Operations Research

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

In This Issue

To cite this article:
(2016) In This Issue. Operations Research 64(2):ii-iv. https://doi.org/10.1287/opre.2016.1492

Operations Research
Vol. 64, No. 2, March–April 2016, pp. ii–iv
ISSN 0030-364X (print) | ISSN 1526-5463 (online)

http://dx.doi.org/10.1287/opre.2016.1492

In This Issue

How to Route Traffic in Networks with Fluctuating Link Capacity?

Wireless networks suffer from weather conditions that constantly affect wireless links, making their transmission rate fluctuating below the maximum values. Such fluctuations, if not explicitly taken into account in network design and operation, can heavily degrade the end-to-end traffic flows. This issue is becoming profound as wireless communications is a key element in voice, video and data transfer today, increasing the need for developing traffic routing mechanisms and associated optimization methods that address networks where partial link availability states are common in everyday operation. In "Optimizing Flow Thinning Protection in Multicommodity Networks with Variable Link Capacity," M. Pióro, Y. Fouquet, D. Nace, and M. Poss propose a simple, implementable traffic control mechanism called Flow Thinning (FT) that uses logical end-to-end tunnels with variable capacity, reacting on-line to capacity fluctuations. The authors present a detailed mathematical model underlying FT operation, solved through mixed-integer programming methods, and compared to other possible mechanisms.

Dynamic Pricing with Reference Price Effects for Heterogeneous Customers

Customers form reference prices through repeated purchases from a seller. There have been many studies of dynamic pricing problems that incorporate reference price effects, but most of them consider homogeneous customer arrivals and the conclusion is that an asymptotically constant pricing policy is optimal. In "Intertemporal Price Discrimination via Reference Price Effects," Z. Wang considers a dynamic pricing problem with reference price effects, where the customers may arrive at heterogeneous time periods. He shows that in this case, the seller's optimal pricing strategy could be a cyclic one, even when the customers are loss-neutral and their demand functions are identical. He further proves the cycle length in certain cases. The result is unique due to the interaction between the heterogeneous arrivals and the reference price effects and helps people understand the role of reference price effect in determining optimal pricing decisions.

Good Path Simulation and Smoothing for Handling Discontinuities in Financial Engineering

Quasi-Monte Carlo (QMC) methods are deterministic versions of Monte Carlo methods and play an increasingly important role in the problems of pricing and hedging of complex financial derivatives. These problems are often of high dimensionality and low smoothness (e.g., discontinuity), both of which may significantly affect the efficiency of QMC methods. In "Handling Discontinuities in Financial Engineering: Good Path Simulation and Smoothing," X. Wang studies what defines a good path simulation method (PSM) for the underlying processes in problems with discontinuities and how to handle the challenging problems of high dimensionality and discontinuity concurrently. It is found that PSM may affect both the effective dimension and the feature of discontinuity. The power to reduce the effective dimension and to realign the discontinuities with the coordinate axes is desirable for a QMC-based PSM. The author develops a two-step procedure that combines dimension reduction and smoothing. The first step is to design a good PSM that has the desirable power just mentioned and the second step is to exploit the nice feature of the functions to remove the discontinuities completely by a simple transformation. Numerical experiments demonstrate the high performance of the proposed method for pricing options and estimating Greeks.

When to Act on Intelligence

Operation Neptune Spear led to the elimination of Osama bin Laden. Intelligence agencies collected some information regarding his whereabouts, and the dilemma was when to act. Raiding a wrong location, based on insufficient information, would cause collateral damage. Waiting too long for more information could result in bin Laden escaping. M. P. Atkinson, M. Kress, and R. J. Lange model this dilemma in "When is Information Sufficient for Action? Search with Unreliable yet Informative Intelligence." The paper considers one target hiding in one of n cells. The searcher does not actively search the cells but instead receives occasional pieces of intelligence that point to a particular cell. The intelligence is imperfect and the searcher needs time to form and substantiate his situational awareness. This time may not be available. This paper addresses the tradeoff between "act now" and "wait and see" and provides decision rules for when to act.

