

Robust Assortment Optimization under the Markov Chain Choice Model

Antoine Désir

Technology and Operations Management Area, INSEAD, antoine.desir@insead.edu

Vineet Goyal

Department of Industrial Engineering and Operations Research, Columbia University, vgoyal@ieor.columbia.edu

Bo Jiang

Research Institute for Interdisciplinary Sciences, School of Information Management and Engineering, Shanghai University of Finance and Economics, isyebojiang@gmail.com

Tian Xie

Research Institute for Interdisciplinary Sciences, School of Information Management and Engineering, Shanghai University of Finance and Economics, xietiansh@gmail.com

Jiawei Zhang

Stern School of Business, New York University, jzhang@stern.nyu.edu

Assortment optimization arises widely in many practical applications such as retailing and online advertising. In this problem, the goal is to select a subset from a universe of substitutable products to offer customers in order to maximize the expected revenue. We study a robust assortment optimization problem under the Markov chain choice model. In this formulation, the parameters of the choice model are assumed to be uncertain and the goal is to maximize the worst-case expected revenue over all parameter values in an uncertainty set. Our main contribution is to prove a min-max duality result when the uncertainty set is row-wise. The result is surprising as the objective function does not satisfy the properties usually needed for known min-max results. Inspired by the duality result, we develop an efficient iterative algorithm for computing the optimal robust assortment under the Markov chain choice model. Moreover, our results yield operational insights into the effect of changing the uncertainty set on the optimal robust assortment. In particular, consistent with previous literature, we find that bigger uncertainty sets always lead to bigger assortments, and a firm should offer larger assortments to hedge against uncertainty.

1. Introduction

Assortment optimization arises widely in many practical applications such as retailing and online advertising. In this problem, the goal is to select a subset of products to offer customers from a universe of substitutable products in order to maximize the expected revenue. The demand of any product depends on the substitution behavior of the customers, which is captured by a choice model. The latter specifies the probability that a random consumer selects a particular product from any given offer set. The objective of the decision maker is to identify an offer set that maximizes her expected revenue. Many parametric choice models have been extensively studied in the literature

in diverse areas including marketing, transportation, economics, and operations management. The multinomial logit (MNL) model is by far the most popular model in practice due to its tractability (McFadden 1978, Talluri and van Ryzin 2004). However, the simplicity of the MNL model comes with commonly recognized limitations such as the independence of irrelevant alternatives (IIA) property (see Ben-Akiva et al. 1985). Informally, the IIA property states that the odds of choosing between two products are not affected by the presence of a third product. This makes the MNL model inadequate for many applications. To alleviate these limitations, Blanchet et al. (2016) propose a Markov chain based-choice model. In this model, customer substitution is captured by a Markov chain, where each product, including a dummy product representing a no-purchase option, corresponds to a state, and substitutions are modeled as transitions in the Markov chain. The authors show that the Markov chain choice model provides good choice probability estimates when the data arises from a wide class of existing choice models. We would like to note that a similar idea had already been used by Zhang and Cooper (2005) in the context of airline revenue management. This alternative way to model customer behavior has recently received a lot of attention. Feldman and Topaloglu (2017) study the network revenue management problem under the Markov chain choice model. Şimşek and Topaloglu (2018) propose a method to estimate the parameters of the model from data. Désir et al. (2020) show that the constrained assortment problem is APX-hard and design efficient approximation algorithms to tackle the constrained variant of the assortment optimization problem. This paper adds to this stream of work by considering a robust variant of the assortment optimization problem under the Markov chain choice model.

In practice, the parameters of the underlying choice model have to be estimated from data. Statistical errors in these parameters are therefore unavoidable. When the decision maker ignores estimation errors, “optimal” decisions based on the point estimators could potentially be suboptimal for the true parameters, especially when the estimates differ from the true parameters. To account for this, we instead propose a robust optimization approach where the uncertainty in the parameter values is explicitly captured by a “confidence set” or an “uncertainty set.” Intuitively, this set includes the true parameters with high confidence based on the statistical estimation procedure. Given such an uncertainty set, the goal of the robust assortment problem is to choose an assortment that maximizes the worst-case expected revenue, where the worst case is taken over all possible parameters values.

1.1. Our contributions

Our main contributions in this work are the following.

Min-max duality. The robust formulation of the assortment optimization problem is a max-min problem. In the outer maximization problem the decision maker chooses the best possible

assortment while in the inner minimization problem an adversary picks the worst model parameters from an uncertainty set in order to minimize the expected revenue of the selected assortment. We prove that the order of the max and min operators in this max-min formulation can be interchanged when the uncertainty set is row-wise, i.e., when the uncertainty sets across the different rows of the transition matrix are unrelated. Under this condition, the optimal objective value stays the same even when the adversary makes the model parameter selection first and then the decision maker chooses the assortment that maximizes the expected revenue for those parameters. This duality result does not follow directly from standard min-max duality results in the literature, which often require assumptions such as quasi-convexity/quasi-concavity of the objective function that are not satisfied by our objective function.

The Markov chain choice model is a strict generalization of the MNL model (see Blanchet et al. 2016). In particular, the assortment optimization under the MNL model is a special case of the assortment optimization under the Markov chain choice model. However, the robust assortment optimization under the MNL model cannot be written as a special case of the robust assortment optimization under the Markov chain choice model with the row-wise uncertainty set. Nevertheless, we introduce a more general formulation that allows us to establish, under certain reasonable assumptions, min-max duality for the robust assortment optimization both under the Markov chain choice model and under the MNL model. Under the MNL model, our result implies a min-max relation without any assumptions on the uncertainty set. This result was previously known in the literature (Rusmevichientong and Topaloglu 2012) and we provide an alternative proof.

Optimal algorithms for robust MC and MNL and insights. We also present efficient algorithms for computing the optimal robust assortment under both Markov chain and MNL choice models. Our algorithms are iterative procedures solving a fixed point equation inspired by the min-max relation. We also present operational insights into the effect of changing the uncertainty set on the optimal robust assortment under the Markov chain choice model. In particular, we find that bigger uncertainty sets always lead to bigger assortments, and a firm should offer larger assortments to hedge against uncertainty.

Numerical experiments. We also present two numerical studies to illustrate the tradeoffs associated with taking a robust approach to assortment optimization under the Markov chain choice model. In particular, we quantify the tradeoff between the capability of hedging against the uncertainty in the model parameters and the conservativeness of the associated optimal robust assortment. In the first study, we solve the robust assortment optimization under the Markov chain choice model and show that the running time of our algorithm nicely scales with the size of the problem instance. Moreover, we illustrate the magnitude of the tradeoff between the expected and worst-case revenues under the robust approach compared to a deterministic approach that does not

account for any uncertainty in the parameters. In the second study, we present a more realistic set of experiments where we assume that we have access only to purchase data generated by an unknown ranking-based model. In this case, we need to learn the parameters of a Markov chain choice model and then construct uncertainty sets from the data. We again compare the performance of the robust approach to that of the deterministic approach that ignores any uncertainty in the parameters. Interestingly, we show in this case that accounting for some uncertainty in the parameters through under the robust approach not only improves the worst-case performance but can also sometimes lead to a better average performance.

1.2. Related literature

Our work is closely related to the choice model and assortment optimization literature as well as the robust optimization literature.

Choice model and assortment optimization. In order to overcome the MNL model limitations and capture a richer class of substitution behaviors, many choice models have been proposed. Most of them increase the model complexity and therefore make the parameters' estimation and assortment optimization significantly more difficult. One of the key challenges in assortment planning is choosing a model that strikes a good balance between its predictability and tractability. The interested reader is referred to Kök et al. (2015) for a comprehensive survey of the literature on the assortment optimization problems.

The MNL model belongs to the very rich class of random utility models where the utility of each product is modeled as the sum of a deterministic component and a random component. The assumption on the joint distribution of the random components specifies the choice model. For instance, when the random components are i.i.d. and follow a Gumbel distribution, this results in the MNL model. In this class of random utility models, generalizations of the MNL model include the nested logit model (Williams 1977) and a mixture of MNL models (McFadden and Train 2000). Under the nested logit model, the products are grouped into nests and products in the same nest have positively correlated random components, which implies that they are closer substitutes. Under the mixture of MNL models, customer heterogeneity is added by considering multiple segments each of which follows a different MNL model. The assortment optimization problem has been studied under both of these extensions. Under the mixture of MNL models, the assortment optimization problem becomes NP-hard, even when the number of models is two (Rusmevichientong et al. 2014). Under the nested logit model, the assortment optimization problem is tractable (Davis et al. 2014), even when additional capacity constraints are added (Gallego and Topaloglu 2014, Désir et al. 2014). One main limitation of the nested logit model is that it depends on a predefined nest structure that is hard to estimate in practice.

Another approach to choice modeling is to represent customer preferences by a distribution over preference lists, i.e., strict orderings of the products. Each customer draws up a preference list and selects among the offered products the highest ranked option. This approach leads to expressive models that can capture very complex substitution behaviors. Farias et al. (2013) show that this approach can lead to more accurate revenue predictions than traditional random utility models. However, the assortment optimization problem becomes intractable under such a general model. In particular, even when the support of the distribution is sparse, there is no polynomial algorithm for the assortment optimization problem with an approximation factor better than $\Omega(1/n^{1-\epsilon})$ for any constant $\epsilon > 0$ unless $P = NP$ (Aouad et al. 2018). Several special cases of this model lead to more tractable optimization problems (Honhon et al. 2010, 2012, Aouad et al. 2021, Jagabathula and Rusmevichientong 2016).

In this paper, we focus on the Markov chain based-choice model, which has recently attracted growing interest (Blanchet et al. 2016, Şimşek and Topaloglu 2018, Feldman and Topaloglu 2017, Désir et al. 2020). The most related work is that of Rusmevichientong and Topaloglu (2012) who study the robust assortment optimization under the MNL model. Since the Markov chain choice model is a strict generalization of the MNL model (see Blanchet et al. 2016), our results strictly generalize Rusmevichientong and Topaloglu (2012). Importantly, the approach taken in Rusmevichientong and Topaloglu (2012) does not extend to the Markov chain choice model and we have to develop new tools to solve the robust assortment optimization problem under the Markov chain choice model.

Robust optimization. Finally, our paper relates to the robust optimization literature (Ben-Tal et al. 2009, Ben-Tal and Nemirovski 2000, Bertsimas and Sim 2003, Gorissen et al. 2014, Xu and Burer 2018) which incorporates the uncertainty in the model parameters into the decision-making process. This paradigm has found application in the literature on operations management and revenue management including airline revenue management (Birbil et al. 2009, Perakis and Roels 2010), pricing (Thiele 2009), portfolio selection (Chen et al. 2011, Zhu and Fukushima 2009), process flexibility (Wang and Zhang 2015), and appointment scheduling (Mak et al. 2014). Gorissen et al. (2014) study a robust linear conic program with column-wise uncertainty in the transpose of the coefficient matrix, and they show that this problem is computationally tractable. Even though the setup is related, their results does not apply in our setting. We give more details in Section 2.

1.3. Outline

The remainder of this paper is organized as follows. In Section 2, we introduce the model and the main min-max result. In Section 3, we introduce a more general formulation that we use to

prove our main result. Section 4 is dedicated to the proof of the min-max result under the general formulation. In Section 5, we discuss some implications of our main result. In particular, we give an efficient algorithm for computing the optimal robust assortment and provide some operational insights. Finally, we conduct some numerical experiments in Section 6 to showcase the benefits of adopting a robust approach.

2. Model and main result

In this section, we first introduce the Markov chain choice model and formulate the corresponding robust assortment optimization problem. We then present the min-max duality of the robust formulation.

2.1. Markov chain choice model under uncertainty

Model parameters. We consider a universe of n products denoted by $\mathcal{N} = \{1, 2, \dots, n\}$. We let product 0 be the no-purchase alternative with the convention that $\mathcal{N}_+ = \mathcal{N} \cup \{0\}$. Each of the n products in \mathcal{N} is associated with a revenue (or price) $r_j \geq 0$. Under the Markov chain choice model, every product is treated as a state of some underlying Markov chain. We assume that the customers arrive at each state i of the Markov chain with some initial arrival probability λ_i , which is the probability that a customer wants to purchase product i when entering the system. Upon arrival, the customer either buys the product if it is in the offer set, denoted by S , or substitutes to another product j according to the underlying transition probabilities ρ_{ij} . The customer continues this random walk until she lands on a product either in the offer set S , at which point she buys the product, or in the no-purchase state, at which point she leaves the system without purchasing anything. In other words, the products included in the offer set S are absorbing states of the Markov chain and the probability that a customer purchases product i , i.e., the choice probability, is equal to the absorption probability of the corresponding state in the Markov chain.

Let $\boldsymbol{\lambda} = [\lambda_1, \dots, \lambda_n] \in \mathbb{R}^n$ be the vector of arrival rates. Similarly, for each i , let $\boldsymbol{\rho}_i = [\rho_{i1}, \dots, \rho_{in}] \in \mathbb{R}^n$ be the vector of outgoing probabilities from product i . We also let $\boldsymbol{\rho}_i^+ = [\rho_{i0}, \dots, \rho_{in}] \in \mathbb{R}^{n+1}$ be the augmented vector of outgoing probabilities. The vector $\boldsymbol{\rho}_i^+$ is a true probability vector, i.e., $\sum_{j=0}^n \rho_{ij} = 1$. Note that once $\boldsymbol{\rho}_i$ is fixed, $\boldsymbol{\rho}_i^+$ is determined as well. We denote by $\boldsymbol{\rho} \in \mathbb{R}^{n \times n}$ the transition matrix whose i^{th} row is given by $\boldsymbol{\rho}_i$.

Expected revenue. For any fixed assortment of products $S \subseteq \mathcal{N}$ and product i , it is useful to introduce an intermediate variable v_i , which denotes the expected revenue from a customer who is currently considering purchasing product i . This customer could be in state i either because it is the first state she visits or because she transitioned to state i after having visited several products

that were not offered. Note that by the Markovian assumption, v_i does not depend on the visit history of the customer. Moreover, \mathbf{v} satisfies the following set of balance equations:

$$\begin{aligned} v_i &= r_i, \forall i \in S, \\ v_i &= \sum_{j \in \mathcal{N}} \rho_{ij} v_j, \forall i \notin S. \end{aligned} \tag{1}$$

In other words, if the state corresponds to an offered product i in the assortment, i.e., $i \in S$, then the customer buys that product and a revenue of r_i is collected. Otherwise, the product does not belong to the assortment, i.e., $i \notin S$, and the customer transitions to another state j with probability ρ_{ij} at which point she generates an expected revenue of v_j . Therefore, given transition matrix $\boldsymbol{\rho}$ and arrival probability vector $\boldsymbol{\lambda}$, if we let \mathbf{v} be the unique solution to the system of equations (1), then the expected revenue achieved by assortment S can be defined as

$$R^{\text{MC}}(S, \boldsymbol{\rho}, \boldsymbol{\lambda}) = \sum_{i \in \mathcal{N}} \lambda_i v_i. \tag{Rev MC}$$

In Section 2.2, we discuss the assumptions that guarantee a unique solution to the system of equations (1).

Assortment optimization. We can now formulate the assortment optimization under the Markov chain choice model as follows:

$$\max_{S \subseteq \mathcal{N}} R^{\text{MC}}(S, \boldsymbol{\rho}, \boldsymbol{\lambda}). \tag{Assort MC}$$

Leveraging the system of linear equations (1), Blanchet et al. (2016) and Feldman and Topaloglu (2017) show that (Assort MC) admits the following dual formulation:

$$\begin{aligned} \min \quad & \sum_{i \in \mathcal{N}} \lambda_i v_i && \text{(Dual Assort MC)} \\ \text{s.t.} \quad & v_i \geq r_i, \forall i \in \mathcal{N}, \\ & v_i \geq \sum_{j \in \mathcal{N}} \rho_{ij} v_j, \forall i \in \mathcal{N}. \end{aligned}$$

The following result from Feldman and Topaloglu (2017) shows that this formulation is valid.

LEMMA 1 (Theorem 2 from Feldman and Topaloglu 2017). (Dual Assort MC) *and* (Assort MC) *share the same optimal objective values.*

Note that (Dual Assort MC) is a linear program unlike the formulation (Assort MC), which is combinatorial in nature since the feasible set is the set of all possible assortments $S \subseteq \mathcal{N}$.

