Assortment Optimization for the Multinomial Logit Model with Repeated Customer Interactions

Chenhao Wang,[a,b] Ningyuan Chen,[c] Pin Gao,[d] Yao Wang,[e,*] Su Zhao[f]

[a] School of Economics and Management, Tongji University, Shanghai 200092, China

[b] Laboratory of High Quality Urban Development and Strategic Decision, Tongji University, Shanghai 200092, China, chenhaowang@tongji.edu.cn

[c] Rotman School of Management, University of Toronto, Toronto, Ontario M5S 3E6, Canada, ningyuan.chen@utoronto.ca

[d] School of Data Science, School of Management and Economics, The Chinese University of Hong Kong, Shenzhen, Guangdong 518172, China, gaopin@cuhk.edu.cn

[e] School of Management, Xi'an Jiaotong University, Xi'an, Shaanxi 710049, China, yao.s.wang@gmail.com

[f] Meituan, Beijing, 100102, China, zhaosu04@meituan.com

This paper presents the multinomial logit (MNL) model with repeated customer interactions. In each period, the same customer selects a product from the recommended assortment or opts out. From the seller's perspective, the choice probability in the current period is updated based on the purchase history of the customer. We derive this conditional choice probability and study the adaptive assortment recommendation strategy. Although the problem is generally intractable, we discover the structures of the optimal assortment when the customer interacts with the seller for two periods and the available products to recommend are identical in the two periods. For a general number of periods, we find that the optimal fixed assortments that are not adapted to the purchase history can achieve 50% of the optimal expected revenue, and the approximation ratio increases to 68.47% when the available products across periods are disjoint. Using real-world datasets, we demonstrate that the model with repeated customer interactions can better predict the purchase behavior and generate higher revenues.

Key words: repeated customer interactions, multinomial logit, sequential recommendation, assortment optimization, discrete choice models

1. Introduction

Assortment planning is a core component in many business models. By planning a set of products to display, the seller aims at extracting a higher revenue or profit. This is usually achieved by modeling the choice behavior of customers over a given assortment with discrete choice models. In the literature, the focus on discrete choice models and assortment optimization has been mostly on a *single interaction*: The seller offers an assortment to a customer, who can be seen as an independent draw from a population and whose choice probabilities follow the discrete choice model. The customer leaves permanently after choosing a product from the assortment or decides not to purchase anything. The objective is the expected revenue or profit for this interaction. (See Section 2 for a discussion of

* Corresponding Author

recent papers that extend the setup.) Arguably, the framework captures a majority of the business scenarios in which assortment planning is required. In most offline settings such as planning the shelf space in a store, the seller cannot personalize the assortment for the arriving customer. Moreover, there is hardly any information for the seller to identify if a customer is a returning or new one. In the former case, the purchase history is usually not recorded. Therefore, the simplifying assumption in the literature that treats customers as independent and identically distributed in a single interaction is fully justifiable and hardly contested.

Recently, thanks to the popularity of online retailing, both presumptions have been challenged. Recommender systems allow sellers to personalize the displayed assortment to customers browsing the website. Cookies and account profiles can be used to record customer purchase history, which largely reveals customer preferences. In other words, the seller could and should take into account the purchase history of individual customers and recommend personalized assortments accordingly. The potential of assortment optimization for customers who interact with the seller repeatedly has been realized by several companies, leading to novel business models. A prime example is the online travel platform Meituan, which tailors hotel options in real-time as users navigate their app based on their booking history. See a detailed case study of Meituan's dataset in Section 7.2. Furthermore, online personal styling services such as *Stitch Fix* and *Wantable* are promoting a *subscription box* business model, which is shown in Figure 1. In summary, subscribers may receive a selection of clothes and accessories every few weeks and decide to keep what they like while returning the rest. Over time, the past purchases of a customer can be used to customize the style of the clothing in future offerings. If a box sent to the customer is treated as a recommended assortment, then the assortment decision is made for the same customer every few weeks. Therefore, as a salient feature of both Stitch Fix and Wantable, individual customers interact with the seller repeatedly. Conceptually, the seller offers an assortment to the customer in each period, from which the customer may or may not choose a product. In the next period, a similar interaction is repeated though the seller may offer a different assortment. This feature brings a new dimension to the model, i.e., the learning of customer preferences over time.

To see how it shifts the paradigm of discrete choice modeling, consider the well-known multinomial logit (MNL) model that assumes the utility of product i for customer j is $u_i + \varepsilon_{ij}$, where u_i is a deterministic product-related quantity known to the seller and ε_{ij} is an idiosyncratic Gumbel random variable unknown to the seller. After an interaction between the seller and the customer, say the customer purchased a product from the assortment in the first period, the information can be used for the seller to infer $\varepsilon_{\cdot j}$ of this very customer. In other words, in the next interaction, the choice probability needs to be updated. Imagine the extreme case: if the same assortment is offered, then with probability one, the same product will be chosen because the utilities for the products remain

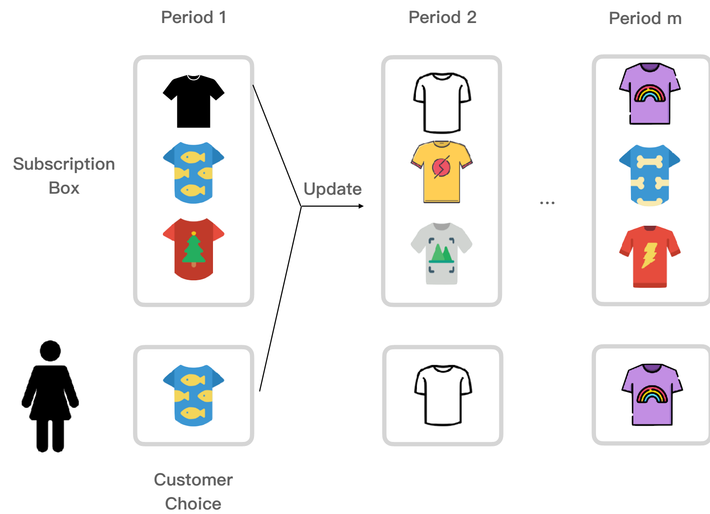

Figure 1 **The subscription box business model exemplified by Stitch Fix and Wantable**

constant over time for the same customer. (We consider changing utilities of purchased products in Section 6.) In particular, the Gumbel random utilities are not redrawn for the products that have been offered before or the no-purchase option. This phenomenon greatly deviates from the standard MNL model with a single interaction.

In this paper, we adapt the MNL model to repeated choices made by the same customer. We update the outcomes of past interactions and calculate the conditional choice probabilities given an assortment in the current period. The aim is to find a recommendation policy to maximize the expected revenue over multiple periods. Note that the recommended assortment needs to be adapted to the outcome in the previous periods, and technically speaking, the problem is a dynamic program. This paper provides structures, approximate algorithms, and performance improvement compared to a fixed assortment recommendation policy (not adapted to the purchase history) in several scenarios.

Our results are summarized as follows:

- When customers interact with the seller in two periods, we analytically derive the conditional choice probability in the second period. The ramification of outcomes in the first period leads to a number of cases, including no-purchase in the first period, the product purchased in period one not available in period two, and the product purchased in period one being available in period two and optimal (or not) to be offered again. The optimal assortment under each case is studied. We show that in some cases the optimization is tractable, while in others it results in NP-hard problems.
- We then study the joint assortment optimization for both periods. Note that the assortment provided in the first period impacts the outcome and the optimal conditional assortment in the second period. We consider two special settings. When the available products in two periods are disjoint, we show the revenue improvement over the optimal fixed policy can be as high as

25% and develop an approximate algorithm. When the set of products in two periods is the same, we discover structural properties and show the joint assortment optimization problem is polynomially solvable. The optimal adaptive policy recommends a revenue-ordered assortment in the first period. If a product is purchased, then in the second period, only this product is recommended again. Otherwise, another revenue-ordered assortment is displayed. In this case, the optimal adaptive policy can extract as high as 50% more revenue compared with the optimal fixed policy.

- With multiple periods having disjoint sets of available products, we provide a closed-form expression for the conditional choice probability given the historical purchase behavior. For the joint assortment optimization problem, we show the fixed recommendation strategy (all assortments are planned beforehand) performs well, and can guarantee at least 68.47% of the expected revenue. Moreover, if all products are popular compared with the no-purchase option, the result can be strengthened. If we allow for overlapping product sets across periods, the fixed recommendation strategy can still guarantee at least 50% of the expected revenue.
- We extend our model to the setting when the utilities of purchased products decrease. This setting is more relevant for durable goods. We focus on two periods and characterize the expression of the conditional choice probability and the performance guarantee of the revenue-ordered policy. For the joint assortment optimization, we show the optimal fixed assortment may not be revenue-ordered and establish the performance of the revenue-ordered assortment.
- Finally, we conduct numerical experiments to study the performance of various assortment policies and quantify the benefits of considering repeated customer interaction with synthetic and real datasets. We generate synthetic data using a Mixed Multinomial Logit (MMNL) model and gradually increase the percentage of customers who have purchased the products in the past. Then we fit the data with our model and the standard MNL model. The result shows that taking into account the repeated interactions can lead to significant improvements of the predictive accuracy and expected revenue, even under model misspecification. For the real data, we calibrate our model and the standard MNL benchmark using real-world hotel booking data from Meituan, where historical bookings are fully observed. Our model achieves significantly better out-of-sample prediction. We further measure the revenue improvement by implementing an adaptive policy that takes into account the past interactions in the Sushi Preference data Kamishima (2003). The result shows an adaptive policy leads to 5.4% higher revenue on average.

2. Literature Review

Assortment optimization has been a central topic for discrete choice models and the literature is mostly focused on a single interaction between a customer and the seller. Here we review a few representative studies. For the MNL model, Talluri and Van Ryzin (2004) show that the optimal assortment is revenue-ordered. A number of studies extend the MNL framework and incorporate various capacity constraints (Rusmevichientong et al. 2010, Sumida et al. 2021, Désir et al. 2022). For the mixed MNL model, Rusmevichientong et al. (2014) show the assortment optimization problem is NP-hard and Bront et al. (2009), Sen et al. (2018) propose solutions based on heuristics or a conic quadratic mixed-integer formulation. Recently, Wang (2022) explores the impact of non-randomness under the MNL model and shows that the assortment optimization is NP-complete when the utility of the no-purchase option is deterministic. Overall, the major contribution of our study relative to this literature is the consideration of repeated interactions between the customer and the seller. Since the outcomes of past interactions can inform the seller about the preferences of the customer, the conditional assortment optimization problem in our framework can be viewed as a special case of personalized assortment optimization, which has been studied recently. Gallego and Berbeglia (2022) study the limit of personalization under various choice models. El Housni and Topaloglu (2023) study a two-stage personalized assortment optimization problem under the mixed MNL choice model. Bernstein and Guo (2023) consider the design of the optimal subscription box for a strategic customer who may not subscribe and choose to visit stores offline. They consider a single interaction between the seller and the customer while we consider multiple interactions. Jagabathula et al. (2022) consider a model where a customer may interact with the seller repeatedly. They use directed acyclic graphs (DAGs) to record the past interactions. DAGs represent the partial order over the products that is consistent with the choice (e.g., prefer product A to B and C in the offered assortment). Since each customer is endowed with an unobserved preference ranking of all the products, DAGs can be used to update the distribution of the rankings of a customer who makes repeated purchases. In Proposition 4, the authors approximate the choice probability conditional on the DAGs observed in the historical data. Based on the MNL model, we provide a closed-form expression for the conditional choice probability, which allows us to further explore the assortment optimization problem.

A number of recent studies extend the framework mentioned in the previous paragraph to multiple purchases from an individual customer. In Tulabandhula et al. (2023), the authors propose a multivariate logit model to allow customers to purchase at most K products, which is exogenously determined. They show the assortment optimization problem is NP-hard and propose several heuristics. Jasin et al. (2024) consider the purchase behavior among multiple categories and show that the optimal assortment consists of revenue-ordered assortments within each category under mild conditions. Lin et al. (2023) extend the rank-based choice model (Farias et al. 2013) to multiple purchases and connect it to the random utility theory. Bai et al. (2024) extend the MNL model to multiple purchases under the random utility maximization framework. Gallego and Wang (2019) assume the customer would purchase all products with utility larger than a product-specific threshold, which is determined by an average budget constraint. Luan et al. (2025) focus on the purchase of multiple units of a single product, assuming the marginal utility decreases quadratically. Zhang et al. (2021) extend the model proposed in Huh and Li (2022) to study the assortment optimization problem when the customer may purchase multiple units of multiple products. Compared to this literature, which typically introduces multiple purchases in a single interaction between the customer and the seller, we model multiple purchases sequentially during multiple interactions.

