



Operations Research

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

In This Issue

To cite this article:

(2020) In This Issue. Operations Research 68(2):iii-vi. <https://doi.org/10.1287/opre.2020.1982>

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 © 2020, 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.2020.1982>

Copyright: © 2020 INFORMS

How Reliable Are Data-Driven Inventory Decisions?

In recent years, there has been much interest in data-driven decision making. Although this can unlock tremendous value across industries, it is very important to remember that data-driven decisions are uncertain quantities with error bars associated with them. Despite the obvious importance, error bars of data-driven decisions have been underinvestigated. The paper “Confidence Intervals for Data-Driven Inventory Policies with Demand Censoring” bridges this gap in knowledge for inventory problems. Specifically, Ban derives approximate analytic formulas for confidence intervals of data-driven solutions to the classical dynamic inventory management problem of Scarf [Scarf H (1959b) The optimality of (s, S) policies in the dynamic inventory problem. Arrow KJ, Karlin S, Suppes P, eds. *Mathematical Methods in the Social Science* (Stanford University Press, Stanford, CA), 196–202.]. Both censored and uncensored data scenarios are considered, and the analyses extend to other commonly studied problems, such as the repeated newsvendor problem and base stock policy problem. Extensive computations on realistic simulated data validate the approximate analytic formulas, which are based on asymptotic theory, establishing their practical value.

Using Callback to Mitigate the Congestion in Call Centers

The callback option can be used as an instrument to effectively mitigate congestion due to temporary and unpredictable bursts of the arrivals in call centers. It works as follows: when the system is congested, the call center manager notifies the arriving customer and presents him with the option to hang up to be called back later within a reasonable time window. When the system congestion decreases, outbound calls are initiated for such customers. In “An Optimal Callback Policy for General Arrival Processes: A Pathwise Analysis,” Ata and Peng derive a simple yet effective policy that determines when the call centers should start offering the callback options and show that it is optimal in the fluid model. The simulation study conducted in the paper also shows that the proposed policies perform well.

A Network Fortification Model to Mitigate Interdiction Risk

Critical infrastructure is inevitably subject to the risk of interdictions, which may potentially cripple its ability to channel resources to fulfill demands at various destinations. Mitigating such risks are therefore important questions faced by the stakeholders. In “Mitigating Interdiction Risk with Fortification,” Hien, Sim, and Xu incorporate the

defender–attacker–defender sequential game model with stochastic programming to propose a preventive approach to achieving network resilience by allocating resources—to build additional capacity or protect extant capacity—and thus fortifying the network against interdiction/disruption. The unique fortification model accounts simultaneously for three components: adversarial disruption, stochastic demand, and the decision maker’s attitude toward risk, and it can be solved efficiently via a robust stochastic approximation approach.

Data-Driven Approach for Estimating Uncertainties Using Judgmental Forecasts with Expert Heterogeneity: Theory and Application to Agriculture Research and Development

In “Estimating Uncertainties Using Judgmental Forecasts with Expert Heterogeneity,” Bansal and Gutierrez develop a data-driven approach to aggregate point forecasts provided by multiple experts into probability distributions while accounting for expert heterogeneity. They consider the commonly encountered business problem of estimating probability distributions when historical data for the underlying uncertainty are not available or relevant. In such situations, firms typically ask a panel of experts to provide their forecasts. The authors develop a new approach to use the experts’ prior judgments to characterize their biases and the consistency of their judgments and use this information to aggregate their forward-looking judgments to obtain probability distributions. This approach quantifies the heterogeneity in expert judgments and incorporates this information into the aggregation protocol. It also establishes that bias and consistency in expert judgments are complements and not substitutes, and therefore choosing an expert with less bias is not necessarily better than choosing an expert with a stronger bias.

Learning about Experts in a Portfolio Choice Problem

The Black–Litterman model provides a framework for combining the forecasts of an equilibrium model and the forward-looking opinions of several experts in a portfolio allocation decision. In “A Generalized Black–Litterman Model,” Chen and Lim propose a generalization of the classical model that accounts for model misspecification and bias in the equilibrium and expert models and show how it can be calibrated using historical view and return data. More generally, this paper shows how multiple experts can be modeled as a Bayesian graphical model and estimated using historical data, which may be of interest in applications that involve the aggregation of expert opinions for the purpose of decision making.

In Highly Congested Networks, Is Selfishness the Problem?

Empirical studies in road networks reveal a fairly surprising property of congestion: When there is too much (or too little) traffic in the network, there is no difference between the best and the fairest traffic allocations (i.e., between the traffic assignment that optimizes the commuters' average travel time versus the one that no commuter would have any incentive to deviate from). In "When Is Selfish Routing Bad? The Price of Anarchy in Light and Heavy Traffic," Colini-Baldeschi, Cominetti, Mertikopoulos, and Scarsini give a theoretical justification to this empirical observation: For a large class of traffic inflow patterns and cost functions (including all polynomials), the gap between social optimality and equilibrium—the network's price of anarchy—converges to 1 in both heavy and light traffic, irrespective of the network topology and the number of origin/destination pairs in the network.