Uncertainty sets and robust assortment optimization. In practice, the parameters are estimated from data and therefore subject to statistical estimation errors. We model this by assuming that the parameters belong to an uncertainty set. Specifically, let \mathcal{U}^ρ and \mathcal{U}^λ be uncertainty sets (possibly nonconvex) that the model parameters $\boldsymbol{\rho}$ and $\boldsymbol{\lambda}$ are adversarially selected from, respectively. The robust assortment optimization problem under the Markov chain choice model can be expressed as

$$\max_{S \subseteq \mathcal{N}} \min_{\boldsymbol{\rho} \in \mathcal{U}^\rho, \boldsymbol{\lambda} \in \mathcal{U}^\lambda} R^{\text{MC}}(S, \boldsymbol{\rho}, \boldsymbol{\lambda}). \quad (\text{Robust Assort MC})$$

In other words, the robust assortment optimization problem consists of choosing an assortment that maximizes the worst-case expected revenue for an adversarial choice of the parameters from the uncertainty sets \mathcal{U}^ρ and \mathcal{U}^λ .

2.2. Min-max result

We show that, under reasonable assumptions, there exists a min-max duality relation for (Robust Assort MC). Let us begin by discussing the assumption on the structure of the uncertainty set \mathcal{U}^ρ for the min-max result.

ASSUMPTION 1 (Row-wise uncertainty set). *We assume that \mathcal{U}^ρ is a row-wise uncertainty set \mathcal{U}^ρ ; i.e., there exist uncertainty sets $\mathcal{U}^{\rho_1}, \dots, \mathcal{U}^{\rho_n}$ such that \mathcal{U}^ρ can be written as a Cartesian product, $\mathcal{U}^\rho = \mathcal{U}^{\rho_1} \times \dots \times \mathcal{U}^{\rho_n}$.*

In other words, for each product $i \in \mathcal{N}$, the transition probabilities $(\rho_{ij})_{j \in \mathcal{N}}$ belong to an uncertainty set \mathcal{U}^{ρ_i} . Note that the transition probabilities cannot change arbitrarily since for each $i \in \mathcal{N}$, $\sum_{j=0}^n \rho_{ij} = 1$. However, under Assumption 1, the rows of the transition matrix are allowed to vary independently of each other. Note that some papers work with the transpose of the transition matrix and Assumption 1 is referred to as column-wise uncertainty (Gorissen et al. 2014, Awasthi et al. 2019). We further make the following assumption.

ASSUMPTION 2. *We assume that $\mathcal{U}^\lambda \subseteq \mathbb{R}_{++}^n$ and that, for each $i \in \mathcal{N}$,*

$$\mathcal{U}^{\rho_i} \subseteq \{\boldsymbol{\rho}_i \geq 0 : \rho_{ii} = 0 \text{ and } \sum_{j \in \mathcal{N}} \rho_{ij} < 1\}.$$

Assumption 2 is standard for the Markov chain choice model (Feldman and Topaloglu 2017) and is typically stated for a fixed $\boldsymbol{\lambda}$ and $\boldsymbol{\rho}$. Given our robust setting, we require those assumptions to hold for all $\boldsymbol{\lambda}$ and $\boldsymbol{\rho}$ in the uncertainty sets. Note that under Assumption 2, the matrix associated with the system of linear equations (1) is invertible as it is strictly diagonally dominant. Therefore, (1) admits a unique solution and (Rev MC) is well defined. We can now state the main result of this section.

THEOREM 1. *Under Assumptions 1 and 2,*

$$\max_{S \subseteq \mathcal{N}} \min_{\rho \in \mathcal{U}^\rho, \lambda \in \mathcal{U}^\lambda} R^{\text{MC}}(S, \rho, \lambda) = \min_{\rho \in \mathcal{U}^\rho, \lambda \in \mathcal{U}^\lambda} \max_{S \subseteq \mathcal{N}} R^{\text{MC}}(S, \rho, \lambda).$$

Moreover, the optimal robust assortment S^* can be characterized as follows:

$$S^* = \{j \in \mathcal{N} \mid v_j^* = r_j\},$$

where \mathbf{v}^* is the unique fixed point of the mapping $f(\mathbf{v}) : \mathbb{R}^n \rightarrow \mathbb{R}^n$ defined for all $\mathbf{v} \in \mathbb{R}^n$ as

$$f(\mathbf{v})_i = \max \left(r_i, \min_{\rho_i \in \mathcal{U}^{\rho_i}} \sum_{j \in \mathcal{N}} \rho_{ij} v_j \right), \quad \forall i \in \mathcal{N}. \quad (2)$$

Our result shows that under a row-wise uncertainty set, the order of the max and min operators is interchangeable in (Robust Assort MC). This means that the optimal expected revenue where the decision maker first chooses the best possible assortment and then an adversary picks the model parameters in order to minimize the expected revenue of the selected assortment is equal to the expected revenue where the adversary makes the model parameter selection first and then the decision maker chooses the assortment that maximizes the expected revenue. We would like to mention that the objective function does not satisfy properties such as convexity and concavity with respect to the minimization and maximization part respectively for which such min-max duality, also known as saddle point results, is known in the literature. Moreover, our result also comes with a characterization of the optimal robust assortment using the fixed point of the mapping $f(\cdot)$. In Section 5.1, we use this characterization to develop an efficient algorithm for finding the optimal robust assortment.

We end this section by relating our result to Gorissen et al. (2014) who study a robust linear conic program with column-wise uncertainty in the transpose of the coefficient matrix. A subtle difference is that they are dealing with constraints of the form of $v_i \geq \max_{\rho} \sum_{i \in \mathcal{N}} \rho_{ij} v_j$ in the robust version of the linear program (Dual Assort MC), whereas as highlighted by the form of $f(\cdot)$ in (2), we have constraints of the form $v_i \geq \min_{\rho} \sum_{j \in \mathcal{N}} \rho_{ij} v_j$. In addition, the uncertainty set in Gorissen et al. (2014) is required to be convex, but our min-max theorem also holds for nonconvex uncertainty sets. Therefore, the results in Gorissen et al. (2014) are not applicable to our setting.

3. A more general formulation

In this section, we prove Theorem 1 by showing a min-max result under a more general formulation that encompasses the Markov chain choice model.

3.1. A general formulation

Although the assortment optimization under the MNL choice model can be viewed as a special case of the assortment optimization under the Markov chain choice model (Blanchet et al. 2016), a similar relation does not necessarily hold for their robust counterpart. Indeed, because of the row-wise assumption, the robust assortment optimization under the MNL choice model cannot be written as a special case of (Robust Assort MC). Nevertheless, we present a more general formulation of (Robust Assort MC), which allows us to provide a unified treatment of the robust assortment optimization under both choice models. In particular, we consider a model that depends on an uncertain parameter \mathbf{u} that is adversarially selected from an uncertainty set \mathcal{U} . We do not make any assumption about the convexity of \mathcal{U} . Given an assortment $S \subseteq \mathcal{N}$ and the model parameter $\mathbf{u} \in \mathcal{U}$, we assume that the *expected revenue* generated by this assortment is given by

$$R^{\text{Gen}}(S, \mathbf{u}) = \sum_{i \in \mathcal{N}} \lambda(\mathbf{u})_i v_i, \quad (\text{Rev General})$$

where v_i is the unique solution to the following system of balance equations:

$$\begin{aligned} v_i &= r_i, \forall i \in S, \\ \sum_{j \in \mathcal{N}} A(\mathbf{u})_{ij} v_j &= b(\mathbf{u})_i, \forall i \notin S, \end{aligned} \quad (3)$$

and where the parameters $b(\mathbf{u})_i \in \mathbb{R}$, $\lambda(\mathbf{u})_i \in \mathbb{R}$ for all $i \in \mathcal{N}$ and $A(\mathbf{u})_{ij} \in \mathbb{R}$ for all $i \in \mathcal{N}$ and $j \in \mathcal{N}$ all depend on the uncertain parameter \mathbf{u} . Note that (Rev General) and (3) generalize (Rev MC) and (1), respectively. We are interested in the following robust assortment optimization problem where the decision maker wants to maximize the worst-case expected revenue:

$$\max_{S \subseteq \mathcal{N}} \min_{\mathbf{u} \in \mathcal{U}} R^{\text{Gen}}(S, \mathbf{u}). \quad (\text{Robust Assort General})$$

As mentioned earlier, our main result is to show that under suitable assumptions, a max-min relation holds, i.e., that the order of the operators can be switched in (Robust Assort General). It will be useful to define the worst-case revenue for a fixed assortment S .

$$\begin{aligned} \min_{\mathbf{u} \in \mathcal{U}} R^{\text{Gen}}(S, \mathbf{u}) &= \min_{\mathbf{u} \in \mathcal{U}, \mathbf{v}} \sum_{i \in \mathcal{N}} \lambda(\mathbf{u})_i v_i \\ \text{s.t. } v_i &= r_i, \forall i \in S, \\ \sum_{j \in \mathcal{N}} A(\mathbf{u})_{ij} v_j &= b(\mathbf{u})_i, \forall i \notin S. \end{aligned} \quad (\text{Worst-case Rev})$$

We next present our min-max result under this more general formulation.

3.2. Min-max result for general formulation

We first present the assumptions that are needed for our main result. The first two assumptions concern the uncertainty set.

ASSUMPTION 3 (Positivity of λ). *For any $\mathbf{u} \in \mathcal{U}$, $\lambda(\mathbf{u}) > 0$.*

ASSUMPTION 4 (Strictly diagonal dominance of the constraint matrix). *For any $\mathbf{u} \in \mathcal{U}$, $A(\mathbf{u})$ is a strictly row diagonally dominant matrix with positive diagonal elements and nonpositive off-diagonal elements, i.e.,*

- For all $i \in \mathcal{N}$, $A(\mathbf{u})_{ii} > 0$.
- For all $i, j \in \mathcal{N}$ such that $i \neq j$, $A(\mathbf{u})_{ij} \leq 0$.
- For all $i \in \mathcal{N}$, $\sum_{j \in \mathcal{N}} A(\mathbf{u})_{ij} > 0$.

Note that we do not assume any structure on \mathcal{U} . Moreover, under the above assumptions, the matrix associated with the linear system (3) is invertible as it is also strictly diagonally dominant. Therefore, the system of linear equations (3) has a unique solution, and thus $R^{\text{Gen}}(S, \mathbf{u})$ is well defined for all $S \subseteq \mathcal{N}$ and $\mathbf{u} \in \mathcal{U}$. To proceed with our discussion, we define the following problem:

$$\begin{aligned} \min_{\mathbf{u} \in \mathcal{U}, \mathbf{v}} \quad & \sum_{i \in \mathcal{N}} \lambda(\mathbf{u})_i v_i \\ \text{s.t.} \quad & v_i \geq r_i, \forall i \in \mathcal{N}, \\ & \sum_{j \in \mathcal{N}} A(\mathbf{u})_{ij} v_j \geq b(\mathbf{u})_i, \forall i \in \mathcal{N}. \end{aligned} \tag{Dual Assort General}$$

Contrasted with (Dual Assort MC), this can be informally interpreted as the dual of (Robust Assort General). As we will see in Proposition 4, (Robust Assort General) and (Dual Assort General) are indeed equivalent when \mathcal{U} is a singleton.

We can now state the last two assumptions that we need. They require that, for both (Dual Assort General) and (Worst-case Rev), the minimization over $\mathbf{u} \in \mathcal{U}$ be done separately in the objective and right-hand side of the constraint for any fixed \mathbf{v} . For the Markov chain choice model, these assumptions are closely related to Assumption 1.

ASSUMPTION 5. (Dual Assort General) *is equivalent to the following optimization problem:*

$$\begin{aligned} \min_{\mathbf{v}} \quad & \min_{\mathbf{u} \in \mathcal{U}} \sum_{i \in \mathcal{N}} \lambda(\mathbf{u})_i v_i \\ \text{s.t.} \quad & v_i \geq r_i, \forall i \in \mathcal{N}, \\ & v_i \geq \min_{\mathbf{u} \in \mathcal{U}} \left[\sum_{j \neq i} \frac{-A(\mathbf{u})_{ij}}{A(\mathbf{u})_{ii}} v_j + \frac{1}{A(\mathbf{u})_{ii}} b(\mathbf{u})_i \right], \forall i \in \mathcal{N}. \end{aligned}$$

ASSUMPTION 6. Given $S \subseteq \mathcal{N}$, (Worst-case Rev) is equivalent to the following optimization problem:

$$\begin{aligned} \min_{\mathbf{v}} \quad & \min_{\mathbf{u} \in \mathcal{U}} \sum_{i \in \mathcal{N}} \lambda(\mathbf{u})_i v_i \\ \text{s.t.} \quad & v_i = r_i, \forall i \in S, \\ & v_i = \min_{\mathbf{u} \in \mathcal{U}} \left[\sum_{j \neq i} \frac{-A(\mathbf{u})_{ij}}{A(\mathbf{u})_{ii}} v_j + \frac{1}{A(\mathbf{u})_{ii}} b(\mathbf{u})_i \right], \forall i \notin S. \end{aligned}$$

We can now state our main result, which is a generalization of Theorem 2.

THEOREM 2. Under Assumptions 3, 4, 5, and 6,

$$\max_{S \subseteq \mathcal{N}} \min_{\mathbf{u} \in \mathcal{U}} R^{\text{Gen}}(S, \mathbf{u}) = \min_{\mathbf{u} \in \mathcal{U}} \max_{S \subseteq \mathcal{N}} R^{\text{Gen}}(S, \mathbf{u}).$$

Furthermore, the optimal solution of the maximization problem can be characterized as follows:

$$S^* = \{j \in \mathcal{N} \mid v_j^* = r_j\},$$

where \mathbf{v}^* is the unique fixed point of the mapping $f(\mathbf{v}) : \mathbb{R}^n \rightarrow \mathbb{R}^n$ defined for all $\mathbf{v} \in \mathbb{R}^n$ as

$$f_i(\mathbf{v}) = \max \left\{ r_i, \min_{\mathbf{u} \in \mathcal{U}} \left[\sum_{j \neq i} \frac{-A(\mathbf{u})_{ij}}{A(\mathbf{u})_{ii}} v_j + \frac{1}{A(\mathbf{u})_{ii}} b(\mathbf{u})_i \right] \right\}, \forall i \in \mathcal{N}. \quad (4)$$

3.3. Proof of Theorem 1 as a corollary of Theorem 2

Before proceeding to the proof of Theorem 2, we show that all its assumptions are indeed satisfied for the Markov chain choice model and thus Theorem 1 is an immediate consequence of Theorem 2. First, note that the Markov chain choice model is indeed a special case of the more general model introduced in (Rev General) if we let parameters $\mathbf{u} = (\boldsymbol{\lambda}, \boldsymbol{\rho})$ and, for all \mathbf{u} , $\boldsymbol{\lambda}(\mathbf{u}) = \boldsymbol{\lambda}$, $\mathbf{A}(\mathbf{u}) = \mathbf{I} - \boldsymbol{\rho}$, and $\mathbf{b}(\mathbf{u}) = \mathbf{0}$. It is immediate to verify that Assumptions 3 and 4 are satisfied under our assumptions on the Markov chain choice model.

PROPOSITION 1. Under Assumptions 1 and 2, both Assumptions 3 and 4 hold.

Under Assumption 1, the rows in the uncertainty set \mathcal{U}^ρ are unrelated, and under Assumption 2, $\lambda_j > 0$ for all $j \in \mathcal{N}$. We use these facts to show that Assumption 5 holds.

