



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(4):ii-iv. <https://doi.org/10.1287/opre.2015.1410>

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

Designing Iterative Auctions for Tree Valuations

Iterative auctions are commonly employed in practice for allocation of multiple items to bidders, despite not guaranteeing efficiency in general. In “Iterative Auction Design for Tree Valuations,” O. Candogan, A. Ozdaglar, and P. Parrilo focus on a tree valuation model, which compactly represents bidders’ valuations for items, while allowing both for value complementarities and substitutabilities. For this class the efficient allocation can be obtained in a tractable way through a solution of a linear program. Moreover, together with its dual this LP also reveals the Walrasian equilibria; thereby identifying a new class of valuations for which Walrasian equilibrium exists, despite the presence of value complementarities. By developing an iterative algorithm for the solution of the aforementioned LP and complementing it with an appropriate payment rule, the paper also provides an efficient iterative auction. This auction relies on a simple pricing rule, compact demand reports, and a novel price update structure.

Should CMS Announce a Threshold Score to Select Bundled Service Providers?

A recent innovation by the Centers for Medicare and Medicaid Services (CMS), called “bundled payments for care improvement” (BPCI) initiative, invites proposers to create bundles by linking certain services for a particular medical condition and specifying bookends of episodes of care. Proposers design service chains and propose target values of quality metrics and payments per episode. Expert panels score proposals based on CMS-announced relative weights. There is no minimum score that will guarantee selection, and no limit on the number of proposers that may be selected. This leads to a non-competitive, but uncertain selection mechanism. In “Bundled Payments for Healthcare Services: Proposer Selection and Information Sharing,” D. Gupta and M. Mehrotra show that uncertain selection can be superior to the optimal fixed-threshold mechanism. Authors construct models to tease out two possible explanations for this phenomenon: information asymmetry and the effect of exclusivity on the total number of beneficiaries attracted to all BPCI providers. The models provide partial justification for CMS’ use of uncertain mechanism.

Reference Point Adaptation and Dynamic Behavioral Portfolio Selection

People make decisions based on comparisons. The benchmark used to compare is called the reference point in behavioral finance. While the majority of the literature assumes a stationary reference point for analytical tractability, the reference point should naturally evolve and adapt to investor’s prior investment experience. In “Dynamic Trading with Reference Point Adaptation and Loss Aversion,” Y. Shi, X. Y. Cui, J. Yao, and D. Li incorporate reference point adaptation into a dynamic portfolio choice model and analyze the derived optimal trading strategy to examine the implications of reference point adaptation for a loss-averse investor. The identified pattern of asymmetric V-shaped stock holding (between gain and loss

positions) offers convincing explanation for several long-standing puzzling phenomena of individuals’ trading behaviors, including house money, break even and disposition effects.

How and Why the Source of Uncertainty, Supply and/or Demand, Matters for A Risk-Averse Decision Maker?

In “Price-Setting Newsvendor Problems with Uncertain Supply and Risk Aversion,” B. Kazaz and S. Webster examine how risk and the source of uncertainty—demand and/or supply—affect tractability and optimal price and quantity decisions. When compared with a jointly concave risk-neutral objective function, the introduction of risk aversion does not alter concavity if the uncertainty stems from demand fluctuations. However, if the source of uncertainty is supply fluctuations, such as in the case agricultural product yields, then the introduction of risk aversion does not necessarily preserve concavity. Their work develops new elasticity conditions that lead to unique optimal decisions. For a given price, a risk-averse newsvendor is known to order less than a risk-neutral newsvendor under demand uncertainty, but she/he will order a greater amount when the source of uncertainty is supply. For a given quantity, a risk-averse newsvendor will price higher than a risk-neutral newsvendor under supply uncertainty, and less than a risk-neutral newsvendor under demand uncertainty.

Operational Constraints for Assortment Planning Under the Nested Logit Model

Assortment problems involve finding the revenue-maximizing set of products to offer when consumers choose among the offered products according to a certain choice model. An important concern in assortment problems is to ensure that the choice model that describes the choice process of the consumers is sophisticated enough to capture the real consumer behavior, while the corresponding optimization problem remains tractable. In “Capacity Constraints Across Nests in Assortment Optimization Under the Nested Logit Model,” J. Feldman and H. Topaloglu develop algorithms to find the revenue-maximizing set of products to offer when consumers choose according to the nested logit model and there are constraints on the number or the space consumption of offered products. When there is a constraint on the number of offered products, the authors show how to obtain the optimal solution. When there is a constraint on the space consumption of offered products, the authors show how to obtain approximate solutions.

