

Technical Note: Capacitated Assortment Optimization under the Multinomial Logit Model with Nested Consideration Sets

Jacob Feldman

Olin Business School, Washington University,
St. Louis, MO 63130, USA
jbfeldman@wustl.edu

Huseyin Topaloglu

School of Operations Research and Information
Engineering, Cornell Tech,
New York, NY 10011, USA
topaloglu@orie.cornell.edu

August 2, 2017

Abstract

We study capacitated assortment problems when customers choose under the multinomial logit model with nested consideration sets. In this choice model, there are multiple customer types and a customer of a particular type is interested in purchasing only a particular subset of products. We use the term consideration set to refer to the subset of products that a customer of a particular type is interested in purchasing. The consideration sets of customers of different types are nested in the sense that the consideration set of one customer type is included in the consideration set of another. The choice process for customers of different types is governed by the same multinomial logit model except for the fact that customers of different types have different consideration sets. Each product, if offered to the customers, occupies a certain amount of space. The sale of each product generates a certain amount of revenue. Given that the customers choose among the offered products according to the multinomial logit model with nested consideration sets, the goal of the assortment problem is to find a set of products to offer to maximize the expected revenue obtained from a customer, while making sure that the total space consumption of the offered products does not exceed a certain limit. We show that this assortment problem is NP-hard, even when there is no limit on the total space consumption of the offered products. Motivated by this complexity result, we give a fully polynomial time approximation scheme for the problem.

Keywords: assortment optimization, customer choice models, approximation algorithms, multinomial logit model

1 Introduction

Incorporating customer choice behavior into revenue management models has been seeing considerable attention. Traditional revenue management models represent the demand for a product by using an exogenous random variable, but this way of representing the demand may be inadequate when there are multiple products that can serve the needs of a customer and the customers choose and substitute among the offered products. When the customers choose and substitute among the offered products, the demand for a particular product depends on what other products are offered, creating complex interactions between the demands for the different products. There are numerous field studies that demonstrate the benefits from taking the customer choice process into consideration when making operational decisions. For example, Chong et al. (2001), Kok and Fisher (2007) and Misra (2008) focus on customers purchasing groceries in different food categories. These

papers indicate that there is significant potential for revenue improvement by explicitly taking the choice process of the customers into consideration, when deciding which assortment of products to make available. Similarly, Vulcano et al. (2010) focus on customers purchasing airline tickets and demonstrate that there is significant potential for revenue improvement when one explicitly accounts for the choice process of the customers among the different itineraries. When taking the choice process of the customers into consideration, a critical tradeoff is that a sophisticated choice model may capture the choice behavior of the customers faithfully, but solving operational problems under a sophisticated choice model may be intractable.

In this paper, we study capacitated assortment problems when customers choose among the offered products according to the multinomial logit model with nested consideration sets. In our assortment problem, a firm has access to a set of products among which it picks an assortment of products to offer to its customers. Each product, if offered, occupies a certain amount of space. The sale of each product generates a certain amount of revenue. Customers choose among the offered assortment of products according to a particular choice model. The goal of the firm is to pick an assortment of products to offer to maximize the expected revenue obtained from a customer, while making sure that the total space consumption of the offered products does not exceed a certain limit. We capture the choice process of the customers by using the multinomial logit model with nested consideration sets. In this choice model, there are multiple customer types. A customer of a particular type is interested in purchasing only a particular subset of products. We use the term consideration set to refer to the subset of products that a customer of a particular type is interested in purchasing. The consideration sets of customers of different types are nested, so that we can index the customer types such that the consideration set of a customer type with a smaller index is included in the consideration set of a customer type with a larger index. The choice process for customers of different types is governed by the same multinomial logit model except for the fact that customers of different types have different consideration sets. A customer observes which of the products in her consideration set are actually included in the offered assortment and she makes a choice only among those products according to the multinomial logit model.

MAIN CONTRIBUTIONS. We show that the problem of finding the assortment of products that maximizes the expected revenue obtained from a customer is NP-hard. This complexity result holds even when there is no limit on the total space consumption of the offered products. Motivated by our complexity result, we give a fully polynomial time approximation scheme (FPTAS) for the

assortment problem. This FPTAS appears to be the first to address assortment problems with consideration sets, while having a running time that is polynomial in the number of customer types. Our FPTAS is based on a dynamic programming formulation of the assortment problem. In particular, since the consideration sets of different customer types are nested, we index the customer types such that the first customer type has the smallest consideration set and the last customer type has the largest consideration set. In our dynamic program, the decision epochs correspond to the customer types in the order of increasing consideration sets. In each decision epoch, we decide whether we offer the product that is in the consideration set of the current customer type, but not in the consideration sets of the previous customer types. To develop an FPTAS, we discretize the state space of the dynamic program by using a geometric grid. Ultimately, we show that if there are n products among which the firm picks an assortment, then we obtain a solution that provides a $1 - \epsilon$ fraction of the optimal expected revenue in $O(n^4 \log(nR_{\max}) \log(nV_{\max}) \log(nR_{\max} V_{\max}/\lambda_{\min}) / \epsilon^3)$ operations, where R_{\max} is the largest value for the product of the revenue and the preference weight of a product, V_{\max} is the largest value for the preference weight of a product and λ_{\min} is the smallest value for the probability of arrival for a particular customer type.

It turns out that if we do not have a limit on the total space consumption of the offered products, then we can solve our dynamic program more efficiently. We show that the running time of our FPTAS given in the previous paragraph reduces to $O(n^3 \log(nR_{\max}) \log(nV_{\max}) / \epsilon^2)$ operations, when we do not have a limit on the total space consumption of the offered products.

MULTINOMIAL LOGIT MODEL WITH NESTED CONSIDERATION SETS. In our choice model, the consideration sets of customers of different types are nested. In particular, assuming that there are m customer types and using N_k to denote the consideration set of customers of type k , we index the customer types $\{1, \dots, m\}$ such that $N_1 \subseteq \dots \subseteq N_m$. Assuming that there are n products, we index the products $\{1, \dots, n\}$ such that products $\{1, \dots, |N_1|\}$ correspond to the products in N_1 , products $\{|N_1| + 1, \dots, |N_2|\}$ correspond to the products in $N_2 \setminus N_1$ and so on. So, without loss of generality, we can assume that the consideration set of customers of a particular type is of the form $\{1, \dots, j\}$ for some $j = 1, \dots, n$. In this case, since the customer types differ in their consideration sets, the number of customer types is at most n . Without loss of generality, we can assume that the number of customer types is n , if necessary, by defining additional dummy customer types.

Consideration sets of the form $\{1, \dots, j\}$ for some $j = 1, \dots, n$ become useful when different customers drop products from consideration based on cutoff values for a particular attribute, such

as price or quality. For example, ordering the products such that product 1 has the lowest price and product n has the highest price, customers may form their consideration sets by focusing on the products within their budget. A customer with a consideration set $\{1, \dots, j\}$ focuses on products whose prices do not exceed the price of product j . Similarly, ordering the products such that product 1 has the highest quality and product n has the lowest quality, a customer with a consideration set $\{1, \dots, j\}$ focuses on products whose quality is no worse than the quality of product j . Goyal et al. (2016) also use such nested consideration sets.

Another feature of our choice model is that the choices of the customers of different types are governed by the same multinomial logit model, except for the fact that customers of different types may have different consideration sets and associate different mean utilities with the no purchase option. This choice model naturally arises when the mean utility that a customer of a particular type associates with a particular product is a *separable function* of the features of the product and the features of the customer type. Under the multinomial logit model, a customer of type k associates the mean utility μ_{jk} with product j . If the set of offered products that are in the consideration set of customer type k is given by S , then a customer of type k purchases product j with probability $e^{\mu_{jk}} / (e^{\mu_{0k}} + \sum_{i \in S} e^{\mu_{ik}})$, where μ_{0k} is the mean utility that a customer of type k associates with the no purchase option. If μ_{jk} is a separable function of the features of the product and the features of the customer type so that $\mu_{jk} = \eta_j + \sigma_k$ for some η_j and σ_k , then the last probability becomes $e^{\eta_j + \sigma_k} / (e^{\mu_{0k}} + \sum_{i \in S} e^{\eta_i + \sigma_k}) = e^{\eta_j} / (e^{\mu_{0k} - \sigma_k} + \sum_{i \in S} e^{\eta_i})$, which is the purchase probability of product j when customers of all types associate the same mean utility η_j with product j , but customers of type k associate the mean utility $\mu_{0k} - \sigma_k$ with the no purchase option.

To give a concrete example for our choice model, we consider the choice process of customers when they shop for canned coffee from a particular store. Different coffee options are characterized by their prices and qualities. We use r_j to denote the price and q_j to denote the quality of coffee option j . Different customer types are characterized by their daily coffee consumption amounts and shopping frequencies at the particular store. We use a_k to denote the daily coffee consumption amount and f_k to denote the shopping frequency of customers of type k . One possible model for the mean utility μ_{jk} that a customer of type k associates with coffee option j is to let $\mu_{jk} = \beta^1 r_j + \beta^2 q_j + \beta^3 a_k + \beta^4 f_k$, where the parameters $\{\beta^1, \dots, \beta^4\}$ are estimated from the past purchase data. Identifying $\beta^1 r_j + \beta^2 q_j$ with σ_j and $\beta^3 a_k + \beta^4 f_k$ with η_k , we obtain the choice model in the previous paragraph. In addition, if different customers drop coffee options from consideration

based on different cutoff values for the price, then we obtain the multinomial logit model with nested consideration sets. Since the customer types are characterized by their daily coffee consumption amounts and shopping frequencies, as well as cutoff values for the price, there can be multiple customer types with the same consideration set, but this does not create a complication for our approach. An advantage of this model is that once we estimate the parameters $\{\beta^1, \dots, \beta^4\}$, given any coffee option with certain price and quality and given any customer type with certain daily coffee consumption amount and shopping frequency, including those that are not in the past purchase data, we can immediately estimate the mean utility that this particular customer attaches to this particular coffee option. A shortcoming of this model is that this model captures the fact that the mean utility that a customer attaches to a coffee option depends on its price and quality, but this dependence is, intuitively speaking, captured in terms of the reaction of an “average” customer, since the parameters β^1 and β^2 do not depend on the type of a customer. One option is to use parameters of the form β_k^1 and β_k^2 , capturing the reaction of customers of type k to price and quality, but we are not able to give a tractable approach for the assortment problem in this case.

