



Operations Research

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

In This Issue

To cite this article:

(2017) In This Issue. Operations Research 65(4):ii-iv. <https://doi.org/10.1287/opre.2017.1664>

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 © 2017, 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

<http://dx.doi.org/10.1287/opre.2017.1664>

Copyright: © 2017 INFORMS

Learning Transmission Network Structure from Electricity Prices

In North American wholesale electricity markets, prices can vary substantially from location to location. The differences in price reflect transmission constraints and line losses which impede the flow of energy from less expensive generating units to demand. These markets can be difficult to analyze because key information about the transmission network is usually not available. In “Inverse Optimization for the Recovery of Market Structure from Market Outcomes,” J. Birge, A. Hortaçsu, and M. Pavlin show that by taking advantage of the fact that the electricity prices are dual prices of an economic dispatch problem solved by the market facilitator, important information about the transmission constraints can be learnt. The inverse optimization algorithm developed by the authors infers parameters which are shown to be useful for reconstructing the pricing mechanism and calculating locational residual demand derivatives using actual data from the Midcontinent ISO energy market.

Modeling Gas Long-Term Contracts in Europe

The drop in European gas demand lead to a massive renegotiation of European long-term contracts to include spot indexation. Various authors have argued that the transition away from long-term contracts is inevitable but the discussion has remained largely verbal and it appears that a portfolio of long and short-term supplies might well be necessary to best accommodate the uncertainties of the market. In “Modeling Gas Markets with Endogenous Long-Term Contracts,” I. Abada, A. Ehrenmann, and Y. Smeers offer a quantitative structuring of that problem by constructing a stochastic equilibrium model of a gas market that involves risk-averse suppliers and mid-streamers that operate through hubs and long-term contracts in an uncertain market. They show how the equilibrium of the market can shift towards more or less long-term supplies depending on the correlation of gas and oil prices and the relative importance of fixed and variable production cost of gas.

Intermediation in Online Advertising

In display advertising, as opposed to acquiring impressions directly from an ad exchange, advertisers often go through intermediaries. These intermediaries bid on behalf of the advertisers at the exchange, and deliver acquired impressions to them. For the service they provide, intermediaries get paid by the advertisers. In “Optimal Contracts for Intermediaries in Online Advertising,” S. Balseiro and O. Candogan study the optimal contracts the intermediaries should offer to advertisers. Under these contracts advertisers find it optimal

to truthfully reveal their type (budgets and targeting criteria), and the intermediary acquires impressions using a simple policy from the exchange. The optimal contracts are obtained by first characterizing the “performance levels” that can be achieved via different policies, and then following a duality based approach. The paper establishes that the presence of the intermediaries, in some settings, can improve the surplus of the advertiser as well as the profits of the ad exchange.

Solving the Dynamic Capacity Allocation Problem for Display Advertising

Display advertising is sold through two markets side-by-side. In the traditional guaranteed market, the publisher commits to deliver a prespecified number of impressions within a fixed time frame through a guaranteed contract. In the spot market, the publisher runs an auction to allocate the impressions every period, and the supply of heterogeneous impressions is highly uncertain and non-storable. Thus, the publisher must solve a dynamic capacity allocation problem of heterogeneous impressions across different contracts and markets, taking into account the uncertainties from both the demand and supply sides. In “Optimal Dynamic Auctions for Display Advertising,” Y-J Chen characterizes the precise tradeoffs between extracting the revenue from the spot markets, materializing the instantaneous benefit shared with the guaranteed advertisers, and releasing the pressure of paying the penalty related to guaranteed contracts. This paper also identifies the dual role of the publisher as a system designer and as a bidder on behalf of the guaranteed advertisers.

Optimizing Large Scale Flexible Production Models for Agribusinesses

What practical and technical challenges exist in exploiting operational flexibilities in practice? In “Product Portfolio Management with Production Flexibility in Agribusiness,” S. Bansal and M. Nagarajan discuss a large-scale production planning problem in the agribusiness domain with two industry characteristics: there exists an operational flexibility of backup production, and both the primary and backup production are subject to random yields. The operational flexibility significantly increases the problem complexity and size. Consequently, traditional stochastic programming based solutions are incapable of providing real-time decision support that was required at the partner firm. They develop a decomposition-aggregation based exact solution to the problem, which can be implemented in manager friendly computer applications. They also establish a new result for the benefit of the flexibility of backup production under random yields: When profit margins are high, backup production

reduces the average total production cost. Conversely, backup production can increase the average total production cost at low profit margins. The optimization protocol has been implemented at the firm, with substantial monetary benefits.

