Operations Research

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

In This Issue

To cite this article:
(2018) In This Issue. Operations Research 66(3):ii-iv. https://doi.org/10.1287/opre.2018.1769

http://pubsonline.informs.org/journal/opre/

OPERATIONS RESEARCH
Vol. 66, No. 3, May–June 2018, pp. ii–iv
ISSN 0030-364X (print), ISSN 1526-5463 (online)

In This Issue

https://doi.org/10.1287/opre.2018.1769

How Much Information Do Online Reviews Communicate?

With the explosion of reviews and ratings for products and services, a central question is how much information reviews communicate to future consumers and if/how customers can infer the underlying quality of a service before making a purchase. In "On Information Distortions in Online Ratings," Besbes and Scarsini analyze this question. The heterogeneity of consumers and the fact that the processes of review generation and purchases are inherently intertwined make the interpretation task of past reviews and quality inference challenging for consumers. The authors establish that, with access to the full history of reviews, rational consumers can eventually learn the true quality. When information is limited to the mean of past reviews, naïve consumers who take the sample mean at face value tend to overestimate the true quality, but minimally sophisticated consumers can recover the true quality. The results point to the strong information content of reviews.

Scrip Systems: Analysis of the Minimum-Scrip Rule

A scrip system refers to a resource-sharing market where scrips play the role of real currency: The beneficiary of the service "pays" scrips to the provider, who in turn can "purchase" services using scrips. In "Analysis of Scrip Systems: On an Open Question in Johnson et al. (2014)," Bo, Dawande, and Janakiraman study the minimum-scrip rule in such a system: When a service request arises, among those who volunteer and who are able to provide service, the one with the least number of scrips is selected to provide the service. Under mild assumptions, the authors show that each player in the system always willing to provide service, regardless of the distribution of scrips among the players, is a Nash equilibrium under the minimum-scrip rule. This suggests that the minimum-scrip rule leads to a high level of social welfare.

Understanding Equilibria in Finite Assignment Markets

Assignment markets such as labor markets allow agents to match with each other and negotiate the payment that one agent must make to the other. In "Convergence of the Core in Assignment Markets," Kanoria, Saban, and Sethuraman study the set of equilibria in such markets as the number of agents grows. Their model incorporates agent types and a stochastic structure and is similar to models used in empirical studies of such markets. The authors show that when the number of agent types is kept fixed, the relative size of the set of equilibria (the "core") vanishes rapidly as the number of agents grows, justifying the typical assumption of a unique core outcome in such markets. The authors also establish convergence to a limit point. The results suggest that, given the market composition, wages are almost uniquely determined in equilibrium.

Prior-free, Budget-Feasible Reverse Auctions

How to make purchases given a fixed budget in a strategic environment solves the buyer's problem in procurement auctions (also known as reverse auctions), given no distributional assumptions on the sellers' reservation prices. In "Budget Feasible Procurement Auctions," Anari, Goel, and Nikzad explain that the designed auction is different from the "cut-off price" auctions that are typically used in procurement problems. This new design features nonlinear allocation and payment rules, providing a better worst-case guarantee than cutoff price auctions. In fact, this new design provides the best worst-case guarantee over all possible truthful and budget-feasible auctions.

Cost-per-Impression Pricing for Display Advertising

In online display advertising, publishers often face uncertain demand from advertisers requesting advertising space to post their ads and uncertain supply of visits from viewers. Although a publisher can guarantee serving an ad with a specific number of impressions, completing the service is highly uncertain as it depends on how many viewers will visit the page where the ad is posted. In such an inherently uncertain environment, pricing is one of the most challenging operational decisions that web publishers face, and, currently, mostly ad hoc approaches are being used. In "Cost-per-Impression Pricing for Display Advertising," Fridgeirsdottir and Najafi-Asadolahi use a queueing system model to determine the publisher's optimal price to charge per impression and show that it can increase in the number of impressions made of each ad. They also show that the optimal cost-per-impression price may increase in the number of ads rotating among slots.

Optimal and Practically Implementable Liver Cancer Surveillance Policies

Liver cancer is one of the deadliest cancers and is the fastest growing cause of cancer-related deaths in the United States. Infection with hepatitis C is the leading cause of primary liver cancer. Although liver cancer surveillance can improve health outcomes, the optimal use still remains unclear. In "Optimal M-Switch Surveillance Policies for Liver Cancer in a Hepatitis C–Infected Population," Chen, Ayer, and Chhatwal systematically examine the most cost-effective surveillance policies by developing a mixed-integer programming—based framework that captures two important practical concerns. First, the framework allows the flexibility of limiting the number of times the screening frequency could change over time (referred to as M-switch policies); second, the framework enables the policies to be tailored to subpopulations based on individuals' observable disease states. By calibrating their model to the best available evidence from med-

ical literature, the authors show that their proposed surveillance policies outperform a one-size-fits-all policy.

