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(4):ii-iv. https://doi.org/10.1287/opre.2019.1898

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

OPERATIONS RESEARCH
Vol. 67, No. 4, July–August 2019, pp. ii–iv
ISSN 0030-364X (print), ISSN 1526-5463 (online)

In This Issue

https://doi.org/10.1287/opre.2019.1898

How to Do Data Analysis While Still Preserving Privacy?

How to protect individual privacy when analysing data has become increasingly important in the era of big data, because many data sets collected are sensitive in areas such as economics, political science, and life science. In "Econometrics with Privacy Preservation," Cai and Kou propose an encryption and recovery (ER) algorithm that allows one to do statistical inference based on the encrypted data, while still preserving individual privacy. The algorithm works even for a colluding majority in the presence of cyberattack. Thus, a general framework for privacy-preserving statistical inference is established in the paper for sensitive data. In addition, their algorithm can be used to address another significant issue—privacy preservation for distributed statistical inference when data are allocated to different parties who are unwilling to share their own data with others.

Matching in Dynamic Markets

Motivated by marketplaces, such as kidney exchange, in "On Matching and Thickness in Heterogeneous Dynamic Markets," Ashlagi, Burq, Jaillet, and Manshadi study dynamic matching in an infinite-horizon stochastic market. While all agents are potentially compatible with each other, some are hard-to-match and others easy-to-match. Agents prefer to be matched as soon as possible and matches are formed either bilaterally or through chains. They focus on effects of both prioritization and market composition on waiting times. They find that when hard-to-match agents arrive less infrequently, bilateral matching is almost as efficient as chains, and prioritizing hard-to-match agents improves their waiting times. When hard-to-match agents arrive more frequently, chains are more efficient and prioritization has no impact. Further, they find that in heterogeneous markets and under bilateral matching, increasing hard-to-match agents' arrival rate has nonmonotone effects on waiting times. This is a fundamental difference with homogeneous markets, where increasing arrival rates always improves waiting time.

Technical Note—Optimizing Foreclosed Housing Acquisitions in Societal Response to Foreclosures

Thousands of nonprofit organizations throughout the United States are engaged in acquiring and redeveloping foreclosed housing units for resale or rental using available resources, such as government grants. Their purpose is the mitigation of the impact of foreclosures by reducing the number of foreclosed homes in neighborhoods, which in turn will result in home value appreciation, thus helping toward the stabilization of the overall economy. In "Optimizing Foreclosed Housing Acquisitions in Societal Response to Foreclosures," Solak, Bayram, Gumus, and Zhuo identify optimal acquisition policies that these nonprofits should implement while considering foreclosed properties for potential acquisition. Working closely with a nonprofit organization in Boston, MA, the authors conclude that in most cases, these organizations should not offer more than the asking price for a foreclosed property and should consider overbidding only when the total available funds are low. Overall, comparisons of the proposed policies with historical acquisition data suggest a potential improvement of approximately 20% in expected total impacts of the acquisitions on nearby property values.

Limiting Errors in Portfolio Construction

Building portfolios that balance risk and reward is a basic problem in finance. The standard approach is to use historical data to learn the relationships between assets and then build a portfolio that optimizes the risk-return trade-off based on the learnt relationships. Unfortunately, in practice, such portfolios are often outperformed even by the simple portfolio that buys equal amounts of each asset. Alternative methods resort to ad hoc modifications of the risk-reward objective, or imposing unjustified constraints on the portfolio. In "Portfolio Construction by Mitigating Error Amplification: The Bounded-Noise Portfolio," Zhao, Chakrabarti, and Muthuraman trace the root cause of poor performance to errors incurred in the learning phase that are then amplified by the optimization phase. Since it is impossible to remove learning errors entirely, they propose a new portfolio construction approach that provably limits the effect of such errors. Such portfolios are shown to achieve good performance in real-world settings.

Achieving Efficient Capacity Allocation Through Dynamic Allocation of Control Rights

In "Long-Term Partnership for Achieving Efficient Capacity Allocation," Liu, Lewis, Song, and Kuribko study a capacity provider and a group of independent buyers sharing a scarce but expensive-to-build capacity over a finite horizon. At the beginning of the horizon, the capacity provider invests in building capacity, and all members invest in increasing their own and possibly other members' market sizes. Each member then observes and updates its private and history-dependent information. The authors propose a multiperiod membership-type agreement as a series of single-period contracts with flexible terms renegotiated in each period. The agreement enforces ex post efficient capacity allocation and ex ante efficient investment.