Simulating the SABR Model

The stochastic alpha-beta-rho (SABR) model has gained great popularity in the financial industry as it can provide good fits to various types of implied volatility curves observed in the marketplace. Nonetheless, no analytical solution to the SABR model exists that can be simulated directly. In “Exact Simulation of the SABR Model,” N. Cai, Y. Song, and N. Chen propose an exact simulation method for the forward price and the volatility of the SABR model in two special but practically interesting cases. In more general cases, the exact simulation method becomes a semi-exact one that is still quite accurate when the time horizon is not long. For long time horizons, they develop a piecewise semi-exact simulation scheme that reduces the biases significantly. For European option pricing under the SABR model, they propose a conditional simulation method that reduces the variance of the plain simulation substantially.

Reserving Capacity in Blocks

When a buyer is uncertain about demand and pays in advance to reserve capacity, it is attractive to have a portfolio of suppliers, including some with a high execution price if their reservation price is low enough. Capacity from the high-priced suppliers will only be used if demand turns out to be very high. In “Supplier Competition with Option Contracts for Discrete Blocks of Capacity,” E. J. Anderson, B. Chen, and L. Shao consider suppliers with different characteristics who compete to supply blocks of capacity. The discreteness of the problem requires a combinatorial approach. The authors show that, when there is a requirement for the buyer to purchase a whole block of capacity, suppliers should offer capacity at execution prices that match their costs, so that the suppliers only make profits from the reservation part of their bids. Moreover, at the equilibrium the buyer chooses the welfare maximizing set of blocks.

How Should a Fashion Retail Chain “Test the Waters” for Demand Prior to the Selling Season?

In a merchandise test, a retailer deploys limited quantities of a new product to its stores in order to learn about demand prior to the main selling season. In “Optimal Merchandise Testing with Limited Inventory,” L. Chen, A. J. Mersereau, and Z. Wang examine the optimal allocation of inventory to stores in a merchandise test. A key tradeoff is between the quantity of stores included in the test and the quality of the demand observations obtained, which can be compromised by demand censoring due to inventory stockouts. The authors find that the retailer’s visibility into the timing of each sales transaction has a pivotal impact on the optimal inventory allocation policy, and they design heuristics that perform near-optimally across a variety of problem instances.

On the Structure of Arc Routing Problems—A Heuristic Perspective

Arc routing is notably different from node routing, due to the necessity of considering precise network information and

deciding on traversal orientations for the services on edges. In node routing algorithms, decisions related to the paths between deliveries are usually concealed within the distance matrix computation. In arc routing applications, especially in urban environments, these decisions depend on service orientations as well as on possible intersection delays and turn restrictions. In “Node, Edge, Arc Routing and Turn Penalties: Multiple Problems—One Neighborhood Extension,” T. Vidal investigates a structural neighborhood decomposition for arc routing heuristics. It shows that memory structures, bidirectional dynamic programming, and lower bounds allow to search a neighborhood involving classical moves on service sequences with optimal service orientations and path choices in amortized $O(1)$ time per move. The approach is generalized to problem variants with “service modes,” and integrated into two classical metaheuristic frameworks, leading to solutions of remarkable quality.

Technology Adoption and Information Gathering with Risk Aversion

Consumers and firms considering adopting a new technology often face uncertainty: Will the benefits of the technology outweigh its costs? Should they adopt immediately or gather more information? How does risk aversion affect these decisions? In “Risk Aversion, Information Acquisition and Technology Adoption,” J. E. Smith and C. Ulu develop a dynamic programming model of this problem where the decision maker’s (DM’s) beliefs are represented by a probability distribution and her risk attitude by a utility function. They establish a number of theoretical results that provide insights into the nature of optimal policies for information gathering, highlighting the effects of risk aversion. In particular, they show that optimal policies for a risk averse DM may not satisfy natural monotonicity properties, if the DM is very risk averse. Establishing these results requires the use of some novel proof techniques that may prove useful in other contexts.

