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(2):ii-iv. https://doi.org/10.1287/opre.2015.1379

OPERATIONS RESEARCH
Vol. 63, No. 2, March–April 2015, pp. ii–iv
ISSN 0030-364X (print) | ISSN 1526-5463 (online)

http://dx.doi.org/10.1287/opre.2015.1379

In This Issue

OR Pioneer Solves World War II Combat Problems

A window onto the origins of operations research is provided by "A glimpse at an Operation Analyst's World War II: Report on the Combat Performance of the Remote Control Turrets of B-29 Aircraft." Alex Green, who later had a long career as an applied physicist, was a 25-year old aerial gunnery expert/Operations Analyst in 1944, when he was sent to the 20th Air Force base in India to help reduce B-29 losses during attacks over occupied China. Long before digital computers, Green used interview data and reports to systematically draw conclusions that could be translated into operational choices. This fascinating story includes interaction with World War II's most famous commander, Curtis LeMay, and flights over the Himalayas ("the Hump"). It reveals how an OR pioneer performed technical analyses under wartime conditions and has parallels to today's OR.

Shipping New Fashion Articles to a Global Store Network

The paper "Initial Shipments for New Products at Zara" by Gallien, Mersereau, Garro, Dapena Mora, and Nóvoa Vidal reports the development and field testing of a new system for optimizing the first lifecycle shipments of new products by Spanish fashion retailer Zara. These decisions have significant strategic importance for Zara because the company rolls out more than 8000 new articles in its stores every year as part of its fast fashion business model. They are also challenging because they involve maximal demand uncertainty, and because Zara's network of more than 1500 stores is both large and heterogeneous. Implementing a novel approach for quantifying the benefits of learning associated with the first week of sales, the forecasting and optimization models presented in the paper were tested as part of a large-scale controlled field experiment organized in 2012, and thus estimated to increase Zara's total average season sales by approximately 2%.

Improving Screening and Staff Planning at Large Health Care Organizations

Routine HIV screening has long been known to be more cost-effective compared to the traditional risk-based screening approach. Yet, very few healthcare organizations have implemented this because it is likely to increase operational costs beyond budgets. In "Planning for HIV Screening, Testing and Care at the Veterans Health Administration," S. Deo, K. Rajaram, S. Rath, U. Karmarkar, and M. Goetz, develop a large scale nonlinear mixed integer programming model which considers budget constraints and determines the optimal screening level in each period as well as the optimum staffing level at each part of the health care system in order to maximize the overall health benefits to patients. They further apply their methodology to the Greater Los Angeles station of the Veterans Health Administration using data from the station as well as academic literature. While routine screening is infeasible for current budget levels, several managerially appealing screening policies provide substantial benefits (20% to 300%) measured in Quality Adjusted Life Years over the current practice of risk-based screening. A modified version of two of these policies has been implemented at this station.

On the Pitfalls of Service Slowdowns

In many service systems, service rate may depend on congestion and is not constant, as is usually assumed in the literature. In the paper "Service System with Slowdowns: Potential Failures and Proposed Solutions," J. Dong, P. Feldman, and G. B. Yom-Tov study the effect of service rate deterioration on system performance. The authors analyze a modification of the classical Erlang-A model, which accounts for service slowdowns, in the Quality-and-Efficiency Driven (QED) regime. The behavior of such systems depends heavily on the rate at which the service rate function deteriorates, referred to as 'load sensitivity" with low and high load sensitivity resulting in very different behavior. The latter exhibits a bi-stability phenomenon, in which the system alternates randomly between QED and Efficiency Driven (ED) performance levels. The authors analyze how the system scale and other model parameters affect bi-stability and discuss policies to avoid ED regime performance.

Modeling Customer Choice Using Product Features: What Assortments to Offer and What Prices to Charge?

When modeling demand in many applications, it is important to capture the substitution between products because when a product is not available, customers may choose something else, creating interactions between the demands for the different products. Discrete-choice models allow us to capture substitution and get a more realistic model of demand. However, as the discrete-choice model becomes more complex, the corresponding operational problems also become more difficult. In "The d-Level Nested Logit Model: Assortment and Price Optimization Problems," G. Li, P. Rusmevichientong, and H. Topaloglu consider the d-level nested logit model of choice, where each product is described by a list of d features. They provide a novel formulation using a tree. Building on this formulation, they develop an efficient method for computing a revenue-maximizing assortment of products to offer, assuming fixed prices. They also consider the price optimization problem where the goal is to find the revenue-maximizing prices.

