



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(4):iii-vi. <https://doi.org/10.1287/opre.2024.inthisissue.v72.n4>

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 © 2024, 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.2024.inthisissue.v72.n4>

Copyright: © 2024 INFORMS

Efficient Learning Algorithms for the Best Capped Base-Stock Policy in Lost Sales Inventory Systems

Periodic review, lost sales inventory systems with lead times are notoriously challenging to optimize. Recently, the capped base-stock policy, which places orders to bring the inventory position up to the order-up-to level subject to the order cap, has demonstrated exceptional performance. In “UCB-Type Learning Algorithms with Kaplan–Meier Estimator for Lost-Sales Inventory Models with Lead Times,” Lyu, Zhang, and Xin propose an upper confidence bound–type learning framework. This framework, which incorporates simulations with the Kaplan–Meier estimator, works with censored demand observations. It can be applied to determine the optimal capped base-stock policy with a tight regret with respect to the planning horizon and the optimal base-stock policy with a regret that matches the best existing result. Both theoretical analysis and extensive numerical experiments demonstrate the effectiveness of the proposed learning framework.

Managing Shipping Emission Control Areas

The design of emission control areas (ECAs) is crucial for reducing global shipping emissions and protecting the environment. In “Shipping Emission Control Area Optimization Considering Carbon Emission Reduction,” Zhuge, Wang, and Zhen focus on the ECA optimization problem for sailing legs with ECAs. First, a case with a no-ECA policy and a case with the current ECA policy are discussed. Then, two new voyage-dependent ECA policies with sulfur limits, designated sailing paths, and speed limits are proposed, under which Stackelberg game models with the ECA regulator and a shipping company are developed. The authors extend the research problem from a sailing leg to a shipping network to improve the practicality of the findings. They also develop a dynamic programming-based algorithm to optimize the ECA policies for the shipping network from the perspective of the ECA regulator. The effectiveness of the proposed policies in reducing social costs is validated by numerical experiments.

A Lyapunov Theory for Finite-Sample Guarantees of Markovian Stochastic Approximation

The stochastic approximation (SA) method stands as the foundational mathematical tool for modern large-scale optimization and machine learning. Therefore, gaining a theoretical understanding of SA algorithms is of fundamental interest. In “A Lyapunov Theory for Finite-Sample Guarantees of Markovian Stochastic Approximation,” Chen, Maguluri, Shakkottai, and Shanmugam present a unified Lyapunov framework for the finite-sample analysis of a Markovian SA algorithm under a contractive operator with respect to an arbitrary norm. The key novelty lies in the construction of a smooth Lyapunov function called the generalized Moreau envelope. The authors demonstrate the effectiveness of their SA results in the context of reinforcement learning (RL), specifically through popular algorithms such as variants of temporal difference (TD) learning and Q-learning. As byproducts, the

results provide theoretical insights into the efficiency of bootstrapping in TD learning with eligibility traces and the bias-variance tradeoff in off-policy learning.

Anytime Valid Comparison of Sequential Forecasters

How do we compare forecasters that each make a probabilistic prediction on a sequence of events (e.g., weather and sports)? In “Comparing Sequential Forecasters,” Choe and Ramdas propose flexible approaches to sequential inference on the mean score difference between any two forecasters. To estimate this time-varying quantity, the authors propose a sequential analog of confidence intervals, called confidence sequences (CSs). These CSs correctly cover the score difference under continuous monitoring, and the evaluator can freely peek at the scores to stop the experiment (“anytime valid”). The authors further develop a complementary anytime valid approach called e-processes, which quantify the evidence against the claim that one forecaster is never better than the other in mean scores. The validity of these methods does not depend on stationarity or other assumptions on how the score differences evolve sequentially. In their paper, the authors showcase CSs and e-processes for comparing real-world baseball and weather forecasters.

Characterizing the Efficiency of Static Pricing Schemes as a Function of the Supply

The problem of selling a supply of k units to a stream of customers constitutes one of the cornerstones in revenue management. Static pricing schemes (that output the same price to all customers) are commonly used because of their simplicity and their many desirable properties; they are anonymous, nonadaptive, and order oblivious. Although the efficiency of those schemes should improve as the supply k increases, prior work has only focused either on algorithms that aim for a constant approximation that is independent of k or on the setting where k becomes really large. In contrast, in “Static Pricing for Multi-unit Prophet Inequalities,” Chawla, Devanur, and Lykouris characterize the efficiency of static pricing schemes as a function of the supply. Our approach stems from identifying a “sweet spot” between selling enough items and obtaining enough utility from customers with high valuations. Subsequent work shows that our pricing scheme is the optimal static pricing for every value of k .

