



Operations Research

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

In This Issue

To cite this article:

(2025) In This Issue. Operations Research 73(5):iii-vii. <https://doi.org/10.1287/opre.2025.inthisissue.v73.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 © 2025, 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.2025.inthisissue.v73.n5>

Copyright: © 2025 INFORMS

How Online Retailers Can Fulfill Orders Faster and Cheaper in Structured RDC–FDC Networks

The global surge in online shopping has driven e-retailers to rapidly expand their warehouse networks for faster deliveries. However, this has significantly increased fulfillment costs. In “Multi-item Online Order Fulfillment in a Two-Layer Network,” Zhao, Wang, and Xin explore efficient ways to reduce these costs by examining how orders are allocated among warehouses in real time. They analyze a simple “myopic” algorithm, where each order is assigned immediately to the cheapest available warehouse without considering future orders. Surprisingly, this straightforward algorithm proves effective in structured warehouse networks consisting of larger regional distribution centers (RDCs) and smaller front distribution centers (FDCs), like those operated by *JD.com*. This contrasts sharply with more flexible fulfillment networks such as Amazon’s, where myopic algorithms usually underperform because of more complex fulfillment options requiring more advanced real-time algorithms. The key takeaway for businesses facing uncertain demand is that a less flexible fulfillment network—such as one consisting of large RDCs and smaller FDCs—enables simple and effective allocation methods.

Smarter Inventory Control Through Policy Mixing

Coordinating inventory across multiple warehouses and retail stores becomes significantly more complex when fixed ordering costs and lost sales are involved. In “Near-Optimal Mixed (s, S) Policy for a Multiwarehouse, Multistore Inventory System with Lost Sales and Fixed Cost,” Miao, Jasin, and Chao introduce a novel solution to this challenge. By leveraging Lagrangian relaxation, the authors develop a mixed (s, S) policy that assigns two carefully selected ordering strategies to each store: one for the early phase of the selling horizon and another for the later phase. This structured mix allows the policy to approximate the optimal solution, satisfying inventory constraints and managing costs effectively. The approach achieves a provable near-optimality guarantee and highlights the power of combining simple policies to tackle complex supply chain problems.

On the Almost Threshold Policy for Multisourcing Under Uncertain Supplies

In “On the Almost Threshold Policy for Multisourcing Under Uncertain Supplies,” Federgruen, Feng, and Shanthikumar examine optimal multisourcing strategies in the presence of supply uncertainty. The authors develop a general framework based on the concept of stochastic midpoint linearity and demonstrate that a simple almost threshold policy is optimal under broad and realistic conditions. Their findings provide both theoretical insights and practical guidance for designing resilient and data-driven multisourcing systems.

Political Districting to Optimize the Polsby-Popper Compactness Score with Application to Voting Rights

When drawing political districts, important criteria include district compactness and the ability of minority groups to convert

votes into seats. In litigation, compactness is usually measured via the Polsby-Popper score; however, this score has largely been absent from the operations research literature, presumably because of its nonlinear nature. In “Political Districting to Optimize the Polsby-Popper Compactness Score with Application to Voting Rights,” Belotti, Buchanan, and Ezazipour propose new mixed-integer second-order cone programs (MISOCPs) to optimize this score, enabling them to find optimally compact districts built from voting precincts. The methods are extended to find compact plans with many majority-minority districts. This is the task faced by plaintiffs in lawsuits brought under the Voting Rights Act of 1965 who must show that an alternative plan exists in which the minority group could achieve better representation. Computational experiments show that the MISOCP-based methods perform as well (or better) than state-of-the-art approaches for this task, scaling to instances with nearly 200,000 census blocks.

Crisis Increases Market Concentration in Oligopolistic Market

In “A Stationary Mean-Field Equilibrium Model of Irreversible Investment in a Two-Regime Economy,” Aid, Basei, and Ferrari study firms size distribution in a mean-field model of Cournot competition in a commodity market, where price follows an inverse power demand function. Firms face irreversible investment decisions and constant depreciation of production capacity. Output is affected by Gaussian productivity shocks, whose volatility and the price function can shift due to rare macroeconomic events modeled by a two-state Markov chain. Firms aim to maximize expected discounted profits, net of investment and operating costs, based on the long-run stationary price. We establish existence and uniqueness of a stationary mean-field equilibrium and characterize it through a barrier-type investment strategy with endogenous thresholds for each economic regime. A quasi-closed form for the stationary distribution of firms’ states is provided. The model generates Pareto-distributed firm sizes, consistent with empirical industry data. It also shows that downturns raise market concentration and that firm performance depends on depreciation rates and the persistence of economic fluctuations.

