



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(1):iii-vi. https://doi.org/10.1287/opre.ITI_72_1

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.IT1_72_1

Copyright: © 2024 INFORMS

Robust Risk Quantification via Shock Propagation in Financial Networks

Despite the significance of risk contagion in financial networks, uncertainties arise in interbank network structures because of limited information. In “Robust Risk Quantification via Shock Propagation in Financial Networks,” Ahn, Chen, and Kim address this by proposing a robust optimization approach to estimate worst-case default probabilities and capital requirements for a specific group of banks (e.g., systemically important financial institutions). By applying this tool, they analyze the impact of different incomplete network information structures and gain regulatory insights into gathering actionable network information.

Demand Estimation Under Uncertain Consideration Sets

In “Demand Estimation Under Uncertain Consideration Sets,” Jagabathula, Mitrofanov, and Vulcano investigate statistical properties of the consider-then-choose (CTC) models, which gained recent attention in the operations literature as an alternative to the classical random utility (RUM) models. The general class of CTC models is defined by a general joint distribution over ranking lists and consideration sets. Starting from the important result that the CTC and RUM classes are equivalent in terms of explanatory power, the authors characterize conditions under which CTC models become identified. Then, they propose expectation-maximization (EM) methods to solve the related estimation problem for different subclasses of CTC models, building from the provably convergent outer-approximation algorithm. Finally, subclasses of CTC models are tested on a synthetic data set and on two real data sets: one from a grocery chain and one from a peer-to-peer (P2P) car sharing platform. The results are consistent in assessing that CTC models tend to dominate RUM models with respect to prediction accuracy when the training data are noisy (i.e., transaction records do not necessarily reflect the physical inventory records) and when there is significant asymmetry between the training data set and the testing data set. These conditions are naturally verified in P2P sharing platforms and in retailers working on long-term forecasts (e.g., semester long) or geographical aggregate forecasts (e.g., forecasts at the distribution center level).

How to Optimally Design Contests in Development Settings for Rapid and Efficient Results

Firms increasingly leverage their expert suppliers to develop innovative products and services while minimizing both time-to-market and expenses. In “Dynamic Development Contests,” Khorasani, Korpeoglu, and Krishnan study how companies can effectively organize contests to stimulate development efforts from competing suppliers to minimize project lead times while keeping the cost of incentives in check. Their characterization of the optimal designs reveals that flexible rewards are most effective when funds are sufficient. However, in cases where resources are

constrained, strategic information disclosure in the form of probabilistic disclosure of any change in the state of competition serves as a valuable incentive instrument that helps contest organizers achieve their lead-time minimization objectives.

A Sequential Model for High-Volume Recruitment Under Random Yields

High-volume recruiting is challenging, as it involves hiring a larger number of people in a short amount of time. In “A Sequential Model for High-Volume Recruitment Under Random Yields,” Du, Li, and Yu model a high-volume recruiting process as a large-scale dynamic program. Their model captures three important features of high-volume recruiting: multiple phases, random yields, and a preset hiring target. They provide a decision tool to answer practical questions about the number of offers to be made in each phase and the number of phases in a recruitment season. To solve the dynamic program, they rely on approximations, and the approximations are asymptotically optimal when the volume is large. Their simulation studies confirm the convergence results. They illustrate how their modeling framework can be put into practice in a case study, which shows that their decision tool can improve the recruiting outcome.

Optimizing Delivery On-Time Performance with Region Partitioning

Managing on-time delivery systems is challenging because of the underlying uncertainties and combinatorial nature of the routing decision. In practice, the efficiency of such systems also hinges on the driver’s familiarity with the local neighborhood. In “Provably Good Region Partitioning for On-Time Last-Mile Delivery,” Carlsson, Liu, Salari, and Yua study a region partitioning policy to minimize the expected delivery time of customer orders in a stochastic and dynamic setting. This policy assigns every driver to a subregion, ensuring that drivers are only dispatched to their territories. The authors characterize the structure of the optimal partitioning policy and show its expected on-time performance converges to that of the flexible dispatching policy in heavy traffic. The optimal characterization features two insightful conditions that are critical to the on-time performance of last-mile delivery systems. Furthermore, the paper develops partitioning algorithms with performance guarantees, leveraging ham sandwich cuts and three-partitions from discrete geometry.

