



Operations Research

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

In This Issue

To cite this article:

(2019) In This Issue. Operations Research 67(5):ii-iv. <https://doi.org/10.1287/opre.2019.1927>

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 © 2019, 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.2019.1927>

Copyright: © 2019 INFORMS

Bayesian Social Learning from Consumer Reviews

When buying a new product online, potential consumers often read online reviews to get an understanding of the quality of the product, and then they decide whether to make a purchase. In turn, people who purchase the product submit reviews to the aggregator site, thus adding new information over time. “Bayesian Social Learning from Consumer Reviews” by Ifrach, Maglaras, Scarsini, and Zseleva is the first to study such a social learning mechanism based on online reviews under the assumption that consumers are Bayesian. It identifies assumptions under which consumers eventually learn the true quality and explores the optimal pricing policy for a seller who participates in such a market.

A New Way to Maximally Spread Online Advertising Impressions Across Targeted Audience Segments

Advertisers engaged in brand-building activities online often purchase significant volumes of impression-based online advertising (e.g., banner ads or video ads) from website publishers or their ad-technology partners who schedule and deliver online advertisements. Such advertising is typically *targeted*, which means that ads can only be shown to specific audience segments chosen by the advertiser (e.g., urban females from any U.S. city). In general, advertisers are not well served by publishers who further constrain an ad campaign’s targeting (e.g., by delivering disproportionately many impressions of urban females from Los Angeles and not from other cities). In “Planning Online Advertising Using Gini Indices,” Lejeune and Turner consider a new objective function for spreading impressions of online advertising across targeted audience segments. The authors show how the Gini Index, which economists use to measure wealth or income inequality, can be deployed within a publisher’s optimization model to minimize the weighted sum of Gini indices across all ad campaigns and, thus, maximally spread impressions across targeted audience segments while limiting demand shortfalls. Key properties and solution structure are compared to a popular existing non-Gini-based ad spreading model developed at Yahoo, and a novel optimization-based decomposition scheme is developed to efficiently solve the Gini-allocation problem. Finally, the authors illustrate how Lorenz curves may be used to visualize a Gini-based spread so that managers can effectively monitor the performance of a publisher’s ad delivery system.

Prepositioning Medical Countermeasures Against Bioterrorism Attacks

To protect civilians in the event of a bioattack, that is, the intentional release of pathogens against them to cause harm,

public authorities maintain stockpiles of antibiotics. Unfortunately, a minute quantity of pathogens is sufficient to infect a large number of people. Furthermore, the treatment time window can be very short, for example, on the order of hours. Therefore, a significant amount of resources goes into maintaining a responsive antibiotics supply chain. In “Designing Response Supply Chain Against Bioattacks,” Simchi-Levi, Trichakis, and Zhang devise a decision-support framework for end-to-end supply chain design that captures logistical costs, transportation and dispensing times, and health effects of treatment delays. The authors model the problem as two-stage robust optimization on a network, which has been shown in the literature to be NP-hard. Using the so-called affinely adjustable robust counterpart, the authors are able to solve this problem on cases with hundreds of thousands of nodes. Furthermore, they prove that affine policies are optimal under certain conditions. Extensive numerical studies demonstrate the efficacy of affine policies for the general problem. An illustrative and high-fidelity case study is performed to examine the biodefense supply chain design for the Strategic National Stockpile maintained by the Centers for Disease Control and Prevention in the United States.

Sharing Rules and Supplier Coalitions in Assembly Systems

Assembly systems are composed of suppliers selling components to an assembler, who configures a final product for sale to the market. This is commonplace in many supply chains, such as healthcare and automobiles. Suppliers in such systems often form coalitions to enhance their bargaining power as well as for economies of scale. In “Dynamic Stable Supplier Coalitions and Invariance in Assembly Systems with Commodity Components,” Nagarajan, Sošić, and Tong show some interesting findings—how payoff allocation among coalition members can affect outcomes and when such outcomes are robust to allocation rules.

Optimizing Product Offering Decisions When Customers Only Consider a Small Number of Products

In “Technical Note—Assortment Optimization with Small Consideration Sets,” Feldman, Paul, and Topaloglu study a customer choice model that captures purchasing behavior when there is a limit on the number of times that a customer will consider purchasing. Under this model, we assume each customer is characterized by a ranked preference list of products and, upon arrival, will purchase the highest-ranking offered product. Because we restrict ourselves to settings in which customers consider a limited number of

products, we assume that these rankings contain at most k products. We call this model the k -product nonparametric choice model. We focus on the assortment-optimization problem under this choice model. In this problem, the retailer wants to find the revenue-maximizing set of products to offer when the buying process of each customer is governed by the k -product nonparametric choice model. We develop a linear programming-based randomized rounding algorithm that gives the best known approximation guarantee.

