Operations Research

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

In This Issue

To cite this article:
(2024) In This Issue. Operations Research 72(6):iii-vi. https://doi.org/10.1287/opre.2024.inthisissue.v72.n6

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

OPERATIONS RESEARCH
Vol. 72, No. 6, November–December 2024, pp. iii–vi
ISSN 0030-364X (print), ISSN 1526-5463 (online)

In This Issue

https://doi.org/10.1287/opre.2024.inthisissue.v72.n6

Finding the Right Business Model for Durable Goods

Technological advances enable new ways for customers to access durable goods such as heavy equipment without the need for ownership. In "Business Model Choice for Heavy Equipment Manufacturers," Blaettchen, Taneri, and Hasija analyze the resulting business models available to manufacturers, considering that manufacturers need to effectively coordinate their sales strategy with their highly profitable after-sales service activities and take into account secondary markets. Employing a game-theoretic framework, the authors show how different equipment characteristics lead to different optimal business model designs. They provide a comprehensive overview of when manufacturers should rely on emerging models based on servicization or peer-to-peer sharing and when it pays to retain a traditional model based on sales. Drawing on a novel framework for analyzing a business model's environmental impact, they further show that emerging models can create win-win situations for the manufacturer and the environment.

Side-Constrained Dynamic Traffic Equilibria

Dynamic flows are a well-studied model for car traffic in road networks. The assumption that every driver chooses her route in such a way as to selfishly minimize her own travel costs leads to the solution concept of dynamic equilibria, that is, dynamic flows wherein every flow particle travels along a cost minimal route (under the travel costs induced by this flow). In practice, however, there are often additional constraints on certain parts of the networks that restrict the options of the individual drivers like traffic limits due to security concerns in tunnels or to keep emission levels in some areas below certain thresholds. In "Side-Constrained Dynamic Traffic Equilibria," Graf and Harks develop a general framework for incorporating such additional side constraints into dynamic flow models. They provide characterization results for the resulting equilibria via (quasi-)variational inequalities and show the first existence result for the nonconvex setting of volume constraints.

DDP: A Duality-Driven Dynamic Programming Algorithm

The curse of dimensionality significantly restricts the use of dynamic programming methods in solving complex problems. Consequently, researchers and practitioners often resort to approximate (suboptimal) control policies that strike a balance between ease of implementation and satisfactory performance. Information relaxation-based duality techniques generate both upper and lower bounds for the true values of stochastic dynamic programming problems, allowing us to evaluate the optimality of approximate policies through the dual gap. However, the literature still lacks guidance on handling cases where the gaps are excessively loose. In "Information Relaxation and a Duality-Driven Algorithm for Stochastic Dynamic Programs," Chen, Ma, Liu, and Yu develop a novel DDP framework to obtain and tighten confidence interval estimates for the true value functions of SDP problems. Leveraging a new finding that the dual operation yields subsolutions, they establish convergence guarantees for DDP. Additionally, a regression-based Monte Carlo method is introduced, aimed at high-dimensional applications. Numerical examples demonstrate that DDP effectively improves dual gaps for various heuristics that are commonly used in the literature.

Understanding Financial Networks Through Incentive-Aware Models

A financial network is a web of contracts between firms. Each firm wants the best possible contracts. However, a contract between two firms requires the cooperation of both firms. This contest between cooperation and competition is studied in "Incentive-Aware Models of Financial Networks" by Jalan, Chakrabarti, and Sarkar. The authors show how contract negotiations lead to a stable network where no firm wants to change contract sizes. In this network, the size of any contract depends on the beliefs of all firms, not just the contract's two parties. Minor news about one firm can affect these beliefs, causing drastic changes in the network. Moreover, under realistic settings, a regulator cannot trace the source of such changes. This research illustrates the importance of firms' beliefs and their implications for network stability. The insights could inform regulatory strategies and financial risk management.

