



Operations Research

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

In This Issue

To cite this article:

(2021) In This Issue. Operations Research 69(3):ii-iv. <https://doi.org/10.1287/opre.2021.2149>

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 © 2021, 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.2021.2149>

Copyright: © 2021 INFORMS

Using Dynamic Personalized Pricing to Mitigate the Quality Variability Penalty

Most service settings involve some degree of variability in the quality of customers' experiences. An understudied mechanism is analyzed by which this variability can reduce firm revenues when customers do not know true service quality but rather learn about that quality over time based on their experiences. Essentially, customers get *stuck* with low-quality opinions and low purchase likelihoods after unusually bad experiences. However, high-quality opinions after unusually good experiences are quickly corrected. In "How Service Quality Variability Hurts Revenue When Customers Learn: Implications for Dynamic Personalized Pricing," DeCroix, Long, and Tong show that dynamic pricing can help mitigate this effect. Specifically, if the firm can set individual prices so that each customer perceives the same surplus, then the variability impact is entirely eliminated. They also demonstrate that simpler heuristics similar to those observed in practice (requiring less information and pricing flexibility) can partially mitigate the revenue loss.

Finding Shortest Paths That Are Reliable and Satisfy Constraints on Time-Dependent Weights

Shortest path problems are used to model and solve a variety of real-world problems. In "The Constrained Reliable Shortest Path Problem in Stochastic Time-Dependent Networks," Ruß, Gust, and Neumann generalize multiple versions of the well-studied shortest path problem: First, the reliable shortest path problem in which travel time is minimized while guaranteeing a certain on-time arrival probability; second, the time-dependent shortest path problem in which edge weights depend on the time at which an edge is used; and third, the constrained shortest path problem in which constraints on additional edge weights need to be satisfied. This generalized problem naturally appears in transportation networks, in which travel times are notoriously uncertain and costs become more and more time dependent. Examples of such time-dependent costs include peak and off-peak fares in public transport, time-dependent congestion charges, and dynamic pricing for shared vehicles.

Lessons from the Founders of OR

"Soft OR and Practice: The Contribution of the Founders of Operations Research," by Dyson, O'Brien, and Shah explores two themes relating to the 43 people whose profiles constitute the Assad and Gass book referred to in the article. First, it explores the direct influence of the founders on the development of soft operations research (OR) and the alignment of their work with the characteristics of soft OR. Second, it explores how many key early developments of OR stemmed from the founders' engagement with practice. It is argued that the affinity of the founders with soft OR and with its emphasis on process in the framing and formulation of problems should therefore be universally recognized as a legitimate branch of OR within its practice and

evidenced within its publications. It is additionally argued that all academics should see engagement with practice as a fertile source of stimulation and ideas.

How to Efficiently Solve the Robust Vehicle Routing Problem with Demand Uncertainty?

Capacitated vehicle routing problems are widely studied combinatorial optimization problems, and branch-and-cut-and-price algorithms can solve instances harder than ever before. These models, however, neglect that demands volumes are often not known with precision when planning the vehicle routes, thus incentivizing decision makers to significantly overestimate the volumes for avoiding coping with infeasible routes. In "Branch-Cut-and-Price for the Robust Capacitated Vehicle Routing Problem with Knapsack Uncertainty," by Pessoa, Poss, Sadykov, and Vanderbeck a robust formulation that models demand uncertainty through a knapsack polytope is considered. A new branch-and-cut-and-price algorithm for the problem is provided, which combines efficient algorithms for the problem with no uncertainty with profound results in robust combinatorial optimization and includes novel heuristics and new valid inequalities. The numerical results illustrate that the resulting approach is two orders of magnitude faster than the best algorithm from the literature, solving twice as many instances to optimality.

Risk Management for Sustainable Sovereign Debt Financing

The sharp increase of sovereign debt internationally, since the 2008 global financial crisis, decisively contributed to several sovereign debt crises. The current COVID-19 pandemic and the fact that public debt remains high globally, have prompted a renewed interest in debt sustainability analysis (DSA) and in policy discussions concerning the most appropriate variables. In "Risk Management for Sustainable Sovereign Debt Financing," Zenios, Consiglio, Athanasopoulou, Moshammer, Gavilan, and Erce develop a normative DSA model to manage tail risk and optimize debt-financing decisions with sustainability conditions on debt stock and flow, under macroeconomic, financial, and fiscal uncertainty. They show that a risk management view alters a government's debt-financing policy to manage tail risk better. Many uncertain variables confound the problem, and portfolio optimization using stochastic programming on scenario trees provides a versatile and effective tool to achieve sustainable debt dynamics. The model is an essential building block of the European Stability Mechanism framework to assess debt sustainability of eurozone member states, including the repayment capacity of crisis countries under €295bn assistance programs.

