

A Greedy Algorithm for the Two-Level Nested Logit Model

Guang Li, Paat Rusmevichientong*

Marshall School of Business, University of Southern California

Abstract

We consider the assortment optimization problem under the classical two-level nested logit model. We establish a necessary and sufficient condition for the optimal assortment and develop a simple and fast greedy algorithm that iteratively removes at most one product from each nest to compute an optimal solution.

Keywords: revenue management, customer choice, assortment optimization, nested logit model, greedy algorithm, optimality condition

1. Introduction

In an assortment optimization problem, a firm wishes to offer the most profitable set of products to satisfy customer demand. As shown in an excellent survey in [1], this problem has many important applications in retail and revenue management. Our paper focuses on assortment optimization in revenue management, where, unlike in the retail setting, the focus is primarily on product selections, without any inventory replenishment.

The seminal work in [2] introduced the first assortment optimization problem under the multinomial logit (MNL) model [3]. Although the MNL model admits a tractable solution, it suffers from the independence of irrelevant alternatives (IIA) property, which may lead to a paradox when alternatives are closely related and result in biased estimates [4]. The deficiency of the MNL model has led many researchers to consider assortment optimization under more complex choice models; see, for example, [5–7]. The nested logit model extends the MNL model by grouping similar alternatives into a nest and allowing differential substitution patterns within and across nests, thereby partially relaxing the IIA restriction [8]. Under the nested logit model, a consumer first chooses a nest of products and then chooses a product from the chosen nest. Existing literature on assortment optimization under the nested logit model makes use of a linear programming framework; see, for example, [9–11].

In this paper, we present an alternative solution to assortment optimization under the two-level nested logit model. We establish a necessary and sufficient condition for an optimal assortment (Theorem 2.3), which is, to our knowledge, the first such optimality condition for this problem. In addition, we reveal a “lumpy” structure of the optimal assortment: surprisingly, within each nest, a certain set of “consecutive” products always appears together in the optimal solution. Moreover, as shown in Theorem 2.4, by looking at a certain index of each product, we can determine *in advance* which products will appear together! Exploiting the optimality condition and the “lumpy” structure, our greedy algorithm iteratively removes (at most) one product from each nest and terminates with an optimal assortment in $O(nm \log m)$ time, for the problem with m nests, with each nest having n products. This is the fastest known running time for this problem.

Moreover, our assortment optimization can be used as a subroutine for the network revenue management problem under the nested logit model. A revenue model over a single-leg flight is first considered in [2]. For extensions of this work to multiple resources, see, for example, [12–14]. The main approach in these papers is to formulate a variety of deterministic linear programming approximations and solve the linear programs using column generation. The column generation subproblem in this setting is precisely the assortment optimization considered in this paper, when customers choose according to a two-level nested logit model.

*Corresponding Author: Marshall School of Business, University of Southern California, Los Angeles, CA 90089, Email: rusmevic@marshall.usc.edu

We perform sensitivity analysis to study the effect of reducing product revenues on the optimal assortment. We show that the entire profile of the optimal revenue and assortment when every product revenue is reduced by δ can be computed simultaneously for all $\delta \in \mathbb{R}_+$ in a single run of the greedy algorithm in $O(nm \log m)$ time. Moreover, we show that the optimal solution to the multi-period capacity allocation problem with total allocation periods T and total capacity C can be computed in $O(TC + nm \log m)$ operations, by applying the greedy algorithm just once.

Problem Description. We have m nests indexed by $i \in \{1, 2, \dots, m\}$, and each nest i has n products indexed by $j \in \{1, 2, \dots, n\}$. All of the analysis easily extends to the case in which some nests have fewer than n products. A product j in nest i is denoted by ji , and in settings in which the nest index is clear from context, we refer to it simply as product j . We denote the no-purchase option by product 0, which is the only product in nest 0. Each product ji has a revenue $r_{ji} \geq 0$, with $r_0 = 0$. Without loss of generality, we assume that the products in each nest are sorted in a descending order of revenues; that is, for each nest $i \in \{1, 2, \dots, m\}$, $r_{1i} \geq r_{2i} \geq \dots \geq r_{ni} \geq 0$.

We assume that the customer is a utility maximizing agent, and the utility that she assigns to each product is:

$$U_0 = \lambda_0 + \epsilon_0 \quad \text{and} \quad U_{ji} = \lambda_{ji} + \epsilon_i + \epsilon_{ji}, \quad i = 1, \dots, m, \quad j = 1, \dots, n,$$

where λ_{ji} and λ_0 are the deterministic components and $\{\epsilon_0, \epsilon_i, \epsilon_{ji} : i = 1, \dots, m, \quad j = 1, \dots, n\}$ are independent random errors. The random variable ϵ_0 has a Gumbel distribution with a location parameter of zero and a scaling parameter of one, which we write as $\epsilon_0 \sim \text{Gumbel}(0, 1)$. For each nest i , $\epsilon_{ji} \sim \text{Gumbel}(0, 1/\tau_i)$ for all j , where $\tau_i \in (0, 1]$ is a nest-specific parameter. Finally, for each i , ϵ_i is independent of $\{\epsilon_{ji} : j = 1, \dots, n\}$, and it is assumed to have a distribution such that $\epsilon_i + \epsilon_{ji} \sim \text{Gumbel}(0, 1)$ for all j ; see [15] for the existence of such a distribution.

Let $\mathbf{S} = (S_1, \dots, S_m)$ denote an assortment of products, where S_i is a subset of the n products $\{1i, 2i, \dots, ni\}$ in nest i , for all $i \in \{1, 2, \dots, m\}$. As the nest index i is explicit, we simply write $S_i \subseteq \{1, 2, \dots, n\}$. For each product ji , let $v_{ji} = e^{\lambda_{ji}/\tau_i}$ denote its preference weight, and let $v_0 = e^{\lambda_0}$. As shown in [16], the probability $Q_{ji}(\mathbf{S})$ that a customer chooses a product $j \in S_i$ under assortment $\mathbf{S}$ is given by:

$$Q_{ji}(\mathbf{S}) = Q_{j|i}(S_i) \times Q_i(\mathbf{S}), \quad \text{where} \quad Q_{j|i}(S_i) = \frac{v_{ji}}{\sum_{\ell \in S_i} v_{\ell i}} \quad \text{and} \quad Q_i(\mathbf{S}) = \frac{(\sum_{\ell \in S_i} v_{\ell i})^{\tau_i}}{v_0 + \sum_{k=1}^m (\sum_{\ell \in S_k} v_{\ell k})^{\tau_k}}.$$

Here, $Q_i(\mathbf{S})$ denotes the probability that a customer chooses a product in nest i , and $Q_{j|i}(S_i)$ denotes the conditional probability of the customer's selecting product ji , given that nest i is chosen.

