http://pubsonline.informs.org/journal/ijds

OPERATIONS RESEARCH
Vol. 00, No. 0, Xxxxx 0000, pp. 000–000
ISSN 2694-4022, EISSN 2694-4030

Submitted to *Operations Research*

Randomized Assortment Optimization

Zhengchao Wang,[a] Heikki Peura,[b] Wolfram Wiesemann[a]
[a]Imperial College Business School, Imperial College London, London, UK, {z.wang19@imperial.ac.uk, ww@imperial.ac.uk}, [b]Aalto University School of Business, Espoo, Finland, heikki.peura@aalto.fi

Abstract. When a firm selects an assortment of products to offer to customers, it uses a choice model to anticipate their probability of purchasing each product. In practice, the estimation of these models is subject to statistical errors, which may lead to significantly suboptimal assortment decisions. Recent work has addressed this issue using robust optimization, where the true parameter values are assumed unknown and the firm chooses an assortment that maximizes its worst-case expected revenues over an uncertainty set of likely parameter values, thus mitigating estimation errors. In this paper, we introduce the concept of *randomization* into the robust assortment optimization literature. We show that the standard approach of deterministically selecting a single assortment to offer is not always optimal in the robust assortment optimization problem. Instead, the firm can improve its worst-case expected revenues by selecting an assortment randomly according to a prudently designed probability distribution. We demonstrate this potential benefit of randomization both theoretically in an abstract problem formulation as well as empirically across three popular choice models: the multinomial logit model, the Markov chain model, and the preference ranking model. We show how an optimal randomization strategy can be determined exactly and heuristically. Besides the superior in-sample performance of randomized assortments, we demonstrate improved out-of-sample performance in a data-driven setting that combines estimation with optimization. Our results suggest that more general versions of the assortment optimization problem—incorporating business constraints, more flexible choice models and/or more general uncertainty sets—tend to be more receptive to the benefits of randomization.

Key words: Assortment optimization; randomization; robust optimization; choice model

1. Introduction

Selecting an assortment of products to offer to customers is a central problem across business operations, with many applications in the travel, hospitality, and retail industries. A firm solving this *assortment optimization problem* seeks to maximize expected revenues or a related objective by selecting a subset of possible products to carry, often subject to constraints such as shelf or display space. With the growth of novel online marketing and e-commerce settings where firms can quickly and flexibly adjust displayed assortments, this classic operations problem is attracting increasing research interest.

The difficulty in this problem lies in accounting for customer behavior: If a prospective customer's preferred product is not part of the offered assortment, she may either substitute it for another one or leave with no purchase. In order to capture this demand substitution, the firm must specify and estimate a customer *choice model*, which describes a customer's probability of choosing each product for any assortment the firm could offer. Since data on customer behavior in the face of the numerous possible assortments is invariably sparse, the estimation of such models is difficult, and estimation errors can lead to significantly suboptimal assortment decisions.

The specification of the choice model involves the familiar bias-variance tradeoff. Most common choice models, such as the multinomial logit (MNL) model (see, *e.g.*, Talluri and van Ryzin 2004), capture customer behavior by imposing strict parametric assumptions. Although this approach results in tractable model estimation and assortment optimization problems that require relatively small amounts of data, imposing restrictive assumptions can introduce a significant bias into estimates (*i.e.*, under-fitting) as well as raise theoretical concerns (such as the independence from irrelevant alternatives, or IIA, property). These pitfalls can be avoided through richer models, such as parametric generalizations of the MNL model, the Markov chain (MC) model (Blanchet et al. 2016), or non-parametric models based on customer preference rankings (Honhon et al. 2012, Bertsimas and Mišić 2019). However, as a result of their added flexibility, more complex models are in turn prone to variance (*i.e.*, over-fitting) unless large amounts of data are available.

The risk of over-fitting the choice model is particularly pernicious when estimates are used to select the best assortment to offer, due to the well-known error-maximization effect of optimization (Smith and Winkler 2006). To address this issue, recent assortment optimization papers have embraced the *robust optimization* paradigm (Rusmevichientong and Topaloglu 2012, Bertsimas and Mišić 2017, Désir et al. 2023). Robust optimization acknowledges the fact that the parameters of the choice model are not known exactly, which is particularly important for complex models with many parameters to be estimated. Instead of point estimates, a decision maker adopting the robust approach specifies an uncertainty set that contains the unknown true parameter values or preferences with a pre-specified confidence. She then chooses an assortment considering the worst parameter setting within this uncertainty set, thus hedging against over-fitting in model estimation.

In this paper, we study *randomized* robust assortment optimization. That is, we show that the standard approach of deterministically selecting a single assortment to offer is not always optimal in the robust assortment optimization problem. Instead, the decision maker can improve her worst-case expected revenues by selecting an assortment randomly according to a prudently designed probability distribution.

We first demonstrate the potential benefit of randomization in a general robust assortment optimization problem whose choice model can represent any (reasonable or unreasonable) customer behavior. We derive conditions that the problem must satisfy such that the decision maker benefits from randomizing her choice, that is, such that the problem becomes *randomization-receptive*. We then leverage these conditions to show the benefit of randomization empirically across three popular choice models: the MNL model, the MC model, and the preference ranking (PR) model.

For the ubiquitous MNL model, we show that its randomization-receptiveness hinges on whether the assortment choice is constrained. In the absence of constraints, the decision maker never benefits from randomizing her choice, that is, the model is *randomization-proof*. By contrast, when the problem is subjected to a cardinality constraint, the MNL model becomes randomization-receptive. We show that not only can the decision maker benefit from choosing an assortment randomly, but the gain from using a randomized strategy over a deterministic one can be arbitrarily large.

The MC model can be randomization-receptive even in the absence of cardinality constraints, but this depends on the characterization of the uncertainty faced by the firm. We show that the MC model is randomization-proof for uncertainty sets that exhibit a specific rectangularity property, which we call product-wise substitution sets. We establish this result through a novel interpretation of the associated assortment optimization problem as a robust Markov decision process (MDP), which allows us to directly apply results from the robust MDP literature. By contrast, for general uncertainty sets the MC model becomes randomization-receptive. Here, the benefits of randomization are bounded in the unconstrained case, but can again become unbounded under a cardinality constraint.

The PR model is also randomization-receptive even in the absence of cardinality constraints. Here again, the benefits of randomization can be unbounded in the constrained setting, whereas they remain bounded in the unconstrained case. We also establish the existence of parsimonious randomization strategies in the PR model: Even though there are exponentially many possible assortments, there are always optimal strategies that randomize between at most $K + 1$ assortments, where K is the number of considered preference rankings. This parsimony is absent in the constrained MNL and MC models, where all optimal randomization strategies can become arbitrarily complex.

For all three choice models, we show how an optimal randomization can be determined exactly (by a column generation scheme) and heuristically (through a local search algorithm). We illustrate the runtimes of these solution schemes, as well as the potential benefits of randomization, on synthetic instances of the cardinality-constrained MNL, MC and PR models. For the MNL model, we also demonstrate how the superior in-sample (worst-case) performance of randomized assortments can translate into an improved out-of-sample (expected) performance in a data-driven setting that combines estimation with optimization. Our synthetic experiments are complemented by a data-driven study on a previously published real-life data set, where we show that randomized assortments significantly outperform their deterministic counterparts. A key insight that emerges from our analysis is that more general versions of the assortment optimization problem—where the generality can be owed to the presence of business constraints, more flexible choice models and/or more general uncertainty sets—tend to be more receptive to the benefits of randomization.

It is important to recognize that randomization may not be practical in all settings, such as for a small brick-and-mortar retailer with a limited number of stores. The natural application for randomized assortment decisions is in e-commerce and other online settings, where different forms of randomization are already being applied (*e.g.*, in A/B testing various aspects of the sales experience). In these contexts, it is easy for a firm to quickly and flexibly vary the assortment displayed to a customer. Even in conventional retail settings, firms are experimenting with novel strategies adopted from online settings such as A/B testing different aspects of the store design.[1] In a similar vein, a randomized assortment choice could be implemented across a chain of otherwise homogeneous retail stores. Our results show how, in such settings, a firm may benefit from choosing an assortment randomly according to a prudently designed probability distribution.

The remainder of the paper proceeds as follows. We review the relevant related literature in Section 2. We next introduce the nominal, deterministic robust and randomized robust assortment optimization problems, as well as our notions of randomization-receptiveness/-proofness, in Section 3, where we also study the randomization-receptiveness of an abstract robust assortment optimization problem. Sections 4–6 study the benefits of randomization, as well as exact and heuristic schemes to determine randomized assortments, under the MNL, MC and PR models. Section 7 presents numerical results for the MNL model using synthetic data and for the PR model using real data, respectively. Section 8 concludes the paper. All proofs, the algorithms to compute randomized assortments, as well as additional numerical experiments on the MC and the PR models are relegated to the e-companion.

Notation. We refer to the sets of non-negative and strictly positive real numbers by $\mathbb{R}_+$ and $\mathbb{R}_{++}$, respectively. We refer to the vector of all ones and the i-th canonical basis vector as $\mathbf{e}$ and $\mathbf{e}_i$, respectively; in both cases, the context will dictate the dimension of these objects. For a finite index set $\mathcal{X}$, we let $\Delta(\mathcal{X}) = \{\boldsymbol{p} \in \mathbb{R}_+^{|\mathcal{X}|} : \sum_{x\in\mathcal{X}} p_x = 1\}$ denote the associated probability simplex. The Hadamard (element-wise) product is denoted by '$\circ$'.

2. Related Literature

This paper is related to the extensive literature on assortment optimization and choice models: Talluri and van Ryzin (2006), Kök et al. (2015), and Gallego and Topaloglu (2019) provide comprehensive overviews of this field. This literature has sought to resolve the twin problems of accurately capturing customer demand substitution using discrete choice models and efficiently finding the corresponding optimal assortment. The most popular choice model is the MNL model dating back to the work of Luce (1959) and Plackett (1975). Although the MNL model is liable to under-fitting data and suffers from the IIA property, it remains popular as both its estimation and the resulting assortment optimization problem can be solved efficiently

[1] See, *e.g.*, https://www.bain.com/insights/successful-a-b-tests-in-retail-hinge-on-these-design-considerations/.

(Talluri and van Ryzin 2004), even under a cardinality constraint on the size of the offered assortment (Rusmevichientong et al. 2010, Davis et al. 2013).

The literature has proposed a number of richer choice models to account for the MNL model's shortcomings. These include generalizations of the MNL model, such as the nested logit (Williams 1977, Davis et al. 2014) and mixture of MNL models (Rusmevichientong et al. 2014), which however come at the cost of more difficult estimation and optimization. Recently, two more general classes of choice models have been proposed. Blanchet et al. (2016) develop a tractable Markov chain model that approximates a number of parametric models; a similar idea was used in a simulation study in Zhang and Cooper (2005). Feldman and Topaloglu (2017) consider the MC model in network revenue management, Désir et al. (2020) study the constrained assortment optimization problem, and Şimşek and Topaloglu (2018) propose a method to estimate its parameters. The second class of models is based on preference rankings, early examples of which include Mahajan and van Ryzin (2001) and Rusmevichientong et al. (2006). This approach considers customer preferences through distributions over preference lists, which allows very general preference structures without imposing a parametric model. Farias et al. (2013) and van Ryzin and Vulcano (2015, 2017) study the estimation of preference ranking models. Although the assortment selection problem is intractable for general PR models (Aouad et al. 2018), special cases (Honhon et al. 2012, Paul et al. 2018, Aouad et al. 2021) can be solved efficiently. Bertsimas and Mišić (2019) consider the closely related problem of product line design under this model and propose a mixed-integer optimization based solution approach.

While more complex choice models reduce the bias in estimates, their added flexibility conversely tends to make them prone to variance (over-fitting). This concern is particularly acute when the estimate feeds into the assortment optimization problem due to the error-maximization effect of optimization (Smith and Winkler 2006), which is well known in finance (Michaud 1989, DeMiguel and Nogales 2009) and machine learning (see, *e.g.*, Bishop 2006, Hastie et al. 2009). The robust optimization approach (Ben-Tal et al. 2009, Bertsimas et al. 2011) explicitly recognizes that estimation should not produce a single point estimate for the choice model parameters but rather an uncertainty set in which the parameters lie with a pre-specified confidence. By selecting the optimal assortment considering the worst parameter setting deemed plausible, the robust approach hedges against over-fitting and can thus be seen as a form of regularization (El Ghaoui and Lebret 1997, Xu et al. 2009). Robust optimization has been applied to a wide array of operational problems in revenue management (Birbil et al. 2009, Perakis and Roels 2010), portfolio selection (Goldfarb and Iyengar 2003, Bertsimas and Sim 2004), inventory management (Bertsimas and Thiele 2006), facility location (Baron et al. 2011), and appointment scheduling (Mak et al. 2015).

