}_t = \frac{1}{Z(h)} \int_{\boldsymbol{\xi} \in \frac{1}{\sqrt{K}}\mathcal{B}_2(Kd)} \exp(\boldsymbol{\xi}^T \mathbf{s}_t - \|\boldsymbol{\xi}\|_{\mathbf{H}_t(\boldsymbol{\theta}_*)}^2) d\boldsymbol{\xi}. \quad (38)$$

Now, let $f(\boldsymbol{\xi}) = \boldsymbol{\xi}^T \mathbf{s}_t - \|\boldsymbol{\xi}\|_{\mathbf{H}_t(\boldsymbol{\theta}_*)}^2$ and $\boldsymbol{\xi}_* = \arg \max_{\|\boldsymbol{\xi}\|_2 \leq \frac{1}{2\sqrt{K}}} f(\boldsymbol{\xi})$. Since

$$f(\boldsymbol{\xi}) = f(\boldsymbol{\xi}_*) + (\boldsymbol{\xi} - \boldsymbol{\xi}_*)^T \nabla f(\boldsymbol{\xi}_*) - (\boldsymbol{\xi} - \boldsymbol{\xi}_*)^T \mathbf{H}_t(\boldsymbol{\theta}_*)(\boldsymbol{\xi} - \boldsymbol{\xi}_*), \quad (39)$$

we have

$$\begin{aligned} \bar{M}_t &= \frac{\exp(f(\boldsymbol{\xi}_*))}{Z(h)} \int_{\mathbb{R}^{Kd}} \mathbb{1}\left\{\|\boldsymbol{\xi}\|_2 \leq \frac{1}{\sqrt{K}}\right\} \exp\left((\boldsymbol{\xi} - \boldsymbol{\xi}_*)^T \nabla f(\boldsymbol{\xi}_*) - \|\boldsymbol{\xi} - \boldsymbol{\xi}_*\|_{\mathbf{H}_t(\boldsymbol{\theta}_*)}^2\right) d\boldsymbol{\xi} \\ &= \frac{\exp(f(\boldsymbol{\xi}_*))}{Z(h)} \int_{\mathbb{R}^{Kd}} \mathbb{1}\left\{\|\boldsymbol{\xi} + \boldsymbol{\xi}_*\|_2 \leq \frac{1}{\sqrt{K}}\right\} \exp\left(\boldsymbol{\xi}^T \nabla f(\boldsymbol{\xi}_*) - \|\boldsymbol{\xi}\|_{\mathbf{H}_t(\boldsymbol{\theta}_*)}^2\right) d\boldsymbol{\xi} \\ &\hspace{20em} \text{(change of variable)} \\ &\geq \frac{\exp(f(\boldsymbol{\xi}_*))}{Z(h)} \int_{\mathbb{R}^{Kd}} \mathbb{1}\left\{\|\boldsymbol{\xi}\|_2 \leq \frac{1}{2\sqrt{K}}\right\} \exp\left(\boldsymbol{\xi}^T \nabla f(\boldsymbol{\xi}_*) - \|\boldsymbol{\xi}\|_{\mathbf{H}_t(\boldsymbol{\theta}_*)}^2\right) d\boldsymbol{\xi} \quad (\|\boldsymbol{\xi}_*\|_2 \leq \frac{1}{2\sqrt{K}}) \\ &= \exp(f(\boldsymbol{\xi}_*)) \frac{Z(g)}{Z(h)} \mathbb{E}_g \left[\exp\left(\boldsymbol{\xi}^T \nabla f(\boldsymbol{\xi}_*)\right) \right] \hspace{10em} \text{(definition of } g(\boldsymbol{\xi})\text{)} \\ &\geq \exp(f(\boldsymbol{\xi}_*)) \frac{Z(g)}{Z(h)} \exp\left(\mathbb{E}_g \left[\boldsymbol{\xi}^T \nabla f(\boldsymbol{\xi}_*) \right]\right) \hspace{10em} \text{(Jensen's inequality)} \\ &= \exp(f(\boldsymbol{\xi}_*)) \frac{Z(g)}{Z(h)}, \hspace{15em} (40) \end{aligned}$$

where in the last equality we used the fact that $\mathbb{E}_g[\boldsymbol{\xi}] = \mathbf{0}$. Now, let $\boldsymbol{\xi}_0 = \frac{\mathbf{H}_t^{-1}(\boldsymbol{\theta}_*)\mathbf{s}_t}{\|\mathbf{s}_t\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta}_*)}} \times \frac{\sqrt{\lambda/K}}{2}$.

Hence

$$\|\boldsymbol{\xi}_0\|_2 = \frac{\sqrt{\lambda/K}}{2} \cdot \frac{\|\mathbf{s}_t\|_{\mathbf{H}_t^{-2}(\boldsymbol{\theta}_*)}}{\|\mathbf{s}_t\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta}_*)}} \leq \frac{\sqrt{\lambda/K}}{2} \cdot \frac{1}{\sqrt{\lambda}} = \frac{1}{2\sqrt{K}}. \quad (41)$$

Combining (36) and (40), we conclude that

$$\begin{aligned} \mathbb{P}\left(\boldsymbol{\xi}_0^T \mathbf{s}_t - \|\boldsymbol{\xi}_0\|_{\mathbf{H}_t(\boldsymbol{\theta}_*)}^2 \geq \log\left(\frac{1}{\delta}\right) + \log\left(\frac{Z(h)}{Z(g)}\right)\right) &\leq \mathbb{P}\left(\max_{\|\boldsymbol{\xi}\|_2 \leq \frac{1}{2\sqrt{K}}} \boldsymbol{\xi}^T \mathbf{s}_t - \|\boldsymbol{\xi}\|_{\mathbf{H}_t(\boldsymbol{\theta}_*)}^2 \geq \log\left(\frac{1}{\delta}\right) + \log\left(\frac{Z(h)}{Z(g)}\right)\right) \\ &= \mathbb{P}\left(\exp(f(\boldsymbol{\xi}_*)) \frac{Z(g)}{Z(h)} \geq \frac{1}{\delta}\right) \\ &\leq \mathbb{P}\left(\bar{M}_t \geq \frac{1}{\delta}\right) \leq \delta. \hspace{10em} (42) \end{aligned}$$

Thus, by the definition of ξ_0 , we have

$$\mathbb{P}\left(\frac{\sqrt{\lambda/K}}{2}\|\mathbf{s}_t\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta}_*)} \geq \frac{\lambda}{4K} + \log\left(\frac{Z(h)}{\delta Z(g)}\right)\right) = \mathbb{P}\left(\|\mathbf{s}_t\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta}_*)} \geq \frac{\sqrt{\lambda/K}}{2} + 2\sqrt{K/\lambda}\log\left(\frac{Z(h)}{\delta Z(g)}\right)\right) \leq \delta, \quad (43)$$

as desired. \square

We use the following lemma borrowed from [Fauy et al. \(2020\)](#) to complete the proof of Theorem 1.

Lemma 9 (Lemma 6 in [Fauy et al. \(2020\)](#)). *It holds that*

$$\log\left(\frac{Z(h)}{\delta Z(g)}\right) \leq \log\left(\frac{\det(\mathbf{H}_t^{1/2}(\boldsymbol{\theta}_*))}{\lambda^{Kd/2}}\right) + Kd \log(2/\delta) \quad (44)$$

A.2 Completing the proof of Theorem 1

We are now ready to complete the proof of Theorem 1 combining the key results presented above as follows:

$$\begin{aligned} 1 - \delta &\leq \mathbb{P}\left(\|\mathbf{s}_t\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta}_*)} + \sqrt{\lambda}S \leq \frac{\sqrt{\lambda/K}}{2} + 2\sqrt{K/\lambda}\log\left(\frac{Z(h)}{\delta Z(g)}\right) + \sqrt{\lambda}S\right) && \text{(Lemma 8)} \\ &\leq \mathbb{P}\left(\|\mathbf{s}_t\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta}_*)} + \sqrt{\lambda}S \leq \frac{\sqrt{\lambda/K}}{2} + 2\sqrt{K/\lambda}\log\left(\frac{\det(\mathbf{H}_t^{1/2}(\boldsymbol{\theta}_*))}{\lambda^{Kd/2}}\right) + \frac{2K^{3/2}d}{\sqrt{\lambda}}\log(2/\delta) + \sqrt{\lambda}S\right) && \text{(Lemma 9)} \\ &\leq \mathbb{P}\left(\left\|\mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t)\right\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta}_*)} \leq \frac{\sqrt{\lambda/K}}{2} + 2\sqrt{K/\lambda}\log\left(\frac{\det(\mathbf{H}_t^{1/2}(\boldsymbol{\theta}_*))}{\lambda^{Kd/2}}\right) + \frac{2K^{3/2}d}{\sqrt{\lambda}}\log(2/\delta) + \sqrt{\lambda}S\right) && \text{(Eqn. (29))} \\ &\leq \mathbb{P}\left(\left\|\mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t)\right\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta}_*)} \leq \frac{K^{3/2}d}{\sqrt{\lambda}}\log\left(1 + \frac{t}{d\lambda}\right) + \frac{\sqrt{\lambda/K}}{2} + \frac{2K^{3/2}d}{\sqrt{\lambda}}\log(2/\delta) + \sqrt{\lambda}S\right), && (45) \end{aligned}$$

where the last inequality follows from

$$\det(\mathbf{H}_t(\boldsymbol{\theta}_*)) = \prod_{i=1}^{Kd} \lambda_i(\mathbf{H}_t(\boldsymbol{\theta}_*)) \leq \left(\text{trace}(\mathbf{H}_t(\boldsymbol{\theta}_*)) / Kd\right)^{Kd}, \quad (46)$$

and

$$\text{trace}(\mathbf{H}_t(\boldsymbol{\theta}_*)) \leq \lambda Kd + \sum_{s=1}^{t-1} \sum_{i=1}^K \mathbf{A}(\mathbf{x}_s, \boldsymbol{\theta}_*)_{ii} \|\mathbf{x}_s\|_2^2 \leq \lambda Kd + Kt. \quad \text{(Assumption 1)}$$

B Proofs of Lemma 1 and Theorem 2

As mentioned in Section 2.3, we organize the proof of Lemma 1 in several key lemmas which we prove below

Conditioning on $\boldsymbol{\theta}_* \in \mathcal{C}_t(\delta)$, $\forall t \in [T]$. Consider the event

$$\mathcal{E} := \{\boldsymbol{\theta}_* \in \mathcal{C}_t(\delta), \forall t \in [T]\}, \quad (47)$$

that $\boldsymbol{\theta}_*$ is inside the confidence sets for all rounds $t \in [T]$. By Theorem 1 the event holds with probability at least $1 - \delta$. Onwards, we condition on this event, and make repeated use of the fact that $\boldsymbol{\theta}_* \in \mathcal{C}_t$ for all $t \in [T]$, without further explicit reference.

B.1 Necessary lemmas

Here, we prove Lemmas 2, 3, and 4 of Section 2.3. For the reader's convenience, we repeat their statements as Lemmas 10, 11, and 13 below.

