CIRRELT

CIRRELT-2022-21

Mitigating Choice Model Ambiguity: A General Framework and its Application to Assortment Optimization

Öykü Naz Attila
Sanjay Dominik Jena
Walter Rei

August 2022

Bureau de Montréal

Université de Montréal
C.P. 6128, succ. Centre-Ville
Montréal (Québec) H3C 3J7
Tél : 1-514-343-7575
Télécopie : 1-514-343-7121

Bureau de Québec

Université Laval,
2325, rue de la Terrasse
Pavillon Palasis-Prince, local 2415
Québec (Québec) G1V 0A6
Tél : 1-418-656-2073
Télécopie : 1-418-656-2624

Mitigating Choice Model Ambiguity: A General Framework and its Application to Assortment Optimization

Öykü Naz Attila[1,2,*], Sanjay Dominik Jena[1,3], Walter Rei[1,3]

1. Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT)
2. Department of Mathematical and Industrial Enginering, Polytechnique Montréal
3. Analytics, Operations, and Information Technologies Department, School of Management, Université du Québec à Montréal

Abstract. Discrete choice models have become a popular tool to accurately predict complex choice behavior. Due to a variety of possible error sources, estimated choice models tend to be subject to ambiguity, inducing different optimal decisions of highly varying quality. This study aims at mitigating choice model ambiguity associated with a given set of models in terms of their ability to yield optimal decisions. We propose a framework and a set of performance metrics to assess the reliability of choice models and their induced decisions. The use of this framework is then exemplified in the context of rank-based choice models for assortment optimization. Extensive sets of numerical results suggest that our proposed approaches indeed allow decision-makers to identify choice models that are likely to produce high quality decisions and may therefore boost confidence in using choice models in practice. While robust optimization on the original set of choice models tends to be rather conservative, ranking the choice models according to the proposed metrics and reducing the ambiguity set allows us to improve the expected assortment quality, as well as the downside risk. Given the practical usefulness of robust optimization in this context, we further propose a decomposition algorithm, solving the optimization problem in a fraction of the original time and revealing that only a few among a large set of choice models are determinant in optimal robust solutions.

Keywords: Discrete choice models, ambiguity, rank-based choice models, assortment optimization.

Acknowledgements. The authors are thankful to Compute Canada for providing the computational resources to carry out the experiments. The second author was supported by the Natural Sciences and Engineering Research Council (NSERC) of Canada under grant 2017-05224, as well as by the Fonds de recherche du Québec – Nature et technologies (FRQNT) under grant 70739202. The third author was the Canada Research Chair in Stochastic Optimization of Transport and Logistics Systems. Additionally, he was supported by the Natural Sciences and Engineering Research Council (NSERC) of Canada under grant 2018-05390. The first author was partially supported by the aforementioned grants.

* Corresponding author: attilaoyku93@gmail.com

Dépôt légal – Bibliothèque et Archives nationales du Québec
Bibliothèque et Archives Canada, 2022

1 Introduction

Modern decision-making support has drastically improved throughout the last decades, allowing decision-makers to tackle planning problems of ever increasing scale, complexity and detail. Access to increasing sources, variety and amounts of data has improved the ability of modeling information and to generate knowledge. "Data-driven" decision support has therefore become increasingly popular and successful. At the scientific forefront, approaches aiming at directly deriving decisions from data have made remarkable progress [see, e.g. 4, 12]. While this remains a worthwhile research direction [see, e.g. 10], currently available approaches have limitations when tackling more complex planning problems. Instead, the predict-then-optimize paradigm still remains the standard for the vast majority of planning problems. To this end, a large variety of predictive models are nowadays readily available, which has led to more effective optimization models in practice. However, a main question that remains for decision-makers using this paradigm is: how reliable is the quality and performance of optimal decisions obtained from an optimization model when its predicted parameters are uncertain or the predictive model itself is ambiguous?

Uncertainty generally refers to the lack of knowledge (here, concerning the parameter values), while ambiguity expresses the possibility of entertaining several interpretations (here, concerning the model) at the same time. When parameter uncertainty is prevalent, a more realistic characterization of the underlying uncertainty has enabled the use of different optimization paradigms, ranging from stochastic to robust optimization, by explicitly accounting for the uncertainty within the models. However, even when parameter uncertainty is relatively small, model ambiguity may prevail. Hence, even a good predictive accuracy may not guarantee high quality decisions induced from such a model, particularly in combinatorial optimization problems, where small parameter variations in the predictive model may lead to structurally different planning solutions. It is therefore the natural desire of the decision-maker to reduce such model ambiguity, and to select a model that guarantees, as much as possible, a reasonable performance in practice. In this paper, we aim at providing a remedy to the ambiguity of Choice Models, a specific class of predictive models, popular in the broad field of Management Sciences.

1.1 Choice model ambiguity and assortment optimization

Choice model ambiguity. Choice Models (CM) aim at explaining and predicting the choice of individuals or entities in complex situations. They have been successfully employed in several domains and are increasingly gaining popularity. Prominent examples are transportation, where the transportation mode is to be predicted among several alternatives [see, e.g. 2], and revenue management, where the product choice of customers is to be predicted among the available alternatives (the related literature is reviewed below). Choice models are either completely, or partially, estimated based on historical observations. For several important applications, choice models outperform existing Machine Learning models in terms of predictive accuracy. More importantly, in contrast to the large majority of Machine Learning models, decision-makers may optimize on choice models.

$$\text{Data} \xrightarrow{\text{Estimation}} \text{Choice model} \xrightarrow{\text{Optimization}} \text{Decisions}$$

Figure 1: Classical two-step decision making approach based on choice models.

The classical two-step approach (see Figure 1) involves estimating a choice model from data in the first step. Then, optimal decisions are derived by optimizing on the choice model. Unfortunately, inheriting the characteristics of predict-then-optimize approaches, even when all choice models are estimated under the same conditions and estimation criteria, they may be subject to *ambiguity*. Generally speaking, ambiguity among choice models (or among predictive models in general) may occur due to multiple choice model optima that satisfy the estimation criteria. Next to data *noise*, which may be unavoidable, common reasons include:

1. *Lack of generalization.* Choice models may be trained on data that is insufficient in quantity or in dimensions. The estimated model may therefore not provide accurate predictions in the *general case* under different data.

2. *Over- and underfitting.* The type of choice model used may have a structure that is too flexible or too restrictive for the given application, which may lead to over- and underfitting, respectively.
3. *Different estimation methods.* Choice models can be estimated using a variety of different estimation methods, each of which may lead to a different calibration of the choice model.
4. *Randomness.* Even when ambiguity due to the aforementioned aspects is minimal, the random nature of estimation methods (and their use of optimization solvers) remain a source of ambiguity. Due to random behavior, the estimation method may converge to different optima.

It is also important to highlight that the concept of handling ambiguity in decision-making has been recognized in the literature. For instance, [21] models the value of information to handle ambiguity and [18] studies the value of information under conditional and unconditional ambiguity.

The issue of ambiguity is not only related to the estimated models themselves and their predictive performance, but more importantly to the optimal decisions they entail. As such, different choice models may imply optimal decisions that are drastically different from one another. To the decision-maker, avoiding induced decisions of low quality is therefore of uttermost importance. Unfortunately, the risk of low quality decisions increases in the case of *discrete* choice models, where individuals do not take continuous, but discrete decisions, and when the planning problem is of a combinatorial nature.

This work therefore aims at addressing the issue of choice model ambiguity, and how it can be mitigated in order to improve the performance of the induced decisions. More specifically, we are interested in identifying techniques to discriminate the quality of choice models, as well as making ideal decisions under choice model ambiguity, instead of basing our decisions on a nominal choice model. In particular, we will propose techniques that allow decision-makers to induce decisions that perform well due to their *stability*, i.e., optimal decisions that stabilize the information from multiple choice models. As a case study, we will use one of the most prominent applications using discrete choice models, namely Assortment Optimization.

Assortment optimization. A fundamental aspect of maximizing profits in most businesses involves the decisions regarding the range of products that should be offered to customers. In general, the set of products that a business offers to its customers can be defined as an *assortment*. Determining a reliable assortment tends to be a critical task, since it naturally impacts the total sales and profitability. Practical examples of this include retailers, as well as online shopping platforms. Given a gamut of products $\mathcal{N}$, each with an estimated revenue $r_i, i \in \mathcal{N}$, the *Assortment Optimization (AO)* problem seeks to identify a subset of products $\mathcal{X} \subseteq \mathcal{N}$ such that the expected revenue generated by these assortment decisions is maximal, while taking into account customer buying behavior, product revenue and other context-specific constraints (for instance, assortment size).

Within the predict-then-optimize paradigm in the context of Assortment Optimization, discrete choice models have been shown to be particularly effective [see, e.g. 13, 23, 14]. As we will review in the next section, different types of choice models have been proposed in order to represent the customer buying behaviour. For each of those choice models, different estimation methods have become available, typically estimating a model from historical sales and assortment data. Each of those aspects are potential sources to provide model ambiguity, even if each model has been estimated on the same data.

1.2 Related literature

We now discuss the literature related to our work. We first review the most popular types of discrete choice models, particularly those used in the context of assortment optimization, along with their estimation methods. For a more detailed review of choice models and their estimations, we refer the interested reader to [16], [23] and [3], as well as the references therein. We then discuss how choice model ambiguity has been handled in the literature.

Discrete choice models and their estimation. Discrete choice models have been successfully used in several domains and undergone an impressive evolution over time. They can roughly be divided into two families: parametric random utility maximization (RUM) models and non-parametric exponential models. Parametric RUM models use a constant number of parameters. Prominent classes of models in this family include the Multinomial Logit (MNL) model, which assigns different utility values to the available choice alternatives for a single customer type that then selects the product with highest utility. The MNL model assumes independence between the available choices [known as *Independence of Irrelevant Alternatives*, 1], which falls short of explaining more complex choice behaviour such as product substitution[1]. A multitude of extensions of the MNL model has since been proposed. For instance, the Nested Logit model [2] provides the possibility of representing choice substitution and the Mixed-MNL model includes several classes of customers. Given that parametric models still rely on a certain market knowledge, their estimation is prone to risks of under- and overfitting, and thus choice model ambiguity.