The stream of literature closest to this paper includes recent studies on assortment optimization over multiple periods. Flores et al. (2019) consider a two-stage MNL model, in which the customer purchases from two disjoint sets of products. The assortments to be offered are predetermined. They show the revenue-ordered assortments are still optimal. Liu et al. (2020) generalize the framework in Flores et al. (2019) to multiple periods. The authors prove the NP-hardness and propose a polynomial-time approximation scheme (PTAS). Feldman and Segev (2022) improve the efficiency of the algorithm in the model proposed in Liu et al. (2020). In these papers, the MNL choice probability in later periods is not updated based on the choice made earlier. Gao et al. (2021) extend the framework by leveraging the following observation: if a customer doesn't make any purchases in a period, the no-purchase option tends to have a larger utility under the conditional distribution. However, since the customer leaves the seller once a purchase is made, the conditional distribution only needs to be updated for the no-purchase option. In contrast, this study allows for multiple purchases over periods. As a result, the potential outcomes from past interactions are exponentially many. It significantly complicates the derivation of the conditional choice probability and the analysis. There are other studies that consider slightly different settings: Xu and Wang (2023) investigate a seller who announces the offered assortments and a customer making purchases by taking into account the expected utility in future periods. Ferreira and Goh (2021) consider assortment rotation over a selling horizon where the customer may purchase multiple products. They show the retailer can strategically earn more revenue by inducing multi-purchase behavior. Fata et al. (2019) study the multi-stage assortment problem with inventory constraints. The seller is able to repeatedly provide the same product through different assortments for a customer. However, they do not consider the updated preferences of customers. More recently, Tang et al. (2023) consider the multi-period assortment optimization problem by incorporating the popularity bias effect and explore the effectiveness of the revenue-ordered policy. Sumida and Zhou (2023) study multi-period assortment optimization with platform disengagement. Customers may interact with the platform repeatedly, but the purchase history is used to calibrate the engagement instead of updating the preference. Finally, Calmon et al. (2021) investigate a repeated revenue management setting where customer goodwill evolves based on past allocation quality. They demonstrate that myopic policies can achieve near-optimal performance under certain regularity conditions. To highlight the key difference between our model and the literature, consider a simple example of two periods and two products with $\mu_1 = \mu_2 = 0$. The standard MNL model states that the choice probability of product 1 in the assortment $S = \{1\}$ is $e^{\mu_1}/(1+e^{\mu_1}) = 1/2$. Suppose a customer visits in the first period and purchases product 1 from S. What is the choice probability of *the same customer* purchasing product 1 from the assortment $\{1,2\}$? Most models considering multiple purchases in the literature use the probability from a standard MNL model $e^{\mu_1}/(1+e^{\mu_1}+e^{\mu_2}) = 1/3$. In our model, conditional on the purchase of the first period, the (updated) choice probability is 2/3, because product 1 has been shown to be preferable by the customer, at least relative to the no-purchase option. This observation and the model built upon it differentiate this paper from the literature.

3. The Model

We consider a seller offering assortments over time to an individual customer. In each period, upon observing the offered assortment, the customer may purchase at most one product. In particular, in period m out of M periods, the offered assortment S_m is chosen from the set of products $\mathcal{N}_m$ available in that period. Note that the available products may or may not be the same across periods and we allow for $\mathcal{N}_i \cap \mathcal{N}_j \neq \emptyset$ for $i \neq j$. We denote the universe of products by $\mathcal{N}$ and thus $\mathcal{N}_k \subseteq \mathcal{N}$. The chosen product in period m is denoted by $c_m \in S_m \cup \{0\}$, where 0 represents a no-purchase. As in the classic MNL model, a customer has utility U_n for product $n \in \mathcal{N}$, where $U_n \sim \text{Gumbel}(\mu_n, 1)$ is an independent (across customers and products) Gumbel random variable with mean μ_n and scale 1. The utility of the outside option is similarly defined as $U_0 \sim \text{Gumbel}(0,1)$. We use $v_n = e^{\mu_n}$ to denote the preference weight of product n.

Next we characterize the choice probability of customers. Unconditionally (before interacting with the customer), the choice probability of a product $i \in S_m$ in period m follows the standard MNL model:

$$\frac{v_i}{1+\sum_{j\in S_m} v_j}.$$

What distinguishes our model from the classic MNL model is the conditional choice probability given the past interactions with the customer. In particular, when S_{m+1} is offered in period $m+1$, the seller observes the offered assortments in the previous periods[1] $S_{[m]} \triangleq \{S_1, \dots, S_m\}$ and the purchase history $\{c_1, \dots, c_m\}$. For subsequent analysis, we record the history before and including period m by three sets: the past periods with a purchase $\mathcal{A}_m \triangleq \{i \in [m] | c_i \neq 0\}$, the purchased products $\mathcal{C}_m \triangleq \{c_i | i \in [m], c_i \neq 0\}$, and past periods without purchases $\mathcal{B}_m \triangleq \{i \in [m] | c_i = 0\}$. Conditional on the purchase history prior to period $m+1$, the choice probability of choosing product $i \in S_{m+1}$ in period $m+1$ is $\mathbb{P}(U_i \geq \max\{U_0, \max_{j\in S_{m+1}} U_j\} | \mathcal{A}_m, \mathcal{C}_m, \mathcal{B}_m, \mathcal{S}_{[m]})$. The probability is taken with respect to the random utilities having Gumbel distributions.

To calculate the conditional choice probability, we need to analyze the joint probability of the purchase history $\mathbb{P}(\mathcal{A}_m, \mathcal{C}_m, \mathcal{B}_m, \mathcal{S}_{[m]})$. We include $\mathcal{S}_{[m]}$ in the event for generality because the seller may choose to offer assortments adaptively and randomly in each period depending on the purchase history. In the following example, we demonstrate the intractability of this probability in general.

EXAMPLE 1. Consider the offered assortment in the first period $S_1 = \{1, 2\}$ and $S_j = \{2j-2, 2j-1, 2j\}$ for the j-th period for $j = \{2, 3, \ldots, M\}$. In this case, the assortments offered in two consecutive periods have one common product $S_i \cap S_{i+1} = \{2i\}$. Suppose the customer chooses $c_i = 2i-1$ in period i. The joint probability of the history $\mathbb{P}(\mathcal{A}_m, \mathcal{C}_m, \mathcal{B}_m, \mathcal{S}_{[m]})$ up to period m can be expressed as:

$$\mathbb{P}(\mathcal{A}_m, \mathcal{C}_m, \mathcal{B}_m, \mathcal{S}_{[m]}) = \mathbb{P}(U_1 \ge \max\{U_2, U_0\}, U_{2i-1} \ge \max\{U_{2i-2}, U_{2i}, U_0\}\ \forall i = 2, \ldots, m). \tag{1}$$

Consider the events $\{U_{2i-1} \ge \max\{U_{2i-2}, U_{2i}, U_0\}\}$ for all i. Even conditional on the value of U_0, the events are interwoven and their intersection cannot be reduced to the CDF of individual U_is. One may condition on the values of $\{U_{2i}\}_{i=0}^m$:

$$\begin{aligned}(1) &= \int \mathbb{P}(U_1 \ge \max\{U_2, U_0\}, \ldots, U_{2m-1} \ge \max\{U_{2m-2}, U_{2m}, U_0\} | U_{2i} = u_{2i},\ i = 0, \ldots, m) df_{\mu_0}(u_0) \cdots df_{\mu_{2m}}(u_{2m}) \\ &= \int \mathbb{P}(U_1 \ge \max\{u_2, u_0\}) \cdot \prod_{i=2}^{m} \mathbb{P}(U_{2i-1} \ge \max\{u_{2i-2}, u_{2i}, u_0\}) df_{\mu_0}(u_0) \cdots df_{\mu_{2m}}(u_{2m}),\end{aligned}$$

where we have used $f_{\mu_i}(\cdot)$ to denote the probability measure of U_i. Because the maximum of $\{u_{2i-2}, u_{2i}, u_0\}$ leads to three possible cases, the probability $\mathbb{P}(U_{2i-1} \ge \max\{u_{2i-2}, u_{2i}, u_0\})$ has three expressions conditional on the maximum of $\{u_{2i-2}, u_{2i}, u_0\}$. As a result, the integrand branches to at least 2^m cases (holding U_0 as constant). Although an individual case can be computed analytically due to the benign properties of the Gumbel distribution, the joint probability is virtually computationally infeasible and does not lend to theoretical analysis. □

Because of this fundamental challenge, in the next section, we start with a two-period model. The seller interacts with the customer for two periods, motivated by the high churn rate in many online media platforms. For example, Hamilton and Singal (2023) analyze the user data from NetEase Cloud Music (Zhang et al. 2022) and find the churn rate to be around 50% for new users. It implies that the assortment and recommendation policy in the first two periods is critical in generating revenues. Then, in Section 5.1, we consider the case when the available products are disjoint across the periods for the same customer. The seller may rotate the products frequently to retain existing customers and attract new ones. For example, the fast fashion company SHEIN releases more than two thousand new clothes daily[2] while each item is produced in a small quantity[3]. In such a situation, it is reasonable to assume that the offered products are never repeated for the same customer. The two special cases cover the two ends of the spectrum of the general model: one on shallow interactions and one on deep interactions with many products. They allow our model and results to be applied to a range of practical problems. Finally, in Section 5.2, we develop a simple algorithm for the general case that achieves a constant ratio of expected revenues. We acknowledge that our model does not incorporate the cancellation of the subscription by customers and the number of periods is exogenously given. Although this is an important factor of subscription boxes and affects how assortments will be offered, we focus on repeated interactions and leave the endogenous cancellation to future research.

4. Assortment Planning for Two Periods

In this section, we consider customers interacting with the seller for two periods, i.e., $M=2$. In this case, the conditional/unconditional choice probability can be expressed in a closed form. In the first period, the customer purchases a product following the standard MNL model, and is given by:

$$\mathbb{P}(i,S_1)=\begin{cases}\frac{v_i}{1+\sum_{k\in S_1}v_k} & \text{if } i\in S_1,\\ \frac{1}{1+\sum_{k\in S_1}v_k} & \text{if } i=0.\end{cases}$$

Given the assortment and the choice in the first period, we consider the choice probability in the second period. Note that there are three cases based on the outcome in the first period: (1) $i=0$, (2) $i\in S_1\cap S_2$, and (3) $i\in S_1\setminus S_2$. Let $V(S)=\sum_{i\in S}v_i$. The next theorem gives the expressions for the conditional choice probability under the three cases.

THEOREM 1. *Given the assortments in the first two periods, S_1 and S_2, and the purchase decision in the first period $i\in S_1\cup\{0\}$, the conditional choice probability of purchasing product $j\in S_2\cup\{0\}$ in period two satisfies*

- *If $i=0$, then*

$$\mathbb{P}^1(j|i,S_1,S_2)=\begin{cases}\frac{1+V(S_1)}{1+V(S_1)+V(S_2\setminus S_1)} & j=0,\\ \frac{v_j}{1+V(S_1)+V(S_2\setminus S_1)} & j\in S_2\setminus S_1,\\ 0 & j\in S_1\cap S_2.\end{cases}$$

- *If $i\in S_1\cap S_2$, then*

$$\mathbb{P}^2(j|i,S_1,S_2)=\begin{cases}\frac{1+V(S_1)}{1+V(S_1)+V(S_2\setminus S_1)} & j=i,\\ \frac{v_j}{1+V(S_1)+V(S_2\setminus S_1)} & j\in S_2\setminus S_1,\\ 0 & j\in S_1\cap S_2,\ j\neq i.\end{cases}$$

- *If $i\in S_1\setminus S_2$, then*

$$\mathbb{P}^3(j|i,S_1,S_2)=\begin{cases}\frac{1+V(S_1)}{1+V(S_2)}\cdot\frac{v_j}{1+V(S_1)+V(S_2\setminus S_1)} & j\in S_1\cap S_2,\\ \frac{1+V(S_1)}{1+V(S_2)}\cdot\frac{1}{1+V(S_1)+V(S_2\setminus S_1)} & j=0,\\ v_j\cdot\left(\frac{1}{1+V(S_1)+V(S_2\setminus S_1)}+\frac{1+V(S_1)}{(1+V(S_2))\cdot(1+V(S_1)+V(S_2\setminus S_1))}\right) & j\in S_2\setminus S_1.\end{cases}$$

In Theorem 1, the expression of the choice probability is similar to the purchase probability in the MNL model when $i \in S_1 \cap S_2$ or $i = 0$. However, if the customer purchases a product only offered in the first period, the conditional choice probability is much more complicated, especially for product $j \in S_2 \setminus S_1$. We can also observe that in this case, given two products $j_1 \in S_1 \cap S_2, j_2 \in S_2 \setminus S_1$ with $v_{j_1} = v_{j_2}$, the customer has a higher choice probability for j_2 than j_1. This is because conditional on the event the customer prefers i over j_1 in the first period, the expected value of U_{j_1} has been distorted downward while the value of U_{j_2} is unaffected. Thus, the customer is more likely to choose j_2.

Given Theorem 1 as well as the choice probability in the first period following the standard MNL model, the seller aims to find S_1 and S_2 to offer to the customer in order to maximize the expected revenue over two periods. Note that the assortment S_2 in the second period may depend on S_1 and the purchased product i in the first period. We refer to such a policy as the *adaptive policy*. In particular, given S_1, the policy recommends an assortment $S_{2,i}$ in the second period when $i \in S_1 \cup \{0\}$ is chosen in period one. Let $R^k(S_2|i, S_1) = \sum_{j \in S_2} r_j \mathbb{P}^k(j|i, S_1, S_2)$ for $k \in \{1, 2, 3\}$ be the conditional expected revenue under assortment S_2 in period two, following the three cases in Theorem 1. The objective for the adaptive policy is as follows:

$$\mathscr{R}_A(S_1, \{S_{2,i}\}_{i \in S_1 \cup \{0\}}) \triangleq \frac{\sum_{i \in S_1} r_i v_i}{1 + V(S_1)} + \sum_{i \in S_1} \frac{v_i}{1 + V(S_1)} \cdot \max\{R^2(S_{2,i}|i, S_1), R^3(S_{2,i}|i, S_1)\} + \frac{R^1(S_{2,0}|0, S_1)}{1 + V(S_1)}. \quad (2)$$

In the objective, the first term is the expected revenue in the first period. The second term is the conditional expected revenue in the second period multiplied by the probability of purchasing i in the first period, in which the maximum is taken over cases two and three regarding whether the same item should be offered in period two. The third term is the expected revenue when no item is purchased in period one.

Note that in order to solve (2), the seller needs to solve the *conditional assortment optimization* problem in period two. That is, given the optimal assortment S_1^* in the first period, tractable expressions need to be derived for

$$\max_{S_{2,i} \subseteq \mathcal{N}_2} R^2(S_{2,i}|i, S_1^*), \ \max_{S_{2,i} \subseteq \mathcal{N}_2} R^3(S_{2,i}|i, S_1^*), \ \max_{S_{2,0} \subseteq \mathcal{N}_2} R^1(S_{2,0}|0, S_1^*).$$

This is the focus of Section 4.1.

