



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(2):iii-vi. <https://doi.org/10.1287/opre.2023.2440>

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.2023.2440>

Copyright: © 2023 INFORMS

Risk-Based Robust Statistical Learning

For the treatment of outliers, in “Risk-Based Robust Statistical Learning by Stochastic Difference-of-Convex Value-Function Optimization,” Liu and Pang propose a risk-based robust statistical learning model. Employing a variant of the conditional value-at-risk risk measure, called the interval conditional value-at-risk (In-CVaR), the model aims to exclude the risks associated with the left and right tails of the loss. The resulting nonsmooth and nonconvex model considers the population In-CVaR risk and distinguishes the upside and downside losses with asymmetric weights. For the solution of the model in both regression and classification, the authors show that the objective function is the difference of two convex functions each being the optimal objective value of a univariate convex stochastic program. A sampling and convex programming-based algorithm is developed with the appropriate control of incremental sample sizes, and its subsequential almost-sure convergence to a critical point is established. Numerical results illustrate the practical performance of the model and methodology.

Resource Allocation in Large-Scale Stochastic Networks

Large-scale stochastic networks are widely used for a variety of systems including telecommunications, patient flows, service and data centers, and so on. Because of their complexity, ensuring the stability of these networks by allocating the required resources needed by each customer class is quite challenging. When there are sufficient resources to serve each customer class, the existence of a policy that stabilizes the system is trivial. One can decouple the network and assign the required resources to each customer class. Conversely, one can anticipate that, when each class is resourceless, it is impossible to stabilize the system independent of the policy used. However, previous analyses did not tackle this question. In “On System-Wide Safety Staffing of Large-Scale Parallel Server Networks,” Hmedi, Arapostathis, and Pang assume that some classes have excess resources, whereas others are deficient. They provide a full characterization of the stability of these networks by identifying a parameter that describes the overall excess or lack of resources in the network.

Generalized Exact Scheduling: A Minimal-Variance Distributed Deadline Scheduler

In “Generalized Exact Scheduling: A Minimal-Variance Distributed Deadline Scheduler,” Nakahira, Ferragut, and Wierman adapt tools from optimization and control theory to characterize the optimal distributed policies in a broad range of settings without any approximation. They show that *Exact Scheduling* minimizes both the stationary mean and variance of the service capacity subject to strict demand and deadline requirements. For more general settings, the authors characterize the minimal-variance distributed policies with soft demand requirements, soft deadline requirements, or both. Moreover, they derive the Pareto-optimality condition for distributed policies that balance

the variance and mean square of the service capacity. The performance of the optimal distributed policies is compared with that of the optimal centralized policy by deriving closed-form bounds. Finally, they discuss a scalable partially centralized algorithm that uses centralized information to boost performance and a method to deal with missing information on service requirements. Their finding can ultimately lead to more efficient power distribution networks and cloud computing services, which optimally match the service capacity to changing demands.

Binary Quadratic Program with Variable Partitioning Constraints

The binary quadratic program with variable partitioning constraints is a very general class of optimization problems that is very difficult to solve because of the nonconvexity and integrality of the variables and is ubiquitous, among others, in network design, computer vision, and transportation and logistics. In “A Convex Reformulation and an Outer Approximation for a Large Class of Binary Quadratic Programs,” Rostami, Errico, and Lodi show how to transform such a nonconvex challenging problem into a convex bilinear program with decomposable structure. The authors develop a branch-and-cut algorithm based on outer approximation cuts, in which the cuts are generated on the fly by efficiently solving separation subproblems. The results of their computational experiments on different problems confirm the efficacy of the solution methods in solving large-scale problem instances.

Computing Exchange Market Equilibrium in Strongly Polynomial Time

Strongly polynomial solvability is rarely the case for nonlinear optimization problems. Interestingly, it’s possible for certain market equilibrium computation problems, including Orlin’s 2010 algorithm for Fisher markets with linear utilities. In Fisher markets, buyers arrive with fixed budgets to be spent on goods. In the more general exchange economy model, goods are brought to the market by agents who will spend their revenues on other goods. This makes the problem more complicated. In “A Strongly Polynomial Algorithm for Linear Exchange Markets” Garg and Végh resolve a longstanding open problem in algorithmic game theory by computing the equilibrium of an exchange economy with linear utilities in strongly polynomial time. Their approach builds on the weakly polynomial algorithm by Duan and Mehlhorn from 2015 but requires several significant new ideas. Along the way, they develop a new technique of approximating a certain linear program by a two-variable per inequality system.