Design a Reward Mechanism to Promote Knowledge Sharing Among Smallholders

Recently, agricultural knowledge-sharing platforms emerged to facilitate knowledge exchange among smallholders so as to enhance farmers' productivity. Despite the active interactions among farmers on the platforms, the amount of knowledge being exchanged is dubious. In a competing environment, are farmers willing to share their knowledge with their peers? In "Knowledge Sharing and Learning Among Smallholders in Developing Economies: Implications, Incentives, and Reward Mechanisms," Xiao, Chen, and Tang find that farmers with valuable knowledge do not have economic incentives to share their knowledge with others. To encourage the sharing of valuable knowledge, we propose a simple "quota-based" reward mechanism that could entice farmers to share knowledge up to the efficient level, maximizing farmer welfare. What is more, the implementation cost of the mechanism can be arbitrarily small. The proposed mechanism can also be applied to many other competing environments beyond agriculture.

Pricing and Assortment Strategies with Product Exchanges

In "Pricing and Assortment Strategies with Product Exchanges," Wagner and Martínez-de-Albéniz identify product returns as a possible opportunity to sell additional goods. They model the consumer search process of trials, returns, and exchanges to derive product choice probabilities, which take the form of an attraction function where the weights are given by product attractiveness, price, and return fees. Using this tractable demand function, they then optimize assortment and pricing decisions, and they characterize the optimal policy: Retailers should hide restocking fees in their prices and offer higher prices to less popular items that appear later in the search sequence. They show that it is best for the retailer to pass all return costs to the consumers, which also maximizes social welfare.

Assignment Mechanisms Under Distributional Constraints

Assigning refugee families to different locations in a host country is a timely challenge due to the global refugee crisis.

In "Assignment Mechanisms Under Distributional Constraints," Ashlagi, Saberi, and Shamel study how to assign refugees to different locations while accommodating various distributional constraints. The social planner seeks to find a constrained efficient assignment with respect to refugees' preferences over different locations. As Pareto-efficient assignments may differ significantly in the number of assigned refugees, a key challenge is to match as many refugees as possible. The paper generalizes the well-known serial dictatorship (SD) and probabilistic serial (PS) mechanisms for assigning indivisible objects to agents to accommodate distributional constraints. The new generalizations of SD and PS maintain their desirable properties while satisfying the distributional constraints with a small error. Both mechanisms assign at least the same number of students as the optimum fractional assignment.

Optimal Pricing in Energy Markets

Motivated by the critical problem of pricing in electricity markets, in "Optimal Pricing in Markets with Nonconvex Costs," Azizan, Su, Krishnamurthy, and Wierman design a new pricing scheme, called equilibrium-constrained (EC) pricing, for jointly finding the optimal allocations and prices in such markets. The proposed method is applicable to a wide range of markets and allows using general parametric price functions. Furthermore, it is both economically and computationally more efficient than the existing methods. The name of this scheme stems from the fact that it directly imposes the equilibrium conditions as constraints in the optimization problem for finding the best allocations, as opposed to adjusting the prices later to make the allocations an equilibrium. The proposed framework also applies to the case of networked markets, which had not been considered in prior literature.

Optimal Signaling of Content Accuracy

Motivated by the proliferation of misinformation in online social media, in "Optimal Signaling of Content Accuracy: Engagement vs. Misinformation," Candogan and Drakopoulos explore how a platform should signal content accuracy in order to influence users' engagement decisions. They take into account the social aspect of engagement decisions, as well as the fact that the available content may contain inaccuracies. They consider two natural objectives for the platform: (i) maximizing user engagement and (ii) minimizing misinformation. They establish that in both cases, the optimal signaling mechanisms admit a simple threshold structure, where the platform recommends users to engage only with content whose inaccuracy level is below a threshold. In the first case, more "central" agents have higher thresholds, and they are exposed to less accurate content in expectation. In the latter case, the optimal thresholds are the same for all agents and independent of the network structure.

Chance-Constrained Covering Problems: A Bicriteria Approximation Algorithm

A chance-constrained optimization problem involves constraints with random data that can be violated with probability bounded from above by a prespecified small risk

parameter. Such constraints are used to model reliability requirements in a variety of application areas, such as finance and energy. Except under very special conditions, chance-constrained problems are extremely difficult. In “Bicriteria Approximation of Chance-Constrained Covering Problems,” Xie and Ahmed study a class of chance-constrained covering problems. They propose a bicriteria approximation scheme. Their scheme finds a solution whose violation probability may be larger than, but is within a constant factor of, the specified risk parameter and whose objective value is within a constant factor of the true optimal value. They extend the approximation results to a setting in which the underlying distribution of the constraint data are not known; that is, they consider distributionally robust chance-constrained covering problems and provide bicriteria approximation results.

