Randomized Assortment Optimization

Zhengchao Wang, Heikki Peura and Wolfram Wiesemann
Imperial College Business School, Imperial College London, London, UK

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 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.

1. Introduction

Selecting an assortment of products to offer to customers is a central problem across business operations, with manifold applications in the travel and hospitality industries, online marketing, as well as brick-and-mortar and online retail. 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. 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 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. 2019). 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 in view of 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 demonstrate this potential benefit of randomization across three popular choice models: the multinomial logit model, the Markov chain model, and the preference ranking model.

We first consider the ubiquitous multinomial logit model with and without a cardinality constraint on the offered assortment. In the unconstrained MNL model, we show that 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. We also show how an optimal randomization can be determined exactly (by a column generation scheme) and heuristically (through a local search algorithm).

The Markov chain (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, and they can again become unbounded in the cardinality-constrained setting.

The preference ranking 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 preference ranking 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 model, where all optimal randomization strategies can become arbitrarily complex. Again, we show how an optimal randomization for the preference ranking model can be determined exactly (by a column generation scheme) and heuristically (through a local search algorithm).

We illustrate the runtimes of our exact and heuristic solution schemes, as well as the potential benefits of randomization, on synthetic instances of the cardinality constrained MNL and preference ranking models. For the former 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. To facilitate reproducibility of our results, the source codes of our algorithms as well as all data sets are made available online.

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.

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. Sections 4–6 study the benefits of randomization, as well as exact and heuristic schemes to determine randomized assortments, under the MNL, the MC and the preference ranking models. Section 7 presents numerical results, and Section 8 concludes the paper. For ease of exposition, all proofs are relegated to the appendix.

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. We let $\Delta = \{\boldsymbol{p} \in \mathbb{R}^m_+ : \sum_{i=1}^m p_i = 1\}$ denote the probability simplex in $\mathbb{R}^m$, where the dimension m will become clear from the context. 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 (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 (2018) 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 preference ranking models (Aouad et al. 2018), special cases (Honhon et al. 2012, Aouad et al. 2015, Paul et al. 2018) 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 in view of the worst parameter setting deemed plausible, the robust approach hedges against overfitting 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. (2019), and Bertsimas and Mišić (2017) study the robust assortment optimization problem under the MNL, Markov chain, and preference ranking 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. (2019) extend this robust approach to the Markov chain model and develop efficient algorithms to solve it. Bertsimas and Mišić (2017) consider the related problem of robust product line design under a preference ranking model with both parameter and structural uncertainty.

In all 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 for randomization to benefit the decision maker arises from applying robust optimization to a discrete optimization problem (Bertsimas et al. 2016, Delage and Saif 2018, Delage et al. 2019). 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 a wider discussion of the potential benefits of randomized strategies in operational problems, we also contribute to the robust optimization literature.

3. The Robust Assortment Optimization Problem

In this section, we define the robust assortment optimization problem, which we subsequently study under three different choice models in the following sections.

We assume that a firm chooses an assortment to offer to customers out of n candidate products $\mathcal{N} = \{1, 2, \cdots, n\}$. A customer either buys one of the offered products or makes no purchase, which we indicate by the index 0. We denote by $\mathcal{N}_0 = \mathcal{N} \cup \{0\}$ the set of all options including the products and the no-purchase alternative. The revenue of product $i \in \mathcal{N}$ is r_i; we adopt the convention that the revenue of the no-purchase option is $r_0 = 0$ and include it in the vector of product revenues $\boldsymbol{r} \in \mathbb{R}_+^{n+1}$.

The firm's expected revenues when offering an assortment $S \subseteq \mathcal{N}$ are $R(S, \boldsymbol{u})$, where $\boldsymbol{u}$ is a vector of parameters of the customer choice model. For ease of notation, we suppress the dependence of R on the product revenues $\boldsymbol{r}$, which we assume to be deterministic throughout the paper. If the parameter values are known to be $\boldsymbol{u}^0$, the firm solves the classical assortment optimization problem

$$R^*_{\text{nom}}(\boldsymbol{u}^0) \;=\; \max_{S \subseteq \mathcal{N}} R(S, \boldsymbol{u}^0). \tag{Nominal}$$

Here and in the following, the set of admissible assortments $S \subseteq \mathcal{N}$ may be further restricted, in particular by a cardinality constraint on the size of the assortment, $|S| \leq C$.

In reality, the choice model parameters $\boldsymbol{u}$ are typically uncertain. In the deterministic robust assortment optimization problem (see, *e.g.*, Rusmevichientong and Topaloglu 2012), the firm estimates a compact uncertainty set $\mathcal{U}$ that contains the unknown true parameter values $\boldsymbol{u}^0$ with high confidence and chooses a *single* assortment that maximizes its worst-case expected revenues over this uncertainty set:

$$R^*_{\text{det}}(\mathcal{U}) \;=\; \max_{S \subseteq \mathcal{N}} \min_{\boldsymbol{u} \in \mathcal{U}} R(S, \boldsymbol{u}). \tag{Deterministic Robust}$$

We denote by $S^*(\mathcal{U})$ an optimal assortment in this problem, which is not necessarily unique. When the uncertainty set $\mathcal{U} = \{\boldsymbol{u}^0\}$ is a singleton, the parameters are known exactly and Deterministic Robust recovers the classical Nominal problem.

Instead of selecting a single assortment to offer, the firm could offer *each* assortment $S \subseteq \mathcal{N}$ with a probability $p_S \geq 0$. The corresponding problem is then

$$R^*_{\text{rand}}(\mathcal{U}) \;=\; \max_{\boldsymbol{p} \in \Delta} \min_{\boldsymbol{u} \in \mathcal{U}} \sum_{S \subseteq \mathcal{N}} p_S \cdot R(S, \boldsymbol{u}), \tag{Randomized Robust}$$

where Δ is the probability simplex of dimension 2^n. We define $\boldsymbol{p}^*$ to be an optimal probability distribution over the assortments in this problem, which again may not be unique.

Below, we compare the solutions to the problems Deterministic Robust and Randomized Robust under different choice models: the multinomial logit model, the Markov chain model, and the preference ranking model. Each of these models implies different definitions for R, $\boldsymbol{u}$ and $\mathcal{U}$. Our goal is to show under what conditions randomization can benefit the firm, that is, when the strict inequality $R^*_{\text{rand}}(\mathcal{U}) > R^*_{\text{det}}(\mathcal{U})$ is satisfied. We say such versions of the problem are receptive to randomization, as per the following definition.

Definition 1 (randomization-receptiveness/proofness). The robust assortment optimization problem under a particular choice model is *randomization-receptive* if there exist instances $(\mathcal{U}, \boldsymbol{r})$ of the problem for which $R^*_{\text{rand}}(\mathcal{U}) > R^*_{\text{det}}(\mathcal{U})$. Otherwise, the problem is *randomization-proof.*

We note that Nominal is randomization-proof by construction. For singleton uncertainty sets $\mathcal{U} = \{\boldsymbol{u}^0\}$, Randomized Robust reduces to a linear program that attains its optimal value at an extreme point of the probability simplex Δ. This in turn corresponds to a randomization that places unit probability onto a single assortment $S \subseteq \mathcal{N}$.

4. Multinomial Logit (MNL) 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 assortment choice. The unconstrained problem is randomization-proof (Section 4.1), while the cardinality-constrained problem is not only randomization-receptive, but the resulting benefit can also be arbitrarily large (Section 4.2).

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

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

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

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

where we identify the parameter vector $\boldsymbol{u}$ with the preference weights $\boldsymbol{v}$. 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. The Unconstrained MNL Model

It is well known that under the unconstrained nominal MNL model, every optimal assortment is revenue-ordered, that is, an optimal assortment always consists of precisely those products whose revenues exceed a certain threshold (Talluri and van Ryzin 2004). Rusmevichientong and Topaloglu (2012) showed that the optimality of revenue-ordered assortments is preserved under the deterministic robust MNL model. In the following, we leverage this result to derive a 'no-regret' property of optimal assortments that will be key to prove the randomization-proofness of the associated assortment optimization problem.

LEMMA 1. *The* DETERMINISTIC ROBUST *variant of the unconstrained MNL model is solved by revenue-ordered assortments, that is, the set of optimal assortments $\mathcal{S}^*(\mathcal{V})$ to this problem contains any assortment $S \subseteq \mathcal{N}$ satisfying*

$$\{i \in \mathcal{N} : r_i > R^*_{det}(\mathcal{V})\} \subseteq S \subseteq \{i \in \mathcal{N} : r_i \geq R^*_{det}(\mathcal{V})\}.$$

Moreover, the ex-ante and ex-post worst-case valuations of any optimal assortment $S^*(\mathcal{V}) \in \mathcal{S}^*(\mathcal{V})$ coincide, that is, we have

$$R(S^*(\mathcal{V}), \boldsymbol{v}^*) = R(S^*(\boldsymbol{v}^*), \boldsymbol{v}^*),$$

where $\boldsymbol{v}^* \in \arg\min_{\boldsymbol{v} \in \mathcal{V}} R(S^*(\mathcal{V}), \boldsymbol{v})$ and $S^*(\boldsymbol{v}^*) \in \arg\max_{S \subseteq \mathcal{N}} R(S, \boldsymbol{v}^*)$.

An immediate consequence of Lemma 1 is the following strong duality property.

COROLLARY 1. *The robust assortment optimization problem under the unconstrained MNL model satisfies strong duality, that is, we have $\max_{S \subseteq \mathcal{N}} \min_{\boldsymbol{v} \in \mathcal{V}} R(S, \boldsymbol{v}) = \min_{\boldsymbol{v} \in \mathcal{V}} \max_{S \subseteq \mathcal{N}} R(S, \boldsymbol{v})$.*

The no-regret property of Lemma 1 shows that by committing to a worst-case optimal assortment $S^*(\mathcal{V}) \in \mathcal{S}^*(\mathcal{V})$, the decision maker does not regret her choice ex post whenever one of the worst-case valuations $\boldsymbol{v}^* \in \arg\min_{\boldsymbol{v} \in \mathcal{V}} R(S^*(\mathcal{V}), \boldsymbol{v})$ materializes. Since the worst-case optimal assortments are not point-wise optimal for every valuation $\boldsymbol{v} \in \mathcal{V}$, however, she may regret her choice if a valuation different from the worst-case valuations materializes. The no-regret property implies (but is in general not implied by) the strong duality in Corollary 1. We note that strong duality for the unconstrained MNL model has also been shown by Désir et al. (2019) through the consideration of the MNL model as a special case of the Markov chain choice model. By contrast, we derive the result directly from the optimality of revenue-ordered assortments in the MNL model.

Using the no-regret property of Lemma 1, we now prove that the robust assortment optimization problem under the unconstrained MNL model is randomization-proof.

Theorem 1. *The robust assortment optimization problem under the unconstrained MNL model is randomization-proof, that is, there is no* $(\mathcal{V}, \boldsymbol{r})$ *such that* $R^*_{rand}(\mathcal{V}) > R^*_{det}(\mathcal{V})$.

The intuition behind the proof of Theorem 1 is as follows. The no-regret property of Lemma 1 implies that the decision maker cannot benefit from randomization. Indeed, under any worst-case valuation $\boldsymbol{v}^* \in \arg\min_{\boldsymbol{v}\in\mathcal{V}} R(S^*(\mathcal{V}), \boldsymbol{v})$ of an optimal deterministic assortment $S^*(\mathcal{V}) \in \mathcal{S}^*(\mathcal{V})$, the expected revenues of $S^*(\mathcal{V})$ weakly dominate the expected revenues of all other assortments $S \subseteq \mathcal{N}$ simultaneously. Thus, randomizing between assortments, which for any fixed $\boldsymbol{v}$ results in convex combinations of the expected revenues of the involved assortments, cannot benefit the decision maker since the worst-case expected revenues are by definition bounded above by the expected revenues under the valuation $\boldsymbol{v}^*$.

4.2. The Cardinality-Constrained MNL Model

The randomization-proofness of the unconstrained MNL model relies on the no-regret property of Lemma 1, which in turn leverages the optimality of revenue-ordered assortments in the deterministic robust MNL problem. As Rusmevichientong et al. (2010) show, revenue-ordered assortments cease to be optimal in the constrained version of the nominal MNL problem where we impose a cardinality constraint $|S| \leq C$ on the size of the offered assortment. The next example demonstrates that as a result of this, the cardinality-constrained robust MNL problem is randomization-receptive.

Example 1. Consider the robust MNL problem with three products, $r_1 = r_2 = r_3 = 10$, the cardinality constraint $|S| \leq 2$, and an uncertainty set $\mathcal{V}$ that comprises the three valuation vectors $\boldsymbol{v}^1 = (1,1,1,2)$, $\boldsymbol{v}^2 = (1,1,2,1)$, and $\boldsymbol{v}^3 = (1,2,1,1)$. For the deterministic assortment $\{1\}$, a worst-case scenario is valuation $\boldsymbol{v}^1$, with expected revenues $1 \cdot 10/(1+1) = 5$. Both $\{2\}$ and $\{3\}$ also have worst-case expected revenues 5, while the two-product assortments $\{1,2\}$, $\{1,3\}$ and $\{2,3\}$ result in $R^*_{\text{det}}(\mathcal{V}) = \frac{20}{3} \approx 6.67$. However, randomizing between $\{1,2\}$, $\{1,3\}$ and $\{2,3\}$ with equal probabilities $1/3$ results in the higher worst-case expected revenues $R^*_{\text{rand}}(\mathcal{V}) = \frac{20}{3} + \frac{5}{18} = \frac{125}{18} \approx 6.94$. In the unconstrained problem it would be optimal to offer the assortment $\{1,2,3\}$, which gives worst-case expected revenues of 8.

We can interpret the perhaps surprising finding of Example 1 from the complementary perspectives of diversification in the presence of estimation errors and game theory.

Under the first interpretation, each of the deterministic assortments exposes the decision maker to significant estimation risk: while the assortment $\{1,2\}$, say, produces high expected revenues under the favorable valuation scenarios $\boldsymbol{v}^2$ and $\boldsymbol{v}^3$, it results in substantially lower expected revenues under the adverse scenario $\boldsymbol{v}^1$. By randomizing between the three assortments, the decision maker can diversify this risk and alleviate the error maximizing effect of optimization. Notice that such hedging is only valuable under a cardinality constraint $|S| \leq C$. In the unconstrained MNL model, the no-regret property of Lemma 1 ensures that the expected revenues of an optimal assortment $S^*(\mathcal{V})$ under any worst-case valuation $\boldsymbol{v}^*$ weakly dominate the expected revenues of all other assortments $S \subseteq \mathcal{N}$, thus precluding any diversification benefits.

Under the second interpretation, we can regard the robust assortment optimization problem as a Stackelberg leader-follower game. In this game, the decision maker (leader) selects an assortment S, after which 'nature' (follower) responds with the most adverse valuation scenario from within the uncertainty set $\mathcal{V}$. A Stackelberg leader may benefit from randomized strategies when the follower is *oblivious* (as shown in the literature on security games, see Korzhyk et al. 2011, Mastin et al. 2015, An et al. 2016 and Bertsimas et al. 2016). That is, rather than observing the actually implemented decision (*i.e.*, the assortment S), the follower can only observe the probabilities with which different decisions are selected (*i.e.*, the randomization weights $\boldsymbol{p}$). Under each deterministic assortment in Example 1, the follower can choose an adverse parameter realization $\boldsymbol{v}$ under which the customer is likely to exercise her no-purchase option, resulting in low worst-case expected revenues. Anticipating these actions, the leader can mitigate them by using a randomized strategy to a significant worst-case benefit.

We now show that the benefit from randomization can indeed be arbitrarily large.

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. $R^*_{det}(\mathcal{V}) = 0$ *while* $R^*_{rand}(\mathcal{V}) > 0$*;*
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 proof of Theorem 2, which is relegated to the Appendix, considers a class of robust constrained MNL instances where all products have equal revenues and the uncertainty set $\mathcal{V}$ comprises all valuation scenarios $\boldsymbol{v}$ 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 scenario $\boldsymbol{v}$ 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 parameter realization.

An immediate consequence of the first statement in Theorem 2 is the following.

Corollary 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 the benefits $R^*_{rand}(\mathcal{V})/R^*_{det}(\mathcal{V})$ from randomization are arbitrarily large.*

Setting $C = n/2$ in the second statement in Theorem 2, we arrive at the following result.

Corollary 3. *For any number of products $n \geq 2$, there are instances of the cardinality-constrained robust MNL problem where the unique optimal randomized assortment strategy randomizes between $\Theta(2^n/\sqrt{n})$ many assortments.*

Corollary 3 shows that the optimal randomization strategy may be very complex in that it requires randomization between an exponentially large number of assortments.

