
Combinatorial Reinforcement Learning with Preference Feedback

Joongkyu Lee¹ Min-hwan Oh¹

Abstract

In this paper, we consider combinatorial reinforcement learning with preference feedback, where a learning agent sequentially offers an action—an assortment of multiple items—to a user, whose preference feedback follows a multinomial logistic (MNL) model. This framework allows us to model real-world scenarios, particularly those involving long-term user engagement, such as in recommender systems and online advertising. However, this framework faces two main challenges: (1) the unknown value of each item, unlike traditional MNL bandits that only address single-step preference feedback, and (2) the difficulty of ensuring optimism while maintaining tractable assortment selection in the combinatorial action space with unknown values. In this paper, we assume a contextual MNL preference model, where the mean utilities are linear, and the value of each item is approximated by a general function. We propose an algorithm, MNL-VQL, that addresses these challenges, making it both computationally and statistically efficient. As a special case, for linear MDPs (with the MNL preference feedback), we establish the first regret lower bound in this framework and show that MNL-VQL achieves nearly minimax-optimal regret. To the best of our knowledge, this is the first work to provide statistical guarantees in combinatorial RL with preference feedback.

1. Introduction

We first formally state the concept of *Combinatorial Reinforcement Learning* (RL), which we refer to as a class of RL problems where the action space is combinatorial, meaning that the agent selects a combination or subset of base actions from a set of possible base actions. Although

some previous studies have addressed problems within this setting—particularly in deep RL (Snehag et al., 2015; He et al., 2016; Swaminathan et al., 2017; Metz et al., 2017; Ryu et al., 2019; Je et al., 2019; Delarue et al., 2020; McInerney et al., 2020; Vlassis et al., 2021; Chaudhari et al., 2024), with less emphasis on theoretical RL—to the best of our knowledge, it appears that no prior work has formally and theoretically defined the concept of combinatorial RL.¹ This framework is especially relevant for real-world applications such as recommender systems and online advertising, where multiple items (base actions) must be selected simultaneously, such as a set of products to recommend or advertisements to display. The challenge in combinatorial RL lies in the exponentially large action space and the need to efficiently optimize the agent’s action selection while balancing exploration and exploitation (a challenge even for single action selection), while considering the long-term effects of these actions.

One of the most widely encountered settings in combinatorial RL is *preference feedback* over combinatorial actions, commonly seen in streaming services, online retail, and similar platforms. Despite the broad applicability of this setting, theoretical studies have predominantly focused on the multinomial logistic (MNL) *bandit* model (Rusmevichientong et al., 2010; Sauré & Zeevi, 2013; Agrawal et al., 2017; 2019; Oh & Iyengar, 2019; 2021; Perivier & Goyal, 2022; Agrawal et al., 2023; Zhang & Sugiyama, 2024; Lee & Oh, 2024). The MNL bandit framework focuses on assortment (a set of items) selection by selecting subsets of items and receiving feedback on chosen items, modeled by the MNL model (McFadden, 1977). However, these studies take a *myopic* approach, optimizing for immediate, *known* rewards without considering the long-term impact on user behavior.

While MNL bandits have been widely studied, the myopic approach is limiting in many real-world scenarios. For example in recommender systems, incorporating the long-term impact of recommendations opens the door to balancing short-term engagement with long-term user satisfaction. For instance, recommending *junk* product or content might lead to high immediate reward but it can decrease user satis-

¹Seoul National University, Seoul, Korea. Correspondence to: Min-hwan Oh <minoh@snu.ac.kr>.

faction over time due to fatigue. This trade-off between immediate and long-term outcomes is not captured by traditional MNL bandit. See Appendix A for a more details.

On the empirical side, several studies have explored long-term user engagement in recommendation systems, particularly using deep RL (Swaminathan et al., 2017; Ie et al., 2019; McInerney et al., 2020; Vlassis et al., 2021; Chaudhari et al., 2024). However, there is a significant gap in the theoretical understanding of combinatorial actions with preference feedback, particularly within the RL framework. To the best of our knowledge, no theoretical work has yet explored this important problem setting.

In this paper, we aim to address this gap by rigorously studying combinatorial RL with preference feedback and developing a provably efficient algorithm that maximizes long-term user engagement by incorporating state transitions (e.g., historical behavior) into decision-making. We consider a setting where the Q -function of an assortment is decomposed into two components: the preference model and the item values, inspired by Ie et al. (2019). Specifically, we focus on the contextual MNL preference model with linear mean utilities (Agrawal & Goyal, 2013; Cheung & Simchi-Levi, 2017; Agrawal et al., 2019; Oh & Iyengar, 2019; 2021; Amani & Thrampoulidis, 2021; Perivier & Goyal, 2022; Agrawal et al., 2023; Zhang & Sugiyama, 2024; Lee & Oh, 2024) and use general function approximation to estimate item values (Jiang et al., 2017; Wang et al., 2020; Jin et al., 2021; Du et al., 2021; Foster et al., 2021; Agarwal et al., 2023; Zhao et al., 2023). The key challenges in this framework are: (1) the unknown long-term value of each item due to the stochastic nature of rewards and transitions, (2) the difficulty of selecting an assortment that ensures optimism while considering tractable assortment optimization in the combinatorial action space, given the unknown values, and (3) achieving a tighter regret bound (for MNL preference model) than simply summing over H MNL bandit regrets.

Technical novelties. To tackle challenge (1), we estimate the optimistic item values by employing point-wise optimism under general function approximation (Agarwal et al., 2023). Based on these *optimistic* item values (which incorporate uncertainty), we then select an assortment that ensures sufficient exploration and guarantees optimism. This step is the most challenging part of our framework. Since the true value of each item is unknown, directly applying techniques from MNL bandits (which assume known true values) is not feasible. Thus, to address challenge (2), we propose a novel method to estimate the *optimistic preference model* by carefully alternating between *optimistic* and *pessimistic* utilities (for the preference model), using the optimistic item values (Equation (7)). Additionally, proving optimism (Lemma D.15) and other related results (Lem-

mas D.5, D.9, and D.13) requires fundamentally more sophisticated analytical techniques. Finally, we avoid naive combinatorial enumeration when selecting the assortment (Equation (8)) by reformulating the optimization problem as a linear program (LP), inspired by Davis et al. (2013). Finally, for challenge (3), instead of naively summing over H MNL bandit regrets, we bound the regret of the MNL preference model in terms of the sum of the variances of value functions and apply the *law of total variance* (Lattimore & Hutter, 2012; Gheshlaghi Azar et al., 2013). This approach reduces MNL regret by a factor of \sqrt{H} compared to directly summing H MNL bandit regrets. Moreover, in the special case of linear MDPs with preference feedback, we achieve a nearly minimax-optimal regret.

Our main contributions are summarized as follows:

- We propose a MNL-VQL, which achieves a regret upper bound of $\tilde{O}(d\sqrt{HK} + \sqrt{\dim(\mathcal{F})KH \log |\mathcal{F}|})$ while maintaining computational efficiency (Theorem 5.1). Here, H is the horizon length, K is the total number of episodes, d is the feature dimension of the MNL preference model, and $\dim(\mathcal{F})$ is the generalized Eluder dimension (see Definition 3.5) of the function class \mathcal{F} . To the best of our knowledge, this is the first theoretical regret guarantee in combinatorial RL with preference feedback.
- For the special case of linear MDPs (with preference feedback), MNL-VQL obtains a regret upper bound of $\tilde{O}(d\sqrt{HK} + d^{\text{lin}}\sqrt{HK})$, where d^{lin} is the feature dimension of the linear MDPs (Theorem 5.2). Furthermore, we establish a matching regret lower bound of $\Omega(d\sqrt{HK} + d^{\text{lin}}\sqrt{HK})$, show the minimax-optimality of our algorithm in linear MDPs (Theorem 5.3).

2. Related Work

MNL bandits. The MNL bandits were initially studied in Rusmevichientong et al. (2010), followed by a line of improvements (Filippi et al., 2010; Rusmevichientong et al., 2010; Agrawal et al., 2017; Oh & Iyengar, 2019; Fauray et al., 2020; Abeille et al., 2021; Fauray et al., 2022; Oh & Iyengar, 2021; Perivier & Goyal, 2022; Agrawal et al., 2023; Lee & Oh, 2024). In MNL bandits, the goal is to offer an assortment that maximizes the expected rewards, which are adaptively learned based on user preference feedback from the offered assortment. However, there are no state transitions, and it is assumed that the value of each item is known, with the value of the outside option fixed at zero. Our study extends this by not only estimating the MNL model but also the long-term item values.

Combinatorial RL with preference feedback. Recently, several studies have demonstrated the empirical success of combinatorial RL with preference feedback (Swaminathan

et al., 2017; Ie et al., 2019; McInerney et al., 2020; Vlassis et al., 2021; Chaudhari et al., 2024), where a set of items is offered to a user, and (relative) choice feedback along with a reward is received, leading to a transition to the next state. However, theoretical results quantifying the benefits of such methods are still few and far between. A closely related work is cascading RL (Du et al., 2024), which also involves selecting a set of items. However, in cascading RL, items are offered to the user one by one, and the user decides only whether to choose the currently offered item. As a result, this framework does not capture relative preference feedback across multiple items. Furthermore, in cascading RL, the probability of choosing each item is independent of the others, which is not the case in our framework.

Another related line of work is preference-based RL (PbRL) (Akroul et al., 2012; Wirth et al., 2017; Christiano et al., 2017; Ouyang et al., 2022; Saha et al., 2023; Zhu et al., 2023; Zhan et al., 2023), where the policy is optimized based on relative, rather than absolute, preference feedback. However, our framework differs from PbRL in that our goal is not to offer just a single item, but to offer multiple items (a combinatorial base action).

RL with non-linear function approximation. With the limitations of the linear models (e.g., as shown in Lee & Oh (2023)), RL under non-linear function approximation has gained attention (Jiang et al., 2017; Wang et al., 2020; Jin et al., 2021; Du et al., 2021; Foster et al., 2021; Ishfaq et al., 2021; Agarwal et al., 2023; Zhao et al., 2023) for modeling complex function spaces like neural networks. Among these, Agarwal et al. (2023); Zhao et al. (2023) achieved the best-known regret guarantees under general function approximation by introducing the concept of generalized Eluder dimension to handle weighted regression. Inspired by their work, we estimate the value of items (referred to as *item-level Q-value*) using general function approximation in this paper.

3. Problem Setting

3.1. Combinatorial MDPs with Preference Feedback

In this paper, we consider an episodic *combinatorial Markov decision processes (MDPs) with preference feedback*, $\mathcal{M}(\mathcal{S}, \mathcal{I}, \mathcal{A}, M, \{\mathcal{P}_h\}_{h=1}^H, \{\mathbb{P}_h\}_{h=1}^H, \{r_h\}_{h=1}^H, H)$. Here, \mathcal{S} is the set of states. Each state $s \in \mathcal{S}$ reflects the user’s status, capturing both relatively static user features (e.g., demographics, interests) and relevant user history or past behavior (e.g., past recommendations, items purchased or clicked). $\mathcal{I} := \{a_1, \dots, a_N, a_0\}$ is the ground set of items (base actions), where a_1, \dots, a_N are items and a_0 refers to the “outside option”, meaning the user has chosen none of the items from the offered set of items (referred to as an “assortment” throughout the paper). It is included in every

assortment *by default*. \mathcal{A} is the set of candidate assortments that always include the outside option a_0 , contain at least one item (other than a_0), and have at most M items (including a_0), i.e., $\mathcal{A} = \{A \subseteq \mathcal{I} : a_0 \in A, 1 \leq |A \setminus \{a_0\}| \leq M\}$, where A is an assortment. For any $(s, A) \in \mathcal{S} \times \mathcal{A}$, we denote $\mathcal{P}_h(a|s, A)$ as the probability of the user choosing on item $a \in A$ (including the outside option a_0). Furthermore, we let $\mathbb{P}_h : \mathcal{S} \times \mathcal{I} \rightarrow \Delta_{\mathcal{S}}$ and $r_h : \mathcal{S} \times \mathcal{I} \times \mathcal{S} \rightarrow \mathbb{R}$ characterize the transition kernel and instantaneous reward, respectively, at a given horizon $h \in [H]$. Throughout this paper, we assume that $\sum_{h=1}^H r_h(s_h, a_h, s_{h+1}) \in [0, 1]$ for all possible sequence $(s_1, a_1, \dots, s_H, a_H, s_{H+1})$. $H \in \mathbb{Z}_+$ is the length of each episode. A policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$ is a mapping from the state space to the assortment space. Since the optimal policy is non-stationary in an episodic MDP, we use π to refer to the H -tuple $\{\pi_h\}_{h=1}^H$.

In each episode $k \in [K]$, an initial state s_1^k is picked *arbitrarily* (e.g., a user arrives at the system). The agent then follows a policy π^k starting from s_1^k . At each step $h \in [H]$, the agent observes the current state s_h^k (e.g., historical behaviors of the user) and offers an assortment $A_h^k = \pi_h^k(s_h^k)$. The user’s preference feedback $a_h^k \in A_h^k$ is then observed, which is drawn based on the choice probability $\mathcal{P}_h(\cdot|s_h^k, A_h^k)$. Next, the system transitions to the next state $s_{h+1}^k \sim \mathbb{P}_h(\cdot|s_h^k, a_h^k)$ and receives a reward $r_h(s_h^k, a_h^k, s_{h+1}^k)$. After H steps, the episode terminates, and the agent proceeds to the next.

For any policy $\pi = \{\pi_h\}_{h=1}^H$, we define the value function of policy π , denoted as $V_h^\pi : \mathcal{S} \rightarrow \mathbb{R}$, as the expected sum of rewards under the policy π until the end of the episode when starting from $s_h = s$, i.e., $V_h^\pi(s) := \mathbb{E} \left[\sum_{h'=h}^H r_{h'}(s_{h'}, a_{h'}, s_{h'+1}) \mid s_h = s \right]$. Moreover, we define the action-value function of policy π , $Q_h^\pi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, as the expected sum of rewards under policy π , starting from step h until the end of the episode after taking action A in state s ; that is, $Q_h^\pi(s, A) := \mathbb{E} \left[\sum_{h'=h}^H r_{h'}(s_{h'}, a_{h'}, s_{h'+1}) \mid s_h = s, A_h = A \right]$. Furthermore, we define the *item-level Q-value function* (also called the *Q-value*) $\bar{Q}_h^\pi(s, a) := \sum_{s'} \mathbb{P}_h(s'|s, a)(r_h(s, a, s') + V_{h+1}^\pi(s'))$. Then, the Bellman equation for assortment RL is denoted as follows:

$$Q_h^\pi(s, A) = \sum_{a \in A} \mathcal{P}_h(a|s, A) \bar{Q}_h^\pi(s, a).$$

Similarly, we define the optimal value function $V_h^*(s) = \sup_{\pi} V_h^\pi(s)$ and the optimal Q-value function as $Q_h^*(s, A) = \sum_{a \in A} \mathcal{P}_h(a|s, A) \bar{Q}_h^*(s, a)$, where $\bar{Q}^* := \sum_{s'} \mathbb{P}_h(s'|s, a)(r_h(s, a, s') + V_{h+1}^*(s'))$ is the *item-level optimal Q-value function*. For any $V : \mathcal{S} \rightarrow \mathbb{R}$ and $h \in [H]$, we define the *item-level Bellman operator* of V as $\mathcal{T}_h V : \mathcal{S} \times \mathcal{I} \rightarrow \mathbb{R}$, such that for all $(s, a) \in \mathcal{S} \times \mathcal{I}$, $\mathcal{T}_h V(s, a) := \mathbb{E}_{s' \sim \mathbb{P}_h(\cdot|s, a)} [r_h(s, a, s') + V(s') \mid s, a]$. The definition of value functions ensures that they satisfy

the equation $\bar{Q}_h^*(s, a) = \mathcal{T}_h V_{h+1}^*(s, a)$. We also define the *second moment item-level Bellman operator* of V as $\mathcal{T}_h^2 V : \mathcal{S} \times \mathcal{I} \rightarrow \mathbb{R}$ such that for all $(s, a) \in \mathcal{S} \times \mathcal{I}$, $\mathcal{T}_h^2 V(s, a) := \mathbb{E}_{s' \sim \mathbb{P}_h(\cdot|s, a)}[(r_h(s, a, s') + V(s'))^2 | s, a]$.

Our goal is to minimize the cumulative regret over K episodes $\text{Regret}(\mathcal{M}, K) := \sum_{k=1}^K V_1^*(s_1^k) - V_1^{\pi_k}(s_1^k)$.

3.2. Multinomial Logistic Preference Model

In this paper, we make a structural assumption about the MDP \mathcal{M} , where the user's choice probability $\{\mathcal{P}_h\}_{h=1}^H$ follows multinomial logistic (MNL) model (McFadden, 1977) parameterized by $\{\theta_h^*\}_{h=1}^H$. We denote $\mathcal{P}_h(\cdot|s, A; \theta_h^*)$ as equivalent to $\mathcal{P}_h(\cdot|s, A)$, explicitly showing the dependence on the parameter θ_h^* . Throughout the paper, we use $\mathcal{P}_h(\cdot|s, A)$ and $\mathcal{P}_h(\cdot|s, A; \theta_h^*)$ interchangeably.

Assumption 3.1 (MNL preference model). Let there exist a known feature map $\phi : \mathcal{S} \times \mathcal{I} \rightarrow \mathbb{R}^d$ such that $\|\phi(s, a)\|_2 \leq 1$, and an unknown $\theta_h^* \in \Theta$ for all $h \in [H]$, where $\Theta = \{\theta \in \mathbb{R}^d : \|\theta\|_2 \leq B = \mathcal{O}(1)\}$. Then, for any $(s, A, H) \in \mathcal{S} \times \mathcal{A} \times [H]$, the probability of choosing any item $a \in A$ is defined as:

$$\mathcal{P}_h(a|s, A) = \mathcal{P}_h(a|s, A; \theta_h^*) = \frac{\exp(\phi(s, a)^\top \theta_h^*)}{\sum_{a' \in A} \exp(\phi(s, a')^\top \theta_h^*)}.$$

Here, without loss of generality, we assume that $\phi(s, a_0) = 0$ for all $s \in \mathcal{S}$ ², which implies that $\exp(\phi(s, a_0)^\top \theta_h^*) = 1$. Thus, the preference model can be equivalently expressed as $\mathcal{P}_h(a|s, A; \theta_h^*) = \frac{\exp(\phi(s, a)^\top \theta_h^*)}{1 + \sum_{a' \in A \setminus \{a_0\}} \exp(\phi(s, a')^\top \theta_h^*)}$.

Following the previous MNL bandits (Oh & Iyengar, 2021; Perivier & Goyal, 2022; Zhang & Sugiyama, 2024; Lee & Oh, 2024), we also introduce the following constant:

Definition 3.2 (Problem-dependent constant). There exist $\kappa > 0$ such that, for any $A \in \mathcal{A}$, $a \in A \setminus \{a_0\}$, $h \in [H]$, we have $\min_{\theta \in \Theta} \mathcal{P}_h(a|s, A, \theta) \mathcal{P}_h(a_0|s, A, \theta) \geq \kappa$.

A small κ indicates a larger deviation from the linear model. Note that $1/\kappa$ can be exponentially large, so it is crucial to avoid any dependency on $1/\kappa$ in our regret bound.

3.3. Generalized Function Approximation for \bar{Q}

We estimate the item-level Q -functions (referred to as \bar{Q} -values) using general function approximation. Specifically, the agent is given a function class $\mathcal{F} := \{\mathcal{F}_h\}_{h=1}^H$, where each set \mathcal{F}_h is composed of functions $f_h : \mathcal{S} \times \mathcal{I} \rightarrow [0, L]$,

²By subtracting $\phi(s, a_0)$ from each $\phi(s, a)$, where $a \in \mathcal{I}$, and defining $\phi'(s, a) := \phi(s, a) - \phi(s, a_0)$, we can ensure that $\phi'(s, a_0) = 0$. This implies that $\exp(\phi'(s, a_0)^\top \theta_h^*) = 1$. This assumption is commonly made in contextual MNL bandits (Oh & Iyengar, 2019; 2021; Perivier & Goyal, 2022; Agrawal et al., 2023; Zhang & Sugiyama, 2024; Lee & Oh, 2024).

where $L = \mathcal{O}(1)$. Since no reward is collected in the $(H + 1)^{\text{th}}$ steps, we set $f_{H+1} = 0$. We denote \mathcal{N} as the maximal size of function class, i.e., $\mathcal{N} = \max_{h \in [H]} |\mathcal{F}_h|$. We assume the completeness and realizability for \mathcal{F} .

Assumption 3.3 (Completeness & Realizability). For each $h \in [H]$ and any $V : \mathcal{S} \rightarrow [0, 1]$, we assume that $\bar{Q}_h^* \in \mathcal{F}_h$, and there exists $f_h, f'_h \in \mathcal{F}_h$ such that, for all $(s, a) \in \mathcal{S} \times \mathcal{I}$,

$$f_h(s, a) = \mathcal{T}_h V(s, a), \quad \text{and} \quad f'_h(s, a) = \mathcal{T}_h^2 V(s, a).$$

Remark 3.4. The completeness and realizability assumptions are standard in RL with general function approximation (Wang et al., 2021; Jin et al., 2021; Agarwal et al., 2023; Zhao et al., 2023). Our assumption is the same as those in Agarwal et al. (2023); Zhao et al. (2023), but stronger than those in Wang et al. (2021); Jin et al. (2021), especially, in terms of the second moment completeness. However, this assumption is essential for using point-wise exploration bonuses and achieving a tighter regret bound. Additionally, it naturally holds for both tabular and linear MDPs.

To capture the complexity of exploration in the MDP, we define the *generalized Eluder dimension*, which is a weighted regression version of the original definition (Russo & Van Roy, 2013).

Definition 3.5 (Generalized Eluder dimension, Agarwal et al. 2023). Let $\rho > 0$, a sequence of state-item pairs $\mathbf{Z}_k = \{z^\tau\}_{\tau=1}^k$, where $z^\tau = (s^\tau, a^\tau)$, and a sequence of positive numbers $\sigma_k = \{\sigma^\tau\}_{\tau=1}^k$. The generalized Eluder dimension of a function class $\mathcal{F} : \mathcal{S} \times \mathcal{I} \rightarrow [0, L]$ with respect to ρ is defined as $\dim_{\nu, K}(\mathcal{F}) := \sum_{\mathbf{Z}_K, \sigma_K : \sigma \geq \nu} \dim(\mathcal{F}, \mathbf{Z}, \sigma)$, where

$$\dim(\mathcal{F}, \mathbf{Z}_K, \sigma_K) := \sum_{k=1}^K \min \left(1, \frac{D_{\mathcal{F}}^2(z^k; \mathbf{Z}_{k-1}, \sigma_{k-1})}{(\sigma^k)^2} \right).$$

$$D_{\mathcal{F}}^2(z^k; \mathbf{Z}_{k-1}, \sigma_{k-1}) = \sup_{f_1, f_2} \frac{(f_1(z) - f_2(z))^2}{\sum_{\tau=1}^{k-1} \frac{(f_1(z^\tau) - f_2(z^\tau))^2}{(\sigma^\tau)^2} + \rho}.$$

We write $d_\nu := \frac{1}{H} \sum_{h=1}^H \dim_{\nu, K}(\mathcal{F}_h)$ for simplicity.

By Theorem 4.6 of Zhao et al. (2023), the generalized Eluder dimension is upper bounded by the standard Eluder dimension (Russo & Van Roy, 2013) up to logarithmic terms.

4. Algorithm

In this section, we introduce an algorithm, which, to the best of our knowledge, is the first to provide statistical guarantees in combinatorial RL with preference feedback while maintaining computational tractability. **Step 1** involves on-line parameter estimation for the MNL preference model, proposed by Lee & Oh (2024). **Steps 2, 3, and 4** implement variance-weighted regression to tighten the regret bound, as outlined in Agarwal et al. (2023). **Step 5**, which is our main

contribution, ensures optimism in a computationally efficient manner, even with uncertainty in item-level Q -values. **Step 6** introduces an exploration step that accounts for estimation errors from the MNL preference model.

Step 1. Online parameter estimation for MNL (Line 6).

At episode k and horizon h , given the user's choice feedback $c_h^k \in A_h^k$, the response for each item $a_{i_m} \in A_h^k$ is defined as $y_h^k(a_{i_m}) := \mathbb{1}(c_h^k = a_{i_m}) \in \{0, 1\}$. Therefore, the response variable $\mathbf{y}_h^k := (y_h^k(a_0), y_h^k(a_{i_1}), \dots, y_h^k(a_{i_l}))$, where $l \leq M - 1$, is sampled from a multinomial distribution: $\mathbf{y}_h^k \sim \text{MNL}\{1, \mathcal{P}_h(a_0|s_h^k, A_h^k; \theta_h^*), \dots, \mathcal{P}_h(a_l|s_h^k, A_h^k; \theta_h^*)\}$, where the parameter 1 indicates that \mathbf{y}_h^k is a single-trial sample, i.e., $y_h^k(a_0) + \sum_{m=1}^l y_h^k(a_{i_m}) = 1$. Then, for any $(k, h) \in [K] \times [H]$, the multinomial logistic loss function is defined as:

$$\ell_h^k(\theta) := - \sum_{a \in A_h^k} y_h^k(a) \log \mathcal{P}_h(a|s_h^k, A_h^k; \theta).$$

Inspired by Zhang & Sugiyama (2024); Lee & Oh (2024), for all $(k, h) \in [K] \times [H]$, we use the online mirror descent algorithm to estimate the true parameter θ_h^* as follows:

$$\theta_h^{k+1} \in \underset{\theta \in \Theta}{\operatorname{argmin}} \langle \nabla \ell_h^k(\theta_h^k), \theta \rangle + \frac{1}{2\eta} \|\theta - \theta_h^k\|_{\mathbf{H}_h^k}^2, \quad (1)$$

where $\eta = \mathcal{O}(\log M)$ is the step-size, $\tilde{\mathbf{H}}_h^k := \mathbf{H}_h^k + \eta \nabla^2 \ell_h^k(\theta_h^k)$, and $\mathbf{H}_h^k := \lambda \mathbf{I}_d + \sum_{\tau=1}^{k-1} \nabla^2 \ell_h^\tau(\theta_h^{\tau+1})$.

Remark 4.1. The computation cost of the optimization problem in (1) is $\mathcal{O}(Md^3)$, which does not scale with k at all.

Then, with high probability, θ_h^* lies within the following confidence interval (Corollary D.2).

$$\mathcal{C}_h^k := \left\{ \theta \in \Theta : \|\theta - \theta_h^k\|_{\mathbf{H}_h^k} \leq \alpha_h^k = \tilde{\mathcal{O}}(\sqrt{d}) \right\}. \quad (2)$$

Step 2. Weighted regression and optimistic estimation for \bar{Q} (Line 7-9). Using the past dataset, we solve the following (weighted) regression problem to fit $\mathcal{T}_h V_{h+1}^k$:

$$\hat{f}_{h,1}^k \in \underset{f_h \in \mathcal{F}_h}{\operatorname{argmin}} \sum_{\tau=1}^{k-1} \frac{\left(f_h(s_h^\tau, a_h^\tau) - r_h^\tau - V_{h+1,1}^k(s_h^\tau, a_h^\tau) \right)^2}{(\bar{\sigma}_h^\tau)^2}, \quad (3)$$

where $(\bar{\sigma}_h^k)^2$ is a variance upper bound, i.e., $(\bar{\sigma}_h^k)^2 \geq \mathbb{V}[r_h + V_{h+1,1}^k(s_{h+1})|s_h^k, a_h^k]$, which will be specified later. Then, an *optimistic \bar{Q} -value* estimate at horizon h is defined as $f_{h,1}^k := \hat{f}_{h,1}^k + b_h^k$, where b_h^k is the optimistic bonus. The bonus is calculated as $b_h^k(s, a) = \max_{f_h \in \mathcal{F}_h} f_h(s, a) - \min_{f_h \in \mathcal{F}_h} f_h(s, a)$. In general, this uncertainty bonus has a high complexity, as the maximizing and minimizing functions can differ arbitrarily for each $(s, a) \in \mathcal{S} \times \mathcal{I}$. To

address this, we use a *low-complexity bonus oracle* \mathcal{B} (Agarwal et al., 2023) that approximately dominates the value obtained from the point-wise maximization over \mathcal{F}_h . With the oracle \mathcal{B} , we can efficiently calculate the bonus b_h^k , with an error of ϵ_b . Due to space constraints, we provide the formal definition of \mathcal{B} in Appendix B (Definition B.1).

Step 3. Overly optimistic/pessimistic estimation for \bar{Q} (Line 10-13). For a sharp analysis of the convergence of the optimistic estimate $f_{h,1}^k$, we define an *overly optimistic \bar{Q} -value* estimate $f_{h,2}^k$, as well as an *overly pessimistic \bar{Q} -value* estimate $f_{h,-2}^k$. Similarly to $f_{h,1}^k$, they are calculated by solving an *unweighted* regression problem (Line 10), and by adding (or subtracting) a bonus function, which is the output of the bonus oracle \mathcal{B} (Line 12-13).

Step 4. Variance estimation (Line 14 and 21). To calculate $\bar{\sigma}_h^k$ introduced in (3), we first estimate the second moment by solving the *unweighted* regression problem:

$$\hat{g}_h^k \in \underset{g_h \in \mathcal{F}_h}{\operatorname{argmin}} \sum_{\tau=1}^{k-1} \left(g_h(s_h^\tau, a_h^\tau) - (r_h^\tau + V_{h+1,1}^k(s_h^\tau, a_h^\tau)) \right)^2.$$

Then, denoting $z_h^k = (s_h^k, a_h^k)$ for simplicity, we calculate the estimated variance as follows (this is an informal description; for the precise formulation, see Equations (D.9) and (D.10) in Appendix):

$$\begin{aligned} (\sigma_h^k)^2 &\simeq \hat{g}_h^k(z_h^k) - \left(\hat{f}_{h,-2}^k(z_h^k) \right)^2 + \mathfrak{D}_{k,h}^1 \cdot \mathcal{O} \left(\sqrt{\log \mathcal{N} \mathcal{N}_b} \right) \\ \bar{\sigma}_h^k &\simeq \max \left\{ \sigma_h^k, \mathcal{O}(\log \mathcal{N} \mathcal{N}_b) \right. \\ &\quad \cdot \left(\max \{ f_{h,2}^k(z_h^k) - f_{h,-2}^k(z_h^k), \mathfrak{D}_{k,h}^\sigma \} \right) \left. \right\}, \end{aligned} \quad (4)$$

where $\mathfrak{D}_{k,h}^1 = D_{\mathcal{F}_h}(z_h^k; \mathbf{Z}_{k-1}, \{\mathbf{1}^\tau\}_{\tau=1}^{k-1})$ and $\mathfrak{D}_{k,h}^\sigma = D_{\mathcal{F}_h}(z_h^k; \mathbf{Z}_{k-1}, \sigma_{k-1})$.

Step 5. Efficient optimistic Q -value estimation based on unknown item values (Line 15-16). In this step, we address our main challenge: selecting an optimistic assortment based on the optimistic \bar{Q} -values, which incorporate uncertainty, while ensuring computational tractability.

To introduce optimism and encourage exploration, we need to solve the following optimization problem using the optimistic (or pessimistic) estimates of the \bar{Q} -values, specifically $f_{h,j}^k$ for $j = 1, \pm 2$:

$$A_{h,j}^k \in \underset{A \in \mathcal{A}}{\operatorname{argmax}} \max_{\theta \in \mathcal{C}_h^k} \sum_{a \in A} \mathcal{P}_h^k(a|s_h^k, A; \theta) f_{h,j}^k(s_h^k, a), \quad (5)$$

where \mathcal{C}_h^k is defined in (2). One naive approach to solving the optimization problem in (5) is to add bonus terms to the estimation for each assortment A and then enumerate all $A \in \mathcal{A}$ to find the maximum. However, this approach results in a computational cost of $\mathcal{O}(|\mathcal{I}|^M)$.

Algorithm 1 MNL-VQL, MNL Preference Model with Variance-weighted Item-level Q-Learning

```

1: Inputs: parameter space  $\Theta$ , function class  $\{\mathcal{F}_h\}_{h=1}^H$ , consistent bonus oracle  $\mathcal{B}$ .
2: Parameters:  $\{\alpha_h^k, \beta_{h,1}^k, \beta_{h,2}^k, \bar{\beta}_h^k\}_{(k,h) \in [H] \times [K]}$ ,  $\{u_k\}_{k=1}^K$ ,  $\rho$ , bonus error  $\epsilon_b$ ,  $\nu$ ,  $\delta$ .
3: Initialize: dataset  $\mathcal{D}_h^0 = \emptyset$  for all  $h \in [H]$ .
4: Generate  $\{\mathcal{D}_h^1\}_{h=1}^H$  from initial state  $s_1^1$  by random policy and set  $\sigma_h^1 = \bar{\sigma}_h^1 = 2$  for all  $h \in [H]$ .
5: for episode  $k = 2, \dots, K$  do
6:   for horizon  $h = H, H-1, \dots, 1$  do
       // CALCULATE OPTIMISTIC AND OVERLY OPTIMISTIC, PESSIMISTIC  $\bar{Q}$ -VALUES
7:    $\hat{f}_{h,1}^k \in \operatorname{argmin}_{f_h \in \mathcal{F}_h} \sum_{\tau=1}^{k-1} \frac{1}{(\bar{\sigma}_h^\tau)^2} \left( f_h(s_h^\tau, a_h^\tau) - r_h^\tau - V_{h+1,1}^k(s_{h+1}^\tau) \right)^2$ .
8:    $b_{h,1}^k \leftarrow \mathcal{B} \left( \{\bar{\sigma}_h^\tau\}_{\tau=1}^{k-1}, \mathcal{D}_h^{k-1}, \mathcal{F}_h, \hat{f}_{h,1}^k, \beta_{h,1}^k, \rho, \epsilon_b \right)$  (see Definition B.1).
9:   Update  $f_{h,1}^k(\cdot, \cdot) \leftarrow \min \left\{ \hat{f}_{h,1}^k(\cdot, \cdot) + b_{h,1}^k(\cdot, \cdot), 1 \right\}$ .
10:   $\hat{f}_{h,j}^k \in \operatorname{argmin}_{f_h \in \mathcal{F}_h} \sum_{\tau=1}^{k-1} \left( f_h(s_h^\tau, a_h^\tau) - r_h^\tau - V_{h+1,j}^k(s_{h+1}^\tau) \right)^2$ ,  $j = \pm 2$ .
11:   $b_{h,2}^k \leftarrow \mathcal{B} \left( \{\bar{\sigma}_h^\tau\}_{\tau=1}^{k-1}, \mathcal{D}_h^{k-1}, \mathcal{F}_h, \hat{f}_{h,2}^k, \beta_{h,2}^k, \rho, \epsilon_b \right)$  (see Definition B.1).
12:  Update  $f_{h,2}^k(\cdot, \cdot) \leftarrow \min \left\{ \hat{f}_{h,2}^k(\cdot, \cdot) + 2b_{h,1}^k(\cdot, \cdot) + b_{h,2}^k(\cdot, \cdot), 1 \right\}$ .
13:  Update  $f_{h,-2}^k(\cdot, \cdot) \leftarrow \max \left\{ \hat{f}_{h,-2}^k(\cdot, \cdot) - b_{h,2}^k(\cdot, \cdot), 0 \right\}$ .
14:   $\hat{g}_h^k \in \operatorname{argmin}_{g_h \in \mathcal{F}_h} \sum_{\tau=1}^{k-1} \left( g_h(s_h^\tau, a_h^\tau) - \left( r_h^\tau + V_{h+1,1}^k(s_{h+1}^\tau) \right) \right)^2$ .
       // UPDATE VALUES
15:  Update  $\tilde{\mathcal{P}}_{h,j}^k(\cdot, \cdot)$  by (7).
16:  Update  $Q_{h,j}^k(\cdot, A) \leftarrow \sum_{a \in A} \tilde{\mathcal{P}}_{h,j}^k(a|\cdot, A) f_{h,j}^k(\cdot, a)$  and  $V_{h,j}^k(\cdot) \leftarrow \max_{A \in \mathcal{A}} Q_{h,j}^k(\cdot, A)$ ,  $j = 1, \pm 2$ .
17:  end for
18:  Receive initial state  $s_1^k$ .
19:  for  $h = 1, 2, \dots, H$  do
20:    Offer  $A_h^k$  by (9) and receive  $a_h^k, r_h^k$ , and  $s_{h+1}^k$ .
21:    Update  $\mathcal{D}_h^k \leftarrow \mathcal{D}_h^{k-1} \cup \{s_h^k, a_h^k, r_h^k, s_{h+1}^k\}$  and update  $\sigma_h^k$  and  $\bar{\sigma}_h^k$  by (4).
22:    Update  $\hat{\theta}_h^{k+1}$  using online mirror descent (1).
23:  end for
24: end for

```

To avoid this exponential computational cost, inspired by Tran-Dinh et al. (2015), we use optimistic MNL utilities instead of directly adding bonus terms to $\sum_{a \in A} \mathcal{P}_h^k(a|s_h^k, A; \theta_h^k) f_{h,j}^k(s_h^k, a)$. However, unlike traditional MNL bandits (Oh & Iyengar, 2019; 2021; Lee & Oh, 2024), simply using optimistic utilities does not always guarantee optimism because of $f_{h,j}^k$ is not the true values. In MNL bandits, the item values are known, and the value of the outside option a_0 is fixed at zero. Therefore, increasing the MNL utilities (using the optimistic utilities) lowers the probability of choosing the outside option, which in turn increases the expected value of the item values.

However, in our setting, using the optimistic utilities can decrease the expected value of $f_{h,j}^k$. To explain why: even if the true value of the outside option, $\bar{Q}_h^*(s, a_0)$, is the lowest, its estimated value, $f_{h,j}^k(s, a_0)$, can be the highest—i.e., $f_{h,j}^k(s, a_0) > f_{h,j}^k(s, a)$ for all $a \in \mathcal{I} \setminus \{a_0\}$ —due to uncertainty. In this case, increasing the MNL utilities results

in a decrease in the expected value of $f_{h,j}^k$. This challenge arises from the unknown item values, \bar{Q}_h^* , which is one of the main difficulties we face in our framework.

To tackle this problem, a more refined approach is required to use the utility based on $f_{h,j}^k$. Given the confidence interval in (2), we define the *optimistic utility* $\tilde{v}_h^k(s, a)$ and the *pessimistic utility* $\check{v}_h^k(s, a)$ as:

$$\begin{aligned} \tilde{v}_h^k(s, a) &:= \phi(s, a)^\top \theta_h^k + \alpha_h^k \|\phi(s, a)\|_{(\mathbf{H}_h^k)^{-1}}, \\ \check{v}_h^k(s, a) &:= \phi(s, a)^\top \theta_h^k - \alpha_h^k \|\phi(s, a)\|_{(\mathbf{H}_h^k)^{-1}}. \end{aligned} \quad (6)$$

We then use the optimistic utility when $f_{h,j}^k(s, a_0)$ is not the highest estimate to calculate the *optimistic choice* probabilities $\tilde{\mathcal{P}}_{h,j}^k$. Formally, let $I_{h,j}^k \in \{1, 0\}$ indicate whether there exists $a \in \mathcal{I} \setminus \{a_0\}$ such that $f_{h,j}^k(s, a) \geq f_{h,j}^k(s, a_0)$

($I_{h,j}^k = 1$ if such an event occurs). Then, we define

$$\tilde{\mathcal{P}}_{h,j}^k(a|s, A) := \begin{cases} \frac{\exp(\tilde{v}_h^k(s, a))}{\sum_{a' \in A} \exp(\tilde{v}_h^k(s, a'))}, & \text{if } I_{h,j}^k = 1 \\ \frac{\exp(\tilde{v}_h^k(s, a))}{\sum_{a' \in A} \exp(\tilde{v}_h^k(s, a'))}, & \text{otherwise.} \end{cases} \quad (7)$$

Equipped with $\tilde{\mathcal{P}}_{h,j}^k$, for any $j = 1, \pm 2$, we select the assortments $A_{h,j}^k$ as follows:

$$A_{h,j}^k \in \operatorname{argmax}_{A \in \mathcal{A}} \underbrace{\sum_{a \in A} \tilde{\mathcal{P}}_{h,j}^k(a|s_h^k, A) f_{h,j}^k(s_h^k, a)}_{=: Q_{h,j}^k(s_h^k, A)}, \quad (8)$$

Here, $Q_{h,j}^k(s_h^k, A) := \sum_{a \in A} \tilde{\mathcal{P}}_{h,j}^k(a|s_h^k, A) f_{h,j}^k(s_h^k, a)$ is the *optimistic* Q -values. This construction can induce sufficient exploration and guarantee optimism (Lemma D.15). Furthermore, by using the optimistic (or pessimistic) utilities for each item, instead of calculating bonus terms for each $A \in \mathcal{A}$, the optimization problem in (8) can be solved efficiently (Davis et al., 2013).

Remark 4.2. The optimization problem in (8) can be transformed into a linear programming (LP), making it solvable in polynomial time with respect to $|\mathcal{I}|$ (see Appendix C).

Step 6. Exploration policy (Line 20). We then offer the assortment A_h^k to the user as follows:

$$A_h^k = \begin{cases} A_{h,1}^k & \text{if } f_{h',1}^k(s_{h'}^k, a_{h'}) \geq f_{h',2}^k(s_{h'}^k, a_{h'}) - u_k, \\ & \forall a_{h'} \in A_{h,1}^k, \forall h' \leq h, \\ A_{h,2}^k & \text{otherwise,} \end{cases} \quad (9)$$

where u_k is a carefully chosen threshold (see Table D.2 for the exact value). When the optimistic sequence $f_{h,1}^k$ and the overly optimistic sequence $f_{h,2}^k$ diverge beyond a certain threshold, we offer the assortment $A_{h,2}^k$, which is selected based on $f_{h,2}^k$. This approach ensures that by occasionally using $f_{h,2}^k$, the variance upper bound $\bar{\sigma}_h^k$, estimated from $f_{h,2}^k$, does not become overly pessimistic.

5. Main Results

5.1. Non-linear Function Approximation for \bar{Q}

Theorem 5.1 (Informal, Regret upper bound of MNL-VQL). *Let $d_\nu = \frac{1}{H} \sum_{h=1}^H \dim_{\nu,K}(\mathcal{F}_h)$ with $\nu = \sqrt{1/KH}$, and set u_k as in Equation (D.20). Suppose Assumptions 3.1 and 3.3 hold. Then, with probability at least $1 - \delta$, the*

regret of MNL-VQL is upper-bounded by:

$$\begin{aligned} \mathbf{Regret}(\mathcal{M}, K) &\lesssim \underbrace{d\sqrt{HK} + \frac{1}{\kappa} d^2 H^2}_{\text{regret from MNL model}} \\ &\quad + \underbrace{\sqrt{d_\nu HK \log \mathcal{N}} + d_\nu H^5 \log \mathcal{N} \log^2(\mathcal{N} \mathcal{N}_b)}_{\text{regret from general function approximation of } \bar{Q}}, \end{aligned}$$

where d is the feature dimension of the MNL preference model, $\mathcal{N} = \max_{h \in [H]} |\mathcal{F}_h|$, and \mathcal{N}_b is the size of the bonus function class, i.e., $\mathcal{N}_b = |\mathcal{W}|$.

Discussion of Theorem 5.1. The proof is deferred to Appendix D. The first two terms arise from the regret of the MNL preference model, while the other two terms come from the regret associated with the general function approximation for item-level Q -values. When $H = 1$, reducing our setting to MNL bandits (though not exactly the traditional MNL bandits, as we consider a more general case where item values are unknown and the value of the outside option can be non-zero), the first two terms of our regret simplify to $\tilde{O}(d\sqrt{K} + \frac{1}{\kappa} d^2)$. This matches the known minimax optimal regret established by Lee & Oh (2024). Note that we avoid the detrimental dependence on κ in our leading term. The last two terms of our regret, incurred from estimating item-level Q -values using general function approximation, similar to Agarwal et al. (2023) and Zhao et al. (2023).

With respect to computational cost, by using the online sensitivity sub-sampling method (Algorithm B.1), we can efficiently implement the bonus oracle \mathcal{B} with $\log |\mathcal{W}| = \log \mathcal{N}_b = \tilde{O}(\max_{h \in [H]} \dim_{\nu,K}(\mathcal{F}_h) \log \mathcal{N} \log |\mathcal{S} \times \mathcal{I}|)$. Furthermore, we can avoid the exponential computational cost required to solve the optimization in (8) (see Remark 4.2). As a result, our algorithm is both computationally tractable and statistically efficient.

5.2. Technical Comparisons to Related Work

Comparison to Lee & Oh (2024). Our framework addresses a strictly more challenging problem than Lee & Oh (2024) because we consider (1) *multiple steps*, (2) *unknown values*, and (3) *nonzero values for the outside option*. The key challenge in tackling these three aspects lies in ensuring optimism while maintaining computational efficiency, which requires a fundamentally different approach—carefully leveraging optimistic and pessimistic utilities. Moreover, in Theorem 5.1, our regret analysis for the MNL preference model goes beyond merely summing over H MNL bandit regrets. Instead, we introduce a novel regret decomposition, bound the regret in terms of the sum of the variances of value functions, and apply the law of total variance (Lattimore & Hutter, 2012; Gheshlaghi Azar et al., 2013). As a result, we achieve a regret reduction by a

factor of \sqrt{H} compared to directly summing over H MNL bandit regrets.

