

Data-Driven Revenue Management: The Interplay of Data, Model, and Decisions

Ningyuan Chen,^{a,*} Ming Hu^a

^aRotman School of Management, University of Toronto, Toronto, Ontario M5S 3E6, Canada

*Corresponding author

Contact: ningyuan.chen@utoronto.ca,  <https://orcid.org/0000-0002-3948-1011> (NC); ming.hu@rotman.utoronto.ca,

 <https://orcid.org/0000-0003-0900-7631> (MH)

Received: January 17, 2023

Accepted: February 6, 2023

Published Online in Articles in Advance:
March 10, 2023

<https://doi.org/10.1287/serv.2023.0322>

Copyright: © 2023 INFORMS

Abstract. Revenue management (RM) is the application of analytical methodologies and tools that predict consumer behavior and optimize product availability and prices to maximize a firm's revenue or profit. In the last decade, data has been playing an increasingly crucial role in business decision making. As firms rely more on collected or acquired data to make business decisions, it brings opportunities and challenges to the RM research community. In this review paper, we systematically categorize the related literature by how a study is "driven" by data and focus on studies that explore the interplay between two or three of the elements: data, model, and decisions, in which the data element must be present. Specifically, we cover five data-driven RM research areas, including inference (data to model), predict then optimize (data to model to decisions), online learning (data to model to decisions to new data in a loop), end-to-end decision making (data directly to decisions), and experimental design (decisions to data to model). Finally, we point out future research directions.

Funding: The research of N. Chen is partly supported by Natural Sciences and Engineering Research Council of Canada Discovery [Grant RGPIN-2020-04038]. The research of M. Hu is in part supported by Natural Sciences and Engineering Research Council of Canada [Grant RGPIN-2021-04295].

Keywords: revenue management • pricing • data-driven • inference • predict then optimize • online learning • end-to-end • experimental design

1. Introduction

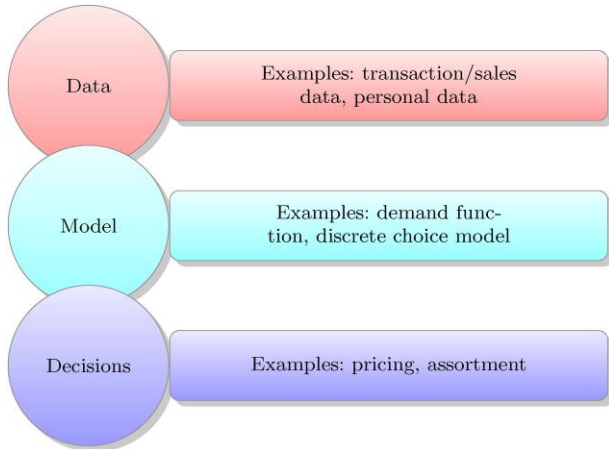
Revenue management (RM) is the application of analytical methodologies and tools that predict consumer behavior and optimize product availability and prices to maximize a firm's revenue or profit. In the last decade, data has been playing an increasingly crucial role in business decision making. As firms rely more on collected or acquired data to make business decisions, it brings opportunities and challenges to the RM research community. On the one hand, because many traditional RM models do not automatically incorporate the data component, the industry trend calls for innovative models, frameworks, algorithms, and policies that are more data-driven, and as a result, the area of data-driven RM blossoms. On the other hand, researchers working on data-driven RM often find it hard to define the area clearly. From a methodological point of view, the boundaries between a number of disciplines, such as operations research, statistics, econometrics, and computer science, have never been blurrier when approaching data-driven research. From the practical point of view, the topic of RM increasingly interweaves with other traditional topics in operations management, such as supply chain and service operations management.

In this review paper, our objective is to systematically categorize the related literature by how a study is "driven"

by data. In particular, in data-driven studies, there are usually three key elements: *data*, *model*, and *decisions* (Bertsimas and Freund 2004, Simchi-Levi 2014), and the interplay among the elements is explored by researchers. For data-driven research in RM, the three elements usually encompass specific contexts related to RM:

- **Data:** Firms collect data to improve the understanding of the market or the quality of the decisions. Depending on the specific setting, the form of the data that is available to the firm or researchers is usually different and a starting point. In RM, for example, the transaction or sales data are a common form of data, recording the sales of products, the prices of the offered assortment, and available promotions. Recently, the personal data of customers, such as age and past transactions, are also used to make personalized recommendations and set personalized promotions.
- **Model:** In the context of RM, the model typically refers to the following two categories that have a significant overlap: a utility-maximization economic framework that explains the behavior of consumers, such as competitive models and strategic behavior, or a statistical framework that explains how the data are generated, such as demand functions and discrete choice models.
- **Decisions:** With data and models, firms are ultimately concerned with making reasonable decisions that benefit

Figure 1. The Three Elements of Data-Driven Research in Revenue Management and Examples



themselves in the short or long term. Pricing and assortment planning are the two primary operational decisions traditionally considered in RM. Recently, there have been a number of new decision regimes inspired by novel business models and available data sources, such as product ranking and personalized recommendations.

A summary of the three elements is illustrated in Figure 1.

1.1. Scope of the Review

In this review, we focus on studies that explore the interplay between two or three of the elements: data, model, and decision, in which the data element must be present. Moreover, we include studies that specialize in an RM application or propose general frameworks that can be directly applied to RM problems. That is, the model element usually involves demand or consumer choices; the decision element is typically related to pricing and assortment although this is not a strict criterion. We do not include studies that focus on inventory management as the primary decision, mainly because there is a large body of literature and already several excellent reviews on inventory and supply chain management (Chen et al. 2022b, chapters 11–13) although inventory is a critical operational lever for firms. We choose not to specifically include empirical studies on RM topics, mainly because of the limited space. We note that there are emerging areas in RM attracting a lot of attention, such as auction, market/platform design, and reusable resources. We include a number of papers from these areas and incorporate them in the framework listed in Figure 1.

For the structure of the review, we take an unusual approach. Instead of categorizing the studies by their applications or topics, we primarily group them by the methods they apply to the interplay of those elements. This treatment can provide an overview of the methodologies for readers who are interested in conducting

data-driven research in a broader area of operations management.

2. Inference

Inference or estimation is a central topic in statistics and predictive analytics. Given a class of models or the parametric/nonparametric structure that generates the data, there are numerous methods developed to find the model in the class that fits the data best and can be used to extrapolate to future patterns. Among them, the least-square and maximum likelihood estimators are probably the most celebrated. Inference is concerned with the interplay of data and model as shown in Figure 2.


Recent developments in RM have seen the benefits of more sophisticated models. These models offer flexibility in explaining the behavior of consumers and may lead to interpretable and efficiently computable optimal decisions. In turn, the general approaches may not work, and they call for specially designed estimation procedures. Next, we review these papers according to their applications.

2.1. Discrete Choice Models

The bulk of the literature that studies the estimation problem in RM is concerned with the estimation of discrete choice models. Discrete choice models provide a framework to describe and explain the choice behavior of consumers when they are faced with a set of products. Train (2009) provides a comprehensive review of commonly used discrete choice models and their estimation, including the multinomial logit (MNL), generalized extreme value, probit, and mixed logit models. The book also discusses the estimation problem when the consumer and product features are present in the data. Although the general framework for estimating these models has been developed, typically using likelihood-based methods, the RM community has been focusing on issues that are not covered in the general framework or the estimation of recently developed discrete choice models.

2.1.1. Practical Considerations. One important practical issue is the effect of stockouts on the choice of consumers. When the inventory of a particular product runs out, discrete choice models, especially the MNL model, provide powerful mechanisms to capture the substitution effects in the purchase behavior. Anupindi et al. (1998) is one of the early papers that focuses on demand estimation in this situation. Kök and Fisher (2007) and Musalem et al. (2010) provide a more formal treatment for the estimation problem using the MNL model. Wan et al. (2018) study the substitution effect not only inside

Figure 2. Inference



a store, but also between stores. Through numerical simulation, they find that the nested logit model (see Train 2009, chapter 4, for more details) provides accurate estimation. Azadeh et al. (2015) use a different framework of mixed-integer nonlinear programming to formulate the estimation problem in order to minimize the estimated and observed bookings in an RM setting.

