

Assortment Optimization under Variants of the Nested Logit Model

James M. Davis

School of Operations Research and Information Engineering,
Cornell University, Ithaca, New York 14853, USA
jmd388@cornell.edu

Guillermo Gallego

Department of Industrial Engineering and Operations Research
Columbia University, New York, New York 10027, USA
gmg2@columbia.edu

Huseyin Topaloglu

School of Operations Research and Information Engineering,
Cornell University, Ithaca, New York 14853, USA
topaloglu@orie.cornell.edu

December 30, 2011

Abstract

We study a class of assortment optimization problems where customers choose among the offered products according to the nested logit model. There is a fixed revenue associated with each product. The objective is to find an assortment of products to offer so as to maximize the expected revenue per customer. We show that the problem is polynomially solvable when the nest dissimilarity parameters of the choice model are between zero and one, and the customers always make a purchase within the selected nest. Relaxing either of these assumptions renders the problem NP-hard. To deal with the NP-hard cases, we develop parsimonious collections of candidate assortments with worst-case performance guarantees. We also formulate a tractable convex program whose optimal objective value is an upper bound on the optimal expected revenue. Thus, we can compare the expected revenue provided by an assortment with the upper bound on the optimal expected revenue to get a feel for the optimality gap of the assortment. By using this approach, our computational experiments demonstrate that the parsimonious collections of assortments we propose yield solutions with optimality gaps less than a fraction of a percent.

Discrete choice models have been used for nearly half a century to understand how customers select among a group of products that vary in terms of price and quality. Of particular interest is how demand for the different products changes as the offer set changes in composition, quality or price. To advance this agenda, researchers have developed discrete choice models based on axioms as in Luce (1959), resulting in the basic attraction model, and based on random utility theory as in McFadden (1974), resulting in the celebrated multinomial logit model. Important extensions include the nested attraction model of which the nested logit model, introduced by Williams (1977), is a special case. Justifications and extensions for the nested logit model are provided in McFadden (1980) and Borsch-Supan (1990). Under the nested logit model, customers first select a nest, and then, a product within the selected nest. The nested logit model was developed primarily to avoid the independence of irrelevant alternatives property suffered by the multinomial logit model.

In this paper, we study a class of assortment optimization problems where the choices of the customers are governed by the nested logit model. Under this model, customers first select a nest, and then, a product within the nest. We assume that there is a fixed revenue associated with each product and that the objective is to find a set of products to offer, or an assortment, that maximizes the expected revenue per customer. This assortment problem is combinatorial in nature and the number of possible assortments can be very large, particularly when there are many potential products to select from in each nest. It is important, therefore, to classify when the problem is polynomially solvable. When not, it is important to find heuristics with worst-case performance guarantees. Our main contribution consists of classifying the complexity of the assortment problem for nested attraction models. We do this along two dimensions. The first dimension is the magnitude of the nest dissimilarity parameters, which characterize the degree of dissimilarity of the products within a nest. The second dimension is the presence or absence of the no purchase alternative within a nest. The only polynomially solvable case is when the nest dissimilarity parameters are all between zero and one, and the no purchase alternative is only available at the time of selecting a nest. If the nest dissimilarity parameters exceed one or the customers can choose a no purchase option after selecting a nest, then we show that the problem is NP-hard. For the NP-hard cases, we provide heuristics with worst-case performance guarantees. In a computational study, we compare the performance of our heuristics with a tractable upper bound on the optimal expected revenue that we develop, and demonstrate that the performance of our heuristics is generally within a fraction of a percent of the optimal.

Research on pricing in the context of the multinomial logit and nested logit models has been fairly active. In that setting, the problem is to choose a set of prices for the products, where the prices of all products jointly determine the probability that a customer purchases a particular product. The objective is to maximize the expected revenue per customer. For the pricing problem, Hanson and Martin (1996) notice that the expected revenue function fails to be concave in prices for the multinomial logit model, but significant progress was made by formulating the pricing problem in terms of market shares, as this results in a concave expected revenue function; see Song and Xue (2007) and Dong et al. (2009). Li and Huh (2011) extend the concavity result to the nested logit model by assuming that the price sensitivities of the products are constant within each nest and the nest dissimilarity parameters are all between zero and one. They show that the expected revenue maximization problem can be reduced to optimizing

over a single variable. Gallego and Wang (2011) relax both of the assumptions in Li and Huh (2011) and extend the analysis to more general nested attraction models. The key result is that the optimal prices add two terms to the unit costs, where the first term is the inverse of the price sensitivities of the products and the second is a nest-dependent constant. The nest-dependent constants are all linked to a single variable that is equal to the optimal expected revenue per customer. Interestingly, optimal prices do not depend directly on quality, implying that products of different quality in a nest that have the same price sensitivity have the same markup. This implies that, everything else being equal, higher quality products sell faster and products with higher price sensitivity sell slower at lower markups.

In many situations, prices of the products are fixed and are not in the control of the decision maker, at least not in the short run. This is true, for example, in the context of revenue management, where a menu of fares is designed to allow the same capacity to be sold at different prices. This is done by differentiating the products by time-of-purchase, traveling restrictions, and the inclusion or exclusion of ancillary services such as luggage handling, mileage accrual and advance seat selection. Revenue managers must dynamically decide which set of products to offer depending on the state of the system, which includes the time-to-departure and the remaining inventory. In retailing, pricing decisions are often centralized, and hence, are fixed in advance, while assortment decisions can be made at the local level. The assortment problem is particularly important at the design stage, where several products could be built based on different design features and prices.

Assortment optimization is an active area of research, and our literature review on assortment optimization focuses on papers that use attraction-based choice models, such as multinomial or nested logit model. We refer the reader to Kok et al. (2008) and Farias et al. (2011) for assortment optimization under other choice models. The paper by van Ryzin and Mahajan (1999) considers an assortment optimization problem where the products both generate revenue and result in operational cost. The objective is to balance the revenue benefit of offering product variety with the overhead cost of carrying a large number of products. Cachon et al. (2005) extend this work to include product search costs. In particular, the authors model the possibility that a customer may find an acceptable product in one store, but still not purchase hoping that another store may carry a more desirable product.

Recent research has demonstrated the importance of incorporating customer choice behavior in revenue management decisions. Talluri and van Ryzin (2004) consider a revenue management problem over a single flight leg. The customers choose among the fare classes that are available for purchase and the objective is to adjust the assortment of available fare classes at each time period to maximize the expected revenue. Gallego et al. (2004), Liu and van Ryzin (2008), Kunnumkal and Topaloglu (2008), Zhang and Adelman (2009) and Bront et al. (2009) extend the model in Talluri and van Ryzin (2004) to a flight network. The fundamental idea behind these papers is to construct various deterministic linear programming approximations. The decision variables in these linear programs correspond to the number of time periods during which a particular subset of origin-destination-fare-class combination is offered. Since there is one decision variable for every possible subset of origin-destination-fare-class combination, the number of decision variables can be quite large. Therefore, it is customary to solve the linear programs by using column generation. The column generation subproblem for these linear

programs precisely corresponds to the assortment optimization problem that we study in this paper whenever the choices of the customers are governed by the nested logit model.

If the customers choose according to the multinomial logit model, then the assortment optimization problem can be solved efficiently, as it can be shown that the optimal assortment includes a certain number of products with the largest revenues. We refer to assortments that include a certain number of products with the largest revenues as nested-by-revenue assortments. The problem becomes more complicated when more general choice models are considered. Rusmevichientong, Shmoys and Topaloglu (2010) study the assortment problem when there are multiple customer types and the customers of different types choose according to different multinomial logit models. They show that the assortment optimization problem is NP-hard even with two customer classes and provide a performance guarantee for nested-by-revenue assortments. Mendez-Diaz et al. (2010) give a branch-and-cut algorithm to find the optimal assortment for the same problem. Rusmevichientong and Topaloglu (2011) study the robust assortment problem when some of the parameters of the model are not known, while Rusmevichientong, Shen and Shmoys (2010) consider constraints on the size of the offered assortment.

There is no work, to our knowledge, on assortment optimization under the nested logit model, although there is a need to expand the scope of tractable solution methods beyond the multinomial logit model. This is particularly the case since the assortment problem we study not only naturally appears in many settings by itself, but it is also the column generation subproblem when we solve linear programming approximations for network revenue management problems. We consider four cases to characterize the situations where our assortment optimization problem can be solved exactly or approximately. The first case considers the situation where the dissimilarity parameters of the nests are between zero and one, and customers always make a purchase within the selected nest. This situation conforms to the standard form of the nested logit model; see Borsch-Supan (1990). For this case, we show that it is optimal to offer a nested-by-revenue assortment within each nest, but this result does not immediately imply that the problem is polynomially solvable since there are exponentially many combinations of nested-by-revenue assortments we can choose for the different nests. We deal with this difficulty by giving a linear program that finds the best combination of nested-by-revenue assortments for each nest. Thus, the problem is tractable under the standard form of the nested logit model.

We show that the assortment optimization problem is NP-hard in all of the remaining cases, but we are able to characterize parsimonious collections of assortments such that if we focus only on these assortments, then we obtain a solution with a certain worst-case performance guarantee. In particular, the second case we consider focuses on the situation where the dissimilarity parameters can take on any value, but the customers always purchase a product within the selected nest. For this case, we show that if we focus only on nested-by-revenue assortments, then we can construct a solution whose worst-case performance guarantee is, roughly speaking, good when the revenues or the attractiveness of the products within a particular nest are not too different from each other. We emphasize that this result allows the revenues or attractiveness of the products to differ arbitrarily when the products are in different nests. In the third case, we consider the situation where the dissimilarity parameters of the nests are between zero and one, but customers may leave a chosen nest without purchasing. We

		NO PURCHASE BEHAVIOR	
		Always Make a Purchase within Selected Nest	May Leave without Purchasing Even After Selecting a Nest
NEST DISSIMILARITY PARAMETERS	Between 0 and 1	Case 1: Section 3 Polynomially-solvable by offering nested-by-revenue assortments within each nest	Case 3: Section 5 Worst-case performance guarantee of two
	Free	Case 2: Section 4 Desirable worst-case performance guarantee when the revenues or attractiveness of the products within a nest are similar	Case 4: Section 6 Desirable worst-case performance guarantee when the dissimilarity parameters of the nests are small

Table 1: Summary of the four cases considered in the paper.

construct a small collection of assortments such that the best assortment within this collection provides an expected revenue that deviates from the optimal expected revenue by no more than a factor of two. Finally, the fourth case considers the most general problem instances with no restrictions on the dissimilarity parameters of the nests and the no purchase behavior. For this case, roughly speaking, we give a collection of assortments such that the best assortment within this collection has a desirable worst-case performance guarantee when the dissimilarity parameters are not too large. Table 1 gives a summary of the four cases and indicates which sections include each one of these cases.

In addition to the worst-case performance guarantees, we formulate a convex program that yields an upper bound on the optimal expected revenue. By comparing the upper bound on the optimal expected revenue with the expected revenue provided by an assortment, we bound the optimality gap of the assortment we obtain for a particular problem instance. We use this approach in our computational experiments to demonstrate that the solutions we obtain by focusing on the collections of assortments mentioned above perform extremely well, providing optimality gaps within a fraction of a percent.

The rest of the paper is organized as follows. In Section 1, we formulate the assortment optimization problem. In Section 2, we give an alternative linear programming formulation of this problem. The number of constraints in this formulation grows exponentially with the number of products within each nest, but we build on the linear program to come up with a general approximation framework for the assortment optimization problem. The following four sections focus on the four cases mentioned above. In Section 3, we assume that the dissimilarity parameters of the nests are between zero and one, and the customers always make a purchase among the products of the selected nest. We show that it is optimal to offer a nested-by-revenue assortment within each nest. In Section 4, we relax the assumption that the dissimilarity parameters are between zero and one, whereas in Section 5, we allow customers to leave a selected nest without purchasing. In these sections, we show that the assortment optimization problem becomes NP-hard under either of these relaxations and we provide worst-case performance guarantees for certain collections of assortments. In Section 6, we focus on the most general instances of the assortment optimization problem with no restrictions on the dissimilarity parameters of the nests and the no purchase behavior of the customers. Similar to the previous two sections, we give a collection of assortments with a certain worst-case performance guarantee. The performance guarantees we give in Sections 4, 5 and 6 are based on the general approximation framework we develop in Section 2. In Section 7, we formulate a convex program that provides an upper bound on the optimal expected revenue. In Section 8, we give our computational experiments. In Section 9, we conclude.

1 Problem Formulation

In this section, we describe the nested attraction model that we use to model customer choice, and then, formulate the assortment optimization problem. There are m nests indexed by $M = \{1, \dots, m\}$. Depending on the application setting, each nest may represent a different category of products, a different sales channel or a different retail store. There are n products that we can offer in each nest. We index the products in each nest by $N = \{1, \dots, n\}$. We use r_{ij} to denote the revenue associated with product j in nest i . We assume that the products in each nest are ordered such that $r_{i1} \geq r_{i2} \geq \dots \geq r_{in}$ for all $i \in M$. We let v_{ij} be the preference weight of product j in nest i and v_{i0} be the preference weight of the no purchase option in nest i . We use $V_i(S_i)$ to denote the total preference weight of all available options when we offer the assortment $S_i \subset N$ in nest i . In other words, we have $V_i(S_i) = v_{i0} + \sum_{j \in S_i} v_{ij}$. Given that a customer decides to purchase a product in nest i , if we offer the assortment S_i in this nest, then the probability that the customer purchases product $j \in S_i$ is given by

$$P_{ij}(S_i) = \frac{v_{ij}}{v_{i0} + \sum_{k \in S_i} v_{ik}} = \frac{v_{ij}}{V_i(S_i)}.$$

We observe that the assumption that each nest includes the same number of products is without loss of generality because if some nest i includes fewer than n products, then we can include additional products j in this nest with preference weight $v_{ij} = 0$ and these products would never be purchased. We emphasize that it is also possible to have $v_{i0} = 0$ for some nest i , in which case, given that a customer decides to purchase a product in nest i , this customer never leaves without purchasing anything. Each nest i has a parameter $\gamma_i \geq 0$ associated with it that characterizes the degree of the dissimilarity of the products in the nest. Furthermore, we use v_0 to denote the preference weight for the option of not choosing any of the nests. If we offer the assortment (S_1, \dots, S_m) over all nests, then a customer chooses nest i with probability

$$Q_i(S_1, \dots, S_m) = \frac{V_i(S_i)^{\gamma_i}}{v_0 + \sum_{l \in M} V_l(S_l)^{\gamma_l}}. \quad (1)$$

We observe that γ_i serves the purpose of dampening or magnifying the total preference weight of the available options within nest i .

If we offer the assortment S_i in nest i , then we can write the expected revenue we obtain from this nest as

$$R_i(S_i) = \sum_{j \in S_i} r_{ij} P_{ij}(S_i) = \frac{\sum_{j \in S_i} r_{ij} v_{ij}}{V_i(S_i)},$$

with the interpretation that $R_i(\emptyset) = 0$. Therefore, if we offer the assortment (S_1, \dots, S_m) over all nests with $S_i \subset N$ for all $i \in M$, then we obtain an expected revenue of

$$\Pi(S_1, \dots, S_m) = \sum_{i \in M} Q_i(S_1, \dots, S_m) R_i(S_i).$$

Our goal is to choose an assortment (S_1, \dots, S_m) that maximizes the expected revenue over all nests, yielding the assortment optimization problem

$$Z^* = \max_{(S_1, \dots, S_m): S_i \subset N, i \in M} \Pi(S_1, \dots, S_m). \quad (2)$$

Throughout this paper, we classify the instances of the assortment optimization problem above along two dimensions. The first dimension is based on the values of the dissimilarity parameters $(\gamma_1, \dots, \gamma_m)$ of the nests. Along this dimension, we separately consider the two cases where (i) we have $\gamma_i \leq 1$ for all $i \in M$, and (ii) there are no restrictions on the dissimilarity parameters. The second dimension of classification is based on the values of the preference weights (v_{10}, \dots, v_{m0}) of the no purchase options within the nests. Along this dimension, we separately consider the two cases where (i) we have $v_{i0} = 0$ for all $i \in M$, and (ii) there are no restrictions on the preference weights of the no purchase options. Since there are two cases to consider along each one of the two dimensions, we study the assortment optimization problem in (2) under four cases. It turns out that while the assortment optimization problem in (2) is polynomially solvable when $\gamma_i \leq 1$ and $v_{i0} = 0$ for all $i \in M$, lifting any one of these restrictions renders the problem NP-hard and we resort to approximation methods. If we have $\gamma_i \leq 1$ for all $i \in M$, but there are no restrictions on the preference weights of the no purchase options, then we can construct a constant-factor approximation algorithm. However, if there are no restrictions on the dissimilarity parameters of the nests, then the performance guarantees of the approximation methods we give depend on the problem data.

Along the first dimension, the case with $\gamma_i \leq 1$ for all $i \in M$ corresponds to the standard form of the nested logit model studied by McFadden (1974), where the author notes that having $\gamma_i \leq 1$ for all $i \in M$ implies that the nested logit model is compatible with a random utility-based choice model. Borsch-Supan (1990) observes that if we use data to estimate the parameters of the nested logit model, then the estimators of the dissimilarity parameters may exceed one and constraining the dissimilarity parameters to lie in the interval $[0, 1]$ may deteriorate the fit. Therefore, it is also useful to study problem (2) without any restrictions on the dissimilarity parameters even though the model may no longer be compatible with random utility-based choice models.

Interestingly, if we allow the dissimilarity parameters of the nests to take on values larger than one, then we can model a useful practical phenomenon through the nested logit model. In particular, under the nested logit model, if we offer the assortment (S_1, \dots, S_m) , then the probability that a customer purchases product $j \in S_i$ in nest i is given by

$$Q_i(S_1, \dots, S_m) P_{ij}(S_i) = \frac{V_i(S_i)^{\gamma_i - 1}}{v_0 + \sum_{l \in M} V_l(S_l)^{\gamma_l}} v_{ij}.$$

From the expression above, we observe that when $\gamma_i \leq 1$, adding a product $k \notin S_i$ to nest i decreases the purchase probability of product $j \in S_i$. Therefore, when the dissimilarity parameters of the nests do not exceed one, the products in a nest always act as competitors to each other and adding a new product to a nest decreases the probability of purchase for the other products in the nest. In practice, this is not always the case. For example, if the nests correspond to different car dealers, then offering a new luxury car may increase the probability of purchase for other cars in the same dealer because the newly offered luxury car may help attract a larger fraction of customers. We observe that if $\gamma_i > 1$ and we add a new product to nest i , then both $V_i(S_i)^{\gamma_i - 1}$ and $\sum_{l \in M} V_l(S_l)^{\gamma_l}$ increase in the expression above. As a result, the probability that a customer purchases product j may increase or decrease. This feature may allow us to model synergies between different products in a nest, and when such synergies exist, it may even be beneficial to include loss leaders in a nest to attract traffic to this nest. Motivated

by this observation, we refer to the case with $\gamma_i \leq 1$ for all $i \in M$ as the case with purely *competitive products*, whereas we refer to the case with no restrictions on the dissimilarity parameters of the nests as the case with *possibly synergistic products*.

Along the second dimension, the case with $v_{i0} = 0$ for all $i \in M$ corresponds to a situation where if a customer decides to make a purchase within a particular nest, then the customer always makes a purchase within the selected nest. In other words, the demand within a nest is fully captured without loss to the no purchase option. We refer to this situation as the case with *fully-captured nests*. On the other hand, if the preference weight of the no purchase option within a nest is strictly positive, then a customer may leave without purchasing anything in the selected nest. We refer to this situation as the case with *partially-captured nests*.

2 Linear Programming Representation

The assortment optimization problem in (2) is of a combinatorial nature. In this section, we present a linear programming formulation of this problem. The linear programming formulation is not too useful directly as a computational tool since its number of constraints grows exponentially with the number of products. However, it turns out that we can build on the linear programming formulation to develop a general approximation result for problem (2). The approximation methods that we propose throughout the paper are tightly related to this general approximation result.