Transshipment: A Mechanism to Mitigate the Mismatch Between the Supply and Demand

In seasonal product markets, retailers/distributors may incur significant losses due to unsold inventory. A prominent example is the flu vaccine market. Each year, new flu vaccines are manufactured in the beginning of February so that they can be used for vaccination in the period from September to January of the next year. The vaccines left unused at the end of this period must be discarded. Because the production of flu vaccines starts six to nine months before the vaccination period, the fluctuation of demand cannot be observed at the time of production; this results in mismatch

between production and consumption. Similar wastage due to such mismatch occurs in other markets such as the fashion industry. In "Optimal Joint Replenishment and Transshipment Policies in a Multi-Period Inventory System with Lost Sales," H. Abouee-Mehrizi, O. Berman, and S. Sharma propose transshipment as a means to mitigate the mismatch between the supply and demand. We consider a multi-period centralized ordering and transshipment decision problem and characterize the structure of optimal ordering and transshipment decisions in each period.

A Clarification on a New Theorem for the Single Machine Total Tardiness Problem

John J. Kanet presented a fourth theorem that can be used in conjunction with Emmons' original three theorems for helping to solve the single-machine total tardiness problem. In "Erratum: One Machine Sequencing to Minimize Total Tardiness: A fourth Theorem for Emmons," J. Kanet, F. Della Croce, C. Koulamas and V. T'kindt clarify that Theorem 4 in Kanet (2014) is not incremental to Theorem 3 of Emmons (1969) for the problem.

Balancing the Risks of Quick and Slow Searches for a Target

In "Trading off quick versus slow actions in optimal search," S. Shechter, F. Ghassemi, Y. Gocgun, and M. Puterman present a discrete search problem in which a searcher decides which grid cell to visit next, as well as which of two types of actions to take upon visiting a cell: a quick or a slow search. A quick search risks damaging the target if it is in the cell (e.g., accidentally blowing up a landmine one is searching for). On the other hand, the time it takes to perform a slower, more careful search exposes the searcher to exogenous risk (e.g., enemy fire). The authors characterize an optimal search policy for an unconstrained-path version of the search, and use it to 1) construct a brand-and-bound solution algorithm for solving smaller instances of, and 2) develop a dynamic programming algorithm for approximately solving larger instances of, a constrained-path search.

Relaxing Stochastic Dominance Rules to Reflect Most Preferences More Effectively

Stochastic dominance enables decision makers to partially rank alternatives with only partial information about utility functions. To avoid situations where most decision makers prefer one alternative to another but stochastic dominance cannot rank them, almost stochastic dominance allows small violations of stochastic dominance rules. Although quite promising, this idea has not caught on in practice due to implementation issues and inconsistencies between integral conditions and their associated utility classes. In the paper, "Generalized Almost Stochastic Dominance," I. Tsetlin, R. L. Winkler, R. J. Huang, and L. Y. Tzeng develop generalized almost second-degree stochastic dominance and almost second-degree risk in terms of appropriate utility classes and corresponding integral conditions, and extend these concepts to higher degrees. They address implementation issues and show that generalized almost stochastic dominance inherits the appealing properties of stochastic dominance. Generalized almost stochastic dominance could be useful in decision analysis, in empirical research (e.g., in finance), and in theoretical analyses of applied situations.

Additive or Multiplicative Cardinal Preferences? Think Again

In the paper "Multiattribute Utility Functions Satisfying Preferential Independence," A. Abbas and Z. Sun present a significant contribution to preferences of a decision maker when uncertainty is present in decisions with multiple objectives. The economics and operations research literature has long advocated the use of additive ordinal preferences for consequences in deterministic multiattribute decisions. They suggest that preferences for consequences in multiattribute decision problems can be represented by an additive function of the attributes. But if preferences for consequences can be represented by such an additive function, what does this tell us about our cardinal preferences for lotteries in multiattribute decision problems? The decision analytic literature has been dominated by additive and multiplicative forms of cardinal utility functions over each of the individual attributes. This paper shows that additive ordinal preferences correspond to a richer class of utility functions than the additive and multiplicative utility forms. It shows equivalence between additive ordinal preferences and Archimedean utility copulas proposed by the first author. The paper also shows how to assess Archimedean utility forms empirically, and also highlights the errors that may result when using arbitrary forms in the analysis.

Coping with Continuous Demand in Planar Facility Location Problems

In many urban applications the number of potential customers can be in the millions, making the representation of individual customers as discrete demand points infeasible. Instead, it may be more accurate to represent demand as continuously distributed over some region. The partition generated by a set of discrete facilities leads to a Voronoi diagram representation – a workhorse of computational geometry and Geographical Information Systems. In "Structural Properties of Voronoi Diagrams in Facility Location Problems with Continuous Demand," I. Averbakh, O. Berman, J. Kalcsics, and D. Krass analyze how the Voronoi diagram changes when a new facility is added to an existing set. They develop efficient geometric solution approaches for several classical facility location problems in which demand is continuously and uniformly distributed over a convex polygon for the case of rectilinear distances. The techniques can be extended to several other facility location problems and, in principle, to other distance metrics.

