



## 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(4):iv–viii. <https://doi.org/10.1287/opre.2025.inthisissue.v73.n4>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact [permissions@informs.org](mailto: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.n4>

Copyright: © 2025 INFORMS

### Dual-Sourcing Inventory Management Under Incomplete Information

The class of dual-sourcing problems is one of the most fundamental challenges in inventory and supply chain management and is a common practice in global supply chain management. Under a dual-sourcing strategy, a firm procures items from a regular source with lower cost and longer lead time and an expedited source with higher cost and shorter lead time. In “Tailored Base-Surge Policies in Dual-Sourcing Inventory Systems with Demand Learning,” Chen and Shi focus on Tailored Base-Surge (TBS) policies, where a constant order is placed with the regular source each period, and the expedited source follows an order-up-to rule. Although existing literature assumes that stochastic demand processes are known inputs to optimization models, decision makers often lack this information in practice and must learn it in real time to maximize expected revenue. The authors develop online learning algorithms for finding the optimal TBS policy, demonstrating strong theoretical and numerical performance.

### Building Flexibility into Energy System Dispatch

The growing use of renewable energy is forcing power system operators to grapple with increasing uncertainty and intermittency in their energy supplies and demands. In “Unit Commitment without Commitment: A Dynamic Programming Approach for Managing an Integrated Energy System Under Uncertainty,” Brown and Smith develop a dynamic programming (DP) framework for balancing supply and demand over time. The approach introduces stochastic Lagrange multipliers as surrogate energy prices, which reduces the system-wide DP to a collection of unit-specific DPs, where each unit is managed to maximize its expected profit over a long time horizon, given these uncertain prices. Real-time dispatch decisions are then made using the unit-specific value functions to capture the longer-term impacts of the dispatch decisions. Using data from the Duke Energy Carolinas and Progress systems, this new approach reduced operational costs by 2% in current systems and 4%–5% in example future scenarios with increased solar and storage capabilities. Strikingly, the proposed methods performed within 0.2%–0.3% of plans based on perfect foresight, across a wide variety of scenarios.

### Tactics to Counter Missile and Drone Threats

In “Hard and Soft Defense Against a Sequence of Aerial Threats,” Atkinson and Kress develop an analytical model for missile defense in which a Blue defender counters sequential attacks by Red threats using both hard defense (HD) interceptors and soft measures (SM), such as jamming and decoys. Their model focuses on a shoot-look-shoot tactic by which Blue fires salvos of interceptors and then assesses the battle damage to decide on further engagement. Their objective is to minimize two performance metrics: the number of threats that manage to penetrate the defenses (leakers) and the number of costly interceptors expended. By establishing efficient frontiers in a two-dimensional performance space, they identify effective firing sequences for various scenarios. Their primary contribution is the development

of a framework that integrates HD with SM, allowing for an examination of how SM affects interceptor allocation and missile defense tactics.

### Increased Use of Split-Liver Transplantation Can Relieve Organ Shortage and Bring More Equitable Outcomes for Patients

In “Split Liver Transplantation: An Analytical Decision Support Model,” Tang, Scheller-Wolf, Tayur, Perito, and Roberts introduce a decision support model for split liver transplantation (SLT). SLT is a procedure that can save two lives with one donated liver, thereby increasing the total benefit derived from the limited number of donated livers available. SLT may also improve equity by giving transplant candidates who are physically smaller, including children, increased access to liver transplants. However, SLT is rarely used in the United States. To help quantify the benefits of increased SLT utilization and provide decision support tools, this research presents a deceased-donor liver allocation model that incorporates both efficiency and fairness objectives. The authors formulate the liver waitlists as a multiqueue fluid system, incorporating specifics of donor-recipient size matching and patients’ dynamically changing health conditions. Leveraging a novel decomposition result, the study finds the exact optimal matching procedure, enabling policy makers to benchmark the performance of different allocation policies against the theoretical optimal. Numerical results, utilizing data from the Organ Procurement and Transplantation Network, show that increased utilization of SLT can significantly reduce patient deaths, increase total quality-adjusted life years, and improve fairness among different patient groups.