PROPOSITION 2. Under Assumption 1, the Markov chain choice model satisfies Assumption 5, i.e., $\theta_1 = \theta_2$ where

$$\begin{aligned} \theta_1 = \min_{\boldsymbol{\lambda} \in \mathcal{U}^\lambda, \boldsymbol{\rho} \in \mathcal{U}^\rho} \min_{\mathbf{v}} \quad & \sum_{i \in \mathcal{N}} \lambda_i v_i \\ \text{s.t.} \quad & v_i \geq r_i, \forall i \in \mathcal{N}, \\ & v_i \geq \sum_{j \in \mathcal{N}} \rho_{ij} v_j, \forall i \in \mathcal{N}. \end{aligned} \quad (5)$$

$$\begin{aligned} \theta_2 = \min_{\mathbf{v}} \min_{\boldsymbol{\lambda} \in \mathcal{U}^\lambda} \quad & \sum_{i \in \mathcal{N}} \lambda_i v_i \\ \text{s.t.} \quad & v_i \geq r_i, \forall j \in \mathcal{N}, \\ & v_i \geq \min_{\boldsymbol{\rho} \in \mathcal{U}^\rho} \sum_{j \in \mathcal{N}} \rho_{ij} v_j, \forall i \in \mathcal{N}. \end{aligned} \quad (6)$$

Proof. Suppose that $(\mathbf{v}^*, \boldsymbol{\rho}^*, \boldsymbol{\lambda}^*)$ and $(\hat{\mathbf{v}}, \hat{\boldsymbol{\lambda}})$ are the optimal solutions to (5) and (6), respectively. For all $i \in \mathcal{N}$, we have

$$v_i^* \geq \sum_{j \in \mathcal{N}} \rho_{ij}^* v_j^* \geq \min_{\boldsymbol{\rho} \in \mathcal{U}^\rho} \sum_{j \in \mathcal{N}} \rho_{ij} v_j^*.$$

This means that $(\mathbf{v}^*, \boldsymbol{\lambda}^*)$ is feasible for (6), and thus $\sum_{i \in \mathcal{N}} \lambda_i^* v_i^* \geq \sum_{i \in \mathcal{N}} \hat{\lambda}_i \hat{v}_i$. On the other hand, for all $i \in \mathcal{N}$, let

$$\hat{\boldsymbol{\rho}}_i = \operatorname{argmin}_{\boldsymbol{\rho}_i \in \mathcal{U}^{\rho_i}} \sum_{j \in \mathcal{N}} \rho_{ij} \hat{v}_j.$$

The tuple $(\hat{\mathbf{v}}, \hat{\boldsymbol{\rho}}, \hat{\boldsymbol{\lambda}})$ is feasible for (5) as

$$\hat{v}_i \geq \min_{\boldsymbol{\rho} \in \mathcal{U}^\rho} \sum_{j \in \mathcal{N}} \rho_{ij} \hat{v}_j = \min_{\boldsymbol{\rho}_i \in \mathcal{U}^{\rho_i}} \sum_{j \in \mathcal{N}} \rho_{ij} \hat{v}_j = \sum_{j \in \mathcal{N}} \hat{\rho}_{ij} \hat{v}_j.$$

Note that we have used Assumption 1 in the above equality since it allows us to construct each $\hat{\rho}_i$ independently. Consequently, $\sum_{i \in \mathcal{N}} \lambda_i^* v_i^* \leq \sum_{i \in \mathcal{N}} \hat{\lambda}_i \hat{v}_i$ and the conclusion follows. \square

We now show that Assumption 6 also holds.

PROPOSITION 3. *Under Assumptions 1 and 2, for given assortment $S \subseteq \mathcal{N}$, we have $\theta_3 = \theta_4$, where*

$$\theta_3 = \min_{\boldsymbol{\lambda} \in \mathcal{U}^\lambda, \boldsymbol{\rho} \in \mathcal{U}^\rho} \min_{\mathbf{v}} \sum_{i \in \mathcal{N}} \lambda_i v_i \quad \left. \begin{array}{l} \text{s.t. } v_i = r_i, \forall i \in S, \\ v_i = \sum_{j \in \mathcal{N}} \rho_{ij} v_j, \forall i \notin S. \end{array} \right\} \quad (7)$$

$$\theta_4 = \min_{\mathbf{v}} \min_{\boldsymbol{\lambda} \in \mathcal{U}^\lambda} \sum_{i \in \mathcal{N}} \lambda_i v_i \quad \left. \begin{array}{l} \text{s.t. } v_i = r_i, \forall i \in S, \\ v_i = \min_{\boldsymbol{\rho} \in \mathcal{U}^\rho} \sum_{j \in \mathcal{N}} \rho_{ij} v_j, \forall i \notin S. \end{array} \right\} \quad (8)$$

Proof. Suppose that $(\mathbf{v}^*, \boldsymbol{\rho}^*, \boldsymbol{\lambda}^*)$ and $(\hat{\mathbf{v}}, \hat{\boldsymbol{\lambda}})$ are the optimal solutions to (7) and (8), respectively. For all $i \notin S$, let $\hat{\boldsymbol{\rho}}_i = \operatorname{argmin}_{\boldsymbol{\rho}_i \in \mathcal{U}^{\rho_i}} \sum_{j \in \mathcal{N}} \rho_{ij} \hat{v}_j$ and let $\hat{\boldsymbol{\rho}} \in \mathbb{R}^{n \times n}$ be the matrix whose i^{th} row is $\hat{\boldsymbol{\rho}}_i$. Again, note that we can construct such $\hat{\boldsymbol{\rho}}$ because of the row-wise structure of the uncertainty set, i.e., Assumption 1. The tuple $(\hat{\mathbf{v}}, \hat{\boldsymbol{\rho}}, \hat{\boldsymbol{\lambda}})$ is feasible for (7), and thus

$$\sum_{i \in \mathcal{N}} \lambda_i^* v_i^* \leq \sum_{i \in \mathcal{N}} \hat{\lambda}_i \hat{v}_i.$$

Next, we show that $v_i^* = \sum_{j \in \mathcal{N}} \rho_{ij}^* v_j^* = \min_{\boldsymbol{\rho}_i \in \mathcal{U}^{\rho_i}} \sum_{j \in \mathcal{N}} \rho_{ij} v_j^*$ for all $i \notin S$. Suppose by contradiction that we have $v_i^* > \min_{\boldsymbol{\rho}_i \in \mathcal{U}^{\rho_i}} \sum_{j \in \mathcal{N}} \rho_{ij} v_j^*$ for some $i \notin S$. In this case, we can decrease the value of v_i^* by a small amount, while maintaining the feasibility of the solution. This, combined with Assumption 2, i.e., $\lambda_i > 0$ for all $i \in \mathcal{N}$, leads to a solution providing a strictly smaller objective value than the optimal solution and therefore a contradiction arises. Consequently, $(\mathbf{v}^*, \boldsymbol{\lambda}^*)$ is also feasible for (8), and

$$\sum_{i \in \mathcal{N}} \lambda_i^* v_i^* \geq \sum_{i \in \mathcal{N}} \hat{\lambda}_i \hat{v}_i,$$

which completes the proof. \square

4. Proof of the general min-max result

This section is devoted to proving Theorem 2.

4.1. Preliminary results

We start by relating (Robust Assort General), our problem of interest, to (Dual Assort General) when there is no uncertainty.

PROPOSITION 4. Fix some $\mathbf{u} \in \mathcal{U}$, and let $\boldsymbol{\lambda} := \boldsymbol{\lambda}(\mathbf{u})$, $\mathbf{A} := \mathbf{A}(\mathbf{u})$, $\mathbf{b} := \mathbf{b}(\mathbf{u})$. Let

$$\begin{aligned} \theta_5 = \max_{\mathbf{v}, S \subseteq \mathcal{N}} \sum_{i \in \mathcal{N}} \lambda_i v_i \\ \text{s.t. } v_i = r_i, \forall i \in S, \\ \sum_{j \in \mathcal{N}} A_{ij} v_j = b_i, \forall i \notin S. \end{aligned} \quad (9)$$

$$\begin{aligned} \theta_6 = \min_{\mathbf{v}} \sum_{i \in \mathcal{N}} \lambda_i v_i \\ \text{s.t. } v_i \geq r_i, \forall i \in \mathcal{N}, \\ \sum_{j \in \mathcal{N}} A_{ij} v_j \geq b_i, \forall i \in \mathcal{N}. \end{aligned} \quad (10)$$

Under Assumptions 3 and 4, $\theta_5 = \theta_6$.

This result can be interpreted as a generalization of Lemma 1 for the general formulation.

Proof. We first show that $\theta_5 \geq \theta_6$. By Assumption 4, for any given $i \in \mathcal{N}$, $A_{ii} > 0$. Therefore, $\sum_{j \in \mathcal{N}} A_{ij} v_j = b_i$ is equivalent to

$$v_i = \sum_{j \neq i} \frac{-A_{ij}}{A_{ii}} v_j + \frac{1}{A_{ii}} b_i.$$

For all $i \in \mathcal{N}$, let

$$\tilde{A}_{ij} = \begin{cases} \frac{-A_{ij}}{A_{ii}}, & \text{if } j \neq i, \\ 0, & \text{if } j = i, \end{cases}$$

and $\tilde{b}_i = \frac{1}{A_{ii}} b_i$, and let $\tilde{\mathbf{A}} := [\tilde{A}_{ij}]$. Under Assumption 4, $0 < \sum_{j \in \mathcal{N}} \tilde{A}_{ij} < 1$ for any $i \in \mathcal{N}$, and hence $\mathbf{I} \pm \tilde{\mathbf{A}}$ are both strictly diagonally dominant. As a consequence, $\mathbf{I} \pm \tilde{\mathbf{A}} \succ \mathbf{0}$, which further implies that $\tilde{\mathbf{A}}$ has a spectral radius that is strictly less than 1. Now define a mapping $g(\mathbf{v}) = [g_1(\mathbf{v}), \dots, g_n(\mathbf{v})]^\top$, where

$$g_i(\mathbf{v}) = \max \left\{ r_i, \sum_{j \in \mathcal{N}} \tilde{A}_{ij} v_j + \tilde{b}_i \right\}, \quad \forall i \in \mathcal{N}.$$

Note that with this notation, (10) can be equivalently rewritten as

$$\begin{aligned} \min_{\mathbf{v}} \sum_{i \in \mathcal{N}} \lambda_i v_i \\ \text{s.t. } \mathbf{v} \geq g(\mathbf{v}). \end{aligned} \quad (11)$$

We next show that the mapping $g(\mathbf{v}) : \mathbb{R}^n \rightarrow \mathbb{R}^n$ has a unique fixed point. To do so, for any $\mathbf{v}, \mathbf{v}' \in \mathbb{R}^n$, let $d(\mathbf{v}, \mathbf{v}') = \|\mathbf{v} - \mathbf{v}'\|_2$ be the ℓ_2 distance between \mathbf{v} and \mathbf{v}' . For all $\mathbf{v}, \mathbf{v}' \in \mathbb{R}^n$, we have

$$\begin{aligned} d(g(\mathbf{v}), g(\mathbf{v}')) &= \|g(\mathbf{v}) - g(\mathbf{v}')\|_2 \\ &\leq \left\| \tilde{\mathbf{A}}(\mathbf{v} - \mathbf{v}') \right\|_2 \\ &< \|\mathbf{v} - \mathbf{v}'\|_2 = d(\mathbf{v}, \mathbf{v}'), \end{aligned}$$

where the first inequality is true because for any $a, b, c \in \mathbb{R}$,

$$|\max\{a, b\} - \max\{a, c\}| \leq |b - c|,$$

and the second inequality follows because for any matrix $\tilde{\mathbf{A}}$ with a spectral radius strictly less than 1,

$$\sup_x \frac{\|\tilde{\mathbf{A}}\mathbf{x}\|_2}{\|\mathbf{x}\|_2} < 1 \Rightarrow \|\tilde{\mathbf{A}}\mathbf{x}\|_2 < \|\mathbf{x}\|_2, \forall \mathbf{x}.$$

Therefore, $g(\cdot)$ is a contraction mapping and has a unique fixed point \mathbf{v}^* that satisfies $\mathbf{v} = g(\mathbf{v})$ by Banach's fixed point theorem. Moreover, \mathbf{v}^* is feasible for (11).

By Assumption 4, $\tilde{A}_{ij} \geq 0$ for all $i, j \in \mathcal{N}$. Therefore, g is monotonically increasing, i.e., for $\mathbf{v} \geq \mathbf{v}'$, $g(\mathbf{v}) \geq g(\mathbf{v}')$. Consequently, $\mathbf{v} \geq g(\mathbf{v})$ implies that for all k ,

$$g(\mathbf{v}) \geq g(g(\mathbf{v})) \geq \dots \geq g^{k-1}(\mathbf{v}) \geq g^k(\mathbf{v}).$$

In addition, $g^k(\mathbf{v}) \geq r$ for all k . Therefore, the sequence $\{g^k(\mathbf{v})\}_{k=1,2,\dots}$ associated with any feasible solution \mathbf{v} of (11) is monotonically decreasing and bounded from below. Consequently, it must converge to the unique fixed point \mathbf{v}^* and

$$\mathbf{v} \geq g(\mathbf{v}) \geq g(g(\mathbf{v})) \geq \dots \geq \mathbf{v}^*.$$

Therefore, the above fact together with Assumption 3 implies that any feasible solution $\mathbf{v} \neq \mathbf{v}^*$ of (11) has a larger objective value than that of \mathbf{v}^* . Consequently, \mathbf{v}^* is the optimum of (10). Moreover, $S = \{i \mid v_i^* = r_j\}$ is feasible for (9) with objective value $\sum_{i \in \mathcal{N}} \lambda_i v_i^* = \theta_6$, which implies that $\theta_5 \geq \theta_6$.

We now prove the reverse inequality, i.e., $\theta_6 \geq \theta_5$. For any assortment S that is feasible for (9), let

$$v_i(S) := \begin{cases} r_i, & i \in S, \\ \sum_{j \in \mathcal{N}} \tilde{A}_{ij} v_j(S) + \tilde{b}_i, & i \notin S. \end{cases}$$

In what follows, we aim to show that

$$\sum_{i \in \mathcal{N}} \lambda_i v_i \leq \sum_{i \in \mathcal{N}} \lambda_i v_i^*, \tag{12}$$

which in turn implies that $\theta_6 \geq \theta_5$ and completes the proof. From the definition of $g(\cdot)$ and the construction of \mathbf{v} , we have $g(\mathbf{v}) \geq \mathbf{v}$. Moreover, since g is monotonically increasing, it holds that

$$g^0(\mathbf{v}) := \mathbf{v} \leq g(\mathbf{v}) \leq g(g(\mathbf{v})) \leq \dots \leq g^{k-1}(\mathbf{v}) \leq g^k(\mathbf{v});$$

i.e., the sequence $\{g^k(\mathbf{v})\}_{k=0,1,\dots}$ is monotonically increasing. Next, we show that $\{g^k(\mathbf{v})\}_{k=0,1,\dots}$ has a uniform upper bound. To this end, note that for given $j \in \mathcal{N}$, $\mathcal{V}_j := \{v_j(S) : S \subseteq \mathcal{N}\}$ is a finite set and hence is bounded. Let $v_{\max} = \max_{j \in \mathcal{N}} \max_{v_j \in \mathcal{V}_j} v_j$. Furthermore, for all $i \in \mathcal{N}$,

$$\begin{aligned} g_i(\mathbf{v}) &= \max \left\{ r_i, \sum_{j \in \mathcal{N}} \tilde{A}_{ij} v_j + \tilde{b}_i \right\} \\ &= \max \left\{ r_i, \sum_{j \in \mathcal{N}} \tilde{A}_{ij} v_j + \left(1 - \sum_{j \in \mathcal{N}} \tilde{A}_{ij} \right) \frac{\tilde{b}_i}{1 - \sum_{j \in \mathcal{N}} \tilde{A}_{ij}} \right\} \\ &\leq \max \left\{ r_i, v_1, \dots, v_n, \frac{\tilde{b}_i}{1 - \sum_{j \in \mathcal{N}} \tilde{A}_{ij}} \right\} \\ &\leq \max \{ v_{\max}, \delta \}, \end{aligned}$$

where

$$\delta = \max \left\{ r_1, \dots, r_n, \frac{\tilde{b}_1}{1 - \sum_{j \in \mathcal{N}} \tilde{A}_{1j}}, \dots, \frac{\tilde{b}_n}{1 - \sum_{j \in \mathcal{N}} \tilde{A}_{nj}} \right\}.$$

Therefore, the sequence $\{g^k(\mathbf{v})\}_{k=0,1,\dots}$ is bounded from above and converges to the unique fixed point \mathbf{v}^* , which shows that any \mathbf{v} defined by the feasible solution S of the maximization problem (9) satisfies $\mathbf{v} \leq \mathbf{v}^*$. By Assumption 3, Equation (12) holds and $\theta_6 \geq \theta_5$, as desired. \square

REMARK 1. The following example shows that our assumptions are necessary. In particular, when \mathbf{A} is a strictly diagonally dominant matrix but with positive off-diagonal elements, Proposition 4 fails. For instance, consider

$$\mathbf{A} = \begin{bmatrix} 1 & -0.8 \\ 0.8 & 1 \end{bmatrix}, \mathbf{b} = \begin{bmatrix} 0.7 \\ 0.2 \end{bmatrix}, \boldsymbol{\lambda} = \begin{bmatrix} 0.9 \\ 0.1 \end{bmatrix}.$$

In this case, the solutions of the following two problems are different.

