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(6):ii-iv. https://doi.org/10.1287/opre.2021.2236

http://pubsonline.informs.org/journal/

OPERATIONS RESEARCH
Vol. 69, No. 6, November–December 2021, pp. ii–iv
ISSN 0030-364X (print), ISSN 1526-5463 (online)

In This Issue

https://doi.org/10.1287/opre.2021.2236

Inventory Management in Stochastic Distribution Systems

In "Technical Note—A Simple Heuristic Policy for Stochastic Distribution Inventory Systems with Fixed Shipment Costs," Zhu, Chen, Hu, and Yang study a continuous-review, two-echelon inventory system with one central warehouse, multiple local facilities, and each facility facing random demand. Local facilities replenish their stock from the central warehouse, which in turn places orders at an outside supplier with ample supply. The optimal policy remains unknown for such a system, and even if it exists, such a policy must be extremely complicated. The authors evaluate a class of easy-to-implement heuristics, referred to as modified echelon (r, Q) policies. The parameters for such a heuristic are obtained by solving a set of independent single-stage systems. They show that the proposed policy is asymptotically optimal, as pairs of system primitives, such as the ratios of the fixed cost of the central facility to those of the local facilities, are scaled up. The authors also show that as the number of retailers grows, the performance bound of the heuristic converges to a primitive-dependent constant.

The Adoption of Portfolio Compression Drives Structural Changes in Financial Networks

Over-the-counter (OTC) derivatives markets amount to several hundred trillion dollars of notional and play a key role in managing systemic risk. In "Compressing Over-the-Counter Markets," D'Errico and Roukny show how the recent introduction of an optimization technology called portfolio compression can deeply transform the size and structure of these markets. In contrast with other multilateral solutions, compression techniques are decentralized; that is, they exploit netting opportunities without requiring the implementation of a clearinghouse. The authors show that the tightly knit and concentrated structure of OTC markets generates large excess in notional that can be efficiently compressed even under conservative scenarios. Using European data on credit-default-swap transactions, the authors find that more than two-thirds of a market total size can be eliminated if all participants engage in a single conservative compression cycle. When considering both central clearing and portfolio compression, the authors find that compressing positions across CCPs by and large eliminates the loss in netting opportunities that is otherwise present when several clearinghouses compete in the system.

Dynamic Routing Games with Adaptive and Nonadaptive Atomic Agents

Dynamic routing games have drawn more and more research attention in recent years. In "Atomic Dynamic Flow Games: Adaptive vs. Nonadaptive Agents," Cao, B. Chen, X. Chen, and Wang propose a new dynamic routing model with selfish atomic agents. They consider two different rationality levels of agents in making routing decisions. If agents are nonadaptive in that they strictly follow the origin–destination paths they choose at the very beginning, the authors constructively prove that a Nash equilibrium (NE) always exists and show that all NEs possess some nice properties. They also provide efficient algorithms for finding an NE. On the other hand, if the agents make adaptive decisions at every nonterminal vertex they reach as to which next edge to take, the authors show that a subgame perfect equilibrium (SPE) always exists, and the SPE set includes all NEs for nonadaptive agents.

Tools From Mathematical Ecology in a Combat Model with Humanitarian Aid Agencies

Conflict models have a long history of taking inspiration from mathematical ecology. In "A Mathematical Model of Humanitarian Aid Agencies in Attritional Conflict Environments," McLennan-Smith, Kalloniatis, Jovanoski, Sidhu, Roberts, Watt, and Towers seek to enrich counterinsurgency (COIN) warfare models to account for modern and future complexities by incorporating nontrophic effects and the functional response from mathematical ecology. The authors consider the application of these ideas in a COIN scenario in which a humanitarian aid agency is present in the conflict environment to support the local population. In this scenario, the aid agency plays the unwilling role of a "hospital shield" whereby it is forced to, or inadvertently, shield combatants or weapons. In contrast to the typical behavior seen in the classic Lanchester system, this model gives rise to limit cycles and bifurcations that the authors interpret through a warfighting application. Finally, through a case study, the authors highlight the importance of the agility of an intervention force in achieving victory when humanitarian aid agencies are present.

Designing Open-bid Procurement Auction with Supplier Qualification Screenings

Manufacturers often use re-sourcing initiatives to keep their suppliers' pricing competitive; e.g., new entrant suppliers are identified and invited to compete with the incumbent supplier for supply contracts in an open-bid auction. To ensure that entrant suppliers have the capability of executing the contract, the conventional approach for many manufacturers is to conduct qualification screenings on all the entrants prior to auction bidding and only allow qualified entrant suppliers to compete in the auction. In "'Now or Later?' When to Deploy Qualification Screening in Open-Bid Auction for Re-Sourcing," Zhang, Chen, and Katok explore an alternative arrangement of this process where all entrant suppliers are invited for bidding first before qualification screenings are selectively conducted afterward to determine the contract winner. This new approach helps reduce the waste of manufacturers' screening efforts on suppliers with uncompetitive bids but in the meantime introduces incentives for less competitive supplier bidding behavior. They provide analytical and numerical evidence that this new approach could be very effective in managing manufacturers' procurement costs.

