



Operations Research

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

In This Issue

To cite this article:

(2023) In This Issue. Operations Research 71(4):iii-vi. <https://doi.org/10.1287/opre.2023.inthisissue.v71.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 © 2023, 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.2023.inthisissue.v71.n4>

Copyright: © 2023 INFORMS

Approximate Submodularity in Network Design Problems

In “Approximate Submodularity in Network Design Problems” DeValve, Pekeć, and Wei study general network design problems where a firm strategically selects its network to better match supply and demand in the future. The authors observe that the arcs in network design problems are cover modular, i.e., approximately substitutes with each other, in the sense that local changes in the objective function can be used to bound global changes. The cover modularity is then applied to prove that a set of simple and intuitive heuristics achieve constant factor approximation guarantees in network design problems. Furthermore, the paper demonstrates cover modularity is also present in a general class of linear programming formulations.

Foundations for Learning Co-Operative Games in The Large Population Regime

Multiagent systems—such as recommendation systems, ride-sharing platforms, food-delivery systems, and data-routing centers—are areas of rapid technology development that require constant improvements to address the lack of efficiency and curse of dimensionality. In “Dynamic Programming Principles for Mean-Field Controls with Learning,” Gu, Guo, Wei, and Xu show that multiagent systems with mean-field approximation and learning can be recast as general forms of reinforcement learning problems, where the state variable is replaced by the probability distribution. This reformulation paves the way for developing efficient value-based and policy-based algorithms for mean-field controls with learning. It is also the first step toward future theoretical development of learning problems with mean-field controls.

Integrated Production Planning and Risk Hedging

Effective management of demand risk is a central problem of production planning. In many applications, it is observed that demand often depends on certain asset prices in the financial market. This calls for a model that jointly optimizes the production decision along with a risk hedging policy, using the latter to compensate for loss of profit in scenarios of weak demand. In “Production Planning with Risk Hedging Under a Conditional Value at Risk Objective,” Wang and Yao develop an integrated production and risk hedging model, where the risk hedging strategy is derived in explicit form and the production quantity is solved as a concave maximization problem. Furthermore, a complete characterization of the efficient frontier is provided, and the improvement over the production-only decision is quantified.

How Product Locations Drive Traffic Throughout a Retail Store

In “Store-Wide Shelf-Space Allocation with Ripple Effects Driving Traffic,” Flamand, Ghoniem, and Maddah develop a framework for deciding where to place products in a store, in addition to apportioning the shelf space among products, in a way that maximizes impulse profit, a phenomenon that may account for 50% of

transactions. By analyzing a large data set of customer receipts from a grocery store in Beirut, the authors develop a regression model that estimates traffic at a shelf based on its location and the “attraction” from products allocated nearby. The traffic model is embedded within a mixed-integer nonlinear program, which they solve via specialized linear approximations. For the store in Beirut, a 65% improvement in impulse profit is anticipated, and the location of products is found to be significantly more important in driving store-wide traffic than the relative shelf-space allocation.

General Equilibrium in a Heterogeneous-Agent Incomplete-Market Economy with Many Consumption Goods and a Risk-Free Bond

General equilibrium models have been extensively studied in economics, with seminal works by Arrow, Debreu, McKenzie, Gale, and Nikaido providing proof for the existence of general equilibrium in a pure exchange economy. However, while the recently popular dynamic incomplete-markets models have attracted attention, the existence results for these models have remained limited. In “General Equilibrium in a Heterogeneous-Agent Incomplete-Market Economy with Many Consumption Goods and a Risk-Free Bond,” Light extends some general equilibrium results to a dynamic incomplete-market economy with many consumption goods and a risk-free bond. Under mild conditions, using a well-known excess demand approach, the article establishes the existence of general equilibrium, showing that the methods used to prove existence in the static pure exchange economy can be used in the much more complicated dynamic incomplete-market setting. The article also establishes uniqueness and comparative statics results for the special case where the agents’ preferences can be represented by a constant elasticity of substitution utility function with an elasticity of substitution greater than or equal to one.

Technical Note—Constrained Assortment Optimization Under the Nested Logit Model

Assortment optimization involves determining the optimal set of products to show customers and is a fundamental problem in retail operations. The nested logit choice model is a popular and widely used choice model to capture customer behavior. In “Technical Note—New Bounds for Cardinality Constrained Assortment Optimization Under the Nested Logit Model,” Kunnumkal presents a new method for making the assortment decisions under the nested logit choice model when there is a constraint on the number of products that can be offered within each nest. Computational experiments reveal that the assortments obtained by the solution method are near optimal, with the average optimality gap being under one percent.