Portfolio Management in a Systemic Context

Ever since the pioneering work of Markowitz on mean-variance analysis, it has been commonly agreed that an investor should diversify their portfolio to reduce risks and maximize payoffs. In “Systemic Portfolio Diversification,” Capponi and Weber argue that this view is limiting if one were to account for the entire financial system rather than the investor in isolation. Diversification can then expose investors and the real economy to high systemic risk during periods of market distress. The authors show that, during a crisis, investors who hold the same underperforming

asset may face the necessity of liquidating their positions simultaneously. This process would result in spillover losses across investors, initiating a self-fulfilling feedback loop which depresses asset prices and reduces the portfolio value of all investors in the economy. They argue that the reduction of portfolio overlaps across different institutions may lower the likelihood of costly widespread sell-offs. Capponi and Weber introduce the new concept of systemic diversification, which prescribes a balance between asset diversification and diversification of fire-sale risk across investors.

A Near-Optimal Capacity and Scheduling Policy for Cloud Users

In “Technical Note—Cloud Cost Optimization: Model, Bounds, and Asymptotics,” Qu, Dawande, and Janakiraman study a long-term, dynamic resource-optimization problem for firms managing various cloud resources to process incoming computing tasks over time. Cloud resources vary in attributes, such as computing speed, memory, accelerators, and storage, tailored to different computing tasks, such as data warehousing, scientific computing, machine learning, and data processing. Firms can choose reserved resources for long-term commitments at lower costs or on-demand resources for flexibility at a higher price. Firms face three key trade-offs: the cost disparity between reserved and on-demand resources, the variation in resource attributes affecting performance and cost, and the challenge of balancing delay and resource costs. The authors propose an asymptotically optimal policy, demonstrating its effectiveness through a detailed numerical study based on Amazon Web Services data.

A Convex Programming Framework for Information Design Under Realistic Human Behavior

Platform markets and services typically have additional relevant information in comparison with their users. Information design studies how the sharing of this information can be leveraged by the platform to influence user behavior and obtain desirable outcomes. Previous research has studied information design assuming that the users act to maximize their expected utility, but this assumption does not always hold in reality. Instead, people often exhibit biases and deviations from expected utility maximization. In “Persuading Risk-Conscious Agents: A Geometric Approach,” Anunrojwong, Iyer, and Lingenbrink study information design with “risk-conscious” agents whose utility functions may depend nonlinearly on their beliefs. They provide a convex programming approach for solving for the optimal persuasion mechanism and establish their structural properties in different settings. They illustrate their approach in an application involving the sharing of waiting-time information in a queueing system. Overall, this work contributes to the study of information design under realistic models of human behavior.

Improved Price of Anarchy Bounds In Scheduling Games Via Resource-Aware Mechanisms

In large distributed systems, ensuring the efficient utilization of the available resources is a very challenging task. Given limited information regarding the state of the system and no centralized control over the outcome, decentralized scheduling mechanisms are unable to enforce optimal utilization. To better understand such systems, some classic papers that introduced game theoretic models used the “price of anarchy” measure to evaluate the system’s performance. In “Resource-Aware Cost-Sharing Methods for Scheduling Games,” Christodoulou, Gkatzelis, and Sgouritsa overcome some of the overly pessimistic results shown in this prior work by enhancing the scheduling mechanisms with access

to some additional information regarding the state of the system: a “resource-aware” mechanism knows what machines are available in the system and uses this information to carefully incentivize the users toward more efficient Nash equilibrium outcomes.