Thus, the total expected revenue $\text{Rev}(\mathbf{S})$ for any assortment $\mathbf{S} = (S_1, \dots, S_m)$ is given by:

$$\text{Rev}(\mathbf{S}) = \sum_{i=1}^m \sum_{j \in S_i} r_{ji} Q_{ji}(\mathbf{S}) = \sum_{i=1}^m Q_i(\mathbf{S}) \sum_{j \in S_i} r_{ji} Q_{j|i}(S_i) = \sum_{i=1}^m Q_i(\mathbf{S}) \text{Rev}_i(S_i),$$

where $\text{Rev}_i(S_i) = \sum_{j \in S_i} r_{ji} Q_{j|i}(S_i) = \frac{\sum_{\ell \in S_i} r_{\ell i} v_{\ell i}}{\sum_{\ell \in S_i} v_{\ell i}}$ denotes the expected revenue from S_i in nest i , and we set it to be zero if S_i is empty. The assortment optimization problem is thus given by:

$$Z^* = \max_{\mathbf{S}=(S_1, \dots, S_m)} \text{Rev}(\mathbf{S}),$$

and we denote an optimal assortment by $\mathbf{S}^* = (S_1^*, \dots, S_m^*)$ and let $Z^* = \text{Rev}(\mathbf{S}^*)$ denote the optimal revenue. If there are ties, we can choose the optimal assortment according to any predetermined tie-breaking rule.

2. Characterization of the Optimal Assortments

Let $\mathcal{N}_+ = \{\{1\}, \{1, 2\}, \dots, \{1, 2, \dots, n\}\}$ denote the collection of revenue-ordered subsets. As shown in the following lemma, at the optimal solution, any nonempty nest is revenue-ordered. The lemma is a restatement of Theorem 4 in [9], and we refer the reader to the paper for the details of the proof.

Lemma 2.1 (Optimal Nest Structure). *For $i = 1, \dots, m$, if $S_i^* \neq \emptyset$, then $S_i^* \in \mathcal{N}_+$.*

Although Lemma 2.1 shows that $S_i^* \in \{\emptyset\} \cup \mathcal{N}_+$ for all i , exhaustive search is not feasible because there are $(n+1)^m$ possible configurations of revenue-ordered subsets among m nests. To develop an efficient algorithm, we

will establish a *necessary and sufficient* condition for an optimal assortment. For any integers k_1 and k_2 , let

$$[k_1, k_2] = \begin{cases} \{k_1, k_1 + 1, \dots, k_2\} & \text{if } k_1 \leq k_2, \\ \emptyset & \text{otherwise.} \end{cases}$$

Also, for each nest i and $1 \leq k_1 \leq k_2 \leq n$, let $G_i(k_1, k_2)$ be defined by:

$$G_i(k_1, k_2) = \text{Rev}_i([1, k_2]) - \frac{\text{Rev}_i([1, k_2]) - \text{Rev}_i([k_1, k_2])}{f_i \left(1 - \sum_{\ell \in [k_1, k_2]} v_{\ell i} / \sum_{\ell \in [1, k_2]} v_{\ell i} \right)},$$

where $f_i : [0, 1] \rightarrow [0, 1]$ is a continuous, strictly increasing function defined by:

$$f_i(x) = \begin{cases} (x^{1-\tau_i} - x)/(1 - x) & \text{if } x \in [0, 1) \\ \tau_i & \text{if } x = 1 \end{cases}.$$

Moreover, we define $0/0 = 0$, thus $G_i(1, k_2) = \text{Rev}_i([1, k_2])$ for all k_2 .

Interpretation of $G_i(\cdot, \cdot)$. Suppose $S_i = [1, k_2]$, and consider $k_1 \in \{1, \dots, k_2\}$. The function $G_i([k_1, k_2])$ can be used to evaluate the net effect of removing $[k_1, k_2]$ from nest i on the total revenue, without any computation on the resultant assortment. This intuition is confirmed in the following lemma, which establishes a precise condition on when we can use $G_i(\cdot, \cdot)$ to determine if a subset can be removed from S_i to improve the total revenue.

Lemma 2.2 (When Is Removing Subsets Beneficial?). *For any assortment $\mathbf{S} = (S_1, \dots, S_m)$, a collection of subsets $\{A_i \subseteq S_i : i = 1, 2, \dots, m\}$ can be removed from $\mathbf{S}$ to achieve a greater revenue if and only if*

$$\frac{\sum_{i=1}^m [V(S_i)^{\tau_i} \text{Rev}_i(S_i) - V(S_i \setminus A_i)^{\tau_i} \text{Rev}_i(S_i \setminus A_i)]}{\sum_{i=1}^m [V(S_i)^{\tau_i} - V(S_i \setminus A_i)^{\tau_i}]} < \text{Rev}(\mathbf{S}),$$

where for any set X , $V(X)$ denotes the sum of the preference weights of the products in X .

Moreover, if $S_i = [1, p_i]$ with $p_i \geq 1$, $A_i = [k_i, p_i]$ with $k_i \leq p_i$, and $A_\ell = \emptyset$ for all $\ell \neq i$, then removing A_i from S_i leads to a greater revenue if and only if $G_i(k_i, p_i) < \text{Rev}(\mathbf{S})$; removing A_i from S_i leads to the same revenue if and only if $G_i(k_i, p_i) = \text{Rev}(\mathbf{S})$.