- The maximizer of

$$\begin{aligned} &\max_{\mathbf{v}, S \subseteq \mathcal{N}} \sum_{i \in \mathcal{N}} \lambda_i v_i \\ &\text{s.t. } v_i = r_i, \forall i \in S, \\ &\quad \sum_j A_{ij} v_j = b_i, \forall i \notin S \end{aligned}$$

is $S^* = \{2\}$ with $\mathbf{v}_1^* = [0.7, 0]^\top$. The corresponding optimal objective value is 0.63.

- The minimizer of

$$\begin{aligned} \min_{\mathbf{v}} \quad & \sum_{i \in \mathcal{N}} \lambda_i v_i \\ \text{s.t.} \quad & v_i \geq r_i, \forall i \in \mathcal{N}, \\ & \sum_j A_{ij} v_j \geq b_i, \forall i \in \mathcal{N} \end{aligned}$$

is $\mathbf{v}_2^* = [0, 0.875]^\top$. The corresponding optimal objective value is 0.0875.

In this case, Proposition 4 fails because the mapping $g(\cdot)$ in Proposition 4 is no longer monotonically increasing.

4.2. Proof of Theorem 2

As with most min-max results, one direction of the inequality is easily verified. Here, we have that for all $S \subseteq \mathcal{N}$ and $\mathbf{u}' \in \mathcal{U}$, $\min_{\mathbf{u} \in \mathcal{U}} R(S, \mathbf{u}) \leq R(S, \mathbf{u}')$. Taking the maximum over $S \subseteq \mathcal{N}$ on both sides yields $\max_{S \subseteq \mathcal{N}} \min_{\mathbf{u} \in \mathcal{U}} R(S, \mathbf{u}) \leq \max_{S \subseteq \mathcal{N}} R(S, \mathbf{u}')$, which then immediately implies

$$\max_{S \subseteq \mathcal{N}} \min_{\mathbf{u} \in \mathcal{U}} R(S, \mathbf{u}) \leq \min_{\mathbf{u} \in \mathcal{U}} \max_{S \subseteq \mathcal{N}} R(S, \mathbf{u}).$$

We now prove that the reverse inequality also holds. We start by reformulating the min-max problem. Denote $w(\mathbf{v}) = \min_{\mathbf{u} \in \mathcal{U}} \sum_{i \in \mathcal{N}} \lambda(\mathbf{u})_i v_i$. Using Proposition 4 and Assumption 5, we can rewrite the min-max problem as

$$\begin{aligned} \min_{\mathbf{u} \in \mathcal{U}} \max_{S \subseteq \mathcal{N}} R(S, \mathbf{u}) &= \min_{\mathbf{v}} w(\mathbf{v}) \\ \text{s.t.} \quad & v_i \geq r_i, \forall i \in \mathcal{N}, \\ & v_i \geq \min_{\mathbf{u} \in \mathcal{U}} \left[\sum_{j \neq i} \frac{-A(\mathbf{u})_{ij}}{A(\mathbf{u})_{ii}} v_j + \frac{1}{A(\mathbf{u})_{ii}} b(\mathbf{u})_i \right], \forall i \in \mathcal{N}. \end{aligned} \tag{13}$$

Formulation (13) can be simplified to

$$\begin{aligned} \min_{\mathbf{v}} \quad & w(\mathbf{v}) \\ \text{s.t.} \quad & \mathbf{v} \geq f(\mathbf{v}), \end{aligned} \tag{14}$$

where the mapping $f(\cdot)$ was previously defined such that for all \mathbf{v} and $i \in \mathcal{N}$,

$$f_i(\mathbf{v}) = \max \left\{ r_i, \min_{\mathbf{u} \in \mathcal{U}} \left[\sum_{j \neq i} \frac{-A(\mathbf{u})_{ij}}{A(\mathbf{u})_{ii}} v_j + \frac{1}{A(\mathbf{u})_{ii}} b(\mathbf{u})_i \right] \right\}.$$

Then we can show that the optimal solution \mathbf{v}^* of the above problem satisfies $\mathbf{v}^* = f(\mathbf{v}^*)$. To see this, we first show that f is monotonically increasing, i.e., $f(\mathbf{v}') \geq f(\mathbf{v}'')$ for any $\mathbf{v}' \geq \mathbf{v}''$ (the

inequality holds component-wise). For any $\mathbf{u} \in \mathcal{U}$, Assumption 4 ensures that $\frac{-A(\mathbf{u})_{ij}}{A(\mathbf{u})_{ii}} \geq 0$ for all $i, j \in \mathcal{N}$. Therefore, for any $\mathbf{v}' \geq \mathbf{v}''$, by denoting $\mathbf{u}' = \arg \min_{\mathbf{u} \in \mathcal{U}} \left[\sum_{j \neq i} \frac{-A(\mathbf{u})_{ij}}{A(\mathbf{u})_{ii}} v'_j + \frac{1}{A(\mathbf{u})_{ii}} b(\mathbf{u})_i \right]$, we have

$$\begin{aligned} \min_{\mathbf{u} \in \mathcal{U}} \left[\sum_{j \neq i} \frac{-A(\mathbf{u})_{ij}}{A(\mathbf{u})_{ii}} v'_j + \frac{1}{A(\mathbf{u})_{ii}} b(\mathbf{u})_i \right] &= \sum_{j \neq i} \frac{-A(\mathbf{u}')_{ij}}{A(\mathbf{u}')_{ii}} v'_j + \frac{1}{A(\mathbf{u}')_{ii}} b(\mathbf{u}')_i \\ &\geq \sum_{j \neq i} \frac{-A(\mathbf{u}')_{ij}}{A(\mathbf{u}')_{ii}} v''_j + \frac{1}{A(\mathbf{u}')_{ii}} b(\mathbf{u}')_i \\ &\geq \min_{\mathbf{u} \in \mathcal{U}} \left[\sum_{j \neq i} \frac{-A(\mathbf{u})_{ij}}{A(\mathbf{u})_{ii}} v''_j + \frac{1}{A(\mathbf{u})_{ii}} b(\mathbf{u})_i \right], \end{aligned}$$

for all $i \in \mathcal{N}$. As a result, for all $i \in \mathcal{N}$, we have

$$\begin{aligned} f_i(\mathbf{v}') &= \max \left\{ r_i, \min_{\mathbf{u} \in \mathcal{U}} \left[\sum_{j \neq i} \frac{-A(\mathbf{u})_{ij}}{A(\mathbf{u})_{ii}} v'_j + \frac{1}{A(\mathbf{u})_{ii}} b(\mathbf{u})_i \right] \right\} \\ &\geq \max \left\{ r_i, \min_{\mathbf{u} \in \mathcal{U}} \left[\sum_{j \neq i} \frac{-A(\mathbf{u})_{ij}}{A(\mathbf{u})_{ii}} v''_j + \frac{1}{A(\mathbf{u})_{ii}} b(\mathbf{u})_i \right] \right\} = f_i(\mathbf{v}''), \end{aligned}$$

which means that f is monotonically increasing as desired. We are now in a position to show, by contradiction, that (14) is equivalent to

$$\begin{aligned} \min_{\mathbf{v}} w(\mathbf{v}) \\ \text{s.t. } \mathbf{v} = f(\mathbf{v}). \end{aligned} \tag{15}$$

Assume that there exists an optimal solution \mathbf{v}^* of (14) such that $v_i^* > f_i(\mathbf{v}^*)$ for some i . Since f is monotonically increasing and $\lambda(\mathbf{u})_i > 0$ for any $i \in \mathcal{N}$ and $\mathbf{u} \in \mathcal{U}$ due to Assumption 3, we can decrease v_i^* by a small amount while strictly decreasing the objective value and not violating other constraints. This contradicts the optimality of \mathbf{v}^* . In summary, we have shown that the min-max problem (13) is equivalent to (15).

We now investigate the max-min problem. Its inner minimization problem $\min_{\mathbf{u} \in \mathcal{U}} R(S, \mathbf{u})$ is

$$\begin{aligned} \min_{\mathbf{u} \in \mathcal{U}, \mathbf{v}} \sum_{i \in \mathcal{N}} \lambda(\mathbf{u})_i v_i \\ \text{s.t. } v_i = r_i, \forall i \in S, \\ \sum_{j \in \mathcal{N}} A(\mathbf{u})_{ij} v_j = b(\mathbf{u})_i, \forall i \notin S, \end{aligned}$$

which by Assumption 6 is equivalent to

$$\begin{aligned} \min_{\mathbf{u} \in \mathcal{U}} R(S, \mathbf{u}) &= \min_{\mathbf{v}} w(\mathbf{v}) \\ \text{s.t. } v_i &= r_i, \forall i \in S \\ v_i &= \min_{\mathbf{u} \in \mathcal{U}} \left[\sum_{j \neq i} \frac{-A(\mathbf{u})_{ij}}{A(\mathbf{u})_{ii}} v_j + \frac{1}{A(\mathbf{u})_{ii}} b(\mathbf{u})_i \right], \forall i \notin S. \end{aligned} \tag{16}$$

Before proceeding, we state two claims that are needed to complete the proof. Their proofs are deferred to Appendices A and B.

CLAIM 1. For any given $S \subseteq \mathcal{N}$, Problem (16) has a unique feasible solution.

CLAIM 2. The mapping $f(\cdot)$ defined in (4) has a unique fixed point.

Let $S^* \subseteq \mathcal{N}$ be the optimal assortment for $\max_{S \subseteq \mathcal{N}} \min_{\mathbf{u} \in \mathcal{U}} R(S, \mathbf{u})$ with \mathbf{v}^* being the corresponding optimum of the inner problem (16). By the feasibility of \mathbf{v}^* , we have

$$v_i^* = \min_{\mathbf{u} \in \mathcal{U}} \left[\sum_{j \neq i} \frac{-A(\mathbf{u})_{ij}}{A(\mathbf{u})_{ii}} v_j^* + \frac{1}{A(\mathbf{u})_{ii}} b(\mathbf{u})_i \right], \forall i \notin S^* \text{ and } v_i^* = r_i, \forall i \in S^*.$$

It then follows that for any $i \in \mathcal{N}$,

$$f_i(\mathbf{v}^*) = \max \left\{ r_i, \min_{\mathbf{u} \in \mathcal{U}} \left[\sum_{j \neq i} \frac{-A(\mathbf{u})_{ij}}{A(\mathbf{u})_{ii}} v_j^* + \frac{1}{A(\mathbf{u})_{ii}} b(\mathbf{u})_i \right] \right\} \geq v_i^*.$$

On the other hand, taking any $\mathbf{u}_0 \in \mathcal{U}$, we must have

$$f_i(\mathbf{v}^*) \leq \max \left\{ r_i, \sum_{j \neq i} \frac{-A(\mathbf{u}_0)_{ij}}{A(\mathbf{u}_0)_{ii}} v_j^* + \frac{\sum_j A(\mathbf{u}_0)_{ij}}{A(\mathbf{u}_0)_{ii}} \frac{b(\mathbf{u}_0)_i}{\sum_j A(\mathbf{u}_0)_{ij}} \right\} \leq B,$$

where $B := \max \left\{ r_1, \dots, r_n, v_1^*, \dots, v_n^*, \frac{\sum_j A(\mathbf{u}_0)_{1j}}{A(\mathbf{u}_0)_{11}} \frac{b(\mathbf{u}_0)_1}{\sum_j A(\mathbf{u}_0)_{1j}}, \dots, \frac{\sum_j A(\mathbf{u}_0)_{nj}}{A(\mathbf{u}_0)_{nn}} \frac{b(\mathbf{u}_0)_n}{\sum_j A(\mathbf{u}_0)_{nj}} \right\}$. Since $f(\cdot)$ is monotonically increasing, we can recursively prove that $f^{k-1}(\mathbf{v}^*) \leq f^k(\mathbf{v}^*) \leq B$ for any k . It follows that the sequence $\{f^k(\mathbf{v}^*)\}_{k=1,2,\dots}$ is monotonically increasing, bounded from above, and converges to a fixed point $\hat{\mathbf{v}}$ of $f(\cdot)$, i.e., $f(\hat{\mathbf{v}}) = \hat{\mathbf{v}} \geq \mathbf{v}^*$. Consider $\hat{S} = \{i \in \mathcal{N} \mid \hat{v}_i = r_i\}$. Note that $\hat{\mathbf{v}}$ is a feasible solution of (16) associated with \hat{S} . Combining this fact with Claim 1 implies that $\hat{\mathbf{v}}$ is also optimal for (16) with respect to \hat{S} . Consequently, we have $w(\hat{\mathbf{v}}) \leq w(\mathbf{v}^*)$. On the other hand, $\hat{\mathbf{v}} \geq \mathbf{v}^*$ implies that

$$w(\hat{\mathbf{v}}) = \min_{\mathbf{u} \in \mathcal{U}} \sum_{i \in \mathcal{N}} \lambda(\mathbf{u})_i \hat{v}_i \geq \min_{\mathbf{u} \in \mathcal{U}} \sum_{i \in \mathcal{N}} \lambda(\mathbf{u})_i v_i^* = w(\mathbf{v}^*).$$

Furthermore, recall that Assumption 3 guarantees that $\lambda(\mathbf{u})_i > 0$ for any $i \in \mathcal{N}$ and $\mathbf{u} \in \mathcal{U}$, which in turn implies that $\mathbf{v}^* = \hat{\mathbf{v}}$. Hence, $\mathbf{v}^* = f(\mathbf{v}^*)$ and \mathbf{v}^* is a feasible solution of (15). Consequently, the optimal value of the max-min problem is no less than that of (15), which yields

$$\max_{S \subseteq \mathcal{N}} \min_{\mathbf{u} \in \mathcal{U}} R(S, \mathbf{u}) = w(\mathbf{v}^*) \geq \min_{\mathbf{u} \in \mathcal{U}} \max_{S \subseteq \mathcal{N}} R(S, \mathbf{u}).$$

Since $\mathbf{v}^* = \hat{\mathbf{v}} = f(\mathbf{v}^*)$ is the unique fixed point of $f(\cdot)$ by Claim 2, the associated optimal assortment is $S^* = \hat{S} = \{j \in \mathcal{N} \mid v_j^* = r_j\}$. This completes the proof. \square

REMARK 2. According to the proof of Theorem 2, the optimal (S^*, \mathbf{v}^*) depends on \mathcal{U}^ρ but not on \mathcal{U}^λ . The latter affects the value of the objective function but not the optimal assortment.

5. Implications

5.1. Algorithm for the robust Markov chain choice model

In the proof of the min-max result in Theorem 1, we show that the robust assortment optimization problem can be formulated as a fixed point problem. Motivated by this formulation, we present an iterative algorithm for computing the unique fixed point of the following mapping:

$$f(\mathbf{v})_i = \max \left(r_i, \min_{\boldsymbol{\rho}_i \in \mathcal{U}^{\rho_i}} \sum_{j \in \mathcal{N}} \rho_{ij} v_j \right), \quad \forall i \in \mathcal{N},$$

for the Markov chain choice model. If $\mathbf{v}^* = f(\mathbf{v}^*)$ is the fixed point, then we can construct an optimal solution to (Robust Assort MC) by letting $S^* = \{i \in \mathcal{N} \mid v_i^* = r_i\}$. Algorithm 1 details this procedure.