Non-parametric choice models are characterized by a flexible structure, which may adapt to the level of estimation precision of the model. Rank-based (RB) models [13] are undoubtedly among the most popular in this class. Here, preference sequences (i.e., ordered lists of the choice alternatives) represent the choice behaviour of a customer type. A customer then selects the product that ranks highest in her list and is also available in the assortment. The choice model itself is composed by a set of different customer types, along with a probability distribution over the different types. A major computational challenge stems from the number of possible preference sequences, which is theoretically exponential large in the number of choices alternatives. As a result, a variety of different estimation methods have been proposed to overcome this challenge. Mainly, [23] introduced the market discovery algorithm, a column-generation based method that iteratively generates preference sequences and therefore gradually improves the model fit. While the market discovery algorithm originally maximizes the log-likelihood probability, different variants have been proposed to instead minimize the L_1 (least absolute deviation) estimation error [e.g. 5, 14]. Those models remain highly attractive for practitioners, given that they are acknowledged for relatively high predictive accuracy [see, e.g. 3], their estimation is computationally tractable [see, e.g. 24, 14] and they can be efficiently optimized on [see, e.g. 7]. However, both their flexibility and their variety of possible estimation mechanisms are a major source of choice model ambiguity. Finally, a more recent branch of research focuses on capturing non-rational choice [see, e.g. 9, 15]. Given that such choice models require an even more flexible structure, they are prone to even more ambiguity.

Choice model uncertainty and ambiguity. The issue of statistical and numerical errors within the estimation procedure leading to ambiguity in the resulting choice models has been generally recognized in the literature. Most works have addressed this issue by explicitly modeling the uncertainty around the parametric structure of the choice models. Particular advancements have been made in the area of Assortment Optimization, where a significant part of the literature focuses on deriving robust assortments that perform best under worst-case choice model parameters. In that sense, [20] consider a robust assortment optimization problem under the well-known Multinomial Logit model, in which the customer utility vector is uncertain. With capacity constraints, the authors solve the problem under the assumption of a box uncertainty. As a direct extension, [19] consider that the entire utility vector of the customer may be uncertain, which leads to an uncertainty set of different customer types. This problem maximizes the expected assortment revenue and is known as the Mixed-MNL model. While this choice model explicitly represents several customer types, ambiguity concerning the entire choice model itself is not considered. [11] propose a more general *min-max* framework for robust assortment optimization, where the uncertainty on choice probabilities is encoded in a system of linear equations. The authors show how to model and solve problems under the Markov-chain and the MNL models under specific conditions on the uncertainty sets. While the framework itself is quite general, it still has to be explored under which conditions more general ambiguity sets and choice models, such as RB models, can be solved.

For RB choice models, [22] recently provided a characterization of an optimal assortment that maximizes the expected revenue under any worst-case choice model that is coherent with historical transaction data [see 13]. The author elaborates on the computational complexity and shows that, under such rather conservative choice model uncertainty, the problem can be solved in polynomial time. For similar choice model uncertainty, [17] maximize the expected assortment revenue under the worst-case order of the preference list in a problem setting where the cumulative effect of unavailable products over time is taken into account. An interesting recent direction in robust assortment optimization is taken by [26]. Instead of returning a single assortment that performs best under worst-case choice model uncertainty, the authors optimize the probability distribution over all feasible assortments such that the expected revenue is maximized under worst-case uncertainty. Computational experiments are provided under the assumption of norm-ball uncertainty for the constrained MNL model, the general Markov Chain model and a RB model.

Two works particularly related to our work can be found in the context of Product Line Design, a

[1]Product substitution refers to the fact that a customer may choose an alternative product, if her preferred product is not available. This also implies that adding a certain product to the offer set may reduce the sales of other products.

combinatorial optimization problem in revenue management that is similar to Assortment Optimization both in its objective and its structure. [25] consider a simpler RB choice model variant, in which all customers have the same probability. The underlying uncertainty ultimately results in a discrete set of customer scenarios. Instead of maximizing revenue under the worst-case scenario, the authors minimize the ex-post regret, i.e., the maximum opportunity loss of selecting a product design different then the one that would be optimal for the specific customer scenario. Also in the context of robust Product Line Design, [6] consider a variety of different uncertainty sets and show that approaches based on a nominal choice model may severely under-perform the robust approaches in the worst case.

Although the studies above have tackled the uncertainty around the parametric structure of choice models, few have addressed the ambiguity among choice models *that have already been estimated*. Indeed, certain parametric uncertainty sets (such as boxes or norm-balls on utility values) may result in overly conservative choice model ambiguity sets. Here, [6] are the exception, given that the authors also propose an ambiguity set composed of a discrete set of choice models. The authors introduce a robust model that maximizes the estimated revenue under the worst-case choice model. Given that the problem remains difficult to solve, the authors also provide a heuristic approach that can tackle the nominal and robust variants of the problem. Representing choice model ambiguity in the form of a discrete set of choice models holds several advantages. In particular, it allows for more flexibility since choice models can easily be included or omitted. This easily enables the decision-maker to adjust the level of conservativeness with regards to the assortment decisions. To the best of our knowledge, our paper is the first attempt to explicitly reduce ambiguity before the actual optimization model is employed. Finally, we note that only few works have considered parameter uncertainty within RB choice models, which subsume RUM models and are the focus of our case study.

1.3 Objective and Contributions

The principal objective of this paper is to propose techniques to mitigate choice model ambiguity. To this end, its main contributions can be summarized as follows:

1. *General evaluation framework to gauge reliability:* We provide a general framework to gauge the reliability of the various choice models. This framework, coupled with a specific performance metric, can hence be used to efficiently reduce the ambiguity set prior to optimization. The use of our evaluation approach thus directly leads to an overall less conservative optimization approach when facing the considered uncertainty.

2. *Performance metrics and classification scheme:* We propose a classification scheme for performance metrics, as well as several concrete instances of such general metrics, which can be used within any predict-then-optimize framework based on choice models to "score" the quality of the decisions induced by a given choice model. These metrics therefore go beyond predictive accuracy of choice models, and directly consider their impact on the quality of the optimal decisions, contributing to the research gap between *data* and *decisions* [see, e.g. 10]. We also give an example of an application-specific performance metric. Both the evaluation framework and the performance metrics are then interpreted in the context of rank-based choice models for Assortment Optimization.

3. *Robust Optimization decomposition method.* We consider the *robust optimization problem* in the context of assortment optimization as introduced by [6] in order to hedge the risk of using ill-suited choice models by providing worst-case protection. We explicitly formulate the problem for general rank-based choice models as a Mixed-integer programming (MIP) model, which can be solved by general-purpose MIP solvers. We provide an efficient decomposition approach to solve this problem to optimality, which reveals that only a small percentage of the choice models are critical in determining robust optimal solutions. Consequently, this reduces the time required to solve the problem by more than 85% on average and allows us to efficiently solve large-scale problems in a matter of minutes.

4. *Extensive numerical experiments.* Finally, we provide extensive computational experiments based on a large set of RB problem instances: (i) we explicitly analyze and quantify the ambiguity stemming from the most basic source of ambiguity: numerical randomness when executing the code of the estimation method; even under such an inevitable error source, the ambiguity can be quite high, leading to a performance variation of up to 60% of the estimated assortment revenue among the models; capturing

an average of up to 93.3% of the optimal revenue for the considered ground truth; it is here worthwhile to mention that we also found that an MMNL generating ground truth, popular in the literature, results in much lower ambiguity; (ii) we compare the expected revenues among the assortments induced by the individual choice models, the choice models recommended by the various metrics and those proposed by the robust optimization model. Our results indicate that the proposed metrics successfully reduce ambiguity and retain choice models of high quality: selecting a single choice model according to any of the decision-driven metrics allows to capture an additional 4%-13% of the optimal revenue on average when compared to the mean revenue performance of assortments optimized on the individual choice models from the ambiguity set. One of the proposed metrics captures an additional 8.5%-12.9% of the optimal revenue among the different problem instance classes. The robust optimization model also successfully hedges against ambiguity. It captures about 2% of additional revenue when compared to the revenue induced by the choice model stemming from the best performing metric. Finally, (iii) we quantify the revenue improvement when using the robust optimization model based on a subset of the choice models recommended by the various metrics, enabling certain models to capture an additional 2%-3% of the optimal revenue when compared to the robust optimization over all available choice models, which corresponds to 95-96% of the optimal revenue on average.

Our developments enable decision-makers to benefit from the advantages of optimizing on a discrete set of choice models in contrast to optimizing on a parametric uncertainty set: (i) it gives the decision-maker more control over the ambiguity set; (ii) it allows the decision-maker to impose other criteria than just fitting historical data on considered choice models; (iii) any ambiguity set can be sufficiently well represented by a discrete set by using an appropriate sample of choice models.

1.4 Organization of the paper

Section 2 introduces the general evaluation framework and four novel metrics that distinguish choice models based on their probabilistic structure and their induced decisions. These metrics are then interpreted in the context of rank-based choice models for assortment optimization. Section 3 introduces the robust assortment optimization problem. To solve the robust problem, we propose a min-max decomposition framework. Section 4 provides extensive computational experiments that thoroughly analyze the quality of the resulting assortments obtained from the deterministic and robust models, as well as from the choice models ranked best among the four metrics. Finally, we conclude in Section 5 and propose future research directions.

