
Product Ranking for Revenue Maximization with Multiple Purchases

Renzhe Xu¹, Xingxuan Zhang¹, Bo Li², Yafeng Zhang³, Xiaolong Chen³, Peng Cui^{1*}

¹Department of Computer Science and Technology, Tsinghua University, Beijing, China

²School of Economics and Management, Tsinghua University, Beijing, China

³Meituan Inc., Beijing, China

xrz199721gmail.com, xingxuanzhang@hotmail.com, libo@sem.tsinghua.edu.cn
yafengz@outlook.com, chenxiaolong13@meituan.com, cuip@tsinghua.edu.cn

Abstract

Product ranking is the core problem for revenue-maximizing online retailers. To design proper product ranking algorithms, various consumer choice models are proposed to characterize the consumers' behaviors when they are provided with a list of products. However, existing works assume that each consumer purchases at most one product or will keep viewing the product list after purchasing a product, which does not agree with the common practice in real scenarios. In this paper, we assume that each consumer can purchase multiple products at will. To model consumers' willingness to view and purchase, we set a random attention span and purchase budget, which determines the maximal amount of products that he/she views and purchases, respectively. Under this setting, we first design an optimal ranking policy when the online retailer can precisely model consumers' behaviors. Based on the policy, we further develop the Multiple-Purchase-with-Budget UCB (MPB-UCB) algorithms with $\tilde{O}(\sqrt{T})$ regret that estimate consumers' behaviors and maximize revenue simultaneously in online settings. Experiments on both synthetic and semi-synthetic datasets prove the effectiveness of the proposed algorithms. The source code is available at <https://github.com/windxrz/MPB-UCB>.

1 Introduction

Online retailing has become increasingly popular over the last decades [17, 28, 52]. The way of product ranking is the crux for online retailers because it determines the consumers' shopping behaviors [17] and thus influences the retailers' revenue [20, 49]. For instance, the probability of consumers' purchasing from a firm or clicking an advertisement is strongly related to the display order [8, 3, 33]. Therefore, it is crucial for online retailers to design proper product ranking algorithms for revenue maximization.

The consumer choice model, when faced with a list of products, is the basis of the problem [17]. Generally speaking, existing works assume that the consumers view the list of products sequentially and select products according to a choice model [43, 42, 32, 38, 18, 26, 25, 37]. However, most of them assume that each consumer purchases at most one product and stop browsing immediately afterward, which does not agree with the common practice since a consumer may expect multiple purchases on a website or an app [34, 35, 16, 28]. Recently, several works consider a multiple-purchase setting while their targets are not revenue maximization [28, 15], or they assume that consumers always continue viewing the product list after purchasing any product, which is not practical in real scenarios [40].

*Corresponding Author

In this paper, we propose a more realistic consumer choice model to characterize consumer behaviors under multiple-purchase settings. Following [13, 17, 40], we assume that the consumers view the product list sequentially and purchase a product when its utility exceeds a certain threshold. Similar to [17], we assume the existence of the attention span for each consumer, which determines the maximal amount of products that the consumer is willing to view. More importantly, instead of forcing the consumer to quit viewing once she purchases one product, we allow each consumer to conduct multiple purchases at will and explore product ranking models for revenue maximization. Specifically, we constrain each consumer with a purchase budget to determine the maximal amount of purchases. Both the attention span and the purchase budget are assumed to obey the geometric distribution [13, 40].

Under this setting, we study the optimal product ranking policy to maximize online retailers' revenue. We first characterize the optimal ranking policy that achieves the maximal revenue when the online retailer knows the consumer parameters, including the purchase probabilities for each product and the distribution of the attention span and purchase budget. Afterward, we investigate the problem in online learning frameworks where consumer parameters are unknown to online retailers. In both non-contextual (*i.e.*, all consumers share the same parameters) and contextual settings (*i.e.*, consumers have personalized parameters), we provide the Multiple-Purchase-with-Budget UCB (MPB-UCB) algorithms that could fit consumers' parameters and maximize revenue simultaneously. We further theoretically prove that the algorithms achieve $\tilde{O}(\sqrt{T})$ regrets in both settings. Finally, we conduct experiments to verify the effectiveness of the algorithms.

To conclude, our contributions are highlighted as follows:

- We propose a novel consumer choice model to deal with multiple-purchase activities of consumers in online retailing scenarios.
- We characterize the optimal ranking policy and further design the MPB-UCB algorithms with $\tilde{O}(\sqrt{T})$ regret based on them in both non-contextual and contextual online settings.
- We conduct experiments on both synthetic and semi-synthetic datasets to prove the effectiveness of our method in the multiple-purchase setting.

2 Related Works

Consumer choice model Traditional choice models adopt the multinomial logit (MNL) [12, 22, 43] models to characterize consumers' behaviors. However, these models can not describe the behaviors when the consumers are provided with a list of products. As a result, several choice models are proposed recently. In detail, Abeliuk et al. [2] add a position bias term in the traditional MNL model. Gallego et al. [30], Aouad and Segev [4] suppose that the products are framed into a set of virtual web pages. Consumers select a product, if any, from these pages following a general choice model (such as MNL). Flores et al. [29] propose the Sequential Multinomial Logit (SML), which assumed a two-stage MNL model and supposed that products are divided into two stages with prior knowledge. Liu et al. [41], Feldman and Segev [27], Gao et al. [31] further extend the setting to multiple stages. These works suppose that each stage contains multiple products while recently several works assume that a product is treated as a single stage [17, 13, 14]. These models are also closely related to the cascade model [12, 36, 38, 54, 19]. Gao et al. [32] further propose the general cascade click models. Following these works, we suppose that each stage contains only one product and the consumers view the products sequentially.

In addition, several works consider the setting where consumers may have multiple interactions with the platforms. However, most of them focus on the click model [15, 16, 34–36] or MNL-based consumer choice models [48, 53, 50]. In addition, Ferreira et al. [28] maximize the number of consumers who engage with the site. Liang et al. [40] aim to maximize the revenue in the multiple-purchase setting while they assume that consumers always continue viewing the product list after purchasing any product, which is not practical in real scenarios.

Online learning and multi-armed bandit Online retailers need to design proper algorithms to model consumers' behaviors and maximize revenue at the same time, which is closely related to the multi-armed bandit framework [47, 45]. Many works [5, 9, 11, 10, 24] have been proposed to deal with demand learning or price experimentation problems with the help of the framework in the field

of revenue management. Recently, more methods [40, 17, 13, 14, 6, 44] adopt similar techniques in the product ranking setting, which inspires the algorithm proposed in this paper.

3 Preliminaries

Let $[N] \triangleq \{1, 2, \dots, N\}$ be the set of all products owned by an online retailer (he). The product k generates revenue r_k when a consumer purchases it. Let $r_{\max} \triangleq \max_{k \in [N]} r_k$ be the maximal revenue for all products and $\mathbf{r} = \{r_k\}_{k=1}^N$. We use the notation $\sigma_t \dashv [N]$ to denote the ranking policy σ_t on the products for the t -th consumer (she). More precisely, $\forall k \in [N], \sigma_{t,k} \in [N]$ represents the product displayed in the k -th position for her. Reversely, $\sigma_{t,k}^{-1}$ denotes the position of the k -th product in the ranking list. Furthermore, for any subset $Q \subseteq [N]$ that represents the possible positions in the ranking list, we use $\sigma_{t,Q}$ to denote the corresponding set of the products *w.r.t.* the position set Q , *i.e.*, $\sigma_{t,Q} = \{\sigma_{t,k} : k \in Q\}$. In addition, we use $\max(\cdot)$ to denote the maximal element in a set.

Consumer choice model for multiple purchases When the t -th consumer arrives, she views the list of products sequentially. We assume that the behaviors of different consumers are independent. The probability of product k being purchased by her is denoted as $\lambda_{t,k}$ and let $\boldsymbol{\lambda}_t = \{\lambda_{t,k}\}_{k=1}^N$. We assume for tractability that consumer interests in any products are independent, which is common in product ranking papers [4, 18, 23, 49]. This indicates that the purchase behavior of any product set does not affect the purchase probabilities of other products.

To characterize the consumers' behaviors under multiple-purchase settings, we assume that the consumer is endowed with the attention span V_t and purchase budget B_t , which determine the maximal number of the products that she will view and purchase, respectively. In reality, the consumer always has random attention span and purchase budget [17]. Therefore, following [13, 40], we assume $V_t \sim \text{Geo}(1 - q_t)$ and $B_t \sim \text{Geo}(1 - s_t)$ are independent geometric distributions with parameters $1 - q_t$ and $1 - s_t$, respectively. As a result, $P(V_t = v) = q_t^{v-1}(1 - q_t)$ for any $v \geq 1$ and $P(B_t = b) = s_t^{b-1}(1 - s_t)$ for any $b \geq 1$. With attention span V_t and purchase budget B_t , the consumer keeps viewing the products until she views V_t or purchases B_t products.

Because of the attention span V_t , the consumer will not purchase products whose positions are greater than V_t in the ranking list. Let $Q \subseteq [V_t]$ denote the possible positions of the purchased products by the consumer. The corresponding indices of purchased products are $\sigma_{t,Q}$. Therefore, the probability that the consumer purchases the product set $\sigma_{t,Q}$ is given by

$$P^{\text{buy}}(Q; V_t, B_t, \boldsymbol{\lambda}_t, \boldsymbol{\sigma}_t) = \begin{cases} \left(\prod_{k \in Q} \lambda_{t, \sigma_{t,k}} \right) \left(\prod_{k \in [V_t] \setminus Q} (1 - \lambda_{t, \sigma_{t,k}}) \right), & \text{if } |Q| < B_t. \\ \left(\prod_{k \in Q} \lambda_{t, \sigma_{t,k}} \right) \left(\prod_{k \in [V_t] \setminus Q, k \leq \max(Q)} (1 - \lambda_{t, \sigma_{t,k}}) \right), & \text{if } |Q| = B_t. \\ 0, & \text{if } |Q| > B_t. \end{cases} \quad (1)$$

Specifically, the probability $P^{\text{buy}}(Q; V_t, B_t, \boldsymbol{\lambda}_t, \boldsymbol{\sigma}_t)$ is calculated based on the cardinality of Q . When $|Q| < B_t$, the consumer does not meet her purchase budget and stops browsing the list after viewing V_t products. When $|Q| = B_t$, the consumer stops browsing as long as she purchases the last product in the list Q . When $|Q| > B_t$, the probability becomes 0 since the number of purchases $|Q|$ exceeds her purchase budget.

Remark. The setting of the purchase budget provides a rational way to characterize the consumers' behaviors with multiple purchases and extends previous works that suppose consumers would purchase at most one product [13, 17]. Furthermore, when $s_t \rightarrow 0$, the consumers would tend to purchase at most one product and our setting degenerates to the setting proposed in [13]. Thus the previous setting can be considered a special case of our setting. In addition, Liang et al. [40] assume that consumers always continue viewing the product list after purchasing any product, which is not practical in real scenarios. By contrast, the purchase budget proposed in our paper leads to a different and more practical consumer choice model. In addition, our model introduces an extra parameter (*i.e.*, the parameter), which makes our problem more challenging.

Revenue optimization The online retailer chooses a ranking policy $\sigma_t \dashv [N]$ on the products for the t -th consumer. For the consumer with fixed purchase probabilities for each product $\boldsymbol{\lambda}_t$, attention

span V_t , and purchase budget B_t , the expected revenue could be written as

$$R(V_t, B_t; \lambda_t, \sigma_t) = \sum_{Q \subseteq [V_t]} P^{\text{buy}}(Q; V_t, B_t, \lambda_t, \sigma_t) \left(\sum_{i \in Q} r_{\sigma_t, i} \right). \quad (2)$$

In reality, the attention span V_t and purchase budget B_t are random and can not be estimated accurately by the online retailer. As a result, the retailer calculates the expected revenue for the t -th consumer by sampling V_t from $\text{Geo}(1 - q_t)$ and sampling B_t from $\text{Geo}(1 - s_t)$, *i.e.*,

$$\text{Revenue}(\sigma_t; \lambda_t, q_t, s_t) = \mathbb{E}_{V_t \sim \text{Geo}(1 - q_t), B_t \sim \text{Geo}(1 - s_t)} [R(V_t, B_t; \lambda_t, \sigma_t)]. \quad (3)$$

To maximize the revenue, the optimal ranking policy is given by

$$\sigma_t^*(\lambda_t, q_t, s_t) = \arg \max_{\sigma_t \rightarrow [N]} \text{Revenue}(\sigma_t; \lambda_t, q_t, s_t). \quad (4)$$

For simplicity, we use σ_t^* to denote $\sigma_t^*(\lambda_t, q_t, s_t)$.

Basic assumption We also need the following assumption to guarantee the tractability of the proposed problem. In reality, consumers always have finite attention span, which indicates that q_t is always smaller than 1.

Assumption 3.1. There exists a known constant $\epsilon_Q > 0$ such that for all $t > 0$, $q_t \leq 1 - \epsilon_Q$.

4 Optimal Ranking Policy Given Consumers' Characteristics

In this section, we introduce the optimal ranking policy when the online retailer exactly knows the consumers' parameters, including the purchase probability for each product λ , random attention span (parametrized with q), and random purchase budget (parametrized with s).

We drop the subscript t in this section for simplicity. The optimal ranking strategy stems from the following theorem.

Theorem 4.1. Under Assumption 3.1, the following holds for the optimal ranking strategy $\sigma^*(\lambda, q, s)$.

$$\forall 1 \leq i < N, \quad \frac{\lambda_{\sigma_i^*} r_{\sigma_i^*}}{1 - q + q(1 - s)\lambda_{\sigma_i^*}} \geq \frac{\lambda_{\sigma_{i+1}^*} r_{\sigma_{i+1}^*}}{1 - q + q(1 - s)\lambda_{\sigma_{i+1}^*}}. \quad (5)$$

Remark. The expected revenue of a single product σ_i^* is $\lambda_{\sigma_i^*} r_{\sigma_i^*}$, which grows as the purchase probability and the revenue generated by the product increase. Yet purchasing a product reduces the purchase probability of the following products, and thus the expected revenue of other products decreases as $\lambda_{\sigma_i^*}$ increases. Furthermore, when $s \rightarrow 0$, the term in Equation (5) becomes the same as that in [13, Theorem 1] if ignoring the marketing fatigue (*i.e.*, overexposure to unwanted marketing messages) term in their result. As a result, Theorem 4.1 consistently extends the results on the single-purchase setting [13] to the multiple-purchase setting.

In addition, Liang et al. [40] assume that the online retailer will stop recommending products with a certain probability $1 - w_t$ to avoid the cost brought by marketing fatigue. As a result, the parameter w_t is a part of the ranking policy instead of consumers' characteristics and should be optimized to achieve maximal revenue. In our paper, we focus on the multiple purchases with budget setting and leave the extension to the case with the marketing fatigue as future work.

According to Theorem 4.1, the online retailer can provide the optimal ranking policy σ^* by sorting the products in descending order with value $(\lambda_k r_k) / (1 - q + q(1 - s)\lambda_k)$ for the k -th product. The time complexity is $O(N \log N)$.

5 Online Learning of the Ranking Policy

In this section, we consider a more realistic scenario where the online retailer has no prior knowledge about consumers' characteristics. At each timestamp, a consumer arrives and the online retailer can only provide proper product ranking policies based on the feedback of previous consumers. As a result, the retailer needs to design online learning algorithms to model consumers' behaviors and maximize revenue in the meantime. We consider two settings (*i.e.*, the non-contextual setting where all consumers share the same parameters (Section 5.1) and the contextual setting where consumers have personalized behaviors (Section 5.2)) and develop the Multiple-Purchase-with-Budget UCB (MPB-UCB) algorithms on both settings. Before introducing the details of the algorithms, we first formally define the notations to characterize the consumers' behaviors.

Notations for consumers' behaviors Let $\Phi_t \in \mathbb{Z}_+$ be the random variable that denotes the number of products viewed by the t -th consumer. Let random variable $\gamma_{t,k} \in \{0, 1\}$ denote whether the consumer purchases the product displayed in the k -th position (*i.e.*, product $\sigma_{t,k}$) if she views the product. Specifically, $\gamma_{t,k} = 1$ if and only if the product is purchased. Let $\boldsymbol{\gamma}_t = \{\gamma_{t,k}\}_{k=1}^{\Phi_t}$.

Let random variable $\eta_{t,k} \in \{0, 1\}$ denote whether the consumer continues viewing products if she views the product displayed in the k -th position but does not buy it. $\eta_{t,k}$ is observed only when $k \leq \Phi_t$ and $\gamma_{t,k} = 0$. Let random variable $\mu_{t,k} \in \{0, 1\}$ denote whether the consumer keeps viewing products if she buys the product in the k -th position. $\mu_{t,k}$ is observed only when $k \leq \Phi_t$ and $\gamma_{t,k} = 1$. Let $\boldsymbol{\eta}_t = \{\eta_{t,k}\}_{k=1}^{\Phi_t}$ and $\boldsymbol{\mu}_t = \{\mu_{t,k}\}_{k=1}^{\Phi_t}$.

Let $\text{Ber}(\cdot)$ denote the Bernoulli distribution and we have the following result.