We remark that the optimal solution to the conditional assortment planning problem highly depends on both i and S_1. For example, if $i = 0$, then all products in S_1 will not be chosen in the second period because they are even less attractive than the no-purchase option. Thus, the seller only needs to consider offering products in $\mathcal{N}_2 \setminus S_1$. On the other hand, if $i \in S_1 \cap \mathcal{N}_2$, then the seller may consider offering i again in period two, which makes it unnecessary to offer any products in $S_1 \setminus \{i\}$ as the customer has shown the preference for i over these products in period one.

4.1. Conditional Assortment Planning in Period Two

In this section, we characterize or provide algorithms to compute the optimal assortment in period two under different cases. Throughout the paper, we let $W(S) = \sum_{i\in S} r_i v_i$, where r_i is the revenue of product i, and $R(S) = W(S)/(1+V(S))$.

Case one: no purchase in period one. If the customer doesn't purchase anything in period one, then by Theorem 1 the conditional expected revenue in period two is

$$R^1(S_2|0,S_1) = \sum_{j\in S_2\setminus S_1} r_j \mathbb{P}^1(j|0,S_1,S_2) = \frac{W(S_2\setminus S_1)}{1+V(S_1)+V(S_2\setminus S_1)}. \tag{3}$$

Notice that (3) is the expected revenue for the standard MNL choice model with an inflated utility for the no-purchase option. That is, $1+V(S_1)$ instead of 1. Thus, the optimal assortment is a revenue-ordered assortment of $\mathcal{N}_2\setminus S_1$ according to Talluri and Van Ryzin (2004). The optimal assortment can be computed with time complexity $O(|\mathcal{N}_2|)$.

Case two: $i \in S_1 \cap \mathcal{N}_2$ purchased in period one and S_2 includes i. Next we derive the optimal assortment when the seller decides to include the purchased product in period one in the assortment offered in period two S_2. Note that this is not a complete case, as the optimal revenue in this case needs to be compared to the alternative when S_2 doesn't include i to obtain the ultimately optimal assortment. In this case, the expected revenue is

$$\begin{aligned} R^2(S_2|i,S_1) &= \sum_{j\in S_2\setminus S_1} r_j \mathbb{P}^2(j|i,S_1,S_2) + r_i \mathbb{P}^2(i|i,S_1,S_2) \\ &= \sum_{j\in S_2\setminus S_1} \frac{r_j v_j}{1+V(S_1)+V(S_2\setminus S_1)} + \frac{r_i(1+V(S_1))}{1+V(S_1)+V(S_2\setminus S_1)} = \frac{W(S_2\setminus S_1)+r_i(1+V(S_1))}{1+V(S_1)+V(S_2\setminus S_1)}. \end{aligned} \tag{4}$$

Since $i \in S_1 \cap S_2$, the seller needs to decide which products to offer among $\mathcal{N}_2\setminus S_1$ in addition to i. In the following lemma, we show that the optimal assortment can be efficiently computed.

Lemma 1. *The optimal assortment S_2 to maximize* (4) *is* $\{i\}$ *in addition to a revenue-ordered assortment in $\mathcal{N}_2\setminus S_1$.*

Due to Lemma 1, the optimal assortment in this case can be computed in linear time $O(|\mathcal{N}_2|)$.

Case three: product i purchased in period one not included in S_2. In this case, the expected revenue is

$$\begin{aligned} R^3(S_2|i,S_1) &= \sum_{j\in S_1\cap S_2} r_j \mathbb{P}^3(j|i,S_1,S_2) + \sum_{j\in S_2\setminus S_1} r_j \mathbb{P}^3(j|i,S_1,S_2) \\ &= \sum_{j\in S_1\cap S_2} \frac{1+V(S_1)}{1+V(S_2)} \cdot \frac{r_j v_j}{1+V(S_1)+V(S_2\setminus S_1)} \end{aligned}$$

$$
\begin{aligned}
&+\sum_{j\in S_2\setminus S_1} r_j v_j\cdot\left(\frac{1}{1+V(S_1)+V(S_2\setminus S_1)}+\frac{1+V(S_1)}{(1+V(S_2))\cdot(1+V(S_1)+V(S_2\setminus S_1))}\right)\\
=&\frac{1+V(S_1)}{1+V(S_2)}\cdot\frac{W(S_1\cap S_2)}{1+V(S_1)+V(S_2\setminus S_1)}\\
&+W(S_2\setminus S_1)\cdot\left(\frac{1}{1+V(S_1)+V(S_2\setminus S_1)}+\frac{1+V(S_1)}{(1+V(S_2))\cdot(1+V(S_1)+V(S_2\setminus S_1))}\right).
\end{aligned}\tag{5}
$$

From (5), we observe that the products in $S_2\setminus S_1$ and $S_2\cap S_1$ contribute to the expected revenue separately. The optimizations of the two subsets are likely to have independent structures. Indeed, we can show that the optimization of the assortment of the products in $\mathcal{N}_2\cap S_1$ is not challenging.

LEMMA 2. *For the optimal assortment S_2 of* (5)*, the displayed products in $S_2\cap S_1$ are a revenue-ordered assortment in $\mathcal{N}_2\cap S_1$.*

Lemma 2 suggests a strategy to solve (5): the seller may enumerate all the revenue-ordered assortments in $S_1\cap\mathcal{N}_2$, and for each optimize the remaining products in $\mathcal{N}_2\setminus S_1$. That is, we convert (5) to a two-stage optimization problem:

$$
\begin{aligned}
\max_{S\subseteq S_1\setminus\{i\}\cap\mathcal{N}_2, S\text{ is RO}}\ \max_{S_2\subseteq\mathcal{N}_2\setminus S_1}&\frac{1+V(S_1)}{1+V(S)+V(S_2)}\cdot\frac{W(S)}{1+V(S_1)+V(S_2)}+\\
&W(S_2)\cdot\left(\frac{1}{1+V(S_1)+V(S_2)}+\frac{1+V(S_1)}{(1+V(S)+V(S_2))\cdot(1+V(S_1)+V(S_2))}\right),
\end{aligned}\tag{6}
$$

where RO is short for revenue-ordered. We employ a slight abuse of notation in the previous expression, where we partition S_2 from (5) into $S_1\cap S_2$ and $S_2\setminus S_1$, denoted as S and S_2 in (6), respectively.

However, it turns out that the optimization regarding the products in $\mathcal{N}_2\setminus S_1$ is computationally challenging. Consider a special case where $S_1\cap\mathcal{N}_2=\emptyset$, i.e., *the products that can be offered in period two are all new.* In this case, we only need to focus on the inner optimization problem in (6). The expected revenue can be simplified as follows:

$$
R^3(S_2|i,S_1)=\frac{1+V(S_1)}{1+V(S_2)}\cdot\frac{W(S_2)}{1+V(S_1)+V(S_2)}+\frac{W(S_2)}{1+V(S_1)+V(S_2)}\tag{7}
$$

$$
=\frac{1+V(S_1)}{V(S_1)}\cdot W(S_2)\cdot\left(\frac{1}{1+V(S_2)}-\frac{1}{(1+V(S_1))\cdot(1+V(S_1)+V(S_2))}\right).\tag{8}
$$

To maximize the expected revenue, the seller needs to solve the optimization problem:

$$
\max_{S_2\subseteq\mathcal{N}_2} W(S_2)\cdot\left(\frac{1}{1+V(S_2)}-\frac{1}{(1+V(S_1))\cdot(1+V(S_1)+V(S_2))}\right).\tag{9}
$$

To solve (9) for case three and $S_1\cap\mathcal{N}_2=\emptyset$, we first establish a sufficient condition for the optimality of the revenue-ordered assortment.

LEMMA 3. *There exists a constant C such that when $V(S_1)\ge C$, the revenue-ordered assortment is optimal for Problem* (9)*.*

The result implies that a sufficiently large first-period preference weight $V(S_1)$ simplifies the problem and makes the revenue-ordered assortment optimal. On the other hand, the following example shows that when $V(S_1)$ is small, the optimal assortment may not be revenue-ordered. This stands in contrast to the classic result in the MNL model.

EXAMPLE 2. Case three and $S_1 \cap \mathcal{N}_2 = \emptyset$: suppose a customer purchases an item from S_1 with the total preference weights $V(S_1) = 0.03$. In the second period, three products are available with preference weights and revenues, respectively, $\boldsymbol{v} = (4, 100, 2)$ and $\boldsymbol{r} = (8.8, 8.4, 8.3)$. The optimal assortment to maximize the objective value in (9) is shown to be $\{1, 3\}$, which is not revenue-ordered. In Example 2, since the total preference weights in the first period are extremely small, the purchase event indicates the utility of the no-purchase option is quite small, leading to a higher conditional purchase probability. Specifically, the conditional probability of purchasing product one is 65.27% under assortment $\{1, 3\}$, significantly higher than 57.14% in the standard MNL model without the purchase history. This shifts the seller's focus to strategically managing product substitution in order to maximize the revenue. The assortment $\{1, 3\}$ is optimal because it avoids cannibalization from a highly popular, but less profitable product 2, thereby increasing the market share and revenue of the most profitable item.

The detailed comparison is quantified and shown in Table 1.

	Prod. 1	Prod. 2	Prod. 3	No-purchase	Revenue
$\{1\}$	95.90%			4.10%	8.439
$\{1,2\}$	3.84%	96.14%		0.02%	8.414
$\{1,3\}$	65.27%		32.63%	2.10%	8.452

Table 1 The conditional choice probability of each product under different assortments and the expected revenue in Example 2

Despite the sufficient condition in Lemma 3, in general problem (9) is hard to solve.

PROPOSITION 1. *The optimization problem* (9) *is NP-hard.*

Summarizing the three cases above, given the purchase decision in the first period, we can analyze the conditional assortment optimization problem in the second period. If the customer purchases nothing in the first period, then the problem is polynomially solvable. On the other hand, if the product $i \in S_1 \setminus \mathcal{N}_2$ is purchased in period one, then finding the optimal assortment in the second period is computationally challenging, especially in case three. Next, we provide a sufficient condition to rule out case three.

LEMMA 4. *If* $r_i \geq \max_{S_2 \subseteq \mathcal{N}_2} R(S_2)$ *for* $i \in S_1 \cap \mathcal{N}_2$ *and the customer purchases product* i *in the first period, then it is optimal to include it in the assortment of the second period.*

Lemma 4 shows that when the customer purchases a profitable product in the first period, the conditional assortment optimization problem in the second period is easy to solve because only case two needs to be considered. The seller offers product i in addition to a revenue-ordered assortment in $\mathcal{N}_2 \setminus S_1$.

Finally, we study the performance of revenue-ordered assortments for problem (6). More precisely, we consider the union of two revenue-ordered assortments: one consisting of products that have not been shown in period one (in $\mathcal{N}_2 \setminus S_1$) and another consisting of products that have been shown (in $S_1 \setminus \{i\} \cap \mathcal{N}_2$). Note that the union may not be revenue-ordered. Since the number of revenue-ordered assortments is $|\mathcal{N}_2|$, we can compute this policy in $O(|\mathcal{N}_2|^2)$ time.

Theorem 2. *For the optimization problem* (6)*, the revenue-ordered assortments can attain at least* $1/3$ *of the optimal revenue. Moreover, it can be improved to* $\frac{1+V(S_1)}{2+V(S_1)}$ *when* $S_1 \cap \mathcal{N}_2 = \emptyset$*.*

Theorem 2 shows the robust performance of revenue-ordered assortments. In the special case where no products in assortment S_1 can be provided repeatedly in the second period, we show that the approximation ratio is at least $\frac{1}{2}$. This policy is nearly optimal when the total preference weights in assortment S_1 are sufficiently large.

4.2. Joint Assortment Planning in Two Periods

Based on the negative result in Proposition 1, we conclude that finding the optimal adaptive policy for the two periods is intractable. To solve the joint assortment planning problem, we resort to the following approaches. First, we focus on two special cases: the available products (1) do not overlap (disjoint) or (2) are identical in the two periods. Second, we develop approximation algorithms, a common approach used in the literature. Third, we consider the fixed recommendation policy. That is, the seller chooses the offered assortments in the two periods in advance *regardless of* the outcome in the first period. The fixed policy can be solved efficiently: for the chosen assortments S_1 and S_2, the expected revenue can be expressed as

$$\mathscr{R}_F(S_1, S_2) \triangleq R(S_1) + R(S_2), \tag{10}$$

where $R(S)$ is the expected revenue of the MNL model for assortment S. The equality is due to the observation that the unconditional choice probability of each period under the fixed policy coincides with the MNL model. Subject to $S_1 \subset \mathcal{N}_1$ and $S_2 \subset \mathcal{N}_2$, they can be optimized separately and the resulting optimal assortment, S_1^{ro} and S_2^{ro}, are revenue-ordered. We analyze the performance of the fixed policy relative to the adaptive policy.

4.2.1. Special Case One: $\mathcal{N}_1 \cap \mathcal{N}_2 = \emptyset$ When the potential products to offer are disjoint in the two periods, the expression $R^3(S_2|i,S_1)$ is independent of the purchased product $i \in S_1$ according to (7). Thus the conditional expected revenue only depends on whether the customer makes a purchase or not. We further simplify the adaptive policy as $\{S_1, S_{2,0}, S_{2,1}\}$ and the joint expected revenue as

$$
\begin{aligned}
&\frac{W(S_1)}{1+V(S_1)} + \frac{V(S_1)}{1+V(S_1)} R^3(S_{2,1}|i,S_1) + \frac{1}{1+V(S_1)} R^1(S_{2,0}|0,S_1)\\
=&\frac{W(S_1)}{1+V(S_1)} + \frac{V(S_1)}{1+V(S_1)} \cdot W(S_{2,1}) \left(\frac{1}{1+V(S_1)+V(S_{2,1})} + \frac{1+V(S_1)}{(1+V(S_{2,1}))(1+V(S_1)+V(S_{2,1}))} \right)\\
&+ \frac{1}{1+V(S_1)} \frac{W(S_{2,0})}{1+V(S_1)+V(S_{2,0})}\\
=&\frac{W(S_1)}{1+V(S_1)} + W(S_{2,1}) \cdot \left(\frac{1}{1+V(S_{2,1})} - \frac{1}{(1+V(S_1)) \cdot (1+V(S_1)+V(S_{2,1}))} \right)\\
&+ \frac{1}{1+V(S_1)} \frac{W(S_{2,0})}{1+V(S_1)+V(S_{2,0})} =: \mathscr{R}_A(S_1, S_{2,0}, S_{2,1}).
\end{aligned} \tag{11}
$$

Then the optimization problem (2) reduces to the following problem

$$
\max_{S_1 \subseteq \mathcal{N}_1, S_{2,1}, S_{2,0} \subseteq \mathcal{N}_2} \mathscr{R}_A(S_1, S_{2,0}, S_{2,1}). \tag{12}
$$

Denote the optimal solution for the optimization problem (12) as $\{S_1^*, S_{2,0}^*, S_{2,1}^*\}$. Due to the difficulty of calculating $S_{2,1}^*$ as shown in Proposition 1, we design a fully polynomial-time approximation scheme (FPTAS) for the joint assortment optimization problem (12), the details are discussed in Appendix **??**.