LITERATURE REVIEW. There is recent work on models where customers form consideration sets. Aouad et al. (2016) work with consider-then-choose choice models, where each customer has a consideration set and a ranking of the products in her consideration set. She purchases the most preferred available product in her consideration set. The authors study the assortment problem for a variety of possible consideration set and ranking structures. Jagabathula and Vulcano (2015) develop a choice model, where the consideration set of a customer includes only the product purchased during the last visit and the products in promotion, if there are any. They focus on estimating the parameters of their choice model. Jagabathula and Rusmevichientong (2015) use a choice model where a customer considers only the products whose prices are below a certain threshold. The authors focus on parameter estimation and pricing problems. Sahin and Wang (2014) develop a choice model that incorporates the product search cost so that the set of products that a customer considers purchasing can be different from what the firm offers. Dai et al. (2014) develop a revenue management model, where a customer may not even consider purchasing some of the offered itineraries due to personal restrictions on, for example, time of departure.

In a mixture of multinomial logit models, there are multiple customer types and customers of different types choose according to different multinomial logit models. The multinomial logit models for the different customer types may have completely different parameters and consideration

sets. McFadden and Train (2000) show that a broad class of choice models can be approximated arbitrarily accurately by using a mixture of multinomial logit models. Our choice model is a special case of a mixture of multinomial logit models, where the consideration sets are nested and customers of different types choose according to the same multinomial logit model except for the fact that their consideration sets are different. Talluri and van Ryzin (2004) and Gallego et al. (2004) show that assortment problems under the multinomial logit model with a single customer type can be solved efficiently. Bront et al. (2009), Desir and Goyal (2014) and Rusmevichientong et al. (2014) give approximation schemes and heuristics for assortment problems under a mixture of multinomial logit models. The running times of their approximation schemes increase exponentially with the number of customer types, whereas our running times increase polynomially.

Network revenue management problems focus on dynamically controlling the availability of products when there are multiple resources and the sale of a product consumes the capacity of a combination of resources. Gallego et al. (2004) give a deterministic linear programming approximation for network revenue management problems, where the number of decision variables increases exponentially with the number of products. Thus, the deterministic approximation is commonly solved by using column generation. The column generation subproblem has the same structure as our assortment problem when customers choose under the multinomial logit model with nested consideration sets. Gallego et al. (2016) discuss that if we can solve the column generation subproblem with a certain optimality gap, then we can also solve the deterministic approximation with a certain optimality gap. Talluri (2011) focuses on the deterministic approximation when the consideration sets of different customer types are disjoint or have small overlaps. Meissner et al. (2012) study a tractable relaxation of the deterministic approximation and add the so called product cuts to tighten it. Stauss and Talluri (2016) use a graph to represent the overlap between the consideration sets of different customer types and focus on solving the deterministic approximation when this graph is a tree. Kunnumkal and Talluri (2014) use a relaxation of the dynamic programming formulation of the network revenue management problem, which yields a tractable approach when the consideration set of each customer type is small.

ORGANIZATION. In Section 2, we give a formulation for our assortment problem and show that our assortment problem is NP-hard. In Section 3, we formulate our assortment problem as a dynamic program and give an approximation to this dynamic program. In Section 4, we show how to use the approximate dynamic program to estimate the optimal expected revenue with a certain

relative gap. In Section 5, we use our analysis of the approximate dynamic program to develop our FPTAS. In Section 6, we give a numerical study for our FPTAS. In Section 7, we conclude.

2 Assortment Problem and Computational Complexity

In our assortment problem, there are n products indexed by $\{1, \dots, n\}$. Associated with product j , we have a revenue r_j and a space consumption c_j . The limit on the total space consumption of the offered products is C . We use the vector $x = (x_1, \dots, x_n) \in \{0, 1\}^n$ to capture the set of products offered to customers, where $x_j = 1$ if product j is offered and $x_j = 0$ if product j is not offered. There are n customer types indexed by $\{1, \dots, n\}$. A customer arriving into the system is of type k with probability λ_k . A customer of type k considers purchasing only the products in the set $\{1, \dots, k\}$. We refer to the set $N_k = \{1, \dots, k\}$ as the consideration set of a customer of type k . A customer of type k chooses among the products in her consideration set according to the multinomial logit model. In particular, using v_j to denote the preference weight of product j and v_0 to denote the preference weight of the no purchase option, if the set of products offered to the customers is given by the vector x , then a customer of type k chooses product $j \in N_k$ with probability $P_{jk}(x) = v_j x_j / (v_0 + \sum_{i \in N_k} v_i x_i)$. Thus, if the set of products offered to the customers is given by the vector x , then we obtain an expected revenue of $\sum_{k=1}^n \lambda_k \sum_{j \in N_k} r_j P_{jk}(x)$ from a customer. Our goal is to find the set of products to offer to maximize the expected revenue obtained from each customer, while making sure that the total space consumption of the offered products does not exceed the limit on the total space consumption, yielding the problem

$$z^* = \max_{\substack{x \in \{0, 1\}^n : \\ \sum_{j=1}^n c_j x_j \leq C}} \left\{ \sum_{k=1}^n \lambda_k \left\{ \sum_{j \in N_k} r_j P_{jk}(x) \right\} \right\} = \max_{\substack{x \in \{0, 1\}^n : \\ \sum_{j=1}^n c_j x_j \leq C}} \left\{ \sum_{k=1}^n \lambda_k \left\{ \frac{\sum_{j=1}^k r_j v_j x_j}{v_0 + \sum_{j=1}^k v_j x_j} \right\} \right\}. \quad (1)$$

In problem (1), we can allow the preference weight v_0 of the no purchase option to depend on the customer type. In particular, if a customer of type k associates the preference weight v_{0k} with the no purchase option, then we can simply replace v_0 in problem (1) with v_{0k} . All of our results continue to hold under this extension with essentially no modifications and we obtain an FPTAS with the same running time.

COMPUTATIONAL COMPLEXITY. In problem (1), the consideration sets are nested and customers of different types associate the same preference weight with a particular product, as long as this product is in their consideration sets. Rusmevichientong et al. (2014) consider assortment problems under a mixture of multinomial logit models, where consideration sets are arbitrary and customers of different types associate arbitrary preference weights with a particular product. The

authors show that the feasibility version of their assortment problem is NP-complete even if there are only two customer types. In the remainder of this section, we show that the feasibility version of problem (1) is also NP-complete, despite the fact this problem has the special structure where the consideration sets are nested and customers of different types associate the same preference weight with a particular product. This complexity result holds even when there is no limit on the total space consumption of the offered products. To discuss the computational complexity of problem (1), we focus on the following feasibility version of the problem.

ASSORTMENT FEASIBILITY. We are given an expected revenue threshold K . The assortment feasibility problem asks whether there exists a vector $x \in \{0, 1\}^n$ that provides an expected revenue of K or more in problem (1).

In the next theorem, we show that the feasibility version of problem (1) is NP-complete. We defer the proof to Appendix A. The proof of this theorem uses a reduction from the partition problem, which is a well-known NP-complete problem; see Garey and Johnson (1979).

Theorem 1 *The assortment feasibility problem is NP-complete.*

The computational complexity result in Theorem 1 motivates us to look for an FPTAS for problem (1). In the next section, we give a dynamic programming formulation for problem (1). This dynamic program forms the starting point for our FPTAS.

3 Dynamic Programming Formulation

To develop an FPTAS for problem (1), we follow three steps. First, we formulate problem (1) as a dynamic program, but there are exponentially many possible values for the state variable in this dynamic programming formulation. So, we give an approximation to the dynamic programming formulation, where there are polynomially many possible values for the state variable. Second, noting that z^* is the optimal expected revenue in problem (1), we show that we can use the value functions in the approximate dynamic program to obtain an estimate \tilde{z} of the optimal expected revenue that satisfies $\tilde{z} \geq (1 - \epsilon)z^*$ for an appropriate choice of $\epsilon \in (0, 1)$. Third, we follow the optimal state and action trajectory in the approximate dynamic program to obtain a feasible solution to problem (1). Also, using REV to denote the expected revenue corresponding to this feasible solution, we show that $\text{REV} \geq \tilde{z}$. In this case, we obtain $\text{REV} \geq \tilde{z} \geq (1 - \epsilon)z^*$, which

bounds the relative gap between the expected revenue from the solution obtained by using the approximate dynamic program and the optimal expected revenue. In this section, we focus on the first step. In the next two sections, we focus on the second and third steps.

