Submitted to *Operations Research*
manuscript OPRE-2019-03-133

Robust Assortment Optimization under the Markov Chain 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 problems arise 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 to customers in order to maximize the expected revenue. We study the robust assortment optimization under the Markov chain choice model. In this formulation, the parameters of the model are assumed to be unknown and the goal is to maximize the worst-case expected revenue over all parameter values in an uncertainty set. Our main contribution is to show a min-max duality result when the uncertainty set is row-wise. To prove this result, we introduce a framework that captures the robust assortment problem under a more general class of choice models. Under this class of models, which includes the Markov chain model, the choice probabilities are given as solutions to a system of linear equations. We show, under certain reasonable assumptions, a min-max duality result for the robust assortment optimization for this class of choice models. This is surprising as the objective function does not satisfy the properties usually needed for known saddle point results. Furthermore, we give an efficient algorithm to compute the optimal robust assortment under the Markov chain model. Inspired by the min-max relation, this algorithm is iterative in nature. Moreover, our results yield operational insights towards the effect of changing the uncertainty set on the optimal robust assortment. In particular, consistent with previous literature, we find that bigger uncertainty set always lead to bigger assortment, and a firm should offer larger assortments to hedge against uncertainty.

1. Introduction

Assortment optimization problems arise 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 to 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 that is captured by a choice model which specifies the probability 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 his 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 this limitation, 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 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 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 problem. This paper adds to this stream of work by considering a robust variant of the assortment problem under the Markov chain model.

In practice, the parameters of the underlying choice model have to be estimated from data. Statistical errors in these parameters are therefore unavoidable. By ignoring this uncertainty on the parameters, “optimal” decisions based on the point estimators could potentially be sub-optimal 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 parameters 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 are as follows.

Min-max duality. The robust formulation of the assortment optimization problem consists of a minimization problem nested inside a maximization problem. Our main result is to show a min-max duality for the robust assortment optimization under the Markov chain model. When the uncertainty set is row-wise, i.e. when the uncertainty across the different rows of the transition matrix are unrelated, the order of the max and min operators in the robust assortment optimization formulation can be interchanged. This implies that the optimal expected revenue where the decision maker first chooses the best possible assortment and then an adversary picks the model parameters 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 to maximize the expected revenue. We would like to mention that our objective function does not satisfy the 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, holds.

A general framework. To prove our main result, we introduce a more general framework that captures the robust assortment optimization problem under a broader class of choice models. In particular, we consider a setting where the choice probabilities, and therefore the revenue function, are given as solutions to a system of linear equations. A special case of this setting is the Markov chain choice model, where the substitution behavior is captured by the transitions in an underlying Markov chain and where the system of linear equations naturally arises as the balance equations in the Markov chain. The matrix defining the system of linear equations depends on a parameter belonging to an uncertainty set. The goal in the robust assortment optimization problem is to select the assortment that achieves the best worst-case expected revenue over all possible parameters in the uncertainty set. We establish, under certain reasonable assumptions, a min-max duality for the robust assortment optimization for this class of choice models which then implies our result for the Markov chain. Interestingly, it also implies a min-max relation for the MNL model under a general uncertainy set. Note that because the MNL model is a special case of the Markov chain model (Blanchet et al. (2016)), we are able to obtain a stronger result, i.e. for general uncertainty set, for the MNL model.

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

Numerical experiments. Finally, we conduct some numerical experiments that help quantify the tradeoffs between the capability of hedging against the uncertainty in the model parameters and the conservativeness of the associated optimal robust assortment. In particular, we present two computational studies. In the first one, we explore a setting where the underlying ground truth is a known Markov chain model with uncertain parameters. We then present a more realistic set of experiments where the ground truth is an unknown ranking-based model. In this case, we assume that we only have access to choice data and propose a data-driven approach where we first learn the parameters of Markov chain and then construct uncertainty sets using a procedure inspired by the bootstrap method. In both cases, we show how choosing the level of uncertainty allows trading off between the worst-case and average performances. Moreover, in the more realistic setting, our results show that taking a robust approach indeed yields a higher worst-case performance and can even generates a higher average performance in many cases. This indicates that our approach is able to provide robustness without sacrificing much optimality in the revenue when there is model mis-specification and insufficient amount of training data such that the estimated parameters is far from the ground truth.

1.2. Related literature

Our work is closely to 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 limitation and capture a richer class of substitution, 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 planing is choosing a model that strikes a good balance between its predictability and tractability. The interested readers are referred to Kök et al. (2015) for a comprehensive background reading on the assortment optimization problems.

The MNL model belongs to the very rich class of random utility models wherein the utility of each product is modeled as the sum of a deterministic component and a random noise. The assumption on the joint distribution of these noises specifies the choice model. For instance, when the noises 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 the mixture of MNL model (McFadden and Train 2000). Under the nested logit model, the products are grouped into nests and products in the same nest have positively correlated noises which implies that they are closer substitutes. Under a mixture of MNL model, customer heterogeneity is added by considering multiple segments with each following a different MNL model. The assortment optimization problem has been studied under both of these extensions. Under the mixture of MNL model, the assortment optimization problem becomes NP-hard, even when the number of mixtures is two (Rusmevichientong et al. (2014)). Under the nested logit model, the assortment optimization problem is tractable (Davis et al. 2014) even when adding additional capacity constraints (Gallego and Topaloglu 2014, Désir et al. 2014). One main limitation of this model is that it depends on a predefined nest structure which 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 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 behavior. 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 untractable 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 $\mathrm{P} = \mathrm{NP}$ (Aouad et al. 2018). Several special cases of this model lead to more tractable optimization problems (Honhon et al. (2010), Honhon et al. (2012), Aouad et al. (2015), Jagabathula and Rusmevichientong (2016)).

In this paper, we focus on the Markov chain based choice model which has recently received a growing interest (Blanchet et al. 2016, Şimşek and Topaloglu 2018, Feldman and Topaloglu 2017, Désir et al. 2020). The most related work is Rusmevichientong and Topaloglu (2012) which studies the robust assortment optimization under the MNL model. Since the Markov chain 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 model and we have to develop new tools to solve the robust assortment optimization problem under the Markov chain based 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 (2016)) which incorporates the uncertainty in the model parameters into the decision making process. Recently, this paradigm has found some application in the operations management and revenue management literature such as 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)), appointment scheduling (Mak et al. (2014)). Gorissen et al. (2014) study a robust linear conic program with column-wise uncertainty on the transpose of the coefficient matrix, and they show that this problem this problem is computationally tractable. Even though the setup is related, their results does not apply in our setting. We give more details after having introduced the notation in Section 2.1.

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 framework which we use to prove our main result. Section 4 is dedicated to the proof of the min-max result under the general framework. In Section 5, we discuss some implications of our main result. In particular, we give an efficient algorithm to compute 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 introduce the Markov chain model and formulate the associated robust assortment optimization problem. We then state our main theorem on the min-max result of the robust assortment model.