Comparison to Agarwal et al. (2023). While the regret analysis for general function approximation largely follows the approach of Agarwal et al. (2023), several technical lemmas (e.g., Lemmas D.23 and D.24) and parameters (e.g., u_k) are revised to accommodate the estimation errors specific to MNL models.

Overall, our results cannot be gleaned by simply piecing together prior techniques. Instead, they arise from an involved analysis, leading to stronger theoretical guarantees in more general and new combinatorial RL settings.

5.3. Linear MDPs with Preference Feedback

As a special case, we also consider linear MDPs (refer Definition E.1) with preference feedback. To show the dependency on parameters, we denote the linear MDPs as \mathcal{M}_{Ξ^*} , where $\Xi^* = \{\{\theta_h^*\}_{h=1}^H, \{\mu_h^*\}_{h=1}^H, \{w_h^*\}_{h=1}^H\}$. Note that the bonus oracle can be easily implemented using the standard elliptical bonus, which satisfies all the necessary properties (refer Appendix E). The proof is deferred to Appendix E.

Theorem 5.2 (Informal, Regret upper bound for linear MDPs). *In linear MDPs, under the same conditions as Theorem 5.1, with probability at least $1 - \delta$, the regret of MNL-VQL is upper-bounded by:*

$$\text{Regret}(\mathcal{M}_{\Xi^*}, K) \lesssim d\sqrt{HK} + \frac{1}{\kappa}d^2H^2 + d^{\text{lin}}\sqrt{HK} + (d^{\text{lin}})^6H^5.$$

We also establish a matching lower bound by constructing a novel multi-layered (linear) MDP (see Figure F.1) with a preference feedback. The proof is deferred to Appendix F.

Theorem 5.3 (Informal, Regret lower bound for linear MDPs). *For any algorithm and sufficiently large K , there exists an episodic linear MDP \mathcal{M}_{Ξ} with MNL preference feedback such that the worst-case expected regret is lower bounded as follows:*

$$\sup_{\Xi} \mathbb{E}_{\Xi} [\text{Regret}(\mathcal{M}_{\Xi}, K)] = \Omega(d\sqrt{HK} + d^{\text{lin}}\sqrt{HK}).$$

Discussion of Theorems 5.2 and 5.3. For sufficiently large K , i.e., $K \geq \tilde{O}(d^2H^3/\kappa^2 + (d^{\text{lin}})^{10}H^9)$, the regret upper bound for linear MDPs scales as $\tilde{O}(d\sqrt{HK} + d^{\text{lin}}\sqrt{HK})$, which matches the lower bound up to logarithmic factors. Note that if we rescale the rewards to be $1/H$ in the lower bounds of Zhou et al. (2021a) to align with our setting, their regret bound matches the second term of our regret bound, $\Omega(d^{\text{lin}}\sqrt{HK})$. To the best of our knowledge, this is the first theoretical result proving minimax-optimality in linear MDPs with preference feedback.

6. Numerical Experiment

In this section, we empirically evaluate the performance of our algorithm, MNL-VQL, in two settings: a synthetic environment (Subsection 6.1) and a real-world dataset (Subsection 6.2).

We compare our algorithm against two baselines: **Myopic** and **LSVI-UCB** (Jin et al., 2020). **Myopic** is a variant of OFU-MNL+ (Lee & Oh, 2024) adapted for *unknown* rewards. It selects assortments based only on immediate rewards, ignoring long-term effects. **LSVI-UCB** (Jin et al., 2020) treats each assortment as a single atomic action, requiring enumeration of all possible assortments. We also include the optimal policy (**Optimal**) as a reference.

6.1. Synthetic Environment

Setup. We consider an *online shopping with budget environment* (Figure G.1), modeled as a linear MDP with an MNL preference model. Let $\mathcal{S} = \{s_1, \dots, s_{|\mathcal{S}|}\}$ denote the set of states and $\mathcal{I} = \{a_1, \dots, a_N, a_0\}$ the set of items, where a_0 represents the outside option (no purchase). Each state $s_j \in \mathcal{S}$ corresponds to a *user's budget level*, with larger indices indicating a higher budget. The initial state is set to the middle budget level, $s_{\lfloor |\mathcal{S}|/2 \rfloor}$. The transition probabilities \mathbb{P}_h , rewards r_h , and preference model \mathcal{P}_h remain constant across all time steps $h \in [H]$, so we omit the subscript h .

At state s_j , the agent offers an assortment $A \in \mathcal{A}$ with a maximum size of M . The user either purchases an item $a_i \in A$ or does not ($a_0 \in A$). If the user buys item a_i , the agent receives a reward $r(s_j, a_i) = \left(\frac{i}{100N} + \frac{j}{|\mathcal{S}|}\right)/H$ and the state transitions according to $\mathbb{P}(s_{\min(j+1, |\mathcal{S}|)} | s_j, a_i) = 1 - \frac{j}{N}$, and $\mathbb{P}(s_{\max(j-1, 0)} | s_j, a_i) = \frac{j}{N}$. If the user does not buy anything (a_0), the reward is $r(s_j, a_0) = 0$, and the state transitions as $\mathbb{P}(s_{\min(j+1, |\mathcal{S}|)} | s_j, a_0) = 1$. For the MNL preference model, the true parameter $\theta^* \in \mathbb{R}^d$ and the feature vector $\phi(s, a) \in \mathbb{R}^d$ are randomly sampled from a d -dimensional uniform distribution for each instance.

Results. Figure 1 demonstrates that our algorithm significantly outperforms other baseline algorithms. Remarkably, **Myopic** converges to a suboptimal solution, highlighting the importance of accounting for long-term values. Moreover, in Appendix Table G.1, we show that our algorithm is much faster than others, especially when the total number of assortments $|\mathcal{A}|$ is large. Due to the extremely slow runtime of LSVI-UCB, we could not include its performance results for $N = 20$ and $N = 40$. For more details, see Appendix G.

6.2. Real-World MovieLens Experiment

Setup. The MovieLens dataset contains 25 million ratings on a 5-star scale for 62,000 movies (base items a) provided by 162,000 users (u). We define the state s as the number




Figure 1: **Synthetic experiment:** Episodic returns averaged over 10 runs. Dotted lines indicate estimated returns for incomplete runs due to excessive runtime. Shading denotes ± 1 standard deviation.




Figure 2: **MovieLens experiment:** The dotted lines represent estimated (virtual) episodic returns for cases that could not be run due to excessively long runtimes. Shaded regions represent ± 1 standard deviation.

of movies a user u has watched after entering the system, denoted by $s = (u, n)$, where $n \in \{0, \dots, H - 1\}$ is the number of movies watched during the session. We interpret the ratings as representing MNL utilities.

In each episode k , a user (u_k) is randomly sampled and arrives at the recommender system, initiating the state $s_1^k = (u_k, 0)$. The agent offers a set of items with a maximum size of M . If the user clicks on an item, they receive a reward of 1 and transition to the next state $s_2^k = (u_k, 1)$. If no item is clicked, the user receives no reward and remains in the current state ($s_2^k = s_1^k$). In addition, certain *junk* items—such as those with a provocative title and poster but poor content—can cause users to leave the system immediately. This is modeled as a transition to an *absorbing state*, where no further rewards are received and the state remains unchanged regardless of future actions. We believe the presence of such junk items is quite natural and reflective of real-world recommendation environments.

For our experiments, we use a subset of the dataset containing 1.1×10^3 users and a varying number of movies, $N \in \{50, 100, 200\}$. To construct MNL features, we follow a similar experimental setup as in Li et al. (2019), employing low-rank matrix factorization. For linear MDP features, we apply the same approach as used in our syn-

thetic data experiments. We set the parameters as follows: $K = 10000$, $H = 3$, $M = 4$, $|\mathcal{S}| = 100 * (H + 1) = 400$ (including the absorbing state), $d = 26$ (MNL feature dimension), $d^{\text{lin}} = 204$ (Linear MDP feature dimension), $N \in \{50, 100, 200\}$ (number of base items) and $|\mathcal{A}| = \sum_{m=1}^{M-1} \binom{N}{m} \in \{20875, 166750, 1333500\}$. The proportion of junk items is set to 30%.

Results. Consistent with the synthetic experiment results, Figure 2 shows that our algorithm substantially outperforms the baseline methods on the real-world dataset. This demonstrates the robustness of our approach and its practical effectiveness in realistic settings.

7. Conclusion

In this work, we study combinatorial RL with preference feedback, extending MNL bandit problems to account for the influence of user states and state transitions in applications like recommendation systems. Under an MNL preference model with linear utilities and general function approximation for item values, we propose an efficient algorithm, MNL-VQL, which, to the best of our knowledge, provides the first statistical guarantee. As a special case, in linear MDPs, we show the minimax-optimality of MNL-VQL by establishing matching upper and lower bounds.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

Acknowledgements

This work was supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (No. RS-2022-NR071853 and RS-2023-00222663) and by AI-Bio Research Grant through Seoul National University.

References

- Abeille, M., Fauray, L., and Calauzènes, C. Instance-wise minimax-optimal algorithms for logistic bandits. In *International Conference on Artificial Intelligence and Statistics*, pp. 3691–3699. PMLR, 2021.
- Agarwal, A., Jin, Y., and Zhang, T. Voq1: Towards optimal regret in model-free rl with nonlinear function approximation. In *The Thirty Sixth Annual Conference on Learning Theory*, pp. 987–1063. PMLR, 2023.
- Agarwal, P., Tulabandhula, T., and Avadhanula, V. A tractable online learning algorithm for the multinomial logit contextual bandit. *European Journal of Operational Research*, 310(2):737–750, 2023.
- Agarwal, S. and Goyal, N. Thompson sampling for contextual bandits with linear payoffs. In *International Conference on Machine Learning*, pp. 127–135. PMLR, 2013.
- Agarwal, S., Avadhanula, V., Goyal, V., and Zeevi, A. Thompson sampling for the mnl-bandit. In *Conference on learning theory*, pp. 76–78. PMLR, 2017.
- Agarwal, S., Avadhanula, V., Goyal, V., and Zeevi, A. Mnl-bandit: A dynamic learning approach to assortment selection. *Operations Research*, 67(5):1453–1485, 2019.
- Akrou, R., Schoenauer, M., and Sebag, M. April: Active preference learning-based reinforcement learning. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2012, Bristol, UK, September 24-28, 2012. Proceedings, Part II* 23, pp. 116–131. Springer, 2012.
- Amani, S. and Thrampoulidis, C. Ucb-based algorithms for multinomial logistic regression bandits. *Advances in Neural Information Processing Systems*, 34:2913–2924, 2021.
- Chaudhari, S., Arbour, D., Theodorou, G., and Vlassis, N. Distributional off-policy evaluation for slate recommendations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 8265–8273, 2024.
- Chen, K. D. and Hausman, W. H. Mathematical properties of the optimal product line selection problem using choice-based conjoint analysis. *Management Science*, 46(2):327–332, 2000.
- Chen, W., Wang, Y., and Yuan, Y. Combinatorial multi-armed bandit: General framework and applications. In *International conference on machine learning*, pp. 151–159. PMLR, 2013.
- Cheung, W. C. and Simchi-Levi, D. Thompson sampling for online personalized assortment optimization problems with multinomial logit choice models. *Available at SSRN 3075658*, 2017.
- Christiano, P. F., Leike, J., Brown, T., Martic, M., Legg, S., and Amodei, D. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30, 2017.
- Combes, R., Talebi Mazraeh Shahi, M. S., Proutiere, A., et al. Combinatorial bandits revisited. *Advances in neural information processing systems*, 28, 2015.
- Cooper, A. C. W. et al. Programming with linear fractional functionals. *Naval Research logistics quarterly*, 9(3): 181–186, 1962.
- Davis, J., Gallego, G., and Topaloglu, H. Assortment planning under the multinomial logit model with totally unimodular constraint structures. department of ieor, columbia university, 2013.
- Delarue, A., Anderson, R., and Tjandraatmadja, C. Reinforcement learning with combinatorial actions: An application to vehicle routing. *Advances in Neural Information Processing Systems*, 33:609–620, 2020.
- Du, S., Kakade, S., Lee, J., Lovett, S., Mahajan, G., Sun, W., and Wang, R. Bilinear classes: A structural framework for provable generalization in rl. In *International Conference on Machine Learning*, pp. 2826–2836. PMLR, 2021.
- Du, Y., Srikant, R., and Chen, W. Cascading reinforcement learning. *arXiv preprint arXiv:2401.08961*, 2024.
- Fauray, L., Abeille, M., Calauzènes, C., and Fercoq, O. Improved optimistic algorithms for logistic bandits. In *International Conference on Machine Learning*, pp. 3052–3060. PMLR, 2020.
- Fauray, L., Abeille, M., Jun, K.-S., and Calauzènes, C. Jointly efficient and optimal algorithms for logistic bandits. In

- International Conference on Artificial Intelligence and Statistics*, pp. 546–580. PMLR, 2022.
- Filippi, S., Cappé, O., Garivier, A., and Szepesvári, C. Parametric bandits: The generalized linear case. In *Proceedings of the 23rd International Conference on Neural Information Processing Systems - Volume 1, NIPS’10*, pp. 586–594, Red Hook, NY, USA, 2010. Curran Associates Inc.
- Foster, D. J., Kakade, S. M., Qian, J., and Rakhlin, A. The statistical complexity of interactive decision making. *arXiv preprint arXiv:2112.13487*, 2021.
- Gheshlaghi Azar, M., Munos, R., and Kappen, H. J. Minimax pac bounds on the sample complexity of reinforcement learning with a generative model. *Machine learning*, 91:325–349, 2013.
- He, J., Ostendorf, M., He, X., Chen, J., Gao, J., Li, L., and Deng, L. Deep reinforcement learning with a combinatorial action space for predicting popular reddit threads. *arXiv preprint arXiv:1606.03667*, 2016.
- Ie, E., Jain, V., Wang, J., Narvekar, S., Agarwal, R., Wu, R., Cheng, H.-T., Chandra, T., and Boutilier, C. Slateq: A tractable decomposition for reinforcement learning with recommendation sets. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pp. 2592–2599, 2019.
- Ishfaq, H., Cui, Q., Nguyen, V., Ayoub, A., Yang, Z., Wang, Z., Precup, D., and Yang, L. Randomized exploration in reinforcement learning with general value function approximation. In *International Conference on Machine Learning*, volume 139, pp. 4607–4616. PMLR, 2021.
- Jiang, N., Krishnamurthy, A., Agarwal, A., Langford, J., and Schapire, R. E. Contextual decision processes with low bellman rank are pac-learnable. In *International Conference on Machine Learning*, pp. 1704–1713. PMLR, 2017.
- Jin, C., Allen-Zhu, Z., Bubeck, S., and Jordan, M. I. Is q-learning provably efficient? In *Advances in Neural Information Processing Systems*, volume 31, pp. 4868–4878, 2018.
- Jin, C., Yang, Z., Wang, Z., and Jordan, M. I. Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*, pp. 2137–2143. PMLR, 2020.
- Jin, C., Liu, Q., and Miryoosefi, S. Bellman eluder dimension: New rich classes of rl problems, and sample-efficient algorithms. *Advances in neural information processing systems*, 34:13406–13418, 2021.
- Kong, D., Salakhutdinov, R., Wang, R., and Yang, L. F. Online sub-sampling for reinforcement learning with general function approximation. *arXiv preprint arXiv:2106.07203*, 2021.
- Kveton, B., Szepesvari, C., Wen, Z., and Ashkan, A. Cascading bandits: Learning to rank in the cascade model. In *International conference on machine learning*, pp. 767–776. PMLR, 2015a.
- Kveton, B., Wen, Z., Ashkan, A., and Szepesvari, C. Combinatorial cascading bandits. *Advances in Neural Information Processing Systems*, 28, 2015b.
- Lattimore, T. and Hutter, M. Pac bounds for discounted mdps. In *Algorithmic Learning Theory: 23rd International Conference, ALT 2012, Lyon, France, October 29-31, 2012. Proceedings 23*, pp. 320–334. Springer, 2012.
- Lee, J. and Oh, M.-h. Demystifying linear mdps and novel dynamics aggregation framework. In *The Twelfth International Conference on Learning Representations*, 2023.
- Lee, J. and Oh, M.-h. Nearly minimax optimal regret for multinomial logistic bandit. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- Li, S., Lattimore, T., and Szepesvári, C. Online learning to rank with features. In *International Conference on Machine Learning*, pp. 3856–3865. PMLR, 2019.
- McFadden, D. Modelling the choice of residential location. 1977.
- McInerney, J., Brost, B., Chandar, P., Mehrotra, R., and Carterette, B. Counterfactual evaluation of slate recommendations with sequential reward interactions. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 1779–1788, 2020.
- Metz, L., Ibarz, J., Jaitly, N., and Davidson, J. Discrete sequential prediction of continuous actions for deep rl. *arXiv preprint arXiv:1705.05035*, 2017.
- Mhammedi, Z., Koolen, W. M., and Van Erven, T. Lipschitz adaptivity with multiple learning rates in online learning. In *Conference on Learning Theory*, pp. 2490–2511. PMLR, 2019.
- Oh, M.-h. and Iyengar, G. Thompson sampling for multinomial logit contextual bandits. *Advances in Neural Information Processing Systems*, 32, 2019.
- Oh, M.-h. and Iyengar, G. Multinomial logit contextual bandits: Provable optimality and practicality. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 9205–9213, 2021.

- Orabona, F. A modern introduction to online learning. *arXiv preprint arXiv:1912.13213*, 2019.
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022.
- Perivier, N. and Goyal, V. Dynamic pricing and assortment under a contextual mnl demand. *Advances in Neural Information Processing Systems*, 35:3461–3474, 2022.
- Rusmevichientong, P., Shen, Z.-J. M., and Shmoys, D. B. Dynamic assortment optimization with a multinomial logit choice model and capacity constraint. *Operations research*, 58(6):1666–1680, 2010.
- Russo, D. and Van Roy, B. Eluder dimension and the sample complexity of optimistic exploration. In *Advances in Neural Information Processing Systems*, pp. 2256–2264, 2013.
- Ryu, M., Chow, Y., Anderson, R., Tjandraatmadja, C., and Boutilier, C. Caql: Continuous action q-learning. *arXiv preprint arXiv:1909.12397*, 2019.
- Saha, A., Pacchiano, A., and Lee, J. Dueling rl: Reinforcement learning with trajectory preferences. In *International Conference on Artificial Intelligence and Statistics*, pp. 6263–6289. PMLR, 2023.
- Sauré, D. and Zeevi, A. Optimal dynamic assortment planning with demand learning. *Manufacturing & Service Operations Management*, 15(3):387–404, 2013.
- Sunehag, P., Evans, R., Dulac-Arnold, G., Zwols, Y., Visentin, D., and Coppin, B. Deep reinforcement learning with attention for slate markov decision processes with high-dimensional states and actions. *arXiv preprint arXiv:1512.01124*, 2015.
- Swaminathan, A., Krishnamurthy, A., Agarwal, A., Dudik, M., Langford, J., Jose, D., and Zitouni, I. Off-policy evaluation for slate recommendation. *Advances in Neural Information Processing Systems*, 30, 2017.
- Tran-Dinh, Q., Li, Y.-H., and Cevher, V. Composite convex minimization involving self-concordant-like cost functions. In *Modelling, Computation and Optimization in Information Systems and Management Sciences: Proceedings of the 3rd International Conference on Modelling, Computation and Optimization in Information Systems and Management Sciences-MCO 2015-Part I*, pp. 155–168. Springer, 2015.
- Vlassis, N., Chandrashekar, A., Amat, F., and Kallus, N. Control variates for slate off-policy evaluation. *Advances in Neural Information Processing Systems*, 34: 3667–3679, 2021.
- Wang, R., Salakhutdinov, R. R., and Yang, L. Reinforcement learning with general value function approximation: Provably efficient approach via bounded eluder dimension. *Advances in Neural Information Processing Systems*, 33, 2020.
- Wang, Y., Wang, R., Du, S. S., and Krishnamurthy, A. Optimism in reinforcement learning with generalized linear function approximation. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=CBmJwzneppz>.
- Wirth, C., Akrou, R., Neumann, G., and Fürnkranz, J. A survey of preference-based reinforcement learning methods. *Journal of Machine Learning Research*, 18(136): 1–46, 2017.
- Yang, L. and Wang, M. Sample-optimal parametric q-learning using linearly additive features. In *International Conference on Machine Learning*, pp. 6995–7004. PMLR, 2019.
- Zhan, W., Uehara, M., Kallus, N., Lee, J. D., and Sun, W. Provable offline preference-based reinforcement learning. *arXiv preprint arXiv:2305.14816*, 2023.
- Zhang, Y.-J. and Sugiyama, M. Online (multinomial) logistic bandit: Improved regret and constant computation cost. *Advances in Neural Information Processing Systems*, 36, 2024.
- Zhao, H., He, J., and Gu, Q. A nearly optimal and low-switching algorithm for reinforcement learning with general function approximation. *arXiv preprint arXiv:2311.15238*, 2023.
- Zhou, D., Gu, Q., and Szepesvari, C. Nearly minimax optimal reinforcement learning for linear mixture markov decision processes. In *Conference on Learning Theory*, pp. 4532–4576. PMLR, 2021a.
- Zhou, D., He, J., and Gu, Q. Provably efficient reinforcement learning for discounted mdps with feature mapping. In *International Conference on Machine Learning*, pp. 12793–12802. PMLR, 2021b.
- Zhu, B., Jordan, M., and Jiao, J. Principled reinforcement learning with human feedback from pairwise or k-wise comparisons. In *International Conference on Machine Learning*, pp. 43037–43067. PMLR, 2023.

Appendix

A. Illustrative Explanation for Combinatorial RL with Preference Feedback

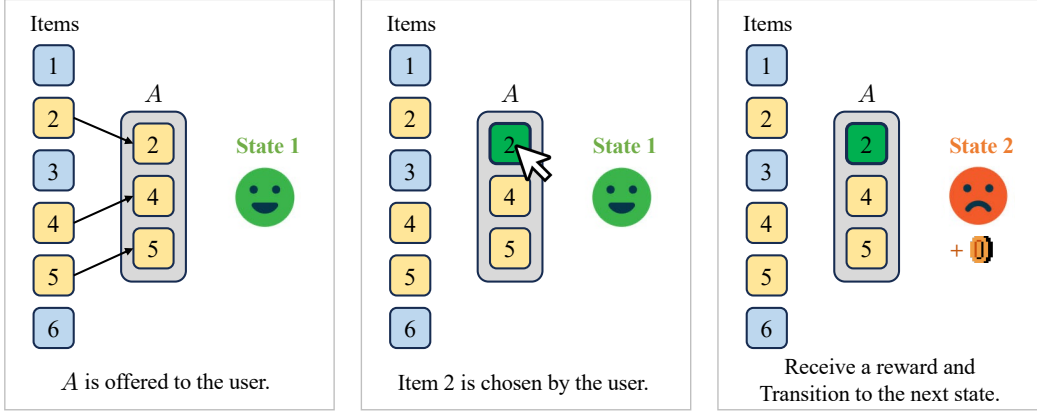


Figure A.1: Illustration of combinatorial RL with preference feedback.

In this section, we provide additional explanation of our framework, combinatorial RL with preference feedback, for better clarity. In this framework, at each episode, as a user arrives at the system (starting in the initial state, e.g., high loyalty), a learning agent selects an assortment A (a set of items) and offers it to the user (the first figure in Figure A.1). The user then chooses an item from the assortment A (the second figure in Figure A.1). The agent receives a reward, along with preference (or choice) feedback, and transitions to the next state (e.g., lower loyalty) (the last figure in Figure A.1). This process repeats until the episode concludes.

The key advantage of this framework is its ability to capture the long-term value of choosing an item by considering state transitions and avoiding myopic decisions. For instance, in Figure A.1, a user may choose a *junk* item that provides a high immediate reward. However, repeatedly recommending such items can lead to user fatigue, resulting in a transition to a state of lower satisfaction or loyalty to the system, ultimately leading to a lower cumulative reward.

We compare our framework with other related works.

vs MNL bandits. Our framework can be considered as a multi-step extension of MNL bandits (Rusmevichientong et al., 2010; Sauré & Zeevi, 2013; Agrawal et al., 2017; 2019; Oh & Iyengar, 2019; 2021; Perivier & Goyal, 2022; Agrawal et al., 2023; Zhang & Sugiyama, 2024; Lee & Oh, 2024). In MNL bandits, there are no state transitions; thus, in Figure F.1, the user exits the system immediately after receiving a reward.

Another important difference is that, in MNL bandits, the value (reward) of choosing an item is assumed to be known, and the value of choosing the outside option a_0 is always assumed to be zero. In contrast, in our framework, the value of choosing an item is unknown due to the stochastic nature of rewards and transition probabilities. Additionally, we allow the value of choosing the outside option a_0 to be non-zero.

vs Cascading RL. In cascading RL (Du et al., 2024), the agent also selects a set of items, and state transitions are taken into account when making decisions. However, these items are offered to the user one by one, and the user decides whether to choose the currently offered item.

Cascading RL fundamentally differs from our framework because the user does not compare multiple items at once, so it does not involve relative preference feedback. Another key distinction is that in cascading RL, the probability of choosing each item is independent of the others in the chosen set of items. In contrast, in our MNL preference model, the choice probability of an item is influenced by the other items in the assortment.

vs PbRL. In preference-based RL (PbRL) (Akrouer et al., 2012; Wirth et al., 2017; Christiano et al., 2017; Ouyang et al.,

2022; Saha et al., 2023; Zhu et al., 2023; Zhan et al., 2023), the agent learns not from explicit numerical rewards, but through preferences as feedback. The user is presented with two (or sometimes multiple) items and chooses a preferred one.

In our framework, if we treat the reward signal generated by the user’s choice as a preference signal instead of a numerical reward, we can learn the policy based on user preferences, similar to PbRL, without relying on explicit rewards. However, our framework differs fundamentally from PbRL because our goal is not just to offer a single item, but multiple items—a combinatorial (base) action—at each timestep.

B. Efficient Bonus Oracle \mathcal{B} using Online-subsampling

The guarantees of Algorithm 1 rely on a consistent bonus oracle, \mathcal{B} , that satisfies Definition B.1.

Definition B.1 (Oracle \mathcal{B} , Agarwal et al. 2023). For any $(h, k) \in [H] \times [K]$, sequence of $\{\bar{\sigma}_h^\tau\}_{\tau=1}^{k-1}$ and dataset $\mathcal{D}_h^{k-1} = \{(s_h^\tau, a_h^\tau, r_h^\tau, s_{h+1}^\tau)\}_{\tau=1}^{k-1}$, function class \mathcal{F}_h with $\hat{f}_h \in \mathcal{F}_h$, $\beta_h, \rho > 0$, error parameter $\epsilon_b \geq 0$, the bonus oracle $\mathcal{B}(\{\bar{\sigma}_h^\tau\}_{\tau=1}^{k-1}, \mathcal{D}_h^{k-1}, \mathcal{F}_h, \hat{f}_h, \beta_h, \rho, \epsilon_b)$ outputs a bonus function $b_h(\cdot)$ such that, for any $z_h = (s_h, a_h) \in \mathcal{S} \times \mathcal{I}$, we have

- $b_h : \mathcal{S} \times \mathcal{I} \rightarrow \mathbb{R}_+$ belongs to a bonus function class \mathcal{W} and denote $\mathcal{N}_b = |\mathcal{W}|$.
- $b_h(z_h) \geq \max \left\{ |f_h(z_h) - \hat{f}_h(z_h)|, f_h \in \mathcal{F}_h : \sum_{\tau=1}^{k-1} \frac{1}{(\bar{\sigma}_h^\tau)^2} \left(f_h(z_h^\tau) - \hat{f}_h^k(z_h^\tau) \right)^2 \leq (\beta_h)^2 \right\}$.
- $b_h(z_h) \leq C \cdot \left(\mathcal{D}_{\mathcal{F}_h}(z_h; \{\bar{\sigma}_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}_h^\tau\}_{\tau=1}^{k-1}) \cdot \sqrt{(\beta_h)^2 + \rho} + \epsilon_b \cdot \beta_h \right)$ with $0 < C < \infty$.

Further we say the oracle \mathcal{B} is consistent if for any $k < k'$ with consistent $\{\bar{\sigma}_h^\tau\}_{\tau=1}^{k-1} \subseteq \{\bar{\sigma}_h^\tau\}_{\tau=1}^{k'-1}$, $\mathcal{D}_h^{k-1} \subseteq \mathcal{D}_h^{k'-1}$, β_h^k non-decreasing in k for each $h \in [H]$ and \hat{f}_h^k as defined in (3), it holds that $\mathcal{B}(\{\bar{\sigma}_h^\tau\}_{\tau=1}^{k-1}, \mathcal{D}_h^{k-1}, \mathcal{F}_h, \hat{f}_h^k, \beta_h^k, \rho, \epsilon_b) \geq \mathcal{B}(\{\bar{\sigma}_h^\tau\}_{\tau=1}^{k'-1}, \mathcal{D}_h^{k'-1}, \mathcal{F}_h, \hat{f}_h^{k'}, \beta_h^{k'}, \rho, \epsilon_b)$ element-wise.

With the oracle \mathcal{B} , we can efficiently calculate the optimistic \bar{Q} -value estimate f_h^k with an error of ϵ_b .

To implement this oracle, we use the online sensitivity sub-sampling approach described by Agarwal et al. (2023), which builds on the original sensitivity sub-sampling method proposed by Kong et al. (2021) and Wang et al. (2020).

For completeness, we include the sub-sampling procedure in Algorithm B.1 and show its guarantees in Proposition B.3. Let $z = (s, a) \in \mathcal{S} \times \mathcal{I}$. We first define the weighted data set \mathcal{Z} , where each element is $(z, \bar{\sigma}(z))$, and introduce the *weighted sensitivity score* as follows:

$$\text{sensitivity}_{\mathcal{Z}, \mathcal{F}, \gamma, \nu}(z) = \min \left\{ \sup_{f, f' \in \mathcal{F}} \frac{\frac{1}{\bar{\sigma}^2(z)} (f(z) - f'(z))^2}{\min \left\{ \sum_{z' \in \mathcal{Z}} \frac{1}{\bar{\sigma}^2(z')} (f(z') - f'(z'))^2, \frac{K(H+1)^2}{\nu^2} \right\}} + \gamma^2, 1 \right\}.$$

Now we introduce the sub-sampling procedure.

Algorithm B.1 Online Sensitivity Sub-sampling with Weights

- 1: **Inputs:** function class \mathcal{F} , current sub-sampled dataset $\hat{\mathcal{Z}} \subseteq \mathcal{S} \times \mathcal{I}$, new state-action pair s, a , parameter γ , threshold $\nu > 0$, failure probability δ .
- 2: **Parameter:** $1 \leq C < \infty$
- 3: Let p_z be the smallest real number such that

$$1/p_z \text{ is an integer and } p_z \geq \min \left\{ 1, C \cdot \text{sensitivity}_{\hat{\mathcal{Z}}, \mathcal{F}, \gamma, \nu}(z) \cdot \log(K\mathcal{N}/\delta) \right\}.$$

- 4: Independently add $1/p_z$ copies of $(z, \bar{\sigma}(z))$ into $\hat{\mathcal{Z}}$ with probability p_z .
 - 5: **Return:** $\hat{\mathcal{Z}}$.
-

For the weighted dataset $\mathcal{Z}_h^{k-1} = \{(s_h^\tau, a_h^\tau), \bar{\sigma}_h^\tau\}_{\tau=1}^{k-1}$, we defined $\|f\|_{\mathcal{Z}_h^{k-1}}^2 = \sum_{z \in \mathcal{Z}_h^{k-1}} \frac{1}{\bar{\sigma}^2(z)} f^2(z)$, i.e., weighted sum of ℓ_2 -norm square. We denote $\hat{\mathcal{Z}}_h^{k-1}$ as the dataset sub-sampled from \mathcal{Z}_h^{k-1} . At every $(k, h) \in [K] \times [H]$, we call Algorithm B.1 with the current sub-sampled dataset $\hat{\mathcal{Z}}_h^{k-1}$ and the new state action-pair $z_h^k = (s_h^k, a_h^k)$ to generate the next sub-sampled dataset $\hat{\mathcal{Z}}_h^k$.

The following proposition shows that the distance of any two functions measured by the historical dataset \mathcal{Z}_h^{k-1} is well approximated by the subsampling dataset $\hat{\mathcal{Z}}_h^{k-1}$. Additionally, it shows that both the number of distinct elements in $|\hat{\mathcal{Z}}_h^{k-1}|$ and its the total size (counting repetitions) do not scale with $\text{poly}(\mathcal{S})$.

Proposition B.2 (Guarantees of online sensitivity sub-sampling, Proposition 13 of Agarwal et al. 2023). *Let $z = (s, a) \in \mathcal{S} \times \mathcal{I}$. When $\bar{\sigma}(z) \geq \nu$ for any z , then with probability at least $1 - \delta$, it holds that*

$$\begin{aligned} \sup_{f_1, f_2: \|f_1 - f_2\|_{\mathcal{Z}_h^{k-1}}^2 \leq \gamma^2} |f_1(z) - f_2(z)| &\leq \sup_{f_1, f_2: \|f_1 - f_2\|_{\hat{\mathcal{Z}}_h^{k-1}}^2 \leq 10^2 \gamma^2} |f_1(z) - f_2(z)| \\ &\leq \sup_{f_1, f_2: \|f_1 - f_2\|_{\mathcal{Z}_h^k}^2 \leq 10^4 \gamma^2} |f_1(z) - f_2(z)|. \end{aligned}$$

Further, for any $(k, h) \in [K] \times [H]$, the number of distinct elements in sub-sampled dataset $\hat{\mathcal{Z}}_h^k$ is always bounded by $\mathcal{O}(\log \frac{KN}{\delta} \cdot \max_{h \in [H]} \dim_{\nu, K}(\mathcal{F}_h))$ and the total size of $\hat{\mathcal{Z}}_h^k$ is bounded by $\mathcal{O}(K^3/\delta)$.

We can assert that the predictive differences between the functions are preserved up to constant factors, while requiring significantly less data. Then, the size of the bonus class \mathcal{W} in Definition B.1 is bounded as follows:

Proposition B.3 (Implementing \mathcal{B} using online-subsampling, Corollary 14 of Agarwal et al. 2023). *There exists an algorithm (see Algorithm B.1) such that, with probability at least $1 - \delta$, implements a consistent bonus oracle \mathcal{B} with $\epsilon_b = 0$ for all $(k, h) \in [K] \times [H]$, where*

$$\log |\mathcal{W}| \leq \mathcal{O} \left(\max_{h \in [H]} \dim_{\nu, K}(\mathcal{F}_h) \cdot \log \frac{KN}{\delta} \log \frac{K|\mathcal{S} \times \mathcal{I}|}{\delta} \right).$$

C. Efficient combinatorial optimization

In this section, we explain how to solve the combinatorial optimization problem in (8), following the method outlined in Davis et al. (2013); Ie et al. (2019).

To find an assortment $A \in \mathcal{A}$ that maximizes the optimistic Q -value, a fundamental step in Q -learning and crucial for inducing exploration, we must solve the following combinatorial optimization problem:

$$\max_{A \in \mathcal{A}} \sum_{a \in A} \tilde{\mathcal{P}}_{h,j}^k(a|s_h^k, A) f_{h,j}^k(s_h^k, a), \quad (\text{C.1})$$

where $\tilde{\mathcal{P}}_{h,j}^k$ is the optimistic choice probability as defined in (7) (also in (D.17)), and $f_{h,j}^k$ is the \bar{Q} -value estimate (item-level Q -values) as defined in (D.15).

Fix $(k, h, s, j) \in [K] \times [H] \times \mathcal{S} \times \{1, 2, -2\}$. For simplicity, we will abbreviate these indices. Accordingly, we denote $w_a = \exp(\tilde{v}_h^k(s_h^k, a))$ or $w_a = \exp(\check{v}_h^k(s_h^k, a))$, depending on the value of j . Additionally, let $\tilde{f}_a = f_{h,j}^k(s_h^k, a)$ for simplicity.

We can then express the optimization problem in (C.1) in terms of w as fractional mixed-integer program (MIP), with binary variables $x_a \in \{0, 1\}$ for each item $a \in \mathcal{I} \setminus \{a_0\}$, indicating whether a is included in the assortment A :

$$\begin{aligned} \max \quad & \frac{w_{a_0} \tilde{f}_{a_0} + \sum_{a \in \mathcal{I} \setminus \{a_0\}} x_a w_a \tilde{f}_a}{w_{a_0} + \sum_{a' \in \mathcal{I} \setminus \{a_0\}} x_{a'} w_{a'}} \\ \text{s.t.} \quad & \sum_{a \in \mathcal{I} \setminus \{a_0\}} x_a \leq M - 1; \\ & x_a \in \{0, 1\}, \quad \forall a \in \mathcal{I} \setminus \{a_0\}. \end{aligned} \quad (\text{C.2})$$

By [Chen & Hausman \(2000\)](#), the binary indicator in the MIP can be relaxed, resulting in the following fractional linear program (LP):

$$\begin{aligned} \max \quad & \frac{w_{a_0} \tilde{f}_{a_0} + \sum_{a \in \mathcal{I} \setminus \{a_0\}} x_a w_a \tilde{f}_a}{w_{a_0} + \sum_{a' \in \mathcal{I} \setminus \{a_0\}} x_{a'} w_{a'}} \\ \text{s.t.} \quad & \sum_{a \in \mathcal{I} \setminus \{a_0\}} x_a \leq M - 1; \\ & 0 \leq x_a \leq 1, \quad \forall a \in \mathcal{I} \setminus \{a_0\}. \end{aligned} \quad (\text{C.3})$$

Since this relaxed problem is a fractional linear program (LP), using the Charnes-Cooper method ([Cooper et al., 1962](#)), it can be transformed into a (non-fractional) LP. To achieve this, we introduce additional variables:

$$t = \frac{1}{w_{a_0} + \sum_{a' \in \mathcal{I} \setminus \{a_0\}} x_{a'} w_{a'}}, \quad y_a = \frac{x_a}{w_{a_0} + \sum_{a' \in \mathcal{I} \setminus \{a_0\}} x_{a'} w_{a'}}.$$

Then, we can obtain the following LP:

$$\begin{aligned} \max \quad & \sum_{a \in \mathcal{I} \setminus \{a_0\}} \tilde{f}_a w_a y_a + \tilde{f}_{a_0} w_{a_0} t \\ \text{s.t.} \quad & \sum_{a \in \mathcal{I} \setminus \{a_0\}} w_a y_a + w_{a_0} t = 1; \\ & \sum_{a \in \mathcal{I} \setminus \{a_0\}} y_a \leq (M - 1)t; \quad t \geq 0. \end{aligned} \quad (\text{C.4})$$

The optimal solution $(y_{a_1}^*, \dots, y_{a_N}^*, t^*)$ to this LP in (C.4) provides the optimal x_i values for the fractional LP in (C.3) by setting $x_a = y_a^*/t^*$. This, in turn, determines the optimal assortment in the original fractional MIP (Equation (C.2)) by including any item where $y_a^* > 0$. Thus, the optimization problem is proven to be solvable in polynomial time.

D. Proof of Theorem 5.1

D.1. Notations and Preliminaries

In this subsection, for easy reference, we introduce notations and definitions used throughout the proof. The key notations are summarized in Table D.1, and the specific parameter choices are listed in Table D.2.