Algorithm 1 Iterative algorithm for computing the optimal robust assortment under the Markov chain choice model

Input: The uncertainty set \mathcal{U}^{ρ_i} for all $i \in \mathcal{N}$

Output: The optimal assortment S^*

```

1: for  $t = 1, 2, \dots$  do
2:   for  $i = 1, 2, \dots, n$  do
3:      $v_i^t \leftarrow \max \left( r_i, \min_{\boldsymbol{\rho}_i \in \mathcal{U}^{\rho_i}} \sum_{j \in \mathcal{N}} \rho_{ij} v_j^{t-1} \right)$ 
4:     if  $v^t = \mathbf{v}^{t-1}$  then return  $S^* = \{i \in \mathcal{N} \mid v_i^t = r_i\}$ 

```

This iterative procedure converges in polynomial time when the no-purchase probability $\rho_{i0} = 1 - \sum_{j \in \mathcal{N}} \rho_{ij}$ is polynomially bounded away from zero for any $\boldsymbol{\rho}_i \in \mathcal{U}^{\rho_i}$ and $i \in \mathcal{N}$. More formally, we need the following to hold true:

$$\delta = \min_{i \in \mathcal{N}} \min_{\boldsymbol{\rho}_i \in \mathcal{U}^{\rho_i}} \rho_{i0} = \Omega(1/n^\alpha),$$

for some constant α .

PROPOSITION 5. *Suppose that $\delta = \Omega(1/n^\alpha)$ for some constant α and the uncertainty set \mathcal{U}^ρ satisfies Assumption 2. Then Algorithm 1 finds an optimal solution to (Robust Assort MC) in polynomially many steps.*

Proof. In the proof of Theorem 2, we have shown that $f(\cdot)$ is monotonically increasing and bounded from above. Therefore, $\{f^t(\mathbf{v}^0)\}_{t=1,2,\dots}$ converges to \mathbf{v}^* for any starting point \mathbf{v}^0 . We now prove that the algorithm terminates in polynomially many steps. Observe that in Algorithm 1, since \mathbf{v}^t is increasing, once v_i^t exceeds r_i for some $i \in \mathcal{N}$, it never goes back to r_i again. Moreover, in each iteration, there is a probability of at least δ of transitioning to state 0. Therefore, after t

steps, the maximum possible expected revenue for i^0 is $(1 - \delta)^t r_{\max}$, and henceforth we use the following notations:

$$r_{\max} = \max_{i \in \mathcal{N}} r_i, \quad r_{\min} = \min_{i \in \mathcal{N}} r_i, \quad i^0 = \arg \min_{i \in \mathcal{N}} r_i.$$

Consequently, the maximum possible iteration number would not be larger than T with $(1 - \delta)^T r_{\max} \geq r_{\min}$. Therefore, the algorithm terminates in at most $\log(r_{\max}/r_{\min})/\delta = \Omega(n^\alpha) \log(r_{\max}/r_{\min})$ steps. Thus, Algorithm 1 converges in polynomially many steps to the fixed point \mathbf{v}^* of $f(\cdot)$. By Theorem 1, this implies that an optimal solution to (Robust Assort MC) is found in polynomially many steps. \square

5.2. Operational insights into the robust Markov chain choice model

In this subsection, we study how the robust optimal assortment changes with respect to change in the uncertainty set and the revenue of each product under the Markov chain choice model. Rusmevichientong and Topaloglu (2012) provide similar results for the robust MNL model and we are able to extend their insights to the Markov chain choice model here.

Recall that we are still working under Assumption 1; i.e., the uncertainty set can be represented as $(\mathcal{U}^\lambda, \mathcal{U}^\rho)$, where $\mathcal{U}^\rho = \mathcal{U}^{\rho_1} \times \dots \times \mathcal{U}^{\rho_n}$. Let $S^*(\mathcal{U}^\lambda, \mathcal{U}^\rho)$ be an optimal assortment for (Robust Assort MC) and $Z^*(\mathcal{U}^\lambda, \mathcal{U}^\rho)$ be the corresponding objective value for given uncertainty sets \mathcal{U}^λ and \mathcal{U}^ρ . We first present a sensitivity analysis with respect to the uncertainty sets.

PROPOSITION 6. *For any $\mathcal{U}^\lambda \subseteq \hat{\mathcal{U}}^\lambda$ and $\mathcal{U}^\rho \subseteq \hat{\mathcal{U}}^\rho$, it holds that $Z^*(\hat{\mathcal{U}}^\lambda, \hat{\mathcal{U}}^\rho) \leq Z^*(\mathcal{U}^\lambda, \mathcal{U}^\rho)$. Moreover, there exists an optimal assortment $S^*(\mathcal{U}^\lambda, \mathcal{U}^\rho)$ and $S^*(\hat{\mathcal{U}}^\lambda, \hat{\mathcal{U}}^\rho)$ such that $S^*(\mathcal{U}^\lambda, \mathcal{U}^\rho) \subseteq S^*(\hat{\mathcal{U}}^\lambda, \hat{\mathcal{U}}^\rho)$.*

Proposition 6, whose proof is presented in Appendix C, states that when the degree of uncertainty increases, the optimal worst-case revenue will decrease. Moreover, the decision maker should offer larger assortments to hedge against greater uncertainty. Increasing product variety helps to hedge against a greater uncertainty in the parameters.

We next provide a characterization of an optimal assortment that relates the optimal robust assortment to the optimal assortments when the parameters are given. In particular, for a given $\boldsymbol{\rho} \in \mathcal{U}^\rho$, let $S_\rho^* = S^*(\mathcal{U}^\lambda, \{\boldsymbol{\rho}\})$. We show that a robust optimal assortment can be constructed by taking the union over all $\{S_\rho^* : \boldsymbol{\rho} \in \mathcal{U}^\rho\}$. As a consequence, the decision maker should focus on the customer types with transition probabilities that lead to a large optimal assortment in order to hedge against the worst-case scenario.

PROPOSITION 7. *For any $(\mathcal{U}^\lambda, \mathcal{U}^\rho)$, $\bigcup_{\boldsymbol{\rho} \in \mathcal{U}^\rho} S_\rho^*$ is an optimal assortment for (Robust Assort MC).*

Proof. Let $S^*(\mathcal{U}^\lambda, \{\boldsymbol{\rho}\})$ be an optimal assortment for (Robust Assort MC). It follows from Proposition 6 that $S_\rho^* = S^*(\mathcal{U}^\lambda, \{\boldsymbol{\rho}\}) \subseteq S^*(\mathcal{U}^\lambda, \mathcal{U}^\rho)$ for any $\boldsymbol{\rho} \in \mathcal{U}^\rho$. Therefore, $\bigcup_{\boldsymbol{\rho} \in \mathcal{U}^\rho} S_\rho^* \subseteq$

$S^*(\mathcal{U}^\lambda, \mathcal{U}^\rho)$. To prove the converse inclusion, let $(\boldsymbol{\rho}^*, \boldsymbol{\lambda}^*)$ denote the optimal solution to $\min_{\boldsymbol{\rho} \in \mathcal{U}^\rho, \boldsymbol{\lambda} \in \mathcal{U}^\lambda} R^{\text{MC}}(S^*(\mathcal{U}^\lambda, \mathcal{U}^\rho), \boldsymbol{\rho}, \boldsymbol{\lambda})$ and let $S_{\boldsymbol{\rho}^*}^* = \arg \max_{S \subseteq \mathcal{N}} R^{\text{MC}}(S, \boldsymbol{\rho}^*, \boldsymbol{\lambda}^*)$. By Theorem 1, we have

$$\begin{aligned} \max_{S \subseteq \mathcal{N}} \min_{\boldsymbol{\rho} \in \mathcal{U}^\rho, \boldsymbol{\lambda} \in \mathcal{U}^\lambda} R^{\text{MC}}(S, \boldsymbol{\rho}, \boldsymbol{\lambda}) &= R^{\text{MC}}(S^*(\mathcal{U}^\lambda, \mathcal{U}^\rho), \boldsymbol{\rho}^*, \boldsymbol{\lambda}^*) \\ &= \min_{\boldsymbol{\rho} \in \mathcal{U}^\rho, \boldsymbol{\lambda} \in \mathcal{U}^\lambda} \max_{S \subseteq \mathcal{N}} R^{\text{MC}}(S, \boldsymbol{\rho}, \boldsymbol{\lambda}) = R^{\text{MC}}(S_{\boldsymbol{\rho}^*}^*, \boldsymbol{\rho}^*, \boldsymbol{\lambda}^*). \end{aligned}$$

This completes the proof. \square

Finally, we present a result showing that the robust optimal assortment shrinks as we decrease the product revenues. This type of result is helpful in revenue management problems that incorporate customer choice behavior. For instance, Feldman and Topaloglu (2017) study revenue management problems where customers choose among offered products according to the Markov chain choice model. In this problem, the decision maker needs to choose, in each period, a static assortment optimization problem where the revenue of each product is reduced by the same amount (see Feldman and Topaloglu 2017). Understanding how the optimal assortment varies when all the revenues are reduced by the same amount is therefore useful.

PROPOSITION 8. *For any $(\mathcal{U}^\lambda, \mathcal{U}^\rho)$, let $S_\eta^*(\mathcal{U}^\lambda, \mathcal{U}^\rho)$ be an optimal robust assortment for the revenues \mathbf{r}^η where $r_i^\eta = r_i + \eta$, for all $i \in \mathcal{N}$. For any $\eta > 0$,*

$$S^*(\mathcal{U}^\lambda, \mathcal{U}^\rho) \subseteq S_\eta^*(\mathcal{U}^\lambda, \mathcal{U}^\rho).$$

In words, additive incremental revenues lead to larger robust assortments. The proof of this result is presented in Appendix D.

5.3. Connection to the MNL model: General uncertainty set

5.3.1. Recovering the MNL choice model. Blanchet et al. (2016) show that the MNL model is a special case of the Markov chain choice model. More specifically, letting $\mathbf{p} = [p_0, p_1, \dots, p_n]$ be the MNL parameters such that $\sum_{i=0}^n p_i = 1$, the expected revenue under the MNL model can be represented as

$$R^{\text{MNL}}(S, \mathbf{p}) = \sum_{i \in \mathcal{N}} p_i v_i, \quad (\text{Rev MNL})$$

where v_i is the unique solution to the following system of equations:

$$\begin{aligned} v_i &= r_i, \forall i \in S, \\ v_i &= \sum_{j \in \mathcal{N}} p_j v_j, \forall i \notin S. \end{aligned}$$

It is straightforward to see that this is again a special case of (Rev General) by letting parameters $\mathbf{u} = \mathbf{p}$, $\boldsymbol{\lambda}(\mathbf{u}) = \mathbf{p}$, $\mathbf{b}(\mathbf{u}) = \mathbf{0}$, and

$$A(\mathbf{u})_{ij} = \begin{cases} 1 - p_j & i = j \\ -p_j & i \neq j \end{cases}.$$

The robust assortment optimization problem under the MNL model can be expressed as

$$\max_{S \subseteq \mathcal{N}} \min_{\mathbf{p} \in \mathcal{U}^p} R^{\text{MNL}}(S, \mathbf{p}). \quad (\text{Robust Assort MNL})$$

In addition, we make the following regularization assumptions.

ASSUMPTION 7. *In the MNL choice model, we assume that for all $\mathbf{p} \in \mathcal{U}^p$, we have $p_i > 0$ for all $i \in \mathcal{N}$ and $\sum_{i \in \mathcal{N}} p_i < 1$.*

Assumption 7 directly implies Assumptions 3 and 4. Moreover,

$$\begin{aligned} & \min_{\mathbf{p} \in \mathcal{U}^p, \mathbf{v}} \sum_{i \in \mathcal{N}} p_i v_i \\ & \text{s.t. } v_i \geq r_i, \forall i \in \mathcal{N}, \\ & v_i \geq \sum_{j \in \mathcal{N}} p_j v_j, \forall i \in \mathcal{N} \end{aligned}$$

is equivalent to

$$\begin{aligned} & \min_{\mathbf{v}} \min_{\mathbf{p} \in \mathcal{U}^p} \sum_{i \in \mathcal{N}} p_i v_i \\ & \text{s.t. } v_i \geq r_i, \forall i \in \mathcal{N}, \\ & v_i \geq \min_{\mathbf{p} \in \mathcal{U}^p} \sum_{j \in \mathcal{N}} p_j v_j, \forall i \in \mathcal{N} \end{aligned}$$

since the terms $\min_{\mathbf{p} \in \mathcal{U}^p} \sum_{i \in \mathcal{N}} p_i v_i$ in the objective and the constraint are identical. As a result, Assumption 5 holds and we can similarly show that Assumption 6 also holds. Consequently, using Theorem 2, we have the following result.

COROLLARY 1. *Under Assumption 7,*

$$\max_{S \subseteq \mathcal{N}} \min_{\mathbf{p} \in \mathcal{U}^p} R^{\text{MNL}}(S, \mathbf{p}) = \min_{\mathbf{p} \in \mathcal{U}^p} \max_{S \subseteq \mathcal{N}} R^{\text{MNL}}(S, \mathbf{p}).$$

Note that because MNL is a special case of the Markov chain choice model, we obtain a stronger result under the MNL model. In particular, Assumption 7 allows for a more general uncertainty set than Assumptions 1 and 2.

5.3.2. Algorithm for finding the optimal robust assortment. A similar procedure to Algorithm 1 also applies to the MNL model. Algorithm 2 details this procedure. The convergence result is similar to that in Algorithm 1, and thus is omitted. This procedure also helps uncover an interesting structural property of the optimal robust assortment. In particular, in every iteration t , $v_i^t = \max(r_i, \gamma^t)$ for all $i \in \mathcal{N}$ where, importantly, γ^t does not depend on i . At every iteration t , let

$$S^t = \{i \in \mathcal{N} \mid v_i^t = r_i\} = \{i \in \mathcal{N} \mid r_i \geq \gamma^t\}.$$

Algorithm 2 Iterative algorithm for computing the optimal robust assortment under the MNL choice model

Input: The uncertainty set \mathcal{U}^p

Output: The optimal assortment S^*

- 1: **for** $t = 1, 2, \dots$ **do**
 - 2: $\gamma^t = \min_{\mathbf{p} \in \mathcal{U}^p} \sum_{i \in \mathcal{N}} p_i v_i^{t-1}$
 - 3: $v_i^t \leftarrow \max(r_i, \gamma^t)$, for all $i \in \mathcal{N}$
 - 4: **if** $\mathbf{v}^t = \mathbf{v}^{t-1}$ **then return** $S^* = \{i \in \mathcal{N} \mid v_i^t = r_i\}$
-

Note that for every t , S^t is therefore a revenue-ordered assortment; i.e., it consists of the highest k revenue products for some k . Since γ^t is an increasing sequence, $\{S^t\}_t$ is a sequence of revenue-ordered assortments such that $S^{t+1} \subseteq S^t$. This implies that S^* as a limit set of $\{S^t\}_t$ is a revenue-ordered assortment, which is consistent with Rusmevichientong and Topaloglu (2012) even though we present a different approach to the problem.

6. Numerical experiments

In this section, we present two computational studies to evaluate the performance and benefits of adopting a robust approach to assortment optimization. In particular, in Section 6.1, we solve (Robust Assort MC) to showcase the running time of our algorithm as well as the magnitude of the tradeoff between the expected and worst-case revenues under our robust approach compared to a deterministic approach that does not account for parameter uncertainty. We then present in Section 6.2 a set of experiments where we generate purchase data according to an *unknown* ranking-based model. In this case, we need to learn the parameters of a Markov chain choice model from the data. We then compare a deterministic approach that directly solves the assortment optimization under the learned parameters and a robust approach that solves a robust assortment optimization problem. We evaluate the performance of these two approaches under the true ranking-based model unknown to the decision maker.

All the experiments are run on a standard desktop computer with a 3.7 GHz Intel Core i5, 16 GB RAM, running Mac OSX Mojave. Moreover, all the mixed-integer programs (MIPs) are solved using Gurobi Optimizer v.9.0.3.

6.1. Solving the robust assortment problem

In this section, we present a numerical study to illustrate the tradeoff between expected performance and worst-case performance when solving (Robust Assort MC). In particular, we show that the running time of our algorithm nicely scales with the number of products.