PROPOSITION 2. *Let* $\mathcal{N} = \mathcal{N}_1 \cup \mathcal{N}_2$, $\gamma = \frac{\max_{i\in\mathcal{N}} r_i v_i}{\min_{j\in\mathcal{N}} r_j v_j}, \alpha = \frac{\max_{i\in\mathcal{N}} v_i}{\min_{j\in\mathcal{N}} v_j}$, *then we can design an FPTAS that returns a* $(1-\epsilon)$*-optimal solution to* (12)*. The running time is* $O\left(\log^3(|\mathcal{N}|\gamma)\log^3(|\mathcal{N}|\alpha)\frac{|\mathcal{N}|^3}{\epsilon^6}\right)$.

The performance of the fixed policy Compared with a fixed recommendation policy, (12) provides slight additional flexibility. The assortment in the second period is adaptively offered based on a binary outcome, whether or not there is a purchase in the first period. It turns out that the benefit of the additional flexibility is worth 25% of the revenue.

THEOREM 3. *When* $\mathcal{N}_1 \cap \mathcal{N}_2 = \emptyset$, *the fixed policy* $\{S_1^{ro}, S_2^{ro}\}$ *satisfies* $\mathscr{R}_F(S_1^{ro}, S_2^{ro}) \geq \frac{4}{5}\mathscr{R}_A(S_1^*, S_{2,0}^*, S_{2,1}^*)$.

Theorem 3 quantifies the revenue improvement by offering adaptive assortments in the second period. It also points to the limitation of personalization, as the extra revenue obtained is capped at 25%. Moreover, we show the bound is tight, as indicated by the following example.

EXAMPLE 3. The first period has one product, whose utility and revenue are $v_1 = 1$ and $r_1 = \epsilon$ for some $\epsilon \in (0, \frac{1}{2})$. The second period has two new products $\{2, 3\}$, whose utilities and revenues are $v_2 = \epsilon, r_2 = 1/\epsilon$ and $v_3 = 1/\epsilon, r_3 = 1-\epsilon$. The optimal adaptive assortments are shown to be $S_1^* = \{1\}, S_{2,1}^* = \{2\}, S_{2,0}^* = \{2,3\}$, and the expected revenue is

$$\mathscr{R}_A(S_1^*, S_{2,0}^*, S_{2,1}^*) = \frac{\epsilon}{2} + \frac{1}{1+\epsilon} - \frac{1}{2(2+\epsilon)} + \frac{1}{2}\cdot\frac{1+(1-\epsilon)/\epsilon}{2+\epsilon+1/\epsilon} \xrightarrow{\epsilon\to 0} \frac{5}{4}.$$

On the other hand, the optimal fixed policy is $S_1^{ro} = \{1\}$ and $S_2^{ro} = \{2\}$, whose expected revenue is $\epsilon/2 + 1/(1+\epsilon) \xrightarrow{\epsilon\to 0} 1$.

4.2.2. Special Case Two: $\mathcal{N}_1 \cap \mathcal{N}_2 = \mathcal{N}$ We next study a common setting when the seller has the same set of products to provide in both periods, that is $\mathcal{N}_1 = \mathcal{N}_2 = \mathcal{N}$. We introduce some necessary notation to streamline the analysis. Given the purchased item i in the first period, we denote the optimal assortments in the second period for the three cases, explained in Section 4.1, as $S_{2,0}^*, S_{2,2,i}^*, S_{2,3,i}^*$. We drop the dependence on S_1 for simplicity.

The optimal adaptive policy At first glance, the joint assortment problem cannot be simplified in this setting. Indeed, for each purchased decision i in the first period, we need to compare the conditional expected revenue from $S_{2,2,i}^*$ and $S_{2,3,i}^*$, which has been shown to be NP-hard. However, we can further strengthen the result in Lemma 4 and identify the condition for $S_{2,2,i}^*$ to be optimal. Let $S^{ro} \triangleq \arg\max_{S\subseteq\mathcal{N}} R(S)$, we first introduce the following classic result from Talluri and Van Ryzin (2004), Rusmevichientong and Topaloglu (2012).

Lemma 5 **(Theorem 3.2 in Rusmevichientong and Topaloglu (2012))**. *Product $i \in S^{ro}$ if and only if $r_i > R(S^{ro})$.*

Combined with Lemma 4, we can see that if the product purchased by the customer in the first period, product i, is in S^{ro}, then $r_i \geq R(S^{ro}) = \max_{S_2\subseteq\mathcal{N}_2} R(S_2)$. By Lemma 4, it is optimal for the seller to include product i in the assortment of the second period. As a result, we always encounter case two described in Section 4.1 and $S_{2,2,i}^*$ can be computed efficiently.

To further explore the structure of $S_{2,2,i}^*$ for $i \in S^{ro}$, note that it is not optimal to include products that are cheaper than r_i, because the seller would make more profit by only offering product i in the second period, which results in a purchase and a revenue of r_i. By Lemma 1, the revenue-ordered assortment in addition to product i in the second period only includes products that sell at a higher price than r_i. Summarizing the observations, we have

Corollary 1. *If the product purchased in the first period $i \in S^{ro}$, then $i \in S_{2,2,i}^* \subseteq S^{ro}$ and $r_i = \min\{r_j : j \in S_{2,2,i}^*\}$. In particular, if S_1 is also revenue-ordered, then $S_{2,2,i}^* = \{i\}$.*

In summary, according to Lemma 4 and Lemma 5, if $i \in S^{ro} \cap S_1$ (recall that S^{ro} is the optimal assortment for the standard MNL model with product set $\mathcal{N}$), then it is optimal to repeatedly recommend product i in the second period. It also implies that $R^2(S_{2,2,i}^*|i, S_1) \geq R^3(S_{2,3,i}^*|i, S_1)$ and

$S^*_{2,2,i}$ can be found in polynomial time. Next we show that this condition that greatly simplifies the optimization, $i \in S^{ro}$, always holds, because it is optimal to offer a revenue-ordered assortment from S^{ro} in the first period. As a result, any product purchased in the first period must belong to S^{ro}.

Proposition 3. *When $\mathcal{N}_1 = \mathcal{N}_2 = \mathcal{N}$, the optimal adaptive policy $\{S^*_1, \{S^*_{2,i}\}_{i \in S^*_1 \cup \{0\}}\}$ can be described as follows:*

1. *$S^*_1 \subseteq S^{ro}$ and it is revenue-ordered.*
2. *$S^*_{2,i} = \{i\}$ for $i \in S^*_1$.*
3. *$S^*_{2,0} \subseteq \mathcal{N} \setminus S^*_1$ and it is revenue-ordered.*

In Proposition 3, the optimal adaptive policy is simple to implement: (1) offer a revenue-ordered assortment in the first period, (2) if the customer makes a purchase, repeatedly provide the same product in the second period, (3) otherwise, display another revenue-ordered assortment. Moreover, finding the optimal adaptive policy only requires solving the following optimization problem. Since $S^*_1, S^*_{2,0}$ are revenue-ordered, it can be computed with time complexity $O(|\mathcal{N}|^2)$.

$$\max_{S_1, S_{2,0} \text{ are RO }, S_{2,0} \subseteq \mathcal{N} \setminus S_1} \frac{2W(S_1)}{1+V(S_1)} + \frac{1}{1+V(S_1)} \cdot \frac{W(S_{2,0})}{1+V(S_1)+V(S_{2,0})}, \tag{13}$$

where RO is short for revenue-ordered.

We remark that the optimal adaptive policy in the second period is not unique. Given the purchased product i in the first period, the conditional probability of purchasing product $j \in S^*_1 \setminus \{i\}$ is 0 in the second period. Thus, the seller may provide these products in $S^*_{2,i}$ without changing the purchasing probability or expected revenues. In practice, offering these products can be beneficial due to, for example, variety seeking behavior that is not explicitly captured in our model.

The performance of the fixed policy When $\mathcal{N}_1 \cap \mathcal{N}_2 = \emptyset$, we have shown that the seller can earn as much as 25% more revenue by recommending assortments adaptively based on whether the customer makes a purchase in the first period. However, when $\mathcal{N}_1 = \mathcal{N}_2$, the conditional assortment in the second period is contingent on more than two outcomes. As a result, the flexibility makes the advantage of adaptive recommendation more salient. The next result shows a weaker performance of the fixed recommendation policy. Note that the optimal assortment is S^{ro} in both periods for the fixed policy.

Theorem 4. *When $\mathcal{N}_1 = \mathcal{N}_2 = \mathcal{N}$, the optimal fixed policy $\{S^{ro}, S^{ro}\}$ satisfies $\mathscr{R}_F(S^{ro}, S^{ro}) \geq \frac{2}{3}\mathscr{R}_A(\{S^*_1, \{S^*_{2,i}\}_{i \in S^*_1 \cup \{0\}}\})$.*

In Theorem 4, we show that the approximation ratio for the optimal fixed policy decreases from 4/5 to 2/3. The ratio turns out to be tight, as demonstrated by the following example.

EXAMPLE 4. Considering two products in $\mathcal{N}$, whose utilities and revenues are $v_1 = \epsilon, r_1 = 1 + 1/\epsilon, v_2 = 1/\epsilon, r_2 = 1$. The optimal revenue-ordered assortment is $S^{ro} = \{1\}$ with expected revenue $R(S^{ro}) = 1$. Therefore, the expected revenue of offering S^{ro} in both periods is 2. The optimal adaptive recommendation policy is $S_1^* = \{1\}$. For the second period, if the customer purchases the product, then offer $\{1\}$ again in the second period; otherwise, offer $\{2\}$. The expected revenue is

$$\frac{1+\epsilon}{1+\epsilon} + \frac{\epsilon}{1+\epsilon}\cdot(1+1/\epsilon) + \frac{1}{1+\epsilon}\cdot\frac{1/\epsilon}{1+\epsilon+1/\epsilon} = 2 + \frac{1}{(1+\epsilon)(1+\epsilon+\epsilon^2)} \xrightarrow{\epsilon\to 0} 3,$$

which leads to a performance ratio of 2/3.

5. Joint Assortment Optimization in Multiple Periods

In this section, we extend the results to multiple periods. As one may expect, the technical challenges in Section 4 for two periods are only amplified for multiple periods. We first consider the case when the products are disjoint before the general case.

5.1. Disjoint Product Sets

In this section, we assume the sets of products offered in different periods are disjoint: $\mathcal{N}_i \cap \mathcal{N}_j = \emptyset$ for all $i \neq j$. Given the historical purchases in the first m periods, we first derive the probability for the event $(\mathcal{A}_m, \mathcal{C}_m, \mathcal{B}_m, \mathcal{S}_{[m]})$, which is defined in Section 3.

PROPOSITION 4. *Suppose in the first m periods, the offered assortments are $\mathcal{S}_{[m]} = \{S_1, \dots, S_m\}$, the customer purchases products $\mathcal{C}_m$ in periods $\mathcal{A}_m$ and makes no purchases in periods $\mathcal{B}_m$. The probability of the event is*

$$\Pr(\mathcal{A}_m, \mathcal{C}_m, \mathcal{B}_m, \mathcal{S}_{[m]}) = \frac{\prod_{p\in\mathcal{A}_m} v_{c_p}}{\prod_{p\in\mathcal{A}_m} V(S_p)} \sum_{T\subseteq\mathcal{A}_m} (-1)^{|T|} \frac{1}{1 + V(\mathcal{B}_m) + \sum_{q\in T} V(S_q)},$$

where $V(\mathcal{B}_m) = \sum_{j\in\mathcal{B}_m} V(S_j)$.

Proposition 4 allows us to characterize the conditional choice behavior of period $m+1$ when S_{m+1} is offered. In particular, the purchase probability of product $j \in S_{m+1}$ is

$$\mathbb{P}(U_j \geq \max\{U_0, \max_{l\in S_{m+1}\setminus\{j\}} U_l\} | \mathcal{A}_m, \mathcal{C}_m, \mathcal{B}_m, \mathcal{S}_{[m]}) = \frac{\mathbb{P}(\mathcal{A}_m \cup \{m+1\}, \mathcal{C}_m \cup \{j\}, \mathcal{B}_m, \mathcal{S}_{[m]} \cup \{S_{m+1}\})}{\mathbb{P}(\mathcal{A}_m, \mathcal{C}_m, \mathcal{B}_m, \mathcal{S}_{[m]})}. \tag{14}$$

Substituting Proposition 4 into the above expression, we observe that the conditional choice probability is independent of the purchased products in the history, as long as the periods $\mathcal{A}_m$ and the set of assortments $\mathcal{S}_{[m]}$ are given. This is because the term $\prod_{i\in\mathcal{A}_m} v_{c_i}$ cancels out in (14). We also observe that $V(S_i)$, rather than the full trajectory $(\mathcal{A}_m, \mathcal{C}_m, \mathcal{B}_m, \mathcal{S}_{[m]})$, plays an important role in the probability.