The robust approach has proved successful in accounting for parameter uncertainty in choice models. The estimation procedure of Farias et al. (2013) is based on obtaining a worst-case revenues estimate among preference distributions. Rusmevichientong and Topaloglu (2012), Désir et al. (2023), and Bertsimas and Mišić (2017) study the robust assortment optimization problem under the MNL, MC, and PR models, respectively. Rusmevichientong and Topaloglu (2012) introduce uncertainty sets over MNL valuations and show that this robust MNL model preserves the feature of the nominal model that revenue-ordered assortments are optimal. Désir et al. (2023) extend this robust approach to the MC model and develop efficient algorithms to solve it. Bertsimas and Mišić (2017) consider the related problem of robust product line design under a PR model with both parameter and structural uncertainty.

In all of the aforementioned models, the decision maker deterministically chooses a single assortment to offer. We extend these models by allowing the decision maker to instead randomly choose an assortment according to a probability distribution, and showing when this may benefit her. From a mathematical perspective, the potential benefit from randomization arises from applying robust optimization to a discrete optimization problem (Bertsimas et al. 2016, Delage et al. 2019, Delage and Saif 2022). We show, however, that the superior performance of randomized assortments under this worst-case objective can translate into improved results under the original expected value objective if the model parameters are estimated from data, as is typically the case in practice. Through this insight, as well as our theoretical analysis of the randomization receptiveness of an abstract assortment optimization problem, we also contribute to the robust optimization literature.

While randomized strategies appear to be new in the context of assortment optimization, they have been successfully applied in a growing number of related revenue management contexts. The most prominent example of this are heuristic admission control strategies employed in network revenue management (Reiman and Wang 2008, Jasin and Kumar 2012, Jasin 2015, Bumpensanti and Wang 2020) to decide which assortment (*e.g.*, of flight legs) to offer to customers over time (or which classes of customers to admit) under uncertainty and resource constraints. These papers show that judiciously resolving a deterministic approximation of the problem can provide guarantees on expected revenues losses, when the firm probabilistically decides which assortment to offer, which is referred to as the allocation control (PAC) heuristic. Ferreira et al. (2018) study a retailer's dynamic pricing problem with inventory constraints and learning about customer demand. The retailer balances an exploration-exploitation tradeoff by selecting a distribution over a discrete set of prices to offer. In a recent paper, Ma (2023) compares the decision to offer an assortment with the use of lotteries where participating buyers do not know which product they will eventually receive. Drawing on the mechanism design literature, he provides conditions for offering an assortment to be optimal compared to such randomized allocation mechanisms. We add to this emerging stream of literature by showing the value of optimal randomized strategies in the classic assortment optimization problem.

3. Value of Randomization in Robust Assortment Optimization

In this section, we define the generic robust assortment optimization problem to study the value of randomized strategies in solving this problem. In Section 3.1 we study *whether* a randomized strategy may benefit a firm solving this generic problem, and derive conditions for this to be the case. Section 3.2 then discusses the intuition for *why* randomization can be beneficial. In the subsequent Sections 4-6, we investigate *when*, under prevalent choice models, randomization offers practical benefits. Section 3.3 summarizes these results.

In the *nominal assortment optimization problem*, a firm chooses an assortment to offer to customers out of N candidate products. A customer then either buys one of the offered products or makes no purchase, according to a discrete choice model. Formally, an instance of the problem is a tuple $(\mathcal{N}, \mathcal{S}, \mathfrak{C}, \boldsymbol{r})$ defined as follows. Let $\mathcal{N} = \{1, \cdots, n\}$ denote the set of products, and let $\mathcal{N}_0 = \mathcal{N} \cup \{0\}$ be the extended set that contains the no-purchase alternative indexed 0. We denote by $\mathcal{S} \subseteq \{S : S \subseteq \mathcal{N}\}$ the set of admissible assortments, which may exclude some subsets of $\mathcal{N}$ due to (*e.g.*, cardinality) constraints. When the problem is unconstrained, $\mathcal{S}$ is the power set of $\mathcal{N}$. The choice model is a mapping $\mathfrak{C} : \mathcal{S} \to \Delta(\mathcal{N}_0)$ such that $\mathfrak{C}(i|S) = 0$ for all $i \in \mathcal{N} \setminus S$, $S \in \mathcal{S}$. The price of product $i \in \mathcal{N}$ is $r_i > 0$; we also include the price of the no-purchase option, $r_0 = 0$, in the vector of product prices $\boldsymbol{r} \in \mathbb{R}_+^{n+1}$. The firm chooses an assortment $S^\star \in \mathcal{S}$ that maximizes its expected revenues $R(S) = \sum_{i \in S} r_i \cdot \mathfrak{C}(i|S)$ among all admissible assortments $S \in \mathcal{S}$:

$$R^\star_{\text{nom}} = \max_{S \in \mathcal{S}} R(S). \tag{Nominal}$$

In reality, the choice model is typically estimated from data and hence uncertain. The *robust assortment optimization problem* takes this uncertainty into account. An instance of this problem is a tuple $(\mathcal{N}, \mathcal{S}, \mathfrak{C}, \mathcal{U}, \boldsymbol{r})$, where the set of products $\mathcal{N}$ and the set of admissible assortments $\mathcal{S}$ are defined as before. The choice model $\mathfrak{C} : \mathcal{S} \times \mathcal{U} \to \Delta(\mathcal{N}_0)$, however, now depends on a realization $u \in \mathcal{U}$ from an uncertainty set $\mathcal{U}$ (*e.g.,* over choice model parameters). To avoid technicalities, we assume that $\mathcal{U}$ is a compact subset of a finite-dimensional space. We stipulate that $\mathfrak{C}(i|S, u) = 0$ for all $i \in \mathcal{N} \setminus S$, $S \in \mathcal{S}$ and $u \in \mathcal{U}$. The *deterministic robust assortment optimization problem* (see, *e.g.*, Rusmevichientong and Topaloglu 2012) then chooses the assortment that maximizes the worst-case expected revenues

$$R^\star_{\text{det}}(\mathcal{U}) = \max_{S \in \mathcal{S}} R^\star(S), \tag{Deterministic Robust}$$

where $R^\star(S) = \min_{u \in \mathcal{U}} R(S, u)$ and $R(S, u) = \sum_{i \in S} r_i \cdot \mathfrak{C}(i|S, u)$. When the uncertainty set $\mathcal{U} = \{u^0\}$ is a singleton, the choice model is known exactly and (Deterministic Robust) recovers the classical (Nominal) problem.

Instead of selecting a single assortment to offer, the firm could offer *each* assortment $S \in \mathcal{S}$ with a probability $p_S \geq 0$. In this paper, we therefore propose to solve the *randomized robust assortment optimization problem*:

$$R^\star_{\text{rand}}(\mathcal{U}) = \max_{\boldsymbol{p} \in \Delta(\mathcal{S})} R^\star(\boldsymbol{p}), \tag{Randomized Robust}$$

where $R^\star(\boldsymbol{p}) = \min_{u \in \mathcal{U}} R(\boldsymbol{p}, u)$ and $R(\boldsymbol{p}, u) = \sum_{S \in \mathcal{S}} p_S \cdot R(S, u)$. The deterministic assortments $S \in \mathcal{S}$ correspond to degenerated randomized assortments $\mathbf{e}_S \in \Delta(\mathcal{S})$ that place all probability mass onto individual assortments, and thus the feasible region of Randomized Robust contains the feasible solutions of DETERMINISTIC ROBUST. Note that depending on its argument, R may refer either to the nominal expected revenues or the expected revenues under the parameter realization $u \in \mathcal{U}$, and that $R^\star$ may refer to the worst-case expected revenues of a single assortment or a randomization strategy.

3.1. Can Randomization Be Beneficial?

We now study whether the decision maker in the robust assortment optimization problem can benefit from randomized strategies for generic choice models $\mathfrak{C}$ and uncertainty sets $\mathcal{U}$. In later sections, we consider the value of randomization under three common choice models: the MNL model, the MC model, and the PR model, each of which implies different definitions for $\mathfrak{C}$ and $\mathcal{U}$.

Since the RANDOMIZED ROBUST problem subsumes the DETERMINISTIC ROBUST problem, clearly $R^\star_{\text{rand}}(\mathcal{U}) \geq R^\star_{\text{det}}(\mathcal{U})$. Our goal is to show under what conditions randomization can benefit the firm, that is, when the strict inequality $R^\star_{\text{rand}}(\mathcal{U}) > R^\star_{\text{det}}(\mathcal{U})$ is satisfied. We say the problem is then receptive to randomization, as per the following definition.

DEFINITION 1 (RANDOMIZATION-RECEPTIVENESS/PROOFNESS). An instance of the robust assortment optimization problem is *randomization-receptive* if $R^\star_{\text{rand}}(\mathcal{U}) > R^\star_{\text{det}}(\mathcal{U})$ and *randomization-proof* otherwise. Likewise, we say that the (un-)constrained robust assortment optimization problem under a particular choice model is randomization-proof if all of its instances are randomization-proof, and it is randomization-receptive otherwise.

We note that NOMINAL is randomization-proof by construction. Indeed, for singleton uncertainty sets $\mathcal{U} = \{u^0\}$, RANDOMIZED ROBUST reduces to a linear program that attains its optimal value at an extreme point of the probability simplex $\Delta(\mathcal{S})$. This in turn corresponds to a randomization that places unit probability onto a single assortment $S \in \mathcal{S}$.

In contrast, the decision maker can benefit from randomization in the robust assortment optimization problem. The following result establishes when this is the case by providing a necessary and sufficient condition for an instance of the problem to be randomization-proof.

THEOREM 1. *Fix a robust assortment optimization instance $(\mathcal{N}, \mathcal{S}, \mathfrak{C}, \mathcal{U}, \boldsymbol{r})$. The instance is randomization-proof if and only if there is $\mathbb{Q} \in \mathcal{P}(\mathcal{U})$ such that $\mathbb{E}_{\mathbb{Q}}[R(S, \tilde{u})] \leq R^\star_{\text{det}}(\mathcal{U})$ for all $S \in \mathcal{S}$, where $\mathcal{P}(\mathcal{U})$ denotes the set of all probability distributions supported on $\mathcal{U}$.*

We illustrate the statement of Theorem 1 with an example. Consider a robust assortment optimization problem with three feasible assortments S_1, S_2 and S_3 as well as two possible uncertainty realizations u and u'. The expected revenues of each assortment are $(R(S_1, u), R(S_2, u), R(S_3, u)) = (2, 0, 1)$ and $(R(S_1, u'), R(S_2, u'), R(S_3, u')) = (1/2, 3/2, 1)$. Figure 1 (left) shows that this problem is *randomization-proof*. The figure depicts the expected revenues for all possible randomizations over the three assortments for u and u', with the worst-case expected revenues shown as solid regions. The optimal deterministic solution is $\mathbf{e}_{S_3} = (0, 0, 1)$, located along the vertical axis, with $R^\star_{\text{det}}(\mathcal{U}) = 1$. Note that any deviation from the assortment S_3 leads to a deterioration of the worst-case expected revenues: there is no direction in which we can move away from $\mathbf{e}_{S_3}$ to improve the worst-case expected revenues. In terms of the condition in Theorem 1, the optimality of the deterministic solution $\mathbf{e}_{S_3}$ is certified by the distribution $\mathbb{Q} = \frac{1}{3} \cdot \delta_u + \frac{2}{3} \cdot \delta_{u'}$, so that the expected value $\mathbb{E}_{\mathbb{Q}}[R(S, \tilde{u})]$ equals $R^\star_{\text{det}}(\mathcal{U}) = 1$ for all three assortments.

Next consider the same example but with the updated expected revenues $(R(S_1, u), R(S_2, u), R(S_3, u)) = (1, 1/2, 3/2)$ and $(R(S_1, u'), R(S_2, u'), R(S_3, u')) = (1/2, 2, 1)$. Figure 1 (right) shows that this problem is *randomization-receptive*. In contrast to the previous case, a direction to improve the worst-case revenues now exists at each deterministic solution $\mathbf{e}_{S_1}$, $\mathbf{e}_{S_2}$ and $\mathbf{e}_{S_3}$. In terms of Theorem 1, any distribution $\mathbb{Q}$ with $\mathbb{Q}[\tilde{u} = u'] < 1$ implies that $\mathbb{E}_{\mathbb{Q}}[R(S_3, \tilde{u})] > R^\star_{\text{det}}(\mathcal{U})$, while the distributions satisfying $\mathbb{Q}[\tilde{u} = u'] = 1$ imply $\mathbb{E}_{\mathbb{Q}}[R(S_2, \tilde{u})] = 2 > R^\star_{\text{det}}(\mathcal{U})$.