Data Uncertainty in Markov Chains: Application to Cost-Effectiveness Analyses of Medical Innovations

Markov chains are widely used in model-based evaluations of healthcare interventions (e.g., screening policies, treatment methods). However, uncertainty in the model parameters may pose a threat to the validity of the conclusions reached from such analyses. Existing methods for assessing the impact of parameter uncertainty are ill-suited to problem contexts where the number of uncertainties is large, as is the case when evaluating medical innovations. In "Data Uncertainty in Markov Chains: Application to Cost-Effectiveness Analyses of Medical Innovations," Goh, Bayati, Zenios, Singh, and Moore describe a new method for assessing the impact of parameter uncertainty that is more suitable to such contexts. The authors show that their approach is amenable to computational implementation and apply it to assess relative cost effectiveness of fecal immunochemical testing, a new screening approach for colorectal cancer.

Improving Patient Distribution to Hospitals in the Aftermath of a Disaster

A challenge in disaster response is choosing the destination hospital for each ambulance transporting a patient from the disaster, when both patients and hospitals are geographically dispersed and only limited information may be available about congestion levels at hospitals. The choice of hospital affects the length of time the ambulance is in use, which impacts its ability to make round trips to pick up additional patients. In "Dynamic Distribution of Patients to Medical Facilities in the Aftermath of a Disaster," Mills, Argon, and Ziya formulate this problem as a Markov decision process and propose two patient-routing heuristics. These dynamic policies are better than going to the nearest hospital, even when information about congestion at hospitals is not available in real time. Via a simulation case study, the authors demonstrate how emergency responders can better understand their capability to respond to specific disaster scenarios.

To Improve Portfolio Performance, Calibrate Your Transaction Costs

Mean-variance portfolios are the workhorse of the investment management industry, but two crucial elements in their implementation are estimation error and transaction costs. In "A Robust Perspective on Transaction Costs in Portfolio Optimization," Olivares-Nadal and DeMiguel show that these two elements are intimately related and propose a class of data-driven portfolios with improved performance. Theoretically, the paper shows that the portfolio problem with transaction costs is equivalent to three different portfolio problems designed to reduce estimation error. Motivated by these results, the authors propose a data-driven approach to portfolio optimization that consists of treating transaction costs as a regularization term to be calibrated. Empirically, the data-driven portfolios perform favorably because they strike an optimal trade-off between rebalancing the portfolio to capture the information in recent historical return data and avoiding the large transaction costs and impact of estimation error associated with excessive trading.

Managing an Inventory Portfolio with Limited Flexibility

What is the combined influence of customer substitution and limited supply capacity on a firm's optimal inventory policy for its product portfolio? In "Multiproduct Inventory Management Under Customer Substitution and Capacity Restrictions," Schlapp and Fleischmann derive the optimal policy for a single-season setup and investigate how this policy changes with market conditions. The authors reveal not only that the firm has to account for direct substitution effects in its decision-making process but that indirect effects may be even more important. These indirect substitution dynamics may imply counterintuitive inventory decisions, such as increasing a product's inventory level when more customers substitute away from it. The authors discuss several product categories for which these surprising effects are particularly strong. The article thus informs inventory managers about the intricate interplay between substitution and capacity limitations and encourages them to challenge their intuition concerning appropriate policies for managing inventory portfolios.

Estimating the Markov Chain Choice Model from Purchase Data

Being able to predict the choice behavior of the customers when they are offered a certain set of alternatives has a substantial value to a seller. A Markov chain choice model can be used to capture the choice process of the customers when the customers choose and substitute among the available products. Although this model is flexible for capturing the choice process, exploiting this flexibility in practice depends on efficiently calibrating the parameters of the model from the past purchase history of the customers. In "An Expectation-maximization Algorithm to Estimate the Parameters of the Markov Chain Choice Model," Şimşek and Topaloglu fill this gap by developing an efficient parameter calibration algorithm based on the expectation-maximization method. The authors show that the Markov chain choice model, coupled with their algorithm, can yield better predictions of customer choice behavior compared with other commonly used alternatives.

How Do Loss-Averse Customers Behave in Service Systems, and What Is the Pricing Strategy?

When customers are loss averse with respect to both price and waiting time in an unobservable queue, what will be their joining strategies? And what is the corresponding pricing strategy? In "Service Pricing with Loss-Averse Customers," Yang, Guo, and Wang develop a framework to consider customer loss-aversion behavior with respect to multiple attributes in a queueing system. The authors show that with customer loss aversion, there may exist three equilibrium joining strategies. The authors also find that loss aversion polarizes queues, making long queues even longer and short queues even shorter, which provides a new explanation for herding behavior in queues. The authors also study pricing competition in a symmetric duopoly market. Interestingly, with loss-averse customers, a firm can obtain a higher profit in a duopoly market than in a monopoly market.

