



Operations Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

In This Issue

To cite this article:

(2021) In This Issue. Operations Research 69(4):ii-iv. <https://doi.org/10.1287/opre.2021.2174>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2021, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

In This Issue

<https://doi.org/10.1287/opre.2021.2174>

Copyright: © 2021 INFORMS

Preservation Results for Proving Additively Convex Value Functions for High-Dimensional Stochastic Optimization Problems

Additive convexity is a key property in avoiding the curse of dimensionality for computing the optimal policy for high-dimensional stochastic optimization problems. In “Technical Note—Preservation of Additive Convexity and Its Applications in Stochastic Optimization Problems,” Gong and Wang establish two preservation results of additive convexity for a class of optimal transformation problems and a class of optimal disposal problems. For both classes of problems, there are multiple resources; their results show that if these resources have different priorities to be transformed/disposed under the optimal policy, then the additive convexity and bounded monotonicity of the objective function are preserved to the value function after optimization. They demonstrate the applications of their results to several high-dimensional stochastic optimization problems in operations management.

How Do Customers React to Inventory Pooling?

Retailers have increasingly pursued initiatives to combine inventory located throughout their enterprise into a single stock from which customers anywhere in their distribution network may purchase. Also known as inventory pooling, this practice is well known to generate operational value by reducing inventory requirements and stock-outs. In “Inventory Integration with Rational Consumers,” Aflaki and Swinney examine a different consequence of pooling: how it influences customer purchasing behavior. They show that integration can lead to behavioral consequences resulting from changing customer purchasing incentives, especially for seasonal goods that are sold in end-of-season clearance sales. These behavioral effects can be negative, and can even outweigh the operational benefits of pooling, if pooling leads to an increase in clearance sale inventory availability that encourages more customers to wait for discounts before buying. Specific conditions that lead to negative (or positive) behavioral value of integration are discussed. Overall, the results illustrate that the ways customers react to inventory pooling can be just as important as its operational consequences.

How to Expand Market Size in Product Assortment Management and Pricing?

The growth of market size is crucially important to firms, although researchers often assume that market size is constant in assortment and pricing management. In “Technical Note—Consumer Choice and Market Expansion: Modeling, Optimization, and Estimation,” Wang develops a model that incorporates the market expansion effects into discrete consumer choice models and investigates various operations problems. Market size, measured by the number of people who are interested in the products from the same category, is largely influenced by firms’ operations strategy, and it also affects

assortment planning and pricing decisions. Failure to account for market expansion effects may lead to substantial losses in demand estimation and operations management. Based on real data, this paper uses an alternating-optimization expectation-maximization method that separates the estimation of consumer choice behavior and market expansion effects to calibrate the new model. The end-to-end solution approach on modeling, operations, and estimation is readily applicable in real business.

Solving Inventory Models with Service-Level Constraints by Using a Deterministic Approximation

Much of the inventory literature in the past has focused on cost-based models because of analytical tractability rather than accurate characterization of reality. However, the cost of unmet demand is often difficult to quantify in practice, and hence, service level is typically used as a more direct metric for evaluating the performance of inventory replenishment policies. An ideal policy should satisfy the service-level requirement while minimizing the total cost over the planning horizon. In “On a Deterministic Approximation of Inventory Systems with Sequential Service-Level Constraints,” Wei, Jasin, and Xin consider two fundamental inventory models (with backorder and lost sales) with sequential service-level constraints, and study the performance of an order-up-to policy whose parameters are determined using the optimal solution of a deterministic approximation of a backorder inventory system. The authors prove that the proposed heuristic is asymptotically optimal for both the backorder and lost-sales systems in the setting with a high service-level requirement. The results give credence to the use of deterministic approximations for solving complex inventory problems in practice, especially for applications in which the targeted service level is sufficiently high.

Heavy-Tail Optimality for Free

Distributionally robust optimization is increasingly becoming a popular methodology to deal with uncertainty in optimization problems. Although the methodology optimizes for the worst-case distribution, a better understanding of the prescriptions from the models important from a practical perspective. In “On the Heavy-Tail Behavior of the Distributionally Robust Newsvendor,” Das, Dhara, and Natarajan provide a novel analysis for the problem in the newsvendor setting with moment information. They show that the distributionally robust newsvendor by planning for the worst possible demand distribution with moment information will remain optimal if the true demand distribution is heavy-tailed. The prescribed optimal solution has a *heavy-tail optimality* property for free.