Together, the two graphs in Figure 1 suggest that a deterministic assortment S maximizes the worst-case expected revenues if and only if the graph $\boldsymbol{p} \mapsto R^\star(\boldsymbol{p})$ does not contain a feasible direction of improvement at the point $\mathbf{e}_S$. This is indeed what the condition in Theorem 1 amounts to: At every optimal deterministic assortment, there is no direction to move without compromising worst-case expected revenues. This intuition can be made precise by considering the super-differential of $R^\star(\mathbf{e}_S)$, and the resulting condition turns out to be equivalent to the condition in Theorem 1. We relegate the formal justification of this statement to Appendix B.

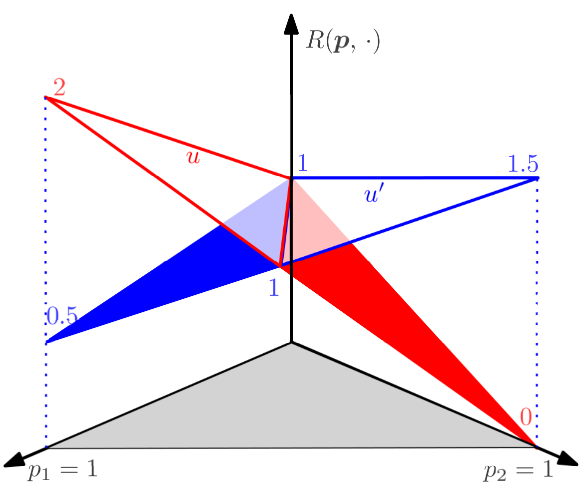

Figure 1 Randomization-proof (left) and randomization-receptive (right) instances of RANDOMIZED ROBUST. In both cases, we show two-dimensional projections of the three-dimensional probability simplex $\Delta(\mathcal{S})$ with $\mathcal{S} = \{S_1, S_2, S_3\}$ (triangles at the bottom of each diagram). The other two planes illustrate the expected revenues under the two uncertainty realizations u and u', respectively; solid (hollow) areas indicate regions where the respective realization constitutes the worst (best) case.

The condition in Theorem 1 closely relates to the concept of strong duality, which we will use in subsequent sections when discussing common choice models. Specifically, strong duality implies the condition in Theorem 1, but the reverse is not true. Consequently, as the next result shows, strong duality implies randomization-proofness.

COROLLARY 1. *Suppose that strong duality holds for the problem* $(\mathcal{N}, \mathcal{S}, \mathfrak{C}, \mathcal{U}, \boldsymbol{r})$*:*

$$\max_{S \in \mathcal{S}} \min_{u \in \mathcal{U}} R(S, u) \; = \; \min_{u \in \mathcal{U}} \max_{S \in \mathcal{S}} R(S, u).$$

Then the problem is randomization-proof.

Intuitively, any optimal deterministic assortment achieves the objective value of the left-hand side of the equation in Corollary 1, while the right-hand side of this equation constitutes an upper bound on the worst-case expected revenues achievable by any (deterministic or randomized) assortment strategy: It records the expected revenues that can be achieved if the decision maker knew the uncertainty realization u prior to making an assortment choice. Since this 'crystal ball' upper bound is achieved by an optimal deterministic assortment, the decision maker cannot improve by randomizing between assortments. The reverse statement, however, is not true: A problem instance may be randomization-proof even when strong duality does not hold. Indeed, the left-hand example in Figure 1 illustrates this: we can verify that the 'crystal ball' solution is u' with expected revenues 1.5 (at S_2), which is greater than $R^\star_{\text{det}}(\mathcal{U}) = 1$. Thus, strong duality does not hold, yet our earlier discussion showed that the instance is randomization-proof.

The equivalence between strong duality and the condition in Theorem 1 holds in the special case where the worst-case uncertainty realization is unique. Theorem 1 then simplifies to the following result: The problem is randomization-proof if and only if, under the worst-case realization, the decision maker cannot improve her expected revenues by deviating from the optimal deterministic assortment $S^\star$.

COROLLARY 2. *Fix a robust assortment optimization problem* $(\mathcal{N}, \mathcal{S}, \mathfrak{C}, \mathcal{U}, \boldsymbol{r})$ *and let* $S^\star$ *be an optimal deterministic assortment. If the worst-case parameter set* $\arg\min_{u \in \mathcal{U}} R(S^\star, u)$ *is a singleton* $\{u^\star\}$*, then the problem is randomization-proof if and only if*

$$R^\star(S^\star) \geq \max_{S \in \mathcal{S}} R(S, u^\star).$$

3.2. Why Is Randomization Beneficial?

Theorem 1 and Corollary 2 provide conditions under which the firm benefits from randomizing its assortment. We now investigate why this is the case. To establish an intuition, consider the following example.

EXAMPLE 1. *A firm wishes to offer an assortment containing no more than two out a product universe of three products, with prices* $r_1 = 10$*,* $r_2 = 9$ *and* $r_3 = 8$*. The customers' purchase behavior follows a multinomial logit model with nominal product valuations* $\boldsymbol{v}^0 = (v^0_0, v^0_1, v^0_2, v^0_3) = (1, 1, 1, 1)$*, where the outside option has value* $v^0_0 = 1$*. The optimal nominal assortment under the product valuations* $\boldsymbol{v}^0$ *comprises the products* $\{1, 2\}$ *and generates expected revenues of* 6.33 *(*cf. *the first row of Table 1).*

Valuation	Assortments						
	$\{1\}$	$\{2\}$	$\{3\}$	$\{1,2\}$	$\{1,3\}$	$\{2,3\}$	$0.476\cdot\{1,2\}+0.524\cdot\{2,3\}$
$\boldsymbol{v}^0$	5.00	4.50	4.00	6.33	6.00	5.67	5.98
$\boldsymbol{v}^1$	1.67	5.09	5.33	5.48	5.63	6.44	5.98
$\boldsymbol{v}^2$	7.5	3.00	3.56	7.67	7.58	4.74	6.13

Table 1 Expected revenues of different deterministic as well as the optimal randomized assortment under the three valuations $\boldsymbol{v}^0$, $\boldsymbol{v}^1$ and $\boldsymbol{v}^2$ in our motivating example. The green and red shadings refer to the nominal and worst-case performance of the optimal nominal assortment, respectively; the purple and cyan shadings refer to the worst-case performances of the optimal deterministic robust and randomized robust assortments, respectively.

Assume now that the product valuations may deviate from $\boldsymbol{v}^0$ *and instead amount to* $\boldsymbol{v}^1 = (v_0^1, v_1^1, v_2^1, v_3^1) = (1, 0.2, 1.3, 2)$ *or* $\boldsymbol{v}^2 = (v_0^2, v_1^2, v_2^2, v_3^2) = (1, 3, 0.5, 0.8)$ *as well. In that case, Table 1 shows that the expected revenues of the optimal nominal assortment* $\{1,2\}$ *decrease by up to* 13.47% *to* 5.48 *under scenario* $\boldsymbol{v}^2$. *A firm using the deterministic robust approach would instead offer the assortment comprising the products* $\{1,3\}$, *generating worst-case expected revenues of* 5.63 *that are* 2.29% *higher than those of the optimal nominal assortment. However, by prudently randomizing between the assortments* $\{1,2\}$ *and* $\{2,3\}$, *the firm can further increase its worst-case expected revenues by* 5.67% *to* 5.98.

The benefits of randomization can be understood from three complementary perspectives. In the example, the firm faces the risk of adverse outcomes due to taking decisions under parameter uncertainty. From a managerial perspective, randomization allows the firm to diversify against this risk and alleviate the error-maximizing effect of optimization. From a mathematical viewpoint, this diversification amounts to enlarging the feasible region of the deterministic robust assortment optimization problem by embedding the discrete assortment choices of the latter into a probability simplex that allows the firm to select convex combinations of multiple assortments in the randomized problem. From a game-theoretic viewpoint, finally, the deterministic robust assortment optimization problem constitutes a Stackelberg leader-follower game where the firm (the leader) chooses an assortment and 'nature' (the follower) picks the worst valuation vector $\boldsymbol{v}^i$ upon observing the selected assortment. In the randomized problem, on the other hand, nature only observes the firm's randomization strategy (as opposed to the randomly selected assortment), which restricts its power. As per Theorem 1 and Corollary 2, the firm then enjoys diversification benefits from randomization as long as a profitable move away from the optimal deterministic assortment exists under the worst-case realization. Indeed, the example also illustrates the randomization-proofness condition in Corollary 2: Since each optimal deterministic assortment in the table has a unique worst-case valuation, the condition in Corollary 2 does not hold, and the problem is randomization-receptive.

3.3. When Is Randomization Beneficial?

In Sections 4-6, we study the benefits of randomization under three prevalent choice models: the multinomial logit model, the Markov chain model, and the preference ranking model. For each model, we investigate *(i)* whether the problem can benefit from randomization; *(ii)* how large the potential benefits are; and

Table 2 Summary of theoretical results on the benefits of randomization.

Choice model	Constrained	Randomization receptiveness	Maximum benefit $R^\star_{\text{rand}}(\mathcal{U})/R^\star_{\text{det}}(\mathcal{U})$	Complexity of randomization
Multinomial Logit Section 4	No	No	—	—
	Yes	Yes Theorem 2	Unbounded Theorem 2	Exponential Theorem 2
Markov Chain Section 5	No	Yes/No Theorem 4, Observation 2	Bounded Observation 3	(Unknown)
	Yes	Yes Theorem 5	Unbounded Theorem 5	Exponential Theorem 5
Preference Ranking Section 6	No	Yes Observation 4	Bounded Observation 5	Linear Theorem 7
	Yes	Yes Theorem 6	Unbounded Theorem 6	Linear Theorem 7

(iii) how complex the optimal randomization strategies may be, *i.e.*, how many different assortments they may involve.

Table 2 summarizes our findings. The randomization-receptiveness of the robust assortment optimization problem depends on the complexity of the choice model and underlying business constraints. Under the parsimonious MNL model, the problem is randomization-proof in the absence of any further constraints on the chosen assortment. But when we either impose a cardinality constraint on the size of the assortment, or use the MC or PR models to capture more complex customer behavior, randomization becomes beneficial. Investigating the magnitude of the potential benefits of randomization, we find that these are bounded in the absence of business constraints, but they can become arbitrarily large under cardinality constraints. Interestingly, under both the MNL and MC models, the optimal randomization strategy may involve exponentially many assortments, whereas the PR model always admits parsimonious optimal randomization strategies that involve a small number of assortments.

4. Multinomial Logit Model

We first study the robust assortment optimization problem under the MNL model. We show that whether the firm can benefit from randomization depends on the existence of a cardinality constraint on the size of the assortment. The unconstrained problem is randomization-proof, while the cardinality-constrained problem is not only randomization-receptive, but the resulting benefit can also be arbitrarily large (Section 4.1). Section 4.2 develops algorithms for optimally solving the latter problem.

The MNL model is parameterized by a vector of customer product valuations $\boldsymbol{v} = (v_0, v_1, \cdots, v_n) \in \mathbb{R}^{n+1}_{++}$. Given these valuations and an assortment $S \in \mathcal{S}$, a customer purchases product $i \in \mathcal{N}$ with probability

$$\psi_i(S, \boldsymbol{v}) = \begin{cases} \dfrac{v_i}{v_0 + \sum_{j \in S} v_j} & \text{if } i \in S, \\ 0 & \text{otherwise,} \end{cases} \tag{1}$$

and the corresponding no-purchase probability is $\psi_0(S, \boldsymbol{v}) = 1 - \sum_{i \in S} \psi_i(S, \boldsymbol{v})$. The expected revenues amount to

$$R(S, \boldsymbol{v}) \;=\; \sum_{i \in S} r_i \cdot \psi_i(S, \boldsymbol{v}) \;=\; \frac{\sum_{i \in S} r_i \cdot v_i}{v_0 + \sum_{i \in S} v_i}, \tag{2}$$

where we identify the parameter vector $\boldsymbol{u}$ with the product valuations $\boldsymbol{v}$. In the corresponding robust assortment optimization problems, we assume that $\boldsymbol{u}$ is only known to be contained in a compact uncertainty set $\mathcal{U} = \mathcal{V} \subseteq \mathbb{R}^{n+1}_{++}$.