Process Flexibility Design in the Dynamic Make-to-Order Environment

The classical process flexibility literature has mostly focused on single-period models, as finding the optimal multiperiod policy is difficult because of the curse of dimensionality. In “Process Flexibility for Multiperiod Production Systems,” Shi, Wei, and Zhong study the design of sparse flexibility in dynamic make-to-order environments. Leveraging theories from applied probability, they propose a notion of “generalized chaining,” which allows one to design effective sparse flexibility structures in regimes in which capacity slacks are low. They then construct a structure that performs asymptotically close to full flexibility with only $m + n$ arcs, where m and n represent the number of plants and product types, while also demonstrating that $m + n$ is tight, that is, constructing systems in which even the best flexibility structure with $m + n - 1$ arcs cannot achieve the same asymptotic performance. These results provide insights into how design principles for process flexibility extend beyond the classical single-period model.

Stochastic Optimization with Decisions Truncated by Dependent Random Variables

In many operations management problems, the decisions are truncated by random variables. Take a dual sourcing inventory management problem as an example: the suppliers may have random capacities, and the actual received quantity from ordering is truncated by this random capacity. Often, the random capacities of different suppliers may be dependent. An interesting challenge is that due to the truncation, the optimization problem may not be convex. In “Technical Note—Stochastic Optimization with Decisions Truncated by Positively Dependent Random Variables,” Chen and Gao propose a transformation technique to convert the original nonconvex minimization problem to an equivalent convex one. They demonstrate the application of their method using an inventory substitution problem with dependent random supply capacities and a two-part fee cost structure. In addition, their method can also incorporate the decision maker’s risk attitude.

New Ambiguity Set for Distributionally Robust Optimization

Distributionally robust optimization has grown into one of the most popular approaches in addressing optimization under uncertainty. As its key ingredient, the ambiguity set specifies several types of distributional information about the ambiguous true distribution. In “Distributionally Robust Optimization with Infinitely Constrained Ambiguity

Sets,” Chen, Sim, and Xu contribute to the existing ambiguity sets in the literature by proposing a new class of ambiguity sets that encompasses a potentially infinite number of expectation constraints. This class has the benefits of, among other things, a tighter approximation of stochastic independence that is ubiquitous in characterizing uncertainty. The authors propose an algorithm involving a greedy improvement procedure to solve the corresponding distributionally robust optimization problem. In their computational study, the authors are able to obtain significantly improved solutions using this algorithm.

Algorithms and Complexity for the Replenishment Schedule to Minimize Peak Storage Problem

The replenishment storage problem (RSP) is to minimize the storage capacity requirement for a multi-item inventory system, where each item has a given reorder size, uniform demand, and cycle length. The discrete-RSP, which allows reorders to take place only at integer time units within the cycle, is NP-hard for constant joint cycle length (the least common multiple of the length of all individual cycles). In “The Replenishment Schedule to Minimize Peak Storage Problem: The Gap Between the Continuous and Discrete Versions of the Problem,” Hochbaum and Rao show that discrete-RSP is weakly NP-hard for constant joint cycle length and strongly NP-hard for nonconstant joint cycle length. For constant joint cycle length discrete-RSP, we present a pseudopolynomial time algorithm that solves the problem optimally and the first known fully polynomial time approximation scheme for the single-cycle RSP. The scheme is utilizing a newly introduced integer programming formulation of the problem. For the strongly NP-hard RSP with nonconstant joint cycle length, we devise a polynomial time approximation scheme that runs in linear time for fixed ϵ . This paper further explores the relationship of discrete-RSP with the continuous-RSP, where reorders can take place at any time within a cycle.

Uncovering Input Information from Output Data

Successful simulation-based analytics rely on the availability of input models that sufficiently resemble real-world processes. In many operational situations, it can be a luxury to have access to the direct observations on these inputs, and only aggregate or output-level data are available. The paper “Optimization-Based Calibration of Simulation Input Models” by Goeva, Lam, Qian, and Zhang studies a statistically principled approach to uncover the inputs from only output data. Their approach harnesses a novel use of data-driven distributionally robust optimization that imposes constraints to match information at the output level. The authors show the statistical and computational guarantees in relation to the function complexity of these constraints and demonstrate the numerical effectiveness of their input reconstruction scheme.