Lemma 10. Recall the definition of the lse function: $\text{lse}(\mathbf{s}) := \log(1 + \sum_{i=1}^K \exp(\mathbf{s}_i))$ and let $\boldsymbol{\mu} : \mathbb{R}^K \rightarrow \mathbb{R}^K$ be defined as

$$\boldsymbol{\mu}(\mathbf{s}) := \nabla \text{lse}(\mathbf{s}) = \frac{1}{1 + \sum_{i=1}^K \exp(\mathbf{s}_i)} \begin{bmatrix} \exp(\mathbf{s}_1) \\ \vdots \\ \exp(\mathbf{s}_K) \end{bmatrix}. \quad (48)$$

Moreover, for $\mathbf{x} \in \mathbb{R}^d$ and $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \mathbb{R}^{Kd}$ define

$$\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2) := \int_0^1 \nabla \boldsymbol{\mu} \left(\begin{bmatrix} v \bar{\boldsymbol{\theta}}_{11}^T \mathbf{x} + (1-v) \bar{\boldsymbol{\theta}}_{21}^T \mathbf{x} \\ v \bar{\boldsymbol{\theta}}_{12}^T \mathbf{x} + (1-v) \bar{\boldsymbol{\theta}}_{22}^T \mathbf{x} \\ \vdots \\ v \bar{\boldsymbol{\theta}}_{1K}^T \mathbf{x} + (1-v) \bar{\boldsymbol{\theta}}_{2K}^T \mathbf{x} \end{bmatrix} \right) dv = \int_0^1 \mathbf{A}(\mathbf{x}, v\boldsymbol{\theta}_1 + (1-v)\boldsymbol{\theta}_2) dv, \quad (49)$$

where the matrix $\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})$ is given in (10). Then, it holds that

$$\mathbf{z}(\mathbf{x}, \boldsymbol{\theta}_1) - \mathbf{z}(\mathbf{x}, \boldsymbol{\theta}_2) = \left[\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \otimes \mathbf{x}^T \right] (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2). \quad (50)$$

Remark 3 (Notation). In the notation of the lemma, notice that our vector function $\mathbf{z}(\mathbf{x}, \boldsymbol{\theta})$ defined in (2) can be expressed in terms of $\boldsymbol{\mu}$ as:

$$\mathbf{z}(\mathbf{x}, \boldsymbol{\theta}) = \boldsymbol{\mu}([\bar{\boldsymbol{\theta}}_1^T \mathbf{x}, \dots, \bar{\boldsymbol{\theta}}_K^T \mathbf{x}]^T).$$

Similarly,

$$\mathbf{A}(\mathbf{x}, \boldsymbol{\theta}) = \nabla \boldsymbol{\mu}([\bar{\boldsymbol{\theta}}_1^T \mathbf{x}, \dots, \bar{\boldsymbol{\theta}}_K^T \mathbf{x}]^T).$$

Proof. According to *mean-value Theorem*, for a differentiable function $\mathbf{f} : \mathbb{R}^P \rightarrow \mathbb{R}^Q$ and differentiable function $\mathbf{q} : \mathbb{R} \rightarrow \mathbb{R}^P$, we have

$$\int_a^b \nabla \mathbf{f}^T(\mathbf{q}(v)) \mathbf{q}'(v) dv = \mathbf{f}(\mathbf{q}(b)) - \mathbf{f}(\mathbf{q}(a)). \quad (51)$$

Now, let $\mathbf{f}(\boldsymbol{\theta}) = \mathbf{z}(\mathbf{x}, \boldsymbol{\theta})$, $a = 0$, $b = 1$ and $\mathbf{q}(v) = v\boldsymbol{\theta}_1 + (1-v)\boldsymbol{\theta}_2$. By the mean-value theorem, we have

$$\mathbf{z}(\mathbf{x}, \boldsymbol{\theta}_1) - \mathbf{z}(\mathbf{x}, \boldsymbol{\theta}_2) = \left[\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \otimes \mathbf{x}^T \right] (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2) = \left[\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \otimes \mathbf{x}^T \right] (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2), \quad (52)$$

where the last equality follows from the fact that $(\mathbf{A} \otimes \mathbf{B})^T = \mathbf{A}^T \otimes \mathbf{B}^T$ for any matrices \mathbf{A} and \mathbf{B} and matrix $\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \in \mathbb{R}^{K \times K}$ being a symmetric matrix. \square

Lemma 11. Define

$$\mathbf{G}_t(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2) := \lambda I_{Kd} + \sum_{s=1}^{t-1} \mathbf{B}(\mathbf{x}_s, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \otimes \mathbf{x}_s \mathbf{x}_s^T. \quad (53)$$

Then, for any $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \mathbb{R}^{Kd}$, we have

$$\mathbf{g}_t(\boldsymbol{\theta}_1) - \mathbf{g}_t(\boldsymbol{\theta}_2) = \mathbf{G}_t(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2) (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2). \quad (54)$$

Proof. By the definition of $\mathbf{g}_t(\boldsymbol{\theta})$ in (11) we have

$$\begin{aligned}
\mathbf{g}_t(\boldsymbol{\theta}_1) - \mathbf{g}_t(\boldsymbol{\theta}_2) &= \lambda(\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2) + \sum_{s=1}^{t-1} [\mathbf{z}(\mathbf{x}_s, \boldsymbol{\theta}_1) - \mathbf{z}(\mathbf{x}_s, \boldsymbol{\theta}_2)] \otimes \mathbf{x}_s \\
&= \lambda(\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2) + \sum_{s=1}^{t-1} \left[\left[\mathbf{B}(\mathbf{x}_s, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \otimes \mathbf{x}_s^T \right] (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2) \right] \otimes \mathbf{x}_s \quad (\text{Eqn. (50)}) \\
&= \lambda(\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2) + \sum_{s=1}^{t-1} \left[\left[\mathbf{B}(\mathbf{x}_s, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \otimes \mathbf{x}_s^T \right] \otimes \mathbf{x}_s \right] (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2) \\
&\hspace{15em} (\text{mixed-product property}) \\
&= \lambda(\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2) + \sum_{s=1}^{t-1} \left[\mathbf{B}(\mathbf{x}_s, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \otimes \left[\mathbf{x}_s^T \otimes \mathbf{x}_s \right] \right] (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2) \\
&= \lambda(\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2) + \sum_{s=1}^{t-1} \left[\mathbf{B}(\mathbf{x}_s, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \otimes \mathbf{x}_s \mathbf{x}_s^T \right] (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2) = \mathbf{G}_t(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)(\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2),
\end{aligned}$$

as desired. \square

Lemma 12. *Define*

$$\tilde{\mathbf{V}}_t := \kappa \lambda I_{Kd} + \sum_{s=1}^{t-1} \sum_{i=1}^K (\mathbf{e}_i \otimes \mathbf{x}_s)(\mathbf{e}_i \otimes \mathbf{x}_s)^T = I_K \otimes \mathbf{V}_t. \quad (55)$$

Then, for any $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \Theta$, we have

$$\frac{1}{\kappa} \tilde{\mathbf{V}}_t \preceq \mathbf{G}_t(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2).$$

Proof. By definition of $\mathbf{G}_t(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$ for any $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \Theta$ in (53), we have

$$\begin{aligned}
\mathbf{G}_t(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2) &= \lambda I_{Kd} + \sum_{s=1}^{t-1} \mathbf{B}(\mathbf{x}_s, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \otimes \mathbf{x}_s \mathbf{x}_s^T \\
&\succeq \frac{1}{\kappa} \left[\kappa \lambda I_{Kd} + \sum_{s=1}^{t-1} I_K \otimes \mathbf{x}_s \mathbf{x}_s^T \right] \quad (\text{Assumption 2}) \\
&= \frac{1}{\kappa} \tilde{\mathbf{V}}_t. \quad (56)
\end{aligned}$$

We derived the inequality above as follows. Assumption 2 and definition of $\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$ in (49) give $\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \succeq \frac{1}{\kappa} I_K$. Now, from standard spectral properties of the Kronecker product, the eigenvalues of the matrix $\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \otimes \mathbf{x}_s \mathbf{x}_s^T - \frac{1}{\kappa} I_K \otimes \mathbf{x}_s \mathbf{x}_s^T = (\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2) - \frac{1}{\kappa} I_K) \otimes \mathbf{x}_s \mathbf{x}_s^T$ are the pairwise products of the individual eigenvalues of $\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2) - \frac{1}{\kappa} I_K$ and of $\mathbf{x}_s \mathbf{x}_s^T$. The desired inequality then follows since both of the latter two matrices are positive semi-definite. \square

Lemma 13 (Generalized self-concordance). *For any $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \Theta$, we have*

$$(1 + 2S)^{-1} \mathbf{H}_t(\boldsymbol{\theta}_1) \preceq \mathbf{G}_t(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \quad \text{and} \quad (1 + 2S)^{-1} \mathbf{H}_t(\boldsymbol{\theta}_2) \preceq \mathbf{G}_t(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2). \quad (57)$$

Proof. Recall the function $\text{lse} : \mathbb{R}^K \rightarrow \mathbb{R}$

$$\text{lse}(\mathbf{s}) := \log \left(1 + \sum_{i=1}^K \exp(\mathbf{s}_i) \right). \quad (58)$$

Then its gradient and Hessian are written as

$$\nabla \text{lse}(\mathbf{s}) = \boldsymbol{\mu}(\mathbf{s}) := \frac{1}{1 + \sum_{i=1}^K \exp(\mathbf{s}_i)} \begin{bmatrix} \exp(\mathbf{s}_1) \\ \vdots \\ \exp(\mathbf{s}_K) \end{bmatrix} \quad \text{and} \quad \nabla^2 \text{lse}(\mathbf{s}) = \text{diag}(\mathbf{s}) - \mathbf{s} \mathbf{s}^T. \quad (59)$$

As shown in Lemma 4 in [Tran-Dinh et al. \(2015\)](#), the function lse is (M_{lse}, ν) -generalized self-concordant with $\nu = 1$ and $M_{\text{lse}} = \sqrt{6}$. Therefore, according to Corollary 2 in [Sun and Tran-Dinh \(2019\)](#), for any $\mathbf{x} \in \mathcal{D}$ and $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \Theta$, we have

$$\frac{1 - \exp(-d(\mathbf{x}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2))}{d(\mathbf{x}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2)} \nabla^2 f \left(\begin{bmatrix} \bar{\boldsymbol{\theta}}_{21}^T \mathbf{x} \\ \bar{\boldsymbol{\theta}}_{22}^T \mathbf{x} \\ \vdots \\ \bar{\boldsymbol{\theta}}_{2K}^T \mathbf{x} \end{bmatrix} \right) \preceq \int_0^1 \nabla^2 f \left(\begin{bmatrix} v \bar{\boldsymbol{\theta}}_{11}^T \mathbf{x} + (1-v) \bar{\boldsymbol{\theta}}_{21}^T \mathbf{x} \\ v \bar{\boldsymbol{\theta}}_{12}^T \mathbf{x} + (1-v) \bar{\boldsymbol{\theta}}_{22}^T \mathbf{x} \\ \vdots \\ v \bar{\boldsymbol{\theta}}_{1K}^T \mathbf{x} + (1-v) \bar{\boldsymbol{\theta}}_{2K}^T \mathbf{x} \end{bmatrix} \right) dv, \quad (60)$$

where $d(\mathbf{x}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2) = \left\| \begin{bmatrix} \bar{\boldsymbol{\theta}}_{11}^T \mathbf{x} - \bar{\boldsymbol{\theta}}_{21}^T \mathbf{x}, \dots, \bar{\boldsymbol{\theta}}_{1K}^T \mathbf{x} - \bar{\boldsymbol{\theta}}_{2K}^T \mathbf{x} \end{bmatrix}^T \right\|_2$. Now, by the definition of $\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})$ and $\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$ in (10) and (49), we conclude from (60) that

$$\begin{aligned} \mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2) &\succeq (1 + d(\mathbf{x}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2))^{-1} \mathbf{A}(\mathbf{x}, \boldsymbol{\theta}_2) && \left(\frac{1 - \exp(-t)}{t} \geq (1+t)^{-1} \right) \\ &\succeq (1 + 2S)^{-1} \mathbf{A}(\mathbf{x}, \boldsymbol{\theta}_2), && (61) \end{aligned}$$

where the last inequality follows from Assumption 1 and the fact that both $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \Theta$. Combining (61) and the definition of $\mathbf{G}_t(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$ in (53) for any $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \Theta$, we have

$$\begin{aligned} \mathbf{G}_t(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2) &= \lambda I_{Kd} + \sum_{s=1}^{t-1} \mathbf{B}(\mathbf{x}_s, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \otimes \mathbf{x}_s \mathbf{x}_s^T \\ &\succeq (1 + 2S)^{-1} \left[\lambda I_{Kd} + \sum_{s=1}^{t-1} \mathbf{A}(\mathbf{x}_s, \boldsymbol{\theta}_2) \otimes \mathbf{x}_s \mathbf{x}_s^T \right] \\ &= (1 + 2S)^{-1} \mathbf{H}_t(\boldsymbol{\theta}_2). && \text{(by definition of } \mathbf{H}_t(\boldsymbol{\theta}_2) \text{ in Eqn. (9))} \end{aligned}$$

By symmetric roles of $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$ in the definition of $\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$, we can similarly prove that $(1 + 2S)^{-1} \mathbf{H}_t(\boldsymbol{\theta}_1) \preceq \mathbf{G}_t(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$. \square

B.2 Completing the proof of Lemma 1

In this section, we complete the proof of Lemma 1 using Lemmas 10, 11, 12 and 13.