4.1. Randomization-Receptiveness of the Multinomial Logit Problem

The unconstrained problem presented above is randomization-proof. Indeed, Désir et al. (2023) show that the problem satisfies strong duality, that is, we have

$$\max_{S \in \mathcal{S}} \min_{\boldsymbol{v} \in \mathcal{V}} R(S, \boldsymbol{v}) \;=\; \min_{\boldsymbol{v} \in \mathcal{V}} \max_{S \in \mathcal{S}} R(S, \boldsymbol{v}).$$

Corollary 1 then implies that randomization over assortments cannot benefit the firm.

Strong duality, however, ceases to hold if we impose a cardinality constraint $|S| \leq C$ on the size of the offered assortment. We have seen in Example 1 that the firm may benefit from randomization under a cardinality constraint. We now show that the benefit from randomization can indeed be arbitrarily large in the constrained problem, but that at the same time the optimal strategy may require randomization between many assortments.

THEOREM 2. *For any number of products $n \geq 2$ and any restriction $|S| \leq C$, $C \in \{1, \cdots, n-1\}$, there are instances of the cardinality-constrained robust MNL problem where*

1. $R^\star_{\text{det}}(\mathcal{V}) = 0$ *while* $R^\star_{\text{rand}}(\mathcal{V}) > 0$*;*
2. *the unique optimal randomized assortment strategy places equal (positive) probability on each assortment S satisfying $|S| = C$ and zero probability on all other assortments.*

The first part of Theorem 2 implies that the benefits $R^\star_{\text{rand}}(\mathcal{V})/R^\star_{\text{det}}(\mathcal{V})$ from randomization are unbounded. The proof of the theorem, which is relegated to the appendix, considers a class of robust constrained MNL instances where all products have equal price and the uncertainty set $\mathcal{V}$ comprises all valuation vectors $\boldsymbol{v}$ in which exactly B valuations are zero and the remaining valuations are 1. For $B \geq C$, where C is the admissible assortment cardinality, any deterministic assortment S results in worst-case expected revenues of zero since the uncertainty set contains a vector $\boldsymbol{v}$ in which $v_i = 0$ for all $i \in S$. The optimal randomized strategy, on the other hand, places equal (positive) probability on all assortments with exactly C products, and zero probability on all other assortments. As long as $B < n$, this randomized strategy raises strictly positive worst-case revenues since there is a strictly positive probability that the offered assortment contains products whose values are strictly positive even under the worst-case valuation realization. The optimal randomization strategy, however, may involve exponentially many assortments, as we can see from the second statement by setting $C = \lfloor n/2 \rfloor$, which results in $\Theta(2^n/\sqrt{n})$ many assortments. Optimal randomization strategies can thus generally be complex. We next turn to the problem of finding such optimal strategies.

4.2. Solving the Randomized Constrained Multinomial Logit Problem

Appendix A presents two algorithms to compute randomization strategies for the constrained robust MNL problem with a binary representable uncertainty set of the form

$$\mathcal{V} = \{\boldsymbol{v} = \boldsymbol{F}\boldsymbol{\xi} \, : \, \boldsymbol{A}\boldsymbol{\xi} \le \boldsymbol{b}, \;\; \boldsymbol{\xi} \in \{0,1\}^m\}, \tag{3}$$

where $\boldsymbol{F} \in \mathbb{R}^{(n+1)\times m}$, $\boldsymbol{A} \in \mathbb{R}^{l\times m}$ and $\boldsymbol{b} \in \mathbb{R}^l$. While any discrete uncertainty set is binary representable, we list some popular uncertainty sets that enjoy compact representations.

(i) **Budget uncertainty sets.** For lower and upper valuation bounds $\underline{\boldsymbol{v}}, \overline{\boldsymbol{v}} \in \mathbb{R}^{n+1}$ and an uncertainty budget $\Gamma \in \mathbb{N}$, we define the uncertainty set

$$\mathcal{V} = \left\{\boldsymbol{v} = \overline{\boldsymbol{v}} - (\overline{\boldsymbol{v}} - \underline{\boldsymbol{v}}) \circ \boldsymbol{\xi} \, : \, \mathbf{e}^\top \boldsymbol{\xi} \le \Gamma, \;\; \boldsymbol{\xi} \in \{0,1\}^{n+1}\right\}.$$

Under the budget uncertainty set, up to Γ valuations v_i can attain their lower bounds $\underline{v}_i$, whereas the remaining valuations v_i attain their upper bounds $\overline{v}_i$.

(ii) **Factor model uncertainty sets.** For a nominal valuation vector $\boldsymbol{v}^0 \in \mathbb{R}^{n+1}$ and a factor loading matrix $\boldsymbol{\Phi} \in \mathbb{R}^{(n+1)\times m}$, we define the uncertainty set

$$\mathcal{V} = \left\{\boldsymbol{v} = \boldsymbol{v}^0 + \boldsymbol{\Phi}(2\boldsymbol{\xi} - \mathbf{e}) \, : \, \boldsymbol{\xi} \in \{0,1\}^m\right\}.$$

Here, the valuations $\boldsymbol{v}$ differ from their nominal values $\boldsymbol{v}^0$ by $\boldsymbol{\Phi}\mathcal{B}_\infty$, where $\mathcal{B}_\infty = \{-1,1\}^m$ contains the extreme points of the unit ∞-norm ball in $\mathbb{R}^m$.

(iii) **Norm uncertainty sets.** For a nominal valuation vector $\boldsymbol{v}^0 \in \mathbb{R}^{n+1}$ and lower and upper valuation bounds $\underline{\boldsymbol{v}}, \overline{\boldsymbol{v}} \in \mathbb{R}^{n+1}$, we define the uncertainty set

$$\mathcal{V} = \left\{\boldsymbol{v} = \boldsymbol{v}^0 + (\overline{\boldsymbol{v}} - \boldsymbol{v}^0) \circ \boldsymbol{\xi}^+ + (\underline{\boldsymbol{v}} - \boldsymbol{v}^0) \circ \boldsymbol{\xi}^- \, : \, \begin{array}{l} \left\|(\overline{\boldsymbol{v}} - \boldsymbol{v}^0) \circ \boldsymbol{\xi}^+ + (\underline{\boldsymbol{v}} - \boldsymbol{v}^0) \circ \boldsymbol{\xi}^-\right\|_p \le \theta \\ \boldsymbol{\xi}^+ + \boldsymbol{\xi}^- \le \mathbf{e}, \;\; \boldsymbol{\xi}^+, \boldsymbol{\xi}^- \in \{0,1\}^{n+1} \end{array}\right\}$$

where $p \in \{0, 1, \infty\}$ and $\theta \in \mathbb{R}_+$. Norm uncertainty sets hedge against valuations $\boldsymbol{v}$ that reside in a θ-neighbourhood of the nominal valuations $\boldsymbol{v}^0$, as measured by the p-norm.

The randomized constrained MNL model can be formulated as a robust linear program. To this end, denote by $\mathcal{S} = \{S \in \mathcal{N} \, : \, |S| \le C\}$ the set of assortments whose cardinality is less or equal to the size restriction C. The problem can then be written as

$$\begin{array}{ll} \text{maximize} & \displaystyle \min_{\boldsymbol{v} \in \mathcal{V}} \sum_{S \in \mathcal{S}} p_S \cdot \frac{\sum_{i \in S} r_i \cdot v_i}{v_0 + \sum_{i \in S} v_i} \\ \text{subject to} & \displaystyle \sum_{S \in \mathcal{S}} p_S = 1 \\ & p_S \ge 0, \; S \in \mathcal{S}. \end{array} \tag{4}$$

This problem is computationally challenging since it typically comprises an exponential number of decision variables, and its objective function contains an embedded optimization problem that minimizes a non-convex function over a discrete uncertainty set $\mathcal{V}$. For later reference, we note that the dual of problem (4) amounts to the robust linear program

$$\begin{array}{ll} \text{minimize} & \displaystyle\max_{S\in\mathcal{S}} \sum_{\boldsymbol{v}\in\mathcal{V}} \kappa_{\boldsymbol{v}} \cdot \frac{\sum_{i\in S} r_i \cdot v_i}{v_0 + \sum_{i\in S} v_i} \\ \text{subject to} & \displaystyle\sum_{\boldsymbol{v}\in\mathcal{V}} \kappa_{\boldsymbol{v}} = 1 \\ & \kappa_{\boldsymbol{v}} \geq 0,\ \boldsymbol{v}\in\mathcal{V}. \end{array} \tag{5}$$

Strong duality between (4) and (5) holds since (4) is feasible by construction. Similar to problem (4), problem (5) is computationally challenging due to the typically exponential number of decision variables as well as the embedded maximization over a discrete 'uncertainty set' $\mathcal{S}$ whose decision variables S appear in a non-convex objective function.

In Appendix A, we propose an exact column generation scheme as well as a heuristic local search method to solve the randomized constrained MNL problem.

5. Markov Chain Model

The MC model proposed by Blanchet et al. (2016) generalizes the MNL model. Under the MC model, the choice behavior of a customer is described through a vector $\boldsymbol{\lambda}\in\mathbb{R}_+^{n+1}$, where λ_i, $i=0,\cdots,n$, denotes the probability of product i being the most preferred choice, as well as a matrix $\boldsymbol{\rho}=(\rho_{ij})_{i,j}\in\mathbb{R}_+^{(n+1)\times(n+1)}$, where ρ_{ij}, $i,j=0,\cdots,n$, characterizes the probability of the customer substituting product i by product j if product i is not available. In other words, an arriving customer attempts to purchase each product $i=0,\cdots,n$ with probability λ_i. If the preferred product, say i, is not offered, then she attempts to purchase each product $j=0,\cdots,n$ with probability ρ_{ij}, and the process continues as if j had been her preferred choice. We require that $\rho_{00}=1$ as well as $\rho_{ii}=0$ and $\rho_{i0}>0$ for all $i\neq 0$.

In the robust assortment optimization problem under the MC model (Désir et al. 2023), the substitution matrices $\boldsymbol{\rho}$ are only known to reside in an uncertainty set

$$\mathcal{U}\subseteq\left\{\boldsymbol{\rho}\in\mathbb{R}_+^{(n+1)\times(n+1)} : \sum_{j=0}^{n}\rho_{ij}=1\ \forall i=0,\cdots,n\right\},$$

which we assume to be compact in order to avoid technicalities. Under this setting, the substitution behavior of the customers can be determined by any substitution matrix $\boldsymbol{\rho}\in\mathcal{U}$, and the decision maker optimizes the expected revenues considering the worst such matrix. We require that all $\boldsymbol{\rho}\in\mathcal{U}$ satisfy $\rho_{00}=1$ as well as $\rho_{ii}=0$ and $\rho_{i0}>0$ for all $i\neq 0$. In addition, we require that there exists a $\boldsymbol{\rho}\in\mathcal{U}$ and a product $i\neq 0$ satisfying that $\rho_{i0}<1$.

We now study whether the robust assortment optimization problem under the MC model benefits from randomization. To do so, we first express the problem as a robust Markov decision process (MDP) in Section 5.1. Leveraging this interpretation, we then show that the benefits of randomization depend on the nature of parameter uncertainty. Specifically, we contrast product-wise substitution sets where no information is available about the dependence of ρ_{ij} and ρ_{kl} whenever $i \neq k$ and general substitution sets (Section 5.2). Section 5.3, finally, discusses the solution of the (un-)constrained robust assortment optimization problem under the MC model.

5.1. A Robust MDP Reformulation for the Unconstrained Markov Chain Problem

Our objective is to reformulate the robust assortment optimization problem under the MC model as an instance of a robust MDP as per the following definition.

DEFINITION 2 (ROBUST MDP). A robust MDP is defined by the tuple $(\mathcal{X}, \mathcal{A}, \mathcal{P}, q, c, \gamma)$, where $\mathcal{X}$ denotes the state space, $\mathcal{A}$ represents the action space, $q(x)$, $x \in \mathcal{X}$, characterize the initial state distribution, $c(x,a)$, $x \in \mathcal{X}$ and $a \in \mathcal{A}$, denote the immediate rewards, and $\gamma \in (0,1)$ is the discount factor. The ambiguity set

$$\mathcal{P} \subseteq \left\{ p : \mathcal{X} \times \mathcal{A} \times \mathcal{X} \to \mathbb{R}_+ \, : \, \sum_{x' \in \mathcal{X}} p(x'|x,a) = 1 \;\; \forall x \in \mathcal{X}, \, \forall a \in \mathcal{A} \right\}$$

contains all transition kernels p that are deemed plausible by the decision maker.