Optimizing Continuous Berth Allocation: New Lower Bound and Exact Method

When vessels arrive at a port, they need to be assigned a mooring time and a berth location on a continuous quay.

Such a problem of continuous berth allocation is a crucial element in port planning. However, the performance of its existing solution methods is far from satisfactory, as the best known lower and upper bounds on the optimal berth allocation have a large gap. In "New Lower Bound and Exact Method for the Continuous Berth Allocation Problem," Xu and Lee introduce a novel relaxation of the problem, which leads to a new and much better lower bound, one that can also be computed efficiently. By utilizing the new lower bound, the authors develop an exact method that significantly reduces the gap between the lower and upper bounds. The bound and method are then extended to a more complicated problem with a quay crane constraint, demonstrating their general applicability.

When to Abandon an Unresolved Project

Academics and corporate innovators spend their workdays pursuing breakthroughs that may never materialize. When working on a project, it can be difficult to decide when to cut your losses. Optimism tells us that just beyond our grasp hovers the solution that will make it all work. On the other hand, the end we have in sight may well be a dead end. Every day we spend chasing a mirage wastes valuable time and resources. Absent a glaring signpost of failure, how does one know when it makes strategic sense to abandon an idea? In "When to Abandon a Research Project and Search for a New One," McCardle, Tsetlin, and Winkler help answer this question systematically. The article presents a prescriptive analytical model designed to mirror a researcher's real-life situation. There is a risky project being worked on, with no reliable information regarding how close to completion the project actually is. The option to drop the project and resume searching for a more worthwhile endeavor is always available. To obtain an accurate-enough estimate of when to move on, a researcher needs to only know two things: the arrival rate of new projects (i.e., the speed at which the search process generates results) and the arrival rate of success (i.e., how quickly a researcher is usually able to finish a project). The timing is relatively insensitive to the probability of successfully completing a project.

Robust Optimization: Certainty in the Face of Uncertainty

Optimization of systems such as inventory control is commonly subject to uncertainties in the form of the (partially) unknown demand probability distribution. Traditionally, starting with the work of Markowitz, knowledge of mean and variance is assumed. In "Robust Optimization with Ambiguous Stochastic Constraints Under Mean and Dispersion Information" by Postek, Ben-Tal, den Hertog, and Melenberg, knowledge of mean and mean absolute deviation of the random variables is assumed, and tight expressions of the robust expected values of the objective are derived based on the 1972 result of Ben-Tal and Hochman, which gives tight approximations of the expectation of general convex (concave) function. To reduce computational effort, the bounds are applied on the aggregated random variables (such as cumulative demand). The authors also show that this information can be used to construct good approximations of chance constraints.

Efficient Algorithms with Performance Guarantees for Approximating the Optimal Solution to a Class of Stochastic Selection Problems

Resource allocation problems often involve a first stage where a subset of options from a larger set of candidates needs to be selected. When uncertainties are present in the magnitude of the resources as well as costs and rewards, the resulting optimization problem is an NP-hard, mixed-integer, nonlinear two-stage stochastic program. In "Approximation Algorithms for a Class of Stochastic Selection Problems with Reward and Cost Considerations," Strinka and Romeijn develop efficiently implementable approximation algorithms that use a combination of rounding (of selection variables), scaling (of continuous variables), and sample average approximation (to deal with uncertainties). Under complete recourse in the second stage and a scalability property of the cost and reward terms in the objective function, these algorithms produce solutions that are, with high probability, provably close to optimality. The results apply in risk-neutral as well as risk-averse settings, under very mild assumptions on the randomness of the problem parameters (including dependence between the options).

Conic Programming Reformulations of Two-Stage Distributionally Robust Linear Programs over Wasserstein Balls

Dynamic decision problems under uncertainty are prevalent in many areas of engineering, science, and economics. These decision problems involve recourse decisions that adapt to the information observed over time and are usually solved approximately by restricting the adaptive decisions to simple parametric decision rules. However, the corresponding approximation error can be substantial. In "Conic Programming Reformulations of Two-Stage Distributionally Robust Linear Programs over Wasserstein Balls," Hanasusanto and Kuhn show that two-stage robust and distributionally robust linear programs can often be reformulated exactly as conic programs that scale polynomially with the problem dimensions. The resulting optimization problems are generically intractable but amenable to tight approximations in semidefinite programming (SDP). Their numerical tests demonstrate that the SDP approximations distinctly outperform the state-of-the-art decision rule approximations in terms of accuracy.

A Statistical Inference Perspective on Inverse Optimization

Inverse optimization refers to the inference of unknown model parameters of an optimization problem based on its optimal solutions. In "Inverse Optimization with Noisy Data," Aswani, Shen, and Siddiq approach inverse optimization through the lens of statistical inference. They propose the first formulation for inverse optimization that produces estimates of the unknown model parameters that are statistically consistent. The authors propose two numerical algorithms for solving the inverse optimization and demonstrate its competitiveness with existing methods using both synthetic and real data sets.