Lemma 5.1. $\{\gamma_{t,k}\}_{t>0, 1 \leq k < N}$ are independent of each other, $\{\eta_{t,k}\}_{t>0, 1 \leq k < N}$ are independent of each other, and $\{\mu_{t,k}\}_{t>0, 1 \leq k < N}$ are independent of each other. In addition, $\forall t > 0$ and $1 \leq k < N$, $\gamma_{t,k} \sim \text{Ber}(\lambda_{\sigma_{t,k}})$, $\eta_{t,k} \sim \text{Ber}(q_t)$, and $\mu_{t,k} \sim \text{Ber}(q_t s_t)$.

Lemma 5.1 is the basis of the algorithms that estimate parameters $\boldsymbol{\lambda}$, q , and s from historical data. In addition, compared with s_t , it is easier to estimate $q_t s_t$ from data and we denote it as $w_t \triangleq q_t s_t$.

5.1 Non-contextual Online Setting

In this subsection, we assume that all consumers share the same parameters, including the purchase probabilities on different products and the parameters on the distribution of the attention span and purchase budget. In detail, we assume that there exist $\boldsymbol{\lambda}$, q , s such that $\boldsymbol{\lambda}_t = \boldsymbol{\lambda}$, $q_t = q$, and $s_t = s$.

A) Estimation of Parameters

To estimate the parameters at timestamp t , we leverage the observed behavior of consumers before timestamp t . We first calculate the following statistics that can help provide unbiased estimators of the consumer parameters.

In detail, let $C_{t,k}$ be the number of times that the product k is observed and $c_{t,k}$ be the number of times the product k is purchased, *i.e.*, $C_{t,k} \triangleq \sum_{u=1}^t \sum_{i=1}^{\Phi_t} \mathbb{I}[\sigma_{t,i} = k]$ and $c_{t,k} \triangleq \sum_{u=1}^t \sum_{i=1}^{\Phi_t} \mathbb{I}[\sigma_{t,i} = k] \mathbb{I}[\gamma_{t,i} = 1]$. In addition, Let D_t^Q be the number of times that consumers view a product and do not purchase it and d_t^Q be the number of times that consumers continue viewing the list after viewing a product without purchasing it, *i.e.*, $D_t^Q \triangleq \sum_{u=1}^t \sum_{i=1}^{\Phi_t} \mathbb{I}[\gamma_{t,i} = 0]$ and $d_t^Q \triangleq \sum_{u=1}^t \sum_{i=1}^{\Phi_t} \mathbb{I}[\gamma_{t,i} = 0] \eta_{t,i}$. Furthermore, let D_t^W be the number of times that consumers purchase a product and d_t^W be the number of times that consumers continue viewing the list after purchasing a product, *i.e.*, $D_t^W \triangleq \sum_{u=1}^t \sum_{i=1}^{\Phi_t} \mathbb{I}[\gamma_{t,i} = 1]$ and $d_t^W \triangleq \sum_{u=1}^t \sum_{i=1}^{\Phi_t} \mathbb{I}[\gamma_{t,i} = 1] \mu_{t,i}$.

With these statistics, the parameters $\boldsymbol{\lambda}$, q , and $w = qs$ are estimated as follows.

$$\hat{\lambda}_{t,k} \triangleq c_{t,k}/C_{t,k}, \quad \hat{q}_t \triangleq d_t^Q/D_t^Q, \quad \hat{w}_t \triangleq d_t^W/D_t^W. \quad (6)$$

We provide the following proposition that gives the error bound of our parameter estimation approach.

Proposition 5.2. For all $t \in [T]$, with probability at least $1 - 6Nt^{-3}$, the following holds:

$$\forall k \in [N], \left| \hat{\lambda}_{t,k} - \lambda_k \right| \leq \sqrt{\frac{2 \log t}{C_{t,k}}}, \quad \left| \hat{q}_t - q \right| \leq \sqrt{\frac{2 \log t}{D_t^Q}}, \quad \text{and} \quad \left| \hat{w}_t - qs \right| \leq \sqrt{\frac{2 \log t}{D_t^W}}. \quad (7)$$

B) Algorithm

Following the classic *optimism in the face of uncertainty principle* [1], we design a UCB-like algorithm that learns consumer behaviors and maximizes revenue simultaneously. In detail, according to Proposition 5.2, we learn the optimistic estimators of the parameters given as follows

$$\tilde{\lambda}_{t,k} \triangleq \min \left\{ 1, \hat{\lambda}_{t,k} + \sqrt{\frac{2 \log t}{C_{t,k}}} \right\}, \quad \tilde{q}_t \triangleq \min \left\{ 1 - \epsilon_Q, \hat{q}_t + \sqrt{\frac{2 \log t}{D_t^Q}} \right\}, \quad \tilde{w}_t \triangleq \min \left\{ \tilde{q}_t, \hat{w}_t + \sqrt{\frac{2 \log t}{D_t^W}} \right\}. \quad (8)$$

Algorithm 1: MPB-UCB (Non-contextual)

- 1 **Input:** Products revenue r and hyper-parameter ϵ_Q .
 - 2 **Initialization:** $\tilde{\lambda}_{0,k} = 1$ for $k \in [N]$, $\tilde{q}_0 = 1 - \epsilon_Q$, $\tilde{w}_0 = 1 - \epsilon_Q$.
 - 3 **for** $t=1:T$ **do**
 - 4 Let σ_t be the optimal offline policy from [Theorem 4.1](#) with $\lambda = \tilde{\lambda}_{t-1}$, $q = \tilde{q}_{t-1}$, and $s = \tilde{w}_{t-1}/\tilde{q}_{t-1}$.
 - 5 Offer ranking policy σ_t and observe $\Phi_t, \gamma_t, \eta_t, \mu_t$.
 - 6 Update statistics $C_{t,k}, c_{t,k}, D_t^Q, d_t^Q, D_t^W$, and d_t^W .
 - 7 Calculate $\tilde{\lambda}_{t,k}, \tilde{q}_t$, and \tilde{w}_t by [Equations \(6\) and \(8\)](#).
 - 8 **end**
-

As shown in [Algorithm 1](#), when the t -th consumer arrives, the online retailer displays the products based on [Theorem 4.1](#) with $\tilde{\lambda}_{t-1}$, \tilde{q}_{t-1} , and $\tilde{w}_{t-1}/\tilde{q}_{t-1}$ shown in Line 4. Afterward, the online retailer updates the statistics according to consumers' feedback as shown in Lines 5 and 6. The estimators are then updated in Line 7.

C) Regret Analysis

To theoretically analyze the performance of [Algorithm 1](#), we first introduce the regret, which measures the total difference between the maximal revenue achieved by σ^* and the cumulative reward by an online algorithm after T rounds.

$$\text{Reg}_T \triangleq \sum_{t=1}^T \text{Revenue}(\sigma^*; \lambda, q, s) - \text{Revenue}(\sigma_t; \lambda, q, s). \quad (9)$$

Reg_T is a random variable and the randomness comes from the uncertainty in consumers' behaviors. As a result, we focus on $\mathbb{E}[\text{Reg}_T]$, which is a common routine in bandit literature [\[46\]](#). The performance of the algorithm is given by the following theorem.

Theorem 5.3. *Under [Assumption 3.1](#) (with parameter ϵ_Q), for any $T > 1$, the regret achieved by [Algorithm 1](#) is bounded by*

$$\mathbb{E}[\text{Reg}_T] \leq C \cdot \frac{r_{\max}}{\epsilon_Q} N \sqrt{T \log T}, \quad (10)$$

where C is a absolute constant (independent of problem parameters) and $r_{\max} = \max_{k \in [N]} r_k$.

Remark. We analyze different parameters' impacts on the regret bound in [Equation \(10\)](#). Firstly, the regret grows at $\tilde{O}(\sqrt{T})$, which is a standard result in online learning algorithms [\[46\]](#). Secondly, the regret bound is linearly related to N , which is similar to the results in cascade bandit models [\[54\]](#) and revenue management literature [\[17\]](#). Finally, the bound depends on r_{\max}/ϵ_Q , which measures the maximal expected revenue the online retailer could achieve from each consumer.

5.2 Contextual Online Setting

We consider a more realistic scenario where the consumers' behaviors are various and depend on their own features. In detail, let $x_t \in \mathbb{R}^{m_x}$ be the feature of the t -th consumer with dimension m_x and $y_{t,k} \in \mathbb{R}^{m_y}$ be the joint feature of the t -th consumer and k -th product with dimension m_y . $y_{t,k}$ can be the concatenation of consumer features and product features or fusion through a transformation such as a matrix multiplication. In this paper, we adopt concatenation for simplicity. Let $\|\cdot\|$ denote the ℓ_2 norm. Following [\[17\]](#), we consider the classic linear bandit setting [\[7, 21, 39\]](#) and the ground-truth data-generating process is given by the following assumptions.

Assumption 5.1. There exist constant vectors $\beta^\Lambda \in \mathbb{R}^{m_y}$, $\beta^Q \in \mathbb{R}^{m_x}$, and $\beta^S \in \mathbb{R}^{m_x}$, such that

$$\lambda_{t,k} = y_{t,k}^\top \beta^\Lambda, \forall k \in [N], \quad q_t = x_t^\top \beta^Q, \quad \text{and} \quad s_t = x_t^\top \beta^S. \quad (11)$$

Assumption 5.2. For any $t > 0$ and $k \in [N]$, $\|x_t\| \leq 1$ and $\|y_{t,k}\| \leq 1$. In addition, there exists a constant $U > 0$ such that $\|\beta^\Lambda\|, \|\beta^Q\|, \|\beta^S\| \leq U$.

Remark. These two assumptions are common in revenue maximization and linear bandits [\[1, 46\]](#). [Assumption 5.1](#) can be satisfied with kernel functions for complex data-generating processes and [Assumption 5.2](#) can be achieved by normalization. Thus the assumptions are rational and easily attainable in real applications.

A) Estimation of Parameters

For the estimation of β^Λ and β^Q , we consider the ridge regression method. Given the observation of consumer's behaviors till the t -th timestamp, the estimation $\hat{\beta}_t^\Lambda, \hat{\beta}_t^Q$ is given by

$$\hat{\beta}_t^\Lambda = \operatorname{argmin}_{\beta \in \mathbb{R}^{m_y}} \sum_{u=1}^t \sum_{k=1}^{\Phi_u} \left(y_{u,\sigma_{u,k}}^\top \beta - \gamma_{u,k} \right)^2 + \alpha_\Lambda \|\beta\|_2^2 = \left(\Sigma_t^\Lambda \right)^{-1} \rho_t^\Lambda, \quad (12)$$

$$\hat{\beta}_t^Q = \operatorname{argmin}_{\beta \in \mathbb{R}^{m_x}} \sum_{u=1}^t \sum_{k=1}^{\Phi_u} \mathbb{I}[\gamma_{u,k} = 0] \left(x_u^\top \beta - \eta_{u,k} \right)^2 + \alpha_Q \|\beta\|_2^2 = \left(\Sigma_t^Q \right)^{-1} \rho_t^Q. \quad (13)$$

Here $\Sigma_t^\Lambda = \sum_{u=1}^t \sum_{k=1}^{\Phi_u} y_{u,\sigma_{u,k}} y_{u,\sigma_{u,k}}^\top + \alpha_\Lambda \mathbf{I}_{m_y}$, $\rho_t^\Lambda = \sum_{u=1}^t \sum_{k=1}^{\Phi_u} \gamma_{u,k} y_{u,\sigma_{u,k}}$, $\Sigma_t^Q = \sum_{u=1}^t \sum_{k=1}^{\Phi_u} \mathbb{I}[\gamma_{u,k} = 0] x_u x_u^\top + \alpha_Q \mathbf{I}_{m_x}$, and $\rho_t^Q = \sum_{u=1}^t \sum_{k=1}^{\Phi_u} \mathbb{I}[\gamma_{u,k} = 0] \eta_{u,k} x_u$. \mathbf{I}_{m_x} and \mathbf{I}_{m_y} are the identity matrices with dimension m_x and m_y , respectively.

For the estimation of β^S , similar to Section 5.1, we focus on regressing $w_t = q_t s_t$ instead. Note that $w_t = q_t s_t = x_t^\top \beta^Q (\beta^S)^\top x_t = \operatorname{vec}(x_t x_t^\top)^\top \operatorname{vec}(\beta^Q (\beta^S)^\top)$, where $\operatorname{vec}(\cdot)$ denotes the vectorization operation that maps a matrix to a vector. Therefore, define $z_t \triangleq \operatorname{vec}(x_t x_t^\top)$ and $\beta^W \triangleq \operatorname{vec}(\beta^Q (\beta^S)^\top)$ and we can get $w_t = z_t^\top \beta^W$. Similarly, we estimate β^W by conducting ridge regression as follows

$$\hat{\beta}_t^W = \operatorname{argmin}_{\beta \in \mathbb{R}^{m_x^2}} \sum_{u=1}^t \sum_{k=1}^{\Phi_u} \mathbb{I}[\gamma_{u,k} = 1] \left(z_u^\top \beta - \mu_{u,k} \right)^2 + \alpha_W \|\beta\|_2^2 = \left(\Sigma_t^W \right)^{-1} \rho_t^W. \quad (14)$$

Here $\Sigma_t^W = \sum_{u=1}^t \sum_{k=1}^{\Phi_u} \mathbb{I}[\gamma_{u,k} = 1] z_u z_u^\top + \alpha_W \mathbf{I}_{m_x^2}$ and $\rho_t^W = \sum_{u=1}^t \sum_{k=1}^{\Phi_u} \mathbb{I}[\gamma_{u,k} = 1] \mu_{u,k} z_u$. $\mathbf{I}_{m_x^2}$ is the identity matrix with dimension m_x^2 .

Let $\|\beta\|_A = \sqrt{\beta^\top A \beta}$ for any positive definite matrix A . Proposition 5.4 gives the estimation error bound of the coefficients.

Proposition 5.4. *Under Assumptions 5.1 and 5.2 (with parameter U), for any $t \geq 1$, with probability at least $1 - 3t^{-2}$, the following holds:*

$$\left\| \hat{\beta}_t^\Lambda - \beta^\Lambda \right\|_{\Sigma_t^\Lambda} \leq \tau(t, m_y, \alpha_\Lambda), \left\| \hat{\beta}_t^Q - \beta^Q \right\|_{\Sigma_t^Q} \leq \tau(t, m_x, \alpha_Q), \text{ and } \left\| \hat{\beta}_t^W - \beta^W \right\|_{\Sigma_t^W} \leq \tau(t, m_x^2, \alpha_W). \quad (15)$$

Here $\tau(t, m, \alpha) = 1/2\sqrt{m \log(1 + tN/(m\alpha))} + 4 \log t + \alpha^{1/2}(U + 1)^2$.

B) Algorithm

Given the estimation of β^Λ , β^Q and β^W we estimate $\lambda_{t,k}$, q_t and w_t as follows.

$$\begin{aligned} \tilde{\lambda}_{t,k} &= \operatorname{Proj}_{[0,1]} \left(y_{t,k}^\top \hat{\beta}_{t-1}^\Lambda + \tau(t-1, m_y, \alpha_\Lambda) \|y_{t,k}\|_{(\Sigma_{t-1}^\Lambda)^{-1}} \right) \\ \tilde{q}_t &= \operatorname{Proj}_{[0,1-\epsilon_Q]} \left(x_t^\top \hat{\beta}_{t-1}^Q + \tau(t-1, m_x, \alpha_Q) \|x_t\|_{(\Sigma_{t-1}^Q)^{-1}} \right) \\ \tilde{w}_t &= \operatorname{Proj}_{[0,\tilde{q}_t]} \left(z_t^\top \hat{\beta}_{t-1}^W + \tau(t-1, m_x^2, \alpha_W) \|z_t\|_{(\Sigma_{t-1}^W)^{-1}} \right) \end{aligned} \quad (16)$$

Here $\operatorname{Proj}_{[a,b]}(\cdot)$ is a function that projects the input into the interval $[a, b]$.

As shown in Algorithm 2, when the t -th consumer arrives, the online retailer observes the consumer features x_t and $y_{t,k}$ for $k \in [N]$ and calculates the intermediate feature z_t . Afterward, he calculates $\tilde{\lambda}_{t,k}$, \tilde{q}_t , and \tilde{w}_t according to Equation (16) in Line 5. Then the ranking policy σ_t is given as shown in Line 6 and the statistics are updated in Lines 7 and 8. Finally, the parameters $\hat{\beta}_t^\Lambda$, $\hat{\beta}_t^Q$, and $\hat{\beta}_t^W$ are estimated in Line 9 for the next iteration.