To formulate problem (2) as a linear program, we first observe that this problem is equivalent to $\min\{x : x \geq \sum_{i \in M} Q_i(S_1, \dots, S_m) R_i(S_i) \forall (S_1, \dots, S_m) \text{ with } S_i \subset N, i \in M\}$. By using the definition of $Q_i(S_1, \dots, S_m)$ in (1), we can write the constraints in this problem as

$$v_0 x \geq \sum_{i \in M} (V_i(S_i)^{\gamma_i} R_i(S_i) - V_i(S_i)^{\gamma_i} x) \quad \forall (S_1, \dots, S_m) \text{ with } S_i \subset N, i \in M,$$

which, in turn, are equivalent to the single constraint

$$v_0 x \geq \max_{(S_1, \dots, S_m): S_i \subset N, i \in M} \left\{ \sum_{i \in M} V_i(S_i)^{\gamma_i} (R_i(S_i) - x) \right\}.$$

The key observation is that the optimization problem on the right side of the constraint above decomposes by the nests and the constraint can be written as

$$v_0 x \geq \sum_{i \in M} \max_{S_i \subset N} V_i(S_i)^{\gamma_i} (R_i(S_i) - x).$$

Therefore, problem (2) is equivalent to

$$\begin{aligned} \min \quad & x \\ \text{s.t.} \quad & v_0 x \geq \sum_{i \in M} \max_{S_i \subset N} V_i(S_i)^{\gamma_i} (R_i(S_i) - x), \end{aligned}$$

where the only decision variable is x . Noting that Z^* is the optimal objective value of problem (2), the discussion so far implies that if x^* is the optimal solution to the problem above, then the optimal

objective value of this problem is also x^* and we have $x^* = Z^*$. The constraint of the problem above is nonlinear, but to linearize this constraint, we can define the decision variables $y = (y_1, \dots, y_m)$ as $y_i = \max_{S_i \subset N} V_i(S_i)^{\gamma_i} (R_i(S_i) - x)$ and write the problem as

$$\begin{aligned} \min \quad & x \\ \text{s.t.} \quad & v_0 x \geq \sum_{i \in M} y_i \\ & y_i \geq V_i(S_i)^{\gamma_i} (R_i(S_i) - x) \quad \forall S_i \subset N, i \in M, \end{aligned} \tag{3}$$

where the decision variables are (x, y) . Problem (3) is a linear program with $1 + m$ decision variables and $1 + m 2^n$ constraints. One of the interesting observations is that the number of possible assortments in problem (2) is 2^{mn} , which increases exponentially in both the number of nests and the number of products in each nest. In contrast, the numbers of decision variables and constraints in problem (3) grow linearly with the number of nests, and problem (3) can be tractable when the number of products in each nest is relatively small, irrespective of the number of nests. When the number of products in each nest is large, a possible solution approach for problem (3) is to use column generation on its dual, but due to the presence of the dissimilarity coefficient in the second set of constraints, the column generation subproblem is nonlinear and this renders the column generation approach intractable.

Although problem (3) is difficult to solve when the number of products in each nest is large, we can build on this problem to develop a general approximation method. Assume that we identify a collection of candidate assortments $\{A_{it} : t \in \mathcal{T}_i\}$ that we may consider offering in nest i , where we have $A_{it} \subset N$ for all $t \in \mathcal{T}_i$. We are interested in finding a combination of these assortments for the different nests so that the combined assortment provides the largest possible expected revenue. In other words, we are interested in finding the assortment that provides the largest expected revenue when we consider assortments of the form (S_1, \dots, S_m) with $S_i \in \{A_{it} : t \in \mathcal{T}_i\}$ for all $i \in M$. This problem can be formulated as the linear program

$$\begin{aligned} \min \quad & x \\ \text{s.t.} \quad & v_0 x \geq \sum_{i \in M} y_i \\ & y_i \geq V_i(S_i)^{\gamma_i} (R_i(S_i) - x) \quad \forall S_i \in \{A_{it} : i \in \mathcal{T}_i\}, i \in M. \end{aligned} \tag{4}$$

The number of decision variables in problem (4) is still $1 + m$. The number of constraints is $1 + \sum_{i \in M} |\mathcal{T}_i|$, which can be reasonable when the collections of candidate assortments are not too large.

We now provide some observations to develop a general approximation result by building on problem (4). The constraints in problem (4) can succinctly be written as

$$v_0 x \geq \sum_{i \in M} \max_{S_i \in \{A_{it} : t \in \mathcal{T}_i\}} V_i(S_i)^{\gamma_i} (R_i(S_i) - x).$$

We will now argue that the constraint above must be satisfied as equality at an optimal solution. Let (\hat{x}, \hat{y}) be an optimal solution to problem (4) and suppose there is a gap. We can then decrease the value of \hat{x} without violating the constraint, thereby obtaining a strictly better solution to problem (4) than \hat{x} ,

establishing the claim. Therefore, letting \hat{S}_i be the solution to the maximization problem on the right side of the constraint above with $x = \hat{x}$, it must be the case that $v_0 \hat{x} = \sum_{i \in M} V_i(\hat{S}_i)^{\gamma_i} (R_i(\hat{S}_i) - \hat{x})$. Solving for \hat{x} , we obtain

$$\hat{x} = \frac{\sum_{i \in M} V_i(\hat{S}_i)^{\gamma_i} R_i(\hat{S}_i)}{v_0 + \sum_{i \in M} V_i(\hat{S}_i)^{\gamma_i}} = \sum_{i \in M} Q_i(\hat{S}_1, \dots, \hat{S}_m) R_i(\hat{S}_i) = \Pi(\hat{S}_1, \dots, \hat{S}_m).$$

Consequently, if (\hat{x}, \hat{y}) is an optimal solution to problem (4) and \hat{S}_i is an optimal solution to the problem $\max_{S_i \in \{A_{it} : t \in \mathcal{T}_i\}} V_i(S_i)^{\gamma_i} (R_i(S_i) - \hat{x})$, then the expected revenue obtained by offering the assortment $(\hat{S}_1, \dots, \hat{S}_m)$ is precisely \hat{x} .

We observe that problem (4) includes only a subset of the constraints in problem (3), which implies that problem (4) is a relaxed version of problem (3). Therefore, if we let (x^*, y^*) and (\hat{x}, \hat{y}) respectively be the optimal solutions to problems (3) and (4), then we have $Z^* = x^* \geq \hat{x}$. Furthermore, for some α and β , if we can show that $(\alpha \hat{x}, \beta \hat{y})$ is a feasible solution to problem (3), then we also obtain $\alpha \hat{x} \geq Z^* = x^* \geq \hat{x}$, which implies that the expected revenue obtained by offering the assortment $(\hat{S}_1, \dots, \hat{S}_m)$ as defined above deviates from the optimal expected revenue by no more than a factor of α . We collect these observations in the following theorem.

Theorem 1 *Let (\hat{x}, \hat{y}) be an optimal solution to problem (4), and for all $i \in M$, let \hat{S}_i be an optimal solution to the problem*

$$\max_{S_i \in \{A_{it} : t \in \mathcal{T}_i\}} V_i(S_i)^{\gamma_i} (R_i(S_i) - \hat{x}). \quad (5)$$

If $(\alpha \hat{x}, \beta \hat{y})$ is a feasible solution to problem (3) for some α and β , then the expected revenue obtained by offering the assortment $(\hat{S}_1, \dots, \hat{S}_m)$ deviates from the optimal objective value of problem (2) by no more than a factor of α . In other words, letting $\hat{Z} = \Pi(\hat{S}_1, \dots, \hat{S}_m)$, we have $\alpha \hat{Z} \geq Z^ \geq \hat{Z}$.*

Theorem 1 provides sufficient conditions under which we can stitch together a good assortment from the collections of candidate assortments $\{A_{it} : t \in \mathcal{T}_i\}$ for $i \in M$. The thought process we used to reach Theorem 1 will be critical throughout the paper. In particular, we will design collections of assortments such that if (\hat{x}, \hat{y}) is the optimal solution to problem (4) with these collections of assortments, then $(\alpha \hat{x}, \beta \hat{y})$ ends up being a feasible solution to problem (3) for some α and β . In that case, we can solve problem (4) with these collections of assortments to obtain the optimal solution (\hat{x}, \hat{y}) . Letting \hat{S}_i be the assortment that solves problem (5), Theorem 1 implies that the expected revenue from the assortment $(\hat{S}_1, \dots, \hat{S}_m)$ deviates from the optimal expected revenue by at most a factor of α .

We use various collections of candidate assortments. One possibility is to use the assortments that include a certain number of products with the largest revenues. This class of assortments is known to be optimal when the customer choices are governed by the multinomial logit model and Rusmevichientong, Shmoys and Topaloglu (2010) give a performance guarantee for such assortments when the underlying choice model is the multinomial logit model with multiple customer types. In the next section, we show that this class of assortments is still optimal under the nested logit model as long as we only have competitive products and fully-captured nests, but in general, we may need to look beyond this class to find good solutions for the assortment optimization problem we are interested in.

3 Competitive Products and Fully-Captured Nests

In this section, we focus on instances of the assortment optimization problem in (2) with $\gamma_i \leq 1$ and $v_{i0} = 0$ for all $i \in M$. For this case, we show that there exists an optimal solution (S_1^*, \dots, S_m^*) to problem (2) such that each one of the assortments S_1^*, \dots, S_m^* is either the empty assortment or is of the form $S_i^* = \{1, 2, \dots, j\}$ for some $j \in N$. Noting that the products in each nest are ordered such that $r_{i1} \geq r_{i2} \geq \dots \geq r_{in}$, this result implies that an optimal assortment in each nest includes a certain number of products with the largest revenues. We call such assortments *nested-by-revenue* assortments. For notational brevity, we use N_{ij} to denote the nested-by-revenue assortment that includes the first j products with the largest revenues in nest i . In other words, we have $N_{ij} = \{1, 2, \dots, j\}$ for all $i \in M, j \in N$. For notational uniformity, we also let $N_{i0} = \emptyset$ and $N_+ = N \cup \{0\}$, in which case, our goal is to show that an optimal solution (S_1^*, \dots, S_m^*) to problem (2) is of the form $S_i^* = N_{ij}$ for some $j \in N_+$. Throughout this section, we assume without loss of generality that $v_0 > 0$. Otherwise, since we have $v_{i0} = 0$ for all $i \in M$, it is optimal to offer only one product with the largest revenue $\max_{i \in M} r_{i1}$ over all nests and this product would be purchased with probability one.

The following proposition shows that if it is optimal to offer a nonempty assortment in a nest, then the expected revenue from this nest should at least be equal to the optimal expected revenue over all nests.

Proposition 2 *If (S_1^*, \dots, S_m^*) is an optimal solution to problem (2) and $S_i^* \neq \emptyset$, then $R_i(S_i^*) \geq Z^*$.*

Proof. To get a contradiction, assume that $S_i^* \neq \emptyset$ and $R_i(S_i^*) < Z^*$. For notational convenience, let $R_l = R_l(S_l^*)$ and $q_l = V_l(S_l^*)^{\gamma_l}$ for all $l \in M$. Thus, we have $Q_i(S_1^*, \dots, S_m^*) = q_i / (v_0 + \sum_{l \in M} q_l)$ and we can write the optimal expected revenue as

$$\begin{aligned} Z^* = \Pi(S_1^*, \dots, S_m^*) &= \frac{\sum_{l \in M} q_l R_l}{v_0 + \sum_{l \in M} q_l} = \frac{v_0 + \sum_{l \in M, l \neq i} q_l}{v_0 + \sum_{l \in M} q_l} \frac{\sum_{l \in M, l \neq i} q_l R_l}{v_0 + \sum_{l \in M, l \neq i} q_l} + \frac{q_i R_i}{v_0 + \sum_{l \in M} q_l} \\ &= \frac{v_0 + \sum_{l \in M, l \neq i} q_l}{v_0 + \sum_{l \in M} q_l} \Pi(S_1^*, \dots, S_{i-1}^*, \emptyset, S_{i+1}^*, \dots, S_m^*) + \frac{q_i}{v_0 + \sum_{l \in M} q_l} R_i. \end{aligned}$$

Noting that $v_0 > 0$ and $q_i = V_i(S_i^*)^{\gamma_i} > 0$, the equality above shows that Z^* is a nontrivial convex combination of $\Pi(S_1^*, \dots, S_{i-1}^*, \emptyset, S_{i+1}^*, \dots, S_m^*)$ and $R_i = R_i(S_i^*)$. So, having $R_i < Z^*$ implies that $\Pi(S_1^*, \dots, S_{i-1}^*, \emptyset, S_{i+1}^*, \dots, S_m^*) > Z^*$ contradicting the fact that Z^* is the optimal expected revenue. \square

In the following lemma, we show that the products with small revenues can be removed from the assortment without degrading the performance.

Lemma 3 *Assume that $Z = \Pi(S_1, \dots, S_m)$ for some assortment (S_1, \dots, S_m) and there exists a product $j \in S_i$ that satisfies $r_{ij} < \gamma_i Z + (1 - \gamma_i) R_i(S_i)$ and $R_i(S_i) \geq Z$. Then, removing product j from S_i yields a strictly larger expected revenue than Z .*

Proof. Let \hat{S}_i be the assortment constructed by removing product j from S_i . We show that the assortment $(S_1, \dots, S_{i-1}, \hat{S}_i, S_{i+1}, \dots, S_m)$ provides an expected revenue of \hat{Z} satisfying $\hat{Z} > Z$. For

notational convenience, let $\hat{R}_i = R_i(\hat{S}_i)$, $\hat{q}_i = V_i(\hat{S}_i)^{\gamma_i}$, $R_l = R_l(S_l)$ and $q_l = V_l(S_l)^{\gamma_l}$ for all $l \in M$. Following an argument similar to the one in the proof of Proposition 2, we can write the expected revenue from the assortment (S_1, \dots, S_m) as

$$\begin{aligned} Z = \Pi(S_1, \dots, S_m) &= \frac{\sum_{l \in M} q_l R_l}{v_0 + \sum_{l \in M} q_l} = \frac{v_0 + \sum_{l \in M, l \neq i} q_l + \hat{q}_i}{v_0 + \sum_{l \in M} q_l} \frac{\sum_{l \in M, l \neq i} q_l R_l + \hat{q}_i \hat{R}_i}{v_0 + \sum_{l \in M, l \neq i} q_l + \hat{q}_i} + \frac{q_i R_i - \hat{q}_i \hat{R}_i}{v_0 + \sum_{l \in M} q_l} \\ &= \frac{v_0 + \sum_{l \in M, l \neq i} q_l + \hat{q}_i}{v_0 + \sum_{l \in M} q_l} \Pi(S_1, \dots, S_{i-1}, \hat{S}_i, S_{i+1}, \dots, S_m) + \frac{q_i - \hat{q}_i}{v_0 + \sum_{l \in M} q_l} \frac{q_i R_i - \hat{q}_i \hat{R}_i}{q_i - \hat{q}_i}. \end{aligned}$$

Therefore, Z is a convex combination of $\hat{Z} = \Pi(S_1, \dots, S_{i-1}, \hat{S}_i, S_{i+1}, \dots, S_m)$ and $(q_i R_i - \hat{q}_i \hat{R}_i)/(q_i - \hat{q}_i)$. In this case, the desired result follows if we can show that $Z > (q_i R_i - \hat{q}_i \hat{R}_i)/(q_i - \hat{q}_i)$. In the rest of the proof, we equivalently show that $q_i R_i - \hat{q}_i \hat{R}_i < (q_i - \hat{q}_i) Z$.

Let $\alpha = V_i(\hat{S}_i)/V_i(S_i)$, so $\hat{q}_i = \alpha^{\gamma_i} q_i$. Using the fact that $v_{ij} = V_i(S_i) - V_i(\hat{S}_i)$, we can write \hat{R}_i as

$$\hat{R}_i = \frac{\sum_{k \in \hat{S}_i} r_{ik} v_{ik}}{V_i(\hat{S}_i)} = \frac{\sum_{k \in S_i} r_{ik} v_{ik} - r_{ij} (V_i(S_i) - V_i(\hat{S}_i))}{\alpha V_i(S_i)} = \frac{1}{\alpha} R_i - \frac{1 - \alpha}{\alpha} r_{ij}.$$

Therefore, $q_i R_i - \hat{q}_i \hat{R}_i < (q_i - \hat{q}_i) Z$ holds if and only of

$$q_i R_i - \alpha^{\gamma_i} q_i \left[\frac{1}{\alpha} R_i - \frac{1 - \alpha}{\alpha} r_{ij} \right] < (q_i - \alpha^{\gamma_i} q_i) Z.$$

Arranging the terms in the expression above, we observe that $q_i R_i - \hat{q}_i \hat{R}_i < (q_i - \hat{q}_i) Z$ holds if and only if $r_{ij} < g(\alpha) Z + (1 - g(\alpha)) R_i$, where $g(\alpha) = (1 - \alpha^{\gamma_i})/(\alpha^{\gamma_i-1} - \alpha^{\gamma_i})$. By the hypothesis, we have $r_{ij} < \gamma_i Z + (1 - \gamma_i) R_i$. So, it is enough to show that $\gamma_i Z + (1 - \gamma_i) R_i \leq g(\alpha) Z + (1 - g(\alpha)) R_i$. Since both sides of the last inequality are convex combinations of Z and R_i and $Z \leq R_i$ by the hypothesis, the inequality holds as long as $g(\alpha) \leq \gamma_i$. However, the last relationship is true because $g(\alpha)$ is increasing in α when $\gamma_i \leq 1$ and by L'Hopital's rule, $g(\alpha) \leq \lim_{\alpha \uparrow 1} g(\alpha) = \gamma_i$. \square

Lemma 3 gives us a mechanism to remove certain products with small revenues without reducing the expected revenue from an assortment. To see a useful implication of Proposition 2 and Lemma 3, assume that (S_1^*, \dots, S_m^*) is an optimal solution to problem (2) with $S_i^* \neq \emptyset$. We must have $R_i(S_i^*) \geq Z^*$ by Proposition 2. In this case, we must have $r_{ij} \geq \gamma_i Z^* + (1 - \gamma_i) R_i(S_i^*)$ for all $j \in S_i^*$ by Lemma 3. Otherwise, we can remove a product from S_i^* and obtain an assortment that provides a strictly larger expected revenue than Z^* . We use this observation in the following theorem to show that nested-by-revenue assortments provide an optimal solution to problem (2).

Theorem 4 *There exists an optimal solution (S_1^*, \dots, S_m^*) to problem (2) such that, for all $i \in M$, we have $S_i^* = N_{ij}$ for some $j \in N_+$.*

Proof. Assume that (S_1^*, \dots, S_m^*) is an optimal solution to problem (2) providing an expected revenue of Z^* and that there is a nest i such that S_i^* contains product j but not product k with $k < j$. Let \hat{S}_i be the assortment constructed by adding product k to S_i^* . Using \hat{Z} to denote the expected revenue from the assortment $(S_1^*, \dots, S_{i-1}^*, \hat{S}_i, S_{i+1}^*, \dots, S_m^*)$, we show that $\hat{Z} \geq Z^*$. Therefore, the assortment

$(S_1^*, \dots, S_{i-1}^*, \hat{S}_i, S_{i+1}^*, \dots, S_m^*)$ must also be optimal. Repeating the argument until the assortments for all nests are of the form $\{1, 2, \dots, j\}$ for some $j \in N_+$ establishes the result.

Similar to the notation in the proof of Lemma 3, let $\hat{R}_i = R_i(\hat{S}_i)$, $R_i = R_i(S_i)$, $\hat{q}_i = V_i(\hat{S}_i)^{\gamma_i}$ and $q_i = V_i(S_i)^{\gamma_i}$. The main idea is to use an argument similar to the one in the proof of Lemma 3 to write $\hat{Z} = \Pi(S_1^*, \dots, S_{i-1}^*, \hat{S}_i, S_{i+1}^*, \dots, S_m^*)$ as a convex combination of Z^* and $(\hat{q}_i \hat{R}_i - q_i R_i)/(\hat{q}_i - q_i)$. In this case, the desired result follows if we can show that $Z^* \leq (\hat{q}_i \hat{R}_i - q_i R_i)/(\hat{q}_i - q_i)$. We equivalently show that $(\hat{q}_i - q_i) Z^* \leq \hat{q}_i \hat{R}_i - q_i R_i$. Let $\alpha = V_i(S_i^*)/V_i(\hat{S}_i)$, so $\hat{q}_i = q_i/\alpha^{\gamma_i}$. Using the fact that $v_{ik} = V_i(\hat{S}_i) - V_i(S_i^*)$, we write \hat{R}_i as

$$\hat{R}_i = \frac{\sum_{j' \in \hat{S}_i} r_{ij'} v_{ij'}}{V_i(\hat{S}_i)} = \alpha \frac{\sum_{j' \in S_i^*} r_{ij'} v_{ij'} + r_{ik} (V_i(\hat{S}_i) - V_i(S_i^*))}{V_i(S_i^*)} = \alpha R_i + (1 - \alpha) r_{ik}.$$

Therefore, $(\hat{q}_i - q_i) Z^* \leq \hat{q}_i \hat{R}_i - q_i R_i$ holds if and only if

$$\left[\frac{q_i}{\alpha^{\gamma_i}} - q_i \right] Z^* \leq \frac{q_i}{\alpha^{\gamma_i}} (\alpha R_i + (1 - \alpha) r_{ik}) - q_i R_i.$$

Arranging the terms in the expression above, we observe that $(\hat{q}_i - q_i) Z^* \leq \hat{q}_i \hat{R}_i - q_i R_i$ holds if and only if $h(\alpha) Z^* + (1 - h(\alpha)) R_i \leq r_{ik}$, where we let $h(\alpha) = (1 - \alpha^{\gamma_i})/(1 - \alpha)$. From the discussion that follows Lemma 3, since $j \in S_i^*$, we know that $\gamma_i Z^* + (1 - \gamma_i) R_i \leq r_{ij} \leq r_{ik}$, where the last inequality follows from the fact that $k < j$. So, it is enough to show that $h(\alpha) Z^* + (1 - h(\alpha)) R_i \leq \gamma_i Z^* + (1 - \gamma_i) R_i$. Since both sides of the last inequality are convex combinations of Z^* and R_i and $Z^* \leq R_i$ by Proposition 2, the inequality holds if and only if $h(\alpha) \geq \gamma_i$. However, the last relationship is true because $h(\alpha)$ is decreasing in α and by L'Hoptial's rule, $h(\alpha) \geq \lim_{\alpha \uparrow 1} h(\alpha) = \gamma_i$. \square

Theorem 4 shows that we can construct an optimal solution to problem (2) by only considering the nested-by-revenue assortments. To find the best combination of such assortments for the different nests, we can make use of problem (4). In particular, we replace the collection of candidate assortments $\{A_{it} : t \in \mathcal{T}_i\}$ in problem (4) with the nested-by-revenue assortments $\{N_{ij} : j \in N_+\}$ and solve problem (4) to find the best combination of nested-by-revenue assortments for the different nests. By Theorem 4, this best combination has to be an optimal solution to problem (2). In this way, we can find an optimal solution to problem (2) by solving a linear program with $1 + m$ decision variables and $1 + m(1 + n)$ constraints.

