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(6):ii-v. https://doi.org/10.1287/opre.2018.1824

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

OPERATIONS RESEARCH
Vol. 66, No. 6, November–December 2018, pp. ii–v
ISSN 0030-364X (print), ISSN 1526-5463 (online)

In This Issue

https://doi.org/10.1287/opre.2018.1824

Efficiently Integrating out Unobserved States in Structural Estimation Problems

A common problem in the estimation of structural economic decision models is that only a subset of the information available to the decision maker is present in the data. Therefore, an econometrician estimating the model needs to "integrate out" these unobserved states. Until recently, all approaches to this problem involved Monte Carlo techniques, making maximum likelihood estimation difficult due to the presence of simulation noise. In "Divide and Conquer: Recursive Likelihood Function Integration for Hidden Markov Models with Continuous Latent Variables," Reich develops an efficient and fully deterministic integration scheme to recursively integrate out unobserved states when computing the likelihood function. Moreover, he proves numerical error bounds and convergence rates that suggest that the method is highly efficient under appropriate smoothness conditions, which is confirmed in several examples. In particular, he estimates an extension of the famous bus engine replacement model of Rust (1987) allowing for serially correlated errors.

Optimal Wildlife Corridor Design

Wildlife corridors connect patches of critical habitat for species in need of protection, especially migratory species. Disconnected and fragmented habitats are inadequate to ensure the survival of sensitive species that need enough space to find food or to breed. Also, long, winding, and narrow corridors might not provide enough protection from the surrounding human-dominated landscape. To create wildlife corridors, conservation organizations strategically acquire rights to land parcels with budget restrictions. Thus, optimal use of scarce funds to assemble corridors with the right geometry is a challenge. In this issue, St John, Tóth, and Zabinsky present a model (in "Optimizing the Geometry of Wildlife Corridors in Conservation Reserve Design") that combines path planning from artificial intelligence with network optimization to provide optimal control of width, length, and other spatial attributes of wildlife corridors. The authors' work is applicable to other fields, such as firebreaks, robotics, drone navigation, and beyond.

Explaining the Emergence of Fashion Cycles from the Consumer Side

Fashion has significant economic and social impacts. A key aspect of fashion is fashion cycle, whose emergence is an intriguing phenomenon. Most previous studies assume that producers play a dominant role. However, the role of consumers is increasingly important in today's flatter and flatter world. In "Fashion and Homophily," Zhang, Cao, Qin and Yang investigate the evolution of fashion from the consumer side using a network game model. Consumers are represented by the nodes of a social network and are classified into conformists and rebels, who prefer to coordinate and anti-coordinate with their neighbors' choices, respectively. The main finding is that a lower homophily index generally promotes the emergence of fashion cycles. In other words, interactions between the same type of consumers (resp. between consumers of different types) tend to inhibit (resp. promote) the emergence of fashion cycles.

Games and Dynamics for Power Control on Random Wireless Networks

The explosive growth of smart devices together with their dense connections via wireless networks has become a ubiquitous feature of smart cities. These devices, which communicate with one another, typically have heterogenous operating characteristics, such as quality-of-service metrics, power costs, etc. As such, it is a both theoretically interesting and practically impactful problem to understand the system behavior as they interact with each other to form a coherent ecosystem. In "Deterministic and Stochastic Wireless Network Games: Equilibrium, Dynamics and Price of Anarchy," Zhou, Bambos, and Glynn take a game-theoretical approach to study the power control aspect of the wireless network. They examine the equilibrium properties in a large, unbounded power regime and provide stochastic stability guarantees of an important class of power control dynamics under random network environments, thereby contributing rich structural insights of the overall network behavior.