By the definition of $\Delta(\mathbf{x}, \boldsymbol{\theta})$, for all $\mathbf{x} \in \mathcal{D}$, $t \in [T]$ and $\boldsymbol{\theta} \in \mathcal{C}_t(\delta)$, we have the following chain of inequalities

$$\begin{aligned}
\Delta(\mathbf{x}, \boldsymbol{\theta}) &= |\boldsymbol{\rho}^T \mathbf{z}(\mathbf{x}, \boldsymbol{\theta}_*) - \boldsymbol{\rho}^T \mathbf{z}(\mathbf{x}, \boldsymbol{\theta})| \\
&\leq R \|\mathbf{z}(\mathbf{x}, \boldsymbol{\theta}_*) - \mathbf{z}(\mathbf{x}, \boldsymbol{\theta})\|_2 && \text{(Assumption 1 and Cauchy-Schwarz inequality)} \\
&= R \left\| \left[\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_*, \boldsymbol{\theta}) \otimes \mathbf{x}^T \right] (\boldsymbol{\theta}_* - \boldsymbol{\theta}) \right\|_2 && \text{(Lemma 10)} \\
&= R \left\| \left[\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_*, \boldsymbol{\theta}) \otimes \mathbf{x}^T \right] \mathbf{G}_t^{-1/2}(\boldsymbol{\theta}_*, \boldsymbol{\theta}) \mathbf{G}_t^{1/2}(\boldsymbol{\theta}_*, \boldsymbol{\theta}) (\boldsymbol{\theta}_* - \boldsymbol{\theta}) \right\|_2 \\
&\leq R \left\| \left[\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_*, \boldsymbol{\theta}) \otimes \mathbf{x}^T \right] \mathbf{G}_t^{-1/2}(\boldsymbol{\theta}_*, \boldsymbol{\theta}) \right\|_2 \|\boldsymbol{\theta}_* - \boldsymbol{\theta}\|_{\mathbf{G}_t(\boldsymbol{\theta}_*, \boldsymbol{\theta})} \\
&&& \text{(Cauchy-Schwarz inequality)} \\
&= R \sqrt{\lambda_{\max} \left(\left[\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_*, \boldsymbol{\theta}) \otimes \mathbf{x}^T \right] \mathbf{G}_t^{-1}(\boldsymbol{\theta}_*, \boldsymbol{\theta}) \left[\mathbf{B}^T(\mathbf{x}, \boldsymbol{\theta}_*, \boldsymbol{\theta}) \otimes \mathbf{x} \right] \right)} \|\mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\boldsymbol{\theta})\|_{\mathbf{G}_t^{-1}(\boldsymbol{\theta}_*, \boldsymbol{\theta})} \\
&&& \text{(Lemma 11)} \\
&= R \sqrt{\lambda_{\max} \left(\mathbf{G}_t^{-1/2}(\boldsymbol{\theta}_*, \boldsymbol{\theta}) \left[\mathbf{B}^T(\mathbf{x}, \boldsymbol{\theta}_*, \boldsymbol{\theta}) \otimes \mathbf{x} \right] \left[\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_*, \boldsymbol{\theta}) \otimes \mathbf{x}^T \right] \mathbf{G}_t^{-1/2}(\boldsymbol{\theta}_*, \boldsymbol{\theta}) \right)} \|\mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\boldsymbol{\theta})\|_{\mathbf{G}_t^{-1}(\boldsymbol{\theta}_*, \boldsymbol{\theta})} \\
&&& \text{(cyclic property of } \lambda_{\max} \text{)} \\
&= R \sqrt{\lambda_{\max} \left(\mathbf{G}_t^{-1/2}(\boldsymbol{\theta}_*, \boldsymbol{\theta}) \left[\mathbf{B}^T(\mathbf{x}, \boldsymbol{\theta}_*, \boldsymbol{\theta}) \mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_*, \boldsymbol{\theta}) \otimes \mathbf{x} \mathbf{x}^T \right] \mathbf{G}_t^{-1/2}(\boldsymbol{\theta}_*, \boldsymbol{\theta}) \right)} \|\mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\boldsymbol{\theta})\|_{\mathbf{G}_t^{-1}(\boldsymbol{\theta}_*, \boldsymbol{\theta})} \\
&&& \text{(mixed-product property)} \\
&\leq RL \sqrt{\lambda_{\max} \left(\mathbf{G}_t^{-1}(\boldsymbol{\theta}_*, \boldsymbol{\theta}) \left[I_K \otimes \mathbf{x} \mathbf{x}^T \right] \right)} \|\mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\boldsymbol{\theta})\|_{\mathbf{G}_t^{-1}(\boldsymbol{\theta}_*, \boldsymbol{\theta})} \\
&&& \text{(Assumption 2 (*))} \\
&= RL \sqrt{\lambda_{\max} \left(\mathbf{G}_t^{-1}(\boldsymbol{\theta}_*, \boldsymbol{\theta}) \left[I_K \otimes \mathbf{x} \right] \left[I_K \otimes \mathbf{x}^T \right] \right)} \|\mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\boldsymbol{\theta})\|_{\mathbf{G}_t^{-1}(\boldsymbol{\theta}_*, \boldsymbol{\theta})} \\
&&& \text{(mixed-product property)} \\
&= RL \sqrt{\lambda_{\max} \left(\left[I_K \otimes \mathbf{x}^T \right] \mathbf{G}_t^{-1}(\boldsymbol{\theta}_*, \boldsymbol{\theta}) \left[I_K \otimes \mathbf{x} \right] \right)} \|\mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\boldsymbol{\theta})\|_{\mathbf{G}_t^{-1}(\boldsymbol{\theta}_*, \boldsymbol{\theta})} \\
&&& \text{(cyclic property of } \lambda_{\max} \text{)} \\
&\leq RL \sqrt{\kappa \lambda_{\max} \left(\left[I_K \otimes \mathbf{x}^T \right] \left[I_K \otimes \mathbf{V}_t^{-1} \right] \left[I_K \otimes \mathbf{x} \right] \right)} \|\mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\boldsymbol{\theta})\|_{\mathbf{G}_t^{-1}(\boldsymbol{\theta}_*, \boldsymbol{\theta})} \\
&&& \text{(Lemma 12 (**))} \\
&= RL \sqrt{\kappa \lambda_{\max} \left(I_K \otimes \|\mathbf{x}\|_{\mathbf{V}_t^{-1}}^2 \right)} \|\mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\boldsymbol{\theta})\|_{\mathbf{G}_t^{-1}(\boldsymbol{\theta}_*, \boldsymbol{\theta})} \text{ (mixed-product property)} \\
&= RL \sqrt{\kappa} \|\mathbf{x}\|_{\mathbf{V}_t^{-1}} \left\| \mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t) + \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t) - \mathbf{g}_t(\boldsymbol{\theta}) \right\|_{\mathbf{G}_t^{-1}(\boldsymbol{\theta}_*, \boldsymbol{\theta})} \\
&\leq RL \sqrt{\kappa} \|\mathbf{x}\|_{\mathbf{V}_t^{-1}} \left[\left\| \mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t) \right\|_{\mathbf{G}_t^{-1}(\boldsymbol{\theta}_*, \boldsymbol{\theta})} + \left\| \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t) - \mathbf{g}_t(\boldsymbol{\theta}) \right\|_{\mathbf{G}_t^{-1}(\boldsymbol{\theta}_*, \boldsymbol{\theta})} \right] \\
&&& \left((I_K \otimes \mathbf{V}_t)^{-1} = I_K \otimes \mathbf{V}_t^{-1} \right) \\
&\leq RL \sqrt{\kappa(1+2S)} \|\mathbf{x}\|_{\mathbf{V}_t^{-1}} \left[\left\| \mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t) \right\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta}_*)} + \left\| \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t) - \mathbf{g}_t(\boldsymbol{\theta}) \right\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta})} \right] \\
&&& \text{(Lemma 13)} \\
&\leq 2RL\beta_t(\delta) \sqrt{\kappa(1+2S)} \|\mathbf{x}\|_{\mathbf{V}_t^{-1}}, && \text{(Theorem 1)}
\end{aligned}$$

as desired. Above, *cyclic property of λ_{\max}* boils down to the fact that for two matrices $\mathbf{M}_1, \mathbf{M}_2 \in \mathbb{R}^{d \times d}$ the eigenvalues of $\mathbf{M}_1 \mathbf{M}_2$ are the same as the eigenvalues of $\mathbf{M}_2 \mathbf{M}_1$ ²; thus, the same is true

²To see why this is true: let λ, \mathbf{v} be an eigenvalue-eigenvector pair of $\mathbf{M}_1 \mathbf{M}_2$, i.e. $\mathbf{M}_1 \mathbf{M}_2 \mathbf{v} = \lambda \mathbf{v}$. Then, $(\mathbf{M}_2 \mathbf{M}_1) \mathbf{M}_2 \mathbf{v} = \lambda \mathbf{M}_2 \mathbf{v}$. Thus, λ is also an eigenvalue of $\mathbf{M}_2 \mathbf{M}_1$.

of the maximum eigenvalue. For completeness, we also detail the derivation of the third in line inequality marked with (*) above. First, from an argument identical to the proof of Lemma 12, we have that $\mathbf{B}^T(\mathbf{x}, \boldsymbol{\theta}_*, \boldsymbol{\theta})\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_*, \boldsymbol{\theta}) \otimes \mathbf{x}\mathbf{x}^T \preceq L^2(I_K \otimes \mathbf{x}\mathbf{x}^T)$. From this and Theorem 7.7.2 in Horn and Johnson (2012) it follows that $\mathbf{G}_t^{-1/2}(\boldsymbol{\theta}_*, \boldsymbol{\theta})[\mathbf{B}^T(\mathbf{x}, \boldsymbol{\theta}_*, \boldsymbol{\theta})\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_*, \boldsymbol{\theta}) \otimes \mathbf{x}\mathbf{x}^T]\mathbf{G}_t^{-1/2}(\boldsymbol{\theta}_*, \boldsymbol{\theta}) \preceq L^2\mathbf{G}_t^{-1/2}(\boldsymbol{\theta}_*, \boldsymbol{\theta})[I_K \otimes \mathbf{x}\mathbf{x}^T]\mathbf{G}_t^{-1/2}(\boldsymbol{\theta}_*, \boldsymbol{\theta})$. The argument then follows directly by the applying the cyclic property of λ_{\max} .