Another important issue is the unobserved lost demand or demand censoring: the consumers who browsed the offered products but did not make a purchase are not observed. Vulcano et al. (2012) tackle the problem by postulating a nonhomogeneous Poisson process for the arrival process of consumers and the MNL model for the choice. They use the expectation–maximization (EM) algorithm (Dempster et al. 1977) to handle the unobserved lost demand. Abdallah and Vulcano (2021) further investigate the theoretical properties, such as the identifiability of this framework, and provide an efficient algorithm, minorization-maximization, for the maximization of the likelihood function. Facing the problem of demand censoring from a different angle (the demand for some alternatives is not observed), Newman et al. (2014) provide a two-step algorithm to maximize the likelihood function. Instead of likelihood-based methods, Subramanian and Harsha (2021) introduce mixed-integer programming to directly minimize a convex loss function between the market share in the data and the model in the presence of demand censoring. Amjad and Shah (2017) employ matrix completion to estimate censored demand. Cho et al. (2023) apply Bayesian inference to data from the hotel industry, which handles unobserved no purchases as well as customer clustering simultaneously. A review can be found in Azadeh et al. (2014).

Wang (2021) investigates the MNL model when the market size may vary based on the total attraction of products in the offered assortment. An estimation procedure based on the EM algorithm is provided. Chen et al. (2022d) study the theoretical properties of maximum likelihood estimation for the MNL model when the data are collected based on personal information. The majority of the studies in this area focus on online learning. See Section 4.1 for more details.

2.1.2. Estimation for Recently Developed Discrete Choice Models. Recently, there have been new discrete choice models proposed to capture more complex consumer behavior to overcome the limitation of random utility models. For example, Farias et al. (2013) propose a rank-based discrete choice model that requires very few structural assumptions of the choice behavior. As a result, the estimation problem given the transaction data of the model becomes very challenging because of the factorial number of parameters: the estimation is typically unidentifiable, and there are numerous sets of parameters that are consistent with the data. Farias et al. (2013) introduce a robust approach and attempt to identify

the parameters that minimize the expected revenue. The resulting optimization problem still has a large number of variables, or equivalently, the dual problem has many constraints. The authors provide approximations to the dual problem by sampling constraints or coming up with efficient representations of the constraints. How to obtain a sparse solution, which takes as few as possible nonzero weights for the rankings, is also discussed in the paper and further explored in Farias et al. (2020). Instead of robust optimization, van Ryzin and Vulcano (2015) take a likelihood-based approach. To maximize the likelihood function and handle the astronomical number of variables, van Ryzin and Vulcano (2015) gradually add new rankings to the current set via column generation. The resulting solution is, thus, sparse. A similar approach is taken in Jena et al. (2020) for partially ordered rankings. As an alternative for likelihood maximization, van Ryzin and Vulcano (2017) use the EM algorithm by treating the actual preference ranking of each customer as missing variables. The issue of demand censoring is also studied in this paper. Ho-Nguyen and Kılınc-Karzan (2021) focus on the estimation of rank-based choice models in a dynamic setting when the observations are collected over time. They formulate a mathematical program to minimize the distance between the implied choice probability from the model and the data. It is cast as a dynamic saddle point problem using a primal–dual framework. Jagabathula and Rusmevichientong (2017) extend the choice model to incorporate price information. Customers are making choices based on the preference ranking subject to a price threshold generated from a common distribution. The estimation is conducted using the EM algorithm due to the unobserved rankings and thresholds. The rank-based model is extended in Jagabathula and Vulcano (2018), which provides a nonparametric framework to estimate the preference of customers from panel data, that is, the purchase history of individual customers as well as the product availability and promotion at the time of purchases. The goal is not to recover the preference ranking but a directed acyclic graph (DAG) that represents the partial order (a set of preferences over pairs of products) of the customer. They use likelihood-based methods to construct the DAG. The DAGs are then clustered among similar customers with conflicting edges removed.

Besides the rank-based model, there are other discrete choice models that are gaining popularity in RM. The Markov chain choice model and its associated assortment optimization problem are studied in Blanchet et al. (2016), which encapsulates the MNL model as a special case. Şimşek and Topaloglu (2018) use the EM algorithm to estimate the parameters of the Markov chain choice model. The parameter identification problem of the model is studied in Gupta and Hsu (2017). Replacing the discrete-time Markov chain with

its continuous-time counterpart, Ragain and Ugander (2016) study a similar but different choice model and its likelihood-based estimation.

One stream of literature builds discrete choice models by solving a consumer welfare maximization problem for a representative consumer when the utilities toward products are random (Natarajan et al. 2009, Mishra et al. 2014). Given the marginal distribution or the first and second moments of the random utilities, the choice probabilities can be solved using convex optimization that maximizes the expected welfare over all eligible joint distributions. This allows for the estimation of the parameters by maximizing the likelihood on top of the choice probabilities.

Recently, a special case of the random utility model, the exponential choice model, is proposed in Alptekinoglu and Semple (2016). The random utility has an exponential distribution, potentially with heteroscedasticity, instead of the Gumbel distribution as in the MNL model. Alptekinoglu and Semple (2016, 2021) and Aouad et al. (2018) demonstrate that, under the exponential choice model, the log-likelihood function is concave in the parameters, and standard convex optimization tools can be leveraged for the maximum likelihood estimation.

Even more flexible than the rank-based choice model, tree-based choice models are proposed in Chen et al. (2019c) and Chen and Mišić (2022). They are shown to be able to represent any discrete choice models. Because of the model's nonparametric nature, its estimation becomes challenging. Chen and Mišić (2022) propose mixed-integer programming to find a sparse set of trees of limited depth. The authors show that a number of "shallow" trees that are logarithmic in the number of assortments are sufficient to fit the data. Chen et al. (2019c) modify random forests, a popular algorithm in machine learning, for the estimation of such a model. The authors establish the consistency of this method, analyze the prediction error, and discuss the practical flexibility of this approach. In Aouad et al. (2022), decision trees are used to segment customers based on the response of customers in various settings, such as choice models or bidding in auctions. The method also demonstrates promising empirical performance. The idea of using machine learning algorithms to estimate discrete choice models is further explored in Cai et al. (2022) and Aouad and Désir (2022), which use neural networks to model and capture the choice behavior. Thanks to the development of deep learning, neural networks can be efficiently estimated with standard programming packages.

2.1.3. Empirical Studies of Discrete Choice Models.

Vulcano et al. (2010) study the empirical performance of the MNL model using data from a major U.S. airline. For the estimation of the model, they provide solutions to two practical issues: demand censoring and the fact that alternatives in the choice set need to be inferred from

the data. Another empirical study in a similar application is conducted in Ratliff et al. (2008). Feldman et al. (2022) compare the empirical performance of an MNL model with product features and a sophisticated machine learning algorithm by conducting a field experiment on Alibaba's marketplaces, and this was adopted by Alibaba as the current practice. The assortment optimization is then solved based on the estimated choice model. The authors find that the MNL-based approach leads to a 28% increase in revenue per customer. Berbeglia et al. (2022) compare a wide range of discrete choice models and their predictive power on a number of synthetic and public real data sets. The authors find that, when the historical data are scarce, the exponential model (Alptekinoglu and Semple 2016) performs the best, followed by the mixed logit, rank-based (see, e.g., Farias et al. 2013), and nested logit (Train 2009) models. With a large data size, the Markov chain model (Blanchet et al. 2016) performs the best empirically, followed by the rank-based and exponential models. They also compare the computational time of the models.

2.2. Other Topics

As with discrete choice models, demand functions are an equally important class of models in RM, and they describe how demand reacts to the price(s) of the offered product(s). However, the estimation of demand functions is not a focus in the literature. This is partly because the estimation of most demand functions, such as the linear demand model, can usually be viewed as special cases of statistical learning problems, such as regression. Therefore, the literature typically focuses on downstream pricing optimization after estimating the demand models. We review papers in this area in Section 3.