2.1. Markov chain model under uncertainty

Model parameters. We consider a universe of n products denoted by $\mathcal{N} = \{1, 2, \ldots, n\}$. We let product 0 be the no-purchase alternative with the convention that $\mathcal{N}_+ = \mathcal{N} \cup \{0\}$. Each of the n products is associated with a revenue (or price) $r_j \geq 0$. Under the Markov chain 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 probabilities λ_i. Upon arrival, the customer either buys the product if it is offered, i.e. $i \in S$, or substitutes to another product j according to the underlying transition probabilities ρ_{ij}. The customer continues this random walk until she either lands on a product in the offer set, 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 associated state in the Markov chain. The model parameters are:

1. An initial arrival probability λ_i for each state $i \in \mathcal{N}$, which is the probability that a customer wants to purchase product i when entering into the system.

2. The transition probabilities ρ_{ij} for all $i \in \mathcal{N}$, $j \in \mathcal{N}_+$, which can be thought of as the probability that the customer transits into considering purchasing product j when her favorite product i is not available.

Let $\boldsymbol{\lambda} = [\lambda_1, \cdots, \lambda_n]^\top$ be the vector of arrival rates. Similarly, for each i, let $\boldsymbol{\rho}_i = [\rho_{i1}, \cdots, \rho_{in}]^\top$ be the vector of out-going probabilities from product i. We also let $\boldsymbol{\rho}_i^+ = [\rho_{i0}, \cdots, \rho_{in}]^\top$ be the augmented vector of out-going probabilities and which a true probability vector, i.e. $\sum_{i=0}^n \rho_{ij} = 1$. Note that once $\boldsymbol{\rho}_i$ is fixed, $\boldsymbol{\rho}_i^+$ is determined as well. We denote by $\boldsymbol{\rho}$ 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 either be in state i 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 Markov assumption, v_i does not depend on the visit history of the customer. Moreover, $\boldsymbol{v}$ satisfies the following set of balance equations

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

That is, if the state corresponds to an offered product i in the assortment, i.e. $i \in S$, then the customer buys that product and generates a revenue of r_i. Otherwise, the product does not belong to the assortment, i.e. $i \notin S$, then the customer transits 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 $\boldsymbol{v}$ be the unique solution to the system of equations (1), then the expected revenue achieved by assortment S can be described as:

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

We discuss in Section 2.2 assumptions that guarantee that the system of equations (1) indeed has a unique solution.

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

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

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

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

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) *correctly computes the optimal value to* (Assort MC).

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

Uncertainty sets and robust assortment optimization. In practice, the parameters are estimated from data and estimation is prone to error. To account for this uncertainty in the parameters, we assume that the parameters belong to some uncertainty sets rather than being fixed. Let $\mathcal{U}^\rho$ and $\mathcal{U}^\lambda$ be uncertainty sets (possibly nonconvex) that the model parameters $\boldsymbol{\rho}$ and $\boldsymbol{\lambda}$ are adversarially selected from. The robust assortment optimization problem under the Markov chain model can be expressed as

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

In other words, the robust assortment optimization problem consists of choosing an assortment that maximizes the worst-case expected revenue over the uncertainty sets $\mathcal{U}^\rho$ and $\mathcal{U}^\lambda$.

2.2. Min-max result

Our main result is to show that, under some assumptions, there exists a min-max duality relation for (Robust Assort MC). We begin by stating an assumption that we need on the structure of the uncertainty set $\mathcal{U}^\rho$.

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

In other words, for each product $i \in \mathcal{N}$, the transition probabilities $(\rho_{ij})_{j\in\mathcal{N}}$ belong to an uncertainty set $\mathcal{U}^{\rho_i}$. 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 from each other. Note that some papers work with the transpose of the transition matrix. In this case, Assumption 1 is referred to as column-wise uncertainty (Gorissen et al. 2014). We further make the following assumption.

Assumption 2. *We assume that the uncertainty sets $\mathcal{U}^{\rho}$ and $\mathcal{U}^{\lambda}$ satisfy the following:*

1. *For every $\boldsymbol{\rho} \in \mathcal{U}^{\rho}$, $\boldsymbol{\rho}$ is irreducible and has a spectral radius which is stricly less than* 1.
2. *For every $\boldsymbol{\rho} \in \mathcal{U}^{\rho}$ and any $i \in \mathcal{N}$, $\rho_{ii} = 0$.*
3. $\mathcal{U}^{\lambda} \subseteq \mathbb{R}^{n}_{++}$.

Assumption 2 is standard for the Markov chain model (Blanchet et al. 2016) 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 system of linear equations (1) admits a unique solution. We can now state the main result of this section.

Theorem 1. *Under Assumptions 1 and 2,*

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

Moreover, the optimal robust assortment can be characterized as follows,

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

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

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

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 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 to maximize the expected revenue. We would like to mention that the objective function does not satisfy the properties like convexity and concavity with respect to the minimization and maximization part respectively for which such min-max duality, also known as saddle point results, holds. Moreover, our result also comes with a characterization of the optimal robust assortment using the fixed point of the mapping $f(\cdot)$. We exploit this characterization to develop an efficient algorithm to find the optimal assortment in Section 5.1.

We end this section by relating our result to Gorissen et al. (2014) which studies a robust linear conic program with column-wise uncertainty on 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), while as highlighted by the form of $f(\cdot)$ in (2), we have constraints of the form $v_i \geq \min_{\boldsymbol{\rho}} \sum_{j\in\mathcal{N}} \rho_{ij} v_j$. In addition, the uncertainty set in Gorissen et al. (2014) is required to be convex, but our min-max theorem also holds true non-convex uncertainty sets. Therefore, the results in Gorissen et al. (2014) are not applicable to our setting.

3. From the Markov chain model to a general framework

In this section, we prove Theorem 1 by showing a min-max result for a more general class of choice models. In particular, we introduce in Section 3.1 a class of models that encompasses the Markov chain model. We state a more general min-max result in Section 3.2 and prove in Section 3.3 how it implies Theorem 1. We defer the proof of the general min-max result to Section 4.

3.1. A more general model

In order to study the robust assortment optimization problem under the Markov chain model, we present a more general framework that captures the robust assortment optimization problem under a broader class of choice models. In particular, we consider a model that depends on an uncertain parameter $\boldsymbol{u}$ which 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 $\boldsymbol{u} \in \mathcal{U}$, we assume that the *expected revenue* generated by this assortment is given by

$$R^{\mathsf{Gen}}(S, \boldsymbol{u}) = \sum_{i\in\mathcal{N}} \lambda(\boldsymbol{u})_i v_i, \tag{Rev General}$$

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

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

and where the parameters $b(\boldsymbol{u})_i \in \mathbb{R}$, $\lambda(\boldsymbol{u})_i \in \mathbb{R}$ for all $i \in \mathcal{N}$ and $A(\boldsymbol{u})_{ij} \in \mathbb{R}$ for all $i \in \mathcal{N}$ and $j \in \mathcal{N}$ all depend on the uncertain parameter $\boldsymbol{u}$. Note how (Rev General) and (3) generalize (Rev MC) and (1) respectively. Consequently, it should be clear that the proposed framework is more general than the Markov chain model.