The Mean-Risk Problem: A Novel Approach to Time Consistency

When dealing with dynamic optimization problems, time consistency is a desirable property as it allows one to solve the problem efficiently through a backward recursion. The mean-

risk problem is known to be time inconsistent when considered in its scalarized form. However, when left in its original bi-objective form, it turns out to satisfy a more general time consistency property that seems better suited to a vector optimization problem. In “Time Consistency of the Mean-Risk Problem,” Kováčová and Rudloff introduce a set-valued version of the famous Bellman principle and show that the bi-objective mean-risk problem does satisfy it. Then, the upper image, a set that contains the efficient frontier on its boundary, recurses backward in time. The authors present conditions under which this recursion can be exploited directly to compute a solution in the spirit of dynamic programming. This opens the door for a new branch in mathematics: dynamic multivariate programming.

Data-Driven Transit Network Design at Scale

Public transit remains the most efficient way to service a densely-packed commuter population. However, increased competition in the transportation space, physical distancing during the COVID-19 pandemic, and tight budget constraints have contributed to declining ridership. Transit authorities around the United States are interested in re-designing bus networks to improve ridership. In “Data-Driven Transit Network Design at Scale,” Bertsimas, Ng, and Yan propose computational methods to support such re-design efforts, and show that their methods are both tractable and effective in a case study in Boston.

Replicating Portfolios for Economic Capital Calculations

Financial services firms, such as banks and insurance companies, are expected by customers, regulators, and capital providers to hold a sufficient amount of capital to absorb the financial losses they may incur in times of financial stress. This capital provides security and protection for customers owed money by such firms, including, for example, those who have taken out life insurance with a life insurance provider. In “A Large-Scale Optimization Model for Replicating Portfolios in the Life Insurance Industry,” Adelman, Fernandez-Arjona, Mayer, and Schmedders describe the replicating portfolio (RP) model used to approximate life insurance liabilities in a large global insurance company active in more than 215 countries. The model uses cash flow matching to calculate a portfolio of financial derivatives that replicates the behavior of the company’s life insurance liabilities with the accuracy necessary to enable the calculation of economic capital. The authors also discuss the critical role played by RPs in the risk-management process of the insurance company.

How to Sell Bundle Sizes?

Can you sell multiple items by providing only prices for different sizes of bundles rather than the different possible combinations of them? In “Large-Scale Bundle-Size Pricing: A Theoretical Analysis,” Abdallah, Asadpour, and Reed provide a framework for understanding “bundle-size pricing” (BSP) where only a menu of bundle sizes and their corresponding prices are offered. Although BSP is commonly used across several industries, little is known about the optimal BSP policy in terms of sizes and prices, along with the theoretical properties of its profit. The authors provide a simple and tractable theoretical framework to analyze the large-scale BSP problem where a multiproduct firm is selling a large number of products. They characterize the circumstances under which such policies perform well by studying the effect of various factors such as marginal cost or customers’ budget on the performance of BSP and identify possible causes of its inefficiency.

Improving the Chilean College Admissions System

In “Improving the Chilean College Admissions System,” Rios, Larroucau, Parra, and Cominetti describe the design and implementation of a new system to solve the Chilean college admissions problem. The authors develop an algorithm that (i) obtains all applicant/program pairs that can be part of a stable allocation when preferences are not strict and when all students tied in the last seat of a program (if any) must be allocated and (ii) efficiently incorporates affirmative action, which is part of the system to correct the inefficiencies that arise from having double-assigned students. By unifying the regular admission with affirmative action, the solution proposed and implemented by the authors has improved the allocation of approximately 2.5% of students assigned every year since 2016, helping to improve the overall efficiency of the system.

Mallows-Smoothed Distribution over Rankings Approach for Modeling Choice

Assortment optimization is an important problem arising in many applications. Key to this problem is a demand model able to specify the expected demand in response to a given offer set. In “Mallows-Smoothed Distribution over Rankings Approach for Modeling Choice,” Désir, Goyal, Jagabathula, and Segev study the smoothing of the existing sparse rank-based choice model using Mallows kernels. Through a detailed case study, the authors show that smoothing can significantly improve the parsimony of sparse rank-based models. Furthermore, despite the exponential support size of the Mallows distribution, an efficient procedure to compute any choice probabilities is developed. This procedure is leveraged to formulate the assortment-optimization problem as a compact mixed integer program, leading to an efficient approach for solving problem instances of practical nature and scale. Finally, to complement this MIP formulation, several near-optimal algorithms are designed by unraveling various structural properties of the Mallows distribution.