Recently, some researchers have started focusing on causal inference and the endogeneity in observational data. For example, historical prices and market demand are observed and used in estimating the demand functions. If the prices are correlated with the demand through observed/unobserved covariates as the firm uses specific pricing policies in the data, then the eventual optimal price based on the estimated demand function can be distorted. This is a common issue in empirical studies. Bertsimas and Kallus (2023) quantify the error when the firm uses the past data to find the downstream optimal price, failing to handle the endogeneity. In Alley et al. (2022), the authors handle this issue in secondary ticket selling by constructing a semiparametric model that explicitly takes account of the endogeneity and using a recently developed tool in causal inference, double machine learning, to estimate the model. Biggs et al. (2021a) focus on the observational data when the firm sets personalized prices. The authors design loss functions, minimizing which directly yields the optimal prices instead of the demand functions. Li and Talluri (2020) propose a generalized method-of-moments

approach to estimate the MNL model with the observational choice data. The method is robust to the firm's policy of assortment selection in the data. Wang et al. (2022) combine the instrumental variable with random forest to obtain a nonparametric estimator that can identify unobserved confounding in observational data.

Bundle pricing is another common selling mechanism when customers may buy multiple items. The literature typically assumes that the customers' valuations of products (or their distributions) are given and focuses on the optimal or near-optimal bundle pricing schemes. There are a few papers studying the estimation of consumers' valuations from bundle sales data. Jedidi et al. (2003) use Bayesian inference to learn the distribution of valuations from the regions. Letham et al. (2014) approximate the region with boxes and derive closed-form estimators for the distribution parameters. Ma and Simchi-Levi (2022) propose a method to extract the linear demand function for individual products from the bundle sales data. Chen et al. (2022a) leverage the fact that a data point of bundle transactions (the choice of a bundle from a price menu) does not reveal the exact valuations, but rather a region in which they reside. They convert the estimation problem to one with region-censored observations.

3. Predict Then Optimize

RM ultimately studies the optimal decisions for firms, such as assortments, pricing, or promotions. Applying optimization tools to the estimated models from the data, the firms can make high-quality decisions supported by the predictive model. This predict-then-optimize framework is a natural extension to Section 2. The papers reviewed in this section typically include estimation and optimization as separate modules, which distinguishes them from those reviewed in Sections 4 and 5. They are usually motivated by real-world problems which, because of the scale or the special structure of the application, call for novel treatments. The paradigm is illustrated in Figure 3.

Building on work with Zara, a fast-fashion retailer, Caro and Gallien (2010, 2012) and Gallien et al. (2015) design forecasting models to predict the demand over a replenishment cycle or a sale season. The demand model is then fed to inventory (price) optimization to determine the optimal order quantities (clearance prices). The framework is designed for the specific application, taking into account many nuances of the business practice, including the age of a product (specific to the fast-


fashion industry), item clustering, and so on. Subsequent field experiments at Zara show the superb performance of their approaches compared with the existing method. Ferreira et al. (2016) propose a demand prediction and price optimization model for Rue La La, an online fashion retailer. Their approach based on regression trees overcomes the challenge of demand prediction for new products and improves the revenue significantly for the company. The tree-based prediction model is also used in Ban et al. (2019) to forecast the demand for new products with features and manage its procurement. Cohen et al. (2017) focus on the promotion optimization problem for the fast-moving consumer goods industry and incorporate business rules of the industry. After estimating the demand functions, the authors propose linear programming to approximately solve the problem. They work with Oracle Retail and apply the framework to a large-scale case study. Fisher et al. (2018) and Schlosser and Boissier (2018) consider the competitive market and propose frameworks to estimate the demand function with competitors' prices and show improved performance based on the estimation. Baardman et al. (2020) propose a model to forecast the customer trend and then plan the optimal promotion targeting policy. Boada-Collado and Martínez-de Albéniz (2020) propose an econometric model to quantify the impact of inventory levels on sales in a fashion retail setting. The authors use the estimated model to provide instructions on how to balance the inventory across products optimally. In Besbes et al. (2020), the authors develop a framework that estimates the demand and then sets prices for a large number of rotatable spare parts for an aircraft manufacturer. Xu et al. (2019) design a demand model and an estimation procedure for Major League Baseball ticket sales, which leads to significantly higher profits after optimizing the price from the model. Arslan et al. (2022) develop a framework to estimate the purchase behavior from customers in a sports ticket market with multiple sales channels, which is then used to optimize the price.

4. Online Learning

Online learning is a framework that integrates estimation/inference and decision making. The decision maker (e.g., the firm) starts with no or little knowledge of the model (e.g., the price elasticity of the market). The data (the historical price and demand) is typically not present either. The firm has to make decisions over time (e.g., setting prices), collect data to learn the model, and gradually approximate the actual optimal decision as if the ground-truth model were known from the beginning. It is a dynamic process as the collection of data, inference of the model and decision making are all interdependent and have to be conducted in a cyclic manner. The diagram illustrating the process is shown in Figure 4.

Figure 3. Predict Then Optimize



Figure 4. Online Learning

Many of the early works using online learning in RM applications are reviewed in den Boer (2015). In this survey, we focus on a number of new directions in this area that have been studied in the last few years.

4.1. Personalized Pricing

Personalization is the business practice by which a firm customizes the displayed products or offered promotions for individual customers. It is enabled by the increasing amount of data firms collect from their customers, including their personal information (often referred to as features), preferences, and how the latter is predicted by the former. For example, a bank can offer personalized rates to a car loaner. Personalized pricing models can also be applied to a setting in which a firm customizes prices for individual products (e.g., Rue La La offerings and Airbnb listings) based on the features of the products with the customers' valuations for different features possibly learned over time. Qiang and Bayati (2016) is one of the earliest papers that incorporates the covariate into the demand function. That is, the demand function is linear in the offered price and the covariate (customer features or market environment) present in a period. The decision maker needs to determine the price and learn the linear coefficients of the covariate simultaneously. The authors show that, under certain conditions, a greedy algorithm achieves $\log(T)$ regret, where T is the length of the learning horizon. A similar linear setting is studied in Ban and Keskin (2021) and Javanmard and Nazerzadeh (2019). In Ban and Keskin (2021), the authors focus on the sparsity structure of the covariate (only s out of d features in the covariate affect the demand function), and their algorithm, leveraging LASSO regularization, can detect the ambient features in the demand model and incur near-optimal regret that mainly depends on s . Sparsity is also considered in Javanmard and Nazerzadeh (2019); however, the algorithm is based on the maximum likelihood instead of the least squares because of the different demand models. Motivated by the same application, Cohen et al. (2020) investigate personalized dynamic pricing for customers whose valuations are linear functions of their features with the same coefficient among customers. However, the basic model does not have any randomness, and the binary outcome depends on whether the valuation is higher than the offered price. Therefore, the algorithm can be seen as a binary search in high dimensions.

There are multiple extensions to the personalized pricing problem with linear covariates. Chen and Gallego (2021) investigate the nonparametric formulation in which the demand model depends on the covariate in an unstructured way. Their algorithm is based on adaptive binning and achieves the optimal regret; however, the regret rate inevitably scales with the dimension of the covariate unlike the linear demand function studied in other papers. Bastani et al. (2022) consider a number of online learning experiments in which the linear coefficient in each experiment is generated from a common prior, and each experiment has the same horizon. This is motivated by the dynamic pricing of different products, in which the learned knowledge of one product may be transferred to others. Because of the increasing concern about the infringement of privacy when the firm collects consumer information or features, Tang et al. (2020) and Chen et al. (2022c) provide privacy-preserving learning algorithms when differential privacy is required. Nambiar et al. (2019) consider the model misspecification when the firm mistakenly uses a demand model linear in the features, whereas the ground-truth model can be nonlinear. Their proposed algorithm achieves the optimal regret against a clairvoyant who uses the optimal price of a linear model that best approximates the actual nonlinear one. Miao et al. (2022) tackle the challenge when many products have low sales. By clustering the sales data of products of similar demand patterns and learning their demand functions on the fly, they provide an algorithm that performs well theoretically and empirically. Shah et al. (2019) provide a semiparametric formulation in which the residual distribution is nonparametric. Wang et al. (2021) relax the independent and identically distributed assumption on the covariates.