Multimarket Multireservoir Hydro Scheduling

Peak/off-peak spreads on European electricity forward and spot markets are eroding due to the ongoing nuclear phaseout in Germany and the steady growth in photovoltaic capacity. The reduced profitability of peak/off-peak arbitrage forces hydro-power producers to recover part of their original profitability on the reserve markets. In “A Planner-Trader Decomposition for Multimarket Hydro Scheduling” Schindler, Rujerapaiboon, Kuhn, and Wiesemann propose a bi-layer stochastic programming framework that jointly optimizes the trading strategies on the spot and reserve markets. The model faithfully accounts for uncertainty in electricity prices, water inflows, and reserve activations, and it ensures that the hydropower producers can fulfill their market commitments under any circumstances. The model is numerically challenging due to the various sources of uncertainty that are revealed at different time scales and that affect the problem’s objective function and constraints, and the authors propose a new planner-trader decomposition and an information restriction for its solution. A case study based on real data from Austria reveals significant benefits of simultaneously participating in the spot and the reserve markets.

Breaking the Sample Size Barrier in Model-Based Reinforcement Learning with a Generative Model

In “Breaking the Sample Size Barrier in Model-Based Reinforcement Learning with a Generative Model,” Li, Wei, Chi, and Chen study a central issue in modern reinforcement learning, the sample efficiency, and makes progress toward solving an idealistic scenario that assumes access to a generative model or a simulator. Despite a large number of prior works tackling this problem, a complete picture of the trade-offs between sample complexity and statistical accuracy has yet to be determined. In particular, all prior results suffer from a severe sample size barrier in the sense that their claimed statistical guarantees hold only when the sample size exceeds some enormous threshold. The current paper overcomes this barrier and fully settles this problem; more specifically, they establish the minimax optimality of the model-based approach for any given target accuracy level. To the best of our knowledge, this work delivers the first minimax-optimal guarantees that accommodate the entire range of sample sizes (beyond which finding a meaningful policy is information theoretically infeasible).

Is Q-Learning Minimax Optimal? A Tight Sample Complexity Analysis

In “Is Q-Learning Minimax Optimal? A Tight Sample Complexity Analysis,” Li, Cai, Chen, Wei and, Chi investigate a model-free algorithm of broad interest in reinforcement learning, namely, Q-learning. Whereas substantial progress had been made toward understanding the sample efficiency of Q-learning in recent years, it remained largely unclear whether Q-learning is sample-optimal and how to sharpen the sample complexity analysis of Q-learning. The authors settle these questions: (1) When there is only a single action, they show that Q-learning (or, equivalently, TD learning) is provably minimax optimal. (2) When there are at least two actions, their theory unveils the strict suboptimality of Q-learning and rigorizes the negative impact of overestimation in Q-learning. The authors theory accommodates both the synchronous case (i.e., the case in which independent samples are drawn)

and the asynchronous case (i.e., the case in which one only has access to a single Markovian trajectory).

Solving Rank Constrained Least Squares via Recursive Importance Sketching

In statistics and machine learning, we sometimes run into the rank-constrained least squares problems, for which the authors need to find the best low-rank fit between sets of data, such as trying to figure out what factors are affecting the data, filling in missing information, or finding connections between different sets of data. In “Recursive Importance Sketching for Rank Constrained Least Squares: Algorithms and High-Order Convergence,” Luo, Huang, Li, and Zhang introduce a new method for solving this problem called the recursive importance sketching algorithm (RISRO), in which the central idea is to break the problem down into smaller, easier parts using a unique technique called “recursive importance sketching.” This new method is not only easy to use, but it is also very efficient and gives accurate results. They prove that RISRO converges in a local quadratic-linear and quadratic rate under some mild conditions. Simulation studies also demonstrate the superior performance of RISRO.