Routing a Vehicle to Collect Data After an Earthquake

In the immediate aftermath of a major earthquake, it is crucial to quickly and accurately assess structural damage throughout the region. It is especially important to identify buildings that have become unsafe in order to prioritize evacuation efforts. Only a very small number of building inspections can be feasibly performed in a narrow time frame; however, their results can then be combined with other data sources to predict damage at other locations that were not inspected. In “D-Optimal Orienteering for Post-earthquake Reconnaissance Planning,” Wang, Xie, Ryzhov, Marković, and Ou present a novel nonlinear integer program that combines vehicle routing with a statistical objective, the goal being to maximize data quality. An exact method based on row and column generation is developed to solve problems with up to 200

buildings. The approach is validated in a realistic case study using real-world building data obtained from a state-of-the-art earthquake simulator.

Optimal Abort Policy for Mission-Critical Systems Under Imperfect Condition Monitoring

Controlling stochastic systems often relies on the assumption of Markovian dynamics. However, this assumption frequently breaks down in mission-critical systems subject to failures—such as drones for power grid inspections—where the system failure rate increases over time. To enhance system survivability, operators may choose to abort missions based on noisy condition-monitoring signals. Yet, determining the optimal abort time in such settings leads to an intractable stopping problem under partial observability and non-Markovian behavior. In “Optimal Abort Policy for Mission-Critical Systems Under Imperfect Condition Monitoring,” Sun, Hu, and Ye introduce a novel Erlang mixture-based approximation that transforms the original non-Markovian process into continuous-time Markov chains. This approximation enables the formulation of partially observable Markov decision processes (POMDPs), whose optimal policies are shown to converge almost surely to the original optimal abort decision rules as the Erlang rate increases. Structural properties of the optimal POMDP policy are established, and a modified point-based value iteration algorithm is proposed to numerically solve the POMDP.

Unveiling Hidden Shifts: Detecting Change Points in Dynamic Networks with Missing Links

Dynamic networks are ubiquitous in a world increasingly driven by interconnected systems from social media platforms to financial transactions, transportation systems, and cybersecurity infrastructures. Unlike static entities, these networks evolve, often undergoing abrupt structural changes. Detecting these changes is essential as they can signal critical events, such as cyberattacks or shifts in user behavior. The challenge becomes even greater when dealing with incomplete data, a common reality in large-scale networks. Missing links, from nonresponses or privacy concerns, create uncertainty, complicating the identification of structural changes. In “Change-Point Detection in Dynamic Networks with Missing Links” Enikeeva and Klopp tackle this issue with a robust framework based on the matrix CUSUM test statistic. This novel methodology is adaptable to expanding networks and resilient to missing data. It detects structural shifts with minimax optimality, ensuring theoretical robustness. Extensive simulations and real-world examples reveal that the proposed approach outperforms existing methods. It provides a valuable tool for monitoring dynamic networks, detecting fraud, and spotting emerging trends.

Better Regularization for Sequential Decision Spaces: Fast Convergence Rates for Nash, Correlated, and Team Equilibria

In “Better Regularization for Sequential Decision Spaces: Fast Convergence Rates for Nash, Correlated, and Team Equilibria,” Farina, Kroer, and Sandholm study the application of first-order methods to the problem of computing equilibria of large-scale extensive-form games. It introduces a new weighted entropy-based distance-generating function for instantiating first-order methods. The new function achieves significantly better strong-convexity properties than existing weight schemes for the dilated entropy while maintaining the same easily implemented closed-form proximal mapping as the prior state of the art. The authors then generalize their new entropy distance function, as well as the whole class of dilated distance functions, to the scaled extension

operator. This yields the first efficiently computable distance-generating function for the decision polytopes capturing correlated and team solution concepts for extensive-form games. By instantiating first-order methods with these regularizers, several new results are achieved, such as the first method for computing ex ante correlated team equilibria with a guaranteed $1/T$ rate of convergence and efficient proximal updates.