More generally, under a general choice model, the expression for the expected revenue function is $R(S) = \sum_{i\in S} r_i \pi(i, S)$ where $\pi(i, S)$ is the choice probability, i.e. the probability that product i

is chosen when the assortment S is offered. Consequently, (Rev General) can be interpreted as an expected revenue where the choice probabilities are implicitly given by a system of linear equations. More precisely, using the Markov chain analogy, for a fixed parameter $\boldsymbol{u}$, it is useful to think of $\lambda(\boldsymbol{u})_i$ as the fraction of customers whose most preferred product is i, i.e., who would purchase product i if it is present in the assortment. Using this clustering of the population, v_i represents the revenue from the customers who want to buy product i. If product i is available in the assortment, v_i is equal to r_i, the revenue of product i. Otherwise, the customers substitute to other products (possibly to the outside option). The revenue generated by these substitutions is captured by the system of linear equations $\sum_{j\in\mathcal{N}} A(\boldsymbol{u})_{ij} v_j = b(\boldsymbol{u})_i$ for all $i \notin S$. Note that in the Markov chain model, the substitution behavior is captured by the transitions in an underlying Markov chain. This general framework allows the substitution patterns to be encoded by any matrix $\boldsymbol{A}(\boldsymbol{u})$. In Section 3.3, we map this general model to the Markov chain model and explicit the dependence on $\boldsymbol{u}$. Note that we also show in Section 5.3 how to map this general model to the MNL model.

We are interested in the following robust assortment optimization problem where the decision maker wants to maximize worst-case revenue.

$$\max_{S\subseteq\mathcal{N}} \min_{\boldsymbol{u}\in\mathcal{U}} R^{\mathsf{Gen}}(S,\boldsymbol{u}). \tag{Robust Assort General}$$

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.

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

We next detail our technical results.

3.2. General min-max result

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

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

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

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

Moreover, for any $\boldsymbol{U} = [\boldsymbol{u}^1, \ldots, \boldsymbol{u}^n] \in \mathcal{U}^n$, *let* $\boldsymbol{A}(\boldsymbol{u})$ *be the matrix whose* i^{th} *row is the* i^{th} *row of* $A(u^i)$ *for* $i = 1, \ldots, n$. *We further assume that* $\boldsymbol{A}(\boldsymbol{u})$ *is irreducible for any* $\boldsymbol{U} = [\boldsymbol{u}^1, \ldots, \boldsymbol{u}^n] \in \mathcal{U}^n$.

Note that we do not assume any structure on $\mathcal{U}$. Rather, Assumption 4 imposes some structure on the constraint matrix $\boldsymbol{A}(\boldsymbol{u})$. In particular, we require not only that $\boldsymbol{A}(\boldsymbol{u})$ is irreducible for any parameter $\boldsymbol{u} \in \mathcal{U}$ but also that $\boldsymbol{A}(\boldsymbol{u})$, which is constructed by picking one row from each of n matrices $\boldsymbol{A}(\boldsymbol{u}^i)$, is irreducible for any family $\boldsymbol{u}^1, \ldots, \boldsymbol{u}^n \in \mathcal{U}$.

Note that, under the above assumptions, the linear system in defining (Rev General) has a unique solution, and thus, $R^{\mathsf{Gen}}(S, \boldsymbol{u})$ is well defined for all $S \subseteq \mathcal{N}$ and $\boldsymbol{u} \in \mathcal{U}$. To proceed with our discussion, we define the following problem:

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

Contrasting this with (Dual Assort MC), it 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. In particular, we need that the optimization problems that we introduced can be reformulated as optimization problems over $\boldsymbol{v}$ alone and that the uncertainty in the parameter $\boldsymbol{u}$ can be captured through the objective function and the constraints for the two related problems. More precisely, the assumptions assumes that the minimization operator over the uncertainty parameter can be pushed around. For the Markov chain model, these assumptions are closely related to Assumption 1.

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

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

Under Assumption 5, the minimization over $\boldsymbol{u} \in \mathcal{U}$ can be done separately over the objective and right-hand side of the constraint for any fixed $\boldsymbol{v}$. In other words, the minimization operand can be pushed into the constraint and objective. We make a similar assumption for (Worst-case Rev).

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

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

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

Theorem 2. *Under Assumptions 3, 4, 5 and 6,*

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

Furthermore, the optimal robust assortment can be characterized as follows,

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

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

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

This result not only establishes the min-max relation for the general framework, it also gives a characterization of the optimal robust assortment. We leverage this to design efficient algorithms in Section 5.

3.3. Proof of Theorem 1

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

Proposition 1. *Under Assumption 1 and 2, both Assumptions 3 and 4 hold.*

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

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

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

and

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

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

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

This means that $(\boldsymbol{v}^*, \boldsymbol{\lambda}^*)$ 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

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

The tuple $(\hat{\boldsymbol{v}}, \hat{\boldsymbol{\rho}}, \hat{\boldsymbol{\lambda}})$ 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. □

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*

$$\begin{aligned}\theta_3 = \min_{\boldsymbol{\lambda}\in\mathcal{U}^\lambda,\boldsymbol{\rho}\in\mathcal{U}^\rho} \quad & \min_{\boldsymbol{v}} \sum_{i\in\mathcal{N}} \lambda_i v_i \\ & s.t.\ v_i = r_i, \forall i \in S, \\ & \qquad v_i = \sum_{j\in\mathcal{N}} \rho_{ij} v_j, \forall i \notin S.\end{aligned} \tag{7}$$

and

$$\begin{aligned}\theta_4 = \min_{v} \quad & \min_{\boldsymbol{\lambda}\in\mathcal{U}^\lambda} \sum_{i\in\mathcal{N}} \lambda_i v_i \\ & s.t.\ v_i = r_i,\ \forall i \in S, \\ & \qquad v_i = \min_{\boldsymbol{\rho}\in\mathcal{U}^\rho} \sum_{j\in\mathcal{N}} \rho_{ij} v_j,\ \forall i \notin S.\end{aligned} \tag{8}$$

Proof. Suppose $(\boldsymbol{v}^*, \boldsymbol{\rho}^*, \boldsymbol{\lambda}^*)$ and $(\hat{\boldsymbol{v}}, \hat{\boldsymbol{\lambda}})$ are the optimal solutions to (7) and (8) respectively. For all $i \notin S$, let $\hat{\rho}_i = \arg\min_{\boldsymbol{\rho}_i \in \mathcal{U}^{\rho_i}} \sum_{j\in\mathcal{N}} \rho_{ij}\hat{v}_j$ and let $\hat{\boldsymbol{\rho}} = [\hat{\rho}_1^\top, \cdots, \hat{\rho}_n^\top]$. 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 $(\hat{\boldsymbol{v}}, \hat{\boldsymbol{\rho}}, \hat{\boldsymbol{\lambda}})$ is feasible for (7) , and thus

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

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 occurs. Consequently, $(\boldsymbol{v}^*, \boldsymbol{\lambda}^*)$ is also feasible for (8), and

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

which completes the proof. □

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 $\boldsymbol{u} \in \mathcal{U}$, let $\boldsymbol{\lambda} := \boldsymbol{\lambda}(\boldsymbol{u}), \boldsymbol{A} := \boldsymbol{A}(\boldsymbol{u}), \boldsymbol{b} := \boldsymbol{b}(\boldsymbol{u})$. Let*

