



Operations Research

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

In This Issue

To cite this article:

(2023) In This Issue. Operations Research 71(1):iii-v. <https://doi.org/10.1287/opre.2022.2433>

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 © 2023, 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.2022.2433>

Copyright: © 2023 INFORMS

Optimal Marketing of Grid Energy Storage

Grid energy storage plays a key role in making carbon-free, renewable energy production a reality. Yet, when it comes to maximizing profit, owners of storage assets still struggle with coordinating their trading activities across time because of the complex nature of multisettlement electricity markets. In “The Value of Coordination in Multimarket Bidding of Grid-Energy Storage,” Löhndorf and Wozabal propose a multistage stochastic programming model for market-oriented optimization of energy storage. To calculate lower and upper bounds on optimal values, they develop novel methods for scenario-tree generation and information relaxation. They show that a coordinated policy that reserves capacity for the short-term markets is optimal and that the gap to a sequential policy, increases with short-term price volatility and market liquidity. The authors find that coordination is beneficial for all considered asset types and that flexible storages with high price impact benefit most. Their findings inform storage owners which markets contribute most value, how to organize trading across time, and how to calculate optimal bidding strategies.

New Metric for Bed Occupancy Can Potentially Reduce Bed Shortages in Hospitals

In “The Analytics of Bed Shortages: Coherent Metric, Prediction, and Optimization,” Xie, Loke, Sim, and Lam propose a risk-adjusted version of bed occupancy rates (BORs) that can be used for patient admission control and bed capacity planning in hospitals. Simulations indicate a potential increase of 11% in elective patient admissions by planning using the proposed bed shortage index (BSI) as opposed to the traditional BOR. The BSI is also illustrated for purposes of bed capacity allocation between departments and across acute and nonacute hospitals.

Predicting Shortages in a Supply Chain Using Simulation

A supply chain shortage is a serious problem that can lead to assembly plant shutdowns. However, predicting such shortages using simulation poses a challenge, because the information necessary to initialize such a supply chain simulation is often only partially observable in many real-world applications. In “Simulation-Based Prediction,” Lim and Glynn investigate this problem. They formulate such a prediction problem as the problem of computing the conditional expectation of the quantity of interest, given the observed state of the system. Simulation can be easily applied to computing such a conditional expectation when the simulation state is fully observed in the real system. The authors propose a new simulation methodology appropriate to the many settings in which the observed current state does not fully determine the simulation’s initial state. With the use of such methods, simulation has the potential to more accurately predict upcoming supply chain bottlenecks and to enhance predictions in the many other problem settings where simulation is commonly used.

Robust Satisficing

Optimization under uncertainty, even in the Era of Big Data, remains a perennial difficulty; without taking into account uncertainty in the data, the *Optimizer’s Curse* dictates that we should also expect inferior results in the actual outcome. To address this issue, in “Robust Satisficing,” Long, Sim and Zhou propose a new *robust satisficing* framework where the modeler specifies an attainable target as a trade-off for the model’s ability to withstand greater uncertainty. The authors introduce robust satisficing models for data-driven linear optimization, combinatorial optimization, and linear optimization problems with recourse, which can be applied to solving a host of practical optimization problems arising in operations management, finance, healthcare, and transportation, among others. In alleviating the Optimizer’s Curse, the authors show computationally that the solutions to the robust satisficing models can be more effective than those obtained by traditional robust optimization approaches.

Importance of Joint Pricing and Matching Optimization in Online Marketplaces

Motivated by applications from gig economy and online marketplaces, the authors study a two-sided queueing system under joint pricing and matching controls. The queueing system is modeled by a bipartite graph, where the vertices represent customer or server types and the edges represent compatible customer-server pairs. In “Dynamic Pricing and Matching for Two-Sided Queues,” Varma, Bumpensanti, Maguluri, and Wang propose a threshold-based two-price policy and queue length-based maximum-weight matching policy and show that it achieves a near-optimal profit. They study the system under the large-scale regime, wherein the arrival rates are scaled up, and under the large-market regime, wherein both the arrival rates and numbers of customer and server types increase. The authors show that two-price policy is a primary driver for optimality in the large-scale regime. They demonstrate the advantage of maximum-weight matching with respect to the number of customer and server types. Concurrently, they show that the interplay of pricing and matching is crucial for optimality in the large-market regime.

