



Operations Research

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

In This Issue

To cite this article:

(2015) In This Issue. Operations Research 63(6):ii-iv. <https://doi.org/10.1287/opre.2015.1467>

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 © 2015, 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

Data-Driven Tenure Decisions

Tenure decisions, key decisions in academic institutions, are primarily based on subjective assessments of candidates. In “Tenure Analytics: Models for Predicting Research Impact,” D. Bertsimas, E. Brynjolfsson, S. Reichman, and J. Silberholz study whether early-career publication data available to tenure committees can predict later-career performance. Using a large-scale bibliometric database from the field of operations research, they train statistical models using publication data from scholars’ first five years of publication to predict long-term publication outcomes and whether scholars will become INFORMS Fellows. Using a dataset of tenure decisions at top operations research programs, the authors show that the proposed models, constrained to tenure the same number of candidates as tenure committees did, made a different decision in 30% of cases, on average selecting scholars with better long-term publication outcomes. These results show that data-driven tools could complement the existing tenure decision-making process in academia.

How Are Prices Set at Peak Demands?

Peak producers of non-storable products, such as electricity, provide crucial flexible operating capacity to respond to infrequent and transient high demand periods. Despite often having significant price-setting power, they need to profit from a limited number of pricing decisions in order to meet financial targets for viability. In “Dynamic Pricing of Peak Production,” H. Peura and D. Bunn show how revenue uncertainty and financial objectives together create dynamic patterns in peak pricing. Furthermore, depending on market conditions, internal financial targets may either benefit or hurt the profitability of the firm. In the electricity context, these insights are useful not only for understanding how purely energy-based revenues can sustain the financial viability of peaking plants, and the dynamic emergence of price spikes, but also in providing the underlying process for pricing derivative contracts that policy makers may encourage or offer for resource adequacy.

Planning Wind Energy Storage and Transmission

Regions with abundant wind resources usually have no ready access to the existing electric grid. However, building transmission lines that instantaneously deliver all geographically distributed wind energy can be costly. Energy storage (ES) systems can help reduce the cost of bridging wind farms and grid, and mitigate the intermittency of wind outputs. In “Joint Planning of Energy Storage and Transmission for Wind Energy Generation,” W. Qi, Y. Liang, and Z. Max Shen develop new models and solution approaches for transmission network planning with co-location of ES systems. These models determine the sizes and sites of ES systems along with transmission lines configuration. The authors find that, in most cases, using even small-sized ES systems can significantly reduce

the total expected cost, but their marginal values diminish faster than those of the transmission lines. The authors also identify the major bottleneck cost factors for different forms of ES technologies.

What Drives the Accuracy of the Sample Average Approximation?

Sample average approximation (SAA) is a widely used method for distribution-free stochastic optimization when the decision-maker has access to a sample drawn from the unknown distribution. In many settings however, there might not be enough data or data is generated from an expensive procedure. An example is inventory planning for a new product launch. Hence, it is important to understand the SAA solution’s underlying accuracy. In “The Data-Driven Newsvendor Problem: New Bounds and Insights,” R. Levi, G. Perakis, and J. Uichanco find that existing theoretical guarantees for SAA applied to the newsvendor problem significantly underestimate the empirical accuracy. In a new analysis, they prove SAA accuracy to be linked to the distribution’s weighted mean spread. Their work also develops a tight distribution-free bound on the weighted mean spread of log-concave distributions. Hence, they prove a theoretical guarantee for the SAA that agrees with empirical accuracy when applied to log-concave distributions.

Pricing for Customers Who Are Willing to Wait

Customers, if presented with a price they judge to be too high for a particular product, may wait for the price to fall to an acceptable level rather than simply deciding not to purchase the product. How should a seller of the product set prices when facing such “patient” customers? Y. Liu and W. L. Cooper address this question in “Optimal Dynamic Pricing with Patient Customers.” They show that the seller can maximize its long-run average revenue by implementing a cyclic pricing policy, wherein a fixed pattern (a “cycle”) of prices is repeated over time. They establish that within an individual cycle, the prices should decrease and they bound the length of such a cycle in terms of the degree of customers’ patience and the size of the set of allowable prices. These results render an otherwise intractable pricing optimization problem amenable to solution with a simple computational algorithm.

Joint Inventory and Markdown Price Decisions when Strategic Consumers Learn

When selling fashionable products, retailers often mark down prices towards the end of a selling season to clear out inventory. In a repeat-purchase market where consumers interact with a retailer season after season, such practice trains strategic consumers to wait for heavy discounts in the future. Particularly, strategic consumers will learn from the retailer’s past prices, form their references of markdown prices, and time their purchases. In “The Reference

Effects on a Retailer's Dynamic Pricing and Inventory Strategies with Strategic Consumers," S. Wu, Q. Liu, and R. Q. Zhang characterize how the retailer's inventory decision complements the markdown price decision to optimally influence the reference price held by strategic consumers. They find that the reference price follows a mean reverting behavior under some conditions; that is, whenever the reference deviates from the mean, it is expected to revert back to the mean in the future. Nevertheless, the reference price may also oscillate among different reverting means under some market conditions.