4.2. Network Revenue Management

Network RM is originally proposed in Gallego and van Ryzin (1997). It studies the pricing problem of multiple products over a finite horizon when each product may consume various resources that are stored at the beginning of the horizon but cannot be replenished. Typical examples include the sales of airline itineraries (products), each of which may include multiple legs of flights (resources). The online learning of network RM is studied in Besbes and Zeevi (2012), in which the demand for the products under a price vector is initially unknown. Their base model studies a finite set of possible price vectors. The regret for continuous prices is higher because of the curse of dimensionality. The model with discrete prices is also studied in Badaniyuru et al. (2013) with a different asymptotic regime, and the optimal regret is obtained by their algorithm. Ferreira et al. (2018) apply Thompson sampling to the problem in Besbes and Zeevi (2012) and prove the optimal regret. Other papers tackle the problem with

continuous prices. Chen et al. (2019a) show that, under an infinite degree of smoothness, the optimal regret can be obtained. Chen and Gallego (2022) use a primal–dual algorithm that achieves the optimal regret under separable demand and a single resource constraint. Chen and Shi (2023) design an algorithm that achieves dimensionless regret, which does not depend on the number of products or resources. In a recent paper, Miao and Wang (2021) develop an algorithm to achieve the optimal regret $O(\sqrt{T})$ for continuous prices under a set of relatively standard assumptions.

4.3. Assortment Optimization

Unlike pricing, the optimal assortment features a discrete optimization problem: the firm chooses to offer a subset of products to customers in order to maximize the expected revenue. The customers typically follow a discrete choice model (see Section 2 for more introduction). When the firm doesn't have the information of the discrete choice model initially, online learning has to be used: the parameters of the discrete choice model are unknown, and the firm needs to learn them by offering different assortments over time. Online learning of assortment optimization is first studied in Rusmevichientong et al. (2010), Ulu et al. (2012), and Sauré and Zeevi (2013). Rusmevichientong et al. (2010) and Sauré and Zeevi (2013) focus on the MNL model and use the explore-then-commit algorithm (see Lattimore and Szepesvári 2020, chapter 6, for more details). That is, the estimation of the model parameters and the optimization based on the estimation are disintegrated, and the firm conducts both modules episodically. Integrated algorithms that have better empirical performances in general online learning problems, such as Thompson sampling and upper confidence bound (UCB), are applied to assortment optimization in Agrawal et al. (2017, 2019), respectively.

One of the main extensions to online learning for assortment optimization is personalization, that is, different customers may have different preferences, and it should be accommodated with personalized assortments. The primary framework to incorporate personalization is the linear MNL model, which adjusts the expected utilities of the products for preferences by a linear function of the contextual information, such as customer covariates. Kallus and Udell (2020) study a case in which both customers and products have high-dimensional features. For the efficient use of data, the authors impose a low-rank structure and leverage the matrix-norm regularization to recover the structure, which effectively lowers the regret. Cheung and Simchi-Levi (2017) and Oh and Iyengar (2019) apply Thompson sampling to the linear MNL model, focusing on Bayesian and worst case regret. Chen et al. (2020c) investigate the nonstationary case when the context information may shift over time. Instead of linearizing the features,

Bernstein et al. (2019) use nonparametric Bayesian methods to cluster customers on the fly and formulate the problem as a dynamic program.

Other than personalized assortment, there are other extensions. Chen et al. (2021b) study the online learning of the nested logit model instead of the MNL model. Miao and Chao (2021) study the online learning of a joint pricing and assortment problem based on Thompson sampling. Dong et al. (2020) consider switching costs when the firm changes the products in the assortment in online learning. In Feng et al. (2022), the objective is not to minimize the regret, but to learn the customer preference rankings as efficiently as possible by offering a sequence of assortments. In that sense, it is closer to the best arm identification problem (see Lattimore and Szepesvári 2020, chapter 33) or experimental design.

4.4. Other Topics

We discuss a few other topics that are studied using the online learning framework. One such topic is product ranking, which can be seen as a more flexible way to display products than assortments without orders discussed in Section 4.3. The firm can not only choose an assortment of products to display to the customers, but also their display orders. It is enabled by the prevalence of online platforms. The display order of the products, thus, has an impact on customers' behavior. The most widely used ranking model is the cascade model (Craswell et al. 2008). In the RM community, there are a number of papers recently tackling the ranking problem, usually with an angle from online learning. In Ferreira et al. (2022), the authors consider heterogeneous customers whose attention span and clicking probabilities are drawn from a distribution. As a result, the optimal off-line ranking is not decreasing in terms of popularity, which is the case for the cascade model. They show a greedy algorithm can guarantee a 1/2-approximation ratio and provide an online learning algorithm that can learn the optimal ranking algorithm with high probability. Gao et al. (2022) incorporate the pricing decision into the cascade model and provide a UCB-based learning algorithm. Cao and Sun (2019) and Chen et al. (2021a) investigate the cascade model with random attention spans to maximize the revenue. Their online learning algorithms are based on UCB and gradually learn the distribution of the attention spans and the clicking probabilities. Golrezaei et al. (2023) consider fake users that may corrupt the data and design algorithms to counter such adversaries.

Revenue management and pricing with reusable resources are also studied in the online learning framework, in which the capacity can be reused, such as in service systems (Jia et al. 2023). It is also a key feature of vehicle-sharing systems as vehicles completing trips can be reused (Banerjee et al. 2022, Benjaafar and Shen

2022). The online learning approach to such systems is studied in Benjaafar et al. (2023).

Another topic that is gaining traction is learning in auctions. Balseiro and Gur (2019) consider a budget-constrained advertiser participating in sequential auctions and learning the value of future auctions. The existence of competing bidders makes the analysis different from single-agent learning models. Kanoria and Nazerzadeh (2021) focus on the interaction between the auctioneer and the bidders. When the auctioneer attempts to learn the distribution of bidders' valuations, it may induce strategic behavior of the buyer and untruthful bids. They show that personalized reserve prices can be used to solve the problem when there are multiple bidders. Amin et al. (2014) and Golrezaei et al. (2021) study contextual second price auctions in which the bidders' valuations may depend on the context and the relationship is unknown to the auctioneer. The authors propose policies that are robust to strategic behavior and achieve sublinear regret. Ye et al. (2023) use online learning to handle the problem when the platform intends to learn the value of new ads within a short time.

Other extensions to the online learning framework in RM include learning algorithms assisted by an off-line data set (Bu et al. 2020), demand learning when customers have reference effects (den Boer and Keskin 2022), learning discontinuous demand functions (den Boer and Keskin 2020), nonstationary market environments (Chen et al. 2019b, 2020; Zhu and Zheng 2020), add-on discounts (Simchi-Levi et al. 2022), and competition (in particular, algorithmic collusion) (Hansen et al. 2021, Meylahn and den Boer 2022).

5. End-to-End Decision Making


End-to-end decision making is a new paradigm that has gained popularity in the last few years. The end-to-end framework deemphasizes the role of models in the process and tries to establish a direct link from data to decision, demonstrated in Figure 5. The term first appears in Donti et al. (2017), which discusses a few operational applications, such as inventory management. A popular formulation coined as "prescriptive analytics" is provided in Bertsimas and Kallus (2020). Bertsimas and Kallus (2020) study a generic optimization problem: consider covariate X , intermediate variable Y , decision z , and an objective function $c(z, Y)$. Given the historical data $\{(X_i, Y_i)\}_{i=1}^n$, the decision maker attempts to minimize the objective $c(z, Y)$ over z after observing a new covariate $X = x$, that is, $\min_z E[c(z, Y) | X = x]$. Although $c(z, Y)$ is known, the relationship between X and Y needs to be

learned from the data. This formulation inspired a number of methods and directions, such as Kallus and Mao (2023) and Elmachtoub and Grigas (2022). A similar framework specialized for the newsvendor problem is presented in Ban and Rudin (2019). This paradigm is also familiar to the optimization community and bears similarity to sample average approximations (Kleywegt et al. 2002) and data-driven distributionally robust optimization (Esfahani and Kuhn 2018).