Job Prioritization When Capacity is Uncertain

Many production/service systems have a variety of dedicated resources that feed a shared resource. Motivated by an oil-field services company, S. Demirel, I. Duenyas, and R. Kapuscinski address how to optimally control such a system. In the article, “Production/Inventory Control for a Make-To-Stock/Calibrate-To-Order System with Dedicated and Shared Resources,” the authors

characterize the optimal production and inventory policy and also develop simple heuristics. The paper provides insights into how to manage the finite shared resource and how its capacity influences how the firm should prioritize work as well as the inventory targets for the parts made on dedicated resources. Interestingly, this influence can be unintuitive: when multiple products share a finite shared capacity, the inventory target for the product having the larger (or less variable) production capacity can be larger. The authors explain the intuition behind this and other surprising behaviors and provide insights into how to optimally manage such systems.

Pricing Products in a Nested Choice Structure

When a firm sells a portfolio of heterogeneous products that are substitutable, it is of interest to the firm to understand how customers make choices among multiple products. This question has been studied empirically in many contexts, and one of the most common models is the nested version of the multinomial logit model called the nested logit model. A question that naturally arises subsequently is how the firm should set the prices to maximize its overall profit. This pricing problem is solved in “Pricing under the Nested Attraction Model with a Multi-Stage Choice Structure” by W. T. Huh and H. Li, where the structure of the optimal policy is characterized and an efficient solution algorithm is proposed.

Integrated Supply Chain Planning Under Demand Uncertainty

Integrated supply chain planning tools have gained more attention during the past decade as a promising solution to achieve significant savings. The production routing problem (PRP) is an example of an integrated problem that combines lot-sizing, inventory management and vehicle routing. Due to the difficulty of solving the PRP, most studies have addressed the deterministic variant of the problem in which demand is assumed to be known. As demand uncertainty is a major issue in supply chain management, Y. Adulyasak, J.-F. Cordeau, and R. Jans describe Benders decomposition methods to solve the stochastic PRP with demand uncertainty in two-stage and multi-stage decision processes in “Benders Decomposition for Production Routing Under Demand Uncertainty.” The decisions in the first stage include production setups and customer visit schedules, while the production and delivery quantities are determined in the subsequent stages. An efficient Benders decomposition algorithm with several enhancements is proposed to solve the two-stage stochastic PRP, while the Benders algorithm for the multi-stage PRP makes use of the solution obtained for the two-stage problem. These approaches allow instances of realistic size to be solved to optimality. The authors also leverage the reoptimization capabilities of Benders decomposition to address the two-stage stochastic PRP with a large number of scenarios in a sample average approximation (SAA) framework and the multi-stage stochastic PRP in a rollout heuristic, which result in significant speedups compared to other alternatives.

The Edge of Optimization over Randomization

Experimentation on statistically equivalent groups is a cornerstone of modern scientific inquiry. Traditionally, researchers almost exclusively use randomization—complete, pairwise, or otherwise—to construct such groups to the point where “randomized controlled trial” has become a single word in most researchers’ lexicon. But

is randomization the only, or indeed the best, way to compose groups for all controlled trials? In “The Power of Optimization over Randomization in the Design of Experiments Involving Small Samples,” D. Bertsimas, M. Johnson, and N. Kallus revisit randomization, highlight its shortcomings when sample size is limited by expense, rarity, or other factors, and propose instead an approach based on mixed integer optimization. The authors demonstrate that their proposed method achieves better balance and, via a novel test based on the bootstrap, greater statistical power than the existing, randomization-based approaches, especially in the context of small experiments. Applications range from fundamental research in medicine to empirical operations management.

A Finite Adaptability Scheme for Two-Stage Robust Optimization Problems with Binary Recourse Decisions

A wide variety of decision problems under uncertainty in engineering, science and economics span across multiple time stages and thus involve recourse decisions that adapt to the information observed over time. Moreover, in many application domains some or all of these recourse decisions must take integer values (e.g. yes/no decisions), which poses severe computational challenges. In the paper “K-Adaptability in Two-Stage Robust Binary Programming,” G. Hanasusanto, D. Kuhn and W. Wiesemann develop a finite adaptability approximation for two-stage decision problems that involve binary recourse decisions. The authors embrace the modern robust optimization paradigm to model the uncertainty, and they derive tractable problem reformulations by restricting the decision maker’s flexibility in the second stage to implementing the best out of K recourse decisions that are selected here-and-now. For problems where only the objective function is affected by uncertainty, they show that the restriction to $K = n$ of the 2^n theoretically possible recourse decisions suffices to recover an optimal solution.