A Novel Class of Robust and Fast Algorithms for Online Allocation Problems

A central problem in operations research is allocating limited resources sequentially to maximize cumulative rewards. Applications abound and include network revenue management and internet advertising among many others. Existing data-driven algorithms are tailored for convex settings with either adversarial or stochastic inputs. Many modern applications of online allocations problems, however, are nonconvex. Furthermore, algorithms for adversarial inputs may be too conservative in practice, whereas algorithms for stochastic inputs can perform poorly when the model is misspecified. In “The Best of Many Worlds: Dual Mirror Descent for Online Allocation Problems,” Balseiro, Lu, and Mirrokni present a novel class of algorithms for nonconvex online allocation problems

that attain good performance simultaneously in stochastic and adversarial input models and also in various nonstationary settings. The resulting algorithms are simple, fast, and robust to noise and corruption in the observations, in contrast to existing methods from the literature.

Dual Bounds for Periodical Stochastic Programs

In “Dual Bounds for Periodical Stochastic Programs,” Shapiro and Cheng discuss a construction of the dual of a periodical formulation of infinite-horizon linear stochastic programs with a discount factor. The dual problem is used for computing a deterministic upper bound for the optimal value of the considered multistage stochastic program. Numerical experiments demonstrate behavior of that upper bound, especially when the discount factor is close to one.

Learning Compact High-Dimensional Models in Noisy Environments

Building compact, interpretable statistical models where the output depends upon a small number of input features is a well-known problem in modern analytics applications. A fundamental tool used in this context is the prominent best subset selection (BSS) procedure, which seeks to obtain the best linear fit to data subject to a constraint on the number of nonzero features. Whereas the BSS procedure works exceptionally well in some regimes, it performs pretty poorly in out-of-sample predictive performance when the underlying data are noisy, which is quite common in practice. In “Subset Selection with Shrinkage: Sparse Linear Modeling When the SNR is Low,” Mazumder, Radchenko, and Dedieu explore this relatively less-understood overfitting behavior of BSS in low-signal noisy environments and propose alternatives that appear to mitigate such shortcomings. They study the theoretical statistical properties of our proposed regularized BSS procedure and show promising computational results on various data sets, using tools from integer programming and first-order methods.

Learning Personalized Treatment Policies from Observational Data

As a result of digitization of the economy, more and more decision makers from a wide range of domains have gained the ability to target products, services, and information provision based on individual characteristics. Examples include selecting offers, prices, advertisements, or emails to send to consumers, choosing a bid to submit in a contextual first-price auctions, and determining which medication to prescribe to a patient. The key to enabling this is to learn a treatment policy from historical observational data in a sample-efficient way, hence uncovering the best personalized treatment choice recommendation. In “Offline Multi-Action Policy Learning: Generalization and Optimization,” Zhou, Athey, and Wager provide a sample-optimal policy learning algorithm that is computationally efficient and that learns a tree-based treatment policy from observational data. In the quest toward fully automated personalization, the work provides a theoretically sound and practically implementable approach.

Combating the COVID-19 Pandemic with Operations Research

“Starting in December 2019, the COVID-19 pandemic is an unprecedented humanitarian crisis with millions of deaths worldwide”. In “Forecasting COVID-19 and Analyzing the Effect of Government Interventions,” Li, Bouardi, Lami, Trikalinos, Trichakis, and Bertsimas proposed a novel epidemiological model, DELPHI, that combined a novel modeling of government interventions, nonlinear optimization, and compartmental epidemiology models to forecast

COVID-19 spread. They used DELPHI to demonstrate how lockdowns reduced the transmission by nearly 80%, whereas earlier societal action could have saved more than 75% of the lives lost in many countries. The authors also created a scenario analysis toolkit that utilized DELPHI’s modeling of interventions to generate “what if” scenarios under different future interventions. Janssen Pharmaceuticals utilized this toolkit to select optimized locations for the Phase III trial of their COVID-19 vaccine, leading to a trial acceleration of 8 weeks and saving thousands of lives.