A robust MDP starts in state $x \in \mathcal{X}$ with known probability $q(x)$. The decision maker can then select any action $a \in \mathcal{A}$, upon which an immediate reward of $c(x,a)$ is earned and the MDP transitions to state $x' \in \mathcal{X}$ with probability $p(x'|x,a)$, where p can be any element of the ambiguity set $\mathcal{P}$. The process then continues in the same fashion, governed by the same transition kernel p, for an infinite length of time. The decision maker wishes to determine a policy $\pi : \mathcal{X} \to \mathcal{A}$, which declares for each state $x \in \mathcal{X}$ which action $a \in \mathcal{A}$ is to be taken in state x, that maximizes the worst-case expected total discounted reward:

$$\max_{\pi \in \Pi} \min_{p \in \mathcal{P}} \mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^t \cdot c(x_t, a_t) \;\middle|\; x_0 \sim q \right]$$

Here, Π denotes the set of all deterministic, memoryless policies $\pi : \mathcal{X} \to \mathcal{A}$, $\{(x_t, a_t)\}_{t=0}^{\infty}$ is the stochastic process induced by the initial probabilities q, the transition probabilities p and the policy π, and $\mathbb{E}$ is the expectation operator with respect to this process.

We next construct a robust MDP for the unconstrained robust assortment optimization problem under the MC model.

DEFINITION 3 (ROBUST MDP REPRESENTATION). For a given MC model, we define the robust MDP $(\mathcal{X}, \mathcal{A}, \mathcal{P}, q, c, \gamma)$ via the state space $\mathcal{X} = \mathcal{N}_0$, the action space $\mathcal{A} = \{\top, \perp\}$, the ambiguity set $\mathcal{P}$ that contains all transition kernels p satisfying

$$p(x'|x,\top) = \begin{cases} 1 & \text{if } x' = x, \\ 0 & \text{otherwise,} \end{cases} \qquad p(x'|x,\perp) = \begin{cases} \rho_{xx'}/\gamma & \text{if } x, x' \neq 0, \\ 1 - \sum_{\chi \in \mathcal{N}} \rho_{x\chi}/\gamma & \text{if } x' = 0, \\ 0 & \text{otherwise} \end{cases} \qquad \forall x, x' \in \mathcal{X} \tag{6}$$

for some $\boldsymbol{\rho} \in \mathcal{U}$, the initial state probabilities $q(x) = \lambda_x$, $x \in \mathcal{X}$, the immediate rewards $c(x, \top) = (1-\gamma) r_x$ and $c(x, \perp) = 0$, $x \in \mathcal{X}$, and $\gamma = 1 - \min\{\rho_{i0} : \boldsymbol{\rho} \in \mathcal{U},\, i \in \mathcal{N}\}$.

Intuitively, the states of the robust MDP describe the different purchase options $x \in \mathcal{N}_0$ of the customer, and the actions $\top$ and $\perp$ characterize the options of the decision maker to include (exclude) any of the products $x \in \mathcal{N}$ in/from the assortment. Note that the virtual product 0 is always available, and hence the selected action does not matter in state 0. If the customer enters a state whose associated purchase option is part of the assortment, then this state keeps generating revenues and is never left. Otherwise, the next purchase option considered by the customer is selected randomly according to some $\boldsymbol{\rho} \in \mathcal{U}$.

We next show that Definition 3 indeed describes a valid robust MDP.

OBSERVATION 1. The robust MDP from Definition 3 is well-defined.

We can now establish the equivalence between the unconstrained robust assortment optimization problem under the MC model and our robust MDP from Definition 3.

THEOREM 3. *The unconstrained robust assortment optimization problem under the MC model is equivalent to the robust MDP from Definition 3 in the following sense:*

(i) For every $S \in \mathcal{S}$, the worst-case expected total discounted reward of any $\pi_S \in \Pi$ satisfying $\pi_S(i) = \top$, $i \in S$, and $\pi_S(i) = \perp$, $i \in \mathcal{N} \setminus S$, coincides with the revenues $R^\star(S)$.

(ii) For every $\pi \in \Pi$, the revenues $R^\star(S)$ of $S = \{i \in \mathcal{N} : \pi(i) = \top\}$ coincide with the worst-case expected total discounted reward of π.

Theorem 3 shows that there is a one-to-many relationship between the assortments $S \in \mathcal{S}$ of the assortment optimization problem and the policies $\pi_S : \mathcal{X} \to \mathcal{A}$ through the relation $i \in S \iff \pi_S(i) = \top$, $i \in \mathcal{N}$. The relationship is not one-to-one since the no-purchase option 0 is always available, irrespective of whether $\pi_S(0) = \top$ or $\pi_S(0) = \perp$. Theorem 3 allows us to apply the rich arsenal of solution methods for robust MDPs to the unconstrained deterministic robust assortment optimization problem under the MC model with product-wise substitution sets (*cf.* Section 5.3). In the next subsection, we will leverage the established theory for robust MDPs to investigate under which conditions the decision maker may benefit from randomizing between multiple assortments, that is, when $R^\star_{\text{rand}}(\mathcal{U}) > R^\star_{\text{det}}(\mathcal{U})$.

REMARK 1 (NOMINAL ASSORTMENT OPTIMIZATION). If the uncertainty set of the robust assortment optimization problem is a singleton, say $\mathcal{U} = \{\boldsymbol{\rho}^0\}$, then the robust MDP from Definition 3 reduces to a nominal MDP. In that case, the conclusions from Observation 1 and Theorem 3 continue to apply.

5.2. Randomization-Receptiveness of the Markov Chain Problem

We saw that under the MNL model, the unconstrained robust assortment optimization problem is randomization-proof. We now investigate whether this remains the case for the more general MC model. It turns out that while certain special cases of this more general problem are randomization-proof, the firm may generally benefit from randomization even in the unconstrained setting, and this carries over to the cardinality-constrained setting.

We first consider the special case of uncertainty sets with *product-wise substitution sets*,

$$\mathcal{U} = \left\{ \boldsymbol{\rho} \in \mathbb{R}_+^{(n+1)\times(n+1)} : \exists \boldsymbol{\rho}^0, \cdots, \boldsymbol{\rho}^n \in \mathcal{U} \text{ such that } \rho_{ij} = \rho_{ij}^i \;\; \forall i, j \in \mathcal{N}_0 \right\}.$$

One readily verifies that this condition is equivalent to requiring that

$$\mathcal{U} = \bigtimes_{i \in \mathcal{N}_0} \mathcal{U}_i, \quad \text{where } \mathcal{U}_i = \left\{ \boldsymbol{\rho}_i = (\rho_{i0}, \cdots, \rho_{in}) \in \mathbb{R}_+^{n+1} : \boldsymbol{\rho} \in \mathcal{U} \right\}.$$

Intuitively, the uncertainty set $\mathcal{U}$ has product-wise substitution sets when knowledge of the uncertain substitution probabilities $\boldsymbol{\rho}_i = (\rho_{i0}, \cdots, \rho_{in})$ for product $i \in \mathcal{N}_0$ does not allow the decision maker to infer anything about the uncertain substitution probabilities $\boldsymbol{\rho}_j = (\rho_{j0}, \cdots, \rho_{jn})$ for any other product $j \in \mathcal{N}_0$, $j \neq i$, beyond the fact that $\boldsymbol{\rho}_j \in \mathcal{U}_j$. In this case, similar to the MNL model, randomization offers no benefits in the unconstrained problem.

THEOREM 4. *If the uncertainty set $\mathcal{U}$ of the unconstrained robust assortment optimization problem under the MC model has product-wise substitution sets, the problem is randomization-proof.*

The proof of Theorem 4 shows that for product-wise substitution sets, the robust MDP from Definition 3 has an (x, a)-rectangular ambiguity set, which implies strong duality:

$$\max_{\pi \in \Pi} \min_{p \in \mathcal{P}} \mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^t \cdot r(x_t, a_t) \,\middle|\, x_0 \sim q \right] \;=\; \min_{p \in \mathcal{P}} \max_{\pi \in \Pi} \mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^t \cdot r(x_t, a_t) \,\middle|\, x_0 \sim q \right].$$

Corollary 1 then implies that the problem is randomization-proof.

Désir et al. (2023) derive a strong duality result for the unconstrained robust assortment optimization problem under the MC model with product-wise substitution sets (termed *row-wise uncertainty* in that work). In contrast to our result, which establishes a connection between the unconstrained assortment optimization problem and robust MDPs and subsequently leverages existing results for robust MDPs, Désir et al. (2023) prove strong duality *ab initio*.

While product-wise substitution sets may appear to constitute an intuitive choice for the uncertainty set $\mathcal{U}$, they are unlikely to arise as a result of a statistical estimation from historical data. Instead, a data-driven estimation approach based on a confidence region formed, for example, by a maximum likelihood estimation, would exhibit an asymptotically elliptical shape under which all rows $\boldsymbol{\rho}_i$ of the uncertain substitution matrix $\boldsymbol{\rho}$ are dependent (*cf.* Billingsley 1961). Under such general substitution sets, strong duality no longer holds for the associated robust MDPs (*cf.* Wiesemann et al. 2013), which opens up the possibility for the problem to be randomization-receptive. This does not automatically imply, however, that randomization is beneficial, as the next result shows.

PROPOSITION 1. *Any instance of the unconstrained robust MC problem with two products, irrespective of the geometry of the uncertainty set $\mathcal{U}$, is randomization-proof.*

Generally, however, the more complex choice behavior captured by the MC model allows the firm to benefit from randomization. Specifically, with three or more products, we can find instances where the unconstrained robust MC problem with general substitution sets is receptive to randomization.

OBSERVATION 2. For any number $n \geq 3$ of products, the unconstrained MC problem is randomization-receptive.

In the constrained MNL problem, we found the potential benefits of randomization to be unbounded. In the unconstrained MC problem, however, this is not the case, as the following result shows.

OBSERVATION 3. Under the unconstrained MC problem with general substitution sets, the benefits of randomization are bounded from above by

$$\frac{R^\star_{rand}(\mathcal{U})}{R^\star_{det}(\mathcal{U})} \leq \frac{\max\{r_i : i \in \mathcal{N}\}}{\min\{r_i : i \in \mathcal{N}\}}.$$

The intuition underlying Observation 3 is as follows. Imagine first that the customers do not exercise the outside option. In that case, the benefits of randomization cannot exceed the bound of Observation 3 since in the unconstrained problem, the assortment containing all products $\mathcal{N}$ is always feasible and generates revenues of at least $\min\{r_i : i \in \mathcal{N}\}$, while no randomized assortment strategy can generate more revenues than $\max\{r_i : i \in \mathcal{N}\}$. Since the probability of a customer exercising the outside strategy under assortment T is at most as large as under any assortment $S \subseteq T$, one can show that the bound of Observation 3 carries over to the case where customers can exercise the outside option.

We now turn to the cardinality-constrained problem. The next result confirms that similar to the MNL problem, the cardinality-constrained MC problem is randomization-receptive under product-wise and general substitution sets. Moreover, the benefits of randomization are again unbounded, and any optimal randomization strategy can be complex.

THEOREM 5. *For any number of products $n \geq 3$ and any restriction $|S| \leq C$, $C \in \{1, \cdots, n-2\}$, there are instances of the cardinality-constrained robust MC problem with product-wise substitution sets where*

- $R^\star_{\text{det}}(\mathcal{U}) = 0$ *while* $R^\star_{\text{rand}}(\mathcal{U}) > 0$.
- *the unique optimal randomized assortment strategy places equal (positive) probability on each assortment S satisfying $|S| = C$ and zero probability on all other assortments.*

5.3. Solving the Randomized Markov Chain Problem

To solve the unconstrained and constrained randomized assortment optimization problem under the MC model with general substitution sets, we stipulate that the uncertainty set $\mathcal{U} \subseteq \mathbb{R}_+^{(n+1)\times(n+1)}$ is a finite set of substitution matrices. In the special case of product-wise substitution sets, we furthermore stipulate that

$\mathcal{U} = \bigtimes_{i \in \mathcal{N}_0} \mathcal{U}_i$ for finite sets $\mathcal{U}_i \subseteq \mathbb{R}^{n+1}$ of substitution vectors for each product $i \in \mathcal{N}_0$. Popular examples of such uncertainty sets include projections of factor model uncertainty sets onto the cross-product of $n+1$ probability simplices $\Delta(\mathcal{N}_0)$. We present in Appendix A a two-layer column generation framework to solve the corresponding assortment optimization problems.