Production Planning in a Multiechelon Supply Chain

In “Multiechelon Lot Sizing: New Complexities and Inequalities,” Zhao and Zhang study a multiechelon lot-sizing problem for a serial supply chain that consists of a production level and several transportation levels. Distinct from the existing literature on lot-sizing models, their model allows demands in the top production echelon as well as any following transportation echelons. Based on this general model, they answer an open question in the literature by showing that the multiechelon lot sizing with intermediate demands is NP-hard, provide efficient polynomial time algorithms for problems with a fixed number of echelons and develop a branch-and-cut algorithm using proposed inequalities to solve large multi-item multiechelon lot-sizing problems effectively.

Search for a Hidden Object

A person may struggle to find a lost key, a task force may deploy to locate a hidden bomb, and a rescue squad may race to find survivors. When searching for an object that is hidden in one of several possible locations, should one skim through each location so as to cover more ground quickly, or should one examine each location carefully to reduce the chance of overlook? In “Fast or Slow: Search in Discrete Locations with Two Search Modes,” Clarkson, Glazebrook, and Lin study the trade-off between speed and accuracy. They first establish theorems to rule out inferior search methods, and then develop effective strategies that deliver near-optimal performance for most practical problems. Their findings have important applications in rescue operations, counterterrorism, and cyber security, as search technology continues to advance today.

A Primal–Dual Lifting Scheme for Two-Stage Robust Optimization

Two-stage robust optimization problems, in which decisions are taken both in anticipation of and in response to the observation of unknown parameters from within an uncertainty set, are notoriously challenging. In the paper “A Primal–Dual Lifting Scheme for Two-Stage Robust

Optimization,” Georghiou, Tsoukalas, and Wiesemann develop convergent hierarchies of primal (conservative) and dual (progressive) bounds for these problems that trade off the competing goals of tractability and optimality. Based on these bounds, the authors propose a primal–dual lifting scheme that accommodates for generic polyhedral uncertainty sets, infeasible problem instances, and the absence of a relatively complete recourse. The incumbent solutions in each step of the algorithm afford rigorous error bounds, and they can be interpreted as piecewise affine decision rules. The performance of the algorithm is demonstrated on illustrative examples.

Bed Management in Intensive Care Units

Intensive care unit (ICU) beds are valuable resources that are typically in short supply and therefore their effective and efficient use is essential particularly during periods when patient demand is high. In “Allocation of Intensive Care Unit Beds in Periods of High Demand,” Ouyang, Argon, and Ziya provide insights into what kind of patient prioritization decisions are likely to improve patient health outcomes by analyzing stylized mathematical formulations that capture the fundamental trade-off involved in ICU bed management. They also propose simple policies that are likely to perform well in practice and test them with a simulation study. The findings suggest that even simple policies are likely to bring significant benefits, although more work is needed to investigate whether there could be benefits to using methods that aim to capture patient health conditions in a manner that is more precise than assumed in the paper.

Submodularity in Conic Quadratic Mixed 0–1 Optimization

Although submodularity has played an important role in developing efficient algorithms for numerous combinatorial optimization problems, its use in mixed 0–1 optimization has been rather limited. In “Submodularity in Conic Quadratic Mixed 0–1 Optimization,” Atamtürk and Gómez derive strong valid inequalities for conic quadratic mixed 0–1 optimization exploiting submodular structures observed in binary restrictions. They show that the exploited structure arises in numerous practical problems, including assortment optimization, chance constraints with uncorrelated as well as correlated random variables, queueing system design, binary linear fractional problems, and Sharpe ratio maximization. They demonstrate the effectiveness of their approach through computational experiments on a broad class of problems.

Taylored MDPs

Brownian approximations often capture structural relationships that are inaccessible in the original, “too” detailed dynamic programming problem. However, they have not been widely used as a way to reduce computational complexity. Inspired by perturbation techniques that were recently developed for the approximation of stationary queues by “Brownian queues,” Braverman, Gurvich, and Huang in

“On the Taylor Expansion of Value Functions” introduce a framework for approximate dynamic programming and apply it to discrete-time-and-state chains with countable action sets. The framework stipulates applying a second-order Taylor expansion to the value function, replacing the Bellman equation with one in continuous space and time in which the

transition matrix is reduced to its first and second moments. Bounds on the optimality gap are developed, and they can be viewed as a conceptual underpinning for the good performance of controls derived from Brownian approximations. The framework leads to an “aggregation” approach with performance guarantees.