4.3. Solving the Randomized Constrained MNL Problem

We next present two algorithms to compute randomization strategies for the constrained robust MNL problem with a binary representable uncertainty set of the form

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

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 representatives 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

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

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

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

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 $\mathcal{N}^C = \{S \in \mathcal{N} : |S| \leq C\}$ the set of assortments whose cardinality is less or equal to the size restriction C. The problem can then be written as

$$\begin{array}{ll} \text{maximize} & \displaystyle\min_{\boldsymbol{v}\in\mathcal{V}} \sum_{S\in\mathcal{N}^C} p_S \cdot \frac{\sum_{i\in S} r_i v_i}{v_0 + \sum_{i\in S} v_i} \\ \text{subject to} & \displaystyle\sum_{S\in\mathcal{N}^C} p_S = 1 \\ & p_S \geq 0,\ S \in \mathcal{N}^C. \end{array} \tag{3}$$

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 (3) amounts to the robust linear program

$$\begin{array}{ll} \text{minimize} & \displaystyle\max_{S\in\mathcal{N}^C} \sum_{\boldsymbol{v}\in\mathcal{V}} \lambda_{\boldsymbol{v}} \cdot \frac{\sum_{i\in S} r_i v_i}{v_0 + \sum_{i\in S} v_i} \\ \text{subject to} & \displaystyle\sum_{\boldsymbol{v}\in\mathcal{V}} \lambda_{\boldsymbol{v}} = 1 \\ & \lambda_{\boldsymbol{v}} \geq 0,\ \boldsymbol{v} \in \mathcal{V}. \end{array} \tag{4}$$

Strong duality between (3) and (4) holds since (3) is feasible by construction. Similar to problem (3), problem (4) is computationally challenging due to the typically exponential number of decision variables as well as the maximization over a discrete 'uncertainty set' $\mathcal{N}^C$ whose decision variables S appear in a non-convex objective function.

In the following, we propose an exact column generation scheme (Section 4.3.1) as well as a heuristic local search method (Section 4.3.2) to solve this problem.

4.3.1. A Column Generation Scheme Our exact solution scheme for the randomized constrained MNL problem iteratively refines conservative approximations of the primal and dual problems (3) and (4). The approach resembles the two-layer column generation scheme in Section 5.3 of Delage and Saif (2018) and is described in the following.

Column generation scheme for problem (3)

1. **Initialization.** Set $\hat{\mathcal{N}}^C$ to any non-empty subset of $\mathcal{N}^C$ (currently considered assortments). Set $\hat{\mathcal{V}}$ to any non-empty subset of $\mathcal{V}$ (currently considered valuations). Set $\mathrm{LB} = -\infty$ and $\mathrm{UB} = +\infty$ (current lower and upper bounds on the optimal value).
2. **Primal Step.** Solve the restricted primal problem (3) that replaces $\mathcal{N}^C$ with $\hat{\mathcal{N}}^C$:
   (a) Set $\mathrm{LB_p} = -\infty$ and $\mathrm{UB_p} = \mathrm{UB}$ (lower and upper bounds on the optimal value of the restricted primal problem). Set $\boldsymbol{p}$ to any point in the probability simplex over $\hat{\mathcal{N}}^C$ (current restricted randomization strategy).
   (b) Solve the subproblem
$$\begin{aligned} \text{minimize} \quad & \sum_{S \in \hat{\mathcal{N}}^C} p_S \cdot \frac{\sum_{i \in S} r_i v_i}{v_0 + \sum_{i \in S} v_i} \\ \text{subject to} \quad & \boldsymbol{v} \in \mathcal{V}. \end{aligned} \tag{5}$$
   Update $\mathrm{LB_p} = \max\{\mathrm{LB_p}, \tau^*\}$, where τ^* is the optimal value of the subproblem, and $\hat{\mathcal{V}} = \hat{\mathcal{V}} \cup \{\boldsymbol{v}^*\}$, where $\boldsymbol{v}^*$ is the optimal solution of the subproblem.
   (c) Solve the doubly restricted primal problem (3) that replaces $\mathcal{N}^C$ with $\hat{\mathcal{N}}^C$ and $\mathcal{V}$ with $\hat{\mathcal{V}}$. Update $\mathrm{UB_p} = \min\{\mathrm{UB_p}, \tau^*\}$, where τ^* is the optimal value of the doubly restricted problem, and set $\boldsymbol{p}$ to the optimal solution of the problem.
   (d) If $\mathrm{UB_p} - \mathrm{LB_p} < \epsilon/2$, set $\mathrm{LB} = \mathrm{LB_p}$ and go to Step 3. Otherwise, go to Step 2(b).
3. **Dual Step.** Solve the restricted dual problem (4) which replaces $\mathcal{V}$ with $\hat{\mathcal{V}}$:
   (a) Set $\mathrm{LB_d} = \mathrm{LB}$ and $\mathrm{UB_d} = +\infty$ (lower and upper bounds on the optimal value of the restricted dual problem). Set $\boldsymbol{\lambda}$ to any point in the probability simplex over $\hat{\mathcal{V}}$ (convex combination of the currently considered valuations).

(b) Solve the subproblem

$$\begin{array}{ll} \text{maximize} & \displaystyle\sum_{\boldsymbol{v}\in\hat{\mathcal{V}}} \lambda_{\boldsymbol{v}} \cdot \frac{\sum_{i\in S} r_i v_i}{v_0 + \sum_{i\in S} v_i} \\ \text{subject to} & S \in \mathcal{N}^C. \end{array} \tag{6}$$

Update $\mathrm{UB_d} = \min\{\mathrm{UB_d}, \tau^*\}$, where τ^* is the optimal value of the subproblem, and $\hat{\mathcal{N}}^C = \hat{\mathcal{N}}^C \cup \{S^*\}$, where S^* is the optimal solution of the subproblem.

(c) Solve the doubly restricted dual problem (4) that replaces $\mathcal{N}^C$ with $\hat{\mathcal{N}}^C$ and $\mathcal{V}$ with $\hat{\mathcal{V}}$. Update $\mathrm{LB_d} = \max\{\mathrm{LB_d}, \tau^*\}$, where τ^* is the optimal value of the doubly restricted problem, and set $\boldsymbol{\lambda}$ to the optimal solution of the problem.

(d) If $\mathrm{UB_d} - \mathrm{LB_d} < \epsilon/2$, set $\mathrm{UB} = \mathrm{UB_d}$ and go to Step 4. Otherwise, go to Step 3(b).

4. **Termination.** Iterate between Steps 2 and 3 until the gap between LB and UB is zero or sufficiently small. Then terminate with the primal-dual pair $(\boldsymbol{p}^*, \boldsymbol{\lambda}^*) \in \mathbb{R}^{|\mathcal{N}^C|} \times \mathbb{R}^{|\mathcal{V}|}$, where $p^*_S = p_S$, $S \in \hat{\mathcal{N}}^C$, $= 0$ otherwise and $\lambda^*_{\boldsymbol{v}} = \lambda_{\boldsymbol{v}}$, $\boldsymbol{v} \in \mathcal{V}$, $= 0$ otherwise.

The 'outer layer' of the algorithm iterates between conservative approximations to the primal and dual problems (3) and (4) that consider all uncertainty realizations $\boldsymbol{v} \in \mathcal{V}$ (in the primal problem) and assortments $S \in \mathcal{N}^C$ (in the dual problem) but that only consider subsets of the admissible decisions p_S, $S \in \hat{\mathcal{N}}^C \subseteq \mathcal{N}^C$ (in the primal problem) and $\lambda_{\boldsymbol{v}}$, $\boldsymbol{v} \in \hat{\mathcal{V}} \subseteq \mathcal{V}$ (in the dual problem). The resulting restrictions remain computationally challenging as they involve embedded non-convex optimization problems, and they are solved iteratively up to a pre-specified tolerance of $\epsilon \geq 0$. To this end, the 'inner layers' of the algorithm iterate between *(i)* determining an optimal solution to the embedded optimization problem while fixing the solution to the outer optimization problems in Steps 2(b) and 3(b) and *(ii)* solving the outer optimization problems for the refined uncertainty sets in Steps 2(c) and 3(c). The subproblems are non-convex sums-of-ratios problems that admit equivalent reformulations as mixed-integer optimization problems, see Appendix A.3.

Our column generation scheme is reminiscent to, but differs in several important aspects from, the two-layer algorithm proposed in Section 5.3 of Delage and Saif (2018). While Delage and Saif (2018) study a two-stage linear robust optimization problem, our randomized constrained MNL problem (3) is a single-stage robust optimization problem that comprises a non-convex objective function. To solve their problem, Delage and Saif (2018) equivalently express their primal and dual subproblems through Karush-Kuhn-Tucker reformulations, whereas our subproblems amount to sums-of-ratios problems that are amenable to mixed-integer reformulations as detailed in Appendix A.3.

The correctness of our algorithm is summarized in the following statement.

THEOREM 3. *The two-layer column generation algorithm presented above converges in a finite number of iterations and terminates with an optimal primal-dual solution pair* $(\boldsymbol{p}^*, \boldsymbol{\lambda}^*)$ *to problems* (3) *and* (4).

4.3.2. A Local Search Heuristic The main computational challenge in the column generation scheme from the previous section arises from the subproblems of Steps 2(b) and 3(b), which amount to discrete sums-of-ratios problems. In the following, we adapt the ADXOpt heuristic of Jagabathula (2014), which was originally designed for nominal assortment optimization problems under general choice models, to these subproblems. Since the subproblem of Step 3(b) is in the form considered by Jagabathula (2014), we focus on the subproblem of Step 2(b) and consider the budget uncertainty set of Section 4.3:

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

The adaptations to other uncertainty sets are straightforward.

Local search heuristic for problem (5)

1. **Initialization.** Set $\boldsymbol{\xi} = \mathbf{0}$ (current worst-case valuation).
2. **Additions.** If $\mathbf{e}^\top \boldsymbol{\xi} < \Gamma$, evaluate the objective function of problem (5) for all $\boldsymbol{\xi}' = \boldsymbol{\xi} + \mathbf{e}_i$, $i = 0, \ldots, n$, for which $\boldsymbol{\xi}' \in \{0,1\}^{n+1}$. If any such $\boldsymbol{\xi}'$ improves the objective function of (5), update $\boldsymbol{\xi}$ to the best such $\boldsymbol{\xi}'$ and repeat Step 2.
3. **Deletions and exchanges.** Evaluate the objective function of problem (5) for all $\boldsymbol{\xi}' = \boldsymbol{\xi} - \mathbf{e}_i$, $i = 0, \ldots, n$, and $\boldsymbol{\xi}' = \boldsymbol{\xi} + \mathbf{e}_i - \mathbf{e}_j$, $i, j = 0, \ldots, n$ and $i \neq j$, for which $\boldsymbol{\xi}' \in \{0,1\}^{n+1}$. If any such $\boldsymbol{\xi}'$ improves the objective function of (5), update $\boldsymbol{\xi}$ to the best such $\boldsymbol{\xi}'$ and go to Step 2. Otherwise, go to Step 4.
4. **Termination.** Terminate and return $\boldsymbol{\xi}$ as worst-case valuation.

Being a local search heuristic, the worst-case valuation returned by our adaptation of ADXOpt may not represent a global minimizer of the primal subproblem (5). As such, we may increase LB_p and, subsequently, LB, excessively. Likewise, the application of ADXOpt to the dual subproblem (6) may decrease UB_d and, subsequently, UB, excessively. Either of these occurrences may cause the two-layer column generation scheme to terminate early with a feasible but suboptimal primal-dual solution pair $(\boldsymbol{p}^*, \boldsymbol{\lambda}^*)$. Our numerical experiments in Section 7 show, however, that the local search heuristic performs remarkably well and returns solutions that are very close—and often identical—to the optimal ones.

5. Markov Chain Model

Under the Markov chain (MC) model proposed by Blanchet et al. (2016), the choice behavior of a customer is described through a vector $\boldsymbol{\lambda} \in \mathbb{R}_+^{n+1}$, where λ_i, $i = 0, \ldots, n$, 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}, $i, j = 0, \ldots, n$, 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 $i = 0, \ldots, n$ with probability λ_i. If the preferred product, say i, is not offered, then the customer attempts to purchase each product $j = 0, \ldots, n$ with probability ρ_{ij}, and the process continues as if j had been the customer's preferred choice. We require that $\rho_{00} = 1$ as well as $\rho_{ii} = 0$ and $\rho_{i0} > 0$ for all $i \neq 0$.

Désir et al. (2019) study a robust assortment optimization problem under the MC model where the substitution matrices $\boldsymbol{\rho}$ are only known to reside in an uncertainty set

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

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 in view of 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 $\rho_{i0} < 1$.

In the following, Section 5.1 first expresses the robust assortment optimization problem under the MC model as a robust Markov decision process (MDP). We then study the benefits of randomization under two different classes of uncertainty sets for the MC model: product-wise substitution sets where no information is available about the dependence of ρ_{ij} and ρ_{kl} whenever $i \neq k$ (Section 5.2) and general substitution sets (Section 5.3).

5.1. A Robust MDP Reformulation for the MC Model

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 $(\mathcal{S}, \mathcal{A}, \mathcal{P}, q, c, \gamma)$, where $\mathcal{S}$ denotes the state space, $\mathcal{A}$ represents the action space, $q(s)$, $s \in \mathcal{S}$, characterize the initial state distribution, $c(s,a)$, $s \in \mathcal{S}$ and $a \in \mathcal{A}$, denote the immediate rewards, and $\gamma \in (0,1)$ is the discount factor. The ambiguity set

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

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

A robust MDP starts in state $s \in \mathcal{S}$ with known probability $q(s)$. The decision maker can then select any action $a \in \mathcal{A}$, upon which an immediate reward of $c(s,a)$ is earned and the MDP transitions to state $s' \in \mathcal{S}$ with probability $p(s'|s,a)$, 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 $\pi : \mathcal{S} \to \mathcal{A}$, which declares for each state $s \in \mathcal{S}$ which action $a \in \mathcal{A}$ is to be taken in state s, that maximizes the worst-case expected total discounted reward:

$$\max_{\pi \in \Pi} \min_{p \in \mathcal{P}} \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \cdot r(s_t, a_t) \,\middle|\, s_0 \sim q \right]$$

Here, Π denotes the set of all deterministic, memoryless policies $\pi : \mathcal{S} \to \mathcal{A}$, $\{(s_t, a_t)\}_{t=0}^{\infty}$ 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 robust assortment optimization problem.

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

$$p(s'|s,\top) = \begin{cases} 1 & \text{if } s' = s, \\ 0 & \text{otherwise,} \end{cases} \qquad p(s'|s,\bot) = \begin{cases} \rho_{ss'}/\gamma & \text{if } s, s' \neq 0, \\ 1 - \sum_{\sigma \in \mathcal{N}} \rho_{s\sigma}/\gamma & \text{if } s' = 0, \\ 0 & \text{otherwise} \end{cases} \qquad \forall s, s' \in \mathcal{S} \tag{7}$$

for some $\boldsymbol{\rho} \in \mathcal{U}$, the initial state probabilities $q(s) = \lambda_s$, $s \in \mathcal{S}$, the immediate rewards $c(s,\top) = (1-\gamma) r_s$ and $c(s,\bot) = 0$, $s \in \mathcal{S}$, 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 $s \in \mathcal{N}_0$ of the customer, and the actions $\top$ and $\bot$ characterize the options of the decision maker to include (exclude) any of the products $s \in \mathcal{N}$ 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 are now ready to establish the equivalence between the robust assortment optimization problem under the MC model and our robust MDP from Definition 3. To this end, denote by $R(S,\mathcal{U}) = \min_{\boldsymbol{u}\in\mathcal{U}} R(S,\boldsymbol{u})$ the worst-case expected revenues of the assortment $S \subseteq \mathcal{N}$.

Theorem 4. *The robust assortment optimization problem under the MC model is equivalent to the robust MDP from Definition 3 in the following sense:*

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

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

Theorem 4 shows that there is a one-to-many relationship between the assortments $S \subseteq \mathcal{N}$ of the assortment optimization problem and the policies $\pi_S : \mathcal{S} \to \mathcal{A}$ through the relation $i \in S \Leftrightarrow \pi_S(s) = \top$, $i \in \mathcal{N}$. The relationship is not one-to-one since the no-purchase option 0 is always available, irrespective of whether $\pi_S(0) = \top$ or $\pi_S(0) = \bot$. In the following two subsections, 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 $R^*_{\text{rand}}(\mathcal{U}) > R^*_{\text{det}}(\mathcal{U})$.

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 4 continue to apply.

5.2. Product-Wise Substitution Sets

We say that the uncertainty set $\mathcal{U}$ of a robust assortment optimization problem under the MC model has *product-wise substitution sets* whenever

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

One readily verifies that this condition is equivalent to requiring that

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

Intuitively speaking, the uncertainty set $\mathcal{U}$ has product-wise substitution sets when knowledge of the uncertain substitution probabilities $\boldsymbol{\rho}_i = (\rho_{i0}, \ldots, \rho_{i,n})$ for product $i \in \mathcal{N}_0$ does not allow the decision maker to infer anything about the uncertain substitution probabilities $\boldsymbol{\rho}_j = (\rho_{j0}, \ldots, \rho_{j,n})$ for any other product $j \in \mathcal{N}_0$, $j \neq i$, beyond the fact that $\boldsymbol{\rho}_j \in \mathcal{U}_j$.

THEOREM 5. *If the uncertainty set $\mathcal{U}$ of the robust assortment optimization problem under the MC model has product-wise substitution sets, the problem is randomization-proof.*

The proof of Theorem 5 shows that for product-wise substitution sets, the robust MDP from Definition 3 has an (s,a)-rectangular ambiguity set, which implies strong duality:

$$\max_{\pi \in \Pi} \min_{p \in \mathcal{P}} \mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^t \cdot r(s_t, a_t) \;\middle|\; s_0 \sim q \right] \;=\; \min_{p \in \mathcal{P}} \max_{\pi \in \Pi} \mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^t \cdot r(s_t, a_t) \;\middle|\; s_0 \sim q \right].$$

Any policy π^* corresponding to an optimal deterministic assortment achieves the objective value on the left-hand side of this equation, whereas 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 substitution matrix $\boldsymbol{\rho}$ 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.