In this section, we review recent papers using this paradigm in the context of RM. In a series of papers (Fu et al. 2015, Babaioff et al. 2018, Huang et al. 2018, Daskalakis and Zampetakis 2020, Allouah et al. 2022), the authors consider a seller maximizing revenue by setting prices when the demand function (the distribution function of consumers' valuations) is unknown. The seller only observes a number of samples from the unknown distribution and aims at mapping the samples to a price directly without estimating the demand function. The performance of the pricing policy is measured by the expected revenue for the worst case distribution among a class, whereas the expectation is taken over the random samples. The studies provide upper and lower bounds for the policies for a different number of samples (one to infinite) and classes of distributions (regular or monotone hazard rate). Chen et al. (2023) consider assortment pricing when a firm needs to set prices for a basket of products to maximize the expected revenue based on transaction data in which the offered prices and purchases of historical customers are observed. Instead of using the transaction data to fit a discrete choice model, the authors directly construct from the data a number of polytopes (one for each historical customer) into which their valuation vector falls and is consistent with the observed transactions. To set the prices, the seller attempts to maximize the revenue when the valuation vector of a new customer is the worst case within a valuation polytope that is drawn uniformly randomly from the past polytopes. In Biggs (2022), the framework is generalized to personalized pricing using a convex surrogate loss function without modeling the personalized demand function.

This end-to-end framework is used frequently in the study of auctions, which has become one of the major topics in RM. Biggs et al. (2021a, b) study personalized pricing policies that directly map customer features to the optimal personalized price without modeling the demand. In Cole and Roughgarden (2014), the authors study the standard revenue-maximizing single-item auction when the distributions of bidders' valuations are unknown. Samples drawn from the distributions are given. The auctioneers, thus, replace the unknown distributions with the empirical distributions of the samples. The sample complexity is analyzed when the optimal revenue can be approximated within a given precision. Derakhshan et al. (2022) study the data-driven reserve

Figure 5. End-to-End Decision Making



price optimization problem and use linear programming to provide a bound that improves those offered in the literature.

6. Experimental Design

Recently, experimental design has attracted the attention of many RM researchers, partly because of the wide adoption of A/B testing in the industry. Experimental design is concerned with the generation of historical data such that the model inference can be performed efficiently. For example, in the RM context, the firm may be interested in learning the attractiveness of recently launched products. When a customer queries a related keyword, the firm can display an assortment of products and record the purchase behavior. The firm is interested in the display of products such that the products' attractiveness can be learned as efficiently as possible. The paradigm can be represented by Figure 6. In particular, the design of experiments studies how to make decisions to affect the collection of data (the first arrow), and it is usually paired with the estimation of quantities of interest (the second arrow). Although experimental design is a classic topic in statistics (Dean et al. 2017), the recent literature focuses on problems emerging from large-scale A/B tests conducted on online platforms. As a result, the papers reviewed in this section may not be directly related to traditional RM topics such as pricing. It remains an exciting direction to explore the design of experiments in specific RM applications.

Bhat et al. (2020) is among the earliest papers that provide a theoretical framework and algorithms to conduct A/B tests in an online fashion in order to achieve statistical efficiency. In two-sided online platforms, Johari et al. (2022) use a mean-field model to study the interaction between the two sides of demand and supply: although the subjects of one side can be divided into treatment and control groups, they are not independent because of the interference through the other side. Therefore, the traditional estimation of the treatment effect has a bias. Although the design problem is not studied, Johari et al. (2022) analyze the bias caused by the interference. Switchback experiments are commonly used by platforms to mitigate such bias. Bojinov et al. (2023) study the optimal design of switchback experiments and provide statistics for inference. To debias the estimated treatment effect under interference, Farias et al. (2022) propose a new estimator inspired by Q-learning in reinforcement learning, which significantly reduces the bias and meanwhile has a

small variance. Zhao and Zhou (2022) study the design of experiments that sequentially allocate subjects with covariates in order to balance the covariate distribution between the treatment and control groups.

7. Concluding Remarks

The field of RM has been shifting from model-based to data-driven approaches in the last decade, thanks to the availability of large-scale data, emerging business models heavily reliant on data analytics, and industry practitioners with an analytics mindset. Despite the recent advances in data-driven RM, the area is far from fully explored, and there are many exciting directions to explore. We list some as follows.

7.1. Computationally Efficient Algorithms

Large-scale data sets introduce new challenges for classic algorithms. For example, the EM algorithm is widely adopted to handle missing variables and is used in RM applications when the lost demand is censored. Unfortunately, it does not scale to data sets of even moderate sizes. The design of algorithms in the large-data regime may have to take into account the calculated trade-off of computational efficiency and performance or theoretical guarantees.

7.2. New Sources and Forms of Data

The last decades have witnessed not only the increasing scale of data but also the increasing variety. The personal features of consumers, the browsing history, the display of the web page upon a query, and new media such as short videos or live streaming all introduce novel forms of data that cannot be tackled by traditional RM methods. How to make good use of them has become an important topic in RM and has inspired new models, frameworks, and algorithms.

7.3. Social Responsibility

Firms are raising awareness of the importance of their social responsibility. They are taking action to reduce the carbon footprint of their operations, improve the working conditions of their employees, and protect the privacy of their customers. Although RM focuses on revenue or profit traditionally, researchers are now exploring other objectives, such as fairness, privacy, and sustainability. Addressing these new objectives in a data-driven fashion remains a critical topic.

7.4. Impact of Artificial Intelligence (AI)

AI has become a transformative force in many industries in the last decade. In the past few years, reinforcement learning, large language models, computer vision, and generative models have far outpaced human beings' expectations of what AI can achieve. The technologies have not quite penetrated the industries in which

Figure 6. Experimental Design



RM scholars are typically interested, such as retailing, fast fashion, airline, and hospitality. As RM researchers, we hope to envision the future of RM with AI and explore the new opportunities and challenges that AI brings to RM.

Acknowledgments

The authors thank Saif Benjaafar (Editor-in-Chief) for his valuable comments and suggestions on the manuscript.