4 Possibly Synergistic Products and Fully-Captured Nests

In this section, we focus on instances of the assortment optimization problem in (2), where we do not have any restrictions on the dissimilarity parameters $(\gamma_1, \dots, \gamma_m)$ of the nests, but we still have $v_{i0} = 0$ for all $i \in M$. We show that allowing the dissimilarity parameters for the nests to take on values larger than one changes the structure of problem (2) drastically. In particular, the result that we establish in the previous section does not necessarily hold when the dissimilarity parameters of the nests can take on arbitrary values and the nested-by-revenue assortments are no longer optimal. In the next section, we first characterize the computational complexity of the problem when we have no restrictions on the nest dissimilarity parameters. Following this result, we give a performance guarantee

for the nested-by-revenue assortments. Throughout this discussion, we assume that $\gamma_i > 1$ for some $i \in M$. Otherwise, nested-by-revenue assortments are optimal by Theorem 4.

4.1 Computational Complexity

We begin by giving an example that shows why nested-by-revenue assortments are no longer optimal when the dissimilarity parameters of the nests can take on values larger than one. This example also demonstrates that nested-by-revenue assortments can perform arbitrarily badly when the revenues and the preference weights of the products in a nest drastically differ from each other. Following this example, we establish that the assortment optimization problem in (2) is NP-hard whenever we allow $\gamma_i > 1$ for some $i \in M$.

To give an example where nested-by-revenue assortments do not perform well, we consider an instance of problem (2) with a single nest. The preference weight for the option of not choosing any of the nests is $v_0 = 1$. The dissimilarity parameter of the nest is $\gamma_1 = 2$. There are three products in the nest. Letting $\varepsilon \leq 1$ be a small positive number, the following table gives the revenues and the preference weights associated with the three products.

Product	1	2	3
Revenue	1	ε^4	0
Preference Weight	ε^2	$3/\varepsilon^2$	$1/\varepsilon$

Since there is only one nest with $\gamma_1 = 2$, the expected revenue from an assortment $S_1 \subset \{1, 2, 3\}$ is

$$\Pi(S_1) = Q_1(S_1) R_1(S_1) = \frac{V_1(S_1)^2}{v_0 + V_1(S_1)^2} \frac{\sum_{j \in S_1} r_{1j} v_{1j}}{V_1(S_1)} = \frac{V_1(S_1) \sum_{j \in S_1} r_{1j} v_{1j}}{v_0 + V_1(S_1)^2}.$$

We compute and bound the expected revenues from the three nested-by-revenue assortments as

$$\begin{aligned} \Pi(\{1\}) &= \frac{\varepsilon^2 \varepsilon^2}{1 + (\varepsilon^2)^2} \leq \varepsilon^4 \\ \Pi(\{1, 2\}) &= \frac{(\varepsilon^2 + 3/\varepsilon^2)(\varepsilon^2 + 3\varepsilon^2)}{1 + (\varepsilon^2 + 3/\varepsilon^2)^2} \leq \frac{(4/\varepsilon^2)4\varepsilon^2}{9/\varepsilon^4} = \frac{16}{9} \varepsilon^4 \\ \Pi(\{1, 2, 3\}) &= \frac{(\varepsilon^2 + 3/\varepsilon^2 + 1/\varepsilon)(\varepsilon^2 + 3\varepsilon^2)}{1 + (\varepsilon^2 + 3/\varepsilon^2 + 1/\varepsilon)^2} \leq \frac{(5/\varepsilon^2)4\varepsilon^2}{9/\varepsilon^4} = \frac{20}{9} \varepsilon^4, \end{aligned}$$

which implies that the expected revenue from the best nested-by-revenue assortment is no larger than $\frac{20}{9} \varepsilon^4$. On the other hand, the expected revenue from the assortment $\{1, 3\}$ is given by

$$\Pi(\{1, 3\}) = \frac{(\varepsilon^2 + 1/\varepsilon)(\varepsilon^2)}{1 + (\varepsilon^2 + 1/\varepsilon)^2} \geq \frac{(1/\varepsilon)\varepsilon^2}{1 + (1/\varepsilon + 1/\varepsilon)^2} \geq \frac{\varepsilon}{1/\varepsilon^2 + (1/\varepsilon + 1/\varepsilon)^2} = \frac{1}{5} \varepsilon^3.$$

Thus, the optimal expected revenue exceeds the expected revenue from the best nested-by-revenue assortment by at least a factor of $(\varepsilon^3/5)/(\frac{20}{9} \varepsilon^4) = 9/(100\varepsilon)$. As $\varepsilon \rightarrow 0$, the performance of the best nested-by-revenue assortment becomes arbitrarily poor. The key observation in this problem instance is that the revenue associated with product 1 is quite large when compared with the other product revenues. Therefore, we would like to be able to sell product 1 with high probability. One can check

that if we offer product 1 by itself, then the probability of purchase for product 1 is between $\varepsilon^4/2$ and ε^4 . On the other hand, if we offer products 1 and 2 together, then the probability of purchase for product 1 is always smaller than $\varepsilon^4/2$. Therefore, if we offer product 2 next to product 1, then the probability of purchase for product 1 goes down. In contrast, if we offer products 1 and 3, then it is possible to check that the probability of purchase for product 1 always exceeds $\varepsilon^3/5$, which is larger than ε^4 for small values of ε . This observation indicates that product 3 acts as a synergistic product to product 1 and offering product 3 next to product 1 increases the probability of purchase for product 1. We also note that even if the revenue of product 3 was not zero but slightly negative, it would still be beneficial to add this product to the offered assortment, justifying a loss leader.

It turns out that if we allow the dissimilarity parameters of the nests to take on values larger than one, then not only the nested-by-revenue assortments cease to be optimal, but problem (2) becomes NP-hard. We devote the rest of this section to showing this result. To show the result we are interested in, we focus on the following decision-theoretic formulation of the assortment optimization problem.

Assortment Feasibility. Given a profit threshold K , is there an assortment (S_1, \dots, S_m) that provides an expected revenue of K or more for problem (2)?

To establish the NP-hardness of problem (2), Theorem 5 below shows that any instance of the partition problem, which is a well-known NP-hard problem as established in Garey and Johnson (1979), can be reduced to an instance of the assortment feasibility problem. Rusmevichientong, Shmoys and Topaloglu (2010) also use a reduction from the partition problem to show the NP-hardness of an assortment optimization problem, but their choice model is the multinomial logit model with multiple customer types, rather than the nested logit model. The partition problem is described as follows.

Partition. Given integer-valued sizes (c_1, \dots, c_n) such that $\sum_{j=1}^n c_j = 2T$ with T integer, can we find a subset $S \subset \{1, \dots, n\}$ such that $\sum_{j \in S} c_j = \sum_{j \in \{1, \dots, n\} \setminus S} c_j = T$?

Theorem 5 *If we allow the dissimilarity parameters for the nests to take on values larger than one, then the assortment feasibility problem is NP-hard.*

Proof. Assume that we are given any instance of the partition problem with sizes (c_1, \dots, c_n) and $\sum_{j=1}^n c_j = 2T$. We define an instance of the assortment feasibility problem as follows. There is only one nest. The preference weight for the option of not choosing any of the nests is $v_0 = (1 + T)^2$. The dissimilarity parameter of the nest is $\gamma_1 = 2$. There are $n + 1$ products in the nest. The revenue associated with the first n products is given by $r_{1j} = 0$ for all $j = 1, \dots, n$. The revenue associated with the last product is $r_{1,n+1} = 2(1 + T)$. The preference weights of the first n products are given by $v_{1j} = c_j$ for all $j = 1, \dots, n$. The preference weight associated with the last product is $v_{1,n+1} = 1$. We set the expected revenue threshold in the assortment feasibility problem as $K = 1$.

In the rest of the proof, we show that there exists an assortment that provides an expected revenue of K or more in the assortment feasibility problem if and only if there exists a subset $S \subset \{1, \dots, n\}$ such that $\sum_{j \in S} c_j = T$. The first observation that if we want to get a positive expected revenue in the

assortment feasibility problem, then we have to offer the last product with revenue $2(1+T)$. Therefore, the only question for the assortment feasibility problem is to choose a subset S among the products with zero revenues that makes sure that we obtain an expected revenue of $K = 1$ or more. If we offer a subset S of the first n products together with the last product, then the expected revenue is $Q_1(S \cup \{n+1\}) R_1(S \cup \{n+1\})$, which evaluates to

$$\frac{(\sum_{j \in S} c_j + 1)^2}{(1+T)^2 + (\sum_{j \in S} c_j + 1)^2} \frac{2(1+T)}{(\sum_{j \in S} c_j + 1)}.$$

Therefore, there exists an assortment with an expected revenue of $K = 1$ or more if and only if

$$\frac{(\sum_{j \in S} c_j + 1) 2(1+T)}{(1+T)^2 + (\sum_{j \in S} c_j + 1)^2} \geq 1.$$

Arranging the terms in the expression above, the inequality above is equivalent to

$$2(1+T) \sum_{j \in S} c_j + 2(1+T) \geq 1 + 2T + T^2 + 1 + 2 \sum_{j \in S} c_j + \left(\sum_{j \in S} c_j \right)^2,$$

which can equivalently be written as

$$\left(\sum_{j \in S} c_j \right)^2 - 2T \sum_{j \in S} c_j + T^2 \leq 0.$$

Since the last inequality is equivalent to $(\sum_{j \in S} c_j - T)^2 \leq 0$, there exists an assortment with an expected revenue of $K = 1$ or more if and only if there exists a subset S with $(\sum_{j \in S} c_j - T)^2 \leq 0$. However, the only way for the last inequality to hold is to have $\sum_{j \in S} c_j = T$. Therefore, finding an assortment that yields an expected revenue of K or more is equivalent to finding a subset S that satisfies $\sum_{j \in S} c_j = T$ and the latter statement is precisely what the partition problem is interested in. \square

4.2 Performance of Nested-by-Revenue Assortments

In the previous section, we show that nested-by-revenue assortments may not perform well when we allow the dissimilarity parameters of the nests to take on values larger than one. Our goal in this section is to develop a performance bound for this class of assortments as a function of the problem data. In particular, recalling that we use N_{ij} to denote the nested-by-revenue assortment that includes the first j products with the largest revenues in nest i , we show that by focusing only on the nested-by-revenue assortments, we can construct a solution to problem (2) whose expected revenue deviates from the optimal expected revenue by no more than a factor of

$$\max_{i \in M, j=2, \dots, n} \left\{ \frac{R_i(N_{i,j-1})}{R_i(N_{ij})} \wedge \frac{R_i(N_{ij})}{R_i(N_{i,j-1})} \frac{V_i(N_{ij})^{\gamma_i}}{V_i(N_{i,j-1})^{\gamma_i}} \right\}, \quad (6)$$

where we let $a \wedge b = \min\{a, b\}$. Before we show this result, it is useful to observe the implications of this performance guarantee.

Assume that the revenues of the products within a nest are balanced in the sense that the largest and the smallest product revenues within a nest differ from each other by at most a factor of δ . Since

the preference weights of the no purchase options within the nests are zero, $R_i(\cdot)$ is always smaller than the largest product revenue in nest i and is always larger than the smallest product revenue. Therefore, the ratio $R_i(N_{ij})/R_i(N_{i,j-1})$ in the expression above cannot exceed δ . This observation implies that we can expect the nested-by-revenue assortments to perform well when the revenues of the products within a particular nest are balanced. Note that the revenues of the products in different nests can still differ from each other arbitrarily.

On the other hand, assume that the preference weights of the products within a nest are balanced in the sense that the largest and the smallest preference weights within a nest differ from each other by at most a factor of δ . In this case, since the assortments N_{ij} and $N_{i,j-1}$ respectively include j and $j-1$ products, the ratio $V_i(N_{ij})/V_i(N_{i,j-1})$ is bounded from above by $(j\delta)/(j-1)$ and the latter expression does not exceed 2δ for any $j = 2, \dots, n$. Furthermore, we observe that the expression in the curly brackets in (6) takes its largest value when the two terms of the minimum operator are equal to each other and this happens when

$$\frac{R_i(N_{i,j-1})}{R_i(N_{ij})} = \frac{V_i(N_{ij})^{\gamma_i/2}}{V_i(N_{i,j-1})^{\gamma_i/2}}.$$

Therefore, the value of the minimum operator in (6) is bounded from above by the expression on the right side above. Noting that we have $V_i(N_{ij})^{\gamma_i/2}/V_i(N_{i,j-1})^{\gamma_i/2} \leq (2\delta)^{\gamma_i/2}$, the performance guarantee we give for nested-by-revenue assortments cannot exceed $\max_{i \in M} (2\delta)^{\gamma_i/2}$. So, we can also expect the nested-by-revenue assortments to perform well when the preference weights of the products within a nest do not differ from each other drastically and the nest dissimilarity parameters are not too large. Similar to the discussion for the product revenues, the preference weights of the products in different nests can still differ from each other arbitrarily.

We make use of Theorem 1 to establish the performance guarantee that we give in (6). Assume that we solve problem (4) after replacing the collection of candidate assortments $\{A_{it} : t \in \mathcal{T}_i\}$ in the second set of constraints with the nested-by-revenue assortments $\{N_{ij} : j \in N_+\}$. Letting (\hat{x}, \hat{y}) be the optimal solution we obtain in this fashion and using α to denote the expression in (6), if we can show that $(\alpha \hat{x}, \alpha \hat{y})$ is a feasible solution to problem (3), then Theorem 1 implies that we can focus only on nested-by-revenue assortments and still obtain an assortment whose expected revenue deviates from the optimal expected revenue by no more than a factor of α .

To pursue this line of reasoning, we note that the second set of constraints in problem (3) can be written as $y_i \geq \max_{S_i \subset N} V_i(S_i)^{\gamma_i} (R_i(S_i) - x)$ for all $i \in M$. Using the decision variables $z_i = (z_{i1}, \dots, z_{in}) \in [0, 1]^n$, we formulate a tighter version of problem (3) as

$$\begin{aligned} \min \quad & x & (7) \\ \text{s.t.} \quad & v_0 x \geq \sum_{i \in M} y_i \\ & y_i \geq \max_{z_i \in [0, 1]^n} \left\{ \left(\sum_{j \in N} v_{ij} z_{ij} \right)^{\gamma_i} \left[\frac{\sum_{j \in N} r_{ij} v_{ij} z_{ij}}{\sum_{j \in N} v_{ij} z_{ij}} - x \right] \right\} \quad \forall i \in M. \end{aligned}$$

Note that if we imposed the constraint $z_i \in \{0, 1\}^n$ in the maximization problem on the right side of the second set of constraints above, then problems (3) and (7) would be equivalent to each other. The way

it is formulated, problem (3) is a relaxed version of problem (7) in the sense that any feasible solution to problem (7) is also a feasible solution to problem (3). In the following lemma, we study the optimal solution to the maximization problem on the right side of the second set of constraints above.

Lemma 6 *There exists an optimal solution z_i^* to the problem*

$$\max_{z_i \in [0,1]^n} \left\{ \left(\sum_{j \in N} v_{ij} z_{ij} \right)^{\gamma_i} \left[\frac{\sum_{j \in N} r_{ij} v_{ij} z_{ij}}{\sum_{j \in N} v_{ij} z_{ij}} - x \right] \right\} \quad (8)$$

such that $z_{i1}^* = 1, z_{i2}^* = 1, \dots, z_{i,k-1}^* = 1, z_{ik}^* \in [0, 1], z_{i,k+1}^* = 0, \dots, z_{in}^* = 0$ for some $k = 1, \dots, n$.

Proof. Assume that \hat{z}_i is an optimal solution to problem (8) and let $C = \sum_{j \in N} v_{ij} \hat{z}_{ij}$. In this case, note that an optimal solution to the continuous knapsack problem

$$\max \left\{ \sum_{j \in N} r_{ij} v_{ij} z_{ij} : \sum_{j \in N} v_{ij} z_{ij} = C, z_i \in [0, 1]^n \right\}$$

is also an optimal solution to problem (8). In the continuous knapsack problem above, the utility of item j is $r_{ij} v_{ij}$ and the space consumption of item j is v_{ij} . Thus, we can solve this problem by sorting the products with respect to their utility-to-space consumption ratios and filling the knapsack starting from the item with the largest utility-to-space consumption ratio. Since the utility-to-space consumption ratio of item j is r_{ij} and the products are ordered such that $r_{i1} \geq r_{i2} \geq \dots \geq r_{in}$, there exists an optimal solution to the continuous knapsack problem above with the form given in the lemma. \square

Except for at most one possible fractional component, Lemma 6 shows that a nested-by-revenue assortment is optimal for the maximization problem on the right side of the second set of constraints in problem (7). Thus, it is not too surprising that if we solve problem (4) after replacing the collection of candidate assortments $\{A_{it} : t \in \mathcal{T}_i\}$ in the second set of constraints with the nested-by-revenue assortments $\{N_{ij} : j \in N_+\}$ and obtain the optimal solution (\hat{x}, \hat{y}) , then (\hat{x}, \hat{y}) is almost feasible to problem (7). Since problem (3) is a relaxed version of problem (7), the solution (\hat{x}, \hat{y}) would be almost feasible to problem (3) as well. We make this intuitive argument precise in the following theorem and show that nested-by-revenue assortments provide the performance guarantee that we give in (6). We defer the proof of this result to the appendix.

Theorem 7 *Let (\hat{x}, \hat{y}) be an optimal solution to problem (4) when we solve this problem after replacing the collection of candidate assortments $\{A_{it} : t \in \mathcal{T}_i\}$ in the second set of constraints with the nested-by-revenue assortments $\{N_{ij} : j \in N_+\}$. Then, using α to denote the expression in (6), $(\alpha \hat{x}, \alpha \hat{y})$ is a feasible solution to problem (3).*

Theorem 7, along with Theorem 1, shows that if we only consider the nested-by-revenue assortments as candidate assortments, then we can construct a solution to problem (2) whose expected revenue deviates from the optimal expected revenue by at most a factor given in (6). Similar to the discussion at

the end of Section 3, to find the best combination of nested-by-revenue assortments, we can solve problem (4) after replacing the collection of candidate assortments $\{A_{it} : t \in \mathcal{T}_i\}$ with the nested-by-revenue assortments $\{N_{ij} : j \in N_+\}$. This amounts to solving a linear program with $1 + m$ decision variables and $1 + m(1 + n)$ constraints.

5 Competitive Products and Partially-Captured Nests

In this section, we consider instances of the assortment optimization problem in (2) with $\gamma_i \leq 1$ for all $i \in M$, but the preference weights of the no purchase options within the nests can take on arbitrary values. In other words, we allow (v_{10}, \dots, v_{m0}) to take on strictly positive values so that a customer can leave without purchasing anything even after this customer chooses a particular nest. For this case, nested-by-revenue assortments are no longer optimal. As a matter of fact, problem (2) turns out to be NP-hard. However, we are able to characterize a small class of assortments such that if we focus on this class of assortments, then we can construct a solution to problem (2) whose expected revenue deviates from the optimal expected revenue by no more than a factor of two.

In the following theorem, we show that the assortment optimization problem in (2) is NP-hard when we have $v_{i0} > 0$ for some $i \in M$. The proof technique we use is similar to the one in Section 4.1. In particular, we consider the assortment feasibility problem as defined in Section 4.1 and show that any instance of the partition problem can be reduced to an instance of the assortment feasibility problem. However, the specifics of the reduction are somewhat more involved. We give the proof of this theorem in the appendix.