For joint assortment optimization, the seller needs to find a sequence of assortments adaptively to maximize the total expected revenue. Recall that given the purchase behavior in the first m

periods $(\mathcal{A}_m, \mathcal{C}_m, \mathcal{B}_m, \mathcal{S}_{[m]})$, from the joint probability (4) we can conclude that the conditional choice probability in (14) doesn't depend on the specific products purchased in the past. It only depends on whether a product is purchased in the previous period. Thus, in the $(m+1)$-th period, the optimal assortment is contingent on 2^m possible outcomes. For simplicity, we can express the purchase behavior in the previous m periods as a binary vector. For example, $(1,0,1,0,\ldots,0)$ represents the trajectory that the customer purchases a product in periods 1 and 3 and nothing in other periods. We further encode the optimal assortment in period $m+1$ as $\mathcal{W}_{m+1} = \{S_0(m+1), \ldots, S_{2^m-1}(m+1)\}$, where, with a little abuse of notation, the subscript represents one of the 2^m outcomes. The joint expected revenue over M periods can be expressed as follows:

$$\mathscr{R}_A(\mathcal{W}_1, \mathcal{W}_2, \ldots, \mathcal{W}_M) = R(S_0(1)) + \sum_{m=2}^{M} \sum_{k=0}^{2^{m-1}-1} P_m^k R_m^k(S_k(m)). \tag{15}$$

where P_m^k denotes the unconditional probability that a customer in period m has purchase behavior k in the past and $R_m^k(S_k(m))$ the corresponding conditional expected revenue when offering $S_k(m)$.

We denote the optimal value in (15) as $\mathscr{R}_A^*$. The number of optimal assortments involved in (15) is $2^M - 1$, which is exponential in the number of periods M. Besides, solving the conditional assortment optimization in a single period is generally NP-hard. Thus, we focus on a more practical objective: characterizing the performance of the optimal fixed recommendation policy relative to $\mathscr{R}_A^*$. Under the optimal fixed policy, the total expected revenue over M periods is

$$\mathscr{R}_F(S_1^{ro}, S_2^{ro}, \ldots, S_M^{ro}) = \sum_{m=1}^{M} R(S_m^{ro}), \tag{16}$$

where $S_m^{ro} = \arg\max_{S \subseteq \mathcal{N}_m} R(S)$ for $m \in [M]$. We denote the optimal expected revenue as $\mathscr{R}_F^*$. The following theorem shows the fixed policy can guarantee a constant ratio of the optimal revenue.

Theorem 5. *When the available products in different periods are disjoint, we have*

$$\mathscr{R}_A^* \leq 1.4603 \cdot \mathscr{R}_F^*. \tag{17}$$

Moreover, there exists an instance such that $\mathscr{R}_A^* \geq \left(1 + \frac{1}{e}\right) \cdot \mathscr{R}_F^*$.

The result shows the benefit of adaptive assortments or personalization is limited. In Theorem 3, we show the extra expected revenue obtained from the optimal adaptive policy is at most 25% in the two-period problem with disjoint sets. When the seller collects more information about the customer's preference for the no-purchase option, the further improvement in the expected revenue is upper bounded by 46.03%. The reason behind this result is that the purchase history does not reflect the preference of the customer for the unseen products in the next period, and the limited information restricts the performance of the adaptive policy. Moreover, even though the upper bound is not necessarily tight, the instance we construct illustrates that the relative gap is bounded by $(1.4603-1-1/e)/(1+1/e)\approx 6.76\%$.

When all products are relatively popular compared to the no-purchase option, we show the revenue loss for the optimal fixed policy can be further bounded. Formally, we have the following proposition.

PROPOSITION 5. *Assume $v_j \geq \frac{1}{s-1}$ for all $j\in\mathcal{N}$, we have $\mathscr{R}_A^* \leq s\cdot\mathscr{R}_F^*$.*

By Proposition 5, if the utility of all products are larger than 10, for example, then the extra revenue obtained by the optimal adaptive policy is at most 10% relative to the revenue of the optimal fixed policy. On the other hand, when the outside option is considerable, the customers are likely to purchase nothing, and it becomes crucial for the seller to infer U_0 from customer purchase behavior. In this case, a personalized recommendation policy can help the seller to garner higher revenues. This implication is corroborated by the next numerical example.

EXAMPLE 5. Consider six products with $\mathcal{N}_1=\{1,2,3\}$ and $\mathcal{N}_2=\{4,5,6\}$. The utility and revenue are given by $\mathbf{v}=\{0.24,1.87,1,8.67,2.1,9.24\}$ and $\boldsymbol{r}=\{1.33,1.29,1.05,1.98,1.31,1.05\}$. We fix the utility of the no-purchase option to 1. Given $\alpha\in[1,4]$, we scale the utility of all products from $\mathbf{v}$ to $\mathbf{v}/\alpha$ and calculate the ratio between the optimal adaptive policy and the fixed policy. The relationship between the ratio and the value of α is shown in Figure 2. With a larger α, the adaptive policy obtains a much higher expected revenue.

However, we point out that this insight is obtained in the scenario of disjoint product sets, in which case learning the outside option is a crucial factor. When the seller may offer the same products over time, learning the preference weights of these products becomes more important, especially when they are expensive and valued differently by customers. In these cases, the adaptive policy may provide unique benefits. Sellers typically invest more effort in maintaining better customer relationships and understanding their preferences. We acknowledge that due to intractability, the performance gap between the optimal adaptive and fixed policies is left for future research.

5.2. Constant Bound for the General Problem

In the general setting, the available products may not be disjoint across periods. The seller is able to provide some products repeatedly to collect more information. Due to the complexity of the optimal adaptive policy, as we have discussed in the disjoint setting, we focus on quantifying the performance of the optimal fixed policy. Similarly, we still denote the optimal expected revenue from the optimal adaptive policy as $\mathscr{R}_A^*$.

Instead of directly analyzing the value of $\mathscr{R}_A^*$, we aim at bounding the expected revenue obtained by an adaptive policy in each period separately. More precisely, in period $m+1$, without loss of generality, we index the available products as $\{1,2,\ldots,n_{m+1}\}$ and assume $r_1\geq r_2\geq\ldots\geq r_{n_{m+1}}$. We group the customers into n_{m+1} types. For the customers of the j-th type, product j is the first

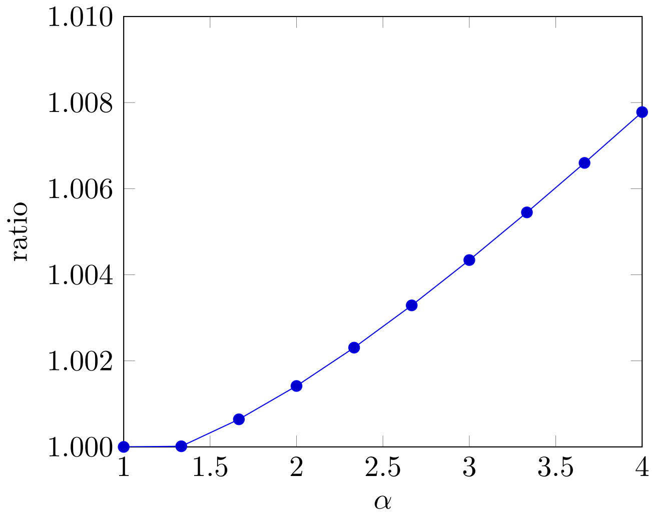

Figure 2 **The revenue ratio of the adaptive policy to the fixed policy.**

product that is preferable to the no-purchase option. That is, the utilities of the products satisfy $\max_{i\le j-1} U_i \le U_0$ and $U_j > U_0$. When the seller recommends an assortment, the maximum revenue obtained from type-j customers is r_j, since the customer would only consider products with utilities higher than U_0. Therefore, an upper bound for the revenue earned by the seller in period $m+1$ is $\sum_{j=1}^{n_{m+1}} r_j \mathbb{P}(\text{type } j)$, in which case the seller is *clairvoyant* and can observe the customer's type exactly. Note that the type of a customer only depends on the distribution of the Gumbel random vector $(U_0, U_1, \ldots, U_{n_{m+1}})$.

On the other hand, given the purchase history $(\mathcal{A}_m, \mathcal{C}_m, \mathcal{B}_m, \mathcal{S}_{[m]})$ in the first m periods, the seller may or may not be able to identify the type of the customer exactly. For example, consider $m=1$ and the seller has three products, $\{1,2,3\}$, available in the second period. In the first period, suppose only product 3 is included in the assortment. In this case, given the purchase history of the first period, the seller cannot possibly distinguish a type-1 customer from a type-2 customer. As a result, the expected revenue will always be less than a clairvoyant seller.

The bound we construct is based on the following observation: the clairvoyant seller can earn a revenue $\sum_{j=1}^{n_{m+1}} r_j \mathbb{P}(\text{type } j)$ in period $m+1$, which is an upper bound for the revenue of the optimal adaptive policy. If we can show that $\sum_{j=1}^{n_{m+1}} r_j \mathbb{P}(\text{type } j) \le 2R_{m+1}^{ro}$, twice the optimal revenue under the MNL model (or the optimal fixed policy according to (16)) in period $m+1$, then one can conclude that in period $m+1$, the optimal fixed policy at least achieves 50% of the revenue of the adaptive policy. This observation allows us to break the problem into a single stage, because both $\mathbb{P}(\text{type } j)$ and $R(S_{m+1}^{ro})$ only depend on the Gumbel random variables $(U_0, \ldots, U_{n_{m+1}})$, not the historical purchases. By analyzing the properties of the Gumbel distribution, we obtain the following performance guarantee for the optimal fixed policy.

Theorem 6. *In the general multi-period setting, we have*

$$\mathscr{R}_A^* \le 2 \cdot \mathscr{R}_F^*. \tag{18}$$

In Theorem 6, we show the optimal fixed policy can maintain at least half of the optimal expected revenue compared with the adaptive policy. Moreover, compared with the optimal adaptive policy which involves at most $O(|\mathcal{N}|^M)$ assortments, the optimal fixed policy can be computed in $O(|\mathcal{N}|M)$. Thus the optimal fixed policy is more computationally efficient and the theoretical result also validates the robustness of the fixed policy.

6. Extension: Discounted Utility After Purchase

In our model, a crucial assumption is that the preference weight for each product remains constant across multiple periods. This assumption may hold for non-durable goods. For durable goods, however, the utility of a second product that has been purchased may be diminishing. In this section, we consider the discounted utility and study the joint assortment optimization problem in two periods. Specifically, the utility of product i will be reduced from U_i to $U_i - d$ once it is purchased, where $d > 0$ is a given constant. In this setting, the conditional choice probability has a different expression when the seller decides to repeatedly offer the purchased product. We present the choice probability in the following lemma.

Lemma 6. *Suppose the customer purchases product $i \in S_1 \cap S_2$ in the first period. The conditional choice probability of product $j \in S_2 \cup \{0\}$ in the second period satisfies*

$$\mathbb{P}_d^2(j|i,S_1,S_2) = \begin{cases} \frac{e^{-d}(1+V(S_1))}{1+V(S_2\setminus\{i\})+e^{-d}(v_i+V(S_1\setminus S_2))} & j=i, \\ \frac{(1-e^{-d})(1+V(S_1\cap S_2)-v_i)+1+V(S_2)+V(S_1\setminus S_2)}{[1+V(S_2\setminus\{i\})+e^{-d}(v_i+V(S_1\setminus S_2))]\cdot[1+V(S_2)+V(S_1\setminus S_2)]} \cdot v_j & j \in S_2 \setminus S_1, \\ \frac{(1-e^{-d})(1+V(S_1))}{[1+V(S_2\setminus\{i\})+e^{-d}(v_i+V(S_1\setminus S_2))]\cdot[1+V(S_2)+V(S_1\setminus S_2)]} \cdot v_j & j \in S_1\cap S_2 \cup \{0\},\ j \ne i. \end{cases} \tag{19}$$

From Lemma 6, we observe that the probability of repeatedly purchasing product i decays exponentially in d. Moreover, the customer may reconsider the products in the subset $S_1 \cap S_2 \setminus \{i\}$, which is not the case in Theorem 1. However, the conditional probability becomes much more complicated. When the discount d goes to infinity, i.e., the customer never purchases a second copy of a product, we can simplify the conditional probability:

$$\mathbb{P}_d^2(j|i,S_1,S_2) = \begin{cases} 0 & j=i, \\ \frac{2+V(S_1)+V(S_2\setminus\{i\})}{[1+V(S_2\setminus\{i\})]\cdot[1+V(S_2)+V(S_1\setminus S_2)]} \cdot v_j & j \in S_2 \setminus S_1, \\ \frac{1+V(S_1)}{[1+V(S_2\setminus\{i\})]\cdot[1+V(S_2)+V(S_1\setminus S_2)]} \cdot v_j & j \in S_1\cap S_2 \cup \{0\},\ j \ne i. \end{cases} \tag{20}$$

The conditional choice probability in (20) is equivalent to the conditional probability in Theorem 1 after changing the preference weight of product i in the second period to 0.

In the discounted utility setting, when product i is excluded from S_2, the problem reduces to Problem (6), which is addressed in Lemma 2 and Theorem 2. Therefore, we focus on the case where product i is included in S_2. The conditional assortment optimization problem is formulated as follows:

$$\max_{S_2 \subseteq \mathcal{N}_2 : i \in S_2} \sum_{j \in S_2} r_j \mathbb{P}_d^2(j|i, S_1, S_2).$$

By investigating the structure of the revenue function, we further extend the results in Lemma 2 and Theorem 2 to the discounted setting.

THEOREM 7. *The optimal conditional assortment S_2^* in the discounted setting satisfies the following property: $S_2^* \cap S_1$ is a revenue-ordered assortment in $\mathcal{N}_2 \cap S_1 \setminus \{i\}$. Moreover, an assortment constructed by the union of a revenue-ordered subset of $\mathcal{N}_2 \cap S_1 \setminus \{i\}$ and a revenue-ordered subset of $\mathcal{N}_2 \setminus S_1$ guarantees an approximation ratio of $\frac{1}{3-e^{-d}}$ compared to S_2^*.*

Notice when $d = +\infty$, we recover the performance guarantee in Theorem 2. The result shows the revenue-ordered policy is quite robust for the conditional assortment optimization problem.