References

- Abdallah T, Vulcano G (2021) Demand estimation under the multinomial logit model from sales transaction data. *Manufacturing Service Oper. Management* 23(5):1196–1216.
- Agrawal S, Avadhanula V, Goyal V, Zeevi A (2017) Thompson sampling for the MNL-bandit. Kale S, Shamir O, eds. *Proc. 2017 Conf. Learn. Theory*, Proceedings of Machine Learning Research Series, vol. 65 (PMLR), 76–78.
- Agrawal S, Avadhanula V, Goyal V, Zeevi A (2019) MNL-bandit: A dynamic learning approach to assortment selection. *Oper. Res.* 67(5):1453–1485.
- Alley M, Biggs M, Hariss R, Herrmann C, Li ML, Perakis G (2022) Pricing for heterogeneous products: Analytics for ticket reselling. *Manufacturing Service Oper. Management* Forthcoming.
- Allouah A, Bahamou A, Besbes O (2022) Pricing with samples. *Oper. Res.* 70(2):1088–1104.
- Alptekinoglu A, Semple JH (2016) The exponential choice model: A new alternative for assortment and price optimization. *Oper. Res.* 64(1):79–93.
- Alptekinoglu A, Semple JH (2021) Heteroscedastic exponential choice. *Oper. Res.* 69(3):841–858.
- Amin K, Rostamizadeh A, Syed U (2014) Repeated contextual auctions with strategic buyers. Ghahramani Z, Welling M, Cortes C, Lawrence N, Weinberger KQ, eds. *Advances in Neural Information Processing Systems*, vol. 27 (Curran Associates, Inc., Red Hook, NY), 622–630. <https://proceedings.neurips.cc/paper/2014/file/66368270ffd51418ec58bd793f2d9b1b-Paper.pdf>.
- Amjad MJ, Shah D (2017) Censored demand estimation in retail. *Proc. ACM Measurement Anal. Comput. Systems* vol. 1, 1–28.
- Anupindi R, Dada M, Gupta S (1998) Estimation of consumer demand with stock-out based substitution: An application to vending machine products. *Marketing Sci.* 17(4):406–423.
- Aouad A, Désir A (2022) Representing random utility choice models with neural networks. Preprint, submitted July 26, <https://arxiv.org/abs/2207.12877>.
- Aouad A, Feldman J, Segev D (2018) The exponential choice model: Algorithmic frameworks for assortment optimization and data-driven estimation case studies. Preprint, submitted June 21, <https://dx.doi.org/10.2139/ssrn.3192068>.
- Aouad A, Elmachtoub AN, Ferreira KJ, McNellis R (2022) Market segmentation trees. *Manufacturing Service Oper. Management* Forthcoming.
- Arslan HA, Easley RF, Wang R, Yilmaz Ö (2022) Data-driven sports ticket pricing for multiple sales channels with heterogeneous customers. *Manufacturing Service Oper. Management* 24(2):1241–1260.
- Azadeh SS, Marcotte P, Savard G (2014) A taxonomy of demand uncensoring methods in revenue management. *J. Revenue Pricing Management* 13(6):440–456.
- Azadeh SS, Marcotte P, Savard G (2015) A non-parametric approach to demand forecasting in revenue management. *Comput. Oper. Res.* 63:23–31.
- Baardman L, Boroujeni SB, Cohen-Hillel T, Panchamgam K, Perakis G (2020) Detecting customer trends for optimal promotion targeting. *Manufacturing Service Oper. Management*, ePub ahead of print October 6, <https://doi.org/10.1287/msom.2020.0893>.
- Babaioff M, Gonczarowski YA, Mansour Y, Moran S (2018) Are two (samples) really better than one? *Proc. 2018 ACM Conf. Econom. Comput.* (Association for Computing Machinery, New York), 175.
- Badanidiyuru A, Kleinberg R, Slivkins A (2013) Bandits with knapsacks. *2013 IEEE 54th Annual Sympos. Foundations Comput. Sci. (IEEE)*, 207–216.
- Balseiro SR, Gur Y (2019) Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Sci.* 65(9):3952–3968.
- Ban G-Y, Keskin NB (2021) Personalized dynamic pricing with machine learning: High-dimensional features and heterogeneous elasticity. *Management Sci.* 67(9):5549–5568.
- Ban G-Y, Rudin C (2019) The big data newsvendor: Practical insights from machine learning. *Oper. Res.* 67(1):90–108.
- Ban G-Y, Gallien J, Mersereau AJ (2019) Dynamic procurement of new products with covariate information: The residual tree method. *Manufacturing Service Oper. Management* 21(4):798–815.
- Banerjee S, Freund D, Lykouris T (2022) Pricing and optimization in shared vehicle systems: An approximation framework. *Oper. Res.* 70(3):1783–1805.
- Bastani H, Simchi-Levi D, Zhu R (2022) Meta dynamic pricing: Transfer learning across experiments. *Management Sci.* 68(3):1865–1881.
- Benjaafar S, Shen X (2022) Pricing in on-demand (and one-way) vehicle sharing networks. Preprint, submitted January 5, <https://dx.doi.org/10.2139/ssrn.3998297>.
- Benjaafar S, Gao X, Shen X, Zhang H (2023) Online learning for pricing in on-demand vehicle sharing networks. Preprint, submitted February 1, <https://dx.doi.org/10.2139/ssrn.4344364>.
- Berbeglia G, Garassino A, Vulcano G (2022) A comparative empirical study of discrete choice models in retail operations. *Management Sci.* 68(6):4005–4023.
- Bernstein F, Modaresi S, Sauré D (2019) A dynamic clustering approach to data-driven assortment personalization. *Management Sci.* 65(5):2095–2115.
- Bertsimas D, Freund RM (2004) *Data, Models, and Decisions: The Fundamentals of Management Science* (Dynamic Ideas, Waltham, MA).
- Bertsimas D, Kallus N (2020) From predictive to prescriptive analytics. *Management Sci.* 66(3):1025–1044.
- Bertsimas D, Kallus N (2023) The power and limits of predictive approaches to observational data-driven optimization: The case of pricing. *INFORMS J. Optim.* Forthcoming.
- Besbes O, Zeevi A (2012) Blind network revenue management. *Oper. Res.* 60(6):1537–1550.
- Besbes O, Elmachtoub AN, Sun Y (2020) Pricing analytics for rotatable spare parts. *INFORMS J. Appl. Anal.* 50(5):313–324.
- Bhat N, Farias VF, Moallemi CC, Sinha D (2020) Near-optimal A-B testing. *Management Sci.* 66(10):4477–4495.
- Biggs M (2022) Convex loss functions for contextual pricing with observational posted-price data. Preprint, submitted February 16, <https://arxiv.org/abs/2202.10944>.
- Biggs M, Gao R, Sun W (2021a) Loss functions for discrete contextual pricing with observational data. Preprint, submitted November 18, <https://arxiv.org/abs/2111.09933>.
- Biggs M, Sun W, Ettl M (2021b) Model distillation for revenue optimization: Interpretable personalized pricing. Meila M, Zhang T, eds. *Proc. 38th Internat. Conf. Machine Learn.* Proceedings of Machine Learning Research Series, vol. 139 (PMLR), 946–956.
- Blanchet J, Gallego G, Goyal V (2016) A Markov chain approximation to choice modeling. *Oper. Res.* 64(4):886–905.
- Boada-Collado P, Martínez-de Albéniz V (2020) Estimating and optimizing the impact of inventory on consumer choices in a fashion retail setting. *Manufacturing Service Oper. Management* 22(3):582–597.