To formulate problem (1) as a dynamic program, assume that we have already decided which of the products in $\{1, \dots, j-1\}$ are offered to the customers and these decisions are given by $(x_1, \dots, x_{j-1}) \in \{0, 1\}^{j-1}$, where $x_k = 1$ if product k is offered and $x_k = 0$ if product k is not offered. We use P_j to denote $\sum_{k=1}^{j-1} r_k v_k x_k$ and S_j to denote $\sum_{k=1}^{j-1} v_k x_k$. In this case, the expected revenue from a customer of type j is given by $\lambda_j \sum_{k=1}^j r_k v_k x_k / (v_0 + \sum_{k=1}^j v_k x_k) = \lambda_j (P_j + r_j v_j x_j) / (v_0 + S_j + v_j x_j)$. Therefore, once we decide whether product j is offered, we can compute the expected revenue from a customer of type j as a function of only P_j , S_j and x_j . Furthermore, we can compute $P_{j+1} = \sum_{k=1}^j r_k v_k x_k$ and $S_{j+1} = \sum_{k=1}^j v_k x_k$ as $P_{j+1} = P_j + r_j v_j x_j$ and $S_{j+1} = S_j + v_j x_j$, which are also functions of only P_j , S_j and x_j . So, given that the decisions for the products in $\{1, \dots, j-1\}$ satisfy $P_j = \sum_{k=1}^{j-1} r_k v_k x_k$ and $S_j = \sum_{k=1}^{j-1} v_k x_k$ and we want to generate an expected revenue of R_j or more from customers of type $\{j, \dots, n\}$, we let $V_j(P_j, S_j, R_j)$ be the minimum total space consumption of the products in $\{j, \dots, n\}$ to do so. We can compute $\{V_j(\cdot, \cdot, \cdot) : j = 1, \dots, n\}$ by using the dynamic program

$$V_j(P_j, S_j, R_j) = \min_{x_j \in \{0, 1\}} \left\{ c_j x_j + V_{j+1} \left(P_j + r_j v_j x_j, S_j + v_j x_j, R_j - \lambda_j \frac{P_j + r_j v_j x_j}{v_0 + S_j + v_j x_j} \right) \right\}, \quad (2)$$

with the boundary condition that $V_{n+1}(\cdot, \cdot, R_{n+1}) = 0$ if $R_{n+1} \leq 0$ and $V_{n+1}(\cdot, \cdot, R_{n+1}) = \infty$ if $R_{n+1} > 0$. The dynamic program above follows from the following reasoning. Given that the decisions for the products in $\{1, \dots, j-1\}$ satisfy $P_j = \sum_{k=1}^{j-1} r_k v_k x_k$ and $S_j = \sum_{k=1}^{j-1} v_k x_k$, as mentioned above, the values of $P_{j+1} = \sum_{k=1}^j r_k v_k x_k$ and $S_{j+1} = \sum_{k=1}^j v_k x_k$ can be computed as $P_{j+1} = P_j + r_j v_j x_j$ and $S_{j+1} = S_j + v_j x_j$. Furthermore, if we want to generate an expected revenue of R_j or more from customers of type $\{j, \dots, n\}$, then after making the decision for product j , the expected revenue that we want to generate from customers of type $\{j+1, \dots, n\}$ is at least $R_j - \lambda_j \frac{P_j + r_j v_j x_j}{v_0 + S_j + v_j x_j}$. The boundary condition follows from the fact that if we make the decisions for all of the products and we still have a strictly positive expected revenue that we want to generate, then there is no set of products that provides the expected revenue that we want to generate. We refer to $V_j(\cdot, \cdot, \cdot)$ as the value function and (P_j, S_j, R_j) as the state variable for decision epoch j . To compute the value functions $\{V_j(\cdot, \cdot, \cdot) : j = 1, \dots, n\}$, we move over the decision epochs in reverse order. For all $j = n, \dots, 1$, given that we know the value function $V_{j+1}(\cdot, \cdot, \cdot)$, we compute the value

function $V_j(\cdot, \cdot, \cdot)$. Note that $V_1(0, 0, R_1)$ is the minimum total space consumption of the products in $\{1, \dots, n\}$ given that we want to generate an expected revenue of R_1 or more from customers of type $\{1, \dots, n\}$. Thus, the optimal objective value of problem (1) is given by

$$z^* = \max_{z \in \mathbb{R}_+} \left\{ z : V_1(0, 0, z) \leq C \right\}, \quad (3)$$

which is the largest expected revenue that we can generate from customers of type $\{1, \dots, n\}$ while ensuring that the total space consumption of the products in $\{1, \dots, n\}$ does not exceed C . It is important to emphasize that the dynamic program in (2) does not provide a tractable solution approach for problem (1). In particular, since there are exponentially many possible sets of products that we can offer to the customers, there are exponentially many possible values for the state variable (P_j, S_j, R_j) as well, which implies that computing the value function $V_j(P_j, S_j, R_j)$ for all possible values for the state variable (P_j, S_j, R_j) is intractable. In the remainder of this section, we give an approximate version of the dynamic program in (2), which is based on discretizing the state space in the original dynamic program. The approximate dynamic program forms the basis for the FPTAS that we develop for problem (1).

To approximate the dynamic program in (2), we choose a value of $\rho > 0$ and focus on the grid points in the domain $\text{Dom}^\rho = \{0\} \cup \{(1 + \rho)^k : k = \dots, -1, 0, 1, \dots\}$, which is a geometric grid of size $1 + \rho$ augmented by an additional point at zero. The specific choice of ρ is determined later, but we hint that the ultimate choice of $\rho = \epsilon/6n$ will yield a performance guarantee of $1 - \epsilon$ for $\epsilon \in (0, 1)$. We define the round down operator $\lfloor \cdot \rfloor$ that rounds its argument down to the closest value in Dom^ρ . In other words, we have $\lfloor x \rfloor = \max\{y \in \text{Dom}^\rho : y \leq x\}$. Similarly, we define the round up operator $\lceil \cdot \rceil$ that rounds its argument up to the closest value in Dom^ρ , which is given by $\lceil x \rceil = \min\{y \in \text{Dom}^\rho : y \geq x\}$. The results of the round up and round down operators $\lceil \cdot \rceil$ and $\lfloor \cdot \rfloor$ depend on the value of ρ , but for notational brevity, we do not make this dependence explicit. We consider an approximate version of the dynamic program in (2) given by

$$\Theta_j(P_j, S_j, R_j) = \min_{x_j \in \{0, 1\}} \left\{ c_j x_j + \Theta_{j+1} \left(\lfloor P_j + r_j v_j x_j \rfloor, \lceil S_j + v_j x_j \rceil, \lceil R_j - \lambda_j \frac{P_j + r_j v_j x_j}{v_0 + S_j + v_j x_j} \rceil \right) \right\}, \quad (4)$$

with the boundary condition that $\Theta_{n+1}(\cdot, \cdot, R_{n+1}) = 0$ if $R_{n+1} \leq 0$ and $\Theta_{n+1}(\cdot, \cdot, R_{n+1}) = \infty$ if $R_{n+1} > 0$. On the right side above, we observe that we round the first component of the state variable down, whereas we round the second and third components of the state variable up, which

ultimately allows us to bound the gap between the value functions $V_j(\cdot, \cdot, \cdot)$ and $\Theta_j(\cdot, \cdot, \cdot)$ in (2) and (4). Noting that the results of the round up and down operators depend on the value of ρ , the value functions $\{\Theta_j(\cdot, \cdot, \cdot) : j = 1, \dots, n\}$ depend on the value of ρ as well, but for notational brevity, we do not make this dependence explicit. We can bound each component of the state variable (P_j, S_j, R_j) in the dynamic program in (4). In particular, without loss of generality, we assume that $r_j \geq 1$ and $v_j \geq 1$ for all $j = 1, \dots, n$. If either one of these conditions does not hold, then we can scale up all revenues or all preference weights by the same factor without changing the optimal solution to problem (1). For notational brevity, we let $R_{\max} = \max\{r_j v_j : j = 1, \dots, n\}$ and $V_{\max} = \max\{v_j : j = 0, \dots, n\}$. In the next lemma, we give upper bounds on each component of the state variable (P_j, S_j, R_j) . The proof of this lemma uses a simple induction over the decision epochs and we defer the proof to Appendix B.

Lemma 2 *For any $(x_1, \dots, x_n) \in \{0, 1\}^n$ and $R_1 \in \mathfrak{R}_+$, assume that $\{(P_j, S_j, R_j) : j = 1, \dots, n\}$ are computed as $P_{j+1} = \lfloor P_j + r_j v_j x_j \rfloor$, $S_{j+1} = \lceil S_j + v_j x_j \rceil$ and $R_{j+1} = \lceil R_j - \lambda_j \frac{P_j + r_j v_j x_j}{v_0 + S_j + v_j x_j} \rceil$ with $P_1 = 0$ and $S_1 = 0$. Then, letting $\gamma = (\exp(n\rho) - 1)/\rho$, we have $P_j \leq nR_{\max}$, $S_j \leq \gamma V_{\max}$ and $R_j \leq R_{n+1} + (\lambda_j + \dots + \lambda_n) R_{\max}$ for all $j = 1, \dots, n$.*