6. Preference Ranking Model

The PR model (see, *e.g.*, Bertsimas and Mišić 2019) is parameterized by K bijective preference rankings $\sigma_k : \mathcal{N}_0 \to \{1, \cdots, n+1\}$, $k \in \mathcal{K} = \{1, \cdots, K\}$, with occurrence probabilities $\lambda_k \in \mathbb{R}_+$ satisfying $\mathbf{e}^\top \boldsymbol{\lambda} = 1$. The probability that a customer is characterized by the preference ranking σ_k, $k \in \mathcal{K}$, is λ_k. Faced with the assortment S as well as the no-purchase option 0, such a customer purchases the product $i \in S \cup \{0\}$ that has the smallest rank in σ_k, that is, $i \in \arg\min\{\sigma_k(j) : j \in S \cup \{0\}\}$. Hence, the expected revenues of the assortment $S \in \mathcal{S}$ under the PR model amount to

$$R(S, \boldsymbol{\lambda}) = \sum_{k \in \mathcal{K}} \lambda_k \cdot R_k(S), \tag{7}$$

where $R_k(S) = r_i$ for the unique $i \in S \cup \{0\}$ that satisfies $\sigma_k(i) < \sigma_k(j)$ for all $j \in S \cup \{0\}$, $j \neq i$. In the corresponding robust assortment optimization problem, we assume that the vector of occurrence probabilities $\boldsymbol{\lambda}$ is only known to be contained in a compact ambiguity set $\mathcal{U} = \Lambda \subseteq \{\boldsymbol{\lambda} \in \mathbb{R}^K_+ : \mathbf{e}^\top \boldsymbol{\lambda} = 1\}$, and we seek to maximize the worst-case expected revenues.

Section 6.1 studies the potential benefits of randomization under the robust PR problem, and Section 6.2 discusses the solution of the (un-)constrained robust assortment optimization problem under the PR model.

6.1. Randomization-Receptiveness of the Preference Ranking Problem

Similar to the MC problem under general uncertainty sets, the robust assortment optimization problem under the PR model is randomization-receptive even in the absence of any further constraints (such as cardinality constraints). In fact, randomization can be beneficial for any number of products, as we show next.

OBSERVATION 4. For any number of products $n \geq 2$, the unconstrained as well as the constrained PR problem are randomization-receptive.

Contrary to the constrained MNL and MC problems (*cf.* Theorems 2 and 5), however, and similar to the unconstrained MC problem with general substitution sets (*cf.* Observation 3), the benefits of randomization are then bounded.

OBSERVATION 5. For any instance of the unconstrained robust PR problem with $n = 2$ products, the benefits of randomization are bounded from above by $R^\star_{\mathit{rand}}(\mathcal{U}) \,/\, R^\star_{\mathit{det}}(\mathcal{U}) \leq 2$. Moreover, for any instance the benefits of randomization are bounded from above by

$$\frac{R^\star_{\mathit{rand}}(\mathcal{U})}{R^\star_{\mathit{det}}(\mathcal{U})} \leq \frac{\max\{r_i : i \in \mathcal{N}\}}{\min\{r_i : i \in \mathcal{N}\}}.$$

When we impose a cardinality constraint $|S| \le C$ on the size of the offered assortment, similar to the cardinality-constrained MNL and MC problems (*cf.* Theorems 2 and 5), the potential benefits of randomization in the cardinality-constrained PR problem can be arbitrarily large.

THEOREM 6. *For any number of products $n \ge 2$ and any restriction $|S| \le C$, $C \in \{1, \cdots, n-1\}$, there are instances of the cardinality-constrained robust PR problem where $R^\star_{\text{det}}(\mathcal{U}) = 0$ while $R^\star_{\text{rand}}(\mathcal{U}) > 0$.*

We have seen that in the cardinality-constrained MNL and MC problems, optimal policies may require randomization between an exponentially large number of assortments (*cf.* Theorem 2 and Theorem 5). By contrast, under the PR model (whether constrained or not), there is always a parsimonious optimal randomization strategy.

THEOREM 7. *Under the PR model (with arbitrary constraints on the set of admissible assortments), there always exists an optimal randomization strategy which places strictly positive weight on no more than $K+1$ assortments. Moreover, when the uncertainty set is an affine map of a polyhedral subset of $\mathbb{R}^m$, there exists an optimal randomization strategy which places strictly positive weight on no more than $m+1$ assortments.*

6.2. Solving the Randomized Preference Ranking Problem

Theorem 7 allows us to formulate the unconstrained and constrained randomized assortment optimization problem under the PR model as robust K-adaptability problems that determine $K+1$ assortments and their randomization weights (Bertsimas and Caramanis 2010, Hanasusanto et al. 2015). While the resulting problems can be expressed as mixed-integer linear programs of compact size that are amenable to solution via standard solvers, our numerical experiments indicate that this approach does not scale to problems of interesting size. Instead, we develop an exact and a heuristic column generation scheme to compute randomization strategies for the PR model. To this end, we assume that the ambiguity set Λ is a polyhedral set of the form

$$\Lambda = \left\{ \boldsymbol{\lambda} = \boldsymbol{F}\boldsymbol{\xi} \, : \, \boldsymbol{A}\boldsymbol{\xi} \le \boldsymbol{b}, \;\; \boldsymbol{\xi} \in \mathbb{R}^m_+ \right\}, \tag{8}$$

where $\boldsymbol{F} \in \mathbb{R}^{K \times m}$, $\boldsymbol{A} \in \mathbb{R}^{l \times m}$ and $\boldsymbol{b} \in \mathbb{R}^l$. We list below two popular choices of such sets.

(i) **Norm ambiguity sets.** For a nominal ranking prevalence vector $\boldsymbol{\lambda}^0 \in \mathbb{R}^K_+$ and a radius $\theta \in \mathbb{R}_+$, we define the ambiguity set

$$\Lambda = \left\{ \boldsymbol{\lambda} : \mathbf{e}^\top \boldsymbol{\lambda} = 1, \;\; \left\| \boldsymbol{\lambda} - \boldsymbol{\lambda}^0 \right\|_p \le \theta, \;\; \boldsymbol{\lambda} \in \mathbb{R}^K_+ \right\},$$

where $p \in \{1, \infty\}$. Norm ambiguity sets hedge against all perturbations of the ranking prevalence vector $\boldsymbol{\lambda}$ that are contained in a p-ball of radius θ around the nominal vector $\boldsymbol{\lambda}^0$. In particular, the choice $p = 1$ recovers the total variation distance, a popular ϕ-divergence (Bayraksan and Love 2015), which allows us to choose the radius θ based on statistical bounds.

(ii) **Approximate ellipsoidal ambiguity sets.** For a nominal ranking prevalence vector $\boldsymbol{\lambda}^0 \in \mathbb{R}^K$ and a symmetric and positive definite matrix $\boldsymbol{P} \in \mathbb{R}^{K \times K}$, we set

$$\Lambda = \left\{ \boldsymbol{\lambda} : \mathbf{e}^\top \boldsymbol{\lambda} = 1,\ \boldsymbol{\lambda} = \boldsymbol{\lambda}^0 + \boldsymbol{P}\boldsymbol{\xi},\ \|\boldsymbol{\xi}\|_1 \leq \sqrt{K},\ \|\boldsymbol{\xi}\|_\infty \leq 1,\ \boldsymbol{\lambda} \in \mathbb{R}_+^K,\ \boldsymbol{\xi} \in \mathbb{R}^K \right\}.$$

Note that $\{\boldsymbol{\xi} \in \mathbb{R}^K : \|\boldsymbol{\xi}\|_2 \leq 1\} \subseteq \{\boldsymbol{\xi} \in \mathbb{R}^K : \|\boldsymbol{\xi}\|_1 \leq \sqrt{K}\} \cap \{\boldsymbol{\xi} \in \mathbb{R}^K : \|\boldsymbol{\xi}\|_\infty \leq 1\}$, and thus the ambiguity set constitutes an outer (conservative) approximation of an ellipsoid with center $\boldsymbol{\lambda}^0$ and semi-axes defined by $\boldsymbol{P}$. Ellipsoidal ambiguity sets recover the Pearson χ^2-divergence, another popular ϕ-divergence (Bayraksan and Love 2015), and they emerge asymptotically as confidence regions of a maximum likelihood estimation.

The randomized robust assortment optimization problem under the PR model can be formulated as the following robust linear program:

$$\begin{array}{ll} \text{maximize} & \displaystyle\min_{\boldsymbol{\lambda} \in \Lambda} \sum_{S \in \mathcal{S}} \sum_{k \in \mathcal{K}} p_S \cdot \lambda_k \cdot R_k(S) \\ \text{subject to} & \displaystyle\sum_{S \in \mathcal{S}} p_S = 1 \\ & p_S \geq 0,\ S \in \mathcal{S} \end{array} \tag{9}$$

Problem (9) is computationally challenging as it involves exponentially many decision variables p_S, $S \in \mathcal{S}$. For our solution schemes, it is useful to consider the dual of problem (9):

$$\begin{array}{ll} \text{minimize} & \displaystyle\max_{S \in \mathcal{S}} \sum_{k \in \mathcal{K}} \lambda_k \cdot R_k(S) \\ \text{subject to} & \boldsymbol{\lambda} \in \Lambda \end{array} \tag{10}$$

Strong duality between (9) and (10) holds since (9) is feasible by construction. Although problem (10) contains only polynomially many decision variables, the optimization problem embedded in its objective function maximizes over a combinatorial set $S \in \mathcal{S}$. Thus, neither the primal problem (9) nor the dual problem (10) is amenable to a solution with an off-the-shelf solver. We present in Appendix A.3 a column generation scheme for problem (9) that can be interpreted as a cutting plane approach for problem (10).

7. Numerical Results

This section aims to elucidate the value of randomization using both synthetic and real data. We first investigate when randomization can help to improve the worst-case expected revenues under the cardinality-constrained MNL problem using synthetic data. We also study the computational price to be paid for exact and heuristic solutions to the deterministic and randomized robust assortment optimization problems, and we demonstrate how the improvement of the worst-case expected revenues can translate into improvements of the out-of-sample expected revenues in a data-driven setting. Finally, Section 7.2 demonstrates the value of randomization on a real-world data set. Appendix F presents numerical results on the value of randomization and computational cost in the MC and PR problems on synthetic data. All solution schemes are

Γ \ C	5%	10%	15%	20%	25%	30%	50%	75%
5%	100.00% 33.50%	98.00% 4.17%	83.60% 0.97%	52.40% 0.35%	20.40% 0.20%	6.80% 0.10%	0.40% 0.00%	0.40% 0.00%
10%	96.40% 20.20%	100.00% 19.99%	96.40% 5.09%	79.60% 1.60%	48.40% 0.54%	17.20% 0.24%	0.00% 0.00%	0.00% 0.00%
15%	74.00% 13.94%	95.20% 12.18%	99.60% 13.12%	92.00% 4.48%	68.00% 1.50%	35.60% 0.54%	0.00% 0.00%	0.00% 0.00%
20%	50.00% 9.49%	80.00% 7.19%	92.40% 7.63%	97.20% 8.67%	82.40% 3.42%	56.80% 1.12%	0.00% 0.00%	0.00% 0.00%
25%	26.40% 6.87%	50.40% 4.84%	70.40% 4.56%	84.40% 4.92%	90.00% 5.72%	70.40% 2.36%	1.60% 0.04%	0.80% 0.01%
30%	11.20% 5.27%	27.60% 3.29%	41.60% 2.85%	61.20% 2.68%	74.80% 3.00%	80.00% 3.46%	1.20% 0.15%	0.40% 0.01%
50%	0.00% 0.00%	0.00% 0.00%	0.00% 0.00%	0.40% 0.26%	1.20% 0.28%	1.60% 0.26%	8.40% 1.07%	0.40% 0.01%
75%	0.00% 0.00%	0.00% 0.00%	0.00% 0.00%	0.00% 0.00%	0.40% 0.01%	0.40% 0.01%	0.80% 0.00%	0.40% 0.01%

Table 3 Benefits of randomization in the cardinality-constrained MNL problem. In the table, the rows (columns) correspond to different uncertainty budgets Γ (assortment cardinalities C).

implemented in C++ and run on Intel Xeon 2.20GHz cluster nodes with 16 GB dedicated main memory in four-core mode. Our data sets and detailed results, together with the source codes of all our algorithms, can be found online.[2]

7.1. The Cardinality-Constrained Multinomial Logit Problem

We first consider the cardinality-constrained MNL model where the product prices r_i are selected uniformly at random from the interval $[0, 10]$. We use a budget uncertainty set (*cf.* Section 4.2) where the lower and upper product valuations $\underline{v}_i$ and $\overline{v}_i$, $i \in \mathcal{N}$, are chosen uniformly at random from the intervals $[0, 4]$ and $[6, 10]$, respectively, whereas the valuation of the no-purchase option is fixed at $v_0 = 5$.