B.3 Completing the proof of Theorem 2

Consider the standard instantaneous regret decomposition as follows:

$$\begin{aligned} r_t &= |\boldsymbol{\rho}^T \mathbf{z}(\mathbf{x}_*, \boldsymbol{\theta}_*) - \boldsymbol{\rho}^T \mathbf{z}(\mathbf{x}_t, \boldsymbol{\theta}_*)| \\ &\leq |\boldsymbol{\rho}^T \mathbf{z}(\mathbf{x}_*, \boldsymbol{\theta}_t) + \epsilon_t(\mathbf{x}_*) - \boldsymbol{\rho}^T \mathbf{z}(\mathbf{x}_t, \boldsymbol{\theta}_*)| \\ &\leq |\boldsymbol{\rho}^T \mathbf{z}(\mathbf{x}_t, \boldsymbol{\theta}_t) - \boldsymbol{\rho}^T \mathbf{z}(\mathbf{x}_t, \boldsymbol{\theta}_*) + \epsilon_t(\mathbf{x}_t)| \\ &\leq 2\epsilon_t(\mathbf{x}_t). \end{aligned} \tag{62}$$

Therefore

$$\begin{aligned} R_T &\leq 2 \sum_{t=1}^T \epsilon_t(\mathbf{x}_t) \\ &= 4RL\beta_T(\delta) \sqrt{\kappa(1+2S)} \sum_{t=1}^T \|\mathbf{x}_t\|_{\mathbf{V}_t^{-1}} \tag{Corollary 1} \\ &\leq 4RL\beta_T(\delta) \sqrt{\kappa(1+2S)} \sqrt{T \sum_{t=1}^T \|\mathbf{x}_t\|_{\mathbf{V}_t^{-1}}^2} \tag{Cauchy-Schwarz inequality} \\ &\leq 4RL\beta_T(\delta) \sqrt{2 \max(1, \frac{1}{\lambda\kappa}) \kappa(1+2S) dT \log \left(1 + \frac{T}{\kappa\lambda d}\right)}. \end{aligned} \tag{63}$$

In the last inequality, we used the standard argument in regret analysis of linear bandit algorithm stated in the following Abbasi-Yadkori et al. (2011) (Lemma 11):

$$\sum_{t=1}^n \min \left(\|\mathbf{y}_t\|_{\mathbf{V}_t^{-1}}^2, 1 \right) \leq 2 \log \frac{\det \mathbf{A}_{n+1}}{\det \mathbf{A}_1} \quad \text{where} \quad \mathbf{A}_n = \mathbf{A}_1 + \sum_{t=1}^{n-1} \mathbf{y}_t \mathbf{y}_t^T. \tag{64}$$

C Miscellaneous Useful results

C.1 A looser confidence set

Lemma 14. Let $\gamma_t(\delta) := 2 \left(\sqrt{\lambda S} + 2 \sqrt{\log(1/\delta) + Kd \log \left(1 + \frac{t}{\kappa\lambda d}\right)} \right)$. Then, it holds that $\boldsymbol{\theta}_* \in \mathcal{E}_t(\delta)$ with probability at least $1 - \delta$.

Proof.

$$\begin{aligned}
\|\boldsymbol{\theta}_* - \tilde{\boldsymbol{\theta}}_t\|_{\tilde{\mathbf{V}}_t} &\leq \sqrt{\kappa} \|\boldsymbol{\theta}_* - \tilde{\boldsymbol{\theta}}_t\|_{\mathbf{G}_t(\boldsymbol{\theta}_*, \tilde{\boldsymbol{\theta}}_t)} && \text{(Lemma 12)} \\
&= \sqrt{\kappa} \|\mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\tilde{\boldsymbol{\theta}}_t)\|_{\mathbf{G}_t^{-1}(\boldsymbol{\theta}_*, \tilde{\boldsymbol{\theta}}_t)} && \text{(Lemma 3)} \\
&\leq \sqrt{\kappa} \left(\|\mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t)\|_{\mathbf{G}_t^{-1}(\boldsymbol{\theta}_*, \tilde{\boldsymbol{\theta}}_t)} + \|\mathbf{g}_t(\hat{\boldsymbol{\theta}}_t) - \mathbf{g}_t(\tilde{\boldsymbol{\theta}}_t)\|_{\mathbf{G}_t^{-1}(\boldsymbol{\theta}_*, \tilde{\boldsymbol{\theta}}_t)} \right) \\
&\leq \kappa \left(\|\mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t)\|_{\tilde{\mathbf{V}}_t^{-1}} + \|\mathbf{g}_t(\hat{\boldsymbol{\theta}}_t) - \mathbf{g}_t(\tilde{\boldsymbol{\theta}}_t)\|_{\tilde{\mathbf{V}}_t^{-1}} \right) && \text{(Lemma 12)} \\
&\leq 2\kappa \|\mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t)\|_{\tilde{\mathbf{V}}_t^{-1}} && \text{(Definition of } \tilde{\boldsymbol{\theta}}_t \text{ and } \boldsymbol{\theta}_* \in \Theta) \\
&\leq 2\kappa \left(\sqrt{\lambda} S + \|\mathbf{s}_t\|_{\tilde{\mathbf{V}}_t^{-1}} \right) && (65) \\
&\leq 2\kappa \left(\sqrt{\lambda} S + 2\sqrt{\log(1/\delta) + Kd \log\left(1 + \frac{t}{\kappa\lambda d}\right)} \right). \\
&\hspace{15em} \text{(Theorem 1 in Abbasi-Yadkori et al. (2011))}
\end{aligned}$$

□

Lemma 15. For all $t \in [T]$, it holds that $\mathcal{C}_t(\delta) \subseteq \tilde{\mathcal{C}}_t(\delta)$.

Proof. Let $\boldsymbol{\theta} \in \mathcal{C}_t(\delta)$.

$$\begin{aligned}
\|\boldsymbol{\theta} - \boldsymbol{\theta}_t\|_{\mathbf{H}_t(\boldsymbol{\theta})} &\leq \sqrt{1 + 2S} \|\boldsymbol{\theta} - \boldsymbol{\theta}_t\|_{\mathbf{G}_t(\boldsymbol{\theta}, \boldsymbol{\theta}_t)} && \text{(Lemma 13)} \\
&= \sqrt{1 + 2S} \|\mathbf{g}_t(\boldsymbol{\theta}) - \mathbf{g}_t(\boldsymbol{\theta}_t)\|_{\mathbf{G}_t^{-1}(\boldsymbol{\theta}, \boldsymbol{\theta}_t)} && \text{(Lemma 11)} \\
&\leq \sqrt{1 + 2S} \left(\|\mathbf{g}_t(\boldsymbol{\theta}) - \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t)\|_{\mathbf{G}_t^{-1}(\boldsymbol{\theta}, \boldsymbol{\theta}_t)} + \|\mathbf{g}_t(\hat{\boldsymbol{\theta}}_t) - \mathbf{g}_t(\boldsymbol{\theta}_t)\|_{\mathbf{G}_t^{-1}(\boldsymbol{\theta}, \boldsymbol{\theta}_t)} \right) \\
&\leq (1 + 2S) \left(\|\mathbf{g}_t(\boldsymbol{\theta}) - \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t)\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta})} + \|\mathbf{g}_t(\hat{\boldsymbol{\theta}}_t) - \mathbf{g}_t(\boldsymbol{\theta}_t)\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta}_t)} \right) \\
&\hspace{15em} \text{(Lemma 13)} \\
&\leq (2 + 4S) \beta_t(\delta). && \text{(Theorem 1)}
\end{aligned}$$

Thus, $\boldsymbol{\theta} \in \tilde{\mathcal{C}}_t(\delta)$ which completes the proof. □

C.2 Proof of (24)

Based on this confidence set and under the event $\{\boldsymbol{\theta}_* \in \mathcal{E}_t(\delta), \forall t \in [T]\}$, we derive a new upper bound on $\Delta(\mathbf{x}, \boldsymbol{\theta})$ for all $\mathbf{x} \in \mathcal{D}$ and $\boldsymbol{\theta} \in \mathcal{E}_t(\delta)$ as follows:

$$\begin{aligned}
\Delta(\mathbf{x}, \boldsymbol{\theta}) &= |\boldsymbol{\rho}^T \mathbf{z}(\mathbf{x}, \boldsymbol{\theta}_*) - \boldsymbol{\rho}^T \mathbf{z}(\mathbf{x}, \boldsymbol{\theta})| \\
&\leq R \|\mathbf{z}(\mathbf{x}, \boldsymbol{\theta}_*) - \mathbf{z}(\mathbf{x}, \boldsymbol{\theta})\|_2 \quad (\text{Assumption 1 and Cauchy-Schwarz inequality}) \\
&= R \left\| \left[\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_*, \boldsymbol{\theta}) \otimes \mathbf{x}^T \right] (\boldsymbol{\theta}_* - \boldsymbol{\theta}) \right\|_2 \quad (\text{Lemma 10}) \\
&= R \left\| \left[\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_*, \boldsymbol{\theta}) \otimes \mathbf{x}^T \right] \tilde{\mathbf{V}}_t^{-1/2} \tilde{\mathbf{V}}_t^{1/2} (\boldsymbol{\theta}_* - \boldsymbol{\theta}) \right\|_2 \\
&\leq R \left\| \left[\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_*, \boldsymbol{\theta}) \otimes \mathbf{x}^T \right] \tilde{\mathbf{V}}_t^{-1/2} \right\|_2 \|\boldsymbol{\theta}_* - \boldsymbol{\theta}\|_{\tilde{\mathbf{V}}_t} \quad (\text{Cauchy-Schwarz inequality}) \\
&= R \sqrt{\lambda_{\max} \left(\left[\mathbf{B}(\mathbf{x}, \boldsymbol{\theta}_*, \boldsymbol{\theta}) \otimes \mathbf{x}^T \right] \tilde{\mathbf{V}}_t^{-1} \left[\mathbf{B}^T(\mathbf{x}, \boldsymbol{\theta}_*, \boldsymbol{\theta}) \otimes \mathbf{x} \right] \right)} \|\boldsymbol{\theta}_* - \boldsymbol{\theta}\|_{\tilde{\mathbf{V}}_t} \\
&\leq RL \sqrt{\lambda_{\max} \left(\left[I_K \otimes \mathbf{V}_t^{-1} \right] \left[I_K \otimes \mathbf{x} \mathbf{x}^T \right] \right)} \|\boldsymbol{\theta}_* - \boldsymbol{\theta}\|_{\tilde{\mathbf{V}}_t} \\
&\hspace{15em} (\text{Assumption 2 and cyclic property of } \lambda_{\max}) \\
&= RL \sqrt{\lambda_{\max} \left(\left[I_K \otimes \mathbf{V}_t^{-1} \right] \left[I_K \otimes \mathbf{x} \right] \left[I_K \otimes \mathbf{x}^T \right] \right)} \|\boldsymbol{\theta}_* - \boldsymbol{\theta}\|_{\tilde{\mathbf{V}}_t} \\
&\hspace{15em} (\text{mixed-product property}) \\
&= RL \sqrt{\lambda_{\max} \left(\left[I_K \otimes \mathbf{x}^T \right] \left[I_K \otimes \mathbf{V}_t^{-1} \right] \left[I_K \otimes \mathbf{x} \right] \right)} \|\boldsymbol{\theta}_* - \boldsymbol{\theta}\|_{\tilde{\mathbf{V}}_t} \\
&\hspace{15em} (\text{cyclic property of } \lambda_{\max}) \\
&= RL \|\mathbf{x}\|_{\mathbf{V}_t^{-1}} \left[\|\boldsymbol{\theta}_* - \tilde{\boldsymbol{\theta}}_t\|_{\tilde{\mathbf{V}}_t} + \|\tilde{\boldsymbol{\theta}}_t - \boldsymbol{\theta}\|_{\tilde{\mathbf{V}}_t} \right] \\
&\leq 2RL\kappa\gamma_t(\delta) \|\mathbf{x}\|_{\mathbf{V}_t^{-1}}. \quad (\text{Lemma 14})
\end{aligned}$$

C.3 A useful diagonally-dominant M -matrix

Lemma 16. For any $\mathbf{x} \in \mathbb{R}^d$ and $\boldsymbol{\theta} \in \mathbb{R}^{Kd}$, matrix $\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})$ defined in (10) is a strictly diagonally dominant M -matrix.