We can use Lemma 2 to give upper and lower bounds on each component of the state variable (P_j, S_j, R_j) as follows. Since $r_j \geq 1$ and $v_j \geq 1$ for all $j = 1, \dots, n$, the smallest nonzero values of P_j and S_j are one. Also, by Lemma 2, we have $P_j \leq nR_{\max}$ and $S_j \leq \gamma V_{\max}$. Therefore, we can assume that $P_j \in \{0\} \cup [1, nR_{\max}]$ and $S_j \in \{0\} \cup [1, \gamma V_{\max}]$ for all $j = 1, \dots, n$. Also, if $R_j > R_{\max}$, then by Lemma 2, we have $R_{n+1} \geq R_j - (\lambda_j + \dots + \lambda_n) R_{\max} > (1 - (\lambda_j + \dots + \lambda_n)) R_{\max} \geq 0$. In this case, noting the boundary condition of the dynamic program in (4), if $R_j > R_{\max}$, then we immediately have $\Theta_j(\cdot, \cdot, R_j) = \infty$. In other words, there is no need to compute $\Theta_j(\cdot, \cdot, R_j)$ explicitly for $R_j > R_{\max}$. On the other hand, we let $\lambda_{\min} = \min\{\lambda_j : j = 1, \dots, n\}$ for notational brevity. By the discussion in this paragraph, we have $P_j \in \{0\} \cup [1, nR_{\max}]$ and $S_j \in \{0\} \cup [1, \gamma V_{\max}]$. In this case, noting that $r_j \geq 1$ and $v_j \geq 1$ for all $j = 1, \dots, n$, if $\lambda_j (P_j + r_j v_j x_j)/(v_0 + S_j + v_j x_j)$ is nonzero, then it satisfies $\lambda_j (P_j + r_j v_j x_j)/(v_0 + S_j + v_j x_j) \geq \lambda_{\min}/((2 + \gamma) V_{\max})$. That is, the quantity $\lambda_j (P_j + r_j v_j x_j)/(v_0 + S_j + v_j x_j)$ is either zero or is at least $\lambda_{\min}/((2 + \gamma) V_{\max})$. This observation implies that if the expected revenue that we want to generate from customer types $\{j, \dots, n\}$ is above zero but below $\lambda_{\min}/((2 + \gamma) V_{\max})$, then we can increase the expected revenue that we want to generate up to $\lambda_{\min}/((2 + \gamma) V_{\max})$, since offering any of the products to any of

the customer types immediately yields an expected revenue of at least $\lambda_{\min}/((2 + \gamma) V_{\max})$. In this case, we can assume that $R_j \in \{0\} \cup [\lambda_{\min}/((2 + \gamma) V_{\max}), R_{\max}]$. On the right side of (4), we need $\Theta_{j+1}(P_{j+1}, S_{j+1}, R_{j+1})$ for the values of $(P_{j+1}, S_{j+1}, R_{j+1}) \in \text{Dom}^\rho \times \text{Dom}^\rho \times \text{Dom}^\rho$. So, noting the bounds for the different components of the state variable (P_j, S_j, R_j) , we need to store the value of $\Theta_j(P_j, S_j, R_j)$ only for $(P_j, S_j, R_j) \in \text{Dom}^\rho \times \text{Dom}^\rho \times \text{Dom}^\rho$ such that $P_j \in \{0\} \cup [1, \lceil nR_{\max} \rceil]$, $S_j \in \{0\} \cup [1, \lceil \gamma V_{\max} \rceil]$ and $R_j \in \{0\} \cup [\lceil \lambda_{\min}/((2 + \gamma) V_{\max}) \rceil, \lceil R_{\max} \rceil]$.

4 Estimating the Optimal Expected Revenue

Noting that z^* is the optimal expected revenue in problem (1), in this section, we show that we can use the approximate dynamic program in (4) to obtain an estimate \tilde{z} of the optimal expected revenue that satisfies $\tilde{z} \geq (1 - \epsilon) z^*$ for an appropriate choice of $\epsilon \in (0, 1)$. In the next lemma, we give elementary properties of the value functions computed through the dynamic program in (4). The first part of the lemma shows that if we decrease P_j , increase S_j or increase R_j , then $\Theta_j(P_j, S_j, R_j)$ increases. The second part of the lemma shows how we can counterbalance a decrease in P_j and an increase in S_j by a decrease in R_j to ensure that the value function $\Theta_j(P_j, S_j, R_j)$ does not increase. In particular, if we decrease P_j by at most a factor of $1 + \rho$ and increase S_j by at most a factor of $1 + \rho$, then we can decrease R_j by at least a factor of $(1 + \rho)^2$ to ensure that $\Theta_j(P_j, S_j, R_j)$ does not increase. We defer the proof of this lemma to Appendix C.

Lemma 3 (a) *Assume that (P_j, S_j, R_j) and $(\hat{P}_j, \hat{S}_j, \hat{R}_j)$ satisfy $\hat{P}_j \leq P_j$, $\hat{S}_j \geq S_j$ and $\hat{R}_j \geq R_j$. Then, we have $\Theta_j(\hat{P}_j, \hat{S}_j, \hat{R}_j) \geq \Theta_j(P_j, S_j, R_j)$ for all $j = 1, \dots, n$.*

(b) *Assume that (P_j, S_j, R_j) and $(\hat{P}_j, \hat{S}_j, \hat{R}_j)$ satisfy $\hat{P}_j \geq P_j/(1 + \rho)$, $\hat{S}_j \leq (1 + \rho)S_j$ and $\hat{R}_j \leq R_j/(1 + \rho)^2$. Then, we have $\Theta_j(\hat{P}_j, \hat{S}_j, \hat{R}_j) \leq \Theta_j(P_j, S_j, R_j)$ for all $j = 1, \dots, n$.*

In the next proposition, we build on Lemma 3 to show that we can decrease the third component of the state variable (P_j, S_j, R_j) by a factor of $(1 + \rho)^{3(n-j)}$ to obtain a lower bound on the value function $V_j(\cdot, \cdot, \cdot)$ by using the value function $\Theta_j(\cdot, \cdot, \cdot)$. This proposition is the critical ingredient in the development of our FPTAS and it ultimately allows us to obtain an estimate \tilde{z} of the optimal expected revenue that satisfies $\tilde{z} \geq z^*/(1 + \rho)^{3n}$.

Proposition 4 *For all $j = 1, \dots, n$, we have $V_j(P_j, S_j, R_j) \geq \Theta_j(P_j, S_j, R_j / (1 + \rho)^{3(n-j)})$.*

Proof. We show the result by using induction over the decision epochs in reverse order. For the decision epoch $n + 1$, by the boundary condition of the dynamic program in (2), we have $V_{n+1}(\cdot, \cdot, R_{n+1}) = 0$ if $R_{n+1} \leq 0$ and $V_{n+1}(\cdot, \cdot, R_{n+1}) = \infty$ if $R_{n+1} > 0$. Similarly, by the boundary condition of the dynamic program in (4), we have $\Theta_{n+1}(\cdot, \cdot, R_{n+1}) = 0$ if $R_{n+1} \leq 0$ and $\Theta_{n+1}(\cdot, \cdot, R_{n+1}) = \infty$ if $R_{n+1} > 0$. We have $R_{n+1} \leq 0$ if and only if $(1 + \rho)^3 R_{n+1} \leq 0$. So, we obtain $V_{n+1}(\cdot, \cdot, R_{n+1}) = \Theta_{n+1}(\cdot, \cdot, (1 + \rho)^3 R_{n+1})$ and the result holds for decision epoch $n + 1$. Assuming that the result holds for decision epoch $j + 1$, we show that the result holds for decision epoch j . For notational brevity, we let $\beta = 1/(1 + \rho)^{3(n-j)}$. We fix an arbitrary value for the state variable (P_j, S_j, R_j) and $x_j \in \{0, 1\}$. For the fixed values of (P_j, S_j, R_j) and $x_j \in \{0, 1\}$, we let $P_{j+1} = P_j + r_j v_j x_j$, $S_{j+1} = S_j + v_j x_j$, $R_{j+1} = R_j - \lambda_j \frac{P_j + r_j v_j x_j}{v_0 + S_j + v_j x_j}$ and $R_{j+1}^\beta = \beta R_j - \lambda_j \frac{P_j + r_j v_j x_j}{v_0 + S_j + v_j x_j}$. By the induction assumption, we have

$$V_{j+1}(P_{j+1}, S_{j+1}, R_{j+1}) \geq \Theta_{j+1}(P_{j+1}, S_{j+1}, R_{j+1}/(1 + \rho)^{3(n-j-1)}). \quad (5)$$

Furthermore, we observe that we can lower bound $R_{j+1}/(1 + \rho)^{3(n-j-1)}$ by $(1 + \rho)^3 R_{j+1}^\beta$, which follows from the chain of inequalities

$$\begin{aligned} \frac{1}{(1 + \rho)^{3(n-j-1)}} R_{j+1} &= (1 + \rho)^3 \beta R_{j+1} = (1 + \rho)^3 \beta \left\{ R_j - \lambda_j \frac{P_j + r_j v_j x_j}{v_0 + S_j + v_j x_j} \right\} \\ &\geq (1 + \rho)^3 \left\{ \beta R_j - \lambda_j \frac{P_j + r_j v_j x_j}{v_0 + S_j + v_j x_j} \right\} = (1 + \rho)^3 R_{j+1}^\beta. \end{aligned}$$

Using the first part of Lemma 3 for decision epoch $j + 1$, $\Theta_{j+1}(\cdot, \cdot, R_{j+1})$ is increasing in R_{j+1} , in which case, noting the inequality above, we obtain

$$\Theta_{j+1}(P_{j+1}, S_{j+1}, R_{j+1}/(1 + \rho)^{3(n-j-1)}) \geq \Theta_{j+1}(P_{j+1}, S_{j+1}, (1 + \rho)^3 R_{j+1}^\beta). \quad (6)$$

Since $\lfloor P_{j+1} \rfloor \geq P_{j+1}/(1 + \rho)$, $\lceil S_{j+1} \rceil \leq (1 + \rho) S_{j+1}$ and $(1 + \rho) R_{j+1}^\beta = (1 + \rho)^3 R_{j+1}^\beta / (1 + \rho)^2$ by the definition of $\lceil \cdot \rceil$ and $\lfloor \cdot \rfloor$, using the second part of Lemma 3 for decision epoch $j + 1$, we get

$$\Theta_{j+1}(P_{j+1}, S_{j+1}, (1 + \rho)^3 R_{j+1}^\beta) \geq \Theta_{j+1}(\lfloor P_{j+1} \rfloor, \lceil S_{j+1} \rceil, (1 + \rho) R_{j+1}^\beta). \quad (7)$$