### Asymptotic Optimality of Nested Fulfillment in Assemble-to-Order Generalized $W$ Systems

Despite their prevalence and considerable benefits, assemble-to-order systems are notoriously difficult to analyze and manage. In “Joint Order Fulfillment and Inventory Management in Assemble-to-Order Generalized  $W$  Systems,” Yu, Zheng, and Chen develop a strategy under which the policy for demand fulfillment is static nested and that for inventory procurement is of the newsvendor type. They show that the strategy is asymptotically optimal in the high-demand regime for assemble-to-order generalized  $W$  systems, in which  $k$  products are assembled from a common component and  $k$  product-specific components.

### Optimizing Liner Shipping Network Design via Simultaneous Column-and-Row Generation

In “A Simultaneous Column-and-Row Generation Solution Method for Liner Shipping Network Design,” Xia, Xu, and Baldacci tackle a challenging liner shipping network design problem, which involves complex rotation structures and interdependent decisions on rotation design, fleet deployment, and cargo routing. The authors propose a new exact solution method based on a simultaneous column-and-row generation framework with novel

acceleration techniques, named the linear programming–based approach and postpricing phase. These techniques leverage dual information to avoid generating unnecessary columns, speeding up the convergence of the solution method. The effectiveness of the proposed method is demonstrated for two variants of the liner shipping network design problem. Beyond maritime transportation, this work contributes to the broader field of mathematical programming by introducing an adaptable solution framework and speedup techniques for solving large-scale integer linear programs with column-dependent rows.

### On Simple Mechanisms for Dependent Items

In “On Simple Mechanisms for Dependent Items,” Cai and Oikonomou study the performance of simple mechanisms—such as item pricing or bundling—in settings where a buyer’s valuations for items are dependent. Previous research has largely focused on the setting with independent item valuations, which, although theoretically appealing, is unrealistic. Other studies have either examined limited forms of dependency or allowed arbitrary dependencies, but in the latter case, the resulting approximation ratios are exponential in the number of items, making them too weak to be practical for real-world applications. To bridge this gap, the authors use Markov random fields (MRFs), a graphical model well suited for representing complex dependencies among random variables, to capture correlations among items. They establish a parametric approximation, showing that the performance of simple mechanisms is independent of the number of items and degrades gracefully with the degree of dependency among items, as characterized by the MRF. This framework extends to general valuation functions, enabling it to account for practical constraints faced by practitioners, such as cardinality constraints.

### Fair and Efficient Allocation Algorithms Do Not Require Knowing Exact Item Values

Food rescue organizations are tasked with allocating often-unpredictable donations to recipients who need it. For a large class of recipient valuation functions, this can be done in a fair and efficient manner as long as each recipient reports their value for each arriving donation. In practice, however, such valuations are rarely elicited. In “Dynamic Fair Division with Partial Information,” Benadè, Halpern, and Psomas ask whether simultaneous fairness and efficiency remain possible when the allocator receives limited information about recipient valuations, even as little as a single binary signal. For recipients with i.i.d. or correlated values, the paper provides an algorithm which is envy-free and  $(1 - \epsilon)$ -welfare-maximizing with high probability. Asymptotically tight results are also established for independent, non-identical agents. This shows that fair and efficient online allocation algorithms do not critically rely on recipients being able to precisely report their utility functions.

### Online Allocation of Reusable Resources: New Algorithms and Analytical Tools

In “Asymptotically Optimal Competitive Ratio for Online Allocation of Reusable Resources,” Goyal, Iyengar, and Udwani develop novel algorithms and analysis techniques for online allocation of reusable resources. Their approach leads to an algorithm with the highest possible competitive ratio, a result that was previously out of reach with the algorithms and techniques that are used in classic settings in which resources are nonreusable. More generally, their linear programming–free analysis approach is useful for analyzing the performance of online algorithms for various other settings in which the standard primal-dual approach fails.