In a recent paper, Désir et al. (2019) derive a strong duality result for the robust assortment optimization problem under the MC model. In contrast to our result, which establishes a connection between the assortment optimization problem and robust MDPs and subsequently leverages existing results for robust MDPs, Désir et al. (2019) prove strong duality *ab initio* for a class of row-rectangular uncertainty sets, which are closely related (but not equivalent) to our product-wise substitution sets.

Table 1 **Expected revenues of different assortments in Example 2 (worst-case scenarios highlighted in bold).**

Scenario	$(\uparrow,\uparrow,\uparrow)$	$(\uparrow,\uparrow,\downarrow)$	$(\uparrow,\downarrow,\uparrow)$	$(\uparrow,\downarrow,\downarrow)$	$(\downarrow,\uparrow,\uparrow)$	$(\downarrow,\uparrow,\downarrow)$	$(\downarrow,\downarrow,\uparrow)$	$(\downarrow,\downarrow,\downarrow)$
$\{1\}$	**7.3**	7.8	7.8	8.1	7.3	7.8	7.8	8.1
$\{2\}$	7.8	**7.3**	7.8	7.3	8.1	7.8	8.1	7.8
$\{3\}$	8.1	8.1	7.8	7.8	7.8	7.8	**7.3**	7.3
$\sum_i \frac{1}{3}\cdot\{i\}$	**7.7**	7.7	7.8	7.7	7.7	7.8	7.7	7.7

While the *unconstrained* robust assortment optimization problem is randomization-proof whenever $\mathcal{U}$ has product-wise substitution sets, the problem becomes randomization-receptive if we impose a cardinality constraint on the size of the admissible assortments.

Example 2. Consider an MC model with three products, product-wise revenues of $r_i = 10$, $i \in \mathcal{N}$, initial choice probabilities $\boldsymbol{\lambda} = (0, 1/3, 1/3, 1/3)$ and the uncertainty set

$$\mathcal{U} = \left\{ \boldsymbol{\rho} \in \mathbb{R}_+^{4\times 4} : \boldsymbol{\rho}_0 = (1,0,0,0) \text{ and } \begin{array}{r} \boldsymbol{\rho}_1 \in \{\boldsymbol{\rho}^{\uparrow} = (0.2, 0, 0.3, 0.5),\ \boldsymbol{\rho}^{\downarrow} = (0.2, 0, 0.5, 0.3)\}, \\ \boldsymbol{\rho}_2 \in \{\boldsymbol{\rho}^{\uparrow} = (0.2, 0.3, 0, 0.5),\ \boldsymbol{\rho}^{\downarrow} = (0.2, 0.5, 0, 0.3)\}, \\ \boldsymbol{\rho}_3 \in \{\boldsymbol{\rho}^{\uparrow} = (0.2, 0.3, 0.5, 0),\ \boldsymbol{\rho}^{\downarrow} = (0.2, 0.5, 0.3, 0)\} \end{array} \right\}.$$

The uncertainty set $\mathcal{U}$ contains $2^3 = 8$ substitution matrices, and one readily verifies that $\mathcal{U}$ has product-wise substitution sets. Assume that only assortments with a single product are allowed, that is, $|S| \leq C = 1$. Table 1 lists the expected revenues of all admissible deterministic assortments, as well as one randomized assortment, under all scenarios of $\mathcal{U}$. The table shows that the randomized assortment outperforms the eligible deterministic assortments in terms of worst-case expected revenues.

We next show that, similar to the constrained MNL model (*cf.* Theorem 2), the potential benefits of randomization in the cardinality-constrained MC model can be arbitrarily large.

Theorem 6. *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 $R^*_{det}(\mathcal{U}) = 0$ while $R^*_{rand}(\mathcal{U}) > 0$.*

Similar to Section 4.2, an immediate consequence of Theorem 6 is the following.

Corollary 4. *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 the benefits $R^*_{rand}(\mathcal{U})/R^*_{det}(\mathcal{U})$ from randomization are arbitrarily large.*

Given that the cardinality-constrained MC model is receptive to randomization, it is natural to ask how optimal randomized assortment strategies can be determined under this model. In fact, even determining an optimal deterministic assortment appears to be challenging in this setting since the cardinality restriction imposes a global constraint on the set of admissible policies for the associated robust MDP. This constraint cannot be accommodated for by a straightforward lifting of the state space (*i.e.*, by introducing an additional state for the assortment space consumed so far) as the resulting robust MDP would violate the non-anticipativity requirement, which stipulates that the assortment has to be selected before the customer's substitution behavior is observed.

5.3. General Substitution Sets

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.*

On the other hand, we can find instances with three or more products where the unconstrained robust MC problem with general substitution sets is receptive to randomization.

EXAMPLE 3. Consider an MC model with three products, revenues $r_1 = 4.66, r_2 = 1, r_3 = 10$, initial choice probabilities $\boldsymbol{\lambda} = (0, 0.37, 0.62, 0.01)$, and the uncertainty set

$$\mathcal{U} = \left\{ \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0.01 & 0 & 0.05 & 0.94 \\ 0.26 & 0.69 & 0 & 0.05 \\ 0.90 & 0.05 & 0.05 & 0 \end{pmatrix}, \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0.90 & 0 & 0.05 & 0.05 \\ 0.01 & 0.44 & 0 & 0.55 \\ 0.01 & 0.94 & 0.05 & 0 \end{pmatrix} \right\}.$$

One readily verifies that $\mathcal{U}$ does not have product-wise substitution sets. Table 2 lists the expected revenues of all deterministic assortments, as well as one randomized assortment,

Table 2 Expected revenues of different assortments in Example 3 (worst-case scenarios highlighted in bold).

Assortment	{1}	{2}	{3}	{1,2}	{1,3}	{2,3}	{1,2,3}	$0.374 \cdot \{3\} + 0.626 \cdot \{1,3\}$
Scenario 1	**3.72**	0.66	7.93	**2.35**	**4.13**	4.22	2.44	5.68
Scenario 2	4.62	**0.64**	**4.02**	2.39	6.57	**0.92**	**2.44**	**5.61**

under both scenarios in $\mathcal{U}$. The table shows that the randomized assortment outperforms the deterministic assortments in terms of worst-case expected revenues.

Intuitively speaking, the adversary 'nature' can benefit from knowing the decision maker's assortment choice if the uncertainty set in the MC model does not have product-wise substitution sets. Since randomizing between different assortments ensures that nature only observes the randomization strategy but not the actual assortment shown to each customer, the decision maker can increase her worst-case expected revenues.

Observation 2. *For any number $n \geq 3$ of products, the unconstrained MC problem is randomization-receptive, that is, there exist instances $(\mathcal{U}, \boldsymbol{r})$ such that $R^*_{rand}(\mathcal{U}) > R^*_{det}(\mathcal{U})$.*

A natural question is whether, similar to the constrained robust MNL and MC models, the potential benefits of randomization can be unbounded in the unconstrained MC model under general substitution sets. As the following result shows, this is not the case.

Proposition 2. *Under the unconstrained MC model with general substitution sets, the benefits of randomization are bounded from above by*

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

We next discuss two immediate consequences of Proposition 2.

Corollary 5. *As long as all products $i \in \mathcal{N}$ carry strictly positive prices $r_i > 0$, the benefits of randomization under the unconstrained MC model with general substitution sets are bounded. Moreover, the problem is randomization-proof if all products have the same price, that is, if $r_i = r_j$ for all $i, j \in \mathcal{N}$.*

Since general substitution sets encompass product-wise substitution sets as a special case, the findings of Theorem 6 and Corollary 4 immediately apply to the cardinality-constrained MC model with general substitution sets as well.

Similar to the cardinality constrained MC model with product-wise substitution sets, determining optimal randomized assortment strategies for the unconstrained MC model with general substitution sets appears to be challenging. In fact, even determining an optimal deterministic assortment appears to be difficult as it would amount to finding an optimal policy for a robust MDP with a non-rectangular uncertainty set, which is a notoriously difficult task (*cf.* Wiesemann et al. 2013). Moreover, the class of randomized policies in the robust MDP appears not to be expressive enough to capture all randomized assortment strategies as the former are isomorphic to the 0-1-hypercube in $\mathbb{R}^{n+1}$, whereas the latter are isomorphic to probability simplex in $\mathbb{R}^{2^n}$.

6. Preference Ranking Model

The preference ranking model (see, *e.g.*, Bertsimas and Mišić 2019) is parameterized by K bijective preference rankings $\sigma_k : \mathcal{N}_0 \to \{1, \ldots, n+1\}$, $k \in \mathcal{K} = \{1, \ldots, K\}$, 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 $S \subseteq \mathcal{N}$ under the preference ranking model amount to

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

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.

In the following, we study the potential benefits of randomization under the unconstrained (Section 6.1) and cardinality-constrained (Section 6.2) robust preference ranking model. Section 6.3, finally, discusses exact and heuristic solution schemes for the robust assortment optimization problem under the (un-)constrained preference ranking model.

6.1. The Unconstrained Preference Ranking Model

We first show that even in the absence of any further constraints (such as cardinality constraints), the preference ranking model is randomization-receptive.

Example 4. Consider the unconstrained robust preference ranking problem with two products, $r_1 = 1$ and $r_2 = 2$ and the following three preference rankings:

$$1 \to 2 \to 0, \qquad 2 \to 1 \to 0, \qquad 1 \to 0 \to 2.$$

Table 3 **Ranking prevalences and expected revenues of different assortments in Example 4.**

Scenario	λ_1	λ_2	λ_3	$\{1\}$	$\{2\}$	$\{1,2\}$	$\frac{1}{3}\cdot\{2\}+\frac{2}{3}\cdot\{1,2\}$
1	1	0	0	**1**	2	**1**	**1.33**
2	0	0.5	0.5	1	**1**	1.5	**1.33**

The associated ranking prevalences and the expected revenues of the deterministic assortments as well as one randomized assortment are listed in Table 3. As the table shows, randomizing between the assortments $\{2\}$ and $\{1,2\}$ results in a 33% improvement of the worst-case expected revenues over the best deterministic assortment.

In fact, randomization can be beneficial for any number of products, as we show next.

Observation 3. *For any number of products $n \geq 2$, the preference ranking problem is randomization-receptive, that is, there exist instances $(\mathcal{U}, \boldsymbol{r})$ such that $R^*_{rand}(\mathcal{U}) > R^*_{det}(\mathcal{U})$.*

Contrary to the constrained MNL and MC models (*cf.* Theorems 2 and 6), however, and similar to the unconstrained MC model with general substitution sets (*cf.* Proposition 2), the benefits of randomization are bounded.

Proposition 3. *For any instance of the robust preference ranking problem with $n = 2$ products, the benefits of randomization are bounded from above by $R^*_{rand}(\mathcal{U}) \,/\, R^*_{det}(\mathcal{U}) \leq 2$. Moreover, for any instance the benefits of randomization are bounded from above by*

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

Example 5 in the appendix shows that the bound for the two-product case is asymptotically tight. Similar to Section 5.3, two consequences of Proposition 3 are immediate.

Corollary 6. *If all products carry strictly positive prices $r_i > 0$, the benefits of randomization under the robust preference ranking model are bounded. Moreover, the problem is randomization-proof if all products have the same price, that is, if $r_i = r_j$ for all $i, j \in \mathcal{N}$.*

We have seen that in the cardinality-constrained MNL problem, optimal policies may require randomization between an exponentially large number of assortments (*cf.* Theorem 2). We now show that under the unconstrained preference ranking model, there is always a parsimonious optimal randomization strategy.

Theorem 7. *Under the preference ranking model, there exists an optimal randomization strategy which places strictly positive weight on no more than $K + 1$ assortments.*

6.2. The Cardinality-Constrained Preference Ranking Model

We now consider a variant of the preference ranking model where all admissible assortments must carry at most C products, that is, any of admissible assortment $S \subseteq \mathcal{N}$ must satisfy $|S| \leq C$. We first show that, similar to the cardinality-constrained MNL and MC models (*cf.* Theorems 2 and 6), the potential benefits of randomization in the cardinality-constrained preference ranking model can be arbitrarily large.

Theorem 8. *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 preference ranking problem where $R^*_{det}(\mathcal{U}) = 0$ while $R^*_{rand}(\mathcal{U}) > 0$.*

As before, an immediate consequence of Theorem 8 is the following.

Corollary 7. *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 preference ranking problem where the benefits $R^*_{rand}(\mathcal{U})/R^*_{det}(\mathcal{U})$ from randomization are arbitrarily large.*

Finally, we show that the existence of parsimonious optimal randomization strategies carries over to the preference ranking model with arbitrary constraints on the set of admissible assortments.

Corollary 8. *Under the preference ranking model with arbitrary constraints on the set of admissible assortments, there exists an optimal randomization strategy which places strictly positive weight on no more than $K+1$ assortments.*

6.3. Solving the Randomized Preference Ranking Model

Theorem 7 and Corollary 8 allow us to formulate the randomized assortment optimization problem under the unconstrained and constrained preference ranking 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 now present an exact and a heuristic column generation scheme to compute randomization strategies for the preference ranking model. To this end, we assume that the ambiguity set Λ is a polyhedral set of the form

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

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

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

Note that $\{\boldsymbol{\xi} \in \mathbb{R}^K : \|\boldsymbol{\xi}\|_2 \le 1\} \subseteq \{\boldsymbol{\xi} \in \mathbb{R}^K : \|\boldsymbol{\xi}\|_1 \le \sqrt{K}\} \cap \{\boldsymbol{\xi} \in \mathbb{R}^K : \|\boldsymbol{\xi}\|_\infty \le 1\}$, 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 preference ranking model can be formulated as the following robust linear program:

$$\begin{aligned} \text{maximize} \quad & \min_{\boldsymbol{\lambda} \in \Lambda} \sum_{S \subseteq \mathcal{N}} \sum_{k \in \mathcal{K}} p_S \cdot \lambda_k \cdot R_k(S) \\ \text{subject to} \quad & \sum_{S \subseteq \mathcal{N}} p_S = 1 \\ & p_S \ge 0,\ S \subseteq \mathcal{N} \end{aligned} \tag{8}$$

Problem (8) is computationally challenging as it involves exponentially many decision variables p_S, $S \subseteq \mathcal{N}$. For our solution schemes, it is useful to consider the dual of problem (8):

$$\begin{aligned} \text{minimize} \quad & \max_{S \subseteq \mathcal{N}} \sum_{k \in \mathcal{K}} \lambda_k \cdot R_k(S) \\ \text{subject to} \quad & \boldsymbol{\lambda} \in \Lambda \end{aligned} \tag{9}$$

Strong duality between (8) and (9) holds since (8) is feasible by construction. Although problem (9) contains only polynomially many decision variables, the optimization problem embedded in its objective function maximizes over a combinatorial set $S \subseteq \mathcal{N}$. Thus, neither the primal problem (8) nor the dual problem (9) is amenable to a solution with an off-the-shelf solver. In the following, we present a column generation scheme for problem (8) that can be interpreted as a cutting plane approach for problem (9).

Column generation scheme for problem (8)

1. **Initialization.** Set $\mathcal{S}$ to any non-empty subset of the power set $2^{\mathcal{N}}$ of $\mathcal{N}$ (currently considered assortments). Set $\text{LB} = -\infty$ and $\text{UB} = +\infty$ (current lower and upper bounds on the optimal value).
2. **Master problem.** Solve the restricted dual problem (9) that replaces $S \subseteq \mathcal{N}$ with $S \in \mathcal{S}$ in the objective function. Update $\text{LB} = \max\{\text{LB}, \tau^*\}$, where τ^* is the optimal value of the restricted dual problem, and set $\boldsymbol{\lambda}$ to the optimal solution of the problem.
3. **Subproblem.** Solve the subproblem

$$\begin{array}{ll} \text{maximize} & \displaystyle\sum_{k \in \mathcal{K}} \lambda_k \cdot R_k(S) \\ \text{subject to} & S \subseteq \mathcal{N}. \end{array}$$

   Update $\text{UB} = \min\{\text{UB}, \tau^*\}$, where τ^* is the optimal value of the subproblem, and $\mathcal{S} = \mathcal{S} \cup \{S^*\}$, where S^* is the optimal solution of the problem.
4. **Termination.** Iterate between Steps 2 and 3 until the gap between LB and UB is zero or sufficiently small. Then terminate with the primal-dual pair $(\boldsymbol{p}^*, \boldsymbol{\lambda}^*) \in \mathbb{R}^{2^n} \times \mathbb{R}^K$, where p^*_S, $S \in \mathcal{S}$, can be derived from the shadow prices of the epigraph reformulation of the restricted dual problem (9), and $p^*_S = 0$ for $S \notin \mathcal{S}$.

The algorithm iterates between solving two problems: Step 2 solves an increasingly accurate progressive approximation of the dual problem (9) to determine a candidate solution $\boldsymbol{\lambda} \in \Lambda$. Step 3 evaluates the exact objective function of problem (9) for $\boldsymbol{\lambda}$ and identifies a cutting plane to add to the approximation from Step 2. Both problems can be expressed as finite (mixed-integer) linear programs that can be solved with standard software (Bertsimas and Mišić 2019). If the set of admissible assortments is subject to further constraints (such as cardinality constraints), then we incorporate those constraints in the problem solved in Step 3.