Developing a Monotone Improving Solution Algorithm for Sequential Decision-Making Problems with a Countably Infinite Number of States

For Markov decision processes (MDPs) with a finite number of states, policy iteration, one of the popular solution methods, finds a sequence of policies improving monotonically in value and converging to optimality. However, for MDPs with a countably infinite number of states, existing solution methods fail to improve monotonically. In “Simplex Algorithm for Countable-State Discounted Markov Decision Processes,” I. Lee, M. Epelman, E. Romeijn, and R. Smith introduce the first algorithm for such MDPs that converges to optimality and improves monotonically, and which is implementable in the sense that each of its iterations requires only a finite amount of data and computation. The algorithm is a simplex-type algorithm for solving a countably-infinite linear programming (CILP) formulation of the MDP. The research also extends our knowledge about infinite LPs by offering an algorithm that solves a new class of CILPs for which existing simplex-like algorithms are inapplicable.

Traveling Salesman Problem: Breaking a Thirty-Year-Old Barrier

The traveling salesman problem (TSP) is one of the most celebrated and intensively studied problems in combinatorial optimization. Throughout time, it has found applications in logistics, planning, manufacturing and testing microchips, as well as DNA sequencing. It has also been used as a benchmark for the performance of novel ideas in optimization such as ant colony optimization, memetic algorithms, or Adleman's molecular computing. In "An $O(\log n / \log \log n)$ -Approximation Algorithm for the Asymmetric Traveling Salesman Problem," A. Asadpour, M. Goemans, A. Mądry, S. Oveis Gharan, and A. Saberi establish a new connection between the approximability of the general (asymmetric) TSP and the notion of "thin trees", and also employ "maximum entropy rounding" (a novel method of dependent randomized rounding of LP relaxations of combinatorial optimization problems) in order to provide an $O(\log n / \log \log n)$ -approximation algorithm for the asymmetric TSP, hence breaking the almost thirty-year-old barrier of $\Theta(\log n)$ -approximation ratio stemming from the work of Frieze et al. (1982).

Simulation Experiment Design for Estimating the Input–Output Relationship

A basic experiment with a simulation model runs the model with the value of the model's input fixed, to estimate an expected value of model output given this input value. However, sometimes it is desired to estimate the function that maps input values to expected values of model output. To achieve good accuracy in estimating this input-output function, traditional experiment designs can require a very high computational cost. In "Multilevel Monte Carlo Metamodeling," I. Rosenbaum and J. Staum show how multilevel experiment designs reduce the computational cost required to achieve desired accuracy.

Using Redundancy to Reduce Response Time: Exact Analysis

Redundancy is an important new strategy for reducing response time in multiserver systems. The idea is to create multiple copies of each job and wait for only the first copy to complete service. While redundancy has been of considerable interest to the theoretical community in the past several years, there is very little analysis of response time in systems with redundancy. In "Redundancy- d : The Power of d Choices for Redundancy," K. Gardner, M. Harchol-Balter, A. Scheller-Wolf, M. Velednitsky, and S. Zbarsky study a particular dispatching policy called the Redundancy- d policy, under which each job creates d copies of itself and sends these copies to d randomly selected queues. They derive the first exact, closed-form expressions for both the mean and the full distribution of response time under Redundancy- d . They use their analysis to study the effect of d on response time.

Production with Risk Hedging

Many production firms have observed demands for their products being significantly impacted by the price of commodities or the state of the economy in general. This motivates the introduction of a hedging strategy, in addition to the usual production quantity decision, to enhance risk mitigation. In "Production with Risk Hedging—Optimal Policy and Efficient Frontier," L. Wang and D. Yao model demand as a stochastic process with two components: in addition to the usual Gaussian component reflecting demand volatility, there is a drift component taking the form of a function of a financial asset price. With this new demand model and under a mean-variance problem formulation, the one-time production quantity decision and a real-time risk-hedging strategy are jointly optimized. In addition, the risk-return profile is completely characterized via an efficient frontier, and risk reduction achieved by the hedging strategy is explicitly quantified.