### Community-Engaged School District Design: A Stream-Based Approach

School district design decisions are a crucial part of public education systems, and operations research has played an important role in informing such decisions since the landmark United States Supreme Court decision on school desegregation in 1954. Design processes today place significant importance on comprehensive community engagement to create equitable and effective enrollment policies reflective of community needs and values. In “Community-Engaged School District Design: A Stream-Based Approach,” Ozel, Smilowitz, and Goldstein revisit the school district design problem with a focus on codesigning with community partners. The authors introduce a new compact formulation that makes multiple interrelated assignment decisions simultaneously with composite decision variables, called “streams.” This formulation is computationally efficient and easily reconfigurable for evolving problem specifications that are endemic to community codesign. These features were essential in the district redesign process in an Illinois public school district, described in the paper, allowing the community to iteratively develop proposals to address inequities in access to education and improve the student assignment process.

### Dynamic Pricing Without Statistical Assumptions on Covariates

In “Technical Note—On Dynamic Pricing with Covariates,” Wang, Talluri, and Li introduce novel methods for setting prices when key market factors, or covariates, vary over time. Although most existing literature assumes these factors follow a fixed distribution, the authors show such assumptions are unnecessary for achieving near-optimal pricing strategies. They provide new pricing algorithms with upper regret bounds that match lower regret bounds (up to a logarithmic factor)—even when covariates exhibit no fixed statistical pattern. In particular, the analysis underscores how these algorithms can efficiently adapt to changing market conditions, enabling businesses to continuously optimize their pricing strategies. In addition, the paper extends these results to a constrained scenario with limited, nonreplenishable inventory, unveiling how a variation metric quantifies the effect of temporal shifts in covariates on pricing performance. Numerical experiments validate the practicality of these algorithms. By relaxing common statistical assumptions, this research offers a more flexible approach for data-driven dynamic pricing solutions.

### Reliable Blockchain Transaction Fee Mechanisms with Revenue Guarantees

In “Bayesian Mechanism Design for Blockchain Transaction Fee Allocation,” Chen, Simchi-Levi, Zhao, and Zhou address the critical challenge of designing reliable blockchain transaction fee mechanisms (TFMs) that ensure truthfulness, collusion-proofness, and desirable miner revenues. By proposing an innovative “auxiliary mechanism method,” the authors establish connections between Bayesian Nash Incentive Compatible (BNIC) and Dominant Strategy Incentive Compatible (DSIC) mechanisms through an “auxiliary-variation decomposition.” This breakthrough allows the proposed mechanism to bypass previous limitations in TFM design via the Bayesian game setting, achieving constant-ratio approximations of optimal miner revenues while providing strong incentive guarantees. The research not only advances the theoretical understanding of blockchain fee mechanisms but also demonstrates the potential versatility of the auxiliary mechanism method for designing and optimizing BNIC mechanisms across broader applications. This work represents a significant contribution to both blockchain economics and mechanism

design, offering a pathway toward more reliable and efficient blockchain ecosystems.

### The Power of Randomization and Prior Information in Clock Auctions

The recently introduced and highly practical class of deferred acceptance clock auctions offers a unique combination of desirable properties, including obvious strategy-proofness, unconditional winner privacy, and transparency. Early excitement around these auctions was dampened due to a result demonstrating that no deterministic and prior-free clock auction can guarantee better than a logarithmic approximation of the optimal social welfare, even in simple settings. In “Bayesian and Randomized Clock Auctions,” Feldman, Gkatzelis, Gravin, and Schoepflin reignite this excitement by demonstrating that one can achieve improved results for a wide class of instances if a clock auction can either use randomization or has access to prior information. The authors provide three different clock auctions that achieve these improved guarantees given access to different amounts of prior information, exhibiting how additional information can further simplify the auction design.