C) Regret Analysis

We analyze the performance of Algorithm 2 by investigating its regret bound. For a sequence of T consumers with features $\mathbf{X}_T = \{x_t\}_{t=1}^T$ and $\mathbf{Y}_T = \{\{y_{t,k}\}_{k=1}^N\}_{t=1}^T$, the regret is given by

$$\operatorname{Reg}_T(\mathbf{X}_T, \mathbf{Y}_T) = \sum_{t=1}^T \operatorname{Revenue}(\sigma_t^*; \lambda_t, q_t, s_t) - \operatorname{Revenue}(\sigma_t; \lambda_t, q_t, s_t). \quad (17)$$

Algorithm 2: MPB-UCB (Contextual)

- 1 **Input:** Products revenue r and hyper-parameter $\epsilon_Q, \alpha_\Lambda, \alpha_Q,$ and α_W .
 - 2 **Initialization:** $\hat{\beta}_0^\Lambda = \mathbf{0}, \hat{\beta}_0^Q = \mathbf{0},$ and $\hat{\beta}_0^W = \mathbf{0}$.
 - 3 **for** $t=1:T$ **do**
 - 4 Observe consumer features x_t and $y_{t,k}$ for $k \in [N]$. Let $z_t = \text{vec}(x_t x_t^\top)$.
 - 5 Calculate $\tilde{\lambda}_{t,k}, \tilde{q}_t,$ and \tilde{w}_t according to Equation (16).
 - 6 Let σ_t be the optimal ranking policy from Theorem 4.1 with $\lambda = \tilde{\lambda}_t, q = \tilde{q}_t,$ and $s = \tilde{w}_t/\tilde{q}_t$.
 - 7 Offer ranking policy σ_t and observe $\Phi_t, \gamma_t, \eta_t, \mu_t$.
 - 8 Calculate statistics $\Sigma_t^\Lambda, \rho_t^\Lambda, \Sigma_t^Q, \rho_t^Q, \Sigma_t^W,$ and ρ_t^W .
 - 9 Calculate estimated parameters $\hat{\beta}_t^\Lambda, \hat{\beta}_t^Q,$ and $\hat{\beta}_t^W$ according to Equations (12), (13) and (14).
 - 10 **end**
-

Similar to the non-contextual setting, $\text{Reg}_T(\mathbf{X}_T, \mathbf{Y}_T)$ is a random variable and we focus on $\mathbb{E}[\text{Reg}_T(\mathbf{X}_T, \mathbf{Y}_T)]$. The performance of the algorithm is given by the following theorem.

Theorem 5.5. *Under Assumptions 3.1 (with parameter ϵ_Q), 5.1, and 5.2 (with parameter U), for any $T > 1$, the regret achieved by Algorithm 2 is bounded by*

$$\mathbb{E}[\text{Reg}_T(\mathbf{X}_T, \mathbf{Y}_T)] \leq C \cdot \frac{r_{\max}}{\epsilon_Q} (\chi(T, m_y, \alpha_\Lambda) + \chi(T, m_x, \alpha_Q) + \chi(T, m_x^2, \alpha_W)). \quad (18)$$

Here C is a absolute constant (independent of problem parameters), $r_{\max} = \max_{k \in [N]} r_k$, and

$$\chi(T, m, \alpha) = \left(\frac{1}{2} \sqrt{m \log \left(1 + \frac{TN}{m\alpha} \right) + 4 \log T + \alpha^{1/2} (U + 1)^2} \right) \sqrt{\frac{TN^2 m \log \left(1 + \frac{TN}{m\alpha} \right)}{\alpha \log(1 + \alpha^{-1})}}. \quad (19)$$

Remark. We analyze different parameters' impacts on the regret bound in Equation (19). Firstly, the regret grows at $\tilde{O}(\sqrt{T})$, a standard result in online learning algorithms [46]. Secondly, the regret depends on $\tilde{O}(N)$, which shares a similar result with the revenue management literature [17]. Thirdly, similar to Theorem 5.3, the bound depends linearly on r_{\max}/ϵ_Q that measures the maximal expected revenue the online retailer could achieve from each consumer. Fourthly, the bound depends on the dimensions of features (i.e., m_x and m_y) by $\tilde{O}(\sqrt{m_y})$ and $\tilde{O}(m_x)$ because the feature dimension for the estimation of $w_t = q_t s_t$ is m_x^2 and then the result is consistent with [17, 54]. Finally, the bound depends on $O(U^2 + U)$ since $\|\beta^W\| \leq U^2$ and $\|\beta^Q\|, \|\beta^\Lambda\| \leq U$, which agrees with [17].

6 Experiments

In this section, we conduct experiments on synthetic data to verify the performances of the online ranking policies. Reports and discussions of more results including the expected revenue and the ratio over the expected revenue of the optimal ranking policy are in Section B.

Baselines We adopt the following baselines. Firstly, Cao and Sun [13] considered the single-purchase setting, which is the simplified version of our multiple-purchase setting. We denote it as **Single Purchase**. Secondly, we adopt the method in [40] (denoted as **Keep Viewing**), which assumes that consumers always keep viewing the ranking list after purchasing products. Furthermore, we implement two explore-then-exploit-based algorithms (denoted as **Explore Then Exploit A** and **Explore Then Exploit B**, respectively) to verify the advantages of our method that could balance exploration and exploitation.

6.1 Synthetic Data

Data-generating processes For the non-contextual setting, we consider $N = 50$ and 300 products and $T = 100,000$ consumers. The revenue for each product r_k is uniformly distributed in $[0, 1]$ and the purchase probabilities for each product are uniformly sampled from $[0, 0.3]$. We set $q = 0.9$ and $s = 0.5$ and 0.8 to test the performance of different models on characterizing consumers' behaviors with multiple purchases.

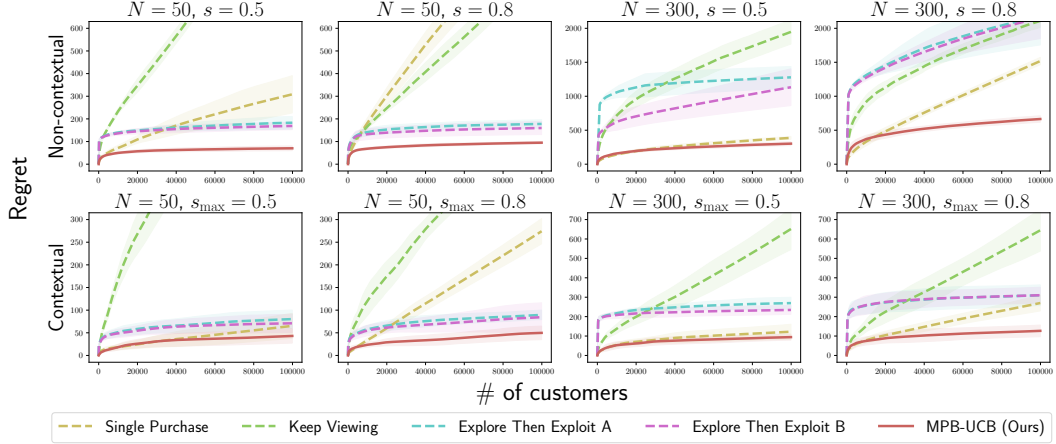


Figure 1: Experimental results on synthetic data. We plot the regret curves for different settings by varying $N = 50/300$ and $s/s_{\max} = 0.5/0.8$ in both non-contextual and contextual scenarios.

For the contextual setting, we consider $N = 50$ and 300 products and $T = 100,000$ consumers. The revenue r_k is also uniformly distributed in $[0, 1]$. Both the dimension of consumer features and product features are set to 5 (*i.e.*, $m_x = 5$). The joint feature $y_{t,k}$ (the feature of the t -th consumer and k -th product) is a concatenate of the corresponding consumer and product feature so that $m_y = 10$. The consumer features and product features are uniformly sampled from $[0.8/\sqrt{m_x}, 1/\sqrt{m_x}]$ and $[0, 1/\sqrt{m_x}]$. Afterward, the coefficients β^A , β^Q , and β^S in Equation (11) are uniformly sampled from $[0, 1]$. To ensure $\lambda_{t,k}$, q_t , s_t have similar ranges in the contextual setting, we normalize the coefficients such that the maximal elements in $\lambda_{t,k}$ and q_t (denoted as λ_{\max} and q_{\max}) are 0.3 and 0.9, respectively. Similarly, by normalizing β^S , we set the maximal element in s_t (denoted as s_{\max}) to 0.5 and 0.8.

Results and analysis We implement the experiments for 5 distinct simulations by resampling consumers' behaviors while their basic characteristics remain unchanged. Specifically, in both the non-contextual and contextual settings, the purchase probabilities (*i.e.*, λ_t) and parameters on the span attention and purchase budget (*i.e.*, q_t and s_t) are the same. However, the same consumer in distinct simulations may purchase different products due to randomness.

We then evaluate our method (MPB-UCB) and baselines via calculating the regret according to Equations (9) and (17). The results are shown in Figure 1 and our method outperforms all baselines in both settings with various parameters. On the one hand, the Single Purchase and Keep Viewing baselines consider different consumer choice models and are not directly applicable here to the more realistic multiple-purchase setting. On the other hand, compared with the explore-then-exploit-based algorithms, our method integrates the advantages of traditional UCB-like algorithms, leading to a better exploration-exploitation trade-off for smaller regret. In addition, as introduced in Section B, Explore Then Exploit B incorporates learning processes during the exploration phase, leading to better performance compared with Explore Then Exploit A.

6.2 Semi-synthetic Data

We utilize the Ad Display/Click Data from Taobao¹, which contains ad display/purchase logs (26 million records) of 1,140,000 anonymized users from the website of Taobao for 8 days. Because the behavior logs only contain the purchase information of the users on a category of a brand, we view a (brand, category) pair in the dataset as a product and the average prices of the ads in this pair as the revenue for the product.

Because we can not obtain consumers' behaviors when we offer a new product ranking policy, following [14], we first estimate the parameters (*i.e.*, λ , q , s in the non-contextual setting and β^A , β^Q , β^S in the contextual setting) with all of the data. Afterward, we use the estimated parameters to simulate consumers' behaviors when we provide them with different ranking lists of products.

¹<https://tianchi.aliyun.com/datalab/dataSet.html?dataId=56>. License: CC BY-NC 4.0.

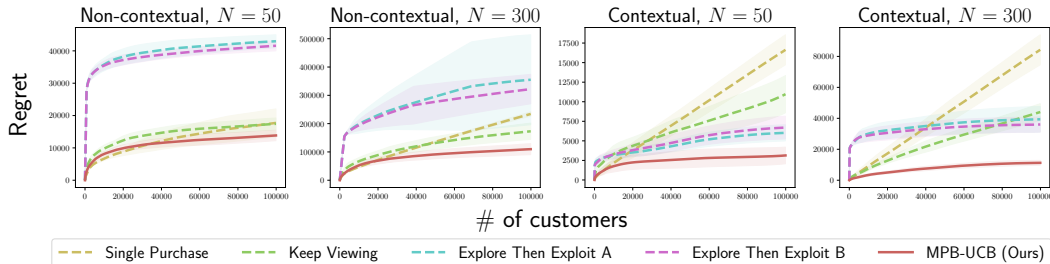


Figure 2: Experimental results on the semi-synthetic data. We plot the regret curves for different settings by varying $N = 50/300$ in both non-contextual and contextual scenarios.

Data processing Since each consumer may launch and leave the platform many times during the 8 days, we need to distinguish different launching activities of the same consumer. Specifically, we suppose that a consumer stops viewing the product list if it has been more than 10 minutes since her last interaction with the platform. Afterward, we can estimate the parameters in both non-contextual and contextual settings.

In the non-contextual setting, we calculate the statistics $C_{t,k}$, $c_{t,k}$, D_t^Q , d_t^Q , D_t^W , and d_t^W from the data and estimate λ , q , s directly according to Lemma 5.1. The estimation is $q = 0.870$ and $s = 0.907$. Afterward, we sample $N = 50$ and 300 products with prices no more than 200 and purchase probabilities at least 0.1. Similar to the synthetic data, we set $T = 100,000$.

In the contextual setting, we set the gender, age level, consumption grade, and shopping level as consumers' features and the number of ads, the logarithm of the average, standard deviation, maximum, and median of the ads' prices in the (brand, category) pair as products' features. With the statistics $C_{t,k}$, $c_{t,k}$, D_t^Q , d_t^Q , D_t^W , and d_t^W , we use the ridge regression with regularization strength 1 to estimate β^Λ , β^Q , and β^S . Afterward, we sample $N = 50$ and 300 products with prices no more than 200 and purchase probabilities at least 0.1. In addition, we sample $T = 100,000$ consumers.

Results and analysis Similar to the synthetic experiments, we implement the experiments for 5 distinct simulations by resampling consumers' behaviors while their basic characteristics remain unchanged. The results of the regret are shown in Figure 2. As shown in the figure, our method outperforms all baselines in different scenarios. Firstly, the Single Purchase and Keep Viewing baselines are not directly applicable here since they consider different consumer choice models. It is worth mentioning that the estimated s is large from the dataset. As a result, the consumers prefer to keep viewing products if they purchase any product, making the Keep Viewing baseline outperforms Single Purchase in this setting. Secondly, our method combines the benefits of existing UCB-like algorithms, resulting in a superior exploration-exploitation trade-off with less regret compared with explore-then-exploit-based methods. In addition, similar to the analysis in the synthetic experiments, Explore Then Exploit B outperforms Explore Then Exploit A in most cases.

7 Conclusion

To conclude, we propose a novel consumer choice model to deal with multiple-purchase activities of consumers in online scenarios. We characterize the optimal ranking policy and further design the MPB-UCB algorithms with $\tilde{O}(\sqrt{T})$ regret in both non-contextual and contextual online settings. We conduct extensive experiments to prove the effectiveness of our method.

Acknowledgments

Peng Cui's research was supported in part by National Key R&D Program of China (No. 2018AAA0102004, No. 2020AAA0106300), National Natural Science Foundation of China (No. U1936219, 62141607), Beijing Academy of Artificial Intelligence (BAAI). Bo Li's research was supported by the National Natural Science Foundation of China (No.72171131), the Tsinghua University Initiative Scientific Research Grant (No. 2019THZWJC11), Technology and Innovation Major Project of the Ministry of Science and Technology of China under Grants 2020AAA0108400 and 2020AAA0108403.

References

- [1] Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 24, 2011. 5, 6, 26, 28
- [2] Andrés Abeliuk, Gerardo Berbeglia, Manuel Cebrian, and Pascal Van Hentenryck. Assortment optimization under a multinomial logit model with position bias and social influence. *4OR*, 14(1):57–75, 2016. 2
- [3] Ashish Agarwal, Kartik Hosanagar, and Michael D Smith. Location, location, location: An analysis of profitability of position in online advertising markets. *Journal of marketing research*, 48(6):1057–1073, 2011. 1
- [4] Ali Aouad and Danny Segev. Display optimization for vertically differentiated locations under multinomial logit preferences. *Management Science*, 67(6):3519–3550, 2021. 2, 3
- [5] Victor F Araman and René Caldentey. Dynamic pricing for nonperishable products with demand learning. *Operations research*, 57(5):1169–1188, 2009. 2
- [6] Arash Asadpour, Rad Niazadeh, Amin Saberi, and Ali Shameli. Ranking an assortment of products via sequential submodular optimization. *arXiv preprint arXiv:2002.09458*, 2020. 3
- [7] Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov):397–422, 2002. 6
- [8] Michael R Baye, J Rupert J Gatti, Paul Kattuman, and John Morgan. Clicks, discontinuities, and firm demand online. *Journal of Economics & Management Strategy*, 18(4):935–975, 2009. 1
- [9] Omar Besbes and Assaf Zeevi. Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations Research*, 57(6):1407–1420, 2009. 2
- [10] Omar Besbes and Assaf Zeevi. Blind network revenue management. *Operations research*, 60(6):1537–1550, 2012. 2
- [11] Josef Broder and Paat Rusmevichientong. Dynamic pricing under a general parametric choice model. *Operations Research*, 60(4):965–980, 2012. 2
- [12] Juan José Miranda Bront, Isabel Méndez-Díaz, and Gustavo Vulcano. A column generation algorithm for choice-based network revenue management. *Operations research*, 57(3):769–784, 2009. 2
- [13] Junyu Cao and Wei Sun. Dynamic learning of sequential choice bandit problem under marketing fatigue. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3264–3271, 2019. 2, 3, 4, 8, 15
- [14] Junyu Cao, Wei Sun, and Zuo-Jun Max Shen. Sequential choice bandits: Learning with marketing fatigue. *Available at SSRN 3355211*, 2019. 2, 3, 9, 15
- [15] Junyu Cao, Wei Sun, Zuo-Jun Max Shen, and Markus Ettl. Fatigue-aware bandits for dependent click models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 3341–3348, 2020. 1, 2
- [16] Olivier Chapelle and Ya Zhang. A dynamic bayesian network click model for web search ranking. In *Proceedings of the 18th international conference on World wide web*, pages 1–10, 2009. 1, 2
- [17] Ningyuan Chen, Anran Li, and Shuoguang Yang. Revenue maximization and learning in products ranking. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 316–317, 2021. 1, 2, 3, 6, 8, 15
- [18] Yuxin Chen and Song Yao. Sequential search with refinement: Model and application with click-stream data. *Management Science*, 63(12):4345–4365, 2017. 1, 3