6.1.1. Experimental setup. We begin by describing the family of random instances being tested in our computational experiments. For each instance, we generate a robust assortment problem as follows. We assume that each product's revenue is uniformly distributed over the interval $[0, 1]$. We then generate a Markov chain by generating $(n + 1)^2$ independent random variables X_{ij} , each picked uniformly over the interval $[0, 1]$. We then set $\rho_{ij}^{\text{modal}} = X_{ij} / \sum_{k=0}^n X_{ik}$ for all $i, j \geq 1$ such that $i \neq j$ and refer to this as the modal Markov chain. We do not allow self-loops, i.e., $\rho_{ii} = 0$ for all i . For the arrival rates, we then generate n independent random variables Y_i , each picked uniformly over the interval $[0, 1]$, and set $\lambda_i^{\text{modal}} = Y_i / \sum_{j=1}^n Y_j$ for all $i \neq 0$. For $\epsilon > 0$, we define a row-wise uncertainty set as follows. For each i , let

$$\mathcal{U}_\epsilon^{\rho_i} = \left\{ \rho_i^{\text{modal}} + \gamma_i \mid \sum_j \gamma_{ij} = 0, \text{ and } \forall j, \max\{(1 - \epsilon)\rho_{ij}^{\text{modal}}, 0\} \leq \rho_{ij}^{\text{modal}} + \gamma_{ij} \leq \min\{(1 + \epsilon)\rho_{ij}^{\text{modal}}, 1\} \right\}. \quad (17)$$

In other words, the uncertainty set we consider is centered around ρ^{modal} and the magnitude of variations around ρ^{modal} is controlled by ϵ . For this known uncertainty set, we compute two assortments. The first one, S^{modal} , is the optimal assortment when the Markov chain parameters are given by ρ^{modal} . In a way, this is the assortment that one would compute if one were taking the average parameters as the real parameters and not accounting for robustness. The second one, $S_\epsilon^{\text{robust}}$, is the optimal robust assortment computed using Algorithm 1, i.e., the assortment that maximizes the worst-case expected revenue over the uncertainty set $\mathcal{U}_\epsilon = \mathcal{U}_\epsilon^{\rho_1} \times \dots \times \mathcal{U}_\epsilon^{\rho_n}$. For each assortment and level of uncertainty ϵ , we compare the expected revenue under the modal parameters, denoted by $R^{\text{modal}}(\cdot)$, and the worst-case expected revenue, denoted by $R_\epsilon^{\text{worst}}(\cdot)$. For the latter, we use (Worst-case Rev) to compute the worst-case expected revenue generated by a given assortment over \mathcal{U}_ϵ .

6.1.2. Results and discussion. We highlight the tradeoff between expected revenue and worst-case expected revenue when accounting for parameter uncertainty. In particular, for $\epsilon \in \{0.05, 0.10, 0.25, 0.5\}$, we compute

$$\Delta_{\text{modal}} = \frac{R^{\text{modal}}(S_\epsilon^{\text{robust}})}{R^{\text{modal}}(S^{\text{modal}})}, \text{ and } \Delta_{\text{worst}} = \frac{R_\epsilon^{\text{worst}}(S_\epsilon^{\text{robust}})}{R_\epsilon^{\text{worst}}(S^{\text{modal}})}.$$

Δ_{modal} is the ratio of expected revenue under the modal parameters and is a proxy for the average performance of the assortment. On the other hand, Δ_{worst} is the ratio of worst-case expected revenue and captures how well the assortment hedges against the uncertainty in the parameters. In Table 1, we report the average and minimum/maximum value of these ratios over 100 randomly generated instances. First, note that all the values are less than one for Δ_{modal} and greater than one for Δ_{worst} . This is not surprising as S^{modal} maximizes $R^{\text{modal}}(\cdot)$ and $S_\epsilon^{\text{robust}}$ maximizes $R_\epsilon^{\text{worst}}(\cdot)$, respectively.

Table 1 Tradeoff between modal and worst-case expected revenues when $n = 20$ and $n = 50$. The average and minimum are taken over 100 instances.

ϵ	$n = 20$				$n = 50$			
	Δ_{modal}		Δ_{worst}		Δ_{modal}		Δ_{worst}	
	Average	Minimum	Average	Maximum	Average	Minimum	Average	Maximum
0.05	0.9999	0.9969	1.0001	1.0022	0.9999	0.9984	1.0001	1.0010
0.10	0.9993	0.9947	1.0004	1.0082	0.9996	0.9974	1.0004	1.0025
0.25	0.9959	0.9744	1.0051	1.0299	0.9971	0.9889	1.0037	1.0154
0.50	0.9825	0.9377	1.0334	1.1354	0.9861	0.9604	1.0240	1.0693

Table 2 Running time of Algorithm 1 for $\epsilon = 0.25$. The average and maximum are taken over 100 instances.

n	Running time (s)		Number of iterations	
	Average	Maximum	Average	Maximum
10	0.32	0.67	20.01	44
20	2.13	3.92	32.72	61
30	5.95	9.19	40.12	61
40	12.94	18.39	48.98	69
50	21.82	32.29	52.81	78

Moreover, as ϵ increases, the uncertainty set we hedge against becomes larger. Therefore, Δ_{modal} decreases as Δ_{worst} increases. In terms of the magnitude of the improvement, we observe that there is almost a linear tradeoff between the expected and worst-case expected revenues with respect to the average performance of all instances. For example, when $\epsilon = 0.50$ and $n = 50$, the robust approach captures on average 98.61% of $R^{\text{modal}}(S^{\text{modal}})$ while the expected revenue $R_{\epsilon}^{\text{worst}}(S_{\epsilon}^{\text{robust}})$ is on average 2.40% higher than $R_{\epsilon}^{\text{worst}}(S^{\text{modal}})$. In terms of extreme performance, the robust approach seems to limit the losses while providing consequent gains under worst-case parameters. For instance, when $\epsilon = 0.50$ and $n = 20$, the minimum value of Δ_{modal} over all the instances is 0.9377 while the maximum value of Δ_{worst} is 1.1354.

In terms of running time, solving the robust assortment problem is more expensive. However, our iterative approach nicely scales in the number of products n , as illustrated in Table 2. The number of iterations grows linearly with the number of products and (Robust Assort MC) can be solved in around 20s on average when $n = 50$.

6.2. Unknown ground truth

In practice, the parameters of the Markov chain choice model are not known and need to be learned from the data. We present in this section a more realistic set of numerical experiments where we learn the parameters of the Markov chain choice model from the data and compare a robust approach to a deterministic one that does not account for any parameter uncertainty. We begin by describing the setting and then the corresponding results.

6.2.1. Experimental setup. In our numerical experiments, we assume that we have access only to purchase data in order to learn the parameters of the Markov chain choice model.

The ground truth choice model. We adopt a ranking-based choice model as the unknown ground truth choice model that governs the customer choice process (Mahajan and van Ryzin 2001, Honhon et al. 2012, van Ryzin and Vulcano 2017, Jagabathula and Rusmevichientong 2016, Farias et al. 2013). In this model, the preferences are described by a probability distribution over rankings or preference lists of products. Each preference list specifies a rank ordering of the products such that lower-ranked products are more preferred. In the experiments, we randomly generate m ranked lists σ^g for $g = 1, \dots, m$. Each list $\sigma^g = (\sigma_1^g, \dots, \sigma_{n+1}^g)$ is an ordering of the products in \mathcal{N}^+ where $\sigma_i^g < \sigma_j^g$ indicates that product σ_i^g is preferred to product σ_j^g under preference list σ^g . We denote by β^g the probability that an arriving customer chooses the ranked list σ^g for $g = 1, \dots, m$. The setup is inspired by Şimşek and Topaloglu (2018) and we use a similar process to generate β^g , which we describe next. For each $g = 1, \dots, m$, we first generate γ^g uniformly over the interval $[0, 1]$ and then set $\beta^g = \gamma^g / \sum_{h=1}^m \gamma^h$. In each choice instance, a customer samples a preference list from the underlying distribution and then chooses the most preferred available product (possibly, the no-purchase option) from her list. Then, given an assortment S , the probability that product i is chosen under the ranking-based choice model is $\sum_{g=1}^m \beta^g \cdot \mathbf{1}\{i = \arg \min_{j \in S} \sigma_j^g\}$. We set $m = 2n$ in our experiments. Finally, we assume that for each product i , there is one ranked list where the most preferred product is product i .

Generating purchase data. Once we have generated a ground truth choice model, we use it to generate some purchase data $\{(S^t, \mathbf{Z}^t) : t = 1, \dots, T\}$. More precisely, for each customer t , S^t denotes the offered assortment and $\mathbf{Z}^t = (Z_1^t, \dots, Z_n^t)$ denotes the purchase decisions; i.e., $Z_i^t = 1$ if and only if the customer purchases product i . Following Şimşek and Topaloglu (2018), for each assortment S^t , the no-purchase option is always available and each of the other products is offered with probability $1/2$.

Benchmark. Similar to the previous section, we compare the robust Markov chain approach with a deterministic Markov chain approach that does not account for parameter uncertainty. For both approaches, we apply the expectation-maximization (EM) algorithm from Şimşek and Topaloglu (2018) to the historical purchase data $(S^t, \mathbf{Z}^t)_t$ in order to compute $\boldsymbol{\rho}^{\text{modal}}$ and $\boldsymbol{\lambda}^{\text{modal}}$, which are the estimated arrival probability vector and transition probability matrix of the Markov chain choice model. For the deterministic approach, we then assume that these parameters are the correct ones, and compute the offer set S^{modal} by solving the corresponding deterministic assortment optimization problem.

A data-driven design of the uncertainty set. For the robust approach, we account for some uncertainty in the estimated parameters. In particular, we propose a data-driven approach to constructing the uncertainty sets inspired by the bootstrap method (Efron and Tibshirani 1986). This method, popular in statistics, is a practical technique that provides approximations to coverage probabilities of confidence intervals by resampling from the data or using a model estimated from the data. Our detailed procedure for constructing a row-wise uncertainty set for the Markov chain choice model is given in Algorithm 3. More specifically, we use the estimated arrival probability vector λ^{modal} and transition probability matrix ρ^{modal} from the deterministic approach as a ground truth model to generate K new sets of purchase data. With each newly set of generated purchase data, we use the EM algorithm again to get a new set of estimated parameters $\rho^{(k)}$. The variations in the estimated coefficient $\rho^{(k)}$ from the bootstrap procedure drive the construction of our uncertainty set. Indeed, for parameters that are close to each other over the different estimations, we construct a smaller uncertainty set. On the other hand, we build a bigger uncertainty set around the parameters that exhibit more variance. In particular, we use the magnitude of the ratio $\rho_{ij}^{(k)} / \rho_{ij}^{\text{modal}}$ to inform the uncertainty we allow around ρ_{ij}^{modal} . As a result, the uncertainty set returned by Algorithm 3 is centered around ρ^{modal} and thus is similar to the previous section. We also scale the uncertainty set uniformly by a parameter $\alpha \in [0, 1]$ in order to control the robustness level and observe the effects of introducing more or less uncertainty. Note that we do not construct the uncertainty set on the arrival probability λ , as the optimal assortment is independent of the uncertainty in λ . For each robustness level α , let S_α^{robust} be the optimal robust assortment computed using Algorithm 1, i.e., the assortment that maximizes the worst-case expected revenue over the uncertainty set \mathcal{U}_α^ρ constructed in Algorithm 3.

6.2.2. Results and discussion. In this setting where the ground truth model is not known, we show that our robust approach can help address two sources of potential errors. First, there are estimation errors arising from the potentially insufficient amount of data. Second, there may be misspecification errors since we are fitting a Markov chain choice model whereas the ground truth choice model is not.

Let $R^{\text{true}}(\cdot)$ be the expected revenue under the ground truth model. We denote by S^{true} the assortment that maximizes the expected revenue under the ground truth model. Despite the assortment optimization problem under the ranking-based choice model being NP-hard (Aouad et al. 2018), we can use a mixed-integer program (Bertsimas and Mišić 2019) to compute S^{true} . For $m \in \{\text{modal}, \text{robust}\}$ and $\alpha \in \{0.05, 0.10, 0.25, 0.50\}$, we compute

$$\Delta_{\text{expected}}^m = \frac{R^{\text{true}}(S^m)}{R^{\text{true}}(S^{\text{true}})},$$

Algorithm 3 Algorithm for constructing the uncertainty set of a Markov chain choice model, based on bootstrap

Input: The purchase data $\{(S^t, \mathbf{Z}^t(S^t)) : t = 1, \dots, T\}$.

Output: The uncertainty set \mathcal{U}_α^ρ .

- 1: Apply EM algorithm to $\{(S^t, \mathbf{Z}^t(S^t))\}$ and compute the estimated arrival probability $\boldsymbol{\lambda}^{\text{modal}}$ and transition probability matrix $\boldsymbol{\rho}^{\text{modal}}$.
- 2: **for** $k = 1, 2, \dots, K$ **do**
- 3: **(Resampling)** Independently draw new purchase data $\{(S^t, \hat{\mathbf{Z}}^t(S^t))\}$ with the ground truth being a Markov chain choice model with parameters $\boldsymbol{\lambda}^{\text{modal}}, \boldsymbol{\rho}^{\text{modal}}$.
- 4: Apply EM algorithm to $\{(S^t, \hat{\mathbf{Z}}^{k,t}(S^t))\}$ to get another estimator $\boldsymbol{\rho}^{(k)}$.
- 5: **(Constructing uncertainty set)** For $i = 1, \dots, n$, compute

$$\mathcal{U}_\alpha^{\rho_i} = \left\{ \boldsymbol{\rho}_i^{\text{modal}} + \boldsymbol{\gamma}_i \mid \sum_j \gamma_{ij} = 0, \text{ and } \forall j, \max\{(1 - \epsilon_{ij})\rho_{ij}^{\text{modal}}, 0\} \leq \rho_{ij}^{\text{modal}} + \gamma_{ij} \leq \min\{(1 + \epsilon_{ij})\rho_{ij}^{\text{modal}}, 1\} \right\}$$

where $\epsilon_{ij} = \alpha \frac{\sum_{k=1}^K |\rho_{ij}^{(k)} / \rho_{ij}^{\text{modal}} - 1|}{K}$ and $0 < \alpha \leq 1$ controls the robustness level of the uncertainty set.

- 6: Return $\mathcal{U}_\alpha^\rho = \mathcal{U}_\alpha^{\rho_1} \times \dots \times \mathcal{U}_\alpha^{\rho_n}$
-

which quantifies how far the assortment is from the ground truth optimal assortment in terms of expected revenue.

To measure the robustness of the different approaches, we look at several other quantities. The first one is the probability of no-purchase under the ground truth model. More precisely, denote $\pi^{\text{true}}(i, S)$ as the choice probability of product i when the offer set is S under the ground truth model. For $m \in \{\text{modal}, \text{robust}\}$, let $p_{\text{purch}} = 1 - \pi^{\text{true}}(0, S^m)$. We use this probability as a proxy for the robustness of the solution. Indeed, if a customer does not purchase any product, this yields zero revenue. A higher probability of purchase therefore implies a better hedge against those bad events. In the same spirit, we also report the q^{th} percentile of the ground truth revenue when considered as a random variable. More precisely, let $X^{\text{true}}(S^m)$ be a random variable such that $X^{\text{true}}(S^m) = r_i$ with probability $\pi^{\text{true}}(i, S^m)$. Then, for $m \in \{\text{modal}, \text{robust}\}$, we report $R_{q^{\text{th}}}^{\text{true}}(S^m)$ defined such that $P(X^{\text{true}}(S^m) \leq R_{q^{\text{th}}}^{\text{true}}(S^m)) = q$. For a given $q \in \{5, 10\}$, a higher value of $R_{q^{\text{th}}}^{\text{true}}(S^m)$ means that the assortment is more robust since it guarantees that with probability $1 - q$, the realized revenue is above $R_{q^{\text{th}}}^{\text{true}}(S^m)$.

Table 3 reports the different metrics, which are averaged over 100 instances for each set of parameters. Note that we present the results for different values of $T \in \{1000, 2500, 5000\}$. We observe that even in this more realistic setting, the robust approach offers a nice tradeoff between

Table 3 Performance of the robust Markov chain approach when $n = 10$ and $n = 20$. The reported metrics are averaged over 100 instances.