2 Framework to Identify High-quality Choice Models

In this section, we first introduce the general framework to quantify the quality of a given choice model. We here assume that the decision-maker has available a finite set of different choice models, each considered to be well fitted to the available training data. This is then followed by a formal introduction of novel metrics that are used to distinguish choice models by their ability to induce reliable decisions.

Notation. We generally use bold faced characters such as $\boldsymbol{x} \in \mathbb{R}^N$ and $\boldsymbol{A} \in \mathbb{R}^{M \times N}$ to represent vectors and matrices. We use x_i to denote the i-th element of vector $\boldsymbol{x}$.

2.1 General evaluation framework and performance measures

We will first introduce the general framework, independent from a specific application. We then interpret the framework in the context of assortment optimization.

2.1.1 Application independent framework.

Let us assume that the decision-maker has access to a discrete and finite set of D choice models, referred to as the *ambiguity set*, given as:

$$\mathcal{D} := \{C_d : \forall d \in \{1, \ldots, D\}\}. \tag{1}$$

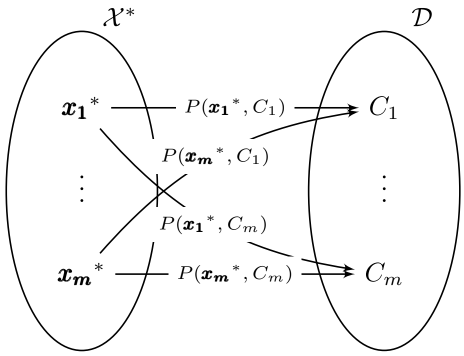

Figure 2: Relative performance measure between assortment decisions and choice models.

We assume that each choice model C_d in set $\mathcal{D}$ satisfies the criteria of sufficient estimation defined by the decision-maker (typically a sufficiently well fit to the training data). The decision-maker therefore faces a problem of choice model ambiguity, and may not know which model to rely on. As outlined in Section 1, this ambiguity may have several sources, such as the use of different estimation methods or simply by obtaining different models at each execution of the same estimation method (also see Section 4). Typically, the available choice models are prone to structural differences. Note, again, that these structurally different models may provide the same accuracy of prediction error on several data sets, or even provide the same choice predictions for all choice offers on certain data sets. However, given the structural difference between the models, this will not hold true on all possible data sets, and models may therefore induce different decisions. In practice, some of these decisions can lead to poor performance, caused by the ambiguity around the quality of the choice models in set $\mathcal{D}$. Let $\boldsymbol{x_d}^*$ be the vector of optimal decisions induced from choice model $C_d \in \mathcal{D}$. We introduce the set of optimal decisions induced by a choice model $C_d \in \mathcal{D}$ as:

$$\mathcal{X}^* := \{\boldsymbol{x_d}^* : \forall d \in \{1, \dots, D\}\}. \quad (2)$$

A decision-driven approach for interpreting the level of reliability of the vectors of optimal decisions in $\mathcal{X}^*$ is by exploiting the probabilistic structure of the predictive models in $\mathcal{D}$. In other words, we are interested in gauging the reliability of the optimal decisions in set $\mathcal{X}^*$ by computing a relative performance measure of a given optimal decision $\boldsymbol{x_d}^* \in \mathcal{X}^*$ under the probability mass function of choice models $C_m \in \mathcal{D},\ d \neq m$. To this end, we will first define a relative performance measure.

Definition 1. ***Relative performance measure.*** *A function $P(\boldsymbol{x_d}, C_m) \to \mathbb{R}$ is called a relative performance measure, if it defines a criteria to estimate the quality of decisions $\boldsymbol{x_d}$ under the choice behaviour dictated by choice model C_m.*

In the context of assortment optimization, such a relative performance measure may be the estimated revenue of assortment decisions $\boldsymbol{x_d}$ under a different choice model C_m. Such a measure first estimates the sales probabilities under the specified assortment and then computes the total revenue based on the revenue of each product.

Figure 2 visualizes the concept of using the relative performance measure $P(\cdot, \cdot)$, where each set of decisions $\boldsymbol{x_d}^*$ induced from choice model C_d can be evaluated under different choice models $C_1, \dots, C_m$, yielding an estimated performance value $P(\boldsymbol{x_d}^*, C_m)$. Loosely speaking, the value $P(\boldsymbol{x_d}^*, C_m)$ quantifies the performance of decisions $\boldsymbol{x_d}^*$ through the lens of C_m. Before we dive into several possibilities prescribing how to exploit this relative performance measure using different performance metrics, we will first apply this general evaluation framework in the context of assortment optimization.

2.1.2 Rank-based choice models for assortment optimization.

While our proposed evaluation framework and performance metrics are not limited to a specific application and type of choice model, we will use assortment optimization based on a rank-based choice model [13] as a case study. A rank-based choice model $C_d := (\boldsymbol{\sigma_d}, \boldsymbol{\lambda_d}),\ d \in \{1, \dots, D\}$, consists of a set $\boldsymbol{\sigma_d}$ of preference sequences, along with a probability mass function $\boldsymbol{\lambda_d}$ over these preference sequences. While the set $\boldsymbol{\sigma_d}$ of preference sequences can theoretically be exponentially large, it is a common practice to only include the set of preference sequences that have a non-zero probability. Specifically, we also define for choice model C_d an index set $\mathcal{K}_d = \{1, \ldots, K_d\}$ of preference sequences, where σ_d^k is the preference sequence of customer type $k \in \mathcal{K}_d$ with $\lambda_d^k > 0$. Given the structural differences among the choice models $C_d \in \mathcal{D}$, the size of $\boldsymbol{\sigma_d}$ and $\boldsymbol{\lambda_d}$ may also be different for each choice model in $\mathcal{D}$.

We now interpret the relative performance measure $P(\cdot, \cdot)$ in the context of assortment optimization, where the performance is typically measured as the estimated revenue of a given assortment, further referred to as the revenue estimate $R(\cdot, \cdot)$. To this end, consider the following binary parameter to represent the choice of customer type $k \in \{1, \ldots, K_d\}$ under a given choice model C_d as:

$$y_i^k(\boldsymbol{x_d}^*, C_d) = \begin{cases} 1 & \text{if product } i \text{ is chosen under assortment decisions } \boldsymbol{x_d}^* \text{ and preference sequence } \sigma_d^k, \\ 0 & \text{otherwise.} \end{cases}$$

Note that, in the specific context of assortment optimization, we define $\boldsymbol{x_d}^* \subseteq \{0,1\}^N$ as the vector of optimal assortment decisions induced from choice model $C_d \in \mathcal{D}$. In particular, N is the number of existing products in the product universe $\mathcal{N}$ and the i^{th} element of $\boldsymbol{x_d}^*$ is 1 if and only if product i is part of the assortment.

The total expected revenue of the assortment defined by decisions $\boldsymbol{x_d}^*$ under the probability mass function of choice model C_d is computed as:

$$R(\boldsymbol{x_d}^*, C_d) = \sum_{k \in \mathcal{K}_d} \sum_{i \in \mathcal{N}} r_i \lambda_d^k y_i^k(\boldsymbol{x_d}^*, C_d), \tag{3}$$

where r_i is the revenue of product i, $i \in \mathcal{N}$. By definition, $R(\boldsymbol{x_d}^*, C_d)$ is equivalent to the optimal objective function value that is obtained from solving the assortment optimization problem under choice model C_d. The relative performance measure of optimal assortment decisions $\boldsymbol{x_d}^*$, here defined as the estimated revenue $R(\cdot, \cdot)$, can also be computed under a different choice model C_m, where $m \neq d$:

$$R(\boldsymbol{x_d}^*, C_m) = \sum_{k \in \mathcal{K}_d} \sum_{i \in \mathcal{N}} r_i \lambda_m^k y_i^k(\boldsymbol{x_d}^*, C_m). \tag{4}$$

The difference between the two revenue estimates given in Equations (3) and (4) is based on the difference of probabilistic structures of C_d and C_m (i.e., $(\boldsymbol{\sigma_d}, \boldsymbol{\lambda_d})$ and $(\boldsymbol{\sigma_m}, \boldsymbol{\lambda_m})$, respectively). While Equation (3) defines the nominal revenue obtained under choice model C_d and its optimal assortment decision $\boldsymbol{x_d}^*$, Equation (4) indicates the total expected revenue of this assortment under the probability mass function of C_m. Loosely speaking, we are asking choice model C_m to estimate the revenue we may expect when implementing assortment $\boldsymbol{x_d}^*$. Since $C_d := (\boldsymbol{\sigma_d}, \boldsymbol{\lambda_d})$ and $C_m := (\boldsymbol{\sigma_m}, \boldsymbol{\lambda_m})$ (where $d \neq m$) represent two unique probability mass functions, $R(\boldsymbol{x_d}^*, C_m)$ is likely to differ from $R(\boldsymbol{x_d}^*, C_d)$. In the next section, we will derive different metrics that gauge the level of reliability of a choice model and its resulting assortment decisions, using Equation (4) to quantify, for example, the opportunity cost of using different choice models under the given decisions.

2.2 Performance metrics

2.2.1 Definition and classification of performance metrics.

With the general objective in mind to interpret *the level of reliability* of the available choice models in $\mathcal{D}$, the relative performance measure $P(\cdot, \cdot)$ defined above evaluates the decisions of one choice model through the lens of another choice model. In order to evaluate the quality of a choice model relative to a set of different choice models, we first formally define a relative performance metric.

Definition 2. ***Relative performance metric.*** *A function $M(C_d, \mathcal{D}) \rightarrow \mathbb{R}$ is called a relative performance metric, if it defines a criteria to estimate the quality of a choice model C_d relative to a set of choice models $\mathcal{D}$.*

We may classify relative performance metrics into three different criteria: which type of information they process (*structural properties*), which choice models they need to consult (*consensus properties*) and the number of choice models whose information is aggregated into the final performance score (*performance source properties*).