Treatment Planning for Mass Casualty Incidents

The existing emergency response guidelines emphasize prioritizing the treatment of victims with initially critical health conditions, but tend to overlook the potential deterioration of less critical victims. Such deterioration can result in prolonged treatment times and irreversible health damages. In “Treatment Planning for Victims with Heterogeneous Time Sensitivities in Mass Casualty Incidents,” Shi, Liu, and Wan draw insights from a unique data set containing timestamps of surgeries conducted in a field hospital established in response to a large-scale earthquake.

They develop scheduling models to enhance treatment planning for mass casualty incidents. They identify conditions under which victims with less critical initial conditions may have higher or lower priority than their counterparts in an optimal schedule, aiming to do the greatest good for the greatest number. Through a counterfactual analysis utilizing their data set, the authors demonstrate that implementing their model could significantly reduce surgical makespan, overdue cases, and victim deterioration compared with the previously employed treatment plan.

The Rationality of Risk-averse Patient Selection by Transplant Program

U.S. health agencies periodically evaluate transplant programs based on their patients' posttransplant survival outcomes, and a program is flagged for review if the number of unsuccessful transplants far exceeds what would be expected based on national averages. Some researchers have expressed concerns that these regulations might cause programs to reject high-risk patients, whereas others have questioned if such a response would be rational. In "Characterizing Rational Transplant Program Response to Outcome-Based Regulation," Mildebrath, Lee, Sinha, Schaefer, and Gaber use chance-constrained optimization to demonstrate that it may, in fact, be rational for a transplant program to become more selective when evaluating transplant candidates for admittance to the waitlist. They also demonstrate that the regulations may unfairly penalize medium-sized programs. Moreover, their model quantifies which patients may be most at risk for adverse selection by programs. Their results provide insights to policymakers by quantitatively characterizing the response of rational programs to outcome-based regulations.

Trade-Offs in Dynamic Allocation Problems

Food rescue organizations often receive donations to allocate to food pantries or families. Donations are unpredictable, and goods are often perishable; as a result, allocations have to be made within a short time frame after arrival without knowledge of future arrivals. It is important that donations go to organizations that are able to use them; at the same time, organizations that serve different communities should be treated equitably. In "Fair and Efficient Online Allocations," Benadè, Kazachkov, Procaccia, Psomas, and Zeng study fairness-efficiency trade-offs in such online allocation problems. Against adversarial arrivals, no algorithm can provide nontrivial guarantees for both these objectives simultaneously. When item values are drawn from (potentially correlated) distributions, there is no trade-off, and a simultaneously fair and efficient algorithm is presented.

Augmenting the Modeling Power of the Multinomial Logit Model

Using choice models to capture customer choice behavior has steadily become the common practice in revenue management. Discrete choice models allow one to capture the fact that customers substitute among the offered products, so if a particular product is not offered, then a portion of the customers interested in this product will substitute into a suitable available alternative. The multinomial logit model is one of the most commonly used choice models to capture customer choice in practice. This choice model is based on the utility maximization framework. A customer associates random utilities with the products, as well as the no-purchase option, in which case, the customer purchases the product with the largest utility, as long as its utility exceeds that of the no-purchase option. In "Assortment Optimization Under the Multinomial Logit Model with Utility-Based Rank Cutoffs," Bai, Feldman, Topaloglu, and Wagner augment the modeling flexibility of the multinomial logit model, where a customer also leaves without a purchase if she

cannot find one of her top few choices. They develop algorithms to find revenue-maximizing assortments.

Working with U.S. Policymakers to Redesign National Organ Allocation

The Organ Procurement & Transplantation Network (OPTN), which manages transplantation activities in the United States, recently partnered with the MIT Operations Research Center to design and implement novel organ allocation policies that are more equitable, efficient, and inclusive. National organ allocation policies need to strike a delicate balance between efficiency and fairness in multiple objectives, reconciling often disparate value judgments and priorities from many different stakeholders. In "Reshaping National Organ Allocation Policy," Papalexopoulos, Alcorn, Bertsimas, Goff, Stewart, and Trichakis introduced a novel optimization- and machine learning-based framework to aid policy design and navigate challenging fairness-efficiency tradeoffs. The authors collaborated with the OPTN to apply the framework to the design of a new national allocation policy for lungs, which was implemented in March 2023 and is anticipated to reduce waitlist mortality by approximately 20%. Based on this success, the authors are now working toward the redesign of the entire U.S. organ allocation system, including kidneys, pancreata, hearts, and livers.