Lagrangian Inference for Ranking Problems

Understanding ranking orders of different items is of great importance in many applications such as sports, online game, and recommendation, among many others. In “Lagrangian Inference for Ranking Problems,” Liu, Fang, and Lu provide a novel approach to inferring the ranking systems. They aim to answer questions like, is item A better than item B? Is item A among the top 10 items? Such inference problems are challenging as they involve combinatorial structures. The key technical innovation is a new Lagrangian inference framework with new bootstrap tools. Strong theoretical guarantees are provided showing the optimality of the proposed method. A novel application of inferring movies’ ranking using a large-scale data set is provided to demonstrate the applicability of the proposed method.

A Robust Spectral Clustering Algorithm for Sub-Gaussian Mixture Models with Outliers

Traditional clustering algorithms such as k -means and vanilla spectral clustering are known to deteriorate significantly in the presence of outliers. Several previous works in literature have proposed robust variants of these algorithms; however, they do not provide any theoretical guarantees. Extending previous clustering literature on Gaussian mixture models, in “A Robust Spectral Clustering Algorithm for Sub-Gaussian Mixture Models with Outliers,” Srivastava, Sarkar, and Hanasusanto developed a new spectral clustering algorithm and provided error bounds for the algorithm under a general sub-Gaussian mixture model setting with outliers. Surprisingly, their derived error bound matches with the best-known bound for semidefinite programs under the same setting without outliers. Numerical experiments on a variety of simulated and real-world data sets further demonstrate that their algorithm is less sensitive to outliers compared with other state-of-the-art algorithms.

Sequential Mechanisms with Ex Post Individual Rationality

Online multiproduct sellers increasingly use interactive selling strategies to customize their offers to individual buyers. For example, a seller may adjust the prices of products dynamically based on user interaction and offer discounts for buying bundles of products. What selling strategy should such a seller use to maximize profit? In “Sequential Mechanisms with Ex Post Individual Rationality,” Ashlagi, Daskalakis, and Haghpour provide a recursive characterization of the optimal selling strategy. This characterization is used to identify conditions under which the ability to bundle products is less profitable for the seller than the ability to adjust prices dynamically.

Partially Recovering a Graph Alignment in the Correlated Erdős–Rényi Model

Given two graphs, how can we partially recover a one-to-one mapping between nodes that maximizes edge overlap? This problem, known as the graph alignment problem, appears in settings such as social network deanonymization and cellular biology. In

“Partial Recovery in the Graph Alignment Problem,” Hall and Massoulié consider a stylized mathematical model of problems of this type: They assume that the input graphs are generated via a probabilistic model, namely, the correlated Erdős–Rényi model with parameters (n, q, s) . The authors provide both necessary and sufficient conditions on (n, q, s) under which partial recovery can be achieved. In particular, they show that partial recovery can be achieved in the $nqs = \Theta(1)$ regime under certain additional assumptions.

Selecting a Parcel Type Portfolio to Reduce Unused Space in Transportation

Wrongly sized parcels lead to unused space and inefficient transportation. With continuously increasing e-commerce and last-mile delivery volumes, available parcel types at a warehouse can significantly impact unused space that is transported. In “A Branch-and-Repair Method for Three-Dimensional Bin Selection and Packing in E-Commerce,” Fontaine and Minner solve the trade-off between cost of unused space and cost of parcel variety through optimizing the portfolio of available parcel types. To solve large instances with millions of binary decision variables, the authors develop an exact decomposition method that allows for relaxing many binary variables, improves branch-and-check by repairing infeasible solutions, and shows how to avoid solving many subproblems. A case study using real data shows the efficiency of the proposed method and the impact of the portfolio on unused transportation space.