Proof. For $i = 1, \dots, m$, let $\hat{S}_i = S_i \setminus A_i$. Assume that $A_i \neq \emptyset$ for some i ; otherwise, the result is trivially true. As $v_{ji} > 0$ for all j , $\sum_{i=1}^m V(\hat{S}_i) < \sum_{i=1}^m V(S_i)$. By definition, $\text{Rev}(\hat{\mathbf{S}}) = \frac{\sum_{i=1}^m V(\hat{S}_i)^{\tau_i} \text{Rev}_i(\hat{S}_i)}{v_0 + \sum_{i=1}^m V(\hat{S}_i)^{\tau_i}}$. Then,

$$\begin{aligned} \text{Rev}(\mathbf{S}) &= \frac{\sum_{i=1}^m V(S_i)^{\tau_i} \text{Rev}_i(S_i)}{v_0 + \sum_{i=1}^m V(S_i)^{\tau_i}} = \frac{v_0 + \sum_{i=1}^m V(\hat{S}_i)^{\tau_i}}{v_0 + \sum_{i=1}^m V(S_i)^{\tau_i}} \times \text{Rev}(\hat{\mathbf{S}}) \\ &+ \frac{\sum_{i=1}^m [V(S_i)^{\tau_i} - V(\hat{S}_i)^{\tau_i}]}{v_0 + \sum_{i=1}^m V(S_i)^{\tau_i}} \times \frac{\sum_{i=1}^m [V(S_i)^{\tau_i} \text{Rev}_i(S_i) - V(\hat{S}_i)^{\tau_i} \text{Rev}_i(\hat{S}_i)]}{\sum_{i=1}^m [V(S_i)^{\tau_i} - V(\hat{S}_i)^{\tau_i}]} \end{aligned}$$

Thus, $\text{Rev}(\mathbf{S})$ is a convex combination of $\text{Rev}(\hat{\mathbf{S}})$ and $\frac{\sum_{i=1}^m [V(S_i)^{\tau_i} \text{Rev}_i(S_i) - V(\hat{S}_i)^{\tau_i} \text{Rev}_i(\hat{S}_i)]}{\sum_{i=1}^m [V(S_i)^{\tau_i} - V(\hat{S}_i)^{\tau_i}]}$, which gives the desired result.

Let $S_i = [1, p_i]$, $A_i = [k_i, p_i] \neq \emptyset$, and $A_\ell = \emptyset$ for all $\ell \neq i$. As $f_i \left(1 - \frac{V(A_i)}{V(S_i)} \right) = \frac{V(S_i)^{\tau_i} - V(\hat{S}_i)^{\tau_i}}{V(\hat{S}_i)^{\tau_i} - V(S_i)^{\tau_i}}$ and $\text{Rev}_i(\hat{S}_i) = \frac{V(S_i)}{V(\hat{S}_i)} \text{Rev}_i(S_i) + \left(1 - \frac{V(S_i)}{V(\hat{S}_i)} \right) \text{Rev}_i(A_i)$, it is easy to verify that

$$V(S_i)^{\tau_i} \text{Rev}_i(S_i) - V(\hat{S}_i)^{\tau_i} \text{Rev}_i(\hat{S}_i) = \left(V(S_i)^{\tau_i} - V(\hat{S}_i)^{\tau_i} \right) G_i(k_i, p_i).$$

Therefore, $\frac{V(S_i)^{\tau_i} \text{Rev}_i(S_i) - V(\hat{S}_i)^{\tau_i} \text{Rev}_i(\hat{S}_i)}{V(S_i)^{\tau_i} - V(\hat{S}_i)^{\tau_i}} \leq \text{Rev}(\mathbf{S})$ if and only if $G_i(k_i, p_i) \leq \text{Rev}(\mathbf{S})$. The inequality is strict if and only if $G_i(k_i, p_i) < \text{Rev}(\mathbf{S})$. $\square$

Lemma 2.2 is consistent with our interpretation of $G_i(\cdot, \cdot)$: for $p_i \geq 1$, if $S_i = [1, p_i]$ and $G_i(k_i, p_i) < \text{Rev}(\mathbf{S})$, then removing $[k_1, k_2]$ from S_i is beneficial. If $S_i^* = [1, p_i]$ in an optimal assortment $\mathbf{S}^*$, then $G_i(k_i, p_i) \geq \text{Rev}(\mathbf{S}^*)$ should hold for all $k_i \leq p_i$. The intuition is confirmed by the main result of this section, which is stated in the following theorem. The theorem establishes a necessary and sufficient condition for an optimal assortment, and it forms the basis for our greedy algorithm in Section 3.

Theorem 2.3 (Optimality Condition). *Consider any assortment $\mathbf{S} = ([1, p_1], \dots, [1, p_m])$ such that $S_i^* \subseteq [1, p_i]$ for all i . Then, $\mathbf{S}$ is optimal if and only if for every nest i such that $p_i \geq 1$,*

$$\min_{j=1, \dots, p_i} G_i(j, p_i) \geq \text{Rev}(\mathbf{S}).$$

Proof. Suppose that $\mathbf{S} = (S_1, \dots, S_m)$ is an optimal solution, where $S_i = [1, p_i]$ for all i . If $p_i \geq 1$, then $S_i \neq \emptyset$. As $\mathbf{S}$ is optimal, removing a subset $[k_i, p_i]$ from S_i cannot improve the revenue. Thus, it follows from Lemma 2.2 that $G_i(k_i, p_i) \geq \text{Rev}(\mathbf{S})$ for all $k_i \in \{1, 2, \dots, p_i\}$, which is the desired result.

To establish sufficiency, consider any assortment $\mathbf{S} = (S_1, \dots, S_m)$ such that $S_i^* \subseteq S_i = [1, p_i]$ for all i , and $\mathbf{S}$ satisfies the condition of the theorem. We will show that $\text{Rev}(\mathbf{S}) = Z^*$. For each i , let $E_i = S_i \setminus S_i^* = [k_i, p_i]$ for some index k_i . By our hypothesis, for every nest i such that $E_i \neq \emptyset$, $G_i(k_i, p_i) \geq \text{Rev}(\mathbf{S})$, thus removing E_i from S_i does not increase the revenue. It then follows from Lemma 2.2 that $\text{Rev}(\mathbf{S}) \leq \frac{V(S_i)^{\tau_i} \text{Rev}_i(S_i) - V(S_i \setminus E_i)^{\tau_i} \text{Rev}_i(S_i \setminus E_i)}{V(S_i)^{\tau_i} - V(S_i \setminus E_i)^{\tau_i}}$, where for any set X , $V(X)$ is the sum of the preference weights of the products in X . Thus,