Online Learning with Sample Selection Bias

Personalized recommendation systems often face the challenge of making optimal decisions when user preferences are unknown, and outcomes are only observed if users engage with the platform (e.g., clicking a recommendation). In “Online Learning with Sample Selection Bias,” Singhvi and Singhvi study this problem in the context of sequential decision making, where the censoring of outcomes leads to selection bias. Ignoring this bias results in suboptimal recommendations and linear regret, even for well-performing existing learning algorithms. To address this, they propose the sample selection bandit (SSB) algorithm, which combines Heckman’s two-step estimator with the “optimism under uncertainty” principle. The authors also demonstrate that SSB achieves a near-optimal regret rate. Extensive numerical experiments using synthetic and real-world donation data confirm that SSB significantly outperforms existing algorithms, effectively addressing selection bias while improving recommendations and outcomes in practical settings.

Injecting Little Variations to Create Large Gains: Rethinking Online Retail Pricing

In today’s fast-changing markets, businesses rely on precise, timely causal insights to guide effective operational decisions. In “Instrumenting While Experimenting: An Empirical Method for Competitive Pricing at Scale,” Jiang and Li introduce an idea of “instrumenting while experimenting.” By introducing small, random variations into routine decision-making processes, they generate causal insights on the impact of business decisions across different competitive dynamics, enabling large-scale operational improvements without disrupting daily business operations. Partnering with a major U.S. e-commerce retailer, the researchers apply this approach to develop a competitive pricing method that enhances response accuracy to competitors’ price changes at scale. Implemented on over 10,000 products, the method delivers significant gains—boosting revenue by more than 15% and profit by over 10%. These gains stem from the successful implementation of the instrumenting while experimenting idea as well as the joint use of experimentation, causal inference, and optimization tools. Moreover, this framework and its implementation provide deeper substantive insights into competitive response strategies and offer a promising blueprint for businesses navigating dynamic online markets.

Balancing Flexibility and Performance in Online Resource Allocation

How do firms optimize resource allocation strategies when frequent adjustments are costly or restricted? In “Blind Network Revenue Management and Bandits with Knapsacks Under Limited Switches,” Simchi-Levi, Xu, and Zhao explore this challenge. The authors investigate the impact of a switching constraint, which limits the number of times a firm can adjust allocations, on dynamic decision making, demand learning, and resource management. By establishing matching upper and lower regret bounds, the authors show how the statistical complexity of online learning changes when both resource and switching constraints are present. Their findings reveal that the optimal regret rate

follows a piecewise-constant function of the switching budget, providing key insights into algorithmic design for constrained decision making. The study's simulations demonstrate that firms can maintain strong performance and significantly reduce adjustments, offering practical implications for industries with operational rigidity.

How to Allocate Limited Resources When People Compete for Them?

Data-driven decision rules are increasingly used to allocate limited resources. However, in domains such as college admissions and job hiring, these policies can be gamed and may perform poorly when people can strategically change their behavior and compete for the resource. Given a fixed selection criterion for scoring individuals, competition causes the threshold for receiving the resource to oscillate and change over time. As a concrete example, the median SAT score of an accepted student at a university can evolve over time because of competition. In "Policy Learning with Competing Agents," Stanford researchers Sahoo and Wager demonstrate that this process can stabilize and reach an equilibrium. They develop an algorithm based on a local experimentation scheme to learn data-driven decision rules that maximize social welfare in the presence of competition.

On Consistency of Signature Using Lasso

The signature transform is a powerful tool for feature extraction in time series analysis, yet its statistical properties remain underexplored. In "On Consistency of Signature Using Lasso," Guo, Wang, Zhang, and Zhao examine the consistency of signature when used with Lasso regression. The study establishes conditions under which Lasso achieves consistent feature selection, both asymptotically and in finite samples. A key finding is that the Itô signature performs better for processes resembling Brownian motion, whereas the Stratonovich signature is more effective for mean-reverting processes. The authors further demonstrate the practical relevance of their results by applying signature-based Lasso regression to option pricing, highlighting its potential in financial modeling. These insights provide a crucial theoretical foundation for improving predictive performance across various domains, including operations research and machine learning.

Markdown Policies for Demand Learning with Forward-Looking Customers

Demand uncertainty and forward-looking customer behavior pose substantial challenges for sellers. In "Markdown Policies for Demand Learning with Forward-Looking Customers," Birge, Chen, and Keskin analyze a markdown pricing problem involving demand model uncertainty and strategic customers. The authors identify that strategic customer behavior creates a strong intertemporal dependence, where early markdowns influence later demand outcomes. They characterize the impact of this intertemporal dependence on demand learning and develop near-optimal policies that judiciously delay markdowns to manage forward-looking customer behavior while ensuring efficient learning.