- [19] Wang Chi Cheung, Vincent Tan, and Zixin Zhong. A thompson sampling algorithm for cascading bandits. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 438–447. PMLR, 2019. 2
- [20] Roberta Comunian, Caroline Chapain, and Nick Clifton. Location, location, location: exploring the complex relationship between creative industries and place. *Creative Industries Journal*, 3(1):5–10, 2010. 1
- [21] Varsha Dani, Thomas P Hayes, and Sham M Kakade. Stochastic linear optimization under bandit feedback. In *21st Annual Conference on Learning Theory*, pages 355–366, 2008. 6
- [22] James M Davis, Guillermo Gallego, and Huseyin Topaloglu. Assortment optimization under variants of the nested logit model. *Operations Research*, 62(2):250–273, 2014. 2
- [23] Babur De los Santos and Sergei Koulayev. Optimizing click-through in online rankings with endogenous search refinement. *Marketing Science*, 36(4):542–564, 2017. 3
- [24] Arnoud V Den Boer. Dynamic pricing and learning: historical origins, current research, and new directions. *Surveys in operations research and management science*, 20(1):1–18, 2015. 2
- [25] Mahsa Derakhshan, Negin Golrezaei, Vahideh Manshadi, and Vahab Mirrokni. Product ranking on online platforms. *Management Science*, 2022. 1
- [26] Daria Dzyabura and John R Hauser. Recommending products when consumers learn their preference weights. *Marketing Science*, 38(3):417–441, 2019. 1
- [27] Jacob Feldman and Danny Segev. Improved approximation schemes for mnl-driven sequential assortment optimization. Available at SSRN 3440645, 2019. 2
- [28] Kris J Ferreira, Sunanda Parthasarathy, and Shreyas Sekar. Learning to rank an assortment of products. *Management Science*, 2021. 1, 2
- [29] Alvaro Flores, Gerardo Berbeglia, and Pascal Van Hentenryck. Assortment optimization under the sequential multinomial logit model. *European Journal of Operational Research*, 273(3):1052–1064, 2019. 2
- [30] Guillermo Gallego, Anran Li, Van-Anh Truong, and Xinshang Wang. Approximation algorithms for product framing and pricing. *Operations Research*, 68(1):134–160, 2020. 2
- [31] Pin Gao, Yuhang Ma, Ningyuan Chen, Guillermo Gallego, Anran Li, Paat Rusmevichientong, and Huseyin Topaloglu. Assortment optimization and pricing under the multinomial logit model with impatient customers: Sequential recommendation and selection. *Operations research*, 69(5):1509–1532, 2021. 2
- [32] Xiangyu Gao, Stefanus Jasin, Sajjad Najafi, and Huanan Zhang. Joint learning and optimization for multi-product pricing (and ranking) under a general cascade click model. *Management Science*, 2022. 1, 2
- [33] Anindya Ghose and Sha Yang. An empirical analysis of search engine advertising: Sponsored search in electronic markets. *Management science*, 55(10):1605–1622, 2009. 1
- [34] Fan Guo, Chao Liu, Anitha Kannan, Tom Minka, Michael Taylor, Yi-Min Wang, and Christos Faloutsos. Click chain model in web search. In *Proceedings of the 18th international conference on World wide web*, pages 11–20, 2009. 1, 2
- [35] Fan Guo, Chao Liu, and Yi Min Wang. Efficient multiple-click models in web search. In *Proceedings of the second acm international conference on web search and data mining*, pages 124–131, 2009. 1
- [36] Sumeet Katariya, Branislav Kveton, Csaba Szepesvari, and Zheng Wen. Dcm bandits: Learning to rank with multiple clicks. In *International Conference on Machine Learning*, pages 1215–1224. PMLR, 2016. 2
- [37] Sergei Koulayev. Search for differentiated products: identification and estimation. *The RAND Journal of Economics*, 45(3):553–575, 2014. 1

- [38] Branislav Kveton, Csaba Szepesvari, Zheng Wen, and Azin Ashkan. Cascading bandits: Learning to rank in the cascade model. In *International Conference on Machine Learning*, pages 767–776. PMLR, 2015. 1, 2
- [39] Lihong Li, Wei Chu, John Langford, and Robert E Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, pages 661–670, 2010. 6
- [40] Yuan Liang, Chunlin Huang, Xiuguo Bao, and Ke Xu. Sequential dynamic event recommendation in event-based social networks: An upper confidence bound approach. *Information Sciences*, 542:1–23, 2021. 1, 2, 3, 4, 8, 15
- [41] Nan Liu, Yuhang Ma, and Huseyin Topaloglu. Assortment optimization under the multinomial logit model with sequential offerings. *INFORMS Journal on Computing*, 32(3):835–853, 2020. 2
- [42] Qing Liu and Neeraj Arora. Efficient choice designs for a consider-then-choose model. *Marketing Science*, 30(2):321–338, 2011. 1
- [43] Siddharth Mahajan and Garrett Van Ryzin. Stocking retail assortments under dynamic consumer substitution. *Operations Research*, 49(3):334–351, 2001. 1, 2
- [44] Rad Niazadeh, Negin Golrezaei, Joshua R Wang, Fransisca Susan, and Ashwinkumar Badanidiyuru. Online learning via offline greedy algorithms: Applications in market design and optimization. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 737–738, 2021. 3
- [45] Herbert Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58(5):527–535, 1952. 2
- [46] Aleksandrs Slivkins. Introduction to multi-armed bandits. *arXiv preprint arXiv:1904.07272*, 2019. 6, 8
- [47] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018. 2
- [48] Theja Tulabandhula, Deeksha Sinha, and Prasoon Patidar. Multi-purchase behavior: Modeling and optimization. *arXiv preprint arXiv:2006.08055*, 2020. 2
- [49] Raluca M Ursu. The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions. *Marketing Science*, 37(4):530–552, 2018. 1, 3
- [50] Mengmeng Wang, Xun Zhang, and Xiaolong Li. Multiple-purchase choice model: Estimation and optimization. *Available at SSRN 4147705*, 2022. 2
- [51] Zheng Wen, Branislav Kveton, Michal Valko, and Sharan Vaswani. Online influence maximization under independent cascade model with semi-bandit feedback. *Advances in neural information processing systems*, 30, 2017. 27
- [52] Renzhe Xu, Xingxuan Zhang, Peng Cui, Bo Li, Zheyang Shen, and Jiazheng Xu. Regulatory instruments for fair personalized pricing. In *Proceedings of the ACM Web Conference 2022*, pages 4–15, 2022. 1
- [53] Heng Zhang, Hossein Piri, Woonghee Tim Huh, and Hongmin Li. Assortment optimization under multiple-discrete customer choices. *Available at SSRN 3988982*, 2021. 2
- [54] Shi Zong, Hao Ni, Kenny Sung, Nan Rosemary Ke, Zheng Wen, and Branislav Kveton. Cascading bandits for large-scale recommendation problems. *arXiv preprint arXiv:1603.05359*, 2016. 2, 6, 8

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes] See Section A.
 - (c) Did you discuss any potential negative societal impacts of your work? [Yes] See Section A.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes] See Assumptions 3.1, 5.1, and 5.2.
 - (b) Did you include complete proofs of all theoretical results? [Yes] See Section C.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See Section B.
 - (b) Did you specify all the training details (e.g., data splits, hyper-parameters, how they were chosen)? [Yes] See Section B.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See Figure 1 and Figure 2 in the main body and other figures in the appendix.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Section B.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [Yes] The dataset url is included in the paper.
 - (b) Did you mention the license of the assets? [Yes] The license is included in the paper.
 - (c) Did you include any new assets either in the supplemental material or as a URL? [No]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes] It is a public dataset with the CC BY-NC 4.0 license.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [Yes] The data has been anonymized.
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Limitations and Societal Impacts

One potential limitation is that we suppose the distribution of the attention span and purchase budget obey the geometric distribution. While the assumption is common [13, 40], we can relax it in future work, similar to [17, 14].

This paper does not have critical negative societal impacts because the purpose of the paper is to design an optimal ranking policy for online retailers. Besides the online retailers, consumers may also benefit because Equation (5) indicates that a product with a higher purchase probability may appear in a higher order in the ranking list. As a result, consumers may find the products shown in the ranking list highly relevant.

B More Experimental Results

B.1 Implementation Details

Baselines

- **Single Purchase** [13]. Cao and Sun [13] suppose that consumers always purchase at most one product. As a result, we implement them by setting $\tilde{s}_t = 0$ for all consumers. The estimations of $\tilde{\lambda}_{t,k}$ and \tilde{q}_t are the same as our method.
- **Keep Viewing** [40]. Liang et al. [40] suppose that consumers always keep viewing the list if they purchase one product. To get the optimal ranking policy with the knowledge of consumers' characteristics, we can sort the products with value $(\lambda_k r_k)/(1 - q - (1 - q)\lambda_k)$ [40, Theorem 1]. Afterward, we use similar techniques to estimate $\tilde{\lambda}_{t,k}$ and \tilde{q}_t and maximize the revenue in the online settings.
- **Explore Then Exploit A**. Following [13], we implement two explore-then-exploit-based algorithms as baselines. For the Explore Then Exploit A method, we first adopt an exploration stage to ensure that each product is displayed to consumers at least $\delta \log T$ times. (Here δ is a hyper-parameter.) To be specific, in the exploration stage, we rank the products according to their current numbers of displays. Afterward, we use the estimated parameters to provide the optimal ranking policy according to Theorem 4.1 in the main body of the paper.
- **Explore Then Exploit B**. This method is a variant of the Explore Then Exploit A method. In detail, in the exploration stage, after ranking the products according to their number of displays, we further rank the products that have been displayed to consumers at least $\delta \log T$ times according to Theorem 4.1 in the main body.

Hyper-parameters We add hyper-parameters ξ_Λ , ξ_Q , and ξ_W to determine the level of exploration for our method and two baselines (Single purchase and Keep viewing). The other two baselines do not include the hyper-parameter ξ because they do not have the dynamic exploration step. In detail, in the non-contextual setting, the parameters $\tilde{\lambda}_{t,k}$, \tilde{q}_t , and \tilde{w}_t of our method are calculated as

$$\tilde{\lambda}_{t,k} \triangleq \min \left\{ 1, \hat{\lambda}_{t,k} + \xi_\Lambda \sqrt{\frac{\log t}{C_{t,k}}} \right\}, \tilde{q}_t \triangleq \min \left\{ 1 - \epsilon_Q, \hat{q}_t + \xi_Q \sqrt{\frac{\log t}{D_t^Q}} \right\}, \tilde{w}_t \triangleq \min \left\{ \hat{q}_t, \hat{w}_t + \xi_W \sqrt{\frac{\log t}{D_t^W}} \right\}. \quad (20)$$

In the contextual setting, the parameters are also calculated according to the hyper-parameters ξ_Λ , ξ_Q , and ξ_W ,

$$\begin{aligned} \tilde{\lambda}_{t,k} &= \text{Proj}_{[0,1]} \left(y_{t,k}^\top \hat{\beta}_{t-1}^\Lambda + \xi_\Lambda \cdot \tau(t-1, m_y, \alpha_\Lambda) \|y_{t,k}\|_{(\Sigma_{t-1}^\Lambda)^{-1}} \right), \\ \tilde{q}_t &= \text{Proj}_{[0,1-\epsilon_Q]} \left(x_t^\top \hat{\beta}_{t-1}^Q + \xi_Q \cdot \tau(t-1, m_x, \alpha_Q) \|x_t\|_{(\Sigma_{t-1}^Q)^{-1}} \right), \\ \tilde{w}_t &= \text{Proj}_{[0,\tilde{q}_t]} \left(z_t^\top \hat{\beta}_{t-1}^W + \xi_W \cdot \tau(t-1, m_z^2, \alpha_W) \|z_t\|_{(\Sigma_{t-1}^W)^{-1}} \right). \end{aligned} \quad (21)$$

For the synthetic dataset, we search the hyper-parameters $\xi_\Lambda \in \{0.1, 0.3, 0.5\}$, $\xi_Q, \xi_W \in \{0.05, 0.1, 0.2\}$, $\delta \in \{0.5, 1, 2, 5, 10\}$, and $\alpha_Q = \alpha_W = \alpha_\Lambda \in \{0.01, 0.1, 1\}$. For the semi-synthetic dataset, we search the hyper-parameters $\xi_\Lambda \in \{0.1, 0.3, 0.5\}$, $\xi_Q, \xi_W \in \{0.05, 0.1, 0.2\}$,

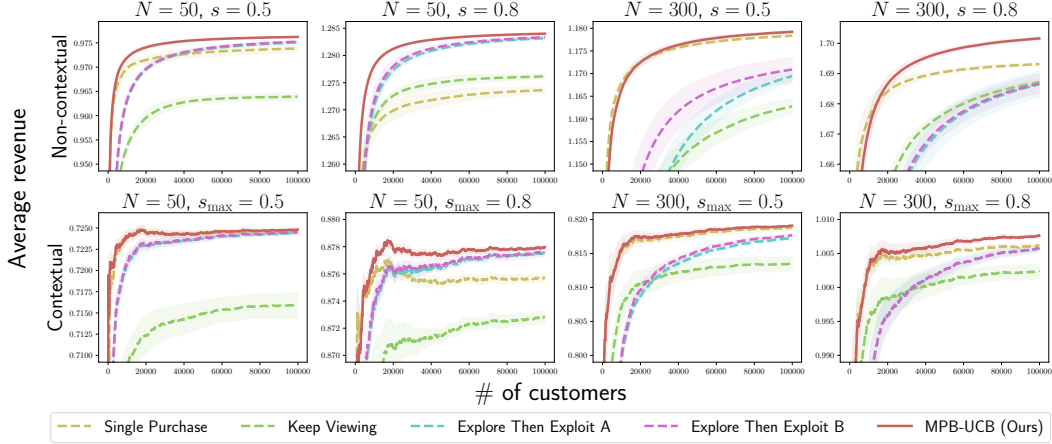


Figure 3: The average revenue metric on the synthetic data.

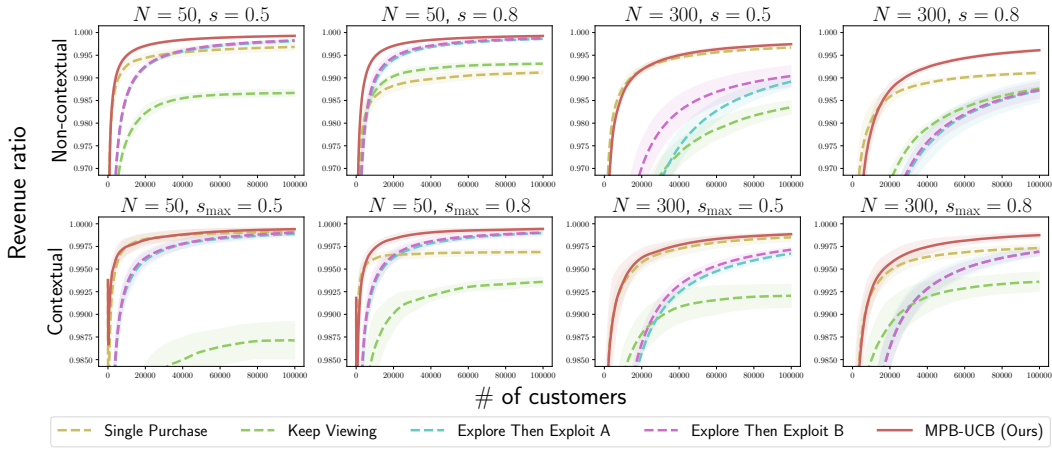


Figure 4: The revenue ratio metric on the synthetic data.

$\delta \in \{0.5, 1, 2, 5, 10\}$, and $\alpha_Q = \alpha_W = \alpha_\Lambda \in \{0.1, 1, 10\}$. The grid search code is based on the NNI package².

B.2 More Results on the Synthetic and Semi-synthetic Data

We further report two more metrics to verify the effectiveness of our method. In detail, Average revenue(T) measures the average revenue achieved by different algorithms till the T -th consumer. Revenue ratio(T) measures the ratio of the achieved revenue over the optimal ranking policy. Formally, the metrics are calculated as follows.

$$\begin{aligned} \text{Average revenue}(T) &= \frac{1}{T} \sum_{t=1}^T \text{Revenue}(\sigma_t; \lambda_t, q_t, s_t), \\ \text{Revenue ratio}(T) &= \frac{\sum_{t=1}^T \text{Revenue}(\sigma_t; \lambda_t, q_t, s_t)}{\sum_{t=1}^T \text{Revenue}(\sigma_t^*; \lambda_t, q_t, s_t)}. \end{aligned} \quad (22)$$

For both metrics, higher values indicate better performances. The results of these two metrics on both the synthetic and Semi-synthetic datasets are shown in Figures 3, 4, 5, and 6. These figures further demonstrate that our method outperforms all baselines in both contextual and non-contextual settings with various parameters (*i.e.*, s , s_{\max} , and N).

²<https://github.com/microsoft/nni>. Licence: MIT License.




Figure 5: The average revenue metric on the semi-synthetic data.




Figure 6: The revenue ratio metric on the semi-synthetic data.