Proof. Recall that $\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})_{ij} := z_i(\mathbf{x}, \boldsymbol{\theta}) (\delta_{ij} - z_j(\mathbf{x}, \boldsymbol{\theta}))$ for all $i, j \in [K]$. Thanks to the MNL probabilistic model, one can easily observe that for each $i \in [K]$ of matrix $\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})$, we have:

$$\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})_{ii} = z_0(\mathbf{x}, \boldsymbol{\theta}) z_i(\mathbf{x}, \boldsymbol{\theta}) + \sum_{j \neq i} |\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})_{ij}| > \sum_{j \neq i} |\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})_{ij}|, \quad (66)$$

$$\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})_{ii} = z_0(\mathbf{x}, \boldsymbol{\theta}) z_i(\mathbf{x}, \boldsymbol{\theta}) + \sum_{j \neq i} |\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})_{ji}| > \sum_{j \neq i} |\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})_{ji}|. \quad (67)$$

Therefore, for any $\mathbf{x} \in \mathbb{R}^d$ and $\boldsymbol{\theta} \in \mathbb{R}^{Kd}$, matrix $\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})$ is strictly diagonally dominant. Furthermore, a strictly diagonally dominant matrix which is also *symmetric* with positive diagonal entries is positive definite (Theorem 6.1.10 in [Horn and Johnson \(2012\)](#)). Thus, $\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})$ is a positive definite matrix with positive eigenvalues, which paired with the fact that all its off-diagonal entries are negative proves that it is also an M -matrix. \square

C.4 Proof of (26)

For convenience, define $\tilde{S} = S/\sqrt{K}$.

Algorithm 2: Improved MNL-UCB

- 1 **for** $t = 1$ **to** T **do**
 - 2 Compute $\boldsymbol{\theta}_t$ as in (73).
 - 3 Compute $\mathbf{x}_t := \arg \max_{\mathbf{x} \in \mathcal{D}} \boldsymbol{\rho}^T \mathbf{z}(\mathbf{x}, \boldsymbol{\theta}_t) + \bar{c}_t(\mathbf{x}, \boldsymbol{\theta}_t)$ with $\boldsymbol{\theta}_t$ and $\bar{c}_t(\mathbf{x}, \boldsymbol{\theta}_t)$ defined in (73) and (78).
 - 4 Play \mathbf{x}_t and observe y_t .
-

An application of Theorem 1.1 in [Tian and Huang \(2010\)](#) that provides upper and lower bounds on the minimum eigenvalue of a strictly diagonally dominant M -matrix, together with Lemma 5 give the following for all \mathbf{x} and $\boldsymbol{\theta}$: $\min_{i \in [K]} \sum_{j=1}^K \mathbf{A}(\mathbf{x}, \boldsymbol{\theta})_{ij} \leq \lambda_{\min}(\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})) \leq \max_{i \in [K]} \sum_{j=1}^K \mathbf{A}(\mathbf{x}, \boldsymbol{\theta})_{ij}$. By the definition of κ in Assumption 2, the previous lower bound and the definitions of $\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})$ and of the multinomial logit probability model we have then that

$$\begin{aligned}
 \frac{1}{\kappa} &:= \min_{\mathbf{x} \in \mathcal{D}, \boldsymbol{\theta} \in \Theta} \lambda_{\min}(\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})) \geq \min_{\mathbf{x} \in \mathcal{D}, \boldsymbol{\theta} \in \Theta, i \in [K]} \sum_{j=1}^K \mathbf{A}(\mathbf{x}, \boldsymbol{\theta})_{ij} \\
 &= \min_{\mathbf{x} \in \mathcal{D}, \boldsymbol{\theta} \in \Theta, i \in [K]} \frac{e^{\bar{\boldsymbol{\theta}}_i^T \mathbf{x}}}{\left(1 + \sum_{j=1}^K e^{\bar{\boldsymbol{\theta}}_j^T \mathbf{x}}\right)^2} = \min_{i \in [n]} \min_{\mathbf{x} \in \mathcal{D}, \boldsymbol{\theta} \in \Theta} \frac{e^{\bar{\boldsymbol{\theta}}_i^T \mathbf{x}}}{\left(1 + \sum_{j=1}^K e^{\bar{\boldsymbol{\theta}}_j^T \mathbf{x}}\right)^2} \\
 &\geq \min_{i \in [n]} \frac{e^{-SX}}{(1 + Ke^{SX})^2} = \frac{e^{-SX}}{(1 + Ke^{SX})^2}. \tag{68}
 \end{aligned}$$

The lower bound in the last line follows since for any $\mathbf{x} \in \mathcal{D}$ and $\boldsymbol{\theta} \in \Theta$ by Cauchy-Schwarz: $-SX \leq \bar{\boldsymbol{\theta}}_i^T \mathbf{x} \leq SX$, for all $i \in [n]$. This proves the advertised upper bound on κ .

For the lower bound, we proceed similarly to find that

$$\begin{aligned}
 \frac{1}{\kappa} &:= \min_{\mathbf{x} \in \mathcal{D}, \boldsymbol{\theta} \in \Theta} \lambda_{\min}(\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})) \leq \min_{\mathbf{x} \in \mathcal{D}, \boldsymbol{\theta} \in \Theta} \max_{i \in [K]} \sum_{j=1}^K \mathbf{A}(\mathbf{x}, \boldsymbol{\theta})_{ij} \\
 &= \min_{\mathbf{x} \in \mathcal{D}, \boldsymbol{\theta} \in \Theta} \max_{i \in [K]} \frac{e^{\bar{\boldsymbol{\theta}}_i^T \mathbf{x}}}{\left(1 + \sum_{j=1}^K e^{\bar{\boldsymbol{\theta}}_j^T \mathbf{x}}\right)^2} \\
 &\leq \max_{i \in [K]} \frac{e^{-\tilde{S}X}}{(1 + Ke^{-\tilde{S}X})^2} = \frac{e^{-\tilde{S}X}}{(1 + Ke^{-\tilde{S}X})^2}.
 \end{aligned}$$

The inequality in the last line follows by choosing feasible \mathbf{x} and $\boldsymbol{\theta}$ as follows. Let $\bar{\boldsymbol{\theta}}_1 = \dots = \bar{\boldsymbol{\theta}}_K = \bar{\boldsymbol{\theta}}$ with $\|\bar{\boldsymbol{\theta}}\|_2 = \tilde{S}$ and $\mathbf{x} = -\frac{X}{\tilde{S}}\bar{\boldsymbol{\theta}}$. The above gives the desired upper bound and concludes the proof.

D Improved Multinomial Logit UCB

In this section, we exploit careful algorithmic designs to entirely drop the regret's dependency on κ (except in logarithmic terms). First, we introduce the following assumption.

Assumption 3 (Problem-dependent constants). *There exist $0 < M < \infty$ and $0 < M' < \infty$ such that for all $\mathbf{x} \in \mathcal{D}, \boldsymbol{\theta} \in \Theta$ and $i, j \in [K]$*

$$\left| \lambda_{\max} \left(\nabla^2 z_i(\mathbf{x}, \boldsymbol{\theta}) \right) \right| \leq M \quad \text{and} \quad \left\| \nabla \mathbf{A}(\mathbf{x}, \boldsymbol{\theta})_{ij} \right\|_2 \leq M', \tag{69}$$

where $z_i(\mathbf{x}, \boldsymbol{\theta})$ is the i -th entry of vector $\mathbf{z}(\mathbf{x}, \boldsymbol{\theta})$ and $\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})_{ij}$ is the ij -th entry of matrix $\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})$.

A close inspection of the proof presented in the previous subsections reveals that the ‘‘extra’’ $\sqrt{\kappa}$ factor enters the regret in Equation (**) in the display in Section B.2. Specifically this follows when

we replace the matrix $\mathbf{G}_t(\boldsymbol{\theta}_*, \boldsymbol{\theta})$ with the simpler matrix $\tilde{\mathbf{V}}_t$ (cf. Lemma 12). In turn, this is possible by replacing the key matrix $\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})$ —the Jacobian of the MNL model—by the smaller—in the sense of the Loewner order—matrix $\frac{1}{\kappa} I_K$ (cf. Assumption 1). In the binary case, the idea of Fauray et al. (2020) to circumvent this step is to introduce a refined “local” lower bound to the derivative of the logistic model (corresponding to our Jacobian above). In the binary case, such a lower bound always exists. However, this is not at all clear in the multinomial case because the Loewner order is only a partial order. This is yet another demonstration that extensions to the multinomial case are challenging on their own right and might require careful treatment. In what follows, we propose a new definition for the estimator $\boldsymbol{\theta}_t$ chosen from a newly defined set $\Theta \cap \mathcal{M}_t(\delta)$ (instead of Θ in (21)) formed based on the minimal elements of partially Loewner-ordered set $\mathcal{C}_t(\delta) \cap \Theta$. Please note that a minimal element of some partially ordered set \mathcal{S} is defined as an element of \mathcal{S} that is not greater than any other element in \mathcal{S} .

At each round $t \in [T]$, let $\{\boldsymbol{\theta}'_{t,i} \in \mathcal{C}_t(\delta) \cap \Theta\}_{i \in [N_t]}$ be a set of N_t vectors such that if for each $i \in [N_t]$, we define the sets

$$\mathcal{C}_{t,i}(\delta) := \{\boldsymbol{\theta} \in \mathcal{C}_t(\delta) \cap \Theta : \mathbf{A}(\mathbf{x}_t, \boldsymbol{\theta}) \succeq \mathbf{A}(\mathbf{x}_t, \boldsymbol{\theta}'_{t,i})\}, \quad (70)$$

then $\mathcal{C}_t(\delta) \cap \Theta = \cup_{i=1}^{N_t} \mathcal{C}_{t,i}(\delta)$. In fact, vectors $\boldsymbol{\theta}'_{t,i}$ are the minimal elements of the partially ordered set $\mathcal{C}_t(\delta) \cap \Theta$. Next, we define the set

$$\mathcal{M}_t(\delta) := \left\{ \boldsymbol{\theta} \in \mathbb{R}^{Kd} : \forall s \in [t-1], \exists i(s) \in [N_s] \text{ such that } \mathbf{A}(\mathbf{x}_s, \boldsymbol{\theta}) \succeq \mathbf{A}(\mathbf{x}_s, \boldsymbol{\theta}'_{s,i(s)}) \right\}. \quad (71)$$

Next, we define the following feasible set of estimators from which the new estimator $\boldsymbol{\theta}_t$ is chosen:

$$\mathcal{W}_t(\delta) := \mathcal{C}_t(\delta) \cap \mathcal{M}_t(\delta), \quad (72)$$

Now, we define the estimator $\boldsymbol{\theta}_t$ computed by the agent at round t :

$$\boldsymbol{\theta}_t := \arg \min_{\boldsymbol{\theta} \in \Theta \cap \mathcal{M}_t(\delta)} \left\| \mathbf{g}_t(\boldsymbol{\theta}) - \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t) \right\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta})}. \quad (73)$$

Under the event \mathcal{E} in (47), $\boldsymbol{\theta}_* \in \mathcal{C}_t(\delta) \cap \Theta$ for all $t > 0$. Thus, for all $t \in [T]$, there exists $i(t) \in [N_t]$ such that $\boldsymbol{\theta}_* \in \mathcal{C}_{t,i(t)}$. This combined with the definition of $\mathcal{M}_t(\delta)$ in (71) guarantees that $\boldsymbol{\theta}_* \in \mathcal{M}_t(\delta)$ and therefore $\boldsymbol{\theta}_* \in \Theta \cap \mathcal{M}_t(\delta)$. This shows that $\boldsymbol{\theta}_t \in \mathcal{W}_t(\delta)$ for all $t \in [T]$ since 1) $\boldsymbol{\theta}_t \in \mathcal{C}_t(\delta)$ as $\left\| \mathbf{g}_t(\boldsymbol{\theta}_t) - \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t) \right\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta}_t)} \leq \left\| \mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t) \right\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta}_*)} \leq \beta_t(\delta)$; and 2) $\boldsymbol{\theta}_t \in \Theta \cap \mathcal{M}_t(\delta)$.