Technical Note—Performance of a Queueing System with Scheduled Arrivals

A scheduled arrival sequence is one in which customers are scheduled to arrive at constant interarrival times, but each customer's actual arrival time is perturbed from her scheduled arrival time by a random perturbation. In "Technical Note—Stability of a Queue Fed by Scheduled Traffic at Critical Loading," Araman and Glynn consider a single server queue with deterministic

service times in which customers arrive following a scheduled arrival process. Unlike a queue fed by renewal traffic, this queue is shown to be stable even when the utilization is equal to one. It is also shown that for finite mean perturbations, a necessary and sufficient condition for stability is when the positive part of the perturbation has bounded support, with no requirement on the negative part of the perturbation. Perhaps surprisingly, this criterion is not reversible, in the sense that such a queue can be stable for a scheduled traffic process in forward time, but unstable for the time-reversal of the same traffic process.

Search in the Dark

In "Search in the Dark: The Case with Recall and Gaussian Learning," Baucells and Zorc address the classic sequential search problem where decision makers sample from a distribution to maximize rewards minus sampling costs. They focus on search with recall, where the reward is the highest sampled value. They solve the long-standing problem of determining optimal stopping rules when the sampling distribution is normal but both its mean and variance are unknown, and hence, they are progressively learned via sampling. The solution reveals that traditional methods, which rely on single reservation prices, are inadequate. The proposed approach offers a practical and efficient way to determine the optimal stopping rule. Relevant applications include job search, the sale of an asset, or technology adoption.

Accelerating Benders Decomposition via Deepest Cuts

In the global economy, billions of dollars of merchandise are routed using software that, at its core, uses optimization technology. Over many decades, researchers have devised different approaches to make algorithms faster, and this is true for Benders decomposition as well. Benders speeds up finding an optimal solution to a problem with millions of variables and constraints by iteratively learning which constraints are important and considering only these constraints. In "Deepest Cuts for Benders Decomposition," Hosseini and Turner selectively choose the constraints that eliminate the largest number of irrelevant solutions at each step would lead to finding the optimal solution in the fewest number of Benders steps. Geometrically, this amounts to choosing so-called deep cuts. Of course, in attempting to minimize the number of steps, we do need to spend more time taking each individual step, but our experimental results on several types of problems arising in supply chain analytics show that this approach makes sense and significantly reduces the solution time.

Capitalizing on the Relationships Between Tree Ensembles and Multilinear Functions

Tree ensembles are machine learning models used for regression and classification that combine the predictions of multiple trees. When such trained models are embedded into optimization models in the form of constraints or objectives, a key question is that of deriving best integer programming formulations for them. In "A Reciprocity Between Tree Ensemble Optimization and Multilinear Optimization," Kim, Richard, and Tawarmalani establish a polynomial-size reduction between the optimization of functions expressed as tree ensembles and the optimization of multilinear functions over a Cartesian product of simplices. This bidirectional reduction permits the derivation of new stronger formulations for tree ensemble optimization problems, including ideal formulations for single trees. It also provides a new framework for the construction of polynomially-sized convex hull descriptions for certain multilinear sets, which permits the generalization of many results from the literature.

Technical Note—Dynamic Duopolistic Competition with Sticky Prices

In “Technical Note—Dynamic Duopolistic Competition with Sticky Prices,” Heston and Hu resolve a puzzle about industrial competition. Previous papers analyze the effect of frictions on competition in continuous time, such as sticky prices or transactions costs. Those papers conclude that miniscule frictions have a large impact on market prices and quantities. However, Heston and Hu show that these conclusions are an artifact of the continuous-time math. Heston and Hu use discrete-time models to analyze the effect of frictions that are small per unit of time. These models show the effect of small frictions is small.

Data-Driven Clustering and Feature-Based Retail Electricity Pricing with Smart Meters

The adoption of smart meters and dynamic pricing programs is rapidly increasing among electric utility companies. In “Data-Driven Clustering and Feature-Based Retail Electricity Pricing with Smart Meters,” Keskin, Li, and Sunar analyze how utility companies should use smart meter data for better pricing decisions. Utility companies typically have access to consumption patterns and high-dimensional features on customer characteristics and exogenous factors. The authors identify that such feature data can exhibit different forms of heterogeneity—over time and over customers. They show that the different forms of feature heterogeneity significantly worsen the best profit performance that can be achieved by a data-driven dynamic pricing policy. The authors also develop a policy based on joint spectral clustering and contextual dynamic pricing and prove that this policy achieves near-optimal profit performance.