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

and

$$\begin{aligned} \theta_6 = \min_{\boldsymbol{v}} \ & \sum_{i\in\mathcal{N}} \lambda_i v_i \\ s.t.\ & v_i \ge r_i, \forall i \in \mathcal{N}, \\ & \sum_{j\in\mathcal{N}} A_{ij} v_j \ge b_i, \forall i \in \mathcal{N}. \end{aligned} \tag{10}$$

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

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

Proof. We first show that $\theta_5 \geq \theta_6$. By Assumption 4, for any given $i \in \mathcal{N}$, $A_{ii} > 0$. We can therefore write

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

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

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

and $\tilde{b}_i = \frac{1}{A_{ii}} b_i$, and let $\tilde{\boldsymbol{A}} := [\tilde{A}_{ij}]$. Under Assumption 4, $\tilde{\boldsymbol{A}}$ has a spectral radius which is strictly less than 1, and $0 < \sum_{j\in\mathcal{N}} \tilde{A}_{ij} < 1$ for any $i \in \mathcal{N}$. Define the following mapping:

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

and $g(\boldsymbol{v}) = [g_1(\boldsymbol{v}), \dots, g_n(\boldsymbol{v})]^\top$. Note that with this notation, (10) can be equivalently rewritten as

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

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

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

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

$$|\max\{a,b\} - \max\{a,c\}| \leq |b - c| \Rightarrow |\max\{a,b\} - \max\{a,c\}|^2 \leq |b - c|^2,$$

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

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

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

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

$$g(\boldsymbol{v}) \geq g(g(\boldsymbol{v})) \geq \ldots \geq g^{k-1}(\boldsymbol{v}) \geq g^k(\boldsymbol{v}).$$

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

$$\boldsymbol{v} \geq g(\boldsymbol{v}) \geq g(g(\boldsymbol{v})) \geq \ldots \geq \boldsymbol{v}^*.$$

Therefore, the above fact together with Assumption 3 implies that any feasible solution $\boldsymbol{v} \neq \boldsymbol{v}^*$ of (11) has a larger objective value than that of $\boldsymbol{v}^*$. Consequently, $\boldsymbol{v}^*$ 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 which 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 the following, we want 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 will imply that $\theta_6 \geq \theta_5$ and complete the proof. From the definition of $g(\cdot)$ and the construction of $\boldsymbol{v}$, we have $g(\boldsymbol{v}) \geq \boldsymbol{v}$. Moreover, since g is monotone increasing, it holds that

$$g^0(\boldsymbol{v}) := \boldsymbol{v} \leq g(\boldsymbol{v}) \leq g(g(\boldsymbol{v})) \leq \ldots \leq g^{k-1}(\boldsymbol{v}) \leq g^k(\boldsymbol{v}),$$

i.e., the sequence $\{g^k(\boldsymbol{v})\}_{k=0,1,\ldots}$ is monotonically increasing. Next, we show that $\{g^k(\boldsymbol{v})\}_{k=0,1,\ldots}$ 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 $i \in \mathcal{N}$,

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

where

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

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

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

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

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

- The maximizer of

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

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

- The minimizer of

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

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

In this case, Proposition 4 fails because the mapping $g(\cdot)$ in Proposition 4 is no longer monotone increasing. Moreover, starting from $\boldsymbol{v}_1^*$, $\lim\limits_{k \to \infty} g^k(\boldsymbol{v}_1^*) = [0.329, 0.463]^\top$ for which the objective is 0.343.

4.2. Proof of Theorem 2

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

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

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

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

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

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

We can then rewrite (13) succinctly as

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

In fact, we can show that the optimal solution $\boldsymbol{v}^*$ of the above problem satisfies $\boldsymbol{v}^* = f(\boldsymbol{v}^*)$. To see this, we first show that f is monotonically increasing, i.e., $f(\boldsymbol{v}') \geq f(\boldsymbol{v}'')$ for any $\boldsymbol{v}' \geq \boldsymbol{v}''$ (the inequality is true componentwise). For any $\boldsymbol{u}\in\mathcal{U}$, Assumption 4 ensures that $\frac{-A(\boldsymbol{u})_{ij}}{A(\boldsymbol{u})_{ii}} \geq 0$ for all $i,j\in\mathcal{N}$. Therefore for any $\boldsymbol{v}'\geq\boldsymbol{v}''$, denoting $\boldsymbol{u}' = \arg\min_{\boldsymbol{u}\in\mathcal{U}}\left[\sum_{j\neq i} \frac{-A(\boldsymbol{u})_{ij}}{A(\boldsymbol{u})_{ii}} v'_j + \frac{1}{A(\boldsymbol{u})_{ii}} b(\boldsymbol{u})_i\right]$, we have for all $i\in\mathcal{N}$,

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

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

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

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

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

By contradiction, assume there exists an optimal solution $\boldsymbol{v}^*$ of (14) such that $v_i^* > f_i(\boldsymbol{v}^*)$ for some i. Since f is monotonically increasing and $\lambda(\boldsymbol{u})_i > 0$ for any $i \in \mathcal{N}$ and $\boldsymbol{u} \in \mathcal{U}$ 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 $\boldsymbol{v}^*$. 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 $\min_{\boldsymbol{u}\in\mathcal{U}} R(S,\boldsymbol{u})$ is

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

Therefore, according to Assumption 6, it can further be written as

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

Before proceeding, we state two claims which are needed to complete the proof. Their proofs are postponed to Appendix A and B.

Claim 1. *For any given $S \subseteq \mathcal{N}$, the 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 to $\max_{S\in\mathcal{N}} \min_{\boldsymbol{u}\in\mathcal{U}} R(S,\boldsymbol{u})$ with $\boldsymbol{v}^*$ being the corresponding optimum of the inner problem (16). By feasibility of $\boldsymbol{v}^*$, we have

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

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

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

On the other hand, take any $\boldsymbol{u}_0 \in \mathcal{U}$, then

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

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

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

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

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

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

Remark 2. According to the proof of Theorem 2, the optimal $(S^*, \boldsymbol{v}^*)$ only depends on $\mathcal{U}^\rho$ but not on $\mathcal{U}^\lambda$. The latter affects the value of the objective function but not the optimal assortment.

5. Implications

5.1. Algorithm for robust Markov chain model

In order to prove the min-max result in Theorem 1, we reformulated the robust assortment optimization problem as a fixed point problem. We present an iterative algorithm that builds on this reformulation and aims to compute the unique fixed point of the following mapping

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

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

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,

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

for some constant α.

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

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

Output: The optimal assortment S^*

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

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

Proof. In the proof of Theorem 2, we have proved that $f(\cdot)$ is monotonically increasing and bounded above. Therefore, $\{f^t(\boldsymbol{v}^0)\}_{t=1,2,\ldots}$ converges to $\boldsymbol{v}^*$ for any starting point $\boldsymbol{v}^0$. We now prove that the algorithm terminates in polynomially many steps. Observe that in Algorithm 1, since $\boldsymbol{v}^t$ is increasing, once v_i^t exceeds r_i for some for some $i \in \mathcal{N}$, it never goes back to r_i again. Denote

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

Observe that in each iteration, there is a probability of at least δ of being absorbed by state 0. Therefore, after t steps, the maximum possible expected revenue for i^0 is $(1-\delta)^t r_{\max}$. 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.