A Dual Perspective to the Robust Screening Problem

Robust screening problem is concerned with the problem of a seller seeking a selling mechanism that maximizes the worst-case revenue obtained from a buyer whose valuation distribution lies in a certain ambiguity set. In "Screening with Limited Information: A Dual Perspective," Chen, Hu, and Wang show that strong duality holds between the problem of finding the optimal robust mechanism and a minimax pricing problem, where the adversary first chooses a worst-case distribution and then the seller decides the best posted price mechanism. The duality result connects prior literature that separately studies the primal (robust screening) and problems related to the dual and offers a unified geometric intuition in solving the problem.

Harvesting Solar Power Fosters Prices in a Vicious Cycle: Breaking the Cycle with Price Mechanisms

In "Harvesting Solar Power Fosters Prices in a Vicious Cycle: Breaking the Cycle with Price Mechanisms," Mamaghani and Çakanyildirim quantify and explore how residential solar power adoption affects the price of electricity and the profit of a utility firm. It reveals the mechanism behind what is known as the "utility (death) spiral." As more residences adopt solar power generation, the price of electricity goes up, whereas the profit of utility firm comes down. The study finds the equilibrium of a game played between consumers deciding to adopt residential solar power and the utility deciding on prices in a regulated market. The equilibrium, involving a potential bankruptcy for the utility firm and a high price for a nonadopter consumer, is desirable for neither. To avoid this, it suggests either decreasing the buy-back price that the utility pays for the excess power of an adopter or introducing a subscription for adopters. These suggestions can inform pricing policies in the power market.

New Flexible Advertising Model Unveiled: Enhancing Advertising Strategies

Researchers have generalized the classical Sethi advertising model (1983) to provide a flexible yet tractable approach to optimizing advertising strategies. Unlike traditional methods, the generalized

Sethi advertising model is designed to adapt to different market conditions in various media and markets, offering advertisers greater flexibility. By assuming a constant-returns-to-scale relationship between advertising expenditure and untapped market potential, the generalized Sethi advertising model provides advertisers with more options without sacrificing tractability. In “Technical Note—The Generalized Sethi Advertising Model,” Kennedy, Sethi, Siu, and Yam have made important breakthroughs by developing practical formulas for determining the best advertising strategies for both individual companies and groups of companies. They have also conducted analyses to better understand how different factors influence these strategies. The generalized Sethi advertising model opens up new opportunities for advertisers to fine-tune their campaigns and reach more customers. This research marks a significant step forward in advertising optimization, giving businesses a valuable tool to achieve better results and make informed decisions in their advertising efforts.

Task-Server Constraints Hardly Hurt Performance

Modern service systems, like cloud computing platforms or data center environments, commonly face a high degree of heterogeneity. This heterogeneity is not only caused by different server speeds but also, by binding task-server relations that must be taken into account when assigning incoming tasks. Unfortunately, there are hardly any theoretical performance guarantees as these systems do not fall within the typical supermarket modeling framework which heavily relies on strong symmetry and homogeneity assumptions. In “Heavy-Traffic Universality of Redundancy Systems with Assignment Constraints,” Cardinaels, Borst, and van Leeuwen provide insight in the performance of these systems operating under redundancy scheduling policies. Surprisingly, when experiencing high demand, these systems exhibit state space collapse and can achieve a similar level of resource pooling and performance as a fully flexible system, even subject to quite strict task-server constraints.

Can Inefficiency be Rational?

Excess resources or slack may serve as a buffer against environmental shocks, help decouple organizations, ease planning and implementation, support innovation, and enable effective responses to competitors. Slack may however also be the result of inefficiency. In “Distinguishing Useful and Wasteful Slack,” Bogetoft and Kerstens propose an approach to separate useful and wasteful slack. If an organization can maintain the same levels of output and slack at lower cost, there is wasteful or nonrationalizable spending. We develop ways to measure the extent to which total spending can be rationalized and show how to statistically estimate and test the usefulness of the available slack using bootstrapping.