T	m	α	$n = 10$				$n = 20$			
			$\Delta_{\text{expected}}^m$	$\Delta_{5^{\text{th}}}^m$	$\Delta_{10^{\text{th}}}^m$	p_{purch}	$\Delta_{\text{expected}}^m$	$\Delta_{5^{\text{th}}}^m$	$\Delta_{10^{\text{th}}}^m$	p_{purch}
1000	modal	-	0.9692	0.0862	0.1523	0.8279	0.9646	0.0741	0.2350	0.8714
		0.05	0.9698	0.0869	0.1563	0.8290	0.9645	0.0595	0.2406	0.8741
	robust	0.10	0.9715	0.0872	0.1532	0.8324	0.9638	0.0792	0.2790	0.8796
		0.25	0.9707	0.0872	0.1530	0.8409	0.9610	0.0830	0.2950	0.8845
		0.50	0.9674	0.0909	0.1776	0.8510	0.9563	0.1123	0.3384	0.8996
2500	modal	-	0.9729	0.1060	0.2350	0.8512	0.9642	0.0560	0.2811	0.8776
		0.05	0.9729	0.1208	0.2427	0.8512	0.9643	0.0555	0.2650	0.8780
	robust	0.10	0.9708	0.1124	0.2419	0.8532	0.9635	0.0555	0.2744	0.8780
		0.25	0.9717	0.1189	0.2503	0.8588	0.9659	0.0743	0.3086	0.8842
		0.50	0.9730	0.1108	0.2716	0.8672	0.9641	0.0883	0.3300	0.8901
5000	modal	-	0.9681	0.0505	0.0968	0.8090	0.9661	0.0357	0.1678	0.8555
		0.05	0.9678	0.0505	0.0968	0.8090	0.9660	0.0442	0.1971	0.8567
	robust	0.10	0.9678	0.0502	0.0965	0.8110	0.9653	0.0357	0.1959	0.8569
		0.25	0.9691	0.0580	0.1118	0.8159	0.9672	0.0522	0.2091	0.8617
		0.50	0.9722	0.0580	0.1178	0.8259	0.9637	0.0521	0.2138	0.8640

expected revenue and worst-case expected revenue. First, in terms of expected revenue, the variations in $\Delta_{\text{expected}}^{\text{robust}}$ are very mild when α increases, i.e., when we are assuming a larger uncertainty set when computing $S_{\alpha}^{\text{robust}}$. Even when $\Delta_{\text{expected}}^{\text{robust}}$ decreases, the loss compared to the deterministic assortment S^{modal} is quite limited. For instance, for $T = 5000$ and $n = 20$, the deterministic S^{modal} captures 96.61% of the expected revenue of the optimal ground truth assortment S^{true} while the robust assortment $S_{\alpha}^{\text{robust}}$ with $\alpha = 0.10$ captures 96.53% of the expected revenue. Moreover, it turns out that in many cases, the robust approach actually outperforms the deterministic approach. For instance, for $T = 1000$ and $n = 10$, the deterministic S^{modal} captures 96.92% of the expected revenue of the optimal ground truth assortment S^{true} while the robust assortment $S_{\alpha}^{\text{robust}}$ with $\alpha = 0.10$ captures 97.15% of the expected revenue. It appears that given the lack of data and the model misspecification, adding some robustness can help hedge against the case where the estimated parameters are far from the ground truth.

We next compare how the two approaches perform in terms of robustness. Both the 5th and 10th percentiles as well as purchase probability increase with α , suggesting a more robust solution. Moreover, the magnitude of the gains in terms of robustness seems to be significant. For instance, for $T = 2500$, $\alpha = 0.50$, and $n = 20$, the expected revenue of the robust approach is very close to that of the deterministic approach. More precisely, $\Delta_{\text{expected}}^{\text{robust}}/\Delta_{\text{expected}}^{\text{modal}} = 0.999$. On the other hand, the average 5th percentile in revenue of the robust approach increases by more than 57% compared to

the deterministic one while the probability of purchase increases from 87.76% to 89.01%, suggesting an increased robustness.

Interestingly, in many instances, taking a robust approach dominates the deterministic approach both in expected revenue and in robustness. This suggests that in more realistic settings, accounting for some uncertainty in the parameters can lead to better average and worst-case performances! We also test the performance of an MNL model. Note that Algorithm 2 allows computing a robust assortment in this case as well. The results, presented in Appendix E, show that the Markov chain approach significantly outperforms the MNL approach.

7. Conclusion

In this paper, we study the robust assortment optimization problem under the Markov chain choice model. Under reasonable assumptions, mainly that the uncertainty sets across the different rows of the transition matrix are unrelated, we show that this problem admits a max-min duality relationship; i.e., the two operators in the robust optimization problem can be swapped. This is surprising as none of the properties for known saddle point results are satisfied. Based on the duality results, we also develop an efficient iterative algorithm for finding the optimal robust assortment.

To prove our main result, we introduced a general formulation for choice models that assumes that the choice probabilities, which capture the substitution behavior of customers, are solutions to a system of linear equations. It would be interesting to see if this approach can unify an even broader class of choice models. For instance, Désir et al. (2021) recently showed that under the Mallows model (a choice model based on a particular probability distribution over preference lists), the choice probabilities can be obtained by solving a system of linear equations. If we can develop general estimation and/or optimization techniques for this class of models, this might give a more parsimonious approach to modeling choice.

Finally, another interesting research direction is to push the results for broader types of uncertainty sets for the Markov chain choice model, in particular by allowing the uncertainty sets across different rows of the Markov chain choice model to be related. This would allow modeling a budget of uncertainty across rows in order to limit the adversary's power.

Acknowledgments

The authors thank the associate editor and the referees for their thoughtful and detailed comments, which greatly improved the presentation of the paper.

References

Aouad, Ali, Vivek Farias, Retsef Levi. 2021. Assortment optimization under consider-then-choose choice models. *Management Science* **67**(6) 3368–3386.

- Aouad, Ali, Vivek F Farias, Retsef Levi, Danny Segev. 2018. The approximability of assortment optimization under ranking preferences. *Operations Research* **66**(6) 1661–1669.
- Awasthi, Pranjal, Vineet Goyal, Brian Y Lu. 2019. On the adaptivity gap in two-stage robust linear optimization under uncertain packing constraints. *Mathematical Programming* **173**(1) 313–352.
- Ben-Akiva, Moshe E, Steven R Lerman, Steven R Lerman. 1985. *Discrete Choice Analysis: Theory and Application to Travel Demand*. MIT Press.
- Ben-Tal, Aharon, Laurent El Ghaoui, Arkadi Nemirovski. 2009. *Robust Optimization*. Princeton University Press.
- Ben-Tal, Aharon, Arkadi Nemirovski. 2000. Robust solutions of linear programming problems contaminated with uncertain data. *Mathematical Programming* **88**(3) 411–424.
- Bertsimas, Dimitris, Velibor V Mišić. 2019. Exact first-choice product line optimization. *Operations Research* **67**(3) 651–670.
- Bertsimas, Dimitris, Melvyn Sim. 2003. Robust discrete optimization and network flows. *Mathematical Programming* **98**(1) 49–71.
- Birbil, Ş İlker, JBG Frenk, Joaquim AS Gromicho, Shuzhong Zhang. 2009. The role of robust optimization in single-leg airline revenue management. *Management Science* **55**(1) 148–163.
- Blanchet, Jose, Guillermo Gallego, Vineet Goyal. 2016. A Markov chain approximation to choice modeling. *Operations Research* **64**(4) 886–905.
- Chen, Li, Simai He, Shuzhong Zhang. 2011. Tight bounds for some risk measures, with applications to robust portfolio selection. *Operations Research* **59**(4) 847–865.
- Davis, James M, Guillermo Gallego, Huseyin Topaloglu. 2014. Assortment optimization under variants of the nested logit model. *Operations Research* **62**(2) 250–273.
- Désir, Antoine, Vineet Goyal, Srikanth Jagabathula, Danny Segev. 2021. Mallows-smoothed distribution over rankings approach for modeling choice. *Operations Research* **69**(4) 1206–1227.
- Désir, Antoine, Vineet Goyal, Danny Segev, Chun Ye. 2020. Constrained assortment optimization under the Markov chain-based choice model. *Management Science* **66**(2) 698–721.
- Désir, Antoine, Vineet Goyal, Jiawei Zhang. 2014. Near-optimal algorithms for capacity constrained assortment optimization. *Available at SSRN 2543309* .
- Efron, Bradley, Robert Tibshirani. 1986. Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Statistical Science* **1**(1) 54–75.
- Farias, Vivek F, Srikanth Jagabathula, Devavrat Shah. 2013. A nonparametric approach to modeling choice with limited data. *Management Science* **59**(2) 305–322.
- Feldman, Jacob B, Huseyin Topaloglu. 2017. Revenue management under the Markov chain choice model. *Operations Research* **65**(5) 1322–1342.

- Gallego, Guillermo, Huseyin Topaloglu. 2014. Constrained assortment optimization for the nested logit model. *Management Science* **60**(10) 2583–2601.
- Gorissen, Bram L, Hans Blanc, Dick den Hertog, Aharon Ben-Tal. 2014. Deriving robust and globalized robust solutions of uncertain linear programs with general convex uncertainty sets. *Operations Research* **62**(3) 672–679.
- Honhon, Dorothée, Vishal Gaur, Sridhar Seshadri. 2010. Assortment planning and inventory decisions under stockout-based substitution. *Operations Research* **58**(5) 1364–1379.
- Honhon, Dorothée, Sreelata Jonnalagedda, Xiajun Amy Pan. 2012. Optimal algorithms for assortment selection under ranking-based consumer choice models. *Manufacturing & Service Operations Management* **14**(2) 279–289.
- Jagabathula, Srikanth, Paat Rusmevichientong. 2016. A nonparametric joint assortment and price choice model. *Management Science* **63**(9) 3128–3145.
- Kök, A Gürhan, Marshall L Fisher, Ramnath Vaidyanathan. 2015. Assortment planning: Review of literature and industry practice. *Retail Supply Chain Management*. Springer, 175–236.
- Mahajan, Siddharth, Garrett van Ryzin. 2001. Stocking retail assortments under dynamic customer substitution. *Operations Research* **49**(3) 334–351.
- Mak, Ho-Yin, Ying Rong, Jiawei Zhang. 2014. Appointment scheduling with limited distributional information. *Management Science* **61**(2) 316–334.
- McFadden, Daniel. 1978. Modeling the choice of residential location. *Transportation Research Record* (673).
- McFadden, Daniel, Kenneth Train. 2000. Mixed MNL models for discrete response. *Journal of Applied Econometrics* **15**(5) 447–470.
- Perakis, Georgia, Guillaume Roels. 2010. Robust controls for network revenue management. *Manufacturing & Service Operations Management* **12**(1) 56–76.
- Rusmevichientong, Paat, David Shmoys, Chaoxu Tong, Huseyin Topaloglu. 2014. Assortment optimization under the multinomial logit model with random choice parameters. *Production and Operations Management* **23**(11) 2023–2039.
- Rusmevichientong, Paat, Huseyin Topaloglu. 2012. Robust assortment optimization in revenue management under the multinomial logit choice model. *Operations Research* **60**(4) 865–882.
- Şimşek, A Serdar, Huseyin Topaloglu. 2018. An expectation-maximization algorithm to estimate the parameters of the Markov chain choice model. *Operations Research* **66**(3) 748–760.
- Talluri, Kalyan, Garrett van Ryzin. 2004. Revenue management under a general discrete choice model of consumer behavior. *Management Science* **50**(1) 15–33.
- Thiele, Aurélie. 2009. Multi-product pricing via robust optimisation. *Journal of Revenue and Pricing Management* **8**(1) 67–80.

- van Ryzin, Garrett, Gustavo Vulcano. 2017. An expectation-maximization method to estimate a rank-based choice model of demand. *Operations Research* **65**(2) 396–407.
- Wang, Xuan, Jiawei Zhang. 2015. Process flexibility: A distribution-free bound on the performance of k -chain. *Operations Research* **63**(3) 555–571.
- Williams, Huw CWL. 1977. On the formation of travel demand models and economic evaluation measures of user benefit. *Environment and Planning A* **9**(3) 285–344.
- Xu, Guanglin, Samuel Burer. 2018. A copositive approach for two-stage adjustable robust optimization with uncertain right-hand sides. *Computational Optimization and Applications* **70** 33–59.
- Zhang, Dan, William L Cooper. 2005. Revenue management for parallel flights with customer-choice behavior. *Operations Research* **53**(3) 415–431.
- Zhu, Shushang, Masao Fukushima. 2009. Worst-case conditional value-at-risk with application to robust portfolio management. *Operations Research* **57**(5) 1155–1168.

This page is intentionally blank. Proper e-companion title page, with INFORMS branding and exact metadata of the main paper, will be produced by the INFORMS office when the issue is being assembled.

Online Appendix

Robust Assortment Optimization under the Markov chain choice model

Appendix A: Proof of Claim 1

We fix some $S \subseteq \mathcal{N}$ and construct a new mapping $h(\cdot)$ defined for all \mathbf{v} by

$$h_i(\mathbf{v}) = r_i, i \in S; h_i(\mathbf{v}) = \min_{\mathbf{u} \in \mathcal{U}} \left[\sum_{j \neq i} \frac{-A(\mathbf{u})_{ij}}{A(\mathbf{u})_{ii}} v_j + \frac{1}{A(\mathbf{u})_{ii}} b(\mathbf{u})_i \right], i \notin S.$$

To prove the claim, it suffices to show this mapping has a unique fixed point. Suppose on the contrary that h has two different fixed points, \mathbf{v}^1 and \mathbf{v}^2 . For $i \notin S$, $h_i(\mathbf{v})$ is the minimum of many affine functions and is therefore concave. Consequently, by letting

$$\tilde{A}(\mathbf{u})_{ij} = \begin{cases} \frac{-A(\mathbf{u})_{ij}}{A(\mathbf{u})_{ii}}, & j \neq i, \\ 0, & j = i, \end{cases}$$

we can find $\mathbf{y}^1 \in \partial h_i(\mathbf{v}^1) \subseteq \text{Conv}\{\tilde{\mathbf{A}}(\mathbf{u})_{i:} \mid \mathbf{u} \in \mathcal{U}\}$ and $\mathbf{y}^2 \in \partial h_i(\mathbf{v}^2) \subseteq \text{Conv}\{\tilde{\mathbf{A}}(\mathbf{u})_{i:} \mid \mathbf{u} \in \mathcal{U}\}$ such that

$$(\mathbf{y}^1)^\top (\mathbf{v}^1 - \mathbf{v}^2) \leq h_i(\mathbf{v}^1) - h_i(\mathbf{v}^2) \leq (\mathbf{y}^2)^\top (\mathbf{v}^1 - \mathbf{v}^2).$$

For $i \in \mathcal{N}$, let

$$\tilde{\mathbf{u}}^i := \arg \max_{\mathbf{u} \in \mathcal{U}} \sum_{j \neq i} \frac{-A(\mathbf{u})_{ij}}{A(\mathbf{u})_{ii}} |v_j^1 - v_j^2|.$$

Consider the matrix $\tilde{\mathbf{A}}(\tilde{\mathbf{U}})$ with $\tilde{\mathbf{U}} = [\tilde{\mathbf{u}}^1, \dots, \tilde{\mathbf{u}}^n]$ whose i^{th} row is equal to $\tilde{\mathbf{A}}(\tilde{\mathbf{u}}^i)_{i:}$. Assumption 4 implies that all components of $\tilde{\mathbf{A}}(\tilde{\mathbf{U}})$ are nonnegative. Then, for all $i \notin S$,

$$|h_i(\mathbf{v}^1) - h_i(\mathbf{v}^2)| \leq \max \left\{ \sum_{j \in \mathcal{N}} y_j^1 |v_j^1 - v_j^2|, \sum_{j \in \mathcal{N}} y_j^2 |v_j^1 - v_j^2| \right\} \leq \sum_{j \in \mathcal{N}} \tilde{A}(\tilde{\mathbf{u}}^i)_{ij} |v_j^1 - v_j^2|.$$

For all $i \in S$, $|h_i(\mathbf{v}^1) - h_i(\mathbf{v}^2)| = 0 \leq \sum_{j \in \mathcal{N}} \tilde{A}(\tilde{\mathbf{u}}^i)_{ij} |v_j^1 - v_j^2|$.