Theorem 8 *If we allow the preference weights of the no purchase options within the nests to take on strictly positive values, then the assortment feasibility problem is NP-hard.*

Motivated by this computational complexity result, we turn our attention to obtaining approximate solutions. In this section, we develop a tractable algorithm that obtains a solution to problem (2) whose expected revenue deviates from the optimal expected revenue by at most a factor of two. To show this, we use an alternative representation of problem (3). Using the decision variables $z_i = (z_{i1}, \dots, z_{in})$, we define $K_i(\epsilon_i)$ as the optimal objective value of the knapsack problem

$$K_i(\epsilon_i) = \max \left\{ \sum_{j \in N} r_{ij} v_{ij} z_{ij} : \sum_{j \in N} v_{ij} z_{ij} \leq \epsilon_i, z_i \in \{0, 1\}^n \right\}. \quad (9)$$

In this case, noting that $V_i(S_i)$ and $R_i(S_i)$ in the second set of constraints in problem (3) are respectively given by $V_i(S_i) = v_{i0} + \sum_{j \in S_i} v_{ij}$ and $R_i(S_i) = \sum_{j \in S_i} r_{ij} v_{ij} / V_i(S_i)$, we consider the problem

$$\begin{aligned} \min \quad & x \\ \text{s.t.} \quad & v_0 x \geq \sum_{i \in M} y_i \\ & y_i \geq \max_{\epsilon_i \geq 0} \left\{ (v_{i0} + \epsilon_i)^{\gamma_i} \left[\frac{K_i(\epsilon_i)}{v_{i0} + \epsilon_i} - x \right] \right\} \quad \forall i \in M. \end{aligned} \quad (10)$$

The following lemma shows that problem (10) is equivalent to problem (3).

Lemma 9 Problems (3) and (10) are equivalent to each other in the sense that an optimal solution to one problem is also an optimal solution to the other.

Proof. Noting that the second set of constraints in problem (3) can equivalently be written as $y_i \geq \max_{S_i \subset N} V_i(S_i)^{\gamma_i} (R_i(S_i) - x)$ for all $i \in M$, the result follows if we can show that

$$\max_{S_i \subset N} V_i(S_i)^{\gamma_i} (R_i(S_i) - x) = \max_{\epsilon_i \geq 0} \left\{ (v_{i0} + \epsilon_i)^{\gamma_i} \left[\frac{K_i(\epsilon_i)}{v_{i0} + \epsilon_i} - x \right] \right\}$$

for any $x \geq 0$. Let ζ_L^* and ζ_R^* respectively be the optimal objective values of the problems on the left and right side above. First, we show that $\zeta_L^* \leq \zeta_R^*$. Assume that S_i^* is an optimal solution to the problem on the left side above and define $\epsilon_i^* = \sum_{j \in S_i^*} v_{ij}$. The solution obtained by setting $z_{ij} = 1$ for all $j \in S_i^*$ and $z_{ij} = 0$ otherwise is a feasible solution to problem (9) with $\epsilon_i = \epsilon_i^*$, which implies that $K_i(\epsilon_i^*) \geq \sum_{j \in S_i^*} r_{ij} v_{ij}$. Thus, if we evaluate the objective value of the problem on the right side above at $\epsilon_i = \epsilon_i^*$, then we obtain at least ζ_L^* , in which case, we obtain $\zeta_R^* \geq \zeta_L^*$.

Second, we show that $\zeta_L^* \geq \zeta_R^*$. Let ϵ_i^* be an optimal solution to the problem on the right side above and solve problem (9) after setting $\epsilon_i = \epsilon_i^*$. Letting z_i^* be the solution we obtain, we observe that we can assume without loss of generality that $\sum_{j \in N} v_{ij} z_{ij}^* = \epsilon_i^*$. To see this claim, if we have $\sum_{j \in N} v_{ij} z_{ij}^* < \epsilon_i^*$, then we can decrease the value of ϵ_i^* to $\hat{\epsilon}_i = \sum_{j \in N} v_{ij} z_{ij}^*$ while still preserving $K_i(\epsilon_i^*) = K_i(\hat{\epsilon}_i)$. In this case, using the fact that $\gamma_i \leq 1$ and $x \geq 0$, we obtain

$$\begin{aligned} (v_{i0} + \epsilon_i^*)^{\gamma_i} \left[\frac{K_i(\epsilon_i^*)}{v_{i0} + \epsilon_i^*} - x \right] &= \frac{K_i(\epsilon_i^*)}{(v_{i0} + \epsilon_i^*)^{1-\gamma_i}} - (v_{i0} + \epsilon_i^*)^{\gamma_i} x \\ &\leq \frac{K_i(\hat{\epsilon}_i)}{(v_{i0} + \hat{\epsilon}_i)^{1-\gamma_i}} - (v_{i0} + \hat{\epsilon}_i)^{\gamma_i} x = (v_{i0} + \hat{\epsilon}_i)^{\gamma_i} \left[\frac{K_i(\hat{\epsilon}_i)}{v_{i0} + \hat{\epsilon}_i} - x \right], \end{aligned}$$

which shows that $\hat{\epsilon}_i$ should also be an optimal solution to the problem on the right side above, establishing the claim. By using the solution z_i^* , we define the assortment S_i^* as $S_i^* = \{j \in N : z_{ij}^* = 1\}$. Since $\sum_{j \in S_i^*} v_{ij} = \sum_{j \in N} v_{ij} z_{ij}^* = \epsilon_i^*$ and $\sum_{j \in S_i^*} r_{ij} v_{ij} = \sum_{j \in N} r_{ij} v_{ij} z_{ij}^* = K_i(\epsilon_i^*)$, the assortment S_i^* provides an objective value of ζ_R^* for the problem on the left side above and we obtain $\zeta_L^* \geq \zeta_R^*$. \square

To exploit the equivalence between problems (3) and (10) in a tractable fashion, we use the continuous relaxation of the knapsack problem in (9), which is given by

$$\hat{K}_i(\epsilon_i) = \max \left\{ \sum_{j \in N} r_{ij} v_{ij} z_{ij} : \sum_{j \in N} v_{ij} z_{ij} \leq \epsilon_i, 0 \leq z_{ij} \leq \mathbf{1}(v_{ij} \leq \epsilon_i) \forall j \in N \right\}, \quad (11)$$

where we use $\mathbf{1}(\cdot)$ to denote the indicator function. Since problem (11) is a relaxation of problem (9), we have $\hat{K}_i(\epsilon_i) \geq K_i(\epsilon_i)$. The problem above is a continuous knapsack problem, where the utility of item j is $r_{ij} v_{ij}$, the space consumption of item j is v_{ij} and we can only consider the items whose space consumptions do not exceed ϵ_i . Noting that the utility-to-space consumption ratio of item j is r_{ij} , we can solve this problem by sorting the products with respect to their revenues and filling the knapsack starting from the product with the largest revenue, as long as we only consider the products whose preference weights do not exceed ϵ_i . We let $\hat{z}_i(\epsilon_i) = (\hat{z}_{i1}(\epsilon_i), \dots, \hat{z}_{in}(\epsilon_i))$ be an optimal solution to problem (11)

that we obtain in this fashion. We observe that $\hat{z}_i(\epsilon_i)$ has at most one fractional component. By using this solution, we define the assortment $\hat{S}_i(\epsilon_i)$ as $\hat{S}_i(\epsilon_i) = \{j \in N : \hat{z}_{ij}(\epsilon_i) = 1\}$, which includes only the strictly positive and integer-valued components of $\hat{z}_i(\epsilon_i)$. We use the assortments $\{\hat{S}_i(\epsilon_i) : \epsilon_i \in [0, \infty]\}$ as a collection of candidate assortments for nest i . We shortly show in this section that this collection of assortments includes no more than $1 + n^2$ assortments and each one of these $1 + n^2$ assortments can be identified in a tractable fashion.

To be able to obtain the performance guarantee of two, we augment the collection of assortments $\{\hat{S}_i(\epsilon_i) : \epsilon_i \in [0, \infty]\}$ for nest i by the collection of singleton assortments $\{\{j\} : j \in N\}$. We solve problem (4) after replacing the collection of candidate assortments $\{A_{it} : t \in \mathcal{T}_i\}$ in the second set of constraints with the collection of assortments $\{\hat{S}_i(\epsilon_i) : \epsilon_i \in [0, \infty]\} \cup \{\{j\} : j \in N\}$. Letting (\hat{x}, \hat{y}) be the optimal solution to problem (4) that we obtain in this fashion, if we can show that $(2\hat{x}, 2\hat{y})$ is a feasible solution to problem (3), then Theorem 1 implies that we can focus only on the assortments $\{\hat{S}_i(\epsilon_i) : \epsilon_i \in [0, \infty]\} \cup \{\{j\} : j \in N\}$ for nest i and still obtain an assortment whose expected revenue deviates from the optimal expected revenue by no more than a factor of two. We pursue this result in the following theorem, but defer the proof of this result to the appendix.

Theorem 10 *Let (\hat{x}, \hat{y}) be an optimal solution to problem (4) when we solve this problem after replacing the collection of candidate assortments $\{A_{it} : t \in \mathcal{T}_i\}$ in the second set of constraints with the assortments $\{\hat{S}_i(\epsilon_i) : \epsilon_i \in [0, \infty]\} \cup \{\{j\} : j \in N\}$. Then, $(2\hat{x}, 2\hat{y})$ is a feasible solution to problem (3).*

The equivalence between problems (3) and (10) given in Lemma 9 lays out the connection between the assortment optimization problem we are interested in and the knapsack problem, as long as we have $\gamma_i \leq 1$ for all $i \in M$. It is well-known that one can construct a solution to the knapsack problem in (9) by using its continuous relaxation in (11) and the objective value of this solution would deviate from the optimal objective value of the knapsack problem by no more than a factor of two; see Williamson and Shmoys (2011). The proof of Theorem 10 implicitly makes use of this result. Similarly, there exists a well-known fully polynomial-time approximation scheme for the knapsack problem, which can be found in Williamson and Shmoys (2011). Building on this fully polynomial-time approximation scheme, it is indeed possible to develop a fully polynomial-time approximation scheme for the assortment optimization problem we are interested in. We do not pursue the fully polynomial-time approximation scheme because this scheme can be developed by using an argument that is very similar to the preceding discussion in this section. The only difference is that instead of constructing approximate solutions to the knapsack problem in (9) by using its continuous relaxation, we construct approximate solutions by using the fully polynomial-time approximation scheme for knapsack problems.

In the remainder of this section, we argue that the collection of assortments $\{\hat{S}_i(\epsilon_i) : \epsilon_i \in [0, \infty]\}$ includes no more than $1 + n^2$ assortments and each one of these $1 + n^2$ assortments can be identified in a tractable fashion. Fix ϵ_i and consider the assortment $\hat{S}_i(\epsilon_i)$. The definition of $\hat{z}_i(\epsilon_i)$ implies that $\hat{S}_i(\epsilon_i)$ is a nested-by-revenue assortment as long as we focus only on the products whose preference weights do not exceed ϵ_i . In other words, letting k be the number of products whose preference weights do not exceed ϵ_i , $\hat{S}_i(\epsilon_i)$ is a nested-by-revenue assortment as long as we focus only on the products with the k

smallest preference weights. Given that we focus only on the products with the k smallest preference weights, we use N_{ij}^k to denote the nested-by-revenue assortment that includes the first j products with the largest revenues in nest i . Therefore, $\hat{S}_i(\epsilon_i)$ must be one of the assortments $\{N_{ij}^k : j = 0, \dots, k\}$, where we let $N_{i0}^k = \emptyset$ for notational uniformity. Since the only possible values for k are $k = 1, \dots, n$, it follows that $\{\hat{S}_i(\epsilon_i) : \epsilon_i \in [0, \infty]\} \subset \{N_{ij}^k : k \in N, j = 0, \dots, k\}$. The latter collection of assortments includes no more than $1 + n^2$ distinct assortments, all of which can be easily be identified.

Theorem 10 shows that if we use the assortments $\{\hat{S}_i(\epsilon_i) : \epsilon_i \in [0, \infty]\} \cup \{\{j\} : j \in N\}$ as a collection of candidate assortments for nest i , then we can stitch together from these candidate assortments a solution to problem (2) whose expected revenue deviates from the optimal expected revenue by at most a factor of two. Since we have $\{\hat{S}_i(\epsilon_i) : \epsilon_i \in [0, \infty]\} \subset \{N_{ij}^k : k \in N, j = 0, \dots, k\}$, using the assortments $\{N_{ij}^k : k \in N, j = 0, \dots, k\} \cup \{\{j\} : j \in N\}$ as candidate assortments for nest i cannot degrade the performance guarantee of two. To find the best combination of these assortments, we can solve problem (4) after replacing the collection of candidate assortments $\{A_{it} : t \in \mathcal{T}_i\}$ with the assortments $\{N_{ij}^k : k \in N, j = 0, \dots, k\} \cup \{\{j\} : j \in N\}$. This amounts to solving a linear program with $1 + m$ decision variables and $1 + m(1 + n + n^2)$ constraints.

6 Possibly Synergistic Products and Partially-Captured Nests

In this section, we consider the most general instances of the assortment optimization problem in (2), where we do not have any restrictions on the dissimilarity parameters $(\gamma_1, \dots, \gamma_m)$ of the nests and the preference weights (v_{10}, \dots, v_{m0}) of the no purchase options within the nests. For these most general instances, we develop a tractable algorithm that obtains a solution to problem (2) whose expected revenue deviates from the optimal expected revenue by no more than a factor of $\max_{i \in M} 2^{2\gamma_i+1}$. Without loss of generality, we assume that $\max_{i \in M} \gamma_i > 1$ throughout this section. Otherwise, we can use the algorithm given in the previous section.

Similar to the previous section, our approach is based on viewing problem (3) from a knapsack perspective. We define $G_i(\epsilon_i)$ as the optimal objective value of the knapsack problem

$$G_i(\epsilon_i) = \max_{S_i \subset N} \left\{ \sum_{j \in S_i} r_{ij} v_{ij} : v_{i0} + \sum_{j \in S_i} v_{ij} = \epsilon_i \right\}, \quad (12)$$

where we let $G_i(\epsilon_i) = -\infty$ when the problem on the right side above is infeasible. The second set of constraints in problem (3) are given by $y_i \geq V_i(S_i)^{\gamma_i} (\sum_{j \in S_i} r_{ij} v_{ij} / V_i(S_i) - x)$ for all $S_i \subset N$, $i \in M$ and problem (12) finds the largest value of $\sum_{j \in S_i} r_{ij} v_{ij}$ while keeping $V_i(S_i)$ constant at ϵ_i . Thus, problem (3) can equivalently be written as

$$\begin{aligned} \min \quad & x & (13) \\ \text{s.t.} \quad & v_0 x \geq \sum_{i \in M} y_i \\ & y_i \geq \epsilon_i^{\gamma_i} \left[\frac{G_i(\epsilon_i)}{\epsilon_i} - x \right] \quad \forall \epsilon_i \geq 0, i \in M, \end{aligned}$$

where we treat $0/0$ as zero in the second set of constraints. Letting $v_i^L = v_{i0} + \min_{j \in N} v_{ij}$ and $v_i^U = v_{i0} + \sum_{j \in N} v_{ij}$, the definition of $G_i(\epsilon_i)$ in (12) implies that $G_i(\epsilon_i) = G_i(v_{i0})$ or $G_i(\epsilon_i) = -\infty$ for all $\epsilon_i < v_i^L$, whereas $G_i(\epsilon_i) = -\infty$ for all $\epsilon_i > v_i^U$. Therefore, it is enough to consider the values of ϵ_i with $\epsilon_i = v_{i0}$ or $\epsilon_i \in [v_i^L, v_i^U]$ in the second set of constraints in problem (13).

Although problem (13) is equivalent to problem (3), it is not yet computationally helpful because computing $G_i(\epsilon_i)$ is difficult. The algorithm we develop in this section is based on constructing tractable approximations to $G_i(\epsilon_i)$. In particular, instead of considering every single possible value of ϵ_i with $\epsilon_i = v_{i0}$ or $\epsilon_i \in [v_i^L, v_i^U]$ in the second set of constraints in problem (13), we consider the values of ϵ_i that are close to the powers of two. To do this, we define l_i^L as $l_i^L = \min\{l \in \mathbb{Z} : 2^l \geq v_i^L\}$ and $l_i^U = \min\{l \in \mathbb{Z} : 2^l \geq v_i^U\}$ so that we have $[v_i^L, v_i^U] \subset [2^{l_i^L-1}, 2^{l_i^U}]$. In this case, whenever ϵ_i lies in the interval $[2^{l-1}, 2^l]$ for some $l = l_i^L, \dots, l_i^U$, we approximate $G_i(\epsilon_i)$ by

$$\hat{G}_{il} = \max_{S_i \subset N} \left\{ \sum_{j \in S_i} r_{ij} v_{ij} : 2^{l-1} \leq v_{i0} + \sum_{j \in S_i} v_{ij} \leq 2^l \right\}. \quad (14)$$

Problem (14) is still difficult because it is a knapsack problem with both upper and lower bounds, but it turns out we can compute approximate solutions to this problem without too much difficulty. In particular, Proposition 14 in the appendix shows that we can easily find an assortment \hat{S}_{il} that satisfies $2^{l-1} \leq v_{i0} + \sum_{j \in \hat{S}_{il}} v_{ij} \leq 2^l$ and $2 \sum_{j \in \hat{S}_{il}} r_{ij} v_{ij} \geq \hat{G}_{il}$. In other words, the assortment \hat{S}_{il} is a feasible solution to problem (14) and the objective value provided by this assortment deviates from the optimal objective value of problem (14) by no more than a factor of two. In this way, we have an approximation to $G_i(\epsilon_i)$ for any $\epsilon_i \in [v_i^L, v_i^U]$. It is, of course, trivial to approximate $G_i(\epsilon_i)$ when $\epsilon_i = v_{i0}$ by noting that we have $G_i(v_{i0}) = 0$ by definition.

We propose using the assortments $\{\hat{S}_{il} : l = l_i^L, \dots, l_i^U\} \cup \{\emptyset\}$ as a collection of candidate assortments for nest i . To find the best combination of such assortments, we solve problem (4) after replacing the collection of candidate assortments $\{A_{it} : t \in \mathcal{T}_i\}$ in the second set of constraints with the collection of assortments $\{\hat{S}_{il} : l = l_i^L, \dots, l_i^U\} \cup \{\emptyset\}$. Letting (\hat{x}, \hat{y}) be the optimal solution to problem (4) that we obtain in this fashion and setting $\bar{\gamma} = \max_{i \in M} \gamma_i$ for notational convenience, if we can show that $(2^{2\bar{\gamma}+1} \hat{x}, 2^{2\bar{\gamma}+1} \hat{y})$ is a feasible solution to problem (3), then Theorem 1 implies that we can focus only on the assortments $\{\hat{S}_{il} : l = l_i^L, \dots, l_i^U\} \cup \{\emptyset\}$ for nest i and still obtain an assortment whose expected revenue deviates from the optimal expected revenue by no more than a factor of $2^{2\bar{\gamma}+1}$. We establish this result in the following theorem. The proof of this theorem is given in the appendix.

Theorem 11 *Let (\hat{x}, \hat{y}) be an optimal solution to problem (4) when we solve this problem after replacing the collection of candidate assortments $\{A_{it} : t \in \mathcal{T}_i\}$ in the second set of constraints with the assortments $\{\hat{S}_{il} : l = l_i^L, \dots, l_i^U\} \cup \{\emptyset\}$. Then, $(2^{2\bar{\gamma}+1} \hat{x}, 2^{2\bar{\gamma}+1} \hat{y})$ is a feasible solution to problem (3).*

Theorem 11 implies that if we use the assortments $\{\hat{S}_{il} : l = l_i^L, \dots, l_i^U\} \cup \{\emptyset\}$ as candidate assortments for nest i , then we can combine these assortments for the different nests to obtain a solution to problem (2) whose expected revenue deviates from the optimal expected revenue by no more than a factor of $2^{2\bar{\gamma}+1}$. To find the best combination of these assortments, we need to solve problem (4) after

replacing the collection of candidate assortments $\{A_{it} : t \in \mathcal{T}_i\}$ in the second set of constraints with the collection of assortments $\{\hat{S}_{il} : l = l_i^L, \dots, l_i^U\} \cup \{\emptyset\}$. Since we have $2^{l_i^L} \geq v_i^L$ and $2^{l_i^U - 1} \leq v_i^U$, we have $l_i^U - l_i^L \leq 1 + \log_2(v_i^U/v_i^L)$. Therefore, for the most general instances of the assortment optimization problem we are interested in, we can solve a linear program with $1 + m$ decision variables and at most $1 + m(2 + \log_2(v_i^U/v_i^L))$ constraints to find an assortment whose expected revenue deviates from the optimal expected revenue by no more than a factor of $2^{2\gamma+1}$.