Contextual Search in the Presence of Adversarial Corruptions

Contextual search is a generalization of binary search, which captures settings such as feature-based dynamic pricing. In this

paradigm, a decision maker repeatedly interacts with a set of agents; in the pricing example, the decision maker first observes the context/ features of an item, then sets a feature-dependent price, and the agent responds by either buying or not. Standard formulations of this problem assume that agents act in accordance with a specific homogeneous response model. In practice, however, some responses may be adversarially corrupted (e.g., when some agents exhibit outlier behavior). Existing algorithms heavily depend on the assumed response model being (approximately) accurate for all agents and have poor performance in the presence of even a few such outliers. In “Contextual Search in the Presence of Adversarial Corruptions,” Krishnamurthy, Lykouris, Podimata, and Schapire initiate the study of contextual search in the presence of such adversarial corruptions. Their algorithms’ performance is near-optimal in the absence of outliers and degrades gracefully with their number.

Optimal Policies for Online Platforms When Social Learning Occurs

Before buying products online, consumers read the reviews written by the previous customers. If they buy the product, they write a review themselves. When the product is of unknown quality, consumers learn it over time; that is, social learning occurs. If consumers have various purchase options of similar products of different brands, the platform that they use may affect this social learning by choosing the order in which the products appear on its website. In “Product Ranking in the Presence of Social Learning,” Maglaras, Scarsini, Shin, and Vaccari compare various policies that the platform may adopt, with the goal of maximizing its revenue collected from commission fees for sold items. The criterion to compare the policies is the worst-case regret with respect to a fully informed platform benchmark.

Product Ranking via Sequential Submodular Maximization

The advents of online retailing and advertising have created new opportunities for online platforms to incorporate algorithmic techniques to improve shoppers’ experience and drive user engagement, which, in return, can help with the long-term growth of these platforms, and also to help with having socially aware operations that consider fairness across different demographic groups. Motivated by the product-ranking problem in online shopping, in “Sequential Submodular Maximization and Applications to Ranking an Assortment of Products” Asadpour, Niazadeh, Saberi, and Shameli introduce and study a new class of combinatorial optimization problems over the space of permutations, which is referred to as “sequential submodular maximization.” Using this class of problems, it provides algorithmic solutions for maximizing users’ engagement and also for balancing the users’ engagement across different demographic groups of shoppers to obtain fairness. In particular, the authors propose an optimal $(1 - 1/e)$ -approximation algorithms for maximizing users’ engagement and a bicriteria $(1 - 1/e^2, 1 - 1/e^2)$ for maximizing users’ engagement subject to group fairness constraints.

Analytics in the Face of Fraudulent Data

In “Learning Product Rankings Robust to Fake Users,” Golrezaei, Manshadi, Schneider, and Sekar present a novel online learning algorithm for identifying optimal product rankings in the presence of fake users and corrupted data. In recent years, e-commerce platforms, such as Amazon, have witnessed a growing number of fake users and click farms. These fraudulent actors seek to boost the position of certain products in the display ordering (i.e., product ranking). Further, platforms’ reliance on data analytics exacerbates the effect of these fake users as machine learning algorithms

leverage user feedback to determine product rankings. In the face of these challenges, the present article departs from the status quo that is based on detecting fake users and instead proposes a robust learning methodology. More specifically, the article presents a robust online learning algorithm that converges to the optimal product ranking even when it is impossible to distinguish between real and fake users in the data.

Assortment Personalization in E-commerce

In the age of big data, online businesses have access to a tremendous amount of browsing and purchasing data from their customers. Consequently, there has been great effort in exploiting this data to offer personalized product assortments to each customer based on what is known about her or his preferences. A personalized assortment can potentially allow enhancing the customer experience as well as improving the revenue of the retailer. In “Joint Assortment Optimization and Customization Under a Mixture of Multinomial Logit Models: On the Value of Personalized Assortments,” El Housni and Topaloglu consider an optimization model that jointly optimizes the assortment of products carried by an online retailer as well as the personalized assortments offered to each customer. They study the value of personalization and develop algorithms with provable theoretical guarantees to solve the optimization problem.