Demand Estimation When Customers Choose from A Product Hierarchy

In “Estimating Large-Scale Tree Logit Models,” Jagabathula, Rusmevichientong, Venkataraman, and Zhao tackle the demand estimation problem under the tree logit model, also known as the nested logit or d-level nested logit model. The model is ideal for scenarios in which products can be grouped naturally based on their attributes into a hierarchy or taxonomy, such as flight itineraries grouped by departure time (morning or evening) and number of stops (nonstop or one stop). The current estimation methods are not practical for real-world applications that can involve hundreds or even thousands of products. The authors develop a fast, iterative method that computes a sequence of parameter estimates using simple closed-form updates by exploiting the structure of the negative log-likelihood objective. Numerical results on both synthetic and real data show that their proposed algorithm outperforms state-of-the-art optimization methods, especially for large-scale tree logit models with thousands of products.

The Attractiveness of Static Pricing for Network Revenue Management

For a variety of practical reasons, including ease of understanding and familiarity for customers, firms use static-pricing algorithms in which product prices remain unchanged over time. In a multi-product setting with no demand substitutability or complementarity between products, prior work has shown that static-pricing algorithms can offer excellent performance. In “Technical Note—A Note on State-Independent Policies in Network Revenue Management,” Manchiraju, Dawande, and Janakiraman extend that work and show that static-pricing algorithms offer a near-optimal performance even in the presence of demand substitutability and complementarity.

Incomplete Information VCG Contracts for Common Agency

The “common agency” model, introduced by Bernheim and Whinston in 1986 combines the fundamental challenge of the principal–agent model with the challenges of coordinating multiple principals. In “Technical Note—Incomplete information VCG contracts for common agency,” Alon, Talgam-Cohen, Lavi, and Shamash show that the class of common agency settings for which there exists a contract that guarantees truthfulness of all

principals, welfare maximization, and the two standard properties from contract theory—limited liability for the agent and individual rationality for the principals—is identifiable by a polynomial-time algorithm. Furthermore, for these settings, the authors design a polynomial-time computable contract: given valuation reports from the principals, it returns, if possible for the setting, a payment scheme for the agent that constitutes a contract with all desired properties.

Solving Competitive Facility Location and Bilevel MINLP

In “Sequential Competitive Facility Location: Exact and Approximate Algorithms,” Qi, Jiang, and Shen consider a competitive facility location problem (CFLP), where two firms sequentially open new facilities within their budgets and maximize their market shares of demand following a probabilistic choice model. They derive an equivalent, single-level mixed-integer nonlinear program (MINLP) reformulation of the bilevel MINLP and exploit the problem structures to derive valid inequalities based on submodularity and concave overestimation. The authors also develop an approximation algorithm with a constant approximation guarantee. They further study several extensions of CFLP that have general facility-opening costs, outside competitors, and diverse facility-planning decisions. Their approaches significantly accelerate the computation of CFLP on large-sized instances that have not been solved optimally or even heuristically by existing methods.

Adaptive Patient Flow Management

Appointment scheduling has significant clinical, operational, and economical impact on healthcare systems. An informed scheduling strategy that can effectively match patient demand and service capacity dynamically is vital for the business of medical providers, quality of care, and patient satisfaction. By regulating patient flow via an appointment system, healthcare providers can mitigate arrival process variability and improve operational performance. The simultaneous consideration of appointment day (interday scheduling) and time of day (intraday scheduling) in dynamic scheduling decisions is an important theoretical and practical problem that has remained open because of its stochastic nature, complex structure, and large dimensionality. In “Dynamic Interday and Intraday Scheduling,” Zacharias, Liu, and Begen fill this critical gap in the literature. They introduce a novel dynamic programming framework, designed with the intention of bridging two independently established streams of literature, and to leverage their latest advances in tackling the joint problem. The authors advance the theory of the field to provide a rigorous and practically implantable solution.