Optimal Auction Duration in Financial Markets

In the considered auction market, market makers fill the order book during a given time period while some other investors send market orders. The clearing price is set to maximize the exchanged volume at the clearing time according to the supply and demand of each market participant. The error made between this clearing price and the efficient price is derived as a function of the auction duration. In "Optimal Auction Duration: A Price Formation Viewpoint," Jusselin, Mastrolia, and Rosenbaum study the impact of the behavior of market takers on this error to minimize their transaction costs. They compute the optimal duration of the auctions for 77 stocks traded on Euronext and compare the quality of the price formation process under this optimal value to the case of a continuous limit order book. Continuous limit order books are usually found to be suboptimal. Order of magnitude of optimal auction durations is from 2–10 minutes.

How Bookies Can Outwit Sophisticated Bettors

Sports-betting markets are based entirely on predictions. A bettor has to pick a winning contestant, and a market maker-a bookie-bets on the opponent. As bookies have to take the other side of every bet, it is of great value to understand the market making problem, that is, how to set the spread lines as "prices" for the bookies. Nevertheless, understanding of this problem is limited. Specifically, sophisticated bettors exist in the market, and a bookie can be manipulated by skillful bettors because of information asymmetry. In "Dynamic Learning and Market Making in Spread Betting Markets with Informed Bettors," Birge, Feng, Keskin, and Schultz study the market-making problem under information asymmetry and market manipulation. They show that, although many popular learning and pricing algorithms, such as Bayesian policies, are effective in learning, they are vulnerable to strategic manipulations. The authors propose a dynamic learning and pricing algorithm, called the inertial policy, that collects information from the market effectively but also protects the bookie from strategic manipulations.

The Pitfalls of Online Learning in the Presence of Strategic Users

Many online platforms—ride-hailing, freelancing, and online advertising platforms are prominent examples—can use past data generated by a user to personalize future offerings. In such environments, the platform could employ a myriad of dynamic strategies to personalize offers based on the data collected from past interactions. How much additional benefit can be derived from such dynamic strategies when the user is strategic? In "On the Futility of Dynamics in Robust Mechanism Design," Balseiro, Kim, and Russo use the language of robust mechanism design to formalize a stark impossibility result. They identify a broad class of problems in which an optimal dynamic mechanism is static without meaningful dynamics. Their work highlights novel challenges faced by platforms when trying to learn and exploit the private information of strategic users: dynamic mechanisms are vulnerable to manipulations from a strategic user who could try to induce future outcomes that are beneficial for him at the expense of the platform.

Information Inundation and Polarization

In "Information Inundation on Platforms and Implications," Allon, Drakopoulos, and Manshadi study the process of learning and opinion formation in the presence of platforms. The paper is motivated by the relatively recent changes in how people access and consume news. Specifically, more and more people access news through social platforms, such as Facebook, Reddit, or Twitter. The authors try to answer a simple question: Why is it that we become more polarized even though we have more information sources than ever? The authors assume that people are almost rational and are not easily biased as a result of psychological biases. Yet the authors show that, even if people are good "statisticians" in the sense that they try to find the information that reduces their ignorance by the most significant amount, they choose to consume information that slows their ability to learn. In fact, they show that screening information sources results in polarization: everyone is sort of "stuck" in their side of the political map. They are essentially illustrating a mechanism that generates confirmation bias even for a seemingly rational reader.

How to Optimize Posted Price Mechanisms?

The sequential posted-price (SPP) mechanism is one of the widely used selling mechanisms in practice. In this mechanism, the seller presents each buyer with a price sequentially and the buyer can either accept or reject the mechanism's offer. Despite the widespread use of the SPP mechanism, the problem of optimizing prices in this mechanism has not been fully addressed. In "Improved Revenue Bounds for Posted-Price and Second-Price Mechanisms," Beyhaghi, Golrezaei, Leme, Pal, and Sivan construct SPP mechanisms by considering the best of two simple pricing rules: one that imitates the optimal mechanism and the other that posts a uniform price (same price for every buyer). Their simple pricing rules can be easily generalized to the setting with multiple units and yield the first improvement over long-established approximation factors.

How to Optimally Set Redemption Hurdles for Reward Programs?

Redemption hurdles, such as finite expiration terms and redemption thresholds, are common for customer reward programs. In "An Analysis of 'Buy X, Get One Free' Reward Programs," Liu, Sun, and Zhang study the economic rationale behind redemption hurdles and how they should be optimally set. They show analytically that redemption hurdles can be used as a price-discriminating vehicle that increases firm profitability. Redemption hurdles can facilitate the firm's price discrimination on consumers whose valuations may vary over time. Redemption threshold alone cannot ensure profitability unless it is coupled with a finite expiration term or a positive transaction utility from the rewarded free product. Optimal design of redemption hurdles is not straightforward, and the interdependence between the two types of redemption hurdles and the price is nontrivial. Optimally set redemption hurdles may not only increase firm profitability but also, increase the welfare of consumers who purchase frequently.