Practical Algorithms for Finding Nash Equilibria in Colonel Blotto Games

The Colonel Blotto game (initially introduced by Borel in 1921) is commonly used for analyzing a wide range of applications from the U.S. Presidential election to innovative technology competitions to advertising, sports, and politics. After around a century,

Ahmadinejad et al. provided the first polynomial-time algorithm for computing the Nash equilibria in Colonel Blotto games. However, their algorithm consists of an exponential-size LP solved by the ellipsoid method, which is highly impractical. In “Fast and Simple Solutions of Blotto Games,” Behnezhad, Dehghani, Derakhshan, Hajighayi, and Seddighin provide the first polynomial-size LP formulation of the optimal strategies for the Colonel Blotto game using linear extension techniques. They use this polynomial-size LP to provide a simpler and significantly faster algorithm for finding optimal strategies of the Colonel Blotto game. They further show this representation is asymptotically tight, which means there exists no other linear representation of the strategy space with fewer constraints.

Novel Connections Between Convex Optimization and Markov Models Lead to Faster Algorithms

Markov decision processes (MDPs) are used to model stochastic systems in many applications, but computing good policies becomes hard when the effective horizon become very large. In “A First-Order Approach to Accelerated Value Iteration,” Goyal and Grand-Clément present a connection between value iteration (VI) algorithms and gradient descent methods from convex optimization and use acceleration and momentum to design faster algorithms, with convergence guarantees for the computation of the value function of a fixed policy for reversible MDP instances. The authors provide a lower bound on the convergence properties of any first-order algorithm for solving MDPs, where no algorithm can converge faster than VI. Finally, they introduce safe accelerated value iteration (S-AVI), which alternates between accelerated updates and value iteration updates. The algorithm S-AVI is worst-case optimal and retains the theoretical convergence properties of VI while exhibiting strong empirical performances and providing significant speedups when compared with classical approaches for a large test bed of MDP instances.

Optimization Models to Tackle Gerrymandering

Gerrymandering has been a fundamental issue in American democracy for more than two centuries, with significant implications for electoral representation. Traditional optimization models for political districting primarily model nonpolitical fairness metrics such as the compactness of districts. In “Multiobjective Optimization for Politically Fair Districting: A Scalable Multilevel Approach,” Swamy, King, and Jacobson propose optimization models that explicitly incorporate political fairness objectives using political data from past elections. These objectives model fundamental fairness principles such as vote-seat proportionality (efficiency gap), partisan (a)symmetry, and competitiveness. They propose a solution strategy, called the multilevel algorithm, that solves large instances of the problem using a series of matching-based graph contractions. A case study on congressional districting in Wisconsin demonstrates that district plans balance the interests of the voters and the political parties.

Online Matching with Post-allocation Stochasticity

In “Online Matching with Stochastic Rewards: Optimal Competitive Ratio via Path-Based Formulation,” Goyal and Udwani develop a novel algorithm analysis approach to address stochastic elements in online matching. The approach leads to several new results that were previously out of reach for a fundamental generalization of online matching. More generally, the approach is useful for analyzing the performance of online algorithms for matching in settings with stochastic uncertainty that manifests after matching decisions are made.

How To Protect Customers’ Privacy in Personalized Pricing

With the rapid development of artificial intelligence and big data, the application of data-driven personalized pricing has been increasingly prevalent in real practices such as finance, insurance, and retailing. However, with the public’s growing concern of the abuse of their personal data, legislative efforts are being taken to guarantee data privacy. In “Differential Privacy in Personalized Pricing with Nonparametric Demand Models,” Chen, Miao, and Wang guarantee customers’ data privacy from the algorithm design of our dynamic personalized pricing policies. Two algorithms are developed with different levels of privacy guarantee. The first algorithm protects customers’ data in a centralized manner, meaning that the data aggregator (the pricing platform) is trusted, and the attacker is unlikely to know customers’ personal information. The second algorithm has a stronger privacy guarantee, which is mathematically proved to be able to protect customers’ data even when the data set is hacked. Besides privacy protection, both of our algorithms are effective in achieving near-optimal revenue maximization.

