



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(5):iii-vii. <https://doi.org/10.1287/opre.2024.inthisissue.v72.n5>

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.n5>

Copyright: © 2024 INFORMS

Analyzing Production-Inventory Systems with General Demand: Cost Minimization and Risk Analytics

Frequent production rate changes are prohibitive because of high setup costs or setup times in producing such items as sugar, glass, computer displays, and cell-free proteins. Thus, constant production rates are deployed for producing these items even when their demands are random. In “Technical Note—Production Management with General Demands and Lost Sales,” Han, Li, Sethi, Siu, and Yam obtain the optimal constant production rate for a production-inventory system with Lévy demand for long-run average and expected discounted cost objectives, explicitly in some cases and numerically in general with a Fourier-cosine scheme they develop. This scheme can help in computing risk analytics of the inventory system, such as stockout probability and expected shortfall. These measures are particularly significant for assessing supply resilience, especially for emergency products or services like medicines and healthcare equipment. This study’s analytical and numerical findings contribute to enhancing efficiency and decision making in production management.

Asymptotic Optimality of Simple Heuristic Policies for Multidimensional Inventory Systems

Stochastic inventory systems with multidimensional state spaces, such as lost-sales system with positive lead time and perishable inventory system, are challenging to manage because of the curse of dimensionality, and their optimal control policies are extremely complex. In “Asymptotic Scaling of Optimal Cost and Asymptotic Optimality of Base-Stock Policy in Several Multidimensional Inventory Systems,” Bu, Gong, and Chao consider three classes of such systems in the regime of large unit penalty cost, and they establish the asymptotic optimality of (modified) simple base-stock policies as well as an explicit expression for the optimal cost rate in each of these systems. These results justify the applications of such policies in real-world applications, and they are achieved by constructing tight newsvendor upper and lower bounds on the systems’ costs and analyzing the asymptotic scaling of newsvendor costs with large unit penalty cost. This approach is expected to be useful in studying other multidimensional stochastic inventory systems.

Using Induced Order Statistics to Construct Optimal Impact Portfolios with General Dependence and Marginals

In “Optimal Impact Portfolios with General Dependence and Marginals,” Lo, Wu, Zhang, and Zhao develop a mathematical framework for constructing optimal impact portfolios and quantifying their financial performance by characterizing the returns of impact-ranked assets using induced order statistics and copulas. The results apply to any joint distribution of impact factors and residual returns, making them broadly applicable to a wide range of contexts. The authors develop significant extensions of the theory of induced order statistics, with which they are able to

characterize the distribution of residual returns of individual assets ranked by the impact factor. Their framework provides a toolkit for practitioners to construct impact portfolios and quantify their performance based on real data. This allows impact investors to achieve higher risk-adjusted returns than those with impact portfolios constructed using simpler heuristics such as negative or positive screening.

Inventory Projection and Asymptotic Optimality

In “Projected Inventory-Level Policies for Lost Sales Inventory Systems: Asymptotic Optimality in Two Regimes,” van Jaarsveld and Arts propose a new policy for the canonical periodic review lost sales inventory problem. Under this policy, orders are placed in each period such that the expected inventory level at the moment of replenishment order arrival is at a fixed level. This single-parameter policy is called the projected inventory-level (PIL) policy. The PIL inherits asymptotic optimality properties from the constant-order policy for long lead time and from the base-stock policy when the cost of losing a sale is large. The PIL policy has lower cost than competing heuristics in a numerical study. The PIL seems to be a promising approach for other inventory systems with a high-dimensional state space, such as the perishable item inventory management problem.

Assessing the Value of Payables Finance

Payables finance, also known as reverse factoring or supply chain finance, is a form of trade finance offered by a bank that provides a supplier with the option to receive a buyer’s payables early while allowing the buyer to extend its payment due date. There has been an increasing adoption of payables finance by various industries in recent years. In “Optimal Cash Management with Payables Finance,” Yan, Chen, and Ding characterize the supplier’s optimal cash policy under the payables finance arrangement with a buyer and a bank. The authors show that it is the supplier’s future cash flow uncertainty, together with the supplier’s risk averseness, that drives the cash liquidity value of payables finance. Numerical results of applying the analysis to data sets obtained from a major U.S. chemical company suggest that adopting payables finance can generate considerable value for the company and its suppliers.