7 Upper Bounds on Optimal Expected Revenue

In the previous three sections, we give collections of candidate assortments such that if we focus on these collections, then we can stitch together an assortment with a certain performance guarantee. The performance guarantees we give reflect the worst-case performance of a given collection of assortments, where the worst-case is taken over all possible problem instances. In this section, our goal is to develop a tractable upper bound on the optimal expected revenue that we can compute for each individual problem instance. We can then use the collections of assortments given in the previous sections to obtain the best possible assortment and compare the expected revenue from this assortment with the problem instance-specific upper bound on the optimal expected revenue. In this way, we can bound the optimality gap of the assortment we obtain for a particular problem instance.

To construct an upper bound on the optimal expected revenue, we use a tighter version of problem (3) that is similar to the one in Section 4.2. In particular, since the second set of constraints in problem (3) can be written as $y_i \geq \max_{S_i \subset N} V_i(S_i)^{\gamma_i} (R_i(S_i) - x)$ for all $i \in M$, using the decision variables $z_i = (z_{i1}, \dots, z_{in}) \in [0, 1]^n$, we formulate a tighter version of problem (3) as

$$\begin{aligned} \min \quad & x & (15) \\ \text{s.t.} \quad & v_0 x \geq \sum_{i \in M} y_i \\ & y_i \geq \max_{z_i \in [0, 1]^n} \left\{ \left(v_{i0} + \sum_{j \in N} v_{ij} z_{ij} \right)^{\gamma_i} \left[\frac{\sum_{j \in N} r_{ij} v_{ij} z_{ij}}{v_{i0} + \sum_{j \in N} v_{ij} z_{ij}} - x \right] \right\} \quad \forall i \in M. \end{aligned}$$

Problem (15) is a tighter version of problem (3) as any feasible solution to problem (15) is a feasible solution to problem (3). Therefore, if we can solve problem (15) in a tractable fashion, then the optimal objective value of this problem provides an upper bound on the optimal expected revenue.

To see how we can solve problem (15) in a tractable fashion, for a fixed value of x , we use $F_i(z_i | x)$ to denote objective function of the maximization problem on the right side of the second set of constraints in problem (15) and let $\hat{F}_i(x) = \max_{z_i \in [0, 1]^n} F_i(z_i | x)$. In this case, problem (15) is equivalent to

$$\begin{aligned} \min \quad & x & (16) \\ \text{s.t.} \quad & v_0 x \geq \sum_{i \in M} y_i \\ & y_i \geq \hat{F}_i(x) \quad \forall i \in M. \end{aligned}$$

If we can show that $\hat{F}_i(x)$ is a convex function of x , then the feasible region of problem (16) ends up being convex. Thus, noting that the objective function is linear, problem (16) becomes a convex optimization

problem. Furthermore, if we can obtain subgradients of $\hat{F}_i(x)$ with respect to x in a tractable fashion, then we can use a standard cutting plane method for convex optimization to solve problem (16); see Ruszczyński (2006). The following proposition follows a standard argument in nonlinear programming to show that $\hat{F}_i(x)$ is indeed a convex function of x and demonstrates how to obtain subgradients of this function with respect to x .

Proposition 12 $\hat{F}_i(x)$ is convex. Furthermore, if we let $\hat{z}_i(x)$ be an optimal solution to the problem $\max_{z_i \in [0,1]^n} F_i(z_i | x)$, then a subgradient of $\hat{F}_i(\cdot)$ at x is given by $-(v_{i0} + \sum_{j \in N} v_{ij} \hat{z}_{ij}(x))^{\gamma_i}$.

Proof. By the definitions of $\hat{F}_i(x)$ and $\hat{z}_i(x)$, it follows that $\hat{F}_i(x) = F_i(\hat{z}_i(x) | x)$ and $\hat{F}_i(x') \geq F_i(\hat{z}_i(x) | x')$. Subtracting the equality from the inequality and noting that $F_i(\hat{z}_i(x) | x') - F_i(\hat{z}_i(x) | x) = -(v_{i0} + \sum_{j \in N} v_{ij} \hat{z}_{ij}(x))^{\gamma_i} (x' - x)$, we obtain

$$\hat{F}_i(x') \geq \hat{F}_i(x) - \left(v_{i0} + \sum_{j \in N} v_{ij} \hat{z}_{ij}(x) \right)^{\gamma_i} (x' - x),$$

which implies that $\hat{F}_i(\cdot)$ satisfies the subgradient inequality at x with a subgradient given by $-(v_{i0} + \sum_{j \in N} v_{ij} \hat{z}_{ij}(x))^{\gamma_i}$. In this case, by Theorem 3.2.6 in Bazaraa et al. (1993), $\hat{F}_i(x)$ is a convex function of x with a subgradient as given in the lemma. \square

Proposition 12 shows that we can obtain a subgradient of $\hat{F}_i(\cdot)$ at x by solving the problem

$$\max_{z_i \in [0,1]^n} F_i(z_i | x) = \max_{z_i \in [0,1]^n} \left\{ \left(v_{i0} + \sum_{j \in N} v_{ij} z_{ij} \right)^{\gamma_i} \left[\frac{\sum_{j \in N} r_{ij} v_{ij} z_{ij}}{v_{i0} + \sum_{j \in N} v_{ij} z_{ij}} - x \right] \right\} \quad (17)$$

at a fixed value of x . Thus, if we can solve the problem above in a tractable fashion, then we can also solve problem (16) through a standard cutting plane method. In the rest of this section, we focus on solving problem (17) at a fixed value of x .

Following the same argument in the proof of Lemma 6, we can show that there exists an optimal solution z_i^* to problem (17) that satisfies $z_{i1}^* = 1, z_{i2}^* = 1, \dots, z_{i,k-1}^* = 1, z_{ik}^* \in [0, 1], z_{i,k+1}^* = 0, \dots, z_{in}^* = 0$ for some $k = 1, \dots, n$. In other words, if we define the vector $\delta^k = (\delta_1^k, \dots, \delta_n^k) \in \mathfrak{R}^n$ such that $\delta_1^k = 1, \dots, \delta_{k-1}^k = 1, \delta_k^k = 0, \dots, \delta_n^k = 0$ and use e^k to denote the k -th unit vector in \mathfrak{R}^n , then an optimal solution to problem (17) is of the form $\delta^k + \rho e^k$ for some $\rho \in [0, 1]$ and $k = 1, \dots, n$. To find the best value for ρ , we can solve the problem $\max_{\rho \in [0,1]} F_i(\delta^k + \rho e^k | x)$, which is a scalar optimization problem. A simple check verifies that the first derivative of $F_i(\delta^k + \rho e^k | x)$ with respect to ρ vanishes at one point and the point at which the first derivative vanishes can be computed in closed-form fashion. Therefore, the optimal objective value of the problem $\max_{\rho \in [0,1]} F_i(\delta^k + \rho e^k | x)$ is attained either at one of the two end points of the interval $[0, 1]$ or at the point where the first derivative of $F_i(\delta^k + \rho e^k | x)$ with respect to ρ vanishes. By checking the objective value of the problem $\max_{\rho \in [0,1]} F_i(\delta^k + \rho e^k | x)$ at these three points, we can easily solve this problem. In this case, the optimal objective value of problem (17) can be obtained by solving the problem $\max_{\rho \in [0,1]} F_i(\delta^k + \rho e^k | x)$ for all $k = 1, \dots, n$ and picking the one that yields the largest optimal objective value.

8 Computational Experiments

We have two goals in our computational experiments. First, we are interested in testing the quality of the assortments we obtain by focusing on various collections of candidate assortments. Second, we want to check the quality of the upper bounds on the optimal expected revenue that we obtain from problem (15). We first describe the collections of candidate assortments we consider and provide the details of our experimental setup. We then give the findings from our computational experiments.

8.1 Candidate Assortments

In our computational experiments, we consider three collections of candidate assortments.

Nested-by-revenue assortments. Nested-by-revenue assortments refer to the collection of assortments $\{N_{ij} : j \in N_+\}$ for nest i . For each nest i , there are $1 + n$ nested-by-revenue assortments. We use this collection of assortments in Sections 3 and 4, where we show that they are optimal for problem instances with competitive products and fully-captured nests and they provide the performance guarantee in (6) for problem instances with possibly synergistic products and fully-captured nests. We are not able to establish performance guarantees for these assortments for more general problem instances. However, since nested-by-revenue assortments are intuitively appealing and easy to implement, it is useful to test their performance for more general problem instances as well.

Nested-by-preference-and-revenue assortments. Nested-by-preference-and-revenue assortments refer to the collection of assortments $\{N_{ij}^k : k \in N, j = 0, \dots, k\} \cup \{\{j\} : j \in N\}$ for nest i . For each nest i , there are no more than $1 + n + n^2$ nested-by-preference-and-revenue assortments. We use this collection of assortments in Section 5, where we show that they provide a performance guarantee of two for problem instances with competitive products and partially-captured nests. We observe that the nested-by-preference-and-revenue assortment N_{ij}^n is the same as the nested-by-revenue assortment N_{ij} . Therefore, the collection of nested-by-preference-and-revenue assortments include nested-by-revenue assortments and it is worthwhile to see how much additional benefit we obtain by enlarging the collection of candidate assortments from nested-by-revenue to nested-by-preference-and-revenue.

Powers-of-two assortments. Powers-of-two assortments correspond to the collection of assortments $\{\hat{S}_{il} : l = l_i^L, \dots, l_i^U\} \cup \{\emptyset\}$ for nest i . For each nest i , there are no more than $2 + \log_2(v_i^U/v_i^L)$ such assortments. Noting that $v_i^U = v_{i0} + \sum_{j \in N} v_{ij}$ and $v_i^L = v_{i0} + \min_{j \in N} v_{ij}$, we observe that the number of powers-of-two assortments grows logarithmically with n , whereas the numbers of nested-by-revenue and nested-by-preference-and-revenue assortments grow polynomially with n . We consider powers-of-two assortments in Section 6 and show that these assortments provide a performance guarantee of $\max_{i \in M} 2^{2\gamma_i+1}$ for the most general instances of the assortment optimization problem.

8.2 Experimental Setup

We provide computational results for test problems with (i) possibly synergistic products and fully-captured nests, (ii) competitive products and partially-captured nests, and (iii) possibly synergistic

products and partially-captured nests. We do not present results for test problems with competitive products and fully-captured nests since nested-by-revenue assortments are optimal for them. In our test problems, the number of nests is $m = 5$ and the number of products in each nest is $n = 20$. For the test problems with competitive products, we generate the dissimilarity parameters of the nests from the uniform distribution over $[0.25, 0.75]$, whereas for the test problems with possibly synergistic products, we generate the dissimilarity parameters from the uniform distribution over $[1.5, 2.5]$.

To come up with the preference weights and the revenues associated with the products, we generate U_{ij} from the uniform distribution over $[0, 1]$ for all $i \in M, j \in N$. We also generate W_{ij} and Y_{ij} from the uniform distribution over $[\zeta^L, \zeta^U]$ for all $i \in M, j \in N$, where ζ^L and ζ^U are parameters we vary. In this case, we set the preference weight associated with product j in nest i as $v_{ij} = 10 U_{ij}^2 W_{ij}$, whereas we set the revenue of product j in nest i as $r_{ij} = 10 (1 - U_{ij})^\kappa Y_{ij}$, where κ is another parameter we vary. The role of U_{ij} is to introduce negative correlation between the preference weights and the revenues so that the products with larger revenues tend to have smaller preference weights. On the other hand, W_{ij} and Y_{ij} introduce some idiosyncratic noise in the preference weights and revenues so that it is not always the case the expensive products have smaller preference weights. The parameter κ skews the distribution of product revenues. As κ gets larger, we have a larger number of products with small revenues.

As far as the preference weights of the no purchase options are concerned, we have $v_{i0} = 0$ for all $i \in M$ for the test problems with fully-captured nests, in which case, we only need to generate v_0 . For the test problems with partially-captured nests, we set $v_0 = 0$ so that we only need to generate v_{i0} for all $i \in M$. To generate the preference weights of the no purchase options, we calibrate v_0 and v_{i0} so that best nested-by-revenue assortments include about half of the products. Our experience is that if a good assortment includes too few products, then we can find a good solution by enumerating all assortments with a limited number of products in it. Similarly, if a good assortment includes too many products, then we can find a good assortment by enumerating all assortments with a large number of products. Thus, test problems that include a moderate number of products in a good assortment appear to be more challenging. For the test problems with possibly synergistic products and fully-captured nests, we set $v_0 = 0.5$. For the test problems with competitive products and partially-captured nests, we set $v_{i0} = 15$ for all $i \in M$. Finally, for the test problems with possibly synergistic products and partially-captured nests, we set $v_{i0} = 0.5$ for all $i \in M$. With these settings, a few setup runs indicated that best nested-by-revenue assortments include about half of the products.

We vary $[\zeta^L, \zeta^U]$ over $\{[1.0, 1.0], [0.8, 1.2], [0.5, 1.5]\}$, corresponding to low, medium and high levels of idiosyncratic noise. We vary κ over $\{1, 2\}$, corresponding to two levels of skewness in the product revenues. For each combination of parameters, we generate 50,000 individual problem instances and test the three collections of assortments described in Section 8.1 on each problem instance.

8.3 Computational Results

We give our computational results in Tables 2, 3 and 4, where Tables 2, 3 and 4 respectively show the performance of the nested-by-revenue, nested-by-preference-and-revenue and powers-of-two

assortments. In these tables, the first column shows the category of the problem indicating whether we focus on the case with (i) possibly synergistic products and fully-captured nests, (ii) competitive products and partially-captured nests, or (iii) possibly synergistic products and partially-captured nests. The second column shows the level of idiosyncratic noise and the level of skewness in the product revenues by using the tuple $([\zeta^L, \zeta^U], \kappa)$. Recall that we generate 50,000 individual problem instances for each combination of problem parameters. For each problem instance, we solve problem (15) to obtain an upper bound on the optimal expected revenue. We use UB^k to denote the upper bound on the optimal expected revenue we obtain for problem instance k . Furthermore, given a collection of candidate assortments, we find the best assortment that we can stitch together by focusing these candidate assortments. In Table 2, the candidate assortments are the nested-by-revenue assortments, whereas in Tables 3 and 4, the candidate assortments are respectively the nested-by-preference-and-revenue and powers-of-two assortments. For problem instance k , we use Best^k to denote the expected revenue provided by the best assortment we can find by focusing on a particular collection of candidate assortments. The third column in Tables 2, 3 and 4 shows the number of problem instances k for which we have $\text{UB}^k > \text{Best}^k$. These problem instances correspond to those where we are not able to establish the optimality of the best assortment we can find. The fourth column focuses on the problem instances for which we cannot establish the optimality of the best assortment we can find, and reports the percent gap between UB^k and Best^k over these problem instances. In other words, using \mathcal{K} to denote the set of problem instances $\{k = 1, \dots, 50,000 : \text{UB}^k > \text{Best}^k\}$, the fourth column gives

$$\frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} 100 \frac{\text{UB}^k - \text{Best}^k}{\text{UB}^k}.$$

The fourth column can be interpreted as the average optimality gap of the best assortment we can find given that we cannot verify the optimality of this assortment. The fifth column gives the 99.9-th percentile of the gaps between UB^k and Best^k over all 50,000 problem instances. That is, the fifth column gives the 99.9-th percentile of the data $\{100(\text{UB}^k - \text{Best}^k)/\text{UB}^k : k = 1, \dots, 50,000\}$. Finally, the sixth column shows the average number of products per nest in the best assortment we find.

The results in Table 2 indicate that nested-by-revenue assortments perform remarkably well. For the test problems with possibly synergistic products and fully-captured nests, we cannot establish the optimality of the best nested-by-revenue assortment for a majority of the problem instances. However, we are able to verify that the average optimality gap of best nested-by-revenue assortments is no larger than 0.114%. In 99.9% of the problem instances we generate, the optimality gap of the best nested-by-revenue assortment is no larger than 0.914%. An important point in Table 2 is that this table indicates that both nested-by-revenue assortments perform quite well and the upper bounds on the optimal expected revenue provided by problem (15) are quite tight. Both of these findings are crucial. Without identifying a class of candidate assortments that perform well, we would not be able to find a good solution to the assortment optimization problem. Without having a tight bound on the optimal expected revenue, we would not be able to verify the quality of the assortments we obtain. The results are even more appealing when we move on to the test problems with competitive products and partially-captured nests or possibly synergistic products and partially-captured nests. For these test problems, we can verify that the best nested-by-revenue assortment is indeed optimal for a majority

of the problem instances. Over the problem instances where we cannot verify optimality, the average optimality gap of best nested-by-revenue assortments is no larger than 0.008%. In 99.9% of the problem instances, the optimality gap of the best nested-by-revenue assortment is no larger than 0.090%. We also note that we tried a variety of strategies to generate our test problems and the experimental setup in Section 8.2 corresponds to the case where the performance of nested-by-revenue assortments turned out to be the least effective. For other classes of randomly generated problem instances, the performance of nested-by-revenue assortments can be even stronger.

Our findings in Table 3 indicate that nested-by-preference-and-revenue assortments provide marginal improvements over nested-by-revenue assortments. The problem instances for which we cannot verify the optimality of nested-by-preference-and-revenue assortments are the same as the problem instances for which we cannot verify the optimality of nested-by-revenue assortments. However, for certain but few problem instances, we can obtain slightly better performance by enlarging the collection of candidate assortments from nested-by-revenue to nested-by-preference-and-revenue. On the other hand, Table 4 shows that the performance of powers-of-two assortments is not as good as the performance of the other two collections of assortments we consider. For almost all of the problem instances, we cannot verify the optimality of the best powers-of-two assortment. The average optimality gap of best powers-of-two assortments can be as high as 6.024%. Noting that the number of possible powers-of-two assortments grows logarithmically with n , whereas the numbers of nested-by-revenue assortments grows linearly with n , this is the price we pay by focusing on a smaller number of possible assortments when we work with the powers-of-two assortments.

To sum up, nested-by-revenue assortments perform quite well in all of our test problems, even for problem instances for which we are not able to give a theoretical performance guarantee for such assortments. The improvement provided by nested-by-preference-and-revenue assortments generally remains on the marginal side. Although we give a performance guarantee for powers-of-two assortments for the most general instances of the assortment optimization problem, such assortments turn out to be too restricted when compared with nested-by-revenue and nested-by-preference-and-revenue assortments. As equally important as these observations, the upper bound on the optimal expected revenue that we obtain by solving problem (15) turns out to be quite tight. Thus, even if there is no theoretical performance guarantee for a certain collection of candidate assortments, we can stitch together the best possible solution by focusing on assortments within this collection. In this case, we can solve problem (15) to get a feel for the optimality gap of the best solution obtained in this fashion. If the optimality gap is reasonably small, then there is no reason to enrich the collection of candidate assortments under consideration. This approach would work not only for the collections of candidate assortments we study in this paper, but for any collection of candidate assortments.

9 Conclusions

In this paper, we studied a class of assortment optimization problems under the nested attraction model. We showed that the problem is polynomially solvable when the dissimilarity parameters of the nests are between zero and one, and the customers always purchase a product within their selected

Prob. Categ.	Prob. Param. $([\zeta^L, \zeta^U], \kappa)$	Sub. Cnt.	Avg. % Gap in Sub.	99.9-th Per. % Gap.	Avg. Ass. Size
Syner. Prod. Full-Cap.	$([1.0, 1.0], 1)$	46,667	0.038	0.282	8.8
	$([0.8, 1.2], 1)$	46,989	0.042	0.329	9.5
	$([0.5, 1.5], 1)$	47,583	0.072	0.617	8.7
	$([1.0, 1.0], 2)$	47,698	0.081	0.573	8.8
	$([0.8, 1.2], 2)$	47,778	0.085	0.615	8.9
	$([0.5, 1.5], 2)$	48,248	0.114	0.914	8.7
Comp. Prod. Partial-Cap.	$([1.0, 1.0], 1)$	10,305	0.004	0.041	8.7
	$([0.8, 1.2], 1)$	9,517	0.004	0.049	8.9
	$([0.5, 1.5], 1)$	8,893	0.006	0.090	8.5
	$([1.0, 1.0], 2)$	6,416	0.003	0.023	9.0
	$([0.8, 1.2], 2)$	5,758	0.002	0.020	9.4
	$([0.5, 1.5], 2)$	5,603	0.003	0.035	9.2
Syner. Prod. Partial-Cap.	$([1.0, 1.0], 1)$	5,893	0.004	0.025	9.5
	$([0.8, 1.2], 1)$	6,602	0.004	0.030	9.9
	$([0.5, 1.5], 1)$	7,956	0.006	0.064	9.6
	$([1.0, 1.0], 2)$	6,347	0.006	0.036	9.3
	$([0.8, 1.2], 2)$	6,606	0.006	0.042	9.9
	$([0.5, 1.5], 2)$	7,617	0.008	0.068	9.8

Table 2: Performance of nested-by-revenue assortments.

nest. Relaxing either one of these assumptions renders the problem NP-hard. To deal with the NP-hard cases, we developed collections of assortments with worst-case performance guarantees. Furthermore, we formulated a tractable convex program whose optimal objective value provides an upper bound on the optimal expected revenue. In this case, we can compare the expected revenue provided by an assortment with the upper bound on the optimal expected revenue to get a feel for the optimality gap of the assortment. By following this approach, our computational experiments showed that we can obtain solutions for the assortment optimization problem whose performance is within a fraction of a percent of the optimal expected revenue.