Table 3 presents results for $n = 20$ products where the uncertainty budgets Γ (rows) and the assortment cardinalities $|S| \leq C$ (columns) are set to various percentages of n. For each table entry, the first (upper) value denotes the percentage of 250 randomly generated instances in which the optimal randomized assortment outperformed the optimal deterministic robust assortment in terms of worst-case expected revenues, while the second (lower) value reports the average outperformance on those instances. The table shows that the benefits of randomization are most significant when Γ is close to C and both quantities are small relative to the number of products. The latter is intuitive as $C = n$ recovers the randomization-proof unconstrained robust MNL problem while $\Gamma = n$ recovers the randomization-proof nominal MNL problem under the valuations $\boldsymbol{v}^0 = \underline{\boldsymbol{v}}$ (*cf.* Example 3.2 of Rusmevichientong and Topaloglu 2012).

We next investigate the benefits of randomization when the number n of products varies. To this end, we select $C = \Gamma = \lfloor \frac{1}{2}\sqrt{n} \rfloor$, which is in line with the central limit theorem-type uncertainty budget sets

[2] www.doc.ic.ac.uk/∼wwiesema/assortment_opt.zip

proposed by Bandi and Bertsimas (2012). The 'randomized exact objective' column of Table 4 shows that the benefits of randomization (measured here in terms of the average outperformance over the deterministic robust assortment over 250 random problem instances), while decreasing with problem size, is significant for all considered instance sizes. Interestingly, our exact column generation scheme for the randomized problem is actually faster than the cutting plane technique that we implemented for the deterministic robust problem, as the columns 'deterministic exact runtime' and 'randomized exact runtime' in Table 4 reveal. This is caused by the fact that the cutting plane technique requires significantly more iterations (median 14 for $n = 25$ and median 47 for $n = 50$, for example) than the column generation scheme (median 5 for $n = 25$ and median 7 for $n = 50$, with 2-5 primal and dual iterations per main iteration). Since each (main) iteration adds a fractional linear term to the problems that requires an individual reformulation resulting in additional auxiliary variables and big-M constraints, the number of (main) iterations is a key performance indicator for both algorithms.

The columns 'deterministic heuristic' and 'randomized heuristic', finally, report the runtimes and objective values of the ADXOpt heuristic for the deterministic and randomized robust problems, respectively. In both cases, the objective values are again measured relative to the worst-case expected revenues of the exact deterministic robust problem. We see that for both problems, the ADXOpt heuristic performs very well, particularly on larger instance sizes. For the deterministic robust problem, the optimality gaps approach 0% as the instance sizes grow, while for the randomized robust problem, the outperformance of the ADXOpt heuristic approaches the outperformance of the exact solution approach. We thus conclude that one should solve the deterministic and randomized robust problems exactly for small problem sizes, in which case the cutting plane and column generation schemes are fast, while one may want to resort to heuristic solutions for larger problem sizes, where the optimality gap of ADXOpt decreases rapidly.

	deterministic			**randomized**			
	exact	**heuristic**		**exact**		**heuristic**	
n	**runtime**	**objective**	**runtime**	**objective**	**runtime**	**objective**	**runtime**
5	0.02	-6.80%	0.00s	37.95%	0.01s	24.97%	0.00s
10	0.05	-1.83%	0.00s	35.15%	0.03s	32.03%	0.00s
15	0.13	-0.96%	0.00s	35.39%	0.08s	34.01%	0.00s
20	2.30	-0.84%	0.00s	21.30%	1.80s	20.11%	0.03s
25	6.94	-0.41%	0.01s	22.72%	3.58s	21.96%	0.05s
30	18.38	-0.38%	0.01s	21.99%	5.88s	21.15%	0.09s
35	27.18	-0.40%	0.02s	21.63%	8.10s	21.04%	0.12s
40	299.18	-0.23%	0.05s	9.41%	64.21s	8.58%	0.76s
45	592.28	-0.12%	0.07s	7.29%	97.56s	6.71%	1.09s
50	937.29	-0.11%	0.09s	6.09%	171.07s	5.66%	1.57s

Table 4 **Exact and heuristic solutions for the cardinality-constrained MNL problem. All percentages are reported relative to the optimal solutions of the deterministic robust problem.**

We consider the cardinality-constrained MNL model where the product prices r_i are selected uniformly at random from the interval $[0, 10]$. We use a budget uncertainty set (*cf.* Section 4.2) where the lower and upper product valuations $\underline{v}_i$ and $\overline{v}_i$, $i \in \mathcal{N}$, are chosen uniformly at random from the intervals $[0,4]$ and $[6,10]$, respectively, whereas the valuation of the no-purchase option is fixed at $v_0 = 5$.

Table 3 presents results for $n = 20$ products where the uncertainty budgets Γ (rows) and the assortment cardinalities $|S| \leq C$ (columns) are set to various percentages of n. For each table entry, the first (upper) value denotes the percentage of 250 randomly generated instances in which the optimal randomized assortment outperformed the optimal deterministic robust assortment in terms of worst-case expected revenues, while the second (lower) value reports the average outperformance on those instances. The table shows that the benefits of randomization are most significant when Γ is close to C and both quantities are small relative to the number of products. The latter is intuitive as $C = n$ recovers the randomization-proof unconstrained robust MNL problem while $\Gamma = n$ recovers the randomization-proof nominal MNL problem under the valuations $\boldsymbol{v}^0 = \underline{\boldsymbol{v}}$ (*cf.* Example 3.2 of Rusmevichientong and Topaloglu 2012).

We close this section with a data-driven experiment. To this end, we fix $n = 10$ products and consider different cardinalities $C \in \{1,2,3,4\}$ for the admissible assortments. The product prices r_i are generated randomly by the same procedure as before. The true customer valuations are unknown and satisfy $v_i^0 = \mathrm{e}^{\beta_i^0}$, where β_i^0 is drawn uniformly at random from -3 to 3, $i \in \mathcal{N}$; we fix $v_0^0 = \mathrm{e}^0$ and assume that this quantity is known. We assume that 5, 10, $\cdots$, 95 historical samples of random assortments of cardinality C are available, together with the purchase choice that each customer made. In the nominal model, we then estimate the valuations $\hat{\boldsymbol{v}}$ from the historical data using a maximum likelihood estimation, and we solve the resulting nominal assortment optimization problem. In the deterministic and randomized robust models, we employ the budget uncertainty set where both $\Gamma \in \{0, \cdots, n\}$ and $(\underline{\boldsymbol{v}}, \overline{\boldsymbol{v}}) = ([1-\gamma]\hat{\boldsymbol{v}}, [1+\gamma]\hat{\boldsymbol{v}})$, $\gamma \in \{0, 0.025, \cdots, 0.5\}$ are selected using 7-fold cross-validation on the available historical data. Since the validation data reports the customer choices for assortments that differ from the assortment computed using the training data, each iteration of our cross-validation proceeds in two steps. We first use the validation set to perform a maximum likelihood estimation of the valuations. We subsequently use the nominal MNL model associated with these estimated valuations to approximate the out-of-sample revenues of the assortment computed using the training data. We then compare the actual out-of-sample expected revenues of the three models under newly generated data using the true valuations $\boldsymbol{v}^0$.

For each cardinality $C \in \{1,2,3,4\}$ and sample size 5, 10, $\cdots$, 95, Figure 2 reports the average optimality gaps (over 100 randomly generated instances) of the out-of-sample expected revenues of each model relative to the expected revenues of the nominal model under the true valuations. The figure shows that these optimality gaps decrease with sample size (*i.e.*, from left to right in each graph) as well as the cardinality of the assortment (*i.e.*, from the leftmost to the rightmost graph). This is expected as in both cases, the estimation problem can rely on more data and hence becomes easier. In all cases, the randomized robust model outperforms the deterministic robust model, which in turn outperforms the nominal model. We emphasize

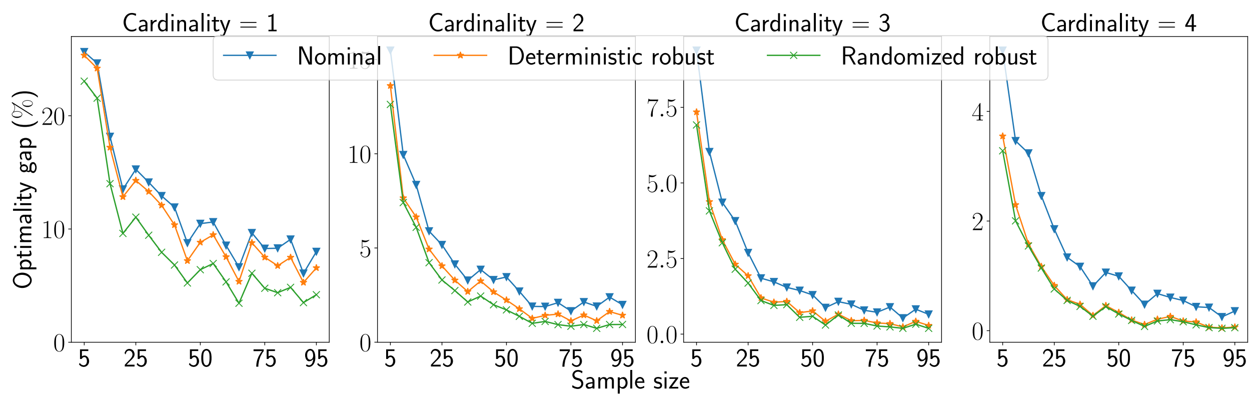

Figure 2 Data-driven experiment for the MNL problem with cardinality constraint $C = 1, 2, 3, 4$, from left to right. All optimality gaps are reported relative to the expected revenues of the clairvoyant model that knows the true customer valuations.

that this is not *a priori* obvious as the deterministic and randomized robust models use the available historical data for both estimation and parameter selection, and hence have less data than the nominal model to estimate $\hat{\boldsymbol{v}}$. Interestingly, for small cardinalities—where the estimation problem is most challenging—the randomized robust model significantly outperforms both the nominal and the deterministic robust model. For larger cardinalities, the performance of the deterministic and randomized robust models are becoming more similar. It is noteworthy, however, that the randomized robust model is never performing worse than the deterministic robust model (in terms of average optimality gap), while it is at the same time easier to solve (*cf.* Table 4).

7.2. Experiments with Real-World Data

In this section, we demonstrate the value of randomized strategies on real data. To this end, we use the data set of Bertsimas and Mišić (2017, 2019), which is derived from a field study on consumer preferences for hypothetical Timbuk2 laptop bag designs by Toubia et al. (2003) (see also Belloni et al. 2008). The data set contains 330 preference rankings for 3,584 bag designs, as well as the revenues associated with each design.[3] Bertsimas and Mišić (2017) use the data to study the value of accounting for uncertainty in deterministic robust assortment optimization under the PR model (among others). Here, we extend their approach to the randomized robust assortment optimization problem.

Bertsimas and Mišić (2017) solve the nominal and deterministic robust problems heuristically for the entire data set. By contrast, we use the approach described in Section 6.2 to solve the nominal as well as the deterministic and randomized robust problems exactly for a subset of available products. Specifically, we pick the $n = 300$ products generating the highest average consumer utilities in the data set and trim the preference ranking lists to only include these products and the no-purchase option. Following Bertsimas and Mišić (2017), we assume that each customer type is equally probable, that is, $\lambda_k = 1/K$ for all $k \in \mathcal{K}$. For

[3] The data set is available at https://github.com/vvmisic/optimalPLD.

		WCL(%)	**RI**$^{\mathbf{D}}$(%)	**RI**$^{\mathbf{R}}$(%)	**RI**$^{\mathbf{RD}}$(%)
θ	**0.1**	6.07	0.15	0.15	0.00
	0.2	12.14	0.34	0.35	0.01
	0.3	18.22	0.56	0.78	0.21
	0.4	24.09	0.76	2.74	1.96
	0.5	29.88	0.88	6.62	5.69
	0.6	35.68	1.02	11.73	10.60

Table 5 Benefits of randomization in the real-world case study over $n = 300$ products. The rows correspond to different radii θ of the uncertainty set.

the deterministic and randomized robust problems, we employ a 1-norm ambiguity set with varying radius θ (*cf.* Section 6.2).

Following Bertsimas and Mišić (2017), we use two metrics to quantify the value of the deterministic and randomized robust approaches: the worst-case loss and the relative improvement. The worst-case loss (WCL) measures the largest possible loss of the optimal nominal assortment S^{N} over the uncertainty set $\mathcal{U}$:

$$\text{WCL} = 100\% \cdot \frac{R^\star_{\text{nom}} - R^\star(S^{\text{N}})}{R^\star_{\text{nom}}},$$

where $R^\star_{\text{nom}}$ is the expected revenues of S^{N} under the nominal occurrence probabilities $\boldsymbol{\lambda}_0$ and $R^\star(S^{\text{N}})$ denotes the worst-case expected revenues of S^{N} over the uncertainty set $\mathcal{U}$. The second metric is the relative improvement (RI) in worst-case expected revenues when using the deterministic or randomized robust approach instead of the nominal assortment:

$$\begin{aligned} \text{RI}^{\text{D}} &= 100\% \cdot \frac{R^\star_{\text{det}}(\mathcal{U}) - R^\star(S^{\text{N}})}{R^\star(S^{\text{N}})} \\ \text{and} \quad \text{RI}^{\text{R}} &= 100\% \cdot \frac{R^\star_{\text{rand}}(\mathcal{U}) - R^\star(S^{\text{N}})}{R^\star(S^{\text{N}})}, \end{aligned}$$

where $R^\star_{\text{det}}(\mathcal{U})$ and $R^\star_{\text{rand}}(\mathcal{U})$ denote the optimal values of the deterministic and randomized robust assortment optimization problems, respectively.