Online parameter update and confidence interval for MNL preference model. We define the multinomial logistic loss function at $(k, h) \in [K] \times [H]$ as follows:

$$\ell_h^k(\theta) := - \sum_{a \in A_h^k} y_h^k(a) \log \mathcal{P}_h(a|s_h^k, A_h^k; \theta). \quad (\text{D.1})$$

To achieve constant-time parameter estimation, we use the online mirror descent algorithm to estimate the true parameter θ_h^* :

$$\theta_h^{k+1} \in \underset{\theta \in \Theta}{\operatorname{argmin}} \langle \nabla \ell_h^k(\theta_h^k), \theta \rangle + \frac{1}{2\eta} \|\theta - \theta_h^k\|_{\tilde{\mathbf{H}}_h^k}^2, \quad \text{where } \Theta = \{\theta \in \mathbb{R}^d : \|\theta\|_2 \leq B\}, \quad (\text{D.2})$$

where $\eta = \frac{1}{2} \log(M + 1) + B + 1$ is the step-size parameter, and the related matrices are defined as:

$$\begin{aligned} \mathbf{H}_h^k &:= \lambda \mathbf{I}_d + \sum_{\tau=1}^{k-1} \nabla^2 \ell_h^\tau(\theta_h^{\tau+1}), \\ \tilde{\mathbf{H}}_h^k &:= \mathbf{H}_h^k + \eta \nabla^2 \ell_h^k(\theta_h^k), \end{aligned} \quad (\text{D.3})$$

where

$$\begin{aligned} \nabla^2 \ell_h^k(\theta) &= \sum_{a \in A_h^k} \mathcal{P}_h(a|s_h^k, A_h^k; \theta) \phi(s_h^k, a) \phi(s_h^k, a)^\top \\ &\quad - \sum_{a \in A_h^k} \sum_{a' \in A_h^k} \mathcal{P}_h(a|s_h^k, A_h^k; \theta) \mathcal{P}_h(a'|s_h^k, A_h^k; \theta) \phi(s_h^k, a) \phi(s_h^k, a')^\top. \end{aligned}$$

Table D.1: Summary of notations

Notation	Meaning	Remark
$\mathcal{S}, \mathcal{A}, \mathcal{I}$	state space, action (assortment) space, item set	
k, h	$k \in [K]$ episode, $h \in [H + 1]$ horizon	
$r_h^k, s_h^k, A_h^k, a_h^k$	reward, state, action and item at k, h	
r_h, s_h, A_h, a_h	random reward, state, action and item h	
z	shorthand for state-item pair (s, a)	
$\bar{Q} : \mathcal{S} \times \mathcal{I} \rightarrow \mathbb{R}$	item-level Q -value function	
$\mathcal{T}_h, \mathcal{T}_h^2$	Bellman operator and second-moment operator	
\mathcal{D}_h^{k-1}	data set $\{(s_h^\tau, a_h^\tau, r_h^\tau, s_{h+1}^\tau)\}_{\tau=1}^{k-1}$	
\mathcal{F}_h	function class for horizon $h \in [H]$	Ass. 3.3
$\mathcal{F}_h^{\text{lin}}$	linear function class for horizon $h \in [H]$	Eqn. E.1
$\mathcal{F}_h^{\text{lin}}(\epsilon_c)$	ϵ_c -cover of linear function class $\mathcal{F}_h^{\text{lin}}$	
\mathcal{W}	bonus function class defined for bonus oracle \mathcal{B}	Def. B.1
ϵ_b	error parameter for bonus oracle	
\mathcal{N}	maximal size of function class, i.e., $\max_{h \in [H]} \mathcal{F}_h $	
\mathcal{N}_b	size of bonus function class $ \mathcal{W} $	
$D_{\mathcal{F}}^2(z; \{z^\tau\}_{\tau=1}^{k-1}, \{\sigma^\tau\}_{\tau=1}^{k-1})$	$:= \sup_{f_1, f_2 \in \mathcal{F}} \frac{(f_1(z) - f_2(z))^2}{\sum_{\tau=1}^{k-1} \frac{1}{(\sigma^\tau)^2} (f_1(z^\tau) - f_2(z^\tau))^2 + \rho}$	ρ param.
$\dim_{\nu, K}(\mathcal{F})$	generalized Eluder dimension defined in Definition 3.5	ν param.
d_ν	$:= \frac{1}{H} \sum_{h=1}^H \dim_{\nu, K}(\mathcal{F}_h)$ (Definition 3.5)	ν param.
$\ell_h^k(\theta)$	$-\sum_{a \in A_h^k} y_h^k(a) \log \mathcal{P}_h(a s_h^k, A_h^k; \theta)$, loss for MNL model at k, h	Eqn. D.1
$\mathbf{H}_h^k, \tilde{\mathbf{H}}_h^k$	$= \lambda \mathbf{I}_d + \sum_{\tau=1}^{k-1} \nabla^2 \ell_h^\tau(\theta_h^{\tau+1})$, $= \mathbf{H}_h^k + \eta \nabla^2 \ell_h^k(\theta_h^k)$, respectively	Eqn. D.3
\mathcal{C}_h^k	confidence interval for MNL model at k, h	Eqn. D.5
$f_{h,1}^k$	optimistic \bar{Q} at k, h	
$\hat{f}_{h,1}^k$	solution of fitting weighted regression at k, h	Eqn. D.7
$\mathcal{F}_{h,1}^k$	version space of optimistic \bar{Q} at k, h	Eqn. D.8
$f_{h,\pm 2}^k$	overly optimistic (pessimistic) \bar{Q} at k, h	
$\hat{f}_{h,\pm 2}^k$	solution of fitting unweighted regression at k, h	Eqn. D.11
$\mathcal{F}_{h,\pm 2}^k$	version space of overly optimistic (pessimistic) \bar{Q} at k, h	Eqn. D.12
\hat{g}_h^k	solution of fitting second-moment regression at k, h	Eqn. D.13
\mathcal{G}_h^k	version space of second-moment estimates at k, h	Eqn. D.14
\mathcal{E}^θ	event that $\{\theta_h^* \in \mathcal{C}_h^k \text{ for all } k \geq 1 \text{ and all } h \in [H]\}$	
\mathcal{E}_h^k	event that $\{\mathcal{T}_h V_{h+1,j}^k \in \mathcal{F}_{h,j}^k \text{ for } j = 1, \pm 2 \text{ and } \mathcal{T}_h^2 V_{h+1,1}^k \in \mathcal{G}_h^k\}$	Ass. 3.3
$\mathcal{E}_{\leq k}$	joint event that $\bigcap_{\tau=1}^k \bigcap_{h=1}^H \mathcal{E}_h^\tau$	
$\tilde{v}_h^k(s_h^k, a), \check{v}_h^k(s_h^k, a)$	optimistic (pessimistic) utility defined in (6)	
$\tilde{\mathcal{P}}_h^k(a s, A)$	optimistic choice probability defined in (7)	
$Q_{h,j}^k(s, A)$	$:= \sum_{a \in A} \tilde{\mathcal{P}}_{h,j}^k(a s, A) f_{h,j}^k(s, a)$ for $j = 1, \pm 2$	
$V_{h,j}^k(s)$	$:= \max_{A \in \mathcal{A}} Q_{h,j}^k(s, A)$ for $j = 1, \pm 2$	
$A_{h,j}^k$	$\in \arg\max_{A \in \mathcal{A}} \sum_{a \in A} \tilde{\mathcal{P}}_{h,j}^k(a s_h^k, A) f_{h,j}^k(s_h^k, a)$ for $j = 1, \pm 2$	
A_h^k	chosen assortment at k, h by assortment selection rule in (9)	
$Q_h^k(s, A), V_h^k(s)$	realized optimistic values determined by (D.18)	Eqn. 9
h_k	random h when first taking action $A_{h,2}^k$ at k , i.e., $A_h^k = A_{h,2}^k$	Eqn. 9
$\mathcal{K}_o, \mathcal{K}_{oo}$	disjoint subsets of $[K]$ when $h_k = H + 1$ or $h_k \in [H]$	Eqn. 9
$b_{h,j}^k$	bonus term obtained in Line 8 and 11 using \mathcal{B}	Def. B.1
$\mathbb{E}_{\mathcal{P}}[\cdot s_h^k, a_h^k], \mathbb{V}_{\mathcal{P}}[\cdot s_h^k, a_h^k]$	$\mathbb{E}_{s_{h+1} \sim \mathbb{P}_h(\cdot s_h^k, a_h^k)}[\cdot s_h^k, a_h^k], \mathbb{V}_{s_{h+1} \sim \mathbb{P}_h(\cdot s_h^k, a_h^k)}[\cdot s_h^k, a_h^k]$	
$\mathbb{E}_{\mathcal{P}}[\cdot s_h^k, A_h^k]$	$\mathbb{E}_{a_h \sim \mathcal{P}_h(\cdot s_h^k, A_h^k)}[\cdot s_h^k, A_h^k]$	

Table D.2: Summary of parameter choices

Notation	Choice	Remark
δ	$\delta \in (0, 1/(H^2 + 15))$	
δ_h^k	$\delta / ((K + 1)(H + 1))$	
η	$\frac{1}{2} \log(M + 1) + B + 1$, step-size parameter for OMD	Eqn. D.2
λ	$84\sqrt{2}d\eta$, regularization parameter	
α_h^k	$\mathcal{O}(\sqrt{d} \log k \log M)$, confidence radius of \mathcal{C}_h^k	Eqn. D.6
ϵ_c	error due to taking covering of function class	
ν	$\sqrt{1/KH}$	Def. 3.5
ρ	1	Def. 3.5
$o(\delta)$	$\sqrt{\log \frac{\mathcal{N}^2(2 \log(4LK/\nu)+2)(\log(8L/\nu^2)+2)}{\delta}}$	
$\iota(\delta)$	$3\sqrt{\log \frac{\mathcal{N}\mathcal{N}_b(2 \log(4LK/\nu)+2)(\log(8L/\nu^2)+2)}{\delta}}$	
$\beta_{h,1}^k$	$\sqrt{(6\sqrt{\rho} + 156) \cdot \log \frac{\mathcal{N}^2(K+1)(H+1)(2 \log \frac{4LK}{\nu}+2)(\log \frac{8L}{\nu^2}+2)}{\delta}}$, confidence radius of $\mathcal{F}_{h,1}^k$	Eqn. D.8
$\iota'(\delta)$	$\sqrt{2 \log \frac{\mathcal{N}\mathcal{N}_b(2 \log(18LK)+2)(\log(18L)+2)}{\delta}}$	
$\beta_{h,2}^k$	$\sqrt{2(24L + 21)(\iota'(\delta_h^k))^2}$, confidence radius of $\mathcal{F}_{h,\pm 2}^k$	Eqn. D.12
$\iota''(\delta)$	$\sqrt{2 \log \frac{\mathcal{N}\mathcal{N}_b(2 \log(32LK)+2)(\log(32L)+2)}{\delta}}$	
$\bar{\beta}_h^k$	$\sqrt{8(11L + 9)(\iota''(\delta_h^k))^2}$, confidence radius of \mathcal{G}_h^k	Eqn. D.14
$(\sigma_h^k)^2$	$\min \left\{ 4, \hat{g}_h^k(z_h^k) - \left(\hat{f}_{h,-2}^k(z_h^k) \right)^2 \right.$ $\left. + D_{\mathcal{F}_h}^2(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\mathbf{1}^\tau\}_{\tau=1}^{k-1}) \cdot \left(\sqrt{(\bar{\beta}_h^k)^2 + \rho} + 2L\sqrt{(\beta_{h,2}^k)^2 + \rho} \right) \right\}$	Eqn. 4
$\bar{\sigma}_h^k$	$\max \left\{ \sigma_h^k, \nu, \sqrt{2}\iota(\delta_h^k)\sqrt{f_{h,2}^k(z_h^k) - f_{h,-2}^k(z_h^k)}, \right.$ $\left. 2 \left(\sqrt{o(\delta_h^k)} + \iota(\delta_h^k) \right) \cdot \sqrt{D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}_h^\tau\}_{\tau=1}^{k-1})} \right\}$	Eqn. D.7
u_k	$C \cdot \left(\sqrt{\log \frac{\mathcal{N}KH}{\nu\delta}} \cdot \left(\log \frac{\mathcal{N}\mathcal{N}_bKH}{\nu\delta} \cdot H^{5/2}\sqrt{d_\nu} + \sqrt{k}H\epsilon_b \right) \right.$ $\left. + dH^{5/2} \log K \log M \sqrt{\log \frac{\mathcal{N}\mathcal{N}_bKH}{\nu\delta}} \right) / \sqrt{k}$ for $C > 0$	Eqn. 9

By a standard online mirror descent formulation (Orabona, 2019), (D.2) can be solved using a single projected gradient step through the following equivalent formula:

$$\bar{\theta}_h^{k+1} = \theta_h^k - \eta \left(\tilde{H}_h^k \right)^{-1} \nabla \ell_h^k(\theta_h^k), \quad \text{and} \quad \theta_h^{k+1} \in \underset{\theta \in \Theta}{\operatorname{argmin}} \|\theta - \bar{\theta}_h^{k+1}\|_{\tilde{H}_h^k}, \quad (\text{D.4})$$

which enjoys a computational cost of only $\mathcal{O}(Md^3)$, completely independent of k (Mhammedi et al., 2019; Zhang & Sugiyama, 2024; Lee & Oh, 2024).

We define the confidence interval at $(k, h) \in [K] \times [H]$ as follows:

$$\mathcal{C}_h^k := \left\{ \theta \in \Theta : \|\theta - \theta_h^k\|_{\mathbf{H}_h^k} \leq \alpha_h^k \right\}, \quad (\text{D.5})$$

where the radius of the confidence interval \mathcal{C}_h^k is as follows:

$$\alpha_h^k = \sqrt{2\eta \left(11 \cdot (3 \log(1 + (M+1)k) + B + 2) \log \left(\frac{2\sqrt{1+2k}}{\delta} \right) + 2 + \frac{7\sqrt{6}}{6} d\eta \log \left(1 + \frac{k+1}{2\lambda} \right) + 2 \right) + 4\lambda B^2} \quad (\text{D.6})$$

Then, we define the optimistic and pessimistic utility as follows:

$$\hat{v}_h^k(s, a) := \phi(s, a)^\top \boldsymbol{\theta}_h^k + \alpha_h^k \|\phi(s, a)\|_{(\mathbf{H}_h^k)^{-1}}, \quad \check{v}_h^k(s, a) := \phi(s, a)^\top \boldsymbol{\theta}_h^k - \alpha_h^k \|\phi(s, a)\|_{(\mathbf{H}_h^k)^{-1}},$$

Regression and confidence intervals for item-level functions. In this paper, we define $\mathcal{N} := \max_{h \in [H]} |\mathcal{F}_h|$ as the maximum size of the function classes $\mathcal{F}_1, \dots, \mathcal{F}_H$ and $\mathcal{N}_b = |\mathcal{W}|$ as the size of the bonus function class \mathcal{W} .

For all $(k, h) \in [K] \times [H]$, the weighted regression problem for fitting the optimistic item-level Q -functions, \bar{Q} , along with the version space of these functions, is defined as:

$$\hat{f}_{h,1}^k \in \operatorname{argmin}_{f_h \in \mathcal{F}_h} \sum_{\tau=1}^{k-1} \frac{1}{(\bar{\sigma}_h^\tau)^2} \left(f_h(s_h^\tau, a_h^\tau) - r_h^\tau - V_{h+1,1}^k(s_{h+1}^\tau) \right)^2, \quad (\text{D.7})$$

$$\mathcal{F}_{h,1}^k := \left\{ f_h \in \mathcal{F}_h : \sum_{\tau=1}^{k-1} \frac{1}{(\bar{\sigma}_h^\tau)^2} \left(f_h(s_h^\tau, a_h^\tau) - \hat{f}_{h,1}^k(s_h^\tau, a_h^\tau) \right)^2 \leq (\beta_{h,1}^k)^2 \right\}. \quad (\text{D.8})$$

Let $z_h^k = (s_h^k, a_h^k)$. The parameters are as follows (for $k \geq 2$):

$$\begin{aligned} (\sigma_h^k)^2 &:= \min \left\{ 4, \hat{g}_h^k(z_h^k) - \left(\hat{f}_{h,-2}^k(z_h^k) \right)^2 \right. \\ &\quad \left. + D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\mathbf{1}^\tau\}_{\tau=1}^{k-1}) \cdot \left(\sqrt{(\bar{\beta}_h^k)^2 + \rho} + \sqrt{(\beta_{h,2}^k)^2 + \rho} \right) \right\}, \end{aligned} \quad (\text{D.9})$$

$$\begin{aligned} \bar{\sigma}_h^k &:= \max \left\{ \sigma_h^k, \nu, \sqrt{2\iota(\delta_h^k)} \sqrt{f_{h,2}^k(z_h^k) - f_{h,-2}^k(z_h^k)}, \right. \\ &\quad \left. 2 \left(\sqrt{o(\delta_h^k)} + \iota(\delta_h^k) \right) \cdot \sqrt{D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}_h^\tau\}_{\tau=1}^{k-1})} \right\}, \end{aligned} \quad (\text{D.10})$$

$$\beta_{h,1}^k := \sqrt{(6\sqrt{\rho} + 156) \cdot \log \frac{\mathcal{N}^2(K+1)(H+1)(2\log \frac{4LK}{\nu} + 2)(\log \frac{8L}{\nu^2} + 2)}{\delta}},$$

$$o(\delta_h^k) := \sqrt{\log \frac{\mathcal{N}^2(2\log(4LK/\nu) + 2)(\log(8L/\nu^2) + 2)}{\delta_h^k}}, \quad \delta_h^k := \frac{\delta}{(K+1)(H+1)},$$

$$\iota(\delta_h^k) := 3\sqrt{\log \frac{\mathcal{N}\mathcal{N}_b(2\log(4LK/\nu) + 2)(\log(8L/\nu^2) + 2)}{\delta_h^k}}.$$

For all $(k, h) \in [K] \times [H]$, the unweighted regression for fitting overly optimistic and overly pessimistic item-level Q -functions, along with the version space of them, is defined as follows:

$$\hat{f}_{h,\pm 2}^k \in \operatorname{argmin}_{f_h \in \mathcal{F}_h} \sum_{\tau=1}^{k-1} \left(f_h(s_h^\tau, a_h^\tau) - r_h^\tau - V_{h+1,\pm 2}^k(s_{h+1}^\tau) \right)^2, \quad (\text{D.11})$$

$$\mathcal{F}_{h,\pm 2}^k := \left\{ f_h \in \mathcal{F}_h : \sum_{\tau=1}^{k-1} \left(f_h(s_h^\tau, a_h^\tau) - \hat{f}_{h,\pm 2}^k(s_h^\tau, a_h^\tau) \right)^2 \leq (\beta_{h,2}^k)^2 \right\}. \quad (\text{D.12})$$

We choose the parameters as follows:

$$\beta_{h,2}^k := \sqrt{2(24L + 21)(\iota'(\delta_h^k))^2},$$

$$\iota'(\delta_h^k) := \sqrt{2 \log \frac{\mathcal{NN}_b(2 \log(18LK) + 2)(\log(18L) + 2)}{\delta_h^k}}, \quad \delta_h^k := \frac{\delta}{(K+1)(H+1)}.$$

For all $(k, h) \in [K] \times [H]$, the unweighted regression for fitting second-moment function values to item-level Q -functions, and their version space, is as follows:

$$\hat{g}_h^k \in \operatorname{argmin}_{g_h \in \mathcal{F}_h} \sum_{\tau=1}^{k-1} \left(g_h(s_h^\tau, a_h^\tau) - (r_h^\tau + V_{h+1,1}^k(s_{h+1}^\tau)) \right)^2, \quad (\text{D.13})$$

$$\mathcal{G}_h^k := \left\{ g_h \in \mathcal{F}_h : \sum_{\tau=1}^{k-1} \left(g_h(s_h^\tau, a_h^\tau) - \hat{g}_h^k(s_h^\tau, a_h^\tau) \right)^2 \leq (\bar{\beta}_h^k)^2 \right\}. \quad (\text{D.14})$$

We choose the parameters as follows:

$$\bar{\beta}_h^k := \sqrt{8(11L + 9)(\iota''(\delta_h^k))^2},$$

$$\iota''(\delta_h^k) := \sqrt{2 \log \frac{\mathcal{NN}_b(2 \log(32LK) + 2)(\log(32L) + 2)}{\delta}}, \quad \delta_h^k := \frac{\delta}{(K+1)(H+1)}.$$

Given the center of the constructed confidence intervals, $\hat{f}_{h,j}^k$ for $j = 1, \pm 2$, we define the optimistic, overly optimistic, and overly pessimistic \bar{Q} -values as follows:

$$\begin{aligned} f_{h,1}^k(\cdot, \cdot) &:= \min \left\{ \hat{f}_{h,1}^k(\cdot, \cdot) + b_{h,1}^k(\cdot, \cdot), 1 \right\}, \\ f_{h,2}^k(\cdot, \cdot) &:= \min \left\{ \hat{f}_{h,2}^k(\cdot, \cdot) + 2b_{h,1}^k(\cdot, \cdot) + b_{h,2}^k(\cdot, \cdot), 1 \right\}, \\ f_{h,-2}^k(\cdot, \cdot) &:= \max \left\{ \hat{f}_{h,-2}^k(\cdot, \cdot) - b_{h,2}^k(\cdot, \cdot), 0 \right\}. \end{aligned} \quad (\text{D.15})$$

Good events. We define the following good events:

$$\begin{aligned} \mathcal{E}^\theta &:= \{ \forall k \geq 1, \forall h \in [H] : \boldsymbol{\theta}_h^* \in \mathcal{C}_h^k \}, \\ \mathcal{E}_{\leq K} &:= \bigcap_{k=1}^K \bigcap_{h=1}^H \mathcal{E}_h^k, \\ \mathcal{E}_h^k &:= \mathcal{E}_{h,1}^k \cap \mathcal{E}_{h,2}^k \cap \mathcal{E}_{h,-2}^k \cap \bar{\mathcal{E}}_h^k, \end{aligned} \quad (\text{D.16})$$

where $\mathcal{E}_{h,j}^k := \{ \mathcal{T}_h V_{h+1,1}^k \in \mathcal{F}_{h,j}^k \}$ for $j = 1, \pm 2$, and $\bar{\mathcal{E}}_h^k := \{ \mathcal{T}_h^2 V_{h+1,1}^k \in \mathcal{G}_h^k \}$.

Optimistic Q -values. For $(k, h, s, A) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A}$ and for $j = 1, \pm 2$, we define the optimistic choice probability as follows:

$$\tilde{\mathcal{P}}_{h,j}^k(a|s, A) := \begin{cases} \frac{\exp(\tilde{v}_h^k(s, a))}{\sum_{a' \in A} \exp(\tilde{v}_h^k(s, a'))}, & \text{if } \exists a \in \mathcal{I} \setminus \{a_0\} \text{ s.t. } f_{h,j}^k(s, a) \geq f_{h,j}^k(s, a_0) \\ \frac{\exp(\tilde{v}_h^k(s, a))}{\sum_{a' \in A} \exp(\tilde{v}_h^k(s, a'))}, & \text{otherwise,} \end{cases} \quad (\text{D.17})$$

where

$$\tilde{v}_h^k(s, a) := \phi(s, a)^\top \boldsymbol{\theta}_h^k + \alpha_h^k \|\phi(s, a)\|_{(\mathbf{H}_h^k)^{-1}}, \quad \check{v}_h^k(s, a) := \phi(s, a)^\top \boldsymbol{\theta}_h^k - \alpha_h^k \|\phi(s, a)\|_{(\mathbf{H}_h^k)^{-1}}.$$

Next, we define the optimistic Q -values for $j = 1, \pm 2$, each constructed using $f_{h,j}^k$ and $\tilde{\mathcal{P}}_{h,j}^k$:

$$Q_{h,j}^k(s, A) = \sum_{a \in A} \tilde{\mathcal{P}}_{h,j}^k(a|s, A) f_{h,j}^k(s, a), \quad V_{h,j}^k(s) = \max_{A \in \mathcal{A}} Q_{h,j}^k(s, A).$$

For convenience, we also define the *realized* optimistic value functions at $(k, h) \in [K] \times [H]$ as follows:

$$Q_h^k(s, A) := \begin{cases} Q_{h,1}^k(s, A) & \text{if } A_h^k = A_{h,1}^k, \\ Q_{h,2}^k(s, A) & \text{otherwise,} \end{cases} \quad V_h^k(s) = \max_{A \in \mathcal{A}} Q_h^k(s, A), \quad (\text{D.18})$$

where A_h^k is the assortment offered to the user by the assortment selection rule in (9) (or equivalently in (D.19)). Therefore, we write $\pi_h^k(s_h^k) = \operatorname{argmax}_{A \in \mathcal{A}} Q_h^k(s_h^k, A)$.

Design of exploration policy. At each episode k , the agent collects data using both $A_{h,1}^k$ and $A_{h,2}^k$, where $A_{h,j}^k \in \operatorname{argmax}_{A \in \mathcal{A}} \sum_{a \in A} \tilde{P}_{h,j}^k(a|s_h^k, A) f_{h,j}^k(s_h^k, a)$ for $j = 1, 2$. Given a sequence of pre-specified $\{u_k\}_{k=1}^K$, at episode k , the agent select an assortment based on the following rule:

$$A_h^k = \begin{cases} A_{h,1}^k & \text{if } f_{h',1}^k(s_{h'}^k, a_{h'}) \geq f_{h',2}^k(s_{h'}^k, a_{h'}) - u_k, \quad \forall a_{h'} \in A_{h',1}^k, \forall h' \leq h. \\ A_{h,2}^k & \text{otherwise,} \end{cases} \quad (\text{D.19})$$

where

$$u_k = \mathcal{O} \left(\frac{\sqrt{\log \frac{NKH}{\nu\delta}} \cdot \left(\log \frac{N\mathcal{N}_bKH}{\nu\delta} \cdot H^{5/2} \sqrt{d_\nu} + \sqrt{k} H \epsilon_b \right) + dH^{5/2} \log K \log M \sqrt{\log \frac{N\mathcal{N}_bKH}{\nu\delta}}}{\sqrt{k}} \right). \quad (\text{D.20})$$

We denote $h_k \in [H+1]$ as the (random) horizon at which the agent first begins offering the assortment $A_{h,2}^k$. More formally, for $h \leq h_k$, the assortment offered is $A_h^k = A_{h,1}^k$, and for $h \geq h_k$, the assortment offered is $A_h^k = A_{h,2}^k$. We then divide the set of episodes $[K]$ into two disjoint subsets: \mathcal{K}_o and \mathcal{K}_{oo} , such that

$$\mathcal{K}_o := \{k \in [K] : h_k = H+1\}, \quad \text{and} \quad \mathcal{K}_{oo} := \{k \in [K] : h_k \leq H\}.$$

Later in the proof, we separately bound the regret for each case.

Other notations. Throughout the proof, we use $z = (s, a)$, $z_h = (s_h, a_h)$ and $z_h^k = (s_h^k, a_h^k)$ interchangeably. We sometimes use $\mathcal{P}_h(\cdot|s, A; \theta_h^*)$ instead of $\mathcal{P}_h(\cdot|s, A)$ to explicitly indicate the dependence on the parameter θ_h^* . For simplicity, we denote $\mathbb{E}_{\mathbb{P}}[\cdot|s_h^k, a_h^k] = \mathbb{E}_{s_{h+1} \sim \mathbb{P}_h(\cdot|s_h^k, a_h^k)}[\cdot|s_h^k, a_h^k]$, $\mathbb{V}_{\mathbb{P}}[\cdot|s_h^k, a_h^k] = \mathbb{V}_{s_{h+1} \sim \mathbb{P}_h(\cdot|s_h^k, a_h^k)}[\cdot|s_h^k, a_h^k]$, and $\mathbb{E}_{\mathcal{P}}[\cdot|s_h^k, A_h^k] = \mathbb{E}_{a_h \sim \mathcal{P}_h(\cdot|s_h^k, A_h^k)}[\cdot|s_h^k, A_h^k]$.

D.2. Confidence Intervals and good events

In this subsection, we show that, given the construction of confidence intervals \mathcal{C}_h^k and $\mathcal{F}_{h,j}^k$ for $j = 1, \pm 2$, the good events \mathcal{E}^θ and $\mathcal{E}_{\leq K}$ occurs with high probability.

Proposition D.1 (Online parameter confidence interval, Lemma 1 of Lee & Oh 2024). *Let $\delta \in (0, 1)$. Under Assumption 3.1, for the confidence interval defined in (2) with*

$$\begin{aligned} \alpha_h^k &= \sqrt{2\eta \left(11 \cdot (3 \log(1 + (M+1)k) + B+2) \log \left(\frac{2\sqrt{1+2k}}{\delta} \right) + 2 + \frac{7\sqrt{6}}{6} d\eta \log \left(1 + \frac{k+1}{2\lambda} \right) + 2 \right) + 4\lambda B^2} \\ &\quad + 16 \left(\log \left(\frac{2\sqrt{1+2k}}{\delta} \right) \right)^2 + \frac{7\sqrt{6}}{6} d\eta \log \left(1 + \frac{k+1}{2\lambda} \right) + 2 + \lambda B^2 \Big]^{1/2} \\ &= \mathcal{O}(\sqrt{d} \log k \log M), \end{aligned}$$

$\eta = \frac{1}{2} \log(M+1) + B+1$ and $\lambda = 84\sqrt{2}d\eta$, and for any $h \in [H]$, we have

$$\Pr[\forall k \geq 1, \theta_h^* \in \mathcal{C}_h^k] \geq 1 - \delta.$$

Now, we define the good event for the preference model \mathcal{E}^θ as follows:

$$\mathcal{E}^\theta := \{\forall k \geq 1, \forall h \in [H] : \theta_h^* \in \mathcal{C}_h^k\}.$$

Then, by applying proposition D.1 and using a union bound over $h \in [H]$, we obtain the following corollary:

Corollary D.2 (Good event for MNL preference model). *Under the same assumption and settings as in Proposition D.1, for $\delta \in (0, 1)$, with probability at least $1 - \delta$, the good event for the preference model \mathcal{E}^θ happens, i.e., $\theta_h^\star \in \mathcal{C}_h^k$ for all $k \geq 1$ and all $h \in [H]$.*

The following proposition shows that $\mathcal{T}_h V_{h+1,j}^k$ for $j = 1, \pm 2$, and $\mathcal{T}_h^2 V_{h+1,1}^k$ lie within their respective confidence intervals.

Proposition D.3 (Good event for general functions, Proposition 33 of Agarwal et al. 2023). *Suppose Algorithm 1 uses a consistent bonus oracle satisfying Definition B.1. Let $\delta \in (0, 1/5)$. Then, with probability at least $1 - 5\delta$, the good event $\mathcal{E}_{\leq K} = \bigcap_{k=1}^K \bigcap_{h=1}^H \mathcal{E}_h^k$ happens, that is, $\mathcal{T}_h V_{h+1,1}^k \in \mathcal{F}_{h,1}^k$, $\mathcal{T}_h V_{h+1,\pm 2}^k \in \mathcal{F}_{h,\pm 2}^k$, and $\mathcal{T}_h^2 V_{h+1,1}^k \in \mathcal{G}_h^k$ for all $(k, h) \in [K] \times [H]$.*

D.3. Bound for MNL Preference Model

In this subsection, we provide proofs for several properties of the MNL preference model.

The following lemma presents both the optimistic and pessimistic utilities.

Lemma D.4. *For any $(k, h, s, a) \in [K] \times [H] \times \mathcal{S} \times \mathcal{I}$, let $\tilde{v}_h^k(s, a) := \phi(s, a)^\top \theta_h^k + \alpha_h^k \|\phi(s, a)\|_{(\mathbf{H}_h^k)^{-1}}$ and $\check{v}_h^k(s, a) = \phi(s, a)^\top \theta_h^k - \alpha_h^k \|\phi(s, a)\|_{(\mathbf{H}_h^k)^{-1}}$. Under the good event \mathcal{E}^θ defined in (D.16), it holds that*

$$0 \leq \tilde{v}_h^k(s, a) - \phi(s, a)^\top \theta_h^\star \leq 2\alpha_h^k \|\phi(s, a)\|_{(\mathbf{H}_h^k)^{-1}},$$

$$\text{and } 0 \leq \phi(s, a)^\top \theta_h^\star - \check{v}_h^k(s, a) \leq 2\alpha_h^k \|\phi(s, a)\|_{(\mathbf{H}_h^k)^{-1}}.$$

Proof of Lemma D.4. Conditioning on the good event \mathcal{E}^θ holds, we have

$$|\phi(s, a)^\top \theta_h^k - \phi(s, a)^\top \theta_h^\star| \leq \|\phi(s, a)\|_{(\mathbf{H}_h^k)^{-1}} \|\theta_h^k - \theta_h^\star\|_{\mathbf{H}_h^k} \leq \alpha_h^k \|\phi(s, a)\|_{(\mathbf{H}_h^k)^{-1}},$$

where the first inequality holds by Hölder's inequality, and the last inequality holds by Corollary D.2. Therefore, it follows that

$$\begin{aligned} \tilde{v}_h^k(s, a) - \phi(s, a)^\top \theta_h^\star &= \phi(s, a)^\top \theta_h^k - \phi(s, a)^\top \theta_h^\star + \alpha_h^k \|\phi(s, a)\|_{(\mathbf{H}_h^k)^{-1}} \\ &\leq 2\alpha_h^k \|\phi(s, a)\|_{(\mathbf{H}_h^k)^{-1}}. \end{aligned}$$

Furthermore, since $\phi(s, a)^\top \theta_h^k - \phi(s, a)^\top \theta_h^\star \geq -\alpha_h^k \|\phi(s, a)\|_{(\mathbf{H}_h^k)^{-1}}$, we also get

$$\check{v}_h^k(s, a) - \phi(s, a)^\top \theta_h^\star = \phi(s, a)^\top \theta_h^k - \phi(s, a)^\top \theta_h^\star + \alpha_h^k \|\phi(s, a)\|_{(\mathbf{H}_h^k)^{-1}} \geq 0.$$

The second statement directly follows from the results mentioned above. \square

Lemma D.5 is useful for proving optimism (Lemma D.15) and bounding the approximation error of the optimistic \bar{Q} (Lemma D.20).

Lemma D.5. *For all $(k, h, s, A) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A}$ and any $j \in \{1, 2\}$, under the good event \mathcal{E}^θ defined in (D.16), there exists a subset $\tilde{A} \subseteq A$ such that $\tilde{A} \in \mathcal{A}$ and*

$$\max \left\{ \sum_{a \in A} \mathcal{P}_h(a|s, A) f_{h,j}^k(s, a), \sum_{a \in A} \tilde{\mathcal{P}}_{h,j'}^k(a|s, A) f_{h,j}^k(s, a) \right\} \leq \sum_{a \in \tilde{A}} \tilde{\mathcal{P}}_{h,j}^k(a|s, \tilde{A}) f_{h,j}^k(s, a),$$

where $j' \neq j$.

Proof of Lemma D.5. First, recall that, without loss of generality, we assumed that $\phi(s, a_0) = 0$ for all $s \in \mathcal{S}$. Therefore,

the true preference model and the optimistic preference model can be written as

$$\begin{aligned} \mathcal{P}_h(a|s, A) &= \frac{\exp(\phi(s, a)^\top \boldsymbol{\theta}_h^*)}{1 + \sum_{a' \in A \setminus \{a_0\}} \exp(\phi(s, a')^\top \boldsymbol{\theta}_h^*)}, \\ \tilde{\mathcal{P}}_{h,j}^k(a|s, A) &= \begin{cases} \frac{\exp(\tilde{v}_h^k(s, a))}{1 + \sum_{a' \in A \setminus \{a_0\}} \exp(\tilde{v}_h^k(s, a'))}, & \text{if } \exists a \in \mathcal{I} \setminus \{a_0\} \text{ s.t. } f_{h,j}^k(s, a) \geq f_{h,j}^k(s, a_0) \\ \frac{\exp(\check{v}_h^k(s, a))}{1 + \sum_{a' \in A \setminus \{a_0\}} \exp(\check{v}_h^k(s, a'))}, & \text{otherwise.} \end{cases} \end{aligned} \quad (\text{D.21})$$

Fix $s \in \mathcal{S}$ and $A \in \mathcal{A}$. We now present the proof by considering two cases: (i) $f_{h,j}^k(s, a_0) > f_{h,j}^k(s, a)$ for all $a \in A$ and (ii) $\exists a \in A \setminus \{a_0\}$ such that $f_{h,j}^k(s, a_0) \leq f_{h,j}^k(s, a)$.

Case (i) $f_{h,j}^k(s, a_0) > f_{h,j}^k(s, a)$ for all $a \in A$.

We denote $\tilde{a} \in \operatorname{argmax}_{a \in A \setminus \{a_0\}} f_{h,j}^k(s, a)$. Let $\tilde{A} = \{a_0, \tilde{a}\}$. Since $f_{h,j}^k(s, a_0) > f_{h,j}^k(s, a)$ for all $a \in A$ and a_0 is always included in A , removing any item $a \in A \setminus \{a_0\}$ from A increases the expected value of $f_{h,j}^k$. Thus, we get

$$\begin{aligned} \sum_{a \in A} \mathcal{P}_h(a|s, A) f_{h,j}^k(s, a) &\leq \sum_{a \in \tilde{A}} \mathcal{P}_h(a|s, \tilde{A}) f_{h,j}^k(s, a) \\ \text{and } \sum_{a \in A} \tilde{\mathcal{P}}_{h,j'}^k(a|s, A) f_{h,j}^k(s, a) &\leq \sum_{a \in \tilde{A}} \tilde{\mathcal{P}}_{h,j'}^k(a|s, \tilde{A}) f_{h,j}^k(s, a). \end{aligned} \quad (\text{D.22})$$

By the definition of $\tilde{\mathcal{P}}_{h,j}^k$ in (7), we use the pessimistic utility $\check{v}_h^k(s, a)$ in this case. Since $\check{v}_h^k(s, a) \leq \phi(s, a)^\top \boldsymbol{\theta}_h^*$ by Lemma D.4, using this utility, $\check{v}_h^k(s, a)$, reduces the probability of selecting \tilde{a} (compared to the true choice probability \mathcal{P}_h). Moreover, we know that $f_{h,j}^k(s, a_0) \geq f_{h,j}^k(s, \tilde{a})$, we have

$$\sum_{a \in \tilde{A}} \mathcal{P}_h(a|s, \tilde{A}) f_{h,j}^k(s, a) \leq \sum_{a \in \tilde{A}} \tilde{\mathcal{P}}_{h,j}^k(a|s, \tilde{A}) f_{h,j}^k(s, a). \quad (\text{D.23})$$

Furthermore, if $\tilde{\mathcal{P}}_{h,j'}^k$ is constructed using the pessimistic utility $\check{v}_h^k(s, a)$, then, $\tilde{\mathcal{P}}_{h,j'}^k = \tilde{\mathcal{P}}_{h,j}^k$. However, if $\tilde{\mathcal{P}}_{h,j'}^k$ is constructed using the optimistic utility $\tilde{v}_h^k(s, a)$, replacing $\check{v}_h^k(s, a)$ with $\tilde{v}_h^k(s, a)$ (which is equivalent to replacing $\tilde{\mathcal{P}}_{h,j'}^k$ with $\tilde{\mathcal{P}}_{h,j}^k$) decreases the probability of choosing \tilde{a} , meaning increase the expected value of $f_{h,j}^k$. Thus, we get

$$\sum_{a \in \tilde{A}} \tilde{\mathcal{P}}_{h,j'}^k(a|s, \tilde{A}) f_{h,j}^k(s, a) \leq \sum_{a \in \tilde{A}} \tilde{\mathcal{P}}_{h,j}^k(a|s, \tilde{A}) f_{h,j}^k(s, a). \quad (\text{D.24})$$

Combining (D.22), (D.23), and (D.24), we have

$$\max \left\{ \sum_{a \in A} \mathcal{P}_h(a|s, A) f_{h,j}^k(s, a), \sum_{a \in A} \tilde{\mathcal{P}}_{h,j'}^k(a|s, A) f_{h,j}^k(s, a) \right\} \leq \sum_{a \in \tilde{A}} \tilde{\mathcal{P}}_{h,j}^k(a|s, \tilde{A}) f_{h,j}^k(s, a).$$

Case (ii) $\exists a \in A \setminus \{a_0\}$ such that $f_{h,j}^k(s, a_0) \leq f_{h,j}^k(s, a)$.

Let $\tilde{A} = \{A' \subseteq A : f_{h,j}^k(s, a) \geq f_{h,j}^k(s, a_0), \forall a \in A'\}$. Note that $|\tilde{A}| \geq 2$ and $a_0 \in \tilde{A}$ by definition of action space \mathcal{A} . By selecting \tilde{A} instead of A , we exclude items with the small values of $f_{h,j}^k(s, a)$, thereby increasing the expected value of $f_{h,j}^k$, i.e.,

$$\begin{aligned} \sum_{a \in A} \mathcal{P}_h(a|s, A) f_{h,j}^k(s, a) &\leq \sum_{a \in \tilde{A}} \mathcal{P}_h(a|s, \tilde{A}) f_{h,j}^k(s, a) \\ \text{and } \sum_{a \in A} \tilde{\mathcal{P}}_{h,j'}^k(a|s, A) f_{h,j}^k(s, a) &\leq \sum_{a \in \tilde{A}} \tilde{\mathcal{P}}_{h,j'}^k(a|s, \tilde{A}) f_{h,j}^k(s, a). \end{aligned} \quad (\text{D.25})$$

By the definition of $\tilde{\mathcal{P}}_{h,j}^k$ in (7), we use the optimistic utility $\tilde{v}_h^k(s, a)$ in this case. Since $\phi(s, a)^\top \boldsymbol{\theta}_h^* \leq \tilde{v}_h^k(s, a)$ by Lemma D.4, this choice of utility increases the probability of choosing item $a \in \tilde{A} \setminus \{a_0\}$ compared to the true \mathcal{P}_h , implying

that

$$\sum_{a \in \tilde{A}} \mathcal{P}_h(a|s, \tilde{A}) f_{h,j}^k(s, a) \leq \sum_{a \in \tilde{A}} \tilde{\mathcal{P}}_{h,j}^k(a|s, \tilde{A}) f_{h,j}^k(s, a). \quad (\text{D.26})$$

Moreover, if $\tilde{\mathcal{P}}_{h,j'}^k$ is constructed using the pessimistic utility $\tilde{v}_h^k(s, a)$, replacing $\tilde{v}_h^k(s, a)$ with $\tilde{v}_h^k(s, a)$ (which is equivalent to replacing $\tilde{\mathcal{P}}_{h,j'}^k$ with $\tilde{\mathcal{P}}_{h,j}^k$) increases the probability of choosing item $a \neq \tilde{A} \setminus \{a_0\}$. However, if $\tilde{\mathcal{P}}_{h,j'}^k$ is constructed using the optimistic utility $\tilde{v}_h^k(s, a)$, we have $\tilde{\mathcal{P}}_{h,j}^k = \tilde{\mathcal{P}}_{h,j'}^k$. To this end, we get

$$\sum_{a \in \tilde{A}} \tilde{\mathcal{P}}_{h,j'}^k(a|s, \tilde{A}) f_{h,j}^k(s, a) \leq \sum_{a \in \tilde{A}} \tilde{\mathcal{P}}_{h,j}^k(a|s, \tilde{A}) f_{h,j}^k(s, a). \quad (\text{D.27})$$

Combining (D.25), (D.26), and (D.27), we have

$$\max \left\{ \sum_{a \in A} \mathcal{P}_h(a|s, A) f_{h,j}^k(s, a), \sum_{a \in A} \tilde{\mathcal{P}}_{h,j'}^k(a|s, A) f_{h,j}^k(s, a) \right\} \leq \sum_{a \in \tilde{A}} \tilde{\mathcal{P}}_{h,j}^k(a|s, \tilde{A}) f_{h,j}^k(s, a).$$

This concludes the proof of Lemma D.5. \square

Denote $J(k, h) \in \{1, 2\}$ as the chosen index of function $f_{h,j}^k$ at (k, h) . Then, we show that the value estimates for the chosen assortment, $f_{h,J(k,h)}^k(s_h^k, a)$ for all $a \in A_h^k$, are greater than equal to $\sum_{a \in A_h^k} \mathcal{P}_{h,J(k,h)}^k(a|s_h^k, A_h^k) f_{h,J(k,h)}^k(s_h^k, a)$.

Lemma D.6. *For any $(k, h) \in [K] \times [H]$, let $J(k, h) : \mathcal{K} \times [H] \rightarrow \{1, 2\}$ be the one-to-one function that maps from $\mathcal{K} \times [H]$ to the index set $\{1, 2\}$ such that $A_h^k = A_{h,J(k,h)}^k \in \operatorname{argmax}_{A \in \mathcal{A}} \sum_{a \in A} \tilde{\mathcal{P}}_{h,J(k,h)}^k(a|s_h^k, A) f_{h,J(k,h)}^k(s_h^k, a)$. Then, under the good event \mathcal{E}^θ defined in (D.16), we have*

$$f_{h,J(k,h)}^k(s_h^k, a) \geq \sum_{a \in A_h^k} \mathcal{P}_h(a|s_h^k, A_h^k) f_{h,J(k,h)}^k(s_h^k, a), \quad \forall a \in A_h^k.$$

Proof of Lemma D.6. By Lemma D.5, there exists a subset $\tilde{A} \subseteq A_h^k$ and $\tilde{A} \in \mathcal{A}$ such that

$$\begin{aligned} \sum_{a \in A_h^k} \mathcal{P}_h(a|s_h^k, A_h^k) f_{h,J(k,h)}^k(s_h^k, a) &\leq \sum_{a \in \tilde{A}} \tilde{\mathcal{P}}_{h,J(k,h)}^k(a|s_h^k, \tilde{A}) f_{h,J(k,h)}^k(s_h^k, a) \\ &\leq \sum_{a \in A_h^k} \tilde{\mathcal{P}}_{h,J(k,h)}^k(a|s_h^k, A_h^k) f_{h,J(k,h)}^k(s_h^k, a), \end{aligned}$$

where the second inequality holds the assortment selection rule. Thus, it is sufficient to show that $f_{h,J(k,h)}^k(s_h^k, a) \geq \sum_{a \in A_h^k} \tilde{\mathcal{P}}_{h,J(k,h)}^k(a|s_h^k, A_h^k) f_{h,J(k,h)}^k(s_h^k, a)$ for all $a \in A_h^k$.

We prove this by contradiction. Suppose there exists an item $a \in A_h^k$ such that

$$f_{h,J(k,h)}^k(s_h^k, a) < \sum_{a \in A_h^k} \tilde{\mathcal{P}}_{h,J(k,h)}^k(a|s_h^k, A_h^k) f_{h,J(k,h)}^k(s_h^k, a).$$

If we remove item a from the assortment A_h^k , it would result in higher value. This contradicts the optimality of A_h^k . Hence, we conclude

$$f_{h,J(k,h)}^k(s_h^k, a) \geq \sum_{a \in A_h^k} \tilde{\mathcal{P}}_{h,J(k,h)}^k(a|s_h^k, A_h^k) f_{h,J(k,h)}^k(s_h^k, a),$$

which completes the proof. \square

Lemma D.7 is an elliptical potential lemma used for bounding the regret incurred from the MNL preference model (Lemma D.10 and D.13).

Lemma D.7 (Elliptical potential lemma, Lemma E.2 and H.3 of Lee & Oh 2024). Assume that $\lambda \geq 2$ and $\phi(s, a_0) = 0$ for all $s \in \mathcal{S}$. For any $(k, h, a) \in [K] \times [H] \times \mathcal{I}$, we define $\tilde{\phi}(s_h^k, a) = \phi(s_h^k, a) - \mathbb{E}_{a' \sim \mathcal{P}_h(\cdot | s_h^k, A_h^k; \theta_h^{k+1})}[\phi(s_h^k, a')]$. Then, for \mathbf{H}_h^k defined in (D.3), and for any $h \in [H]$, the following statements hold true:

$$\begin{aligned} \sum_{\tau=1}^k \sum_{a \in A_h^\tau} \mathcal{P}_h(a | s_h^\tau, A_h^\tau; \theta_h^{\tau+1}) \mathcal{P}_h(a_0 | s_h^\tau, A_h^\tau; \theta_h^{\tau+1}) \|\phi(s_h^\tau, a)\|_{(\mathbf{H}_h^\tau)^{-1}}^2 &\leq 2d \log \left(1 + \frac{k}{d\lambda}\right), \\ \sum_{\tau=1}^k \sum_{a \in A_h^\tau} \mathcal{P}_h(a | s_h^\tau, A_h^\tau; \theta_h^{\tau+1}) \|\tilde{\phi}(s_h^\tau, a)\|_{(\mathbf{H}_h^\tau)^{-1}}^2 &\leq 2d \log \left(1 + \frac{k}{d\lambda}\right), \\ \sum_{\tau=1}^k \max \left\{ \max_{a \in A_h^\tau} \|\phi(s_h^\tau, a)\|_{(\mathbf{H}_h^\tau)^{-1}}^2, \max_{a \in A_h^\tau} \|\tilde{\phi}(s_h^\tau, a)\|_{(\mathbf{H}_h^\tau)^{-1}}^2 \right\} &\leq \frac{2}{\kappa} d \log \left(1 + \frac{k}{d\lambda}\right). \end{aligned}$$

Lemma D.8 is used to derive the tight bound for the second-order regret term of the MNL preference model (Lemma D.13).

Lemma D.8 (Lemma E.3 of Lee & Oh 2024). Let $M \in \mathbf{Z}^+$. Define $R : \mathbb{R}^M \rightarrow \mathbb{R}$, such that for any $\mathbf{v} = (v_1, \dots, v_M) \in \mathbb{R}^M$, $R(\mathbf{v}) = \sum_{m=1}^M \frac{\exp(v_m)}{1 + \sum_{l=1}^M \exp(v_l)}$. Let $p_m(\mathbf{v}) = \frac{\exp(v_m)}{1 + \sum_{l=1}^M \exp(v_l)}$. Then, for all $m \in [M]$, we have

$$\left| \frac{\partial^2 R}{\partial m \partial n} \right| \leq \begin{cases} 3p_m(\mathbf{v}) & \text{if } m = n, \\ 2p_m(\mathbf{v})p_n(\mathbf{v}) & \text{if } m \neq n. \end{cases}$$

Lemma D.9 is crucial for deriving the κ -improved bound for the MNL preference model (Lemma D.13), enabling the analysis.

Lemma D.9 (Overly optimistic choice probability). We define

$$\tilde{\mathcal{P}}_{h,j}^k(a | s, A) = \begin{cases} \frac{\exp(\phi(s, a)^\top \theta_h^* + 2\alpha_h^k \|\phi(s, a)\|_{(\mathbf{H}_h^k)^{-1}})}{\sum_{a' \in A} \exp(\phi(s, a')^\top \theta_h^* + 2\alpha_h^k \|\phi(s, a')\|_{(\mathbf{H}_h^k)^{-1}})}, & \text{if } \exists a \in \mathcal{I} \setminus \{a_0\} \text{ s.t.} \\ & f_{h,j}^k(s, a) \geq f_{h,j}^k(s, a_0), \\ \frac{\exp(\phi(s, a)^\top \theta_h^* - 2\alpha_h^k \|\phi(s, a)\|_{(\mathbf{H}_h^k)^{-1}})}{\sum_{a' \in A} \exp(\phi(s, a')^\top \theta_h^* - 2\alpha_h^k \|\phi(s, a')\|_{(\mathbf{H}_h^k)^{-1}})}, & \text{otherwise.} \end{cases} \quad (\text{D.28})$$

Let $A_{h,j}^k \in \arg\max_{A \in \mathcal{A}} \sum_{a \in A} \tilde{\mathcal{P}}_{h,j}^k(a | s_h^k, A) f_{h,j}^k(s_h^k, a)$, where $j \in \{1, 2\}$. Then, under the good event \mathcal{E}^θ , for all $(k, h, j) \in [K] \times [H] \times \{1, 2\}$, we have

$$\sum_{a \in A_{h,j}^k} \tilde{\mathcal{P}}_{h,j}^k(a | s_h^k, A_{h,j}^k) f_{h,j}^k(s_h^k, a) \leq \sum_{a \in A_{h,j}^k} \tilde{\mathcal{P}}_{h,j}^k(a | s_h^k, A_{h,j}^k) f_{h,j}^k(s_h^k, a).$$

Proof of Lemma D.9. Fix $(k, h, j) \in [K] \times [H] \times \{1, 2\}$. We consider the two cases: (i) $f_{h,j}^k(s_h^k, a_0) > f_{h,j}^k(s_h^k, a)$ for all $a \in \mathcal{I}$ and (ii) $\exists a \in \mathcal{I} \setminus \{a_0\}$ such that $f_{h,j}^k(s_h^k, a_0) \leq f_{h,j}^k(s_h^k, a)$.