We hope that the techniques developed in this paper will be useful to deal with more general problems with a variety of constraints on the offered assortment. These constraints may limit the cardinality of the assortment or may impose precedence relationships between the products that can be offered. In addition, we are working on designing attraction-based choice models to capture customer choice behavior that is not encompassed by the nested logit model, while keeping the corresponding assortment optimization problems tractable.

References

- Bazaraa, M. S., Sherali, H. D. and Shetty, C. M. (1993), *Nonlinear Programming: Theory and Algorithms*, second edn, John Wiley & Sons, Inc., New York.
- Borsch-Supan, A. (1990), ‘On the compatibility of nested logit models with utility maximization’, *Journal of Econometrics* **43**, 373–388.
- Bront, J. J. M., Mendez-Diaz, I. and Vulcano, G. (2009), ‘A column generation algorithm for choice-based network revenue management’, *Operations Research* **57**(3), 769–784.

Prob. Categ.	Prob. Param. $([\zeta^L, \zeta^U], \kappa)$	Sub. Cnt.	Avg. % Gap in Sub.	99.9-th Per. % Gap.	Avg. Ass. Size
Syner. Prod. Full-Cap.	$([1.0, 1.0], 1)$	46,667	0.038	0.282	8.8
	$([0.8, 1.2], 1)$	46,989	0.041	0.319	9.5
	$([0.5, 1.5], 1)$	47,583	0.065	0.540	8.7
	$([1.0, 1.0], 2)$	47,698	0.081	0.573	8.8
	$([0.8, 1.2], 2)$	47,778	0.084	0.601	8.9
	$([0.5, 1.5], 2)$	48,248	0.106	0.847	8.7
Comp. Prod. Partial-Cap.	$([1.0, 1.0], 1)$	10,305	0.004	0.041	8.7
	$([0.8, 1.2], 1)$	9,517	0.004	0.049	8.9
	$([0.5, 1.5], 1)$	8,893	0.006	0.083	8.5
	$([1.0, 1.0], 2)$	6,416	0.003	0.023	9.0
	$([0.8, 1.2], 2)$	5,758	0.002	0.020	9.4
	$([0.5, 1.5], 2)$	5,603	0.003	0.034	9.2
Syner. Prod. Partial-Cap.	$([1.0, 1.0], 1)$	5,893	0.004	0.025	9.5
	$([0.8, 1.2], 1)$	6,602	0.004	0.030	9.9
	$([0.5, 1.5], 1)$	7,956	0.006	0.061	9.6
	$([1.0, 1.0], 2)$	6,347	0.006	0.036	9.3
	$([0.8, 1.2], 2)$	6,606	0.006	0.041	9.9
	$([0.5, 1.5], 2)$	7,617	0.008	0.068	9.8

Table 3: Performance of nested-by-preference-and-revenue assortments.

Prob. Categ.	Prob. Param. $([\zeta^L, \zeta^U], \kappa)$	Sub. Cnt.	Avg. % Gap in Sub.	99.9-th Per. % Gap.	Avg. Ass. Size
Syner. Prod. Full-Cap.	$([1.0, 1.0], 1)$	50,000	3.049	7.218	7.9
	$([0.8, 1.2], 1)$	50,000	2.895	7.103	8.5
	$([0.5, 1.5], 1)$	50,000	2.999	7.956	8.2
	$([1.0, 1.0], 2)$	49,998	4.413	10.476	8.2
	$([0.8, 1.2], 2)$	49,998	3.539	9.330	8.4
	$([0.5, 1.5], 2)$	49,998	3.256	9.167	8.2
Comp. Prod. Partial-Cap.	$([1.0, 1.0], 1)$	49,998	3.306	7.753	7.9
	$([0.8, 1.2], 1)$	50,000	3.062	7.724	8.2
	$([0.5, 1.5], 1)$	50,000	3.659	9.264	8.4
	$([1.0, 1.0], 2)$	49,995	5.280	13.092	8.8
	$([0.8, 1.2], 2)$	49,991	4.688	11.938	8.9
	$([0.5, 1.5], 2)$	49,998	6.024	16.668	8.8
Syner. Prod. Partial-Cap.	$([1.0, 1.0], 1)$	49,991	3.858	9.076	8.4
	$([0.8, 1.2], 1)$	50,000	3.777	9.493	9.8
	$([0.5, 1.5], 1)$	50,000	3.750	10.320	9.3
	$([1.0, 1.0], 2)$	49,999	5.215	12.423	9.0
	$([0.8, 1.2], 2)$	49,999	5.005	12.305	9.2
	$([0.5, 1.5], 2)$	49,999	5.673	14.082	9.0

Table 4: Performance of powers-of-two assortments.

- Cachon, G., Terwiesch, C. and Xu, Y. (2005), ‘Assortment planning in the presence of consumer search’, *Manufacturing & Service Operations Management* **7**(4), 330–346.
- Dong, L., Kouvelis, P. and Tian, Z. (2009), ‘Dynamic pricing and inventory control of substitute products’, *Manufacturing & Service Operations Management* **11**(2), 317–339.
- Farias, V. F., Jagabathula, S. and Shah, D. (2011), ‘A non-parametric approach to modeling choice with limited data’, *Management Science* (to appear).
- Gallego, G., Iyengar, G., Phillips, R. and Dubey, A. (2004), Managing flexible products on a network, Computational Optimization Research Center Technical Report TR-2004-01, Columbia University.
- Gallego, G. and Wang, R. (2011), Multi-product price optimization and competition under the nested attraction model, Technical report, Columbia University.
- Garey, M. and Johnson, D. (1979), *Computers and Intractability: A Guide to the Theory of NP-Completeness*, Freeman, San Francisco, CA.
- Hanson, W. and Martin, K. (1996), ‘Optimizing multinomial logit profit functions’, *Management Science* **42**(7), 992–1003.
- Kok, A. G., Fisher, M. L. and Vaidyanathan, R. (2008), Assortment planning: Review of literature and industry practice, in ‘Retail Supply Chain Management’, Kluwer Publishers.
- Kunnumkal, S. and Topaloglu, H. (2008), ‘A refined deterministic linear program for the network revenue management problem with customer choice behavior’, *Naval Research Logistics Quarterly* **55**(6), 563–580.
- Li, H. and Huh, W. T. (2011), ‘Pricing multiple products with the multinomial logit and nested models: Concavity and implications’, *Manufacturing & Service Operations Management* (to appear).
- Liu, Q. and van Ryzin, G. (2008), ‘On the choice-based linear programming model for network revenue management’, *Manufacturing & Service Operations Management* **10**(2), 288–310.
- Luce, R. D. (1959), *Individual Choice Behavior: A Theoretical Analysis*, Wiley, New York, NY.
- McFadden, D. (1974), Conditional logit analysis of qualitative choice behavior, in P. Zarembka, ed., ‘*Frontiers in Economics*’, Academic Press, pp. 105–142.
- McFadden, D. (1980), ‘Econometric models for probabilistic choice among products’, *The Journal of Business* **53**(3), S13–29.
- Mendez-Diaz, I., Bront, J. J. M., Vulcano, G. and Zabala, P. (2010), ‘A branch-and-cut algorithm for the latent-class logit assortment problem’, *Discrete Applied Mathematics* (to appear).
- Rusmevichientong, P., Shen, Z.-J. M. and Shmoys, D. B. (2010), ‘Dynamic assortment optimization with a multinomial logit choice model and capacity constraint’, *Operations Research* **58**(6), 1666–1680.
- Rusmevichientong, P., Shmoys, D. B. and Topaloglu, H. (2010), Assortment optimization with mixtures of logits, Technical report, Cornell University, School of Operations Research and Information Engineering.
Available at <http://legacy.orie.cornell.edu/~huseyin/publications/publications.html>.
- Rusmevichientong, P. and Topaloglu, H. (2011), Robust assortment optimization in revenue management under the multinomial logit choice model, Technical report, Cornell University, School of Operations Research and Information Engineering.
Available at <http://legacy.orie.cornell.edu/~huseyin/publications/publications.html>.
- Ruszczynski, A. (2006), *Nonlinear Optimization*, Princeton University Press, Princeton, New Jersey.
- Song, J.-S. and Xue, Z. (2007), Demand management and inventory control for substitutable products, Technical report, Duke University, Durham, NC.

- Talluri, K. and van Ryzin, G. (2004), ‘Revenue management under a general discrete choice model of consumer behavior’, *Management Science* **50**(1), 15–33.
- van Ryzin, G. and Mahajan, S. (1999), ‘On the relationship between inventory costs and variety benefits in retail assortments’, *Management Science* **45**(11), 1496–1509.
- Williams, H. C. W. L. (1977), ‘On the formation of travel demand models and economic evaluation measures of user benefit’, *Environment and Planning* **9**(3), 285–344.
- Williamson, D. P. and Shmoys, D. B. (2011), *The Design of Approximation Algorithms*, Cambridge University Press, New York.
- Zhang, D. and Adelman, D. (2009), ‘An approximate dynamic programming approach to network revenue management with customer choice’, *Transportation Science* **42**(3), 381–394.

A Appendix: Omitted Proofs

A.1 Proof of Theorem 7

Noting that problem (3) is a relaxed version of problem (7), it is enough to show that $(\alpha \hat{x}, \alpha \hat{y})$ is a feasible solution to problem (7). We observe that since (\hat{x}, \hat{y}) is an optimal solution to problem (4) after replacing the collection of assortments $\{A_{it} : t \in \mathcal{T}_i\}$ in the second set of constraints with the nested-by-revenue assortments $\{N_{ij} : j \in N_+\}$, this solution satisfies the second set of constraints for nest i and the nested-by-revenue assortment $N_{i0} = \emptyset$. Noting that $V_i(\emptyset) = 0$, it follows that $\hat{y}_i \geq 0$ for all $i \in M$. In this case, the first constraint in problem (4) implies that $\hat{x} \geq 0$.

We fix an arbitrary nest i and let \hat{z}_i be an optimal solution to the maximization problem on the right side of the second set of constraints in problem (7) when this maximization problem is solved at $x = \alpha \hat{x}$. By Lemma 6, \hat{z}_i is of the form $\hat{z}_{i1} = 1, \hat{z}_{i2} = 1, \dots, \hat{z}_{i,k-1} = 1, \hat{z}_{ik} \in [0, 1], \hat{z}_{i,k+1} = 0, \dots, \hat{z}_{in} = 0$ for some $k = 1, \dots, n$. We define ρ as $\rho = \hat{z}_{ik} \in [0, 1]$ and consider two cases.

Case 1. Assume that $k \geq 2$. We branch into two subcases.

Case 1.a. Noting that (\hat{x}, \hat{y}) is the optimal solution to problem (4) after replacing the collection of assortments $\{A_{it} : t \in \mathcal{T}_i\}$ in the second set of constraints with the nested-by-revenue assortments $\{N_{ij} : j \in N_+\}$, this solution satisfies the second set of constraints in problem (4) for nest i and the nested-by-revenue assortment $N_{ik} = \{1, 2, \dots, k\}$. Thus, it holds that

$$\hat{y}_i \geq \left(\sum_{j=1}^k v_{ij} \right)^{\gamma_i} \left[\frac{\sum_{j=1}^k r_{ij} v_{ij}}{\sum_{j=1}^k v_{ij}} - \hat{x} \right].$$

For notational convenience, let $R_{ik'} = \sum_{j=1}^{k'} r_{ij} v_{ij}$ and $q_{ik'} = \sum_{j=1}^{k'} v_{ij}$ for all $k' = 1, \dots, n$. Multiplying the inequality above by $\frac{R_{i,k-1} + r_{ik} v_{ik} \rho}{q_{i,k-1} + v_{ik} \rho} \frac{q_{ik}}{R_{ik}}$, we obtain

$$\frac{R_{i,k-1} + r_{ik} v_{ik} \rho}{q_{i,k-1} + v_{ik} \rho} \frac{q_{ik}}{R_{ik}} \hat{y}_i \geq q_{ik}^{\gamma_i} \left[\frac{R_{i,k-1} + r_{ik} v_{ik} \rho}{q_{i,k-1} + v_{ik} \rho} - \frac{R_{i,k-1} + r_{ik} v_{ik} \rho}{q_{i,k-1} + v_{ik} \rho} \frac{q_{ik}}{R_{ik}} \hat{x} \right]. \quad (18)$$

It is simple to check that the first derivative of $\frac{R_{i,k-1} + r_{ik} v_{ik} \rho}{q_{i,k-1} + v_{ik} \rho}$ with respect to ρ has the same sign as $r_{ik} q_{i,k-1} - R_{i,k-1}$ and we have $r_{ik} q_{i,k-1} - R_{i,k-1} = \sum_{j=1}^{k-1} (r_{ik} - r_{ij}) v_{ij} \leq 0$, where the last inequality is by the fact that $r_{i1} \geq r_{i2} \geq \dots \geq r_{in}$. Thus, it follows that $\frac{R_{i,k-1} + r_{ik} v_{ik} \rho}{q_{i,k-1} + v_{ik} \rho}$ is decreasing in ρ so that it is bounded from above by $R_{i,k-1}/q_{i,k-1}$. Also, noting the definitions of $R_i(S_i)$, $R_{ik'}$ and $q_{ik'}$, we have $R_{ik'}/q_{ik'} = R_i(N_{ik'})$. In this case, we can bound the expression that multiplies \hat{y}_i and \hat{x} in (18) as

$$\frac{R_{i,k-1} + r_{ik} v_{ik} \rho}{q_{i,k-1} + v_{ik} \rho} \frac{q_{ik}}{R_{ik}} \leq \frac{R_{i,k-1}}{q_{i,k-1}} \frac{q_{ik}}{R_{ik}} = \frac{R_i(N_{i,k-1})}{R_i(N_{ik})}.$$

Letting $\alpha_{ik}^1 = R_i(N_{i,k-1})/R_i(N_{ik})$ for notational brevity, using the upper bound above in (18) and noting that $\hat{y}_i \geq 0$ and $\hat{x} \geq 0$, the inequality in (18) implies that

$$\alpha_{ik}^1 \hat{y}_i \geq q_{ik}^{\gamma_i} \left[\frac{R_{i,k-1} + r_{ik} v_{ik} \rho}{q_{i,k-1} + v_{ik} \rho} - \alpha_{ik}^1 \hat{x} \right].$$

Finally, if we multiply the right side of the inequality above by $(q_{k-1} + v_{ik} \rho)^{\gamma_i} / q_k^{\gamma_i} \leq 1$, but not the left side, then the inequality is above still preserved and we have

$$\alpha_{ik}^1 \hat{y}_i \geq (q_{i,k-1} + v_{ik} \rho)^{\gamma_i} \left[\frac{R_{i,k-1} + r_{ik} v_{ik} \rho}{q_{i,k-1} + v_{ik} \rho} - \alpha_{ik}^1 \hat{x} \right]. \quad (19)$$

Case 1.b. Noting that (\hat{x}, \hat{y}) is the optimal solution to problem (4) after replacing the collection of assortments $\{A_{it} : t \in \mathcal{T}_i\}$ in the second set of constraints with the nested-by-revenue assortments $\{N_{ij} : j \in N_+\}$, this solution also satisfies the second set of constraints in problem (4) for nest i and the nested-by-revenue assortment $N_{i,k-1} = \{1, 2, \dots, k-1\}$. Thus, we obtain

$$\hat{y}_i \geq \left(\sum_{j=1}^{k-1} v_{ij} \right)^{\gamma_i} \left[\frac{\sum_{j=1}^{k-1} r_{ij} v_{ij}}{\sum_{j=1}^{k-1} v_{ij}} - \hat{x} \right] = q_{i,k-1}^{\gamma_i} \left[\frac{R_{i,k-1}}{q_{i,k-1}} - \hat{x} \right].$$

Multiplying the inequality above by $\frac{(q_{i,k-1} + v_{ik} \rho)^{\gamma_i}}{q_{i,k-1}^{\gamma_i}} \frac{R_{i,k-1} + r_{ik} v_{ik} \rho}{q_{i,k-1} + v_{ik} \rho} \frac{q_{i,k-1}}{R_{i,k-1}}$ and arranging the terms, we have

$$\frac{R_{i,k-1} + r_{ik} v_{ik} \rho}{(q_{i,k-1} + v_{ik} \rho)^{1-\gamma_i}} \frac{q_{i,k-1}^{1-\gamma_i}}{R_{i,k-1}} \hat{y}_i \geq (q_{i,k-1} + v_{ik} \rho)^{\gamma_i} \left[\frac{R_{i,k-1} + r_{ik} v_{ik} \rho}{q_{i,k-1} + v_{ik} \rho} - \frac{R_{i,k-1} + r_{ik} v_{ik} \rho}{q_{i,k-1} + v_{ik} \rho} \frac{q_{i,k-1}}{R_{i,k-1}} \hat{x} \right].$$

Since $\hat{x} \geq 0$, if we make the expression that multiplies \hat{x} in the inequality above even larger by multiplying it by $\frac{(q_{i,k-1} + v_{ik} \rho)^{\gamma_i}}{q_{i,k-1}^{\gamma_i}} \geq 1$, then the inequality above is still preserved and we have

$$\frac{R_{i,k-1} + r_{ik} v_{ik} \rho}{(q_{i,k-1} + v_{ik} \rho)^{1-\gamma_i}} \frac{q_{i,k-1}^{1-\gamma_i}}{R_{i,k-1}} \hat{y}_i \geq (q_{i,k-1} + v_{ik} \rho)^{\gamma_i} \left[\frac{R_{i,k-1} + r_{ik} v_{ik} \rho}{q_{i,k-1} + v_{ik} \rho} - \frac{R_{i,k-1} + r_{ik} v_{ik} \rho}{(q_{i,k-1} + v_{ik} \rho)^{1-\gamma_i}} \frac{q_{i,k-1}^{1-\gamma_i}}{R_{i,k-1}} \hat{x} \right]. \quad (20)$$

It is simple to check that whenever the first derivative of $\frac{R_{i,k-1} + r_{ik} v_{ik} \rho}{(q_{i,k-1} + v_{ik} \rho)^{1-\gamma_i}}$ with respect to ρ vanishes, the second derivative takes a positive value. Therefore, this expression is maximized at either $\rho = 0$ or $\rho = 1$. In this case, we can bound the expression that multiplies \hat{y}_i and \hat{x} in (20) as

$$\frac{R_{i,k-1} + r_{ik} v_{ik} \rho}{(q_{i,k-1} + v_{ik} \rho)^{1-\gamma_i}} \frac{q_{i,k-1}^{1-\gamma_i}}{R_{i,k-1}} \leq 1 \vee \frac{R_{ik}}{q_{ik}^{1-\gamma_i}} \frac{q_{i,k-1}^{1-\gamma_i}}{R_{i,k-1}} = 1 \vee \frac{R_i(N_{ik})}{R_i(N_{i,k-1})} \frac{V_i(N_{ik})^{\gamma_i}}{V_i(N_{i,k-1})^{\gamma_i}},$$

where we use $a \vee b = \max\{a, b\}$. The two terms in the maximum operator on the right side of the first inequality are obtained by evaluating the expression on the left side of the inequality at $\rho = 0$ and $\rho = 1$. The equality above follows by noting that $R_{ik'}/q_{ik'} = R_i(N_{ik'})$ and $q_{ik'} = V_i(N_{ik'})$. Letting $\alpha_{ik}^2 = \frac{R_i(N_{ik})}{R_i(N_{i,k-1})} \frac{V_i(N_{ik})^{\gamma_i}}{V_i(N_{i,k-1})^{\gamma_i}}$ for notational convenience, using the upper bound above in (20) and noting that $\hat{y}_i \geq$ and $\hat{x} \geq 0$, the inequality in (20) implies that

$$(1 \vee \alpha_{ik}^2) \hat{y}_i \geq (q_{i,k-1} + v_{ik} \rho)^{\gamma_i} \left[\frac{R_{i,k-1} + r_{ik} v_{ik} \rho}{q_{i,k-1} + v_{ik} \rho} - (1 \vee \alpha_{ik}^2) \hat{x} \right]. \quad (21)$$

Putting Cases 1.a and 1.b together, we observe that \hat{y}_i and \hat{x} satisfy both (19) and (21), in which case, they must also satisfy

$$[\alpha_{ik}^1 \wedge (1 \vee \alpha_{ik}^2)] \hat{y}_i \geq (q_{i,k-1} + v_{ik} \rho)^{\gamma_i} \left[\frac{R_{i,k-1} + r_{ik} v_{ik} \rho}{q_{i,k-1} + v_{ik} \rho} - [\alpha_{ik}^1 \wedge (1 \vee \alpha_{ik}^2)] \hat{x} \right]. \quad (22)$$