- Bojinov I, Simchi-Levi D, Zhao J (2023) Design and analysis of switch-back experiments. *Management Sci.* Forthcoming.
- Bu J, Simchi-Levi D, Xu Y (2020) Online pricing with offline data: Phase transition and inverse square law. Daumé III H, Aarti S, eds. *Proc. 37th Internat. Conf. Machine Learn.* Proceedings of Machine Learning Research Series, vol. 119 (PMLR), 1202–1210.
- Cai Z, Wang H, Talluri K, Li X (2022) Deep learning for choice modeling. Preprint, submitted August 19, <https://arxiv.org/abs/2208.09325>.
- Cao J, Sun W (2019) Dynamic learning of sequential choice bandit problem under marketing fatigue. *Proc. Conf. AAAI Artificial Intelligence* (AAAI Press, Palo Alto, CA), 33:3264–3271.
- Caro F, Gallien J (2010) Inventory management of a fast-fashion retail network. *Oper. Res.* 58(2):257–273.
- Caro F, Gallien J (2012) Clearance pricing optimization for a fast-fashion retailer. *Oper. Res.* 60(6):1404–1422.
- Chen N, Gallego G (2021) Nonparametric pricing analytics with customer covariates. *Oper. Res.* 69(3):974–984.
- Chen N, Gallego G (2022) A primal-dual learning algorithm for personalized dynamic pricing with an inventory constraint. *Math. Oper. Res.* 47(4):2585–2613.
- Chen N, Farajollahzadeh S, Wang G (2022a) Learning consumer preferences from bundle sales data. Preprint, submitted September 11, <https://arxiv.org/abs/2209.04942>.
- Chen N, Gallego G, Tang Z (2019c) The use of binary choice forests to model and estimate discrete choices. Preprint, submitted August 3, 2019, <https://arxiv.org/abs/1908.01109>.
- Chen N, Li A, Yang S (2021a) Revenue maximization and learning in products ranking. *Proc. 22nd ACM Conf. Econom. Comput.*, 316–317.
- Chen N, Wang C, Wang L (2020) Learning and optimization with seasonal patterns. Preprint, submitted May 16, <https://arxiv.org/abs/2005.08088>.
- Chen N, Cire A, Hu M, Lagzi S (2023) Model-free assortment pricing with transaction data. *Management Sci.* Forthcoming.
- Chen Q, Jasin S, Duenyas I (2019a) Nonparametric self-adjusting control for joint learning and optimization of multiproduct pricing with finite resource capacity. *Math. Oper. Res.* 44(2):601–631.
- Chen X, Jasin S, Shi C, eds. (2022b) *The Elements of Joint Learning and Optimization in Operations Management*, 1st ed., Springer Series in Supply Chain Management, vol. 18 (Springer International Publishing).
- Chen X, Simchi-Levi D, Wang Y (2022c) Privacy-preserving dynamic personalized pricing with demand learning. *Management Sci.* 68(7):4878–4898.
- Chen X, Wang Y, Zhou Y (2020c) Dynamic assortment optimization with changing contextual information. *J. Machine Learn. Res.* 21(1):8918–8961.
- Chen X, Owen Z, Pixton C, Simchi-Levi D (2022d) A statistical learning approach to personalization in revenue management. *Management Sci.* 68(3):1923–1937.
- Chen X, Shi C, Wang Y, Zhou Y (2021b) Dynamic assortment planning under nested logit models. *Production Oper. Management* 30(1):85–102.
- Chen Y, Shi C (2023) Network revenue management with online inverse batch gradient descent method. *Production Oper. Management* Forthcoming.
- Chen Y, Wen Z, Xie Y (2019b) Dynamic pricing in an evolving and unknown marketplace. Preprint, submitted June 6, <https://dx.doi.org/10.2139/ssrn.3382957>.
- Chen Y-C, Mišić VV (2022) Decision forest: A nonparametric approach to modeling irrational choice. *Management Sci.* 68(10):7090–7111.
- Cheung WC, Simchi-Levi D (2017) Thompson sampling for online personalized assortment optimization problems with multinomial logit choice models. Preprint, submitted November 27, <https://dx.doi.org/10.2139/ssrn.3075658>.
- Cho S, Ferguson M, Pekgün P, Vakhutinsky A (2023) Estimating personalized demand with unobserved no-purchases using a mixture model: An application in the hotel industry. *Manufacturing Service Oper. Management* Forthcoming.
- Cohen MC, Lobel I, Leme RP (2020) Feature-based dynamic pricing. *Management Sci.* 66(11):4921–4943.
- Cohen MC, Leung N-HZ, Panchamgam K, Perakis G, Smith A (2017) The impact of linear optimization on promotion planning. *Oper. Res.* 65(2):446–468.
- Cole R, Roughgarden T (2014) The sample complexity of revenue maximization. *Proc. 46th Annual ACM Sympos. Theory Comput.* (Association for Computing Machinery, New York), 243–252.
- Craswell N, Zoeter O, Taylor M, Ramsey B (2008) An experimental comparison of click position-bias models. *Proc. 2008 Internat. Conf. Web Search Data Mining*, (Association for Computing Machinery, New York), 87–94.
- Daskalakis C, Zampetakis M (2020) More revenue from two samples via factor revealing SDPs. *Proc. 21st ACM Conf. Econom. Comput.* (Association for Computing Machinery, New York), 257–272.
- Dean A, Voss D, Draguljić D (2017) *Design and Analysis of Experiments*. Springer Texts in Statistics (Springer International Publishing).
- Dempster AP, Laird NM, Rubin DB (1977) Maximum likelihood from incomplete data via the EM algorithm. *J. Roy. Statist. Soc. B* 39(1):1–22.
- den Boer AV (2015) Dynamic pricing and learning: Historical origins, current research, and new directions. *Surveys Oper. Res. Management Sci.* 20(1):1–18.
- den Boer AV, Keskin NB (2020) Discontinuous demand functions: Estimation and pricing. *Management Sci.* 66(10):4516–4534.
- den Boer AV, Keskin NB (2022) Dynamic pricing with demand learning and reference effects. *Management Sci.* 68(10):7112–7130.
- Derakhshan M, Golrezaei N, Leme RP (2022) Linear program-based approximation for personalized reserve prices. *Management Sci.* 68(3):1849–1864.
- Dong K, Li Y, Zhang Q, Zhou Y (2020) Multinomial logit bandit with low switching cost. Daumé H III, Singh A, eds. *Proc. 37th Internat. Conf. Machine Learn.*, vol. 119 (PMLR), 2607–2615.
- Donti PL, Amos B, Kolter JZ (2017) Task-based end-to-end model learning in stochastic optimization. *Proc. 31st Internat. Conf. Neural Inform. Processing Systems*, vol. 30 (NIPS) (Curran Associates Inc., Red Hook, NY), 5490–5500.
- Elmachtoub AN, Grigas P (2022) Smart “predict, then optimize.” *Management Sci.* 68(1):9–26.
- Esfahani PM, Kuhn D (2018) Data-driven distributionally robust optimization using the wasserstein metric: Performance guarantees and tractable reformulations. *Math. Programming* 171(1):115–166.
- Farias V, Jagabathula S, Shah D (2020) Inferring sparse preference lists from partial information. *Stochastic Systems* 10(4):335–360.
- Farias VF, Jagabathula S, Shah D (2013) A nonparametric approach to modeling choice with limited data. *Management Sci.* 59(2):305–322.
- Farias VF, Li AA, Peng T, Zheng AT (2022) Markovian interference in experiments. Preprint, submitted June 6, <https://arxiv.org/abs/2206.02371>.
- Feldman J, Zhang DJ, Liu X, Zhang N (2022) Customer choice models vs. machine learning: Finding optimal product displays on Alibaba. *Oper. Res.* 70(1):309–328.
- Feng Y, Caldentey R, Ryan CT (2022) Robust learning of consumer preferences. *Oper. Res.* 70(2):918–962.
- Ferreira KJ, Bin HAL, Simchi-Levi D (2016) Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing Service Oper. Management* 18(1):69–88.
- Ferreira KJ, Parthasarathy S, Sekar S (2022) Learning to rank an assortment of products. *Management Sci.* 68(3):1828–1848.
- Ferreira KJ, Simchi-Levi D, Wang H (2018) Online network revenue management using Thompson sampling. *Oper. Res.* 66(6):1586–1602.
- Fisher M, Gallino S, Li J (2018) Competition-based dynamic pricing in online retailing: A methodology validated with field experiments. *Management Sci.* 64(6):2496–2514.