Data Analytics for Optimal Detection of Metastatic Prostate Cancer

In “OR Practice–Data Analytics for Optimal Detection of Metastatic Prostate Cancer,” Merdan, Barnett, Denton, Montie, and Miller used data-analytics approaches to develop, calibrate, and validate predictive models, using machine learning, to help urologists make more accurate decisions about ordering diagnostic bone and CT scans for their patients with prostate cancer.

The models were validated using statistical methods based on bootstrapping and evaluation on out-of-sample data in advance of implementation. These models were used to design and implement guidelines that optimally weigh the benefits and harms of radiological imaging for detection of metastatic prostate cancer. The models were also implemented on a publicly available website called askMUSIC.

The Michigan Urological Surgery Improvement Collaborative implemented these guidelines across the State of Michigan, which were predicted to reduce the number of patients who underwent diagnostic bone and CT scans by more than 40 percent while slightly improving detection rate of cancer that had spread, or metastasized, during the study. The reduced scanning also saved patients and insurers an estimated \$275,000 over a one-year period.

Procurement Mechanisms for Assortments of Products

Many procurement agencies around the world construct assortments of differentiated products from which consumers can buy from. A leading example is framework agreements, a type of procurement mechanism commonly used by governments. This type of practice is studied head on. “Procurement Mechanisms for Assortments of Differentiated Products,” by Saban and Weintraub introduces a mechanism design formulation of the procurement agency’s problem and solves it under progressively more realistic implementation constraints. The results show how restricting entry of close-substitute products into the assortment can increase price competition, reducing spending significantly, without much damage to the variety offered to consumers. Furthermore, the results have practical implications that can be used by procurement agencies to increase consumer surplus and have already been used to redesign FAs in the Chilean government.

Online Allocation and Pricing: Constant Regret via Bellman Inequalities

In “Online Allocation and Pricing: Constant Regret via Bellman Inequalities,” Vera, Banerjee, and Gurvich introduce a framework for designing simple and efficient policies for a family of online allocation and pricing problems that includes online packing, budget-constrained probing, dynamic pricing, and online contextual bandits with knapsacks. The policy performance is evaluated in terms of its regret (i.e., additive gap) relative to an offline controller that is endowed with more information than the online controller. The framework is based on Bellman inequalities, which decompose the loss of an algorithm into two distinct sources of error: (1) those arising from computational tractability issues, and (2) those arising from the (mis)estimation of random trajectories. Balancing these errors guides the choice of benchmarks, and leads to tractable policies with strong performance guarantees. In all the examples, the proposed policies only require resolving a linear program in each period, followed by a simple greedy action-selection rule; thus, the authors policies are practical as well as provably near optimal.

Modeling Choices with Different Variabilities

In “Heteroscedastic Exponential Choice,” Alptekinoglu and Semple investigate a discrete choice model that can handle choices with utilities having different variances. This new model, the Heteroscedastic Exponential Choice (HEC) model, nests the Exponential Choice (EC) model as a special case. Like EC, HEC has excellent structure and permits the choice probabilities, demand elasticities, and consumer surplus to be expressed in an analytically convenient (closed) form. Determining optimal monopoly prices for products in an assortment as well as determining equilibrium prices for an oligopoly of single-product firms are both tractable problems. Moreover, fitting the choice model to real data is straightforward, given the shape (log-concavity) of the likelihood function for a given variance structure. The HEC model performed well against MNL (multinomial logit) and EC models in empirical tests on household panel data for 30 categories of grocery products. In particular, brand variance was statistically significant in virtually every product category tested.

Managing Advertisers’ Budgets in Online Advertising Platforms

Most online advertising platforms sell ad placements using repeated auctions. Platforms often provide budget-management strategies that allow advertisers to control their cumulative expenditures: Advertisers declare the maximum daily amount they are willing to pay, and the platform adjusts allocations and payments to guarantee that cumulative expenditures do not exceed budgets. Many strategies can be used to accommodate budget constraints, and each one, when applied to all budget-constrained advertisers simultaneously, drives the system toward a different equilibrium. For example, platforms can modify the advertisers’ bids or the auctions’ allocation and payment rules. In “Budget-Management Strategies in Repeated Auctions,” Balseiro, Kim, Mahdian, and Mirrokni compare the equilibrium outcome of a range of budget-management strategies used in practice. They study the impact of budget-management strategies on the trade-off between the seller’s profit and buyers’ utility and shed light on the incentive properties of budget-management strategies.

A Theoretical Foundation for Value Functions in Sports

Value functions provide the basis for analyzing how sports teams gain competitive advantage. However, up to now, they lack a theoretical basis. In “Points Gained in Football: Using Markov Process-Based Value Functions to Assess Team Performance,” Chan, Fernandes, and Puterman provide a rigorous justification for using a dynamic programming approach to estimate value functions in the context of North American football. They formulate a Markov reward process model for a game between two asymmetric teams to show that the value function is the expected number of points acquired or lost from a particular down and yard line over and above the steady state value per play. Using more than 160,000 plays from the 2013–2016 National Football League seasons, they derive an empirical value function that provides the basis for a “points gained” performance assessment metric. Points gained quantifies the value created or lost on any play relative to expected performance. The authors show how this metric provides new insight into factors that distinguish outstanding performance. For example, passing plays generate the most points gained, whereas running plays tend to generate negative value. Short passes account for the majority of the top teams’ success and the worst teams’ poor performance.