How Can Firms Compete under Incomplete Information?

Competition is an important factor in shaping firms' strategies. Often, firms are uncertain about their competitive capabilities and strategies. In "Non-Linear Pricing Competition with Private Capacity Information," H. Nazerzadeh and G. Perakis study pricing strategies of competing capacity-constrained firms. They show that under incomplete information, non-linear pricing emerges at the equilibrium which may create inefficiencies. A by-product of their work is providing a framework for equilibrium analysis of games of incomplete information using techniques originally developed for auction design.

Using Last-Minute Advance Demand Information to Better Allocate Limited Supply to Sales Markets

Using advance demand information such as feedback from customers or data from pre-sales can considerably improve decision-making in supply chains. In the paper "Supply Allocation under Sequential Advance Demand Information," F. Papier analyzes how an agricultural seed company can benefit from advance demand information when having to decide on how much of its limited supply to allocate to the different national markets in Europe. The paper shows that the value of advance demand information is

greatest when the supply level is neither too high nor too low but similar to the overall market forecast. Furthermore, markets with an early sales period should receive less supply than comparable markets with a later sales period, because the decision to supply them has to be made when little information about the later markets is available. The author proposes a new planning method for allocating supply that leverages advance demand information to improve the profitability of the company.

Learning and Optimization in Multiproduct Inventory with Censored Demand

Most, if not all, of the papers on stochastic multi-product inventory systems assume that inventory decisions are made with full knowledge of the distribution of future demand. However, in practice, the demand distribution is usually not known a priori and the firm only has access to past sales data (often referred to as censored demand data). The selection of the most appropriate distribution and its parameters remains difficult. In "Nonparametric Data-Driven Algorithms for Multi-Product Inventory Systems," C. Shi, W. Chen, and I. Duenyas develop a nonparametric data-driven algorithm for the management of stochastic periodic-review multi-product inventory systems with a warehouse-capacity constraint. They show that the average expected cost of the proposed algorithm converges to the optimal cost at a square root convergence rate, the best convergence rate one can expect from online learning theory.

A Carrot-and-Stick Approach to Environmental Disclosure

A traditional regulatory instrument to detect harmful environmental hazards from private industries is inspection. Over the past few decades, environmental regulation agencies in many countries have been exploring more innovative ways to discover hazards such as these. A prominent example is the US EPA's Audit Policy, which offers rewards (in the form of reduced penalty or remediation subsidies/assistantship) to firms that promptly self-report an environmental hazard. In "Inducing Environmental Disclosures: A Dynamic Mechanism Design Approach," S. Wang, P. Sun and F. de Véricourt provide theoretical guidelines to design efficient regulation policies to induce voluntary hazard disclosure by leveraging both inspections and disclosure rewards. They formulate the policy design problem as a dynamic mechanism design problem with costly state verification. Their analysis suggests that efficient regulations have a perhaps surprisingly simple cyclic structure, which they fully characterize in closed form.

Creating Resilient Supply Chains: Which Suppliers to Back Up?

Harnessing economy of scope, supply chains may benefit from the investment in flexible backup infrastructure as a mechanism to mitigate disruption risks. Such an investment should be combined with backing up the appropriate suppliers, and need to be made in light of the ability maintain inventory safeguards in response to dynamic risks. In "Compensating for Dynamic Supply Disruptions: Backup Flexibility Design," S. Saghafian and M. Van Oyen study effective backup flexibility designs in such an environment. They link the dynamics of the inventory shortfall process to queueing and dam/storage processes, and characterize the base-stock levels and costs. They develop easy-to-compute yet effective indices that guide a supply chain manager in deciding which supplier to back up. Their findings shed light on the importance of considering (i) the duration (and not just frequency) of disruptions, and (ii) the variance (and not just the mean) of the demand uncertainty.