Our column generation scheme is reminiscent of the single-layer algorithm proposed in Section 5.4 of Delage and Saif (2018). Its correctness is summarized next.

THEOREM 9. *The column generation algorithm presented above converges in a finite number of iterations and terminates with an optimal primal-dual solution pair* $(\boldsymbol{p}^*, \boldsymbol{\lambda}^*)$ *to problems* (8) *and* (9).

The main computational challenge in the above column generation scheme arises from the subproblem of Step 3, which amounts to a nominal assortment optimization problem over the preference ranking model. Instead of solving this problem exactly, we can solve it heuristically with the ADXOpt heuristic of Jagabathula (2014). In doing so, we may miss the global maximizer of the subproblem and therefore decrease UB excessively, which in turn can cause the column generation scheme to terminate early with a feasible but suboptimal primal-dual solution pair $(\boldsymbol{p}^*, \boldsymbol{\lambda}^*)$. Similar to the cardinality-constrained MNL model, however, our numerical experiments in Section 7 show that the local search heuristic performs remarkably well and generally returns solutions of excellent quality. Moreover, Jagabathula (2014) explains how additional constraints (such as cardinality constraints) on the admissible assortments can be accounted for in the solution scheme.

7. Numerical Results

This section presents numerical results for the cardinality-constrained MNL model (Section 7.1) and the unconstrained preference ranking model (Section 7.2). Our results aim to elucidate when randomization can help to improve worst-case expected revenues on synthetic data sets, and 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. We also investigate the computational price to be paid for exact and heuristic solutions to the deterministic and randomized robust assortment optimization problems. All solution schemes are 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.[1]

7.1. Cardinality-Constrained MNL Model

We first consider the cardinality-constrained MNL model where the product revenues r_i are selected uniformly at random from the interval $[0, 10]$. We use a buget uncertainty set (*cf.* Section 4.3) where the lower and upper product valuations $\underline{v}_i$ and $\overline{v}_i$, $i \in \mathcal{N}$, are

[1] `www.doc.ic.ac.uk/~wwiesema/assortment_opt.zip`

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

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 4 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 the 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 (*cf.* Theorem 1) 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 proposed by Bandi and Bertsimas (2012). The 'randomized exact objective' column of Table 5 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 5 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 robust and the randomized problem, 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 problem, the outperformance of the ADXOpt heuristic approaches the outperformance of the exact solution approach. We thus conclude that one should solve the deterministic robust and randomized 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 5 **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 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 revenues 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, ..., 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 robust and the randomized model, we employ the budget uncertainty set where both $\Gamma \in \{0,\ldots,n\}$ and $(\underline{\boldsymbol{v}}, \overline{\boldsymbol{v}}) = ([1-\gamma]\hat{\boldsymbol{v}}, [1+\gamma]\hat{\boldsymbol{v}})$, $\gamma \in \{0, 0.025, \ldots, 0.5\}$ are selected using 7-fold cross-validation on the available historical data. We then compare the out-of-sample expected revenues of the three models under the true valuations $\boldsymbol{v}^0$.

For each cardinality $C \in \{1,2,3,4\}$ and sample size 5, 10, ..., 95, Figure 1 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 model outperforms the deterministic robust model, which in turn outperforms the nominal model. We emphasize that this is not *a priori* obvious as the deterministic robust and the randomized model 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 model significantly outperforms both the nominal and the deterministic robust model. For larger cardinalities, the performance of the deterministic robust and the randomized model are becoming more similar. It is noteworthy, however, that the randomized 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 5).

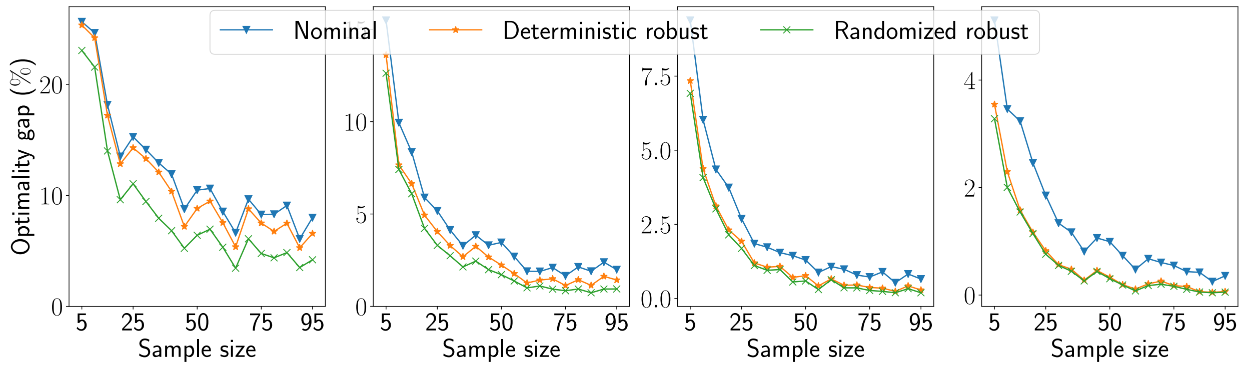

Figure 1 **Data-driven experiment for the cardinality-constrained MNL problem. All optimality gaps are reported relative to the expected revenues of the clairvoyant model that knows the true customer valuations.**

7.2. Preference Ranking Model

We next consider the preference ranking model where the product revenues are again selected uniformly at random from the interval $[0, 10]$. We use a 1-norm ambiguity set (*cf.* Section 6.3) where the nominal ranking prevalence vector $\boldsymbol{\lambda}^0$ is selected uniformly at random from the K-dimensional probability simplex.

In our first experiment, we fix the number of products to $n = 20$ and vary the radius $\theta \in \{0.25, 0.5, \ldots, 2\}$ of the ambiguity set as well as the number of preference rankings $K \in \{10, 100, 200, 500, 1{,}000\}$. Table 6 presents the percentages of instances (across 250 randomly generated instances) where the randomized model improved upon the deterministic robust one in terms of worst-case expected revenues (first number) as well as the average relative outperformance across those instances (second number). The table shows that the benefits of randomization are most pronounced when the ambiguity is large but not covering the entire probability simplex (which is the case when $\theta = 2$). This is intuitive since under full ambiguity, nature can choose a Dirac distribution that places all probability mass on the least favorable (assortment-dependent) preference ranking, in which case randomization does not help. Likewise, if the ambiguity is small, we approach a non-robust problem where randomization does not help either.

We now investigate the benefits of randomization when the number n of products varies. To this end, Table 7 fixes the radius of ambiguity to $\theta = 1.5$ and reports the outperformance of the exact (column 'randomized exact objective') and heuristic (column 'randomized heuristic objective') solutions to our randomized model relative to the optimal solution to the deterministic robust model. The table shows that while the benefits of randomization

ρ \ K	**10**	**100**	**200**	**500**	**1,000**
0.00	0.00% 0.00%	0.00% 0.00%	0.00% 0.00%	0.00% 0.00%	0.00% 0.00%
0.25	3.00% 0.60%	13.30% 0.12%	13.30% 0.06%	15.45% 0.06%	12.88% 0.03%
0.50	30.04% 1.21%	60.09% 0.29%	58.80% 0.21%	50.64% 0.19%	62.23% 0.23%
0.75	25.75% 0.63%	78.11% 0.55%	80.69% 0.51%	83.26% 0.49%	78.97% 0.46%
1.00	33.48% 1.99%	88.41% 1.06%	94.85% 1.00%	93.56% 1.08%	94.85% 0.96%
1.25	29.18% 3.08%	95.71% 2.37%	97.00% 2.50%	96.14% 1.91%	97.85% 1.84%
1.50	19.74% 2.63%	98.71% 3.91%	98.71% 3.47%	100.00% 3.38%	100.00% 3.04%
1.75	0.00% 0.00%	93.56% 10.81%	99.57% 9.22%	100.00% 6.49%	99.57% 7.17%
2.00	0.00% 0.00%	0.00% 0.00%	0.00% 0.00%	0.00% 0.00%	0.00% 0.00%

Table 6 **Benefits of randomization in the preference ranking problem. In the table the rows (columns) correspond to different ambiguity radii θ (numbers of preference rankings K).**

decrease with the number of products n, they remain non-trivial for all considered problem sizes. Moreover, the performance of the ADXOpt heuristic is more or less independent of the number of products n as well as improving with the number of preference rankings K.

Table 3 also reports the runtimes of the mixed-integer programs for the deterministic robust model (column 'deterministic exact runtime') as well as our column generation scheme (column 'randomized exact runtime') and the ADXOpt heuristic (column 'randomized heuristic runtime') for the randomized model. The table shows that the randomized exact problem is significantly more challenging to solve than deterministic robust one. Indeed, while the exact deterministic robust problem can be solved as a monolithic MILP, the exact randomized problem is solved in a cutting plane fashion and thus suffers from similar scalability issues as the cutting plane scheme for the deterministic robust MNL problem (*cf.* Table 5). The number of cutting plane iterations grows with the problem size (median 6 for $n = 20$, $K = 10$ vs median 49 for $n = 50$, $K = 1{,}000$), and each subproblem of the algorithm requires the solution of a MILP whose number of decision variables scales in K. The ADXOpt heuristic, on the other hand, provides close-to-optimal solutions while requiring substantially less computational resources.

We have also attempted to solve the randomized assortment optimization problem under the preference ranking model as a K-adaptability problem (*cf.* Section 6.3). It turns out, however, that the resulting mixed-integer optimization problems grow quickly in size and become difficult to solve for instances with more than 10 products. We therefore do not report computational results for this solution scheme.

		deterministic	**randomized**			
		exact	**exact**		**heuristic**	
n	K	**runtime**	**objective**	**runtime**	**objective**	**runtime**
10	**10**	0.01s	0.54%	0.00s	0.51%	0.00s
10	**100**	0.05s	4.03%	0.12s	3.64%	0.03s
10	**200**	0.10s	4.95%	0.32s	4.10%	0.07s
10	**500**	0.41s	3.11%	0.90s	2.74%	0.18s
10	**1000**	0.68s	3.44%	2.57s	2.92%	0.49s
20	**10**	0.01s	0.44%	0.01s	-0.15%	0.00s
20	**100**	0.16s	3.83%	0.81s	3.16%	0.08s
20	**200**	0.34s	5.32%	2.48s	4.32%	0.18s
20	**500**	1.41s	3.38%	7.49s	2.94%	0.54s
20	**1000**	5.02s	2.99%	20.59s	2.61%	1.09s
30	**10**	0.01s	0.41%	0.02s	0.11%	0.00s
30	**100**	0.33s	3.82%	3.11s	2.89%	0.16s
30	**200**	0.78s	3.69%	7.38s	3.13%	0.40s
30	**500**	4.43s	2.97%	31.24s	2.61%	1.03s
30	**1000**	16.30s	2.91%	102.33s	2.62%	2.16s
40	**10**	0.03s	0.14%	0.03s	-0.01%	0.00s
40	**100**	0.49s	4.83%	6.43s	3.81%	0.33s
40	**200**	1.66s	3.23%	22.06s	2.52%	0.71s
40	**500**	8.73s	2.40%	82.31s	2.13%	1.81s
40	**1000**	43.61s	2.32%	362.83s	2.07%	3.83s
50	**10**	0.03s	0.08%	0.03s	-0.08%	0.01s
50	**100**	0.92s	2.17%	8.19s	1.74%	0.45s
50	**200**	2.56s	3.95%	53.55s	3.20%	1.22s
50	**500**	17.28s	2.71%	219.30s	2.33%	2.91s
50	**1000**	75.45s	2.46%	1011.82s	2.24%	6.99s

Table 7 **Exact and heuristic solutions for the preference ranking problem. All percentages are reported relative to the optimal solutions of the deterministic robust problem.**

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 Markov chain 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. Another direction for future research is finding optimal randomization strategies for the Markov chain model (Section 5). While we were able to develop efficient algorithms for the (constrained) MNL and preference ranking models, finding optimal randomizations for the MC model appears much more challenging.

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 and EP/T024712/1.

References

An B, Tambe M, Sinha A. 2016. Stackelberg security games (SSG): basics and application overview. Abbas A, Tambe M, von Winterfeldt D, eds., *Improving Homeland Security Decisions*. Cambridge University Press, Cambridge, 485–507.

Aouad A, Farias V, Levi R. 2015. Assortment optimization under consider-then-choose choice models. Working paper, Massachusetts Institute of Technology, Cambridge.

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.

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

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

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

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

Bertsimas D, Mišić VV. 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.

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.

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.

Delage E, Saif A. 2018. The value of randomized solutions in mixed-integer distributionally robust optimization problems. Working paper, HEC Montréal, Montréal.

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

Désir A, Goyal V, Jiang B, Zhang J. 2019. A nonconvex min-max framework and its applications to robust assortment optimization. Working paper, INSEAD, Fontainebleau.

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.

Feldman J, Topaloglu H. 2018. Technical note: Capacitated assortment optimization under the multinomial logit model with nested consideration sets. *Operations Research* **66**(2) 380–391.

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.

Iyengar GN. 2005. Robust dynamic programming. *Mathematics of Operations Research* **30**(2) 257–280.

Jagabathula S. 2014. Assortment optimization under general choice. Working paper, New York University, New York.

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*. Springer, Boston, MA, 175–236.

Korzhyk D, Yin Z, Kiekintveld C, Conitzer V, Tambe M. 2011. Stackelberg vs. Nash in security games: An extended investigation of interchangeability, equivalence, and uniqueness. *Journal of Artificial Intelligence Research* **41**(2) 297–327.

Li H-L. 1994. A global approach for general 0–1 fractional programming. *European Journal of Operational Research* **73**(3) 590–596.

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

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

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

Mastin A, Jaillet P, Chin S. 2015. Randomized minmax regret for combinatorial optimization under uncertainty. Elbassioni K, Makino K, eds., *Algorithms and Computation*. Springer, 491–501.

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

Nilim A, El Ghaoui L. 2005. Robust control of Markov decision processes with uncertain transition matrices. *Operations Research* **53**(5) 780–798.

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.

Rusmevichientong P, Shen Z-JM, 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.

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.

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

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.

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*. Curran Associates, Inc., 1801–1808.

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

Appendix A: Auxiliary Results and Proofs for Section 4

A.1. The Unconstrained MNL Model

Proof of Lemma 1. The first part of the statement is similar to Theorem 3.2 of Rusmevichientong and Topaloglu (2012). The key difference is that Rusmevichientong and Topaloglu (2012) focus on the smallest optimal assortment $\underline{S}(\mathcal{V}) = \{i \in \mathcal{N} \ : \ r_i > R^*_{\text{det}}(\mathcal{V})\}$, whereas we consider all optimal assortments that emerge from adding any product set $S \subseteq \{i \in \mathcal{N} \ : \ r_i = R^*_{\text{det}}(\mathcal{V})\}$ to $\underline{S}(\mathcal{V})$. Since the proof of this extension closely follows the proof of Theorem 3.2 of Rusmevichientong and Topaloglu (2012), we omit it for brevity.

For later use, we first show that all optimal assortments share the same set of worst-case valuations. To this end, fix $S, S' \in \mathcal{S}(\mathcal{V})$ and any $\boldsymbol{v}^* \in \arg\min_{\boldsymbol{v} \in \mathcal{V}} R(S, \boldsymbol{v})$. We then have that

$$R(S,\boldsymbol{v}^*) \;=\; \frac{\sum_{i \in \underline{S}(\mathcal{V})} r_i v_i^* + \sum_{i \in S \setminus \underline{S}(\mathcal{V})} r_i v_i^*}{v_0 + \sum_{i \in \underline{S}(\mathcal{V})} v_i^* + \sum_{i \in S \setminus \underline{S}(\mathcal{V})} v_i^*} \;=\; R^*_{\text{det}}(\mathcal{V}) \;=\; \frac{\sum_{i \in \underline{S}(\mathcal{V})} r_i v_i^* + \sum_{i \in S' \setminus \underline{S}(\mathcal{V})} r_i v_i^*}{v_0 + \sum_{i \in \underline{S}(\mathcal{V})} v_i^* + \sum_{i \in S' \setminus \underline{S}(\mathcal{V})} v_i^*} \;=\; R(S',\boldsymbol{v}^*),$$

where the first and last identities follow from the fact that $\underline{S}(\mathcal{V}) \subseteq S^*(\mathcal{V})$ for any worst-case optimal assortment $S^*(\mathcal{V}) \in \mathcal{S}^*(\mathcal{V})$, whereas the second and third identities are due to the fact that the revenues r_i of all products $i \in (S \cup S') \setminus \underline{S}(\mathcal{V})$ are equal to $R^*_{\text{det}}(\mathcal{V})$ by the first part of the lemma, and the fact that $\frac{a}{b} = c$ implies that $\frac{a+\alpha c}{b+\alpha} = c$ for all $a, c, \alpha \geq 0$ and $b > 0$.