In the next step towards completing the design of the algorithm, in Lemma 18, we define an improved exploration bonus $\epsilon_t(\mathbf{x}, \boldsymbol{\theta}_t)$ for any $\mathbf{x} \in \mathcal{D}$ and $\boldsymbol{\theta}_t$ defined in (73). First, we prove the following key Lemma 17, whose results will be used in the proof of Lemma 18.

Lemma 17. *For all $i \in [K]$, $\mathbf{x} \in \mathcal{D}$, $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \mathcal{W}_t(\delta)$, and $t \in [T]$, with probability at least $1 - \delta$, it holds that*

$$z_i(\mathbf{x}, \boldsymbol{\theta}_1) - z_i(\mathbf{x}, \boldsymbol{\theta}_2) \leq \left[\mathbf{x}^T (\bar{\boldsymbol{\theta}}_{11} - \bar{\boldsymbol{\theta}}_{21}), \dots, \mathbf{x}^T (\bar{\boldsymbol{\theta}}_{1K} - \bar{\boldsymbol{\theta}}_{2K}) \right] \mathbf{A}(\mathbf{x}, \boldsymbol{\theta}_2)_i + 4\kappa M \beta_t^2(\delta) \|\mathbf{x}\|_{\mathbf{V}_t^{-1}}^2, \quad (74)$$

and therefore

$$\mathbf{z}(\mathbf{x}, \boldsymbol{\theta}_1) - \mathbf{z}(\mathbf{x}, \boldsymbol{\theta}_2) \leq \left[\mathbf{A}(\mathbf{x}, \boldsymbol{\theta}_2) \otimes \mathbf{x}^T \right] (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2) + 4\kappa M (1 + 2S) \beta_t^2(\delta) \|\mathbf{x}\|_{\mathbf{V}_t^{-1}}^2 \mathbf{1}, \quad (75)$$

Proof. According to Taylor’s theorem, for a continuous twice differential function $f : \mathbb{R} \rightarrow \mathbb{R}$, we have

$$f(b) = f(a) + f'(a)(b-a) + \int_a^b f''(v)(b-v)dv. \quad (76)$$

Let $\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})_i$ be the i -th column of $\mathbf{A}(\mathbf{x}, \boldsymbol{\theta})$, $a = 0$, $b = 1$, $f(v) = z_i(\mathbf{x}, \boldsymbol{\theta}_2 + v(\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2))$, and $\mathbf{M}_i(\mathbf{x}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2, v) := \nabla^2 z_i(\mathbf{x}, v\boldsymbol{\theta}_1 + (1-v)\boldsymbol{\theta}_2)$ for any $v \in [0, 1]$, $i \in [K]$, $\mathbf{x} \in \mathcal{D}$, and $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \Theta$. By the Taylor's theorem, for any $i \in [K]$, we have

$$\begin{aligned} z_i(\mathbf{x}, \boldsymbol{\theta}_1) - z_i(\mathbf{x}, \boldsymbol{\theta}_2) &= \left[\mathbf{x}^T(\bar{\boldsymbol{\theta}}_{11} - \bar{\boldsymbol{\theta}}_{21}), \dots, \mathbf{x}^T(\bar{\boldsymbol{\theta}}_{1K} - \bar{\boldsymbol{\theta}}_{2K}) \right] \mathbf{A}(\mathbf{x}, \boldsymbol{\theta}_2)_i \\ &\quad + \int_{v=0}^1 (1-v) \left\| \begin{bmatrix} \mathbf{x}^T(\bar{\boldsymbol{\theta}}_{11} - \bar{\boldsymbol{\theta}}_{21}) \\ \vdots \\ \mathbf{x}^T(\bar{\boldsymbol{\theta}}_{1K} - \bar{\boldsymbol{\theta}}_{2K}) \end{bmatrix} \right\|_{\mathbf{M}_i(\mathbf{x}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2, v)}^2 dv \end{aligned} \quad (77)$$

Now, we bound the second term for any $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \mathcal{W}_i(\delta)$ as follows:

$$\begin{aligned} \int_{v=0}^1 (1-v) \left\| \begin{bmatrix} \mathbf{x}^T(\bar{\boldsymbol{\theta}}_{11} - \bar{\boldsymbol{\theta}}_{21}) \\ \vdots \\ \mathbf{x}^T(\bar{\boldsymbol{\theta}}_{1K} - \bar{\boldsymbol{\theta}}_{2K}) \end{bmatrix} \right\|_{\mathbf{M}_i(\mathbf{x}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2, v)}^2 dv &\leq M \left\| \begin{bmatrix} \mathbf{x}^T(\bar{\boldsymbol{\theta}}_{11} - \bar{\boldsymbol{\theta}}_{21}) \\ \vdots \\ \mathbf{x}^T(\bar{\boldsymbol{\theta}}_{1K} - \bar{\boldsymbol{\theta}}_{2K}) \end{bmatrix} \right\|_2^2 && \text{(Assumption 3)} \\ &= M \left\| \begin{bmatrix} I_K \otimes \mathbf{x}^T \end{bmatrix} (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2) \right\|_2^2 \\ &= M \left\| \begin{bmatrix} I_K \otimes \mathbf{x}^T \end{bmatrix} \mathbf{G}_t^{-1/2}(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \mathbf{G}_t^{1/2}(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2) (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2) \right\|_2^2 \\ &\leq M \left\| \begin{bmatrix} I_K \otimes \mathbf{x}^T \end{bmatrix} \mathbf{G}_t^{-1/2}(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \right\|_2^2 \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_{\mathbf{G}_t(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)}^2 \\ &\quad \text{(Cauchy-Schwarz inequality)} \\ &= M \lambda_{\max} \left(\mathbf{G}_t^{-1}(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \begin{bmatrix} I_K \otimes \mathbf{x} \mathbf{x}^T \end{bmatrix} \right) \|\mathbf{g}_t(\boldsymbol{\theta}_1) - \mathbf{g}_t(\boldsymbol{\theta}_2)\|_{\mathbf{G}_t^{-1}(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)}^2 \\ &\quad \text{(mixed-product property, cyclic property of } \lambda_{\max}, \text{ Lemma 11)} \\ &\leq \kappa M (1 + 2S) \lambda_{\max} \left(\begin{bmatrix} I_K \otimes \mathbf{V}_t^{-1} \end{bmatrix} \begin{bmatrix} I_K \otimes \mathbf{x} \mathbf{x}^T \end{bmatrix} \right) \|\mathbf{g}_t(\boldsymbol{\theta}_1) - \mathbf{g}_t(\boldsymbol{\theta}_2)\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta}_1)}^2 \\ &\quad \text{(Lemmas 12 and 13)} \\ &= \kappa M (1 + 2S) \|\mathbf{x}\|_{\mathbf{V}_t^{-1}}^2 \left\| \mathbf{g}_t(\boldsymbol{\theta}_1) - \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t) + \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t) - \mathbf{g}_t(\boldsymbol{\theta}_2) \right\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta}_1)}^2 \\ &\quad \text{(cyclic property of } \lambda_{\max}) \\ &\leq 4\kappa M (1 + 2S) \beta_t^2(\delta) \|\mathbf{x}\|_{\mathbf{V}_t^{-1}}^2, \quad \text{(Theorem 1)} \end{aligned}$$

as desired. \square

Now, we are ready to state the improved exploration bonus in the following Lemma

Lemma 18 (Improved exploration bonus). *For all $\mathbf{x} \in \mathcal{D}$ and $t \in [T]$, with probability at least $1 - \delta$, we have*

$$\Delta(\mathbf{x}, \boldsymbol{\theta}_t) \leq \bar{\epsilon}_t(\mathbf{x}, \boldsymbol{\theta}_t) := R(2 + 4S) \beta_t(\delta) \left\| \begin{bmatrix} \mathbf{A}(\mathbf{x}, \boldsymbol{\theta}_t) \otimes \mathbf{x}^T \end{bmatrix} \mathbf{H}_t^{-1/2}(\boldsymbol{\theta}_t) \right\|_2 + 4\kappa M \left(\sum_{i=1}^K \rho_i \right) (1 + 2S) \beta_t^2(\delta) \|\mathbf{x}\|_{\mathbf{V}_t^{-1}}^2, \quad (78)$$

where $\boldsymbol{\theta}_t$ is chosen as in (73).

Proof. As a direct conclusion of Lemma 17, for all $\mathbf{x} \in \mathcal{D}$, $\boldsymbol{\theta} \in \mathcal{W}_t(\delta)$ and $t \in [T]$, with probability at least $1 - \delta$ it holds that

$$\begin{aligned}
\Delta(\mathbf{x}, \boldsymbol{\theta}) &\leq R \left\| \left[\mathbf{A}(\mathbf{x}, \boldsymbol{\theta}) \otimes \mathbf{x}^T \right] (\boldsymbol{\theta}_* - \boldsymbol{\theta}) \right\|_2 + 4\kappa M \left(\sum_{i=1}^K \rho_i \right) (1 + 2S) \beta_t^2(\delta) \|\mathbf{x}\|_{\mathbf{V}_t^{-1}}^2 \\
&= R \left\| \left[\mathbf{A}(\mathbf{x}, \boldsymbol{\theta}) \otimes \mathbf{x}^T \right] \mathbf{G}_t^{-1/2}(\boldsymbol{\theta}_*, \boldsymbol{\theta}) \right\|_2 \left\| \mathbf{g}_t(\boldsymbol{\theta}_*) - \mathbf{g}_t(\boldsymbol{\theta}) \right\|_{\mathbf{G}_t^{-1}(\boldsymbol{\theta}_*, \boldsymbol{\theta})} + 4\kappa M \left(\sum_{i=1}^K \rho_i \right) (1 + 2S) \beta_t^2(\delta) \|\mathbf{x}\|_{\mathbf{V}_t^{-1}}^2 \\
&\leq 2R \left(\sqrt{1 + 2S} \right) \beta_t(\delta) \left\| \left[\mathbf{A}(\mathbf{x}, \boldsymbol{\theta}) \otimes \mathbf{x}^T \right] \mathbf{G}_t^{-1/2}(\boldsymbol{\theta}_*, \boldsymbol{\theta}) \right\|_2 + 4\kappa M \left(\sum_{i=1}^K \rho_i \right) (1 + 2S) \beta_t^2(\delta) \|\mathbf{x}\|_{\mathbf{V}_t^{-1}}^2 \\
&\hspace{15em} \text{(Lemma 13 and Theorem 1)} \\
&\leq R(2 + 4S) \beta_t(\delta) \left\| \left[\mathbf{A}(\mathbf{x}, \boldsymbol{\theta}) \otimes \mathbf{x}^T \right] \mathbf{H}_t^{-1/2}(\boldsymbol{\theta}) \right\|_2 + 4\kappa M \left(\sum_{i=1}^K \rho_i \right) (1 + 2S) \beta_t^2(\delta) \|\mathbf{x}\|_{\mathbf{V}_t^{-1}}^2, \\
&\hspace{15em} \text{(Lemma 13)}
\end{aligned}$$

which completes the proof of Lemma 18 as $\boldsymbol{\theta}_t \in \mathcal{W}_t(\delta)$ at all rounds $t \in [T]$. \square