Assumption 4 also implies that $0 < \sum_{j \in \mathcal{N}} \tilde{A}(\tilde{\mathbf{u}}^i)_{ij} < 1$ for any $i \in \mathcal{N}$, and hence $\mathbf{I} \pm \tilde{\mathbf{A}}(\tilde{\mathbf{U}})$ are both strictly diagonally dominant. As a consequence, $\mathbf{I} \pm \tilde{\mathbf{A}}(\tilde{\mathbf{U}}) \succ \mathbf{0}$ and, moreover, $\tilde{\mathbf{A}}(\tilde{\mathbf{U}})$ has a spectral radius that is strictly less than 1. Let \mathbf{z} be the left eigenvector of $\tilde{\mathbf{A}}(\tilde{\mathbf{U}})$ associated with the largest absolute eigenvalue τ . It follows from Perron–Frobenius theorem that $\mathbf{z} > \mathbf{0}$ and $0 < \tau < 1$. As a result,

$$\sum_{i \in \mathcal{N}} z_i |v_i^1 - v_i^2| = \sum_{i \in \mathcal{N}} z_i |h_i(\mathbf{v}^1) - h_i(\mathbf{v}^2)|$$

$$\begin{aligned}
&\leq \sum_{i \in \mathcal{N}} z_i \sum_{j \in \mathcal{N}} \tilde{A}(\tilde{\mathbf{u}}^i)_{ij} |v_j^1 - v_j^2| \\
&= \sum_{j \in \mathcal{N}} |v_j^1 - v_j^2| \left(\sum_{i \in \mathcal{N}} z_i \tilde{A}(\tilde{\mathbf{u}}^i)_{ij} \right) \\
&= \sum_{j \in \mathcal{N}} |v_j^1 - v_j^2| \tau z_j \\
&= \tau \cdot \sum_{j \in \mathcal{N}} z_j |v_j^1 - v_j^2|,
\end{aligned}$$

which contradicts the fact that $\mathbf{v}^1 \neq \mathbf{v}^2$ and completes the proof. \square

Appendix B: Proof of Claim 2

Define the mapping \tilde{h} such that for all \mathbf{v} ,

$$\tilde{h}_i(\mathbf{v}) = \min_{\mathbf{u} \in \mathcal{U}} \left[\sum_{j \neq i} \frac{-A(\mathbf{u})_{ij}}{A(\mathbf{u})_{ii}} v_j + \frac{1}{A(\mathbf{u})_{ii}} b(\mathbf{u})_i \right], \forall i \in \mathcal{N}.$$

With this notation, we have $f_i(\mathbf{v}) = \max\{r_i, \tilde{h}_i(\mathbf{v})\}$ for all \mathbf{v} and $i \in \mathcal{N}$. Suppose by contradiction that f has two different fixed points, \mathbf{v}^1 and \mathbf{v}^2 . In the proof of Claim 1, we have shown that $|\tilde{h}_i(\mathbf{v}^1) - \tilde{h}_i(\mathbf{v}^2)| \leq \tau |v_i^1 - v_i^2|$ with $0 < \tau < 1$. Using this together with the following inequality:

$$|\max\{a, b\} - \max\{a, c\}| \leq |b - c|,$$

we have that

$$|f_i(\mathbf{v}^1) - f_i(\mathbf{v}^2)| \leq |\tilde{h}_i(\mathbf{v}^1) - \tilde{h}_i(\mathbf{v}^2)| \leq \tau |v_i^1 - v_i^2|.$$

This is a contradiction and concludes the proof. \square

Appendix C: Proof of Proposition 6

Given $(\mathcal{U}^\lambda, \mathcal{U}^\rho)$ and $(\hat{\mathcal{U}}^\lambda, \hat{\mathcal{U}}^\rho)$, denote by \mathbf{v}^* and $\hat{\mathbf{v}}$ the corresponding fixed point defined in Theorem 1, i.e.,

$$v_i^* = \max \left\{ r_i, \min_{\rho_i \in \mathcal{U}^{\rho_i}} \sum_{j \in \mathcal{N}} \rho_{ij} v_j^* \right\} \text{ and } \hat{v}_i = \max \left\{ r_i, \min_{\rho_i \in \hat{\mathcal{U}}^{\rho_i}} \sum_{j \in \mathcal{N}} \rho_{ij} \hat{v}_j \right\} \text{ for any } i \in \mathcal{N}.$$

Let $f(\cdot)$ be the mapping associated with \mathcal{U}^ρ , i.e.,

$$f(\mathbf{v})_i = \max \left\{ r_i, \min_{\rho_i \in \mathcal{U}^{\rho_i}} \sum_{j \in \mathcal{N}} \rho_{ij} v_j \right\}, i \in \mathcal{N}.$$

Since $\mathcal{U}^\rho \subseteq \hat{\mathcal{U}}^\rho$ and using the monotonicity of $f(\cdot)$, we have

$$\hat{\mathbf{v}} \leq f(\hat{\mathbf{v}}) \leq f^2(\hat{\mathbf{v}}) \leq \dots \leq f^d(\hat{\mathbf{v}}) \rightarrow \mathbf{v}^*.$$

Therefore, $\hat{\mathbf{v}} \leq \mathbf{v}^*$. Moreover, since $\mathcal{U}^\lambda \subseteq \hat{\mathcal{U}}^\lambda$, we obtain

$$Z^*(\hat{\mathcal{U}}^\lambda, \hat{\mathcal{U}}^\rho) = \min_{\lambda \in \hat{\mathcal{U}}^\lambda} \sum_{i \in \mathcal{N}} \lambda_i \hat{v}_i \leq \min_{\lambda \in \mathcal{U}^\lambda} \sum_{i \in \mathcal{N}} \lambda_i v_i^* = Z^*(\mathcal{U}^\lambda, \mathcal{U}^\rho).$$

Additionally,

$$\min_{\rho_i \in \mathcal{U}^{\rho_i}} \sum_{j \in \mathcal{N}} \rho_{ij} v_j^* \geq \min_{\rho_i \in \mathcal{U}^{\rho_i}} \sum_{j \in \mathcal{N}} \rho_{ij} \hat{v}_j \geq \min_{\rho_i \in \hat{\mathcal{U}}^{\rho_i}} \sum_{j \in \mathcal{N}} \rho_{ij} \hat{v}_j.$$

Consequently, it follows from Theorem 1 that

$$\begin{aligned} S^*(\mathcal{U}^\lambda, \mathcal{U}^\rho) &= \left\{ i \in \mathcal{N} : v_i^* = r_i \geq \min_{\rho_i \in \mathcal{U}^{\rho_i}} \sum_{j \in \mathcal{N}} \rho_{ij} v_j^* \right\} \\ &\subseteq \left\{ i \in \mathcal{N} : \hat{v}_i = r_i \geq \min_{\rho_i \in \hat{\mathcal{U}}^{\rho_i}} \sum_{j \in \mathcal{N}} \rho_{ij} \hat{v}_j \right\} = S^*(\hat{\mathcal{U}}^\lambda, \hat{\mathcal{U}}^\rho). \end{aligned}$$

This completes the proof. \square

Appendix D: Proof of Proposition 8

Let \mathbf{v}^η be the fixed point associated with $S_\eta^*(\mathcal{U}^\lambda, \mathcal{U}^\rho)$. We define $\tilde{\mathbf{v}}$ by letting $\tilde{v}_i = v_i^\eta - \eta$ for all $i \in \mathcal{N}$. For all $i \in S_\eta^*(\mathcal{U}^\lambda, \mathcal{U}^\rho)$, we have $\tilde{v}_i = r_i$. For $i \notin S_\eta^*(\mathcal{U}^\lambda, \mathcal{U}^\rho)$, we have

$$\min_{\rho_i \in \mathcal{U}^{\rho_i}} \sum_{j \in \mathcal{N}} \rho_{ij} \tilde{v}_j = \min_{\rho_i \in \mathcal{U}^{\rho_i}} \sum_{j \in \mathcal{N}} \rho_{ij} (v_j^\eta - \eta) > \min_{\rho_i \in \mathcal{U}^{\rho_i}} \sum_{j \in \mathcal{N}} \rho_{ij} v_j^\eta - \eta = v_i^\eta - \eta = \tilde{v}_i > r_i,$$

where the first inequality holds since $0 < \sum_j \rho_{ij} < 1$ for all $\rho_i \in \mathcal{U}^{\rho_i}$. Consider the mapping $f(\cdot)$ such that for all \mathbf{v} ,

$$f_i(\mathbf{v}) = \max \left\{ r_i, \min_{\rho_i \in \mathcal{U}^{\rho_i}} \sum_{j \in \mathcal{N}} \rho_{ij} v_j \right\}, \quad i \in \mathcal{N}.$$

Let \mathbf{v}^* be the unique fixed point of $f(\cdot)$. Using the monotonicity of $f(\cdot)$, we have

$$\tilde{\mathbf{v}} \leq f(\tilde{\mathbf{v}}) \leq f^2(\tilde{\mathbf{v}}) \leq \dots \leq f^d(\tilde{\mathbf{v}}) \rightarrow \mathbf{v}^*.$$

Therefore, for any $i \notin S_\eta^*(\mathcal{U}^\lambda, \mathcal{U}^\rho)$, $v_i^* \geq \tilde{v}_i > r_i$. Consequently, $i \notin S^*(\mathcal{U}^\lambda, \mathcal{U}^\rho)$ by Theorem 1. This implies that $\{i \notin (S_\eta^*(\mathcal{U}^\lambda, \mathcal{U}^\rho))\} \subseteq \{i \notin (S^*(\mathcal{U}^\lambda, \mathcal{U}^\rho))\}$ and in turn that $S^*(\mathcal{U}^\lambda, \mathcal{U}^\rho) \subseteq S_\eta^*(\mathcal{U}^\lambda, \mathcal{U}^\rho)$. \square

Appendix E: MNL benchmark

For the experiments described in Section 6.2, we also test an approach that is based on an MNL model. As for the Markov chain approach, we compute a deterministic MNL and a robust MNL solution. For the deterministic approach, we use a standard EM algorithm (e.g., Talluri and van Ryzin 2004) to estimate the parameter of the MNL model $\mathbf{p}^{\text{modal}}$. We then compute the optimal assortment S^{modal} under the estimated parameters.

For the robust approach, we account for some uncertainty in the estimated parameters and construct the uncertainty sets by the bootstrap method described in Algorithm 4. More specifically,

Table EC.1 Performance of the robust MNL approach when $n = 10$ and $n = 20$. The metrics are computed over 100 instances.

T	m	α	$n = 10$				$n = 20$			
			$\Delta_{\text{expected}}^m$	$\Delta_{5^{\text{th}}}^m$	$\Delta_{10^{\text{th}}}^m$	p_{purch}	$\Delta_{\text{expected}}^m$	$\Delta_{5^{\text{th}}}^m$	$\Delta_{10^{\text{th}}}^m$	p_{purch}
1000	modal	-	0.9546	0.0641	0.1384	0.8183	0.9536	0.0079	0.1542	0.8423
		0.05	0.9539	0.0641	0.1384	0.8186	0.9536	0.0166	0.1542	0.8423
	robust	0.10	0.9539	0.0574	0.1384	0.8186	0.9536	0.0079	0.1622	0.8426
		0.25	0.9543	0.0641	0.1384	0.8205	0.9537	0.0166	0.1525	0.8438
		0.50	0.9532	0.0574	0.1384	0.8205	0.9553	0.0161	0.1828	0.8482
2500	modal	-	0.9500	0.0469	0.1425	0.8180	0.9555	0.0400	0.1828	0.8560
		0.05	0.9500	0.0469	0.1427	0.8180	0.9555	0.0400	0.1828	0.8560
	robust	0.10	0.9497	0.0469	0.1431	0.8180	0.9555	0.0400	0.1908	0.8560
		0.25	0.9508	0.0469	0.1428	0.8202	0.9553	0.0400	0.1908	0.8570
		0.50	0.9504	0.0469	0.1489	0.8211	0.9565	0.0396	0.1980	0.8601
5000	modal	-	0.9550	0.0399	0.0939	0.7943	0.9562	0.0227	0.1155	0.8410
		0.05	0.9550	0.0312	0.0937	0.7943	0.9561	0.0227	0.1239	0.8410
	robust	0.10	0.9555	0.0399	0.1012	0.7953	0.9559	0.0227	0.1073	0.8412
		0.25	0.9540	0.0399	0.0937	0.7961	0.9558	0.0227	0.1239	0.8415
		0.50	0.9566	0.0510	0.1044	0.8023	0.9562	0.0154	0.1230	0.8423

we use the estimated parameter $\mathbf{p}^{\text{modal}}$ of the MNL model from the deterministic approach as a ground truth model to generate K new sets of purchase data. With each newly generated set of purchase data, we use the EM algorithm again to get a new set of estimated parameters $\mathbf{p}^{(k)}$. Then, we use the magnitude of the ratio $p_j^{(k)}/p_j^{\text{modal}}$ to inform the uncertainty we allow around p_j^{modal} . We also scale the uncertainty set uniformly by a parameter $\alpha \in [0, 1]$ in order to control the robustness level and observe the effects of introducing more or less uncertainty. For each $\alpha \in \{0.05, 0.10, 0.25, 0.50\}$, let S_α^{robust} be the optimal robust assortment computed using Algorithm 2, i.e., the assortment that maximizes the worst-case expected revenue over the uncertainty set \mathcal{U}_α^p . Table EC.1 reports the same metrics when using an MNL model instead of a Markov chain choice model.

We find that the Markov chain approach outperforms the MNL approach. Consistent with existing literature (Blanchet et al. 2016), the deterministic Markov chain outperforms the deterministic MNL with respect to the expected performance metric $\Delta_{\text{expected}}^{\text{modal}}$. We find that the robust assortment using the Markov chain also outperforms the robust assortment using the MNL model for almost all the parameters. In particular, not only does the Markov chain approach dominate the MNL approach with respect to the expected revenue metric, it also dominates the MNL approach with respect to the robustness metric. More precisely, for all uncertainty levels α , the 5th percentile metric and the purchase probability are higher under the Markov chain approach. Moreover, it seems that the variations are much smaller in the case of the MNL model, suggesting that adding

robustness does not provide a lot of benefits in our experiments. On the other hand, adopting a robust Markov chain approach can simultaneously improve the expected revenue as well as the robustness of the solution.

Algorithm 4 Algorithm for constructing an uncertainty set of a MNL model, based on bootstrap

Input: The purchase data $\{(S^t, \mathbf{Z}^t(S^t)) : t = 1, \dots, T\}$

Output: Uncertainty set \mathcal{U}_ϵ^v

- 1: Apply EM algorithm to $\{(S^t, \mathbf{Z}^t(S^t))\}$ to get an estimated probability $\mathbf{p}^{\text{modal}}$ of the MNL model.
- 2: **for** $k = 1, 2, \dots, K$ **do**
- 3: **(Resampling)** Independently draw new purchase data $\{(S^t, \hat{\mathbf{Z}}^t(S^t))\}$ with the ground truth being a MNL model with parameters $\mathbf{p}^{\text{modal}}$.
- 4: Apply EM algorithm to $\{(S^t, \hat{\mathbf{Z}}^t(S^t))\}$ and get another estimator $\mathbf{p}^{(k)}$.
- 5: **(Constructing the uncertainty set)** Compute

$$\mathcal{U}_\alpha^p = \left\{ \mathbf{p}^{\text{modal}} + \boldsymbol{\gamma} \mid \sum_j \gamma_j = 0 \text{ and } \forall j, \max\{(1 - \epsilon_j)p_j^{\text{modal}}, 0\} \leq p_j^{\text{modal}} + \gamma_j \leq \min\{(1 + \epsilon_j)p_j^{\text{modal}}, 1\} \right\}$$

where $\epsilon_j = \alpha \frac{\sum_{k=1}^K |p_j^{(k)} / p_j^{\text{modal}} - 1|}{K}$ and $0 < \alpha \leq 1$ controls the robustness level of the uncertainty set.