We next prove the second part of the lemma, which states that the ex-ante and ex-post worst-case valuations of any optimal assortment $S^*(\mathcal{V}) \in \mathcal{S}^*(\mathcal{V})$ coincide. Since $S^*(\boldsymbol{v}^*)$ is an optimal assortment for the scenario $\boldsymbol{v}^*$, we have $R(S^*(\boldsymbol{v}^*), \boldsymbol{v}^*) \geq R(S^*(\mathcal{V}), \boldsymbol{v}^*)$, and this inequality trivially holds with equality whenever $\mathcal{V} = \{\boldsymbol{v}^*\}$. For the case where $\mathcal{V}$ contains other scenarios as well, we show the equality by contradiction. Indeed, assume that $R(S^*(\mathcal{V}), \boldsymbol{v}^*) < R(S^*(\boldsymbol{v}^*), \boldsymbol{v}^*)$. Since revenue-ordered assortments are optimal under the scenario $\boldsymbol{v}^*$, we have $r_i \geq R(S^*(\boldsymbol{v}^*), \boldsymbol{v}^*) > R(S^*(\mathcal{V}), \boldsymbol{v}^*)$ for all $i \in S^*(\boldsymbol{v}^*)$. The first part of the lemma then implies that $S^*(\boldsymbol{v}^*) \subseteq \underline{S}(\mathcal{V}) = \{i \in \mathcal{N} : r_i > R(S^*(\mathcal{V}), \boldsymbol{v}^*)\}$, that is, either $S^*(\boldsymbol{v}^*) = \underline{S}(\mathcal{V})$ or $S^*(\boldsymbol{v}^*) \subset \underline{S}(\mathcal{V})$. In the following, we show that either case leads to a contradiction.

Assume first that $S^*(\boldsymbol{v}^*) = \underline{S}(\mathcal{V})$. Then $R(\underline{S}(\mathcal{V}), \boldsymbol{v}^*) = R(S^*(\boldsymbol{v}^*), \boldsymbol{v}^*) > R(S^*(\mathcal{V}), \boldsymbol{v}^*)$, where the equality follows from $S^*(\boldsymbol{v}^*) = \underline{S}(\mathcal{V})$ and the inequality is due to our earlier assumption in the previous paragraph. Since $S^*(\mathcal{V}), \underline{S}(\mathcal{V}) \in \mathcal{S}^*(\mathcal{V})$, however, this is not possible, as all worst-case optimal assortments share the same set of worst-case scenarios.

Assume now that $S^*(\boldsymbol{v}^*) \subset \underline{S}(\mathcal{V})$. We show that in this case, there is a product $i_1 \in \mathcal{N}$ such that $R(S^*(\boldsymbol{v}^*) \cup \{i_1\}, \boldsymbol{v}^*) > R(S^*(\boldsymbol{v}^*), \boldsymbol{v}^*)$, which contradicts the optimality of $S^*(\boldsymbol{v}^*)$ for scenario $\boldsymbol{v}^*$. To see this, let $S' = \{i_1, \cdots, i_J\} = \underline{S}(\mathcal{V}) \setminus S^*(\boldsymbol{v}^*)$, and assume without loss of generality that $r_{i_1} \geq \cdots \geq r_{i_J}$. We show via induction that $r_{i_j} > R(S^*(\boldsymbol{v}^*) \cup \{i_1, \cdots, i_j\}, \boldsymbol{v}^*)$ for all $j = 1, \cdots, J$. The case $j = 1$ then implies that $r_{i_1} > R(S^*(\boldsymbol{v}^*) \cup \{i_1\}, \boldsymbol{v}^*)$, and Lemma 3.1 (b) and (c) of Rusmevichientong and Topaloglu (2012) then further imply that $R(S^*(\boldsymbol{v}^*) \cup \{i_1\}, \boldsymbol{v}^*) > R(S^*(\boldsymbol{v}^*), \boldsymbol{v}^*)$ as desired.

The base case $j = J$ of our induction states that $r_{i_J} > R(S^*(\boldsymbol{v}^*) \cup \{i_1, \cdots, i_J\}, \boldsymbol{v}^*)$. This is indeed the case since $S^*(\boldsymbol{v}^*) \cup \{i_1, \cdots, i_J\} = \underline{S}(\mathcal{V})$ and $r_{i_J} > R(\underline{S}(\mathcal{V}), \boldsymbol{v}^*) = R^*_{\text{det}}(\mathcal{V})$ follows from the fact that i_J is an element of $\underline{S}(\mathcal{V})$ and the definition of $\underline{S}(\mathcal{V})$. For the induction step, assume that $r_{i_{j+1}} > R(S^*(\boldsymbol{v}^*) \cup \{i_1, \cdots, i_{j+1}\}, \boldsymbol{v}^*)$. Lemma 3.1 (a) and (c) of Rusmevichientong and Topaloglu (2012) then imply that $r_{i_{j+1}} > R(S^*(\boldsymbol{v}^*) \cup \{i_1, \cdots, i_j\}, \boldsymbol{v}^*)$. Then $r_{i_j} \geq r_{i_{j+1}}$ implies $r_{i_j} > R(S^*(\boldsymbol{v}^*) \cup \{i_1, \cdots, i_j\}, \boldsymbol{v}^*)$ as well. □

Proof of Corollary 1. The inequality $\max_{S\subseteq\mathcal{N}} \min_{\boldsymbol{v}\in\mathcal{V}} R(S,\boldsymbol{v}) \leq \min_{\boldsymbol{v}\in\mathcal{V}} \max_{S\subseteq\mathcal{N}} R(S,\boldsymbol{v})$ directly follows from weak duality. To see the reverse inequality, fix $S^*(\boldsymbol{v}^*)$ and $S^*(\mathcal{V})$ as in Lemma 1, and note that

$$\min_{\boldsymbol{v}\in\mathcal{V}} \max_{S\subseteq\mathcal{N}} R(S,\boldsymbol{v}) \;\leq\; \max_{S\subseteq\mathcal{N}} R(S,\boldsymbol{v}^*) \;=\; R(S^*(\boldsymbol{v}^*),\boldsymbol{v}^*) \;=\; R(S^*(\mathcal{V}),\boldsymbol{v}^*) \;=\; \max_{S\subseteq\mathcal{N}} \min_{\boldsymbol{v}\in\mathcal{V}} R(S,\boldsymbol{v}).$$

Here, the inequality follows from the fact that $\boldsymbol{v}^*\in\mathcal{V}$, and the second equality uses part 2 of Lemma 1. □

Proof of Theorem 1. Fix any probability vector $\boldsymbol{p}\in\Delta$ in RANDOMIZED ROBUST as well as any optimal assortment $S^*(\mathcal{V})\in\mathcal{S}^*(\mathcal{V})$ in DETERMINISTIC ROBUST. We then have

$$\begin{aligned}
\min_{\boldsymbol{v}\in\mathcal{V}} \sum_{S\subseteq\mathcal{N}} p_S\cdot R(S,\boldsymbol{v}) \;&\leq\; \sum_{S\subseteq\mathcal{N}} p_S\cdot R(S,\boldsymbol{v}^*) && \text{for } \boldsymbol{v}^*\in\arg\min_{\boldsymbol{v}\in\mathcal{V}} R(S^*(\mathcal{V}),\boldsymbol{v})\\
&\leq\; \sum_{S\subseteq\mathcal{N}} p_S\cdot R(S^*(\boldsymbol{v}^*),\boldsymbol{v}^*) && \text{for } S^*(\boldsymbol{v}^*)\in\arg\max_{S\subseteq\mathcal{N}} R(S,\boldsymbol{v}^*)\\
&=\; R(S^*(\boldsymbol{v}^*),\boldsymbol{v}^*) \;=\; R(S^*(\mathcal{V}),\boldsymbol{v}^*),
\end{aligned}$$

where the first inequality holds since $\boldsymbol{v}^*\in\mathcal{V}$ and thus $\boldsymbol{v}^*$ is considered in the minimization on the left-hand side. Similarly, the second inequality follows since each $S\subseteq\mathcal{N}$ and thus each S is considered in the maximization that defines $S^*(\boldsymbol{v}^*)$. The two identities, finally, follow from the fact that $\sum_{S\subseteq\mathcal{N}} p_S = 1$ as well as from the no-regret property of Lemma 1, respectively. The statement of the theorem follows since the choice of the probability vector $\boldsymbol{p}$ was arbitrary. □

A.2. The Cardinality-Constrained MNL Model

For our analysis of the cardinality-constrained MNL model, we consider for any number of products $n\geq 2$ and any cardinality restriction $|S|\leq C$, $C\in\{1,\cdots,n-1\}$ the instance where the price of each product i, $i=1,\cdots,n$, is $r_i=1$, while the uncertainty set is $\mathcal{V}=\{(1,\boldsymbol{v}) : \sum_{i\in\mathcal{N}} v_i = n-C,\;\; \boldsymbol{v}\in\{0,1\}^n\}$. In the remainder of this section, we refer to this model as the *reference model*.

The proof of Theorem 2 relies on the following auxiliary results, which we state and prove first.

LEMMA 2. *Under the reference model, the worst-case expected revenues of any feasible solution to the deterministic robust assortment optimization problem is* 0.

Proof of Lemma 2. Fix any feasible solution $S\subseteq\mathcal{N}$, $|S|\leq C$, to the deterministic robust assortment optimization problem. Under any valuation scenario $\boldsymbol{v}=(1,\boldsymbol{v}')\in\mathcal{V}$ with $v'_i=0$ if $i\in S$, the expected revenues of S vanish. The statement now follows from the fact that the worst-case expected revenues are always non-negative. □

LEMMA 3. *Under the reference model, the randomized assortment strategy $\boldsymbol{p}^*\in\Delta$, where $p^*_S = 1/|\{S\subseteq\mathcal{N} : |S|=C\}|$ if $|S|=C$ and $p^*_S=0$ otherwise, strictly outperforms any other feasible randomized assortment strategy $\boldsymbol{p}$ that only places positive probability $p_S>0$ mass on assortments S of full cardinality $|S|=C$.*

Proof of Lemma 3. Assume to the contrary that there is a feasible randomized assortment strategy $\boldsymbol{p}\neq\boldsymbol{p}^*$ that *(i)* places positive probability $p_S>0$ only on assortment S of cardinality $|S|=C$ and that *(ii)* has worst-case revenues larger than or equal to those of $\boldsymbol{p}^*$. Since $\boldsymbol{p}\neq\boldsymbol{p}^*$, there must be two components i and j of $\boldsymbol{p}$—without loss of generality $i=1$ and $j=2$—such that $p_i>p_j$ and the associated assortments

S_i and S_j satisfy $|S_i| = |S_j| = C$. Let Π denote the group of all permutations $\pi : \mathbb{R}^{2^n} \to \mathbb{R}^{2^n}$ that permute those components of a randomization strategy in Δ that correspond to assortments S of cardinality $|S| = C$. Moreover, assume that $\pi_1 \in \Pi$ is the identity and that $\pi_2 \in \Pi$ reverses the order of the first two components but leaves all other components unchanged.

Consider now the randomization strategy $\boldsymbol{p}'$ defined through

$$\begin{aligned}\boldsymbol{p}' &= \frac{1}{|\Pi \setminus \{\pi_1, \pi_2\}|} \sum_{\pi \in \Pi \setminus \{\pi_1, \pi_2\}} \pi(\boldsymbol{p}) + \left(\frac{1}{|\Pi|} + \frac{\epsilon}{p_1 - p_2}\right) \cdot \pi_1(\boldsymbol{p}) + \left(\frac{1}{|\Pi|} - \frac{\epsilon}{p_1 - p_2}\right) \cdot \pi_2(\boldsymbol{p}) \\ &= \boldsymbol{p}^* + \epsilon(+1, -1, 0, \cdots, 0)^\top.\end{aligned}$$

For $\epsilon > 0$ sufficiently small, $\boldsymbol{p}'$ is a convex combination of the permutations $\pi(\boldsymbol{p})$, $\pi \in \Pi$, and as such $\boldsymbol{p}' \in \Delta$. Moreover, we have $p'_S > 0$ only for assortments $S \subseteq \mathcal{N}$ with $|S| = C$. By the symmetry of the problem instance, any permutation $\pi(\boldsymbol{p})$, $\pi \in \Pi$, has the same worst-case expected revenues as $\boldsymbol{p}$. The concavity of the worst-case expected revenues then implies that $\boldsymbol{p}'$ has worst-case expected revenues that are at least as large as those of $\boldsymbol{p}$ and thus—by our earlier assumption—of $\boldsymbol{p}^*$. We next show that this leads to a contradiction, which will complete the proof.

To compare the worst-case expected revenues of $\boldsymbol{p}'$ and $\boldsymbol{p}^*$, we note that

$$\begin{aligned}&\min_{\boldsymbol{v} \in \mathcal{V}} \sum_{S \subseteq \mathcal{N}} p'_S \cdot R(S, \boldsymbol{v}) - \min_{\boldsymbol{v} \in \mathcal{V}} \sum_{S \subseteq \mathcal{N}} p^*_S \cdot R(S, \boldsymbol{v}) \\ = \;& \min_{\boldsymbol{v} \in \mathcal{V}} \sum_{S \subseteq \mathcal{N}} [p'_S - p^*_S] \cdot R(S, \boldsymbol{v}) \\ = \;& \min_{\boldsymbol{v} \in \mathcal{V}} \epsilon \cdot [R(S_1, \boldsymbol{v}) - R(S_2, \boldsymbol{v})] \\ = \;& \min_{\boldsymbol{v} \in \mathcal{V}} \epsilon \cdot \left[\frac{\sum_{i \in S_1} v_i}{v_0 + \sum_{i \in S_1} v_i} - \frac{\sum_{i \in S_2} v_i}{v_0 + \sum_{i \in S_2} v_i}\right] \\ = \;& \min_{\boldsymbol{v} \in \mathcal{V}} \epsilon \cdot \left[\frac{\sum_{i \in S_1 \cap S_2} v_i + \sum_{i \in S_1 \setminus S_2} v_i}{v_0 + \sum_{i \in S_1 \cap S_2} v_i + \sum_{i \in S_1 \setminus S_2} v_i} - \frac{\sum_{i \in S_1 \cap S_2} v_i + \sum_{i \in S_2 \setminus S_1} v_i}{v_0 + \sum_{i \in S_1 \cap S_2} v_i + \sum_{i \in S_2 \setminus S_1} v_i}\right]. \qquad (10)\end{aligned}$$

Here, the first identity holds since the symmetry of $\boldsymbol{p}^*$ implies that any valuation scenario $\boldsymbol{v} \in \mathcal{V}$ is a worst-case valuation for $\boldsymbol{p}^*$. The second identity assumes that S_1 and S_2 are the assortments corresponding to the first two components of $\boldsymbol{p}'$ and $\boldsymbol{p}^*$, and it follows from the construction of $\boldsymbol{p}'$. By choosing $\boldsymbol{v} \in \mathcal{V}$ such that as many components $i \in S_1 \setminus S_2$ of $\boldsymbol{v}$ as possible are 0 and as many components $i \in S_2 \setminus S_1$ of $\boldsymbol{v}$ as possible are 1, we can see that the expression (10) must be strictly negative. Since this implies that the worst-case revenues of $\boldsymbol{p}'$ are strictly lower than those of $\boldsymbol{p}^*$, we obtain the desired contradiction. □

Lemma 4. *Under the reference model, any optimal randomization strategy $\boldsymbol{p} \in \Delta$ satisfies $p_S = 0$ for all $S \subseteq \mathcal{N}$ with $|S| \neq C$.*

Proof of Lemma 4. Note first that any optimal randomization strategy $\boldsymbol{p}$ has to be feasible and hence place zero probability on any assortment $S \subseteq \mathcal{N}$, $|S| > C$, violating the cardinality constraint. Assume to the contrary that there is an optimal randomization strategy $\boldsymbol{p} \in \Delta$ with $p_S > 0$ for some $S \subseteq \mathcal{N}$ satisfying $|S| < C$. We show that $\boldsymbol{p}$ can be transformed into another optimal randomization strategy $\boldsymbol{p}' \in \Delta$ satisfying $p'_S = 0$ for all $S \subseteq \mathcal{N}$ with $|S| < C$ and $\boldsymbol{p}' \neq \boldsymbol{p}^*$, where $\boldsymbol{p}^*$ is defined in Lemma 3. Since this contradicts the statement of Lemma 3, we obtain the desired contradiction.

We construct $\boldsymbol{p}'$ as follows. For any assortment $S \subseteq \mathcal{N}$ with $|S| \leq C$, let $\theta(S) \subseteq \mathcal{N}$ be an assortment satisfying $S \subseteq \theta(S)$ and $|\theta(S)| = C$; for all other assortments S, we define $\theta(S) = S$. We set $p'_T = \sum_{S \subseteq \mathcal{N}\,:\,\theta(S)=T} p_S$ for all $T \subseteq \mathcal{N}$ with $|T| = C$, and $p'_T = 0$ otherwise. For later use we note that there are multiple ways to define the mapping θ, and that multiple different probability vectors $\boldsymbol{p}'$ can be obtained through different θ. We have $\boldsymbol{p}' \in \Delta$, and the worst-case expected revenues of $\boldsymbol{p}' \in \Delta$ are at least as large as those of $\boldsymbol{p}$ since for any valuation scenario $\boldsymbol{v} \in \mathcal{V}$, we have

$$\begin{aligned}\sum_{S \subseteq \mathcal{N}} p'_S \cdot R(S, \boldsymbol{v}) &= \sum_{\substack{S \subseteq \mathcal{N}: \\ |S| = C}} p'_S \cdot R(S, \boldsymbol{v}) = \sum_{\substack{S \subseteq \mathcal{N}: \\ |S| = C}} p_S \cdot R(S, \boldsymbol{v}) + \sum_{\substack{S \subseteq \mathcal{N}: \\ |S| < C}} p_S \cdot R(\theta(S), \boldsymbol{v}) \\ &\geq \sum_{\substack{S \subseteq \mathcal{N}: \\ |S| = C}} p_S \cdot R(S, \boldsymbol{v}) + \sum_{\substack{S \subseteq \mathcal{N}: \\ |S| < C}} p_S \cdot R(S, \boldsymbol{v}) = \sum_{S \subseteq \mathcal{N}} p_S \cdot R(S, \boldsymbol{v}).\end{aligned}$$

Here, the inequality holds since for every $\boldsymbol{v} \in \mathcal{V}$, we have

$$R(S, \boldsymbol{v}) = \frac{\sum_{i \in S} v_i}{v_0 + \sum_{i \in S} v_i} \leq \frac{\sum_{i \in S} v_i + \sum_{i \in \theta(S) \setminus S} v_i}{v_0 + \sum_{i \in S} v_i + \sum_{i \in \theta(S) \setminus S} v_i} = R(\theta(S), \boldsymbol{v}), \tag{11}$$

where the inequality follows from the fact that $a/b \leq (a+c)/(b+c)$ whenever $0 \leq a < b$ and $c \geq 0$.