$$\begin{aligned} \text{Rev}(\mathbf{S}) &\leq \min_{i: E_i \neq \emptyset} \frac{V(S_i)^{\tau_i} \text{Rev}_i(S_i) - V(S_i \setminus E_i)^{\tau_i} \text{Rev}_i(S_i \setminus E_i)}{V(S_i)^{\tau_i} - V(S_i \setminus E_i)^{\tau_i}} \\ &\leq \frac{\sum_{i: E_i \neq \emptyset} [V(S_i)^{\tau_i} \text{Rev}_i(S_i) - V(S_i \setminus E_i)^{\tau_i} \text{Rev}_i(S_i \setminus E_i)]}{\sum_{i: E_i \neq \emptyset} [V(S_i)^{\tau_i} - V(S_i \setminus E_i)^{\tau_i}]}, \end{aligned}$$

where the second inequality follows because for any $\mathbf{x} \in \mathbb{R}^k$ and $\mathbf{y} \in \mathbb{R}_+^k$, $\frac{\sum_{i=1}^k x_i}{\sum_{i=1}^k y_i} \geq \min_{i=1, \dots, k} \frac{x_i}{y_i}$, and for all i , $V(S_i)^{\tau_i} - V(S_i \setminus E_i)^{\tau_i} \geq 0$. Note that if $E_i = \emptyset$, then $S_i = S_i^*$. It then follows from Lemma 2.2 that $\text{Rev}(\mathbf{S}) \geq \text{Rev}(S_1 \setminus E_1, S_2 \setminus E_2, \dots, S_m \setminus E_m) = \text{Rev}(S_1^*, \dots, S_m^*) = Z^*$, which is the desired result. $\square$

Surprisingly, as shown in the following theorem, the optimal assortment in each nest is “lumpy,” with certain consecutive products always appearing together. Moreover, for each nest i , by looking at the index $G_i(j, j)$ of each product j , we can determine *in advance* which products will appear together!

Theorem 2.4 (Lumpiness of the Optimal Assortments). *For every nest i , if there exist products j and k , such that $j < k$ and*

$$G_i(j, j) < G_i(j+1, j+1) < \dots < G_i(k, k),$$

then either $\{j, j+1, \dots, k\} \subseteq S_i^$ or $\{j, j+1, \dots, k\} \cap S_i^* = \emptyset$.*

Proof. It suffices to prove the result when $k = j+1$; that is, we need to show that if $G_i(j, j) < G_i(j+1, j+1)$, then either $\{j, j+1\} \subseteq S_i^*$ or $\{j, j+1\} \cap S_i^* = \emptyset$. There are two cases to consider.

Case 1: $j+1 \in S_i^*$. In this case, $\{j, j+1\} \subseteq S_i^*$ because $S_i^* \in \mathcal{N}_+$ by Lemma 2.1.

Case 2: $j+1 \notin S_i^*$. Suppose that $j \in S_i^*$. Then, $S_i^* = [1, j]$. Let $\hat{\mathbf{S}} = (S_1^*, \dots, [1, j+1], \dots, S_m^*)$. By definition, $\text{Rev}(\hat{\mathbf{S}}) \leq \text{Rev}(\mathbf{S}^*)$, and as $\mathbf{S}^*$ is obtained from $\hat{\mathbf{S}}$ by removing product $j+1$ from nest i , it follows from Lemma 2.2 that $G_i(j+1, j+1) \leq \text{Rev}(\hat{\mathbf{S}})$. So, we have that $G_i(j, j) < G_i(j+1, j+1) \leq \text{Rev}(\hat{\mathbf{S}}) \leq \text{Rev}(\mathbf{S}^*)$, which implies that $\text{Rev}(S_1^*, \dots, S_i^* \setminus \{j\}, \dots, S_m^*) > \text{Rev}(\mathbf{S}^*)$. This contradicts the optimality of $\mathbf{S}^*$. Thus, $j \notin S_i^*$. $\square$

Note that the “lumpiness” property is nest-specific and is independent of v_0 . We can view the occurrence of “lumpiness” as the consequence of the property of $G_i(\cdot, \cdot)$ and the revenue-ordered property of the optimal assortment. Suppose $S_i \supseteq S_i^*$, and $G_i(j, j) < G_i(j+1, j+1)$ for products j and $j+1$ in S_i . If $G_i(j+1, j+1) < \text{Rev}(\mathbf{S})$, then it is beneficial to remove both products j and $j+1$ from nest i . Otherwise, the revenue-ordered property implies that if $j+1 \in S_i^*$, then $j \in S_i^*$. Thus, products j and $j+1$ always appear together in the optimal solution; that is, they can be merged into a single product when computing the optimal assortment, as illustrated in the next section.

3. A Greedy Algorithm

Our proposed GREEDY ALGORITHM generates a sequence of assortments $\{\mathcal{S}^t : t = 0, 1, \dots\}$, terminating with an assortment that satisfies the optimality condition in Theorem 2.3. In Stage 1 of the algorithm, we exploit the lumpy structure of the optimal assortments shown in Theorem 2.4, by combining products that will always appear together into a single group. In Stage 2, we make use of the optimality condition and iteratively remove at most one product from each nest until an optimal assortment is obtained. A formal description of the algorithm is given as follows.

Stage 1 (Lumping): Compute the index $G_i(j, j)$ for every product ji . For $i = 1, \dots, m$, if $G_i(j, j) < G_i(j+1, j+1)$, then it follows from Theorem 2.4 that either $\{j, j+1\} \subseteq S_i^*$ or $\{j, j+1\} \cap S_i^* = \emptyset$. Thus, we can combine the two products, and replace products j and $j+1$ with a single “new” product with a revenue $\text{Rev}_i([j, j+1])$ and a preference weight $v_{ji} + v_{j+1, i}$. Assign the new product the index j , and calculate the new $G_i(j, j)$. Repeat this process until we obtain a list of indices $G_i(j, j)$ that is non-increasing in j . Without loss of generality, assume that at the end of Stage 1, each nest i has n products such that $G_i(1, 1) \geq G_i(2, 2) \geq \dots \geq G_i(n, n)$.