Because of the complexity, we focus on the setting $\mathcal{N}_1 = \mathcal{N}_2 = \mathcal{N}$ for tractability and study the performance of the fixed policy. Since the preference weight of the purchased product decreases in the second period, the total expected revenue under the fixed policy (S_1, S_2) can be expressed as follows:

$$\begin{aligned} \mathscr{R}_F(S_1, S_2) \triangleq & R(S_1) + \sum_{i \in S_1 \cap S_2} \frac{v_i}{1 + V(S_1)} \cdot \sum_{j \in S_2} \mathbb{P}_d^2(j|i, S_1, S_2) r_j \\ & + \sum_{i \in S_1 \setminus S_2} \frac{v_i}{1 + V(S_1)} \cdot \sum_{j \in S_2} \mathbb{P}^3(j|i, S_1, S_2) r_j + \frac{R^1(S_2|0, S_1)}{1 + V(S_1)}. \end{aligned} \tag{21}$$

In contrast to the base model, in (21), the total expected revenue in the second period depends on the assortment in the first period. Thus, we cannot optimize the assortments in the two periods separately. The following example shows that the assortment in the first period under the optimal fixed policy may not be revenue-ordered. Moreover, even if we restrict our focus to fixed policies that offer the same assortment in both periods, the revenue-ordered assortment may still be sub-optimal.

EXAMPLE 6. Consider four products with the corresponding preference weights and revenues $\boldsymbol{v} = (3.32, 2.97, 1.12, 8.63)$, $\boldsymbol{r} = (5.93, 4.82, 4.07, 4.06)$. When $d \to \infty$, the optimal fixed policy is to provide $\{2, 3, 4\}$ in the first period and $\{1, 2\}$ in the second period. If the same assortment has to be offered in both periods, then the optimal assortment is $\{1, 2, 4\}$.

When the same assortment has to be offered in both periods and $d \to \infty$, Bai et al. (2024) investigate the optimal assortment problem. In particular, for the assortments $S_1 = S_2 = S$ in the two periods, the total expected revenue is:

$$\mathscr{R}_F(S, S) \triangleq \sum_{i \in S} \left[\sum_{j \in S \setminus \{i\}} \frac{(1 - e^{-d}) v_j}{(1 + V(S) - (1 - e^{-d}) v_j)(1 + V(S))} + \frac{e^{-d}}{1 + V(S) - (1 - e^{-d}) v_i} + \frac{1}{1 + V(S)} \right] \cdot v_i r_i. \tag{22}$$

Furthermore, when $d=\infty$, we can further simplify the expression:

$$\mathscr{R}_F(S,S) \triangleq \sum_{i\in S}\left[\sum_{j\in S\setminus\{i\}} \frac{v_j}{(1+V(S)-v_j)(1+V(S))} + \frac{1}{1+V(S)}\right]\cdot v_i r_i. \tag{23}$$

The revenue-ordered assortment is shown to achieve an approximation ratio 1/2 for the optimization problem (23) in Bai et al. (2024), although the hardness is left as an open question. Note that the result in Bai et al. (2024) is regarding the gap between the optimal revenue-ordered policy and the optimal fixed policy, both with the same assortment in two periods. Next we prove that the same revenue-ordered assortment can guarantee a constant approximation ratio compared with the optimal adaptive policy even when d is finite.

Theorem 8. *When $\mathcal{N}_1=\mathcal{N}_2=\mathcal{N}$, there exists a revenue-ordered assortment S such that*

$$\frac{1}{2+e^{-d}}\cdot\mathscr{R}_A^* \leq \mathscr{R}_F(S,S).$$

Moreover, the approximation ratio can be improved to $\frac{1}{2}$ if one of the following conditions is satisfied:

1. $v_i \geq 1 \quad \forall i\in\mathcal{N}$;
2. $e^{-d} \leq \min_{i\in\mathcal{N}} \frac{v_i}{1+v_i-v_i^2}$.

In Theorem 8, we show the fixed policy that offers the same revenue-ordered assortment in the two periods guarantees a constant approximation ratio compared with the optimal adaptive policy. When the products are popular or the discount parameter d is sufficiently large, we can improve the approximation ratio to 1/2. Thus, we strengthen the results in Bai et al. (2024) by establishing the same approximation ratio 1/2 relative to a stronger (adaptive) policy. It demonstrates the power of revenue-ordered assortments even in the adaptive setting.

7. Numerical Experiment

In this section, we first design synthetic experiments to illustrate the benefits of taking into account the repeated interactions when estimating the model. It allows us to compare to the ground-truth model that generates the data. We then consider two real datasets to demonstrate the benefits compared to the standard MNL model.

7.1. Synthetic Experiments

To assess the robustness of our approach to model misspecification, we generate the ground-truth data using a Mixed Multinomial Logit (MMNL) model with k customer segments. For each segment $g\in\{1,\dots,k\}$, we randomly sample the preference weights $\boldsymbol{v}^g$ uniformly from $[0,5]$. The market share of each segment, denoted by α_g, is generated by normalizing k independent draws from a uniform distribution $U(0,1)$. Note that when $k=1$, the ground-truth model reduces to the standard MNL model without misspecification.

Given the ground-truth model, we simulate transaction data for a set of customers $\mathcal{D}$. For some customers $\mathcal{D}_1 \subseteq \mathcal{D}$, we record the choice behavior for three consecutive periods. For the rest $\mathcal{D}_2 = \mathcal{D} \setminus \mathcal{D}_1$, we only record the transaction of a single period. The transaction data can be represented as $\mathcal{T} = \{(S_t, i_t) : t \in \mathcal{D}_2, (S_t^1, i_t^1, S_t^2, i_t^2, S_t^3, i_t^3) : t \in \mathcal{D}_1\}$, where S_t is the assortment offered for the t-th customer of a single transaction and $i_t \in S_t \cup \{0\}$ is the purchase decision (0 represents the no-purchase option). Similarly, (S_t^1, S_t^2, S_t^3) are the assortments displayed to customer $t \in \mathcal{D}_1$ and (i_t^1, i_t^2, i_t^3) are the purchase decision. We denote $\rho = \frac{3|\mathcal{D}_1|}{3|\mathcal{D}_1|+|\mathcal{D}_2|}$ as the percentage of the transaction data from repeated customers.

Given the transaction data, we use three different methods to estimate the parameters μ_i.

1. MNL: Treat $(S_t^1, i_t^1, S_t^2, i_t^2, S_t^3, i_t^3)$ as three independent transactions $\{(S_t^1, i_t^1), (S_t^2, i_t^2), (S_t^3, i_t^3)\}$ and fit the standard MNL model without repeated customers.
2. RCI-2: Treat $(S_t^1, i_t^1, S_t^2, i_t^2, S_t^3, i_t^3)$ as two independent customers $\{(S_t^1, i_t^1, S_t^2, i_t^2), (S_t^3, i_t^3)\}$. In this case, we partially use the historical purchases of a customer.
3. RCI-3: We use all the historical information and fit our model.

In summary, the three methods leverage the repeated interactions progressively. The first method completely ignores such information; the second method uses it partially; and the last method considers all the past interactions. The three methods fit the same set of parameters $\boldsymbol{v}$ using the maximum-likelihood estimation (MLE). We present the details in Appendix **??**. We denote the fitted parameters as $\boldsymbol{v}^{\mathrm{MNL}}$, $\boldsymbol{v}^{\mathrm{RCI\text{-}2}}$ and $\boldsymbol{v}^{\mathrm{RCI\text{-}3}}$ and use the root mean squared error (RMSE) to measure the prediction accuracy:

$$\mathrm{RMSE}\,(\boldsymbol{v}) = \sqrt{\frac{\sum_{S\subseteq\mathcal{N}} \sum_{i\in S\cup\{0\}} \left(\mathbb{P}_{\boldsymbol{v}}(i \mid S) - \mathbb{P}_{\mathrm{true}}(i|S)\right)^2}{\sum_{S\subseteq\mathcal{N}}(|S|+1)}},$$

where $\mathbb{P}_{\boldsymbol{v}}(i \mid S) = \frac{v_i}{1+\sum_{j\in S} v_j}$ represents the choice probability under the estimated MNL model and $\mathbb{P}_{\mathrm{true}}(i|S) = \sum_{g=1}^k \alpha_g \frac{v_i^g}{1+\sum_{j\in S} v_j^g}$ is the ground-truth choice probability under the MMNL model. This metric is computed over all possible assortments. The RMSE is frequently used in measuring the prediction of choice models (Berbeglia et al. 2022). The lower value indicates the fitted model is close to the ground-truth model. We calculate the percentage improvement over the MNL model as

$$\frac{\mathrm{RMSE}\,(\boldsymbol{v}^{\mathrm{MNL}}) - \mathrm{RMSE}\,(\boldsymbol{v}^{\mathrm{RCI}})}{\mathrm{RMSE}\,(\boldsymbol{v}^{\mathrm{MNL}})} \times 100.$$

We also calculate the downstream impact of the prediction accuracy on the expected revenue. More precisely, given the estimated preference weights $\boldsymbol{v}$ and revenue $\boldsymbol{r}$, we can calculate the optimal assortment under the standard MNL choice model, denoted as $S_{\boldsymbol{v}}$. Thus, we are comparing the three assortments $S_{\boldsymbol{v}^{\mathrm{MNL}}}$, $S_{\boldsymbol{v}^{\mathrm{RCI\text{-}2}}}$ and $S_{\boldsymbol{v}^{\mathrm{RCI\text{-}3}}}$. For a given assortment S and the revenues $\boldsymbol{r}$ of products, we calculate the expected revenue under the ground truth $R(S, \boldsymbol{r}) = \sum_{i\in S} r_i \mathbb{P}_{\boldsymbol{v}^{\mathrm{true}}}(i|S)$. We sample

10000 revenue vector uniformly from the range $[0,10]$. To capture the improvement of our model, we focus on the instances where the optimized assortments differ, denoted by the set of indices $\mathcal{J} = \{j : S_{\boldsymbol{v}^{\mathrm{RCI}}} \neq S_{\boldsymbol{v}^{\mathrm{MNL}}}\}$. We calculate the average percentage revenue improvement relative to method one over these samples:

$$\sum_{j \in \mathcal{J}} \frac{R(S_{\boldsymbol{v}^{\mathrm{RCI}}}, \boldsymbol{r}^j) - R(S_{\boldsymbol{v}^{\mathrm{MNL}}}, \boldsymbol{r}^j)}{R(S_{\boldsymbol{v}^{\mathrm{MNL}}}, \boldsymbol{r}^j)} \times 100.$$

In the experiment, we vary the number of customer segments $k \in \{1,3,5,8\}$, the total number of assortments in the transaction data $T \in \{300,600,1200,3000\}$, the ratio $\rho \in \{0.1,0.3,0.8\}$, and the total number of products $n \in \{10,15,20\}$. Overall, we have 144 parameter settings. For each setting, we repeat 1000 trials and report the average percentage improvement in RMSE and revenue in Table 2 and 3.

The results demonstrate that incorporating repeated interactions consistently improves the predictive accuracy, with the magnitude of improvement depending on the interplay between the fraction of repeat customers and the level of heterogeneity. Our method yields the most substantial gains in the absence of model misspecification ($k=1$). For example, with $n=10$, $\rho=0.8$, and $T=300$, RCI-3 reduces RMSE by 1.858%, indicating that explicitly modeling interaction history enables a more precise estimation of utility parameters. Moreover, this predictive advantage persists even under model misspecification ($k>1$), particularly in settings with a moderate proportion of repeat customers ($\rho=0.3$). In these cases, the RCI model continues to provide statistically significant improvements over the standard MNL; for instance, with $n=10$, the RMSE improves by approximately 0.5%–0.7% across varying k. This suggests that repeated interactions remain a valuable signal for future customer behavior, even when the ground-truth model is imperfectly specified.