If $\boldsymbol{p}' \neq \boldsymbol{p}^*$, where $\boldsymbol{p}^*$ is defined in Lemma 3, then we have the desired contradiction to Lemma 3. Otherwise, if $\boldsymbol{p}' = \boldsymbol{p}^*$, then we can exercise our aforementioned flexibility in the choice of θ to derive another optimal randomization strategy $\boldsymbol{p}'' \neq \boldsymbol{p}'$ that results in the desired contradiction. □

Proof of Theorem 2. The first statement follows immediately from Lemmas 2 and 3 and the fact that for any valuation scenario $\boldsymbol{v} \in \mathcal{V}$, the optimal randomized strategy $\boldsymbol{p}^*$ places strictly positive probability on assortments S that include products $i \in \mathcal{N}$ with valuation $v_i = 1$. The second statement, on the other hand, follows immediately from Lemmas 3 and 4. □

Proof of Corollary 2. The proof of this corollary follows directly from the first part of Theorem 2. □

Proof of Corollary 3. We consider our reference model for an even number of products n and set $C = n/2$. (An analogous argument can be made for the case where n is odd.) In this case, the unique optimal randomization strategy $\boldsymbol{p}^*$ defined in Lemma 3 places strictly positive probability on

$$\binom{n}{C} = \frac{n!}{(n-C)!C!} = \frac{n!}{\frac{n}{2}!\frac{n}{2}!} \approx \frac{\sqrt{2\pi n}(\frac{n}{e})^n}{\left(\sqrt{2\pi \frac{n}{2}}(\frac{n}{2e})^{\frac{n}{2}}\right)^2} = \sqrt{\frac{2}{\pi}} n^{-\frac{1}{2}} 2^n$$

different assortments, where the approximation is due to Stirling's formula. □

A.3. Solving the Randomized Constrained MNL Problem

In the following, we present mixed-integer linear programs (MILPs) for the subproblems (5) and (6) arising in our column generation scheme for problem (3). Our reformulations resemble the MILP fornulations for sum-of-ratio problems, see Li (1994).

For the subproblem (5) in the primal step of our column generation scheme, we first replace the decision variables $\boldsymbol{v} \in \mathcal{V}$ with their definition and introduce auxiliary variables d_S, $S \in \hat{\mathcal{N}}^C$, that equate to the (inverses of the) denominators in the objective function:

$$\begin{array}{lll} \text{minimize} & \displaystyle\sum_{S \in \hat{\mathcal{N}}^C} p_S \cdot d_S \cdot \sum_{i \in S} r_i \cdot \boldsymbol{F}_i^\top \boldsymbol{\xi} & \\ \text{subject to} & d_S \cdot \left[\displaystyle\sum_{i \in S \cup \{0\}} \boldsymbol{F}_i^\top \boldsymbol{\xi} \right] = 1 & \forall S \in \hat{\mathcal{N}}^C \\ & \boldsymbol{A}\boldsymbol{\xi} \leq \boldsymbol{b} & \\ & \boldsymbol{\xi} \in \{0,1\}^m, \;\; d_S \in \mathbb{R}_+, \, S \in \hat{\mathcal{N}}^C. & \end{array}$$

Here, $\boldsymbol{F}_i^\top$ denotes the i-th row of the matrix $\boldsymbol{F}$ (where the counting starts from zero). For all $S \in \hat{\mathcal{N}}^C$ and $j \in \{1, \cdots, m\}$, we now apply the following exact linearization of the bilinear terms:

$$z_{Sj} = d_S \cdot \xi_j \quad \Longleftrightarrow \quad \begin{bmatrix} z_{Sj} \leq d_S, & z_{Sj} \leq \mathrm{M} \cdot \xi_j \\ z_{Sj} \geq 0, & z_{Sj} \geq d_S - \mathrm{M} \cdot (1 - \xi_j) \end{bmatrix}$$

Here, M denotes a sufficiently large positive constant that can be determined from the geometry of the uncertainty set (Li 1994). This results in the following MILP.

$$\begin{array}{lll} \text{minimize} & \displaystyle\sum_{S \in \hat{\mathcal{N}}^C} \sum_{i \in S} \sum_{j=1}^{m} p_S \cdot r_i \cdot F_{ij} \cdot z_{Sj} & \\ \text{subject to} & \displaystyle\sum_{i \in S \cup \{0\}} \sum_{j=1}^{m} F_{ij} z_{Sj} = 1 & \forall S \in \hat{\mathcal{N}}^C \\ & z_{Sj} \leq d_S, \;\; z_{Sj} \leq \mathrm{M} \cdot \xi_j & \forall S \in \hat{\mathcal{N}}^C, \, \forall j \in \{1, \cdots, m\} \\ & z_{Sj} \geq d_S - \mathrm{M} \cdot (1 - \xi_j) & \forall S \in \hat{\mathcal{N}}^C, \, \forall j \in \{1, \cdots, m\} \\ & \boldsymbol{A}\boldsymbol{\xi} \leq \boldsymbol{b} & \\ & \boldsymbol{\xi} \in \{0,1\}^m, \;\; d_S \in \mathbb{R}_+, \, S \in \hat{\mathcal{N}}^C, \;\; z_{Sj} \in \mathbb{R}_+, \, S \in \hat{\mathcal{N}}^C \text{ and } j \in \{1, \cdots, m\}. & \end{array}$$

In view of the subproblem (6) in the dual step of our algorithm, we first introduce the binary decision variable $x_i \in \{0,1\}$, $i \in \mathcal{N}$, to denote whether product i is contained in the assortment S and introduce auxiliary variables $d_{\boldsymbol{v}}$, $\boldsymbol{v} \in \hat{\mathcal{V}}$, that equate to the (inverses of the) denominators in the objective function:

$$\begin{array}{lll} \text{maximize} & \displaystyle\sum_{\boldsymbol{v} \in \hat{\mathcal{V}}} \lambda_{\boldsymbol{v}} \cdot d_{\boldsymbol{v}} \cdot \sum_{i \in \mathcal{N}} r_i v_i x_i & \\ \text{subject to} & d_{\boldsymbol{v}} \cdot \left[v_0 + \displaystyle\sum_{i \in \mathcal{N}} v_i x_i \right] = 1 & \forall \boldsymbol{v} \in \hat{\mathcal{V}} \\ & \displaystyle\sum_{i \in \mathcal{N}} x_i \leq C & \\ & \boldsymbol{x} \in \{0,1\}^n, \;\; d_{\boldsymbol{v}} \in \mathbb{R}_+, \, \boldsymbol{v} \in \hat{\mathcal{V}} & \end{array}$$

For all $\boldsymbol{v} \in \hat{\mathcal{V}}$ and $i \in \mathcal{N}$, we now apply the following exact linearization of the bilinear terms:

$$z_{\boldsymbol{v}i} = d_{\boldsymbol{v}} \cdot x_i \quad \Longleftrightarrow \quad \begin{bmatrix} z_{\boldsymbol{v}i} \leq d_{\boldsymbol{v}}, & z_{\boldsymbol{v}i} \leq \mathrm{M} \cdot x_i \\ z_{\boldsymbol{v}i} \geq 0, & z_{\boldsymbol{v}i} \geq d_{\boldsymbol{v}} - \mathrm{M} \cdot (1 - x_i) \end{bmatrix}$$

Here, M denotes a sufficiently large positive constant that can be determined from the geometry of the uncertainty set (Li 1994). This results in the following MILP.

$$\begin{array}{lll} \text{maximize} & \displaystyle\sum_{\boldsymbol{v} \in \hat{\mathcal{V}}} \sum_{i \in \mathcal{N}} \lambda_{\boldsymbol{v}} \cdot r_i \cdot v_i \cdot z_{\boldsymbol{v}i} & \\ \text{subject to} & v_0 \cdot d_{\boldsymbol{v}} + \displaystyle\sum_{i \in \mathcal{N}} v_i \cdot z_{\boldsymbol{v}i} = 1 & \forall \boldsymbol{v} \in \hat{\mathcal{V}} \\ & z_{\boldsymbol{v}i} \leq d_{\boldsymbol{v}}, \;\; z_{\boldsymbol{v}i} \leq \mathrm{M} \cdot x_i & \\ & z_{\boldsymbol{v}i} \geq d_{\boldsymbol{v}} - \mathrm{M} \cdot (1 - x_i) & \\ & \displaystyle\sum_{i \in \mathcal{N}} x_i \leq C & \\ & \boldsymbol{x} \in \{0,1\}^n, \;\; d_{\boldsymbol{v}} \in \mathbb{R}_+, \, \boldsymbol{v} \in \hat{\mathcal{V}}, \;\; z_{\boldsymbol{v}i} \in \mathbb{R}_+, \, \boldsymbol{v} \in \hat{\mathcal{V}} \text{ and } i \in \mathcal{N} & \end{array}$$

Proof of Theorem 3. The proof largely follows similar arguments as the proof of Theorem 2 of Delage and Saif (2018) and is thus omitted for the sake of brevity. □

Appendix B: Proofs for Section 5

B.1. A Robust MDP Reformulation for the MC Model

Proof of Observation 1. As per Definition 2, we need to verify that the initial state distribution q is a probability distribution, that for all transition kernels $p \in \mathcal{P}$ and all $s \in \mathcal{S}$ and $a \in \mathcal{A}$, $p(\cdot|s,a)$ is a probability distribution, and that the discount factor γ is contained in the open interval $(0,1)$.

Since $\boldsymbol{\lambda}$ is a probability distribution over $\mathcal{N}_0$, so is q. The discount factor γ is strictly less than 1 since $\rho_{i0} > 0$ for all $\boldsymbol{\rho} \in \mathcal{U}$ and $i \in \mathcal{N}$. At the same time, γ is strictly positive since $\rho_{i0} < 1$ for at least one $\boldsymbol{\rho} \in \mathcal{U}$ and $i \in \mathcal{N}$. Next, fix any transition kernel $p \in \mathcal{P}$ and state $s \in \mathcal{S}$. For the action $a = \top$, there is precisely one state $s' \in \mathcal{S}$ for which $p(s'|s,a) = 1$, while $p(s''|s,a) = 0$ for all other states $s'' \in \mathcal{S}$. For the action $a = \perp$, we note that for any state $s' \neq 0$, we have

$$0 \;\leq\; \rho_{ss'}/\gamma \;=\; p(s'|s,a) \;\leq\; \sum_{s'' \in \mathcal{N}} p(s''|s,a) \;=\; \sum_{s'' \in \mathcal{N}} \rho_{ss''}/\gamma \;=\; (1-\rho_{s0})/\gamma \;\leq\; 1,$$

where the first two inequalities hold since $\rho_{ss''} \geq 0$ for all $s'' \neq 0$ and $\gamma > 0$ by construction, the last equality holds since $(\rho_{s0}, \cdots, \rho_{sn})$ is a probability distribution over $\mathcal{N}_0$, and the last inequality follows from the construction of γ. We thus have $p(s'|s,a) \in [0,1]$ for all $s' \neq 0$ as well as $\sum_{s'' \in \mathcal{N}} p(s''|s,a) \in [0,1]$. It then follows directly from our construction in (7) that $p(0|s,a) \in [0,1]$ as well as $\sum_{s'' \in \mathcal{S}} p(s''|s,a) = 1$. □

Proof of Theorem 4. For any assortment $S \subseteq \mathcal{N}$, the worst-case expected revenues $R(S,\mathcal{U})$ are given by equation (13) of Désir et al. (2019):

$$\begin{array}{lll}
\text{minimize} & \displaystyle\sum_{i \in \mathcal{N}_0} \lambda_i g_i & \\
\text{subject to} & g_i = r_i & \forall i \in S \\
& g_i = \displaystyle\sum_{j \in \mathcal{N}} \rho_{ij} g_j & \forall i \in \mathcal{N} \setminus S \\
& \boldsymbol{\rho} \in \mathcal{U},\;\; \boldsymbol{g} \in \mathbb{R}^{n+1}_+. &
\end{array}$$

Similarly, for a policy $\pi_S \in \Pi$, the worst-case expected total discounted reward is given by equation (28) of Nilim and El Ghaoui (2005):

$$\begin{array}{lll}
\text{minimize} & \displaystyle\sum_{s \in \mathcal{S}} q_s v_s & \\
\text{subject to} & v_s \geq (1-\gamma) r_s + \gamma v_s & \forall s \in \mathcal{S} : \pi_S(s) = \top \\
& v_s \geq 0 + \gamma \displaystyle\sum_{s' \in \mathcal{S}} p(s'|s,\perp) v_{s'} & \forall s \in \mathcal{S} : \pi_S(s) = \perp \\
& p \in \mathcal{P},\;\; \boldsymbol{v} \in \mathbb{R}^{n+1}_+. &
\end{array}$$

Lemma 2 of Nilim and El Ghaoui (2005) implies that there is an optimal solution to this problem that satisfies all inequalities as equalities. We can thus replace all inequalities with equalities, and substituting p with its definition yields

$$\begin{array}{lll}
\text{minimize} & \displaystyle\sum_{s \in \mathcal{S}} q_s v_s & \\
\text{subject to} & v_s = r_s & \forall s \in \mathcal{S} : \pi_S(s) = \top \\
& v_s = \displaystyle\sum_{s' \in \mathcal{S}} \rho_{ss'} v_{s'} & \forall s \in \mathcal{S} : \pi_S(s) = \perp \\
& \boldsymbol{\rho} \in \mathcal{U},\;\; \boldsymbol{v} \in \mathbb{R}^{n+1}_+. &
\end{array}$$

The statement now follows from the fact that $\boldsymbol{q} = \boldsymbol{\lambda}$ and $\mathcal{S} = \mathcal{N}_0$. □

B.2. Product-Wise Substitution Sets

Proof of Theorem 5. We first show that a product-wise substitution set in the MC model implies that the ambiguity set of the associated robust MDP in Definition 3 is (s,a)-rectangular. We can then use established results from the robust MDP literature to prove the statement of the theorem.

A robust MDP has an (s,a)-rectangular ambiguity set whenever

$$\mathcal{P} = \{p : \mathcal{S}\times\mathcal{A}\times\mathcal{S}\to\mathbb{R}_+ \,:\, \forall s\in\mathcal{S},\ a\in\mathcal{A}\ \exists p^{sa}\in\mathcal{P} \text{ such that } p(\cdot|s,a) = p^{sa}(\cdot|s,a)\},$$

see Iyengar (2005), Nilim and El Ghaoui (2005) and Wiesemann et al. (2013). In other words, for any selection of conditional distributions p^{sa}, $s\in\mathcal{S}$ and $a\in\mathcal{A}$, we need to show that the composite distribution p satisfying $p(\cdot|s,a) = p^{sa}(\cdot|s,a)$ for all $s\in\mathcal{S}$ and $a\in\mathcal{A}$ is also contained in the ambiguity set $\mathcal{P}$. Fix any such selection of conditional distributions p^{sa}, $s\in\mathcal{S}$ and $a\in\mathcal{A}$, which by Definition 3 must correspond to a selection of substitution matrices $\boldsymbol{\rho}^{sa}\in\mathcal{U}$, $s\in\mathcal{S}$ and $a\in\mathcal{A}$. Consider the substitution matrix $\boldsymbol{\rho}$ whose i-th row coincides with the i-th row of $\boldsymbol{\rho}^{i\perp}$, $i=0,\cdots,n$. Our definition of product-wise substitution sets guarantees that $\boldsymbol{\rho}\in\mathcal{U}$. Moreover, since $p(\cdot|s,\top)$ does not depend on the substitution matrix, the transition kernel p associated with $\boldsymbol{\rho}$ satisfies $p(\cdot|s,a) = p^{sa}(\cdot|s,a)$ for all $s\in\mathcal{S}$ and $a\in\mathcal{A}$ as desired. We thus conclude that the ambiguity set of the associated robust MDP in Definition 3 is (s,a)-rectangular.