Finding Robust Financial Network with Limited Availability of Data

In “Robust Financial Networks,” Hu, Mitchell, and Tompaidis study networks of financial institutions where only aggregate information on liabilities is available. The authors introduce the robust liability network, that is, the network consistent with the available information that exhibits the worst expected losses. They provide an algorithm to identify the robust liability network and, using aggregate data provided by bank holding companies to the Federal Reserve in form FR Y-9C, determine robust liability networks for U.S. banks under various network configurations. They show that the robust liability network is sparse, with links

between institutions that hold highly correlated portfolios. They illustrate the methodology in two applications. (1) They look at how robust liability networks changed around the onset of the COVID-19 pandemic. (2) They evaluate the impact of a potential regulation that limits risk-taking based on each institution's conditional value-at-risk. Their results can be used by regulators to monitor systemic risk in financial networks.

Random Audits to Counteract Evasion Incentives

When companies or organizations can evade audits on their harmful incidents, how should the affected entities design their audit and penalty policies? In "Audit and Remediation Strategies in the Presence of Evasion Capabilities," Wang, de Véricourt, and Sun find random audits may be needed in the optimal policy. Specifically, the optimal policy alternates between ascending monetary penalties (without any audits) and random audits at a constant rate (when the penalty reach its maximum level). Only when the evasion is ineffective or the self-correction is too costly do deterministic audits become optimal. They tackle the problem in a continuous-time principal-agent framework with both adverse selection and moral hazard.

Near-Optimal Bayesian Online Assortment of Reusable Resources

Motivated by rental services in e-commerce, the authors consider revenue maximization in the online assortment of reusable resources for different types of arriving consumers. In "Technical Note—Near-Optimal Bayesian Online Assortment of Reusable Resources," Feng, Niazadeh, and Saberi design competitive online algorithms compared to the optimal online policy in the Bayesian setting, where consumer types are drawn independently from known heterogeneous distributions over time. In scenarios with large initial inventories, their main result is a near-optimal competitive algorithm for reusable resources. The algorithm relies on an expected linear programming (LP) benchmark, solves this LP, and simulates the solution through independent randomized rounding. The main challenge is achieving inventory feasibility efficiently using these simulation-based algorithms. To address this, the authors design discarding policies for each resource, balancing inventory feasibility and revenue loss. Discarding a unit of a resource impacts future consumption of other resources, so the authors introduce post-processing assortment procedures to design and analyze their discarding policies. Additionally, the authors present an improved competitive algorithm for nonreusable resources and evaluate our algorithms using numerical simulations on synthetic data.

Improving Operating Room Efficiency

After each surgery in a hospital, every reusable instrument that has entered the operating room must be sterilized before it can be used again. Hospitals spend several million dollars a year on this sterilization process, instrument tray assembly, and instrument repurchase costs. However, 70%–80% of instruments that enter the operating room are not used at a large majority of hospitals resulting in wastage of millions of dollars. For a medium-sized hospital, approximately \$4–\$7 million are spent on the sterilization process. Across more than 5,000 hospitals in the United States, several billion dollars per year are wasted on unnecessary sterilization. In "Data-Driven Surgical Tray Optimization to Improve Operating Room Efficiency," Deshpande, Mundru, Rath, Knowles, Rowe, and Wood at the UNC Kenan-Flagler Business School, with partners from healthcare technology firm Operative Flow Technologies, have developed an algorithm that uses detailed data of instruments used to create optimized instrument tray configuration. The algorithm creates trays informed by

historical usage of instruments to reduce unused instruments in trays significantly. This algorithm scales for hundreds of trays, thousands of instruments, and thousands of surgeries and can be applied to major healthcare systems. At one implementation, the number of unused instruments was reduced by 54%. Their solution has subsequently been implemented at multiple healthcare systems leading to several million dollars of cost savings for the hospitals.

A Simple Monotonic Dynamic Pricing Policy

In "Simple Monotonic Readjustment Policies with Applications to Markdown Pricing and Pricing in the Presence of Strategic Customers," Chen and Jasin consider a canonical revenue management setting in which a monopolist seeks to maximize expected total revenues by selling a fixed inventory of product to customers who arrive sequentially over time. It is known in the literature that a simple dynamic pricing policy can perform reasonably well in such setting. However, the literature typically assumes that price can move freely over time (i.e., it can either go up or down). In practice, firms often impose certain restrictions on price movement. In this paper, the authors consider the setting in which price can only move monotonically over time (either only decreasing or only increasing). The authors devise a simple provably near-optimal dynamic pricing policy. Aside from markdown applications, the policy can also be applied to the setting with strategic customers.