Case (i) $f_{h,j}^k(s_h^k, a_0) > f_{h,j}^k(s_h^k, a)$ for all $a \in \mathcal{I}$.

Recall that, by the definition of $\tilde{\mathcal{P}}_{h,j}^k$ in (7), we use the pessimistic utility $\check{v}_h^k(s, a)$ to construct $\tilde{\mathcal{P}}_{h,j}^k$; in this case. Note that the outside option a_0 must be included in the assortment, i.e., $a_0 \in A_h^k$. Moreover, under the event \mathcal{E}^θ , by Lemma D.4, we have

$$\check{v}_h^k(s_h^k, a) \geq \phi(s_h^k, a)^\top \theta_h^* - 2\alpha_h^k \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}}.$$

Thus, since we assume, without loss of generality, that $\phi(s_h^k, a_0) = 0$ (refer (D.21)), using $\tilde{\mathcal{P}}_{h,j}^k$ instead of $\tilde{\mathcal{P}}_{h,j}^k$ decreases the probability of choosing any item $a \in A_h^k \setminus \{a_0\}$. As a result, the expected value of $f_{h,j}^k$ increases, since $f_{h,j}^k(s, a_0) \geq$

$f_{h,j}^k(s, a)$ for all $a \in A_h^k$. Formally, we have

$$\sum_{a \in A_{h,j}^k} \tilde{\mathcal{P}}_{h,j}^k(a|s_h^k, A_{h,j}^k) f_{h,j}^k(s_h^k, a) \leq \sum_{a \in A_{h,j}^k} \tilde{\mathcal{P}}_{h,j}^k(a|s_h^k, A_{h,j}^k) f_{h,j}^k(s_h^k, a).$$

Case (ii) $\exists a \in \mathcal{I} \setminus \{a_0\}$ such that $f_{h,j}^k(s_h^k, a_0) \leq f_{h,j}^k(s_h^k, a)$.

First, we show that for all $a \in A_h^k \setminus \{a_0\}$, we have $f_{h,j}^k(s_h^k, a) \geq \sum_{a \in A_{h,j}^k} \tilde{\mathcal{P}}_{h,j}^k(a|s_h^k, A_{h,j}^k) f_{h,j}^k(s_h^k, a)$. Suppose that there exists $a \in A_h^k \setminus \{a_0\}$ for which $f_{h,j}^k(s_h^k, a) < \sum_{a \in A_{h,j}^k} \tilde{\mathcal{P}}_{h,j}^k(a|s_h^k, A_{h,j}^k) f_{h,j}^k(s_h^k, a)$. Then, removing item a from the assortment A_h^k results in the increase in the expected value of $f_{h,j}^k$. Consequently, this contradicts the optimality of A_h^k . Hence, we get

$$f_{h,j}^k(s_h^k, a) \geq \sum_{a \in A_{h,j}^k} \tilde{\mathcal{P}}_{h,j}^k(a|s_h^k, A_{h,j}^k) f_{h,j}^k(s_h^k, a), \quad \forall a \in A_h^k \setminus \{a_0\}.$$

On the other hand, recall that, by the definition of $\tilde{\mathcal{P}}_{h,j}^k$ in (7), we use the pessimistic utility $\tilde{v}_h^k(s, a)$ to construct $\tilde{\mathcal{P}}_{h,j}^k$ in this case. Furthermore, by Lemma D.4, we know that

$$\tilde{v}_h^k(s_h^k, a) \leq \phi(s_h^k, a)^\top \boldsymbol{\theta}_h^* + 2\alpha_h^k \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}}.$$

If we increase $\tilde{v}_h^k(s_h^k, a)$ to $\phi(s_h^k, a)^\top \boldsymbol{\theta}_h^* + 2\alpha_h^k \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}}$ for all $a \in A_h^k \setminus \{a_0\}$, the probability of choosing the outside option decreases (because $\phi(s_h^k, a_0) = 0$). In other words, the sum of probabilities of choosing $a \in A_h^k \setminus \{a_0\}$ increases. Since $f_{h,j}^k(s_h^k, a) \geq \sum_{a \in A_{h,j}^k} \tilde{\mathcal{P}}_{h,j}^k(a|s_h^k, A_{h,j}^k) f_{h,j}^k(s_h^k, a)$ for all $a \in A_h^k \setminus \{a_0\}$, the expected value of $f_{h,j}^k$ increases. Formally, we get

$$\sum_{a \in A_{h,j}^k} \tilde{\mathcal{P}}_{h,j}^k(a|s_h^k, A_{h,j}^k) f_{h,j}^k(s_h^k, a) \leq \sum_{a \in A_{h,j}^k} \tilde{\mathcal{P}}_{h,j}^k(a|s_h^k, A_{h,j}^k) f_{h,j}^k(s_h^k, a).$$

This concludes the proof. \square

Lemma D.10 will be used to carefully bound the sum of $b_{h,1}^k$ (Lemma D.23). Note that the following MNL bandit regret improves upon the one proposed in (Oh & Iyengar, 2021) by a factor of $1/\sqrt{\kappa}$, which can be exponentially large.

Lemma D.10 (Crude bound for MNL bandits). *For any $h \in [H]$, $j = \{1, 2, -2\}$ and subset $\mathcal{K} \in [K]$, under the good event \mathcal{E}^θ defined in (D.16), we have*

$$\sum_{k \in \mathcal{K}} \left| \sum_{a \in A_h^k} \left(\tilde{\mathcal{P}}_{h,j}^k(a|s_h^k, A_h^k) - \mathcal{P}_h(a|s_h^k, A_h^k) \right) f_{h,j}^k(s_h^k, a) \right| \leq \mathcal{O} \left(\frac{1}{\sqrt{\kappa}} d \sqrt{|\mathcal{K}|} \cdot (\log K)^{3/2} \log M \right),$$

where M is the maximum size of the assortment.

Proof of Lemma D.10. We denote M_h^k as the size of the assortment at horizon h in episode k , i.e., $M_h^k = |A_h^k|$. For any $j \in \{1, 2, -2\}$, we define a function $R_j : \mathbb{R}^{M_h^k} \rightarrow \mathbb{R}$ such that, for all $\mathbf{v} \in \mathbb{R}^{M_h^k}$, $R_j(\mathbf{v}) = \sum_{m=1}^{M_h^k} \frac{\exp(v_m) f_{h,j}^k(s_h^k, a_{im})}{1 + \sum_{l=1}^{M_h^k} \exp(v_l)}$.

For simplicity, we denote $v_{h,j}^k(s, a)$ as the utility, which can represent either the optimistic utility $\tilde{v}_h^k(s, a)$ or the pessimistic utility $\check{v}_h^k(s, a)$, as determined by (7), depending on $f_{h,j}^k$. Let $\mathbf{v}_{h,j}^k(s_h^k) = \left(v_{h,j}^k(s_h^k, a) \right)_{a \in A_h^k} \in \mathbb{R}^{M_h^k}$ and $\mathbf{v}_h^*(s_h^k) = \left(\phi(s_h^k, a)^\top \boldsymbol{\theta}_h^* \right)_{a \in A_h^k} \in \mathbb{R}^{M_h^k}$. Then, by the mean value theorem, there exists a vector $\bar{\mathbf{v}}_{h,j}^k(s_h^k)$, which is a convex combination of $\mathbf{v}_{h,j}^k(s_h^k)$ and $\mathbf{v}_h^*(s_h^k)$, such that

$$\begin{aligned} \sum_{k \in \mathcal{K}} \left| \sum_{a \in A_h^k} \left(\tilde{\mathcal{P}}_{h,j}^k(a|s_h^k, A_h^k) - \mathcal{P}_h(a|s_h^k, A_h^k) \right) f_{h,j}^k(s_h^k, a) \right| &= \sum_{k \in \mathcal{K}} |R_j(\mathbf{v}_{h,j}^k(s_h^k)) - R_j(\mathbf{v}_h^*(s_h^k))| \\ &= \sum_{k \in \mathcal{K}} \left| \nabla R_j(\bar{\mathbf{v}}_{h,j}^k(s_h^k))^\top (\mathbf{v}_{h,j}^k(s_h^k) - \mathbf{v}_h^*(s_h^k)) \right|. \end{aligned}$$

Therefore, we get

$$\begin{aligned}
 & \sum_{k \in \mathcal{K}} \left| \nabla R_j \left(\bar{v}_{h,j}^k(s_h^k) \right)^\top \left(v_{h,j}^k(s_h^k) - v_h^*(s_h^k) \right) \right| \\
 &= \sum_{k \in \mathcal{K}} \left| \sum_{a \in A_h^k} \frac{\exp \left(\bar{v}_{h,j}^k(s_h^k, a) \right) f_{h,j}^k(s_h^k, a)}{\sum_{a'' \in A_h^k} \exp \left(\bar{v}_{h,j}^k(s_h^k, a'') \right)} \left(v_{h,j}^k(s_h^k, a) - \phi(s_h^k, a)^\top \theta_h^* \right) \right. \\
 &\quad \left. - \sum_{a \in A_h^k} \sum_{a' \in A_h^k} \frac{\exp \left(\bar{v}_{h,j}^k(s_h^k, a) \right) f_{h,j}^k(s_h^k, a) \exp \left(\bar{v}_{h,j}^k(s_h^k, a') \right)}{\left(\sum_{a'' \in A_h^k} \exp \left(\bar{v}_{h,j}^k(s_h^k, a'') \right) \right)^2} \left(v_{h,j}^k(s_h^k, a') - \phi(s_h^k, a')^\top \theta_h^* \right) \right| \\
 &= \sum_{k \in \mathcal{K}} \left| \sum_{a \in A_h^k} \mathcal{P}_h(a | s_h^k, A_h^k; \bar{v}_{h,j}^k(s_h^k)) \left(v_{h,j}^k(s_h^k, a) - \phi(s_h^k, a)^\top \theta_h^* \right) \right. \\
 &\quad \left. \cdot \left(f_{h,j}^k(s_h^k, a) - \mathbb{E}_{a' \sim \mathcal{P}_h(\cdot | s_h^k, A_h^k; \bar{v}_{h,j}^k(s_h^k))} [f_{h,j}^k(s_h^k, a')] \right) \right| \\
 &\leq 2 \sum_{k \in \mathcal{K}} \sum_{a \in A_h^k} \mathcal{P}_h(a | s_h^k, A_h^k; \bar{v}_{h,j}^k(s_h^k)) \left| \left(v_{h,j}^k(s_h^k, a) - \phi(s_h^k, a)^\top \theta_h^* \right) \right|, \tag{D.29}
 \end{aligned}$$

where the inequality is from $f_{h,j}^k \leq 1$. Recall that $v_{h,j}^k(s_h^k, a)$ can be either $\tilde{v}_{h,j}^k(s_h^k, a)$ or $\check{v}_{h,j}^k(s_h^k, a)$. Then, by Lemma D.4, we have $\left| \left(v_{h,j}^k(s_h^k, a) - \phi(s_h^k, a)^\top \theta_h^* \right) \right| \leq 2\alpha_h^k \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}}$. Hence, we can further bound the right-hand side of (D.29).

$$\begin{aligned}
 & 2 \sum_{k \in \mathcal{K}} \sum_{a \in A_h^k} \mathcal{P}_h(a | s_h^k, A_h^k; \bar{v}_{h,j}^k(s_h^k)) \left| \left(v_{h,j}^k(s_h^k, a) - \phi(s_h^k, a)^\top \theta_h^* \right) \right| \\
 &\leq 4\alpha_h^K \sum_{k \in \mathcal{K}} \sum_{a \in A_h^k} \mathcal{P}_h(a | s_h^k, A_h^k; \bar{v}_{h,j}^k(s_h^k)) \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}} \\
 &\leq 4\alpha_h^K \sqrt{\sum_{k \in \mathcal{K}} \sum_{a \in A_h^k} \mathcal{P}_h(a | s_h^k, A_h^k; \bar{v}_{h,j}^k(s_h^k))} \sqrt{\sum_{k \in \mathcal{K}} \sum_{a \in A_h^k} \mathcal{P}_h(a | s_h^k, A_h^k; \bar{v}_{h,j}^k(s_h^k)) \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}}^2} \\
 &\leq 4\alpha_h^K \sqrt{|\mathcal{K}|} \cdot \sqrt{\sum_{k=1}^K \sum_{a \in A_h^k} \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}}^2} \\
 &= 4\alpha_h^K \sqrt{|\mathcal{K}|} \cdot \sqrt{\sum_{k=1}^K \sum_{a \in A_h^k} \frac{\mathcal{P}_h(a | s_h^k, A_h^k; \theta_h^{k+1}) \mathcal{P}_h(a_0 | s_h^k, A_h^k; \theta_h^{k+1})}{\mathcal{P}_h(a | s_h^k, A_h^k; \theta_h^{k+1}) \mathcal{P}_h(a_0 | s_h^k, A_h^k; \theta_h^{k+1})} \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}}^2} \\
 &\leq 4\alpha_h^K \sqrt{|\mathcal{K}|} \cdot \sqrt{\frac{1}{\kappa} \cdot \sum_{k=1}^K \sum_{a \in A_h^k} \mathcal{P}_h(a | s_h^k, A_h^k; \theta_h^{k+1}) \mathcal{P}_h(a_0 | s_h^k, A_h^k; \theta_h^{k+1}) \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}}^2} \\
 &\leq 4\alpha_h^K \sqrt{|\mathcal{K}|} \cdot \sqrt{\frac{1}{\kappa} \cdot 2d \log \left(1 + \frac{K}{d\lambda} \right)}, \tag{D.30}
 \end{aligned}$$

where the first inequality holds because $\alpha_h^1 \leq \dots \leq \alpha_h^K$, the second inequality follows from the Cauchy-Schwarz inequality, the second-to-the last inequality holds due to the definition of κ , and the last inequality holds by Lemma D.7.

Combining (D.29) and (D.30), and plugging in the value of α_h^k , we derive that

$$\sum_{k \in \mathcal{K}} \left| \sum_{a \in A_h^k} \left(\tilde{\mathcal{P}}_{h,j}^k(a | s_h^k, A_h^k) - \mathcal{P}_h(a | s_h^k, A_h^k) \right) f_{h,j}^k(s_h^k, a) \right| = \mathcal{O} \left(\frac{1}{\sqrt{\kappa}} d \sqrt{|\mathcal{K}|} \cdot (\log K)^{3/2} \log M \right).$$

□

Lemma D.11. For any $(k, h) \in [K] \times [H]$, $\theta_1, \theta_2 \in \mathcal{C}_h^k$, and $\omega_h^k(a) \geq 0$, under the event \mathcal{E}^θ defined in (D.16), we have

$$\sum_{a \in A_h^k} |\mathcal{P}_h(a|s_h^k, A_h^k; \theta_1) - \mathcal{P}_h(a|s_h^k, A_h^k; \theta_2)| \omega_h^k(a) \leq 4\alpha_h^k \max_{a \in A_h^k} \omega_h^k(a) \max_{a \in A_h^k} \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}}.$$

Proof of Lemma D.11. By the mean value theorem, there exists $\xi = (1 - c)\theta_1 + c\theta_2$ for some $c \in (0, 1)$ such that

$$\begin{aligned} \sum_{a \in A_h^k} |\mathcal{P}_h(a|s_h^k, A_h^k; \theta_1) - \mathcal{P}_h(a|s_h^k, A_h^k; \theta_2)| \omega_h^k(a) &= \sum_{a \in A_h^k} |\nabla \mathcal{P}_h(a|s_h^k, A_h^k; \xi)^\top (\theta_1 - \theta_2)| \omega_h^k(a) \\ &= \sum_{a \in A_h^k} \left| \left(\mathcal{P}_h(a|s_h^k, A_h^k; \xi) \phi(s_h^k, a) - \mathcal{P}_h(a|s_h^k, A_h^k; \xi) \sum_{a' \in A_h^k} \mathcal{P}_h(a'|s_h^k, A_h^k; \xi) \phi(s_h^k, a') \right)^\top (\theta_1 - \theta_2) \right| \omega_h^k(a) \\ &\leq \sum_{a \in A_h^k} \mathcal{P}_h(a|s_h^k, A_h^k; \xi) |\phi(s_h^k, a)^\top (\theta_1 - \theta_2)| \omega_h^k(a) + \sum_{a \in A_h^k} \mathcal{P}_h(a|s_h^k, A_h^k; \xi) \omega_h^k(a) \sum_{a' \in A_h^k} \mathcal{P}_h(a'|s_h^k, A_h^k; \xi) |\phi(s_h^k, a')^\top (\theta_1 - \theta_2)| \\ &\leq 2\alpha_h^k \sum_{a \in A_h^k} \mathcal{P}_h(a|s_h^k, A_h^k; \xi) \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}} \omega_h^k(a) + \max_{a \in A_h^k} \omega_h^k(a) \sum_{a \in A_h^k} \mathcal{P}_h(a|s_h^k, A_h^k; \xi) \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}} \\ &\leq 4\alpha_h^k \max_{a \in A_h^k} \omega_h^k(a) \max_{a \in A_h^k} \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}}, \end{aligned}$$

where the second-to-last inequality holds under the good event \mathcal{E}^θ defined in (D.16). \square

Lemma D.12 pertains to the *law of total variance* (Lattimore & Hutter, 2012; Gheshlaghi Azar et al., 2013) that the variance of the value function is smaller than its magnitude by a factor \sqrt{H} .

Lemma D.12 (Total variance lemma, Lemma C.5 of Jin et al. 2018). Let $f_{h,j}^k \in [0, 1]$. Then, with probability at least $1 - \delta$, we have

$$\sum_{k=1}^K \sum_{h=1}^H [\mathbb{V}_h f_{h,j}^k](s_h^k) = \mathcal{O}(K + H \log(1/\delta)).$$

Lemma D.13 is crucial for obtaining a κ -independent regret in our leading term. While the proof is largely inspired by Lee & Oh (2024), extending their result to our setting is non-trivial because the unknown item values $f_{h,j}^k$ add complexity to the analysis. Moreover, thanks to Lemma D.12, we can obtain a tight bound by a factor of \sqrt{H} , instead of naively summing over H MNL bandit regrets.

Lemma D.13 (κ -improved bound for MNL bandits). For any subset $\mathcal{K} \in [K]$, let $J(k, h) : \mathcal{K} \times [H] \rightarrow \{1, 2\}$ be the one-to-one function that maps from $\mathcal{K} \times [H]$ to the index set $\{1, 2\}$ such that $A_h^k = A_{h, J(k, h)}^k \in \arg\max_{A \in \mathcal{A}} \sum_{a \in A} \tilde{\mathcal{P}}_{h, J(k, h)}^k(a|s_h^k, A) f_{h, J(k, h)}^k(s_h^k, a)$. Then, under the good event \mathcal{E}^θ defined in (D.16), with probability at least $1 - \delta$, we have

$$\begin{aligned} \sum_{k \in \mathcal{K}} \sum_{h=1}^H \sum_{a \in A_h^k} \left(\tilde{\mathcal{P}}_{h, J(k, h)}^k(a|s_h^k, A_h^k) - \mathcal{P}_h(a|s_h^k, A_h^k) \right) f_{h, J(k, h)}^k(s_h^k, a) \\ = \mathcal{O} \left(d\sqrt{H|\mathcal{K}|} (\log K)^{3/2} \log M + \frac{1}{\kappa} d^2 H (\log K)^3 (\log M)^2 \right). \end{aligned}$$

Proof of Lemma D.13. We begin by defining $\tilde{\mathcal{P}}_{h,j}^k(a|s, A)$ as given in D.28. Then, by Lemma D.9, we have

$$\begin{aligned} \sum_{k \in \mathcal{K}} \sum_{h=1}^H \sum_{a \in A_h^k} \left(\tilde{\mathcal{P}}_{h, J(k, h)}^k(a|s_h^k, A_h^k) - \mathcal{P}_h(a|s_h^k, A_h^k) \right) f_{h, J(k, h)}^k(s_h^k, a) \\ \leq \sum_{k \in \mathcal{K}} \sum_{h=1}^H \sum_{a \in A_h^k} \left(\tilde{\mathcal{P}}_{h, J(k, h)}^k(a|s_h^k, A_h^k) - \mathcal{P}_h(a|s_h^k, A_h^k) \right) f_{h, J(k, h)}^k(s_h^k, a). \end{aligned}$$

We denote M_h^k as the size of the assortment at horizon h in episode k , i.e., $M_h^k = |A_h^k|$. We define a function $\tilde{R} : \mathbb{R}^{M_h^k} \rightarrow \mathbb{R}$ such that for all $\mathbf{v} \in \mathbb{R}^{M_h^k}$, $\tilde{R}(\mathbf{v}) = \sum_{m=1}^{M_h^k} \frac{\exp(v_m) f_{h,J(k,h)}^k(s_h^k, a_{im})}{1 + \sum_{l=1}^{M_h^k} \exp(v_l)}$.

For any $(k, h) \in \mathcal{K} \times [H]$ and all $a \in \mathcal{I}$, we denote $v_h^k(s_h^k, a)$ as the utility, which can be either $\phi(s_h^k, a)^\top \boldsymbol{\theta}_h^* + 2\alpha_h^k \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}}$ or $\phi(s_h^k, a)^\top \boldsymbol{\theta}_h^* - 2\alpha_h^k \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}}$, determined deterministically based on the history up to (k, h) :

$$v_h^k(s_h^k, a) = \begin{cases} \phi(s_h^k, a)^\top \boldsymbol{\theta}_h^* + 2\alpha_h^k \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}}, & \text{if } \exists a \in \mathcal{I} \setminus \{a_0\} \text{ s.t. } f_{h,J(k,h)}^k(s_h^k, a) \geq f_{h,J(k,h)}^k(s_h^k, a_0) \\ \phi(s_h^k, a)^\top \boldsymbol{\theta}_h^* - 2\alpha_h^k \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}}, & \text{if } \forall a \in \mathcal{I} \setminus \{a_0\} \text{ } f_{h,J(k,h)}^k(s_h^k, a) < f_{h,J(k,h)}^k(s_h^k, a_0). \end{cases}$$

Let $\mathbf{v}_h^k(s_h^k) = (v_h^k(s_h^k, a))_{a \in A_h^k} \in \mathbb{R}^{M_h^k}$ and $\mathbf{v}_h^*(s_h^k) = (\phi(s_h^k, a)^\top \boldsymbol{\theta}_h^*)_{a \in A_h^k} \in \mathbb{R}^{M_h^k}$. Thanks to exact second-order Taylor expansion, we obtain that

$$\begin{aligned} & \sum_{k \in \mathcal{K}} \sum_{h=1}^H \sum_{a \in A_h^k} \left(\tilde{\mathcal{P}}_{h,J(k,h)}^k(a|s_h^k, A_h^k) - \mathcal{P}_h(a|s_h^k, A_h^k) \right) f_{h,J(k,h)}^k(s_h^k, a) \\ &= \sum_{k \in \mathcal{K}} \sum_{h=1}^H \tilde{R}(\mathbf{v}_h^k(s_h^k)) - \tilde{R}(\mathbf{v}_h^*(s_h^k)) \\ &= \underbrace{\sum_{k \in \mathcal{K}} \sum_{h=1}^H \nabla \tilde{R}(\mathbf{v}_h^*(s_h^k))^\top (\mathbf{v}_h^k(s_h^k) - \mathbf{v}_h^*(s_h^k))}_{(\text{A})} + \underbrace{\frac{1}{2} \sum_{k \in \mathcal{K}} \sum_{h=1}^H (\mathbf{v}_h^k(s_h^k) - \mathbf{v}_h^*(s_h^k))^\top \nabla^2 \tilde{R}(\bar{\mathbf{v}}_h^k(s_h^k)) (\mathbf{v}_h^k(s_h^k) - \mathbf{v}_h^*(s_h^k))}_{(\text{B})}, \end{aligned} \quad (\text{D.31})$$

where $\bar{\mathbf{v}}_h^k(s_h^k) = (\bar{v}_h^k(s_h^k, a))_{a \in A_h^k} \in \mathbb{R}^{M_h^k}$ is the convex combination of $\mathbf{v}_h^k(s_h^k)$ and $\mathbf{v}_h^*(s_h^k)$.

We first bound the term (A) in (D.31).

$$\begin{aligned} & \sum_{k \in \mathcal{K}} \sum_{h=1}^H \nabla \tilde{R}(\mathbf{v}_h^*(s_h^k))^\top (\mathbf{v}_h^k(s_h^k) - \mathbf{v}_h^*(s_h^k)) \\ &= \sum_{k \in \mathcal{K}} \sum_{h=1}^H \left(\sum_{a \in A_h^k} \frac{\exp(\phi(s_h^k, a)^\top \boldsymbol{\theta}_h^*) f_{h,J(k,h)}^k(s_h^k, a)}{\sum_{a'' \in A_h^k} \exp(\phi(s_h^k, a'')^\top \boldsymbol{\theta}_h^*)} (v_h^k(s_h^k, a) - \phi(s_h^k, a)^\top \boldsymbol{\theta}_h^*) \right. \\ & \quad \left. - \sum_{a \in A_h^k} \sum_{a' \in A_h^k} \frac{\exp(\phi(s_h^k, a)^\top \boldsymbol{\theta}_h^*) f_{h,J(k,h)}^k(s_h^k, a) \exp(\phi(s_h^k, a')^\top \boldsymbol{\theta}_h^*)}{\left(\sum_{a'' \in A_h^k} \exp(\phi(s_h^k, a'')^\top \boldsymbol{\theta}_h^*) \right)^2} (v_h^k(s_h^k, a') - \phi(s_h^k, a')^\top \boldsymbol{\theta}_h^*) \right) \\ &= \sum_{k \in \mathcal{K}} \sum_{h=1}^H \sum_{a \in A_h^k} \mathcal{P}_h(a|s_h^k, A_h^k; \boldsymbol{\theta}_h^*) f_{h,J(k,h)}^k(s_h^k, a) \\ & \quad \cdot \left((v_h^k(s_h^k, a) - \phi(s_h^k, a)^\top \boldsymbol{\theta}_h^*) - \sum_{a' \in A_h^k} \mathcal{P}_h(a'|s_h^k, A_h^k; \boldsymbol{\theta}_h^*) (v_h^k(s_h^k, a') - \phi(s_h^k, a')^\top \boldsymbol{\theta}_h^*) \right). \end{aligned} \quad (\text{D.32})$$

We bound the right-hand side of (D.32) by examining two separate cases. For any fixed $h \in [H]$, let $\mathcal{K}_{(i)}$ denote the set of episodes where **Case (i)** holds, and $\mathcal{K}_{(ii)}$ denote the set of episodes where **Case (ii)** holds. More formally, we define:

$$\mathcal{K}_{(i)} = \left\{ k \in \mathcal{K} : v_h^k(s_h^k, a) = \phi(s_h^k, a)^\top \boldsymbol{\theta}_h^* + 2\alpha_h^k \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}} \right\} \quad (\text{Case (i)})$$

$$\mathcal{K}_{(ii)} = \left\{ k \in \mathcal{K} : v_h^k(s_h^k, a) = \phi(s_h^k, a)^\top \boldsymbol{\theta}_h^* - 2\alpha_h^k \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}} \right\}. \quad (\text{Case (ii)})$$

Case (i) For $(k, h) \in \mathcal{K} \times [H]$ such that $v_h^k(s_h^k, a) = \phi(s_h^k, a)^\top \boldsymbol{\theta}_h^* + 2\alpha_h^k \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}}$.

Denoting $\mathbb{E}_{\theta}[\cdot] = \mathbb{E}_{a' \sim \mathcal{P}_h(\cdot|s_h^k, A_h^k; \theta)}[\cdot]$ and $\bar{\alpha}_K := \max_{h \in [H]} \alpha_h^K$ for simplicity, we get

$$\begin{aligned}
 & \sum_{k \in \mathcal{K}_{(i)}} \sum_{h=1}^H \nabla \tilde{R}(\mathbf{v}_h^*(s_h^k))^\top (\mathbf{v}_h(s_h^k) - \mathbf{v}_h^*(s_h^k)) \\
 &= \sum_{k \in \mathcal{K}_{(i)}} \sum_{h=1}^H 2\alpha_h^k \sum_{a \in A_h^k} \mathcal{P}_h(a|s_h^k, A_h^k; \theta_h^*) f_{h,J(k,h)}^k(s_h^k, a) \left(\|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}} - \mathbb{E}_{\theta_h^*} [\|\phi(s_h^k, a')\|_{(\mathbf{H}_h^k)^{-1}}] \right) \\
 &\leq 2\bar{\alpha}_K \sum_{k \in \mathcal{K}_{(i)}} \sum_{h=1}^H \sum_{a \in A_h^k} \mathcal{P}_h(a|s_h^k, A_h^k; \theta_h^*) f_{h,J(k,h)}^k(s_h^k, a) \left(\|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}} - \mathbb{E}_{\theta_h^*} [\|\phi(s_h^k, a')\|_{(\mathbf{H}_h^k)^{-1}}] \right) \\
 &= 2\bar{\alpha}_K \sum_{k \in \mathcal{K}_{(i)}} \sum_{h=1}^H \mathbb{E}_{\theta_h^*} \left[\left(f_{h,J(k,h)}^k(s_h^k, a) - \mathbb{E}_{\theta_h^*} [f_{h,J(k,h)}^k(s_h^k, a')] \right) \left(\|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}} - \mathbb{E}_{\theta_h^*} [\|\phi(s_h^k, a')\|_{(\mathbf{H}_h^k)^{-1}}] \right) \right] \\
 &\leq 2\bar{\alpha}_K \sum_{k \in \mathcal{K}_{(i)}} \sum_{h=1}^H \underbrace{\mathbb{E}_{\theta_h^*} \left[\left(f_{h,J(k,h)}^k(s_h^k, a) - \mathbb{E}_{\theta_h^*} [f_{h,J(k,h)}^k(s_h^k, a')] \right) \right]}_{\geq 0} \left(\|\phi(s_h^k, a) - \mathbb{E}_{\theta_h^*} [\phi(s_h^k, a')]\|_{(\mathbf{H}_h^k)^{-1}} \right), \quad (\text{D.33})
 \end{aligned}$$

where, in the first inequality, we use the fact that α_h^K is non-decreasing with respect to k , and by Lemma D.6, we have

$$\begin{aligned}
 & \sum_{a \in A_h^k} \mathcal{P}_h(a|s_h^k, A_h^k; \theta_h^*) f_{h,J(k,h)}^k(s_h^k, a) \left(\|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}} - \mathbb{E}_{\theta_h^*} [\|\phi(s_h^k, a')\|_{(\mathbf{H}_h^k)^{-1}}] \right) \\
 &= \sum_{a \in A_h^k} \mathcal{P}_h(a|s_h^k, A_h^k; \theta_h^*) \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}} \left(f_{h,J(k,h)}^k(s_h^k, a) - \underbrace{\mathbb{E}_{\theta_h^*} [f_{h,J(k,h)}^k(s_h^k, a')]}_{=\sum_{a' \in A_h^k} \mathcal{P}_h(a'|s_h^k, A_h^k) f_{h,J(k,h)}^k(s_h^k, a')} \right) \\
 &\geq 0.
 \end{aligned}$$

And the last inequality of Equation (D.33) holds because

$$\begin{aligned}
 \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}} - \mathbb{E}_{\theta_h^*} [\|\phi(s_h^k, a')\|_{(\mathbf{H}_h^k)^{-1}}] &\leq \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}} - \|\mathbb{E}_{\theta_h^*} [\phi(s_h^k, a')]\|_{(\mathbf{H}_h^k)^{-1}} \\
 &\leq \|\phi(s_h^k, a) - \mathbb{E}_{\theta_h^*} [\phi(s_h^k, a')]\|_{(\mathbf{H}_h^k)^{-1}},
 \end{aligned}$$

where the first inequality holds by Jensen's inequality and the last inequality holds due to the fact that $\|\mathbf{a}\| = \|\mathbf{a} - \mathbf{b} + \mathbf{b}\| \leq \|\mathbf{a} - \mathbf{b}\| + \|\mathbf{b}\|$ for any vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^d$.

We further decompose the right-hand side of (D.33) as follows:

$$\begin{aligned}
 & \sum_{k \in \mathcal{K}_{(i)}} \sum_{h=1}^H \mathbb{E}_{\theta_h^*} \left[\left(f_{h,J(k,h)}^k(s_h^k, a) - \mathbb{E}_{\theta_h^*} [f_{h,J(k,h)}^k(s_h^k, a')] \right) \|\phi(s_h^k, a) - \mathbb{E}_{\theta_h^*} [\phi(s_h^k, a')]\|_{(\mathbf{H}_h^k)^{-1}} \right] \\
 &= \sum_{k \in \mathcal{K}_{(i)}} \sum_{h=1}^H \sum_{a \in A_h^k} \sqrt{\mathcal{P}_h(a|s_h^k, A_h^k; \theta_h^*) \mathcal{P}_h(a|s_h^k, A_h^k; \theta_h^{k+1})} \left(f_{h,J(k,h)}^k(s_h^k, a) - \mathbb{E}_{\theta_h^*} [f_{h,J(k,h)}^k(s_h^k, a')] \right) \\
 &\quad \cdot \|\phi(s_h^k, a) - \mathbb{E}_{\theta_h^{k+1}} [\phi(s_h^k, a')]\|_{(\mathbf{H}_h^k)^{-1}} \\
 &+ \sum_{k \in \mathcal{K}_{(i)}} \sum_{h=1}^H \sum_{a \in A_h^k} \left(\sqrt{\mathcal{P}_h(a|s_h^k, A_h^k; \theta_h^*)} - \sqrt{\mathcal{P}_h(a|s_h^k, A_h^k; \theta_h^{k+1})} \right) \sqrt{\mathcal{P}_h(a|s_h^k, A_h^k; \theta_h^*)} \\
 &\quad \cdot \left(f_{h,J(k,h)}^k(s_h^k, a) - \mathbb{E}_{\theta_h^*} [f_{h,J(k,h)}^k(s_h^k, a')] \right) \|\phi(s_h^k, a) - \mathbb{E}_{\theta_h^{k+1}} [\phi(s_h^k, a')]\|_{(\mathbf{H}_h^k)^{-1}} \\
 &+ \sum_{k \in \mathcal{K}_{(i)}} \sum_{h=1}^H \sum_{a \in A_h^k} \mathcal{P}_h(a|s_h^k, A_h^k; \theta_h^*) \left(f_{h,J(k,h)}^k(s_h^k, a) - \mathbb{E}_{\theta_h^*} [f_{h,J(k,h)}^k(s_h^k, a')] \right) \\
 &\quad \cdot \left(\|\phi(s_h^k, a) - \mathbb{E}_{\theta_h^*} [\phi(s_h^k, a')]\|_{(\mathbf{H}_h^k)^{-1}} - \|\phi(s_h^k, a) - \mathbb{E}_{\theta_h^{k+1}} [\phi(s_h^k, a')]\|_{(\mathbf{H}_h^k)^{-1}} \right). \quad (\text{D.34})
 \end{aligned}$$

For simplicity, let $\bar{\phi}(s_h^k, a) = \phi(s_h^k, a) - \mathbb{E}_{\theta_h^*} [\phi(s_h^k, a')]$ and $\tilde{\phi}(s_h^k, a) = \phi(s_h^k, a) - \mathbb{E}_{\theta_h^{k+1}} [\phi(s_h^k, a')]$. Now, we bound the terms on the right-hand side of (D.34) individually. For the first term, with probability at least $1 - \delta$, we get

$$\begin{aligned}
 & \sum_{k \in \mathcal{K}_{(i)}} \sum_{h=1}^H \sum_{a \in A_h^k} \sqrt{\mathcal{P}_h(a|s_h^k, A_h^k; \theta_h^*) \mathcal{P}_h(a|s_h^k, A_h^k; \theta_h^{k+1})} \left(f_{h,J(k,h)}^k(s_h^k, a) - \mathbb{E}_{\theta_h^*} [f_{h,J(k,h)}^k(s_h^k, a')] \right) \|\tilde{\phi}(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}} \\
 & \leq \sqrt{\sum_{k \in \mathcal{K}} \sum_{h=1}^H \sum_{a \in A_h^k} \mathcal{P}_h(a|s_h^k, A_h^k; \theta_h^*) \left(f_{h,J(k,h)}^k(s_h^k, a) - \mathbb{E}_{\theta_h^*} [f_{h,J(k,h)}^k(s_h^k, a')] \right)^2} \\
 & \quad \quad \quad =: [\mathbb{V}_h f_{h,J(k,h)}^k](s_h^k) \\
 & \quad \cdot \sqrt{\sum_{k \in \mathcal{K}} \sum_{h=1}^H \sum_{a \in A_h^k} \mathcal{P}_h(a|s_h^k, A_h^k; \theta_h^{k+1}) \|\tilde{\phi}(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}}^2} \\
 & \leq \sqrt{\sum_{k \in \mathcal{K}} \sum_{h=1}^H [\mathbb{V}_h f_{h,J(k,h)}^k](s_h^k)} \sqrt{2dH \log \left(1 + \frac{K}{d\lambda} \right)} \\
 & = \mathcal{O} \left(\sqrt{|\mathcal{K}| + H \log(1/\delta)} \right) \cdot \sqrt{2dH \log \left(1 + \frac{K}{d\lambda} \right)}, \tag{D.35}
 \end{aligned}$$

where the first inequality follows from the Cauchy-Schwarz inequality, the second-to-last inequality holds by Lemma D.7, and the last equality holds by Lemma D.12. Additionally, the second term in (D.34) can be bounded as follows:

$$\begin{aligned}
 & \sum_{k \in \mathcal{K}_{(i)}} \sum_{h=1}^H \sum_{a \in A_h^k} \left(\sqrt{\mathcal{P}_h(a|s_h^k, A_h^k; \theta_h^*)} - \sqrt{\mathcal{P}_h(a|s_h^k, A_h^k; \theta_h^{k+1})} \right) \sqrt{\mathcal{P}_h(a|s_h^k, A_h^k; \theta_h^*)} \\
 & \quad \cdot \left(f_{h,J(k,h)}^k(s_h^k, a) - \mathbb{E}_{\theta_h^*} [f_{h,J(k,h)}^k(s_h^k, a')] \right) \|\tilde{\phi}(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}} \\
 & \leq \sum_{k \in \mathcal{K}_{(i)}} \sum_{h=1}^H \sum_{a \in A_h^k} \frac{|\mathcal{P}_h(a|s_h^k, A_h^k; \theta_h^*) - \mathcal{P}_h(a|s_h^k, A_h^k; \theta_h^{k+1})|}{\sqrt{\mathcal{P}_h(a|s_h^k, A_h^k; \theta_h^*)} + \sqrt{\mathcal{P}_h(a|s_h^k, A_h^k; \theta_h^{k+1})}} \sqrt{\mathcal{P}_h(a|s_h^k, A_h^k; \theta_h^*)} \|\tilde{\phi}(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}} \\
 & \leq \sum_{k \in \mathcal{K}_{(i)}} \sum_{h=1}^H \sum_{a \in A_h^k} |\mathcal{P}_h(a|s_h^k, A_h^k; \theta_h^*) - \mathcal{P}_h(a|s_h^k, A_h^k; \theta_h^{k+1})| \|\tilde{\phi}(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}} \\
 & \leq 4\bar{\alpha}_K \sum_{k \in \mathcal{K}_{(i)}} \sum_{h=1}^H \max_{a \in A_h^k} \|\tilde{\phi}(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}} \max_{a \in A_h^k} \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}} \\
 & \leq 4\bar{\alpha}_K \sqrt{\sum_{k=1}^K \sum_{h=1}^H \max_{a \in A_h^k} \|\tilde{\phi}(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}}^2} \sqrt{\sum_{k=1}^K \sum_{h=1}^H \max_{a \in A_h^k} \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}}^2} \\
 & \leq \frac{4}{\sqrt{\kappa}} \bar{\alpha}_K \sqrt{\sum_{k=1}^K \sum_{h=1}^H \sum_{a \in A_h^k} \mathcal{P}_h(a|s_h^k, A_h^k; \theta_h^{k+1}) \|\tilde{\phi}(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}}^2} \sqrt{\sum_{k=1}^K \sum_{h=1}^H \max_{a \in A_h^k} \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}}^2} \\
 & \leq \frac{8}{\sqrt{\kappa}} \bar{\alpha}_K dH \log \left(1 + \frac{K}{d\lambda} \right), \tag{D.36}
 \end{aligned}$$

where the third inequity holds by Lemma D.11 and by the definition of $\bar{\alpha}_K = \max_{h \in [H]} \alpha_h^K$, the second-to-last inequality holds by the definition of κ , and the last inequality holds by Lemma D.7.

Finally, we bound the last term in (D.34). Using the inequality $\|\mathbf{a}\| - \|\mathbf{b}\| \leq \|\mathbf{a} - \mathbf{b}\|$ for any vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^d$, we have

$$\begin{aligned}
 & \sum_{k \in \mathcal{K}_{(i)}} \sum_{h=1}^H \sum_{a \in A_h^k} \mathcal{P}_h(a|s_h^k, A_h^k; \boldsymbol{\theta}_h^*) \left(\underbrace{f_{h,J(k,h)}^k(s_h^k, a) - \mathbb{E}_{\boldsymbol{\theta}_h^*} [f_{h,J(k,h)}^k(s_h^k, a')]}_{\geq 0} \right) \left(\|\tilde{\phi}(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}} - \|\bar{\phi}(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}} \right) \\
 & \leq \sum_{k \in \mathcal{K}_{(i)}} \sum_{h=1}^H \sum_{a \in A_h^k} \mathcal{P}_h(a|s_h^k, A_h^k; \boldsymbol{\theta}_h^*) \left\| \sum_{a' \in A_h^k} (\mathcal{P}_h(a'|s_h^k, A_h^k; \boldsymbol{\theta}_h^*) - \mathcal{P}_h(a'|s_h^k, A_h^k; \boldsymbol{\theta}_h^{k+1})) \phi(s_h^k, a') \right\|_{(\mathbf{H}_h^k)^{-1}} \\
 & \leq \sum_{k=1}^K \sum_{h=1}^H \sum_{a \in A_h^k} |\mathcal{P}_h(a|s_h^k, A_h^k; \boldsymbol{\theta}_h^*) - \mathcal{P}_h(a|s_h^k, A_h^k; \boldsymbol{\theta}_h^{k+1})| \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}} \\
 & \leq 4\bar{\alpha}_K \sum_{k=1}^K \sum_{h=1}^H \max_{a \in A_h^k} \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}}^2 \\
 & \leq \frac{8}{\kappa} \bar{\alpha}_K dH \log \left(1 + \frac{K}{d\lambda} \right), \tag{D.37}
 \end{aligned}$$

where the third inequity holds by Lemma D.11 and by the definition of $\bar{\alpha}_K = \max_{h \in [H]} \alpha_h^K$ and the last inequality holds by Lemma D.7.

By plugging (D.35), (D.36), and (D.37) into (D.34), and combining the result with (D.33), we obtain

$$\begin{aligned}
 & \sum_{k \in \mathcal{K}_{(i)}} \sum_{h=1}^H \nabla \tilde{R}(\mathbf{v}_h^*(s_h^k))^\top (\mathbf{v}_h^k(s_h^k) - \mathbf{v}_h^*(s_h^k)) \\
 & \leq \mathcal{O} \left(\sqrt{|\mathcal{K}| + H \log(1/\delta)} \right) \cdot 2\bar{\alpha}_K \sqrt{2dH \log \left(1 + \frac{K}{d\lambda} \right) + \frac{32}{\kappa} \bar{\alpha}_K^2 dH \log \left(1 + \frac{K}{d\lambda} \right)} \\
 & = \mathcal{O} \left(d\sqrt{H|\mathcal{K}|} (\log K)^{3/2} \log M + \frac{1}{\kappa} d^2 H (\log K)^3 (\log M)^2 \right). \tag{D.38}
 \end{aligned}$$

Now, we consider the second case to bound the term (A) in Equation (D.31).

Case (ii) For $(k, h) \in \mathcal{K} \times [H]$ such that $v_h^k(s_h^k, a) = \phi(s_h^k, a)^\top \boldsymbol{\theta}_h^* - 2\alpha_h^k \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}}$.