Theorem 4 of Nilim and El Ghaoui (2005) shows that strong duality holds for robust MDPs with (s,a)-rectangular ambiguity sets, that is,

$$\max_{\pi\in\Pi}\min_{p\in\mathcal{P}} \mathbb{E}\left[\sum_{t=0}^{\infty}\gamma^t\cdot r(s_t,a_t)\,\middle|\, s_0\sim q\right] \;=\; \min_{p\in\mathcal{P}}\max_{\pi\in\Pi} \mathbb{E}\left[\sum_{t=0}^{\infty}\gamma^t\cdot r(s_t,a_t)\,\middle|\, s_0\sim q\right].$$

Theorem 4 of our paper implies that

$$\max_{S\subseteq\mathcal{N}} R(S,\mathcal{U}) \;=\; \max_{\pi\in\Pi}\min_{p\in\mathcal{P}} \mathbb{E}\left[\sum_{t=0}^{\infty}\gamma^t\cdot r(s_t,a_t)\,\middle|\, s_0\sim q\right].$$

Moreover, the same theorem implies that

$$\min_{p\in\mathcal{P}}\max_{\pi\in\Pi} \mathbb{E}\left[\sum_{t=0}^{\infty}\gamma^t\cdot r(s_t,a_t)\,\middle|\, s_0\sim q\right] \;=\; \min_{\boldsymbol{\rho}\in\mathcal{U}}\max_{S\subseteq\mathcal{N}} R(S,\boldsymbol{\rho})$$

since Definition 3 implies the existence of a one-to-one correspondence between transition kernels $p\in\mathcal{P}$ and substitution matrices $\boldsymbol{\rho}\in\mathcal{U}$, and Theorem 4 of our paper can be applied individually to each singleton uncertainty set $\{\boldsymbol{\rho}\}$ constructed from any substitution matrix $\boldsymbol{\rho}\in\mathcal{U}$. We have thus established strong duality for our robust assortment optimization problem, that is,

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

The right-hand side of this equation bounds from above the worst-case expected revenues achievable by any (deterministic or randomized) assortment strategy. Since this bound is attained by the optimal deterministic assortment on the left-hand side of this equation, we cannot improve by adopting a randomized assortment strategy. □

Proof of Theorem 6. For any $n \geq 3$ and $C \leq n-2$, we construct the following instance of the cardinality-constrained robust MC problem. The product-wise revenues are $r_1 = 0$ and $r_i = 1$, $i \in \{2, \cdots, n\}$, the initial choice probabilities satisfy $\lambda_1 = 1$ and $\lambda_i = 0$, $i \neq 1$, and the uncertainty set $\mathcal{U}$ contains all substitution matrices $\boldsymbol{\rho}$ satisfying $\boldsymbol{\rho}_0 = (1, 0, \cdots, 0)$, $\boldsymbol{\rho}_1 = (0, \boldsymbol{\rho}')$ where $\boldsymbol{\rho}'$ is any element of the probability simplex in $\mathbb{R}^n$, as well as $\boldsymbol{\rho}_i$ being any element of the probability simplex in $\mathbb{R}^{n+1}$, $i \in \{2, \cdots, n\}$. Note that by construction, this uncertainty set exhibits product-wise substitution sets.

The worst-case expected revenues of every deterministic assortment $S \subseteq \mathcal{N}$ satisfying $|S| \leq C$ is 0. This is clearly the case if $1 \in S$. If $1 \notin S$, fix any product $i \in \mathcal{N} \setminus (S \cup \{1\})$, which exists since $C \leq n-2$. For any substitution matrix $\boldsymbol{\rho}^i$ satisfying $\rho^i_{1i} = \rho^i_{i0} = 1$, we have $R(S, \boldsymbol{\rho}^i) = 0$.

We now consider any randomized assortment strategy that places a strictly positive probability on every assortment $S \subseteq \mathcal{N} \setminus \{1\}$ satisfying $|S| \leq C$ and that places zero probability on all other assortments. Fix any substitution matrix $\boldsymbol{\rho} \in \mathcal{U}$. Since $\rho_{1i} > 0$ for at least one $i \in \{2, \ldots, n\}$ and since this product i will be offered with strictly positive probability, the randomized assortment strategy must generate strictly positive expected revenues under the substitution matrix $\boldsymbol{\rho}$. Since $\boldsymbol{\rho}$ was arbitrary, the worst-case expected revenues of the randomized assortment strategy must be strictly positive as well. □

Proof of Corollary 4. The proof of this corollary follows directly from Theorem 6. □

B.3. General Substitution Sets

Proof of Proposition 1. For $n = 2$, we can offer three different assortments: $\{1\}$, $\{2\}$ and $\{1, 2\}$. Fix any worst-case substitution matrices $\boldsymbol{\rho}^1 \in \mathcal{U}$ and $\boldsymbol{\rho}^2 \in \mathcal{U}$ for the two assortments $\{1\}$ and $\{2\}$, respectively (they may not be unique). The worst-case expected revenues for the three assortments amount to

$$R(\{1\}, \boldsymbol{\rho}^1) = \lambda_1 r_1 + \lambda_2 \rho^1_{21} r_1, \quad R(\{2\}, \boldsymbol{\rho}^2) = \lambda_1 \rho^2_{12} r_2 + \lambda_2 r_2 \quad \text{and} \quad R(\{1, 2\}, \boldsymbol{\rho}) = \lambda_1 r_1 + \lambda_2 r_2;$$

note in particular that the expected revenues of the assortment $\{1, 2\}$ do not depend on the realized substitution matrix $\boldsymbol{\rho} \in \mathcal{U}$. We show that the problem is randomization-proof by considering two cases: *(i)* $R(\{1, 2\}, \boldsymbol{\rho}) \leq \max\{R(\{1\}, \boldsymbol{\rho}^1), R(\{2\}, \boldsymbol{\rho}^2)\}$ and *(ii)* $R(\{1, 2\}, \boldsymbol{\rho}) > \max\{R(\{1\}, \boldsymbol{\rho}^1), R(\{2\}, \boldsymbol{\rho}^2)\}$.

Case *(i)*. We assume that $R(\{1, 2\}, \boldsymbol{\rho}) \leq R(\{1\}, \boldsymbol{\rho}^1)$; the case where $R(\{1, 2\}, \boldsymbol{\rho}) \leq R(\{2\}, \boldsymbol{\rho}^2)$ is symmetric. Since $R(\{1, 2\}, \boldsymbol{\rho}) \leq R(\{1\}, \boldsymbol{\rho}^1)$, we have $\lambda_1 r_1 + \lambda_2 r_2 \leq \lambda_1 r_1 + \lambda_2 \rho^1_{21} r_1$ and thus $\lambda_2 = 0$ or $r_2 \leq \rho^1_{21} r_1$. Since only $R(\{2\}, \boldsymbol{\rho})$ depends on $\boldsymbol{\rho}$ if $\lambda_2 = 0$, it is easy to see that the problem is randomization proof. Assume now that $r_2 \leq \rho^1_{21} r_1$. For any substitution matrix $\boldsymbol{\rho} \in \mathcal{U}$, we then have

$$R(\{2\}, \boldsymbol{\rho}) \;=\; \lambda_1 \rho_{12} r_2 + \lambda_2 r_2 \;\leq\; \lambda_1 \rho_{12} \rho^1_{21} r_1 + \lambda_2 \rho^1_{21} r_1 \;\leq\; \lambda_1 r_1 + \lambda_2 \rho^1_{21} r_1 \;=\; R(\{1\}, \boldsymbol{\rho}^1), \tag{12}$$

where the first inequality holds since $r_2 \leq \rho^1_{21} r_1$, while the second inequality is due to the fact that $\rho_{12}, \rho^1_{21} \leq 1$. Now consider an arbitrary randomized assortment strategy $\alpha \cdot \{1\} + \beta \cdot \{2\} + \gamma \cdot \{1, 2\}$ with $\alpha, \beta, \gamma \geq 0$ and $\alpha + \beta + \gamma = 1$, and let $\boldsymbol{\rho}^{\mathrm{r}}$ be a worst-case substitution matrix for this strategy. We then have that

$$\alpha R(\{1\}, \boldsymbol{\rho}^{\mathrm{r}}) + \beta R(\{2\}, \boldsymbol{\rho}^{\mathrm{r}}) + \gamma R(\{1, 2\}, \boldsymbol{\rho}^{\mathrm{r}}) \;\leq\; \alpha R(\{1\}, \boldsymbol{\rho}^1) + \beta R(\{1\}, \boldsymbol{\rho}^1) + \gamma R(\{1\}, \boldsymbol{\rho}^1) \;=\; R(\{1\}, \boldsymbol{\rho}^1),$$

where the first inequality holds since $R(\{2\},\boldsymbol{\rho}^{\mathrm{r}}) \le R(\{1\},\boldsymbol{\rho}^1)$ by (12) and $R(\{1,2\},\boldsymbol{\rho}^{\mathrm{r}}) \le R(\{1\},\boldsymbol{\rho}^1)$ by assumption. We thus conclude that the problem is randomization-proof.

Case *(ii)*. We prove this case by contradiction. We first show that if a problem instance is randomization-receptive, then it is optimal to randomize between the singleton assortments $\{1\}$ and $\{2\}$. We then derive a necessary and sufficient condition for the randomization between $\{1\}$ and $\{2\}$ to be beneficial. Finally, we show that this condition contradicts our assumption that $R(\{1,2\},\boldsymbol{\rho}) > \max\{R(\{1\},\boldsymbol{\rho}^1),\, R(\{2\},\boldsymbol{\rho}^2)\}$.

Assume that the problem is randomization-receptive, and let $\alpha\cdot\{1\}+\beta\cdot\{2\}+\gamma\cdot\{1,2\}$ with $\alpha,\beta\ge 0,\gamma>0$ and $\alpha+\beta+\gamma=1$ be an optimal randomization strategy. Denote by $\boldsymbol{\rho}^{\mathrm{r}}$ one of its worst-case substitution matrices. We then have $\alpha R(\{1\},\boldsymbol{\rho}^{\mathrm{r}})+\beta R(\{2\},\boldsymbol{\rho}^{\mathrm{r}})+\gamma R(\{1,2\},\boldsymbol{\rho}^{\mathrm{r}}) > R(\{1,2\},\boldsymbol{\rho}^{\mathrm{r}})$, and hence $(1-\gamma)R(\{1,2\},\boldsymbol{\rho}^{\mathrm{r}}) < \alpha R(\{1\},\boldsymbol{\rho}^{\mathrm{r}})+\beta R(\{2\},\boldsymbol{\rho}^{\mathrm{r}})$. This implies that

$$\begin{aligned}\alpha R(\{1\},\boldsymbol{\rho}^{\mathrm{r}})+\beta R(\{2\},\boldsymbol{\rho}^{\mathrm{r}})+\gamma R(\{1,2\},\boldsymbol{\rho}^{\mathrm{r}}) \;&<\; \left[1+\frac{\gamma}{1-\gamma}\right](\alpha R(\{1\},\boldsymbol{\rho}^{\mathrm{r}})+\beta R(\{2\},\boldsymbol{\rho}^{\mathrm{r}}))\\ &=\; \frac{\alpha}{\alpha+\beta}R(\{1\},\boldsymbol{\rho}^{\mathrm{r}})+\frac{\beta}{\alpha+\beta}R(\{2\},\boldsymbol{\rho}^{\mathrm{r}}),\end{aligned}$$

where the inequality uses the fact that $(1-\gamma)R(\{1,2\},\boldsymbol{\rho}^{\mathrm{r}}) < \alpha R(\{1\},\boldsymbol{\rho}^{\mathrm{r}})+\beta R(\{2\},\boldsymbol{\rho}^{\mathrm{r}})$, while the equality holds since $\gamma = 1-\alpha-\beta$. Since $R(\{1,2\},\boldsymbol{\rho})$ does not depend on $\boldsymbol{\rho}$, we have $\boldsymbol{\rho}^{\mathrm{r}} \in \arg\min_{\boldsymbol{\rho}\in\mathcal{U}} \frac{\alpha}{\alpha+\beta}R(\{1\},\boldsymbol{\rho}) + \frac{\beta}{\alpha+\beta}R(\{2\},\boldsymbol{\rho})$. It is therefore optimal to set $\gamma=0$ and randomize only between the assortments $\{1\}$ and $\{2\}$.

The problem of randomizing between the assortments $\{1\}$ and $\{2\}$ can be written as

$$\begin{aligned}&\underset{\beta\in[0,1]}{\text{maximize}}\left[\min_{\boldsymbol{\rho}\in\mathcal{U}}\{\beta R(\{1\},\boldsymbol{\rho})+(1-\beta)R(\{2\},\boldsymbol{\rho})\}\right]\\ =\;&\underset{\beta\in[0,1]}{\text{maximize}}\left[\min_{\boldsymbol{\rho}\in\mathcal{U}}\{\lambda_1 r_1+\lambda_2 r_2+\beta\lambda_2(\rho_{21}r_1-r_2)+(1-\beta)\lambda_1(\rho_{12}r_2-r_1)\}\right].\end{aligned}$$

For the problem to be randomization-receptive, it thus needs to satisfy

$$\begin{aligned}&\underset{\beta\in[0,1]}{\text{maximize}}\left[\min_{\boldsymbol{\rho}\in\mathcal{U}}\{\lambda_1 r_1+\lambda_2 r_2+\beta\lambda_2(\rho_{21}r_1-r_2)+(1-\beta)\lambda_1(\rho_{12}r_2-r_1)\}\right] \;>\; \lambda_1\rho_{12}r_2+\lambda_2 r_2\\ \iff\;&\exists\beta\in(0,1) \text{ such that } \beta\lambda_2(\rho_{21}r_1-r_2)+(1-\beta)\lambda_1(\rho_{12}r_2-r_1)>0\quad\forall\boldsymbol{\rho}\in\mathcal{U},\end{aligned}\tag{13}$$

where we use our case assumption that the assortment $\{1,2\}$ generates the highest worst-case expected revenues. We now show that (13) contradicts the assumption that $R(\{1,2\},\boldsymbol{\rho}) > \max\{R(\{1\},\boldsymbol{\rho}^1),\, R(\{2\},\boldsymbol{\rho}^2)\}$.

By our case assumption, we have $\lambda_1 r_1+\lambda_2 r_2 > \lambda_1 r_1+\lambda_2\rho^1_{21}r_1$ and $\lambda_1 r_1+\lambda_2 r_2 > \lambda_2 r_2+\lambda_1\rho^2_{12}r_2$. These inequalities imply that $\lambda_1,\lambda_2>0$ as well as $r_2>\rho^1_{21}r_1$ and $r_1>\rho^2_{12}r_2$. For notational convenience, we introduce the constants a, b and c defined via:

$$r_2 = ar_1 \text{ for some } a\in(\rho^1_{21},\,1/\rho^2_{12}),\quad \lambda_2=b\lambda_1 \text{ for some } b\in\mathbb{R}_{++},\quad 1-\beta=c\beta \text{ for some } c\in\mathbb{R}_{++}.$$

Using this notation, we can re-express the inequality in (13) as

$$\begin{aligned}&\beta\lambda_2(\rho_{21}r_1-r_2)+(1-\beta)\lambda_1(\rho_{12}r_2-r_1)>0\\ \iff\;&\beta b\lambda_1(\rho_{21}r_1-ar_1)+c\beta\lambda_1(\rho_{12}ar_1-r_1)>0\\ \iff\;&b\rho_{21}-ba>c-ca\rho_{12}.\end{aligned}$$

Here, the first equivalence results from the definitions of a, b and c, whereas the second equivalence follows from dividing the expressions by $\beta\lambda_1 r_1>0$. Since this inequality has to hold for all $\boldsymbol{\rho}\in\mathcal{U}$, it has to hold in particular for $\boldsymbol{\rho}^1$ and $\boldsymbol{\rho}^2$, which respectively result in the conditions

$$\begin{aligned}&b\rho^1_{21}-ba \;>\; c-ca\rho^1_{12} \;\Longrightarrow\; b\rho^1_{21}-ba \;>\; c-ca \;\Longrightarrow\; b(\rho^1_{21}-a) \;>\; c(1-a) \;\Longrightarrow\; a \;>\; 1\\ \text{and}\quad &b\rho^2_{21}-ba \;>\; c-ca\rho^2_{12} \;\Longrightarrow\; b-ba \;>\; c-ca\rho^2_{12} \;\Longrightarrow\; b(1-a) \;>\; c(1-a\rho^2_{12}) \;\Longrightarrow\; a \;<\; 1.\end{aligned}$$

Here, the respective first implications hold since $\rho^1_{12}, \rho^2_{21} \in [0,1]$, and the last ones follow from the fact that $a \in (\rho^1_{21},\ 1/\rho^2_{12})$ and $b, c > 0$. This is a contradiction, and we thus conclude that the problem must be randomization-proof under case *(ii)* as well. □

Proof of Observation 2. We extend the randomization-receptive instance from Example 3 to $n > 3$ products as follows. The product-wise revenues are $\boldsymbol{r} = (0, 4.66, 1.00, 10.00, 1, \cdots, 1)$, the initial choice probabilities are $\boldsymbol{\lambda} = (0, 0.37, 0.62, 0.01, 0, \cdots, 0)$, and the uncertainty set is

$$\mathcal{U} = \left\{ \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0.01 & 0 & 0.05 & 0.94 & 0 & \cdots & 0 \\ 0.26 & 0.69 & 0 & 0.05 & 0 & \cdots & 0 \\ 0.90 & 0.05 & 0.05 & 0 & 0 & \cdots & 0 \\ 1 & 0 & 0 & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 0 & 0 & 0 & 0 & \cdots & 0 \end{pmatrix}, \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0.90 & 0 & 0.05 & 0.05 & 0 & \cdots & 0 \\ 0.01 & 0.44 & 0 & 0.55 & 0 & \cdots & 0 \\ 0.01 & 0.94 & 0.05 & 0 & 0 & \cdots & 0 \\ 1 & 0 & 0 & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 0 & 0 & 0 & 0 & \cdots & 0 \end{pmatrix} \right\}.$$

