Operations Research

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

In This Issue

To cite this article:
(2018) In This Issue. Operations Research 66(5):ii-iv. https://doi.org/10.1287/opre.2018.1809

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

OPERATIONS RESEARCH
Vol. 66, No. 5, September–October 2018, pp. ii–iv
ISSN 0030-364X (print), ISSN 1526-5463 (online)

In This Issue

https://doi.org/10.1287/opre.2018.1809

Dynamic Modeling of Electricity Price Risk

Accurate representation of price risk is a crucial element in many short-term decisions in the electricity industry. However, electricity price distributions exhibit a variety of shapes, hour by hour, and their dynamics have become even more complicated with the increasing penetration of renewable energy resources. As a consequence, general analytical models have tended to be oversimplifications and inaccurate for practice. In "A Stochastic Latent Moment Model for Electricity Price Formation," Gianfreda and Bunn show that a general four-parameter density specification is sufficiently flexible to fit the wide range of shapes displayed in the German hourly prices during 2006–2016. Using generalized linear multivariate estimation for the four moments and relating these to the key drivers, such as wind and solar, provide a transparent, credible, and sufficiently flexible model to capture the dynamic switches in the upside and downside risks. The approach was extensively back-tested.

Asymmetries in Networks Cause Inefficiencies

Generating a surplus may require an agreement among three or more players, but some of them may not be able to cooperate directly for informational, legal, political, or physical reasons. To study the role of such cooperation restrictions, in "Multilateral Bargaining in Networks: On the Prevalence of Inefficiencies," Lee introduces a noncooperative bargaining model in which players can bargain only with their directly connected players in a given network. As players in different positions have different bargaining powers, they can take a better position in future bargaining by strategically excluding some of their neighbors from the bargaining partners. Such strategic behaviors cause delay in reaching an agreement. The paper characterizes a condition for efficiency—an efficient equilibrium exists for all discount factors if and only if the network is either complete or circular—in which the players are all symmetric in terms of their position not only in the initial network but also in any induced network.

Launching Social Goods Sequentially and Pricing Them Wisely

Network effects are widely documented in various products, as a customer's consumption decision can have direct impacts on other customers' satisfaction about the same product. Firms in the markets of such "social goods" have strived to experiment with novel selling schemes to create and capture values from the customers' social interactions. As a common practice in the entertainment industry, a vital strategy is sequential launch: potential customers for the same product are served at different times. This is precisely the business tactic that Zhou and Chen address in "Optimal Pricing with Sequential Consumption in Networks." In their paper, Zhou and Chen show that the optimal pricing under arbitrary launch sequence is independent of customers' network positions, the launch sequence, and the underlying social interaction relations among customers; additionally, the optimal sequence is a chain structure. Moreover, Zhou and Chen establish the pecking order by which the firm shall approach customers following the sequence of descending valuation-cost margins.

Generating Random Networks Without Short Cycles

Random graph generation is an important tool for studying large, complex networks. Despite an abundance of random graph models, constructing models with application-driven constraints is poorly understood. To advance this area, in "Generating Random Graphs Without Short Cycles," Bayati, Montanari, and Saberi focus on a stylized family of graphs and propose an algorithm for randomly generating them. The authors propose the first polynomial-time algorithm for randomly generating such graphs that works by sequentially adding edges via a carefully constructed sampling scheme and provide a new method for analyzing such algorithms.

Customer Segmentation in the Big Data Era

In "A Model-Based Embedding Technique for Segmenting Customers," Jagabathula, Subramanian, and Venkataraman revisit the classical problem of customer segmentation but in the context of the current big data era, which poses new and unique challenges in identifying the segments. For instance, the diversity of the types of customer actions (e.g., clicks, purchases, ratings, etc.) captured as well as the high sparsity (customers interact with typically a tiny fraction of the item universe) make it difficult to assess similarity between different customers. To overcome these challenges, the authors propose a model-based embedding technique that extends extant clustering techniques in machine learning to handle categorical observations from diverse data sources that have (many) missing entries. Using a real-world case study on eBay data, the authors show that their proposed segmentation technique scales to large datasets and can be used to generate high-quality, personalized recommendations for the different segments.

Characterizing Acceptance Sets That Are Surplus-Invariant, Law-Invariant, and Conic

In "Surplus-Invariant, Law-Invariant, and Conic Acceptance Sets Must be the Sets Induced by Value-at-Risk," Xue Dong He and Xianhua Peng provide complete characterization of acceptance sets that are surplus-invariant, law-invariant, and conic. Acceptance sets of capital positions are specified by regulators for testing firms' capital adequacy. Because the test aims to protect firms' liability holders, it is desirable for the acceptance sets to satisfy surplus-invariance, i.e., whether

a capital position is acceptable does not depend on the surplus for the equity holder (i.e., the positive part of the position). They prove that any surplus-invariant, law-invariant, and conic acceptance set must be the set of capital positions whose value-at-risk at a given level is less than zero. They also show that the result still holds if conicity is replaced by numeraire-invariance, a property stipulating that whether a firm passes the test should not depend on the currency used to denominate its assets.