On the Optimality of Three-Part Tariff Plans

Businesses and organizations use many methods setting prices for different levels of services. A popular choice is two-part tariffs, combining a fee and a marginal price for each unit. For example, OpenTable charges restaurants a monthly fee and a per-diner fee. Three-part tariffs combine an allowance with a fixed access fee and a marginal price for above-allowance consumption. For example, SurveyMonkey charges a flat fee of $26 per month for up to 1,000 responses in a month and $0.15 for each response above 1,000. In "On the Optimality of Three-Part Tariff Plans: When Does Free Allowance Matter," Bhargava and Gangwar simplify the optimal three-part tariff design and the impact on market outcomes and profit. In most situations, a two-part tariff can produce the same outcome and profit as the optimal three-part tariff. Three-part tariffs are needed when consumers have biases, such as overconfidence or overestimation of their consumption, or if the market has segments with highly sharp across-segment differences.

Closed-Form Solutions for Worst-Case Risk Measures and Distributions

A major challenge in risk measurement is that the key parameter, namely the loss distribution, is notoriously difficult

to specify. As an alternative, the worst-case approach that relies only on the mean and variance information to find the worst possible risk level has received much attention. It is known that for popular risk measures, such as Value-at-Risk (VaR) and Conditional Value-at-Risk (CVaR), their worst-case values can be calculated in closed-form. In "Closed-Form Solutions for Worst-Case Law Invariant Risk Measures With Application to Robust Portfolio Optimization," Li shows that such closed-form solutions, in fact, exist for any (law invariant) coherent risk measure. The author further provides worst-case distributions characterized in terms of risk spectrums and presents the applications of the closed-form results in portfolio optimization.

What Determines Margin Requirements?

Margin requirements provide protection against counterparty defaults, and have been known to serve as a conduit to tightening credit and leverage spirals during crisis periods. In "Clearinghouse Margin Requirements," Agostino Capponi and Cheng analyze the theoretical underpinnings of clearinghouse margin and fee requirements for centrally cleared trades. The authors explicitly relate margin levels to market variables, including market risk, credit risk, and funding costs, thus identifying a novel channel through which margins can tighten. The analysis provided in the paper highlights the crucial role of margin requirements in the presence of asymmetric information, displays a rich interplay between fee and margin levels, explains variation in margins for futures contracts with different tenors, and yields insights for systemic risk analysis via the margin shortfall.

Pricing Under Generalized Extreme Value (GEV) Models

In pricing applications, it is important to capture the substitution among products because when the price of a product increases, customers may choose other items, creating interactions between the demands across different products. Discrete choice models can capture the substitution and provide more realistic models of demand. Generalized Extreme Value (GEV) models encompass a large class of discrete choice models. In "Multi-Product Pricing Under the Generalized Extreme Value Models with Homogeneous Price Sensitivity Parameters," Zhang, Rusmevichientong, and Topaloglu study the multi-product price optimization problem under an arbitrary GEV model, and all products have the same price sensitivity parameter. With no constraint, the optimal prices exhibit a constant-markup property, and there is an explicit formula for the optimal markup in terms of the Lambert-W function. With convex constraints on the expected sales, the problem can be transformed to an equivalent convex program that can be solved efficiently.

Adjacency-Clustering: A New Clustering Method that Exploits Neighborhood Effects

In integrated circuit manufacturing, it is known that defects tend to be clustered on the wafer and a chip is likely to be defective if its neighbors are defective. This Neighborhood Effect is not well captured in traditional yield modeling approaches. In "Adjacency-Clustering and Its Application for Yield Prediction in Integrated Circuit Manufacturing," Hochbaum and Liu propose a new clustering method, called adjacency-clustering, that captures neighborhood effects, and demonstrate how to use it as a new yield prediction model. The authors compare the performance of the adjacency-clustering model to that of leading yield prediction models including the Poisson, the negative binomial, the Poisson regression, and negative binomial regression models, on real datasets and on simulated datasets. The results demonstrate that the adjacency-clustering model captures the neighborhood effect and delivers superior prediction accuracy. Moreover, the concept and methodology of adjacency-clustering are applicable in any context where a neighborhood effect is present, such as disease risk mapping and energy consumption prediction.