In this case, we know that $f_{h,J(k,h)}^k(s_h^k, a) < f_{h,J(k,h)}^k(s_h^k, a_0)$ for all $a \in \mathcal{I} \setminus \{a_0\}$. This implies that $|A_h^k| = 2$, since adding any item $a \in \mathcal{I} \setminus \{a_0\}$ to the set $\{a_0\}$ always decreases the expected value of $f_{h,J(k,h)}^k$. Furthermore, since we assume $\phi(s_h^k, a_0) = 0$ (which also implies $v_h^k(s_h^k, a_0) = 0$), and denoting $A_h^k = \{a_0, \tilde{a}_h^k\}$, we have:

$$\begin{aligned}
 & \sum_{k \in \mathcal{K}_{(ii)}} \sum_{h=1}^H \nabla \tilde{R}(\mathbf{v}_h^*(s_h^k))^\top (\mathbf{v}_h^k(s_h^k) - \mathbf{v}_h^*(s_h^k)) \\
 & = \sum_{k \in \mathcal{K}_{(ii)}} \sum_{h=1}^H 2\alpha_h^k \mathcal{P}_h(\tilde{a}_h^k|s_h^k, A_h^k; \boldsymbol{\theta}_h^*) f_{h,J(k,h)}^k(s_h^k, \tilde{a}_h^k) \left(\|\phi(s_h^k, \tilde{a}_h^k)\|_{(\mathbf{H}_h^k)^{-1}} - \mathcal{P}_h(\tilde{a}_h^k|s_h^k, A_h^k; \boldsymbol{\theta}_h^*) \left[\|\phi(s_h^k, \tilde{a}_h^k)\|_{(\mathbf{H}_h^k)^{-1}} \right] \right) \\
 & \leq 2\bar{\alpha}_K \sum_{k \in \mathcal{K}_{(ii)}} \sum_{h=1}^H \mathbb{E}_{\boldsymbol{\theta}_h^*} \left[\left(f_{h,J(k,h)}^k(s_h^k, a) - \mathbb{E}_{\boldsymbol{\theta}_h^*} [f_{h,J(k,h)}^k(s_h^k, a')] \right) \|\phi(s_h^k, a) - \mathbb{E}_{\boldsymbol{\theta}_h^*} [\phi(s_h^k, a')]\|_{(\mathbf{H}_h^k)^{-1}} \right], \tag{D.39}
 \end{aligned}$$

where in the inequality, we use the definition $\bar{\alpha}_K := \max_{h \in [H]} \alpha_h^K$. The rest of the analysis is similar to that in **Case (i)**. Therefore, we derive

$$\sum_{k \in \mathcal{K}_{(ii)}} \sum_{h=1}^H \nabla \tilde{R}(\mathbf{v}_h^*(s_h^k))^\top (\mathbf{v}_h^k(s_h^k) - \mathbf{v}_h^*(s_h^k)) = \mathcal{O} \left(d\sqrt{H|\mathcal{K}|} (\log K)^{3/2} \log M + \frac{1}{\kappa} d^2 H (\log K)^3 (\log M)^2 \right). \tag{D.40}$$

Now, we bound the term (B) in (D.31). Let $p_a(\bar{\mathbf{v}}_h^k(s_h^k)) = \frac{\exp(\bar{\mathbf{v}}_h^k(s_h^k, a))}{1 + \sum_{a'' \in A_h^k} \exp(\bar{\mathbf{v}}_h^k(s_h^k, a''))}$.

$$\begin{aligned}
 & \frac{1}{2} \sum_{k \in \mathcal{K}} \sum_{h=1}^H (\mathbf{v}_h^k(s_h^k) - \mathbf{v}_h^*(s_h^k))^\top \nabla^2 \tilde{R}(\bar{\mathbf{v}}_h^k(s_h^k)) (\mathbf{v}_h^k(s_h^k) - \mathbf{v}_h^*(s_h^k)) \\
 &= \frac{1}{2} \sum_{k \in \mathcal{K}} \sum_{h=1}^H \sum_{a \in A_h^k} \sum_{a' \in A_h^k} (\mathbf{v}_h^k(s_h^k, a) - \phi(s_h^k, a)^\top \boldsymbol{\theta}_h^*) \frac{\partial^2 \tilde{R}(\bar{\mathbf{v}}_h^k(s_h^k))}{\partial a \partial a'} (\mathbf{v}_h^k(s_h^k, a') - \phi(s_h^k, a')^\top \boldsymbol{\theta}_h^*) \\
 &= \frac{1}{2} \sum_{k \in \mathcal{K}} \sum_{h=1}^H \sum_{a \in A_h^k} \sum_{\substack{a' \in A_h^k \\ a' \neq a}} (\mathbf{v}_h^k(s_h^k, a) - \phi(s_h^k, a)^\top \boldsymbol{\theta}_h^*) \frac{\partial^2 \tilde{R}(\bar{\mathbf{v}}_h^k(s_h^k))}{\partial a \partial a'} (\mathbf{v}_h^k(s_h^k, a') - \phi(s_h^k, a')^\top \boldsymbol{\theta}_h^*) \\
 &+ \frac{1}{2} \sum_{k \in \mathcal{K}} \sum_{h=1}^H \sum_{a \in A_h^k} (\mathbf{v}_h^k(s_h^k, a) - \phi(s_h^k, a)^\top \boldsymbol{\theta}_h^*)^2 \frac{\partial^2 \tilde{R}(\bar{\mathbf{v}}_h^k(s_h^k))}{\partial a \partial a} \\
 &\leq \sum_{k \in \mathcal{K}} \sum_{h=1}^H \sum_{a \in A_h^k} \sum_{\substack{a' \in A_h^k \\ a' \neq a}} |v_h^k(s_h^k, a) - \phi(s_h^k, a)^\top \boldsymbol{\theta}_h^*| p_a(\bar{\mathbf{v}}_h^k(s_h^k)) p_{a'}(\bar{\mathbf{v}}_h^k(s_h^k)) |v_h^k(s_h^k, a') - \phi(s_h^k, a')^\top \boldsymbol{\theta}_h^*| \\
 &+ \frac{3}{2} \sum_{k \in \mathcal{K}} \sum_{h=1}^H \sum_{a \in A_h^k} (\mathbf{v}_h^k(s_h^k, a) - \phi(s_h^k, a)^\top \boldsymbol{\theta}_h^*)^2 p_a(\bar{\mathbf{v}}_h^k(s_h^k)), \tag{D.41}
 \end{aligned}$$

where the inequality holds by Lemma D.8 and $f_{j,J(k,h)}^k \leq 1$. To bound the first term in (D.41), by applying the AM-GM inequality, we get

$$\begin{aligned}
 & \sum_{k \in \mathcal{K}} \sum_{h=1}^H \sum_{a \in A_h^k} \sum_{\substack{a' \in A_h^k \\ a' \neq a}} |v_h^k(s_h^k, a) - \phi(s_h^k, a)^\top \boldsymbol{\theta}_h^*| p_a(\bar{\mathbf{v}}_h^k(s_h^k)) p_{a'}(\bar{\mathbf{v}}_h^k(s_h^k)) |v_h^k(s_h^k, a') - \phi(s_h^k, a')^\top \boldsymbol{\theta}_h^*| \\
 &\leq \sum_{k \in \mathcal{K}} \sum_{h=1}^H \sum_{a \in A_h^k} \sum_{a' \in A_h^k} |v_h^k(s_h^k, a) - \phi(s_h^k, a)^\top \boldsymbol{\theta}_h^*| p_a(\bar{\mathbf{v}}_h^k(s_h^k)) p_{a'}(\bar{\mathbf{v}}_h^k(s_h^k)) |v_h^k(s_h^k, a') - \phi(s_h^k, a')^\top \boldsymbol{\theta}_h^*| \\
 &\leq \frac{1}{2} \sum_{k \in \mathcal{K}} \sum_{h=1}^H \sum_{a \in A_h^k} \sum_{a' \in A_h^k} (\mathbf{v}_h^k(s_h^k, a) - \phi(s_h^k, a)^\top \boldsymbol{\theta}_h^*)^2 p_a(\bar{\mathbf{v}}_h^k(s_h^k)) p_{a'}(\bar{\mathbf{v}}_h^k(s_h^k)) \\
 &+ \frac{1}{2} \sum_{k \in \mathcal{K}} \sum_{h=1}^H \sum_{a \in A_h^k} \sum_{a' \in A_h^k} p_a(\bar{\mathbf{v}}_h^k(s_h^k)) p_{a'}(\bar{\mathbf{v}}_h^k(s_h^k)) (\mathbf{v}_h^k(s_h^k, a') - \phi(s_h^k, a')^\top \boldsymbol{\theta}_h^*)^2 \\
 &= \sum_{k \in \mathcal{K}} \sum_{h=1}^H \sum_{a \in A_h^k} (\mathbf{v}_h^k(s_h^k, a) - \phi(s_h^k, a)^\top \boldsymbol{\theta}_h^*)^2 p_a(\bar{\mathbf{v}}_h^k(s_h^k)). \tag{D.42}
 \end{aligned}$$

Plugging (D.42) into (D.41), we have

$$\begin{aligned}
 & \frac{1}{2} \sum_{k \in \mathcal{K}} \sum_{h=1}^H (\mathbf{v}_h^k(s_h^k) - \mathbf{v}_h^*(s_h^k))^\top \nabla^2 \tilde{R}(\bar{\mathbf{v}}_h^k(s_h^k)) (\mathbf{v}_h^k(s_h^k) - \mathbf{v}_h^*(s_h^k)) \\
 & \leq \frac{5}{2} \sum_{k \in \mathcal{K}} \sum_{h=1}^H \sum_{a \in A_h^k} p_a(\bar{\mathbf{v}}_h^k(s_h^k)) (v_h^k(s_h^k, a) - \phi(s_h^k, a)^\top \boldsymbol{\theta}_h^*)^2 \\
 & \leq 10 (\bar{\alpha}_K)^2 \sum_{k \in \mathcal{K}} \sum_{h=1}^H \sum_{a \in A_h^k} p_a(\bar{\mathbf{v}}_h^k(s_h^k)) \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}} \\
 & \leq 10 (\bar{\alpha}_K)^2 \sum_{k \in \mathcal{K}} \sum_{h=1}^H \max_{a \in A_h^k} \|\phi(s_h^k, a)\|_{(\mathbf{H}_h^k)^{-1}} = \mathcal{O} \left(\frac{1}{\kappa} d^2 H (\log K)^3 (\log M)^2 \right), \tag{D.43}
 \end{aligned}$$

where the last inequality holds by Lemma D.7.

Combining (D.38), (D.40) and (D.43), we obtain

$$\begin{aligned}
 & \sum_{k \in \mathcal{K}} \sum_{h=1}^H \sum_{a \in A_h^k} \left(\tilde{\mathcal{P}}_{h,J(k,h)}^k(a|s_h^k, A_h^k) - \mathcal{P}_h(a|s_h^k, A_h^k) \right) f_{h,J(k,h)}^k(s_h^k, a) \\
 & = \mathcal{O} \left(d \sqrt{H|\mathcal{K}|} (\log K)^{3/2} \log M + \frac{1}{\kappa} d^2 H (\log K)^3 (\log M)^2 \right).
 \end{aligned}$$

This conclude the proof of Lemma D.13. \square

D.4. Optimism

In this subsection, we prove the optimism of our value estimates V_h^k .

Lemma D.14 (Point-wise monotonicity, Lemma 31 of Agarwal et al. 2023). *Suppose Algorithm 1 uses a consistent bonus oracle satisfying Definition B.1. For any fixed $(k, h) \in [K] \times [H]$, conditioning on events $\mathcal{E}_{\leq k-1} \cap \left(\bigcap_{h'=h}^H \mathcal{E}_{h'}^k \right)$, for all $(s_h, a_h) \in \mathcal{S} \times \mathcal{I}$, we have*

1. $\bar{Q}_h^*(s_h, a_h) \leq f_{h,1}^k(s_h, a_h);$
2. $f_{h,-1}^k(s_h, a_h) \leq \bar{Q}_h^*(s_h, a_h);$
3. $f_{h,2}^\tau(s_h, a_h) \geq \max \left\{ \mathcal{T}_h V_{h+1,1}^k(s_h, a_h), f_{h,1}^k(s_h, a_h) \right\}, \quad \forall \tau \in [k].$

Lemma D.15 (Optimism). *Let V_h^k be the realized optimistic value function defined in (D.18). Suppose Algorithm 1 uses a consistent bonus oracle satisfying Definition B.1. On the even conditioning on the good event $\mathcal{E}^\theta \cap \mathcal{E}_{\leq K}$, for all $(k, h) \in [K] \times [H]$, we have*

$$V_h^k(s_h^k) \geq V_h^*(s_h^k).$$

Proof of Lemma D.15. We denote $A_h^{k,*} \in \operatorname{argmax}_A \sum_{a \in A} \mathcal{P}_h(a|s_h^k, A) \bar{Q}_h^*(s_h^k, a)$. If $A_h^k = A_{h,1}^k$, by the definition of the

optimal value function V_h^* , we have

$$\begin{aligned}
 V_h^*(s_h^k) &= \max_{A \in \mathcal{A}} \sum_{a \in A} \mathcal{P}_h(a|s_h^k, A) \overline{Q}_h^*(s_h^k, a) \\
 &= \sum_{a \in A_h^{k,*}} \mathcal{P}_h(a|s_h^k, A_h^{k,*}) \overline{Q}_h^*(s_h^k, a) \\
 &\leq \sum_{a \in A_h^{k,*}} \mathcal{P}_h(a|s_h^k, A_h^{k,*}) f_{h,1}^k(s_h^k, a) \\
 &\leq \sum_{a \in \tilde{A}_h^k} \tilde{\mathcal{P}}_{h,1}^k(a|s_h^k, \tilde{A}_h^k) f_{h,1}^k(s_h^k, a) \\
 &\leq \sum_{a \in A_{h,1}^k} \tilde{\mathcal{P}}_{h,1}^k(a|s_h^k, A_{h,1}^k) f_{h,1}^k(s_h^k, a) \\
 &= \max_{A \in \mathcal{A}} \sum_{a \in A} \tilde{\mathcal{P}}_{h,1}^k(a|s_h^k, A) f_{h,1}^k(s_h^k, a) = V_h^k(s_h^k),
 \end{aligned}$$

where the first inequality holds by Lemma D.14, in the second inequality, we use the fact that, by Lemma D.5, there exists a subset $\tilde{A}_h^k \subseteq A_h^{k,*}$ with $\tilde{A}_h^k \in \mathcal{A}$ such that the inequality holds, and the last inequality holds by the definition of $A_{h,1}^k$.

The case where $A_h^k = A_{h,2}^k$ can be proven using the same reasoning. \square

D.5. Variances

In this subsection, we present properties related to variances.

Lemma D.16 (Upper bound of variance estimator, Lemma 34 of Agarwal et al. 2023). *Let $z_h^k = (s_h^k, a_h^k)$. We denote $\mathbb{E}_{\mathbb{P}}[\cdot|s_h^k, a_h^k] = \mathbb{E}_{s_{h+1} \sim \mathbb{P}_h(\cdot|s_h^k, a_h^k)}[\cdot|s_h^k, a_h^k]$ and $\mathbb{V}_{\mathbb{P}}[\cdot|s_h^k, a_h^k] = \mathbb{V}_{s_{h+1} \sim \mathbb{P}_h(\cdot|s_h^k, a_h^k)}[\cdot|s_h^k, a_h^k]$, where the expectation is only taken over s_{h+1} due to the model transition for shorthand. Suppose Algorithm 1 uses a consistent bonus oracle satisfying Definition B.1. For any episode $k \geq 2$ conditioning on the good event $\mathcal{E}_{\leq K}$, the variance estimator σ_h^k satisfies*

$$\begin{aligned}
 (\sigma_h^k)^2 &\leq \mathbb{V}[r_h + V_{h+1,1}^k(s_{h+1}) | z_h^k] + 4(f_{h,2}^k(z_h^k) - f_{h,-2}^k(z_h^k)) \\
 &\quad + 4 \min \left\{ 1, D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\mathbf{1}^\tau\}_{\tau=1}^{k-1}) \cdot \left(2\sqrt{(\bar{\beta}_h^k)^2 + \rho} + 4L\sqrt{(\beta_{h,2}^k)^2 + \rho} \right) \right\}.
 \end{aligned}$$

Lemma D.17 (Sum of variances, Corollary 50 of Agarwal et al. 2023). *Let $z_h^k = (s_h^k, a_h^k)$. We denote $\mathbb{E}_{\mathbb{P}}[\cdot|s_h^k, a_h^k] = \mathbb{E}_{s_{h+1} \sim \mathbb{P}_h(\cdot|s_h^k, a_h^k)}[\cdot|s_h^k, a_h^k]$ and $\mathbb{V}_{\mathbb{P}}[\cdot|s_h^k, a_h^k] = \mathbb{V}_{s_{h+1} \sim \mathbb{P}_h(\cdot|s_h^k, a_h^k)}[\cdot|s_h^k, a_h^k]$, where the expectation is only taken over s_{h+1} due to the model transition for shorthand. When $L = \mathcal{O}(1)$, with probability at least $1 - \delta$, we have*

$$\begin{aligned}
 &\sum_{k=1}^K \sum_{h=1}^H \mathbb{V}[r_h + V_{h+1,1}^k(s_{h+1}) | z_h^k] \\
 &\leq \mathcal{O} \left(H \sum_{k=1}^K \sum_{h=1}^H (f_{h,2}^k(z_h^k) - f_{h,-2}^k(z_h^k)) + K + KH^2\delta + H^2|\mathcal{K}_{oo}| + H^4 \log^2 \frac{KH}{\delta} \right).
 \end{aligned}$$

D.6. Approximation Error of Optimistic, Overly Optimistic (Pessimistic) \overline{Q} -values

In this section, we provide some inequalities for bounding the optimistic values, overly optimistic values, and overly pessimistic values sequence, which are useful for the proofs in Subsection D.7.

Lemma D.18 (Approximation error of overly pessimistic \overline{Q}). *Suppose Algorithm 1 uses a consistent bonus oracle satisfying*

Definition B.1. Conditioning on the good event $\mathcal{E}_{\leq K}$, for any $(k, h) \in [K] \times [H]$, it holds that

$$\begin{aligned} (f_{h,-2}^k - \bar{Q}_h^{\pi_k})(s_h^k, a_h^k) &\geq \sum_{h'=h+1}^H \sum_{a' \in A_{h'}^k} \left(\tilde{\mathcal{P}}_{h',-2}^k(a'|s_{h'}^k, A_{h'}^k) - \mathcal{P}_{h'}(a'|s_{h'}^k, A_{h'}^k) \right) f_{h',-2}^k(s_{h'}^k, a') \\ &\quad - 2 \sum_{h'=h}^H b_{h',2}^k(s_{h'}^k, a_{h'}^k) + \sum_{h'=h+1}^H \zeta_{h',-2}^k + \sum_{h'=h+1}^H \dot{\zeta}_{h',-2}^k, \end{aligned}$$

where $\zeta_{h,-2}^k := \mathbb{E}_{\mathbb{P}} \left[(V_{h,-2}^k - V_h^{\pi_k})(s_h) \mid s_{h-1}^k, a_{h-1}^k \right] - (V_{h,-2}^k - V_h^{\pi_k})(s_h^k)$ and $\dot{\zeta}_{h,-2}^k := \mathbb{E}_{\mathcal{P}} \left[\left(f_{h,-2}^k - \bar{Q}_h^{\pi_k} \right)(s_h^k, a_h^k) \mid s_h^k, A_h^k \right] - \left(f_{h,-2}^k - \bar{Q}_h^{\pi_k} \right)(s_h^k, a_h^k)$.

Proof of Lemma D.18. Under the event $\mathcal{E}_{\leq K}$, we have

$$\begin{aligned} (f_{h,-2}^k - \bar{Q}_h^{\pi_k})(s_h^k, a_h^k) &= (f_{h,-2}^k - \mathcal{T}_h V_{h+1,-2}^k)(s_h^k, a_h^k) + (\mathcal{T}_h V_{h+1,-2}^k - \bar{Q}_h^{\pi_k})(s_h^k, a_h^k) \\ &\geq -2b_{h,2}^k(s_h^k, a_h^k) + \mathbb{E}_{\mathbb{P}} \left[(r_h - r_h) + (V_{h+1,-2}^k - V_{h+1}^{\pi_k})(s_{h+1}) \mid s_h^k, a_h^k \right] \\ &= -2b_{h,2}^k(s_h^k, a_h^k) + (V_{h+1,-2}^k - V_{h+1}^{\pi_k})(s_{h+1}^k) + \zeta_{h+1,-2}^k \\ &\geq -2b_{h,2}^k(s_h^k, a_h^k) + (Q_{h+1,-2}^k - Q_{h+1}^{\pi_k})(s_{h+1}^k, A_{h+1}^k) + \zeta_{h+1,-2}^k \\ &= \sum_{a' \in A_{h+1}^k} \left(\tilde{\mathcal{P}}_{h+1,-2}^k(a'|s_{h+1}^k, A_{h+1}^k) f_{h+1,-2}^k(s_{h+1}^k, a') - \mathcal{P}_{h+1}(a'|s_{h+1}^k, A_{h+1}^k) \bar{Q}_{h+1}^{\pi_k}(s_{h+1}^k, a') \right) \\ &\quad - 2b_{h,2}^k(s_h^k, a_h^k) + \zeta_{h+1,-2}^k \\ &= \sum_{a' \in A_{h+1}^k} \left(\tilde{\mathcal{P}}_{h+1,-2}^k(a'|s_{h+1}^k, A_{h+1}^k) - \mathcal{P}_{h+1}(a'|s_{h+1}^k, A_{h+1}^k) \right) f_{h+1,-2}^k(s_{h+1}^k, a') \\ &\quad + \mathbb{E}_{\mathcal{P}} \left[\left(f_{h+1,-2}^k - \bar{Q}_{h+1}^{\pi_k} \right)(s_{h+1}^k, a_{h+1}^k) \mid s_{h+1}^k, A_{h+1}^k \right] - 2b_{h,2}^k(s_h^k, a_h^k) + \zeta_{h+1,-2}^k \\ &= \sum_{a' \in A_{h+1}^k} \left(\tilde{\mathcal{P}}_{h+1,-2}^k(a'|s_{h+1}^k, A_{h+1}^k) - \mathcal{P}_{h+1}(a'|s_{h+1}^k, A_{h+1}^k) \right) f_{h+1,-2}^k(s_{h+1}^k, a') \\ &\quad + \left(f_{h+1,-2}^k - \bar{Q}_{h+1}^{\pi_k} \right)(s_{h+1}^k, a_{h+1}^k) - 2b_{h,2}^k(s_h^k, a_h^k) + \zeta_{h+1,-2}^k + \dot{\zeta}_{h+1,-2}^k, \end{aligned}$$

where the first inequality holds because $\mathcal{T}_h V_{h+1,-2}^k \in \mathcal{F}_{h,-2}^k$ under the event $\mathcal{E}_{\leq K}$ and definition of $b_{h,2}^k$, and the last inequality holds since $V_{h+1,-2}^k(s_{h+1}^k) \geq Q_{h+1,-2}^k(s_{h+1}^k, A_{h+1}^k)$.

Hence, by recursion we obtain that

$$\begin{aligned} (f_{h,-2}^k - \bar{Q}_h^{\pi_k})(s_h^k, a_h^k) &\geq \sum_{h'=h+1}^H \sum_{a' \in A_{h'}^k} \left(\tilde{\mathcal{P}}_{h',-2}^k(a'|s_{h'}^k, A_{h'}^k) - \mathcal{P}_{h'}(a'|s_{h'}^k, A_{h'}^k) \right) f_{h',-2}^k(s_{h'}^k, a') \\ &\quad - 2 \sum_{h'=h}^H b_{h',2}^k(s_{h'}^k, a_{h'}^k) + \sum_{h'=h+1}^H \zeta_{h',-2}^k + \sum_{h'=h+1}^H \dot{\zeta}_{h',-2}^k. \end{aligned}$$

□

Lemma D.19 (Approximation error of overly optimistic \bar{Q}). Suppose Algorithm 1 uses a consistent bonus oracle satisfying Definition B.1. Conditioning on the good event $\mathcal{E}_{\leq K}$, for any $k \in [K]$ and any $h \geq h_k$, it holds that

$$\begin{aligned} (f_{h,2}^k - \bar{Q}_h^{\pi_k})(s_h^k, a_h^k) &\leq \sum_{h'=h+1}^H \sum_{a' \in A_{h'}^k} \left(\tilde{\mathcal{P}}_{h',2}^k(a'|s_{h'}^k, A_{h'}^k) - \mathcal{P}_{h'}(a'|s_{h'}^k, A_{h'}^k) \right) f_{h',2}^k(s_{h'}^k, a') \\ &\quad + 2 \sum_{h'=h}^H b_{h',1}^k(s_{h'}^k, a_{h'}^k) + 2 \sum_{h'=h}^H b_{h',2}^k(s_{h'}^k, a_{h'}^k) + \sum_{h'=h+1}^H \zeta_{h',2}^k + \sum_{h'=h+1}^H \dot{\zeta}_{h',2}^k, \end{aligned}$$

where $\zeta_{h,2}^k := \mathbb{E}_{\mathbb{P}} \left[(V_{h,2}^k - V_h^{\pi_k})(s_h) \mid s_{h-1}^k, a_{h-1}^k \right] - (V_{h,2}^k - V_h^{\pi_k})(s_h^k)$ and $\dot{\zeta}_{h,2}^k := \mathbb{E}_{\mathcal{P}} \left[\left(f_{h,2}^k - \overline{Q}_h^{\pi_k} \right) (s_h^k, a_h) \mid s_h^k, A_h^k \right] - \left(f_{h,2}^k - \overline{Q}_h^{\pi_k} \right) (s_h^k, a_h^k)$.

Proof of Lemma D.19. Under the event $\mathcal{E}_{\leq K}$, at $h \geq h_k$, we have

$$\begin{aligned}
 (f_{h,2}^k - \overline{Q}_h^{\pi_k})(s_h^k, a_h^k) &= (f_{h,2}^k - \mathcal{T}_h V_{h+1,2}^k)(s_h^k, a_h^k) + (\mathcal{T}_h V_{h+1,2}^k - \overline{Q}_h^{\pi_k})(s_h^k, a_h^k) \\
 &\leq 2b_{h,1}^k(s_h^k, a_h^k) + 2b_{h,2}^k(s_h^k, a_h^k) + \mathbb{E} \left[(V_{h+1,2}^k - V_{h+1}^{\pi_k})(s_{h+1}) \mid s_h^k, a_h^k \right] \\
 &= 2b_{h,1}^k(s_h^k, a_h^k) + 2b_{h,2}^k(s_h^k, a_h^k) + (V_{h+1,2}^k - V_{h+1}^{\pi_k})(s_{h+1}^k) + \zeta_{h+1,2}^k \\
 &= 2b_{h,1}^k(s_h^k, a_h^k) + 2b_{h,2}^k(s_h^k, a_h^k) + (Q_{h+1,2}^k - Q_{h+1}^{\pi_k})(s_{h+1}^k, A_{h+1}^k) + \zeta_{h+1,2}^k \\
 &= \sum_{a' \in A_{h+1}^k} \left(\tilde{\mathcal{P}}_{h+1,2}^k(a' \mid s_{h+1}^k, A_{h+1}^k) f_{h+1,2}^k(s_{h+1}^k, a') - \mathcal{P}_{h+1}(a' \mid s_{h+1}^k, A_{h+1}^k) \overline{Q}_{h+1}^{\pi_k}(s_{h+1}^k, a') \right) \\
 &\quad + 2b_{h,1}^k(s_h^k, a_h^k) + 2b_{h,2}^k(s_h^k, a_h^k) + \zeta_{h+1,2}^k \\
 &= \sum_{a' \in A_{h+1}^k} \left(\tilde{\mathcal{P}}_{h+1,2}^k(a' \mid s_{h+1}^k, A_{h+1}^k) - \mathcal{P}_{h+1}(a' \mid s_{h+1}^k, A_{h+1}^k) \right) f_{h+1,2}^k(s_{h+1}^k, a') \\
 &\quad + \mathbb{E}_{\mathcal{P}} \left[\left(f_{h+1,2}^k - \overline{Q}_{h+1}^{\pi_k} \right) (s_{h+1}^k, a_{h+1}) \mid s_{h+1}^k, A_{h+1}^k \right] + 2b_{h,1}^k(s_h^k, a_h^k) + 2b_{h,2}^k(s_h^k, a_h^k) + \zeta_{h+1,2}^k \\
 &= \sum_{a' \in A_{h+1}^k} \left(\tilde{\mathcal{P}}_{h+1,2}^k(a' \mid s_{h+1}^k, A_{h+1}^k) - \mathcal{P}_{h+1}(a' \mid s_{h+1}^k, A_{h+1}^k) \right) f_{h+1,2}^k(s_{h+1}^k, a') \\
 &\quad + \left(f_{h+1,2}^k - \overline{Q}_{h+1}^{\pi_k} \right) (s_{h+1}^k, a_{h+1}^k) + 2b_{h,1}^k(s_h^k, a_h^k) + 2b_{h,2}^k(s_h^k, a_h^k) + \zeta_{h+1,2}^k + \dot{\zeta}_{h+1,2}^k,
 \end{aligned}$$

where the first inequality holds based on the assumption that $\mathcal{T}_h V_{h+1,2}^k \in \mathcal{F}_{h,2}^k$ and definition of $b_{h,2}^k$, and the third equality holds because for $h \geq h_k$, we know that $A_{h+1}^k \in \arg\max_A Q_{h+1}^k(s_{h+1}^k, A) = \arg\max_A \sum_{a \in A} \tilde{\mathcal{P}}_{h+1,2}^k(a \mid s_{h+1}^k, A) f_{h+1,2}^k(s_{h+1}^k, a)$ by the data collection policy in (9).

Therefore, by recursion we get

$$\begin{aligned}
 (f_{h,2}^k - \overline{Q}_h^{\pi_k})(s_h^k, a_h^k) &\leq \sum_{h'=h+1}^H \sum_{a' \in A_{h'}^k} \left(\tilde{\mathcal{P}}_{h',2}^k(a' \mid s_{h'}^k, A_{h'}^k) - \mathcal{P}_{h'}(a' \mid s_{h'}^k, A_{h'}^k) \right) f_{h',2}^k(s_{h'}^k, a') \\
 &\quad + 2 \sum_{h'=h}^H b_{h',1}^k(s_{h'}^k, a_{h'}^k) + 2 \sum_{h'=h}^H b_{h',2}^k(s_{h'}^k, a_{h'}^k) + \sum_{h'=h+1}^H \zeta_{h',2}^k + \sum_{h'=h+1}^H \dot{\zeta}_{h',2}^k.
 \end{aligned}$$

□

Lemma D.20 (Approximation error of optimistic \overline{Q}). *Suppose Algorithm 1 uses a consistent bonus oracle satisfying Definition B.1. Conditioning on the good event $\mathcal{E}^\theta \cap \mathcal{E}_{\leq K}$, for any $k \in [K]$ and any $h \leq h_k$, we have*

$$\begin{aligned}
 (f_{h,1}^k - \overline{Q}_h^{\pi_k})(s_h^k, a_h^k) &\leq \sum_{h'=h+1}^{h_k-1} \sum_{a' \in A_{h'}^k} \left(\tilde{\mathcal{P}}_{h',1}^k(a' \mid s_{h'}^k, A_{h'}^k) - \mathcal{P}_{h'}(a' \mid s_{h'}^k, A_{h'}^k) \right) f_{h',1}^k(s_{h'}^k, a') \\
 &\quad + \sum_{h'=h_k}^H \sum_{a' \in A_{h'}^k} \left(\tilde{\mathcal{P}}_{h',2}^k(a' \mid s_{h'}^k, A_{h'}^k) - \mathcal{P}_{h'}(a' \mid s_{h'}^k, A_{h'}^k) \right) f_{h',2}^k(s_{h'}^k, a') \\
 &\quad + 2 \sum_{h'=h}^H b_{h',1}^k(s_{h'}^k, a_{h'}^k) + 2 \sum_{h'=h_k}^H b_{h',2}^k(s_{h'}^k, a_{h'}^k) + \sum_{h'=h+1}^{h_k-1} (\zeta_{h',1}^k + \dot{\zeta}_{h',1}^k) + \sum_{h'=h_k}^H (\zeta_{h',2}^k + \dot{\zeta}_{h',2}^k),
 \end{aligned}$$

where

$$\begin{aligned}\zeta_{h,1}^k &:= \mathbb{E}_{\mathbb{P}} \left[(V_{h,1}^k - V_h^{\pi_k})(s_h) \mid s_{h-1}^k, a_{h-1}^k \right] - (V_{h,1}^k - V_h^{\pi_k})(s_h^k), \\ \dot{\zeta}_{h,1}^k &:= \mathbb{E}_{\mathcal{P}} \left[\left(f_{h,1}^k - \overline{Q}_h^{\pi_k} \right) (s_h^k, a_h) \mid s_h^k, A_h^k \right] - \left(f_{h,1}^k - \overline{Q}_h^{\pi_k} \right) (s_h^k, a_h^k), \\ \zeta_{h,2}^k &:= \mathbb{E}_{\mathbb{P}} \left[(V_{h,2}^k - V_h^{\pi_k})(s_h) \mid s_{h-1}^k, a_{h-1}^k \right] - (V_{h,2}^k - V_h^{\pi_k})(s_h^k), \\ \dot{\zeta}_{h,2}^k &:= \mathbb{E}_{\mathcal{P}} \left[\left(f_{h,2}^k - \overline{Q}_h^{\pi_k} \right) (s_h^k, a_h) \mid s_h^k, A_h^k \right] - \left(f_{h,2}^k - \overline{Q}_h^{\pi_k} \right) (s_h^k, a_h^k).\end{aligned}$$

Here, following the convention, we use the empty sum notation, i.e., $\sum_{i=a}^b x_i = 0$, when $b \leq a$.

Proof of Lemma D.20. Under the event $\mathcal{E}_{\leq K}$, by Lemma D.14 it holds that $f_{h,1}^k(s, a) \leq f_{h,2}^k(s, a)$ for all $(s, a) \in \mathcal{S} \times \mathcal{I}$. Therefore, for $h = h_k$, by Lemma D.19, we get

$$\begin{aligned}(f_{h,1}^k - \overline{Q}_h^{\pi_k})(s_h^k, a_h^k) &\leq (f_{h,2}^k - \overline{Q}_h^{\pi_k})(s_h^k, a_h^k) \\ &\leq \sum_{h'=h+1}^H \sum_{a' \in A_{h'}^k} \left(\tilde{\mathcal{P}}_{h',2}^k(a' | s_{h'}^k, A_{h'}^k) - \mathcal{P}_{h'}(a' | s_{h'}^k, A_{h'}^k) \right) f_{h',2}^k(s_{h'}^k, a') \\ &\quad + 2 \sum_{h'=h}^H b_{h',1}^k(s_{h'}^k, a_{h'}^k) + 2 \sum_{h'=h}^H b_{h',2}^k(s_{h'}^k, a_{h'}^k) + \sum_{h'=h+1}^H \zeta_{h',2}^k + \sum_{h'=h+1}^H \dot{\zeta}_{h',2}^k.\end{aligned}\tag{D.44}$$

For $h = h_k - 1$, we have

$$\begin{aligned}(f_{h,1}^k - \overline{Q}_h^{\pi_k})(s_h^k, a_h^k) &= (f_{h,1}^k - \mathcal{T}_h V_{h+1,1}^k)(s_h^k, a_h^k) + (\mathcal{T}_h V_{h+1,1}^k - \overline{Q}_h^{\pi_k})(s_h^k, a_h^k) \\ &\leq 2b_{h,1}^k(s_h^k, a_h^k) + \mathbb{E} \left[(V_{h+1,1}^k - V_h^{\pi_k})(s_{h+1}) \mid s_h^k, a_h^k \right] \\ &\leq 2b_{h,1}^k(s_h^k, a_h^k) + \mathbb{E} \left[(V_{h+1,2}^k - V_h^{\pi_k})(s_{h+1}) \mid s_h^k, a_h^k \right] \\ &= 2b_{h,1}^k(s_h^k, a_h^k) + (V_{h+1,2}^k - V_h^{\pi_k})(s_{h+1}^k) + \zeta_{h+1,2}^k \\ &= 2b_{h,1}^k(s_h^k, a_h^k) + (Q_{h+1,2}^k - \overline{Q}_h^{\pi_k})(s_{h+1}^k, A_{h+1}^k) + \zeta_{h+1,2}^k \\ &\leq \sum_{a' \in A_{h+1}^k} \left(\tilde{\mathcal{P}}_{h+1,2}^k(a' | s_{h+1}^k, A_{h+1}^k) f_{h+1,2}^k(s_{h+1}^k, a') - \mathcal{P}_{h+1}(a' | s_{h+1}^k, A_{h+1}^k) \overline{Q}_{h+1}^{\pi_k}(s_{h+1}^k, a') \right) \\ &\quad + 2b_{h,1}^k(s_h^k, a_h^k) + \zeta_{h+1,2}^k \\ &= \sum_{a' \in A_{h+1}^k} \left(\tilde{\mathcal{P}}_{h+1,2}^k(a' | s_{h+1}^k, A_{h+1}^k) - \mathcal{P}_{h+1}(a' | s_{h+1}^k, A_{h+1}^k) \right) f_{h+1,2}^k(s_{h+1}^k, a') \\ &\quad + \mathbb{E}_{\mathcal{P}} \left[\left(f_{h+1,2}^k - \overline{Q}_{h+1}^{\pi_k} \right) (s_{h+1}^k, a_{h+1}^k) \mid s_{h+1}^k, A_{h+1}^k \right] + 2b_{h,1}^k(s_h^k, a_h^k) + \zeta_{h+1,2}^k \\ &= \sum_{a' \in A_{h+1}^k} \left(\tilde{\mathcal{P}}_{h+1,2}^k(a' | s_{h+1}^k, A_{h+1}^k) - \mathcal{P}_{h+1}(a' | s_{h+1}^k, A_{h+1}^k) \right) f_{h+1,2}^k(s_{h+1}^k, a') \\ &\quad + \left(f_{h+1,2}^k - \overline{Q}_{h+1}^{\pi_k} \right) (s_{h+1}^k, a_{h+1}^k) + 2b_{h,1}^k(s_h^k, a_h^k) + \zeta_{h+1,2}^k + \dot{\zeta}_{h+1,2}^k,\end{aligned}\tag{D.45}$$

where the first inequality holds based on the assumption that $\mathcal{T}_h V_{h+1,1}^k \in \mathcal{F}_{h,1}^k$ and definition of $b_{h,1}^k$, and the second inequality holds because for any $s_{h+1} \in \mathcal{S}$, we have

$$\begin{aligned}V_{h+1,1}^k(s_{h+1}) &= \sum_{a' \in A_{h+1,1}} \tilde{\mathcal{P}}_{h+1,1}^k(a' | s_{h+1}, A_{h+1,1}) f_{h+1,1}^k(s_{h+1}, a') \\ &\leq \sum_{a' \in A_{h+1,1}} \tilde{\mathcal{P}}_{h+1,1}^k(a' | s_{h+1}, A_{h+1,1}) f_{h+1,2}^k(s_{h+1}, a') \\ &\leq \sum_{a' \in \tilde{A}_{h+1,1}} \tilde{\mathcal{P}}_{h+1,2}^k(a' | s_{h+1}, \tilde{A}_{h+1,1}) f_{h+1,2}^k(s_{h+1}, a') \\ &\leq \sum_{a' \in A_{h+1,2}^k} \tilde{\mathcal{P}}_{h+1,2}^k(a' | s_{h+1}, A_{h+1,2}^k) f_{h+1,2}^k(s_{h+1}, a') = V_{h+1,2}^k(s_{h+1}),\end{aligned}$$

where in the first equality, we denote $A_{h+1,1} \in \operatorname{argmax}_A \sum_{a' \in A} \tilde{\mathcal{P}}_{h+1,1}^k(a'|s_{h+1}, A) f_{h+1,1}^k(s_{h+1}, a')$, for the first inequality, we use the fact that $f_{h+1,1}^k(s, a) \leq f_{h+1,2}^k(s, a)$ (Lemma D.14), the second inequality holds for some $\tilde{A}_{h+1,1} \subseteq A_{h+1,1}$ (Lemma D.5), and the last inequality follows from the definition of $A_{h+1,2}^k$. Moreover, the third equality of (D.45) holds because for $h+1 = h_k$, we know that $A_{h+1}^k = A_{h+1,2}^k$ by the data collection policy in (9).

Therefore, by recursion, we have

$$\begin{aligned} & (f_{h,1}^k - \bar{Q}_h^{\pi_k})(s_h^k, a_h^k) \\ & \leq \sum_{h'=h+1}^H \sum_{a' \in A_{h'}^k} \left(\tilde{\mathcal{P}}_{h',2}^k(a'|s_{h'}^k, A_{h'}^k) - \mathcal{P}_{h'}(a'|s_{h'}^k, A_{h'}^k) \right) f_{h',2}^k(s_{h'}^k, a') \\ & \quad + 2 \sum_{h'=h}^H b_{h',1}^k(s_{h'}^k, a_{h'}^k) + 2 \sum_{h'=h+1}^H b_{h',2}^k(s_{h'}^k, a_{h'}^k) + \sum_{h'=h+1}^H \zeta_{h',2}^k + \sum_{h'=h+1}^H \dot{\zeta}_{h',2}^k \end{aligned} \quad (\text{D.46})$$

Finally, we consider the case where $h < h_k - 1$.

$$\begin{aligned} & (f_{h,1}^k - \bar{Q}_h^{\pi_k})(s_h^k, a_h^k) = (f_{h,1}^k - \mathcal{T}_h V_{h+1,1}^k)(s_h^k, a_h^k) + (\mathcal{T}_h V_{h+1,1}^k - \bar{Q}_h^{\pi_k})(s_h^k, a_h^k) \\ & \leq 2b_{h,1}^k(s_h^k, a_h^k) + \mathbb{E}[(V_{h+1,1}^k - V_{h+1}^{\pi_k})(s_{h+1})|s_h^k, a_h^k] \\ & = 2b_{h,1}^k(s_h^k, a_h^k) + (V_{h+1,1}^k - V_{h+1}^{\pi_k})(s_{h+1}) + \zeta_{h+1,1}^k \\ & = 2b_{h,1}^k(s_h^k, a_h^k) + (Q_{h+1,1}^k - Q_{h+1}^{\pi_k})(s_{h+1}, A_{h+1}^k) + \zeta_{h+1,1}^k \\ & = \sum_{a' \in A_{h+1}^k} \left(\tilde{\mathcal{P}}_{h+1,1}^k(a'|s_{h+1}^k, A_{h+1}^k) f_{h+1,1}^k(s_{h+1}^k, a') - \mathcal{P}_{h+1}(a'|s_{h+1}^k, A_{h+1}^k) \bar{Q}_{h+1}^{\pi_k}(s_{h+1}^k, a') \right) \\ & \quad + 2b_{h,1}^k(s_h^k, a_h^k) + \zeta_{h+1,1}^k \\ & = \sum_{a' \in A_{h+1}^k} \left(\tilde{\mathcal{P}}_{h+1,1}^k(a'|s_{h+1}^k, A_{h+1}^k) - \mathcal{P}_{h+1}(a'|s_{h+1}^k, A_{h+1}^k) \right) f_{h+1,1}^k(s_{h+1}^k, a') \\ & \quad + \mathbb{E}_{\mathcal{P}} \left[\left(f_{h+1,1}^k - \bar{Q}_{h+1}^{\pi_k} \right) (s_{h+1}^k, a_{h+1}) | s_{h+1}^k, A_{h+1}^k \right] + 2b_{h,1}^k(s_h^k, a_h^k) + \zeta_{h+1,1}^k \\ & = \sum_{a' \in A_{h+1}^k} \left(\tilde{\mathcal{P}}_{h+1,1}^k(a'|s_{h+1}^k, A_{h+1}^k) - \mathcal{P}_{h+1}(a'|s_{h+1}^k, A_{h+1}^k) \right) f_{h+1,1}^k(s_{h+1}^k, a') \\ & \quad + \left(f_{h+1,1}^k - \bar{Q}_{h+1}^{\pi_k} \right) (s_{h+1}^k, a_{h+1}^k) + 2b_{h,1}^k(s_h^k, a_h^k) + \zeta_{h+1,1}^k + \dot{\zeta}_{h+1,1}^k, \end{aligned}$$

where the first inequality holds based on the assumption that $\mathcal{T}_h V_{h+1,1}^k \in \mathcal{F}_{h,1}^k$ and definition of $b_{h,1}^k$ and the third equality holds because for $h+1 < h_k$, we have $A_{h+1}^k = A_{h+1,1}^k$ by the data collection policy in (9).

Hence, by recursion we have

$$\begin{aligned} & (f_{h,1}^k - \bar{Q}_h^{\pi_k})(s_h^k, a_h^k) \leq \sum_{h'=h+1}^{h_k-1} \sum_{a' \in A_{h'}^k} \left(\tilde{\mathcal{P}}_{h',1}^k(a'|s_{h'}^k, A_{h'}^k) - \mathcal{P}_{h'}(a'|s_{h'}^k, A_{h'}^k) \right) f_{h',1}^k(s_{h'}^k, a') \\ & \quad + \sum_{h'=h_k}^H \sum_{a' \in A_{h'}^k} \left(\tilde{\mathcal{P}}_{h',2}^k(a'|s_{h'}^k, A_{h'}^k) - \mathcal{P}_{h'}(a'|s_{h'}^k, A_{h'}^k) \right) f_{h',2}^k(s_{h'}^k, a') \\ & \quad + 2 \sum_{h'=h}^H b_{h',1}^k(s_{h'}^k, a_{h'}^k) + 2 \sum_{h'=h_k}^H b_{h',2}^k(s_{h'}^k, a_{h'}^k) + \sum_{h'=h+1}^{h_k-1} (\zeta_{h',1}^k + \dot{\zeta}_{h',1}^k) \\ & \quad + \sum_{h'=h_k}^H (\zeta_{h',2}^k + \dot{\zeta}_{h',2}^k). \end{aligned} \quad (\text{D.47})$$

Combining (D.44), (D.46), and (D.47), we conclude the proof. \square

D.7. Bounds on bonuses and $|\mathcal{K}_{oo}|$

In this subsection, we provide proofs for the bounds on the sum of bonuses (Lemma D.21, Lemma D.22, and Lemma D.23) as well as the bound on the size of \mathcal{K}_{oo} (Lemma D.24).