C Omitted Proofs

C.1 Proof of Theorem 4.1

Proof. We drop all the subscript t for simplicity. We first write the expected revenue function.

$$\begin{aligned}
 & \text{Revenue}(\boldsymbol{\sigma}; \boldsymbol{\lambda}, q, s) \\
 &= \mathbb{E}_{V \sim \text{Geo}(1-q), B \sim \text{Geo}(1-s)} [R(V, B; \boldsymbol{\lambda}, \boldsymbol{\sigma})] \\
 &= \sum_{v=1}^N \sum_{b=1}^N P(V=v)P(B=b) \left(\sum_{Q \subseteq [v], |Q| \leq b} P^{\text{buy}}(Q; v, b, \boldsymbol{\lambda}, \boldsymbol{\sigma}) \left(\sum_{k \in Q} r_{\sigma_k} \right) \right) \\
 &= \sum_{k=1}^N r_{\sigma_k} \left(\sum_{v=1}^N \sum_{b=1}^N P(V=v)P(B=b) \left(\sum_{Q \subseteq [N], k \in Q} P^{\text{buy}}(Q; v, b, \boldsymbol{\lambda}, \boldsymbol{\sigma}) \right) \right) \\
 &= \sum_{k=1}^N r_{\sigma_k} \left(\sum_{v=k}^N \sum_{b=1}^N P(V=v)P(B=b) \left(\sum_{Q \subseteq [N], k \in Q} P^{\text{buy}}(Q; v, b, \boldsymbol{\lambda}, \boldsymbol{\sigma}) \right) \right).
 \end{aligned} \tag{23}$$

For any $Q_k \subseteq [N]$ that satisfies $k \in Q_k$, define $Q_k^- \triangleq \{i \in Q_k : i < k\}$ and $Q_k^+ \triangleq \{i \in Q_k : i > k\}$. We further define

$$G_k^-(Q_k^-) \triangleq \left(\prod_{i \in Q_k^-} \lambda_{\sigma_i} \right) \left(\prod_{1 \leq i < k, i \notin Q_k^-} (1 - \lambda_{\sigma_i}) \right). \tag{24}$$

As a result,

$$\begin{aligned}
& \sum_{v=k}^N \sum_{b=1}^N P(V=v)P(B=b) \left(\sum_{Q \subseteq [N], k \in Q} P^{\text{buy}}(Q; v, b, \boldsymbol{\lambda}, \boldsymbol{\sigma}) \right) \\
&= \sum_{v=k}^N \sum_{b=1}^N P(V=v)P(B=b) \left(\sum_{Q_k^- \subseteq [k-1]} \sum_{Q_k^+ \subseteq [N] \setminus [N]} P^{\text{buy}}(Q; v, b, \boldsymbol{\lambda}, \boldsymbol{\sigma}) \right) \\
&= \lambda_{\sigma_k} \sum_{v=k}^N \sum_{b=1}^N P(V=v)P(B=b) \left(\sum_{Q_k^- \subseteq [k-1]} \sum_{Q_k^+ \subseteq [N] \setminus [N]} G_k^-(Q_k^-) P_{k+1}^{\text{buy}}(Q_k^+; v, b - |Q_k^-| - 1, \boldsymbol{\lambda}, \boldsymbol{\sigma}) \right). \tag{25}
\end{aligned}$$

Here $P_{k+1}^{\text{buy}}(Q_k^+; v, b - |Q_k^-| - 1, \boldsymbol{\lambda}, \boldsymbol{\sigma})$ in the second equality denotes the probability of purchasing Q_k^+ with attention span v and purchase budget b after purchasing $|Q_k^-| + 1$ products, *i.e.*,

$$P_{k+1}^{\text{buy}}(Q_k^+; v, b - |Q_k^-| - 1, \boldsymbol{\lambda}, \boldsymbol{\sigma}) = \begin{cases} \left(\prod_{i \in Q_k^+} \lambda_{\sigma_i} \right) \left(\prod_{k < i \leq v} (1 - \lambda_{\sigma_i}) \right), & \text{if } |Q_k^+| < b - |Q_k^-| - 1. \\ \left(\prod_{i \in Q_k^+} \lambda_{\sigma_i} \right) \left(\prod_{k < i \leq \max(Q_k^+)} (1 - \lambda_{\sigma_i}) \right), & \text{if } |Q_k^+| = b - |Q_k^-| - 1. \\ 0, & \text{if } |Q_k^+| > b - |Q_k^-| - 1. \end{cases} \tag{26}$$

As a result, when $|Q_k^-|, v, b, \boldsymbol{\lambda}$, and $\boldsymbol{\sigma}$ are fixed, by definition, we have

$$\sum_{Q_k^+ \subseteq [N] \setminus [N]} P_{k+1}^{\text{buy}}(Q_k^+; v, b - |Q_k^-| - 1, \boldsymbol{\lambda}, \boldsymbol{\sigma}) = 1 \tag{27}$$

Hence,

$$\begin{aligned}
& \lambda_{\sigma_k} \sum_{v=k}^N \sum_{b=1}^N P(V=v)P(B=b) \left(\sum_{Q_k^- \subseteq [k-1]} \sum_{Q_k^+ \subseteq [N] \setminus [N]} G_k^-(Q_k^-) P_{k+1}^{\text{buy}}(Q_k^+; v, b - |Q_k^-| - 1, \boldsymbol{\lambda}, \boldsymbol{\sigma}) \right) \\
&= \lambda_{\sigma_k} \sum_{v=k}^N \sum_{b=1}^N P(V=v)P(B=b) \left(\sum_{u=0}^{b-1} \left(\sum_{Q_k^- \subseteq [k-1], |Q_k^-|=u} G_k^-(Q_k^-) \right) \left(\sum_{Q_k^+ \subseteq [N] \setminus [N]} P_{k+1}^{\text{buy}}(Q_k^+; v, b - u - 1, \boldsymbol{\lambda}, \boldsymbol{\sigma}) \right) \right) \\
&= \lambda_{\sigma_k} \sum_{v=k}^N \sum_{b=1}^N P(V=v)P(B=b) \left(\sum_{u=0}^{b-1} \left(\sum_{Q_k^- \subseteq [k-1], |Q_k^-|=u} G_k^-(Q_k^-) \right) \right) \\
&= \lambda_{\sigma_k} P(V \geq k) \sum_{b=1}^N P(B=b) \left(\sum_{u=0}^{b-1} \left(\sum_{Q_k^- \subseteq [k-1], |Q_k^-|=u} G_k^-(Q_k^-) \right) \right) \\
&= \lambda_{\sigma_k} P(V \geq k) \sum_{u=0}^{k-1} P(B \geq u+1) \left(\sum_{Q_k^- \subseteq [k-1], |Q_k^-|=u} G_k^-(Q_k^-) \right). \tag{28}
\end{aligned}$$

For any $k \geq 1$ and $0 \leq u \leq k-1$, define

$$H_{\boldsymbol{\sigma}}(k, u; \boldsymbol{\lambda}) \triangleq \sum_{Q_k^- \subseteq [k-1], |Q_k^-|=u} G_k^-(Q_k^-). \tag{29}$$

Intuitively, $H_{\boldsymbol{\sigma}}(k, u; \boldsymbol{\lambda})$ measures the probability that a consumer purchases exactly u products from the position index 1 to $k-1$ if ignoring the attention span and purchase budget. As a result,

$$\begin{aligned}
\text{Revenue}(\boldsymbol{\sigma}; \boldsymbol{\lambda}, q, s) &= \sum_{k=1}^N r_{\sigma_k} \lambda_{\sigma_k} P(V \geq k) \left(\sum_{u=0}^{k-1} P(B \geq u+1) H_{\boldsymbol{\sigma}}(k, u; \boldsymbol{\lambda}) \right) \\
&= \sum_{k=1}^N r_{\sigma_k} \lambda_{\sigma_k} q^{k-1} \left(\sum_{u=0}^{k-1} s^u H_{\boldsymbol{\sigma}}(k, u; \boldsymbol{\lambda}) \right). \tag{30}
\end{aligned}$$

In addition, it is easy to verify that $H_\sigma(k, u; \lambda)$ has the following recurrence relation

$$\begin{aligned}
H_\sigma(1, 0; \lambda) &= 1. \\
H_\sigma(k, 0; \lambda) &= \prod_{i=1}^{k-1} (1 - \lambda_{\sigma_i}), \quad H_\sigma(k, k-1; \lambda) = \prod_{i=1}^{k-1} \lambda_{\sigma_i}, \quad \forall k \geq 2. \\
H_\sigma(k, u; \lambda) &= (1 - \lambda_{\sigma_{k-1}}) H_\sigma(k-1, u; \lambda) + \lambda_{\sigma_{k-1}} H_\sigma(k-1, u-1; \lambda), \quad \forall k \geq 2, 0 < u < k-1.
\end{aligned} \tag{31}$$

From the definition of G_k^- and H_σ , it is easy to check that the value of $H_\sigma(k, u; \lambda)$ does not depend on the order of $\sigma_1, \sigma_2, \dots, \sigma_{k-1}$. As a result, for any $1 \leq i < N$, let σ be a ranking policy that differs from σ^* only in the i -th and $(i+1)$ -th position, *i.e.*,

$$\sigma_j = \begin{cases} \sigma_{i+1}^*, & \text{if } j = i. \\ \sigma_i^*, & \text{if } j = i+1. \\ \sigma_j^*, & \text{otherwise.} \end{cases} \tag{32}$$

Due to the optimality of σ^* , we have

$$\begin{aligned}
0 &\leq \text{Revenue}(\sigma^*; \lambda, q, s) - \text{Revenue}(\sigma; \lambda, q, s) \\
&= r_{\sigma_i^*} \lambda_{\sigma_i^*} q^{i-1} \left(\sum_{u=0}^{i-1} s^u H_{\sigma^*}(i, u; \lambda) \right) + r_{\sigma_{i+1}^*} \lambda_{\sigma_{i+1}^*} q^i \left(\sum_{u=0}^i s^u H_{\sigma^*}(i+1, u; \lambda) \right) \\
&\quad - r_{\sigma_i} \lambda_{\sigma_i} q^{i-1} \left(\sum_{u=0}^{i-1} s^u H_\sigma(i, u; \lambda) \right) + r_{\sigma_{i+1}} \lambda_{\sigma_{i+1}} q^i \left(\sum_{u=0}^i s^u H_\sigma(i+1, u; \lambda) \right) \\
&= r_{\sigma_i^*} \lambda_{\sigma_i^*} q^{i-1} \left(\sum_{u=0}^{i-1} s^u H_{\sigma^*}(i, u; \lambda) \right) - r_{\sigma_i} \lambda_{\sigma_i} q^{i-1} \left(\sum_{u=0}^{i-1} s^u H_\sigma(i, u; \lambda) \right) \\
&\quad + r_{\sigma_{i+1}^*} \lambda_{\sigma_{i+1}^*} q^i \left((1 - \lambda_{\sigma_i^*}) \sum_{u=0}^{i-1} s^u H_{\sigma^*}(i, u; \lambda) + \lambda_{\sigma_i^*} \sum_{u=1}^i s^u H_{\sigma^*}(i, u-1; \lambda) \right) \\
&\quad - r_{\sigma_{i+1}} \lambda_{\sigma_{i+1}} q^i \left((1 - \lambda_{\sigma_i}) \sum_{u=0}^{i-1} s^u H_\sigma(i, u; \lambda) + \lambda_{\sigma_i} \sum_{u=1}^i s^u H_\sigma(i, u-1; \lambda) \right) \\
&= q^{i-1} \left(\sum_{u=0}^{i-1} s^u H_\sigma(i, u; \lambda) \right) \cdot \\
&\quad \left(r_{\sigma_i^*} \lambda_{\sigma_i^*} (1 - q + q(1-s)\lambda_{\sigma_{i+1}^*}) - r_{\sigma_{i+1}} \lambda_{\sigma_{i+1}} (1 - q + q(1-s)\lambda_{\sigma_i}) \right).
\end{aligned} \tag{33}$$

Hence,

$$\frac{r_{\sigma_i^*} \lambda_{\sigma_i^*}}{1 - q + q(1-s)\lambda_{\sigma_i^*}} \geq \frac{r_{\sigma_{i+1}} \lambda_{\sigma_{i+1}}}{1 - q + q(1-s)\lambda_{\sigma_{i+1}}}. \tag{34}$$

Now the claim follows. \square

C.2 Proof of Lemma 5.1

Proof. Random variables $\gamma_{t,k}$ are independent of each other by the assumptions that the behaviors of different consumers are independent and consumer interests in any two products are independent. For $\eta_{t,k}$, suppose the consumer purchases u products after viewing the k -th product, then

$$\begin{aligned}
P(\eta_{t,k} = 1) &= P(V_t > k, B_t > u \mid V_t \geq k, B_t > u) \\
&= \frac{P(V_t > k, B_t > u)}{P(V_t \geq k, B_t > u)} = \frac{P(V_t > k) P(B_t > u)}{P(V_t \geq k) P(B_t > u)} = \frac{q_t^k}{q_t^{k-1}} = q_t.
\end{aligned} \tag{35}$$

As a result, $P(\eta_{t,k} = 1)$ does not depend on the historical purchase activities and is always q_t . By the assumption that the behaviors of different consumers are independent, $\eta_{t,k}$ are independent Bernoulli random variables and $\eta_{t,k} \sim \text{Ber}(q_t)$.

Similarly, for $\mu_{t,k}$, suppose the consumer purchases u products after viewing the k -th product, then

$$\begin{aligned} P(\mu_{t,k} = 1) &= P(V_t > k, B_t > u \mid V_t \geq k, B_t \geq u) \\ &= \frac{P(V_t > k, B_t > u)}{P(V_t \geq k, B_t \geq u)} = \frac{P(V_t > k) P(B_t > u)}{P(V_t \geq k) P(B_t \geq u)} = \frac{q_t^k s_t^u}{q_t^{k-1} s_t^{u-1}} = q_t s_t. \end{aligned} \quad (36)$$

As a result, $P(\eta_{t,k} = 1)$ does not depend on the historical purchase activities. Hence, by the assumption that the behaviors of different consumers are independent, $\mu_{t,k}$ are independent Bernoulli random variables and $\mu_{t,k} \sim \text{Ber}(q_t s_t)$. \square

C.3 Proof of Proposition 5.2

Proof. For each timestamp t , define the favorable event as

$$\xi_t \triangleq \left\{ |q - \hat{q}_t| \leq \sqrt{\frac{2 \log t}{D_t^Q}}, |qs - \hat{w}_t| \leq \sqrt{\frac{2 \log t}{D_t^W}}, |\lambda_k - \hat{\lambda}_{t,k}| \leq \sqrt{\frac{2 \log t}{C_{t,k}}}, \forall k \in [N] \right\}. \quad (37)$$

Let $\bar{\xi}_t$ denote the complement event of ξ_t and $P(\xi_t) = 1 - P(\bar{\xi}_t)$. To estimate $P(\bar{\xi}_t)$, we need some further notations. Let ξ_t^Q denote the event $\left\{ |q - \hat{q}_t| > \sqrt{2 \log t / D_t^Q} \right\}$, ξ_t^W denote the event $\left\{ |qs - \hat{w}_t| > \sqrt{2 \log t / D_t^W} \right\}$, and ξ_t^Λ denote the event $\left\{ \exists k \in [N], |\lambda_k - \hat{\lambda}_{t,k}| > \sqrt{2 \log t / C_{t,k}} \right\}$. As a result, $\bar{\xi}_t = \xi_t^Q \cup \xi_t^W \cup \xi_t^\Lambda$ and

$$P(\bar{\xi}_t) \leq P(\xi_t^Q) + P(\xi_t^W) + P(\xi_t^\Lambda) \quad (38)$$

by the union bound. We first estimate $P(\xi_t^Q)$. Let q'_s be the s -th observation of q and $\hat{q}'_s = \sum_{u=1}^s q'_u / s$ be the average of the top s observations. According to Lemma 5.1 and Hoeffding's inequality, we can get that for any $s \in [Nt]$,

$$P\left(|q - \hat{q}'_s| > \sqrt{\frac{2 \log t}{s}}\right) \leq 2t^{-4}. \quad (39)$$

As a result, by the union bound,

$$\begin{aligned} P(\xi_t^Q) &= P\left(|q - \hat{q}'_{D_t^Q}| > \sqrt{\frac{2 \log t}{D_t^Q}}\right) \leq \sum_{s=1}^{tN} P\left(|q - \hat{q}'_s| > \sqrt{\frac{2 \log t}{s}}, D_t^Q = s\right) \\ &\leq \sum_{s=1}^{tN} P\left(|q - \hat{q}'_s| > \sqrt{\frac{2 \log t}{s}}\right) \leq \sum_{s=1}^{tN} 2t^{-4} = 2Nt^{-3}. \end{aligned} \quad (40)$$

Using similar techniques, we can get that

$$P(\xi_t^W) \leq 2Nt^{-3} \quad \text{and} \quad P(\xi_t^\Lambda) \leq 2Nt^{-3}. \quad (41)$$

Now the claim follows by combining Equations (38), (40), (41). \square

C.4 Proof of Theorem 5.3

We first define

$$O_{t,k}^\Lambda \triangleq \mathbb{I}[k \leq \Phi_t], \quad O_{t,k}^Q \triangleq \mathbb{I}[k \leq \Phi_t, \gamma_{t,k} = 0], \quad O_{t,k}^W \triangleq \mathbb{I}[k \leq \Phi_t, \gamma_{t,k} = 1]. \quad (42)$$

We also need the following two propositions. The subscript t are omitted in the two propositions for simplicity.