Structural properties are related to the structure of the estimated models and/or the optimal decisions induced by these models. According to their structural properties, metrics can be classified as:

- *Model-driven metrics.* Model-driven metrics analyze characteristics of the structure of the choice model itself. Such metrics may perform well in the case of a strong correlation between the structural properties of the choice model and the quality of its induced decisions. By definition, they do not require to actually identify the actual optimal decisions.

- *Decision-driven metrics.* Decision-driven metrics measure the quality of a choice model C_d based on the optimal decisions $\boldsymbol{x_d}^*$ that are induced by the same choice model. This may imply the analysis of the structure of induced decisions, or the evaluation of these decisions through the lens of other choice models. Such metrics may be appropriate when the correlation between choice model structure and the quality of the decisions is weak. Clearly, computing such metrics may require additional computational effort to identify the optimal set of decisions for the given choice model.

Consensus properties refer to the number of choice models the metric needs to access in order to compute the performance score for a specific choice model $C_d \in \mathcal{D}$. We classify the most common cases as:

- *Only the choice model itself.* If the performance score is solely based on information from the choice model itself, as well as on information derived from it (e.g. the induced decisions), access to other choice models is not required.

- *Some other, or all choice models.* On the other extreme, the performance metric may require to access all other choice models as well, e.g. to compare to their structure or to conclude a general consensus.

Finally, performance source properties relate to the set of choice models whose information explicitly impacts the final performance score for a specific choice model C_d. Again, the common cases are:

- *Performance scores stemming from a single choice model.* The final performance score is exclusively derived from one single choice model (typically C_d itself). This does not exclude the possibility that the metric requires access to several choice models in order to select the one whose information will be used to generate the score.

- *Performance scores aggregating information from several choice models.* The final performance score is an aggregation of information from several choice models, for example an average. Naturally, this implies that the metric has accessed several choice models.

It is natural that performance metrics with certain combinations of properties are more computationally expensive than others. Decision-driven metrics require the optimal decisions, which may involve the solution of an optimization problem. Similarly, metrics that require access to several choice models or even incorporate information from different choice models may require more computing time than metrics solely based on the single same choice model. However, the additional (computational) effort may be worthwhile, if such performance scores allow decision-makers to reduce ambiguity and identify assortments that are more reliable for their purpose.

2.2.2 Instances of performance metrics.

It becomes immediate that performance metrics may extend the concept of *model selection* tools for predictive models in general. As such, one may define several performance metrics based on criteria typically used to identify a promising model. For example, it is quite common to use the training error on the training set, but most models may be equally well estimated (as it may be our case, causing model ambiguity). Therefore, other criteria such as the *predictive accuracy on a specific test set* (or, alternatively, the AIC/BIC) are typically used to discriminate the quality of the various available models. The performance, however, depends on the definition of the test set, and may lead to different conclusions when the test set changes.

Further, as already pointed out, such criteria may not sufficiently distinguish the quality of the different choice models, given that they ignore the impact on the induced decisions, which often is a major source of ambiguity.

We will now present four performance metrics: three decision-driven metrics that can be used in the general application-independent framework and include information from other choice models (and therefore aim at consensus) in order to hedge against a possible misrepresentation of the real-world choice behaviour; and one model-driven metric, tailored to the case of rank-based choice models. While the first three metrics can be written using the relative performance measure $P(\cdot,\cdot)$, for the sake of simplicity, we will interpret all four metrics in the terminology of assortment optimization and directly use the relative revenue estimate $R(\cdot,\cdot)$.

Minimum revenue metric. As a classical performance measure used in robust optimization, [6] propose to gauge the worst-case performance of an assortment $\boldsymbol{x}^*$ by estimating its revenue through all other choice models within the ambiguity set. In the same spirit, we here propose to evaluate the assortments stemming from each of the choice models in the same way. Specifically, we ask the other choice models in $C_m \in \mathcal{D}' \subseteq \mathcal{D}$ for their estimate on decisions $\boldsymbol{x_d}^*$, using the relative performance measure $R(\boldsymbol{x_d}^*, C_m)$. This leads to the following performance metric, which we will refer to as the *minimum revenue* metric:

$$M_{MinRev}(C_d, \mathcal{D}') = \min\{R(\boldsymbol{x_d}^*, C_m) : m \in \mathcal{D}'\}. \tag{5}$$

The performance score attributed to C_d is therefore equivalent to the smallest revenue estimate for $\boldsymbol{x_d}^*$ from any of the other choice models, and therefore tends to be a conservative (or *robust*) revenue estimate. Since our interest is to find the choice model that induces the best assortment decision with respect to the remaining choice models, we set $\mathcal{D}' = \mathcal{D} \setminus \{C_d\}$, where higher values of $M_{MinRev}(C_d, \mathcal{D}')$ help us distinguish choice models that perform better under a worst-case scenario setting.

Average revenue metric. While the previous metric aims at providing a conservative estimate of the revenue, we may also average over the different revenue estimates for $\boldsymbol{x_d}^*$ given by all other choice models, leading to the *average revenue* metric:

$$M_{AvRev}(C_d, \mathcal{D}') = \frac{\sum_{m \in \mathcal{D}'} R(\boldsymbol{x_d}^*, C_m)}{|\mathcal{D}'|}. \tag{6}$$

Again, higher values of $M_{AvRev}(C_d, \mathcal{D}')$ indicate an overall better performance for the assortment decisions $\boldsymbol{x_d}^*$ induced by choice model C_d.

Average absolute opportunity cost. The previous two metrics indicate absolute performance, i.e., they refer to an estimate of the estimated revenue for a given set of decisions $\boldsymbol{x_d}^*$. We now look at a relative measure, namely the opportunity loss, which quantifies the difference between the revenue estimated by C_d and the one estimated by another choice model $C_m \in \mathcal{D}'$. This metric, referred to as the *average absolute opportunity cost* metric, is defined as:

$$M_{AvOpp}(C_d, \mathcal{D}') = \frac{\sum_{C_m \in \mathcal{D}'} \frac{|R(\boldsymbol{x_d}^*, C_m) - R(\boldsymbol{x_d}^*, C_d)|}{R(\boldsymbol{x_d}^*, C_d)}}{|\mathcal{D}'|}. \tag{7}$$

Here, $M_{AvOpp}(C_d, \mathcal{D}')$ represents the average absolute relative distance based on the opportunity cost that is associated with the assortment decision induced by choice model C_d. Note that, while the previous two metrics require that the revenue estimates for the assortment decisions are high in order to score well compared to the other assortment decisions, the average absolute opportunity cost metric does not rely on the estimated revenue amount. Instead, it rather identifies how close the revenue estimate is to the estimates through the lens of other choice models. As such, higher values of $M_{AvOpp}(C_d, \mathcal{D}')$ indicate that the revenue associated with the assortment decision induced by choice model C_d deviates strongly under other choice models $C_m \in \mathcal{D}'$. Therefore, we are interested in identifying the choice models for which the value for $M_{AvOpp}(C_d, \mathcal{D}')$ is lowest, which indicates that the total expected revenue tends to be consistent across all other choice models in $\mathcal{D}'$.

		Consensus		Performance source	
Structure	Metric	Single model	Several models	Single model	Several models
	M_{MinRev}		✓	✓	
	M_{AvRev}		✓		✓
	M_{AvOpp}		✓		✓
Model driven	$M_{WeigLen}$	✓		✓	

Table 1: Classification of the metrics according to their properties

While the metric above measures the potential loss associated to the revenue estimation of $\boldsymbol{x_d}^*$ by different choice models, one may also consider the opportunity loss defined as the revenue loss estimated by C_d when a different assortment $\boldsymbol{x_m}^*, m \neq d$ is selected: $\frac{R(\boldsymbol{x_d}^*, C_d) - R(\boldsymbol{x_m}^*, C_d)}{R(\boldsymbol{x_d}^*, C_d)}$. Preliminary numerical experiments carried out on the various metrics have shown that the score provided by a metric using this type of opportunity loss has a slightly lower correlation ($R^2 = 0.17$) with the GT revenue of its corresponding assortment than the one provided by M_{AvOpp} ($R^2 = 0.31$). For the remainder of our study, we therefore focus on M_{AvOpp}.

On a further note, this alternative definition of opportunity loss is also similar to the one used in the objective function of the robust optimization model presented by [25], who maximize the minimum of this opportunity loss. Using the minimum (instead of the average) opportunity loss is yet another possibility of defining a performance metric, from which, for the sake of concision, we refrain from in this work.

Weighted preference sequence length. Finally, we also define the model-driven *weighted preference sequence length* metric, which provides insights on the weighted length of the preference sequences of a given choice model, represented as:

$$M_{WeigLen}(C_d) = \sum_{k \in \mathcal{K}_d} \lambda_d^k l_d^k, \tag{8}$$

where λ_d^k represents the probability associated to preference sequence $k \in \mathcal{K}_d$ in C_d and l_d^k indicates the length of preference sequence k. $M_{WeigLen}(C_d)$ therefore represents the weighted preference sequence length for a given choice model C_d. The intuition behind this metric lies in the tendancy that preference sequences with smaller numbers of ranked products may be less prone to overfitting and generalize better. While this may be advantageous for the predictive accuracy, it is by no means an indication of its capacity to produce high quality solutions. Given that we are interested in selecting a choice model with higher predictive accuracy, we here prioritize lower values of $M_{WeigLen}(C_d)$.

Table 1 locates the four proposed metrics, as well as the worst-case objective metric within the classification scheme proposed in Section 2.2.1. In order to compute the performance score for choice model C_d, metric $M_{WeigLen}$ only requires access to the same choice model and, consequently, also computes the final performance score based on information from this single choice model. In contrast, metrics M_{MinRev}, M_{AvRev} and M_{AvOpp} make use of all available choice models and aim at achieving some sort of consensus among them.