Lemma 13 below shows that $\alpha \geq 1$ for the value of α given in (6). The proof of that lemma also shows that $\alpha_{ik}^1 = R_i(N_{ik})/R_i(N_{i,k-1}) \geq 1$. By the definitions of α , α_{ik}^1 and α_{ik}^2 , we obtain $\alpha \geq \alpha_{ik}^1 \wedge \alpha_{ik}^2$ as well. In this case, we have

$$\alpha_{ik}^1 \wedge (1 \vee \alpha_{ik}^2) = (\alpha_{ik}^1 \wedge 1) \vee (\alpha_{ik}^1 \wedge \alpha_{ik}^2) = 1 \vee (\alpha_{ik}^1 \wedge \alpha_{ik}^2) \leq \alpha.$$

Thus, replacing the expression that multiplies \hat{y}_i and \hat{x} in (22) with an even larger expression α , the inequality is still preserved and we obtain

$$\alpha \hat{y}_i \geq (q_{i,k-1} + v_{ik} \rho)^{\gamma_i} \left[\frac{R_{i,k-1} + r_{ik} v_{ik} \rho}{q_{i,k-1} + v_{ik} \rho} - \alpha \hat{x} \right]. \quad (23)$$

Case 2. Assume that $k = 1$. Since (\hat{x}, \hat{y}) is the optimal solution to problem (4) after replacing the collection of assortments $\{A_{it} : t \in \mathcal{T}_i\}$ in the second set of constraints with the nested-by-revenue assortments $\{N_{ij} : j \in N_+\}$, this solution satisfies the second set of constraints in problem (4) for nest i and the nested-by-revenue assortment $N_{i1} = \{1\}$. Therefore, we have $\hat{y}_i \geq v_{i1}^{\gamma_i} \left[\frac{r_{i1} v_{i1}}{v_{i1}} - \hat{x} \right]$. Since $\alpha \geq 1$, $\hat{y}_i \geq 0$ and $\hat{x} \geq 0$, this inequality yields $\alpha \hat{y}_i \geq v_{i1}^{\gamma_i} \left[\frac{r_{i1} v_{i1}}{v_{i1}} - \alpha \hat{x} \right]$. If we multiply the right side of the last inequality by $\rho^{\gamma_i} \leq 1$, but not the left side, then the inequality is still preserved and we obtain

$$\alpha \hat{y}_i \geq (v_{i1} \rho)^{\gamma_i} \left[\frac{r_{i1} v_{i1} \rho}{v_{i1} \rho} - \alpha \hat{x} \right]. \quad (24)$$

Putting Cases 1 and 2 together, we succinctly write the inequalities in (23) and (24) as

$$\begin{aligned} \alpha \hat{y}_i &\geq \left(\sum_{j=1}^{k-1} v_{ij} \hat{z}_{ij} + v_{ik} \hat{z}_{ik} \right)^{\gamma_i} \left[\frac{\sum_{j=1}^{k-1} r_{ij} v_{ij} \hat{z}_{ij} + r_{ik} v_{ik} \hat{z}_{ik}}{\sum_{j=1}^{k-1} v_{ij} \hat{z}_{ij} + v_{ik} \hat{z}_{ik}} - \alpha \hat{x} \right] \\ &= \max_{z_i \in [0,1]^n} \left\{ \left(\sum_{j \in N} v_{ij} z_{ij} \right)^{\gamma_i} \left[\frac{\sum_{j \in N} r_{ij} v_{ij} z_{ij}}{\sum_{j \in N} v_{ij} z_{ij}} - \alpha \hat{x} \right] \right\}, \end{aligned}$$

where the equality follows from the definition of \hat{z}_i . Since the choice of nest i is arbitrary, the inequality above holds for all $i \in M$, which implies that the solution $(\alpha \hat{x}, \alpha \hat{y})$ satisfies the second set of constraints in problem (7). Since (\hat{x}, \hat{y}) is an optimal solution to problem (4), we have $v_0 \hat{x} \geq \sum_{i \in M} \hat{y}_i$, which implies that $v_0 \alpha \hat{x} \geq \sum_{i \in M} \alpha \hat{y}_i$. Therefore, the solution $(\alpha \hat{x}, \alpha \hat{y})$ satisfies the first constraint in problem (7) as well and we obtain the desired result. \square

Lemma 13 *Using α to denote the expression in (6), if we have $\gamma_i > 1$ for some $i \in M$, then $\alpha \geq 1$.*

Proof. Noting that $R_i(N_{ij}) = \sum_{k=1}^j r_{ik} v_{ik} / \sum_{k=1}^j v_{ik}$, $R_{ij}(N_{ij})$ is the weighted average of the revenues of the first j products in nest i . Since $r_{i1} \geq r_{i2} \geq \dots \geq r_{in}$, it follows that $R_i(N_{i,j-1}) \geq R_i(N_{ij})$. On the other hand, since $\gamma_i > 1$, we have

$$R_i(N_{ij}) V_i(N_{ij})^{\gamma_i} = \frac{\sum_{k=1}^j r_{ik} v_{ik}}{\sum_{k=1}^j v_{ik}} \left(\sum_{k=1}^j v_{ik} \right)^{\gamma_i} \geq \frac{\sum_{k=1}^{j-1} r_{ik} v_{ik}}{\sum_{k=1}^{j-1} v_{ik}} \left(\sum_{k=1}^{j-1} v_{ik} \right)^{\gamma_i} = R_i(N_{i,j-1}) V_i(N_{i,j-1})^{\gamma_i}.$$

Therefore, both terms of the minimum operator in the expression in (6) for nest i are at least one, which implies that α is also at least one. \square

A.2 Proof of Theorem 8

Assume that we are given any instance of the partition problem with sizes (c_1, \dots, c_n) and $\sum_{j=1}^n c_j = 2T$. We define an instance of the assortment feasibility problem as follows. There are two nests. The preference weight for the option of not choosing any of the nests is $v_0 = 0$. The dissimilarity parameters of the two nests are $\gamma_1 = \gamma_2 = 1/2$. For the first nest, the preference weight of the no purchase option is $v_{10} = 2$. This nest has only one product in it. The revenue and the preference weight associated with this product are $r_{11} = 2(T+1)(T+3)$ and $v_{11} = 2(2T+1)$. For the second nest, the preference weight of the no purchase option is $v_{20} = 1$. The second nest has n products in it. The revenues of the products in the second nest are identical and they are given by $r_{2j} = (T+1)(2T+1)$ for all $j = 1, \dots, n$. The preference weights of the products in the second nest are given by $v_{2j} = c_j$ for all $j = 1, \dots, n$. We set the expected revenue threshold in the assortment feasibility problem as $K = (T+2)(2T+1)$.

The first observation that if we only offer the product in the first nest, then the expected revenue we generate from the first nest is

$$R_1(\{1\}) = \frac{r_{11} v_{11}}{v_{01} + v_{11}} = \frac{2(T+1)(T+3)2(2T+1)}{2+2(2T+1)} = (T+3)(2T+1),$$

which is larger than the revenues of the products in the second nest. Thus, if we want to get the largest possible expected revenue, then it is always optimal to offer the product in the first nest. Therefore, the only question for the assortment feasibility problem is to choose a subset S among the products in the second nest that makes sure that we obtain an expected revenue of $K = (T+2)(2T+1)$ or more. If we offer a subset S of the products in the second nest together with the product in the first nest, then the expected revenue is $Q_1(\{1\}, S) R_1(\{1\}) + Q_2(\{1\}, S) R_2(S)$, which evaluates to

$$\begin{aligned} & \frac{\sqrt{2+2(2T+1)}}{\sqrt{2+2(2T+1)} + \sqrt{1+\sum_{j \in S} c_j}} \frac{2(T+1)(T+3)2(2T+1)}{2+2(2T+1)} \\ & + \frac{\sqrt{1+\sum_{j \in S} c_j}}{\sqrt{2+2(2T+1)} + \sqrt{1+\sum_{j \in S} c_j}} \frac{(T+1)(2T+1)\sum_{j \in S} c_j}{1+\sum_{j \in S} c_j}. \end{aligned}$$

Thus, arranging the terms in the expression above, the assortment feasibility problem asks the question of whether there is a subset S such that

$$\frac{\frac{2(T+1)(T+3)2(2T+1)}{\sqrt{2+2(2T+1)}} + (T+1)(2T+1) \frac{\sum_{j \in S} c_j}{\sqrt{1+\sum_{j \in S} c_j}}}{\sqrt{2+2(2T+1)} + \sqrt{1+\sum_{j \in S} c_j}} \geq (T+2)(2T+1).$$

If we cancel the terms in the first fraction in the numerator on the left side and move the denominator to the right, then the inequality above is equivalent to

$$\begin{aligned} & 2(T+3)(2T+1)\sqrt{T+1} + (T+1)(2T+1) \frac{\sum_{j \in S} c_j}{\sqrt{1+\sum_{j \in S} c_j}} \\ & \geq 2(T+2)(2T+1)\sqrt{T+1} + (T+2)(2T+1)\sqrt{1+\sum_{j \in S} c_j}. \end{aligned}$$

Canceling the term $2T + 1$ from both sides of the inequality above, multiplying the inequality by $\sqrt{1 + \sum_{j \in S} c_j}$ and adding and subtracting one from the term $\sum_{j \in S} c_j$, the inequality above can be written as

$$\begin{aligned} 2(T+3)\sqrt{T+1}\sqrt{1 + \sum_{j \in S} c_j} + (T+1)\left(1 + \sum_{j \in S} c_j\right) - (T+1) \\ \geq 2(T+2)\sqrt{T+1}\sqrt{1 + \sum_{j \in S} c_j} + (T+2)\left(1 + \sum_{j \in S} c_j\right). \end{aligned}$$

Finally, collecting all of the terms to the right, the last inequality becomes

$$\left(1 + \sum_{j \in S} c_j\right) - 2\sqrt{T+1}\sqrt{1 + \sum_{j \in S} c_j} + (T+1) \leq 0.$$

Since the last inequality is equivalent to $(\sqrt{1 + \sum_{j \in S} c_j} - \sqrt{T+1})^2 \leq 0$, there exists an assortment with an expected revenue of $K = (T+2)(2T+1)$ or more if and only if there exists a subset S with $(\sqrt{1 + \sum_{j \in S} c_j} - \sqrt{T+1})^2 \leq 0$. However, the only way for the last inequality to hold is to have $\sum_{j \in S} c_j = T$. Therefore, finding an assortment that yields an expected revenue of K or more is equivalent to finding a subset S that satisfies $\sum_{j \in S} c_j = T$ and the latter statement is precisely what the partition problem is interested in. \square

A.3 Proof of Theorem 10

Since problem (10) is equivalent to problem (3), it is enough to show that $(2\hat{x}, 2\hat{y})$ is a feasible solution to problem (10). First, note that $\hat{x} \geq 0$. To see this claim, if $\hat{x} < 0$, then the right sides of the second set of constraints in problem (4) are strictly positive for nonempty assortments so that $\hat{y}_i > 0$ for all $i \in M$. In this case, (\hat{x}, \hat{y}) cannot satisfy the first constraint in problem (4), establishing the claim.

We fix an arbitrary nest i and let $\hat{\epsilon}_i$ be the optimal solution to the maximization problem on the right side of the second set of constraints in problem (10) when this maximization problem is solved at $x = 2\hat{x}$. Finally, let $\hat{z}_i(\hat{\epsilon}_i)$ be the optimal solution to problem (11) when this continuous knapsack problem is solved at $\epsilon_i = \hat{\epsilon}_i$. We consider two cases.

Case 1. Assume that the solution $\hat{z}_i(\hat{\epsilon}_i)$ has exactly one fractional component. We denote this fractional component by $k \in N$. Since (\hat{x}, \hat{y}) is the optimal solution to problem (4) after replacing the collection assortments $\{A_{it} : t \in \mathcal{T}_i\}$ in the second set of constraints with $\{\hat{S}_i(\epsilon_i) : \epsilon_i \in [0, \infty]\} \cup \{\{j\} : j \in N\}$, the solution (\hat{x}, \hat{y}) satisfies the second set of constraints in problem (4) for nest i and the assortment $\hat{S}_i(\hat{\epsilon}_i)$ and we obtain

$$\begin{aligned} \hat{y}_i &\geq V_i(\hat{S}_i(\hat{\epsilon}_i))^{\gamma_i} (R_i(\hat{S}_i(\hat{\epsilon}_i)) - \hat{x}) \\ &= \frac{\sum_{j \in \hat{S}_i(\hat{\epsilon}_i)} r_{ij} v_{ij}}{(v_{i0} + \sum_{j \in \hat{S}_i(\hat{\epsilon}_i)} v_{ij})^{1-\gamma_i}} - \left(v_{i0} + \sum_{j \in \hat{S}_i(\hat{\epsilon}_i)} v_{ij}\right)^{\gamma_i} \hat{x} \geq \frac{\sum_{j \in \hat{S}_i(\hat{\epsilon}_i)} r_{ij} v_{ij}}{(v_{i0} + \hat{\epsilon}_i)^{1-\gamma_i}} - (v_{i0} + \hat{\epsilon}_i)^{\gamma_i} \hat{x}, \end{aligned} \quad (25)$$

where the second inequality above follows by $\gamma_i \leq 1$ and noting that we have $\sum_{j \in \hat{S}_i(\hat{\epsilon}_i)} v_{ij} \leq \sum_{j \in N} v_{ij} \hat{z}_{ij}(\hat{\epsilon}_i) \leq \hat{\epsilon}_i$ by the definitions of $\hat{S}_i(\hat{\epsilon}_i)$ and $\hat{z}_i(\hat{\epsilon}_i)$. Similarly, the solution (\hat{x}, \hat{y}) satisfies the

second set of constraints in problem (4) for nest i and the singleton assortment $\{k\}$ so that

$$\hat{y}_i \geq V_i(\{k\})^{\gamma_i} (R_i(\{k\}) - \hat{x}) = \frac{r_{ik} v_{ik}}{(v_{i0} + v_{ik})^{1-\gamma_i}} - (v_{i0} + v_{ik})^{\gamma_i} \hat{x} \geq \frac{r_{ik} v_{ik}}{(v_{i0} + \hat{\epsilon}_i)^{1-\gamma_i}} - (v_{i0} + \hat{\epsilon}_i)^{\gamma_i} \hat{x}, \quad (26)$$

where the second inequality above follows from the fact that we must have $v_{ik} \leq \hat{\epsilon}_i$ for $\hat{z}_{ik}(\hat{\epsilon}_i)$ to take a strictly positive value. Since $\hat{S}_i(\hat{\epsilon}_i)$ includes all strictly positive and integer-valued components of $\hat{z}_i(\hat{\epsilon}_i)$ and k is the only component of $\hat{z}_i(\hat{\epsilon}_i)$ that takes a fractional value, we have $\sum_{j \in \hat{S}_i(\hat{\epsilon}_i)} r_{ij} v_{ij} + r_{ik} v_{ik} \geq \sum_{j \in N} r_{ij} v_{ij} \hat{z}_{ij}(\hat{\epsilon}_i) = \hat{K}_i(\hat{\epsilon}_i)$, where the equality follows by the definition of $\hat{z}_i(\hat{\epsilon}_i)$. Using this relationship and adding (25) and (26), we have

$$\begin{aligned} 2 \hat{y}_i &\geq \frac{\sum_{j \in \hat{S}_i(\hat{\epsilon}_i)} r_{ij} v_{ij} + r_{ik} v_{ik}}{(v_{i0} + \hat{\epsilon}_i)^{1-\gamma_i}} - (v_{i0} + \hat{\epsilon}_i)^{\gamma_i} 2 \hat{x} \geq \frac{\hat{K}_i(\hat{\epsilon}_i)}{(v_{i0} + \hat{\epsilon}_i)^{1-\gamma_i}} - (v_{i0} + \hat{\epsilon}_i)^{\gamma_i} 2 \hat{x} \\ &\geq \frac{K_i(\hat{\epsilon}_i)}{(v_{i0} + \hat{\epsilon}_i)^{1-\gamma_i}} - (v_{i0} + \hat{\epsilon}_i)^{\gamma_i} 2 \hat{x} = \max_{\epsilon_i \geq 0} \left\{ (v_{i0} + \epsilon_i)^{\gamma_i} \left[\frac{K_i(\epsilon_i)}{v_{i0} + \epsilon_i} - 2 \hat{x} \right] \right\}, \end{aligned} \quad (27)$$

where the last inequality follows from the fact that problem (11) is a relaxation of problem (9) and the last equality follows from the definition of $\hat{\epsilon}_i$.

Case 2. Assume that the solution $\hat{z}_i(\hat{\epsilon}_i)$ does not have any fractional components. In this case, $\hat{S}_i(\hat{\epsilon}_i)$ includes all strictly positive components of $\hat{z}_i(\hat{\epsilon}_i)$ and we obtain $\sum_{j \in \hat{S}_i(\hat{\epsilon}_i)} r_{ij} v_{ij} = \sum_{j \in N} r_{ij} v_{ij} \hat{z}_{ij}(\hat{\epsilon}_i) = \hat{K}_i(\hat{\epsilon}_i)$. Using this relationship and following the same argument that we used to obtain (25) in the first case, we have

$$\hat{y}_i \geq \frac{\sum_{j \in \hat{S}_i(\hat{\epsilon}_i)} r_{ij} v_{ij}}{(v_{i0} + \hat{\epsilon}_i)^{1-\gamma_i}} - (v_{i0} + \hat{\epsilon}_i)^{\gamma_i} \hat{x} \geq \frac{K_i(\hat{\epsilon}_i)}{(v_{i0} + \hat{\epsilon}_i)^{1-\gamma_i}} - (v_{i0} + \hat{\epsilon}_i)^{\gamma_i} \hat{x}.$$

Multiplying the inequality above by two, we obtain

$$\begin{aligned} 2 \hat{y}_i &\geq 2 \frac{K_i(\hat{\epsilon}_i)}{(v_{i0} + \hat{\epsilon}_i)^{1-\gamma_i}} - (v_{i0} + \hat{\epsilon}_i)^{\gamma_i} 2 \hat{x} \\ &\geq \frac{K_i(\hat{\epsilon}_i)}{(v_{i0} + \hat{\epsilon}_i)^{1-\gamma_i}} - (v_{i0} + \hat{\epsilon}_i)^{\gamma_i} 2 \hat{x} = \max_{\epsilon_i \geq 0} \left\{ (v_{i0} + \epsilon_i)^{\gamma_i} \left[\frac{K_i(\epsilon_i)}{v_{i0} + \epsilon_i} - 2 \hat{x} \right] \right\}, \end{aligned} \quad (28)$$

where the last inequality follows from the definition of $\hat{\epsilon}_i$.

Collecting (27) and (28) in the two cases, the solution $(2 \hat{x}, 2 \hat{y})$ satisfies the second set of constraints for nest i in problem (10). Noting that our choice of nest i is arbitrary, the second set of constraints in problem (10) is satisfied by the solution $(2 \hat{x}, 2 \hat{y})$. Finally, since the solution (\hat{x}, \hat{y}) is optimal to problem (4), we have $v_0 \hat{x} \geq \sum_{i \in M} \hat{y}_i$, which implies that $v_0 2 \hat{x} \geq \sum_{i \in M} 2 \hat{y}_i$. Therefore, the solution $(2 \hat{x}, 2 \hat{y})$ satisfies the first constraint in problem (10) as well and we obtain the desired result. \square

A.4 Knapsack Problems with Upper and Lower Bound Constraints

In this section, we give a tractable method to obtain approximate solutions to problem (14). For notational brevity, we omit the subscripts for the nest and use the decision variables $z = (z_1, \dots, z_n)$ to focus on the problem

$$\hat{G}_l = \max \left\{ \sum_{j=1}^n r_j v_j z_j : 2^{l-1} \leq v_0 + \sum_{j=1}^n v_j z_j \leq 2^l, z \in \{0, 1\}^n \right\}. \quad (29)$$

We are interested in finding a feasible solution to the problem above whose objective value deviates from the optimal objective value by no more than a factor of two. We begin by classifying the products in the problem above into two categories. A product j satisfying $v_j \leq 2^{l-1}$ is called a small product, whereas a product j satisfying $v_j > 2^{l-1}$ is called a large product. We observe that an optimal solution to problem (29) cannot include two large products. Without loss of generality, we assume that the small products are $1, \dots, n_s$ and the large products are $n_s + 1, \dots, n$. Furthermore, we order the small products such that $r_1 \geq r_2 \geq \dots \geq r_{n_s}$. Finally, we assume that we have $v_0 + v_j \leq 2^l$ for all $j = 1, \dots, n$. If some product j does not satisfy this inequality, then this product is never used in an optimal solution to problem (29) and we can immediately drop this product from consideration.

The following proposition shows that we can check the objective values provided by at most $3(1+n)$ possible solutions to problem (29) and obtain a solution whose objective value deviates from the optimal objective value by at most a factor of two.