Consequently, Algorithm 1 converges in polynomially many steps to the fixed point $\boldsymbol{v}^*$ of $f(\cdot)$. By Theorem 1, this implies that S^* is an optimal solution to (Robust Assort MC). □

5.2. Operational insights of robust Markov chain model

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

Recall that we still work under Assumption 1, that is the uncertainty set can be represented as $(\mathcal{U}^\lambda, \mathcal{U}^\rho)$ where $\mathcal{U}^\rho = \mathcal{U}^{\rho_1} \times \cdots \times \mathcal{U}^{\rho_n}$. Let $S^*(\mathcal{U}^\lambda, \mathcal{U}^\rho)$ denote the largest robust optimal assortment and $Z^*(\mathcal{U}^\lambda, \mathcal{U}^\rho)$ the corresponding objective value for given uncertainty sets $\mathcal{U}^\lambda$ and $\mathcal{U}^\rho$. 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* $\mathcal{U}^\rho \subseteq \hat{\mathcal{U}}^\rho$,

$$Z^*(\hat{\mathcal{U}}^\lambda,\hat{\mathcal{U}}^\rho) \le Z^*(\mathcal{U}^\lambda,\mathcal{U}^\rho) \quad \text{and} \quad S^*(\mathcal{U}^\lambda,\mathcal{U}^\rho) \subseteq S^*(\hat{\mathcal{U}}^\lambda,\hat{\mathcal{U}}^\rho).$$

Proposition 6, whose proof is presented in Appendix C, states that when the degree of uncertainty increases, the optimal worse case revenue will decrease. Moreover, the decision maker should offer larger assortments to protect against larger uncertainty. Increasing product variety helps hedging against a larder uncertainty in the parameters. We next provide an alternative characterization of $S^*(\mathcal{U}^\lambda,\mathcal{U}^\rho)$ which 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 the robust optimal assortment is the largest optimal assortment among $\{S^*_\rho : \boldsymbol{\rho} \in \mathcal{U}^\rho\}$.

Proposition 7. *For any* $(\mathcal{U}^\lambda,\mathcal{U}^\rho)$, $S^*(\mathcal{U}^\lambda,\mathcal{U}^\rho) = \bigcup_{\boldsymbol{\rho}\in\mathcal{U}^\rho} S^*_\rho$.

Proof If 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 $\min_{\boldsymbol{\rho}\in\mathcal{U}^\rho,\boldsymbol{\lambda}\in\mathcal{U}^\lambda} R^{\mathrm{MC}}(S^*(\mathcal{U}^\lambda,\mathcal{U}^\rho),\rho,\lambda)$ and let $S^*_{\rho^*} = \arg\max_{S\subseteq\mathcal{N}} R^{\mathrm{MC}}(S,\boldsymbol{\rho}^*,\boldsymbol{\lambda}^*)$. By Theorem 1, we have

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

Due to the uniqueness of the largest optimal assortment, we have $S^*_{\rho^*} = S^*(\mathcal{U}^\lambda,\mathcal{U}^\rho)$. This completes the proof. □

As a consequence of Proposition 7, the decision marker should focus on the customer types with transition probabilities that lead to the largest optimal assortment to protect against worst case scenario. Finally, we present a result showing that the robust optimal assortment shrinks as we decrease the product revenues. This type of result is particularly helpful in the single resource revenue management problem. Indeed, in each time period, the decision maker needs to solve 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 helpful is characterizing the structure of the optimal policy for this problem.

Proposition 8. *For any* $(\mathcal{U}^\lambda, \mathcal{U}^\rho)$, *let* $S^*_\eta(\mathcal{U}^\lambda,\mathcal{U}^\rho)$ *be the optimal robust assortment for the revenues* $\mathbf{r}^\eta$ *where* $r^\eta_i = r_i + \eta$, *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 MNL model: general uncertainty set

5.3.1. Recovering the MNL choice model Let $\boldsymbol{p} = [p_0, p_1, \ldots, p_n]$ be the MNL parameters such that $\sum_{i=0}^{n} p_i = 1$. Blanchet et al. (2016) show that the MNL model is a special case of the Markov chain model. Consequently, we can write the expected revenue under the MNL model as the following system of linear equations.

$$R^{\mathsf{MNL}}(S, \boldsymbol{p}) = \sum_{i \in \mathcal{N}} p_i v_i, \tag{Rev MNL}$$

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 $\boldsymbol{u} = \boldsymbol{p}$, $\boldsymbol{\lambda}(\boldsymbol{u}) = \boldsymbol{p}$, $\boldsymbol{b}(\boldsymbol{u}) = 0$ and

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

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

$$\max_{S \subseteq \mathcal{N}} \min_{\boldsymbol{p} \in \mathcal{U}^p} R^{\mathsf{MNL}}(S, p). \tag{Robust Assort MNL}$$

In addition, we make the following regularization assumptions.

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

Assumption 7 directly implies Assumptions 3 and Assumption 4. Moreover,

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

is equivalent to

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

since the terms $\min_{\boldsymbol{p} \in \mathcal{U}^p} \sum_{i \in \mathcal{N}} p_i v_i$ in the objective and 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 Assumptions 7,*

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

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

5.3.2. Algorithm to find 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 some

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

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

Output: The optimal assortment S^*

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

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

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

Note that for every t, S^t is therefore a revenue ordered assortment, i.e. consists of the highest k revenues products for some k. Since γ^t is an increasing sequence, $\{S^t\}_t$ is a sequence of revenue ordered assortment such that $S^{t+1}\subseteq S^t$. This implies that S^* as a limit point 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 the robust approaches to assortment optimization. In particular, in Section 6.1, we explore a setting where the underlying ground truth is a known Markov chain model with uncertain parameters. We use this ideal setting to showcase the running time of our algorithm as well as the magnitude of the trade-off between the expected and worst-case revenue when comparing to a deterministic approach which does not account for parameters uncertainty. We then present in Section 6.2 a more realistic set of experiments where the ground truth is an unknown ranking-based model. In this case, we assume that we only have access to choice data and propose a data-driven approach where we first learn the parameters of a Markov chain model and then construct some uncertainty sets. We illustrate the benefits of adopting a robust approach in this case as well.