Theorem 3 (Regret of Improved MNL-UCB). *Fix $\delta \in (0, 1)$. Let Assumptions 1 and 2 hold. Then, with probability at least $1 - \delta$, it holds that*

$$R_T \leq 12C_1 \sqrt{2 \max(1, \frac{1}{\lambda\kappa}) K T d \log \left(1 + \frac{T}{d\lambda} \right)} + 4C_2 \max(1, \frac{1}{\lambda\kappa}) d \log \left(1 + \frac{T}{\kappa\lambda d} \right), \quad (79)$$

where

$$C_1 = R(2 + 4S) \beta_T(\delta) \quad \text{and} \quad C_2 = 2RM'(2 + 4S) \kappa K \sqrt{(1 + 2S) \beta_T^2(\delta)} + 4\kappa M \left(\sum_{i=1}^K \rho_i \right) (1 + 2S) \beta_T^2(\delta). \quad (80)$$

Furthermore, if $\lambda = Kd \log(T)$, then

$$R_T = \mathcal{O} \left(K^{3/2} d \sqrt{T} \log(T) + \kappa \max(M, M') K^2 d \log(T) \right). \quad (81)$$

Proof. Recall the definition of $\boldsymbol{\theta}'_{t,i}$ in the line above (70). A first order Taylor expansion gives that for all $\mathbf{x} \in \mathcal{D}$, $t > 0$, $h \in [N_t]$, and $i, j \in [K]$

$$\mathbf{A}(\mathbf{x}, \boldsymbol{\theta}_t)_{ij} = \mathbf{A}(\mathbf{x}, \boldsymbol{\theta}'_{t,h})_{ij} + \left[\mathbf{x}^T (\bar{\boldsymbol{\theta}}_{t1} - \bar{\boldsymbol{\theta}}'_{t,h1}), \dots, \mathbf{x}^T (\bar{\boldsymbol{\theta}}_{tK} - \bar{\boldsymbol{\theta}}'_{t,hK}) \right] \int_{v=0}^1 \nabla \mathbf{A}(\mathbf{x}, v\boldsymbol{\theta}_t + (1-v)\boldsymbol{\theta}'_{t,h})_{ij} dv.$$

For any $\mathbf{x} \in \mathbb{R}^d$ and $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \mathbb{R}^{Kd}$, define $\mathbf{U}(\mathbf{x}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \in \mathbb{R}^{K \times K}$ with $\mathbf{U}(\mathbf{x}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2)_{ij} = \left[\mathbf{x}^T (\bar{\boldsymbol{\theta}}_{11} - \bar{\boldsymbol{\theta}}_{21}), \dots, \mathbf{x}^T (\bar{\boldsymbol{\theta}}_{1K} - \bar{\boldsymbol{\theta}}_{2K}) \right] \int_{v=0}^1 \nabla \mathbf{A}(\mathbf{x}, v\boldsymbol{\theta}_1 + (1-v)\boldsymbol{\theta}_2)_{ij} dv$. Thus, we have

$$\mathbf{A}(\mathbf{x}, \boldsymbol{\theta}_t) = \mathbf{A}(\mathbf{x}, \boldsymbol{\theta}'_{t,h}) + \mathbf{U}(\mathbf{x}, \boldsymbol{\theta}_t, \boldsymbol{\theta}'_{t,h}). \quad (82)$$

Now, we upper bound $\lambda_{\max}(\mathbf{U}(\mathbf{x}, \boldsymbol{\theta}_t, \boldsymbol{\theta}'_{t,h}))$ as follows

$$\begin{aligned}
\lambda_{\max}(\mathbf{U}(\mathbf{x}, \boldsymbol{\theta}_t, \boldsymbol{\theta}'_{t,h})) &= \left\| \mathbf{U}(\mathbf{x}, \boldsymbol{\theta}_t, \boldsymbol{\theta}'_{t,h}) \right\|_2 \\
&\leq \left\| \mathbf{U}(\mathbf{x}, \boldsymbol{\theta}_t, \boldsymbol{\theta}'_{t,h}) \right\|_F \\
&= \sqrt{\sum_{i,j \in [K]} \mathbf{U}(\mathbf{x}, \boldsymbol{\theta}_t, \boldsymbol{\theta}'_{t,h})_{ij}^2} \\
&\leq \sqrt{\sum_{i,j \in [K]} \left\| [I_K \otimes \mathbf{x}] (\boldsymbol{\theta}_t - \boldsymbol{\theta}'_{t,h}) \right\|_2^2 \left\| \int_{v=0}^1 \nabla \mathbf{A}(\mathbf{x}, v\boldsymbol{\theta}_t + (1-v)\boldsymbol{\theta}'_{t,h})_{ij} dv \right\|_2^2} \\
&\hspace{15em} \text{(Cauchy-Schwarz inequality)} \\
&\leq M'K \left\| [I_K \otimes \mathbf{x}] (\boldsymbol{\theta}_t - \boldsymbol{\theta}'_{t,h}) \right\|_2 \quad (\boldsymbol{\theta}_t, \boldsymbol{\theta}'_{t,h} \in \Theta \text{ and Assumption 3}) \\
&\leq M'K \left\| [I_K \otimes \mathbf{x}] \mathbf{G}_t^{-1/2}(\boldsymbol{\theta}_t, \boldsymbol{\theta}'_{t,h}) \right\|_2 \left\| \mathbf{g}_t(\boldsymbol{\theta}_t) - \mathbf{g}_t(\boldsymbol{\theta}'_{t,h}) \right\|_{\mathbf{G}_t^{-1}(\boldsymbol{\theta}_t, \boldsymbol{\theta}'_{t,h})} \\
&\hspace{15em} \text{(Lemma 11 and Cauchy-Schwarz inequality)} \\
&\leq 2M'K \sqrt{\kappa(1+2S)} \beta_t(\delta) \|\mathbf{x}\|_{\mathbf{V}_t^{-1}}. \quad \text{(Lemma 12 and } \boldsymbol{\theta}_t, \boldsymbol{\theta}'_{t,h} \in \mathcal{C}_t(\delta))
\end{aligned}$$

Combining this upper bound on $\lambda_{\max}(\mathbf{U}(\mathbf{x}, \boldsymbol{\theta}_t, \boldsymbol{\theta}'_{t,h}))$ and (82), we have

$$\begin{aligned}
\left\| [\mathbf{A}(\mathbf{x}, \boldsymbol{\theta}_t) \otimes \mathbf{x}^T] \mathbf{H}_t^{-1/2}(\boldsymbol{\theta}_t) \right\|_2 &= \left\| \left[(\mathbf{A}(\mathbf{x}, \boldsymbol{\theta}'_{t,h}) + \mathbf{U}(\mathbf{x}, \boldsymbol{\theta}_t, \boldsymbol{\theta}'_{t,h})) \otimes \mathbf{x}^T \right] \mathbf{H}_t^{-1/2}(\boldsymbol{\theta}_t) \right\|_2 \\
&\leq \left\| [\mathbf{A}(\mathbf{x}, \boldsymbol{\theta}'_{t,h}) \otimes \mathbf{x}^T] \mathbf{H}_t^{-1/2}(\boldsymbol{\theta}_t) \right\|_2 + \left\| [\mathbf{U}(\mathbf{x}, \boldsymbol{\theta}_t, \boldsymbol{\theta}'_{t,h}) \otimes \mathbf{x}^T] \mathbf{H}_t^{-1/2}(\boldsymbol{\theta}_t) \right\|_2 \\
&\hspace{15em} \text{(Triangle inequality)} \\
&\leq \left\| [\mathbf{A}(\mathbf{x}, \boldsymbol{\theta}'_{t,h}) \otimes \mathbf{x}^T] \mathbf{H}_t^{-1/2}(\boldsymbol{\theta}_t) \right\|_2 + 2M'K \sqrt{\kappa(1+2S)} \beta_t(\delta) \|\mathbf{x}\|_{\mathbf{V}_t^{-1}} \left\| [I_K \otimes \mathbf{x}^T] \mathbf{H}_t^{-1/2}(\boldsymbol{\theta}_t) \right\|_2 \\
&\leq \left\| [\mathbf{A}(\mathbf{x}, \boldsymbol{\theta}'_{t,h}) \otimes \mathbf{x}^T] \mathbf{H}_t^{-1/2}(\boldsymbol{\theta}_t) \right\|_2 + 2M'\kappa K \sqrt{(1+2S)} \beta_t(\delta) \|\mathbf{x}\|_{\mathbf{V}_t^{-1}} \left\| [I_K \otimes \mathbf{x}^T] \tilde{\mathbf{V}}_t^{-1/2} \right\|_2 \\
&\hspace{15em} (\mathbf{H}_t(\boldsymbol{\theta}_t) \succeq \kappa^{-1} \tilde{\mathbf{V}}_t) \\
&= \left\| [\mathbf{A}(\mathbf{x}, \boldsymbol{\theta}'_{t,h}) \otimes \mathbf{x}^T] \mathbf{H}_t^{-1/2}(\boldsymbol{\theta}_t) \right\|_2 + 2M'\kappa K \sqrt{(1+2S)} \beta_t(\delta) \|\mathbf{x}\|_{\mathbf{V}_t^{-1}}^2.
\end{aligned} \tag{83}$$

Combining (78) and (83), for all $\mathbf{x} \in \mathcal{D}$, $t > 0$, and $h \in [N_t]$ we have

$$\bar{\epsilon}_t(\mathbf{x}, \boldsymbol{\theta}_t) \leq C_1 \left\| [\mathbf{A}(\mathbf{x}, \boldsymbol{\theta}'_{t,h}) \otimes \mathbf{x}^T] \mathbf{H}_t^{-1/2}(\boldsymbol{\theta}_t) \right\|_2 + C_2 \|\mathbf{x}\|_{\mathbf{V}_t^{-1}}^2, \tag{84}$$

where

$$C_1 = R(2+4S) \beta_T(\delta) \quad \text{and} \quad C_2 = 2RM'(2+4S) \kappa K \sqrt{(1+2S)} \beta_T^2(\delta) + 4\kappa M \left(\sum_{i=1}^K \rho_i \right) (1+2S) \beta_T^2(\delta). \tag{85}$$