Stage 2 (Removal): Let $\mathcal{S}^1 = ([1, n], \dots, [1, n])$. For iteration $t \geq 1$, given $\mathcal{S}^t = ([1, J_1^t], [1, J_2^t], \dots, [1, J_m^t])$, let

$$\text{Min}^t = \min_{i: J_i^t \geq 1} G_i(J_i^t, J_i^t) \quad \text{and} \quad \text{Index}^t = \arg \min_{i: J_i^t \geq 1} G_i(J_i^t, J_i^t).$$

If $\text{Min}^t \geq \text{Rev}(\mathcal{S}^t)$, the algorithm terminates and outputs $\mathcal{S}^t$. Otherwise, if $\text{Min}^t < \text{Rev}(\mathcal{S}^t)$, then the algorithm generates a new assortment $\mathcal{S}^{t+1} = (S_1^{t+1}, \dots, S_m^{t+1})$ as follows:

$$S_i^{t+1} = \begin{cases} S_i^t & \text{if } i \neq \text{Index}^t \\ S_i^t \setminus \{J_i^t\} = [1, J_i^t - 1] & \text{if } i = \text{Index}^t. \end{cases}$$

Thus, in each iteration, we remove J_i^t from S_i^t if $G_i(J_i^t, J_i^t) = \text{Min}^t$ **and** the product J_i^t violates the optimality condition given in Theorem 2.3; that is, $G_i(J_i^t, J_i^t) < \text{Rev}(\mathcal{S}^t)$.

The following lemma shows that $\mathcal{S}^t$ always contains the optimal assortment.

Lemma 3.1 (Containment). *For all t , $\mathcal{S}^* \subseteq \mathcal{S}^t$.*

Proof. We will prove the result by induction. It is true for $t = 1$ by our construction. Suppose that the lemma is true for $t > 1$; that is, $\mathcal{S}^* \subseteq \mathcal{S}^t$. We will show that $S_i^* \subseteq S_i^{t+1}$ for all i . Consider an arbitrary nest i . There are two cases to consider: $S_i^* = S_i^t$ and $S_i^* \subset S_i^t$.

Case 1: Suppose that $S_i^* = S_i^t = [1, J_i^t]$. If $J_i^t = 0$, then $i \neq \text{Index}^t$ and $S_i^{t+1} = S_i^t$ by our construction. So, suppose that $J_i^t \geq 1$. As $\text{Rev}(\mathcal{S}^*) \geq \text{Rev}(S_1^*, \dots, S_i^* \setminus \{J_i^t\}, \dots, S_m^*)$, $G_i(J_i^t, J_i^t) \geq \text{Rev}(\mathcal{S}^*) \geq \text{Rev}(\mathcal{S}^t)$ follows from Lemma 2.2 and the optimality of $\mathcal{S}^*$. As $G_i(J_i^t, J_i^t) \geq \text{Rev}(\mathcal{S}^t)$, $S_i^{t+1} = S_i^t$, which is the desired result.

Case 2: Suppose that $S_i^* \subset S_i^t = [1, J_i^t]$. Then, $J_i^t \notin S_i^*$ and $J_i^t \geq 1$. As $S_i^* \in \mathcal{N}_+$, $S_i^* \subseteq [1, J_i^t - 1]$. By our construction, we will either remove J_i^t from S_i^t or do nothing, and in both cases, we have $S_i^* \subseteq S_i^{t+1}$. $\square$

The main result of this section is stated in the following theorem.

Theorem 3.2 (Correctness). *The GREEDY ALGORITHM terminates with an optimal assortment in $O(nm \log m)$ time.*

Proof. We will first establish the running time of the algorithm. Note that $\text{Rev}_i([1, j])$ is a convex combination of $\text{Rev}_i([1, j-1])$ and r_{ji} , and by definition, $G_i(j, j) = \text{Rev}_i([1, j]) - \frac{\text{Rev}_i([1, j]) - r_{ji}}{f_i(1 - v_{ji}/\sum_{\ell=1}^j v_{\ell i})}$. For simplicity, assume that we can compute the function $f_i(x)$ for each x and the index $G_i(j, j)$ for every product ji in $O(1)$ time. For each nest i , we will show that the “lumping” process takes $O(n)$ operations. The lumping process requires two types of operations: 1) comparison between two indices and 2) merging of two products. Whenever we merge two products to create a new one, the total number of products is reduced by one. As we start with n products, the number of mergings is at most $n - 1$. We perform two types of comparisons. Starting from product 1, if the ordering of indices is correct up to product j , we compare $G_i(j, j)$ *forward* with $G_i(j+1, j+1)$. If no violation occurs, then $G_i(j, j) \geq G_i(j+1, j+1)$, and we proceed to product $j+1$. If a violation occurs, then we must merge products j and $j+1$ to create a new product j . To ensure the correct ordering of indices, we need to compare the new index $G_i(j, j)$ *backward* with $G_i(j-1, j-1)$. Given n products, there are $O(n)$ forward comparisons. As a backward

comparison only occurs after each merging, and there are $O(n)$ mergings, the number of backward comparisons is also $O(n)$. Thus, the total number of operations is $O(n)$. With m nests, the total running time for Stage 1 is $O(nm)$.

In Stage 2, the algorithm will continue to run as long as a single product is removed from any nest. Given nm products, the algorithm will terminate in $O(nm)$ iterations. We will show that each iteration takes $O(\log m)$ time.

At the beginning of Stage 2, we create a *self-balancing binary search tree* (SB-BST) with m nodes, where for $i = 1, 2, \dots, m$, node i in the tree corresponds to the index $G_i(n, n)$. This takes $O(m \log m)$ operations; see Chapter 6 in [17] for more details. We use the SB-BST as our data structure because such a tree always maintains a height of $O(\log m)$, allowing for efficient search operations. In each iteration $t \geq 1$, we perform the following three operations:

SEARCH: The algorithm searches for the nest with the minimum index $G_i(J_i^t, J_i^t)$ in $O(\log m)$ time under SB-BST.

DELETE: If the minimum index $G_i(J_i^t, J_i^t)$ is greater than or equal to the revenue of the current assortment, the algorithm terminates. Otherwise, the algorithm removes the product J_i^t from S_i^t in nest i and also removes node i from the tree. The removal is done in $O(\log m)$ operations because the tree may need to re-balance its heights.

