



Operations Research

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

In This Issue

To cite this article:

(2020) In This Issue. Operations Research 68(4):ii-v. <https://doi.org/10.1287/opre.2020.1998>

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 © 2020, 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

<https://doi.org/10.1287/opre.2020.1998>

Copyright: © 2020 INFORMS

A Unified Approach to Regime Classification and Stock Loan Valuation

For traditional perpetual American put options under regime-switching models with positive risk-free interest rates, optimal stopping usually can occur in any regime. Nonetheless, if the risk-free interest rates are allowed to equal zero (the interest rate may drop to zero sometimes in reality), there may exist “continuation regimes” within which optimal stopping can never occur, that is, within which stopping is never optimal. A natural problem is “regime classification,” that is, determination of all continuation regimes. In “Regime Classification and Stock Loan Valuation,” Ning Cai and Wei Zhang develop a unified, fixed point approach to solving this regime classification problem under general regime-switching exponential Levy models with any finite numbers of regimes and general Levy types. Applying this result, they also provide a unified framework for the valuation of infinite maturity stock loans under general regime-switching exponential Levy models.

An Operational Mechanism Design for Fleet Management Coordination in Humanitarian Operations

International humanitarian organizations (IHOs) run decentralized operations. While IHOs place their headquarters close to donors in developed countries, these organizations manage disperse long-term development programs such as basic healthcare in developing countries. Field vehicle fleets fulfill the stochastic transportation needs of programs in the field. But geographic dispersion creates information asymmetry about program needs, which are unknown to the headquarters. Earmarked funding of programs complicates this problem; earmarking forces headquarters to own the fleets and imposes constraints on the ability of headquarters to relocate fleets between programs. The incentive alignment issue is complex because conventional financial rewards and penalties are not viable in this system. In “Fleet Coordination in Decentralized Humanitarian Operations Funded by Earmarked Donations,” Pedraza-Martinez et al. propose a novel mechanism based on a capacity allocation rule (fleet size) to coordinate incentives in this setting. Their implementable mechanism respects the exogenous constraints characteristic of the humanitarian context.

Planning for End-User Substitution in Agribusiness

Saurabh Bansal and James S. Dyer study a common problem in the commercial agribusiness market, where farmers have a preference for a farm input such as a seed based on a fit with their geographical location but are also willing

to accept a closely related substitute. Such consumer-driven choices may not be adequately represented by traditional models that maximize the profit of a firm that seeks to make substitutions while maximizing its profit. They use a set of recent results for evaluations of moments over polyhedra to determine the exact inventory levels a firm should keep of substitutable products. Using proprietary data from a large firm in this domain, they highlight the role of geographical and climate-related factors that affect product substitution in the agribusiness industry and identify specific regions in the United States where product substitution is a source of substantial revenue for firms.

Learning and Earning in High Dimensions: Assortment Personalization Using Low-Rank Preferences

The proliferation of products and users in modern e-commerce presents a substantial challenge. How can a retailer learn each consumer’s preference for so many products—and, more importantly, maximize revenue—from sparse transactional data? In “Dynamic Assortment Personalization in High Dimensions,” Kallus and Udell show how to use the low rank structure of the matrix of preferences to dynamically choose individualized assortments to offer each consumer. This approach can increase revenue by orders of magnitude by learning across users to avoid excessive exploration. Their new method relies on convex optimization to recover hidden low-dimensional structure from sparse, high-dimensional data.

Price Discrimination and Patient Consumers

What does intertemporal price discrimination look like when consumers are patient, where patient consumers are people who are myopic with respect to prices and linger in the system? In “Technical Note—Dynamic Pricing with Heterogeneous Patience Levels,” Lobel studies a new dynamic programming framework to analyze this problem that allows for arbitrary joint distributions of valuations and patience levels. The paper finds that this problem is polynomial-time solvable and that optimal policies often take a form that it calls *incomplete cyclic policies*. Such policies are similar to cyclic decreasing policies, except some of the cycles are cut short.

Should Newsvendors Cooperate If the Demand Distribution is Unknown?

Cooperative game theory provides a framework for retailers to profit from and share the benefit of collaboration in the newsvendor application. The traditional framework, however, relies on the assumption that the joint demand

distribution is completely known. In the paper entitled “Technical Note—Robust Newsvendor Games with Ambiguity in Demand Distributions,” Doan and Nguyen propose a novel framework based on the principles of robust optimization to handle newsvendor games with ambiguity in the demand distribution. Using the proposed framework, the authors are able to analyze whether newsvendors should cooperate or not when only some marginal demand distributions are known.