Throughout this paper, we will make use of the proposed performance metrics in two ways: first, we will identify the model that performs best according to the computed performance score for the case where the decision-maker has to decide which single choice model to use; second, we will exclude choice models which underperform according to their performance score in order to reduce the set of choice models, and hence the associated ambiguity. In the next section, we will discuss the optimization model for a nominal choice model, and propose optimization models that are capable of operating on a set of choice models.

3 Robust Optimizing on Ambiguity Sets

Although the previously mentioned metrics (see Section 2.2) provide a tractable procedure to identify reliable choice models, the decision-maker may still prefer to use an ambiguity set $\mathcal{D}$ to hedge the overall risk of using a single ill-suited choice models. We therefore introduce a robust approach that identifies the assortment that performs best under the worst-performing choice model. As such, the robust model allows the decision-maker to *hedge against choice model ambiguity*

Each choice model $C_d \in \mathcal{D}$, has a unique set of preference sequences and respective probabilities, such that $C_d := (\boldsymbol{\sigma_d}, \boldsymbol{\lambda_d})$. For a given choice model C_d, we formally define the complete set of preference sequences and their relative probabilities as $\sigma_d := \{\sigma_d^1, \sigma_d^2, ..., \sigma_d^{K_d}\}$ and $\lambda_d := \{\lambda_d^1, \lambda_d^2, ..., \lambda_d^{K_d}\}$, respectively, where K_d is the total number of preference sequences present in choice model C_d.

While the model produces a single assortment decision $\boldsymbol{x}$ [similar to the deterministic problem variant proposed by 7], decision variables $\boldsymbol{y}$ are defined such that they capture the preferences among different preference sequences that belong to different choice models in $\mathcal{D}$. Specifically, binary variable y_{id}^k takes value 1, if customer type k chooses product i under preference sequences $\boldsymbol{\sigma_d}$ and assortment decisions $\boldsymbol{x}$, and 0 otherwise. Based on these decision, the assortment revenue estimated by a choice model $d \in \mathcal{D}$ is computed as:

$$\mathcal{R}(C_d) = \sum_{k \in \mathcal{K}_d} \sum_{i \in \mathcal{N}} r_i \lambda_d^k y_{id}^k. \tag{9}$$

Under this setting, the robust counterpart proposes an assortment that performs best under the worst-case choice model within the ambiguity set. In other words, the robust assortment optimization problem AO_{rob} aims at finding an assortment that has highest worst-case revenue among the choice models in $\mathcal{D}$:

$$\max_{\boldsymbol{x},\boldsymbol{y}} \min_{C_d \in \mathcal{D}} \mathcal{R}(C_d) \tag{AO_{rob}}$$

$$\text{s.t.} \sum_{i \in \mathcal{N} \cup \{0\}} y_{id}^k = 1 \qquad d \in \mathcal{D}, k \in \mathcal{K}_d \tag{10}$$

$$y_{id}^k \leq x_i \qquad d \in \mathcal{D},\ k \in \mathcal{K}_d,\ i \in \mathcal{N} \tag{11}$$

$$\sum_{j:\sigma_d^k(j) > \sigma_d^k(i)} y_{jd}^k \leq 1 - x_i \qquad d \in \mathcal{D},\ k \in \mathcal{K}_d,\ i \in \mathcal{N} \tag{12}$$

$$\sum_{j:\sigma_d^k(j) > \sigma_d^k(0)} y_{jd}^k = 0 \qquad d \in \mathcal{D},\ k \in \mathcal{K}_d \tag{13}$$

$$x_i \in \{0, 1\} \qquad i \in \mathcal{N} \tag{14}$$

$$y_{id}^k \geq 0 \qquad d \in \mathcal{D},\ k \in \mathcal{K}_d,\ i \in \mathcal{N} \cup \{0\}, \tag{15}$$

where $\mathcal{R}(C_d)$ is the assortment revenue as defined in (9). Given that robust optimization prepares for the worst-case outcome, the optimal decisions induces may tend to be rather conservative, especially when the worst-case scenario is statistically unlikely to occur. In the context of ambiguity, however, we can be more optimistic about the utility of such an approach, since all choice models in $\mathcal{D}$ are assumed to be a realistic representation of the customer buying behaviour.

By considering statistical outliers (i.e., choice models that are unlikely to be good representatives for the real choice behaviour), AO_{rob} may, however, result in solutions which are overly conservative and provide suboptimal revenues on average. In practice, it may therefore be beneficial to restrict the set of considered choice models used in AO_{rob} as stated in Proposition 1.

Proposition 1. *$z_{rob}^*(\bar{\mathcal{D}}) \geq z_{rob}^*(\mathcal{D})$, where $\bar{\mathcal{D}} \subset \mathcal{D}$.*

Proof. Let $P(\bar{\mathcal{D}}) := \{(\boldsymbol{x}, \boldsymbol{y}) : (10) \text{ - } (15)\}$ indicate the feasible region of AO_{rob} under a given uncertainty set $\bar{\mathcal{D}}$, where Constraints (10) - (15) are written for a subset of scenarios indicated as $\bar{\mathcal{D}}$ instead of the complete uncertainty set $\mathcal{D}$. By definition, the set of feasible region of $P(\mathcal{D})$ is subsumed by the feasible region of $P(\bar{\mathcal{D}})$, since $|\mathcal{D}| \geq |\bar{\mathcal{D}}|$ and $\bar{\mathcal{D}} \subset \mathcal{D}$. $\square$

Based on this proposition, it seems appropriate to eliminate statistical outliers from the used ambiguity set and to concentrate on choice models which are statistically more likely to well represent the underlying choice behaviour. This idea will be numerically explored in Section 4.3.

A decomposition method for AO_{rob}.

Although it is possible to formulate AO_{rob} by including all choice models in $\mathcal{D}$, the resulting model may quickly become intractable if $\mathcal{D}$ is large. We here propose to solve AO_{rob} by means of a Benders type min-max decomposition framework [see, e.g. 8]. The problem is split into a *master problem* and a *subproblem* AO_{rob}^{Sub}. The master problem AO_{rob}^{MP} has a similar structure as the original problem AO_{rob}, but operates on a restricted ambiguity set $\tilde{\mathcal{D}} \subseteq \mathcal{D}$. Solving the master problem at iteration t yields an optimal assortment $\boldsymbol{x}_{\tilde{\mathcal{D}}}$ with worst-case revenue estimate $MP^t(\tilde{\mathcal{D}})$. This optimal assortment is then used in the subproblem which searches for a new choice model $C_d \in \mathcal{D} \setminus \tilde{\mathcal{D}}$ that has the lowest revenue estimate for assortment $\boldsymbol{x}_{\tilde{\mathcal{D}}}$. In the case that this revenue estimate is lower than MP^t, the corresponding choice model is added to the restricted ambiguity set $\tilde{\mathcal{D}}$. This procedure is repeated until no other choice model in $\mathcal{D} \setminus \tilde{\mathcal{D}}$ worsens the total expected revenue, indicating that $MP^t(\tilde{\mathcal{D}})$ corresponds to the highest worst-case revenue among the original ambiguity set $\mathcal{D}$. A detailed explanation of the decomposition procedure is given in Section 3 of the Online Supplement.

To formulate AO_{rob}, and identically, AO_{rob}^{MP}, we outer linearize the inner optimization problem by adding new constraints. The master problem can then be written as:

$$MP^t(\tilde{\mathcal{D}}) = \max_{\boldsymbol{x},\boldsymbol{y}} \pi \qquad (AO_{rob}^{MP})$$

$$\text{s.t. } \pi \leq \mathcal{R}(C_d) \qquad d \in \tilde{\mathcal{D}} \qquad (16)$$

$$\sum_{i \in \mathcal{N} \cup \{0\}} y_{id}^k = 1 \qquad d \in \tilde{\mathcal{D}},\ k \in \mathcal{K}_d \qquad (17)$$

$$y_{id}^k \leq x_i \qquad d \in \tilde{\mathcal{D}},\ k \in \mathcal{K}_d,\ i \in \mathcal{N} \qquad (18)$$

$$\sum_{j:\sigma_d^k(j) > \sigma_d^k(i)} y_{jd}^k \leq 1 - x_i \qquad d \in \tilde{\mathcal{D}},\ k \in \mathcal{K}_d,\ i \in \mathcal{N} \qquad (19)$$

$$\sum_{j:\sigma_d^k(j) > \sigma_d^k(0)} y_{jd}^k = 0 \qquad d \in \tilde{\mathcal{D}},\ k \in \mathcal{K}_d \qquad (20)$$

$$x_i \in \{0, 1\} \qquad i \in \mathcal{N} \qquad (21)$$

$$y_{id}^k \geq 0 \qquad d \in \tilde{\mathcal{D}},\ k \in \mathcal{K}_d,\ i \in \mathcal{N} \cup \{0\}. \qquad (22)$$

The objective function along with Constraints (16) maximize the worst-case revenue among all choice models within the restricted ambiguity set. Constraints (17)–(20) enforce customer preferences among the choice models in $\tilde{\mathcal{D}}$, and Constraints (21)–(22) enforce binary and non-negativity conditions.