Timing It Right: Balancing Inpatient Congestion vs. Readmission Risk at Discharge

One of the most important decisions a hospitalist makes at the intersection of cost and quality of care is when to discharge a patient from the hospital. Keeping patients longer (shorter) increases (decreases) overcrowding and hospital costs but also decreases (increases) readmission risk. In "Timing It Right: Balancing Inpatient Congestion vs. Readmission Risk at Discharge," by Shi, Helm, Diglise-Hawkinson, and Pan, the authors developed a long-run average cost optimization problem for determining on each day who and how many patients to discharge. They combined structural properties of the model with an analytical solution for a special cost structure to approximately solve the high-dimensional Markov decision process. This transformed the originally intractable problem into a

simple univariate optimization problem that can be solved efficiently yet allowed capture of time nonstationarity and fully heterogeneous inpatient populations, where each patient has a personalized risk trajectory. Moreover, the authors took one step beyond theory and implemented their discharge decision support tool in a partner hospital. For the tool to be properly parametrized and implementable, the authors developed a model to predict readmission risk as a function of length of stay that integrated several statistical methods in a novel manner. The resulting implementation was described as a showcase, demonstrating the tool's applicability for integration with general hospital data systems and workflows.

To Pool or Not to Pool: Queueing Design for Large-Scale Service Systems

In large-scale service systems, it is a common practice to organize customers with similar service requirements into a single queue served by a group of servers. This pooled queue structure is deemed highly efficient because the servers' idleness will be minimized. In "To Pool or Not to Pool: Queueing Design for Large-Scale Service Systems," Cao, He, Huang, and Liu demonstrate that the dedicated queue structure, under which each server has her own queue, could be more advantageous for improving the system's service level. Moreover, the servers' additional idleness induced by the dedicated queue structure will be negligible when the system scale is large. By solving a staffing problem, this study also intends to help service system designers answer the following question: To achieve a specified service-level objective in a more efficient manner, should the servers have a common queue or separate queues?

Axiomatic Foundations for Nonadditive Multiattribute Portfolio Utility Functions

Selecting which portfolio of project candidates (e.g., products, infrastructure investments, or policy options) to implement with the available resources is an important decision problem faced frequently by organizations. Such decisions are commonly supported with additive portfolio models in which portfolio utility is computed as the sum of the projects' multiattribute utilities. Nevertheless, such models lack decision-theoretic justification and can therefore result in decision recommendations that are misaligned with the decision maker's preferences. In "Nonadditive Multiattribute Utility Functions for Portfolio Decision Analysis," Liesiö and Vilkkumaa establish the axiomatic foundations of a more general class of multilinear portfolio utility functions, which includes additive and multiplicative portfolio utility functions as special cases. They also develop preference elicitation techniques to assess these portfolio utility functions and optimization models that enable the use of standard off-the-shelf optimization software to identify the most preferred portfolio in view of resource and other constraints. The results of this paper enable the use of portfolio decision analysis in practical applications even when the restrictive preference assumptions underlying the additive model do not hold.

Performance of the Smallest-Variance-First Rule in Appointment Sequencing

Setting up appointment schedules plays an important role in healthcare and various other domains. An important goal is to determine appointment times for the patients in order to keep both patient waiting times and doctor idle time small; this is captured quantitatively by minimizing a cost function involving both of these objectives. One aspect of this problem is to decide the order in which the patients should arrive. Often the smallest-variance-first (SVF) rule is used, placing the patients in order of increasing variance of their service durations. Although SVF is not necessarily optimal, it does well in practice and simulation. In "Performance of the Smallest-Variance-First Rule in Appointment Sequencing," de Kemp, Mandjes, and Olver provide theoretical justification for this rule by obtaining worst-case bounds on its performance compared with the best possible schedule as well as showing that the rule approaches optimality as the number of patients grows large.

Dynamic Scheduling Under the Shift Constraint

Many service systems operate with staff who work in shifts. In "Dynamic Server Assignment in Multiclass Queues with Shifts, with Application to Nurse Staffing in Emergency Departments," Chan, Huang, and Sarhangian study the dynamic assignment of a fixed number of servers to different classes of a multiclass queueing system, under the restriction that the assignments can only be adjusted at the beginning of discrete shifts. The primary motivation is the assignment of nurses to different areas of the emergency department (ED) of a hospital, based on the observed congestion levels at the beginning of each shift. By studying associated fluid control problems, the authors provide insights on the structure of "good" policies and propose heuristic policies that perform well, even in the presence of time-varying arrivals and customer abandonment. An important observation is that, in a parameter regime relevant to the ED application, this partial flexibility with respect to server assignment can significantly reduce the expected waiting costs incurred in the system.