A Machine Learning Approach to Inventory-Constrained Dynamic Pricing

Thompson sampling is a randomized Bayesian machine learning method, whose original motivation was to sequentially evaluate treatments in clinical trials. In recent years, this method has drawn wide attention, as Internet companies have successfully implemented it for online ad display. In "Online Network Revenue Management Using Thompson Sampling," Ferreira, Simchi-Levi, and Wang propose using Thompson sampling for a revenue management problem where the demand function is unknown. A main challenge to adopt Thompson sampling for revenue management is that the original method does not incorporate inventory constraints. However, the authors show that Thompson sampling can be naturally combined with a linear program formulation to include inventory constraints. The result is a dynamic pricing algorithm that incorporates domain knowledge and has strong theoretical performance guarantees as well as promising numerical performance results. Interestingly, the authors demonstrate that Thompson sampling achieves poor performance when it does not take into account domain knowledge. Finally, the proposed dynamic pricing algorithm is highly flexible and is applicable in a range of industries, from airlines and internet advertising all the way to online retailing.

Wasserstein Distance and the Distributionally Robust TSP

One of the most common strategies used in large-scale logistics problems is *districting*, in which one divides a service region into smaller pieces. This is particularly useful when one does not know the true locations of demand, but has only limited information, such as a probability distribution or a small set of sample data. In "Wasserstein Distance and the Distributionally Robust TSP," Carlsson et al. show how to partition a region so as to balance the traveling salesman tours of a set of unknown points, using a distributionally robust method defined by a bound on the Wasserstein distance to a set of landmark points.

Intercepting the Unlawful Vehicles in Transportation Networks

In a transportation network, the regulatory authority locates law enforcement facilities to intercept unlawful vehicles, whereas these vehicles may deviate from their route up to a certain tolerance factor to avoid being intercepted. Taking this noncooperative behavior into account, the objective of the

Operations Research, 2018, vol. 66, no. 6, pp. ii–v, © 2018 INFORMS

"evasive flow capturing problem" is to determine the location of the law enforcement facilities. The problem has applications in transportation, revenue management, and security management. In "Exact Solution of the Evasive Flow Capturing Problem," Arslan, Jabali, and Laporte develop a bilevel model and reformulate it using polyhedral theory to accelerate the solution process and to expand the size of the solvable instances of this highly complex problem.

Performance Analysis of a Simple Policy for Stochastic Scheduling Problems

Scheduling problems have a deep and well-developed literature in both operations research and computer science. Stochastic scheduling problems with many jobs are notoriously difficult, and optimal policies may require constant tracking of the elapsed time of all jobs in process. When the processors (or "machines") are identical, a policy that simply schedules jobs in fixed order of job weight to expected processing time—the WSEPT rule—is asymptotically optimal when the number of jobs is large. Much less understood is the case with specialized (or "unrelated") machines (i.e., each job's processing distribution may vary across machines). In "Static Routing in Stochastic Scheduling: Performance Guarantees and Asymptotic Optimality," Balseiro, Brown, and Chen study stochastic scheduling with unrelated machines. The authors study a simple static routing policy that (i) assigns jobs to machines up front and (ii) schedules the jobs on each machine in the WSEPT order. This static routing policy depends only on job processing times through their expected values and is easy to compute with a single convex optimization problem. The authors explicitly characterize the performance loss of this static routing policy relative to an optimal scheduling policy; this result implies that this static routing policy is asymptotically optimal in the regime of many jobs.

Tight Approximations for Assortment Optimization Under Ranking Preferences

Assortment optimization has received significant attention in recent revenue management and combinatorial optimization literature. In "The Approximability of Assortment Optimization Under Ranking Preferences," Aouad, Farias, Levi, and Segev provide best-possible approximability bounds for this problem under an almost general model specification, where preferences are expressed as a distribution over rankings. This paper shows how this optimization problem relates to the computational task of detecting large independent sets in graphs, allowing the establishment of strong complexity lower bounds with respect to various problem parameters. These findings are complemented by a number of algorithms that attain essentially best-possible approximation factors, proving that the hardness results are tight up to lower-order terms. Surprisingly, their results imply that a simple and widely studied policy, known as revenue-ordered assortments, achieves the best possible performance guarantee with respect to prices.