Managing Perishable Inventory Under Censor Demand

Perishable products, such as fresh food, pharmaceuticals, and blood banks, are ubiquitous and indispensable parts of our society. However, managing perishable inventory systems is notoriously difficult in both theory and computation. The existing literature on this topic assumes that the stochastic demand processes are given as an input to the models and the inventory replenishment decisions are made with full knowledge of the demand distribution. However, in practice, the underlying demand distribution may not be available to the firm a priori. This raises an important research question as to how to devise an efficient and effective learning algorithm that only uses the observed sales data collected over time, so as to minimize the total expected cost. In "Perishable Inventory Problems: Convexity Results for Base-Stock Policies and Learning Algorithms Under Censored Demand," Zhang, Chao, and Shi develop nonparametric learning algorithms for perishable stochastic inventory systems with censored demand and establish a tight square-root convergence rate. The authors also demonstrate the efficacy of the proposed algorithms through an extensive numerical study.

Taking Advance Purchase Commitments to the Operations-Finance Interface

Literature on operations management already identifies many benefits of advance purchase commitments. In "Purchase Order Financing: Credit, Commitment, and Supply Chain Consequences," Reindorp, Tanrisever, and Lange extend this topic to the operations-finance interface: a purchase commitment can serve to mitigate the operational impact of capital market frictions. A commitment removes uncertainty about a supplier's demand prospects, extending her access to capital. The authors quantify the relevant frictions, determine the buyer's optimal commitment, and show how operations and profits are conditioned by the financial context. Key parameters are the supplier's ex ante credit limit, her informational transparency, and the demand characteristics of the final market. Strategic interaction between buyer and supplier entails that the latter will sometimes benefit from a lower ex ante credit limit or lower informational transparency. This study provides a novel perspective on capital market frictions in supply chains.

Merchant Trading of Energy in Networks of Storage and Transport Assets

Merchant companies trade energy in wholesale markets connected by networks of storage and transport assets that they operate. In "Merchant Energy Trading in a Network," Nadarajah and Secomandi broaden to a multistorage-asset setting the applicability of approximate dynamic programming methods developed for more restrictive situations and, in part, available in commercial software. The authors' research relies critically on dynamic programming duality, linear optimization, Monte Carlo simulation, and regression. Use of the dual bound demonstrates that the proposed heuristics lead to effective asset-operating and energy-trading policies for realistic instances. In addition to the research community, these techniques and finding are relevant to both merchant energy trading firms and developers of software that supports their operations.

On the Surprising Relevance of Greedy-Like Algorithms in Choice-Based Inventory Management

A fundamental issue in operations is to allocate limited inventory across a variety of substitutable products. In such settings, customers' choices respond dynamically to product stock-outs. Despite significant attention in operations management, the antecedent literature bears few positive results about the computational aspects of choice-based inventory optimization models. In "Greedy-Like Algorithms for Dynamic Assortment Optimization Under Multinomial Logit Preferences," Aouad, Levi, and Segev propose the first provably good algorithms under the popular multinomial logit choice model. Through structural transformations, the authors show that the problem is amenable to submodular-like maximization, where a weak notion of submodularity comes into play. While the proposed algorithms rely on greedy procedures, the approach draws structural differences vis-à-vis the standard newsvendor model that overlooks product substitution. The article establishes a constant-factor, worst-case guarantee under a broad class of demand distributions and demonstrates practical improvements against existing methods in terms of performance and speed.

A Near-Optimal Algorithm for Multidimensional Binary Search

In "Multidimensional Binary Search for Contextual Decision Making," Lobel, Paes Leme, and Vladu consider a version of multidimensional binary search that emerges from contextual decision problems. Nature chooses a state of the world and a sequence of directions. A player is given access to the directions but not the state. The player wishes to predict the inner product of the state and the direction at each period and receives as feedback only whether the prediction was too high or too low. The authors construct an algorithm for solving this problem, and it proves its optimality with the time horizon and near optimality (up to a logarithmic factor) with respect to the dimension of the state space. Application areas of this work include contextual pricing and personalized medicine.

Achieving Global Collaboration by Carrot and Stick

In an unbalanced cooperative game, the grand cooperation will not be stable if the game has an empty core, implying that the benefit generated from a global collaboration is not high enough to be shared by all players. In "Simultaneous Penalization and Subsidization for Stabilizing Grand Cooperation," Liu, Qi, and Xu propose a new scheme to stabilize the grand cooperation for an unbalanced game. The scheme integrates two existing methods, where an outside party can provide some external resource to subsidize the grand cooperation and, at the same time, charge a penalty to any subset of players who may deviate from the grand cooperation. This carrot-and-stick scheme allows the outside party to make a

trade-off between the subsidization and penalty levels, hence providing high flexibility to the outside party. The authors propose models to formalize the scheme and develop algorithms for computation.