Dynamic Data-Driven Estimation of Nonparametric Choice Models

Choice models are prevalent in several application areas, and their nonparametric estimation was introduced to alleviate unreasonable assumptions in traditional parametric models. Existing literature focuses only on the static observational setting where all of the observations are given up front and lacks algorithms that provide explicit convergence rate guarantees or an a priori analysis for the model accuracy versus sparsity trade-off on the actual estimated model returned. In contrast, in “Technical Note—Dynamic Data-Driven Estimation of Nonparametric Choice Models,” Ho-Nguyen and Kılınç-Karzan focus on estimating a nonparametric choice model from observational data in a dynamic setting, where observations are obtained over time. They show that this estimation problem can be cast as a convex-concave saddle point joint estimation and optimization problem and provide an online convex optimization-based primal-dual framework for deriving algorithms for it. By tailoring this framework carefully to the choice model estimation problem, the authors provide low-cost algorithms that come with provable convergence guarantees, explicit theoretical bounds on the sparsity of the estimated model, and a superior empirical performance.

An OR Look Into Zonal Integration of Electricity Markets

Interconnected electricity systems can integrate their forward markets at multiple levels. Zonal integration represents an

intermediate level, where each zonal system operator retains control over its zone, and trades with neighboring zones are limited only in total volume. This type of integration is currently being championed in Europe, but it has emerged in several regions of the globe. A crucial challenge in studying zonal integration is the convoluted, distinct, and discretionary nature of the existing calculation methodologies for limits on interzonal trades. In “Transmission Capacity Allocation in Zonal Electricity Markets,” Aravena, Lété, Papavasiliou, and Smeers focus on the European case and study the regulatory requirements to be met by calculation methodologies, finding that these requirements imply a unique mathematical definition of the limits on interzonal trades. The authors use this definition to simulate and analyze the performance of zonal integration, considering the equivalent of 100 years of operating conditions.

A Framework for Studying Scarce Resource Allocation

When allocating scarce resources such as organs or public housing units, the policy maker needs to carefully balance two conflicting objectives: maximizing the reward from matching and minimizing inequity across different types of candidates. The authors consider an implementable class of policies that ranks the candidates by the sum of their waiting score and matching score. Similar policies had been proposed for allocating deceased-donor kidneys to patients on the transplant waitlist. In “A Fluid Model for One-Sided Bipartite Matching Queues with Match-Dependent Rewards,” Ding, McCormick, and Nagarajan provide a modeling framework to characterize the waitlist system under this type of ranking policy. This framework allows the policy maker to predict and compare the allocation outcome under different ranking policies. When the efficiency and fairness measurements take certain forms, the authors derive a closed-form scoring formula that optimizes the outcome of the system in the long run.

A New Phase Methodology for Analyzing Partially Observable Semi-Markovian Failing Systems

In “Optimal Control of Partially Observable Semi-Markovian Failing Systems: An Analysis Using a Phase Methodology,” Khaleghi and Kim study a maintenance control problem as a partially observable semi-Markov decision process (POSMDP), a problem class that is typically computationally intractable and not amenable to structural analysis. The authors develop a new

approach based on a phase methodology where the idea is to view the intractable POSMDP as the limiting problem of a sequence of tractable POMDPs. They show that the optimal control policy can be represented as a control limit policy which monitors the estimated conditional reliability at each decision epoch, and, by exploiting this structure, an efficient computational approach to solve for the optimal control limit and corresponding optimal value is developed.

Debt-Constrained Asset Selling

Financing asset acquisition by debt and selling the asset to repay the debt is a common practice in many industries. In “Asset Selling Under Debt Obligations,” Ahn, Wang, and Wu study the problem of selling a divisible asset under debt obligations with and without selling capacity constraint. The authors show that, in the presence of debt, the optimal asset-selling policy must take into account two opposing forces: an incentive to sell part of the asset early to secure debt payment and an incentive to delay selling the asset to capture revenue potential under limited liability. The seller delays or expedites selling the asset, depending on the relative strengths of these two forces, which are attenuated by the selling capacity constraint. The paper also establishes the condition under which selling a divisible asset with capacity constraint is equivalent to selling multiple nondivisible units, and it applies the model to the case of selling natural gas.

Queueing Across Multiple Limit Order Book Markets

In modern equity markets, participants have a choice of many exchanges at which to trade, each of which typically operates as an electronic limit order book. In “Queueing Dynamics and State Space Collapse in Fragmented Limit Order Book Markets,” Maglaras, Moallemi, and Zheng analyze this setting as a multi-class queueing system. Taking into account the effect of investors’ order-routing decisions across exchanges, they find that the equilibrium of this decentralized market exhibits a state space collapse property, whereby the queues at different exchanges are coupled, and the overall behavior of the market is captured through a one-dimensional process that can be viewed as a weighted aggregate queue length across all exchanges. The key driver of this coupling phenomenon is anticipated delay as opposed to the queue lengths themselves.