Proposition C.1. *Suppose $\tilde{\lambda} \succeq \lambda$, $\tilde{q} \geq q$, and $\tilde{s}\tilde{q} \geq sq$, then*

$$\text{Revenue}(\sigma^*; \lambda, q, s) \leq \text{Revenue}(\sigma^*; \tilde{\lambda}, \tilde{q}, \tilde{s}). \quad (43)$$

Proof of Proposition C.1. For any $\sigma \vdash [N]$, $1 \leq n \leq k$, $0 \leq u \leq k - n$, define

$$J_\sigma(n, k, u; \lambda) \triangleq \sum_{Q \subseteq [k-1] \setminus [n-1], |Q|=u} \left(\prod_{i \in Q} \lambda_{\sigma_i} \right) \left(\prod_{n \leq i < k, i \notin Q} \lambda_{\sigma_i} \right). \quad (44)$$

Intuitively, $J_\sigma(n, k, u; \lambda)$ measures the probability that a consumer purchases exactly u products from the position index n to $k - 1$ if ignoring the attention span and purchase budget. As a result, $J_\sigma(1, k, u; \lambda) = H_\sigma(k, u; \lambda)$. In addition, it is easy to check that

$$J_\sigma(n, n, 0; \lambda) = 1, \quad \forall 1 \leq n \leq N.$$

$$J_\sigma(n-1, k, u; \lambda) = \begin{cases} (1 - \lambda_{\sigma_{n-1}})J_\sigma(n, k, u; \lambda), & \text{if } u = 0. \\ \lambda_{\sigma_{n-1}}J_\sigma(n, k, u-1; \lambda), & \text{if } u = k - n + 1. \\ (1 - \lambda_{\sigma_{n-1}})J_\sigma(n, k, u; \lambda) + \lambda_{\sigma_{n-1}}J_\sigma(n, k, u-1; \lambda), & \text{if } 0 < u < k - n + 1. \end{cases} \quad (45)$$

In addition, for any $n \geq 1$, define

$$\tilde{R}_\sigma(n; \lambda, q, s) \triangleq \sum_{k=n}^N r_{\sigma_k} \lambda_{\sigma_k} q^{k-n} \left(\sum_{u=0}^{k-n} s^u J_\sigma(n, k, u; \lambda) \right). \quad (46)$$

Intuitively, $\tilde{R}_\sigma(n; \lambda, q, s)$ measures the expected revenue if the top $n - 1$ products in the ranking list are not purchased. It is easy to verify that $\tilde{R}_\sigma(1; \lambda, q, s) = \text{Revenue}(\sigma; \lambda, q, s)$. Furthermore,

$$\begin{aligned} & \tilde{R}_\sigma(n-1; \lambda, q, s) \\ &= \sum_{k=n-1}^N r_{\sigma_k} \lambda_{\sigma_k} q^{k-n+1} \left(\sum_{u=0}^{k-n+1} s^u J_\sigma(n-1, k, u; \lambda) \right) \\ &= r_{\sigma_{n-1}} \lambda_{\sigma_{n-1}} + \sum_{k=n}^N r_{\sigma_k} \lambda_{\sigma_k} q^{k-n+1} \left(\sum_{u=0}^{k-n+1} s^u J_\sigma(n-1, k, u; \lambda) \right) \\ &= r_{\sigma_{n-1}} \lambda_{\sigma_{n-1}} + \sum_{k=n}^N r_{\sigma_k} \lambda_{\sigma_k} q^{k-n+1} \left(\sum_{u=0}^{k-n} s^u (1 - \lambda_{\sigma_{n-1}}) J_\sigma(n, k, u; \lambda) + \sum_{u=1}^{k-n+1} s^u \lambda_{\sigma_{n-1}} J_\sigma(n, k, u-1; \lambda) \right) \\ &= r_{\sigma_{n-1}} \lambda_{\sigma_{n-1}} + q \left(1 - (1-s)\lambda_{\sigma_{n-1}} \right) \left(\sum_{k=n}^N r_{\sigma_k} \lambda_{\sigma_k} q^{k-n} \left(\sum_{u=0}^{k-n} s^u J_\sigma(n, k, u; \lambda) \right) \right) \\ &= r_{\sigma_{n-1}} \lambda_{\sigma_{n-1}} + (q - q(1-s)\lambda_{\sigma_{n-1}}) \tilde{R}_\sigma(n; \lambda, q, s). \end{aligned} \quad (47)$$

Now we prove $\tilde{R}_{\sigma^*}(n; \lambda, q, s) \leq \tilde{R}_{\sigma^*}(n; \tilde{\lambda}, \tilde{q}, \tilde{s})$, $\forall n \in [N]$ by induction. Firstly, it is easy to check that $\tilde{R}_{\sigma^*}(N; \lambda, q, s) \leq \tilde{R}_{\sigma^*}(N; \tilde{\lambda}, \tilde{q}, \tilde{s})$. Suppose the condition holds for $\tilde{R}_{\sigma^*}(n; \lambda, q, s) \leq \tilde{R}_{\sigma^*}(n; \tilde{\lambda}, \tilde{q}, \tilde{s})$. Consider the case for $n - 1$. Notice that $\tilde{R}_{\sigma^*}(n-1; \lambda, q, s) \geq \tilde{R}_{\sigma^*}(n; \lambda, q, s)$ (otherwise it will increase the revenue if removing the $(n - 1)$ -th product in the list), we can get that

$$\left(r_{\sigma_{n-1}} - q(1-s)\tilde{R}_{\sigma^*}(n; \lambda, q, s) \right) \lambda_{\sigma_{n-1}} = (1-q)\tilde{R}_{\sigma^*}(n; \lambda, q, s) \geq 0. \quad (48)$$

Hence $r_{\sigma_{n-1}} - q(1-s)\tilde{R}_{\sigma^*}(n; \lambda, q, s) \geq 0$. As a result,

$$\begin{aligned} & \tilde{R}_{\sigma^*}(n-1; \lambda, q, s) \\ &= r_{\sigma_{n-1}} \lambda_{\sigma_{n-1}} + (q - q(1-s)\lambda_{\sigma_{n-1}}) \tilde{R}_{\sigma^*}(n; \lambda, q, s) \\ &= q\tilde{R}_{\sigma^*}(n; \lambda, q, s) + \left(r_{\sigma_{n-1}} - q(1-s)\tilde{R}_{\sigma^*}(n; \lambda, q, s) \right) \lambda_{\sigma_{n-1}} \\ &\leq q\tilde{R}_{\sigma^*}(n; \lambda, q, s) + \left(r_{\sigma_{n-1}} - q(1-s)\tilde{R}_{\sigma^*}(n; \lambda, q, s) \right) \tilde{\lambda}_{\sigma_{n-1}} \\ &= r_{\sigma_{n-1}} \tilde{\lambda}_{\sigma_{n-1}} + \left(q - q(1-s)\tilde{\lambda}_{\sigma_{n-1}} \right) \tilde{R}_{\sigma^*}(n; \lambda, q, s) \\ &= r_{\sigma_{n-1}} \tilde{\lambda}_{\sigma_{n-1}} + \left(q \left(1 - \tilde{\lambda}_{\sigma_{n-1}} \right) + qs\tilde{\lambda}_{\sigma_{n-1}} \right) \tilde{R}_{\sigma^*}(n; \lambda, q, s) \\ &\leq r_{\sigma_{n-1}} \tilde{\lambda}_{\sigma_{n-1}} + \left(\tilde{q} \left(1 - \tilde{\lambda}_{\sigma_{n-1}} \right) + \tilde{q}\tilde{s}\tilde{\lambda}_{\sigma_{n-1}} \right) \tilde{R}_{\sigma^*}(n; \lambda, q, s) \\ &\leq r_{\sigma_{n-1}} \tilde{\lambda}_{\sigma_{n-1}} + \left(\tilde{q} - \tilde{q}(1-\tilde{s})\tilde{\lambda}_{\sigma_{n-1}} \right) \tilde{R}_{\sigma^*}(n; \tilde{\lambda}, \tilde{q}, \tilde{s}) = \tilde{R}_{\sigma^*}(n-1; \tilde{\lambda}, \tilde{q}, \tilde{s}). \end{aligned} \quad (49)$$

As a result,

$$\text{Revenue}(\boldsymbol{\sigma}^*; \boldsymbol{\lambda}, q, s) = \tilde{R}_{\boldsymbol{\sigma}^*}(1; \boldsymbol{\lambda}, q, s) \leq \tilde{R}_{\boldsymbol{\sigma}^*}(1; \tilde{\boldsymbol{\lambda}}, \tilde{q}, \tilde{s}) = \text{Revenue}(\boldsymbol{\sigma}^*; \tilde{\boldsymbol{\lambda}}, \tilde{q}, \tilde{s}). \quad (50)$$

□

Proposition C.2. *Under Assumption 3.1, suppose $\tilde{\boldsymbol{\lambda}} \succeq \boldsymbol{\lambda}$, $1 - \epsilon_Q \geq \tilde{q} \geq q$, and $\tilde{s}\tilde{q} \geq sq$, then*

$$\begin{aligned} & \text{Revenue}(\boldsymbol{\sigma}_t; \tilde{\boldsymbol{\lambda}}, \tilde{q}, \tilde{s}) - \text{Revenue}(\boldsymbol{\sigma}_t; \boldsymbol{\lambda}, q, s) \\ & \leq \frac{r_{\max}}{\epsilon_Q} \sum_{k=1}^N \left(|q - \tilde{q}| \mathbb{E}[O_{t,k}^Q] + 3 \left| \lambda_{\sigma_{t,k}} - \tilde{\lambda}_{\sigma_{t,k}} \right| \mathbb{E}[O_{t,k}^A] + |qs - \tilde{q}\tilde{s}| \mathbb{E}[O_{t,k}^W] \right). \end{aligned} \quad (51)$$

Proof of Proposition C.2. For simplicity, we use $\boldsymbol{\sigma}$ to denote $\boldsymbol{\sigma}_t$ in this proof. According to Equation (30),

$$\text{Revenue}(\boldsymbol{\sigma}; \boldsymbol{\lambda}, q, s) = \sum_{k=1}^N r_{\sigma_k} \lambda_{\sigma_k} q^{k-1} \left(\sum_{u=0}^{k-1} s^u H_{\boldsymbol{\sigma}}(k, u; \boldsymbol{\lambda}) \right), \quad (52)$$

and $H_{\boldsymbol{\sigma}}(k, u; \boldsymbol{\lambda})$ is defined in Equation (29). As a result,

$$\begin{aligned} & \text{Revenue}(\boldsymbol{\sigma}; \tilde{\boldsymbol{\lambda}}, \tilde{q}, \tilde{s}) - \text{Revenue}(\boldsymbol{\sigma}; \boldsymbol{\lambda}, q, s) \\ & = \sum_{k=1}^N r_{\sigma_k} \tilde{\lambda}_{\sigma_k} \tilde{q}^{k-1} \left(\sum_{u=0}^{k-1} \tilde{s}^u H_{\boldsymbol{\sigma}}(k, u; \tilde{\boldsymbol{\lambda}}) \right) - \sum_{k=1}^N r_{\sigma_k} \lambda_{\sigma_k} q^{k-1} \left(\sum_{u=0}^{k-1} s^u H_{\boldsymbol{\sigma}}(k, u; \boldsymbol{\lambda}) \right) \\ & = \sum_{k=1}^N r_{\sigma_k} \left(\tilde{\lambda}_{\sigma_k} \tilde{q}^{k-1} \left(\sum_{u=0}^{k-1} \tilde{s}^u H_{\boldsymbol{\sigma}}(k, u; \tilde{\boldsymbol{\lambda}}) \right) - \lambda_{\sigma_k} q^{k-1} \left(\sum_{u=0}^{k-1} s^u H_{\boldsymbol{\sigma}}(k, u; \boldsymbol{\lambda}) \right) \right). \end{aligned} \quad (53)$$

Note that

$$\mathbb{E}[O_{t,k}^A] = P(V_t \geq k) \sum_{u=0}^{k-1} P(B_t \geq u+1) H_{\boldsymbol{\sigma}}(k, u; \boldsymbol{\lambda}) = q^{k-1} \sum_{u=0}^{k-1} s^u H_{\boldsymbol{\sigma}}(k, u; \boldsymbol{\lambda}). \quad (54)$$

$$\mathbb{E}[O_{t,k}^Q] = \mathbb{E}[O_{t,k}^A] P(\gamma_{t,k} = 0) = (1 - \lambda_{\sigma_k}) q^{k-1} \sum_{u=0}^{k-1} s^u H_{\boldsymbol{\sigma}}(k, u; \boldsymbol{\lambda}). \quad (55)$$

$$\mathbb{E}[O_{t,k}^W] = \mathbb{E}[O_{t,k}^A] P(\gamma_{t,k} = 1) = \lambda_{\sigma_k} q^{k-1} \sum_{u=0}^{k-1} s^u H_{\boldsymbol{\sigma}}(k, u; \boldsymbol{\lambda}). \quad (56)$$

For any $1 \leq k \leq N$,

$$\begin{aligned} & \left| \tilde{\lambda}_{\sigma_k} \tilde{q}^{k-1} \left(\sum_{u=0}^{k-1} \tilde{s}^u H_{\boldsymbol{\sigma}}(k, u; \tilde{\boldsymbol{\lambda}}) \right) - \lambda_{\sigma_k} q^{k-1} \left(\sum_{u=0}^{k-1} s^u H_{\boldsymbol{\sigma}}(k, u; \boldsymbol{\lambda}) \right) \right| \\ & \leq \left| \tilde{\lambda}_{\sigma_k} - \lambda_{\sigma_k} \right| \mathbb{E}[O_{t,k}^A] + \tilde{\lambda}_{\sigma_k} \left| \tilde{q}^{k-1} \left(\sum_{u=0}^{k-1} \tilde{s}^u H_{\boldsymbol{\sigma}}(k, u; \tilde{\boldsymbol{\lambda}}) \right) - q^{k-1} \left(\sum_{u=0}^{k-1} s^u H_{\boldsymbol{\sigma}}(k, u; \boldsymbol{\lambda}) \right) \right| \\ & \leq \left| \tilde{\lambda}_{\sigma_k} - \lambda_{\sigma_k} \right| \mathbb{E}[O_{t,k}^A] + \left| \tilde{q}^{k-1} \left(\sum_{u=0}^{k-1} \tilde{s}^u H_{\boldsymbol{\sigma}}(k, u; \tilde{\boldsymbol{\lambda}}) \right) - q^{k-1} \left(\sum_{u=0}^{k-1} s^u H_{\boldsymbol{\sigma}}(k, u; \boldsymbol{\lambda}) \right) \right|. \end{aligned} \quad (57)$$

Here the first inequality is due to the triangle inequality of $|\cdot|$ and the second inequality is according to the fact that $\tilde{\lambda}_{\sigma_k} \leq 1$. Now for any $1 \leq k \leq N$, define

$$A_{\boldsymbol{\sigma}}(k; q, s, \boldsymbol{\lambda}) \triangleq q^{k-1} \left(\sum_{u=0}^{k-1} s^u H_{\boldsymbol{\sigma}}(k, u; \boldsymbol{\lambda}) \right). \quad (58)$$

We can get

$$\begin{aligned}
& A_{\sigma}(k; q, s, \boldsymbol{\lambda}) \\
&= q^{k-1} \left(\sum_{u=0}^{k-1} s^u H_{\sigma}(k, u; \boldsymbol{\lambda}) \right) \\
&= q^{k-1} \left(\sum_{u=0}^{k-2} s^u (1 - \lambda_{\sigma_{k-1}}) H_{\sigma}(k-1, u; \boldsymbol{\lambda}) + \sum_{u=1}^{k-1} s^u \lambda_{\sigma_{k-1}} H_{\sigma}(k-1, u-1; \boldsymbol{\lambda}) \right) \quad (59) \\
&= q^{k-1} \left(\sum_{u=0}^{k-2} s^u (1 - \lambda_{\sigma_{k-1}}) H_{\sigma}(k-1, u; \boldsymbol{\lambda}) + s \sum_{u=0}^{k-2} s^u \lambda_{\sigma_{k-1}} H_{\sigma}(k-1, u; \boldsymbol{\lambda}) \right) \\
&= q (1 - \lambda_{\sigma_{k-1}} + s \lambda_{\sigma_{k-1}}) A_{\sigma}(k-1; q, s, \boldsymbol{\lambda}).
\end{aligned}$$

As a result, for any $1 < k \leq N$,

$$\begin{aligned}
& \left| A_{\sigma}(k; q, s, \boldsymbol{\lambda}) - A_{\sigma}(k; \tilde{q}, \tilde{s}, \tilde{\boldsymbol{\lambda}}) \right| \\
&= \left| q (1 - \lambda_{\sigma_{k-1}} + s \lambda_{\sigma_{k-1}}) A_{\sigma}(k-1; q, s, \boldsymbol{\lambda}) - \tilde{q} (1 - \tilde{\lambda}_{\sigma_{k-1}} + \tilde{s} \tilde{\lambda}_{\sigma_{k-1}}) A_{\sigma}(k-1; \tilde{q}, \tilde{s}, \tilde{\boldsymbol{\lambda}}) \right| \quad (60) \\
&\leq \left| q (1 - \lambda_{\sigma_{k-1}} + s \lambda_{\sigma_{k-1}}) - \tilde{q} (1 - \tilde{\lambda}_{\sigma_{k-1}} + \tilde{s} \tilde{\lambda}_{\sigma_{k-1}}) \right| A_{\sigma}(k-1; q, s, \boldsymbol{\lambda}) \\
&\quad + \tilde{q} (1 - \tilde{\lambda}_{\sigma_{k-1}} + \tilde{s} \tilde{\lambda}_{\sigma_{k-1}}) \left| A_{\sigma}(k-1; q, s, \boldsymbol{\lambda}) - A_{\sigma}(k-1; \tilde{q}, \tilde{s}, \tilde{\boldsymbol{\lambda}}) \right|.
\end{aligned}$$