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. Known ground truth

In this section, we present a numerical study where the underlying ground truth is a known Markov chain model with uncertain parameters. Note that this is the ideal setting that our theory has been developed for. We illustrate the trade-off between expected and worst-case performance induced by our approach and 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 modal Markov chain by generating $(n+1)^2$ independent random variables X_{ij}, each picked uniformly over the interval $[0,1]$. We then set $\rho_{ij}^{\mathsf{modal}} = X_{ij}/\sum_{k=0}^{n} X_{ik}$ for all $i,j \geq 1$ such that $i \neq j$. 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 $\lambda_i^{\mathsf{modal}} = Y_i/\sum_{j=1}^{n} Y_j$ for all $i \neq 0$. For $\epsilon > 0$, we define a row-wise uncertainty set as follows. For each i, let

$$\mathcal{U}_{\epsilon}^{\rho_i} = \left\{ \boldsymbol{\rho}_i^{\mathsf{modal}} + \boldsymbol{\gamma}_i \;\middle|\; \sum_j \gamma_{ij} = 0, \text{ and } \forall j, \max\{(1-\epsilon)\rho_{ij}^{\mathsf{modal}}, 0\} \leq \rho_{ij}^{\mathsf{modal}} + \gamma_{ij} \leq \min\{(1+\epsilon)\rho_{ij}^{\mathsf{modal}}, 1\} \right\}. \tag{17}$$

In other words, the uncertainty set we consider is centered around $\boldsymbol{\rho}^{\mathsf{modal}}$ and the magnitude of variations around $\boldsymbol{\rho}^{\mathsf{modal}}$ are 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 $\boldsymbol{\rho}^{\mathsf{modal}}$. In a way, this is the assortment that one would compute if not accounting for robustness and taking the average parameters as the real parameters. The second one, $S_{\epsilon}^{\mathsf{robust}}$, is the optimal robust assortment computed using Algorithm 1, i.e. the assortment which maximizes the worst-case expected revenue over the uncertainty set $\mathcal{U}_{\epsilon} = \mathcal{U}_{\epsilon}^{\rho_1} \times \cdots \times \mathcal{U}_{\epsilon}^{\rho_n}$. For each assortment

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

	$n=20$				$n=50$			
	Δ_{modal}		Δ_{worst}		Δ_{modal}		Δ_{worst}	
ϵ	Average	Minimum	Average	Minimum	Average	Minimum	Average	Minimum
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

and level of uncertainty ϵ, we compare the expected revenue under the modal parameters, $R^{\mathsf{modal}}(\cdot)$, and the worst-case expected revenue, $R^{\mathsf{worst}}_{\epsilon}(\cdot)$. For the latter, we use (Worst-case Rev) to compute the worst-case expected revenue generated by a given assortment over $\mathcal{U}_{\epsilon}$.

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

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

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 protects against the uncertainty in parameters. In Table 1, we report the average and minimum value of these ratios over 100 randomly generated instances. First, note that all the values are less than one for Δ_{modal} and more than one for Δ_{worst}. This is not surprising as S^{modal} maximizes $R^{\mathsf{modal}}(\cdot)$ and $S^{\mathsf{robust}}_{\epsilon}$ maximizes $R^{\mathsf{worst}}_{\epsilon}(\cdot)$ respectively. Moreover, as ϵ increases, the uncertainty set we protect against increases. Therefore, Δ_{modal} decreases while Δ_{worst} increases. In terms of magnitude of the improvement, we observe that there is almost a linear trade-off between expected and worst-case expected revenue on the average performance of all instances. For example, when $\epsilon = 0.50$ and $n = 50$, the robust approach captures on average 98.61% of the modal assortment expected revenue under the modal parameters while the expected revenue is 2.40% more under the worst-case parameters. 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

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

	Running time (s)		Number of iterations	
n	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

solved in around 20s on average when $n = 50$. In practice, the parameters of the Markov chain model are not known and need to be learned from data. We present in the next section a more realistic set of numerical experiments where we learn the parameters of the Markov chain model from data and compare a robust approach to a deterministic one that does not account for any parameters uncertainty.

6.2. Unknown ground truth

In this section, we present a more realistic numerical study to show the benefits of using a robust approach. We begin by describing the setting and then the corresponding results.

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

The ground truth choice model. We adopt a ranking-based choice model as the 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 $g = 1, \ldots, m$. Each list $\sigma^g = (\sigma^g_1, \cdots, \sigma^g_{n+1})$ is an ordering of the products in $\mathcal{N}^+$ where $\sigma^g_i < \sigma^g_j$ indicates that product σ^g_i is preferred to product σ^g_j under preference list σ^g. We denote by β^g the probability that an arriving customer chooses the ranked list σ^g for $g = 1, \ldots, m$. 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 $g = 1, \ldots, m$, 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 $\sum_{g=1}^m \beta^g \cdot \mathbf{1}\{i = \arg\min_{j \in S} \sigma^g_j\}$. We set $m = 2n$ in our experiments. Finally, we assume that for each product, 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 $\{(S^t, \boldsymbol{Z}^t) : t = 1, \ldots, T\}$. More precisely, for each customer t, S^t denotes the offered assortment and $\boldsymbol{Z}^t = (Z_1^t, \ldots, Z_n^t)$ 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 which does not account for parameter uncertainty. In this approach, we apply the expectation-maximization (EM) algorithm from Şimşek and Topaloglu (2018) to the historical purchase data $(S^t, \boldsymbol{Z}^t)_t$ to compute $\boldsymbol{\rho}^{\mathsf{modal}}$ and $\boldsymbol{\lambda}^{\mathsf{modal}}$, which are the estimated arrival probability vector and transition probability matrix of the Markov chain choice model. We then assume that these parameters are the correct ones, and compute the offer set S^{modal} by solving the corresponding 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 construct the uncertainty sets inspired by the boostrap 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 to construct a row-wise uncertainty set for the Markov chain model is given in Algorithm 3. More specifically, we use the estimated arrival probability vector $\boldsymbol{\lambda}^{\mathsf{modal}}$ and transition probability matrix $\boldsymbol{\rho}^{\mathsf{modal}}$ from the deterministic approach as a ground truth model to generate K new sets of purchase data. With each newly generated purchase data, we use the EM algorithm again to get a new set of estimated parameter $\boldsymbol{\rho}^{(k)}$. The variations in the estimated coefficient $\boldsymbol{\rho}^{(k)}$ from the boostrap 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 $\rho_{ij}^{(k)}/\rho_{ij}^{\mathsf{modal}}$ to inform the uncertainty we allow around $\rho_{ij}^{\mathsf{modal}}$. As a result, the uncertainty set returned by Algorithm 3 is centered around $\boldsymbol{\rho}^{\mathsf{modal}}$ and thus is similar to the previous section. We also scale the uncertainty set uniformly by a parameter $\alpha \in [0, 1]$ to control the robustness level and observe the effects of introducing more or less uncertainty. Note that we do not construct uncertainty set on the arrival probability $\boldsymbol{\lambda}$, as the optimal assortment is independent of the uncertainty of $\boldsymbol{\lambda}$. For each robustness level α, let $S_\alpha^{\mathsf{robust}}$ be the optimal robust assortment computed using Algorithm 1, i.e. the assortment which maximizes the worst-case expected revenue over the uncertainty set $\mathcal{U}_\alpha^\rho$ constructed in Algorithm 3.

Algorithm 3 Construct Uncertainty Set of Markov Chain Model Based on Bootstrap

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

Output: The uncertainty set $\mathcal{U}_\alpha^\rho$.

1: Apply EM algorithm to $\{(S^t, \boldsymbol{Z}^t(S^t))\}$ and compute the estimated arrival probability $\boldsymbol{\lambda}^{\mathsf{modal}}$ and transition probability matrix $\boldsymbol{\rho}^{\mathsf{modal}}$.

2: **for** $k = 1, 2, \ldots, K$ **do**

3: (**Resampling**) Independently draw new purchase data $\{(S^t, \hat{\boldsymbol{Z}}^t(S^t))\}$ with the ground truth being a Markov chain model with parameters $\boldsymbol{\lambda}^{\mathsf{modal}}$, $\boldsymbol{\rho}^{\mathsf{modal}}$.