When and Why do Policy Gradient Methods Find Globally Optimal Policies?

Policy gradient methods, which have powered a lot of recent success in reinforcement learning, search for an optimal policy in a parameterized policy class by performing stochastic gradient descent on the cumulative expected cost-to-go under some initial state distribution. Although widely used, these methods lack theoretical guarantees as the optimization objective is typically nonconvex even for simple control problems, and hence are understood to only converge to a stationary point. In "Global Optimality Guarantees for Policy Gradient Methods," Bhandari and Russo identify structural properties of the underlying Markov Decision Process (MDP) that guarantee that despite nonconvexity, the optimization objective has no suboptimal stationary points, ensuring asymptotic convergence of policy gradient methods to globally optimal policies. Under stronger conditions, authors show the policy gradient objective to satisfy a Polyak-Lojasiewicz (gradient dominance) condition that yields fast convergence rates. In addition, when some of the said conditions are relaxed, authors provide bounds on the suboptimality gap of any stationary point. The results rely on a key connection with policy iteration, a classic dynamic programming algorithm which solves a single period optimization problem at every step. The authors show how structure in the single period optimization problems solved by policy iteration translate into nice properties of the multiperiod policy gradient objective, making it amenable for first-order methods to find globally optimal solutions. The authors also instantiate their framework for several classical control problems including tabular and linear MDPs, linear quadratic control, optimal stopping, and finite-horizon inventor control problems.

Optimizing Revenue in Two-Sided Online Markets through Strategic Quality Selection and Information Disclosure

Ensuring profitable operations in two-sided online marketplaces demands an astute analysis of seller quality and strategic information sharing with buyers. In "Quality Selection in Two-Sided Markets: A Constrained Price Discrimination Approach," Light,

Johari, and Weintraub delve into the nuances of this operation. The authors explore the challenge that platforms encounter in deciding which sellers to allow and how much quality information to share with buyers in order to enhance platform revenue. Utilizing two distinct two-sided market models, the paper unveils conditions under which adopting straightforward information structures, such as excluding certain sellers or not differentiating among participating sellers, proves to be a revenue-maximizing strategy. This study utilizes a constrained price discrimination problem to reveal specific strategies platforms can use to adjust information structures in diverse market scenarios, providing insights for digital platforms aiming to navigate the marketplace more effectively.

Managing Volunteer Engagement

Nonprofit organizations that provide food, shelter, and other services to people in need, rely on volunteers to deliver their services. Unlike paid labor, nonprofit organizations have less control over unpaid volunteers' schedules, efforts, and reliability. However, these organizations can invest in volunteer engagement activities to strive for a steady and adequate supply of volunteer labor. In "A Dynamic Model for Managing Volunteer Engagement," Ata, Tongarlak, Lee, and Field model the volunteer management process. The paper derives a dynamic control policy that judiciously uses the nonprofit's resources for volunteer engagement activities depending on explicit congestion thresholds, thereby lowering the overall cost of the nonprofit operation.

A Pareto Dominance Principle for Data-Driven Optimization

In "A Pareto Dominance Principle for Data-Driven Optimization," Sutter, Van Parys, and Kuhn propose an effective way to make decisions based on data for uncertain situations. In simple terms, a data-driven decision is just a choice we make by looking at the available data. The authors express this choice as the best one according to a model they create from the data. The quality of this decision is judged by how well it performs in situations not seen during training. The authors also consider how often it disappoints in those situations. The challenge is that the authors do not know the exact probability of generating the data. An ideal data-driven decision should work well for any possible probability. However, such ideal decisions are usually not possible. Therefore, the authors look for decisions that work well on unseen data, considering the chances of disappointment. The authors prove that such effective decisions exist under certain conditions, allowing for practical applications. This approach holds regardless of whether the original problem is simple or complex, and it works even when the data are not uniformly collected. The study also uncovers how the characteristics of the data-generating process influence the optimal decision-making model.