Solving the master problem yields an optimal assortment $\boldsymbol{x}_{\tilde{\mathcal{D}}}$. The specific product preferences of each customer type $k \in \mathcal{K}_d,\ \forall d \in \tilde{\mathcal{D}}$ under the current assortment $\boldsymbol{x}_{\tilde{\mathcal{D}}}$ can then be derived as:

$$y_{id}^{k*} = \begin{cases} 1 \text{ , if } \sigma_d^k(i) < \sigma_d^k(j),\ \forall j \in \boldsymbol{x}_{\tilde{\mathcal{D}}},\ i \neq j \\ 0 \text{ , otherwise.} \end{cases} \qquad (23)$$

We use y_{id}^{k*} as an input to the subproblem to define customer preferences under the current optimal assortment. The subproblem, seeking for the worst-case choice model C_d in $\mathcal{D} \setminus \tilde{\mathcal{D}}$ for the current assortment, can be formulated as:

$$SP(\boldsymbol{y}^*, \mathcal{D} \setminus \tilde{\mathcal{D}}) = \min_s \sum_{d \in \mathcal{D} \setminus \tilde{\mathcal{D}}} \sum_{k \in \mathcal{K}_d} \sum_{i \in \mathcal{N}} r_i y_{id}^{k*} \lambda_d^k s_d, \qquad (AO_{rob}^{Sub})$$

$$\text{s.t.} \sum_{d \in \mathcal{D} \setminus \tilde{\mathcal{D}}} s_d = 1 \qquad (24)$$

$$s_d \in \{0, 1\} \qquad \forall d \in \mathcal{D} \setminus \tilde{\mathcal{D}}. \qquad (25)$$

The subproblem seeks for the single choice model s_d such that the total revenue for the current assortment is minimal. If $SP(\boldsymbol{y}^*, \mathcal{D} \setminus \tilde{\mathcal{D}}) < MP(\tilde{\mathcal{D}})$, the choice model s_d indeed estimates the revenue lower than any of the choice models in the current restricted ambiguity set $\tilde{\mathcal{D}}$, and is hence added:

$$\tilde{\mathcal{D}} = \tilde{\mathcal{D}} \cup \{d\} : s_d = 1.$$

The algorithm then iteratively proceeds, i.e., the master problem is resolved with the updated set $\tilde{\mathcal{D}}$ and a new choice models is added. The algorithm stops once no remaining choice model worsens the objective function for the current assortment, i.e., $SP(\boldsymbol{y^*}, \mathcal{D} \setminus \tilde{\mathcal{D}}) \geq MP(\tilde{\mathcal{D}})$. At this stage, we can also confirm that $MP(\mathcal{D}) = MP(\tilde{\mathcal{D}})$, and $\boldsymbol{x}^*_D$ is indeed optimal for Problem AO_{rob}.

We finally note that, instead of maximizing the worst-case revenue in the objective function of AO_{rob}, the robust optimization model could also minimize the opportunity loss [see, e.g. 25]. The decomposition method proposed above can easily be adapted to such a case. As mentioned in Section 2.2.2, such an opportunity loss can also be defined in terms of a performance metric. Interestingly, the equivalent of the opportunity cost as defined by performance metric M_{AvOpp} in Equation (7) cannot be directly used[2] within the objective function of AO_{rob}, which highlights an interesting advantage of using performance metrics to reduce the ambiguity set (i.e., by ranking the choice models using their corresponding performance metric values and then discarding those whose rank are lowest on the list), prior to applying an optimization model.

4 Computational Study

We now discuss a series of empirical results stemming from an extensive set of computational experiments. We first explain the general experimental set-up. Section 4.1 quantifies the magnitude of encountered CM ambiguity regarding the expected GT of the induced assortment decisions. We also demonstrate the effectiveness of the proposed performance metrics to identify high-quality choice models. Section 4.2 then focuses on the quality of the assortments produced by the robust optimization model. Finally, Section 4.3 illustrates the benefits when using a smaller ambiguity set of choice models selected according to the proposed performance metrics.

General experimental set-up and source of ambiguity. The behaviour of the market participants will be represented by a ground-truth (GT) rank-based choice model. Computational experiments are carried out for $N = 30$ products in the product set (as well as the no-purchase option 0), which helps us demonstrate a practical number of alternatives that customers may prefer to choose from in practical settings. The choice model ambiguity set will be composed of choice models that are calibrated by the same estimation procedure. Here, we consider two popular column-generation based estimation procedures: an L_1-*error (least absolute deviation) minimization* [see, e.g. 5, 14] and the log-likelihood maximizing *market discovery* algorithm [23]. Choice models will be estimated on sales data stemming from the GT choice model on a total of 20 assortments (called the *training set*), where each assortment contains 8 randomly selected products (in addition to the no-purchase option). Once the choice models are estimated, the revenue of the resulting optimized assortments will be estimated on the original GT on 100 assortments (called the *test set*). While, in practice, there may be a variety of potential sources of choice model ambiguity, we here focus on the most elementary, and practically inevitable source of ambiguity[3]: random behaviour in the estimation procedures due to different random seeds. This implies that we assume to have access to perfect historical sales information and any type of additional error source (such as sampling errors of sales transactions) may result in even higher choice model ambiguity.

Choice model problem instances. A problem instance corresponds to a GT rank-based choice model composed of $K = 10$ different customer types, each of which has a preference defined over 10 (among the 30) products before the no-purchase option 0 is listed. In order to assess the validity of our results across a variety of different types of market behaviour, GT choice models are designed with different degrees of product preferences. Choice models using market structure **S1** have a random ranking of products in their preference lists, corresponding to a highly diversified clientele. Two other market structures assume that half of the products are highly popular products, whereas the other half are less popular products. In **S2**, the number of highly popular products, randomly selected and ranked first in a preference sequence is given by a random number uniformly chosen from $[0, 9]$. In **S3**, the first 5 products are highly popular products (randomly selected), whereas the following 5 are less popular products (randomly selected).

[2] Performance metric M_{AvOpp} uses $|R(\boldsymbol{x_d}^*, C_m) - R(\boldsymbol{x_d}^*, C_d)|$, where $\boldsymbol{x_d}^*$ is fixed and hence cannot be optimized. Using $\boldsymbol{x}$ instead would minimize the difference between the two revenue estimates, but the reference to the optimal revenue of $\boldsymbol{x_d}^*$ would be lost.

[3] Note that the number of available assortments in the training is, theoretically, also a potential source of ambiguity. However, given the reasonable size of the product universe ($N = 30$), 20 training assortments is assumed to be sufficiently large. Furthermore, more training assortments may be difficult to obtain in practice.

The probability distribution $\boldsymbol{\lambda}$ is uniformly selected at random for each choice model and then scaled, while also taking into consideration that the probabilities allow for accurate transaction sampling that represents perfect sales information (which is required for the market discovery estimation procedures, which takes as input a set of transactions).

The optimization of the assortments will be explored under two different revenue settings: in setting **R1**, all product revenues are uniformly selected at random $r_i \in U[200, 1000]$, while setting **R2** assumes increasing product revenues such that product i is chosen from $r_i \in U[0, 200i]$ as proposed by [11].

The detailed description for each setting and their generation procedure, as well as details regarding the generation of $\boldsymbol{\lambda}$ can be found in Section 1 of the Online Supplement.

Computing environment. All the experiments provided were solved on the Cedar cluster provided by Compute Canada using CPLEX 12.10.0 (Python 3.7 API). Estimation of rank-based choice models were performed under a limit of 40 CPUs and 256MB RAM (per core). The assortment optimization problem was solved under a maximum of 40 CPUs for deterministic and robust models. The detailed characteristics of CPUs can be found at `https://docs.computecanada.ca/wiki/Cedar`.

4.1 Optimization on a nominal choice model and magnitude of ambiguity

Magnitude of ambiguity under different estimation procedures. Choice model ambiguity may be crucial to the decision-maker on two levels: the predictive accuracy of the choice model and the quality of the induced assortment decisions. We quantify both dimensions. Specifically, we generate a total of 50 unique ground truth models for each instance class S1, S2 and S3. For each problem instance, a total of 50 choice models are estimated by each of the estimation procedures: the L_1-error minimization and the log-likelihood maximizing *market discovery* algorithm. Note, again, that for the estimation of the different choice models, we only modify the initial random seed, but assume perfect sales information (i.e., the exact values for $v_{i,m}$) to keep ambiguity as low as possible. The optimal assortment for each choice model is then obtained by solving to optimality the corresponding deterministic MIP formulation [as introduced by 7].

We then evaluate the predictive accuracy $\epsilon_{test}^{C_d}$ of the choice models on $M^{test} = 100$ test assortments, defined as the L_1-error and normalized by the number of assortments: $\epsilon_{test}^{C_d} = \frac{|A_{C_d}^{test}\lambda_{C_d} - v_{C_d}|}{2M^{test}}$, where $A_{C_d}^{test}$ determines the choices under the preference sequences in choice model C_d among the test assortments. The quality of the induced assortment decisions will be measured as the revenue percentage captured of the revenue given by the optimal assortment, computed based on the original GT model (which, in practice, is unknown).

	L_1-error minimization					Market Discovery				
	Min Avg.	Mean Avg.	Max Avg.	Median	Avg. St.Dev.	Min Avg.	Mean Avg.	Max Avg.	Median	Avg. St.Dev.
Predictive Accuracy	0.226	0.247	0.269	0.247	0.009	0.227	0.242	0.257	0.242	0.007
Captured GT%	58.9	82.1	97.2	83.6	10.0	57.9	80.4	96.1	80.6	10.5

Table 2: Avg. GT % captured and avg. predictive accuracy of choice models estimated under L_1-error minimization vs. market discovery algorithm (instance class S3R1, 50 instances).

Table 2 shows the average (over the 50 instances for instance class S3R1) minimum, mean, median and maximum predictive test errors and GT percentages among the 50 choice models estimated for each instance. The average standard deviation column refers to the average (over 50 instances) standard deviation among the 50 choice models in each instance class. The choice models are estimated by two different estimation methods: the L_1-error minimization and the log-likelihood maximizing *market discovery* algorithm.

Both estimation methods result in choice models that generalize to new assortments similarly well with similar average errors (from 0.226 to 0.269) and small standard deviation. Any CM will therefore provide a similar predictive accuracy. In contrast, the quality of the induced decisions varies strongly, ranging from about 58% to 97% of the optimal assortment revenue with a standard deviation of 10% of the optimal revenue. Unfortunately, we found that there is no direct correlation between the predictive accuracy and the assortment quality. However, this is not surprising, since all choice models have been trained under the same conditions (to the same threshold of training error). As a result, the decision-maker has no effective tools to identify the right choice model and assortment, and may end up choosing any of those 50 assortments.