Enhancing CoVaR Estimation Efficiency Through Monte-Carlo Simulation

CoVaR is an important measure of financial systemic risk due to its ability to capture tail dependence between the losses of different portfolios and its capacity to predict financial crises. Estimating CoVaR is challenging because its definition involves a zero-probability event, which is unobservable in the data. The existing model-based methods address this issue by assuming simplified structural models, which introduce biases that are difficult to eliminate. In "Monte Carlo Estimation of CoVaR," Huang, Lin, and Hong propose using Monte Carlo methods to estimate CoVaR, leveraging the modeling flexibility of Monte Carlo simulation. Specifically, they introduce a batching estimator applicable to a wide range of financial models and prove that its best rate of convergence is $n^{-1/3}$, where n is the sample size. Under the widely used delta-gamma approximation model, they further introduce an importance sampling–inspired estimator and prove that its best rate of convergence can be improved to $n^{-1/2}$.

Random Consideration Set Model

In "A Random Consideration Set Model for Demand Estimation, Assortment Optimization, and Pricing," Gallego and Li operationalize a microfounded consumer choice model—the random consideration set (RCS) choice model of Manzini and Mariotti [Manzini P, Mariotti M (2014) Stochastic choice and consideration sets. *Econometrica* 82(3):1153–1176]—that captures the limited attention of consumers, assuming that purchases are based on fixed preference orderings with consideration sets formed from independent attentions. The authors provide a condition for uniquely identifying model parameters and design an efficient algorithm for model parameters estimation. The authors offer a greedy-like algorithm for assortment optimization, adaptable for optimal assortment subject to cardinality constraint or discovering efficient sets. The authors extend the model to consider random product preferences, with a $1/2$ performance-guaranteed approximation algorithm. Using data from a major U.S. airline, the authors find that the RCS model outperforms the mixed multinomial logit model in approximately half of the markets, particularly with smaller, less varied data sets.

Dynamic Assignment Algorithms to Boost Refugee Outcomes

Amid record-breaking forced displacement in recent years, researchers and policymakers alike have become increasingly interested in the idea of algorithmically matching refugees to geographic localities in order to optimize their employment or other integration outcomes. In "Outcome-Driven Dynamic Refugee Assignment with Allocation Balancing," Bansak and Paulson propose new dynamic assignment algorithms for this context designed to maximize a given outcome while meeting the operational needs of refugee resettlement and asylum agencies. Using resettlement data from both the United States and Switzerland, the authors demonstrate how one algorithm (currently being piloted in Switzerland) can achieve near-optimal results compared with a hindsight-optimal matching. They also show that, because of nonstationarities in the arrival process, outcome maximization (even when subject to location capacity constraints) can result in an imbalanced allocation to localities over time, putting periodic strains on limited local resettlement resources. They account for this problem in a second algorithm that achieves near-perfect balance over time with only a small loss in average outcomes compared with the first algorithm.

To Interfere or Not To Interfere: Information Revelation and Price-Setting Incentives in a Multiagent Learning Environment

Demand uncertainty and seller competition are substantial challenges for online platforms. In "To Interfere or Not To Interfere: Information Revelation and Price-Setting Incentives in a Multiagent Learning Environment," Birge, Chen, Keskin, and Ward analyze whether and how an online platform should offer demand information or price incentives to the sellers participating on the platform. The authors show that, when facing uncertain demand, the platform could be better off by doing nothing—that is, not providing any information or incentives to the sellers. They also develop a *strategic reveal-and-incentivize policy* for the platform to choose when to start sharing information and offering rewards to coordinate the sellers' pricing. They prove that the strategic reveal-and-incentivize policy achieves near-optimal profit performance for the platform.

Optimal Auction Design with Deferred Inspection and Reward

In a variety of auction contexts, particularly those facilitated by digital marketplaces, auctioneers are now empowered to review the valuations submitted by purchasers and modify the final payments accordingly. In "Optimal Auction Design with Deferred Inspection and Reward," Alaei, Belloni, Makhdoumi, and Malekian formulate and study the optimal design of such auctions, demonstrating that both the analysis and implementation vastly depart from classical settings. In particular, the analysis requires tools from convex analysis and the calculus of variations, which could be of independent interest in other mechanism design questions.