- Fu H, Immorlica N, Lucier B, Strack P (2015) Randomization beats second price as a prior-independent auction. *Proc. 16th ACM Conf. Econom. Comput.* (Association for Computing Machinery, New York), 323.
- Gallego G, van Ryzin G (1997) A multiproduct dynamic pricing problem and its applications to network yield management. *Oper. Res.* 45(1):24–41.
- Gallien J, Mersereau AJ, Garro A, Dapena Mora A, Nóvoa Vidal M (2015) Initial shipment decisions for new products at Zara. *Oper. Res.* 63(2):269–286.
- Gao X, Jasin S, Najafi S, Zhang H (2022) Joint learning and optimization for multi-product pricing (and ranking) under a general cascade click model. *Management Sci.* 68(10):7362–7382.
- Golrezaei N, Javanmard A, Mirrokni V (2021) Dynamic incentive-aware learning: Robust pricing in contextual auctions. *Oper. Res.* 69(1):297–314.
- Golrezaei N, Manshadi V, Schneider J, Sekar S (2023) Learning product rankings robust to fake users. *Oper. Res.* Forthcoming.
- Gupta A, Hsu D (2017) Parameter identification in Markov chain choice models. Hanneke S, Reyzin L, eds. *Pro. 28th Internat. Conf. Algorithmic Learn. Theory Proc. Machine Learn. Res.*, vol. 76 (PMLR), 330–340.
- Hansen KT, Misra K, Pai MM (2021) Frontiers: Algorithmic collusion: Supra-competitive prices via independent algorithms. *Marketing Sci.* 40(1):1–12.
- Ho-Nguyen N, KilinKç-Karzan F (2021) Dynamic data-driven estimation of nonparametric choice models. *Oper. Res.* 69(4):1228–1239.
- Huang Z, Mansour Y, Roughgarden T (2018) Making the most of your samples. *SIAM J. Comput.* 47(3):651–674.
- Jagabathula S, Rusmevichientong P (2017) A nonparametric joint assortment and price choice model. *Management Sci.* 63(9):3128–3145.
- Jagabathula S, Vulcano G (2018) A partial-order-based model to estimate individual preferences using panel data. *Management Sci.* 64(4):1609–1628.
- Javanmard A, Nazerzadeh H (2019) Dynamic pricing in high dimensions. *J. Machine Learn. Res.* 20(1):315–363.
- Jedidi K, Jagpal S, Manchanda P (2003) Measuring heterogeneous reservation prices for product bundles. *Marketing Sci.* 22(1):107–130.
- Jena SD, Lodi A, Palmer H, Sole C (2020) A partially ranked choice model for large-scale data-driven assortment optimization. *INFORMS J. Optim.* 2(4):297–319.
- Jia H, Shi C, Shen S (2023) Online learning and pricing for service systems with reusable resources. *Oper. Res.* Forthcoming.
- Johari R, Li H, Liskovich I, Weintraub GY (2022) Experimental design in two-sided platforms: An analysis of bias. *Management Sci.* 68(10):7069–7089.
- Kallus N, Mao X (2023) Stochastic optimization forests. *Management Sci.* Forthcoming.
- Kallus N, Udell M (2020) Dynamic assortment personalization in high dimensions. *Oper. Res.* 68(4):1020–1037.
- Kanoria Y, Nazerzadeh H (2021) Incentive-compatible learning of reserve prices for repeated auctions. *Oper. Res.* 69(2):509–524.
- Kleywegt AJ, Shapiro A, Homem-de Mello T (2002) The sample average approximation method for stochastic discrete optimization. *SIAM J. Optim.* 12(2):479–502.
- Kök AG, Fisher ML (2007) Demand estimation and assortment optimization under substitution: Methodology and application. *Oper. Res.* 55(6):1001–1021.
- Lattimore T, Szepesvári C (2020) *Bandit Algorithms* (Cambridge University Press, Cambridge, UK).
- Letham B, Sun W, Sheopuri A (2014) Latent variable copula inference for bundle pricing from retail transaction data. Xing EP, Jebara T, eds. *Proc. 31st Internat. Conf. Machine Learn.* Proceedings of Machine Learning Research Series, vol. 32 (PMLR), 217–225.
- Li A, Talluri K (2020) Estimating demand with unobserved no-purchases on revenue-managed data. Preprint, submitted March 7, <https://dx.doi.org/10.2139/ssrn.3525773>.
- Ma W, Simchi-Levi D (2022) Constructing demand curves from a single observation of bundle sales. *Web Internet Econom. 18th Internat. Conf.*, (Springer-Verlag, Berlin, Heidelberg), 150–166.
- Meylahn JM, den Boer AV (2022) Learning to collude in a pricing duopoly. *Manufacturing Service Oper. Management* 24(5):2577–2594.
- Miao S, Chao X (2021) Dynamic joint assortment and pricing optimization with demand learning. *Manufacturing Service Oper. Management* 23(2):525–545.
- Miao S, Wang Y (2021) Network revenue management with non-parametric demand learning: \sqrt{T} -regret and polynomial dimension dependency. Preprint, submitted October 25, <https://dx.doi.org/10.2139/ssrn.3948140>.
- Miao S, Chen X, Chao X, Liu J, Zhang Y (2022) Context-based dynamic pricing with online clustering. *Production Oper. Management* 31(9):3559–3575.
- Mishra VK, Natarajan K, Padmanabhan D, Teo C-P, Li X (2014) On theoretical and empirical aspects of marginal distribution choice models. *Management Sci.* 60(6):1511–1531.
- Musalem A, Olivares M, Bradlow ET, Terwiesch C, Corsten D (2010) Structural estimation of the effect of out-of-stocks. *Management Sci.* 56(7):1180–1197.
- Nambiar M, Simchi-Levi D, Wang H (2019) Dynamic learning and pricing with model misspecification. *Management Sci.* 65(11):4980–5000.
- Natarajan K, Song M, Teo C-P (2009) Persistency model and its applications in choice modeling. *Management Sci.* 55(3):453–469.
- Newman JP, Ferguson ME, Garrow LA, Jacobs TL (2014) Estimation of choice-based models using sales data from a single firm. *Manufacturing Service Oper. Management* 16(2):184–197.
- Oh M-H, Iyengar G (2019) Thompson sampling for multinomial logit contextual bandits. Wallach H, Larochelle H, Beygelzimer A, d'Alché Buc F, Fox E, Garnett R, eds. *Proc. 33rd Internat. Conf. Neural Inform. Processing Systems* (Curran Associates, Inc., Red Hook, NY).
- Qiang S, Bayati M (2016) Dynamic pricing with demand covariates. Preprint, submitted April 25, <https://arxiv.org/abs/1604.07463>.
- Ragain S, Ugander J (2016) Pairwise choice Markov chains. Lee D, Sugiyama M, Luxburg U, Guyon I, Garnett R, eds. *Proc. 30th Internat. Conf. Neural Inform. Processing Systems* (Curran Associates, Inc., Red Hook, NY), 3206–3214.
- Ratliff RM, Venkateshwara Rao B, Narayan CP, Yellepeddi K (2008) A multi-flight recapture heuristic for estimating unconstrained demand from airline bookings. *J. Revenue Pricing Management* 7(2):153–171.
- Rusmevichientong P, Shen ZJM, Shmoys DB (2010) Dynamic assortment optimization with a multinomial logit choice model and capacity constraint. *Oper. Res.* 58(6):1666–1680.
- Sauré D, Zeevi A (2013) Optimal dynamic assortment planning with demand learning. *Manufacturing Service Oper. Management* 15(3):387–404.
- Schlosser R, Boissier M (2018) Dynamic pricing under competition on online marketplaces: A data-driven approach. *Proc. 24th ACM SIGKDD Internat. Conf. Knowledge Discovery Data Mining* (Association for Computing Machinery, New York), 705–714.
- Shah V, Johari R, Blanchet J (2019) Semi-parametric dynamic contextual pricing. Wallach H, Larochelle H, Beygelzimer A, d'Alché Buc F, Fox E, Garnett R, eds. *Adv. Neural Inform. Processing Systems*, vol. 32 (Curran Associates, Inc.).
- Simchi-Levi D (2014) OM forum—OM research: From problem-driven to data-driven research. *Manufacturing Service Oper. Management* 16(1):2–10.
- Simchi-Levi D, Sun R, Zhang H (2022) Online learning and optimization for revenue management problems with add-on discounts. *Management Sci.* 68(10):7402–7421.
- Şimşek AS, Topaloglu H (2018) An expectation-maximization algorithm to estimate the parameters of the Markov chain choice model. *Oper. Res.* 66(3):748–760.
- Subramanian S, Harsha P (2021) Demand modeling in the presence of unobserved lost sales. *Management Sci.* 67(6):3803–3833.

- Tang W, Ho C-J, Liu Y (2020) Differentially private contextual dynamic pricing. *Proc. 19th Internat. Conf. Autonomous Agents MultiAgent Systems* (International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC), 1368–1376.
- Train KE (2009) *Discrete Choice Methods with Simulation*, 2nd ed. (Cambridge University Press, Cambridge, UK).
- Ulu C, Honhon D, Alptekinoğlu A (2012) Learning consumer tastes through dynamic assortments. *Oper. Res.* 60(4):833–849.
- van Ryzin G, Vulcano G (2015) A market discovery algorithm to estimate a general class of nonparametric choice models. *Management Sci.* 61(2):281–300.
- van Ryzin G, Vulcano G (2017) An expectation-maximization method to estimate a rank-based choice model of demand. *Oper. Res.* 65(2):396–407.
- Vulcano G, van Ryzin G, Chaar W (2010) OM practice—Choice-based revenue management: An empirical study of estimation and optimization. *Manufacturing Service Oper. Management* 12(3):371–392.
- Vulcano G, van Ryzin G, Ratliff R (2012) Estimating primary demand for substitutable products from sales transaction data. *Oper. Res.* 60(2):313–334.
- Wan M, Huang Y, Zhao L, Deng T, Fransoo JC (2018) Demand estimation under multi-store multi-product substitution in high density traditional retail. *Eur. J. Oper. Res.* 266(1):99–111.
- Wang G, Li J, Hopp WJ (2022) An instrumental variable forest approach for detecting heterogeneous treatment effects in observational studies. *Management Sci.* 68(5):3399–3418.
- Wang H, Talluri K, Li X (2021) On dynamic pricing with covariates. Preprint, submitted December 25, <https://arxiv.org/abs/2112.13254>.
- Wang R (2021) Technical note—Consumer choice and market expansion: Modeling, optimization, and estimation. *Oper. Res.* 69(4):1044–1056.
- Xu JJ, Fader PS, Veeraraghavan S (2019) Designing and evaluating dynamic pricing policies for major league baseball tickets. *Manufacturing Service Oper. Management* 21(1):121–138.
- Ye Z, Zhang D, Zhang H, Zhang RP, Chen X, Xu Z (2023) Cold start to improve market thickness on online advertising platforms: Data-driven algorithms and field experiments. *Management Sci.* Forthcoming.
- Zhao J, Zhou Z (2022) Pigeonhole design: Balancing sequential experiments from an online matching perspective. Preprint, submitted January 30, <https://arxiv.org/abs/2201.12936>.
- Zhu F, Zheng Z (2020) When demands evolve larger and noisier: Learning and earning in a growing environment. *Proc. 37th Internat. Conf. Machine Learn.* Proceedings of Machine Learning Research Series, vol. 119 (PMLR), 11629–11638.