Deciding Who Gets an Interview, and in Which Order

Hiring the right person for the job is crucial for the success of any organization. In "Selection and Ordering Policies for Hiring Pipelines via Linear Programming," Epstein and Ma study several problems motivated by a firm that is carrying out a recruitment process. Restricted to a finite time budget, the firm must decide who will be interviewed out of a pool of applicants and who will receive offers among the interviewed applicants. They develop approximation algorithms with constant factor guarantees that approach optimality when the number of vacant positions grows large. The algorithms they develop are nonadaptive: they fix a subset of candidates and an order to conduct the interviews and the order remains unchanged, independent of the outcomes of

other interviews. Thus, they establish bounds on the adaptivity gap: a worst-case measure of how poorly nonadaptive algorithms can perform with respect to their adaptive counterparts.

Improving Disaster Preparedness Through Mutual Catastrophe Insurance

In "A Mutual Catastrophe Insurance Framework for Horizontal Collaboration in Prepositioning Strategic Reserves," Zbib, Balcik, Rancourt, and Laporte present an innovative approach to collaborative disaster preparedness. The novel framework considers a risk-averse mutual insurer offering multiyear insurance contracts with coverage deductibles and limits to a portfolio of risk-averse policyholders. It is designed to foster horizontal collaboration among policyholders for joint disaster preparedness by effectively integrating operational and financial functions. The problem is modeled as a large-scale nonlinear multistage stochastic program and solved by using an effective Benders decomposition algorithm. The framework is validated with real data from 18 Caribbean countries focusing on hurricane preparedness. Given the predicted impacts of climate change, the proposed multiyear mutual catastrophe insurance framework promises to reshape global disaster preparedness and make a profound societal impact by providing a transparent disaster financing plan to protect vulnerable regions. The study's findings stress the importance of long-term cooperation, prenegotiation of indemnification policies, and strategic setting of deductibles and limits by taking into account the correlation between policyholders.

Navigating Uncertainty: The Surprising Benefits of Randomized Product Assortments

When a firm selects an assortment of products to offer to customers, it uses a choice model to anticipate their probability of purchasing each product. In practice, the estimation of these models is subject to statistical errors, which may lead to significantly suboptimal assortment decisions. In "Randomized Assortment Optimization," Wang, Peura, and Wiesemann show that the standard approach of deterministically selecting a single assortment to offer is not always optimal in this setting: Instead, the firm can do better by selecting an assortment randomly according to a prudently designed probability distribution. The authors show how an optimal randomization strategy can be determined for common choice models, improving performance in realistic data-driven settings. The results suggest that more general versions of the assortment optimization problem—incorporating business constraints, more flexible choice models and/or more general forms of uncertainty—tend to be more receptive to the benefits of randomization.

Computing the Solution of Blotto Games, a Century-Year-Old Open Problem

The Blotto game was introduced by Borel more than a century ago in a seminal game theory paper. Surprisingly, characterizing and computing the optimal solutions (or the Nash equilibria) of this game in general was still open! In "An Algorithmic Solution to the Blotto Game Using Multimarginal Couplings," Perchet, Rigollet, and Le Gouic combine theoretical tools from statistics and auction theory to provide an almost explicit characterization of the strategies that are then computed using techniques from optimal transport (and other algorithmic ideas). The complexity of the final algorithm is even polynomial in the number of battlefields (a key parameter of a Blotto game) and the approximation level.

On Hiring Secretaries with Stochastic Departures

In "Technical Note—On Hiring Secretaries with Stochastic Departures," Kesselheim, Psomas, and Vardi study generalization

of the secretary problem, where decisions do not have to be made immediately upon applicants' arrivals. After arriving, each applicant stays in the system for some (random) amount of time and then leaves, whereupon the algorithm has to decide irrevocably whether to select this applicant or not. The arrival and waiting times are drawn from known distributions, and the decision maker's goal is to maximize the probability of selecting the best applicant overall. The paper characterizes the optimal policy for this setting, showing that when deciding whether to select an applicant, it suffices to know only the time and the number of applicants that have arrived so far. Furthermore, the policy is monotone nondecreasing in the number of applicants seen so far, and, under certain natural conditions, monotone nonincreasing in time. Furthermore, when the number of applicants is large, a single threshold policy is almost optimal.