Asymptotic Analysis of Constant-Order Policies for Lost-Sales Inventory Models with Positive Lead Times and Random Supply Functions

Lost-sales inventory models with positive lead times are a classic although notoriously difficult class of inventory models. The optimal control policy is intractable even with moderate lead times because of the curse of dimensionality. It was recently proved that the optimality gap of the best constant-order policy (COP) converges exponentially fast to zero as the lead time increases. However, this result requires the supply to be perfectly reliable. In “Technical Note—Constant-Order Policies for Lost-Sales Inventory Models with Random Supply Functions: Asymptotics and Heuristic,” Bu, Gong, and Yao extend this result to the setting of a general random supply function, including random capacity and random yield as special cases. Further, the authors analyze the asymptotic properties of the best COP with large penalty costs and construct a simple and effective heuristic COP. Finally, the authors provide a numerical study to derive further insights into the performance of the best COP.

Prior-Independent Auctions Through More Competition

In “Robust Auctions for Revenue via Enhanced Competition,” Roughgarden, Talgam-Cohen, and Yan revisit the classic Bulow–Klemperer result. This result compares the revenues of two well-known auction formats: the welfare-maximizing Vickrey auction and the revenue-maximizing Myerson auction. It shows that, with an extra bidder competing for the item, the Vickrey auction becomes as good as the Myerson auction in terms of revenue while maintaining independence from prior distributional information about bidders’ valuations. Unfortunately, Myerson’s toolbox for revenue-optimal auction design does not extend to combinatorial auctions with multiple heterogeneous items, for which optimizing revenue remains a challenge—especially if we want auction designs that are simple and robust enough to use in practice. This paper extends the Bulow–Klemperer result to multiple heterogeneous items by showing that a prior-independent, simple, welfare-maximizing auction with additional competing bidders achieves as much revenue as the ill-understood optimal auction.

Selling Passes to Strategic Customers

Many service providers offer a prepaid package of credits that can be redeemed for future use, often called passes, in conjunction with regular individual sales. In dynamic pricing situations, customers can strategize on the purchase, redemption, and renewal of the pass by optimizing

the timing and choices between passes and individual items based on future prices and their own changing needs. In “Selling Passes to Strategic Customers,” Wang, Levin, and Nediak integrate dynamic choice modeling and optimal control theory to study how to jointly price the passes and individual items in a dynamic setting. The authors endogenize an individual customer’s purchase/redemption decisions in their model and find that the seemingly complex problem has a simple (approximate) solution. The optimal prices remain nearly constant most of the time, except near the beginning and end of the sales horizon, exhibiting so-called turnpike properties. The pass, as a form of advance purchase, allows the seller to capitalize on the customer’s forward-looking behavior by exploiting the uncertainty of customer valuations.

Should Consumers Be Strategic?

Strategic consumer behavior under multiperiod pricing—when consumers anticipate future price changes and, potentially, delay a purchase to get a product at a lower price—is a phenomenon that has received significant recent research attention in the operations and revenue management literatures. In “Becoming Strategic: Endogenous Consumer Time Preferences and Multiperiod Pricing,” Aflaki, Feldman, and Swinney examine this problem from a different angle: They ask whether being strategic actually benefits consumers, whether consumers are willing to exert costly effort to become strategic, and how the answers to these questions impact the optimal decisions of firms pricing over multiple periods. The authors find that much of the conventional wisdom regarding strategic consumer behavior is upended when the decision of consumers to become strategic in the first place is considered, including qualitative features of the firm’s optimal pricing policy and, possibly, the firm’s preference between dynamic and committed pricing strategies.

Bayesian Incentive-Compatible Bandit Exploration

As self-interested individuals (“agents”) make decisions over time, they utilize information revealed by other agents in the past, and produce information that may help agents in the future. This happens in a wide range of recommendation systems, as well as in medical decisions. Each agent would like to *exploit*: select the best action given the current information, but would prefer the previous agents to *explore*: try out various alternatives to collect information. A social planner, by means of a well-designed recommendation policy, can incentivize the agents to balance the exploration and exploitation so as to maximize social welfare. The recommendation policy can be modeled as a multiarm bandit algorithm under Bayesian incentive-compatibility (BIC) constraints.

This line of work has received considerable attention in the “economics and computation” community. Mansour, Slivkins, and Syrgkanis contribute fundamental results on regret minimization: They design BIC bandit algorithms with optimal regret, characterize the dependence and assumptions on the Bayesian priors, and furthermore provide a black-box reduction from an arbitrary bandit algorithm

to a BIC one. Underlying their results is a simple but very powerful idea of carefully hiding a little exploration in a lot of exploitation.

Secrecy versus Transparency in Sales of Network Goods

Amazon and Apple, which sell tablet devices, have adopted different implicit information policies and developed distinct “reputations” about their tablets’ sales volume release. With Amazon, “even a number as basic, and presumably impressive, as how many Kindle e-readers the company sells is never released.” With Apple, iPhone and iPad sales numbers are always released, even if they are disappointing. In the paper “Information Disclosure and Pricing Policies for Sales of Network Goods,” Hu, Wang, and Feng study the sales information release policy, disclosure versus nondisclosure, for selling network goods subject to market size uncertainty. The authors identify two countervailing effects: (1) a prodisclosure “Matthew effect” and (2) an antidisclosure saturation effect, that drive the firms’ sales information disclosure policies. In addition, the authors also study the situation where the firm can decide on an all-or-nothing information disclosure policy together with endogenized prices, including state-independent pricing, contingent preannounced pricing, and contingent pricing without commitment.