For the moment, assume that $R_{j+1}^\beta > 0$. Since $(1 + \rho) R_{j+1}^\beta \geq \lceil R_{j+1}^\beta \rceil$ and $\Theta_{j+1}(\cdot, \cdot, R_{j+1})$ is increasing in R_{j+1} by the first part of Lemma 3, we get $\Theta_{j+1}(\lfloor P_{j+1} \rfloor, \lceil S_{j+1} \rceil, (1 + \rho) R_{j+1}^\beta) \geq$

$\Theta_{j+1}(\lfloor P_{j+1} \rfloor, \lceil S_{j+1} \rceil, \lceil R_{j+1}^\beta \rceil)$. Next, assume that $R_{j+1}^\beta \leq 0$. The third component of the state variable in the dynamic program in (4) can only decrease as the decision epochs progress and it is the sign of this component that affects the boundary condition. Thus, if \hat{R}_{j+1} and \tilde{R}_{j+1} both satisfy $\hat{R}_{j+1} \leq 0$ and $\tilde{R}_{j+1} \leq 0$, then we have $\Theta_{j+1}(\cdot, \cdot, \hat{R}_{j+1}) = 0 = \Theta_{j+1}(\cdot, \cdot, \tilde{R}_{j+1})$. So, if $(1 + \rho) R_{j+1}^\beta \leq 0$, then we have $\lceil R_{j+1}^\beta \rceil \leq 0$ as well and we obtain $\Theta_{j+1}(\lfloor P_{j+1} \rfloor, \lceil S_{j+1} \rceil, (1 + \rho) R_{j+1}^\beta) = 0 = \Theta_{j+1}(\lfloor P_{j+1} \rfloor, \lceil S_{j+1} \rceil, \lceil R_{j+1}^\beta \rceil)$. Thus, under both assumptions, we obtain

$$\Theta_{j+1}(\lfloor P_{j+1} \rfloor, \lceil S_{j+1} \rceil, (1 + \rho) R_{j+1}^\beta) \geq \Theta_{j+1}(\lfloor P_{j+1} \rfloor, \lceil S_{j+1} \rceil, \lceil R_{j+1}^\beta \rceil). \quad (8)$$

Noting the definitions of P_{j+1} , S_{j+1} , R_{j+1} and R_{j+1}^β and combining the inequalities in (5), (6), (7) and (8), it follows that

$$\begin{aligned} V_{j+1}\left(P_j + r_j v_j x_j, S_j + v_j x_j, R_j - \lambda_j \frac{P_j + r_j v_j x_j}{v_0 + S_j + v_j x_j}\right) &= V_{j+1}(P_{j+1}, S_{j+1}, R_{j+1}) \\ &\geq \Theta_{j+1}(\lfloor P_{j+1} \rfloor, \lceil S_{j+1} \rceil, \lceil R_{j+1}^\beta \rceil) = \Theta_{j+1}\left(\lfloor P_j + r_j v_j x_j \rfloor, \lceil S_j + v_j x_j \rceil, \lceil \beta R_j - \lambda_j \frac{P_j + r_j v_j x_j}{v_0 + S_j + v_j x_j} \rceil\right). \end{aligned}$$

Since our choice of $x_j \in \{0, 1\}$ is arbitrary, the inequality above holds for any $x_j \in \{0, 1\}$. Adding $c_j x_j$ to both sides of the inequality above and taking the minimum over $x_j \in \{0, 1\}$, we get

$$\begin{aligned} V_j(P_j, S_j, R_j) &= \min_{x_j \in \{0, 1\}} \left\{ c_j x_j + V_{j+1}\left(P_j + r_j v_j x_j, S_j + v_j x_j, R_j - \lambda_j \frac{P_j + r_j v_j x_j}{v_0 + S_j + v_j x_j}\right) \right\} \\ &\geq \min_{x_j \in \{0, 1\}} \left\{ c_j x_j + \Theta_{j+1}\left(\lfloor P_j + r_j v_j x_j \rfloor, \lceil S_j + v_j x_j \rceil, \lceil \beta R_j - \lambda_j \frac{P_j + r_j v_j x_j}{v_0 + S_j + v_j x_j} \rceil\right) \right\} = \Theta_j(P_j, S_j, \beta R_j), \end{aligned}$$

where the two equalities above use (2) and (4). Since $\beta = 1/(1 + \rho)^{3(n-j)}$, the chain of inequalities above yields $V_j(P_j, S_j, R_j) \geq \Theta_j(P_j, S_j, R_j/(1 + \rho)^{3(n-j)})$, which is the desired result. \square

In the next corollary, we build on Proposition 4 to show that we can use $\{\Theta_j(\cdot, \cdot, \cdot) : j = 1, \dots, n\}$ to obtain an estimate \tilde{z} of the optimal expected revenue that satisfies $\tilde{z} \geq z^*/(1 + \rho)^{3n}$.

Corollary 5 *Let $\tilde{z} = \max\{z \in \text{Dom}^\rho : \Theta_1(0, 0, z) \leq C, z \in [\lfloor \lambda_{\min}/((2 + \gamma) V_{\max}) \rfloor, \lceil R_{\max} \rceil]\}$ and z^* be the optimal objective value of problem (1). Then, we have $\tilde{z} \geq z^*/(1 + \rho)^{3n}$.*

Proof. First, assume that $\lfloor z^*/(1 + \rho)^{3(n-1)} \rfloor \geq \lfloor \lambda_{\min}/((2 + \gamma) V_{\max}) \rfloor$. It is simple to check that the optimal objective value of problem (1) satisfies $z^* \leq R_{\max}$ so that $\lfloor z^*/(1 + \rho)^{3(n-1)} \rfloor \leq z^* \leq \lceil R_{\max} \rceil$. Note that z^* is also given by the optimal objective value of problem (3). Thus,

we have $V_1(0, 0, z^*) \leq C$. So, it follows that $C \geq V_1(0, 0, z^*) \geq \Theta_1(0, 0, z^*/(1 + \rho)^{3(n-1)}) \geq \Theta_1(0, 0, \lfloor z^*/(1 + \rho)^{3(n-1)} \rfloor)$, where the second inequality is by Proposition 4 and the third inequality is by the first part of Lemma 3. This discussion shows that $\lfloor z^*/(1 + \rho)^{3(n-1)} \rfloor$ is a feasible solution to the problem in the corollary providing an objective value of $\lfloor z^*/(1 + \rho)^{3(n-1)} \rfloor$. Since \tilde{z} is the optimal objective value of this problem, we get $\tilde{z} \geq \lfloor z^*/(1 + \rho)^{3(n-1)} \rfloor \geq z^*/(1 + \rho)^{3n}$, where the second inequality is by the fact that $\lfloor (1 + \rho)x \rfloor \geq x$ for $x \in \mathfrak{R}_+$. Second, assume that $\lfloor z^*/(1 + \rho)^{3(n-1)} \rfloor < \lfloor \lambda_{\min}/((2 + \gamma)V_{\max}) \rfloor$. Since \tilde{z} is a feasible solution to the problem in the corollary, we have $\tilde{z} \geq \lfloor \lambda_{\min}/((2 + \gamma)V_{\max}) \rfloor > \lfloor z^*/(1 + \rho)^{3(n-1)} \rfloor \geq z^*/(1 + \rho)^{3n}$. \square

Corollary 5 does not yet yield a solution to problem (1) with a performance guarantee, since due to the round up and down operators in the dynamic program in (4), the estimate \tilde{z} of the optimal expected revenue may not correspond to the expected revenue from a solution.

5 Fully Polynomial Time Approximation Scheme

In this section, we show that we can follow the optimal state and action trajectory in the dynamic program in (4) to obtain a feasible solution to problem (1). Also, the expected revenue from this solution is no worse than \tilde{z} given in Corollary 5. Using these results, we give our FPTAS. The next algorithm follows the optimal state and action trajectory in the dynamic program in (4).

Step 0. Using the dynamic program in (4), for all $j = 1, \dots, n$, compute $\Theta_j(P_j, S_j, R_j)$ for all $(P_j, S_j, R_j) \in \text{Dom}^\rho \times \text{Dom}^\rho \times \text{Dom}^\rho$ such that $P_j \in \{0\} \cup [1, \lceil nR_{\max} \rceil]$, $S_j \in \{0\} \cup [1, \lceil \gamma V_{\max} \rceil]$ and $R_j \in \{0\} \cup [\lfloor \lambda_{\min}/((2 + \gamma)V_{\max}) \rfloor, \lceil R_{\max} \rceil]$. Set $j = 1$, $\tilde{P}_1 = 0$, $\tilde{S}_1 = 0$ and $\tilde{R}_1 = \max\{z \in \text{Dom}^\rho : \Theta_1(0, 0, z) \leq C, z \in [\lfloor \lambda_{\min}/((2 + \gamma)V_{\max}) \rfloor, \lceil R_{\max} \rceil]\}$.

Step 1. Given the values of \tilde{P}_j , \tilde{S}_j and \tilde{R}_j , to decide whether to offer product j , let \tilde{x}_j be an optimal solution to the problem

$$\min_{x_j \in \{0,1\}} \left\{ c_j x_j + \Theta_{j+1} \left(\lfloor \tilde{P}_j + r_j v_j x_j \rfloor, \lceil \tilde{S}_j + v_j x_j \rceil, \lceil \tilde{R}_j - \lambda_j \frac{\tilde{P}_j + r_j v_j x_j}{v_0 + \tilde{S}_j + v_j x_j} \rceil \right) \right\}. \quad (9)$$

Set $\tilde{P}_{j+1} = \lfloor \tilde{P}_j + r_j v_j \tilde{x}_j \rfloor$, $\tilde{S}_{j+1} = \lceil \tilde{S}_j + v_j \tilde{x}_j \rceil$ and $\tilde{R}_{j+1} = \lceil \tilde{R}_j - \lambda_j \frac{\tilde{P}_j + r_j v_j \tilde{x}_j}{v_0 + \tilde{S}_j + v_j \tilde{x}_j} \rceil$.