Lemma D.21 (Crude bound on $b_{h,1}^k$, Lemma 39 of Agarwal et al. 2023). *Let $z_h^k = (s_h^k, a_h^k)$. Given $b_{h,1}^k(\cdot) \leq C \cdot \left(D_{\mathcal{F}_h}(\cdot; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}_h^\tau\}_{\tau=1}^{k-1}) \cdot \sqrt{(\beta_{h,1}^k)^2 + \rho} + \epsilon_b \beta_{h,1}^k \right)$, when $\rho = 1, \nu \leq 1$, it holds that for any subset $\mathcal{K} \in [K]$, we have*

$$\begin{aligned} & \sum_{k \in \mathcal{K}} \sum_{h=1}^H \min \{1 + L, b_{h,1}^k(z_h^k)\} \\ &= \mathcal{O} \left(\sqrt{\log \frac{\mathcal{N}KH}{\nu\delta}} \cdot \left(\sqrt{\log \frac{\mathcal{N}\mathcal{N}_bKH}{\nu\delta}} \cdot H\sqrt{d_\nu|\mathcal{K}|} + \log \frac{\mathcal{N}\mathcal{N}_bKH}{\nu\delta} \cdot d_\nu H + |\mathcal{K}|H\epsilon_b \right) \right). \end{aligned}$$

Lemma D.22 (Crude bound on $b_{h,2}^k$, Lemma 38 of Agarwal et al. 2023). *Let $z_h^k = (s_h^k, a_h^k)$. Given $b_{h,2}^k(\cdot) \leq C \cdot \left(D_{\mathcal{F}_h}(\cdot; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\mathbf{1}^\tau\}_{\tau=1}^{k-1}) \cdot \sqrt{(\beta_{h,2}^k)^2 + \rho} + \epsilon_b \beta_{h,1}^k \right)$, when $\rho = 1, \nu \leq 1$, it holds that for any subset $\mathcal{K} \in [K]$, we have*

$$\sum_{k \in \mathcal{K}} \sum_{h=1}^H \min \{1 + L, b_{h,2}^k(z_h^k)\} = \mathcal{O} \left(\sqrt{\log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta}} \cdot \left(H\sqrt{d_\nu|\mathcal{K}|} + d_\nu H + |\mathcal{K}|H\epsilon_b \right) \right).$$

Lemma D.23 (Fine-grained bound on $b_{h,1}^k$). *Let $z_h^k = (s_h^k, a_h^k)$. Recall that the bonus oracle \mathcal{B} outputs a bonus function such that $b_{h,1}^k(\cdot) \leq C \cdot \left(D_{\mathcal{F}_h}(\cdot; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}_h^\tau\}_{\tau=1}^{k-1}) \cdot \sqrt{(\beta_{h,1}^k)^2 + \rho} + \epsilon_b \beta_{h,1}^k \right)$. When $\rho = 1, \nu = 1/\sqrt{KH}, \delta < (0, 1/7)$ and the event $\mathcal{E}_{\leq K}$ holds, with probability at least $1 - 7\delta$, we have*

$$\begin{aligned} & \sum_{k=1}^K \sum_{h=1}^H \min \{1 + L, b_{h,1}^k(z_h^k)\} \\ &= \mathcal{O} \left(\sqrt{d_\nu HK} \cdot \log \frac{\mathcal{N}KH}{\delta} + \frac{1}{\sqrt{\kappa}} dH^{7/2} \sqrt{d_\nu} (\log K)^{3/2} \log M \cdot \log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta} \cdot \sqrt{\log \frac{\mathcal{N}KH}{\delta}} \right) \\ &+ \mathcal{O} \left(d_\nu H^{7/2} \log \frac{\mathcal{N}KH}{\delta} \cdot \left(\log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta} \right)^{3/2} + \sqrt{\log \frac{\mathcal{N}KH}{\delta}} \cdot \left(KH\epsilon_b + \sqrt{d_\nu KH^3\delta} \right) \right) \\ &+ \mathcal{O} \left(\sqrt{\log \frac{\mathcal{N}KH}{\delta}} \log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta} \cdot \sqrt{d_\nu H} \cdot \left(\sqrt{H^2 \sum_{k \in \mathcal{K}_o} u_k} + \sqrt{H^2 |\mathcal{K}_{oo}|} \right) \right). \end{aligned}$$

Proof of Lemma D.23. By the definition of the oracle \mathcal{B} (Definition B.1), we have

$$\begin{aligned} & \sum_{k=1}^K \sum_{h=1}^H \min \{1 + L, b_{h,1}^k(z_h^k)\} \\ &= \mathcal{O} \left(\sum_{k=1}^K \sum_{h=1}^H \min \left\{ 1, D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}_h^\tau\}_{\tau=1}^{k-1}) \cdot \sqrt{(\beta_{h,1}^k)^2 + \rho} \right\} + KH\epsilon_b \cdot \max_{k,h} \beta_{h,1}^k \right) \\ &= \mathcal{O} \left(\sqrt{\log \frac{\mathcal{N}KH}{\delta}} \cdot \sum_{k=1}^K \sum_{h=1}^H \min \{1, D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}_h^\tau\}_{\tau=1}^{k-1})\} + KH\epsilon_b \right), \end{aligned} \tag{D.48}$$

where the last equality holds by the definition of $\beta_{h,1}^k$.

Now, we bound the summation terms

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \min \{1, D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}\}_{\tau=1}^{k-1})\} \\ = \sum_{k=1}^K \sum_{h=1}^H \min \left\{1, \bar{\sigma}_h^k \cdot (\bar{\sigma}_h^k)^{-1} D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}\}_{\tau=1}^{k-1})\right\} \end{aligned}$$

by dividing into the following cases:

$$\begin{aligned} \mathcal{I}_1 &= \{(k, h) \in [K] \times [H] : (\bar{\sigma}_h^k)^{-1} D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}\}_{\tau=1}^{k-1}) \geq 1\}, \\ \mathcal{I}_2 &= \{(k, h) \in [K] \times [H] : \bar{\sigma}_h^k = \nu, (k, h) \neq \mathcal{I}_1\}, \\ \mathcal{I}_3 &= \{(k, h) \in [K] \times [H] : \bar{\sigma}_h^k = 2 \left(\sqrt{o(\delta_h^k)} + \iota(\delta_h^k) \right) \cdot \sqrt{D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}\}_{\tau=1}^{k-1})}, \\ &\quad (k, h) \neq \mathcal{I}_1\}, \\ \mathcal{I}_4 &= \{(k, h) \in [K] \times [H] : \bar{\sigma}_h^k = \sqrt{2}\iota(\delta_h^k) \sqrt{f_{h,2}^k(z_h^k) - f_{h,-2}^k(z_h^k)}, (k, h) \neq \mathcal{I}_1\} \\ \mathcal{I}_5 &= \{(k, h) \in [K] \times [H] : \bar{\sigma}_h^k = \sigma_h^k, (k, h) \neq \mathcal{I}_1\}, . \end{aligned}$$

For the case of \mathcal{I}_1 , we have

$$\begin{aligned} \sum_{(k,h) \in \mathcal{I}_1} \min \left\{1, \bar{\sigma}_h^k \cdot (\bar{\sigma}_h^k)^{-1} D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}\}_{\tau=1}^{k-1})\right\} \\ \leq \sum_{(k,h) \in \mathcal{I}_1} (\bar{\sigma}_h^k)^{-1} D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}\}_{\tau=1}^{k-1}) \leq \sum_{h=1}^H \dim_{\nu,K}(\mathcal{F}_h) = d_\nu H. \end{aligned} \quad (\text{D.49})$$

For \mathcal{I}_2 , we use the Cauchy-Schwarz inequality to get

$$\begin{aligned} \sum_{(k,h) \in \mathcal{I}_2} \min \left\{1, \bar{\sigma}_h^k \cdot (\bar{\sigma}_h^k)^{-1} D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}\}_{\tau=1}^{k-1})\right\} \\ \leq \sqrt{\nu^2 K H} \cdot \sqrt{\sum_{(k,h) \in \mathcal{I}_2} (\bar{\sigma}_h^k)^{-2} D_{\mathcal{F}_h}^2(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}\}_{\tau=1}^{k-1})} \leq \sqrt{\sum_{h=1}^H \dim_{\nu,K}(\mathcal{F}_h)} = \sqrt{d_\nu H}. \end{aligned} \quad (\text{D.50})$$

For \mathcal{I}_3 , we have

$$\begin{aligned} \sum_{(k,h) \in \mathcal{I}_3} \min \left\{1, \bar{\sigma}_h^k \cdot (\bar{\sigma}_h^k)^{-1} D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}\}_{\tau=1}^{k-1})\right\} \\ \leq \sum_{(k,h) \in \mathcal{I}_3} (8o(\delta_h^k) + \iota^2(\delta_h^k)) \cdot \min \left\{1, (\bar{\sigma}_h^k)^{-2} D_{\mathcal{F}_h}^2(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}\}_{\tau=1}^{k-1})\right\} \\ = \mathcal{O} \left(\left(\sqrt{\log \frac{\mathcal{N}KH}{\delta}} + \log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta} \right) \cdot \sum_{h=1}^H \dim_{\nu,K}(\mathcal{F}_h) \right) \\ = \mathcal{O} \left(\left(\sqrt{\log \frac{\mathcal{N}KH}{\delta}} + \log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta} \right) d_\nu H \right), \end{aligned} \quad (\text{D.51})$$

where the inequality holds because, by dividing both sides of $\bar{\sigma}_h^k = 2 \left(\sqrt{o(\delta_h^k)} + \iota(\delta_h^k) \right) \cdot \sqrt{D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}\}_{\tau=1}^{k-1})}$ by $\sqrt{\bar{\sigma}_h^k}$ and rearranging terms, we get:

$$\bar{\sigma}_h^k \leq (8o(\delta_h^k) + \iota^2(\delta_h^k)) D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}\}_{\tau=1}^{k-1}).$$

We also use the property that $(\bar{\sigma}_h^k)^{-1} D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}\}_{\tau=1}^{k-1}) \leq 1$ for $(k, h) \in \mathcal{I}_3$, which follows directly from the definition of \mathcal{I}_3 .

For \mathcal{I}_4 , we have

$$\begin{aligned}
 & \sum_{(k,h) \in \mathcal{I}_4} \min \left\{ 1, \bar{\sigma}_h^k \cdot (\bar{\sigma}_h^k)^{-1} D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}\}_{\tau=1}^{k-1}) \right\} \\
 & \leq \sum_{(k,h) \in \mathcal{I}_4} \bar{\sigma}_h^k \cdot (\bar{\sigma}_h^k)^{-1} D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}\}_{\tau=1}^{k-1}) \\
 & = \sum_{(k,h) \in \mathcal{I}_4} \sqrt{2} \iota(\delta_h^k) \sqrt{f_{h,2}^k(z_h^k) - f_{h,-2}^k(z_h^k)} \cdot (\bar{\sigma}_h^k)^{-1} D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}\}_{\tau=1}^{k-1}) \\
 & \leq \mathcal{O} \left(\sqrt{\log \frac{\mathcal{N} \mathcal{N}_b K H}{\delta}} \sqrt{\sum_{k=1}^K \sum_{h=1}^H f_{h,2}^k(z_h^k) - f_{h,-2}^k(z_h^k)} \cdot \sqrt{d_\nu H} \right), \tag{D.52}
 \end{aligned}$$

where the last inequality holds by the Cauchy-Schwarz inequality together with the definition of $\iota(\delta_h^k)$.

Lastly, restricting on \mathcal{I}_5 , if the event $\mathcal{E}_{\leq K}$ holds, we have

$$\begin{aligned}
 & \sum_{(k,h) \in \mathcal{I}_5} \min \left\{ 1, \bar{\sigma}_h^k \cdot (\bar{\sigma}_h^k)^{-1} D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}\}_{\tau=1}^{k-1}) \right\} \\
 & \leq \sum_{(k,h) \in \mathcal{I}_5} \bar{\sigma}_h^k \cdot (\bar{\sigma}_h^k)^{-1} D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}\}_{\tau=1}^{k-1}) \\
 & \leq \sqrt{\sum_{(k,h) \in \mathcal{I}_5} (\bar{\sigma}_h^k)^2} \cdot \sqrt{\sum_{(k,h) \in \mathcal{I}_5} (\bar{\sigma}_h^k)^{-2} D_{\mathcal{F}_h}^2(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}\}_{\tau=1}^{k-1})} \\
 & \leq \mathcal{O} \left(\sqrt{\sum_{k,h} \mathbb{V}[r_h + V_{h+1,1}^k(s_{h+1}) \mid z_h^k]} + \sum_{k,h} (f_{h,2}^k(z_h^k) - f_{h,-2}^k(z_h^k)) \cdot \sqrt{d_\nu H} \right) \\
 & + \mathcal{O} \left(\sqrt{\sum_{k,h} \min \{1, D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\mathbf{1}^\tau\}_{\tau=1}^{k-1})\}} \sqrt{\log \frac{\mathcal{N} \mathcal{N}_b K H}{\delta}} \cdot \sqrt{d_\nu H} \right), \tag{D.53}
 \end{aligned}$$

where the second inequality holds by the Cauchy-Schwarz inequality, the last inequality holds by Lemma D.16 and the definition of Eluder dimension.

To further bound the first term on the right-hand side of (D.53), we apply Lemma D.17. Therefore, with probability at least $1 - \delta$, we have

$$\begin{aligned}
 & \sum_{(k,h) \in \mathcal{I}_5} \min \left\{ 1, \bar{\sigma}_h^k \cdot (\bar{\sigma}_h^k)^{-1} D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}\}_{\tau=1}^{k-1}) \right\} \\
 & \leq \mathcal{O} \left(\sqrt{H \sum_{k,h} (f_{h,2}^k(z_h^k) - f_{h,-2}^k(z_h^k)) + K + K H^2 \delta + H^2 |\mathcal{K}_{00}| + H^4 \log^2 \frac{K H}{\delta}} \cdot \sqrt{d_\nu H} \right) \\
 & + \mathcal{O} \left(\sqrt{\sum_{k,h} (f_{h,2}^k(z_h^k) - f_{h,-2}^k(z_h^k)) \cdot \sqrt{d_\nu H}} \right) + \mathcal{O} \left(\sqrt{(d_\nu H + H \sqrt{d_\nu K}) \log \frac{\mathcal{N} \mathcal{N}_b K H}{\delta}} \cdot \sqrt{d_\nu H} \right) \\
 & \leq \mathcal{O} \left(\sqrt{K + K H^2 \delta + H^2 |\mathcal{K}_{00}| + H^4 \log^2 \frac{K H}{\delta}} \cdot \sqrt{d_\nu H} + \sqrt{H \sum_{k,h} (f_{h,2}^k(z_h^k) - f_{h,-2}^k(z_h^k)) \cdot \sqrt{d_\nu H}} \right) \\
 & + \mathcal{O} \left(d_\nu H^{1.5} \sqrt{\log \frac{\mathcal{N} \mathcal{N}_b K H}{\delta}} \right), \tag{D.54}
 \end{aligned}$$

where the first inequality holds by the fact that

$$\begin{aligned} \sum_{k,h} \min \{1, D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\mathbf{1}^\tau\}_{\tau=1}^{k-1})\} &\leq d_\nu H \\ \sum_{k,h} \min \{1, D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\mathbf{1}^\tau\}_{\tau=1}^{k-1})\} &\leq \sqrt{KH} \sqrt{\sum_{k,h} D_{\mathcal{F}_h}^2(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\mathbf{1}^\tau\}_{\tau=1}^{k-1})} \leq H\sqrt{d_\nu K}, \end{aligned}$$

and the last inequality holds by the AM-GM inequality such that

$$H\sqrt{K \cdot d_\nu \log \frac{\mathcal{N}\mathcal{N}_b KH}{\delta}} \leq K + H^2 d_\nu \log \frac{\mathcal{N}\mathcal{N}_b KH}{\delta}.$$

Combining Equation D.49, D.50, D.51, D.52, and D.54, we get

$$\begin{aligned} &\sum_{k=1}^K \sum_{h=1}^H \min \{1, D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}\}_{\tau=1}^{k-1})\} \\ &\leq \mathcal{O} \left(\sqrt{\log \frac{\mathcal{N}\mathcal{N}_b KH}{\delta}} \cdot \sqrt{d_\nu H} \cdot \sqrt{H \sum_{k=1}^K \sum_{h=1}^H f_{h,2}^k(z_h^k) - f_{h,-2}^k(z_h^k)} \right) \\ &\quad + \mathcal{O} \left(\sqrt{K + KH^2\delta + H^2|\mathcal{K}_{00}| + H^4 \log^2 \frac{KH}{\delta}} \cdot \sqrt{d_\nu H} + d_\nu H^{1.5} \sqrt{\log \frac{\mathcal{N}\mathcal{N}_b KH}{\delta}} \right). \end{aligned} \quad (\text{D.55})$$

Now, we bound the term $\sum_{k,h} (f_{h,2}^k(z_h^k) - f_{h,-2}^k(z_h^k))$. For $k \in \mathcal{K}_{00}$, we have

$$\sum_{k \in \mathcal{K}_{00}} \sum_{h=1}^H (f_{h,2}^k(z_h^k) - f_{h,-2}^k(z_h^k)) = \mathcal{O}(|\mathcal{K}_{00}|H).$$

Otherwise, for episodes $k \in \mathcal{K}_0$, we know that it holds true that $f_{h,1}^k(z_h^k) \geq f_{h,2}^k(z_h^k) - u_k$ by (9). Therefore, under the event $\mathcal{E}_{\leq K}$, we have

$$\begin{aligned} &\sum_{k \in \mathcal{K}_0} \sum_{h=1}^H (f_{h,2}^k(z_h^k) - f_{h,-2}^k(z_h^k)) \leq \sum_{k \in \mathcal{K}_0} \sum_{h=1}^H (f_{h,1}^k(z_h^k) - f_{h,-2}^k(z_h^k)) + H \sum_{k \in \mathcal{K}_0} u_k \\ &= \sum_{k \in \mathcal{K}_0} \sum_{h=1}^H \left((f_{h,1}^k - \bar{Q}_h^{\pi_k})(z_h^k) + (\bar{Q}_h^{\pi_k} - f_{h,-2}^k)(z_h^k) \right) + H \sum_{k \in \mathcal{K}_0} u_k \\ &\leq \sum_{k \in \mathcal{K}_0} \sum_{h=1}^H \sum_{h'=h+1}^H \sum_{a' \in A_{h'}^k} \left(\tilde{\mathcal{P}}_{h',1}^k(a'|s_{h'}^k, A_{h'}^k) - \mathcal{P}_{h'}(a'|s_{h'}^k, A_{h'}^k) \right) f_{h',1}^k(s_{h'}^k, a') \\ &\quad - \sum_{k \in \mathcal{K}_0} \sum_{h=1}^H \sum_{h'=h+1}^H \sum_{a' \in A_{h'}^k} \left(\tilde{\mathcal{P}}_{h',-2}^k(a'|s_{h'}^k, A_{h'}^k) - \mathcal{P}_{h'}(a'|s_{h'}^k, A_{h'}^k) \right) f_{h',-2}^k(s_{h'}^k, a') \\ &\quad + 2 \sum_{k \in \mathcal{K}_0} \sum_{h=1}^H \min \left\{ 1 + L, \sum_{h'=h}^H b_{h',1}^k(s_{h'}^k, a_{h'}^k) \right\} + 2 \sum_{k \in \mathcal{K}_0} \sum_{h=1}^H \min \left\{ 1 + L, \sum_{h'=h}^H b_{h',2}^k(s_{h'}^k, a_{h'}^k) \right\} \\ &\quad + \underbrace{\sum_{k \in \mathcal{K}_0} \sum_{h=1}^H \sum_{h'=h+1}^H \left(\zeta_{h',1}^k + \dot{\zeta}_{h',1}^k - \zeta_{h',-2}^k - \dot{\zeta}_{h',-2}^k \right)}_{\text{martingale difference sequences (MDSs)}} + H \sum_{k \in \mathcal{K}_0} u_k, \end{aligned} \quad (\text{D.56})$$

where the second inequality holds by Lemma D.18 and D.20.

To further the right-hand side of (D.56), we apply Lemma D.10 (which holds with probability at least $1 - 2\delta$) to the first and the second terms, Lemma D.21 to the third term, Lemma D.22 to the forth term, and we bound the fifth term using the

Azuma-Hoeffding inequality (which holds with probability at least $1 - 4\delta$). As a result, absorbing the low-order terms, we obtain that

$$\begin{aligned}
 & \sum_{k=1}^K \sum_{h=1}^H (f_{h,2}^k(z_h^k) - f_{h,-2}^k(z_h^k)) \\
 & \leq \mathcal{O} \left(\frac{1}{\sqrt{\kappa}} dH^2 \sqrt{K} \cdot (\log K)^{3/2} \log M \right) \\
 & + \mathcal{O} \left(\sqrt{\log \frac{\mathcal{N}KH}{\delta}} \cdot \left(\sqrt{\log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta}} \cdot H^2 \sqrt{d_\nu K} + \log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta} \cdot d_\nu H^2 \right) \right) \\
 & + \mathcal{O} \left(\sqrt{\log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta}} \cdot KH^2 \epsilon_b + H \sqrt{KH \log \frac{KH}{\delta}} + |\mathcal{K}_{\text{oo}}|H + H \sum_{k \in \mathcal{K}_o} u_k \right), \tag{D.57}
 \end{aligned}$$

where the last inequality holds by the AM-GM inequality.

Plugging (D.57) to (D.55), we have

$$\begin{aligned}
 & \sum_{k=1}^K \sum_{h=1}^H \min \left\{ 1, \bar{\sigma}_h^k \cdot (\bar{\sigma}_h^k)^{-1} D_{\mathcal{F}_h}(z_h^k; \{z_h^\tau\}_{\tau=1}^{k-1}, \{\bar{\sigma}\}_{\tau=1}^{k-1}) \right\} \\
 & \leq \mathcal{O} \left(\sqrt{\log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta}} \cdot \sqrt{d_\nu H} \cdot \left(\sqrt{\frac{1}{\sqrt{\kappa}} dH^3 (\log K)^{3/2} \log M \cdot \sqrt{K}} + H^2 \sum_{k \in \mathcal{K}_o} u_k \right) \right) \\
 & + \mathcal{O} \left(\left(\log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta} \right)^{3/4} \cdot \sqrt{d_\nu H} \cdot \sqrt{\log \frac{\mathcal{N}KH}{\delta}} \cdot \left(H^3 \sqrt{d_\nu K} + \sqrt{\log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta}} \cdot d_\nu H^3 \right) \right) \\
 & + \mathcal{O} \left(\sqrt{\log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta}} \cdot \sqrt{d_\nu H} \cdot \sqrt{\log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta}} \cdot KH^3 \epsilon_b + H^2 \sqrt{KH \log \frac{KH}{\delta}} + |\mathcal{K}_{\text{oo}}|H^2 \right) \\
 & + \mathcal{O} \left(\sqrt{K + KH^2\delta + H^4 \log^2 \frac{KH}{\delta}} \cdot \sqrt{d_\nu H} + d_\nu H^{1.5} \sqrt{\log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta}} \right) \\
 & \leq \mathcal{O} \left(\sqrt{\log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta}} \cdot \sqrt{d_\nu H} \cdot \sqrt{\frac{K}{\log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta}} + \frac{1}{\kappa} d^2 H^6 ((\log K)^{3/2} \log M)^2 \cdot \log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta}} \right) \\
 & + \mathcal{O} \left(\sqrt{\log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta}} \cdot \sqrt{d_\nu H} \cdot \sqrt{d_\nu H^6 \log \frac{\mathcal{N}KH}{\delta} \left(\log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta} \right)^2 + H^2 \sum_{k \in \mathcal{K}_o} u_k} \right) \\
 & + \mathcal{O} \left(\sqrt{\log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta}} \cdot \sqrt{d_\nu H} \cdot \sqrt{H^2 |\mathcal{K}_{\text{oo}}|} + KH \epsilon_b + \sqrt{d_\nu H} \cdot \sqrt{KH^2 \delta} \right) \\
 & = \mathcal{O} \left(\sqrt{d_\nu HK} + \frac{1}{\sqrt{\kappa}} dH^{7/2} \sqrt{d_\nu} (\log K)^{3/2} \log M \cdot \log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta} \right) \\
 & + \mathcal{O} \left(d_\nu H^{7/2} \sqrt{\log \frac{\mathcal{N}KH}{\delta}} \cdot \left(\log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta} \right)^{3/2} \right) \\
 & + \mathcal{O} \left(\sqrt{\log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta}} \cdot \sqrt{d_\nu H} \cdot \left(\sqrt{H^2 \sum_{k \in \mathcal{K}_o} u_k} + \sqrt{H^2 |\mathcal{K}_{\text{oo}}|} \right) + KH \epsilon_b + \sqrt{d_\nu KH^3 \delta} \right), \tag{D.58}
 \end{aligned}$$

where the second inequality holds by applying the AM-GM inequality and absorbing the lower-order terms.

Finally, plugging (D.58) to (D.48), we derive that

$$\begin{aligned}
 & \sum_{k=1}^K \sum_{h=1}^H \min \{1 + L, b_{h,1}^k(z_h^k)\} \\
 & \leq \mathcal{O} \left(\sqrt{d_\nu H K} \cdot \log \frac{NKH}{\delta} + \frac{1}{\sqrt{\kappa}} dH^{7/2} \sqrt{d_\nu} (\log K)^{3/2} \log M \cdot \log \frac{N\mathcal{N}_b KH}{\delta} \cdot \sqrt{\log \frac{NKH}{\delta}} \right) \\
 & + \mathcal{O} \left(d_\nu H^{7/2} \log \frac{NKH}{\delta} \cdot \left(\log \frac{N\mathcal{N}_b KH}{\delta} \right)^{3/2} + \sqrt{\log \frac{NKH}{\delta}} \cdot (KH\epsilon_b + \sqrt{d_\nu KH^3 \delta}) \right) \\
 & + \mathcal{O} \left(\sqrt{\log \frac{NKH}{\delta} \log \frac{N\mathcal{N}_b KH}{\delta}} \cdot \sqrt{d_\nu H} \cdot \left(\sqrt{H^2 \sum_{k \in \mathcal{K}_o} u_k} + \sqrt{H^2 |\mathcal{K}_{oo}|} \right) \right).
 \end{aligned}$$

This concludes the proof of Lemma D.23. \square

Lemma D.24 (Bounding size of \mathcal{K}_{oo}). *Suppose $\nu \leq 1$ and we set*

$$u_k \geq C \cdot \left(\frac{\sqrt{\log \frac{NKH}{\nu\delta}} \cdot \left(\log \frac{N\mathcal{N}_b KH}{\nu\delta} \cdot H^{5/2} \sqrt{d_\nu} + \sqrt{k} H \epsilon_b \right)}{\sqrt{k}} + \frac{dH^{5/2} (\log K)^{3/2} \log M \sqrt{\log \frac{N\mathcal{N}_b KH}{\nu\delta}}}{\sqrt{k}} \right),$$

for some large enough constant $0 < C < \infty$, when the event $\mathcal{E}_{\leq K}$ holds true, then with probability at least $1 - 2\delta$, it holds that

$$|\mathcal{K}_{oo}| \leq \mathcal{O} \left(\frac{K}{H^3 \log \frac{N\mathcal{N}_b KH}{\nu\delta}} + \frac{d_\nu}{H^3} + \frac{d^4 ((\log K)^{3/2} \log M)^4}{\kappa^2 d_\nu H^3 \cdot \log \frac{NKH}{\nu\delta} (\log \frac{N\mathcal{N}_b KH}{\nu\delta})^2} \right).$$

Proof of Lemma D.24. By the definition of h_k , for each $k \in \mathcal{K}_{oo}$, we have $f_{h_k,2}^k(s_{h_k}^k, a_{h_k}^k) \geq f_{h_k,1}^k(s_{h_k}^k, a_{h_k}^k) + u_k$, which implies that

$$\begin{aligned}
 & \sum_{k \in \mathcal{K}_{oo}} (f_{h_k,2}^k - f_{h_k,1}^k)(s_{h_k}^k, a_{h_k}^k) \geq \sum_{k \in \mathcal{K}_{oo}} u_k \\
 & \geq C \cdot \left(\sqrt{\log \frac{NKH}{\nu\delta}} \cdot \left(\log \frac{N\mathcal{N}_b KH}{\nu\delta} \cdot H^{5/2} \sqrt{d_\nu} \cdot \frac{|\mathcal{K}_{oo}|}{\sqrt{K}} + |\mathcal{K}_{oo}| H \epsilon_b \right) \right. \\
 & \quad \left. + dH^{5/2} (\log K)^{3/2} \log M \sqrt{\log \frac{N\mathcal{N}_b KH}{\nu\delta}} \cdot \frac{|\mathcal{K}_{oo}|}{\sqrt{K}} \right). \tag{D.59}
 \end{aligned}$$

Furthermore, under the event $\mathcal{E}_{\leq K}$, by Lemma D.14, it holds that $f_{h,1}^k(s_h, a_h) \geq \bar{Q}_h^*(s_h, a_h) \geq \bar{Q}_h^{\pi_k}(s_h, a_h)$ for all $(s_h, a_h) \in \mathcal{S} \times \mathcal{I}$. Thus, we get

$$\begin{aligned}
 & \sum_{k \in \mathcal{K}_{oo}} (f_{h_k,2}^k - f_{h_k,1}^k)(s_{h_k}^k, a_{h_k}^k) \leq \sum_{k \in \mathcal{K}_{oo}} (f_{h_k,2}^k - \bar{Q}_{h_k}^{\pi_k})(s_{h_k}^k, a_{h_k}^k) \\
 & \leq \sum_{k \in \mathcal{K}_{oo}} \sum_{h'=h_k+1}^H \sum_{a' \in A_{h',2}^k} \left(\tilde{\mathcal{P}}_{h',2}^k(a'|s_{h'}^k, A_{h',2}^k) - \mathcal{P}_{h'}(a'|s_{h'}^k, A_{h',2}^k) \right) f_{h',2}^k(s_{h'}^k, a') \\
 & + 2 \sum_{k \in \mathcal{K}_{oo}} \sum_{h'=h_k}^H \min \{1 + L, b_{h',1}^k(s_{h'}^k, a_{h'}^k)\} + 2 \sum_{k \in \mathcal{K}_{oo}} \sum_{h'=h_k}^H \min \{1 + L, b_{h',2}^k(s_{h'}^k, a_{h'}^k)\} + \sum_{k \in \mathcal{K}_{oo}} \sum_{h'=h_k+1}^H (\zeta_{h',2}^k + \dot{\zeta}_{h',2}^k) \\
 & \leq \mathcal{O} \left(dH \sqrt{|\mathcal{K}_{oo}|} \cdot (\log K)^{3/2} \log M + \frac{1}{\kappa} d^2 H \left((\log K)^{3/2} \log M \right)^2 + \sqrt{|\mathcal{K}_{oo}| H \log \frac{KH}{\delta}} \right) \\
 & + \mathcal{O} \left(\sqrt{\log \frac{NKH}{\nu\delta}} \cdot \left(\sqrt{\log \frac{N\mathcal{N}_b KH}{\nu\delta}} \cdot H \sqrt{d_\nu |\mathcal{K}_{oo}|} + \log \frac{N\mathcal{N}_b KH}{\nu\delta} \cdot d_\nu H + |\mathcal{K}_{oo}| H \epsilon_b \right) \right), \tag{D.60}
 \end{aligned}$$

where in the first inequality, we note that $A_{h'}^k = A_{h',2}^k$ for $h' \geq h_k$, the second inequality follows from D.19. And for the third inequality, we apply Lemma D.13 to the first term, Lemma D.21 to the second term, Lemma D.22 to the third term. Finally, we bound the last term using the Azuma-Hoeffding inequality with probability at least $1 - 2\delta$.

Thus, in order for the two inequalities D.59 and D.60 to hold simultaneously, the following condition must be satisfied:

$$|\mathcal{K}_{\text{oo}}| \leq \mathcal{O} \left(\max \left\{ \frac{K}{H^3 \log \frac{\mathcal{N}\mathcal{N}_bKH}{\nu\delta}}, \frac{\left(d_\nu \sqrt{\log \frac{\mathcal{N}KH}{\nu\delta}} \log \frac{\mathcal{N}\mathcal{N}_bKH}{\nu\delta} + \frac{1}{\kappa} d^2 ((\log K)^{3/2} \log M)^2 \right) \cdot \sqrt{K}}{H^{3/2} \sqrt{\log \frac{\mathcal{N}\mathcal{N}_bKH}{\nu\delta}} \left(\sqrt{d_\nu \log \frac{\mathcal{N}KH}{\nu\delta}} \log \frac{\mathcal{N}\mathcal{N}_bKH}{\nu\delta} + d(\log K)^{3/2} \log M \right)} \right\} \right).$$

Using the AM-GM inequality, we can further bound the second term inside the max operation.

$$\begin{aligned} & \mathcal{O} \left(\frac{\left(d_\nu \sqrt{\log \frac{\mathcal{N}KH}{\nu\delta}} \log \frac{\mathcal{N}\mathcal{N}_bKH}{\nu\delta} + \frac{1}{\kappa} d^2 ((\log K)^{3/2} \log M)^2 \right) \cdot \sqrt{K}}{H^{3/2} \sqrt{\log \frac{\mathcal{N}\mathcal{N}_bKH}{\nu\delta}} \left(\sqrt{d_\nu \log \frac{\mathcal{N}KH}{\nu\delta}} \log \frac{\mathcal{N}\mathcal{N}_bKH}{\nu\delta} + d(\log K)^{3/2} \log M \right)} \right) \\ & \leq \mathcal{O} \left(\frac{K}{H^3 \log \frac{\mathcal{N}\mathcal{N}_bKH}{\nu\delta}} + \left(\frac{d_\nu \sqrt{\log \frac{\mathcal{N}KH}{\nu\delta}} \log \frac{\mathcal{N}\mathcal{N}_bKH}{\nu\delta} + \frac{1}{\kappa} d^2 ((\log K)^{3/2} \log M)^2}{H^{3/2} \sqrt{d_\nu \log \frac{\mathcal{N}KH}{\nu\delta}} \log \frac{\mathcal{N}\mathcal{N}_bKH}{\nu\delta}} \right)^2 \right) \\ & \leq \mathcal{O} \left(\frac{K}{H^3 \log \frac{\mathcal{N}\mathcal{N}_bKH}{\nu\delta}} + \left(\sqrt{\frac{d_\nu}{H^3}} + \frac{d^2 ((\log K)^{3/2} \log M)^2}{\kappa \cdot H^{3/2} \sqrt{d_\nu \log \frac{\mathcal{N}KH}{\nu\delta}} \log \frac{\mathcal{N}\mathcal{N}_bKH}{\nu\delta}} \right)^2 \right) \\ & \leq \mathcal{O} \left(\frac{K}{H^3 \log \frac{\mathcal{N}\mathcal{N}_bKH}{\nu\delta}} + \frac{d_\nu}{H^3} + \frac{d^4 ((\log K)^{3/2} \log M)^4}{\kappa^2 d_\nu H^3 \cdot \log \frac{\mathcal{N}KH}{\nu\delta} (\log \frac{\mathcal{N}\mathcal{N}_bKH}{\nu\delta})^2} \right), \end{aligned}$$

where the second inequality also follow from the AM-GM inequality and the last inequality holds due to the fact that $(a + b)^2 \leq 2a^2 + 2b^2$ for any $a, b \in \mathbb{R}^+$. This concludes the proof. \square

D.8. Proof of Theorem 5.1

Now, we are ready to provide the proof of Theorem 5.1. To start, we formally restate the theorem.

Theorem D.25 (Restatement of Theorem 5.1, Regret upper bound of MNL-VQL). *Suppose Assumptions 3.1 and 3.3 hold. We assume that we have the generalized Eluder dimension $\dim_{\nu,K}(\mathcal{F}_h)$, for $h \in [H]$, as defined in Definition 3.5 with $\rho = 1$, and access to a consistent bonus oracle \mathcal{B} satisfying Definition B.1 with $\epsilon_b = \mathcal{O}(1/KH)$. Let $d_\nu = \frac{1}{H} \sum_{h=1}^H \dim_{\nu,K}(\mathcal{F}_h)$ with $\nu = \sqrt{1/KH}$, and set $u_k = \mathcal{O}(\sqrt{\log \mathcal{N}} \cdot (\log \mathcal{N}\mathcal{N}_b \cdot H^{5/2} \sqrt{d_\nu} + dH^{5/2} \sqrt{\log \mathcal{N}\mathcal{N}_b})/\sqrt{K})$. Then, for any $\delta < 1/(H^2 + 15)$, with probability at least $1 - \delta$, the regret of MNL-VQL is upper-bounded by:*

$$\begin{aligned} \text{Regret}(\mathcal{M}, K) &= \mathcal{O} \left(d\sqrt{HK}(\log K)^{3/2} \log M + \sqrt{d_\nu HK \cdot \log \frac{\mathcal{N}KH}{\delta}} + \frac{1}{\kappa} d^2 H^2 (\log K)^3 (\log M)^2 \right) \\ &\quad + \mathcal{O} \left(d_\nu H^5 \log \frac{\mathcal{N}KH}{\delta} \cdot \left(\log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta} \right)^2 + \sqrt{\log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta}} \cdot KH\epsilon_b \right), \end{aligned}$$

where d is the feature dimension of the MNL preference model, \mathcal{N} is the maximum size of the function class, i.e., $\mathcal{N} = \max_{h \in [H]} |\mathcal{F}_h|$, and \mathcal{N}_b is the size of the bonus function class, i.e., $\mathcal{N}_b = |\mathcal{W}|$.

Proof of Theorem 5.1. When the event $\mathcal{E}^\theta \cap \mathcal{E}_{\leq K}$ happens (with probability at least $1 - 2\delta$), we can bound the regret as

follows:

$$\begin{aligned}
 \mathbf{Regret}(\mathcal{M}, K) &= \sum_{k=1}^K (V_1^* - V_1^{\pi_k})(s_1^k) \leq \sum_{k=1}^K (V_1^k - V_1^{\pi_k})(s_1^k) = \sum_{k=1}^K (Q_1^k - Q_1^{\pi_k})(s_1^k, A_1^k) \\
 &\leq \mathcal{O}(1) + \sum_{k=2}^K (Q_1^k - Q_1^{\pi_k})(s_1^k, A_1^k) \\
 &= \mathcal{O}(1) + \sum_{k \in \mathcal{K}_o \setminus \{1\}} (Q_1^k - Q_1^{\pi_k})(s_1^k, A_1^k) + \sum_{k \in \mathcal{K}_{oo} \setminus \{1\}} (Q_1^k - Q_1^{\pi_k})(s_1^k, A_1^k), \tag{D.61}
 \end{aligned}$$

where the first inequality holds by Lemma D.15.

For $k \in \mathcal{K}_o \setminus \{1\}$, recall that $Q_h^k(s, A_h^k) = \sum_{a_1 \in A_h^k} \tilde{\mathcal{P}}_{h,1}^k(a_1|s, A_h^k) f_{h,1}^k(s_h^k, a_1)$ for all $h \in [H]$, as defined in (D.18). Therefore, we have

$$\begin{aligned}
 (Q_1^k - Q_1^{\pi_k})(s_1^k, A_1^k) &= \sum_{a_1 \in A_1^k} \tilde{\mathcal{P}}_{1,1}^k(a_1|s_1^k, A_1^k) f_{1,1}^k(s_h^k, a_1) - \sum_{a_1 \in A_1^k} \mathcal{P}_1(a_1|s_1^k, A_1^k) \bar{Q}_1^{\pi_k}(s_1^k, a_1) \\
 &= \sum_{a_1 \in A_1^k} \left(\tilde{\mathcal{P}}_{1,1}^k(a_1|s_1^k, A_1^k) - \mathcal{P}_1(a_1|s_1^k, A_1^k) \right) f_{1,1}^k(s_h^k, a_1) \\
 &\quad + \sum_{a_1 \in A_1^k} \mathcal{P}_1(a_1|s_1^k, A_1^k) \left(f_{1,1}^k - \bar{Q}_1^{\pi_k} \right)(s_1^k, a_1) \\
 &= \sum_{a_1 \in A_1^k} \left(\tilde{\mathcal{P}}_{1,1}^k(a_1|s_1^k, A_1^k) - \mathcal{P}_1(a_1|s_1^k, A_1^k) \right) f_{1,1}^k(s_h^k, a_1) + \left(f_{1,1}^k - \bar{Q}_1^{\pi_k} \right)(s_1^k, a_1) \\
 &\quad + \mathbb{E}_{\mathcal{P}} \left[\left(f_{1,1}^k - \bar{Q}_1^{\pi_k} \right)(s_1^k, a_1) \mid s_1^k, A_1^k \right] - \left(f_{1,1}^k - \bar{Q}_1^{\pi_k} \right)(s_1^k, a).
 \end{aligned}$$

Then, by applying Lemma D.20 with $h_k = H + 1$, we have

$$\begin{aligned}
 (Q_1^k - Q_1^{\pi_k})(s_1^k, A_1^k) &\leq \sum_{h=1}^H \sum_{a_h \in A_h^k} \left(\tilde{\mathcal{P}}_{h,1}^k(a_h|s_h^k, A_h^k) - \mathcal{P}_h(a_h|s_h^k, A_h^k) \right) f_{h,1}^k(s_h^k, a_h) \\
 &\quad + 2 \sum_{h=1}^H b_{h,1}^k(s_h^k, a_h^k) + \sum_{h=1}^H \dot{\zeta}_{h,1}^k + \sum_{h=2}^H \zeta_{h,1}^k, \tag{D.62}
 \end{aligned}$$

where $\zeta_{h,1}^k = \mathbb{E}_{\mathbb{P}} \left[(V_{h,1}^k - V_h^{\pi_k})(s_h) \mid s_{h-1}^k, a_{h-1}^k \right] - (V_{h,1}^k - V_h^{\pi_k})(s_h^k)$ and $\dot{\zeta}_{h,1}^k = \mathbb{E}_{\mathcal{P}} \left[\left(f_{h,1}^k - \bar{Q}_h^{\pi_k} \right)(s_h^k, a_h) \mid s_h^k, A_h^k \right] - \left(f_{h,1}^k - \bar{Q}_h^{\pi_k} \right)(s_h^k, a_h^k)$.

Now, we consider the case where $k \in \mathcal{K}_{oo} \setminus \{1\}$. In this cases, note that $h_k \in [H]$. Similar to the above analysis, by Lemma D.20, we get

$$\begin{aligned}
 (Q_1^k - Q_1^{\pi_k})(s_1^k, A_1^k) &\leq \sum_{h=1}^{h_k-1} \sum_{a_h \in A_h^k} \left(\tilde{\mathcal{P}}_{h,1}^k(a_h|s_h^k, A_h^k) - \mathcal{P}_h(a_h|s_h^k, A_h^k) \right) f_{h,1}^k(s_h^k, a_h) \\
 &\quad + \sum_{h=h_k}^H \sum_{a_h \in A_h^k} \left(\tilde{\mathcal{P}}_{h,2}^k(a_h|s_h^k, A_h^k) - \mathcal{P}_h(a_h|s_h^k, A_h^k) \right) f_{h,2}^k(s_h^k, a_h) \\
 &\quad + 2 \sum_{h=1}^H b_{h,1}^k(s_h^k, a_h^k) + 2 \sum_{h=h_k}^H b_{h,2}^k(s_h^k, a_h^k) + \sum_{h=1}^{h_k-1} \dot{\zeta}_{h,1}^k + \sum_{h=2}^{h_k-1} \zeta_{h,1}^k + \sum_{h=h_k}^H \dot{\zeta}_{h,2}^k + \sum_{h=h_k}^H \zeta_{h,2}^k, \tag{D.63}
 \end{aligned}$$

where $\zeta_{h,2}^k = \mathbb{E}_{\mathbb{P}} \left[(V_{h,2}^k - V_h^{\pi_k})(s_h) \mid s_{h-1}^k, a_{h-1}^k \right] - (V_{h,2}^k - V_h^{\pi_k})(s_h^k)$ and $\dot{\zeta}_{h,2}^k = \mathbb{E}_{\mathcal{P}} \left[\left(f_{h,2}^k - \bar{Q}_h^{\pi_k} \right)(s_h^k, a_h) \mid s_h^k, A_h^k \right] - \left(f_{h,2}^k - \bar{Q}_h^{\pi_k} \right)(s_h^k, a_h^k)$.