Near-Optimal Order Fulfillment Heuristic for Online Retailer

There is no doubt that the popularity of online selling will continue to rise. For online retailer with multiple warehouses and multiple customer regions, this opportunity also creates a new technical challenge: How should the seller fulfill the order, should he split the order by sending different items from different warehouses? This is a fundamental issue faced by most online retailers; and yet, little has been done in the literature. In "An LP-Based Correlated Rounding Scheme for Multi-Item Ecommerce Order Fulfillment," S. Jasin and A. Sinha introduce an LP-based heuristic that first decouples the fulfillment decision for different items within the same order (in order to reduce the curse of dimensionality) and then inject an "artificial dependency" to consolidate as many items as possible to reduce total shipping costs. Through intensive numerical experiments, the authors show that the proposed heuristic performs significantly better than the commonly used myopic fulfillment policy which simply fulfills an item from the nearest warehouse.

Reducing the Size of Linear Programming Formulations of Approximate Dynamic Programming

In "Reductions of Approximate Linear Programs for Network Revenue Management," T. Vossen and D. Zhang show that the approximate linear programs for solving the high-dimensional dynamic programming problem from network revenue management can be dramatically reduced in size for two important classes of approximation architecture, the affine approximation and the separable piecewise linear approximation. Their work paved the way for more efficient solution procedures, as approximate linear programs that previously rely on specialized and computationally intensive algorithms to solve can now be directly fed to commercial linear programming solvers. Computational tests show orders of magnitude improvement in solution speed. Theoretical development in the paper is a novel application of the well-known Dantzig-Wolfe decomposition principle developed in 1960s. The generality of the Dantzig-Wolfe decomposition theory makes it promising to develop parallel results in other application context.

An Investigation of Dynamic Upgrading Decisions

Upgrading, where higher-quality products can be used to satisfy demand for a lower-quality product that is sold out, is a widely adopted strategy in practice in both manufacturing and service industries. It is particularly critical for settings in which no inventory replenishment is possible due to long ordering lead time. Yet, upgrading leads to challenging optimization problems as firms need to make delicate allocation tradeoffs both within the same time period and across different periods before observing future demand. In "Dynamic Capacity Management with General Upgrading,"

Y. Yu, X. Chen, and F. Zhang study a general dynamic upgrading problem and show that a simple, sequential upgrading policy is optimal. Based on the sequential policy, an efficient and accurate heuristic is developed for solving the dynamic upgrading problem. Through numerical analysis, they find that deriving the optimal upgrading policy should receive a higher priority than calculating the optimal initial capacity.

Flight Scheduling and Airport Operations Are Interdependent, and Can Be Jointly Optimized

Airport congestion mitigation interventions are generally concerned with improving airport operations at the tactical level (over each day of operations) or managing demand to limit over—capacity scheduling at the strategic level (months in advance of the day of operations). In "An Integrated Scheduling and Operations Approach to Airport Congestion Mitigation," A. Jacquillat and A. Odoni develop a new optimization framework that addresses these two problems in a unified framework that jointly optimizes tactical airport capacity utilization and strategic flight scheduling interventions. It relies on an original modeling architecture and solution algorithm that integrate (i) a Stochastic Queuing Model of airport congestion, (ii) a Dynamic Programming model of airport capacity utilization, and (iii) an Integer Programming model of flight scheduling. Results suggest that large delay reductions could be achieved through limited scheduling interventions and efficient airport operations, and that integrating these two sets of decisions performs better than treating them separately.

Non-Additive Ordinal Functions

Ordinal functions are useful for ranking complex outcomes. They are widely used in decision analysis and economics, but also in psychology and physics. Ordinal functions are usually assumed to be additive in their variables because of a dearth of alternative approaches. Additivity requires that preferences over a subset of attributes do not depend on the fixed levels of the others. In "Ordinal One-Switch Utility Functions," A. Abbas and D. Bell show how to proceed if additivity does not hold. They propose functional forms that, while not additive, do satisfy their "one-switch rule," which requires that preferences on a subset can change, but at most once, as the level of other attributes vary. The paper also provides tests that can check whether a given function satisfies this property.

Present Values to Justify Intertemporal Optimization Functions

In optimization problems, preference conditions serve to justify the use of particular goal functions. For instance, constant discounting of future payoffs is justified by Koopmans' stationarity. In "Discounted Utility and Present Value—A Close Relation," H. Bleichrodt, U. Keskin, K. Rohde, V. Spinu, and P. Wakker introduce new preference conditions for discounting that are more intuitive, mathematically more general, and shorter to write, than existing ones. Their conditions impose independence of present values from sets of other variables. For example, general (including hyperbolic) discounted utility is appropriate if the present value of an extra future outcome is independent of the outcomes in all other periods. Their conditions can be tested qualitatively, as all preference conditions can be, but, unlike others, their conditions can also be directly tested quantitatively, in regressions and ANOVAs for instance. The authors extend their conditions from intertemporal

optimization to other contexts, such as uncertainty. They then give the most efficient justification presently available of a market pricing system that is arbitrage-free and time consistent.