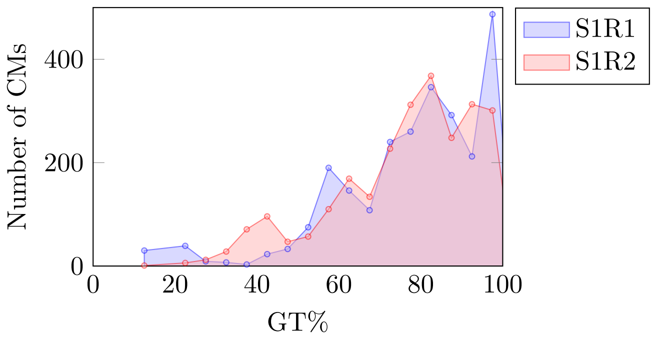

Figure 3: Ground truth revenue distribution of instance classes S1R1 and S1R2, each including 2500 choice models.

Figure 3 shows the GT revenue distribution of the 2500 different assortments (induced by 50 choice models using the L_1-error estimation method for each of the 50 problem instances) for instance classes S1R1 and S1R2. For the two instances classes, the GT% of the optimal revenue captured has been found to be 77.8% and 76.2%, respectively on average. The minimum has been found to be 14.9% and 12.5%, respectively, and the maximum 100% for both instance classes. In practice, ambiguity may be even higher, since other error sources (such as imperfect sales information) may result in less reliable choice models. While, based on Figure 3, there is a reasonable chance that a high-quality assortment is selected, the risk of selecting an assortment of low quality is unreasonably high.

A similar analysis has been carried out for an MMNL generating GT model , which has been a popular approach in the literature [see, e.g. 5, 14]. While the CM ambiguity appears to be lower than a RB generating GT model, revenue estimates still varied by up to 50% among the optimal assortments. We refer to Section 4.1 of the Online Supplement for more details. The forthcoming computational results are therefore derived under a RB generating GT model. Further, given that both estimation methods result in ambiguity of similar magnitude, we will focus on L_1-error minimization.

Use of performance metrics to select a nominal choice model. We now explore the magnitude of choice model ambiguity within the other instance classes (S1, S2 and S3, coupled with R1 and R2), as well as the ability of the performance metrics to identify reliable choice models that tend to produce high-quality decisions.

		R1			R2		
	GT % (Avg.)	S1	S2	S3	S1	S2	S3
Deterministic	Min	59.8	62.9	58.9	57.6	60.0	58.8
	Mean	77.8	82.0	82.1	76.2	80.0	79.9
	Median	76.6	82.3	82.5	75.5	80.4	80.2
	Max	95.0	97.8	97.2	94.6	97.2	97.3
	St. Dev.	9.6	9.3	10.0	9.4	9.5	10.3
Metric (best)	M_{AvOpp}	**89.0**	**90.5**	91.0	**89.1**	**89.4**	**88.4**
	M_{AvRev}	83.4	89.7	89.2	80.7	81.8	85.0
	M_{MinRev}	88.9	87.7	**91.5**	85.8	87.7	87.7
	$M_{WeigLen}$	82.8	84.7	84.1	78.0	83.2	83.8

Table 3: Avg. GT% captured by the deterministic model and the metrics among different instance classes ($\mathcal{D} = 50$, 50 instances per class).

The upper part of Table 3 (indicated by "Deterministic") presents the average (over 50 instances) of the minimum, mean, median and maximum GT percentages, as well as the average standard deviation (as defined for Table 2) among the 50 choice models estimated for each instance across different instance classes. The overall choice model ambiguity remains significant among all six instance classes. In particular, some assortments only capture 57.6% of the GT optimal revenue, while others capture 97.8%. The last five rows in Table 3 present the average performance of the assortments induced by the choice models that have scored best according to each of the five performance metrics (see Section 2.2.2). Interestingly, all five metrics remarkably well avert low quality assortments and consistently present higher average revenues than the mean revenue of the assortments stemming from the individual assortments. The proposed decision-driven metrics seem to perform particularly well, being able to capture approximately an additional 4%-13% of the optimal revenue. In particular, the assortments identified by metrics M_{AvOpp} and M_{MinRev} capture an additional 8.5%–12.9% of the optimal revenue when compared to the mean performance of the separate CMs.

Impact of information loss. Most of the experiments here reported assume that perfect sales information is available. However, the amount of transaction data available from historical sales typically impacts the quality of the choice models that have been estimated, which, in turn can impact the predictive accuracy. We therefore also carried out experiments with imperfect sales information, specifically, with 500, 50 and 20 sales transaction per assortment. The results indicate that, generally, the solution quality degrades by 3–4% of GT revenue, while the standard deviation also tends to slightly increase. Indeed, no approach seems to be immune to the loss of information. Detailed results can be found in Section 4.3 of the Online Supplement.

4.2 Robust optimization on a set of choice models

We now explore the effectiveness of the robust optimization model AO_{rob}, optimizing an assortment on an ambiguity set of choice models (see Section 3). The robust model has been solved[4] on the entire ambiguity set of 50 choice models within each problem instance class.

	R1			R2		
GT %	S1	S2	S3	S1	S2	S3
Min	60.8	49.3	64.3	43.1	63.7	67.5
Mean	89.7	92.9	93.3	89.1	91.0	90.0
Median	93.0	95.7	96.4	92.0	94.0	92.8
Max	100.0	100.0	100.0	100.0	100.0	100.0

Table 4: Avg. GT% captured by AO_{rob} among different instance classes. ($\mathcal{D} = 50$, 50 instances per class).

Table 4 presents the average GT% result obtained from the robust model, indicating the minimum, mean and maximum values among the 50 instances under each instance class. The robust model considerably improves the mean performance of the resulting assortments, capturing an additional 10%-12% of the optimal revenue and outperform even the best performance metrics.

Finally, it is worth noting that the minimum and maximum GT% values found among the 50 instances in Table 4 cannot be directly compared to those in Table 3. For example, the minimum GT% of 43.1% based on the assortment from AO_{rob} in instance class S1R2 may refer to an instance GT that is particularly difficult to capture. For the deterministic model, the corresponding GT% in Table 3 is 57.6%, which refers to the average over the minima (as opposed to the absolute minimum). In fact, the smallest GT% captured by an assortment optimized on an individual assortment (as shown in Figure 3) is 12.5%.

It is, however, quite natural that the performance of the robust optimization model is negatively impacted by unreliable choice models within the ambiguity set. In the next section, we therefore aim at reducing the size of the ambiguity set by including only those choice models that have been considered more reliable by the proposed performance metrics.

Performance of the decomposition method. The here proposed decomposition framework significantly reduces the required computing time to solve the robust model. For the instance class S3R1, we found that the average time required to solve the robust problem under the full ambiguity set (i.e., without applying decomposition) is 308.1 seconds (with a minimum of 58.4 seconds and a maximum of 860.1 seconds). The computing time is substantially lower using the decomposition framework, with an average of 43.6 seconds (and a minimum and maximum of 2.8 and 242.7 seconds, respectively). Figure 4 illustrates the impact of implementing the decomposition framework in terms of computational time requirements. Note that the decomposition algorithm requires on average 7 iterations to reach and prove optimality (with a minimum of 3 and a maximum of 12 iterations). As a result, the size of the ambiguity set is reduced substantially, given that the original ambiguity set considered here includes 50 CMs.

[4]To solve the robust optimization model, we have used the decomposition algorithm with $\epsilon = 10^{-6}$ as termination criteria throughout the computational experiments (see Section 3 in the Online Supplement).

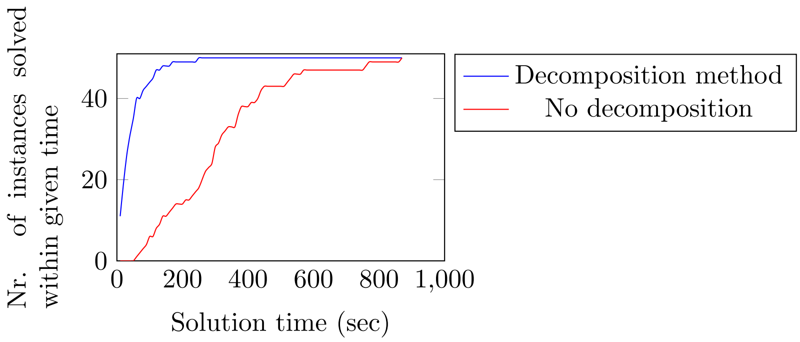

Figure 4: Number of instances solved within given time (in seconds).

4.3 Amending ambiguity sets according to metrics

Given that the robust model is likely to be sensitive to unreliable (i.e., low quality) choice models, we now aim at removing such statistic outliers, using the proposed performance metrics to select a subset $\overline{\mathcal{D}} \subseteq \mathcal{D}$ of CMs for the robust model. In the following experiments, we explore the impact on the model performance under varying size of the available initial ambiguity set $\mathcal{D}$ and the number k (such that $|\overline{\mathcal{D}}| = k$) of choice models to be selected according to a metric.

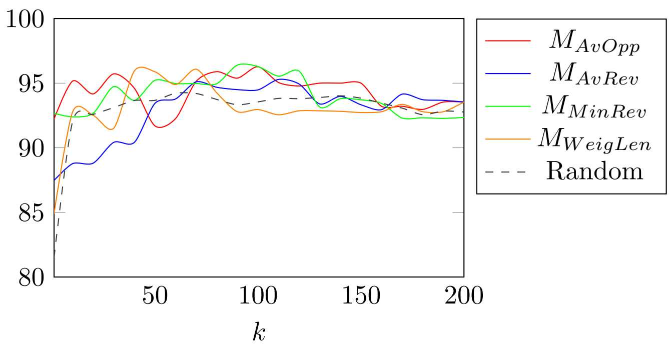

Figure 5: Avg. GT% captured by AO_{rob} under best k choice models according to various metrics (instance class S3R1, 10 instances, 200 choice models, 5 random orderings).