Step 2. Increase j by 1. If $j \leq n$, then go to Step 1. Otherwise, stop and return $(\tilde{x}_1, \dots, \tilde{x}_n)$.

Since the results of the round up and down operators depend on the value of ρ , the output of the algorithm above depends on the value of ρ as well and we refer to this algorithm

as the APPRX(ρ) algorithm, where APPRX stands for approximation and ρ emphasizes its dependence on ρ . The APPRX(ρ) algorithm follows the optimal state and action trajectory in the dynamic program in (4) starting from the initial state $(\tilde{P}_1, \tilde{S}_1, \tilde{R}_1)$, where $\tilde{P}_1 = 0$, $\tilde{S}_1 = 0$ and $\tilde{R}_1 = \max\{z \in \text{Dom}^\rho : \Theta_1(0, 0, z) \leq C, z \in [\lfloor \lambda_{\min}/((2 + \gamma)V_{\max}) \rfloor, \lceil R_{\max} \rceil]\}$. The product offer decisions from the algorithm are $(\tilde{x}_1, \dots, \tilde{x}_n)$. Our goal is to show that these product offer decisions are feasible to problem (1) and they provide a performance guarantee. In the next lemma, we show that the product offer decisions from the APPRX(ρ) algorithm are feasible to problem (1).

Lemma 6 *Assume that $(\tilde{x}_1, \dots, \tilde{x}_n)$ are generated by the APPRX(ρ) algorithm. Then, we have $\sum_{j=1}^n c_j \tilde{x}_j \leq C$.*

Proof. We have $\tilde{R}_1 = \max\{z \in \text{Dom}^\rho : \Theta_1(0, 0, z) \leq C, z \in [\lfloor \lambda_{\min}/((2 + \gamma)V_{\max}) \rfloor, \lceil R_{\max} \rceil]\}$ in Step 0 of the APPRX(ρ) algorithm, which implies that $\Theta_1(0, 0, \tilde{R}_1) \leq C$. In this case, since $\tilde{P}_1 = 0$ and $\tilde{S}_1 = 0$ in the APPRX(ρ) algorithm, we obtain $\Theta_1(\tilde{P}_1, \tilde{S}_1, \tilde{R}_1) \leq C$. Noting the definitions of \tilde{P}_{j+1} , \tilde{S}_{j+1} and \tilde{R}_{j+1} in Step 1 of the APPRX(ρ) algorithm and using the fact that \tilde{x}_j is the optimal solution to problem (9), if we compare problems (4) and (9), then it follows that $\Theta_j(\tilde{P}_j, \tilde{S}_j, \tilde{R}_j) = c_j \tilde{x}_j + \Theta_{j+1}(\tilde{P}_{j+1}, \tilde{S}_{j+1}, \tilde{R}_{j+1})$ for all $j = 1, \dots, n$. Adding these equalities over all $j = 1, \dots, n$, we obtain $\Theta_1(\tilde{R}_1, \tilde{S}_1, \tilde{R}_1) = \sum_{j=1}^n c_j \tilde{x}_j + \Theta_{n+1}(\tilde{P}_{n+1}, \tilde{S}_{n+1}, \tilde{R}_{n+1})$. Thus, noting that $\Theta_1(\tilde{P}_1, \tilde{S}_1, \tilde{R}_1) \leq C$, we have $\sum_{j=1}^n c_j \tilde{x}_j + \Theta_{n+1}(\tilde{P}_{n+1}, \tilde{S}_{n+1}, \tilde{R}_{n+1}) \leq C$. By the boundary condition of the dynamic program in (4), $\Theta_{n+1}(\tilde{P}_{n+1}, \tilde{S}_{n+1}, \tilde{R}_{n+1})$ is either zero or infinity. If $\Theta_{n+1}(\tilde{P}_{n+1}, \tilde{S}_{n+1}, \tilde{R}_{n+1}) = \infty$, then we get a contradiction to the inequality $\sum_{j=1}^n c_j \tilde{x}_j + \Theta_{n+1}(\tilde{P}_{n+1}, \tilde{S}_{n+1}, \tilde{R}_{n+1}) \leq C$. Therefore, we must have $\Theta_{n+1}(\tilde{P}_{n+1}, \tilde{S}_{n+1}, \tilde{R}_{n+1}) = 0$. In this case, since we have $\sum_{j=1}^n c_j \tilde{x}_j + \Theta_{n+1}(\tilde{P}_{n+1}, \tilde{S}_{n+1}, \tilde{R}_{n+1}) \leq C$, we get $\sum_{j=1}^n c_j \tilde{x}_j \leq C$. \square

In the next lemma, we show that the product offer decisions from the APPRX(ρ) algorithm provide a performance guarantee for problem (1). A large part of the proof of this lemma shows that the expected revenue from the product offer decisions generated by the APPRX(ρ) algorithm is no worse than the optimal expected revenue estimate $\tilde{R}_1 = \max\{z \in \text{Dom}^\rho : \Theta_1(0, 0, z) \leq C, z \in [\lfloor \lambda_{\min}/((2 + \gamma)V_{\max}) \rfloor, \lceil R_{\max} \rceil]\}$.

Lemma 7 *Let REV be the expected revenue from the product offer decisions generated by the APPRX(ρ) algorithm and z^* be the optimal objective value of problem (1). Then, we have $\text{REV} \geq z^*/(1 + \rho)^{3n}$.*

Proof. Assume that $(\tilde{P}_1, \dots, \tilde{P}_{n+1})$, $(\tilde{S}_1, \dots, \tilde{S}_{n+1})$ and $(\tilde{R}_1, \dots, \tilde{R}_{n+1})$ are generated by the APPRX(ρ) algorithm. Noting the definition of \tilde{R}_1 in Step 0 of APPRX(ρ) algorithm, by Corollary 5, we have $\tilde{R}_1 \geq z^*/(1+\rho)^{3n}$. We let $(\tilde{x}_1, \dots, \tilde{x}_n)$ be the product offer decisions generated by the APPRX(ρ) algorithm. By the expected revenue expression in (1), we have $\text{REV} = \sum_{j=1}^n \lambda_j \frac{\sum_{k=1}^j r_k v_k \tilde{x}_k}{v_0 + \sum_{k=1}^j v_k \tilde{x}_k}$. In the rest of the proof, we show that $\text{REV} \geq \tilde{R}_1$, in which case, the desired result follows by the fact that $\text{REV} \geq \tilde{R}_1 \geq z^*/(1+\rho)^{3n}$. Since $\tilde{P}_{k+1} = \lceil \tilde{P}_k + r_k v_k \tilde{x}_k \rceil$ in Step 1 of the APPRX(ρ) algorithm, we have $\tilde{P}_{k+1} \leq \tilde{P}_k + r_k v_k \tilde{x}_k$. Adding this inequality over all $k = 1, \dots, j-1$ and noting that $\tilde{P}_1 = 0$, we get $\tilde{P}_j \leq \sum_{k=1}^{j-1} r_k v_k \tilde{x}_k$. A similar argument yields $\tilde{S}_j \geq \sum_{k=1}^{j-1} v_k \tilde{x}_k$ as well. Since $\tilde{R}_{j+1} = \lceil \tilde{R}_j - \lambda_j \frac{\tilde{P}_j + r_j v_j \tilde{x}_j}{v_0 + \tilde{S}_j + v_j \tilde{x}_j} \rceil$ in Step 1 of the APPRX(ρ) algorithm, we get

$$\tilde{R}_{j+1} \geq \tilde{R}_j - \lambda_j \frac{\tilde{P}_j + r_j v_j \tilde{x}_j}{v_0 + \tilde{S}_j + v_j \tilde{x}_j} \geq \tilde{R}_j - \lambda_j \frac{\sum_{k=1}^j r_k v_k \tilde{x}_k}{v_0 + \sum_{k=1}^j v_k \tilde{x}_k},$$

where the second inequality is by the fact that $\tilde{P}_j \leq \sum_{k=1}^{j-1} r_k v_k \tilde{x}_k$ and $\tilde{S}_j \geq \sum_{k=1}^{j-1} v_k \tilde{x}_k$ for all $j = 1, \dots, n$. Adding the chain of inequalities above over all $j = 1, \dots, n$, we get $\tilde{R}_{n+1} \geq \tilde{R}_1 - \sum_{j=1}^n \lambda_j \frac{\sum_{k=1}^j r_k v_k \tilde{x}_k}{v_0 + \sum_{k=1}^j v_k \tilde{x}_k} = \tilde{R}_1 - \text{REV}$. By using the same argument at the end of the proof of Lemma 6, we have $\Theta_{n+1}(\tilde{P}_{n+1}, \tilde{S}_{n+1}, \tilde{R}_{n+1}) = 0$, but $\Theta_{n+1}(\tilde{P}_{n+1}, \tilde{S}_{n+1}, \tilde{R}_{n+1}) = 0$ if and only if $\tilde{R}_{n+1} \leq 0$. Thus, we must have $\tilde{R}_{n+1} \leq 0$. In this case, we obtain $0 \geq \tilde{R}_{n+1} \geq \tilde{R}_1 - \text{REV}$. \square