Selecting Good Solutions from Millions of Options Using Parallel Simulation Optimization

Ranking and selection (R&S) procedures in simulation optimization simulate every feasible solution to provide global statistical error control, often selecting a single solution in finite time that is optimal or near-optimal with high probability. By exploiting parallel computing advancements, large-scale problems with hundreds of thousands and even millions of feasible solutions are suitable for R&S. Naively parallelizing existing R&S methods originally designed for a serial computing setting is generally ineffective, however, as many of these conventional methods uphold family-wise error guarantees that suffer from multiplicity and require pairwise comparisons that present a computational bottleneck. Parallel adaptive survivor selection (PASS) is a new framework specifically designed for large-scale parallel R&S. In “Parallel Adaptive Survivor Selection,” Pei, Nelson, and Hunter

compare systems to an adaptive “standard” that is learned as the algorithm progresses, PASS eliminates inferior solutions with false elimination rate control and with computationally efficient aggregate comparisons rather than pairwise comparisons. PASS satisfies desirable theoretical properties and performs effectively on realistic problems.

Designing Selling Mechanisms when Buyers are Imperfect Optimizer

An assumption that is pervasive in revenue management and economics is that buyers are perfect optimizers. However, in practice, buyers may be limited by their computational capabilities or lack of information and may not be able to perfectly optimize their response to a selling mechanism. This has motivated the introduction of approximate incentive compatibility as a solution concept for practical mechanisms. In “Mechanism Design under Approximate Incentive Compatibility,” Balseiro, Besbes, and Castro study, for the first time, the problem of designing optimal selling mechanisms when buyers are imperfect optimizers. Their work characterizes structural properties of approximate incentive compatible mechanisms and establishes fundamental bounds on how much revenue can be garnered by moving from exact to approximate incentive constraints. Their work brings a new perspective to the theory of mechanism design by shedding light on a novel class of optimization problems, techniques, and challenges that emerge when relaxing incentive constraints.

Approximate Dynamic Programming for Vector Repeated Games

A seminal study showed repeated games with vector payoffs, which has had wide-ranging applications to decision making under uncertainty. Although this study characterized the achievable guarantees in such games under the long-run average payoff criterion, the characterization and computation of these guarantees under the discounted payoff criterion have remained a significant gap in the literature. In “An Approximate Dynamic Programming Approach to Repeated Games with Vector Losses,” Kamble, Loiseau, and Walrand develop a set-valued dynamic programming approach along with an approximation framework to fully address this problem. This theory results in new algorithms that beat the state of the art for applications such as exact

regret minimization in the well-known problem of prediction under expert advice.

Assessing the Impact Of Horizontal Mergers A-Priori For Both Regulators and Firms

In “Identifying Merger Opportunities: The Case of Air Traffic Control,” Adler, Olesen, and Volta propose a model to identify an optimal horizontal merger configuration at the level of an industry or firm with multiple branches. Assuming that each firm operates within a catchment area or owns part of a network, they extend the model to consider feasible mergers that cover a contiguous area, should network effects be a consideration. An application to the European air traffic control system suggests that four contiguous air navigation service providers should replace the current 29 providers and the nine functional airspace blocks proposed in the Single European Skies initiative. The technological developments in air traffic management in which regulators on both sides of the Atlantic have invested heavily, namely SESAR and NextGen, are unlikely to be used without a concomitant reduction in operating costs through economies of scale. They find that the politically oriented solution may save around one third of current costs, but an optimal solution will save closer to 46%.

Data-Driven Chance Constrained Programs over Wasserstein Balls

In the era of modern business analytics, data-driven optimization has emerged as a popular modeling paradigm to transform data into decisions. By constructing an ambiguity set of the potential data-generating distributions and subsequently hedging against all member distributions within this ambiguity set, data-driven optimization effectively combats the ambiguity with which real-life data sets are plagued. In “Technical Note—Data-Driven Chance Constrained Programs over Wasserstein Balls,” Chen, Kuhn, and Wiesemann study data-driven, chance-constrained programs in which a decision has to be feasible with high probability under every distribution within a Wasserstein ball centered at the empirical distribution. The authors show that the problem admits an exact deterministic reformulation as a mixed-integer conic program and demonstrate (in numerical experiments) that the reformulation compares favorably to several state-of-the-art data-driven optimization schemes.