Plugging (D.62) and (D.63) into (D.61), and denoting $J(k, h) : [K] \times [H] \rightarrow \{1, 2\}$ as the one-to-one function that maps from $[K] \times [H]$ to the index set $\{1, 2\}$ such that $A_h^k = A_{h, J(k, h)}^k \in \operatorname{argmax}_{A \in \mathcal{A}} \sum_{a \in \mathcal{A}} \tilde{\mathcal{P}}_{h, J(k, h)}^k(a | s_h^k, A) f_{h, J(k, h)}^k(s_h^k, a)$, we obtain that

$$\begin{aligned} & \mathbf{Regret}(\mathcal{M}, K) \\ & \leq \mathcal{O}(1) + \sum_{k=2}^K \sum_{h=1}^H \sum_{a_h \in A_h^k} \left(\tilde{\mathcal{P}}_{h, J(k, h)}^k(a_h | s_h^k, A_h^k) - \mathcal{P}_h(a_h | s_h^k, A_h^k) \right) f_{h, J(k, h)}^k(s_h^k, a_h) \\ & + 2 \sum_{k=2}^K \sum_{h=1}^H \min \{1 + L, b_{h,1}^k(s_h^k, a_h^k)\} + 2 \sum_{k \in \mathcal{K}_{\text{oo}}} \sum_{h=h_k}^H \min \{1 + L, b_{h,2}^k(s_h^k, a_h^k)\} \\ & + \sum_{k=2}^K \sum_{h=1}^{h_k-1} \zeta_{h,1}^k + \sum_{k=2}^K \sum_{h=2}^{h_k-1} \zeta_{h,1}^k + \sum_{k=2}^K \sum_{h=h_k}^H \zeta_{h,2}^k + \sum_{k=2}^K \sum_{h=h_k}^H \zeta_{h,2}^k. \end{aligned}$$

Now, by applying the results from Lemma D.13 to bound the first term (which holds with probability at least $1 - 2\delta$), Lemma D.23 for the second term (which holds with probability at least $1 - 7\delta$), Lemma D.22 for the third term, and applying the Azuma-Hoeffding inequality to the remaining terms (which holds with probability at least $1 - 4\delta$), we get

$$\begin{aligned} & \mathbf{Regret}(\mathcal{M}, K) \\ & \leq \mathcal{O} \left(d\sqrt{HK} |(\log K)^{3/2} \log M + \frac{1}{\kappa} d^2 H (\log K)^3 (\log M)^2 \right) \\ & + \mathcal{O} \left(\sqrt{d_\nu H K} \cdot \log \frac{\mathcal{N}KH}{\delta} + \frac{1}{\sqrt{\kappa}} dH^{7/2} \sqrt{d_\nu} (\log K)^{3/2} \log M \cdot \log \frac{\mathcal{N}\mathcal{N}_b KH}{\delta} \cdot \sqrt{\log \frac{\mathcal{N}KH}{\delta}} \right) \\ & + \mathcal{O} \left(d_\nu H^{7/2} \log \frac{\mathcal{N}KH}{\delta} \cdot \left(\log \frac{\mathcal{N}\mathcal{N}_b KH}{\delta} \right)^{3/2} + \sqrt{\log \frac{\mathcal{N}KH}{\delta}} \cdot \left(KH\epsilon_b + \sqrt{d_\nu KH^3 \delta} \right) \right) \\ & + \mathcal{O} \left(\sqrt{\log \frac{\mathcal{N}KH}{\delta} \log \frac{\mathcal{N}\mathcal{N}_b KH}{\delta}} \cdot \sqrt{d_\nu H} \cdot \left(\sqrt{H^2 \sum_{k \in \mathcal{K}_o} u_k} + \sqrt{H^2 |\mathcal{K}_{\text{oo}}|} \right) \right) \\ & + \mathcal{O} \left(\sqrt{\log \frac{\mathcal{N}\mathcal{N}_b KH}{\delta}} \cdot \left(H\sqrt{d_\nu |\mathcal{K}_{\text{oo}}|} + |\mathcal{K}_{\text{oo}}| H\epsilon_b \right) \right) \\ & \leq \mathcal{O} \left(d\sqrt{HK} |(\log K)^{3/2} \log M + \frac{1}{\kappa} d^2 H (\log K)^3 (\log M)^2 \right) \\ & + \mathcal{O} \left(\sqrt{d_\nu H K} \cdot \log \frac{\mathcal{N}KH}{\delta} + \frac{1}{\sqrt{\kappa}} dH^{7/2} \sqrt{d_\nu} (\log K)^{3/2} \log M \cdot \log \frac{\mathcal{N}\mathcal{N}_b KH}{\delta} \cdot \sqrt{\log \frac{\mathcal{N}KH}{\delta}} \right) \\ & + \mathcal{O} \left(d_\nu H^{7/2} \log \frac{\mathcal{N}KH}{\delta} \cdot \left(\log \frac{\mathcal{N}\mathcal{N}_b KH}{\delta} \right)^{3/2} + \sqrt{\log \frac{\mathcal{N}\mathcal{N}_b KH}{\delta}} \cdot \left(KH\epsilon_b + \sqrt{d_\nu KH^3 \delta} \right) \right) \\ & + \mathcal{O} \left(\sqrt{\log \frac{\mathcal{N}KH}{\delta} \log \frac{\mathcal{N}\mathcal{N}_b KH}{\delta}} \cdot \sqrt{d_\nu H} \cdot \sqrt{H^2 \sum_{k \in \mathcal{K}_o} u_k} \right), \end{aligned} \tag{D.64}$$

where the second inequality holds by Lemma D.24 with probability at least $1 - 2\delta$, and use the fact that $|\mathcal{K}_{\text{oo}}| H\epsilon_b \leq KH\epsilon_b$. Now, we apply the AM-GM inequality to the term $\mathcal{O} \left(\frac{1}{\sqrt{\kappa}} dH^{7/2} \sqrt{d_\nu} (\log K)^{3/2} \log M \cdot \log \frac{\mathcal{N}\mathcal{N}_b KH}{\delta} \cdot \sqrt{\log \frac{\mathcal{N}KH}{\delta}} \right)$, thus we get

$$\begin{aligned} & \mathcal{O} \left(\frac{1}{\sqrt{\kappa}} dH^{7/2} \sqrt{d_\nu} (\log K)^{3/2} \log M \cdot \log \frac{\mathcal{N}\mathcal{N}_b KH}{\delta} \cdot \sqrt{\log \frac{\mathcal{N}KH}{\delta}} \right) \\ & \leq \mathcal{O} \left(\frac{1}{\kappa} d^2 H^2 (\log K)^3 (\log M)^2 + d_\nu H^5 \left(\log \frac{\mathcal{N}\mathcal{N}_b KH}{\delta} \right)^2 \cdot \log \frac{\mathcal{N}KH}{\delta} \right). \end{aligned} \tag{D.65}$$

Furthermore, by substituting the chosen values of u_k and applying the AM-GM inequality, we get

$$\begin{aligned}
 & \mathcal{O} \left(\sqrt{\log \frac{\mathcal{N}KH}{\delta} \log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta}} \cdot \sqrt{d_\nu H} \cdot \sqrt{H^2 \sum_{k \in \mathcal{K}_o} u_k} \right) \\
 &= \mathcal{O} \left(\sqrt{d_\nu HK \cdot \log \frac{\mathcal{N}KH}{\delta}} + d^2 H^2 (\log K)^3 (\log M)^2 + d_\nu H^5 \log \frac{\mathcal{N}KH}{\delta} \cdot \left(\log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta} \right)^2 \right) \\
 &+ \mathcal{O} \left(\sqrt{\log \frac{\mathcal{N}KH}{\delta}} \cdot KH \epsilon_b \right). \tag{D.66}
 \end{aligned}$$

Then, by plugging (D.65) and (D.66) into (D.64), and setting $\delta < \frac{1}{H^2+15}$, we derive that

$$\begin{aligned}
 \mathbf{Regret}(\mathcal{M}, K) &= \mathcal{O} \left(d\sqrt{HK} (\log K)^{3/2} \log M + \sqrt{d_\nu HK \cdot \log \frac{\mathcal{N}KH}{\delta}} + \frac{1}{\kappa} d^2 H^2 (\log K)^3 (\log M)^2 \right) \\
 &+ \mathcal{O} \left(d_\nu H^5 \log \frac{\mathcal{N}KH}{\delta} \cdot \left(\log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta} \right)^2 + \sqrt{\log \frac{\mathcal{N}\mathcal{N}_bKH}{\delta}} \cdot KH \epsilon_b \right).
 \end{aligned}$$

We conclude the proof of Theorem 5.1. \square

E. Proof of Theorem 5.2

In this section, we introduce several properties of linear function class. We formally define the linear MDP as follows:

Definition E.1 (Linear MDPs, Yang & Wang 2019; Jin et al. 2020). An MDP \mathcal{M} is a linear MDP if we have a known feature mapping $\psi : \mathcal{S} \times \mathcal{I} \rightarrow \mathbb{R}^{d^{\text{lin}}}$, and there exist d^{lin} unknown (signed) measures $\boldsymbol{\mu}_h^* = (\mu_h^{(1)}, \dots, \mu_h^{(d^{\text{lin}})})$ over \mathcal{S} and unknown vector $\mathbf{w}_h^* \in \mathbb{R}^{d^{\text{lin}}}$, such that for any $(s, a) \in \mathcal{S} \times \mathcal{I}$, we have $\mathbb{P}_h(\cdot | s, a) = \langle \psi(s, a), \boldsymbol{\mu}_h^*(\cdot) \rangle$ and $r_h(s, a) = \langle \psi(s, a), \mathbf{w}_h^* \rangle$. We assume that $\sup_{(s,a) \in \mathcal{S} \times \mathcal{I}} \|\psi(s, a)\|_2 \leq 1$, $\max\{\|\sum_{s \in \mathcal{S}} |\boldsymbol{\mu}_h^*(s)|\|_2, \|\mathbf{w}_h^*\|_2\} \leq \sqrt{d^{\text{lin}}}$ for all $h \in [H]$.

In this proof, to explicitly indicate the dependency on parameters, we denote the linear MDPs as $\mathcal{M}_{\boldsymbol{\theta}^*, \boldsymbol{\mu}^*, \mathbf{w}^*}$, where $\boldsymbol{\theta}^* = \{\boldsymbol{\theta}_h^*\}_{h=1}^H$, $\boldsymbol{\mu}^* = \{\boldsymbol{\mu}_h^*\}_{h=1}^H$, and $\mathbf{w}^* = \{\mathbf{w}_h^*\}_{h=1}^H$.

We also assume that $\sum_{h=1}^H r_h \in [0, 1]$. Proposition 2.3 of Jin et al. (2020) shows that linear MDPs satisfy Assumption 3.3 under the linear function class $\mathcal{F}_h^{\text{lin}}$ defined as follows:

$$\mathcal{F}_h^{\text{lin}} := \left\{ \langle \psi(\cdot, \cdot), \boldsymbol{\omega}_h \rangle : \boldsymbol{\omega}_h \in \mathbb{R}^{d^{\text{lin}}}, \|\boldsymbol{\omega}_h\|_2 \leq 2\sqrt{d^{\text{lin}}} \right\}, \quad \text{for any } h \in [H]. \tag{E.1}$$

For linear MDPs, let $\mathcal{F}_h^{\text{lin}}(\epsilon_c)$ be an ϵ_c -cover of $\mathcal{F}_h^{\text{lin}}$ under the ℓ_∞ norm, so that

$$\log |\mathcal{F}_h^{\text{lin}}(\epsilon_c)| = \mathcal{O} \left(d^{\text{lin}} \log \frac{2\sqrt{d^{\text{lin}}}}{\epsilon_c} \right) = \tilde{\mathcal{O}}(d^{\text{lin}}). \tag{E.2}$$

Then, the definition of generalized Eluder dimension for the linear function class $\mathcal{F}_h^{\text{lin}}$ can be expressed as:

Lemma E.2 (Lemma 3 of Agarwal et al. 2023). For the class $\mathcal{F}_h^{\text{lin}}$ defined in (E.1), letting $\mathcal{F}_h^{\text{lin}}(\epsilon_c)$ be the ϵ_c -cover of $\mathcal{F}_h^{\text{lin}}$ for some $\epsilon_c > 0$, we have

$$\dim_{\nu, K}(\mathcal{F}_h^{\text{lin}}(\epsilon_c)) \leq \dim_{\nu, K}(\mathcal{F}_h^{\text{lin}}) = \mathcal{O} \left(d^{\text{lin}} \log \left(1 + \frac{K}{\nu^2 \rho} \right) \right) = \tilde{\mathcal{O}}(d^{\text{lin}}).$$

The bonus oracle for linear MDPs can be easily instantiated using the standard elliptical bonus, and, as demonstrated in the next lemma, satisfies all the required properties for a bonus oracle.

Lemma E.3 (Bonus oracle \mathcal{B} for linear MDPs, Lemma 7 of Agarwal et al. 2023). *Given $K, H \in \mathbb{Z}_+$, suppose all $\beta_h^k \leq \beta$ and β_h^k is non-decreasing in $k \in [K]$ for each $h \in [H]$. For any $k \geq 1, h \in [H]$, variances $\{\bar{\sigma}_h^\tau\}_{\tau=1}^h$ satisfying $\bar{\sigma}_h^\tau \geq \nu$ for some $\nu > 0$, dataset $\mathcal{D}_h^{k-1} = \{\psi(s_h^\tau, a_h^\tau), a_h^\tau, r_h^\tau, \psi(s_{h+1}^\tau, a_{h+1}^\tau)\}_{\tau=1}^{k-1}$, function class \mathcal{F}_h^k and $\hat{f}_h^k \in \mathcal{F}_h^k$ defined via weighted regression in (3), and parameters $\rho, \epsilon_c > 0$, let $\mathcal{B}(\{\bar{\sigma}_h^\tau\}_{\tau=1}^h, \mathcal{D}_h^{k-1}, \mathcal{F}_h^k, \hat{f}_h^k, \beta_h^k, \rho, \epsilon_c) = \|\psi(s, a)\|_{(\Sigma_h^k)^{-1}} \sqrt{(\beta_h^k)^2 + \rho}$, where $\Sigma_h^k = \frac{\rho}{16d}I + \sum_{\tau=1}^{k-1} \frac{1}{(\bar{\sigma}_h^\tau)^2} \psi(s_h^\tau, a_h^\tau) \psi(s_h^\tau, a_h^\tau)^\top$. For any choice of covering radius $\epsilon_c \leq \nu \sqrt{\rho/8K}$, the oracle satisfies all the properties of Definition B.1 with*

$$\log \mathcal{N}_b = \log |\mathcal{W}| = \mathcal{O} \left((d^{\text{lin}})^2 \log \left(1 + d^{\text{lin}} \sqrt{d^{\text{lin}} \beta / (\rho \epsilon_c^2)} \right) \right) = \tilde{\mathcal{O}}((d^{\text{lin}})^2).$$

Theorem E.4 (Formal version of Theorem 5.2, Regret upper bound of MNL-VQL for linear MDPs). *Under the same conditions with Theorem 5.1, suppose that the underlying MDP has linear transition probabilities and rewards, so that the function class for linear MDPs, $\mathcal{F}_h^{\text{lin}}$, satisfies Assumption 3.3. Let $\mathcal{F}_h^{\text{lin}}(\epsilon_c)$ be an ϵ_c -cover of $\mathcal{F}_h^{\text{lin}}$ under ℓ_∞ norm. We set $\rho = 1, u_k = \tilde{\Theta}((d^{\text{lin}})^3 H^{5/2} + d(d^{\text{lin}})^{3/2} H^{5/2})/\sqrt{K}$, $\nu = \sqrt{1/HK}$, $\epsilon_b = \epsilon_c \leq 1/(8HK)$ and $\delta < 1/(H^2 + 15)$. Then, with probability at least $1 - \delta$, the cumulative regret of MNL-VQL, with bonus oracle defined in Lemma E.3, is upper-bounded by*

$$\text{Regret}(\mathcal{M}_{\theta^*, \mu^*, \mathbf{w}^*}, K) = \tilde{\mathcal{O}} \left(\underbrace{d\sqrt{HK} + \frac{1}{\kappa} d^2 H^2}_{\text{regret from MNL model}} + \underbrace{d^{\text{lin}} \sqrt{HK} + (d^{\text{lin}})^6 H^5}_{\text{regret from linear MDPs}} \right).$$

Proof of Theorem 5.2. We apply the above results to linear MDPs $\mathcal{M}_{\theta^*, \mu^*, \mathbf{w}^*}$ with function class $\mathcal{F}_h^{\text{lin}}(\epsilon_c)$, $h \in [H]$, and bonus oracle \mathcal{B} . From (E.2), we know that $\mathcal{N} = \tilde{\mathcal{O}}(d^{\text{lin}})$. Additionally, Lemma E.2 shows that $d_\nu = \tilde{\mathcal{O}}(d^{\text{lin}})$. Therefore, by combining these results with Theorem 5.1 and Lemma E.3, we can establish the upper regret bounds for linear MDPs.

$$\text{Regret}(\mathcal{M}_{\theta^*, \mu^*, \mathbf{w}^*}, K) = \tilde{\mathcal{O}} \left(d\sqrt{HK} + d^{\text{lin}} \sqrt{HK} + \frac{1}{\kappa} d^2 H^2 + (d^{\text{lin}})^6 H^5 \right),$$

where we set $\rho = 1, u_k = \tilde{\Theta}((d^{\text{lin}})^3 H^{5/2} + d(d^{\text{lin}})^{3/2} H^{5/2})/\sqrt{K}$, $\nu = \sqrt{1/HK}$, $\epsilon_b = \epsilon_c \leq 1/(8HK)$ and $\delta < 1/(H^2 + 15)$. \square

F. Proof of Theorem 5.3

In this section, we provide a regret lower bound for linear MDPs with preference model. We construct a hard instance $\mathcal{M}(\mathcal{S}, \mathcal{I}, \mathcal{A}, M, \{\mathcal{P}_h\}_{h=1}^H, \{\mathbb{P}_h\}_{h=1}^H, \{r_{h=1}^H\}, H)$, illustrated as in Figure F.1. This instance is based on an $H + 1$ -layered structure, where each layer is a variation of the hard-to-learn MDPs introduced in Zhou et al. (2021b).

Without loss of generality, we assume that $d^{\text{lin}} \geq 6$ and that $d^{\text{lin}} - 5$ is divisible by 2.³ Let $i \in [H + 2]$ represent the layer index. For each layer $i \in [H + 2]$, there are $H - i + 3$ states, denoted as $x_i^{(1)}, \dots, x_{H+2}^{(i)}$, where $x_{H+2}^{(i)}$ is the absorbing state. Furthermore, there is a *global* absorbing state x_0 , which can only be reached at any state and horizon through the user's choice of the outside option \mathbf{a}_0 (not choosing any item in the assortment). Thus, there are $(H + 1)(H + 2)/2 + 1$ states in total in the set of states \mathcal{S} . There are $2^{(d^{\text{lin}}-5)/2} + 1$ items, so the item set is $\mathcal{I} = \{-1, 1\}^{(d^{\text{lin}}-5)/2} \cup \{\mathbf{a}_0\}$. The set of candidate assortments follows the definition in Section 3, i.e., $\mathcal{A} = \{A \subseteq \mathcal{I} : \mathbf{a}_0 \in A, 1 \leq |A \setminus \{\mathbf{a}_0\}| \leq M\}$.

F.1. Construction of linear transitions and rewards

At each episode $k \in [K]$, the agent starts from the fixed initial state $x_1^{(1)}$. We define \mathbf{a}_h^* as an item such that $\mathbf{a}_h^* \in \arg\max_{\mathbf{a} \in \mathcal{I} \setminus \{\mathbf{a}_0\}} \langle \mu_h, \mathbf{a} \rangle$, where $\mu_h \in \{-\Delta, \Delta\}^{(d^{\text{lin}}-5)/2}$ with $\Delta = \sqrt{\delta/K}/(4\sqrt{2})$ and $\delta = 1/H$.

If the state is $x_h^{(i)}$ with $i \in [H + 1]$ and $h \in [i, H + 1]$, and the user chooses the item \mathbf{a}_h^* , the agent remains in the same layer i and receives a reward of γ^{i-1}/H , where $\gamma = \frac{H}{1+H}$. The next state will be either $x_{h+1 \wedge H+1}^{(i)}$ or $x_{H+2}^{(i)}$, with probabilities $1 - (\delta + \langle \mu_h, \mathbf{a} \rangle)$ and $\delta + \langle \mu_h, \mathbf{a} \rangle$, respectively. If the user chooses an item $\mathbf{a} \neq \mathbf{a}_0, \mathbf{a}_h^*$ in the state $x_h^{(i)}$ with $i \in [H + 1]$ and $h \in [i, H + 1]$, the agent obtains a reward of γ^i/H and transitions to $x_{h+1 \wedge H+1}^{(i+1)}$ or $x_{H+2}^{(i+2 \wedge H+2)}$, with probabilities

³If $d^{\text{lin}} - 5$ is not divisible by 2, we can set $d^{\text{lin}} \leftarrow d^{\text{lin}} + 1$ by adding zero padding.

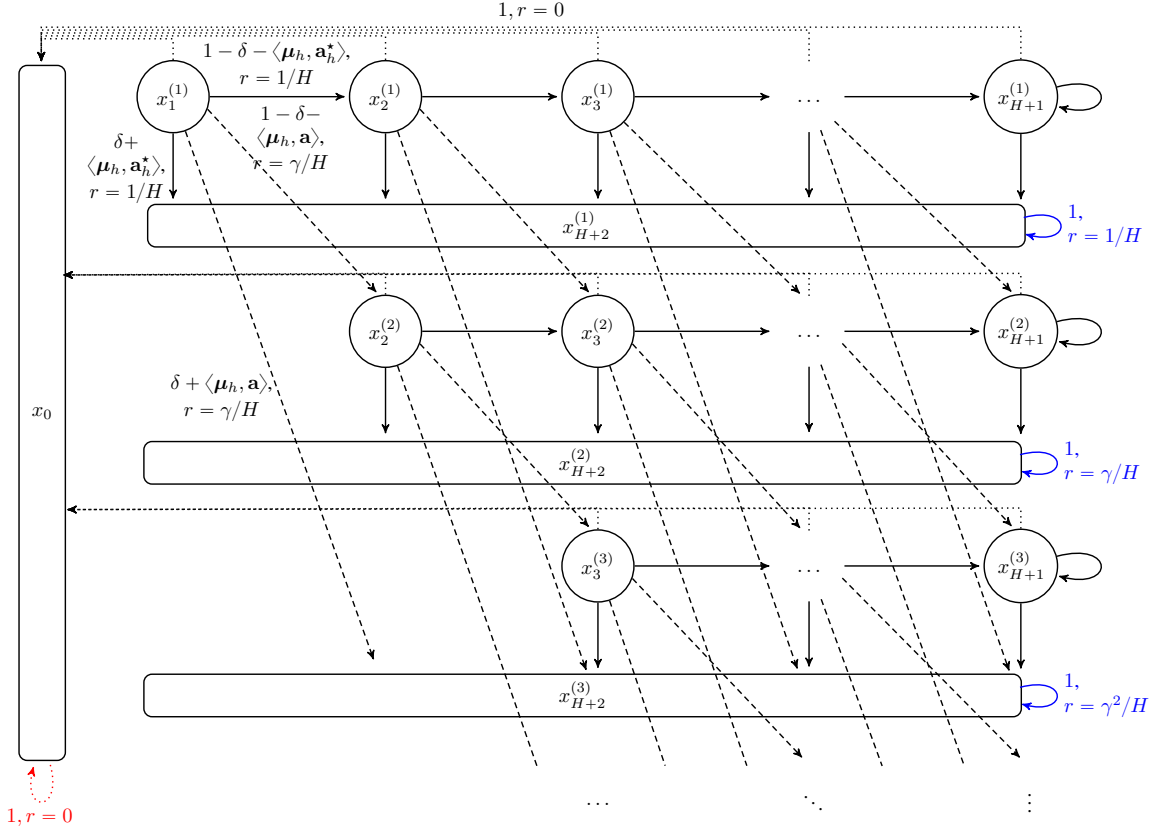


Figure F.1: Inhomogeneous, hard-to-learn linear MDPs with MNL preference model. The solid line indicates the transition caused by the user choosing the item \mathbf{a}_h^* (with a reward of $r_h = \gamma^{i-1}/H$), the dashed line shows the transition caused by the user choosing any item $\mathbf{a} \neq \mathbf{a}_h^*, \mathbf{a}_0$ (with a reward of $r_h = \gamma^i/H$), and the dotted line represents the transition caused by the user choosing the outside option \mathbf{a}_0 (with a reward of $r_h = 0$). The blue solid line indicates a transition from the absorbing state back to itself, caused by the user choosing any item (with a reward of $r_h = \gamma^{i-1}/H$), and the red dotted line indicates a transition from the *global* absorbing state back to itself, caused by the user choosing any item (with a reward of $r_h = 0$).

$1 - (\delta + \langle \boldsymbol{\mu}_h, \mathbf{a} \rangle)$ and $\delta + \langle \boldsymbol{\mu}_h, \mathbf{a} \rangle$, respectively. If the user does not choose any item, i.e., chooses the outside option \mathbf{a}_0 , in the state $x_h^{(i)}$ with $i \in [H+1]$ and $h \in [i, H+1]$, the agent will deterministically transition to the global absorbing state x_0 and receive no reward.

If the agent is in any of the absorbing states $x_{H+2}^{(i)}$ for $i \in [H+2]$ —the agent will remain in the same state and receive a reward of γ^{i-1}/H , regardless of which item (including the outside option) the user chooses.

Formally, we construct transition probabilities $\mathbb{P}_h(s'|s, \mathbf{a}) = \langle \psi(s, \mathbf{a}), \boldsymbol{\mu}_h^*(s') \rangle$, with

$$\psi(s, \mathbf{a}) = \begin{cases} (\alpha, \beta \mathbf{a}^\top, 0, 0, 0, 0, \frac{\gamma^{i-1}}{\sqrt{2}})^\top, & s = x_h^{(i)}, \mathbf{a} = \mathbf{a}_h^*, i \in [H+1], h \in [i, H+1]; \\ (0, 0, \alpha, \beta \mathbf{a}^\top, 0, 0, \frac{\gamma^i}{\sqrt{2}})^\top, & s = x_h^{(i)}, \mathbf{a} \neq \mathbf{a}_h^*, \mathbf{a}_0, i \in [H+1], h \in [i, H+1]; \\ (0, 0^\top, 0, 0^\top, 0, 1, 0)^\top, & s = x_h^{(i)}, \mathbf{a} = \mathbf{a}_0, i \in [H+1], h \in [i, H+1]; \\ (0, 0^\top, 0, 0^\top, \frac{1}{\sqrt{2}}, 0, \frac{\gamma^{i-1}}{\sqrt{2}})^\top, & s = x_{H+2}^{(i)}, i \in [H+2], \end{cases} \quad (\text{F.1})$$

and

$$\boldsymbol{\mu}_h^*(s') = \begin{cases} (\frac{1-\delta}{\alpha}, -\frac{\boldsymbol{\mu}_h^\top}{\beta}, 0, \mathbf{0}, 0, 0)^\top, & s' = x_{h+1 \wedge H+1}^{(i)}; \\ (\frac{\delta}{\alpha}, \frac{\boldsymbol{\mu}_h^\top}{\beta}, 0, \mathbf{0}, \sqrt{2}, 0)^\top, & s' = x_{H+2}^{(i)}; \\ (0, \mathbf{0}, \frac{1-\delta}{\alpha}, -\frac{\boldsymbol{\mu}_h^\top}{\beta}, 0, 0)^\top, & s' = x_{h+1}^{(i+1)}; \\ (0, \mathbf{0}, \frac{\delta}{\alpha}, \frac{\boldsymbol{\mu}_h^\top}{\beta}, 0, 0)^\top, & s' = x_{H+2}^{(i+2 \wedge H+2)}; \\ (0, \mathbf{0}^\top, 0, \mathbf{0}, 0, 1)^\top, & s' = x_0; \\ (0, \mathbf{0}^\top, 0, \mathbf{0}, 0, 0)^\top, & \text{otherwise,} \end{cases} \quad (\text{F.2})$$

where we denote $\mathbf{0} \in \mathbb{R}^{(d^{\text{lin}}-5)/2}$ as the zero vector of dimension $(d^{\text{lin}} - 5)/2$, and set $\gamma = \frac{H}{H+1}$ as the discount factor for transitioning to the next layer. Additionally, we choose $\delta = 1/H$, $\boldsymbol{\mu}_h \in \{-\Delta, \Delta\}^{(d^{\text{lin}}-5)/2}$ with $\Delta = \sqrt{\delta/K}/(4\sqrt{2})$, $\alpha = \sqrt{1/(2 + \Delta \cdot (d^{\text{lin}} - 5))}$, and $\beta = \sqrt{\Delta/(2 + \Delta \cdot (d^{\text{lin}} - 5))}$.

And the parameter vectors for the linear rewards $r_h(s, \mathbf{a}) = \langle \psi(s, \mathbf{a}), \mathbf{w}_h^* \rangle$ are as follows:

$$\mathbf{w}_h^* = (0, \mathbf{0}^\top, 0, \mathbf{0}, 0, 0, \sqrt{2}/H)^\top,$$

which ensures that the reward function satisfies:

$$r_h(s, \mathbf{a}) = \begin{cases} \gamma^{i-1}/H, & s = x_h^{(i)}, \mathbf{a} = \mathbf{a}_h^*, i \in [H+1], h \in [i, H+1]; \\ \gamma^i/H, & s = x_h^{(i)}, \mathbf{a} \neq \mathbf{a}_h^*, \mathbf{a}_0, i \in [H+1], h \in [i, H+1]; \\ 0, & s = x_h^{(i)}, \mathbf{a} = \mathbf{a}_0, i \in [H+1], h \in [i, H+1]; \\ \gamma^{i-1}/H, & s = x_{H+2}^{(i)}, i \in [H+2], \end{cases}$$

where $0 < \gamma \leq \frac{H}{H+1}$ is the discount factor for transitioning to the next layer.

This parameter setting satisfies the boundedness assumption of linear MDPs (refer Definition E.1). First, we show that $\|\psi(s, \mathbf{a})\|_2 \leq 1$:

$$\|\psi(s, \mathbf{a})\|_2^2 \leq \alpha^2 + \frac{d^{\text{lin}} - 5}{2} \beta^2 + \frac{1}{2} = 1, \quad (\text{the first and second cases of (F.1)}),$$

$$\|\psi(s, \mathbf{a})\|_2^2 = 1, \quad (\text{the third case of (F.1)}),$$

$$\|\psi(s, \mathbf{a})\|_2^2 \leq \frac{1}{2} + \frac{1}{2} = 1, \quad (\text{the fourth case of (F.1)}),$$

Moreover, since $d^{\text{lin}} \geq 6$ and $K \geq 13(d^{\text{lin}} - 5)^2/H$, we ensure that $\max\{\|\sum_{s \in S} \boldsymbol{\mu}_h(s)\|_2, \|\mathbf{w}_h^*\|_2\} \leq \sqrt{d^{\text{lin}}}$.

$$\begin{aligned} \left\| \sum_{s \in S} \boldsymbol{\mu}_h(s) \right\|_2^2 &= \frac{2(1-\delta)^2 + 2\delta^2}{\alpha^2} + \frac{\|\boldsymbol{\mu}_h\|_2}{\beta^2} + 3 \\ &\leq 2(2 + \Delta \cdot (d^{\text{lin}} - 5)) + 2\Delta \cdot (d^{\text{lin}} - 5)(2 + \Delta \cdot (d^{\text{lin}} - 5)) \\ &\leq (2 + 2\Delta \cdot (d^{\text{lin}} - 5))^2 \leq d^{\text{lin}}, \\ \text{and } \|\mathbf{w}_h^*\|_2^2 &\leq \frac{2}{H^2} \leq d^{\text{lin}}. \end{aligned}$$

F.2. Construction of MNL preference model

Inspired by the lower bound proposed in Lee & Oh (2024), we construct an adversarial setting for the MNL preference model.

We assume that $d \geq 2$ and that $d - 1$ is divisible by 4 (without loss of generality). Let $\epsilon \in (0, \frac{1}{(d-1)\sqrt{d-1}})$ be a small positive parameter. Throughout the proof, we set $\epsilon = \sqrt{\frac{d-1}{144C \cdot K} \cdot \frac{(H+1)^2}{H}}$, for some $C > 0$. For every subset $W \subseteq [d-1]$, we define the corresponding parameter $\boldsymbol{\theta}_W \in \mathbb{R}^{d-1}$ as $[\boldsymbol{\theta}_W]_j = \epsilon$ for all $j \in W$, and $[\boldsymbol{\theta}_W]_j = 0$ for all $j \notin W$.

Next, for any $h \in [H]$, we define the parameter set as:

$$\begin{aligned}\boldsymbol{\theta}_h^* &\in \Theta := \{(\boldsymbol{\theta}_W^\top, -\log H)^\top : W \in \mathcal{W}_{(d-1)/4}\} \\ &= \{(\boldsymbol{\theta}_W^\top, -\log H)^\top : W \subseteq [d-1], |W| = (d-1)/4\},\end{aligned}$$

where \mathcal{W}_k denotes the class of all subsets of $[d-1]$ of size k .

The feature vector $\phi(s, \mathbf{a})$ is invariant across the state s . For each $U \in \mathcal{W}_{(d-1)/4}$, we define vectors $z_U \in \mathbb{R}^{d-1}$ as follows:

$$[z_U]_j = 1/\sqrt{d-1} \quad \text{for } j \in U; \quad [z_U]_j = 0 \quad \text{for } j \notin U.$$

Let $\mathcal{Z} := \{z_U : U \in \mathcal{W}_{(d-1)/4}\}$. We define the function $Z : \mathcal{I} \rightarrow \mathcal{Z}$, so that $Z(\mathbf{a}) \in \mathcal{Z}$. Then, the feature vector $\phi(s, \mathbf{a})$ is constructed as follows:

$$\phi(s, \mathbf{a}) = \begin{cases} (Z(\mathbf{a})^\top, 0)^\top, & \mathbf{a} \neq \mathbf{a}_0; \\ (\mathbf{0}, 1)^\top, & \mathbf{a} = \mathbf{a}_0, \end{cases}$$

where $\mathbf{0} \in \mathbb{R}^{d-1}$. For all $V \in \mathcal{V}_{d/4}$ and $(s, \mathbf{a}) \in \mathcal{S} \times \mathcal{I}$, it can be verified that $\boldsymbol{\theta}_V$ and $\phi(s, \mathbf{a})$ satisfy the boundedness in Assumption 3.1 as follows:

$$\begin{aligned}\|\phi(s, \mathbf{a})\|_2 &\leq \sqrt{(d-1) \cdot 1/(d-1)} = 1, \\ \|\boldsymbol{\theta}_h^*\|_2 &\leq \sqrt{(d-1)\epsilon^2 + (-\log H)^2} \leq \sqrt{2} \log H =: B.\end{aligned}$$

Let \mathbf{a}_h^* (defined in the previous subsection) also have the maximum utility, i.e., $\mathbf{a}_h^* \in \operatorname{argmax}_{\mathbf{a} \in \mathcal{I} \setminus \{\mathbf{a}_0\}} \langle \boldsymbol{\theta}_h^*, \phi(s, \mathbf{a}) \rangle$ (note that $\phi(\cdot, \cdot)$ is identical for all $s \in \mathcal{S}$).

F.3. Proof of Theorem 5.3

A good policy is one that quickly reaches the state $x_{H+2}^{(i)}$ while remaining in the lower layers (i.e., with lower i). Recall that the item \mathbf{a}_h^* has the highest utility and, therefore, the highest choice probability. It also has the best chance of quickly reaching the state $x_{H+2}^{(i)}$ while staying within the same layer. In other words, a good policy encourages the user to frequently select the item $\mathbf{a}_h^* \in \operatorname{argmax}_{\mathbf{a} \in \mathcal{I} \setminus \{\mathbf{a}_0\}} \langle \boldsymbol{\mu}_h, \mathbf{a} \rangle = \operatorname{argmax}_{\mathbf{a} \in \mathcal{I} \setminus \{\mathbf{a}_0\}} \langle \boldsymbol{\theta}_h^*, \phi(s, \mathbf{a}) \rangle$. Note that \mathbf{a}_h^* is unique due to the way the action space and transition probabilities are constructed.

We formally restate Theorem 5.3 as follows.

Theorem F.1 (Restatement of Theorem 5.3, Regret lower bound for linear MDPs with preference feedback). *Suppose that $d \geq 2$, $d^{\text{lin}} \geq 6$, $H \geq 3$, and $K \geq \max\{C \cdot (d^{\text{lin}} - 5)^2 H(H+1)^2, C' \cdot (d-1)^4(1+H)/H\}$ for some constant $C, C' > 0$. Then, for any algorithm, there exists an episodic linear MDP $\mathcal{M}_{\boldsymbol{\theta}, \boldsymbol{\mu}, \mathbf{w}}$ with MNL preference feedback such that the worst-case expected regret is lower bounded as follows:*

$$\sup_{\boldsymbol{\theta}, \boldsymbol{\mu}, \mathbf{w}} \mathbb{E}_{\boldsymbol{\theta}, \boldsymbol{\mu}, \mathbf{w}} [\text{Regret}(\mathcal{M}_{\boldsymbol{\theta}, \boldsymbol{\mu}, \mathbf{w}}, K)] = \Omega\left(d\sqrt{HK} + d^{\text{lin}}\sqrt{HK}\right).$$

Proof of Theorem 5.3. Fix $\boldsymbol{\theta}$ and $\boldsymbol{\mu}$ so that we can omit the parameter dependency of \mathbb{P} and \mathcal{P} throughout the proof. Based on the construction of the hard instance \mathcal{M} discussed in the previous subsections, the following lemma shows that the optimal assortment at horizon $h \in [H]$ is $\{\mathbf{a}_0, \mathbf{a}_h^*\}$.

Lemma F.2. *For any $h \in [H]$, we have $A_h^* = \{\mathbf{a}_0, \mathbf{a}_h^*\}$.*

Furthermore, we can bound the expected value of \bar{Q}^* for any assortment as follows:

Lemma F.3. *For any $(A, i, h) \in \mathcal{A} \times [H] \times [H]$, let $\tilde{\mathbf{a}}_h^{(i)} \in \operatorname{argmax}_{\mathbf{a} \in A \setminus \{\mathbf{a}_0\}} \phi(x_h^{(i)}, \mathbf{a})^\top \boldsymbol{\theta}_h^*$, $\tilde{A}_h^{(i)} = \{\tilde{\mathbf{a}}_h^{(i)}, \mathbf{a}_0\}$, and $\bar{\mathbf{a}}_h^{(i)} \in \operatorname{argmax}_{\mathbf{a} \in A \setminus \{\mathbf{a}_0\}} \bar{Q}_h^\pi(x_h^{(i)}, \mathbf{a})$. For any $\mathbf{a}' \neq \mathbf{a}_0$, we define*

$$\begin{aligned}\tilde{Q}_h^\pi(x_h^{(i)}, \mathbf{a}_h^*, \mathbf{a}') &:= \begin{cases} \frac{\gamma^{i-1}}{H} + \mathbb{P}_h(x_{h+1}^{(i)} | x_h^{(i)}, \mathbf{a}_h^*) V_{h+1}^\pi(x_{h+1}^{(i)}) + \mathbb{P}_h(x_{H+2}^{(i)} | x_h^{(i)}, \mathbf{a}') \frac{(H-h)\gamma^{i-1}}{H}, & \mathbf{a}' = \mathbf{a}_h^*, \\ \frac{\gamma^{i-1}}{H} + \mathbb{P}_h(x_{h+1}^{(i)} | x_h^{(i)}, \mathbf{a}_h^*) V_{h+1}^\pi(x_{h+1}^{(i)}) + \mathbb{P}_h(x_{H+2}^{(i)} | x_h^{(i)}, \mathbf{a}') \frac{(H-h)\gamma^{i-1}}{H}, & \mathbf{a}' \neq \mathbf{a}_h^*. \end{cases}\end{aligned}$$

Then, for any policy π , if $K \geq 4(d^{\text{lin}} - 5)^2 H(H + 1)^2$, we have

$$\sum_{\mathbf{a} \in A} \mathcal{P}_h(\mathbf{a} | x_h^{(i)}, A) \bar{Q}_h^\pi(x_h^{(i)}, \mathbf{a}) \leq \mathcal{P}_h(\tilde{\mathbf{a}}_h^{(i)} | x_h^{(i)}, \tilde{A}_h^{(i)}) \tilde{Q}_h^\pi(x_h^{(i)}, \mathbf{a}_h^*, \bar{\mathbf{a}}_h^{(i)}).$$

Now, we are ready to provide the proof of Theorem 5.3.

For any $h \in [H]$ and any $A_h \in \mathcal{A}$, let $\tilde{\mathbf{a}}_h \in \arg\max_{\mathbf{a} \in A_h \setminus \{\mathbf{a}_0\}} \phi(x_h^{(i)}, \mathbf{a})^\top \boldsymbol{\theta}_h^*$. We also denote $\tilde{A}_h = \{\tilde{\mathbf{a}}_h, \mathbf{a}_0\}$ and $\bar{\mathbf{a}}_h \in \arg\max_{\mathbf{a} \in A \setminus \{\mathbf{a}_0\}} \bar{Q}_h^\pi(x_h, \mathbf{a})$. Recall that the index i can be omitted for $\tilde{\mathbf{a}}_h$ because the transition and choice probabilities are identical across all $x_h^{(1)}, \dots, x_h^{(H)}$ given $\mathbf{a} \in \mathcal{I}$. The change in layer i only affects the scaling of rewards, and consequently the \bar{Q} -values, but the item that maximizes $\bar{Q}(x_h^{(i)}, \mathbf{a})$ remains the same across layers.