Optimizing School Operations, Holistically

School districts in the United States face a variety of operational problems, often treated in isolation due to their inherent complexity: for instance, school assignment and school transportation are rarely considered jointly. In “Policy Analytics in Public School Operations,” Bertsimas and Delarue develop an optimization-based approach to three key problems in school operations: school assignment, school bus routing, and school start time selection. Their methodology improves upon the state of the art in two ways: by leveraging a simplifying assumption of fixed route arrival times to tractably optimize school bus schedules and school start times simultaneously, and by proposing a postimprovement heuristic to jointly optimize assignment, bus routing, and scheduling. They evaluate their approach on a practical case study.

Explaining and Resolving Delays in Projects

Project management is responsible for 30% of the world’s economic activity, with an annual value of \$27 trillion. Yet, despite half a century of research and the training of millions of project managers, many projects are delivered late. This is typically attributed to Parkinson’s Law, meaning the expansion of work to fill available time. However, in “Work More Tomorrow: Resolving Present Bias in Project Management,” Shi, Hall, and Cui identify and demonstrate the alternative explanation of time-inconsistent behavior, that is, present bias. Under present bias, a decision maker values immediate costs and rewards more than future ones. The authors show that this behavioral issue is responsible for procrastination by project workers and overall project delay. Borrowing concepts from popular employee savings schemes, they develop an incentive scheme that mitigates present bias and significantly enhances project performance, as measured by on-time frequency and expected project tardiness.

Dynamic Pricing of a multiclass Make-to-Stock Queue

In “An Approximate Analysis of Dynamic Pricing, Outsourcing, and Scheduling Policies for a Multiclass Make-to-Stock Queue in the Heavy Traffic Regime,” Ata and Barjesteh propose an effective joint dynamic pricing, outsourcing, and scheduling policy for a multiclass make-to-stock manufacturing system. The authors approximate this control problem with a Brownian control problem. They solve the Brownian problem explicitly by exploiting the solution to a Riccati equation and propose a policy for the manufacturing system based on its solution. The proposed policy is a two-sided barrier policy. Outsourcing and idling processes are used to maintain the workload above the lower and below the upper barriers, respectively. Dynamic prices are used to control the workload process between the two barriers. The authors show using a simulation study that the optimality gap of the proposed policy is small, and their proposed policy outperforms a long list of static pricing policies. Moreover, the gap between their proposed policy and the static pricing policies increases with the server utilization and the outsourcing cost.

Optimal Capacity Rationing Under Service Level Constraints

In “Achieving High Individual Service-Levels Without Safety Stock? Optimal Rationing Policy of Pooled Resources,” Jiang, Wang, and Zhang analyze a resource rationing problem with service level constraints. They present a general framework to study the two-stage problem when customers require individual and possibly different service levels: (1) the capacity level of pooled resources in anticipation of random demand of multiple customers and (2) how the capacity should be allocated to fulfill customer demands after demand realization. The modeling framework generalizes and unifies many existing models in the literature and includes second-stage allocation costs. The authors propose a simple randomized rationing policy for any fixed feasible capacity level and show the optimality of this policy for very general service level constraints, including type I and type II constraints and beyond. They also discuss the optimality of index policies.

Minimizing the Service Requirement Violation Risk in the Inventory Routing Problem with Uncertain Demand

In the inventory routing problem, the supplier acts as a central decision maker who determines the replenishment quantities and also, the delivery times and routes to all retailers. In “Inventory Routing Problem Under Uncertainty,” Cui, Long, Qi, and Zhang develop a novel framework for the uncertain inventory routing problem and allow ambiguity in the probability distribution of each retailer’s uncertain demand. Adopting a service-level viewpoint, they minimize the risk of uncertain inventory levels violating a prespecified acceptable range. They quantify that risk using a new decision criterion, the service violation index, that accounts for how often and how severely the inventory requirement is violated. The solutions proposed here are adaptive in the sense that they vary with the realization of uncertain demand. They provide algorithms to solve the problem exactly and then demonstrate the superiority of their solutions by comparing them with several benchmarks.