Since the products $4, \cdots, n$ are never purchased and hence do not affect the revenues, the reasoning of Example 3 immediately carries over to this problem. □

Proof of Proposition 2. We show that $R^*_{\text{det}}(\mathcal{U}) \geq \sum_{i \in \mathcal{N}} \lambda_i \min\{r_i : i \in \mathcal{N}\}$ and $R^*_{\text{rand}}(\mathcal{U}) \leq \sum_{i \in \mathcal{N}} \lambda_i \max\{r_i : i \in \mathcal{N}\}$, which together imply the statement. To see that $R^*_{\text{det}}(\mathcal{U}) \geq \sum_{i \in \mathcal{N}} \lambda_i \min\{r_i : i \in \mathcal{N}\}$, consider the deterministic assortment $S = \mathcal{N}$. The worst-case expected revenues of this assortment amount to $\sum_{i \in \mathcal{N}} \lambda_i r_i \geq \sum_{i \in \mathcal{N}} \lambda_i \min\{r_i : i \in \mathcal{N}\}$, and hence $R^*_{\text{det}}(\mathcal{U}) \geq \sum_{i \in \mathcal{N}} \lambda_i \min\{r_i : i \in \mathcal{N}\}$. To see that $R^*_{\text{rand}}(\mathcal{U}) \leq \sum_{i \in \mathcal{N}} \lambda_i \max\{r_i : i \in \mathcal{N}\}$, on the other hand, note that any customer will exercise the no-purchase option with a probability of at least $\lambda_0 = 1 - \sum_{i \in \mathcal{N}} \lambda_i$. Thus, even if a randomized assortment strategy would achieve expected revenues of $\max\{r_i : i \in \mathcal{N}\}$ whenever a customer purchases any product, its worst-case expected revenues are bounded above by $\sum_{i \in \mathcal{N}} \lambda_i \max\{r_i : i \in \mathcal{N}\}$. □

Proof of Corollary 5. The proof of this corollary follows directly from Proposition 2. □

Appendix C: Auxiliary Results and Proofs for Section 6

C.1. The Unconstrained Preference Ranking Model

Proof of Observation 3. We extend the randomization-receptive instance from Example 4 to $n \geq 3$ products as follows. The product-wise revenues are $r_1 = 1$, $r_2 = 2$ and $r_i = 1$ for $i \geq 3$. The preference rankings are

$$1 \to 2 \to 0 \to 3 \to 4 \to \cdots \to n, \quad 2 \to 1 \to 0 \to 3 \to 4 \to \cdots \to n, \quad 1 \to 0 \to 2 \to 3 \to 4 \to \cdots \to n.$$

We assume that the ambiguity set Γ contains the two occurrence probability scenarios $\boldsymbol{\lambda}^1 = (1,0,0)$ and $\boldsymbol{\lambda}^2 = (0, 0.5, 0.5)$. The same computations as in Example 4 show that this problem is randomization-receptive. $\square$

Proof of Proposition 3. The bound for a generic number of products n can be derived using similar arguments as in the proof of Proposition 2; we therefore only prove the bound for $n = 2$ products. To avoid trivial cases, we assume that $r_1 \neq r_2$ and $r_1, r_2 > 0$. Without loss of generality, we further assume that $r_1 > r_2$. In this case, including product 1 into an assortment weakly increases the expected revenues under any occurrence probability scenario, and therefore we can restrict ourselves to deterministic and randomized assortments that include product 1. We then consider the preference rankings

$$\begin{array}{lll} 1: & 1 \to 0 & \text{with probability } \lambda_1 \\ 2: & 2 \to 1 \to 0 & \text{with probability } \lambda_2 \\ 3: & 2 \to 0 & \text{with probability } \lambda_3, \end{array}$$

where $\lambda_1, \lambda_2, \lambda_3 \geq 0$ and $\lambda_1 + \lambda_2 + \lambda_3 = 1$. Note that we do not consider any preference rankings that contain the no-purchase option 0 as highest ranked choice; this is without loss of generality as such a preference ranking affects every (deterministic and randomized) assortment in the same way. We now compute the expected revenues of the deterministic and randomized assortments of interest:

$$\begin{array}{rl} \{1\}: & (\lambda_1 + \lambda_2) \cdot r_1 \;=\; (1 - \lambda_3) \cdot r_1 \\ \{1,2\}: & \lambda_1 \cdot r_1 + (\lambda_2 + \lambda_3) \cdot r_2 \;=\; r_2 + \lambda_1 \cdot (r_1 - r_2) \\ \alpha \cdot \{1\} + \beta \cdot \{1,2\}: & \alpha \cdot (1 - \lambda_3) \cdot r_1 + \beta \cdot [r_2 + \lambda_1 \cdot (r_1 - r_2)] \end{array}$$

Fix an optimal randomization strategy (α^*, β^*) as well as the worst-case distributions (not necessarily unique) $\boldsymbol{\lambda}^1 = (\lambda_1^1, \lambda_2^1, \lambda_3^1) \in \mathcal{U}$, $\boldsymbol{\lambda}^{12} = (\lambda_1^{12}, \lambda_2^{12}, \lambda_3^{12}) \in \mathcal{U}$ and $\boldsymbol{\lambda}^{\mathrm{r}} = (\lambda_1^{\mathrm{r}}, \lambda_2^{\mathrm{r}}, \lambda_3^{\mathrm{r}}) \in \mathcal{U}$ for the deterministic assortments $\{1\}$ and $\{1,2\}$ as well as the randomized assortment $\alpha^* \cdot \{1\} + \beta^* \cdot \{1,2\}$, respectively.

The benefit of randomization can now be computed as

$$R \;=\; \frac{\alpha^* \cdot (1 - \lambda_3^{\mathrm{r}}) \cdot r_1 + \beta^* \cdot [r_2 + \lambda_1^{\mathrm{r}} \cdot (r_1 - r_2)]}{\max\{(1 - \lambda_3^1) \cdot r_1,\; r_2 + \lambda_1^{12} \cdot (r_1 - r_2)\}}.$$

Assume first that $(1 - \lambda_3^1) \cdot r_1 \;\geq\; r_2 + \lambda_1^{12} \cdot (r_1 - r_2)$ in the denominator of R. We then have

$$\begin{aligned} R &= \frac{\alpha^* \cdot (1 - \lambda_3^{\mathrm{r}}) \cdot r_1 + \beta^* \cdot [r_2 + \lambda_1^{\mathrm{r}} \cdot (r_1 - r_2)]}{(1 - \lambda_3^1) \cdot r_1} && \leq \frac{\alpha^* \cdot (1 - \lambda_3^1) \cdot r_1 + \beta^* \cdot [r_2 + \lambda_1^1 \cdot (r_1 - r_2)]}{(1 - \lambda_3^1) \cdot r_1} \\ &= 1 + \frac{-\beta^* \cdot (1 - \lambda_3^1) \cdot r_1 + \beta^* \cdot [r_2 + \lambda_1^1 \cdot (r_1 - r_2)]}{(1 - \lambda_3^1) \cdot r_1} && = 1 + \beta^* \cdot \frac{-(1 - \lambda_3^1) \cdot r_1 + r_2 + \lambda_1^1 \cdot (r_1 - r_2)}{(1 - \lambda_3^1) \cdot r_1} \\ &\leq 1 + \beta^* \cdot \frac{-(1 - \lambda_3^1) \cdot r_1 + r_2 + (1 - \lambda_3^1) \cdot (r_1 - r_2)}{(1 - \lambda_3^1) \cdot r_1} && = 1 + \beta^* \cdot \frac{\lambda_3^1 \cdot r_2}{(1 - \lambda_3^1) \cdot r_1} \;\leq\; 1 + \beta^* \;\leq\; 2, \end{aligned}$$

where the first equality is due to $(1 - \lambda_3^1) \cdot r_1 \;\geq\; r_2 + \lambda_1^{12} \cdot (r_1 - r_2)$, the first inequality holds since the occurrence probabilities $\boldsymbol{\lambda}^{\mathrm{r}}$ minimize the expected revenues of the randomized assortment, the second equality follows from $\alpha^* + \beta^* = 1$, the second inequality holds since $\lambda_1^1, \lambda_2^1, \lambda_3^1 \geq 0$ and $\lambda_1^1 + \lambda_2^1 + \lambda_3^1 = 1$, and the third inequality follows from the fact that $(1 - \lambda_3^1) \cdot r_1 \geq r_2 + \lambda_1^{12} \cdot (r_1 - r_2) \geq r_2$ since $r_1 > r_2$ by assumption.

Assume next that $(1 - \lambda_3^1) \cdot r_1 < r_2 + \lambda_1^{12} \cdot (r_1 - r_2)$ in the denominator of R. We then have

$$\begin{aligned}
R &= \frac{\alpha^* \cdot (1 - \lambda_3^{\mathrm{r}}) \cdot r_1 + \beta^* \cdot [r_2 + \lambda_1^{\mathrm{r}} \cdot (r_1 - r_2)]}{r_2 + \lambda_1^{12} \cdot (r_1 - r_2)} \leq \frac{\alpha^* \cdot (1 - \lambda_3^1) \cdot r_1 + \beta^* \cdot [r_2 + \lambda_1^1 \cdot (r_1 - r_2)]}{r_2 + \lambda_1^{12} \cdot (r_1 - r_2)} \\
&= \frac{\alpha^* \cdot (1 - \lambda_3^1) \cdot r_1 + \beta^* \cdot [r_2 + \lambda_1^1 \cdot (r_1 - r_2)]}{(1 - \lambda_3^1) \cdot r_1} \cdot \frac{(1 - \lambda_3^1) \cdot r_1}{r_2 + \lambda_1^{12} \cdot (r_1 - r_2)} \\
&\leq \left(1 + \beta^* \cdot \frac{\lambda_3^1 \cdot r_2}{(1 - \lambda_3^1) \cdot r_1}\right) \cdot \frac{(1 - \lambda_3^1) \cdot r_1}{r_2 + \lambda_1^{12} \cdot (r_1 - r_2)} = \frac{(1 - \lambda_3^1) \cdot r_1}{r_2 + \lambda_1^{12} \cdot (r_1 - r_2)} + \beta^* \cdot \frac{\lambda_3^1 \cdot r_2}{r_2 + \lambda_1^{12} \cdot (r_1 - r_2)} \leq 2,
\end{aligned}$$

where the calculations follow the same ideas as in the previous case. Note that our derivations assume that $(1 - \lambda_3^1) \cdot r_1 \neq 0$. Indeed, if $(1 - \lambda_3^1) \cdot r_1 = 0$, then $\lambda_3^1 = 1$ since $r_1 > 0$ by assumption. In that case, however, the problem is readily verified to be randomization-proof. □

EXAMPLE 5. Consider the unconstrained robust preference ranking problem with two products, $r_1 = \mathrm{M}$ (with M being a large positive number) and $r_2 = 1$ as well as the preference rankings

$$\begin{array}{lll}
1: & 1 \rightarrow 0 & \text{with probability } \lambda_1 \\
2: & 2 \rightarrow 1 \rightarrow 0 & \text{with probability } \lambda_2 \\
3: & 2 \rightarrow 0 & \text{with probability } \lambda_3.
\end{array}$$

We assume that the ambiguity set $\mathcal{U}$ contains the two occurrence probability scenarios $\boldsymbol{\lambda}^1 = (0, 1, 0)$ and $\boldsymbol{\lambda}^2 = (1/\mathrm{M}, 0, 1 - 1/\mathrm{M})$. Following the same reasoning as in the proof of Proposition 3, we can without loss of generality restrict ourselves to the consideration of the following assortments with the associated worst-case expected revenues:

$$\begin{array}{rl}
\{1\}: & \min\{\mathrm{M}, 1\} = 1 \\
\{1,2\}: & \min\{1, 2 - 1/\mathrm{M}\} = 1 \\
\alpha \cdot \{1\} + \beta \cdot \{1,2\}: & \min\{\alpha \cdot \mathrm{M} + \beta, 1 + \beta \cdot (1 - 1/\mathrm{M})\}
\end{array}$$

Clearly, $R^*_{\mathrm{det}}(\mathcal{V}) = 1$ in this instance. The worst-case expected revenues of the randomized assortment strategy $(\alpha^*, \beta^*) = \left(\frac{1-1/\mathrm{M}}{\mathrm{M}-1/\mathrm{M}}, \frac{\mathrm{M}-1}{\mathrm{M}-1/\mathrm{M}}\right)$, on the other hand, evaluate to

$$\min\left\{\frac{1 - 1/\mathrm{M}}{\mathrm{M} - 1/\mathrm{M}} \cdot \mathrm{M} + \frac{\mathrm{M} - 1}{\mathrm{M} - 1/\mathrm{M}}, 1 + \frac{\mathrm{M} - 1}{\mathrm{M} - 1/\mathrm{M}} \cdot (1 - 1/\mathrm{M})\right\} = \frac{2\mathrm{M} - 2}{\mathrm{M} - 1/\mathrm{M}},$$

and the last expression evidently converges to 2 as $\mathrm{M} \longrightarrow \infty$.

Proof of Corollary 6. The proof this result follows directly from Proposition 3. □

Proof of Theorem 7. Fix any optimal randomization strategy $\boldsymbol{p}^*$ for the robust preference ranking problem. If at most $K + 1$ components of $\boldsymbol{p}^*$ are positive, then there is nothing to prove. Otherwise, note that the worst-case revenues of $\boldsymbol{p}^*$ evaluate to

$$\min_{\boldsymbol{\lambda} \in \Lambda} \sum_{S \subseteq \mathcal{N}} p_S^* \cdot R(S, \boldsymbol{\lambda}) = \min_{\boldsymbol{\lambda} \in \Lambda} \sum_{S \subseteq \mathcal{N}} \sum_{k \in \mathcal{K}} p_S^* \cdot \lambda_k \cdot R_k(S) = \min_{\boldsymbol{\lambda} \in \Lambda} \boldsymbol{\lambda}^\top \left[\sum_{S \subseteq \mathcal{N}} p_S^* \cdot \boldsymbol{R}(S)\right], \tag{14}$$

where $\boldsymbol{R}(S) = (R_1(S), \cdots, R_K(S))^\top \in \mathbb{R}^K$. Here, the summation $\sum_{S \subseteq \mathcal{N}} p_S^* \cdot \boldsymbol{R}(S)$ can be interpreted as a convex combination of the 2^n expected revenue vectors $\boldsymbol{R}(S) \in \mathbb{R}^K$, $S \subseteq \mathcal{N}$. Carathéodory's theorem then implies that there are alternative weights q_S^*, $S \subseteq \mathcal{N}$, such that at most $K+1$ components of $\boldsymbol{q}^*$ are positive and that satisfy

$$\sum_{S \subseteq \mathcal{N}} q_S^* \cdot \boldsymbol{R}(S) \;=\; \sum_{S \subseteq \mathcal{N}} p_S^* \cdot \boldsymbol{R}(S).$$

Substituting the expression on the left-hand side of this identity into equation (14), we conclude that

$$\min_{\boldsymbol{\lambda} \in \Lambda} \sum_{S \subseteq \mathcal{N}} q_S^* \cdot R(S, \boldsymbol{\lambda}) \;=\; \min_{\boldsymbol{\lambda} \in \Lambda} \sum_{S \subseteq \mathcal{N}} p_S^* \cdot R(S, \boldsymbol{\lambda}),$$

which implies that the parsimonious randomization strategy $\boldsymbol{q}^*$ attains the same objective value as the optimal randomization strategy $\boldsymbol{p}^*$. □

C.2. The Cardinality-Constrained Preference Ranking Model

Proof of Theorem 8. Consider the cardinality-constrained preference ranking problem with product-wise revenues $r_1 = \cdots = r_n = 1$ and the n preference rankings

$$1 \to 0 \to \cdots, \qquad 2 \to 0 \to \cdots, \qquad \cdots, \qquad n \to 0 \to \cdots$$

with an associated ambiguity set $\mathcal{U}$ that contains all probability distributions $\boldsymbol{\lambda} \in \mathbb{R}^n$ contained in the n-dimensional probability simplex. Any deterministic assortment $S \subseteq \mathcal{N}$, $|S| \leq C$, generates zero revenues under the Dirac distribution that places all probability mass on a preference ranking containing a product $i \in \mathcal{N} \setminus S$ as highest ranked product. Consider now any randomized assortment strategy that places a strictly positive probability on every assortment $S \subseteq \mathcal{N}$ satisfying $|S| \leq C$ and that places zero probability on all other assortments. Since every product is offered with a strictly positive probability, this strategy generates strictly positive expected revenues under any occurrence probability scenario $\boldsymbol{\lambda} \in \mathcal{U}$. We thus conclude that the worst-case expected revenues of this strategy are also strictly positive. □

Proof of Corollary 7. This corollary follows directly from Theorem 8. □

Proof of Corollary 8. The proof of this corollary follows the same steps as in the proof of Theorem 7, and is hence omitted. □

C.3. Solving the Randomized Preference Ranking Model

Proof of Theorem 9. The proof largely follows similar arguments as the proof of Theorem 2 of Delage and Saif (2018) and is thus omitted for the sake of brevity. □