Optimal Pricing and Lead Time Quotation for Contract and Spot Customers

Make-to-order manufacturers or service providers often face the problem of allocating their capacity between contract customers who have contractual agreements defining prices and delivery times and spot purchasers who make occasional purchases without a contractual agreement. While contract customers provide long-term stability in business and the management of capacity, spot purchasers offer opportunities to utilize excess capacity in a profitable way. In "Price and Lead Time Quotation for Contract and Spot Customers," A. Hafizoglu, E. Gel and P. Keskinocak analyze strategies for quoting spot prices and lead times with the goal of maximizing long term profits. They show that the optimal policy often follows a Newsvendor-like structure, balancing marginal benefit and cost, and depends on the price and lead time sensitivity of the customers. They also explore the optimal mix of contract versus spot customers, and show that it is especially important to maintain a good balance of customer mix when contract customers are diverse in their willingness to pay.

Is Long Chain Always Optimal?

Process flexibility is an important operational strategy to enable capacity pooling and better fulfill the uncertain demand. A classical observation in process flexibility is the long chain design, a design that only requires a small amount of flexibility, but achieves almost the same benefit as fully flexible systems. In "Sparse Process Flexibility Designs: Is Long Chain really optimal?," A. Désir, V. Goyal, Y. Wei, and J. Zhang find that contrary to the conventional wisdom, the long chain is not always optimal compared to other designs with the same number of flexibility arcs. To further understand the effectiveness of the long chain, they establish its optimality among all connected designs. This result reveals two important managerial insights. First, structural swap lemmas used in the proof can be applied to the performance of general flexibility designs; and second, the long chain is optimal if exactly one degree of flexibility is added to each plant.

Managing Perishable Inventory with Fixed Costs

Perishable products, such as fresh food produce, pharmaceuticals, are indispensable part of our lives. However, the analysis of dynamic perishable inventory systems is notoriously difficult in both theory and computation due to the multi-dimensional nature. The existing literature does not model positive setup costs (mainly due to its mathematical intractability), despite the fact that fixed cost is often unavoidable in many practical settings. As a result, little is known about the structure of the optimal policies for this class of problems with setup costs. Thus, finding provably near-optimal control policies has been an open challenge. In "Approximation Algorithms for Perishable Inventory Systems with Setup Costs," H. Zhang, C. Shi, and X. Chao develop a constant-factor approximation algorithm for perishable stochastic inventory systems with setup costs under correlated demands. They demonstrate through an extensive numerical study that their algorithms perform consistently near-optimal.

Pricing for Products with Network Effects

A product is said to exhibit network effects if individual customers value it more when more other customers purchase it. How should

a seller set prices for a collection of substitutable products when those products exhibit network effects? C. Du, W. L. Cooper and Z. Wang address this question in "Optimal Pricing for a Multinomial Logit Choice Model with Network Effects." For settings with homogeneous products, they establish that there is a revenue-maximizing pricing policy with a remarkably simple structure. If the network effect is strong enough, then the price of one of the products should be set to a low value and the prices of all other products should be set to a single high value. Otherwise, if the network effect is not strong, all products should be priced identically. This structure indicates that in the presence of significant network effects, the seller can benefit by boosting sales of just a single product. The analysis in the paper reveals that similar insights apply in settings with heterogeneous products as well.

Strategically Selecting Customers to Optimize the Supply Chain

Recent strategies in make-to-order supply chains aim to better match demand to supply capabilities by proactively selecting the 'right' customers or shaping customers' demands. This may be done by diverting customers to third party partners, altering prices and due dates of requested demands, or even bypassing on potential revenue that involves prohibitive costs. Understanding the profitability of each individual customer becomes challenging in supply chains where common resources are used to satisfy many customers. While this could be estimated accurately if customer demands are known advance, unfortunately, knowledge of the exact demand is difficult to obtain, especially since customers arrive sequentially and often require a prompt response. In "Supply Chain Management with Online Customer Selection", A. N. Elmachtoub and R. Levi describe a general framework and algorithms to make such customer selection decisions. The algorithms have strong performance guarantees for several common cost structures, and rely on novel uses of reoptimization.