Epic Math Battles: Nash vs. Pareto

What is the relation between the notion of Nash equilibria and Pareto-optimal points? It is well known that Nash equilibria do not need to be Pareto optimal, and Pareto points do not need to be Nash equilibria. However, in "Technical Note—Characterizing and Computing the Set of Nash Equilibria via Vector Optimization," Feinstein and Rudloff take a deeper look at the relation. It is shown that it is possible to characterize the set of all Nash equilibria as the set of all Pareto-optimal solutions of a certain vector optimization problem. This is accomplished by carefully designing the objective function and the ordering cone of the vector optimization problem such that both notions coincide. This characterization holds for all noncooperative games (nonconvex, convex, linear). It opens up a new way of computing Nash equilibria, as one can now use techniques and algorithms from vector optimization to compute the set of all Nash equilibria, which is in contrast to the classical fixed-point iterations that find just a single Nash equilibrium.

Real-Time Spatial–Intertemporal Pricing and Relocation in a Ride-Hailing Network: Near-Optimal Policies and The Value of Dynamic Pricing

In "Real-Time Spatial–Intertemporal Pricing and Relocation in a Ride-Hailing Network: Near-Optimal Policies and the Value of Dynamic Pricing," Chen, Lei, and Jasin consider a dynamic pricing problem faced by a ride-hailing service provider who manages a fixed number of servers and serves price-sensitive customers within a network. Servers serve arriving customers by relocating from the requested origins to destinations within a certain travel time. The authors first propose a static pricing policy based on the optimal solution to a deterministic relaxation of the original stochastic problem. They show that the proposed static policy matches the best possible asymptotic performance of any static policy. The authors further propose a dynamic pricing policy that adaptively changes the prices in a way that reduces the impact of past demand randomness on the balance of future distributions of servers and customers across the network. They show that the dynamic pricing policy achieves significantly better asymptotical performance. The proposed policies and their performance guarantees are further extended to a case where the firm jointly decides the relocation of vacant servers.

New Wine into Old Bottles: A New Ranking and Selection Approach Promising to Overcome the Curse of Dimensionality

The curse of dimensionality has long been one of the biggest challenges in solving large-scale simulation ranking and selection (R&S) problems. As the number of systems grows, existing

approaches to R&S relying on the Bonferroni correction become increasingly conservative, rendering them overachieving in error control and consuming more computational resources than necessary. In "Bonferroni-Free and Indifference-Zone-Flexible Sequential Elimination Procedures for Ranking and Selection," Wang, Wan, and Chen develop Bonferroni-free and indifference-zone-optional ranking and selection procedures to deliver the prescribed probabilistic guarantee without overshooting. Their approach is to conduct always valid and fully sequential hypothesis tests that enable continuous monitoring of each candidate system and control the probability of correct selection. In addition, the indifference-zone parameter becomes dispensable in their procedures; however, when provided appropriately, it could improve the procedures' computational and statistical efficiency.

A Robust Optimization Perspective on Minkowski Centers

Properly defining the center of a set has been a longstanding question in applied mathematics, with implications in numerical geometry, physics, and optimization algorithms. Minkowski centers are one such definition, whose theoretical benefits are numerous and well documented. In "Minkowski Centers via Robust Optimization: Computation and Applications," den Hertog, Pau-philet, and Soali revisit the advantages of Minkowski centers from a computational, rather than theoretical, perspective. First, the authors show that Minkowski centers are solutions to a robust optimization problem. Under this lens, the authors then provide computationally tractable reformulations or approximations for a series of sets, including polyhedra, polyhedral projections, and intersections of ellipsoids. Computationally, the authors illustrate that Minkowski centers are viable alternatives to other centers, such as Chebyshev or analytic centers, and can speed up the convergence of numerical algorithms like hit-and-run and cutting-plane methods. The authors hope their work sheds new and practical light on Minkowski centers and exposes their potential benefits as a computational tool.