### Design and Pricing of a Single Bundle

Product bundling is a widely used selling strategy among multi-product firms, yet designing and pricing bundles optimally remain a complex challenge. In “Partition and Prosper: Design and Pricing of a Single Bundle,” Sun, Li, and Teo show that the selection and pricing of a single bundle from a range of products under multivariate normal valuations is polynomial-time solvable, provided that the associated covariance matrix can be decomposed into a positive diagonal matrix minus a positive semidefinite matrix of (small) fixed rank. Interestingly, they also show that if the individual product prices are predetermined, the optimization problem becomes  $\mathcal{NP}$ -hard even if customer valuations are independent. The authors use a Bayesian optimization algorithm combined with a novel conic integer programming reformulation to solve the problem under general valuation distributions and demonstrate the superior numerical performance of the algorithm via extensive simulations.

### Stronger Connections and Uncovered Nuance for Queues with Batch Arrivals

Many modern queueing systems face the challenge that arrivals occur in batches of customers or jobs. Such simultaneity in the arrival stream may place a particularly acute stress on the service system. Recently, researchers have identified that shot-noise processes, a family of stochastic models for which the dynamics are deterministic between arrival epochs, can be used to approximate batch-arrival queues as justified through batch scaling limits. In “Establishing Convergence of Infinite-Server Queues with Batch Arrivals to Shot-Noise Processes,” Daw, Fralix, and Pender both elevate and generalize these limits, moving from point-wise convergence in distribution for Markovian models to almost sure convergence on a process level (uniformly on compact sets) for general arrival and service processes, allowing for possible non-stationarity and dependence between arrivals. Additionally, the authors establish these batch-scaling limits for both the queue-length and workload processes, and this leads to an unexpected insight. By comparison with the Markovian case, in which the queue-length and workload limits are proportionally identical, for general service distributions, the batch-scaled queue-length and workload processes can differ significantly, revealing that the workload can persist even as the queue length dwindles.

### Innovative Car Relocation Policy Enhances Efficiency in Car-Sharing Networks

In “Dynamic Relocations in Car-Sharing Networks,” Hosseini, Milner, and Romero develop a new dynamic car relocation policy aimed at improving the efficiency of car-sharing networks. The policy addresses the challenge of unpredictable and uneven demand across the network by projecting the classical fluid model onto the simpler dimension of relocation decisions. This projection leads to a clearer understanding of the fluid model as a set of linear programs. The effectiveness of this new dynamic policy has been rigorously tested against traditional static policies through extensive simulations. The results are promising, showing that the dynamic policy significantly outperforms its static counterpart, reducing the optimality gap by more than 23% on average. This breakthrough offers substantial potential benefits for car-sharing operations, enhancing their ability to respond adaptively to changing demand and thus optimizing resource allocation and user satisfaction.

### The Role of Boundaries in the Spreading of Solar

Peer effects by neighbors play a key role in the spreading of residential solar. Thus, people are more likely to install a solar system on their roof if some of their neighbors have already done so. Because people who live near the municipality boundary have fewer neighbors, does this imply that they are less likely to adopt solar? In “Boundary Effects in the Diffusion of New Products on Cartesian Networks,” Fibich, Levin, and Gillingham analyze this problem analytically using the Bass model on two-dimensional networks and empirically using data on installations of solar systems. They show that boundaries have a significant impact on the adoption of residential units near the municipality boundary. Their effect on the aggregate adoption in the municipality, however, is negligible.

### How to Sell to a Customer Who Searches for Alternative Products

Customer searching behavior has been ubiquitous in modern marketplaces. Dynamic pricing, such as giving promotional prices to customers who purchase earlier, has been commonly adopted in practice to discourage customers from searching. In “Optimal Dynamic Mechanism Under Customer Search,” Hu and Xiao characterize that, in the optimal selling mechanism, the seller should offer a menu of American option contracts. Under such a mechanism, the customer is given all the freedom in conducting the search but must pay an up-front deposit from the chosen option contract in order to enjoy the right to purchase the seller’s product anytime during the search at a prespecified strike price. Through deposits, the seller is able to extract the customer’s additional surplus from searching regardless of whether the customer makes the final purchase or not. As a result, the seller can often obtain a much higher profit than dynamic pricing mechanisms that aim to deter searching.