Model-Based Reinforcement Learning for Offline Zero-Sum Markov Games

In "Model-Based Reinforcement Learning for Offline Zero-Sum Markov Games," Yan, Li, Chen, and Fan make progress toward learning Nash equilibria in two-player, zero-sum Markov games from offline data. Despite a large number of prior works tackling this problem, the state-of-the-art results suffer from the curse of multiple agents in the sense that their sample complexity bounds scale linearly with the total number of joint actions. The authors propose a new model-based algorithm, which provably finds an approximate Nash equilibrium with a sample complexity that scales linearly with the total number of individual actions. The authors also develop a matching minimax lower bound, demonstrating the minimax optimality of the proposed algorithm for a broad regime of interest. An appealing feature of the result lies in algorithmic simplicity, which reveals the unnecessity of sophisticated variance reduction and sample splitting in achieving sample optimality.

How to Manage Horizontally Differentiated Products When Customer Preferences and Demand Distributions Are Unknown

In "Pricing and Positioning of Horizontally Differentiated Products with Incomplete Demand Information," den Boer, Chen, and Wang consider the problem of determining the optimal prices and product configurations of horizontally differentiated products when customers purchase according to a locational choice model and where the problem parameters are initially unknown to the decision maker. The authors propose a data-driven algorithm that learns the optimal prices and product configurations from accumulating sales data, and the authors show that their regret—the expected cumulative loss caused by not using optimal decisions—after T time periods is $O(T^{1/2+o(1)})$. The authors accompany this result by proving an almost-matching lower bound of regret, implying that their algorithms are asymptotically near optimal. In an extension, the authors show how their algorithm can be adapted for the case of fixed locations.

Blockbuster or Niche? Competitive Strategy Under Network Effects

The "long tail" theory was celebrated by *BusinessWeek* as the biggest idea of 2004, soon after the book *The Long Tail* by Chris Anderson was published. The long tail theory calls for applying a low-budget strategy—producing a (relatively) large number of products with (relatively) low investment levels. However, some other cultural industries may tell a different story. The concentration of the most popular titles in the video game industry is growing, a phenomenon known as the blockbuster phenomenon. This phenomenon suggests that firms may adopt a high-budget strategy—producing a (relatively) small number of products with (relatively) high investment levels. In "Market Entry and Competition Under Network Effects," Feng and Hu analytically study the impact of a network effect on entry decisions and investment strategies (i.e., the high-budget versus low-budget strategies) adopted by competing firms based on which they further provide a theory that links the ex post sales volume concentration with the ex ante product variety in a market under network effects.

Dynamic Routing of Queues with Heterogeneous Server Pools

In "The Generalized c/μ Rule for Queues with Heterogeneous Server Pools," Long, H. Zhang, J. Zhang, and Z. G. Zhang study the optimal control of queueing systems with heterogeneous server pools and a single customer class. The goal is to balance the holding cost of the queue with the operating costs of the server pools. The authors introduce the target-allocation policy, the Gc/μ rule, and the fixed priority policy for systems with general, convex, and concave cost functions, respectively. The authors also consider an extension to minimize operating costs and maintain a service-level target for customers waiting in the queue. Moreover, the authors show that their asymptotically optimal routing policies coincide with several classic policies in the literature in special cases.

Scoring in Multi-Event Tournaments

How much is a first place worth? Is the athlete who came first once and third twice better than the athlete who came second three times? Different sports value these positions differently: (60, 54, 48) in IBU biathlon, (25, 18, 15) in F1 racing, (8, 7, 6) in Diamond League athletics. Are these choices based on anything? Is there even a rational way to choose a scoring vector of arbitrary length to rank any number of athletes? In "How Should We Score Athletes and Candidates: Geometric Scoring Rules," Kondratev, Ianovski, and Nesterov investigate two approaches to how the problem of choosing a scoring vector for a tournament can be reduced to the choice of a single parameter: an axiomatic approach based on eliminating spoilers and an optimization approach based on maximizing the expected quality of the winner. Intriguingly, the vectors generated by the second approach are uncannily similar to those used in real sporting events.

Learning and Leveraging Similarities Between Choices

In large-scale simulation optimization, it is impossible to exhaustively simulate every choice. However, there are often inherent similarities between choices: for example, two similar sets of input settings to a simulation model can reasonably be expected to produce similar output. The information gained from simulating one choice can thus be used to infer the values of other similar choices, enabling learning more from a relatively small number of samples. In "Sequential Learning with a Similarity Selection Index," Zhou, Fu, and Ryzhov develop a new similarity model to improve the final selection decision after all samples have been collected. The new "similarity indices" are complementary to all existing information collection procedures, which do not focus on the final decision. At the same time, the new model allows a tractable theoretical treatment of an optimal procedure, which can be efficiently approximated.