The Stability of MNL-Based Demand Under Dynamic Customer Substitution and Its Algorithmic Implications

In “The Stability of MNL-Based Demand Under Dynamic Customer Substitution and Its Algorithmic Implications,” Aouad and Segev study the dynamic assortment planning problem under the widely utilized multinomial logit choice model (MNL). In this single-period assortment optimization and inventory management problem, the retailer jointly decides on an assortment, that is, a subset of products to be offered, as well as on the inventory levels of these products, aiming to maximize the expected revenue subject to a capacity constraint on the total number of units stocked. The demand process is formed by a stochastic stream of arriving customers, who dynamically substitute between products according to the MNL model. Although this dynamic setting is extensively studied, the best known approximation algorithm guarantees an expected revenue of at least 0.139 times the optimum, assuming that the demand distribution has an increasing failure rate. The authors establish novel stochastic inequalities showing that, for any given inventory level, the expected demand of each offered product is “stable” under basic algorithmic operations, such as scaling the MNL preference weights and shifting inventory across comparable products. They exploit this sensitivity analysis to devise the first approximation scheme for dynamic assortment planning under the MNL model.

Measuring the “True” Quality by Rival Bidder Qualities in Score Auctions

In a procurement auction, where price and quality are evaluated by a score, the assessment criteria partially account for the “true” quality offered. The joint information of all the bids can increase the reliability of the measurement process of quality, giving better accuracy of estimation compared with the reported quality. In “Technical Note-Bidding in Multidimensional Auctions When the Qualities of All Bidders Matter,” Lorentziadis examines auctions where the score is adjusted by the qualities of all bidders, who exhibit different production efficiencies and additive separable costs. At equilibrium, the total procurement cost remains the same in the first-score and the second-score auctions. For different adjustments of the score—for example, using a weighted average

of all the qualities—the author finds that the standard unadjusted score auction brings the highest total project cost to the buyer. Procurement managers should give consideration to score rules that account for all bidder qualities.

A Novel and Promising Approximation for Network Revenue Management

In “Product-Based Approximate Linear Programs for Network Revenue Management,” Zhang, Samiedaluae, and Zhang propose a novel separable piecewise linear (SPL) approximation for the network revenue management problem. The coefficients of the proposed SPL approximation can be interpreted as each product’s revenue contribution to the value of each resource in a given period, which provides more granular information compared with the existing resource-based SPL approximation in the literature. The new approximation provides more flexibility for policy construction. Furthermore, the new approximation opens the opportunity to derive a set of valid inequalities to further improve the computational performance and achieve additional gains in the expected revenue. Computational experiments with instances of various network structures and parameters demonstrate its efficacy: The new approximation leads to bid-price policies generating higher expected revenues and demonstrates better performance in terms of both computational efficiency and numerical stability.

Simple Scale-Shift Peer Grading Scheme to Improve Accuracy and Learning Outcomes in MOOCs

A critical issue in operating massive open online courses (MOOCs) is the scalability of providing feedback. Because it is not feasible for instructors to grade a large number of students’ assignments, MOOCs use peer grading systems. In “Economic Behavior of Information Acquisition: Impact on Peer Grading in Massive Open Online Courses,” Yoo and Zhan investigate the efficacy of that practice when student graders are considered rational economic agents. Using an economic model that characterizes the behavior of student graders, they analyse the accuracy of current peer grading scheme. Interestingly, the authors identify a systematic grading bias toward the mean, which discourages students from learning. To improve current practice, they propose a simple scale-shift grading scheme, which can simultaneously improve grading accuracy and adjust grading bias. They discuss how it can be readily implemented in practice with moderate involvement of the instructors and MOOCs.

Packing Three-Dimensional Irregular Objects

Because of its many applications in practice, the cutting and packing literature is extensive and well established. It is mostly concerned with problems in one and two dimensions or with problems where some regularity of the pieces is assumed (e.g., packing boxes). However, the rise of applications in the realm of three-dimensional printing and additive manufacturing has created a demand for efficient packing of three-dimensional irregular objects. In “Voxel-Based Solution Approaches to the Three-Dimensional Irregular Packing Problem,” Lamas-Fernandez, Bennell, and Martinez-Sykora propose a series of tools to tackle this problem using voxels. These include geometric tools, a mathematical model, local search neighborhoods, and details on implementation of metaheuristic algorithms. These tools are tested extensively, and computational results provided show their effectiveness compared with state-of-the-art literature.

Accelerated Algorithms for Ranking

Assigning ranking scores to items based on observed comparison data (e.g., paired comparisons, choice, and full ranking outcomes) has been of continued interest in a wide range of applications,

including information search, aggregation of social opinions, electronic commerce, online gaming platforms, and more recently, evaluation of machine learning algorithms. The key problem is to compute ranking scores, which are of interest for quantifying the strength of skills, relevancies, or preferences, and prediction of ranking outcomes. One of the most popular statistical models of ranking outcomes is the Bradley–Terry model for paired comparisons and its extensions to choice and full ranking outcomes. In “Accelerated MM Algorithms for Inference of Ranking Scores from Comparison Data,” Vojnović, Yun, and Zhou show that a popular MM algorithm for inference of ranking scores for generalized Bradley–Terry ranking models suffers a slow convergence issue, and they propose a new accelerated algorithm that resolves this shortcoming and can yield substantial convergence speedups.