Two-Stage Matching and Pricing in Ride-Hailing Platforms

Matching and pricing are two critical levers in two-sided marketplaces to connect demand and supply. The platform can produce more efficient matching and pricing decisions by batching the demand requests. In “Two-Stage Stochastic Matching and Pricing with Applications to Ride Hailing,” Feng, Niazadeh, and Saberi initiate the study of the two-stage stochastic matching problem with or without pricing to enable the platform to make improved decisions in a batch with an eye toward the imminent future demand requests. This problem is motivated in part by applications in online marketplaces, such as ride-hailing platforms. We design online competitive algorithms for vertex-weighted (or unweighted) two-stage stochastic matching for maximizing supply efficiency and two-stage joint matching and pricing for

maximizing market efficiency. Using various techniques, such as introducing convex programming-based matching and graph decompositions, submodular maximization, and factor-revealing linear programs, we obtain either optimal competitive or improved approximation algorithms compared with naïve solutions. We enrich our theoretical study by data-driven numerical simulations using DiDi’s ride-sharing data sets.

Robust Assortment Optimization Under the Markov Chain Choice Model

Assortment optimization arises widely in many practical applications. In this problem, the goal is to select products to offer customers in order to maximize the expected revenue. In “Robust Assortment Optimization Under the Markov Chain Choice Model,” Désir, Goyal, Jiang, Xie, and Zhang study a robust assortment-optimization problem under the Markov chain choice model, in which the parameters of the choice model are assumed to be uncertain, and the goal is to maximize the worst case expected revenue over all parameter values in an uncertainty set. Our main contribution is to prove a min-max duality result when the uncertainty set is row-wise. The result is surprising as the objective function does not satisfy the properties usually needed for known min-max results. Inspired by the duality result, we develop an efficient iterative algorithm for computing the optimal robust assortment under the Markov chain choice model. Moreover, our results yield operational insights into the effect of changing the uncertainty set on the optimal robust assortment.

A Comprehensive Set of Asymptotic Properties for a Meaningful Aggregation of Malmquist Indices

The Malmquist productivity index (MPI) has become one of the most widely used tools for analyzing dynamic performance of decision-making units. Whereas accounting for economic weights of individual units in aggregations of indices is emphasized in the literature, statistical theory for constructing confidence intervals and performing hypothesis tests based on weighted aggregation of the MPI are still unavailable. In “Statistical Inference for Aggregation of Malmquist Productivity Indices,” Pham, Simar, and Zelenyuk use a novel approach (based on the uniform delta method) to develop new asymptotic theory (including new central limit theorems) for aggregate MPIs as the basis for the statistical inference and test. They also verify the finite-sample performance of their approach via extensive Monte Carlo experiments and provide an illustration using real-world data.

Dial M for Simulation

For years, systems of stochastic differential equations (SDEs) were simulated by discretization, inevitably introducing a bias, which can be difficult to quantify accurately. To circumvent this, some attempts have been made to simulate exactly various models from the SDE solution. These approaches prove capable of producing accurate results. A serious drawback of such an approach, nevertheless, is the implicit need to use extensive numerical methods, which make the entire simulation computationally heavy and quite impracticable. In “Unified Moment-Based Modeling of Integrated Stochastic Processes,” Kyriakou, Brignone, and Fusai present a methodological framework based on M(oments) for the simulation of such models that overcomes earlier limitations. Theoretical results and extensive numerical experiments show that the proposed approach allows accurate simulation of complex stochastic models with low computational effort.

New Exact Approaches Tailored for Kidney Exchange Programs with Hierarchical Objectives

Kidney exchange programs increase the rate of living donor kidney transplantation. Whereas effective integer programming models aimed at maximizing the total number of transplants have been proposed in the literature, these cannot always be extended to handle a hierarchy of objectives, which is often a requirement in practice. In “New Algorithms for Hierarchical Optimization in Kidney Exchange Programs,” Delorme, García, Gondzio, Kalcsics, Manlove, and Pettersson introduce a new integer programming framework to solve large-size instances of kidney exchange programs. The authors use an ad hoc preprocessing and a reduced-cost variable fixing algorithm to dramatically decrease the size of the models, and they devise a diving algorithm that exploits the hierarchical structure of the problem to significantly reduce the number of integer programs that need to be solved. They also show that it is possible to transition between models as different layers are traversed in the hierarchy, allowing each layer to be solved with the most effective model. Experiments on three different European kidney exchange programs show that running times can be reduced by up to three orders of magnitude.