Solving Operational Problems Under Mixtures of Discrete Choice Models

Discrete choice models have recently attracted significant attention to model demand in revenue-management applications, as they can capture the fact that if a product is unavailable, then some customers substitute for this product, whereas others leave the system without a purchase. Although a more sophisticated choice model may capture the choice process of the customers more faithfully, a simpler choice model may result in tractable optimization problems when finding the optimal assortment of products to offer or prices to charge. One approach for coming up with sophisticated choice models is to mix existing ones, where the different segments of customers choose under the different choice models in the mixture. In “Revenue Management Under a Mixture of Independent Demand and Multinomial Logit Models,” Cao, Rusmevichientong, and Topaloglu demonstrate that mixing the independent demand and multinomial logit models can significantly increase the modeling flexibility of each of these choice models, while keeping the corresponding operational assortment optimization problems tractable.

How to Avoid Polarization in Social Networks

In order to facilitate social interactions and the design of deliberative and efficient institutions, a social planner should ensure that diverse opinions can be sustained and debated and ensuring social cohesion. In “On the Design of Public Debate in Social Networks,” Grabisch, Mandel, and Rusinowska provide a theory of the efficient design of public debate. They develop a model of the coevolution of opinions and social relations that allow them to frame this problem in a formal setting. This model of opinion dynamics accounts for the persistence of heterogeneous opinions in society (“strong diversity”) and the dynamic interactions between opinions and social connections. The social planner faces a trade-off between fostering the convergence of opinions in society and increasing the risk of polarization and instability. To resolve this trade-off, the social planner must account for both structural and behavioral characteristics: how fragile is the social network and to what extent individuals tolerate disagreement with their peers.

Reliable Machine Learning via Structured Distributionally Robust Optimization

Data sets used to train machine learning (ML) models often suffer from sampling biases and underrepresent marginalized groups.

Standard machine learning models are trained to optimize average performance and perform poorly on tail subpopulations. In “Distributionally Robust Losses for Latent Covariate Mixtures,” Duchi, Hashimoto, and Namkoong formulate a DRO approach for training ML models to perform uniformly well over subpopulations. They design a worst-case optimization procedure over structured distribution shifts salient in predictive applications: shifts in (a subset of) covariates. The authors propose a convex procedure that controls worst case subpopulation performance and provides finite-sample (nonparametric) convergence guarantees. Empirically, they demonstrate their worst-case procedure on lexical similarity, wine quality, and recidivism prediction tasks and observe significantly improved performance across unseen subpopulations.

Discriminatory Pricing and Introduction Timing of Upgrades in Subscription-Based Services

Many technologies that fuel subscription-based services improve over time. Examples are mobile phones, software suites (Microsoft Office, Adobe Creative Cloud), subscription services (Netflix), and cloud service providers. Additionally, modern subscription services are increasingly personalized to individual subscribers, and as a result, discriminatory pricing is ever-present in the marketplace, allowing firms to reward existing customers with discounts and special offers on upgrades. However, customers may be averse to switching to improved services because of costs related to redesigning business processes, downtime, or customer inertia. In “Optimal Pricing and Introduction Timing of Technology Upgrades in Subscription-Based Services,” Kash, Key, and Zoumpoulis propose a model of technology upgrades featuring discriminatory pricing based on customers’ upgrade experience. They characterize the optimal pricing policy for the service provider and develop an efficient algorithm for computing optimal prices. The authors also characterize the optimal timing of technology introductions and show that it is generally optimal to introduce new technologies in periodic intervals after some time.

Theory and Application of an Infinite-Dimensional Fisher Market Model

Fisher market equilibrium models have long been a central topic in economics and computation. Recently, they have been widely used in the design and implementation of Internet marketplaces. Although the classical models are well studied and can be solved via tractable optimization characterizations, they only allow a finite number of items and thus face scalability issues when the item space is huge or even continuous. In “Infinite-Dimensional Fisher Markets and Tractable Fair Division,” Gao and Kroer propose infinite-dimensional convex programs and show that they capture market equilibria for infinite and possibly continuous item spaces, extending the classical Eisenberg-Gale framework. Using these results, the authors show that a challenging cake-cutting problem for piecewise linear agent valuations is equivalent to finding a market equilibrium and admits a tractable convex optimization characterization. Thus, it can be solved in polynomial time in theory and highly efficiently by numerical optimization software.