Table 5 reports the worst-case losses and relative improvements for different radii θ. The table shows that the worst-case loss from solving the nominal problem grows with the radius of the uncertainty set, as expected. The deterministic robust approach (column RI^{D}), consistently with Bertsimas and Mišić (2017), benefits the decision maker in terms of the relative improvement, and the magnitude of the improvement increases with the radius of the uncertainty set. The relative improvement achieved by the randomized robust strategy (RI^{R}) is significantly higher, however, demonstrating the additional benefit of randomization in hedging against uncertainty in the preference rankings. The difference in relative improvements between the two robust approaches is reported in the column RI^{RD}, which shows that the additional gains of the randomized approach are substantial and grow with the size of the uncertainty set. The results thus indicate the potential promise of employing randomization in robust assortment optimization on real-world data.

8. Concluding Remarks

Our results call for more research into the importance of randomization for revenue management. We have seen that a firm can benefit from using randomization—instead of the standard approach of deterministically offering a single assortment—in the robust assortment optimization problem. However, it is not *a priori* obvious which problem settings may benefit from randomization: two similar models, such as the MC model with product-wise and general uncertainty sets, may lead to starkly different conclusions. While our analysis suggests that more general versions of the assortment optimization problem tend to be more receptive to the benefits of randomization, more research is needed into this potential benefit under other choice models and constraints, as well as other revenue management problems. Also, all of the computational approaches developed in this paper study the static problem formulation where the decision maker does not update her knowledge about the uncertain problem parameters. As a next step, it would be instructive to study optimal dynamic randomization strategies where the decision maker learns about the uncertain problem parameters over time.

Acknowledgments

The first author gratefully acknowledges funding from the Imperial College President's PhD Scholarship programme. The last author gratefully acknowledges funding from the Engineering and Physical Sciences Research Council under the grants EP/R045518/1, EP/T024712/1 and EP/W003317/1. All three authors are indebted to the area editor (Gustavo Vulcano) as well as the anonymous associate editor and referees, whose insightful comments have led to substantial improvements of the manuscript.

Biography

Zhengchao Wang is a Ph.D. candidate in Analytics & Operations at Imperial College Business School, focusing on optimization under uncertainty, statistical learning, and their applications in revenue management and operations for social goods. Before joining Imperial College, he earned a master's degree in energy management and sustainability, specializing in renewable energy utilization.

Heikki Peura is an Associate Professor in the information and service management department at Aalto University School of Business. Before joining Aalto, he was a faculty member at Imperial College Business School. He received a PhD in management science and operations from London Business School and a MSc in engineering physics and mathematics from Aalto University. His research focuses on analytics for operations management with applications in energy, health care, supply chain management and revenue management.

Wolfram Wiesemann is a professor of Analytics & Operations at Imperial College Business School. His research interests evolve around decision-making under uncertainty, with applications to supply chain management, healthcare and energy. Wolfram is an elected member of the boards of the Mathematical Optimization Society and the Stochastic Programming Society, and he currently serves as editor in-chief of Operations Research Letters as well as a department co-editor for Management Science.

References

Aouad A, Farias V, Levi R (2021) Assortment optimization under consider-then-choose choice models. *Management Science* 67(6):3368–3386.

Aouad A, Farias V, Levi R, Segev D (2018) The approximability of assortment optimization under ranking preferences. *Operations Research* 66(6):1661–1669.

Bandi C, Bertsimas D (2012) Tractable stochastic analysis in high dimensions via robust optimization. *Mathematical Programming* 134:23–70.

Baron O, Milner J, Naseraldin H (2011) Facility location: A robust optimization approach. *Production and Operations Management* 20(5):772–785.

Bayraksan G, Love DK (2015) Data-driven stochastic programming using phi-divergences. *INFORMS Tutorials in Operations Research* 1–19.

Belloni A, Freund R, Selove M, Simester D (2008) Optimizing product line designs: Efficient methods and comparisons. *Management Science* 54(9):1544–1552.

Ben-Tal A, El Ghaoui L, Nemirovski A (2009) *Robust Optimization* (Princeton University Press).

Bertsimas D, Brown DB, Caramanis C (2011) Theory and applications of robust optimization. *SIAM Review* 53(3):464–501.

Bertsimas D, Caramanis C (2010) Finite adaptability in multistage linear optimization. *IEEE Transactions on Automatic Control* 55(12):2751–2766.

Bertsimas D, Mišić V (2017) Robust product line design. *Operations Research* 65(1):19–37.

Bertsimas D, Mišić V (2019) Exact first-choice product line optimization. *Operations Research* 67(3):651–670.

Bertsimas D, Nasrabadi E, Orlin JB (2016) On the power of randomization in network interdiction. *Operations Research Letters* 44(1):114–120.

Bertsimas D, Sim M (2004) The price of robustness. *Operations Research* 52(1):35–53.

Bertsimas D, Thiele A (2006) A robust optimization approach to inventory theory. *Operations Research* 54(1):150–168.

Billingsley P (1961) *Statistical Inference for Markov Processes* (The University of Chicago Press).

Birbil Şİ, Frenk J, Gromicho JA, Zhang S (2009) The role of robust optimization in single-leg airline revenue management. *Management Science* 55(1):148–163.

Bishop CM (2006) *Pattern Recognition and Machine Learning* (Springer).

Blanchet J, Gallego G, Goyal V (2016) A Markov chain approximation to choice modeling. *Operations Research* 64(4):886–905, ISSN 0030-364X.

Bumpensanti P, Wang H (2020) A re-solving heuristic with uniformly bounded loss for network revenue management. *Management Science* 66(7):2993–3009.

Davis J, Gallego G, Topaloglu H (2013) Assortment planning under the multinomial logit model with totally unimodular constraint structures, working paper, Cornell University, New York.

Davis J, Gallego G, Topaloglu H (2014) Assortment optimization under variants of the nested logit model. *Operations Research* 62(2):250–273, ISSN 0030-364X.

Delage E, Kuhn D, Wiesemann W (2019) "Dice"-sion–making under uncertainty: When can a random decision reduce risk? *Management Science* 65(7):3282–3301, ISSN 0025-1909.

Delage E, Saif A (2022) The value of randomized solutions in mixed-integer distributionally robust optimization problems. *INFORMS Journal on Computing* 34(1):333–353.

DeMiguel V, Nogales FJ (2009) Portfolio selection with robust estimation. *Operations Research* 57(3):560–577.

Désir A, Goyal V, Jiang B, Xie T, Zhang J (2023) Robust assortment optimization under the Markov chain choice model. *Operations Research* Forthcoming.

Désir A, Goyal V, Segev D, Ye C (2020) Constrained assortment optimization under the Markov chain–based choice model. *Management Science* 66(2):698–721.

El Ghaoui L, Lebret H (1997) Robust solutions to least-squares problems with uncertain data. *SIAM Journal on Matrix Analysis and Applications* 18(4):1035–1064.

Farias V, Jagabathula S, Shah D (2013) A nonparametric approach to modeling choice with limited data. *Management Science* 59(2):305–322, ISSN 0025-1909.

Feldman JB, Topaloglu H (2017) Revenue management under the Markov chain choice model. *Operations Research* 65(5):1322–1342, ISSN 0030-364X.

Ferreira KJ, Simchi-Levi D, Wang H (2018) Online network revenue management using Thompson sampling. *Operations Research* 66(6):1586–1602.

Gallego G, Topaloglu H (2019) *Revenue Management and Pricing Analytics* (Springer).

Goldfarb D, Iyengar G (2003) Robust portfolio selection problems. *Mathematics of Operations Research* 28(1):1–38.

Hanasusanto G, Kuhn D, Wiesemann W (2015) K-adaptability in two-stage robust binary programming. *Operations Research* 63(4):877–891.

Hastie T, Tibshirani R, Friedman J (2009) *The Elements of Statistical Learning: Data Mining, Inference, and Prediction* (Springer).

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

Jasin S (2015) Performance of an LP-based control for revenue management with unknown demand parameters. *Operations Research* 63(4):909–915.

Jasin S, Kumar S (2012) A re-solving heuristic with bounded revenue loss for network revenue management with customer choice. *Mathematics of Operations Research* 37(2):313–345.

Kök AG, Fisher ML, Vaidyanathan R (2015) Assortment planning: Review of literature and industry practice. Agrawal N, Smith S, eds., *Retail Supply Chain Management: Quantitative Models and Empirical Studies*, 175–236 (Springer), ISBN 978-1-4899-7562-1.

Luce RD (1959) *Individual choice behavior: A theoretical analysis* (Courier Corporation).

Ma W (2023) When is assortment optimization optimal? *Management Science* Forthcoming.

Mahajan S, van Ryzin G (2001) Stocking retail assortments under dynamic consumer substitution. *Operations Research* 49(3):334–351.

Mak HY, Rong Y, Zhang J (2015) Appointment scheduling with limited distributional information. *Management Science* 61(2):316–334.

Michaud RO (1989) The Markowitz optimization enigma: Is 'optimized' optimal? *Financial Analysts Journal* 45(1):31–42.

Paul A, Feldman J, Davis J (2018) Assortment optimization and pricing under a nonparametric tree choice model. *Manufacturing & Service Operations Management* 20(3):550–565.

Perakis G, Roels G (2010) Robust controls for network revenue management. *Manufacturing & Service Operations Management* 12(1):56–76.

Plackett RL (1975) The analysis of permutations. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 24(2):193–202.

Reiman MI, Wang Q (2008) An asymptotically optimal policy for a quantity-based network revenue management problem. *Mathematics of Operations Research* 33(2):257–282.

Rusmevichientong P, Shen ZJM, Shmoys D (2010) Dynamic assortment optimization with a multinomial logit choice model and capacity constraint. *Operations Research* 58(6):1666–1680.

Rusmevichientong P, Shmoys D, Tong C, Topaloglu H (2014) Assortment optimization under the multinomial logit model with random choice parameters. *Production and Operations Management* 23(11):2023–2039, ISSN 10591478.

Rusmevichientong P, Topaloglu H (2012) Robust assortment optimization in revenue management under the multinomial logit choice model. *Operations Research* 60(4):865–882.

Rusmevichientong P, Van Roy B, Glynn PW (2006) A nonparametric approach to multiproduct pricing. *Operations Research* 54(1):82–98.

Şimşek AS, Topaloglu H (2018) An expectation-maximization algorithm to estimate the parameters of the Markov chain choice model. *Operations Research* 66(3):748–760.

Smith JE, Winkler RL (2006) The optimizer's curse: Skepticism and postdecision surprise in decision analysis. *Management Science* 52(3):311–322.

Talluri K, van Ryzin G (2004) Revenue management under a general discrete choice model of consumer behavior. *Management Science* 50(1):15–33, ISSN 0025-1909.

Talluri K, van Ryzin G (2006) *The Theory and Practice of Revenue Management* (Springer).

Toubia O, Simester DI, Hauser JR, Dahan E (2003) Fast polyhedral adaptive conjoint estimation. *Marketing Science* 22(3):273–303.

van Ryzin G, Vulcano G (2015) A market discovery algorithm to estimate a general class of nonparametric choice models. *Management Science* 61(2):281–300.

van Ryzin G, Vulcano G (2017) An expectation-maximization method to estimate a rank-based choice model of demand. *Operations Research* 65(2):396–407.

Wiesemann W, Kuhn D, Rustem B (2013) Robust Markov decision processes. *Mathematics of Operations Research* 38(1):153–183, ISSN 0364-765X.

Williams HC (1977) On the formation of travel demand models and economic evaluation measures of user benefit. *Environment and Planning A* 9(3):285–344.

Xu H, Caramanis C, Mannor S (2009) Robust regression and Lasso. Koller D, Schuurmans D, Bengio Y, Bottou L, eds., *Advances in Neural Information Processing Systems 21*, 1801–1808 (Curran Associates, Inc.).

Zhang D, Cooper WL (2005) Revenue management for parallel flights with customer-choice behavior. *Operations Research* 53(3):415–431, ISSN 0030-364X.