Real-Time Resource Allocation: Beyond Stochastic Demand Modeling

Motivated by the practical concern that algorithm designs based on stochastic modeling of demand may perform poorly because of unpredictable patterns, “Online Resource Allocation Under Partially Predictable Demand,” by Hwang, Jaillet, and Manshadi proposes a new arrival model with both unpredictable and predictable components. Under the proposed model—which requires no demand forecasting—the authors design novel online algorithms that take advantage of the limited information the data reveal, despite the presence of an unpredictable component. Their algorithms are adjustable to the relative size of the stochastic component; their analysis reveals that as the portion of the stochastic component grows, the loss due to making online decisions decreases. This highlights the value of (even partial) predictability in online resource allocation. This paper serves as a first step in bridging the long-standing gap between the two well-studied approaches to the design and analysis of online algorithms based on (1) adversarial models and (2) stochastic ones.

Optimal Design of Machine Repair and Maintenance Contract

Maintenance outsourcing is quite common in industries that rely on complex and critical equipment. Instead of investing in the maintenance facilities, firms outsource maintenance activities to specialized companies. However, it may be hard for firms (i.e., *principal*) to observe whether maintenance companies (i.e., *agent*) put sufficient resources into providing the best service, which gives rise to agency issues. In a dynamic environment in which an agent is responsible for both maintenance and repair of a critical machine, how the principal uses payments and termination to tackle agency issues is a challenging problem. In “Optimal Contract for Machine Repair and Maintenance,” Tian, Sun, and Duenyas provide theoretical guidance on designing the optimal contract to induce efforts from an agent to efficiently operate a machine. Although they consider the very general contract forms, the optimal contracts demonstrate simple and intuitive structures, making them easy to describe and implement in practice.

Non-Asymptotic Analysis of Temporal Difference Learning

Temporal difference learning (TD) is a simple iterative algorithm widely used for policy evaluation in Markov reward processes. In “A Finite Time Analysis of Temporal Difference Learning with Linear Function Approximation,” Bhandari, Russo, and Singal prove finite time convergence rates for TD learning with linear function approximation. The analysis follows using a key insight that establishes rigorous connections between TD updates and those of online gradient descent. In a model where observations are corrupted by i.i.d. noise, convergence

results for TD follow by essentially mirroring the analysis for online gradient descent. Using an information-theoretic technique, the authors also provide results for the case when TD is applied to a single Markovian data stream where the algorithm’s updates can be severely biased. Their analysis seamlessly extends to the study of TD learning with eligibility traces and Q-learning for high-dimensional optimal stopping problems.

Nonparametric Pricing Analytics with Customer Covariates

Personalized pricing analytics is becoming an essential tool in retailing. Upon observing the profile of each arriving customer, the firm needs to set a price accordingly based on the observed personalized information, such as income, education background, and past purchasing history, to extract more revenue. For new entrants of the business, the lack of historical data may severely limit the power and profitability of personalized pricing. In “Nonparametric Pricing Analytics with Customer Covariates,” Chen and Gallego recommend a pricing policy to firms that simultaneously learns the preference of customers based on the profiles and maximizes the profit. The pricing policy doesn’t depend on any prior assumptions on how the personalized information affects consumers’ preferences. Instead, it adaptively clusters customers based on their profiles and preferences, offering similar prices for customers who belong to the same cluster trading off granularity and accuracy. The authors prove that the regret of the proposed policy cannot be improved by any other policy.

Sample Out-of-Sample Inference Based on Wasserstein Distance

Financial institutions make decisions according to a model of uncertainty. At the same time, regulators often evaluate the risk exposure of these institutions using a model of uncertainty, which is often different from the one used by the institutions. How can one incorporate both views into a single framework? In “Sample Out-of-Sample Inference Based on Wasserstein Distance,” Blanchet and Kang provide such a framework. They quantify the impact of the misspecification inherent to the financial institution data-driven model via the introduction of an adversarial player. The adversary replaces the institution’s generated scenarios by the regulator’s scenarios subject to a budget constraint and a cost that measures the distance between the two sets of scenarios (using what in statistics is known as the Wasserstein distance). The authors also harnesses statistical theory to make inference about the size of the estimated error when the sample sizes (both of the institution and the regulator) are large. The framework is explained more broadly in the context of distributionally robust optimization (a class of perfect information games, in which decisions are taken against an adversary that perturbs a baseline distribution).