Unified Analysis of the Piecewise Linear Relaxations and Restrictions for the Pooling Problem

The pooling problem is a classical NP-hard bilinear optimization problem motivated from multicommodity network flows. The problem can be computationally challenging to solve on large networks and global optimization algorithms routinely use mixed integer programming based lower and upper bounds. In "Analysis of MILP Techniques for the Pooling Problem," S. Dey and A. Gupte formally prove the approximation quality of these bounds. Both the approximation ratios are shown to depend only on the number of output nodes and are attained in the worst-case. The analysis also leads to an approximation algorithm and a heuristic, which is empirically observed to provide solutions that are often orders of magnitude better than those from commercial optimization solvers.

How to Upper and Lower Bound the Optimal Policy of a POMDP by Myopic Policies?

Computing the optimal policy of a POMDP is in general intractable. The paper, "Myopic Bounds for Optimal Policy of POMDPs: An extension of Lovejoy's structural results" by Krishnamurthy and Pareek, gives sufficient conditions under which judicious myopic policies can be constructed via linear programming so as to provably lower and upper bound the optimal policy. Furthermore, these myopic policy bounds are optimized to maximize the volume of belief states where they coincide with the optimal policy. Numerical examples are given to illustrate these myopic policies and their performance. The electronic companion paper shows how the results of the paper can be extended to controlled-sensing problems. The conditions and results in the paper extend Lovejoy's 1987 paper on structural results for Markov decision processes.

Estimating Sensitivity of Tail Performance Measures in Simulations

Tail performance measures such as quantiles are more appropriate than classical performance measure such as mean and variance when characterizing the extreme properties of a stochastic system. Estimating quantile sensitivities in simulation is an important problem first addressed by Hong (2009), who showed that infinitesimal perturbation analysis (IPA) estimators are weakly consistent. In "Technical Note: On Estimating Quantile Sensitivities via Infinitesimal Perturbation Analysis," G. Jiang and M. C. Fu present an alternative more direct derivation of the IPA estimators and establish conditions under which unbatched and batched estimators are strongly consistent.

Optimal Investment Timing and Dynamic Learning with Multiple Modes

In "Investment Timing with Incomplete Information and Multiple Means of Learning," J. M. Harrison and N. Sunar consider a firm that can dynamically choose one of several costly learning modes to reduce uncertainty about the unknown value of a project. Each learning mode incurs cost at a particular rate and provides information of a particular quality. The firm must also decide when to stop learning and either invest or abandon the project. Using a continuous-time Bayesian framework, the authors prove that (i) in the undiscounted case, the optimal learning policy is to choose the mode that has the smallest cost per signal quality, and (ii) in the discounted case, an optimal learning and investment policy can be summarized by a small number of critical values, and the firm only uses learning modes that lie on a certain convex envelope in cost-rate-versus-signal-quality space.

A Simple Heuristic for Reorder Intervals in Distribution Inventory Systems

Distribution inventory system is one important building block in a complex supply chain network. When managing a distribution system, firms often replenish inventory at each facility up to a base-stock level in every fixed time interval. In "Optimizing Reorder Intervals for Two-Echelon Distribution Systems with Stochastic Demand," K. Shang, Z. Tao, and S. X. Zhou study such inventory management policies in a one-warehouse-multi-retailer distribution system. They develop an evaluation scheme and provide a method to obtain the optimal base-stock levels and reorder intervals that minimize the total system cost. The procedure of finding the optimal reorder intervals is based on allocating the total system cost to each facility and constructing a simple lower bound to the allocated facility cost. Meanwhile, a simple and effective heuristic is developed. One interesting observation is that the optimal reorder intervals tend to satisfy integer-ratio relationships and that the suggested heuristic can generate effective integer-ratio policies for large systems.

Cooperative Games for Pooled M/M/c Queues

When several firms collaborate by sharing their resources (machines, inventories, servers, etc.) for efficiency, a cost allocation problem arises: Who pays for the shared resources or shortages, and how much? In fact, does a fair cost allocation even exist? In "Resource Pooling and Cost Allocation among Independent Service Providers," F. Karsten, M. Slikker, and G.-J. van Houtum study this problem in the context of pooled M/M/c queues. They formulate cooperative games in which any coalition of players merges their resources and individual demand streams into a joint M/M/c queue. They look at both the situation with a given, fixed number of servers per player and the situation with an optimized number of servers per coalition. The authors provide sufficient conditions for their games to admit cost allocations that are fair in the sense that no subset of firms can improve by breaking away.