Robust Optimization with First-Order Algorithms

Robust optimization (RO) is a technique to tractably model uncertain parameters in optimization problems. More recently, it has attracted interest in applications from machine learning and statistics. These recent applications present algorithmic challenges in which the scalability of RO algorithms with problem dimension becomes crucial. The traditional solution method for RO is to transform it into an equivalent—yet more complex—deterministic problem. The alternate solution technique is to iteratively solve sequences of the underlying deterministic model with different values of the uncertain parameters. However, such iterative approaches have been rather prohibitive in practice, especially when solving even the underlying deterministic model is expensive. In "Online First-Order Framework for Robust Convex Optimization," by analyzing the structure of an underlying convex–nonconcave saddle point problem, Ho-Nguyen and Kılınç-Karzan develop an iterative framework for RO in which the cost of each iteration can be remarkably reduced. In particular, they show that, without scarifying from the guarantees on the number of iterations needed, it is possible to use cheap first-order updates instead of deterministic optimization solvers in each iteration.

Ordinal Optimization with Limited Distributional Information

In "Tractable Sampling Strategies for Ordinal Optimization," Shin, Broadie, and Zeevi analyze a problem of ordinal optimization where the objective is to select the best of several competing systems, when the probability distributions governing each system's performance are not known but can be learned via sampling. The objective is to dynamically allocate samples within a finite sampling budget to maximize the likelihood of identifying the best system. An exact solution to this problem over any finite time horizon is difficult to characterize. In lieu of that, the authors introduce a family of practically implementable sampling policies and characterize the set of problem instances over which their performance (over a long time horizon) is essentially the best possible. Furthermore, the authors show via numerical testing that the proposed policies perform well compared with other benchmark policies over finite time horizons.

Sample Size One Is Sufficient

Finding optimal policies when the objective function can only be observed with a noisy error has traditionally required repeated observations to average out the noise. In Single Observation Adaptive Search for Continuous Simulation Optimization, Kiatsupaibul, Smith, and Zabinsky prove for a very large class of continuous stochastic optimization problems and search algorithms that a single observation for each candidate solution suffices to completely eliminate the error. This surprising result can best be understood through the metaphor of facing a continuum of slot machines and observing the reward for a single pull of the arm for each machine as the authors search for a machine delivering the best return. By the impossibility of systems, the average winnings must approach the expected winnings regardless of how the authors select the sequence of machines to test. This same martingale property is inherited for simulation optimization so that a single observation also suffices for convergence to the global optimum. The resulting savings in simulation time can be quite significant.

Performance of Fork-Join Queueing Networks at Scale

Cloud computing and big data analytics are revolutionary technologies benefiting numerous enterprises

and organizations. The underlying parallel and distributed processing systems are essential to achieving high performance. Although engineering efforts have long been made to build such systems of increasing size and complexity, the mathematical foundation toward the understanding of performance limits in large scale remains challenging. A critical issue concerns throughput scalability: whether throughput decreases to zero as these systems scale in size and capabilities. In "Throughput Scalability Analysis of Fork-Join Queueing Networks," Zeng et al. model parallel and distributed processing systems as fork-join queueing networks with blocking and establish necessary and/or sufficient conditions on throughput scalability of such networks. Technology advances in storage space and processing speed are also analyzed with their potential in mitigating throughput degradation. The results provide great insights and useful guidelines for designing next-generation parallel and distributed processing systems at scale.

Fear of the Market or Fear of the Competitor? Ambiguity in a Real Options Game

In "Fear of the Market or Fear of the Competitor? Ambiguity in a Real Options Game," Hellmann and Thijssen study the effects of managerial ambiguity aversion on investment decisions in a duopoly. Ambiguity is over the investment project's expected growth rate (the "trend"). In their continuous-time model they assume that investment commands a first-mover advantage, which, in equilibrium, leads to pre-emption. They show that ambiguity reduces pre-emptive pressure. As a consequence firm value may be increasing in the degree of ambiguity. They also find that the worst-case prior (under which managers compute firm value) is not always given by the lowest trend, but can switch over time between the lowest and the highest trend. Finally, if one of the firms is not ambiguity averse, then its value is increasing in the other firm's degree of ambiguity.