Figure 5 plots the GT% (averaged over the 10 problem instances) captured by the robust model, using a varying number of k best performing choice models according to the M_{AvOpp}, M_{AvRev}, M_{MinRev} and $M_{WeigLen}$ metrics. The first data point, at $k = 1$, corresponds to the case where the single best performing choice model according to a metric is used. The performance improves once more than 1 choice model is used. The performance of all metrics then tends to continue to improve when selecting 50 to 100 out of the 200 choice models, and tends to degrade afterwards. Both metrics M_{AvOpp} and M_{MinRev} seem to outperform the other metrics, which is consistent with the findings in Section 4.1. Generally, selecting the 25% - 50% best performing choice models seems to be appropriate. Note that, when using all choice models (i.e., $k = 200$), the results should theoretically be identical, but, empirically, they may slightly differ due to multiple optima in the assortment optimization models. The figure also reports the average performance over 5 random sequences (as opposed to a sequence proposed by a metric). The results indicate that the performance remains relatively stable with varying k, which suggests that the proposed metrics are indeed effective to discriminate choice models by their reliability (or the quality of the induced decisions).

The impact of using a number of random choice models is further illustrated in Figure 6. The figure shows the average GT% captured by the assortments stemming from k random choice models (averaged over 50 random orderings) and k choice models selected by metric M_{MinRev}. The shaded areas indicate the GT% range obtained under each approach. Selecting 25%–50% of the initial choice models according to M_{MinRev} exhibits a clear advantage in terms of worst-case revenue observed, which also directly impacts

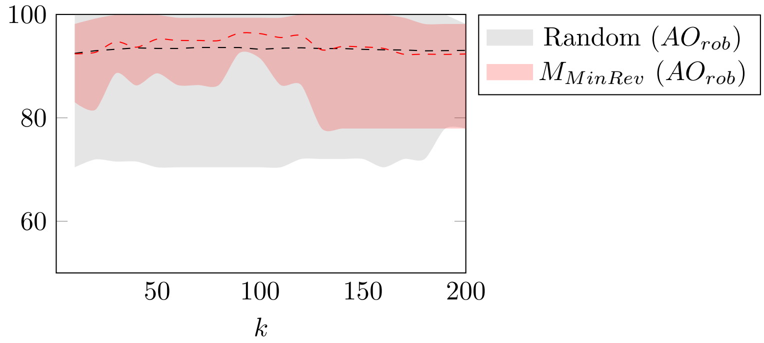

Figure 6: Robust model performance for different k, 50 random orderings, and metric M_{MinRev} over 10 instances.

the average performance. Additional experiments with fixed k and varying $|\mathcal{D}|$, as well as varying k and $|\mathcal{D}|$ confirmed that selecting the highest-scoring 25%–50% of the available assortments tends to be a beneficial choice. Details can be found in Section 4.4 of the Online Supplement.

Based on our computational results, using metrics M_{AvOpp} and M_{MinRev} with 25%-50% of the available choice models within AO_{rob} provides a solid performance, capturing about 96% of the optimal revenue on average. From the decision-makers perspective, this is an impressive performance boost, capturing an additional 13% of the optimal revenue when compared to the mean revenue induced by individual choice models (82.1%), and an additional 2-3% when compared to using AO_{rob} on all choice models (93.3%).

5 Conclusions

In this study, we address the issue of choice model ambiguity within the predict-then-optimize paradigm, which may lead to a varying degree of quality among the induced decisions. We propose a general evaluation framework and a variety of performance metrics to gauge the reliability of choice models. Focusing on the case of rank-based choice models in the popular field of assortment optimization, we then consider the robust optimization problem that optimizes on a finite set of choice models and propose an efficient decomposition algorithm, capable of reducing the required computing time by 85.9% on average.

The proposed approach has several advantages. From a manager's perspective, reducing choice model ambiguity prior to optimization (as opposed to directly optimizing on a parametric uncertainty certain for a choice model as proposed in the literature) allows for more control over which choice models to use within the optimization model. Further, the evaluation of the choice models via the proposed performance metrics are practically viable, i.e., it does not require access to the ground truth model and is solely based on the "opinion" of the other choice models.

Numerically, our results are highly encouraging. While we observed that choice model ambiguity in the case of rank-based choice models resulted in estimated assortment revenues that differed by as much as 60% of the optimal revenue, the proposed metrics used within a robust optimization model were indeed effective to reduce both ambiguity and improve the captured revenue. Moreover, for a specific instance class, the mean revenue induced by individual choice models was found to be 82.1% of the optimal revenue. Compared to this mean performance: (i) the assortments induced by the choice model that has scored best according to a specific metric captured about an additional 9% of the optimal revenue; (ii) the robust model optimized on all choice models captured about an additional 11% of the optimal revenue; (iii) using the metric to select a subset of the choice models within the robust model captured about an additional 13% of the optimal revenue, which corresponds to assortments that capture 96% of the optimal revenue.

It may be particularly appealing to practitioners that our framework is compatible with other types of choice models and other estimation methods [e.g. 24, 14, 15]. For example, using our framework is also likely to be beneficial under an MMNL generating GT. For future work, the evaluation framework and performance metrics can be used in other applications that involve scenario identification to help reduce ambiguity. Likewise, the range of performance metrics involved in this study can further be extended by considering different problem characteristics.

Acknowledgements

The authors are thankful to Compute Canada for providing the computational resources to carry out the experiments. The second author was supported by the Natural Sciences and Engineering Research Council (NSERC) of Canada under grant 2017-05224, as well as by the Fonds de recherche du Québec – Nature et technologies (FRQNT) under grant 70739202. The third author was the Canada Research Chair in Stochastic Optimization of Transport and Logistics Systems. Additionally, he was supported by the Natural Sciences and Engineering Research Council (NSERC) of Canada under grant 2018-05390. The first author was partially supported by the aforementioned grants.

References

[1] Kenneth J Arrow. *Social choice and individual values*. Yale university press, 2012.

[2] Moshe Ben-Akiva. The structure of travel demand models. *PhD thesis, MIT*, 1973.

[3] Gerardo Berbeglia, Agustín Garassino, and Gustavo Vulcano. A comparative empirical study of discrete choice models in retail operations. *Management Science*, 68(6):4005–4023, 2022.

[4] Dimitris Bertsimas and Nathan Kallus. From predictive to prescriptive analytics. *Management Science*, 66(3):1025–1044, 2020.

[5] Dimitris Bertsimas and Velibor V Mišic. Data-driven assortment optimization. Working paper, MIT, 2016.

[6] Dimitris Bertsimas and Velibor V Mišić. Robust product line design. *Operations Research*, 65(1):19–37, 2017.

[7] Dimitris Bertsimas and Velibor V Mišić. Exact first-choice product line optimization. *Operations Research*, 67(3):651–670, 2019.

[8] Daniel Bienstock and Nuri Özbay. Computing robust basestock levels. *Discrete Optimization*, 5(2):389–414, 2008.

[9] Yi-Chun Chen and Velibor V Mišić. Decision forest: A nonparametric approach to modeling irrational choice. *Management Science*, 2022.

[10] Dick den Hertog and Krzysztof Postek. Bridging the gap between predictive and prescriptive analytics-new optimization methodology needed. *Tilburg Univ, Tilburg, The Netherlands*, 2016.

[11] Antoine Désir, Vineet Goyal, Bo Jiang, Tian Xie, and Jiawei Zhang. A nonconvex min-max framework and its applications to robust assortment optimization, 2019.

[12] Adam N Elmachtoub and Paul Grigas. Smart "predict, then optimize". *Management Science*, 68(1):9–26, 2022.

[13] Vivek F Farias, Srikanth Jagabathula, and Devavrat Shah. A nonparametric approach to modeling choice with limited data. *Management science*, 59(2):305–322, 2013.

[14] Sanjay Dominik Jena, Andrea Lodi, Hugo Palmer, and Claudio Sole. A partially ranked choice model for large-scale data-driven assortment optimization. *INFORMS Journal on Optimization*, 2(4):297–319, 2020.

[15] Sanjay Dominik Jena, Andrea Lodi, and Claudio Sole. On the estimation of discrete choice models to capture irrational customer behaviors. *INFORMS Journal on Computing*, 2022.

[16] A Gürhan Kök, Marshall L Fisher, and Ramnath Vaidyanathan. Assortment planning: Review of literature and industry practice. *Retail supply chain management*, pages 99–153, 2008.

[17] Saharnaz Mehrani and Jorge A Sefair. Robust assortment optimization under a ranking-based choice model with product unavailability effect. *Working paper*, 2020.

[18] Diego C Nocetti. Ambiguity and the value of information revisited. *The Geneva Risk and Insurance Review*, 43(1):25–38, 2018.

[19] Paat Rusmevichientong, David Shmoys, Chaoxu Tong, and Huseyin Topaloglu. Assortment optimization under the multinomial logit model with random choice parameters. *Production and Operations Management*, 23(11):2023–2039, 2014.

[20] Paat Rusmevichientong and Huseyin Topaloglu. Robust assortment optimization in revenue management under the multinomial logit choice model. *Operations research*, 60(4):865–882, 2012.

[21] Arthur Snow. Ambiguity and the value of information. *Journal of Risk and Uncertainty*, 40(2):133–145, 2010.

[22] Bradley Sturt. The value of robust assortment optimization under ranking-based choice models. *Working paper*, 2021.

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

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

[25] Tan Wang and Genaro Gutierrez. Robust product line design by protecting the downside while minding the upside. *Production and Operations Management*, 2021.

[26] Zhengchao Wang, Heikki Peura, and Wolfram Wiesemann. Randomized assortment optimization. *Working paper*, 2020.