Enhancing the Efficiency and Accuracy of Inverse Optimization

Inverse optimization (IO) is used to model the behavior of decision-making agents who solve optimization problems in response to external signals. Inspired by the geometry of IO problems, in “Learning in Inverse Optimization: Incenter Cost, Augmented Suboptimality Loss, and Algorithms,” Zattoni Scroccaro, Atasoy, and Mohajerin Esfahani propose the “incenter” concept to solve IO problems, which unlike previously proposed approaches, can be used to derive computationally tractable solutions to this modeling problem. Moreover, they also propose a novel loss function for IO problems and a tailored optimization algorithm to optimize it. Extensive numerical experiments showcase the improved efficiency and accuracy of the proposed IO formulations and algorithm.

Novel Stochastic Gradient Estimators for Black-Box Optimization under Distributional Constraints

Constructing stochastic gradient estimators and descent algorithms is a core problem in stochastic optimization. In the black-box setting, namely when only noisy function evaluations are available, such gradient estimators are typically constructed via finite-differencing. In this regard, there have been schemes that aim to obtain gradient estimators for potentially many dimensions simultaneously via only few sample observations or simulation runs. However, for problems with probability simplex constraints, which arise in a range of applications from distributionally robust analysis to inverse model calibration, these schemes run into challenges one way or another when attempting to balance bias-variance in a constraint-compatible manner. In “Distributionally Constrained Black-Box Stochastic Gradient Estimation and Optimization,” Lam and Zhang were motivated to create a new design framework for random perturbation generators and estimation schemes that bypass these challenges. The

authors’ culminated class of estimators, which is based on the Dirichlet mixtures, is demonstrably effective in distributionally constrained gradient estimation and optimization under various black-box settings.

Context-Sensitive Simulation-Based Decisions When There Is No Time to Simulate

Stochastic simulation is a powerful tool for discovering system design decisions that are the best possible (optimal) when averaged over real-world uncertainty. However, in applications such as personalized medicine and web content optimization, even better decisions can be made if they are tailored to specific, contemporaneous covariate information, such as patient health history and user reading habits. Unfortunately, in these and similar applications, there is no time to perform a refined simulation optimization. In “Ranking and Contextual Selection,” Keslin, Nelson, Pagnoncelli, Plumlee, and Rahimian use off-the-shelf simulation optimization methods to create a database of covariates and associated decisions that form a covariate-to-decision classifier and an upper confidence bound on its optimality gap when applied to covariates not in the database. A realistic example of web page assortment optimization is presented using a data set from Yahoo!.

Refining Robust Decision-Making Integrated with Predictions

In “The Analytics of Robust Satisficing: Predict, Optimize, Satisfice, Then Fortify,” Sim, Tang, Zhou, and Zhu introduce a novel approach to decision making under uncertainty. Their method, termed “estimation-fortified robust satisficing,” leverages advanced predictive and prescriptive analytics to optimize decisions where traditional models falter due to risk ambiguity and estimation uncertainties. This approach not only enhances the resilience of decisions against unforeseen variations but also consistently outperforms conventional predictive methods in scenarios characterized by sparse data. This significant advancement promises to fortify decision-making processes in critical sectors such as finance and operations management, offering a new paradigm in handling the inherent uncertainties of real-world systems.

Computing Nash Equilibria as Pareto Points

Recent work by Feinstein and Rudloff [Feinstein Z, Rudloff B (2024) Technical Note—Characterizing and computing the set of Nash equilibria via vector optimization. *Oper. Res.* 72(5):2082–2096] proved 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. In “Approximating the Set of Nash Equilibria for Convex Games,” Feinstein, Hey, and Rudloff expand on this result to demonstrate that a comparable relation holds between the set of all approximate Nash equilibria and approximate Pareto solutions of a specific vector optimization problem. This characterization holds for all noncooperative games but opens a new way of computing Nash equilibria using techniques and algorithms from convex vector optimization. A sandwich property is proven in which the computed set of approximate Pareto solutions is bounded between the set of true Nash equilibria and the set of all approximate Nash equilibria with controlled error.

Near-Optimal Adaptive Policies for Serving Stochastically Departing Customers

In “Near-Optimal Adaptive Policies for Serving Stochastically Departing Customers,” Segev considers a multistage stochastic optimization problem originally introduced by Cygan et al.