Next, we focus on the number of operations required to run the APPRX(ρ) algorithm. The number of operations required to compute the value functions $\{\Theta_j(\cdot, \cdot, \cdot) : j = 1, \dots, n\}$ in Step 0 of the APPRX(ρ) algorithm dominates the number of operations required to run the other steps. In Step 0 of the APPRX(ρ) algorithm, we need to compute $\Theta_j(P_j, S_j, R_j)$ for all $j = 1, \dots, n$ and $(P_j, S_j, R_j) \in \text{Dom}^\rho \times \text{Dom}^\rho \times \text{Dom}^\rho$ such that $P_j \in \{0\} \cup [1, \lceil nR_{\max} \rceil]$, $S_j \in \{0\} \cup [1, \lceil \gamma V_{\max} \rceil]$ and $R_j \in \{0\} \cup [\lfloor \lambda_{\min}/((2+\gamma)V_{\max}) \rfloor, \lceil R_{\max} \rceil]$. By the definition of Dom^ρ , there are

$$O\left(\frac{\log(nR_{\max})}{\log(1+\rho)} \times \frac{\log(\gamma V_{\max})}{\log(1+\rho)} \times \frac{\log\left(\frac{(1+\gamma)R_{\max}V_{\max}}{\lambda_{\min}}\right)}{\log(1+\rho)}\right) = O\left(\frac{\log(nR_{\max}) \log(\gamma V_{\max}) \log\left(\frac{(1+\gamma)R_{\max}V_{\max}}{\lambda_{\min}}\right)}{\rho^3}\right)$$

possible values of $(P_j, S_j, R_j) \in \text{Dom}^\rho \times \text{Dom}^\rho \times \text{Dom}^\rho$ such that $P_j \in \{0\} \cup [1, \lceil nR_{\max} \rceil]$, $S_j \in \{0\} \cup [1, \lceil \gamma V_{\max} \rceil]$ and $R_j \in \{0\} \cup [\lfloor \lambda_{\min}/((2+\gamma)V_{\max}) \rfloor, \lceil R_{\max} \rceil]$. Thus, we need to compute $\Theta_j(P_j, S_j, R_j)$ for $O(\log(nR_{\max}) \log(\gamma V_{\max}) \log((1+\gamma)R_{\max}V_{\max}/\lambda_{\min})/\rho^3)$ different values of (P_j, S_j, R_j) . We compute the value functions $\{\Theta_j(\cdot, \cdot, \cdot) : j = 1, \dots, n\}$ by moving over the decision epochs in reverse order in the dynamic program in (4) and computing $\Theta_j(P_j, S_j, R_j)$ for a certain value of (P_j, S_j, R_j) takes $O(1)$ operations. So, since there are n decision epochs, we can run

Step 0 of the APPRX(ρ) algorithm in $O(n \log(nR_{\max}) \log(\gamma V_{\max}) \log((1 + \gamma) R_{\max} V_{\max}/\lambda_{\min})/\rho^3)$ operations. In the next theorem, we use this observation to give an FPTAS for problem (1).

Theorem 8 *Letting z^* be the optimal objective value of problem (1), there exists an algorithm such that for any $\epsilon \in (0, 1)$, the algorithm runs in $O(n^4 \log(nR_{\max}) \log(nV_{\max}) \log(nR_{\max} V_{\max}/\lambda_{\min})/\epsilon^3)$ operations and generates product offer decisions that are feasible to problem (1), providing an expected revenue of at least $(1 - \epsilon) z^*$.*

Proof. We run the APPRX(ρ) algorithm with $\rho = \epsilon/(6n)$ and let REV be the expected revenue from the product offer decisions generated by this algorithm. By Lemma 6, the product offer decisions are feasible to problem (1). By Lemma 7, we get $z^* \leq (1 + \epsilon/(6n))^{3n} \text{REV}$. Since $(1 + x/n)^n \leq \exp(x) \leq (1 + 2x)$ for $x \in (0, 1/2)$, we get $(1 + \epsilon/(6n))^{3n} \leq \exp(\epsilon/2) \leq 1 + \epsilon$ for $\epsilon \in (0, 1)$. In this case, we obtain $\text{REV} \geq z^*/(1 + \epsilon/(6n))^{3n} \geq z^*/(1 + \epsilon) \geq (1 - \epsilon) z^*$. So, the expected revenue from the product offer decisions is at least $(1 - \epsilon) z^*$. If we set $\rho = \epsilon/(6n)$, then we have $n\rho \leq 1/6$ so that $\gamma = (\exp(n\rho) - 1)/\rho \leq (1 + 2n\rho - 1)/\rho = 2n$ by the inequalities at the beginning of the proof. By the discussion right before the theorem, we can run the APPRX(ρ) algorithm in $O(n \log(nR_{\max}) \log(\gamma V_{\max}) \log((1 + \gamma) R_{\max} V_{\max}/\lambda_{\min})/\rho^3)$ operations. Replacing ρ with $\epsilon/6n$ and γ with the upper bound of $2n$ in the last expression yields the desired running time. \square

If we do not have a limit on the total space consumption of the offered products, then we can solve the dynamic program in (4) more efficiently, in which case, the running time of our FPTAS reduces to $O(n^3 \log(nR_{\max}) \log(nV_{\max})/\epsilon^2)$ operations. In Appendix D, we analyze the running time of our FPTAS for the uncapacitated version of the assortment problem.

6 Numerical Study

In our numerical study, we test the performance of our FPTAS on a large number of problem instances. We generate each problem instance as follows. We generate the revenue r_j of product j from the uniform distribution over $[1, \bar{R}]$, where \bar{R} is a parameter that we vary. We reindex (r_1, \dots, r_n) such that $r_1 \leq \dots \leq r_n$. Thus, noting that the consideration set of a customer of type k is $\{1, \dots, k\}$, a customer of type k is interested in the k products with the lowest prices. To come up with the arrival probability λ_k of customer type k , we generate β_j from the uniform distribution over $[0, 1]$ for all $j = 1, \dots, n$ and set $\lambda_k = \beta_k / \sum_{j=1}^n \beta_j$. The preference weight of the no purchase

option is $v_0 = 10$. We use two approaches to generate the preference weights of the products. In the first approach, we simply generate the preference weight v_j of product j from the uniform distribution over $[1, 50]$. In the second approach, we generate the preference weight v_j of product j from the uniform distribution over $[1, 50]$ and reindex (v_1, \dots, v_n) such that $v_1 \geq \dots \geq v_n$. In the first approach, referred to as NR, the preference weight of a product does not have any relationship with the price of the product. In the second approach, referred to as PD, the products with higher prices have lower preference weights. In all of our problem instances, the number of products is $n = 36$ and there is no limit on the total space consumption of the offered products.

Using $\text{PR} \in \{\text{NR}, \text{PD}\}$ to denote the approach used to generate the preference weights, we vary (\bar{R}, PR) over $\{5, 25, 50\} \times \{\text{NR}, \text{PD}\}$, yielding six combinations. In each combination, we generate 100 problem instances by using the approach above. For each problem instance, we run the $\text{APPRX}(\rho)$ algorithm to get a solution providing at least $1/2$, $3/4$ or $9/10$ of the optimal expected revenue. By the proof of Theorem 8, we can set $\rho = \epsilon/(6n)$ to get a solution providing at least a $1 - \epsilon$ fraction of the optimal expected revenue. In Appendix E, we also give a dynamic program to get an upper bound on the optimal expected revenue. We compare the expected revenues from the solutions obtained by the $\text{APPRX}(\rho)$ algorithm with the optimal expected revenue upper bounds.

We give our numerical results in Table 1. On the left side of the table, we show the combination (\bar{R}, PR) . In the rest of the table, there are three blocks of five columns. The three blocks focus on the cases where we use the performance guarantees of $1/2$, $3/4$ and $9/10$. The first, second, third and fourth columns in each block respectively show the average, maximum, 90th percentile and 75th percentile of the percent gaps between the upper bound on the optimal expected revenue and the expected revenue from the solution obtained by the $\text{APPRX}(\rho)$ algorithm, where these summary statistics are computed over the 100 problem instances in a particular combination of (\bar{R}, PR) . So, for problem instance ℓ , letting UPBN^ℓ be the upper bound on the optimal expected revenue and REV^ℓ be the expected revenue from the solution obtained by the $\text{APPRX}(\rho)$ algorithm, the first, second, third and fourth columns respectively show the average, maximum, 90th percentile and 75th percentile of the data $\{100 \times \frac{\text{UPBN}^\ell - \text{REV}^\ell}{\text{UPBN}^\ell} : \ell = 1, \dots, 100\}$. The fifth column shows the average CPU seconds for the $\text{APPRX}(\rho)$ algorithm over the 100 problem instances.

The results in Table 1 indicate that the practical performance of the $\text{APPRX}(\rho)$ algorithm can be quite strong. The maximum optimality gap is 3.89% when we use a performance guarantee of $1/2$ and 2.16% when we use a performance guarantee of $9/10$. The average CPU seconds is 5.28

| Combin. (\bar{R} , PR) | Perf. Guar. of $1 - \epsilon = 1/2$ | | | | | Perf. Guar. of $1 - \epsilon = 3/4$ | | | | | Perf. Guar. of $1 - \epsilon = 9/10$ | | | | |
|------------------------------|-------------------------------------|------|------|------|-------------|-------------------------------------|------|------|------|-------------|--------------------------------------|------|------|------|-------------|
| | % Gap with Upp. Bnd | | 90th | 75th | CPU Sec. | % Gap with Upp. Bnd | | 90th | 75th | CPU Sec. | % Gap with Upp. Bnd | | 90th | 75th | CPU Sec. |
| | Avg. | Max. | | | | Avg. | Max. | | | | Avg. | Max. | | | |
| (5, NR) | 1.86 | 2.91 | 2.42 | 2.13 | 4.47 | 1.40 | 1.86 | 1.65 | 1.55 | 19.91 | 1.25 | 1.74 | 1.45 | 1.37 | 152.22 |
| (5, PD) | 2.14 | 3.89 | 3.16 | 2.61 | 5.01 | 1.52 | 2.83 | 1.89 | 1.70 | 22.27 | 1.32 | 1.84 | 1.55 | 1.47 | 170.50 |
| (25, NR) | 1.55 | 2.82 | 2.13 | 1.80 | 5.11 | 1.12 | 1.78 | 1.42 | 1.27 | 23.12 | 1.01 | 1.50 | 1.24 | 1.11 | 179.53 |
| (25, PD) | 1.65 | 3.23 | 2.43 | 2.04 | 5.68 | 1.22 | 2.10 | 1.55 | 1.43 | 25.39 | 1.14 | 1.63 | 1.47 | 1.26 | 196.59 |
| (50, NR) | 1.50 | 2.65 | 2.18 | 1.72 | 5.44 | 1.08 | 1.71 | 1.37 | 1.26 | 24.78 | 0.97 | 1.42 | 1.20 | 1.06 | 198.34 |
| (50, PD) | 1.56 | 3.60 | 2.29 | 1.82 | 5.99 | 1.15 | 1.84 | 1.53 | 1.34 | 27.21 | 1.08 | 2.16 | 1.42 | 1.18 | 217.57 |