Dealing with Degeneracy in the Primal Simplex Method

In the past decade, we have seen variants of the primal simplex algorithm that aim at better dealing with degeneracy. Classical pivoting rules—selecting a variable to enter the basis—use reduced costs that are computed from a fixed dual solution in each iteration. This measures the potential improvement in the objective function but not the actual step size. In "Vector Space Decomposition for Solving Large-Scale Linear Programs," Gauthier, Desrosiers, and Lübbecke suggest a framework in which only part of the dual solution is kept fixed and the remainder belongs to an optimization problem. The rationale behind aiming at (local) dual optimality is to ensure (local) primal feasibility in the resulting pivot step, that is, avoiding a zero-step direction. Several diverse algorithms are natural concretizations of the framework, not only the primal simplex algorithm but, for example, also the minimum mean cycle-canceling algorithm for network flow problems. The authors believe that their framework will be particularly useful in the context of column generation.

Risk-Averse, Two-Stage Stochastic Program with Distributional Ambiguity

Starting from a given set of historical data, in "Risk-Averse, Two-Stage Stochastic Program with Distributional Ambiguity," Jiang and Guan formulate a risk-averse, two-stage stochastic program from the perspective of distributional robustness by hedging against the worst-case distribution within an ambiguity set and considering the corresponding expected total cost. An equivalent reformulation is successfully derived and indicates that each worst-case expectation over an L1-norm–based ambiguity set reduces to a convex combination of a Conditional Value at Risk and an essential supremum. This reformulation explicitly shows how additional data can help reduce the conservatism of the problem from the traditional two-stage robust optimization to the traditional risk-neutral, two-stage stochastic program. Further studies and insights are also explored for ambiguity sets based on L∞-norm, joint L1- and L∞-norm, and joint L1-norm and the first two moments. The convergence results show that the ambiguity aversion vanishes as the data size increases to infinity.

A New and Efficient Method for Simulation Calibration

Calibration poses a considerable challenge to the successful implementation of any simulation model. Essentially, given multiple models, or multiple settings for a single model, we need to identify the one that provides the closest match (on average) to the real-world system being simulated, using only the simulation output to guide our decision. Although this problem belongs to the broad domain of simulation-based optimization, standard methodological approaches from this field may underperform or fail entirely when applied to the calibration problem. In "The Local Time Method for Targeting and Selection," Ryzhov proposes a new and fundamentally different approach to the calibration problem. Through an analogy to the local time of a Brownian motion (a classical probabilistic structure that measures proximity to a given target), the author derives an algorithm that is highly computationally efficient, provably achieves fast convergence rates, and exhibits superior and robust numerical performance.

There's No Place Like Ohm

Many service systems process customers in first-come, first-served (FCFS) fashion, largely because this is perceived widely as a fair method for allocating customers to servers. However, what happens with FCFS when a service system has multiple customer and server types, but customers of a given type can only be processed by servers from a subset of all server types? Such systems abound in settings as diverse as public housing, organ transplantation, and even the adoption of children by prospective parents. An important and natural question in such situations is: for customers of each type, what fraction of customers will receive service from servers of each type? While this is a deceptively simple problem to pose, it is extremely difficult to solve. Remarkably, in an elegant paper, Ivo Adan and Gideon Weiss derived an exact solution to this problem. Unfortunately, with the exception of problems with only a few customer and server types, their solution proves prohibitively difficult to compute. In "Approximating the FCFS Stochastic Matching Model with Ohm's Law," Fazel-Zarandi and Kaplan present an approximation based on a physical analogy with electric circuits: treating customer and server types as voltage sources and sinks, the matching rates of customers to servers are viewed as the current flowing through an electric circuit network, with feasible customer/server type pairings contributing links to the circuit graph. The problem of determining the distribution of current over a network with given current in- and outflows is governed by repeated application of Ohm's law, which dictates that the current flow along any link equals the voltage drop along that link divided by the link resistance. Recasting the FCFS matching problem in terms of Ohm's law, the authors deduce a simple system of linear equations, one for each customer and server type. For certain customer/server type matching topologies, these equations yield the exact solution. Via simulation and numerical study, the authors show that the approximation is quite accurate. Given the simplicity of this approach, the authors have indeed shown that when it comes to the FCFS stochastic matching problem, there is indeed no place like Ohm!

The Impact of Customers' Contact History on Their Abandonment Behavior in Call Centers

In "Customer Learning in Call Centers from Previous Waiting Experiences," Emadi and Swaminathan use a series of Cox regressions to show that callers' abandonment behavior may differ based on their contact history and changes across their different contacts. As a result, callers' beliefs about the waiting time distribution may not match the actual distribution in the call center, and the equilibrium condition in the rational expectation equilibrium assumption may not hold. To model customers' learning process, the authors use a structural estimation approach in a Bayesian learning framework. The estimation results show that new callers who do not have any prior experience with the call center are optimistic about their delays. They also show that the Bayesian learning model not only has a better fit to the data set compared with the rational expectation equilibrium model but also outperforms it in out-of-sample tests.