Table 3 reports the percentage revenue increase for instances where the RCI model yields a different optimal assortment than the standard MNL. The results confirm that the predictive accuracy gains identified earlier translate directly into higher expected revenue through assortment optimization under the estimated parameters. Crucially, the revenue advantage remains robust and often increases under model misspecification ($k>1$). Our model frequently outperforms the baseline by a wider margin as heterogeneity increases. For instance, with $n=10$ and $\rho=0.3$, the revenue improvement for RCI-3 rises from 0.227% at $k=1$ to 0.424% at $k=8$. This trend implies that in complex markets, interaction history effectively proxies for unobserved heterogeneity, allowing the seller to capture value that aggregate models overlook. Furthermore, RCI-3 consistently yields higher gains than RCI-2, confirming that leveraging the full historical information is essential for maximizing the expected revenue.

n	ρ	T	RCI-2 $k=1$	RCI-2 $k=3$	RCI-2 $k=5$	RCI-2 $k=8$	RCI-3 $k=1$	RCI-3 $k=3$	RCI-3 $k=5$	RCI-3 $k=8$
10	0.1	300	0.100%**	0.130%	0.146%	0.113%	0.183%**	0.183%	0.313%	0.260%
10	0.1	600	0.177%	0.153%	0.118%	0.128%	0.314%	0.267%	0.215%	0.242%
10	0.1	1200	0.153%	0.077%**	0.096%	0.141%	0.342%	0.106%*	0.209%	0.234%
10	0.1	3000	0.178%	0.107%	0.064%**	0.096%	0.241%	0.189%	0.127%	0.230%
10	0.3	300	0.247%	0.235%	0.279%	0.350%	0.778%	0.552%	0.739%	0.618%
10	0.3	600	0.446%	0.414%	0.296%	0.295%	0.975%	0.776%	0.646%	0.576%
10	0.3	1200	0.541%	0.272%	0.263%	0.153%**	0.934%	0.470%	0.624%	0.533%
10	0.3	3000	0.221%**	0.196%	0.097%**	0.123%**	0.637%	0.504%	0.313%	0.404%
10	0.8	300	0.632%	0.288%**	0.514%	0.466%	1.858%	1.090%	1.079%	1.259%
10	0.8	600	0.303%**	0.509%	0.381%	0.290%	1.364%	0.977%	0.816%	0.601%
10	0.8	1200	0.452%	0.288%	$-0.004\%^{\dagger}$	$-0.150\%^{\dagger}$	1.400%	0.560%	0.203%†	$-0.214\%^{\dagger}$
10	0.8	3000	0.530%	-0.114%†	-0.432%†	-0.715%†	1.618%	-0.609%†	-1.210%†	-1.971%†
15	0.1	300	0.047%*	0.061%**	0.084%	0.106%	0.133%	0.234%	0.181%	0.206%
15	0.1	600	0.112%	0.097%	0.087%	0.110%	0.191%	0.141%	0.210%	0.307%
15	0.1	1200	0.048%*	0.086%	0.092%	0.075%	0.151%	0.159%	0.188%	0.134%
15	0.1	3000	0.031%†	0.103%	0.062%	0.069%	0.122%**	0.204%	0.097%	0.129%
15	0.3	300	0.212%	0.204%	0.325%	0.134%	0.495%	0.426%	0.481%	0.429%
15	0.3	600	0.169%	0.105%**	0.158%	0.193%	0.383%	0.240%	0.333%	0.296%
15	0.3	1200	0.295%	0.202%	0.262%	0.178%	0.605%	0.454%	0.428%	0.300%
15	0.3	3000	0.162%	0.146%	0.118%	0.122%	0.325%	0.217%	0.224%	0.297%
15	0.8	300	0.332%	0.249%	0.244%	0.247%	0.872%	0.843%	0.782%	0.776%
15	0.8	600	0.317%	0.310%	0.148%**	0.233%	0.883%	0.715%	0.534%	0.491%
15	0.8	1200	0.314%	0.114%*	-0.000%†	0.114%*	0.874%	0.304%	0.115%†	0.210%**
15	0.8	3000	0.422%	-0.166%†	-0.184%†	-0.276%†	0.923%	-0.304%†	-0.479%†	-0.749%†
20	0.1	300	0.092%	0.089%	0.079%	0.070%	0.142%	0.211%	0.128%	0.107%
20	0.1	600	0.044%**	0.043%**	0.072%	0.065%	0.120%	0.098%	0.118%	0.129%
20	0.1	1200	0.035%*	0.074%	0.054%	0.080%	0.081%**	0.121%	0.094%	0.208%
20	0.1	3000	0.045%**	0.053%	0.058%	0.041%**	0.144%	0.108%	0.113%	0.092%
20	0.3	300	0.122%	0.165%	0.157%	0.107%	0.329%	0.320%	0.380%	0.253%
20	0.3	600	0.188%	0.162%	0.111%	0.133%	0.309%	0.236%	0.251%	0.287%
20	0.3	1200	0.207%	0.100%	0.097%	0.094%	0.424%	0.177%	0.263%	0.192%
20	0.3	3000	0.161%	0.062%**	0.112%	0.021%†	0.388%	0.163%	0.235%	0.134%
20	0.8	300	0.172%	0.252%	0.267%	0.242%	0.517%	0.535%	0.393%	0.516%
20	0.8	600	0.287%	0.212%	0.185%	0.175%	0.716%	0.565%	0.542%	0.450%
20	0.8	1200	0.254%	0.164%	0.140%	0.211%	0.659%	0.416%	0.335%	0.425%
20	0.8	3000	0.281%	0.085%**	-0.034%†	-0.055%†	0.747%	0.099%*	-0.116%†	-0.298%†

Table 2 Percentage improvement in prediction accuracy (RMSE) over the MNL model when the ground truth is an MMNL model. All estimates are statistically significant at the $p < 0.01$ level, except for those marked with ** ($p < 0.05$), * ($p < 0.1$) or † (not significant at the 0.1 level).

7.2. Hotel Booking Data from Meituan

We use a private dataset from Meituan, one of the largest online travel agencies in China. It contains detailed user-level browsing and booking histories in a prominent university district of Shanghai during April, 2025. When a user searches on Meituan, a list of hotel options are presented.

			RCI-2				RCI-3			
n	ρ	T	$k=1$	$k=3$	$k=5$	$k=8$	$k=1$	$k=3$	$k=5$	$k=8$
10	0.1	300	0.205%	0.373%	0.322%	0.359%	0.153%	0.368%	0.373%	0.380%
10	0.1	600	0.043%**	0.272%	0.308%	0.289%	0.043%**	0.230%	0.213%	0.187%
10	0.1	1200	0.053%	0.142%	0.207%	0.221%	0.035%	0.152%	0.219%	0.172%
10	0.1	3000	0.048%	0.156%	0.195%	0.154%	0.029%	0.141%	0.171%	0.156%
10	0.3	300	0.166%	0.342%	0.330%	0.386%	0.227%	0.259%	0.400%	0.424%
10	0.3	600	0.132%	0.270%	0.360%	0.283%	0.132%	0.297%	0.338%	0.285%
10	0.3	1200	0.105%	0.256%	0.267%	0.202%	0.092%	0.213%	0.302%	0.234%
10	0.3	3000	0.046%	0.248%	0.289%	0.307%	0.040%	0.297%	0.315%	0.312%
10	0.8	300	0.168%	0.356%	0.428%	0.398%	0.187%	0.405%	0.543%	0.466%
10	0.8	600	0.120%	0.314%	0.342%	0.379%	0.145%	0.367%	0.419%	0.465%
10	0.8	1200	0.079%	0.310%	0.374%	0.375%	0.104%	0.373%	0.474%	0.462%
10	0.8	3000	0.055%	0.316%	0.393%	0.400%	0.056%	0.381%	0.467%	0.480%
15	0.1	300	0.170%	0.417%	0.529%	0.576%	0.088%	0.286%	0.350%	0.335%
15	0.1	600	0.051%	0.114%	0.120%	0.161%	0.067%	0.166%	0.147%	0.227%
15	0.1	1200	0.057%	0.077%	0.148%	0.132%	0.050%	0.074%	0.143%	0.112%
15	0.1	3000	0.015%**	0.080%	0.093%	0.135%	0.021%	0.078%	0.111%	0.126%
15	0.3	300	0.161%	0.262%	0.319%	0.361%	0.209%	0.263%	0.236%	0.380%
15	0.3	600	0.067%	0.162%	0.185%	0.203%	0.075%	0.169%	0.151%	0.197%
15	0.3	1200	0.109%	0.139%	0.155%	0.183%	0.089%	0.140%	0.164%	0.192%
15	0.3	3000	0.030%	0.122%	0.114%	0.150%	0.021%	0.127%	0.135%	0.161%
15	0.8	300	0.152%	0.218%	0.228%	0.248%	0.166%	0.297%	0.318%	0.290%
15	0.8	600	0.082%	0.160%	0.211%	0.214%	0.099%	0.185%	0.230%	0.237%
15	0.8	1200	0.060%	0.134%	0.175%	0.205%	0.063%	0.167%	0.205%	0.231%
15	0.8	3000	0.042%	0.140%	0.194%	0.198%	0.049%	0.169%	0.223%	0.229%
20	0.1	300	0.287%	0.458%	0.489%	0.501%	0.196%	0.260%	0.245%	0.309%
20	0.1	600	0.092%	0.111%	0.159%	0.130%	0.093%	0.105%	0.173%	0.119%
20	0.1	1200	0.033%	0.107%	0.104%	0.127%	0.027%**	0.076%	0.088%	0.082%
20	0.1	3000	0.019%	0.046%	0.067%	0.054%	0.017%	0.044%	0.061%	0.074%
20	0.3	300	0.095%	0.235%	0.183%	0.154%	0.156%	0.269%	0.228%	0.259%
20	0.3	600	0.139%	0.135%	0.129%	0.172%	0.127%	0.110%	0.115%	0.179%
20	0.3	1200	0.081%	0.117%	0.105%	0.110%	0.087%	0.106%	0.133%	0.126%
20	0.3	3000	0.026%	0.070%	0.084%	0.091%	0.026%	0.058%	0.084%	0.089%
20	0.8	300	0.143%	0.171%	0.204%	0.201%	0.156%	0.158%	0.226%	0.242%
20	0.8	600	0.045%	0.140%	0.133%	0.151%	0.087%	0.140%	0.162%	0.189%
20	0.8	1200	0.048%	0.082%	0.118%	0.130%	0.062%	0.105%	0.144%	0.159%
20	0.8	3000	0.033%	0.081%	0.111%	0.111%	0.043%	0.095%	0.129%	0.135%

Table 3 Percentage improvement in revenue over the MNL model when the ground truth is an MMNL model. All estimates are statistically significant at the $p < 0.01$ level, except for those marked with ** ($p < 0.05$).

The user may either book one of the hotels or leave without placing an order. The dataset comprises $334,808$ interaction records from $37,871$ users and 683 hotels. Each record contains the date, user ID, hotel ID, and an indicator of whether the user booked the hotel.

Since users may browse hotels over multiple days before making a final booking decision, we define a decision period as the sequence of hotel views leading up to and including the booking day. The assortment for each decision period is constructed by aggregating all hotels viewed by the user within that period, under the assumption that users consider all viewed hotels before booking. After constructing decision periods, we find that 97.81% of users appear in only one period, and an additional 1.52% appear in two, accounting for 99.33% of all users. Since over 99% of users complete their booking within one or two periods, we restrict our analysis to this majority group and exclude the small fraction of users with more than two periods. This two-period setting serves two purposes. First, it is sufficient to demonstrate the distinction between our method and the traditional MNL model without considering the past interaction. Second, incorporating more periods would make the estimation of choice probabilities computationally expensive for such a dataset. We also exclude hotels with fewer than 10 bookings during the month, as such infrequent choices provide limited information for estimating preference parameters. After these filters, the final dataset consists of $35,236$ users and 170 hotels.

We randomly split users into training (60%) and testing sets (40%). We split the data 20 times, thus creating 20 folds of training/testing pairs. For each pair, we estimate the discounted version of our model in Section 6 (denoted as RCI-MNL) and the MNL model from the training data. Then we evaluate the two fitted models using two metrics in the testing data: the log-likelihood and the root mean squared error, defined as

$$\mathrm{RMSE}\,(\mathcal{T}) = \sqrt{\frac{\sum_{S_t\in\mathcal{T}}\sum_{i\in S_t\cup\{0\}}\left(\mathbb{P}^{\mathrm{MNL/RCI}}(i\mid S_t) - \mathbb{1}\{i_t = i\}\right)^2}{\sum_{S_t\in\mathcal{T}}(|S_t|+1)}},$$

where $\mathcal{T}$ represents the testing data, and i_t is the customer choice in the assortment S_t. The RMSE is defined using the actual choice instead of the ground truth, different from Section 7.1. We divide $\mathcal{T}$ into $\mathcal{T}^s$ and $\mathcal{T}^t$, where $\mathcal{T}^s$ includes the customers who only interact with the seller once and $\mathcal{T}^t$ includes the customers with repeated interactions. We compare the RMSE in $\mathcal{T}^s$ and $\mathcal{T}^t$ separately.

	MNL	RCI-MNL	Improvement
Log-likelihood	-12011	-11688	2.69%
RMSE($\mathcal{T}^s$)	0.2013	0.2011	0.10%
RMSE($\mathcal{T}^t$)	0.4187	0.3515	10.76%

Table 4 The average performance of MNL and our model (RCI-MNL) in the testing data

Table 4 reports the out-of-sample performance of the two models in terms of log-likelihood and RMSE. Despite the fact that only a subset of users exhibit repeated bookings, incorporating past interactions substantially improves the model's predictive accuracy. The improvement is particularly

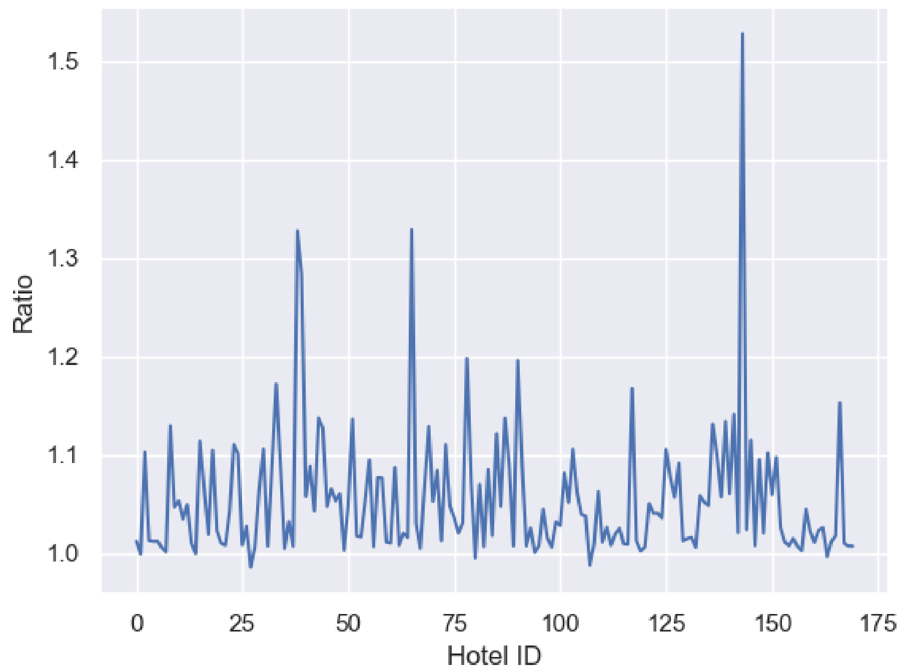

Figure 3 **Ratio of estimated preference weights (MNL over RCI-MNL) across hotels**

pronounced among users with repeated purchases, as reflected by a larger reduction in RMSE. These findings provide strong empirical evidence that leveraging customers' historical behavior enhances the model's predictive power for future booking decisions.