A (Slightly) Improved Approximation Algorithm for Metric TSP

In "A (Slightly) Improved Approximation Algorithm for Metric TSP," Karlin, Klein, and Oveis Gharan design the first improvement over the classical 1.5 approximation algorithm of Christofides-Serdyukov after more than 40 years. Their algorithm first chooses a random spanning tree from the maximum entropy distribution of spanning trees with marginals equal to the optimum LP solution of TSP, and then, similar to Christofides' algorithm, it adds the minimum cost matching on the odd degree vertices of the tree. To analyze their simple algorithms, the authors prove and exploit new tools from the theory of strongly Rayleigh distributions.

Modeling Heterogeneity of Subpopulations in Sequential Decision-Making

When applying decision models, we often estimate input parameters using data. In healthcare and some other applications, data are collected from a population of different entities, such as patients. Thus, one faces a modeling question of whether to estimate different models for subpopulations (called stratifying). The potential benefit of stratifying comes from the heterogeneity of subpopulations. For example, patients who progress faster than others require a separate model and a tailored treatment plan. In "Is Separately Modeling Subpopulations Beneficial for Sequential Decision-Making?," Lee provides theoretical results and empirical methods for deciding whether to stratify subpopulations. The author also presents how to use its results to select the best stratification among many. Improving medical decisions by tailoring to each subpopulation is a building block of precision medicine, and thus, this work aligns closely with the precision medicine paradigm.

Using Adaptive Importance Sampling to Solve Circular Challenge in Stochastic Root Finding

Stochastic root-finding problems are fundamental in the fields of operations research and data science. However, when the root-finding problem involves rare events, crude Monte Carlo can be prohibitively inefficient. Importance sampling (IS) is a commonly used approach, but selecting a good IS parameter requires knowledge of the problem's solution, which creates a circular challenge. In "Adaptive Importance Sampling for Efficient Stochastic Root Finding and Quantile Estimation," He, Jiang, Lam, and Fu propose an adaptive IS approach to untie this circularity. The adaptive IS simultaneously estimates the root and the IS parameters, and can be embedded in sample average approximation–type algorithms and stochastic approximation–type algorithms. The authors provide theoretical analysis on strong consistency and asymptotic normality of the resulting estimators, and show the benefit of adaptivity from a worst-case perspective. They also provide specialized analyses on extreme quantile estimation under milder conditions.

Updating a Classic: Assortment Optimization Under the Multi-Purchase MNL Model

In "Assortment Optimization Under the Multi-Purchase Multinomial Logit Choice Model," by Bai, Feldman, Segev, Topaloglu, and Wagner, the authors' primary contribution resides in proposing the first multi-purchase choice model that can be fully operationalized. The authors' main algorithmic results consist of two distinct polynomial time approximation schemes (PTAS); the first, and simpler of the two, caters to a setting where each customer may buy only a constant number of products, whereas the second, more nuanced algorithm applies to the authors' multi-purchase model in its general form. Additionally, the authors study the revenue potential of making assortment decisions that account for multi-purchase behavior in comparison with those that overlook this phenomenon. In particular, the authors relate both the structure and revenue performance of the optimal assortment under a traditional single-purchase model to that of the optimal assortment in the multi-purchase setting.

A Method of Computational Analytics for Mutual Reserve Optimization with an International Perspective

As well studied in the operations research literature, optimization of a mutual reserve system (e.g., federal reserves) and a nonmutual one such as regular inventory systems requires solving simultaneous

In This Issue
Operations Research, 2024, vol. 72, no. 6, pp. iii–vi, © 2024 INFORMS

systems of quasi-variational inequalities, of which analytical solutions in closed form remain unattainable and computational solutions are still intractable. Thus far, the studies of reserve optimization are of intra-nations (e.g., central bank reserves) as opposed to internations (e.g., COVID vaccine reserves of the United Nations). In "A Splitting Method for Band Control of Brownian Motion: With Application to Mutual Reserve Optimization," Bensoussan, Liu, and Yuan advance a method of computational analytics for mutual reserve optimization, with an international perspective in response to the intensifying challenge on global medical reserves during the COVID pandemic. A solution algorithm is developed in the context of maritime mutual insurances (a long existent international mutual reserve system) and then tested through comprehensive numerical experiments.