By applying Lemma F.3, the value of policy π in state $x_1^{(1)}$ can be bounded as follows:

$$V_1^\pi(x_1^{(1)}) = \sum_{\mathbf{a} \in A_1} \mathcal{P}_1(\mathbf{a} | x_1^{(1)}, A_1) \bar{Q}_1^\pi(x_1^{(1)}, \mathbf{a}) \leq \mathcal{P}_1(\tilde{\mathbf{a}}_1 | x_1^{(1)}, \tilde{A}_1) \tilde{Q}_1^\pi(x_1^{(1)}, \mathbf{a}_1^*, \bar{\mathbf{a}}_1). \quad (\text{F.3})$$

Moreover, according to Lemma F.2, the optimal assortment for horizon $h \in [H]$ is $A_h^* = \{\mathbf{a}_0, \mathbf{a}_h^*\}$. Thus, the optimal value function in state $x_1^{(1)}$ can be written as follows:

$$V_1^*(x_1^{(1)}) = \sum_{\mathbf{a} \in A_1^*} \mathcal{P}_1(\mathbf{a} | x_1^{(1)}, A_1^*) \bar{Q}_1^*(x_1^{(1)}, \mathbf{a}) = \mathcal{P}_1(\mathbf{a}_1^* | x_1^{(1)}, A_1^*) \bar{Q}_1^*(x_1^{(1)}, \mathbf{a}_1^*),$$

where the last equality holds because $\bar{Q}_h^*(x_h^{(i)}, \mathbf{a}_0) = 0$. We denote s_{H+2} can be either $x_{H+2}^{(1)}$ or $x_{H+2}^{(3)}$, depending on whether the item (for transition) is \mathbf{a}_h^* or any $\mathbf{a} \neq \mathbf{a}_h^*, \mathbf{a}_0$. Then, we have

$$\begin{aligned} (V_1^* - V_1^\pi)(x_1^{(1)}) &\geq \mathcal{P}_1(\mathbf{a}_1^* | x_1^{(1)}, A_1^*) \bar{Q}_1^*(x_1^{(1)}, \mathbf{a}_1^*) - \mathcal{P}_1(\tilde{\mathbf{a}}_1 | x_1^{(1)}, \tilde{A}_1) \tilde{Q}_1^\pi(x_1^{(1)}, \mathbf{a}_1^*, \bar{\mathbf{a}}_1) \\ &= \left(\mathcal{P}_1(\mathbf{a}_1^* | x_1^{(1)}, A_1^*) - \mathcal{P}_1(\tilde{\mathbf{a}}_1 | x_1^{(1)}, \tilde{A}_1) \right) \bar{Q}_1^*(x_1^{(1)}, \mathbf{a}_1^*) \\ &\quad + \mathcal{P}_1(\tilde{\mathbf{a}}_1 | x_1^{(1)}, \tilde{A}_1) \left(\bar{Q}_1^*(x_1^{(1)}, \mathbf{a}_1^*) - \tilde{Q}_1^\pi(x_1^{(1)}, \mathbf{a}_1^*, \bar{\mathbf{a}}_1) \right) \\ &= \left(\mathcal{P}_1(\mathbf{a}_1^* | x_1^{(1)}, A_1^*) - \mathcal{P}_1(\tilde{\mathbf{a}}_1 | x_1^{(1)}, \tilde{A}_1) \right) \bar{Q}_1^*(x_1^{(1)}, \mathbf{a}_1^*) \\ &\quad + \mathcal{P}_1(\tilde{\mathbf{a}}_1 | x_1^{(1)}, \tilde{A}_1) \left(\frac{1}{H} + \mathbb{P}_1(x_2^{(1)} | x_1^{(1)}, \mathbf{a}_1^*) V_2^*(x_2^{(1)}) + \mathbb{P}_1(x_{H+2}^{(1)} | x_1^{(1)}, \mathbf{a}_1^*) \frac{(H-1)}{H} \right. \\ &\quad \left. - \left(\frac{1}{H} + \mathbb{P}_1(x_2^{(1)} | x_1^{(1)}, \mathbf{a}_1^*) V_2^\pi(x_2^{(1)}) + \mathbb{P}_1(s_{H+2} | x_1^{(1)}, \bar{\mathbf{a}}_1) \frac{(H-1)}{H} \right) \right) \\ &= \left(\mathcal{P}_1(\mathbf{a}_1^* | x_1^{(1)}, A_1^*) - \mathcal{P}_1(\tilde{\mathbf{a}}_1 | x_1^{(1)}, \tilde{A}_1) \right) \bar{Q}_1^*(x_1^{(1)}, \mathbf{a}_1^*) \\ &\quad + \mathcal{P}_1(\tilde{\mathbf{a}}_1 | x_1^{(1)}, \tilde{A}_1) \mathbb{P}_1(x_2^{(1)} | x_1^{(1)}, \mathbf{a}_1^*) (V_2^* - V_2^\pi)(x_2^{(1)}) \\ &\quad + \mathcal{P}_1(\tilde{\mathbf{a}}_1 | x_1^{(1)}, \tilde{A}_1) \left(\mathbb{P}_1(x_{H+2}^{(1)} | x_1^{(1)}, \mathbf{a}_1^*) - \mathbb{P}_1(s_{H+2} | x_1^{(1)}, \bar{\mathbf{a}}_1) \right) \frac{(H-1)}{H}, \end{aligned} \quad (\text{F.4})$$

where the first inequality holds by (F.3). Note that, by construction, for any $h \in [H]$, we have

$$\begin{aligned} \bar{Q}_h^*(x_h^{(1)}, \mathbf{a}_h^*) &= \frac{H-h+1}{H}, \\ \mathcal{P}_h(\tilde{\mathbf{a}}_h | x_h^{(1)}, \tilde{A}_h) &\geq \frac{1}{1/H+1} = \frac{H}{1+H}, \\ \mathbb{P}_h(x_{h+1}^{(1)} | x_h^{(1)}, \mathbf{a}_h^*) &= 1 - \delta - (d^{\text{lin}} - 5)\Delta, \\ \mathbb{P}_h(x_{H+2}^{(1)} | x_h^{(1)}, \mathbf{a}_h^*) - \mathbb{P}_h(s_{H+2} | x_h^{(1)}, \bar{\mathbf{a}}_h) &= (d^{\text{lin}} - 5)\Delta - \langle \boldsymbol{\mu}_h, \bar{\mathbf{a}}_h \rangle. \end{aligned} \quad (\text{F.5})$$

Hence, by plugging (F.5) into (F.4) and applying recursion, we get

$$\begin{aligned}
 & (V_1^\star - V_1^\pi)(x_1^{(1)}) \\
 & \geq \sum_{h=1}^H \left(\mathcal{P}_h(\mathbf{a}_h^\star | x_h^{(1)}, A_h^\star) - \mathcal{P}_h(\tilde{\mathbf{a}}_h | x_h^{(1)}, \tilde{A}_h) \right) \frac{H-h+1}{H} \cdot \left(\frac{H}{H+1} \right)^{h-1} \cdot ((1-\delta - (d^{\text{lin}} - 5)\Delta))^{h-1} \\
 & + \sum_{h=1}^H ((d^{\text{lin}} - 5)\Delta - \langle \boldsymbol{\mu}_h, \tilde{\mathbf{a}}_h \rangle) \frac{H-h}{H} \cdot \left(\frac{H}{H+1} \right)^h \cdot ((1-\delta - (d^{\text{lin}} - 5)\Delta))^{h-1} \\
 & - \sum_{h=1}^H \frac{(H-h)}{H\sqrt{K}} \cdot \left(\frac{H}{H+1} \right)^h \cdot ((1-\delta - (d^{\text{lin}} - 5)\Delta))^{h-1}.
 \end{aligned} \tag{F.6}$$

Furthermore, since $H \geq 3$ and $3(d^{\text{lin}} - 5)\Delta \leq \delta = 1/H$, we have

$$\begin{aligned}
 \left(\frac{H}{H+1} \right)^h & \geq \left(\frac{H}{H+1} \right)^{H+1} \geq \frac{3}{10}, \\
 ((1-\delta - (d^{\text{lin}} - 5)\Delta))^{h-1} & \geq \left(1 - \frac{4\delta}{3} \right)^H \geq \frac{1}{3}.
 \end{aligned} \tag{F.7}$$

Therefore, by substituting (F.7) into (F.6), and considering the terms where $h \geq H/2$, we obtain

$$\begin{aligned}
 (V_1^\star - V_1^\pi)(x_1^{(1)}) & \geq \frac{1}{20} \sum_{h=1}^{H/2} \left(\mathcal{P}_h(\mathbf{a}_h^\star | x_h^{(1)}, A_h^\star) - \mathcal{P}_h(\tilde{\mathbf{a}}_h | x_h^{(1)}, \tilde{A}_h) \right) + \frac{1}{20} \sum_{h=1}^{H/2} ((d^{\text{lin}} - 5)\Delta - \langle \boldsymbol{\mu}_h, \tilde{\mathbf{a}}_h \rangle) \\
 & = \frac{1}{20} \sum_{h=1}^{H/2} \underbrace{\left(\mathcal{P}_h(\mathbf{a}_h^\star | x_h^{(1)}, A_h^\star) - \mathcal{P}_h(\tilde{\mathbf{a}}_h | x_h^{(1)}, \tilde{A}_h) \right)}_{\text{MNL bandit regret}} + \frac{1}{20} \sum_{h=1}^{H/2} \underbrace{\left(\max_{\mathbf{a} \in \mathcal{I}} \langle \boldsymbol{\mu}_h, \mathbf{a} \rangle - \langle \boldsymbol{\mu}_h, \tilde{\mathbf{a}}_h \rangle \right)}_{\text{linear bandit regret}}.
 \end{aligned} \tag{F.8}$$

On the right-hand side of (F.8), the first term corresponds to an MNL bandit problem. Recall that $|A_h^\star| = |\tilde{A}_h| = 2$ and, by construction, we have

$$\mathcal{P}_h(\mathbf{a}_h^\star | x_h^{(1)}, A_h^\star) = \frac{\exp(\phi(x_h^{(1)}, \mathbf{a}_h^\star)^\top \boldsymbol{\theta}_h^\star)}{1/H + \exp(\phi(x_h^{(1)}, \mathbf{a}_h^\star)^\top \boldsymbol{\theta}_h^\star)}, \quad \mathcal{P}_h(\tilde{\mathbf{a}}_h | x_h^{(1)}, \tilde{A}_h) = \frac{\exp(\phi(x_h^{(1)}, \tilde{\mathbf{a}}_h)^\top \boldsymbol{\theta}_h^\star)}{1/H + \exp(\phi(x_h^{(1)}, \tilde{\mathbf{a}}_h)^\top \boldsymbol{\theta}_h^\star)}.$$

Hence, this corresponds to an MNL bandit problem with a maximum assortment size of $M = 2$, where the *attraction parameter for the outside option* (the constant in the denominator) is $1/H$.

Furthermore, the second term on the right-hand side of (F.8) represents a linear bandit problem. To sum up, the learning problem is not harder than minimizing the regret on $\Omega(H/2)$ MNL and linear bandit problems.

To bound each term of (F.8), we introduce the following propositions:

Proposition F.4 (Regret lower bound of MNL bandits, Lee & Oh 2024). *Let v_0 denote the attraction parameter for the outside option. Let d be divisible by 4. Suppose $K \geq C \cdot d^4 M / (M - 1)$ for some constant $C > 0$. Then, in the uniform reward setting (where rewards are identical) with the reward for the outside option being zero, for any policy and the MNL preference model parameterized by $\boldsymbol{\theta}$, there exists a worst-case problem instance such that the worst-case expected regret is lower bounded as follows:*

$$\sup_{\boldsymbol{\theta}} \mathbb{E}_{\boldsymbol{\theta}} [\text{MNLBanditRegret}(\boldsymbol{\theta}, K)] = \Omega \left(\frac{\sqrt{v_0(M-1)}}{v_0 + M - 1} \cdot d\sqrt{K} \right).$$

Proposition F.5 (Lemma C.8 in Zhou et al. 2021a). *Fix $0 < \delta < 1/3$. Consider the linear bandit problem parameterized with a vector $\boldsymbol{\mu} \in \{-\Delta, \Delta\}^d$ and action set $\mathcal{I} = \{-1, 1\}^d$. And the reward distribution for taking action $\mathbf{a} \in \mathcal{I}$ is a Bernoulli distribution denoted as $B(\delta + \langle \boldsymbol{\mu}, \mathbf{a} \rangle)$. Let K be the number of time steps playing this bandit problem. Assume $K \geq d^2 / (2\delta)$*

and $\Delta = \sqrt{\delta/K}/(4\sqrt{2})$. Then, for any bandit algorithm \mathcal{B} , there exists μ such that the expected pseudo-regret of \mathcal{B} over K steps is lower bounded as follows:

$$\mathbb{E}_\mu[\mathbf{LinearBanditRegret}(\mu, K)] \geq \frac{d\sqrt{K\delta}}{8\sqrt{2}}.$$

where the expectation is with respect to the reward distribution that depends on μ .

Now, by using Proposition F.4 and F.5, we can bound the regret as follows:

$$\begin{aligned} \sup_{\theta, \mu, \mathbf{w}} \mathbb{E}_{\theta, \mu, \mathbf{w}} [\mathbf{Regret}(\mathcal{M}_{\theta, \mu, \mathbf{w}}, K)] &\geq \frac{1}{20} \sum_{h=1}^{H/2} \sup_{\theta} \mathbb{E}_{\theta} \left[\sum_{k=1}^K \left(\mathcal{P}_h(\mathbf{a}_h^* | x_h^{(1)}, A_h^*) - \mathcal{P}_h(\tilde{\mathbf{a}}_h | x_h^{(1)}, \tilde{A}_h) \right) \right] \\ &\quad + \frac{1}{20} \sum_{h=1}^{H/2} \sup_{\mu} \mathbb{E}_{\mu} \left[\sum_{k=1}^K \left(\max_{\mathbf{a} \in \mathcal{I}} \langle \mu_h, \mathbf{a} \rangle - \langle \mu_h, \bar{\mathbf{a}}_h \rangle \right) \right] \\ &= \Omega \left(d\sqrt{HK} + d^{\text{lin}}\sqrt{HK} \right), \end{aligned}$$

where, in the last equality, we use $v_0 = 1/H$, $M = 2$, and $\delta = 1/H$. This concludes the proof of Theorem 5.3. \square

F.4. Proof of Lemmas for Theorem 5.3

F.4.1. PROOF OF LEMMA F.2

Proof of Lemma F.2. For any $i \in [H]$, we can write the optimal \bar{Q} -value in state $x_h^{(i)}$ at horizon $h \in [H]$ as follows:

$$\bar{Q}_h^*(x_h^{(i)}, \mathbf{a}) = \begin{cases} \frac{\gamma^{i-1}}{H} + \mathbb{P}_h(x_{h+1}^{(i)} | x_h^{(i)}, \mathbf{a}) V_{h+1}^*(x_{h+1}^{(i)}) + \mathbb{P}_h(x_{H+2}^{(i)} | x_h^{(i)}, \mathbf{a}) \frac{(H-h)\gamma^{i-1}}{H}, & \mathbf{a} = \mathbf{a}_h^*; \\ \frac{\gamma^i}{H} + \mathbb{P}_h(x_{h+1}^{(i+1)} | x_h^{(i)}, \mathbf{a}) V_{h+1}^*(x_{h+1}^{(i+1)}) + \mathbb{P}_h(x_{H+2}^{(i+2)} | x_h^{(i)}, \mathbf{a}) \frac{(H-h)\gamma^{i+1}}{H}, & \mathbf{a} = \mathbf{a}_h^*, \mathbf{a}_0; \\ 0, & \mathbf{a} = \mathbf{a}_0. \end{cases}$$

First, we show that for any $(i, h) \in [H] \times [H]$, we have

$$\bar{Q}_h^*(x_h^{(i)}, \mathbf{a}) \geq \sum_{\mathbf{a}' \in A_h^*} \mathcal{P}_h(\mathbf{a}' | x_h^{(i)}, A_h^*) \bar{Q}_h^*(x_h^{(i)}, \mathbf{a}'), \quad \forall \mathbf{a} \in A_h^* \setminus \{\mathbf{a}_0\}. \quad (\text{F.9})$$

We prove this by contradiction. Suppose there exists $\mathbf{a} \in A_h^*$ such that $\bar{Q}_h^*(x_h^{(i)}, \mathbf{a}) < \sum_{\mathbf{a}' \in A_h^*} \mathcal{P}_h(\mathbf{a}' | x_h^{(i)}, A_h^*) \bar{Q}_h^*(x_h^{(i)}, \mathbf{a}')$. In that case, removing the item \mathbf{a} from the assortment A_h^* results in a higher expected value of \bar{Q}_h^* . This contradicts the optimality of A_h^* . Therefore, (F.9) must hold.

By the definition of $\bar{Q}_h^*(x_h^{(i)}, \mathbf{a})$, for any $\mathbf{a} \in \mathcal{I} \setminus \{\mathbf{a}_h^*\}$, we have

$$\begin{aligned} \bar{Q}_h^*(x_h^{(i)}, \mathbf{a}) &\leq \frac{\gamma^{i-1}}{H} + \mathbb{P}_h(x_{h+1}^{(i+1)} | x_h^{(i)}, \mathbf{a}) V_{h+1}^*(x_{h+1}^{(i)}) + \mathbb{P}_h(x_{H+2}^{(i+2)} | x_h^{(i)}, \mathbf{a}) \frac{(H-h)\gamma^{i-1}}{H} \\ &\leq \frac{\gamma^{i-1}}{H} + \mathbb{P}_h(x_{h+1}^{(i)} | x_h^{(i)}, \mathbf{a}_h^*) V_{h+1}^*(x_{h+1}^{(i)}) + \mathbb{P}_h(x_{H+2}^{(i)} | x_h^{(i)}, \mathbf{a}_h^*) \frac{(H-h)\gamma^{i-1}}{H} \\ &= \bar{Q}_h^*(x_h^{(i)}, \mathbf{a}_h^*), \end{aligned}$$

where the first inequality holds since $V_{h+1}^*(x_{h+1}^{(i+1)}) \leq V_{h+1}^*(x_{h+1}^{(i)})$, and the second inequality holds due to the fact that $V_{h+1}^*(x_{h+1}^{(i)}) \leq \frac{(H-h)\gamma^{i-1}}{H}$ and $\mathbb{P}_h(x_{H+2}^{(i+2)} | x_h^{(i)}, \mathbf{a}) \leq \mathbb{P}_h(x_{H+2}^{(i)} | x_h^{(i)}, \mathbf{a}_h^*)$.

Since $\bar{Q}_h^*(x_h^{(i)}, \mathbf{a}_h^*)$ has the highest value among all items, the optimal assortment A_h^* should include \mathbf{a}_h^* . Thus, we have $\mathbf{a}_h^*, \mathbf{a}_0 \in A_h^*$. In other words, when $A_h^* = \{\mathbf{a}_h^*, \mathbf{a}_0\}$, the condition in (F.9) is satisfied. Thus, we begin with $A_h^* = \{\mathbf{a}_h^*, \mathbf{a}_0\}$ and check if there exist an item $\mathbf{a} \neq \mathbf{a}_h^*, \mathbf{a}_0$ that can increase the expected value of \bar{Q}_h^* . To this end, for $A_h^* = \{\mathbf{a}_h^*, \mathbf{a}_0\}$, we get

$$\sum_{\mathbf{a}' \in A_h^*} \mathcal{P}_h(\mathbf{a}' | x_h^{(i)}, A_h^*) \bar{Q}_h^*(x_h^{(i)}, \mathbf{a}') = \mathcal{P}_h(\mathbf{a}_h^* | x_h^{(i)}, A_h^*) \bar{Q}_h^*(x_h^{(i)}, \mathbf{a}_h^*) \geq \gamma \bar{Q}_h^*(x_h^{(i)}, \mathbf{a}_h^*), \quad (\text{F.10})$$

where the equality holds since $\overline{Q}_h^*(x_h^{(i)}, a_0) = 0$, and the inequality holds by the definition of γ :

$$\begin{aligned} \gamma &= \frac{H}{1+H} = \frac{1}{1/H+1} \leq \min_{h \in [H]} \min_{s \in S} \min_{A \in \mathcal{A}} \min_{\mathbf{a} \in A \setminus \{a_0\}} \frac{\exp(\phi(s, \mathbf{a})^\top \boldsymbol{\theta}_h^*)}{1/H + \exp(\phi(s, \mathbf{a})^\top \boldsymbol{\theta}_h^*)} \\ &\leq \frac{\exp(\phi(x_h^{(i)}, \mathbf{a}_h^*)^\top \boldsymbol{\theta}_h^*)}{1/H + \exp(\phi(x_h^{(i)}, \mathbf{a}_h^*)^\top \boldsymbol{\theta}_h^*)} = \mathcal{P}_h(\mathbf{a}_h^* | x_h^{(i)}, A_h^*). \end{aligned} \quad (\text{F.11})$$

Here, we rely on the fact that the sigmoid function is monotonically increasing to establish the inequalities.

On the other hand, for any item $\mathbf{a} \neq \mathbf{a}_h^*, \mathbf{a}_0$, we have

$$\begin{aligned} \overline{Q}_h^*(x_h^{(i)}, \mathbf{a}) &= \frac{\gamma^i}{H} + \mathbb{P}_h(x_{h+1}^{(i+1)} | x_h^{(i)}, \mathbf{a}) V_{h+1}^*(x_{h+1}^{(i+1)}) + \mathbb{P}_h(x_{H+2}^{(i+2)} | x_h^{(i)}, \mathbf{a}) \frac{(H-h)\gamma^{i+1}}{H} \\ &\leq \frac{\gamma^i}{H} + \mathbb{P}_h(x_{h+1}^{(i+1)} | x_h^{(i)}, \mathbf{a}) V_{h+1}^*(x_{h+1}^{(i+1)}) + \mathbb{P}_h(x_{H+2}^{(i+2)} | x_h^{(i)}, \mathbf{a}) \frac{(H-h)\gamma^i}{H} \\ &\leq \frac{\gamma^i}{H} + \mathbb{P}_h(x_{h+1}^{(i+1)} | x_h^{(i)}, \mathbf{a}_h^*) V_{h+1}^*(x_{h+1}^{(i+1)}) + \mathbb{P}_h(x_{H+2}^{(i+2)} | x_h^{(i)}, \mathbf{a}_h^*) \frac{(H-h)\gamma^i}{H} \\ &= \gamma \cdot \left(\frac{\gamma^{i-1}}{H} + \mathbb{P}_h(x_{h+1}^{(i)} | x_h^{(i)}, \mathbf{a}_h^*) V_{h+1}^*(x_{h+1}^{(i)}) + \mathbb{P}_h(x_{H+2}^{(i)} | x_h^{(i)}, \mathbf{a}_h^*) \frac{(H-h)\gamma^{i-1}}{H} \right) \\ &= \gamma \overline{Q}_h^*(x_h^{(i)}, \mathbf{a}_h^*), \end{aligned} \quad (\text{F.12})$$

where the second inequality holds because $V_{h+1}^*(x_{h+1}^{(i+1)}) \leq \frac{(H-h)\gamma^i}{H}$, $\mathbb{P}_h(x_{H+2}^{(i+2)} | x_h^{(i)}, \mathbf{a}) \leq \mathbb{P}_h(x_{H+2}^{(i+2)} | x_h^{(i)}, \mathbf{a}_h^*)$, and the second equality follows from the fact that $\gamma V_{h+1}^*(x_{h+1}^{(i+1)}) = V_{h+1}^*(x_{h+1}^{(i+1)})$ by construction.

Combining (F.10) and (F.12), when $A_h^* = \{\mathbf{a}_h^*, \mathbf{a}_0\}$, for any item $\mathbf{a} \neq \mathbf{a}_h^*, \mathbf{a}_0$, we get

$$\sum_{\mathbf{a}' \in A_h^*} \mathcal{P}_h(\mathbf{a}' | x_h^{(i)}, A_h^*) \overline{Q}_h^*(x_h^{(i)}, \mathbf{a}') \geq \overline{Q}_h^*(x_h^{(i)}, \mathbf{a}).$$

Since $\overline{Q}_h^*(x_h^{(i)}, \mathbf{a})$ for $\mathbf{a} \neq \mathbf{a}_h^*, \mathbf{a}_0$ is not greater than the expected value of \overline{Q}_h^* for A_h^* , adding any item $\mathbf{a} \neq \mathbf{a}_h^*, \mathbf{a}_0$ to A_h^* does not increase the expected value of \overline{Q}_h^* . This confirms the optimality of A_h^* . \square

F.4.2. PROOF OF LEMMA F.3

Proof of Lemma F.3. For any $i \in [H]$, we can write the \overline{Q} -value for the policy π in state $x_h^{(i)}$ at horizon $h \in [H]$ as follows:

$$\overline{Q}_h^\pi(x_h^{(i)}, \mathbf{a}) = \begin{cases} \frac{\gamma^{i-1}}{H} + \mathbb{P}_h(x_{h+1}^{(i)} | x_h^{(i)}, \mathbf{a}) V_{h+1}^\pi(x_{h+1}^{(i)}) + \mathbb{P}_h(x_{H+2}^{(i)} | x_h^{(i)}, \mathbf{a}) \frac{(H-h)\gamma^{i-1}}{H}, & \mathbf{a} = \mathbf{a}_h^*; \\ \frac{\gamma^i}{H} + \mathbb{P}_h(x_{h+1}^{(i+1)} | x_h^{(i)}, \mathbf{a}) V_{h+1}^\pi(x_{h+1}^{(i+1)}) + \mathbb{P}_h(x_{H+2}^{(i+2)} | x_h^{(i)}, \mathbf{a}) \frac{(H-h)\gamma^{i+1}}{H}, & \mathbf{a} = \mathbf{a}_h^*, \mathbf{a}_0; \\ 0, & \mathbf{a} = \mathbf{a}_0. \end{cases}$$

We provide a proof by considering the following cases:

Case (i) $\mathbf{a}_h^* \in A$.

Recall that, by (F.11), we have

$$\gamma = \frac{H}{1+H} \leq \frac{\exp(\phi(x_h^{(i)}, \mathbf{a}_h^*)^\top \boldsymbol{\theta}_h^*)}{1/H + \exp(\phi(x_h^{(i)}, \mathbf{a}_h^*)^\top \boldsymbol{\theta}_h^*)}. \quad (\text{F.13})$$

By multiplying $\tilde{Q}_h^\pi(x_h^{(i)}, \mathbf{a}_h^*, \mathbf{a}')$ on both sides of (F.13), we get

$$\begin{aligned}
 \gamma \cdot \tilde{Q}_h^\pi(x_h^{(i)}, \mathbf{a}_h^*, \mathbf{a}') &\leq \frac{\exp\left(\phi(x_h^{(i)}, \mathbf{a}_h^*)^\top \boldsymbol{\theta}_h^*\right) \tilde{Q}_h^\pi(x_h^{(i)}, \mathbf{a}_h^*, \mathbf{a}')}{1/H + \exp\left(\phi(x_h^{(i)}, \mathbf{a}_h^*)^\top \boldsymbol{\theta}_h^*\right)} \\
 &\Leftrightarrow \left(\sum_{\mathbf{a} \in A \setminus \{\mathbf{a}_h^*, \mathbf{a}_0\}} \exp\left(\phi(x_h^{(i)}, \mathbf{a})^\top \boldsymbol{\theta}_h^*\right) \right) \gamma \cdot \tilde{Q}_h^\pi(x_h^{(i)}, \mathbf{a}_h^*, \mathbf{a}') \left(1/H + \exp\left(\phi(x_h^{(i)}, \mathbf{a}_h^*)^\top \boldsymbol{\theta}_h^*\right)\right) \\
 &\leq \left(\sum_{\mathbf{a} \in A \setminus \{\mathbf{a}_h^*, \mathbf{a}_0\}} \exp\left(\phi(x_h^{(i)}, \mathbf{a})^\top \boldsymbol{\theta}_h^*\right) \right) \exp\left(\phi(x_h^{(i)}, \mathbf{a}_h^*)^\top \boldsymbol{\theta}_h^*\right) \tilde{Q}_h^\pi(x_h^{(i)}, \mathbf{a}_h^*, \mathbf{a}') \\
 &\Leftrightarrow \left(\exp\left(\phi(x_h^{(i)}, \mathbf{a}_h^*)^\top \boldsymbol{\theta}_h^*\right) + \sum_{\mathbf{a} \in A \setminus \{\mathbf{a}_h^*, \mathbf{a}_0\}} \gamma \cdot \exp\left(\phi(x_h^{(i)}, \mathbf{a})^\top \boldsymbol{\theta}_h^*\right) \right) \tilde{Q}_h^\pi(x_h^{(i)}, \mathbf{a}_h^*, \mathbf{a}') \\
 &\quad \cdot \left(1/H + \exp\left(\phi(x_h^{(i)}, \mathbf{a}_h^*)^\top \boldsymbol{\theta}_h^*\right)\right) \\
 &\leq \left(1/H + \exp\left(\phi(x_h^{(i)}, \mathbf{a}_h^*)^\top \boldsymbol{\theta}_h^*\right) + \sum_{\mathbf{a} \in A \setminus \{\mathbf{a}_h^*, \mathbf{a}_0\}} \exp\left(\phi(x_h^{(i)}, \mathbf{a})^\top \boldsymbol{\theta}_h^*\right) \right) \exp\left(\phi(x_h^{(i)}, \mathbf{a}_h^*)^\top \boldsymbol{\theta}_h^*\right) \\
 &\quad \cdot \tilde{Q}_h^\pi(x_h^{(i)}, \mathbf{a}_h^*, \mathbf{a}') \\
 &\Leftrightarrow \frac{\exp\left(\phi(x_h^{(i)}, \mathbf{a}_h^*)^\top \boldsymbol{\theta}_h^*\right) \tilde{Q}_h^\pi(x_h^{(i)}, \mathbf{a}_h^*, \mathbf{a}') + \sum_{\mathbf{a} \in A \setminus \{\mathbf{a}_h^*, \mathbf{a}_0\}} \exp\left(\phi(x_h^{(i)}, \mathbf{a})^\top \boldsymbol{\theta}_h^*\right) \gamma \tilde{Q}_h^\pi(x_h^{(i)}, \mathbf{a}_h^*, \mathbf{a}')}{1/H + \sum_{\mathbf{a} \in A \setminus \{\mathbf{a}_0\}} \exp\left(\phi(x_h^{(i)}, \mathbf{a})^\top \boldsymbol{\theta}_h^*\right)} \\
 &\leq \frac{\exp\left(\phi(x_h^{(i)}, \mathbf{a}_h^*)^\top \boldsymbol{\theta}_h^*\right) \tilde{Q}_h^\pi(x_h^{(i)}, \mathbf{a}_h^*, \mathbf{a}')}{1/H + \exp\left(\phi(x_h^{(i)}, \mathbf{a}_h^*)^\top \boldsymbol{\theta}_h^*\right)}. \tag{F.14}
 \end{aligned}$$

On the other hand, by the definition of $\tilde{Q}_h^\pi(x_h^{(i)}, \mathbf{a}_h^*, \mathbf{a}')$, for any $\mathbf{a}' \neq \mathbf{a}_h^*, \mathbf{a}_0$, we have

$$\begin{aligned}
 \gamma \tilde{Q}_h^\pi(x_h^{(i)}, \mathbf{a}_h^*, \mathbf{a}') &= \gamma \cdot \left(\frac{\gamma^{i-1}}{H} + \mathbb{P}_h(x_{h+1}^{(i)} | x_h^{(i)}, \mathbf{a}_h^*) V_{h+1}^\pi(x_{h+1}^{(i)}) + \mathbb{P}_h(x_{H+2}^{(i+2)} | x_h^{(i)}, \mathbf{a}') \frac{(H-h)\gamma^{i-1}}{H} \right) \\
 &= \frac{\gamma^i}{H} + \mathbb{P}_h(x_{h+1}^{(i)} | x_h^{(i)}, \mathbf{a}_h^*) V_{h+1}^\pi(x_{h+1}^{(i+1)}) + \mathbb{P}_h(x_{H+2}^{(i+2)} | x_h^{(i)}, \mathbf{a}') \frac{(H-h)\gamma^i}{H} \\
 &\geq \frac{\gamma^i}{H} + \mathbb{P}_h(x_{h+1}^{(i)} | x_h^{(i)}, \mathbf{a}') V_{h+1}^\pi(x_{h+1}^{(i+1)}) + \mathbb{P}_h(x_{H+2}^{(i+2)} | x_h^{(i)}, \mathbf{a}') \frac{(H-h)\gamma^{i+1}}{H} \\
 &= \bar{Q}_h^\pi(x_h^{(i)}, \mathbf{a}'), \tag{F.15}
 \end{aligned}$$

where the second equality holds since $\gamma V_{h+1}^\pi(x_{h+1}^{(i)}) = V_{h+1}^\pi(x_{h+1}^{(i+1)})$, and the inequality holds because, for $K \geq 4(d^{\text{lin}} - 5)^2 H(H+1)^2$, the following inequality holds:

$$\begin{aligned}
 &\mathbb{P}_h(x_{h+1}^{(i)} | x_h^{(i)}, \mathbf{a}') V_{h+1}^\pi(x_{h+1}^{(i+1)}) + \mathbb{P}_h(x_{H+2}^{(i+2)} | x_h^{(i)}, \mathbf{a}') \frac{(H-h)\gamma^{i+1}}{H} \\
 &\leq \mathbb{P}_h(x_{h+1}^{(i)} | x_h^{(i)}, \mathbf{a}_h^*) V_{h+1}^\pi(x_{h+1}^{(i+1)}) + \mathbb{P}_h(x_{H+2}^{(i+2)} | x_h^{(i)}, \mathbf{a}') \frac{(H-h)\gamma^i}{H} \\
 &\Leftrightarrow \left(\mathbb{P}_h(x_{h+1}^{(i)} | x_h^{(i)}, \mathbf{a}') - \mathbb{P}_h(x_{h+1}^{(i)} | x_h^{(i)}, \mathbf{a}_h^*) \right) V_{h+1}^\pi \leq \mathbb{P}_h(x_{H+2}^{(i+2)} | x_h^{(i)}, \mathbf{a}') \frac{(H-h)}{H} (\gamma^i - \gamma^{i+1}).
 \end{aligned}$$

Specifically, if the upper bound of the left-hand side is less than or equal to the lower bound of the right-hand side, the inequality holds. To demonstrate this, we have:

$$\left(\mathbb{P}_h(x_{h+1}^{(i)} | x_h^{(i)}, \mathbf{a}') - \mathbb{P}_h(x_{h+1}^{(i)} | x_h^{(i)}, \mathbf{a}_h^*) \right) V_{h+1}^\pi \leq 2(d^{\text{lin}} - 5)\Delta \cdot \frac{(H-h)}{H}, \tag{F.16}$$

and, since $\gamma^i = \left(\frac{H}{H+1}\right)^i \geq \left(\frac{H}{H+1}\right)^{H+1} \geq \frac{3}{10}$, we get

$$\begin{aligned} \mathbb{P}_h(x_{H+2}^{(i+2)} | x_h^{(i)}, \mathbf{a}') \frac{(H-h)}{H} (\gamma^i - \gamma^{i+1}) &\geq (\delta - (d^{\text{lin}} - 5)\Delta) \frac{(H-h)}{H} \gamma^i (1 - \gamma) \\ &\geq \left(\frac{1}{H} - (d^{\text{lin}} - 5)\Delta\right) \frac{(H-h)}{H} \cdot \frac{3}{10} \cdot \frac{1}{H+1}. \end{aligned} \quad (\text{F.17})$$

Combining (F.16) and (F.17), and rearranging the terms, we get

$$(d^{\text{lin}} - 5)\Delta \cdot \left(2 + \frac{3}{10(H+1)}\right) \leq \frac{3}{10H(H+1)},$$

which holds when $K \geq 4(d^{\text{lin}} - 5)^2 H(H+1)^2$. This explains how the inequality in (F.15) is satisfied.

Let $\bar{\mathbf{a}}_h^{(i)} \in \operatorname{argmax}_{\mathbf{a} \in A_h \setminus \{\mathbf{a}_0\}} \bar{Q}_h^\pi(x_h^{(i)}, \mathbf{a}) = \mathbf{a}_h^*$. Note that $\bar{\mathbf{a}}_h^{(i)}$ is unique due to the way the action space and transition probabilities are constructed. Then, by combining (F.14) and (F.15), and using the fact that $\bar{Q}_h^\pi(x_h^{(i)}, \mathbf{a}_0) = 0$, we obtain that

$$\begin{aligned} \sum_{\mathbf{a} \in A} \mathcal{P}_h(\mathbf{a} | x_h^{(i)}, A) \bar{Q}_h^\pi(x_h^{(i)}, \mathbf{a}) &\leq \sum_{\mathbf{a} \in A} \mathcal{P}_h(\mathbf{a} | x_h^{(i)}, A) \bar{Q}_h^\pi(x_h^{(i)}, \bar{\mathbf{a}}_h^{(i)}) \\ &\leq \frac{\exp\left(\phi(x_h^{(i)}, \mathbf{a}_h^*)^\top \boldsymbol{\theta}_h^*\right) \tilde{Q}_h^\pi(x_h^{(i)}, \mathbf{a}_h^*, \bar{\mathbf{a}}_h^{(i)})}{1/H + \exp\left(\phi(x_h^{(i)}, \mathbf{a}_h^*)^\top \boldsymbol{\theta}_h^*\right)} \\ &= \frac{\max_{\mathbf{a} \in A \setminus \{\mathbf{a}_0\}} \exp\left(\phi(x_h^{(i)}, \mathbf{a})^\top \boldsymbol{\theta}_h^*\right) \tilde{Q}_h^\pi(x_h^{(i)}, \mathbf{a}_h^*, \bar{\mathbf{a}}_h^{(i)})}{1/H + \max_{\mathbf{a} \in A \setminus \{\mathbf{a}_0\}} \exp\left(\phi(x_h^{(i)}, \mathbf{a})^\top \boldsymbol{\theta}_h^*\right)}, \end{aligned}$$

where the first inequality holds since $\bar{\mathbf{a}}_h^{(i)}$ is the action that maximizes the \bar{Q} -value. The second inequality follows from (F.14) and (F.15), and from the fact that $\bar{Q}_h^\pi(x_h^{(i)}, \bar{\mathbf{a}}_h^{(i)}) = \bar{Q}_h^\pi(x_h^{(i)}, \mathbf{a}_h^*) = \tilde{Q}_h^\pi(x_h^{(i)}, \mathbf{a}_h^*, \bar{\mathbf{a}}_h^{(i)})$. Finally, the last equality holds by the definition of \mathbf{a}_h^* .

Case (ii) $\mathbf{a}_h^* \notin A$.

Again, by (F.11), for any $A \in \mathcal{A}$, we have

$$\begin{aligned} \gamma &\leq \min_{\mathbf{a} \in A \setminus \{\mathbf{a}_0\}} \frac{\exp\left(\phi(x_h^{(i)}, \mathbf{a})^\top \boldsymbol{\theta}_h^*\right)}{1/H + \exp\left(\phi(x_h^{(i)}, \mathbf{a})^\top \boldsymbol{\theta}_h^*\right)} \\ &\leq \frac{\max_{\mathbf{a} \in A \setminus \{\mathbf{a}_0\}} \exp\left(\phi(x_h^{(i)}, \mathbf{a})^\top \boldsymbol{\theta}_h^*\right)}{1/H + \max_{\mathbf{a} \in A \setminus \{\mathbf{a}_0\}} \exp\left(\phi(x_h^{(i)}, \mathbf{a})^\top \boldsymbol{\theta}_h^*\right)} \\ &\leq \frac{\max_{\mathbf{a} \in A \setminus \{\mathbf{a}_0\}} \exp\left(\phi(x_h^{(i)}, \mathbf{a})^\top \boldsymbol{\theta}_h^*\right)}{1/H + \max_{\mathbf{a} \in A \setminus \{\mathbf{a}_0\}} \exp\left(\phi(x_h^{(i)}, \mathbf{a})^\top \boldsymbol{\theta}_h^*\right)} \cdot \frac{1/H + \sum_{\mathbf{a} \in A \setminus \{\mathbf{a}_0\}} \exp\left(\phi(x_h^{(i)}, \mathbf{a})^\top \boldsymbol{\theta}_h^*\right)}{\sum_{\mathbf{a} \in A \setminus \{\mathbf{a}_0\}} \exp\left(\phi(x_h^{(i)}, \mathbf{a})^\top \boldsymbol{\theta}_h^*\right)}, \end{aligned}$$

where the second inequality holds since the sigmoid function is a monotonically increasing function.

We denote $\bar{\mathbf{a}}_h^{(i)} \in \operatorname{argmax}_{\mathbf{a} \in A_h \setminus \{\mathbf{a}_0\}} \bar{Q}_h^\pi(x_h^{(i)}, \mathbf{a})$. Then, multiplying $\tilde{Q}_h^\pi(x_h^{(i)}, \mathbf{a}_h^*, \bar{\mathbf{a}}_h^{(i)})$ (note that $\bar{\mathbf{a}}_h^{(i)} \neq \mathbf{a}_h^*, \mathbf{a}_0$) on both sides and rearranging terms, we get

$$\begin{aligned} &\frac{\sum_{\mathbf{a} \in A \setminus \{\mathbf{a}_0\}} \exp\left(\phi(x_h^{(i)}, \mathbf{a})^\top \boldsymbol{\theta}_h^*\right) \gamma \cdot \tilde{Q}_h^\pi(x_h^{(i)}, \mathbf{a}_h^*, \bar{\mathbf{a}}_h^{(i)})}{1/H + \sum_{\mathbf{a} \in A \setminus \{\mathbf{a}_0\}} \exp\left(\phi(x_h^{(i)}, \mathbf{a})^\top \boldsymbol{\theta}_h^*\right)} \\ &\leq \frac{\max_{\mathbf{a} \in A \setminus \{\mathbf{a}_0\}} \exp\left(\phi(x_h^{(i)}, \mathbf{a})^\top \boldsymbol{\theta}_h^*\right) \tilde{Q}_h^\pi(x_h^{(i)}, \mathbf{a}_h^*, \bar{\mathbf{a}}_h^{(i)})}{1/H + \max_{\mathbf{a} \in A \setminus \{\mathbf{a}_0\}} \exp\left(\phi(x_h^{(i)}, \mathbf{a})^\top \boldsymbol{\theta}_h^*\right)}. \end{aligned}$$

Recall that for any $\mathbf{a}' \neq \mathbf{a}_h^*, \mathbf{a}_0$, we have $\gamma \cdot \tilde{Q}_h^\pi(x_h^{(i)}, \mathbf{a}_h^*, \mathbf{a}') \geq \bar{Q}_h^\pi(x_h^{(i)}, \mathbf{a}')$ by (F.15). Thus, we get

$$\begin{aligned} \sum_{\mathbf{a} \in A} \mathcal{P}_h(\mathbf{a} | x_h^{(i)}, A) \bar{Q}_h^\pi(x_h^{(i)}, \mathbf{a}) &\leq \sum_{\mathbf{a} \in A} \mathcal{P}_h(\mathbf{a} | x_h^{(i)}, A) \bar{Q}_h^\pi(x_h^{(i)}, \bar{\mathbf{a}}_h^{(i)}) \\ &\leq \sum_{\mathbf{a} \in A} \mathcal{P}_h(\mathbf{a} | x_h^{(i)}, A) \gamma \cdot \tilde{Q}_h^\pi(x_h^{(i)}, \mathbf{a}_h^*, \bar{\mathbf{a}}_h^{(i)}) \\ &\leq \frac{\max_{\mathbf{a} \in A \setminus \{\mathbf{a}_0\}} \exp\left(\phi(x_h^{(i)}, \mathbf{a})^\top \boldsymbol{\theta}_h^*\right) \tilde{Q}_h^\pi(x_h^{(i)}, \mathbf{a}_h^*, \bar{\mathbf{a}}_h^{(i)})}{1/H + \max_{\mathbf{a} \in A \setminus \{\mathbf{a}_0\}} \exp\left(\phi(x_h^{(i)}, \mathbf{a})^\top \boldsymbol{\theta}_h^*\right)}. \end{aligned}$$

This concludes the proof of Lemma F.3. \square

G. Numerical Experiments




Figure G.1: The “online shopping with budget” environment with $|\mathcal{S}| = 5$. Each state represents the user’s budget level of 1, 2, 3, 4, or 5. The solid line indicates the transition when the user purchases an actual item a_i (with a reward of $(i/100N + j/|\mathcal{S}|)/H$), and the dashed line shows the transition when the user does not purchase any item (with a reward of 0). The initial state is s_3 .

In this section, we empirically evaluate the performance of our algorithm, MNL-VQL, in linear MDPs. We consider an *online shopping with budget* (refer Figure G.1) environment under linear MDPs and an MNL user preference model. We denote the set of states as $\mathcal{S} = \{s_1, \dots, s_{|\mathcal{S}|}\}$ and the set of items as $\mathcal{I} = \{a_1, \dots, a_N, a_0\}$ (a_0 denotes the outside option). Each state $s_j \in \mathcal{S}$ corresponds to a *user’s budget level*, where a larger index j indicates a higher budget (e.g., $s_{|\mathcal{S}|}$ represents the state with the largest budget). The initial state is set to the medium budget state $s_{\lfloor |\mathcal{S}|/2 \rfloor}$. Furthermore, we let the transition probabilities \mathbb{P}_h , rewards r_h , and preference model \mathcal{P}_h be the same for all $h \in [H]$, and thus we omit the subscript h .

At state s_j , the agent offers an assortment $A \in \mathcal{A}$ with a maximum size of M . The user then either purchases an item $a_i \in A$ or opts not to buy anything, represented by the outside option $a_0 \in A$. Then, the reward is defined as follows:

- If the user purchases an item $a_i \in A$, the reward is: $r(s_j, a_i) = \left(\frac{i}{100N} + \frac{j}{|\mathcal{S}|}\right)/H$.
- If the user does not buy anything (a_0), the reward is: $r(s_j, a_0) = 0$.

The reward can be regarded as the *user’s rating* of the purchased item. It is reasonable to assume that, at higher budget states, users tend to be more generous in their ratings, leading to higher ratings (rewards). And the transition probability is defined as follows:

- If the user purchases an item $a_i \in A$, the transition probability is:

$$\mathbb{P}(s_{\min(j+1, |\mathcal{S}|)} | s_j, a_i) = 1 - \frac{i}{N}, \quad \text{and} \quad \mathbb{P}(s_{\max(j-1, 0)} | s_j, a_i) = \frac{i}{N}.$$

- If the user does not buy anything (a_0), the transition probability is:

$$\mathbb{P}(s_{\min(j+1, |\mathcal{S}|)} | s_j, a_0) = 1$$

If the user does purchase an item, the budget level decreases with a certain probability that depends on the chosen item. Conversely, if the user does not purchase any item (a_0), the budget level increases deterministically.

	Myopic	LSVI-UCB	MNL-VQL (ours)
$N = 10, \mathcal{A} = 637$	0.089 s	0.136 s	0.463 s
$N = 20, \mathcal{A} = 21,699$	0.097 s	4.861 s	0.526 s
$N = 40, \mathcal{A} = 760,098$	0.113 s	453.641 s	0.620 s

 Table G.1: Average runtime (seconds) per episode for $M = 6$.

We construct the feature map $\psi(s, a)$ (for linear MDPs) using SVD. Specifically, the transition kernel $\mathbb{P}(\cdot|\cdot, \cdot) \in \mathbb{R}^{|\mathcal{S}||\mathcal{I}| \times |\mathcal{S}|}$ has at most $|\mathcal{S}|$ singular values, and the reward vector $r(\cdot, \cdot) \in \mathbb{R}^{|\mathcal{S}||\mathcal{I}|}$ has one singular value. Consequently, the feature map $\psi(s, a) \in \mathbb{R}^{d_{lin}}$ lies in a space of dimension $|\mathcal{S}| + 1$, i.e., $d_{lin} = |\mathcal{S}| + 1$.

For MNL preference model, the true parameter $\theta^* \in \mathbb{R}^d$, and the feature $\phi(s, a) \in \mathbb{R}^d$ (for MNL preference model) are randomly sampled from a d -dimensional uniform distribution in each instance.

We set $K = 30000$, $H = 5$, $M = 6$, $|\mathcal{S}| = 5$, $d = 5$ (feature dimension for MNL preference model), $d^{lin} = 6$ (feature dimension for linear MDP), $N \in \{10, 20, 40\}$ (the number of items), and $|\mathcal{A}| = \sum_{m'=1}^{M-1} \binom{N}{m} \in \{637, 21699, 760098\}$ (the number of assortments). Moreover, for simplicity, we set $\bar{\sigma}_h^k = 1$ in our algorithm. As a result, we use unweighted regression to estimate the \bar{Q} -values.

We compare our algorithm with two baselines: **Myopic** and **LSVI-UCB** (Jin et al., 2020). **Myopic** is a variant of OFU-MNL+ (Lee & Oh, 2024) adapted for *unknown* rewards. It is a myopic algorithm that selects assortments based only on immediate rewards, ignoring state transitions. **LSVI-UCB** (Jin et al., 2020) treats each assortment as a single, atomic (holistic) action, requiring enumeration of all possible assortments. To demonstrate the effectiveness of our approach, we also include the performance of the optimal policy (**Optimal**) to highlight that our algorithm is converging toward optimality. We run the algorithms on 10 independent instances and report the episodic return across all episodes.

Figure 1 demonstrates that our algorithm significantly outperforms other baseline algorithms. And Table G.1 shows that our algorithm maintains robust runtime performance even as the total number of assortments $|\mathcal{A}|$ increases. Although the runtime of **Myopic** is approximately 5.3 times faster than ours, its performance is substantially worse, converging to a suboptimal solution. This underscores a key limitation of the myopic strategy—it can completely fail in certain environments, highlighting the importance of accounting for long-term outcomes. Additionally, the runtime of **LSVI-UCB** increases exponentially as N grows, because it requires enumerating all possible assortments. Due to the extremely slow runtime of **LSVI-UCB**, we did not include its performance results for $N = 20$ and $N = 40$. Instead, for these cases, we used dotted lines to represent the average episodic return observed for $N = 10$. Even for the smaller case of $N = 10$, **LSVI-UCB** demonstrated the worst performance. Based on this observation, we suspect that its performance is unlikely to improve as N increases.