Predicting Catastrophic Events in Supply Networks

Catastrophic events can dramatically disrupt distribution networks and supply chains; think, for example, about the threat of a hurricane, which causes significant spikes in

demands for gasoline, water, and other products. What are the weak links in such a supply network, or what is the most likely way in which such a network can be disrupted? In “Rare Event Simulation for Distribution Networks,” Blanchet, Li, and Nakayama devise optimal simulation algorithms for estimating the occurrence of disruptive events in complex distribution networks. This development is noteworthy because it can be used to mitigate the impact in populations that are susceptible to these types of events, which can incur tremendous costs to society.

How to Inform Customers of Their Wait Times?

In many systems with limited service capacity, customers must wait in a queue for service. When customers cannot observe the queue, how should a revenue-maximizing service provider convey information about wait times? In “Optimal Signaling Mechanisms in Unobservable Queues,” Lingenbrink and Iyer study this problem and characterize the structure of the optimal signaling mechanism. To signal optimally, the service provider uses two possible signals, “short” and “long,” to tell customers the queue length is short when below a threshold and long when above it. For the specific case of linear waiting costs, the authors explicitly compute this threshold. Furthermore, they show that for an optimally chosen fixed service price, optimal signaling produces the same expected revenue as a pricing mechanism that sets prices based on the number of customers waiting. This suggests that in settings where one cannot dynamically update prices, signaling can be effective in generating revenue.

Efficient Allocation of Resources Without Money

How should a planner allocate a single resource to multiple requesters efficiently when monetary transfers are not feasible? This question naturally arises in many relevant settings ranging from healthcare and antipoverty programs to cloud computing systems. In these settings, resource requests occur repeatedly, and requesters’ private values might change over time. In “Multiagent Mechanism Design Without Money,” Balseiro, Gurkan, and Sun propose a mechanism that asymptotically achieves the first-best efficient allocation (the welfare-maximizing allocation as if values are publicly observable) as requesters become more patient. Furthermore, the authors provide sharp characterizations of convergence rates to first best as a function of the discount factor. In the case of two agents, the authors prove that the convergence rate of their mechanism is optimal—that is, no other mechanism can converge faster to first best.

Understanding the Fundamentals of Empty-Car Routing in Ridesharing Systems

How can empty cars be efficiently routed in ridesharing systems? In “Empty-Car Routing in Ridesharing Systems,” Braverman, Dai, Liu, and Ying introduce a novel model

based on closed queueing networks and propose an optimization framework to optimize empty-car routing for maximizing system-wide utility functions. We propose a fluid-based optimal routing policy by solving the optimization problem in a large market regime. We establish both process-level and steady-state convergence of the closed queueing network to the fluid-limit and prove the optimal network utility obtained from the fluid-based optimization is an upper bound on the utility in the finite car system for any routing policy under which the closed queueing network has a stationary distribution. This upper bound is achieved asymptotically under the fluid-based optimal routing policy.

An Optimal Dynamic Assortment Selection

In “MNL-Bandit: A Dynamic Learning Approach to Assortment Selection,” Agrawal, Avadhanula, Goyal, and Zeevi consider a dynamic assortment selection problem where consumers choose according to a multinomial logit (MNL) choice model. The decision maker observes this choice, and the objective is to dynamically learn the model parameters while optimizing cumulative revenues over a selling horizon. We refer to this as the MNL-Bandit problem, which arises in several important application domains such as online retail, recommendation systems, and display-based advertising. The authors present an efficient algorithm that judiciously combines exploration of the combinatorial option space and exploitation of that information to maximize the cumulative reward. They show that our algorithm achieves near-optimal performance and has attractive numerical properties. This is the first policy with theoretical guarantees that does not require any knowledge about underlying parameters, making this approach universal in its scope.

Handling Long-Term Constraints in Multiarmed Bandit Problems

Multiarmed bandit (MAB) is a classic model for capturing the exploration–exploitation trade-off inherent in many sequential decision-making problems. The classic MAB framework, however, only allows “local” constraints on decisions and “sum of rewards” as objective. In many real-world applications, there are multiple complex constraints on resources that are consumed during the entire decision process, and performance may be evaluated through nonlinear utility functions on aggregate rewards. In “Bandits with Global Convex Constraints and Objective,” Agrawal and Devanur present a new MAB framework that allows such “global” convex constraints and concave objective functions along with new algorithmic techniques with provably near-optimal performance bounds. The authors discuss applications in several domains, such as network revenue management, crowdsourcing, and pay-per-click advertising, that benefit from the new, more general framework by admitting richer models and more efficient risk-averse solutions.