Here the last inequality is due the triangle inequality of $|\cdot|$. Note that

$$\begin{aligned}
& \left| q (1 - \lambda_{\sigma_{k-1}} + s \lambda_{\sigma_{k-1}}) - \tilde{q} (1 - \tilde{\lambda}_{\sigma_{k-1}} + \tilde{s} \tilde{\lambda}_{\sigma_{k-1}}) \right| A_{\sigma}(k-1; q, s, \boldsymbol{\lambda}) \\
&\leq \left(\left| q (1 - \lambda_{\sigma_{k-1}}) - \tilde{q} (1 - \tilde{\lambda}_{\sigma_{k-1}}) \right| + \left| qs \lambda_{\sigma_{k-1}} - \tilde{q} \tilde{s} \tilde{\lambda}_{\sigma_{k-1}} \right| \right) A_{\sigma}(k-1; q, s, \boldsymbol{\lambda}) \\
&\leq \left(|q - \tilde{q}| (1 - \lambda_{\sigma_{k-1}}) + \tilde{q} \left| \lambda_{\sigma_{k-1}} - \tilde{\lambda}_{\sigma_{k-1}} \right| + |qs - \tilde{q} \tilde{s}| \lambda_{\sigma_{k-1}} + \tilde{q} \tilde{s} \left| \lambda_{\sigma_{k-1}} - \tilde{\lambda}_{\sigma_{k-1}} \right| \right) A_{\sigma}(k-1; q, s, \boldsymbol{\lambda}) \\
&\leq \left(|q - \tilde{q}| (1 - \lambda_{\sigma_{k-1}}) + 2 \left| \lambda_{\sigma_{k-1}} - \tilde{\lambda}_{\sigma_{k-1}} \right| + |qs - \tilde{q} \tilde{s}| \lambda_{\sigma_{k-1}} \right) A_{\sigma}(k-1; q, s, \boldsymbol{\lambda}) \\
&= |q - \tilde{q}| \mathbb{E} \left[O_{t, k-1}^Q \right] + 2 \left| \lambda_{\sigma_{k-1}} - \tilde{\lambda}_{\sigma_{k-1}} \right| \mathbb{E} \left[O_{t, k-1}^{\Lambda} \right] + |qs - \tilde{q} \tilde{s}| \mathbb{E} \left[O_{t, k-1}^W \right]. \quad (61)
\end{aligned}$$

Here the first two inequalities are due to the triangle inequality of $|\cdot|$ and the third inequality is due to that fact that $\tilde{q}, \tilde{s} \leq 1$. Combining [Equations \(60\)](#) and [\(61\)](#), we can get

$$\begin{aligned}
& \left| A_{\sigma}(k; q, s, \boldsymbol{\lambda}) - A_{\sigma}(k; \tilde{q}, \tilde{s}, \tilde{\boldsymbol{\lambda}}) \right| \\
&\leq \sum_{u=1}^{k-1} \left(\left(|q - \tilde{q}| \mathbb{E} \left[O_{t, u}^Q \right] + 2 \left| \lambda_{\sigma_u} - \tilde{\lambda}_{\sigma_u} \right| \mathbb{E} \left[O_{t, u}^{\Lambda} \right] + |qs - \tilde{q} \tilde{s}| \mathbb{E} \left[O_{t, u}^W \right] \right) \left(\prod_{i=u+1}^{k-1} \tilde{q} (1 - \tilde{\lambda}_{\sigma_i} + \tilde{s} \tilde{\lambda}_{\sigma_i}) \right) \right) \\
&\leq \sum_{u=1}^{k-1} \left(\left(|q - \tilde{q}| \mathbb{E} \left[O_{t, u}^Q \right] + 2 \left| \lambda_{\sigma_u} - \tilde{\lambda}_{\sigma_u} \right| \mathbb{E} \left[O_{t, u}^{\Lambda} \right] + |qs - \tilde{q} \tilde{s}| \mathbb{E} \left[O_{t, u}^W \right] \right) \left(\prod_{i=u+1}^{k-1} \tilde{q} \right) \right) \\
&= \sum_{u=1}^{k-1} \left(|q - \tilde{q}| \mathbb{E} \left[O_{t, u}^Q \right] + 2 \left| \lambda_{\sigma_u} - \tilde{\lambda}_{\sigma_u} \right| \mathbb{E} \left[O_{t, u}^{\Lambda} \right] + |qs - \tilde{q} \tilde{s}| \mathbb{E} \left[O_{t, u}^W \right] \right) \tilde{q}^{k-u-1}. \quad (62)
\end{aligned}$$

Combining the above inequality with [Equations \(53\) and \(57\)](#), we can get that

$$\begin{aligned}
& \text{Revenue}(\boldsymbol{\sigma}; \tilde{\boldsymbol{\lambda}}, \tilde{q}, \tilde{s}) - \text{Revenue}(\boldsymbol{\sigma}; \boldsymbol{\lambda}, q, s) \\
& \leq \sum_{k=1}^N r_{\max} \left| \tilde{\lambda}_{\sigma_k} - \lambda_{\sigma_k} \right| \mathbb{E} \left[O_{t,k}^\Lambda \right] \\
& \quad + \sum_{k=2}^N r_{\sigma_k} \tilde{\lambda}_{\sigma_k} \left(\sum_{u=1}^{k-1} \left(|q - \tilde{q}| \mathbb{E} \left[O_{t,u}^Q \right] + 2 \left| \lambda_{\sigma_u} - \tilde{\lambda}_{\sigma_u} \right| \mathbb{E} \left[O_{t,u}^\Lambda \right] + |qs - \tilde{q}\tilde{s}| \mathbb{E} \left[O_{t,u}^W \right] \right) \tilde{q}^{k-u-1} \right) \\
& = \sum_{k=1}^N r_{\max} \left| \tilde{\lambda}_{\sigma_k} - \lambda_{\sigma_k} \right| \mathbb{E} \left[O_{t,k}^\Lambda \right] \\
& \quad + \sum_{k=1}^{N-1} \left(|q - \tilde{q}| \mathbb{E} \left[O_{t,k}^Q \right] + 2 \left| \lambda_{\sigma_k} - \tilde{\lambda}_{\sigma_k} \right| \mathbb{E} \left[O_{t,k}^\Lambda \right] + |qs - \tilde{q}\tilde{s}| \mathbb{E} \left[O_{t,k}^W \right] \right) \left(\sum_{i=k+1}^N r_{\sigma_i} \tilde{\lambda}_{\sigma_i} \tilde{q}^{i-k-1} \right) \\
& \leq \sum_{k=1}^N r_{\max} \left| \tilde{\lambda}_{\sigma_k} - \lambda_{\sigma_k} \right| \mathbb{E} \left[O_{t,k}^\Lambda \right] \\
& \quad + \sum_{k=1}^{N-1} r_{\max} \left(|q - \tilde{q}| \mathbb{E} \left[O_{t,k}^Q \right] + 2 \left| \lambda_{\sigma_k} - \tilde{\lambda}_{\sigma_k} \right| \mathbb{E} \left[O_{t,k}^\Lambda \right] + |qs - \tilde{q}\tilde{s}| \mathbb{E} \left[O_{t,k}^W \right] \right) \left(\sum_{i=k+1}^N \tilde{q}^{i-k-1} \right) \\
& \leq \sum_{k=1}^N r_{\max} \left| \tilde{\lambda}_{\sigma_k} - \lambda_{\sigma_k} \right| \mathbb{E} \left[O_{t,k}^\Lambda \right] \\
& \quad + \sum_{k=1}^{N-1} r_{\max} \left(|q - \tilde{q}| \mathbb{E} \left[O_{t,k}^Q \right] + 2 \left| \lambda_{\sigma_k} - \tilde{\lambda}_{\sigma_k} \right| \mathbb{E} \left[O_{t,k}^\Lambda \right] + |qs - \tilde{q}\tilde{s}| \mathbb{E} \left[O_{t,k}^W \right] \right) \frac{1}{\epsilon_Q} \\
& \leq \frac{r_{\max}}{\epsilon_Q} \sum_{k=1}^N \left(|q - \tilde{q}| \mathbb{E} \left[O_{t,k}^Q \right] + 3 \left| \lambda_{\sigma_k} - \tilde{\lambda}_{\sigma_k} \right| \mathbb{E} \left[O_{t,k}^\Lambda \right] + |qs - \tilde{q}\tilde{s}| \mathbb{E} \left[O_{t,k}^W \right] \right). \tag{63}
\end{aligned}$$

Now the claim follows. \square

Proof of [Theorem 5.3](#). For each timestamp t , define the favorable event as

$$\xi_t \triangleq \left\{ |q - \hat{q}_t| \leq \sqrt{\frac{2 \log t}{D_t^Q}}, |qs - \hat{w}_t| \leq \sqrt{\frac{2 \log t}{D_t^W}}, |\lambda_k - \hat{\lambda}_{t,k}| \leq \sqrt{\frac{2 \log t}{C_{t,k}}}, \forall k \in [N] \right\}. \tag{64}$$

Let $\bar{\xi}_t$ denote the complement event of ξ_t .

$$\begin{aligned}
\mathbb{E}[\text{Reg}_T] & = \sum_{t=1}^T \mathbb{E} \left[(\text{Revenue}(\boldsymbol{\sigma}^*; \boldsymbol{\lambda}, q, s) - \text{Revenue}(\boldsymbol{\sigma}_t; \boldsymbol{\lambda}, q, s)) \mathbb{I}[\xi_{t-1}] \right] \\
& \quad + \sum_{t=1}^T \mathbb{E} \left[(\text{Revenue}(\boldsymbol{\sigma}^*; \boldsymbol{\lambda}, q, s) - \text{Revenue}(\boldsymbol{\sigma}_t; \boldsymbol{\lambda}, q, s)) \mathbb{I}[\bar{\xi}_{t-1}] \right]. \tag{65}
\end{aligned}$$

On the one hand, consider the first term in the RHS of [Equation \(65\)](#), when the event ξ_{t-1} is satisfied, we have $q \leq \tilde{q}_{t-1}$, $qs \leq \tilde{w}_{t-1}$, and $\lambda_k \leq \tilde{\lambda}_{t,k}$, $\forall k \in [N]$. As a result, according to [Proposition C.1](#),

$$\text{Revenue}(\boldsymbol{\sigma}^*; \boldsymbol{\lambda}, q, s) \leq \text{Revenue}(\boldsymbol{\sigma}^*; \tilde{\boldsymbol{\lambda}}_{t-1}, \tilde{q}_{t-1}, \tilde{s}_{t-1}) \leq \text{Revenue}(\boldsymbol{\sigma}_t; \tilde{\boldsymbol{\lambda}}_{t-1}, \tilde{q}_{t-1}, \tilde{s}_{t-1}). \tag{66}$$

Hence, according to [Proposition C.2](#),

$$\begin{aligned}
& \sum_{t=1}^T \mathbb{E} [(\text{Revenue}(\boldsymbol{\sigma}^*; \boldsymbol{\lambda}, q, s) - \text{Revenue}(\boldsymbol{\sigma}_t; \boldsymbol{\lambda}, q, s)) \mathbb{I}[\xi_{t-1}]] \\
& \leq \sum_{t=1}^T \mathbb{E} [(\text{Revenue}(\boldsymbol{\sigma}_t; \bar{\boldsymbol{\lambda}}_{t-1}, \bar{q}_{t-1}, \bar{s}_{t-1}) - \text{Revenue}(\boldsymbol{\sigma}_t; \boldsymbol{\lambda}, q, s)) \mathbb{I}[\xi_{t-1}]] \\
& \leq \frac{r_{\max}}{\epsilon_Q} \sum_{t=1}^T \mathbb{E} \left[\sum_{k=1}^N (|q - \bar{q}_{t-1}| \mathbb{E}[O_{t,k}^Q] + 3 |\lambda_{\sigma_{t,k}} - \bar{\lambda}_{t-1, \sigma_{t,k}}| \mathbb{E}[O_{t,k}^\Lambda] + |qs - \bar{q}_{t-1} \bar{s}_{t-1}| \mathbb{E}[O_{t,k}^W]) \mathbb{I}[\xi_{t-1}] \right] \\
& \leq \frac{2r_{\max}}{\epsilon_Q} \sum_{t=1}^T \mathbb{E} \left[\sum_{k=1}^N \left(\min \left\{ \sqrt{\frac{2 \log t}{D_{t-1}^Q}}, 1 \right\} O_{t,k}^Q + 3 \min \left\{ \sqrt{\frac{2 \log t}{C_{t-1, \sigma_{t,k}}}}, 1 \right\} O_{t,k}^\Lambda + \min \left\{ \sqrt{\frac{2 \log t}{D_{t-1}^W}}, 1 \right\} O_{t,k}^W \right) \right] \\
& \leq \frac{2r_{\max} \sqrt{2 \log T}}{\epsilon_Q} \sum_{t=1}^T \mathbb{E} \left[\sum_{k=1}^N \left(\min \left\{ \sqrt{\frac{1}{D_{t-1}^Q}}, 1 \right\} O_{t,k}^Q + 3 \min \left\{ \sqrt{\frac{1}{C_{t-1, \sigma_{t,k}}}}, 1 \right\} O_{t,k}^\Lambda + \min \left\{ \sqrt{\frac{1}{D_{t-1}^W}}, 1 \right\} O_{t,k}^W \right) \right]. \tag{67}
\end{aligned}$$

Let $\kappa_t^Q \triangleq \sum_{k=1}^N O_{t,k}^Q \leq N$ and we have $D_t^Q = D_{t-1}^Q + \kappa_t^Q$. As a result,

$$\begin{aligned}
& \mathbb{E} \left[\sum_{t=1}^T \sum_{k=1}^N \min \left\{ \sqrt{\frac{1}{D_{t-1}^Q}}, 1 \right\} O_{t,k}^Q \right] \\
& = \mathbb{E} \left[\sum_{t=1}^T \min \left\{ \sqrt{\frac{1}{D_{t-1}^Q}}, 1 \right\} \kappa_t^Q \right] \\
& \leq 1 + N + \sum_{t=2}^T N \sqrt{\frac{1}{(t-1)N}} \\
& \leq 1 + N + \sqrt{N} \left(1 + \int_{t=1}^{T-1} \sqrt{\frac{1}{t}} dt \right) \leq C_1 N + C_2 \sqrt{TN}, \tag{68}
\end{aligned}$$

for some constants C_1 and C_2 . Similarly, we have

$$\mathbb{E} \left[\sum_{t=1}^T \sum_{k=1}^N \min \left\{ \sqrt{\frac{1}{D_{t-1}^W}}, 1 \right\} O_{t,k}^W \right] \leq C_1 N + C_2 \sqrt{TN}. \tag{69}$$

In addition, let $o_{t,k}^\Lambda \triangleq O_{t, \sigma_{t,k}}^\Lambda$ represent whether the product k is viewed in the t -th timestamp. As a result, $C_{t,k} = C_{t-1,k} + o_{t,k}^\Lambda$. Therefore,

$$\begin{aligned}
& \mathbb{E} \left[\sum_{t=1}^T \sum_{k=1}^N \min \left\{ \sqrt{\frac{1}{C_{t-1, \sigma_{t,k}}}}, 1 \right\} O_{t,k}^\Lambda \right] \\
& = \mathbb{E} \left[\sum_{k=1}^N \sum_{t=1}^T \min \left\{ \sqrt{\frac{1}{C_{t-1,k}}}, 1 \right\} o_{t,k}^\Lambda \right] \\
& \leq \mathbb{E} \left[\sum_{k=1}^N \left(1 + \sum_{t=1}^T \sqrt{\frac{1}{t}} \right) \right] \leq C_3 N \sqrt{T}, \tag{70}
\end{aligned}$$

for a constant C_3 . On the other hand, according to [Proposition 5.2](#),

$$\begin{aligned}
& \sum_{t=1}^T \mathbb{E} [(\text{Revenue}(\boldsymbol{\sigma}^*; \boldsymbol{\lambda}, q, s) - \text{Revenue}(\boldsymbol{\sigma}_t; \boldsymbol{\lambda}, q, s)) \mathbb{I}[\bar{\xi}_{t-1}]] \\
& \leq \text{Revenue}(\boldsymbol{\sigma}^*; \boldsymbol{\lambda}, q, s) \sum_{t=1}^T \mathbb{E} [\mathbb{I}[\bar{\xi}_{t-1}]] \leq \frac{r_{\max}}{\epsilon_Q} \sum_{t=1}^T \mathbb{E} [\mathbb{I}[\bar{\xi}_{t-1}]] \\
& \leq \frac{r_{\max}}{\epsilon_Q} \left(1 + \sum_{t=1}^T \frac{6N}{t^3} \right) \leq C_4 \cdot \frac{Nr_{\max}}{\epsilon_Q}. \tag{71}
\end{aligned}$$

for a constant C_4 . Now the claim follows after combining [Equations \(67\), \(68\), \(69\), \(70\), and \(71\)](#). \square

C.5 Proof of Proposition 5.4

Proof. The claim follows from standard routines in linear bandit methods. We first consider the $\hat{\beta}_t^\Lambda - \beta^\Lambda$ term. We can view the tuple (t, k) as the new timestamp. Let $\epsilon_{t,k} \triangleq y_{t,\sigma_{t,k}}^\top \beta^\Lambda - \gamma_{t,k}$. According to Lemma 5.1, $\mathbb{E}[\epsilon_{t,k} | y_{t,\sigma_{t,k}}, \mathcal{H}_{t,k}] = 0$ where $\mathcal{H}_{t,k}$ denotes the historical observations before the t -th consumer views the k -th product. In addition, it is bounded in $[q_t - 1, q_t]$ and hence sub-gaussian with $R = 1/2$. As a result, according to [1, Theorem 2], under Assumption 5.2 with parameter U , for any $\delta \in (0, 1)$, with probability at least $1 - \delta$, for any $t \geq 0$,

$$\begin{aligned} \left\| \hat{\beta}_t^\Lambda - \beta^\Lambda \right\|_{\Sigma_t^\Lambda} &\leq \frac{1}{2} \sqrt{2 \log \frac{\det(\Sigma_t^\Lambda)^{1/2} \det(\alpha_\Lambda \mathbf{I}_{m_y})^{-1/2}}{\delta}} + \alpha_\Lambda^{1/2} U \\ &= \frac{1}{2} \sqrt{\log \frac{\det(\Sigma_t^\Lambda)}{\delta^2 \alpha_\Lambda^{m_y}}} + \alpha_\Lambda^{1/2} U. \end{aligned} \quad (72)$$