INSERT: If the product J_i^t in nest i is removed and $J_i^t > 1$, the algorithm adds a new node with a corresponding index $G_i(J_i^t - 1, J_i^t - 1)$ into the tree. Again, the insertion of a new node in SB-BST takes $O(\log m)$ operations.

Thus, each iteration takes $O(\log m)$ time. As we have $O(nm)$ iterations, the total running time is $O(nm \log m)$.

Let $\mathbf{S} = ([1, p_1], \dots, [1, p_m])$ denote the assortment at the termination of the algorithm. For every nest i such that $p_i \geq 1$, by our construction, $G_i(p_i, p_i) \geq \text{Rev}(\mathbf{S})$. To complete the proof, we will show that $\mathbf{S}$ satisfies the optimality condition in Theorem 2.3; that is, $\min_{j=1,2,\dots,p_i} G_i(j, p_i) \geq \text{Rev}(\mathbf{S})$ for each i such that $p_i \geq 1$. Suppose on the contrary that $\min_{j=1,2,\dots,p_i} G_i(j, p_i) < \text{Rev}(\mathbf{S})$. Let $k_i \in \{1, 2, \dots, p_i - 1\}$ denote the largest index at which the optimality condition is violated; that is, $G_i(k_i, p_i) < \text{Rev}(\mathbf{S}) \leq G_i(k_i + 1, p_i)$. Then it follows from Lemma 2.2 that $\text{Rev}(S_1, \dots, S_i \setminus [k_i, p_i], \dots, S_m) > \text{Rev}(\mathbf{S}) \geq \text{Rev}(S_1, \dots, S_i \setminus [k_i + 1, p_i], \dots, S_m)$. As $S_i \setminus [k_i, p_i] = [1, k_i - 1]$ and $S_i \setminus [k_i + 1, p_i] = [1, k_i]$, by Lemma 2.2, $G_i(k_i, k_i) < \text{Rev}(S_1, \dots, S_i \setminus [k_i + 1, p_i], \dots, S_m) \leq \text{Rev}(\mathbf{S})$. As we lump the products in Stage 1 and $k_i < p_i$, we have $G_i(p_i, p_i) \leq G_i(k_i, k_i) < \text{Rev}(\mathbf{S})$, which contradicts our hypothesis on p_i ! Therefore, $\min_{j=1,2,\dots,p_i} G_i(j, p_i) \geq \text{Rev}(\mathbf{S})$. $\square$

4. Sensitivity Analysis

In this section, we investigate how the optimal assortment changes with product revenues. Our goal is to tractably compute the optimal assortment that maximizes the total revenue when the revenue of every product is reduced by δ , for all $\delta \in \mathbb{R}_+$ simultaneously. That is, we want to solve the following optimization problem:

$$Z(\delta) = \max_{\mathbf{S}=(S_1,\dots,S_m)} \left\{ \sum_{i=1}^m \sum_{j \in S_i} Q_{ji}(\mathbf{S})(r_{ji} - \delta) \right\} = \max_{\mathbf{S}=(S_1,\dots,S_m)} \left\{ \text{Rev}(\mathbf{S}) - [1 - Q_0(\mathbf{S})]\delta \right\},$$

for all $\delta \in \mathbb{R}_+$, where $Q_0(\mathbf{S})$ is the probability that a customer chooses the no-purchase option.

Before we present the main result of this section, we first describe the structural properties of $Z(\delta)$ in the following lemma. The result is standard because $\{Z(\delta) : \delta \in \mathbb{R}_+\}$ is the maximum of linear functions, each of which is decreasing in δ , and we omit the detail of the proof.

Lemma 4.1. *For all $\delta \in \mathbb{R}_+$, the function $\delta \mapsto Z(\delta)$ is convex, decreasing and piecewise linear.*

Let $\mathbf{X}(\delta)$ represent the corresponding optimal assortment associated with the optimization problem for $Z(\delta)$. Throughout this section, we assume that $\mathbf{X}(\delta)$ corresponds to the optimal assortment obtained by running the GREEDY ALGORITHM with reduced revenue $r_{ji} - \delta$ for every product ji . Clearly, a naive application of the GREEDY ALGORITHM to compute $Z(\delta)$ for every single $\delta \in \mathbb{R}_+$ is infeasible. We will make use of the structural properties of $Z(\delta)$ to develop an efficient algorithm that computes the entire profile of $Z(\delta)$ and $\mathbf{X}(\delta)$, simultaneously for all $\delta \in \mathbb{R}_+$, with a running time of $O(nm \log m)$. The main result of this section is stated in the following proposition.

Proposition 4.2 (Computing $Z(\delta)$). *For all $\delta \in \mathbb{R}_+$, the function $\delta \mapsto Z(\delta)$ has at most nm breakpoints. Moreover, the entire profile of $Z(\delta)$, $\mathbf{X}(\delta)$ and the breakpoints can be computed simultaneously for all $\delta \in \mathbb{R}_+$ in a **single run** of the GREEDY ALGORITHM with a running time of $O(nm \log m)$.*

Before we prove Proposition 4.2, it is easy to verify that if we reduce the revenue of every product by δ , the index of each product ji is given by $G_i(j, j) - \delta$, which is simply the difference between the original index $G_i(j, j)$ defined in Section 2 and δ . Moreover, both the sequence of lumping and the ordering of the index of each product after lumping remain the same as we apply the GREEDY ALGORITHM with reduced revenue $r_{ji} - \delta$ for every product ji .

The following lemma shows the relationship between $\mathbf{X}(\delta)$ and $\mathbf{S}^*$, where $\mathbf{S}^*$ is the optimal assortment associated with the assortment optimization problem for Z^* as defined in Section 1.