Proposition 14 *There exist no more than $3(1+n)$ solutions to problem (29) such that if we check the objective value provided by each one of these solutions and pick the feasible solution that provides the largest objective value, then the objective value provided by this solution deviates from the optimal objective value of problem (29) by no more than a factor of two. Furthermore, each one of these $3(1+n)$ solutions can be obtained by a single sorting operation.*

Proof. Since the optimal solution to problem (29) includes at most one large product, we focus on the cases where the optimal solution includes only small products or it includes one large product.

Case 1. Assume that the optimal solution to problem (29) includes only small products. If we have $v_0 + \sum_{j=1}^{n_s} v_j \leq 2^l$, then we do not violate the upper bound constraint in problem (29) even when we include all small products in the solution. Therefore, if we have $v_0 + \sum_{j=1}^{n_s} v_j \leq 2^l$, then the optimal solution to problem (29) includes all small products.

In the rest of this case, we assume that $v_0 + \sum_{j=1}^{n_s} v_j > 2^l$. Focusing only on the small products, we solve the continuous knapsack problem

$$\max \left\{ \sum_{j=1}^{n_s} r_j v_j z_j : v_0 + \sum_{j=1}^{n_s} v_j z_j \leq 2^l, z \in [0, 1]^{n_s} \right\}. \quad (30)$$

The problem above is a continuous knapsack problem, where the utility-to-space consumption ratio of product j is r_j . Noting that $r_1 \geq r_2 \geq \dots \geq r_{n_s}$, there exists an optimal solution $z^* = (z_1^*, \dots, z_{n_s}^*)$ to this problem such that $z_1^* = 1, z_2^* = 1, \dots, z_{k-1}^* = 1, z_k^* \in [0, 1], z_{k+1}^* = 0, \dots, z_{n_s}^* = 0$ for some $k = 1, \dots, n_s$. We define one possible solution $\hat{z} = (\hat{z}_1, \dots, \hat{z}_n)$ to problem (29) as

$$\hat{z}_1 = 1, \hat{z}_2 = 1, \dots, \hat{z}_{k-1} = 1, \hat{z}_k = 0, \hat{z}_{k+2} = 0, \dots, \hat{z}_n = 0,$$

which is obtained by focusing on the strictly positive and integer-valued components of z^* . The solution \hat{z} is feasible to problem (29). To see this claim, since we have $v_0 + \sum_{j=1}^{n_s} v_j > 2^l$, the constraint in problem

(30) is satisfied as an equality at the optimal solution, in which case, we obtain

$$2^l = v_0 + \sum_{j=1}^{n_s} v_j z_j^* = v_0 + \sum_{j=1}^{k-1} v_j \hat{z}_j + v_k z_k^* \leq v_0 + \sum_{j=1}^{k-1} v_j \hat{z}_j + 2^{l-1},$$

where the inequality follows from the fact product k corresponds to a small product. Subtracting 2^{l-1} from the chain of inequalities above, we have $2^{l-1} \leq v_0 + \sum_{j=1}^{k-1} v_j \hat{z}_j$ so that the solution \hat{z} satisfies the lower bound constraint in problem (29). We also observe that $v_0 + \sum_{j=1}^n \hat{z}_j \leq v_0 + \sum_{j=1}^{n_s} z_j^* \leq 2^l$, where the second inequality follows from the fact that z^* is an optimal solution to problem (30). Therefore, \hat{z} satisfies the upper bound constraint in problem (29) as well, establishing the claim.

By focusing on only the potentially fractional component of z^* , we define another possible solution $\tilde{z} = (\tilde{z}_1, \dots, \tilde{z}_n)$ to problem (29) as $\tilde{z}_1 = 0, \tilde{z}_2 = 0, \dots, \tilde{z}_{k-1} = 0, \tilde{z}_k = 1, \tilde{z}_{k+1} = 0, \dots, \tilde{z}_n = 0$. It is not necessarily the case that \tilde{z} is a feasible solution to problem (29). We branch on two subcases.

Case 1.a. Assume that the solution \tilde{z} is feasible to problem (29). Since problem (30) is a relaxation of problem (29), its optimal objective value is an upper bound on the optimal objective value of problem (29). Therefore, we have

$$\hat{G}_l \leq \sum_{j=1}^{n_s} r_j v_j z_j^* = \sum_{j=1}^{k-1} r_j v_j \hat{z}_j + r_k v_k z_k^* \leq \sum_{j=1}^{k-1} r_j v_j \hat{z}_j + r_k v_k \tilde{z}_k \leq 2 \max \left\{ \sum_{j=1}^{k-1} r_j v_j \hat{z}_j, r_k v_k \tilde{z}_k \right\},$$

which implies that if we check the objective values provided by the two feasible solutions \hat{z} and \tilde{z} for problem (29) and pick the best one of these solutions, then the objective value provided by this solution deviates from the optimal objective value of problem (29) by no more than a factor of two.

Case 1.b Assume that the solution \tilde{z} is not feasible to problem (29). In this case, we have $v_0 + v_k < 2^{l-1}$, but since the solution \hat{z} is feasible, we also have $2^{l-1} \leq v_0 + \sum_{j=1}^{k-1} v_j$ and the last two inequalities yield $v_k \leq \sum_{j=1}^{k-1} v_j$. Since the revenues of the products satisfy $r_1 \geq r_2 \geq \dots \geq r_k$, the last inequality implies that $r_k v_k \leq \sum_{j=1}^{k-1} r_j v_j$. In this case, we can follow an argument similar to the one in Case 1.a and observe that since problem (30) is a relaxation of problem (29), its optimal objective value is an upper bound on the optimal objective value of problem (29). Therefore, we have

$$\hat{G}_l \leq \sum_{j=1}^{n_s} r_j v_j z_j^* = \sum_{j=1}^{k-1} r_j v_j + r_k v_k z_k^* \leq \sum_{j=1}^{k-1} r_j v_j + r_k v_k \leq 2 \sum_{j=1}^{k-1} r_j v_j = 2 \sum_{j=1}^{k-1} r_j v_j \hat{z}_j,$$

which implies that the objective value provided by the feasible solution \hat{z} for problem (29) deviates from the optimal objective value by no more than a factor of two.

Putting our observations in Case 1 together, there are three possible solutions we consider in this case. These are the solution that includes all small products and the solutions \hat{z} and \tilde{z} . If we know that the optimal solution to problem (29) includes only small products, then we can check whether each one of these three solutions is feasible to problem (29) and find the feasible solution that provides the best objective value. The objective value provided by this best solution deviates from the optimal objective value of problem (29) by no more than a factor of two.

Case 2. Assume that the optimal solution to problem (29) includes the large product ℓ . We assume without loss of generality that $v_0 + v_\ell + v_j \leq 2^l$ for all $j = 1, \dots, n_s$. If some small product does not satisfy this inequality, then this product is not included in an optimal solution to problem (29) and we can immediately drop this product from consideration.

If we have $v_0 + v_\ell + \sum_{j=1}^{n_s} v_j \leq 2^l$, then we do not violate the upper bound constraint in problem (29) even when we include all small products in the solution along with the large product ℓ . Thus, if $v_0 + v_\ell + \sum_{j=1}^{n_s} v_j \leq 2^l$, then the optimal solution to problem (29) includes all small products along with the large product ℓ . In the rest of this case, we assume that $v_0 + v_\ell + \sum_{j=1}^{n_s} v_j > 2^l$. Making sure that we include the large product ℓ in the solution, we solve the continuous knapsack problem

$$\max \left\{ r_\ell v_\ell + \sum_{j=1}^{n_s} r_j v_j z_j : v_0 + v_\ell + \sum_{j=1}^{n_s} v_j z_j \leq 2^l, z \in [0, 1]^{n_s} \right\}. \quad (31)$$

Similar to Case 1, there exists an optimal solution $z^* = (z_1^*, \dots, z_{n_s}^*)$ to problem (31) such that $z_1^* = 1, z_2^* = 1, \dots, z_{k-1}^* = 1, z_k^* \in [0, 1), z_{k+1}^* = 0, \dots, z_{n_s}^* = 0$ for some $k = 1, \dots, n_s$. We define one possible solution $\hat{z}^\ell = (\hat{z}_1^\ell, \dots, \hat{z}_n^\ell)$ to problem (29) as

$$\hat{z}_1^\ell = 1, \dots, \hat{z}_{k-1}^\ell = 1, \hat{z}_k^\ell = 0, \dots, \hat{z}_{\ell-1}^\ell = 0, \hat{z}_\ell^\ell = 1, \hat{z}_{\ell+1}^\ell = 0, \dots, \hat{z}_n^\ell = 0,$$

which is obtained by focusing on the strictly positive and integer-valued components of z^* along with the large product ℓ . It is straightforward to check that the solution \hat{z}^ℓ is feasible to problem (29). On the other hand, focusing only on the potentially fractional component of z^* and the large product ℓ , we define another solution $\tilde{z}^\ell = (\tilde{z}_1^\ell, \dots, \tilde{z}_n^\ell)$ to problem (29) as

$$\tilde{z}_1^\ell = 0, \dots, \tilde{z}_{k-1}^\ell = 0, \tilde{z}_k^\ell = 1, \tilde{z}_{k+1}^\ell = 0, \dots, \tilde{z}_{\ell-1}^\ell = 0, \tilde{z}_\ell^\ell = 1, \tilde{z}_{\ell+1}^\ell = 0, \dots, \tilde{z}_n^\ell = 0.$$

It is possible to check that \tilde{z}^ℓ is also a feasible solution to problem (29). Furthermore, following the same argument in Case 1.a, we can show that if we check the objective values provided by the two feasible solutions \hat{z}^ℓ and \tilde{z}^ℓ for problem (29) and pick the best one of these solutions, then the objective value provided by this best solution deviates from the optimal objective value of problem (29) by no more than a factor of two.

Putting our observations in Case 2 together, there are three possible solutions we consider in this case. These are the solution that includes all small products along with the large product ℓ and the solutions \hat{z}^ℓ and \tilde{z}^ℓ . If we know that the optimal solution to problem (29) includes the large product ℓ , then we can check whether each one of these three solutions is feasible to problem (29) and find the feasible solution that provides the best objective value. The objective value provided by this best solution deviates from the optimal objective value of problem (29) by no more than a factor of two.

We do not a priori know whether an optimal solution to problem (29) includes only small products or it includes a large product. However, we can generate the three solutions in Case 1, along with the three solutions in Case 2 for all large products. This results in no more than $3(1+n)$ possible solutions. Among these solutions, we can choose the feasible one for problem (29) that provides the best objective value. The objective value provided by this best solution would deviate from the optimal objective value by no more than a factor of two. It is clear that each one of these $3(1+n)$ solutions can be obtained by a single sorting operation. \square

A.5 Proof of Theorem 11

By using the same argument at the beginning of the proof of Theorem 10, it follows that $\hat{x} \geq 0$. Fix an arbitrary nest i . Choose any assortment $S_i \subset N$ within this nest. First, we consider the case where $S_i \neq \emptyset$. Fix $l = l_i^L, \dots, l_i^U$ such that $2^{l-1} \leq v_{i0} + \sum_{j \in S_i} v_{ij} \leq 2^l$. Since (\hat{x}, \hat{y}) is the optimal solution to problem (4) after replacing the collection of assortments $\{A_{it} : t \in \mathcal{T}_i\}$ in the second set of constraints with the assortments $\{\hat{S}_{i\ell} : \ell = l_i^L, \dots, l_i^U\} \cup \{\emptyset\}$, this solution satisfies the second set of constraints in problem (4) for nest i and the assortment \hat{S}_{il} . Therefore, we have

$$\hat{y}_i \geq V_i(\hat{S}_{il})^{\gamma_i} (R_i(\hat{S}_{il}) - \hat{x}) = V_i(\hat{S}_{il})^{\gamma_i-1} \left(\sum_{j \in \hat{S}_{il}} r_{ij} v_{ij} \right) - V_i(\hat{S}_{il})^{\gamma_i} \hat{x}.$$

Multiplying both sides of this inequality by $2^{\bar{\gamma}+1}$, we obtain

$$2^{\bar{\gamma}+1} \hat{y}_i \geq 2^{\bar{\gamma}} V_i(\hat{S}_{il})^{\gamma_i-1} \left(2 \sum_{j \in \hat{S}_{il}} r_{ij} v_{ij} \right) - 2^{\bar{\gamma}+1} V_i(\hat{S}_{il})^{\gamma_i} \hat{x}. \quad (32)$$

We proceed to bound each one of the terms $2 \sum_{j \in \hat{S}_{il}} r_{ij} v_{ij}$, $V_i(\hat{S}_{il})^{\gamma_i-1}$ and $V_i(\hat{S}_{il})^{\gamma_i}$ in the inequality above. By the definition of \hat{S}_{il} , we have

$$2 \sum_{j \in \hat{S}_{il}} r_{ij} v_{ij} \geq \hat{G}_{il} = \max_{S'_i \subset N} \left\{ \sum_{j \in S'_i} r_{ij} v_{ij} : 2^{l-1} \leq v_{i0} + \sum_{j \in S'_i} v_{ij} \leq 2^l \right\} \geq \sum_{j \in S_i} r_{ij} v_{ij}, \quad (33)$$

where the second inequality follows by noting that l is chosen such that $2^{l-1} \leq v_{i0} + \sum_{j \in S_i} v_{ij} \leq 2^l$. The definition of \hat{S}_{il} also implies that $2^{l-1} \leq v_{i0} + \sum_{j \in \hat{S}_{il}} v_{ij} = V_i(\hat{S}_{il}) \leq 2^l$. In this case, if we have $\gamma_i \leq 1$, then $V_i(\hat{S}_{il})^{\gamma_i-1} \geq (2^l)^{\gamma_i-1}$. If, on the other hand, we have $\gamma_i > 1$, then $V_i(\hat{S}_{il})^{\gamma_i-1} \geq (2^{l-1})^{\gamma_i-1} = (2^l)^{\gamma_i-1} 2^{-(\gamma_i-1)}$. So, combining the two cases, we bound $V_i(\hat{S}_{il})^{\gamma_i-1}$ by

$$V_i(\hat{S}_{il})^{\gamma_i-1} \geq (2^l)^{\gamma_i-1} 2^{-[\gamma_i-1]^+},$$

where we use $[a]^+ = \max\{a, 0\}$. Noting also that $V_i(\hat{S}_{il})^{\gamma_i} \leq (2^l)^{\gamma_i}$, using these bounds on $V_i(\hat{S}_{il})^{\gamma_i-1}$ and $V_i(\hat{S}_{il})^{\gamma_i}$ together with the inequality in (33) in (32), we obtain

$$\begin{aligned} 2^{\bar{\gamma}+1} \hat{y}_i &\geq 2^{\bar{\gamma}} (2^l)^{\gamma_i-1} 2^{-[\gamma_i-1]^+} \left(\sum_{j \in S_i} r_{ij} v_{ij} \right) - 2^{\bar{\gamma}+1} (2^l)^{\gamma_i} \hat{x} \\ &= 2^{\bar{\gamma}} (2^l)^{\gamma_i-1} 2^{-[\gamma_i-1]^+} \left(\sum_{j \in S_i} r_{ij} v_{ij} \right) - 2^{\bar{\gamma}+\gamma_i+1} (2^{l-1})^{\gamma_i} \hat{x}. \end{aligned} \quad (34)$$

Our choice of l at the beginning of the proof implies that $2^{l-1} \leq v_{i0} + \sum_{j \in S_i} v_{ij} = V_i(S_i) \leq 2^l$. In this case, if we have $\gamma_i \leq 1$, then $V_i(S_i)^{\gamma_i-1} \leq (2^{l-1})^{\gamma_i-1} = (2^l)^{\gamma_i-1} 2^{1-\gamma_i}$. If, on the other hand, we have $\gamma_i > 1$, then $V_i(S_i)^{\gamma_i-1} \leq (2^l)^{\gamma_i-1}$. Combining the two cases yields $V_i(S_i)^{\gamma_i-1} \leq (2^l)^{\gamma_i-1} 2^{[1-\gamma_i]^+}$ so that we can bound $(2^l)^{\gamma_i-1}$ by

$$(2^l)^{\gamma_i-1} \geq 2^{-[1-\gamma_i]^+} V_i(S_i)^{\gamma_i-1}.$$

Furthermore, noting that $(2^{l-1})^{\gamma_i} \leq V_i(S_i)^{\gamma_i}$, we use these bounds on $(2^l)^{\gamma_i-1}$ and $(2^{l-1})^{\gamma_i}$ in the two terms on the right side of (34) to obtain

$$2^{\bar{\gamma}+1} \hat{y}_i \geq 2^{\bar{\gamma}} 2^{-[\gamma_i-1]^+} 2^{-[1-\gamma_i]^+} V_i(S_i)^{\gamma_i-1} \left(\sum_{j \in S_i} r_{ij} v_{ij} \right) - 2^{\bar{\gamma}+\gamma_i+1} V_i(S_i)^{\gamma_i} \hat{x}.$$

If we have $\gamma_i \leq 1$, then $\bar{\gamma} - [\gamma_i - 1]^+ - [1 - \gamma_i]^+ = \bar{\gamma} - 1 + \gamma_i \geq 0$, where we use the fact that $\bar{\gamma} > 1$. On the other hand, if we have $\gamma_i > 1$, then $\bar{\gamma} - [\gamma_i - 1]^+ - [1 - \gamma_i]^+ = \bar{\gamma} - \gamma_i + 1 \geq 0$ since $\bar{\gamma} \geq \gamma_i$. Therefore, $2^{\bar{\gamma}} 2^{-[\gamma_i - 1]^+} 2^{-[1 - \gamma_i]^+} \geq 1$. We also have $2^{\bar{\gamma} + \gamma_i + 1} \leq 2^{2\bar{\gamma} + 1}$. Thus, the last inequality above yields

$$2^{\bar{\gamma} + 1} \hat{y}_i \geq V_i(S_i)^{\gamma_i - 1} \left(\sum_{j \in S_i} r_{ij} v_{ij} \right) - 2^{2\bar{\gamma} + 1} V_i(S_i)^{\gamma_i} \hat{x}.$$

Since $\sum_{j \in S_i} r_{ij} v_{ij} / V_i(S_i) = R_i(S_i)$, the inequality above shows that the solution $(2^{2\bar{\gamma} + 1} \hat{x}, 2^{\bar{\gamma} + 1} \hat{y})$ satisfies the second set of constraints in problem (3) for the assortment S_i and nest i , as long as $S_i \neq \emptyset$.

Second, we consider the case where $S_i = \emptyset$. The solution (\hat{x}, \hat{y}) satisfies the second set of constraints in problem (4) for the empty assortment within nest i , in which case, we obtain $\hat{y}_i \geq V_i(\emptyset)^{\gamma_i} (R_i(\emptyset) - \hat{x})$. Multiplying this inequality by $2^{\bar{\gamma} + 1}$ and noting that $R_i(\emptyset) = 0$, we have $2^{\bar{\gamma} + 1} \hat{y}_i \geq 2^{\bar{\gamma} + 1} V_i(\emptyset)^{\gamma_i} R_i(\emptyset) - 2^{\bar{\gamma} + 1} V_i(\emptyset)^{\gamma_i} \hat{x} = V_i(\emptyset)^{\gamma_i} R_i(\emptyset) - 2^{\bar{\gamma} + 1} V_i(\emptyset)^{\gamma_i} \hat{x}$. Replacing the term $2^{\bar{\gamma} + 1}$ on the right side of the last inequality with an even larger term $2^{2\bar{\gamma} + 1}$, it follows that

$$2^{\bar{\gamma} + 1} \hat{y}_i \geq V_i(\emptyset)^{\gamma_i} R_i(\emptyset) - 2^{2\bar{\gamma} + 1} V_i(\emptyset)^{\gamma_i} \hat{x}.$$

Therefore, the solution $(2^{2\bar{\gamma} + 1} \hat{x}, 2^{\bar{\gamma} + 1} \hat{y})$ satisfies the second set of constraints in problem (3) for assortment $S_i = \emptyset$ and nest i . Combining the two cases above and noting that our choice of nest i and assortment S_i is arbitrary, we conclude that the solution $(2^{2\bar{\gamma} + 1} \hat{x}, 2^{\bar{\gamma} + 1} \hat{y})$ satisfies the second set of constraints in problem (3).

Since the solution (\hat{x}, \hat{y}) is optimal to problem (4), we have $v_0 \hat{x} \geq \sum_{i \in M} \hat{y}_i$. This implies that $v_0 2^{2\bar{\gamma} + 1} \hat{x} \geq v_0 2^{\bar{\gamma} + 1} \hat{x} \geq \sum_{i \in M} 2^{\bar{\gamma} + 1} \hat{y}_i$, in which case, the solution $(2^{2\bar{\gamma} + 1} \hat{x}, 2^{\bar{\gamma} + 1} \hat{y})$ satisfies the first constraint in problem (3) as well and we obtain the desired result. \square