According to [1, Lemma 10], we have

$$\det(\Sigma_t^\Lambda) \leq \left(\alpha_\Lambda + \frac{Nt}{m_y} \right)^{m_y}. \quad (73)$$

As a result, by letting $\delta = t^{-2}$, for all $t > 0$, with probability at least $1 - t^{-2}$,

$$\left\| \hat{\beta}_t^\Lambda - \beta^\Lambda \right\|_{\Sigma_t^\Lambda} \leq \frac{1}{2} \sqrt{m_y \log \left(1 + \frac{Nt}{m_y \alpha_\Lambda} \right)} + 4 \log t + \alpha_\Lambda^{1/2} U \leq \tau(t, m_y, \alpha_\Lambda). \quad (74)$$

Similarly, we have that with probability at least $1 - t^{-2}$,

$$\left\| \hat{\beta}_t^Q - \beta^Q \right\|_{\Sigma_t^Q} \leq \frac{1}{2} \sqrt{m_x \log \left(1 + \frac{Nt}{m_x \alpha_Q} \right)} + 4 \log t + \alpha_Q^{1/2} U \leq \tau(t, m_x, \alpha_Q). \quad (75)$$

In addition, because

$$\|x_t\| \leq 1 \implies \|z_t\| \leq 1 \quad \text{and} \quad \|\beta^S\|, \|\beta^Q\| \leq U \implies \|\beta^W\| \leq U^2, \quad (76)$$

we have that with probability at least $1 - t^{-2}$,

$$\left\| \hat{\beta}_t^W - \beta^W \right\|_{\Sigma_t^W} \leq \frac{1}{2} \sqrt{m_x^2 \log \left(1 + \frac{Nt}{m_x^2 \alpha_W} \right)} + 4 \log t + \alpha_W^{1/2} U^2 \leq \tau(t, m_x^2, \alpha_W). \quad (77)$$

Now the claim follows. \square

C.6 Proof of Theorem 5.5

Proof. Similar to the proof of Theorem 5.3, for each timestamp t , define the favorable event as

$$\xi_t \triangleq \left\{ \left\| \hat{\beta}_t^\Lambda - \beta^\Lambda \right\|_{\Sigma_t^\Lambda} \leq \tau(t, m_y, \alpha_\Lambda), \left\| \hat{\beta}_t^Q - \beta^Q \right\|_{\Sigma_t^Q} \leq \tau(t, m_x, \alpha_Q), \text{ and } \left\| \hat{\beta}_t^W - \beta^W \right\|_{\Sigma_t^W} \leq \tau(t, m_x^2, \alpha_W) \right\}. \quad (78)$$

According to Proposition 5.4, we have $P(\xi_t) \geq 1 - 3t^{-2}$. Let $\bar{\xi}_t$ denote the complement event of ξ_t .

$$\begin{aligned} \mathbb{E}[\text{Reg}_T | \mathbf{X}_T, \mathbf{Y}_T] &= \sum_{t=1}^T \mathbb{E}[(\text{Revenue}(\sigma_t^*; \boldsymbol{\lambda}_t, q_t, s_t) - \text{Revenue}(\sigma_t; \boldsymbol{\lambda}_t, q_t, s_t)) \mathbb{I}[\xi_{t-1}]] \\ &\quad + \sum_{t=1}^T \mathbb{E}[(\text{Revenue}(\sigma_t^*; \boldsymbol{\lambda}_t, q_t, s_t) - \text{Revenue}(\sigma_t; \boldsymbol{\lambda}_t, q_t, s_t)) \mathbb{I}[\bar{\xi}_{t-1}]]. \end{aligned} \quad (79)$$

On the one hand, consider the first term in the RHS of Equation (79), when the event ξ_t is satisfied, we can show that $\tilde{q}_t \geq q_t$. On the one hand, if $\tilde{q}_t = 1 - \epsilon_Q$, it is obvious that $\tilde{q}_t \geq q_t$. On the other

hand, otherwise,

$$\begin{aligned}
\tilde{q}_t - q_t &= x_t^\top \left(\hat{\beta}_{t-1}^Q - \beta^Q \right) + \tau(t-1, m_x, \alpha_Q) \|x_t\|_{(\Sigma_{t-1}^Q)^{-1}} \\
&= x_t^\top \left(\Sigma_{t-1}^Q \right)^{-1/2} \left(\Sigma_{t-1}^Q \right)^{1/2} \left(\hat{\beta}_{t-1}^Q - \beta^Q \right) + \tau(t-1, m_x, \alpha_Q) \|x_t\|_{(\Sigma_{t-1}^Q)^{-1}} \\
&= \left(\left(\Sigma_{t-1}^Q \right)^{-1/2} x_t \right)^\top \left(\Sigma_{t-1}^Q \right)^{1/2} \left(\hat{\beta}_{t-1}^Q - \beta^Q \right) + \tau(t-1, m_x, \alpha_Q) \|x_t\|_{(\Sigma_{t-1}^Q)^{-1}} \\
&\geq - \|x_t\|_{(\Sigma_{t-1}^Q)^{-1}} \left\| \hat{\beta}_{t-1}^Q - \beta^Q \right\|_{\Sigma_{t-1}^Q} + \tau(t-1, m_x, \alpha_Q) \|x_t\|_{(\Sigma_{t-1}^Q)^{-1}} \geq 0.
\end{aligned} \tag{80}$$

Similarly, we have that $w_t \leq \tilde{w}_t$ and $\lambda_{t,k} \leq \tilde{\lambda}_{t,k}, \forall k \in [N]$. As a result, according to [Proposition C.1](#),

$$\text{Revenue}(\sigma_t^*; \lambda_t, q_t, s_t) \leq \text{Revenue}(\sigma_t^*; \tilde{\lambda}_t, \tilde{q}_t, \tilde{s}_t) \leq \text{Revenue}(\sigma_t; \tilde{\lambda}_t, \tilde{q}_t, \tilde{s}_t). \tag{81}$$

Hence, according to [Proposition C.2](#),

$$\begin{aligned}
&\sum_{t=1}^T \mathbb{E} \left[\left(\text{Revenue}(\sigma_t^*; \lambda_t, q_t, s_t) - \text{Revenue}(\sigma_t; \lambda_t, q_t, s_t) \right) \mathbb{I}[\xi_t] \right] \\
&\leq \sum_{t=1}^T \mathbb{E} \left[\left(\text{Revenue}(\sigma_t; \tilde{\lambda}_t, \tilde{q}_t, \tilde{s}_t) - \text{Revenue}(\sigma_t; \lambda_t, q_t, s_t) \right) \mathbb{I}[\xi_t] \right] \\
&\leq \frac{r_{\max}}{\epsilon_Q} \sum_{t=1}^T \mathbb{E} \left[\sum_{k=1}^N \left(|q_t - \tilde{q}_t| \mathbb{E} \left[O_{t,k}^Q \right] + 3 \left| \lambda_{t,\sigma_{t,k}} - \tilde{\lambda}_{t,\sigma_{t,k}} \right| \mathbb{E} \left[O_{t,k}^\Lambda \right] + |w_t - \tilde{w}_t| \mathbb{E} \left[O_{t,k}^W \right] \right) \mathbb{I}[\xi_t] \right] \\
&\leq \frac{2r_{\max}}{\epsilon_Q} \sum_{t=1}^T \mathbb{E} \left[\sum_{k=1}^N \left(\tau(t-1, m_x, \alpha_Q) \|x_t\|_{(\Sigma_{t-1}^Q)^{-1}} O_{t,k}^Q + 3\tau(t-1, m_y, \alpha_\Lambda) \|y_{t,\sigma_{t,k}}\|_{(\Sigma_{t-1}^\Lambda)^{-1}} O_{t,k}^\Lambda \right. \right. \\
&\quad \left. \left. + \tau(t-1, m_x^2, \alpha_W) \|z_t\|_{(\Sigma_{t-1}^W)^{-1}} O_{t,k}^W \right) \mathbb{I}[\xi_{t-1}] \right].
\end{aligned} \tag{82}$$

We consider the $\sum_{t=1}^T \mathbb{E} \left[\sum_{k=1}^N \tau(t-1, m_y, \alpha_\Lambda) \|y_{t,\sigma_{t,k}}\|_{(\Sigma_{t-1}^\Lambda)^{-1}} O_{t,k}^\Lambda \right]$ term first and this part is similar to the proof of [\[51, Lemma 1\]](#). Define $g_{t,k} = \|y_{t,\sigma_{t,k}}\|_{(\Sigma_{t-1}^\Lambda)^{-1}}$ for all $t \in [T], k \in [N]$ that satisfy $O_{t,k}^\Lambda = 1$. Because

$$\Sigma_t^\Lambda = \sum_{u=1}^t \sum_{k=1}^{\Phi_u} y_{u,\sigma_{u,k}} y_{u,\sigma_{u,k}}^\top + \alpha_\Lambda \mathbf{I}_{m_y} = \Sigma_{t-1}^\Lambda + \sum_{k=1}^{\Phi_t} y_{t,\sigma_{t,k}} y_{t,\sigma_{t,k}}^\top, \quad \Sigma_0^\Lambda = \alpha_\Lambda \mathbf{I}_{m_y}, \tag{83}$$

we have for all $1 \leq k \leq \Phi_t$,

$$\det(\Sigma_t^\Lambda) \geq \det\left(\Sigma_{t-1}^\Lambda + y_{t,\sigma_{t,k}} y_{t,\sigma_{t,k}}^\top\right) = \det(\Sigma_{t-1}^\Lambda) (1 + g_{t,k}^2). \tag{84}$$

As a result,

$$\left(\frac{\det(\Sigma_t^\Lambda)}{\det(\Sigma_{t-1}^\Lambda)} \right)^N \geq \left(\frac{\det(\Sigma_t^\Lambda)}{\det(\Sigma_{t-1}^\Lambda)} \right)^{\Phi_t} = \prod_{k=1}^{\Phi_t} (1 + g_{t,k}^2). \tag{85}$$

Hence,

$$(\det(\Sigma_t^\Lambda))^N \geq \alpha_\Lambda^{N m_y} \prod_{u=1}^t \prod_{k=1}^{\Phi_u} (1 + g_{u,k}^2). \tag{86}$$

On the other hand, due to [Assumption 5.2](#),

$$\begin{aligned}
\text{trace}(\Sigma_t^\Lambda) &= \text{trace} \left(\sum_{u=1}^t \sum_{k=1}^{\Phi_u} y_{u,\sigma_{u,k}} y_{u,\sigma_{u,k}}^\top + \alpha_\Lambda \mathbf{I}_{m_y} \right) = m_y \alpha_\Lambda + \sum_{u=1}^t \sum_{k=1}^{\Phi_u} \|y_{u,\sigma_{u,k}}\|^2 \\
&\leq m_y \alpha_\Lambda + tN.
\end{aligned} \tag{87}$$

According to the [1, Lemma 10], we have $\text{trace}(\Sigma_t^\Lambda)/m_y \geq (\det(\Sigma_t^\Lambda))^{1/m_y}$. As a result,

$$\left(\frac{m_y \alpha_\Lambda + tN}{m_y}\right)^{Nm_y} \geq \alpha_\Lambda^{Nm_y} \prod_{u=1}^t \prod_{k=1}^{\Phi_t} (1 + g_{u,k}^2). \quad (88)$$

By taking the logarithm, we can get that

$$\sum_{u=1}^t \sum_{k=1}^{\Phi_t} \log(1 + g_{u,k}^2) \leq Nm_y \log\left(1 + \frac{tN}{m_y \alpha_\Lambda}\right). \quad (89)$$

Since

$$\begin{aligned} g_{t,k}^2 &= \|y_{t,\sigma_{t,k}}\|_{(\Sigma_{t-1}^\Lambda)^{-1}}^2 = y_{t,\sigma_{t,k}}^\top (\Sigma_{t-1}^\Lambda)^{-1} y_{t,\sigma_{t,k}} \\ &\leq y_{t,\sigma_{t,k}}^\top (\Sigma_0^\Lambda)^{-1} y_{t,\sigma_{t,k}} = \alpha_\Lambda^{-1} \|y_{t,\sigma_{t,k}}\|^2 \leq \alpha_\Lambda^{-1}, \end{aligned} \quad (90)$$

we have $g_{t,k}^2 \leq \alpha_\Lambda^{-1} \log(1 + g_{t,k}^2) / \log(1 + \alpha_\Lambda^{-1})$. As a result,

$$\sum_{u=1}^t \sum_{k=1}^{\Phi_t} g_{u,k}^2 \leq \frac{1}{\alpha_\Lambda \log(1 + \alpha_\Lambda^{-1})} \sum_{u=1}^t \sum_{k=1}^{\Phi_t} \log(1 + g_{u,k}^2) \leq \frac{Nm_y \log\left(1 + \frac{tN}{m_y \alpha_\Lambda}\right)}{\alpha_\Lambda \log(1 + \alpha_\Lambda^{-1})}. \quad (91)$$

Hence, according to the Cauchy–Schwarz inequality, we have

$$\begin{aligned} \sum_{t=1}^T \sum_{k=1}^N \|y_{t,\sigma_{t,k}}\|_{(\Sigma_{t-1}^\Lambda)^{-1}} O_{t,k}^\Lambda &\leq \sqrt{\left(\sum_{t=1}^T \sum_{k=1}^{\Phi_t} g_{t,k}^2\right) \left(\sum_{t=1}^T \sum_{k=1}^{\Phi_t} O_{t,k}^\Lambda\right)} \\ &\leq \sqrt{\frac{TN^2 m_y \log\left(1 + \frac{TN}{m_y \alpha_\Lambda}\right)}{\alpha_\Lambda \log(1 + \alpha_\Lambda^{-1})}}. \end{aligned} \quad (92)$$

As a result, we have

$$\begin{aligned} &\sum_{t=1}^T \sum_{k=1}^N \tau(t-1, m_y, \alpha_\Lambda) \|y_{t,k}\|_{(\Sigma_{t-1}^\Lambda)^{-1}} O_{t,k}^\Lambda \\ &\leq \tau(T-1, m_y, \alpha_\Lambda) \sum_{t=1}^T \sum_{k=1}^N \|y_{t,k}\|_{(\Sigma_{t-1}^\Lambda)^{-1}} O_{t,k}^\Lambda \\ &\leq \tau(T-1, m_y, \alpha_\Lambda) \sqrt{\frac{TN^2 m_y \log\left(1 + \frac{TN}{m_y \alpha_\Lambda}\right)}{\alpha_\Lambda \log(1 + \alpha_\Lambda^{-1})}} \leq \chi(T, m_y, \alpha_\Lambda). \end{aligned} \quad (93)$$

Using similar techniques, we can prove that

$$\begin{aligned} \sum_{t=1}^T \sum_{k=1}^N \tau(t-1, m_x, \alpha_Q) \|x_t\|_{(\Sigma_{t-1}^Q)^{-1}} O_{t,k}^Q &\leq \tau(T-1, m_x, \alpha_Q) \sqrt{\frac{TN^2 m_x \log\left(1 + \frac{TN}{m_x \alpha_Q}\right)}{\alpha_Q \log(1 + \alpha_Q^{-1})}}, \\ &\leq \chi(T, m_x, \alpha_Q), \\ \sum_{t=1}^T \sum_{k=1}^N \tau(t-1, m_x^2, \alpha_W) \|z_t\|_{(\Sigma_{t-1}^Q)^{-1}} O_{t,k}^W &\leq \tau(T-1, m_x^2, \alpha_W) \sqrt{\frac{TN^2 m_x^2 \log\left(1 + \frac{TN}{m_x^2 \alpha_Q}\right)}{\alpha_W \log(1 + \alpha_W^{-1})}}, \\ &\leq \chi(T, m_x^2, \alpha_W). \end{aligned} \quad (94)$$

On the other hand,

$$\begin{aligned}
& \sum_{t=1}^T \mathbb{E} [(\text{Revenue}(\boldsymbol{\sigma}_t^*; \boldsymbol{\lambda}_t, q_t, s_t) - \text{Revenue}(\boldsymbol{\sigma}_t; \boldsymbol{\lambda}_t, q_t, s_t)) \mathbb{I}[\bar{\xi}_{t-1}]] \\
& \leq \text{Revenue}(\boldsymbol{\sigma}_t^*; \boldsymbol{\lambda}_t, q_t, s_t) \sum_{t=1}^T \mathbb{E} [\mathbb{I}[\bar{\xi}_{t-1}]] \\
& \leq \frac{r_{\max}}{\epsilon_Q} \sum_{t=1}^T \frac{3}{t^2} \leq \frac{C_1 r_{\max}}{\epsilon_Q},
\end{aligned} \tag{95}$$

for a constant C_1 . Now the claim follows after combining [Equations \(82\), \(93\), \(94\), and \(95\)](#). \square