We also visualize the ratio of the estimated preference weights under the two models in Figure 3. Specifically, the ratio is computed as the preference weight estimated by the MNL model divided by that from the RCI-MNL model for each hotel. The figure reveals a consistent pattern: when repeated interactions are not accounted for, the MNL model tends to assign inflated preference weights to the hotels. This inflation arises because the standard MNL model cannot distinguish between isolated bookings and consistent rebookings that signal long-term interest. It highlights the importance of capturing users' historical decision patterns and persistent preferences in choice modeling, especially when historical interactions are available.

For the revenue comparison, we note that the estimated discount factor d is only 0.01, suggesting that the utility of the booking event remains effectively unchanged across periods. Under this assumption, we compare the expected revenue generated by the optimal fixed policy and the optimal adaptive policy by solving the joint optimization problem defined in Equation (13). The results show that the optimal adaptive policy yields a 23.41% improvement in revenue over two periods, increasing from 504 RMB under the fixed policy to 622 RMB under the adaptive one. This finding demonstrates that explicitly modeling repeated user interactions can lead to significantly improved revenue outcomes.

7.3. Revenue Performance on Sushi Preference Data

To complement the theoretical analysis for the benefit of the adaptive policy relative to the fixed policy, in this section we use a numerical experiment to test the robustness of the benefit when the ground-truth model is not based on the MNL model. In particular, we utilize a public dataset that includes the preference rankings (e.g., prefer A to C to B) of 5000 people for 10 popular sushi items (Kamishima 2003). Because of the format, the dataset can be readily used to construct the rank-based choice model (Farias et al. 2013), which is nonparametric and incorporates a wider range of behavior than the MNL model. We then use the rank-based model to simulate repeated interactions of customers and test the performance of the optimal adaptive policy recommended by our model and the optimal fixed policy solved from the MNL model, both of which misspecify the choice behavior because they are based on Gumbel random utilities instead of preference rankings. The experiment is used to demonstrate the robustness of the main insight of this study: incorporating the repeated interactions can significantly improve the revenue, even when the underlying choice behavior is more sophisticated than the MNL assumption.

Experimental setup We first construct the ground-truth rank-based choice model from the data with a similar approach to Aouad et al. (2023). More precisely, we sample K rankings uniformly from the 5000 preference rankings. For each sampled ranking σ_k, we insert the no-purchase option into a random position. To capture the price sensitivity, we follow the idea in Jagabathula and Rusmevichientong (2017) and use a threshold b_k. Specifically, the customer ignores the products with price larger than b_k and chooses the most preferred product from the qualified products. Since the largest price for the sushi is 4.485, we uniformly sample b_k in the range $[0,6]$. We also generate the probability of the rankings in the same fashion as Aouad et al. (2023). The number of customer types K is chosen from $\{100, 250, 500\}$. For each K, we construct 20 different ground-truth models.

Given the ground-truth model, we generate the transaction data by a similar approach to Jagabathula and Rusmevichientong (2017): we randomly sample 30 price vectors (following a uniform distribution between 0 and the listed price) and sample 1000 random assortments (each product is included independently with probability 0.5). To generate the two-period transaction data, we first pair any two assortments at random, which represents the assortments in the first period and second period. Then we combine the 500 pairs with the 30 price vectors, resulting in $15,000$ combinations. For each combination, we sample the customer choice from the ground-truth model. For the preference weight of product i under the standard MNL model, we set it as $v_i = \exp(\mu_i - \beta_i r_i)$, where r_i is the price for product i. With the maximum likelihood estimator, we fit the parameters $\{\boldsymbol{\mu}, \boldsymbol{\beta}\}$ from the data.

With the fitted preference weights, we derive the optimal fixed policy and adaptive policy under the MNL choice model over two periods based on our model. Moreover, we also compute the optimal assortment for the ground-truth rank-based model in a single period, which is the actual *optimal fixed* policy. To calculate the expected revenues of the three policies under the ground-truth model, we sample 1000 price vectors, each of which leads to realized revenues under the policies. The average revenue then serves as the policy performance. Note that it is unclear what is the structure of or how to compute the optimal adaptive recommendation policy under the rank-based choice model. Thus, we do not include it as a benchmark and leave it for future research. Moreover, given the price vector, we compare the prediction ability of MNL and our model by root mean squared error.

Computational results In Table 5, we first show the average revenue improvement by implementing the adaptive recommendation policy for $K \in \{100, 250, 500\}$. The result confirms that even when our model misspecifies the ground truth, incorporating the sequential choice behavior in repeated interactions can significantly boost the profitability by learning the preference of customers in previous interactions. The improvement in prediction shows adding repeated interaction in estimation benefits the prediction, even in the model misspecification setting.

	$K = 100$	$K = 250$	$K = 500$
Revenue Improvement Over MNL Fixed	11.83%	11.72%	11.26%
Revenue Improvement Over Optimal Fixed	3.87%	5.50%	6.74%
Prediction Improvement Over MNL	0.34%	0.97%	1.03%

Table 5 **Improvement of the adaptive policy recommended by our model in the sushi dataset**

8. Conclusion

Motivated by new business models such as Stitch Fix and Wantable, we propose a model in which the seller offers a sequence of assortments when the customer interacts with the seller repeatedly. We study the multi-period assortment optimization problem. For two periods, we show the conditional assortment optimization in the second period is NP-hard to solve when the customer has made a purchase in the first period. We then consider the joint assortment optimization problem and establish the performance guarantee for the fixed recommendation policy. For the multi-period setting, we focus on analyzing the performance of the fixed recommendation policy, which is more practical and efficient to implement. The results show the fixed policy is quite robust and can guarantee at least 50% of the expected revenue compared with the optimal policy. We extend our result to the discounted setting that the utility of the purchased product will decrease by a positive constant and show the fixed policy could also guarantee a constant approximation ratio. The numerical experiments based on real datasets confirm the benefits of incorporating the repeated interactions.

We conclude by pointing to a number of future directions. First, in the subscription business model, the firm typically collects a subscription fee upfront, not modeled in the current study. It effectively creates a "sunk cost" for consumers, potentially eliminating the no-purchase option and fundamentally reshaping the trade-off between learning customer preference and maximizing the expected revenue. Future research could formalize this setting to determine how such forced-choice dynamics alter the optimal adaptive policy. Second, while our model focuses on the demand side, inventory management remains a critical operational constraint. It is valuable to develop efficient heuristics to jointly optimize learning-based assortment planning and inventory replenishment. Third, in reality, due to the space limitation, the seller faces the cardinality constraint in the assortment. It would be interesting to explore the benefit of the adaptive policy and develop efficient algorithms to approximate the optimal policy in this case. Finally, extending our framework to more complex choice models such as nested logit model and markov chain choice model, or incorporating pricing decision alongside assortment optimization, would further enhance the practical applicability of sequential learning models.

Acknowledgments

The authors thank the department editor, associate editor, and anonymous reviewers for insightful comments and constructive suggestions that greatly improved the paper. The authors gratefully acknowledge the funding support from Meituan. C. Wang's research was supported by the National Science Foundation of China [Grants 72495133, 72172104] and China Postdoctoral Science Foundation [Grant 2025M780760]. P. Gao's research was supported by the National Natural Science Foundation of China [Grants 72522026, 72201234, and 72192805], Collaborative Research Funding [Grant C6032-21G] of the Hong Kong Research Grants Council, and the Guangdong Provincial Key Laboratory of Mathematical Foundations for Artificial Intelligence [Grant 2023B1212010001]. Y. Wang's research was supported by the National Natural Science Foundation of China under Grant [12371513], and the Major Program of National Fund of Philosophy and Social Science of China under Grant [23&ZD135].

Notes

1. We follow the standard notation in the literature and use $[n] \triangleq \{1, \ldots, n\}$ for a positive integer n.

2. `https://www.businessofapps.com/data/shein-statistics/`

3. `https://www.sheingroup.com/how-we-do-business/`

References

Aouad A, Feldman J, Segev D (2023) The exponomial choice model for assortment optimization: an alternative to the mnl model? *Management Science* 69(5):2814–2832.

Bai Y, Feldman J, Segev D, Topaloglu H, Wagner L (2024) Assortment optimization under the multi-purchase multinomial logit choice model. *Operations Research* 72(6):2631–2664.

Berbeglia G, Garassino A, Vulcano G (2022) A comparative empirical study of discrete choice models in retail operations. *Management Science* 68(6):4005–4023.

Bernstein F, Guo Y (2023) Managing customer search: Assortment planning for a subscription box service. *Manufacturing & Service Operations Management* 25(5):1623–1642.

Bront JJM, Méndez-Díaz I, Vulcano G (2009) A column generation algorithm for choice-based network revenue management. *Operations research* 57(3):769–784.

Calmon AP, Ciocan FD, Romero G (2021) Revenue management with repeated customer interactions. *Management Science* 67(5):2944–2963.

Désir A, Goyal V, Zhang J (2022) Capacitated assortment optimization: Hardness and approximation. *Operations Research* 70(2):893–904.

El Housni O, Topaloglu H (2023) Joint assortment optimization and customization under a mixture of multinomial logit models: On the value of personalized assortments. *Operations research* 71(4):1197–1215.

Farias VF, Jagabathula S, Shah D (2013) A nonparametric approach to modeling choice with limited data. *Management science* 59(2):305–322.

Fata E, Ma W, Simchi-Levi D (2019) Multi-stage and multi-customer assortment optimization with inventory constraints. Working Paper, `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3443109`.

Feldman J, Segev D (2022) The multinomial logit model with sequential offerings: Algorithmic frameworks for product recommendation displays. *Operations Research* 70(4):2162–2184.

Ferreira KJ, Goh J (2021) Assortment rotation and the value of concealment. *Management Science* 67(3):1489–1507.

Flores A, Berbeglia G, Van Hentenryck P (2019) Assortment optimization under the sequential multinomial logit model. *European Journal of Operational Research* 273(3):1052–1064.

Gallego G, Berbeglia G (2022) Bounds, heuristics, and prophet inequalities for assortment optimization. Working Paper, `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3934096`.

Gallego G, Wang R (2019) Threshold utility model with applications to retailing and discrete choice models. Working Paper, `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3420155`.

Gao P, Ma Y, Chen N, Gallego G, Li A, Rusmevichientong P, Topaloglu H (2021) Assortment optimization and pricing under the multinomial logit model with impatient customers: Sequential recommendation and selection. *Operations Research* 69(5):1509–1532.

Hamilton M, Singal R (2023) Churning while experimenting: On the long-term optimality of short-term churn minimization. Working Paper, `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3871915`.

Huh WT, Li H (2022) Optimal pricing under multiple-discrete customer choices and diminishing return of consumption. *Operations Research* 70(2):905–917.

Jagabathula S, Mitrofanov D, Vulcano G (2022) Personalized retail promotions through a dag-based representation of customer preferences. *Operations Research.* 70(2):641–655.

Jagabathula S, Rusmevichientong P (2017) A nonparametric joint assortment and price choice model. *Management Science* 63(9):3128–3145.

Jasin S, Lyu C, Najafi S, Zhang H (2024) Assortment optimization with multi-item basket purchase under multivariate mnl model. *Manufacturing & Service Operations Management* 26(1):215–232.

Kamishima T (2003) Nantonac collaborative filtering: recommendation based on order responses. *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, 583–588.

Lin H, Li X, Wu L (2023) Express: Intra-category multi-choice preferences learning and assortment recommendation in e-commerce. *Production and Operations Management* 10591478251350853.

Liu N, Ma Y, Topaloglu H (2020) Assortment optimization under the multinomial logit model with sequential offerings. *INFORMS Journal on Computing* 32(3):835–853.

Luan S, Wang R, Xu X, Xue W (2025) Joint assortment and price optimization with multiple purchases. *Production and Operations Management* 34(2):187–204.

Rusmevichientong P, Shen ZJM, Shmoys DB (2010) Dynamic assortment optimization with a multinomial logit choice model and capacity constraint. *Operations Research* 58(6):1666–1680.

Rusmevichientong P, Shmoys D, Tong C, Topaloglu H (2014) Assortment optimization under the multinomial logit model with random choice parameters. *Production and Operations Management* 23(11):2023–2039.

Rusmevichientong P, Topaloglu H (2012) Robust assortment optimization in revenue management under the multinomial logit choice model. *Operations research* 60(4):865–882.

Sen A, Atamtürk A, Kaminsky P (2018) A conic integer optimization approach to the constrained assortment problem under the mixed multinomial logit model. *Operations Research* 66(4):994–1003.

Sumida M, Gallego G, Rusmevichientong P, Topaloglu H, Davis J (2021) Revenue-utility tradeoff in assortment optimization under the multinomial logit model with totally unimodular constraints. *Management Science* 67(5):2845–2869.

Sumida M, Zhou A (2023) Optimizing and learning assortment decisions in the presence of platform disengagement. Working Paper, `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4537925`.

Talluri K, Van Ryzin G (2004) Revenue management under a general discrete choice model of consumer behavior. *Management Science* 50(1):15–33.

Tang Z, Jasin S, Duenyas I (2023) On the performance of myopic and revenue-ordered policies for a multi-period assortment optimization under an mnl model with popularity bias. Working Paper, `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4454399`.

Tulabandhula T, Sinha D, Karra SR, Patidar P (2023) Multi-purchase behavior: Modeling, estimation, and optimization. *Manufacturing & Service Operations Management* 25(6):2298–2313.

Wang R (2022) What is the impact of nonrandomness on random choice models? *Manufacturing & Service Operations Management* 24(1):485–503.

Xu Y, Wang Z (2023) Assortment optimization for a multistage choice model. *Manufacturing & Service Operations Management* 25(5):1748–1764.

Zhang DJ, Hu M, Liu X, Wu Y, Li Y (2022) Netease cloud music data. *Manufacturing & Service Operations Management* 24(1):275–284.

Zhang H, Piri H, Huh WT, Li H (2021) Assortment optimization under multiple-discrete customer choices. Working Paper, `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3988982`.