4: Apply EM algorithm to $\{(S^t, \hat{\boldsymbol{Z}}^{k,t}(S^t))\}$ and get another estimator $\boldsymbol{\rho}^{(k)}$.

5: (**Constructing uncertainty set**)For $i = 1, \ldots, n$, compute

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

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

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

6.2.2. Results and discussion. In this setting, the ground truth model is not known and does not come with any uncertainty in the parameters. However, we show that our robust approach can help address two sources of potential errors. First, there are estimation errors coming from the potentially insufficient amount of data. Second, there may be mis-specification errors since we are fitting a Markov chain model whereas the ground truth is not.

Let $R^{\mathsf{true}}(\cdot)$ 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 $\mathsf{m} \in \{\mathsf{modal}, \mathsf{robust}\}$ and $\alpha \in \{0.05, 0.10, 0.25, 0.50\}$, we compute

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

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 $\pi^{\mathsf{true}}(i, S)$ as the choice probability of product i when the offer set is S under the ground truth model. For $\mathsf{m} \in \{\mathsf{modal}, \mathsf{robust}\}$, let $p_{\mathsf{purch}} = 1 - \pi^{\mathsf{true}}(0, S^m)$. We use this probability as a proxy for the robustness of the solution. Indeed, if a customer does not purchase any product, this yields

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

			$n = 10$				$n = 20$			
T	m	α	$\Delta^{\mathsf{m}}_{\mathsf{expected}}$	$\Delta^{\mathsf{m}}_{\mathsf{5th}}$	$\Delta^{\mathsf{m}}_{\mathsf{10th}}$	p_{purch}	$\Delta^{\mathsf{m}}_{\mathsf{expected}}$	$\Delta^{\mathsf{m}}_{\mathsf{5th}}$	$\Delta^{\mathsf{m}}_{\mathsf{10th}}$	p_{purch}
1,000	modal	-	0.9692	0.0862	0.1523	0.8279	0.9646	0.0741	0.2350	0.8714
	robust	0.05	0.9698	0.0869	0.1563	0.8290	0.9645	0.0595	0.2406	0.8741
		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
2,500	modal	-	0.9729	0.1060	0.2350	0.8512	0.9642	0.0560	0.2811	0.8776
	robust	0.05	0.9729	0.1208	0.2427	0.8512	0.9643	0.0555	0.2650	0.8780
		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
5,000	modal	-	0.9681	0.0505	0.0968	0.8090	0.9661	0.0357	0.1678	0.8555
	robust	0.05	0.9678	0.0505	0.0968	0.8090	0.9660	0.0442	0.1971	0.8567
		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

zero revenue. A higher probability of purchase therefore implies a better protection 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 $X^{\mathsf{true}}(S^m)$ be a random variable such that $X^{\mathsf{true}}(S^m) = r_i$ with probability $\pi^{\mathsf{true}}(i, S^m)$. Then, for $\mathsf{m} \in \{\mathsf{modal}, \mathsf{robust}\}$, we report $R^{\mathsf{true}}_{q^{th}}(S^m)$ defined such that $P(X^{\mathsf{true}}(S^m) \leq R^{\mathsf{true}}_{q^{th}}(S^m)) = q$. For a given $q \in \{5, 10\}$, a higher value of $R^{\mathsf{true}}_{q^{th}}(S^m)$ means that the assortment is more robust since it guarantees that with probability $1 - q$, the realized revenue is above $R^{\mathsf{true}}_{q^{th}}(S^m)$.

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 $T \in \{1,000,\ 2,500,\ 5,000\}$. We observe that even in this more realistic setting, the robust approach offers a nice trade-off between expected revenue and worst-case expected revenue. First, in terms of expected revenue, as α increases and we are assuming a larger uncertainty set when computing $S^{\mathsf{robust}}_{\alpha}$, the variations in $\Delta^{\mathsf{robust}}_{\mathsf{expected}}$ are very mild. Even when $\Delta^{\mathsf{robust}}_{\mathsf{expected}}$ decreases, the loss compared to the deterministic assortment S^{modal} is quite limited. For instance, for $T = 5,000$ 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 $S^{\mathsf{robust}}_{\alpha}$ 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 $T = 1,000$ 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 $S^{\mathsf{robust}}_{\alpha}$ with $\alpha = 0.10$

captures 97.15 % of the expected revenue. It appears that given the the lack of data and the model mis-specification, adding some robustness can help hedge against the case where the estimated parameters is far from the ground truth.

We next compare how the two approaches perform in terms of robustness. Both the 5^{th} and 10^{th} 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 $T = 2,500$, $\alpha = 0.50$ and $n = 20$, the robust approach expected revenue is very close to the deterministic one. More precisely, $\Delta^{\mathsf{robust}}_{\mathsf{expected}}/\Delta^{\mathsf{modal}}_{\mathsf{expected}} = 0.999$. On the other hand, the average 5^{th} 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 on the expected revenue and robustness. This suggests that in more realistic settings, accounting for some uncertainty in the parameters can lead to higher average and worst-case performance! 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 model. Under reasonable assumptions, mainly that the uncertainty 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. Inspired by the duality results, we also develop efficient iterative algorithms to find the optimal robust assortment.

To prove our main result, we introduced a general framework for choice models assumes that the choice probabilities, capturing the substitution behavior of customers, are solutions to a system of linear equations. It would be interesting to see if that 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 type of uncertainty sets for the Markov chain model, in particular allowing the uncertainty across different rows of the Makov chain model to be related. Having a budget constraint across rows to limit the adversary power would be one way to correlates the uncertainty sets.

References

Aouad, Ali, Vivek F Farias, Retsef Levi. 2015. Assortment optimization under consider-then-choose choice models .

Aouad, Ali, Vivek F Farias, Retsef Levi, Danny Segev. 2018. The approximability of assortment optimization under ranking preferences. *Operations Research* **66**(6) 661–1669.

Ben-Akiva, Moshe E, Steven R Lerman, Steven R Lerman. 1985. *Discrete choice analysis: theory and application to travel demand*, vol. 9. 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* .

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* 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* .

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, D., S. onnalagedda, X.A. Pan. 2012. Optimal algorithms for assortment selection under ranking-based consumer choice models. *Manufacturing & Service Operations Management* **14**(2) 279–289.

Honhon, Dorothée, Vishal Gaur, Sridhar Seshadri. 2010. Assortment planning and inventory decisions under stockout-based substitution. *Operations research* **58**(5) 1364–1379.

Jagabathula, Srikanth, Paat Rusmevichientong. 2016. A nonparametric joint assortment and price choice model. *Management Science* .

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, S., G. 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* 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, G., G. 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, G., S. Burer. 2016. A copositive approach for two-stage adjustable robust optimization with uncertain right-hand sides. *arXiv:1609.07402* .