Patient Flow Management in Emergency Department

Patient flow control is a key challenge in healthcare systems such as emergency departments (EDs). Specifically, after triage, patients must be examined by a doctor within a time window (depending, for example, on severity). Patients that are in-process (IP), on the other hand, impose congestion costs. ED physicians must hence balance between triage and IP patients, so as to minimize costs while adhering to time-till-first-service constraints. In “Control of Patient Flow in Emergency Departments, or Multiclass Queues with Deadlines and Feedback,” J. Huang, B. Carmeli, and A. Mandelbaum model this prioritization problem as a multiclass queue with deadlines and feedbacks and, for such a queue, we identify an asymptotically optimal control (prioritization) policy. This policy is easy to implement and, from a practical view point, it has the appealing property that the information it requires is readily available in EDs. Numerical experiments illustrate the robustness and extent to which our policy improves ED operational efficiency and service quality.

Impact of Re-Optimization in the Presence of Estimation Error

The last twenty years have witnessed the explosion of research in revenue management. Motivated by the complexity of the problem, many works have focused on the construction of easily

implementable heuristics assuming that the demand rate is known. This has obvious limitations. First, the true demand rate is rarely known and must be estimated from the data; second, it is not clear whether the heuristics constructed in the known rate setting are sufficiently robust in the unknown rate setting. In “Performance of an LP-based Control for Revenue Management with Unknown Demand Parameters,” S. Jasin studies re-optimization-based controls and shows that while some heuristics are reasonably robust in the sense that they do not blow up the impact of estimation error on revenue loss, others are extremely robust because they actually manage to *shrink* the impact of estimation error on revenue loss.

Appointment Capacity Planning in the Presence of Arbitrary Referral and Clinic Cancellation Distributions, and Delay-Dependent No-Show Behavior

A critical step in capacity planning for specialty clinics is deciding how much appointment capacity is needed for routine referrals so that waiting time targets are achieved. Mathematical analysis of this problem is challenging due to highly variable referral rates, random cancellation of appointments by clinics, and delay-dependent patient no-show and rescheduling behaviours. In “Appointment Capacity Planning in Specialty Clinics: A Queueing Approach,” N. Izady develops a discrete-time bulk service queueing model that efficiently calculates performance metrics incorporating all these features collectively. Real data has been used for validating the model as well as demonstrating its application in practice. Using an example, the author also shows how the model can be used in clinics such as primary care clinics where patient panel size rather than appointment capacity is the major decision variable.

A New Procedure for Learning Unknown Correlations between Decisions

Managers and policy-makers often have to choose a decision from a set of competing alternatives whose values are unknown, and can only be estimated through noisy simulations or field experiments. The power of a single experiment may be increased by exploiting correlations, or similarities, between alternatives, such that experimenting with one alternative provides additional information about many others. However, these similarities are

often difficult to quantify. In “Sequential selection with unknown correlation structures,” H. Qu, I.O. Ryzhov, M.C. Fu, and Z. Ding develop a new statistical model, based on approximate Bayesian inference, that is able to simultaneously learn both unknown values and unknown correlations from experiments with individual alternatives. The model is computationally efficient and easy to integrate with value of information methodology, a widely-used algorithmic strategy for adaptive learning. Substantial benefits can be obtained by incorporating unknown correlations into both the statistical model and the algorithmic procedure.

When to stop? The answer is simpler than it looks

We are surrounded by stop-or-go type decisions: buying stocks, shutting down network for security concern, sending patient to operating room, overhauling machine to prevent disastrous failure, etc. Such decisions can be formulated as a partially observable Markov decision process, but the computational complexity often makes the decision problem intractable. This is especially the case when the system state is only partially known. In “Multi-state Bayesian Control Chart over a Finite Horizon,” J. Wang and C.-G. Lee showed that the optimal policy takes different forms depending on the number of periods left in the horizon. As expected, “go” is always better than “stop” toward the end of horizon, whereas at the beginning the optimal policy has a simple control limit structure. Also examined in the paper are two special cases: the phase-type transition time model and multiple absorbing out-of-control states model.

Optimal Dynamic Pricing with Demand Uncertainty

In “Dynamic Pricing and Learning with Finite Inventories,” A. den Boer and B. Zwart consider the problem of optimally pricing a finite inventory that is sold during a finite time horizon, with parametric uncertainty on the demand function. They propose a pricing policy based on the idea of always choosing the selling price that is optimal with respect to available parameter estimates. It is shown that this simple policy achieves near-optimal asymptotic performance. The underlying mechanism of this result is an endogenous learning property satisfied by the system that ensures that no active price experimentation is necessary to eventually learn the optimal selling prices.