Online Learning in Queueing Systems

Most queueing models have no analytic solutions, so previous research often resorts to heavy-traffic analysis for performance analysis and optimization, which requires the system scale (e.g., arrival and service rate) to grow to infinity. In "An Online Learning Approach to Dynamic Pricing and Capacity Sizing in Service Systems," Chen, Liu, and Hong develop a new "scale-free" online learning framework designed for optimizing a queueing system, called gradient-based online learning in queue (GOLiQ). GOLiQ prescribes an efficient procedure to obtain improved decisions in successive cycles using newly collected queueing data (e.g., arrival counts, waiting times, and busy times). Besides its robustness in the system scale, GOLiQ is advantageous when focusing on performance optimization in the long run because its data-driven nature enables it to constantly produce improved solutions which will eventually reach optimality. Effectiveness of GOLiQ is substantiated by theoretical regret analysis (with a logarithmic regret bound) and simulation experiments.

Adaptive Simulation Can Find Nash Equilibria in Queueing Games

The common setting of a queueing-game model consists of a stochastic stream of customers arriving at a queueing system one by one, each customer strategically chooses an action that may depend on information they receive regarding the system state. The aggregate customer decision profile gives rise to a system steady state, and, provided customers arrive at said steady state, if their decision is utility maximizing (ex ante), then this aggregate decision profile constitutes a Nash equilibrium. However, expressing the steady-state distribution for a given decision profile is very often a difficult task, and in such a case, an attempt to find a Nash equilibrium via direct analysis is futile. In "Stochastic Approximation of Symmetric Nash Equilibria in Queueing Games," Ravner and Snitkovsky suggest a novel stochastic algorithm that learns the Nash equilibrium in a class of queueing games, based on a single adaptive simulation. The method is robust and is easy to implement, offering a practical solution to queueing-game models that classical queueing-analytic methods prove inadequate.

Dynamic Pricing with Limited Price Changes

Dynamic pricing with demand learning is a very common problem that retailers face. In this problem, the retailer aims to maximize cumulative revenue collected over a finite time horizon by balancing two objectives: learning demand and maximizing revenue. In "Dynamic Pricing with Unknown Nonparametric Demand and Limited Price Changes," Perakis and Singhvi study this problem when the retailer makes no parametric assumption on the demand and seeks to reduce the amount of price experimentation because of the potential costs associated with price changes. The authors construct a pricing policy that uses second order approximations of the unknown demand function and establish when the proposed policy achieves near-optimal rate of regret while making very limited price changes. They also perform extensive numerical experiments to show that their proposed policy substantially improves over existing methods in terms of the total price changes, with comparable performance on the cumulative regret metric.

Decisions Are Often Subject to Uncertainty

In "Conditional Distributionally Robust Functionals," Shapiro and Pichler address decision making in multiple stages, where prior information is available and where consecutive and successive decisions are made. Risk measures assess the random outcome by taking various candidate probability measures into account. To justify decisions in multiple stages, it is essential to have conditional risk measures available, which respect the information, which was already revealed in the past. The authors address different variants of risk measures and discuss their properties in the specific context and their implications in multistage decision making. Various examples of risk measures on simple probability spaces with finite support illustrate the content. The Wasserstein and nested distance are involved to make decision making with numerous scenarios numerically tractalbe.

Reinforcement Learning for Merchant Energy Production

Modeling as real options the operations of energy production companies that operate in wholesale markets gives rise to a challenging Markov decision process. In "Least Squares Monte Carlo and Pathwise Optimization for Merchant Energy Production," Yang, Nadarajah, and Secomandi study the performance of two reinforcement learning techniques that can be used to determine feasible operating policies and optimality bounds for this model, namely least squares Monte Carlo and pathwise optimization, extending the applicability of the latter method beyond optimal stopping by using principal component analysis and block coordinate descent. The authors find that both approaches lead to near optimal policies, but pathwise optimization outperforms least squares Monte Carlo in terms of dual bounds at the expense of more sizable computational requirements. These findings have potential relevance for managers of energy production assets that use analytics to optimize their operations and researchers interested in broadening the scope of pathwise optimization.