### Optimal Clearing Control of Backlogs

Processing systems, such as make-to-order production or service systems, are often faced with backlogged demand, which results in a prolonged period in which the system is congested, although it has sufficient processing capacity to handle all newly arriving demand. In “Asymptotically Optimal Clearing Control of Backlogs in Multiclass Processing Systems,” Yu, Iravani, and Perry consider a processing system, modeled by a multiclass queueing model, that faces the problem of optimally clearing a large backlog from several classes of customers (or orders). For the special case of two classes, the authors prove that a static priority policy following a discounted  $c\mu/\theta$  rule is asymptotically optimal. When there are more than two classes of customers, the authors show

that any admissible control that follows the best-effort rule becomes asymptotically optimal after a relatively short time. An extensive numerical study shows that these proposed policies are effective and provides guidance on when to choose among the policies in practice.

### **Fairness and Its Effects on Online Resource Allocation**

Online resource allocation is challenging when the decision maker aims to achieve fairness, load balancing, diversity, and so on. In “Optimal Regularized Online Allocation by Adaptive Re-Solving,” Ma, Cao, Tsang, and Xia examine the online allocation problem equipped with a nonseparable regularizer, which is commonly adopted to promote fairness and beyond. Under certain regularity, smoothness, and nondegeneracy assumptions of the dual problem, the authors show that a dual-based optimization strategy with adaptively updated resource constraints and history-based re-solving can achieve optimal logarithmic regret, even without the notorious log-log factor. The authors also demonstrate that an optimal regret upper bound can be achieved using infrequent re-solving and that an even faster algorithm is available if nearly optimal regret performance is acceptable. Their results demonstrate that optimal logarithmic regret is achievable under the max-min fairness or load-balancing constraints.

### **How to Dynamically Allocate Limited Capacity to Service Requests?**

The problem studied in “Dynamic Allocation of Reusable Resources: Logarithmic Regret in Overloaded Networks,” by Xie, Gurvich, and Küçükyavuz, is common in service applications, such as hotels, car rentals, and consulting services. These applications have limited capacity that must be allocated among incoming service requests. Different requests may require resources for varying durations, and some requests might yield higher rewards than others when fulfilled. The decision maker, who controls this capacity, must decide upon each request’s arrival whether to accept it (thus committing resources for a certain period) or reject it in order to reserve capacity for potentially more valuable future requests. This paper expands our understanding of such resource allocation problems by developing an appealingly simple dynamic allocation algorithm with proven effectiveness. In an appropriate operating regime, the suboptimal gap of this algorithm is at most logarithmic in the maximum achievable reward. In other words, for large-scale systems, the suboptimality translates to a negligible percentage deviation from optimal performance.

### **Risk-Adaptive Local Decision Rules**

The decision rules developed in “Risk-Adaptive Local Decision Rules,” by Roysset and Lejeune, help managers recognize relationships between data and decisions and prescribe courses of actions that are guaranteed near-optimal as well as transparent and interpretable. Transparency and interpretability are prerequisites for any decision support tool that needs to be conveyed, understood, and justified broadly. An understanding of the relationship between data and decisions is important for managers as they identify the critical factors, prioritize their focus, guide the resource allocation, and develop mitigation strategies for handling data uncertainty. The relationship between data and decision can be especially convoluted for combinatorial optimization problems. The authors consider simple decisions rules that ensure transparency and interpretability, but also allow nearly any continuous decision rule within a unified framework of analysis. The preferred adaptive minimum decision rule prescribes decisions

with 1% relative suboptimality gap, on average, even when facing distributional shifts in out-of-sample testing.