A New Approach for Optimization-Based Scenario Reduction for Two-Stage Stochastic Optimization

In the field of data-driven optimization under uncertainty, scenario reduction is a commonly used technique for computing a smaller number of scenarios to improve computational tractability and interpretability. However traditional approaches do not consider the decision quality when computing these scenarios. In “Optimization-Based Scenario Reduction for Data-Driven Two-Stage Stochastic Optimization,” Bertsimas and Mundru present a novel optimization-based method that explicitly considers the objective and problem structure for reducing the number of scenarios needed for solving two-stage stochastic optimization problems. This new proposed method is generally applicable and has significantly better performance when the number of reduced scenarios is 1%–2% of the full sample size compared with other state-of-the-art optimization and randomization methods, which suggests this improves both tractability and interpretability.

Robust Dynamic Pricing with Demand Learning in the Presence of Outlier Customers

Dynamic pricing is a core problem in revenue management. Most existing literature assumes that the demand follows a probabilistic model, with an unknown demand curve as the mean. However, in practice, customers may not always behave according to such a model. In “Robust Dynamic Pricing with Demand Learning in the Presence of Outlier Customers,” Chen and Wang study the dynamic pricing problem under model misspecification. To characterize the behavior of outlier customers, an ϵ -contamination model—the most fundamental model in robust statistics and machine learning, is adopted. The challenges brought by the presence of outlier customers are mainly due to the fact that arrivals of outliers and their exhibited demand behaviors are completely arbitrary. To address these challenges, the authors propose robust dynamic pricing policies that can handle any outlier arrival and demand patterns. The proposed policies are fully adaptive without requiring prior knowledge of the outlier proportion parameter.

How to Optimally Schedule Customers with Different Resource Requirements?

In many service systems, customers from different classes may have very different resource requirements. These differences include not only the duration of the service but also the units of resource required. This is especially prominent in healthcare settings, where patients with different severity levels can require a different level of medical attention/supervision translating into varying demands of nurse staffing capacity. Motivated by these applications, in “Managing Queues with Different Resource Requirements,” Zychlinski, Chan, and Dong study the optimal scheduling policy of a multiserver queue in which customers from different classes may require different units of servers.

When customers have heterogeneous resource requirements, in addition to the cost of waiting and resource requirement, we also need to take policy-induced idleness into account. A good policy needs to carefully balance the myopic cost reduction rate and the more forward-looking system throughput. An index-based policy, referred to as the *idle-avoid $c\mu/m$ rule*, is developed to balance these factors. Theoretical analysis along with numerical experiments provide support for good and robust performance of the proposed policy.

Technical Note—An Approximate Dynamic Programming Approach to the Incremental Knapsack Problem

Integer packing problems have traditionally been some of the most fundamental and well-studied computational questions in discrete optimization. In “Technical Note—An Approximate Dynamic Programming Approach to the Incremental Knapsack Problem,” Aouad and Segev study the incremental knapsack problem, where one wishes to sequentially pack items into a knapsack whose capacity expands over a finite planning horizon, with the objective of maximizing time-averaged profits. Although various approximation algorithms were developed under mitigating structural assumptions, obtaining nontrivial performance guarantees

for this problem in its utmost generality has remained an open question thus far. The authors devise the first polynomial-time approximation scheme for general instances of the incremental knapsack problem, which is the strongest guarantee possible given existing hardness results. Their approach synthesizes various techniques related to approximate dynamic programming, including problem decompositions, counting arguments, and efficient rounding methods, which may be of broader interest.

A General Elliptical Potential Lemma

In sequential learning and decision-making problems, the elliptical potential lemma is a key technique to quantify the decrease in the uncertainty of the model as more observations are obtained. However, it requires the observation noise and prior distribution of the unknown parameters to be Gaussian. In “Technical Note—The Elliptical Potential Lemma for General Distributions with an Application to Linear Thompson Sampling,” Hamidi and Bayati introduce a general version of the elliptical potential lemma that relaxes the Gaussian assumption. They also apply their general lemma to prove a minimax optimal Bayesian regret bound for the well-known Thompson sampling algorithm in stochastic linear bandits with changing action sets where prior and noise distributions are general.