The set of interpartner transfers in the agreement makes each member a residual claimant to the surplus it creates and hence induces truthful reports. The agreement is also budget balanced and voluntary. A new solution concept is hence developed for dynamic collective action mechanisms.

Designing Medicare Contracts Using Predictive Analytics

The Medicare Shared Savings Program (MSSP) was created under the Patient Protection and Affordable Care Act to incentivize providers to reduce costs while maintaining quality of care. In "Data-Driven Incentive Design in the Medicare Shared Savings Program," Aswani, Shen, and Siddiq propose a predictive analytics approach to redesigning the MSSP contract with the goal of better aligning incentives and improving financial outcomes from the MSSP. The authors leverage a data set containing the financial performance of providers enrolled in the MSSP, which together accounts for 7 million beneficiaries and over $70 billion in Medicare spending. It is estimated that introducing performance-based subsidies to the MSSP can boost Medicare savings by up to 40% without compromising provider participation in the MSSP. This work highlights the promise that data analytics hold in the design of incentive contracts.

Dynamic Control for Two Substitutable Products with Lead Times

In "Dynamic Optimal Policy for an Inventory System of Two Substitutable Products with Positive Replenishment Lead Time", Feng, Xu, and Zheng analyze a periodic-review inventory system of two products with positive replenishment lead times and finite capacities. The authors aim to characterize the structures of the optimal policies for product replenishment and substitution that minimizes the expected total discounted cost over a finite horizon. A major challenge of this problem is that one constraint does not meet the requirement of a sublattice, which makes the preservation of the structural property (namely, $L^\natural$-convexity) unclear. By an appropriate relaxation of constraints on the substitution quantity, the authors show that under some conditions, the objective functions for ordering and substitution decisions are $L^\natural$-convex, which results in monotone structures for the optimal ordering and substitution policies.

Dynamic Pricing and Inventory Management Through Demand Learning

Because of uncertainty in customer demand and lack of understanding in customer reactions to price changes, it is a challenge for many companies, such as manufacturers and retailers, to match supply and demand. Most of the models in the operations literature, however, have focused on the case in which the underlying customer demand information is known a priori, which is not true in many applications. In "Coordinating Pricing and Inventory Replenishment with Nonparametric Demand Learning," Chen, Chao, and Ahn develop a data-driven algorithm for pricing and inventory decisions that learns the demand and customer information from sales data on the fly, and they show that the profit generated from the algorithm converges to the clairvoyant optimal profit at the quickest possible rate.

On the Polyhedral Structure of Network Models with Unsplittable Node Flow

Network models for unit train scheduling contain the challenging combinatorial requirement that the incoming and outgoing flow of selected nodes cannot be split or merged. In "Network Models with Unsplittable Node Flows with Application to Unit Train Scheduling," Davarnia, Richard, Içyüz-Ay, and Taslimi study the polyhedral structure of such network models and develop various solution methods. In particular, they propose a perfect extended formulation capturing the "no-split no-merge" requirement for single-node relaxations. This formulation leads to a cut-generating linear program that can be used to tighten the linear programming relaxation in the space of original variables. With an aim to improving the efficiency of the cutting plane approach, the authors develop a class of valid inequalities through a polynomial lifting method. They evaluate the performance of the proposed solution methods through computational experiments that show that the developed techniques compare favorably with column generation approaches

How to Design a Product Line Under Cost Uncertainty

Changing regulations, macroeconomic factors, and technological innovations cause substantial uncertainty in the cost of providing services. In their study "Dynamic Selling Mechanisms for Product Differentiation and Learning," Keskin and Birge analyze the design of vertically differentiated product lines in the face of cost uncertainty. Keskin and Birge show that, although cost uncertainty makes a seller shrink its product line, the opportunity to learn about costs encourages expanding the product line. Keskin and Birge also show that imposing minimum quality standards on the product line can facilitate learning and consequently yield good profit performance.

Combating Overconservativeness in Data-driven Distributionally Robust Optimization

Distributionally robust optimization (DRO), a recent methodology to handle stochastic optimization problems in the presence of data, is based on robustifications of stochastic constraints that are enforced to hold over suitably constructed sets of underlying probability distributions. Although DRO enjoys valid feasibility guarantees, it often leads to overconservative solutions. The paper, "Recovering Best Statistical Guarantees via the Empirical Divergence-based Distributionally Robust Optimization," by Lam studies a calibration method for distributional sets to combat conservativeness via a new interpretation of DRO through the statistical angle of empirical likelihood and empirical processes. The proposed method targets achieving precise confidence level guarantees that lead to superior performances over previous approaches.