### **Wasserstein DRO Revisited: A Short Yet General Duality Proof**

Wasserstein distributionally robust optimization has emerged as a recent topic with broader applications in operations research and machine learning. Various proofs have been presented in the literature, each differing in assumptions and levels of generality. In “A Short and General Duality Proof for Wasserstein Distributionally Robust Optimization,” Zhang, Yang, and Gao present a novel elementary proof that not only shortens existing frameworks but also offers surprising generalizations. Leveraging classical Legendre–Fenchel duality, they demonstrate that strong duality is contingent on a certain interchangeability principle. Moreover, they extend this duality result to encompass risk-averse optimization and globalized distributionally robust counterparts.

### **Ironclad Method Boosts Blockchain Consistency and Speeds up Block Production**

Researchers have developed a novel method called Ironclad to address the scalability limitations of blockchains based on the Nakamoto consensus protocol. Although these blockchains have shown promise in various applications, the consensus properties of the protocol impose inherent scalability constraints. A key property, known as consistency, poses a challenge by presenting a trade-off between block production speed and the system’s security against adversarial attacks. To overcome this, in “Improving Blockchain Consistency Bound by Assigning Weights to Random Blocks,” Gong, Q. Zhang, Li, and J. Zhang propose the Ironclad method, which introduces a unique approach by assigning different weights to randomly selected blocks. By applying the Ironclad method to the original Nakamoto protocol and conducting rigorous analysis of its consensus properties, the researchers demonstrate a significant improvement in the consistency bound. This advancement allows for a much faster block production rate compared with the original protocol while maintaining the same level of security guarantees. The Ironclad method presents a promising solution to enhance blockchain scalability, overcoming the limitations of the Nakamoto consensus protocol. This breakthrough has the potential to unlock new possibilities for blockchain applications, paving the way for improved transaction speeds and increased efficiency in various industries.

### **On the Difficulty of Pricing Routes for Stochastic Vehicle Routing Problems**

Many approaches exist for dealing with the uncertainty in the vehicle routing problem with stochastic demands (VRPSD), but the most popular approach models the VRPSD as a two-stage stochastic program where a recourse policy prescribes actions that handle when the realized demands exceed the vehicle capacity. Similarly to other VRP variants, some state-of-the-art algorithms for the VRPSD use set-partitioning formulations that generate variables (routes) via a pricing problem. All of these algorithms, however, have strong assumptions on the probability distribution of customer demands, a simplification that might not be realistic in some applications. In “Hardness of Pricing Routes for Two-Stage Stochastic Vehicle Routing Problems with Scenarios,” Ota and Fukasawa examine the challenges associated with solving the pricing problem of the VRPSD when the customer demands are given by scenarios. They demonstrate that the VRPSD pricing problem is strongly  $\mathcal{NP}$ -hard for a wide variety of recourse policies and route relaxations. This highlights the difficulty of

developing efficient pricing algorithms for the VRPSD with scenario-based demand models.

### Tight Bounds for the Regret in Online Allocation Problems

In 2019, Alessandro Arlotto and Itai Gurvich proved that the expected regret in the multisecretary problem is uniformly bounded as the number of jobs to fill goes to infinity. However, they modeled the secretary valuation distribution as discrete and known, and they concluded their article by explaining that they did not know whether their bound would hold with continuous or unknown valuation distributions. In “Logarithmic Regret in Multisecretary and Online Linear Programs with Continuous Valuations,” Bray provides an answer, furnishing lower and upper bounds that demonstrate that the regret in multisecretary problem grows like the logarithm of the number of positions to fill when the distribution of secretary valuations continuous or unknown. The author then develops the argument to establish analogous bounds for the more general online linear program, defined by Xiaocheng Li and Yinyu Ye. The author tightens Li and Ye’s upper regret bound and develops a corresponding lower regret bound. To create these bounds, the author develops several new convergence results for the online linear program’s dual variables.