Mixed-Integer Formulations for Power Production Problems

The unit commitment problem is a complex mixed-integer nonlinear program that originates in the field of power production. Although it arises in a monopolistic system, there is still great attention to this problem even in a free-market regime, where it constitutes only a subproblem of larger ones. Historically, it was usually solved by Lagrangian relaxation methods. However, the advances achieved by commercial solvers of mixed-integer (linear and convex) programming problems have made such approaches an attractive option. In "New Mixed-Integer Non-linear Programming Formulations for the Unit Commitment Problems with Ramping Constraints," Bacci, Frangioni, Gentile, and Tavlaridis-Gyparakis present the first mixed-integer nonlinear programming formulation with a polynomial number of both variables and constraints that describes the convex hull of the feasible solutions of the unit commitment problem with a single thermal generation unit, comprising all typical constraints and convex power generation costs. Proving that the formulation is exact requires a new result about the convex envelope of specially structured functions that can have independent interest. This new formulation for a single power generation unit is used to derive three new formulations for the general unit commitment problem whose effectiveness has been tested against the state-of-art formulation.

Dynamic Resource Constrained Reward Collection Problems: Unified Model and Analysis

Dynamic resource allocation problems arise under a variety of settings. In “Survey of Dynamic Resource-Constrained Reward Collection Problems: Unified Model and Analysis,” Balseiro, Besbes, and Pizarro introduce a unifying model for a large class of dynamic optimization problems dubbed *dynamic resource-constrained reward collection* (DRC^2) problems. Surveying the literature, they show that this class encompasses a variety of disparate and classical problems typically studied separately, such as dynamic pricing with capacity constraints, dynamic bidding with budgets, network revenue management, online matching, or order fulfillment. Furthermore, they establish that the DRC^2 class is amenable to analysis by characterizing the performance of a central, certainty-equivalent heuristic. Notably, they provide a novel unifying analysis that isolates the drivers of performance, recovers as corollaries some existing specialized results, generalizes other existing results by weakening the assumptions required, and yields new results in specialized settings for which no such characterization was available.

Novel Optimality Cuts for Two-Stage Stochastic Mixed-Integer Programs

The applicability and use of two-stage stochastic mixed-integer programs is well-established, thus calling for efficient decomposition algorithms to solve them. Such algorithms typically rely on optimality cuts to approximate the expected second stage cost function from below. In “A Converging Benders’ Decomposition Algorithm for Two-Stage Mixed-Integer Recourse Models,” van der Laan and Romeijnders derive a new family of optimality cuts that is sufficiently rich to identify the optimal solution of two-stage stochastic mixed-integer programs in general. That is, general mixed-integer decision variables are allowed in both stages, and all data elements are allowed to be random. Moreover, these new optimality cuts require computations that decompose by scenario, and thus, they can be computed efficiently. The authors demonstrate the potential of their approach on a range of problem instances, including the dynamic capacity acquisition and assignment problem from the stochastic integer programming test problem library.

Exponential Cone Programming (ECP), a Novel Approach to Solving Large-Scale Optimization of Electric Vehicle Charging and Beyond

In “An Exponential Cone Programming Approach for Managing Electric Vehicle Charging,” Chen, He, and Zhou propose a novel ECP approach to solving large-scale optimization of electric vehicle charging in public stations such as EVgo. Other than the stochastic arrivals of customers with different arrival/departure times and charging requirements, charging stations routinely incur high demand charges, costs related to the highest per-period total electricity used in a billing cycle, which can be as high as 70% of the total electricity bill. For the case with unlimited chargers, the authors characterize the theoretical performances of the ECP approach. For the case with limited chargers, the authors construct an ECP leveraging the idea from distributionally robust optimization and show in a data-calibrated numerical study that it performs better than common approaches, considering many practical implementation issues. The authors’ method of constructing ECPs can be potentially applicable to approximate more general two-stage stochastic programs.

Utility Preference Robust Optimization with Moment-Type Information Structure

In some decision-making problems, information on the true utility function of the decision maker (DM) may be incomplete, which may bring potential modeling risk. In “Utility Preference Robust Optimization with Moment-Type Information Structure,” Guo, Xu, and Zhang propose a maximin utility preference robust optimization model where information about the DM’s preference is constructed by moment-type conditions. The authors propose a piecewise linear approximation approach to tackle the maximin problem, reformulate the approximate problem as a single mixed integer linear program, and derive error bounds for the approximate ambiguity set, the optimal value, and the optimal solutions. To examine the performance of the model and the computational schemes, they carry out extensive numerical tests and demonstrate the effectiveness of the model and efficiency of the computational methods.