A New Variation of the Nonnegative Least-Squares Problem

In "The Weighted Nonnegative Least-Squares Problem with Implicitly Characterized Points," Fattahi, Dasu, and Ahmadi study a new variation of nonnegative least-squares problems when the discrete points are implicitly known and there are an exponentially large number of them. This problem can have applications in manufacturing, machine learning, clustering, pattern recognition, and high-dimensional statistics. They design an effective solution approach with a lower bound and establish the convergence rate of the lower and upper bounds.

Options' Forecasting in Mass Customization

In "Mass Customization and 'Forecasting Options' Penetration Rates Problem,'" Fattahi, Dasu, and Ahmadi study a new problem motivated by a large auto manufacturer that produces an extremely large number of feasible configurations (end products). Because it is impossible to forecast the demand of individual configurations, companies forecast on options level. The current forecasting approach ignores rules for selecting options, and as a result, forecasts are frequently infeasible, which results in excess inventories, shortages, and customer dissatisfaction. An approach is presented that finds the best feasible forecast. The theoretical properties of this approach are analyzed, and its effectiveness is examined on a set of real instances.

Plant Capacity and Attainability

In the contribution "Plant Capacity and Attainability: Exploration and Remedies," Kerstens, Sadeghi, and Van de Woestyne explore the way in which output-oriented plant capacity is estimated using a ratio of output-oriented efficiency measures to determine by how much outputs can be expanded if variable inputs can be adjusted at will. This traditional definition ignores the issue of attainability: this concept assumes that the amount of variable inputs needed to reach the maximum outputs is available at either the firm or the industry level. The authors argue a priori that neither is guaranteed. Therefore, they propose an attainable output-oriented plant capacity ratio that accounts for the limitations on the amounts of variable inputs available. If such amounts are not available, then an alternative input-oriented plant capacity concept may eventually be employed that does not suffer from the issue of attainability.

Reducing Sample Size Requirements in Clinical Trials

The decision of how to allocate subjects to treatment groups is of great importance in experimental clinical trials for novel investigational drugs, a multibillion-dollar industry. The statistical power of the experiment, its ability to detect a positive treatment effect when one exists, depends in part on the similarity of the groups in terms of measurable covariates that affect the treatment response. In "Covariate-Adaptive Optimization in Online Clinical Trials," Bertsimas, Korolko, and Weinstein present a novel algorithm for online allocation that leverages robust mixed-integer optimization. This new allocation approach shows great promise in computation experiments: matching or exceeding the statistical power of state-of-the-art randomization approaches. The authors suggest the algorithm could substantially reduce both the duration and operating costs of clinical trials in certain settings by achieving desired statistical power at a smaller sample size.

Efficiently Estimating Quantiles in Steady-State Simulations

In simulation-based evaluation of the performance or risk of a complex system, point and confidence-interval (CI) estimates are often required for a given quantile of a steady-state response—for example, the 95th percentile of an arbitrary job's cycle time in a proposed manufacturing system. In "Sequest: A Sequential Procedure for Estimating Quantiles in Steady-State Simulations," Alexopoulos, Goldsman, Mokashi, Tien, and Wilson develop an automated sequential procedure for estimating a given steady-state quantile of a simulation-generated output process. The final point estimator is approximately free of bias due to atypical initial conditions. The final CI estimator incorporates adjustments to compensate for dependence or skewness in the output process. Sequest determines a sample size sufficient to yield a CI satisfying user-specified requirements on the CI's coverage probability and its absolute or relative precision. In a set of difficult test problems, Sequest outperformed its competitors, requiring substantially smaller sample sizes while delivering closer conformance to the CI coverage and precision requirements.

Demand Frequency and Customer Prioritization

Service firms, such as amusement parks, ski resorts, and cultural institutions, sell season passes and memberships that provide prioritized entry, while also offering non-prioritized pay-per-use access. If all customers value their time similarly, it is not clear why some customers get priority service. In "Pricing and Prioritizing Time-Sensitive Customers with Heterogeneous Demand Rates," Afèche, Baron, Milner, and Roet-Green focus on the problem of designing price/lead-time menus and the corresponding priority policy for a profit-maximizing service provider serving customers with private information on their preferences. The key novelty is that the paper studies settings in which customers have demand for multiple uses with different frequencies. They show that prioritizing customers may be optimal when they differ in their demand rates, even if they have the same delay cost. Furthermore, doing so can generate significant profit gains compared with FIFO service.