### Nonadaptive Stochastic Score Classification

Sequential testing problems involve a system with several components, each of which is working with some independent probability. The working/failed status of each component can be determined by performing a test, which is usually expensive. So, the goal is to perform tests in a carefully chosen sequence until the overall system status can be evaluated. These problems arise in a variety of applications, such as healthcare, manufacturing, and telecommunication. A common task in these applications is to categorize the system into one of several classes that correspond to the system status being poor, fair, good, excellent, etc. In “Nonadaptive Stochastic Score Classification and Explainable Half-Space Evaluation,” Ghuge, Gupta, and Nagarajan provide the first constant-factor approximation algorithm for this problem. Moreover, the resulting policy is nonadaptive, which results in significant savings in computational time. The authors also validate their theoretical results via computational experiments, where they observe that their algorithm’s cost is on average at most 50% more than an information-theoretic lower bound.

### Strong Optimal Classification Trees

Decision trees are among the most interpretable and popular machine learning models that are used routinely in applications ranging from revenue management to medicine. Traditional heuristic methods, although fast, lack modeling flexibility for incorporating constraints such as fairness and do not guarantee optimality. Recent efforts aim to overcome these limitations using mixed-integer optimization (MIO) for better modeling flexibility and optimality, but speed remains an issue. In “Strong Optimal Classification Trees,” Aghaei, Gómez, and Vayanos use integer optimization and polyhedral theory to create an MIO-based formulation with a strong LO relaxation resulting in a 29% speed-up in training

time compared with state-of-the-art MIO-based formulations, as well as up to an 8% improvement in out-of-sample accuracy.

### Dynamic Control of Service Systems with Returns

Postdischarge interventions (e.g., follow-up phone calls, outpatient appointments) are commonly used to reduce unplanned hospital readmissions. Such interventions have been shown to be effective (to varying degrees) in reducing the probability of readmission (or return). The interventions are, however, costly, and hence, their benefits need to be carefully balanced with the costs. In “Dynamic Control of Service Systems with Returns: Application to Design of Postdischarge Hospital Readmission Prevention Programs,” Chan, Huang, and Sarhangian investigate how post-service interventions should be dynamically allocated, accounting for their costs as well as their benefits in terms of reducing returns and congestion costs. To this end, they study a transient queueing control problem and examine the structure of the optimal policy by analyzing associated fluid-control problems. The structural results motivate the design of intuitive surge protocols whereby different intensities of interventions are provided based on the congestion level of the system.

### Distributionally Robust Chance Constraints Made Convex

Chance constraints, under either known or ambiguous distributions, yield nonconvex feasible regions in general. In “Convex Chance-Constrained Programs with Wasserstein Ambiguity,” Shen and Jiang identify sufficient conditions that lead to *convex* feasible regions of chance constraints with Wasserstein ambiguity. Notably, it generalizes the seminal work of Prékopa, which established the convexity of joint chance constraints under log-concave probability distributions, to the distributionally robust and distributionally optimistic settings.

### New Insights into Off-line Estimation for Controlled Markov Chains Unveiled

A team of researchers from Purdue and Northwestern Universities have unveiled new findings in off-line estimation for controlled Markov chains, addressing challenges in analyzing complex data generated under arbitrary dynamics. In “Off-line Estimation of Controlled Markov Chains: Minimality and Sample Complexity,” Banerjee, Honnappa, and Rao introduce a nonparametric estimator for transition probabilities, showcasing its robustness even in nonstationary, non-Markovian environments. The team developed precise sample complexity bounds, revealing a delicate interplay between mixing properties of the logging policy and data set size. Their analysis highlights how achieving optimal statistical risk depends on this trade-off, broadening the scope of off-line estimation under diverse conditions. Examples include ergodic and weakly ergodic chains as well as controlled chains with episodic or greedy controls. Significantly, this research confirms that the widely used estimator, which calculates state-action transition ratios, is minimax optimal, ensuring its reliability in general scenarios. This advancement paves the way for improved evaluation of stationary Markov control policies, marking a breakthrough in understanding complex off-line systems.