Running Stores on Wheels for Spatially Queued Customers

Selling products in public spaces with wheeled vending stalls can potentially become ubiquitous in our future cities. Transition into such a “stall economy” paradigm is being spurred by the rapidly advancing self-driving technologies. In “Stall Economy: The Value of Mobility in Retail on Wheels,” Cao and Qi provide models, theory, and insights of spatial queuing systems, in which one

server moves around to meet mobile customers and in which the “last 100 meters” are expensive. They derive the dependence of customer waiting and stall repositioning on two key decisions: the service zone size and the walking distance imposed on customers to meet a stall. The authors also propose a stylized joint truck-stall routing model to capture the inventory replenishment operations. They find that the stall economy potentially profits more than stationary retail, not only because of the mobility of stalls for providing proximity to customers but also, because of its operational flexibilities that allow for avoiding the “last 100 meters” and pooling demands.

A Stackelberg Actor–Critic Approach to Optimal Market Making and Incentives Design with Dark Pools

In “Market Making and Incentives Design in the Presence of a Dark Pool: A Stackelberg Actor–Critic Approach,” Baldacci, Manziuk, Mastrolia, and Rosenbaum consider the issue of a market maker acting at the same time in the lit and dark pools of an exchange. The exchange wishes to establish a suitable make–take fee policy to attract transactions on its venues. They first solve the stochastic control problem of the market maker without the intervention of the exchange. Then, the authors derive the equations defining the optimal contract to be set between the market maker and the exchange. This contract depends on the trading flows generated by the market maker’s activity on the two venues. In both cases, they show existence and uniqueness, in the viscosity sense, of the solutions of the Hamilton–Jacobi–Bellman equations associated to the market maker and exchange’s problems. The authors finally design an actor–critic algorithm inspired by deep reinforcement learning methods, enabling us to approximate efficiently the optimal controls of the market maker and the optimal incentives to be provided by the exchange.

How Can We Build Resilient and Robust Supply Chain Systems in The Face of Large Shocks?

The recent coronavirus disease 2019 pandemic has shown that shortages and supply chain disruptions can have catastrophic effects on the real economy. These observations bring about reflections and first-order questions. How can we design supply chain networks that are robust and resilient to demand and supply shocks? Can we quantify the indirect effects caused by buyers’ and suppliers’ defaults in the network? Is it always cost effective to steer the system toward higher buyers’ and suppliers’ diversification? In “Disruption and Rerouting in Supply Chain Networks,” Birge, Capponi, and Chen argue that in highly capitalized networks, diversifying demand and supply across a larger number of counterparties may result in a more fragile network. Single-sourcing strategies are optimal for a firm only if the firm’s supplier default probability is low, but they perform worse than multiple-sourcing strategies otherwise.

Online Passenger Flow Control in Metro Lines

Crowds management during peak commuting hours is a key challenge facing metro systems worldwide, which results in serious safety concerns and unfair public transit service for commuters on different origin–destination (o–d) pairs. In “Online Passenger Flow Control in Metro Lines,” Liang, Lyu, Teo, and Gao investigate the impact of online decision making on the value of passenger flow control solution methodologies. The authors formulate the problem as a stochastic dynamic program with a fairness (fill rate) constraint and exploit Blackwell’s approachability theorem and Fenchel duality to characterize the attainable service level of each o–d pair. They use these insights to develop online policies that can enable more passengers boarding a train (efficiency) as well as

ensure equitable service level (fairness) provided to each o-d pair. Numerical experiments on a set of transit data from Beijing show that this approach performs well compared with existing benchmarks in the literature.

Does Subjective Evaluation of Probability Impact Asset Prices?

The Nobel Prize-winning capital asset pricing model (CAPM) predicts that expected excess return of any asset is positively proportional to its exposure to the overall market: the beta, leading to an upward-sloping security market line. However, this prediction is contradicted by empirical studies that the return-beta slope is often flat or even downward-sloping, a puzzle called the “beta

anomaly.” The CAPM is premised upon the notion that market participants are all rational, including that they are able to objectively evaluate probabilities. However, evidence abounds that individuals are often unable to do so, examples being purchase of lottery tickets and insurance products, in which the extremely small probabilities of winning or losing big are exaggerated. This phenomenon of distorting probabilities at both tails is called “probability weighting” (PW), which is a key component of modern behavioral finance. In “Beta and Coskewness Pricing: Perspective from Probability Weighting” Shi, Cui, and Zhou approach the beta anomaly through PW. They offer an explanation of the beta anomaly via a new theoretical CAPM involving PW and an extensive empirical study.