Stochastic Optimization for Integrating Staffing and Scheduling with a Healthcare Application

Workforce planning under demand uncertainty requires staffing, scheduling and short-term adjustment considerations. In “A Two-Stage Stochastic Integer Programming Approach to Integrated Staffing and Scheduling with Application to Nurse Management,” K. Kim and S. Mehrotra integrate these decisions in a model under demand uncertainty. The model is formulated as a two-stage stochastic integer program with mixed-integer recourse. Parametric mixed-integer rounding inequalities are shown to generate a convex hull representation of the second-stage feasible region as a function of the first-stage variable, hence, making the problem tractable. A multicut aggregation approach and a prioritized branching strategy are developed within integer L-shaped method to achieve a further 10-fold solution time reduction. Empirical study is performed with realistic problem instances (each with more than 1.3 million integer and 4 billion continuous variables) using 3.5 years of patient volume data from Northwestern Memorial Hospital. Results establish the value of stochastic modeling, and show the promise of solution methodology.

Managing Elective Admission in a Public Hospital

Beds are a critical resource in hospital operations. Highly trained personnel are needed to man these beds. On days with high bed occupancy, patients’ waiting times for admission could be excruciating. In contrast, beds are under-utilized on days with low bed occupancy. Moreover, emergency admissions are unpredictable, while elective admissions are scheduled by the hospital. In “A Robust Optimization Model for Managing Elective Admission in a Public Hospital,” F. W. Meng, J. Qi, M. L. Zhang, J. Ang, S. Chu, and M. Sim propose a robust optimization model to reduce incidents of bed shortfall by enforcing quotas on the elective admissions. They maximize the level of uncertainty that the admission system can withstand without violating the expected bed shortfall constraint. The proposed model is tested in simulations based on real hospital admission data and favorable results for adopting these robust optimization models are reported.

How Does Dependency of Random Returns Affect Decisions in Portfolio Optimization?

In “Robustness to Dependency in Portfolio Optimization using Overlapping Marginals,” X. V. Doan, X. Li, and K. Natarajan develop a tractable distributionally robust optimization model for the portfolio optimization problem taking into account the dependency structure of random returns using overlapping marginals. They also propose heuristics to extract dependency structures from historical data and demonstrate the model efficiency with real financial data from different time periods which cover the recent financial crisis.

Approximate Dynamic Programming when More Is Better

It is often the case that sequential decision problems from operations research exhibit an inherent monotonicity property, i.e., “more is better,” in the optimal value function. An obvious example is resource management with no holding costs, where from the perspective of a fixed point in time, one would always prefer to own more of a particular resource. A less apparent example is when environmental variables influence the stochastic evolution of a primary state variable in a monotonic fashion (e.g., extreme weather can lead to increased expected travel times; high natural gas prices can lead to higher electricity spot prices). In “An Approximate Dynamic Programming Algorithm for Monotone Value Functions,” D. R. Jiang and W. B. Powell analyze an approximate dynamic programming technique to exploit this important structural property. They show that near-optimal solutions can be obtained for various problems using the proposed method when exact approaches are computationally intractable.

A Novel Solution for Stochastic Network Design Problem

The problem of finding the cost effective built in capacities for single commodity stochastic networks where the constraint that all demands should be met on a prescribed probability level (reliability constraint) is challenging. In “Single Commodity Stochastic Network Design under Probabilistic Constraint with Discrete Random Variables,” A. Prekopa and M. Unuvar consider node demands as random and work with set of probabilistic constraints. The proposed solution starts with eliminating the redundant probabilistic constraints and works with the set of remaining ones. The solution method utilizes p -efficient point concept along with LP relaxation and multiple choice knapsack algorithm. This method can be applied in planning flood control networks, designing evacuation routes, capacity planning of cloud computing datacenter networks etc.

Monitoring Drugs for Unknown Side Effects

Postmarketing drug surveillance refers to the process of monitoring drugs among consumers for possible unknown side effects, and is a vital component of comprehensive drug safety. Electronic databases that capture detailed patient-level data are potentially fruitful data sources for such surveillance. However, existing methods that have been proposed for such surveillance do not fully account for the complexities stemming from its open-ended and exploratory nature, which in particular requires simultaneous monitoring of multiple, possibly interdependent, adverse events. In “Active Postmarketing Drug Surveillance for Multiple Adverse Events,” J. Goh, M. V. Bjarnadóttir, M. Bayati, and S. A. Zenios propose a new method for drug surveillance that analyzes sequentially-arriving data and explicitly captures interdependencies between adverse events. Using numerical simulations, they demonstrate the efficacy of their method, showing its robustness to various distributional and parameter assumptions and that it can deliver Types I and II errors that are below prespecified levels.