Ex-Ante and Ex-Post Fairness in Resource Allocation

Consider the problem of allocating indivisible goods among agents with additive valuations, where monetary payments are not allowed. When randomization is allowed, it is possible to achieve compelling notions of fairness such as EV, which states that no agent should prefer any other agent’s allocation to their own. When allocations must be deterministic, achieving exact fairness is impossible but approximate notions such as EV up to one good can be guaranteed. In “Best of Both Worlds: Ex Ante and Ex Post Fairness in Resource Allocation,” Aziz, Freeman, Shah, and Vaish ask whether it is possible to achieve both types of guarantees simultaneously. More specifically, they ask whether there exists a probability distribution over deterministic allocations such that every deterministic allocation is envy-free up to one good and the distribution is exactly envy-free in expectation. The main result of the paper answers this question in the affirmative, showing that ex ante and ex post fairness need not be in conflict.

Quantify the Uncertainty to Decide and Explore Better

In statistical inference, large-sample behavior and confidence interval construction are fundamental in assessing the error and reliability of estimated quantities with respect to the data noises. In “Uncertainty Quantification and Exploration for Reinforcement Learning,” Zhu, Dong, and Lam study the large sample behavior in the classic setting of reinforcement learning. They derive appropriate large-sample asymptotic distributions for the state-action value function (Q-value) and optimal value function estimations when data are collected from the underlying Markov chain. This allows one to evaluate the assertiveness of performances among different decisions. The tight uncertainty quantification also facilitates the development of a pure exploration policy by maximizing the worst-case relative discrepancy among the estimated Q-values (ratio of the mean squared difference to the variance).

This exploration policy aims to collect informative training data to maximize the probability of learning the optimal reward collecting policy, and it achieves good empirical performance.

Computing Child Maintenance for Networks of Divorced Parents

When Dutch parents with children divorce, a mediator compiles a matrix with the financial needs of the children and the financial capacities of parents to meet these needs. Moreover, in case parents have children from previous marriages or are prepared to contribute to stepchildren, a bipartite graph shows which parent is financially responsible for which child. The Dutch high court ruled that the final contributions should be proportionally consistent, implying that shortages for children should be prevented if possible and any remaining parental capacities should be proportionally divided among parents responsible for the same child. Finding by hand this proportional solution is difficult for realistic court cases, as these can include several (step)parents and children. In “On Proportionally Consistent Solutions to the Divorced-Parents Problem,” Romeijnders, Van Foreest, and Wijngaard show that the final unique solution can be found when parents start court cases iteratively and provides efficient algorithms that can deal with large (country-size) networks.

Minimizing Regret in Multiperiod Decision-Making Problems

Regret minimization has gained popularity in a wide range of decision-making problems under uncertainty because of its capacity to identify more opportunistic solutions than worst-case value optimization. Unfortunately, the rigidity of current worst-case regret models and scarcity of tractable solution methods have been serious obstacles in multistage applications. In “Technical Note—Risk-Averse Regret Minimization in Multistage Stochastic Programs,” Poursoltani, Delage, and Georghiou consider a multistage stochastic programming setting with a discrete scenario tree. They introduce the notion of the Δ -regret model, which bridges between the ex ante and ex post regret minimization paradigms that are currently used in the regret minimization literature for single-stage problems. The notion of Δ -regret minimization is investigated for the first time both theoretically and numerically in order to better understand its behavior under a set of popular risk measures.

Omnichannel Service Operations: A Queueing Approach

Both walk-in and online customers appear in many service systems such as restaurants. Online customers are often given a target time for pick up, and the quality of the food/beverage can degrade if it is ready before the arrival of online customers. This distinctive feature brings essential difficulties in the analysis and control of such systems. In “Technical Note—Asymptotically Optimal Control of Omnichannel Service Systems with Pick-up Guarantees,” Gao, Huang, and Zhang study the optimal control problem of such service systems based on a two-class single server queueing model. The paper proposes a nearly-optimal policy that keeps the server idle when the queue of online customers drops below a threshold and there are no walk-in customers.