Table 1: Performance of the APPRX(ρ) algorithm with different performance guarantees.

with a performance guarantee of $1/2$ and 185.79 with a performance guarantee of $9/10$. We carried out our numerical study in Java 1.7.0 with 2.8 GHz Intel Core i7 CPU and 16 GB RAM. We report the average CPU seconds, but the CPU seconds vary among the problem instances in a particular (\bar{R}, PR) combination by no more than 30%. If have 72 products, then the average CPU seconds is 42.80 with a performance guarantee of $1/2$ and 1575.08 with a performance guarantee $9/10$, but we do not report detailed results for 72 products. As expected from the discussion after Theorem 8, doubling the number of products increases the CPU seconds by about a factor of eight.

Meissner et al. (2012) give a linear programming formulation for assortment problems under a mixture of multinomial logit models. This formulation relaxes the assortment problem by assuming that we can offer different subsets of products to different customer types, but uses the so called product cuts to tighten the relaxation. This formulation also provides an upper bound on the optimal expected revenue. In Appendix F, we give a detailed comparison of our approach with the linear programming formulation of Meissner et al. (2012). Our numerical results indicate that our approach compares quite favorably with the linear programming formulation. To give a feel for our numerical results, the largest percent gap between our upper bound on the optimal expected revenue and the expected revenue from the solution obtained by our approach is 2.36%. The same percent gap for the linear programming formulation can be as high as 19.07%.

7 Conclusions and Future Research

One research direction is to work with other consideration set structures. In our FPTAS, customers of different types associate the same preference weight with a particular product, as long as this product is in their consideration sets. Also, they have nested consideration sets. In Appendix G, we show that if customers of different types associate the same preference weight with a particular product, but they have arbitrary consideration sets, then the assortment problem is NP-hard to

approximate within a factor of $O(n^{1-\epsilon})$ for any $\epsilon > 0$. To show this result, we build on Desir and Goyal (2014) and Aouad, Farias, Levi and Segev (2015), but these papers do not immediately imply our result. Desir and Goyal (2014) use a mixture of multinomial logit models, but customers of two different types may associate different preference weights with a particular product when this product is in the consideration sets of both customer types. Aouad, Farias, Levi and Segev (2015) use a choice model where customers of a certain type rank the products according to a certain preference order and purchase their most preferred available product. The complexity result in Appendix G indicates that the case where customers of different types associate the same preference weight with a particular product but they have arbitrary consideration sets is unlikely to provide adequate structure to give an FPTAS for the corresponding assortment problem.

To impose some structure on the consideration sets, we can work with the so called interval consideration sets of the form $\{k, \dots, j\}$ with $k \leq j$. Aouad, Levi and Segev (2015) study dynamic assortment problems where multiple customers are served with limited inventories. In their choice model, each customer of a particular type ranks the products according to a preference order of the form (k, \dots, j) with $k \leq j$, where product i is preferred over product $i + 1$. She purchases her most preferred available product. Letting n be the number of products, the authors group the customer types into $O(\log n)$ groups. The dynamic assortment problem that focuses on each group ends up being tractable to approximate. Picking the best of the solutions from the dynamic assortment problems for different groups, the authors get a solution whose expected revenue deviates from the optimal by at most a factor of $O(\log n)$. In Appendix H, we show that if we have interval consideration sets, then we can use the idea in Aouad, Levi and Segev (2015) to group the customer types into $O(\log n)$ groups, in which case, we can approximate the assortment problem that focuses on each group by using nested consideration sets. So, we can use our FPTAS to approximate the assortment problem for each group. Running our FPTAS $O(n)$ times, we get a solution whose expected revenue deviates from the optimal by at most a factor of $O(\log n)$. Furthermore, letting p be the ratio between the largest and smallest product revenues, under the assumption that the consideration sets are intervals of the form $\{k, \dots, j\}$ and the products are indexed in the order of increasing revenues, Aouad, Levi and Segev (2015) also get a solution for the dynamic assortment problem whose expected revenue deviates from the optimal by at most a factor of $O(\log \log p)$. If we have interval consideration sets and the products are indexed in the order of increasing revenues, then we can use a similar approach to obtain a solution for our assortment problem whose expected

revenue deviates from the optimal by at most a factor of $O(\log \log p)$. Nevertheless, these approaches do not provide an expected revenue arbitrarily close to the optimal. Therefore, giving an FPTAS for the assortment problem under interval consideration sets is one research direction.

Lastly, the running time of our FPTAS depends on R_{\max} , V_{\max} and λ_{\min} . An interesting question is whether one can come up with a strongly polynomial time algorithm, whose running time does not depend on these quantities.

ACKNOWLEDGEMENTS. We thank the associate editor and two anonymous referees whose comments improved our paper substantially and unified our results.

AUTHOR BIOS. **Jacob B. Feldman** is an assistant professor at Olin Business School at Washington University in St. Louis. His research interests include customer choice modeling and approximation algorithms. **Huseyin Topaloglu** is a professor at the School of Operations Research and Information Engineering at Cornell Tech. His research interests include revenue management and pricing.

References

- Aouad, A., Farias, V. and Levi, R. (2016), Assortment optimization under consider-then-choose choice models, Technical report, MIT, Massachusetts, MA.
- Aouad, A., Farias, V., Levi, R. and Segev, D. (2015), The approximability of assortment optimization under ranking preferences, Technical report, MIT, Cambridge, MA.
- Aouad, A., Levi, R. and Segev, D. (2015), Approximation algorithms for dynamic assortment optimization models, Technical report, MIT, Massachusetts, MA.
- Bront, J. J. M., Diaz, I. M. and Vulcano, G. (2009), ‘A column generation algorithm for choice-based network revenue management’, *Oper. Res.* **57**(3), 769–784.
- Chong, J.-K., Ho, T.-H. and Tang, C. S. (2001), ‘A modeling framework for category assortment planning’, *M&SOM* **3**(3), 191–210.
- Dai, J., Ding, W., Kleywegt, A. J., Wang, X. and Zhang, Y. (2014), Choice based revenue management for parallel flights, Technical report, Georgia Tech, Atlanta, GA.
- Desir, A. and Goyal, V. (2014), Near-optimal algorithms for capacity constrained assortment optimization, Technical report, Columbia University, New York, NY.
- Gallego, G., Iyengar, G., Phillips, R. and Dubey, A. (2004), Managing flexible products on a network, Technical report, Columbia University, New York, NY.
- Gallego, G., Li, A., Truong, V.-A. and Wang, X. (2016), Online personalized resource allocation with customer choice, Technical report, Columbia University, New York, NY.
- Garey, M. and Johnson, D. (1979), *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman and Company, New York, NY.
- Goyal, V., Levi, R. and Segev, D. (2016), ‘Near-optimal algorithms for the assortment planning problem under dynamic substitution and stochastic demand’, *Oper. Res.* **64**(1), 219–235.
- Jagabathula, S. and Rusmevichientong, P. (2015), A nonparametric joint assortment and price choice model, Technical report, New York University, New York, NY.
- Jagabathula, S. and Vulcano, G. (2015), A model to estimate individual preferences using panel data, Technical report, New York University, New York, NY.
- Kok, A. G. and Fisher, M. L. (2007), ‘Demand estimation and assortment optimization under substitution: Methodology and application’, *Oper. Res.* **55**(6), 1001–10021.
- Kunnumkal, S. and Talluri, K. (2014), On the tractability of the piecewise-linear approximation

- for general discrete-choice network revenue management, Technical report, Universitat Pompeu Fabra, Barcelona, Spain.
- McFadden, D. and Train, K. (2000), ‘Mixed MNL models for discrete response’, *Journal of Applied Economics* **15**, 447–470.
- Meissner, J., Strauss, A. and Talluri, K. (2012), ‘An enhanced concave program relaxation for choice network revenue management’, *Production and Operations Management* **22**(1), 71–87.
- Misra, K. (2008), Understanding retail assortments in competitive markets, Technical report, Northwestern University, Evanston, IL.
- Rusmevichientong, P., Shmoys, D. B., Tong, C. and Topaloglu, H. (2014), ‘Assortment optimization under the multinomial logit model with random choice parameters’, *POM* **23**(11), 2023–2039.
- Sahin, O. and Wang, R. (2014), The impact of consumer search cost on assortment planning and pricing, Technical report, Johns Hopkins University, Baltimore, MD.
- Stauss, A. K. and Talluri, K. (2016), Tractable consideration set structures for network revenue management, Technical report, Imperial College, London, UK.
- Talluri, K. (2011), A randomized concave programming method for choice network revenue management, Technical report, Universitat Pompeu Fabra, Barcelona, Spain.
- Talluri, K. and van Ryzin, G. (2004), ‘Revenue management under a general discrete choice model of consumer behavior’, *Management Sci.* **50**(1), 15–33.
- Vulcano, G., van Ryzin, G. and Chahr, W. (2010), ‘OM practice – Choice-based revenue management: An empirical study of estimation and optimization’, *ME&SOM* **12**(3), 371–392.