Computing the Price of Robustness in Problems with Piecewise Linear Costs

Although robust optimization is a methodology that has gained a lot of attention in the recent years, its application becomes a real computational challenge when an uncertain cost is expressed as the sum of piecewise linear functions. This occurs for instance in a number of inventory management problems. In "Robust Optimization of Sums of Piecewise Linear Functions with Application to Inventory Problems," A. Ardestani-Jaafari and E. Delage identify two conservative approximation schemes that respectively reduce the robust optimization problem to a linear program or a semi-definite program. The two schemes are numerically effective and actually exact under a set of interesting conditions. Valuable implications are described for newsvendor and inventory problems. An extensive set of empirical results also illustrate the quality of the approximate solutions that are obtained using these two schemes on randomly generated instances of the latter problem.

How and Why We Sample Lévy Bridges: for Stable and Tempered Stable Processes

Motivated from Brownian bridge and its wide applications, Lévy bridge has attracted researchers to investigate its distribution and simulation method. Simulating paths of a Lévy process conditioned on the pregenerated values, so-called bridge sampling, gives us a control over the coarseness of paths construction. In "Simulation of tempered stable Lévy bridges and its applications," K. Kim and S. Kim propose a bridge sampling scheme for stable and tempered stable subordinators. They obtain an approximate conditional distribution through double saddlepoint technique, and develop an acceptance-rejection algorithm based upon the gamma and inverse Gaussian bridges. The proposed method has proven its efficiency and accuracy from extensive numerical experiments. More importantly, it enables adaptive sampling for the valuation of path-dependent options, which provides significant computational efficiency. Furthermore, considerable variance reduction is achieved via stratification.

Modeling the Arrival Rate of Customers to a System

Customer arrivals to service systems such as call centers, emergency systems, or retail stores, are stochastic and difficult to model in a realistic way. Arrivals can be assumed to be Poisson, but with time-dependent and stochastic arrival rates. The arrival rates for different time periods of the same day are also dependent random variables, but not perfectly correlated, and must be modeled via some multivariate distribution. In "Rate-Based Daily Arrival Process Models with Application to Call Centers," B. Oreshkin, N. Régnard, and P. L'Ecuyer propose new models for this type of situation and develop parameter estimation methods for these models. (Estimation is difficult because the arrival rates themselves are unobserved.) They provide examples showing that their models match real call center data much better than previously proposed models and that the choice of model may have a large impact on performance measures (such as averages waiting times or proportion of large waits) estimated by simulation. These arrival models can be appropriate and useful for many other types of systems.

Decomposability and Time Consistency of Multistage Distributionally Robust Stochastic Optimization Problems

Distributionally robust stochastic optimization deals with situations when probability distributions of the involved random data cannot be accurately estimated. The approach is to consider a worst distribution in a specified uncertainty set of probability measures. In multistage settings this may result in violation of Bellman's principle of conditional optimality, which is closely related to time consistency of considered problems. In "Rectangular sets of probability measures," A. Shapiro discusses multistage decomposability of distributionally robust stochastic optimization problems, and its connection to the modern theory of risk averse stochastic optimization. The discussion is presented in frameworks of Stochastic Programming, Optimal Control and Markov Decision Processes.

New Optimization via Simulation Algorithm Uses Pairwise Comparisons to Improve Efficiency

In "Bayesian Optimization via Simulation with Pairwise Sampling and Correlated Prior Beliefs," J. Xie, P. Frazier and S. Chick develop a new algorithm for solving optimization problems whose objective functions can only be evaluated by a time-consuming stochastic simulator. This algorithm uses a technique from machine learning, Gaussian process regression, to model the objective function, and then recommends points to sample next that provide the most useful information. By considering pairs of points, evaluated simultaneously with the same stream of random numbers fed into the simulator, along with similarities in mean rewards at nearby design points, this algorithm is able to offer performance that improves over the state of the art in discrete stochastic optimization via simulation.