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 Model

Appendix A: Proof of Claim 1

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

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

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 $\boldsymbol{v}^1$ and $\boldsymbol{v}^2$. For $i \notin S$, $h_i(\boldsymbol{v})$ is the minimum of many affine functions and is therefore concave. Consequently, by letting

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

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

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

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

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

Consider the matrix $\tilde{\boldsymbol{A}}(\tilde{\boldsymbol{U}})$ with $\tilde{\boldsymbol{U}} = [\tilde{\boldsymbol{u}}^1, \cdots, \tilde{\boldsymbol{u}}^n]$ and its row i equals to $\tilde{\boldsymbol{A}}(\tilde{\boldsymbol{u}}^i)_{i:}$. Assumption 4 implies that all components of $\tilde{\boldsymbol{A}}(\tilde{\boldsymbol{U}})$ are nonnegative. Then, for all $i \notin S$,

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

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

Assumption 4 also implies that of $\tilde{\boldsymbol{A}}(\tilde{\boldsymbol{U}})$ is irreducible as the off-diagonal elements of $\tilde{\boldsymbol{A}}(\tilde{\boldsymbol{U}})$ and $\boldsymbol{A}(\tilde{\boldsymbol{U}})$ considered in Assumption 4 are simultaneously to be zeros or nonzeros, and thus its spectral radius is less than 1. Let z be the left eigenvector of $\tilde{\boldsymbol{A}}(\tilde{\boldsymbol{U}})$ associated with the largest absolute eigenvalue τ. It follows from Perron-Frobenius Theorem that $\boldsymbol{z} > 0$ and $0 < \tau < 1$. As a result,

$$\begin{aligned} \sum_{i\in\mathcal{N}} z_i |v_i^1 - v_i^2| &= \sum_{i\in\mathcal{N}} z_i |h_i(\boldsymbol{v}^1) - h_i(\boldsymbol{v}^2)| \\ &\leq \sum_{i\in\mathcal{N}} z_i \sum_{j\in\mathcal{N}} \tilde{A}(\tilde{\boldsymbol{u}}^i)_{ij} |v_j^1 - v_j^2| \end{aligned}$$

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

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

Appendix B: Proof of Claim 2

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

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

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

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

we have that

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

This is a contradiction and concludes the proof. □

Appendix C: Proof of Proposition 6

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

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

Let $f(\cdot)$ be the mapping associated with $\mathcal{U}^\rho$,

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

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

$$
\hat{\boldsymbol{v}} \leq f(\hat{\boldsymbol{v}}) \leq f^2(\hat{\boldsymbol{v}}) \leq \ldots \leq f^d(\hat{\boldsymbol{v}}) \to \boldsymbol{v}^*.
$$

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

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

Additionally,

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

Consequently, it follows from Theorem 1 that

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

This conclues the proof. □

Appendix D: Proof of Proposition 8

Let $\boldsymbol{v}^\eta$ be the fixed point associated with $S_\eta^*(\mathcal{U}^\lambda, \mathcal{U}^\rho)$. We define $\tilde{\boldsymbol{v}}$ 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

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

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

$$f_i(\boldsymbol{v}) = \max\left\{ r_i, \min_{\boldsymbol{\rho}_i \in \mathcal{U}^{\rho_i}} \sum_{i\in\mathcal{N}} \rho_{ij} v_j \right\}, \; i \in \mathcal{N}.$$

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

$$\tilde{\boldsymbol{v}} \leq f(\tilde{\boldsymbol{v}}) \leq f^2(\tilde{\boldsymbol{v}}) \leq \ldots \leq f^d(\tilde{\boldsymbol{v}}) \to \boldsymbol{v}^*.$$

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

Appendix E: MNL benchmark

For the experiments described in Section 6.2, we also test an approach that would use an MNL model. Similar to the Markov chain approach, we compute a deterministic MNL and a robust MNL solution. For the deterministic approach, we use a standard EM algorithm (see for instance Talluri and Van Ryzin (2004)) to estimate the parameter of the MNL model $\boldsymbol{p}^{\mathsf{modal}}$. 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 boostrap method described in Algorithm 4. More specifically, we use the estimated parameter $\boldsymbol{p}^{\mathsf{modal}}$ 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 purchase data, we use the EM algorithm again to get a new set of estimated parameter $\boldsymbol{p}^{(k)}$. Then, we use

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

			$n=10$				$n=20$			
T	m	α	$\Delta^{\mathsf{m}}_{\mathsf{expected}}$	$\Delta^{\mathsf{m}}_{\mathsf{5th}}$	$\Delta^{\mathsf{m}}_{\mathsf{10th}}$	p_{purch}	$\Delta^{\mathsf{m}}_{\mathsf{expected}}$	$\Delta^{\mathsf{m}}_{\mathsf{5th}}$	$\Delta^{\mathsf{m}}_{\mathsf{10th}}$	p_{purch}
1,000	modal	-	0.9546	0.0641	0.1384	0.8183	0.9536	0.0079	0.1542	0.8423
	robust	0.05	0.9539	0.0641	0.1384	0.8186	0.9536	0.0166	0.1542	0.8423
		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
2,500	modal	-	0.9500	0.0469	0.1425	0.8180	0.9555	0.0400	0.1828	0.8560
	robust	0.05	0.9500	0.0469	0.1427	0.8180	0.9555	0.0400	0.1828	0.8560
		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
5,000	modal	-	0.9550	0.0399	0.0939	0.7943	0.9562	0.0227	0.1155	0.8410
	robust	0.05	0.9550	0.0312	0.0937	0.7943	0.9561	0.0227	0.1239	0.8410
		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

the magnitude of the ratio $p_j^{(k)}/p_j^{\mathsf{modal}}$ to inform the uncertainty we allow around p_j^{modal}. We also scale the uncertainty set uniformly by a parameter $\alpha \in [0,1]$ 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 $S_\alpha^{\mathsf{robust}}$ be the optimal robust assortment computed using Algorithm 2, i.e. the assortment which maximizes the worst-case expected revenue over the uncertainty set $\mathcal{U}_\alpha^p$. Table EC.1 reports the same metrics when using an MNL model instead of a Markov chain 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 on the expected performance metric $\Delta^{\mathsf{modal}}_{\mathsf{expected}}$. 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 dominates the MNL approach on the expected revenue metric, it also dominates the MNL approach on the robustness metrics. More precisely, for all uncertainty levels α, the 5^{th} percentile metric and the purchase probability are higher with 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 the robustness of the solution.

Algorithm 4 Construct Uncertainty Set of MNL Model Based on Bootstrap

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

Output: Uncertainty set $\mathcal{U}_\epsilon^v$

1: Apply EM algorithm to $\{(S^t, \boldsymbol{Z}^t(S^t))\}$ and get an estimated probability $\boldsymbol{p}^{\mathsf{modal}}$ of the MNL model.

2: **for** $k = 1, 2, \ldots, K$ **do**

3: (**Resampling**) Independently draw new purchase data $\{(S^t, \hat{\boldsymbol{Z}}^t(S^t))\}$ with the ground truth being a MNL model with parameters $\boldsymbol{p}^{\mathsf{modal}}$.

4: Apply EM algorithm to $\{(S^t, \hat{\boldsymbol{Z}}^t(S^t))\}$ and get another estimator $\boldsymbol{p}^{(k)}$.

5: (**Constructing the uncertainty set**) Compute

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

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