Lemma 4.3 (Optimal Assortment for $Z(\delta)$). *For all $\delta \geq 0$, $\mathbf{X}(\delta) \subseteq \mathbf{X}(0) = \mathbf{S}^*$.*

Proof. The optimization problem for $Z(0)$ is identical to the one for Z^* ; therefore, $\mathbf{X}(0) = \mathbf{S}^*$. Suppose the GREEDY ALGORITHM that computes $\mathbf{S}^*$ terminates in T iterations. For each $t = 1, \dots, T$, let $\mathbf{S}^t = ([1, J_1^t], [1, J_2^t], \dots, [1, J_m^t])$ denote the assortment in iteration t . Then, the optimality condition implies that $\min_{i: J_i^t \geq 1} G_i(J_i^t, J_i^t) < \text{Rev}(\mathbf{S}^t)$, for $t < T$. Now compare this algorithm with the GREEDY ALGORITHM that computes $\mathbf{X}(\delta)$ with reduced revenue $r_{ji} - \delta$ for every product ji . As the ordering of the index of each product is constant for all $\delta \geq 0$, the sequence of products removed remains the same before the termination of either algorithm. Moreover, in each iteration $t < T$,

$$\min_{i: J_i^t \geq 1} (G_i(J_i^t, J_i^t) - \delta) < \text{Rev}(\mathbf{S}^t) - \delta < \text{Rev}(\mathbf{S}^t) - [1 - Q_0(\mathbf{S}^t)]\delta,$$

implying that the GREEDY ALGORITHM that computes $\mathbf{X}(\delta)$ terminates in iteration T or later. Thus, $\mathbf{X}(\delta) \subseteq \mathbf{S}^*$. $\square$

We denote the breakpoints of the function $\delta \mapsto Z(\delta)$ by $\{\delta^t : t = 0, 1, \dots\}$, $0 = \delta^0 \leq \delta^1 \leq \dots$ and denote the optimal assortment for $\delta \in [\delta^t, \delta^{t+1}]$ by $\mathbf{X}^t = \mathbf{X}(\delta)$. Recall that the optimal assortment $\mathbf{X}^t$ for $\delta \in [\delta^t, \delta^{t+1}]$ corresponds to the optimal assortment obtained by running the GREEDY ALGORITHM with reduced revenue $r_{ji} - \delta$ for every product ji . Thus, every product in $\mathbf{X}^t$ has undergone the lumping process, and the index $G_i(j, j) - \delta$ in every nest i is non-increasing in j . Exploiting this fact, the next lemma shows that given $\mathbf{X}^{t-1}$, we can compute $\mathbf{X}^t$ simply by removing a product with the lowest index.

Lemma 4.4 (Computing an Optimal Assortment). *Suppose $\mathbf{X}^{t-1}$ is optimal for $\delta \in [\delta^{t-1}, \delta^t]$; an optimal assortment $\mathbf{X}^t$ for $\delta \in [\delta^t, \delta^{t+1}]$ can be obtained by removing a product with the minimum index from $\mathbf{X}^{t-1}$.*

Proof. As δ^t is a breakpoint, $\mathbf{X}^t \neq \mathbf{X}^{t-1}$. Thus by Lemma 4.3, $\mathbf{X}^t \subset \mathbf{X}^{t-1}$. Suppose the last product k in nest ℓ has the minimum index value in $\mathbf{X}^{t-1}$. By continuity, $Z(\delta^t) = \text{Rev}(\mathbf{X}^{t-1}) - [1 - Q_0(\mathbf{X}^{t-1})]\delta^t = G_\ell(k, k) - \delta^t$, where the second equality follows from the optimality condition. As $G_\ell(k, k) - \delta^t$ is equal to $Z(\delta^t)$, which is the total revenue of $\mathbf{X}^{t-1}$ with reduced revenue $r_{ji} - \delta^t$ for every product ji , it follows from Lemma 2.2 that the removal of product $k\ell$ leaves the total revenue unchanged. That is, $Z(\delta^t) = \text{Rev}(\mathbf{X}^{t-1} \setminus \{k\ell\}) - [1 - Q_0(\mathbf{X}^{t-1} \setminus \{k\ell\})]\delta^t$. Thus, $\mathbf{X}^{t-1} \setminus \{k\ell\}$ is also optimal at δ^t . Letting $\mathbf{X}^t = \mathbf{X}^{t-1} \setminus \{k\ell\}$ gives the desired result. $\square$

Lemma 4.4 implies that starting with an optimal assortment $\mathbf{X}^0 = \mathbf{S}^*$, we can compute $\mathbf{X}^t$, for $t = 1, 2, \dots$ sequentially *without* any knowledge on the breakpoints δ^t . The following lemma shows that given two consecutive optimal assortments $\mathbf{X}^{t-1}$ and $\mathbf{X}^t$, we can compute the breakpoint δ^t .

Lemma 4.5 (Computing Breakpoints). *Suppose $\mathbf{X}^{t-1}$ is an optimal assortment for $\delta \in [\delta^{t-1}, \delta^t]$, and $\mathbf{X}^t$ is an optimal assortment for $\delta \in [\delta^t, \delta^{t+1}]$. Then,*

$$\delta^t = \frac{\text{Rev}(\mathbf{X}^{t-1}) - \text{Rev}(\mathbf{X}^t)}{Q_0(\mathbf{X}^t) - Q_0(\mathbf{X}^{t-1})}.$$

Proof. By continuity, we have that $Z(\delta^t) = \text{Rev}(\mathbf{X}^{t-1}) - [1 - Q_0(\mathbf{X}^{t-1})]\delta^t = \text{Rev}(\mathbf{X}^t) - [1 - Q_0(\mathbf{X}^t)]\delta^t$, as both $\mathbf{X}^{t-1}$ and $\mathbf{X}^t$ give the same revenue $Z(\delta^t)$ at the breakpoint δ^t . Solving for δ^t gives us the desired result. $\square$

Here is the proof of Proposition 4.2.

Proof. Starting with the full assortment, we first apply the GREEDY ALGORITHM to compute $\mathbf{X}^0 = \mathbf{X}(0) = \mathbf{S}^*$. Given $\mathbf{X}^0$, we obtain $\mathbf{X}^t$ for each $t \in \{1, 2, \dots\}$ by applying Lemma 4.4. Then, given $\mathbf{X}^{t-1}$ and $\mathbf{X}^t$, we apply Lemma 4.5 to calculate the breakpoint δ^t . Finally, given $\mathbf{X}^{t-1}$ for $\delta \in [\delta^{t-1}, \delta^t]$, the optimal revenue is computed as $Z(\delta) = \text{Rev}(\mathbf{X}^{t-1}) - [1 - Q_0(\mathbf{X}^{t-1})]\delta$. There are at most nm products in $\mathbf{X}^0$. As we remove a single product with the minimum index in each iteration, there are at most nm removals, corresponding to at most nm breakpoints of the function $\delta \mapsto Z(\delta)$. The running time of the algorithm is $O(nm \log m)$ by Theorem 3.2. $\square$