When and Why Service Systems in Monopoly are Under-Exploited

Standard economic theory suggests that monopolies result in outputs lower and prices higher than socially desirable. In service systems, customers are often reluctant to join overly crowded systems because their service valuation decreases with system congestion. Thus a high service price is associated with better service through low congestion levels, that is, low system output. But can a monopolist profit more by providing lots of customers with poor service for a very low price? In “Social and Monopoly Optimization in Observable Queues,” Hassin and Snitkovsky introduce a unified approach, relying on the concept of observable queues, for studying the phenomena of monopoly overpricing in service systems. The authors explain why, in most observable queue models, the monopolist tends to underexploit capacity by overcharging its service. Yet they discuss cases in which the monopolist may prefer to attract demand by charging less than the socially optimal price.

Cutting Planes Need Not be Valid in Stochastic Integer Optimization

Cutting planes need not be valid in stochastic integer optimization. Many practical problems under uncertainty, for example, in energy, logistics, and healthcare, can be modeled as mixed-integer stochastic programs (MISPs). However, such problems are notoriously difficult to solve. In “Pseudo-Valid Cutting Planes for Two-Stage Mixed-Integer Stochastic Programs with Right-Hand-Side Uncertainty,” Romeijnnders and van der Laan introduce a novel approach to solve two-stage MISPs. Instead of using exact cuts that are always valid, they propose to use pseudo-valid

cutting planes for the second-stage feasible regions that may cut away feasible integer second-stage solutions for some scenarios and may be overly conservative for others. The advantage of using such cutting planes is that the approximating problem remains convex in the first-stage decision variables and thus can be solved efficiently. Moreover, the performance of these cutting planes is good if the variability of the random parameters in the model is large enough.

Dynamic Scheduling of Multiclass Many-Server Queues with Abandonment

In “Dynamic Scheduling of Multiclass Many-Server Queues with Abandonment: The Generalized $c\mu/h$ Rule,” Long, Shimkin, Zhang, and Zhang propose three scheduling policies to cope with any general cost functions and general patience-time distributions. Their first contribution is to introduce the target-allocation policy, which assigns higher priority to customer classes with larger deviation from the desired allocation of the service capacity and prove its optimality for any general queue-length cost functions and patience-time distributions. The $Gc\mu/h$ rule, which extends the well-known $Gc\mu$ rule by taking abandonment into account, is shown to be optimal for the case of convex queue-length costs and nonincreasing hazard rates of patience. For the case of concave queue-length costs but nondecreasing hazard rates of patience, it is optimal to apply a fixed-priority policy, and a knapsack-like problem is developed to determine the optimal priority order efficiently.

Burden-Free Seasonality in Markov Decision Processes

Markov decision processes are commonly used to model forward-looking behavior. However, cyclic terms, including seasonality, are often omitted from these models because of the increase in computational burden. In “Technical Note—Cyclic Variables and Markov Decision Processes,” Haviv develops a cyclic value function iteration (CVFI), an adjustment to the standard value function iteration. By updating states in a specific order, CVFI allows cyclic variables to be included in the state space with no increase in the computational cost. This result is proved theoretically and shown to hold closely in Monte Carlo simulations.

Determining the Right Size of a Blended Workforce

The rise of the blended workforce, which is identified as one of the top current workplace trends, is prompting firms to reevaluate their staffing strategies. A blended workforce melds as a deliberate business strategy flexible workers (for example, independent contractors or freelancers) with full-time employees. Because flexible workers are free to determine their own work schedules, the supply (total number of workers) is uncertain. In “Managing Supply in the On-Demand Economy: Flexible Workers, Full-Time Employees, or Both?,” Dong and Ibrahim examine the optimal staffing strategy for flexible workers and full-time employees to effectively balance operating costs, time variability in

customer demands, and supply-side uncertainty while not compromising on the quality of service offered to customers. This work gleans insights on the appropriateness of alternative workforce models, which are especially relevant for the gig economy.

How Overloaded Queues Behave When Customers Are Patient

The analysis of queues with multiple servers is typically challenging when the service time distribution is general. Such analysis usually involves an infinite-dimensional process for tracking service ages or residual service times.

In “Diffusion Approximation for Efficiency-Driven Queues When Customers Are Patient,” He demonstrates from a macroscopic perspective that, if customers are relatively patient and the system is overloaded, the dynamics of a many-server queue could be as simple as the dynamics of a single-server queue. In particular, the virtual waiting time process can be captured by a one-dimensional diffusion process, which enables us to obtain simple formulas for performance measures, such as service levels and effective abandonment fractions. To justify this diffusion model, a functional central limit theorem is established for the superposition of stationary renewal processes.