[Cygan M, Englert M, Gupta A, Mucha M, Sankowski P (2013) Catch them if you can: How to serve impatient users. *Proc. 4th Innovations Theoretical Comput. Sci. Conf.*, 485–494], studying how a single server should prioritize stochastically departing customers. In this setting, the objective is to determine an adaptive service policy that maximizes the expected total reward collected along a discrete planning horizon, in the presence of customers who are independently departing between one stage and the next with known stationary probabilities. The paper's main contribution resides in proposing a quasi-polynomial-time approximation scheme for serving impatient customers. Specifically, letting n be the number of underlying customers, our algorithm identifies in $O(n^{O(\log^2 n)})$ time a service policy whose expected reward is within factor $-\epsilon$ of the optimal adaptive reward. The method for deriving this approximation scheme synthesizes various stochastic analyses in order to investigate how the adaptive optimum is affected by alterations to several instance parameters, including the reward values, the departure probabilities, and the collection of customers itself.

Advancing Risk Assessment: New Ways To Compute Quantile Aggregation

"Convolution Bounds on Quantile Aggregation" by Blanchet, Lam, Liu, and Wang is a pivotal study on quantile aggregation amid dependence uncertainty, an area critical to finance, risk management, and statistics. The authors introduce "convolution bounds," derived from a recent inf-convolution formula of quantiles and related risk measures. The obtained analytical tools unify existing results and enhance the understanding of quantile methods by providing general, sharp, and computationally efficient solutions. The results offer insights into the extremal dependence structures, with several implications in risk management and economic analysis applications.

Robust Stability in Multiclass Queueing Networks: A New Approach

In "A Hierarchical Approach to Robust Stability of Multiclass Queueing Networks," Zhao, Gurvich, and Hasenbein introduce a framework to identify sufficient conditions under which a network's stability is robust to the distributed choices of resources about their (local) prioritization of jobs. The framework produces sufficient conditions for such stability by relating it to robust optimization problems where the collection of priority policies plays the role of the uncertainty set. Interestingly, within the studied family of policies, robust stability for any policy is inherited from the stability of the special "corner" policies, which are none other than simple static-priority policies.

Robustifying Conditional Portfolio Decisions via Optimal Transport

How can portfolio managers make reliable investment choices when financial data are noisy and market conditions shift unpredictably? In "Robustifying Conditional Portfolio Decisions via

Optimal Transport," Nguyen, Zhang, Wang, Blanchet, Delage, and Ye introduce a novel data-driven framework that integrates side information, conditional estimation, and distributional robustness into portfolio optimization. Leveraging optimal transport-based ambiguity sets, the authors reformulate the distributionally robust conditional portfolio allocation problem into tractable convex programs. Their approach not only unifies prediction and decision-making but also avoids restrictive parametric assumptions, offering flexibility for both mean-variance and mean-CVaR criteria. Extensive empirical tests on U.S. equity markets demonstrate improved out-of-sample performance compared to benchmark models, highlighting the framework's ability to mitigate estimation errors while adapting to contextual information. Beyond finance, the methodology has potential applications in supply chains, energy planning, and other decision-making domains where robustness under uncertainty is critical. This contribution advances both the theory and practice of robust optimization for conditional decisions.

Breakthrough in Measuring Efficiency with Respect to Nonconvex Technology

The variable returns to scale (VRS) frontier, commonly used in data envelopment analysis, has a convex technology set and follows a specific "regular VRS" structure in economics. This structure demonstrates increasing, constant, and then, decreasing returns to scale, and it is the standard in data envelopment analysis. When the convexity assumption is relaxed, modeling regular VRS becomes challenging, and no satisfactory solution currently exists for multioutput production. In "Regular Variable Returns to Scale Production Frontier and Efficiency Measurement," Li, Tsang, Lee, and He introduce a new framework for analyzing regular variable returns to scale and propose an innovative empirical production frontier. This new frontier can more accurately measure technical efficiency without the convexity assumption. The implications of this research are extensive, impacting fields like manufacturing, agriculture, healthcare, banking, etc., with crucial findings for informed decision making and effective policy implementation.

Smarter Network Control with Less Communication

Managing networks often involves extensive communication between all parts of the system, which can be inefficient and raise privacy concerns. In "A Robust Optimization Approach to Network Control Using Local Information Exchange," Darivianakis, Georghiou, Shafiee, and Lygeros explore an alternative approach where each part of the system only communicates with its closest neighbors. By simplifying the communication structure, the system can operate with less computational effort while maintaining privacy. Applied to areas like energy management and supply chains, this method has the potential to provide similar outcomes to traditional models but with fewer communication requirements. The results suggest that this approach could offer a more efficient and privacy-conscious way to manage networks.