Recall the definition of $\mathcal{M}_t(\delta)$ in (35) and for $s \in [t-1]$, let $i(s) \in [N_s]$ be such that $\mathbf{A}(\mathbf{x}_s, \boldsymbol{\theta}_t) \succeq \mathbf{A}(\mathbf{x}_s, \boldsymbol{\theta}'_{s,i(s)})$. Thus, by the definition of $\mathbf{H}_t(\boldsymbol{\theta})$ in (9), we have

$$\begin{aligned}
\mathbf{H}_t(\boldsymbol{\theta}_t) &= \lambda I_{Kd} + \sum_{s=1}^{t-1} \mathbf{A}(\mathbf{x}_s, \boldsymbol{\theta}_t) \otimes \mathbf{x}_s \mathbf{x}_s^T \\
&\succeq \lambda I_{Kd} + \sum_{s=1}^{t-1} \mathbf{A}(\mathbf{x}_s, \boldsymbol{\theta}'_{s,i(s)}) \otimes \mathbf{x}_s \mathbf{x}_s^T && \text{(Definition of } \boldsymbol{\theta}_t \text{ in (73))} \\
&\succeq \lambda I_{Kd} + \sum_{s=1}^{t-1} \mathbf{A}^2(\mathbf{x}_s, \boldsymbol{\theta}'_{s,i(s)}) \otimes \mathbf{x}_s \mathbf{x}_s^T && (\lambda_i(\mathbf{A}(\mathbf{x}_s, \boldsymbol{\theta}'_{s,i(s)})) \leq 1 \text{ for all } i \in [K]) \\
&= \lambda I_{Kd} + \sum_{s=1}^{t-1} \left[\mathbf{A}(\mathbf{x}_s, \boldsymbol{\theta}'_{s,i(s)}) \otimes \mathbf{x}_s \right] \left[\mathbf{A}(\mathbf{x}_s, \boldsymbol{\theta}'_{s,i(s)}) \otimes \mathbf{x}_s^T \right] && \text{(mixed-product property)} \\
&= \lambda I_{Kd} + \sum_{s=1}^{t-1} \sum_{i=1}^K \left[\mathbf{A}(\mathbf{x}_s, \boldsymbol{\theta}'_{s,i(s)})_i \otimes \mathbf{x}_s \right] \left[\mathbf{A}(\mathbf{x}_s, \boldsymbol{\theta}'_{s,i(s)})_i^T \otimes \mathbf{x}_s^T \right] \\
&= \lambda I_{Kd} + \sum_{s=1}^{t-1} \sum_{i=1}^K \tilde{\mathbf{x}}_{i,s} \tilde{\mathbf{x}}_{i,s}^T := \mathbf{L}_t, && (86)
\end{aligned}$$

where $\mathbf{A}(\mathbf{x}_s, \boldsymbol{\theta}'_{s,i(s)})_i$ is the i -th column of $\mathbf{A}(\mathbf{x}_s, \boldsymbol{\theta}'_{s,i(s)})$ and $\tilde{\mathbf{x}}_{i,s} := \mathbf{A}(\mathbf{x}_s, \boldsymbol{\theta}'_{s,i(s)})_i \otimes \mathbf{x}_s$.

Now, we bound the cumulative sum of the first term in (84) for $h = i(t)$, as follows:

$$\begin{aligned}
\sum_{t=1}^T \left\| \left[\mathbf{A}(\mathbf{x}_t, \boldsymbol{\theta}'_{t,i(t)}) \otimes \mathbf{x}_t^T \right] \mathbf{H}_t^{-1/2}(\boldsymbol{\theta}_t) \right\|_2 &= \sum_{t=1}^T \sqrt{\sum_{i=1}^K \left\| \mathbf{A}(\mathbf{x}_t, \boldsymbol{\theta}'_{t,i(t)})_i \otimes \mathbf{x}_t \right\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta}_t)}^2} \\
&\leq \sqrt{T \sum_{t=1}^T \sum_{i=1}^K \left\| \tilde{\mathbf{x}}_{i,t} \right\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta}_t)}^2} && \text{(Cauchy-Schwarz inequality)} \\
&\leq \sqrt{T \sum_{t=1}^T \sum_{i=1}^K \left\| \tilde{\mathbf{x}}_{i,t} \right\|_{\mathbf{L}_t^{-1}}^2}. && (87)
\end{aligned}$$

For any positive semi-definite matrices \mathbf{A} , \mathbf{B} , and \mathbf{C} such that $\mathbf{A} = \mathbf{B} + \mathbf{C}$, and for any $\mathbf{x} \neq 0$, we have ([Lemma 12] in Abbasi-Yadkori et al. (2011))

$$\frac{\|\mathbf{x}\|_{\mathbf{A}}^2}{\|\mathbf{x}\|_{\mathbf{B}}^2} \leq \frac{\det(\mathbf{A})}{\det(\mathbf{B})} \quad \text{and} \quad \frac{\|\mathbf{x}\|_{\mathbf{B}^{-1}}^2}{\|\mathbf{x}\|_{\mathbf{A}^{-1}}^2} \leq \frac{\det(\mathbf{A})}{\det(\mathbf{B})}. \quad (88)$$

Thus

$$\left\| \tilde{\mathbf{x}}_{i,t} \right\|_{\mathbf{L}_t^{-1}}^2 \leq \left\| \tilde{\mathbf{x}}_{i,t} \right\|_{\mathbf{L}_{t+1}^{-1}}^2 \frac{\det \mathbf{L}_{t+1}}{\det \mathbf{L}_t} \quad (89)$$

Note that for any $t \leq 1$, $\det \mathbf{L}_1 = \lambda^{Kd}$ and $\det \mathbf{L}_T \leq (\text{trace}(\mathbf{L}_T)/Kd)^{Kd} \leq \left(\lambda + \frac{T}{d} \right)^{Kd}$, and consequently:

$$\frac{\det \mathbf{L}_T}{\det \mathbf{L}_1} = \prod_{t=1}^{T-1} \frac{\det \mathbf{L}_{t+1}}{\det \mathbf{L}_t} \leq \left(1 + \frac{T}{d\lambda} \right)^{Kd}. \quad (90)$$

Since $1 \leq \frac{\det \mathbf{L}_{t+1}}{\det \mathbf{L}_t}$ for all $t > 0$, we have $e \leq \frac{\det \mathbf{L}_{t+1}}{\det \mathbf{L}_t}$ for at most $Kd \log \left(1 + \frac{T}{d\lambda} \right)$ pairs of $(i, t) \in [K] \times [T]$, and therefore

$$\left\| \tilde{\mathbf{x}}_{i,t} \right\|_{\mathbf{L}_t^{-1}}^2 \leq e \left\| \tilde{\mathbf{x}}_{i,t} \right\|_{\mathbf{L}_{t+1}^{-1}}^2 \quad (91)$$

is true for all pairs of $(i, t) \in [K] \times [T]$ but at most $\phi(K, d, \lambda, T) := Kd \log \left(1 + \frac{T}{d\lambda}\right)$ of them. We define the set of good pairs by $\mathcal{B}_{\text{good}} := \left\{ (i, t) \in [K] \times [T], \|\tilde{\mathbf{x}}_{i,t}\|_{\mathbf{L}_t^{-1}}^2 \leq e \|\tilde{\mathbf{x}}_{i,t}\|_{\mathbf{L}_{t+1}^{-1}}^2 \right\}$. We have

$$\begin{aligned} \sum_{t=1}^T \sum_{i=1}^K \|\tilde{\mathbf{x}}_{i,t}\|_{\mathbf{L}_t^{-1}}^2 &= \sum_{(i,t) \notin \mathcal{B}_{\text{good}}} \|\tilde{\mathbf{x}}_{i,t}\|_{\mathbf{L}_t^{-1}}^2 + \sum_{(i,t) \in \mathcal{B}_{\text{good}}} \|\tilde{\mathbf{x}}_{i,t}\|_{\mathbf{L}_t^{-1}}^2 \\ &\leq \phi(K, d, \lambda, T)/\lambda + e \sum_{t=1}^T \sum_{i=1}^K \|\tilde{\mathbf{x}}_{i,t}\|_{\mathbf{L}_{t+1}^{-1}}^2 \\ &\leq \phi(K, d, \lambda, T)/\lambda + 2e\phi(K, d, \lambda, T) \end{aligned} \quad (92)$$

Now, combining (78), (87) and (92), we bound R_T in the following:

$$\begin{aligned} R_T &\leq 2 \sum_{t=1}^T \bar{\epsilon}_t(\mathbf{x}_t, \boldsymbol{\theta}_t) \\ &\leq 2C_1 \sqrt{T \sum_{t=1}^T \sum_{i=1}^K \|\tilde{\mathbf{x}}_{i,t}\|_{\mathbf{L}_t^{-1}}^2} + 2C_2 \sum_{t=1}^T \|\mathbf{x}_t\|_{\mathbf{V}_t^{-1}}^2 \\ &\leq 12C_1 \sqrt{2 \max(1, \frac{1}{\lambda\kappa}) dKT \log \left(1 + \frac{T}{d\lambda}\right)} + 4C_2 \max(1, \frac{1}{\lambda\kappa}) d \log \left(1 + \frac{T}{\kappa\lambda d}\right), \end{aligned}$$

as desired. \square

E Additional Experiments

The ‘‘local dependence’’ of the confidence set $\mathcal{C}_t(\delta)$. We highlight the superiority of using $\epsilon_t(\mathbf{x})$ vs $\tilde{\epsilon}_t(\mathbf{x})$ introduced in (22) and (24), respectively. In Figure 2, we give an illustration of how $\mathcal{C}_t(\delta)$ compares to $\mathcal{E}_t(\delta)$ from which $\epsilon_t(\mathbf{x})$ and $\tilde{\epsilon}_t(\mathbf{x})$ are derived for different values of κ . In this figure, instead of $\mathcal{C}_t(\delta)$, we considered a slightly *larger* confidence set $\tilde{\mathcal{C}}_t(\delta)$ in the more familiar (from linear bandits) form of

$$\tilde{\mathcal{C}}_t(\delta) := \left\{ \boldsymbol{\theta} \in \Theta : \left\| \boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_t \right\|_{\mathbf{H}_t^{-1}(\boldsymbol{\theta})} \leq (2 + 4S)\beta_t(\delta) \right\},$$

where $\boldsymbol{\theta}$ is a ‘‘good’’ estimate of $\boldsymbol{\theta}_*$ based on the weighted norm of $\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_t$ rather than $\mathbf{g}_t(\boldsymbol{\theta}) - \mathbf{g}_t(\hat{\boldsymbol{\theta}}_t)$ in (12). Technical details proving $\mathcal{C}_t(\delta) \subseteq \tilde{\mathcal{C}}_t(\delta)$ at all rounds $t \in [T]$ are deferred to Appendix C. For the sake of visualization, we find it instructive to depict the confidence sets in a 1-dimensional space for a realization with $K = 2$. The curves of $\tilde{\mathcal{C}}_t(\delta)$ effectively capture the local dependence of $\mathbf{H}_t(\boldsymbol{\theta})$ on $\boldsymbol{\theta}$. As a result, unlike $\mathcal{E}_t(\delta)$, the confidence set $\tilde{\mathcal{C}}_t(\delta)$ is not an ellipsoid and estimators in different directions are penalized in different ways. Furthermore, the appearance of κ in the radius of the confidence ellipsoid $\mathcal{E}_t(\delta)$ (see (23)), which is caused by the use of the matrix $\tilde{\mathbf{V}}_t$ as a global bound on \mathbf{G}_t (see Appendix C for details), results in a larger confidence set compared to $\tilde{\mathcal{C}}_t(\delta)$ that employs the local metric $\mathbf{H}_t(\boldsymbol{\theta})$. This is also consistent with the observation that as κ grows, the difference between the diameters of $\tilde{\mathcal{C}}_t(\delta)$ and $\mathcal{E}_t(\delta)$ increases.

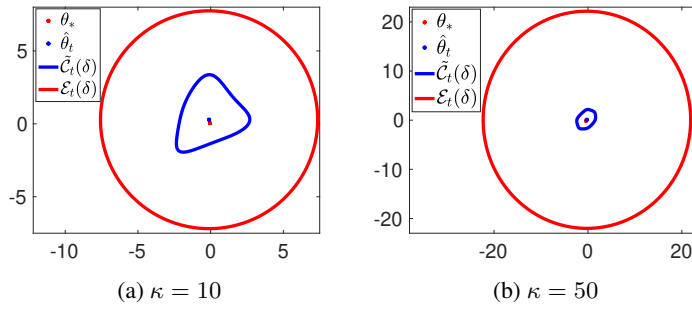


Figure 2: Illustration of $\tilde{\mathcal{C}}_t(\delta)$ and $\mathcal{E}_t(\delta)$ at round $t = 500$ for different values of κ . Note the “local dependence” of $\tilde{\mathcal{C}}_t(\delta)$ which results in a non-ellipsoidal shape.