Application to the Multi-Period Capacity Allocation Problem. We apply our GREEDY ALGORITHM to solve the multi-period capacity allocation problem described in [2] under the two-level nested logit model. We denote the total capacity of seats on a flight leg by C and the total number of allocation periods by T . There are m groups of fare classes, with each group containing n fare classes. Each fare class j in group i has revenue r_{ji} , where $j = 1, 2, \dots, n$ and $i = 1, 2, \dots, m$. Let $\mathbf{S} = (S_1, S_2, \dots, S_m)$ denote the assortment of fare classes. In each period, a single customer chooses a fare class according to a two-level nested logit model. Let $J_t(x)$ denote the maximum expected revenue when we have x seats and t periods left. Then, $J_t(x)$ satisfies the following dynamic programming equation:

$$J_t(x) = \max_{\mathbf{S}=(S_1, \dots, S_m)} \left\{ \sum_{i=1}^m \sum_{j \in S_i} Q_{ji}(\mathbf{S})(r_{ji} - \Delta J_{t-1}(x)) \right\} + J_{t-1}(x),$$

where $\Delta J_{t-1}(x) = J_{t-1}(x) - J_{t-1}(x-1)$, and $J_0 \equiv 0$. It is a standard result that $\Delta J_t(x) \geq 0$ for all x and t . Then, the first term on the right-hand side is exactly the optimization problem for $Z(\delta)$ with $\delta = \Delta J_{t-1}(x)$.

The traditional approach to solve the multi-period capacity allocation problem is by dynamic programming, in which $J_t(x)$ is computed for every single x and t , requiring a total running time of $O(TCnm \log m)$ if we apply the GREEDY ALGORITHM in each subproblem. As shown in the following proposition, we can do much better.

Proposition 4.6 (Solving Multi-Period Capacity Allocation). *Under the two-level nested logit model, the multi-period capacity allocation problem with total allocation periods T and total capacity C can be solved using a single run of the GREEDY ALGORITHM with a total running time of $O(TC + nm \log m)$.*

Proof. Computing $\{\mathbf{X}(\delta) : \delta \in \mathbb{R}_+\}$ and breakpoints $\{\delta^t : t = 1, 2, \dots\}$ takes $O(nm \log m)$ time by Proposition 4.2. We can then build a hash table to store the intervals of breakpoints and the corresponding optimal assortments in $O(nm)$ time. Suppose that for every $\delta \in \mathbb{R}_+$, we can look up in the hash table for $\mathbf{X}(\delta)$ in $O(1)$ time. As x can take $C+1$ values, for each t , the value function $J_t(\cdot)$ can be computed in $O(C)$ operations, simply by looking up the stored values in the hash table. There are T such value functions; thus it takes $O(TC)$ time to determine $J_t(\cdot)$ for all $t \in \{1, \dots, T\}$. Therefore, the total running time is bounded by $O(TC + nm \log m)$. $\square$

Acknowledgments

We would like to thank Huseyin Topaloglu for introducing us to the surprising connection between the assortment optimization problem under the two-level nested logit model and its equivalent linear programming formulation. We would like to thank the area editor, Sridhar Seshadri, the associate editor and the referee for their insightful and detailed comments. Their suggestions greatly improve the presentation and quality of our paper.

References

- [1] A. G. K  k, M. L. Fisher, R. Vaidyanathan, Assortment planning: Review of literature and industry practice, in: Retail Supply Chain Management, Springer US, 2009, pp. 99–153.
- [2] K. Talluri, G. van Ryzin, Revenue management under a general discrete choice model of consumer behavior, Mgmt. Sci. 50 (1) (2004) 15–33.
- [3] D. McFadden, Conditional logit analysis of qualitative choice behavior, in: P. Zarembka (Ed.), Frontiers in Econometrics, Academic Press, 1974, pp. 105–142.
- [4] K. Train, Discrete Choice Methods with Simulation, Cambridge University Press, 2009.
- [5] V. Gaur, D. Honhon, Assortment planning and inventory decisions under a locational choice model, Mgmt. Sci. 52 (10) (2006) 1528–1543.
- [6] D. Honhon, V. Gaur, S. Seshadri, Assortment planning and inventory decisions under stockout-based substitution, Oper. Res. 58 (5) (2010) 1364–1379.
- [7] V. F. Farias, S. Jagabathula, D. Shah, Assortment optimization under general choice, Working Paper, MIT, Cambridge, MA (2011).
- [8] D. McFadden, Modelling the choice of residential location, in: A. Karlqvist, L. Lundqvist, F. Snickars, J. Weibull (Eds.), Spatial Interaction Theory and Planning Models, North-Holland, Amsterdam, 1978, pp. 75–96.
- [9] J. M. Davis, G. Gallego, H. Topaloglu, Assortment optimization under variants of the nested logit model, to appear in Oper. Res. (2014).
- [10] P. Rusmevichientong, Z.-J. M. Shen, D. B. Shmoys, A PTAS for capacitated sum-of-ratios optimization, OR Letters 37 (4) (2009) 230–238.
- [11] G. Gallego, H. Topaloglu, Constrained assortment optimization for the nested logit model, to appear in Mgmt. Sci. (2014).
- [12] G. Gallego, G. Iyengar, R. Phillips, A. Dubey, Managing flexible products on a network, Computational Optimization Research Center Technical Report TR-2004-01, Columbia University, New York, NY (2004).
- [13] Q. Liu, G. van Ryzin, On the choice-based linear programming model for network revenue management, M&SOM 10 (2) (2008) 288–310.
- [14] S. Kunnumkal, Randomization approaches for network revenue management with customer choice behavior, to appear in POM (2013).
- [15] N. S. Cardell, Variance components structures for the extreme-value and logistic distributions with application to models of heterogeneity, Econometric Theory 13 (02) (1997) 185–213.
- [16] M. Ben-Akiva, S. Lerman, Discrete Choice Analysis: Theory and Application to Travel Demand, MIT Press, Cambridge, MA, 1985.
- [17] D. E. Knuth, The Art of Computer Programming, 2nd Edition, Vol. 3, Addison-Wesley Professional, 1998.