



Operations Research

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

In This Issue

To cite this article:

(2026) In This Issue. Operations Research 74(1):iii-vii. <https://doi.org/10.1287/opre.2026.inthisissue.v74.n1>

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 © 2026, 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.2026.inthisissue.v74.n1>

Copyright: © 2026 INFORMS

A New Perspective on Resilience in Flexible Supply Systems

Supply disruptions and demand uncertainty are persistent threats in today's global operations landscape. In "Process Flexibility: A Distribution-Free Approach to Long Chain Resilience," Chen, Chou, and Sun revisit the long-standing concept of long chain flexibility and demonstrate its resilience in the face of supply-side risks. By deriving a closed-form bound on expected sales relative to full flexibility, the study offers a robust, distribution-free guarantee that highlights the long chain's ability to hedge against disruptions. The authors further generalize their findings using a novel moment decomposition approach, extending applicability to a broader range of service metrics and capacity-demand scenarios via semidefinite programming. Their results not only reaffirm the demand-pooling power of long chains but also position them as a highly resilient configuration in disrupted environments, offering practical insights for capacity planning in uncertain supply chains.

New Analytical Insights on Inventory Allocation

How should a limited inventory of a single resource be allocated across multiple customer groups with distinct rewards and arrival rates, especially when each group has a finite population? In "Optimal Allocation of Limited Inventory Among Multiclass Customers with Finite Populations," Ge, Kulkarni, and Swaminathan develop a stochastic framework and formulate the problem as a Markov decision process. By analyzing the structure of the optimal value function, they reveal a new insight; rather than gradually expanding access to lower-priority groups over time, it is in fact optimal to progressively restrict access to these groups. The authors also introduce a fluid model with an explicit solution, which provides a good approximation when the system size is large. These findings offer both theoretical and practical contributions to inventory allocation problems, with potential applications in healthcare and humanitarian resource management as well as commercial product sales.

Incorporating Expert Views into Investment Models

How can investors best incorporate forward-looking expert views into dynamic portfolio strategies? In "Dynamic Black-Litterman," Abdelhakmi and Lim offer a novel solution. The authors generalize the classic Black-Litterman model to a dynamic setting, allowing for continuous trading and expert views of events over varying time horizons that arrive over time. They derive the explicit dynamics of asset prices after incorporating these views, uncovering a surprising and elegant connection to multidimensional Brownian bridges. Remarkably, the paper provides a closed-form solution for the optimal dynamic portfolio policy, a significant finding for a model of this complexity. This includes a specific hedging component to protect against changes in expert views. Numerical experiments demonstrate that this dynamic approach consistently outperforms strategies that repeatedly apply

the single-period model, leading to higher returns and lower portfolio turnover.

Optimizing Package Bids in Day-Ahead Electricity Markets

Day-ahead electricity auctions allow market participants to trade power for delivery the following day. In Europe, these auctions are designed as combinatorial auctions, enabling agents to submit package bids ("block bids") that span multiple time periods rather than bidding separately for each hour. However, power exchanges impose limits on the number of package bids an agent can submit, creating a complex decision problem: Which packages should an agent bid on to best represent their preferences? In "Package Bids in Combinatorial Electricity Auctions: Selection, Welfare Losses, and Alternatives," Hübner and Hug study this selection problem and propose decision-support algorithms that optimize bid choice under uncertainty. They provide theoretical bounds on welfare loss due to bid limits and validate their methods with simulations involving generators, storage systems, and flexible demand. Their findings offer actionable insights for both auctioneers and bidders.

Tracking the Tracker: Optimal Execution Amid Endogenous Order Flow

How should a large investor trade when their actions influence—and are influenced by—others in the market? In "Optimal Trade Execution Under Endogenous Order Flow," Chen, Horst, and Tran investigate optimal execution in financial markets where order flow is endogenous and governed by self-exciting dynamics. Market order arrivals are modeled using a Hawkes process, with intensity shaped by the trader's own activity—capturing a feedback loop between execution and market response. The study considers both risk-neutral and risk-averse investors under market impact, deriving closed-form and semi-closed-form optimal strategies. In the risk-averse case, the solution skews execution toward earlier periods to mitigate inventory risk. The model is also extended to more general Hawkes kernels, enhancing practical applicability. The findings shed light on how sophisticated traders can minimize execution costs while accounting for the risk of being tracked in increasingly transparent markets. This work offers actionable insights for algorithmic execution under realistic microstructure dynamics.

Dynamic Portfolio Allocation Under Market Incompleteness and Wealth Effects

In "Dynamic Portfolio Allocation Under Market Incompleteness and Wealth Effects," Shen, Li, Scaillet, and Jiang develop a novel decomposition of an optimal dynamic portfolio under general incomplete-market models and the wealth-dependent hyperbolic absolute risk aversion (HARA) utility. The authors show that with hedgeable interest rate risk, the optimal portfolio consists of two parts: a pure constant relative risk aversion optimal portfolio and a financing bond portfolio for investor future subsistence

requirements. Under such a structure, the wealth growth rate is always higher for HARA investors with more initial wealth, leading to increased wealth inequality regardless of the underlying model dynamics and realized market scenario. Using the decomposition, the authors solve the HARA optimal policy in closed form under an incomplete-market model with both stochastic interest rate and volatility. The wealth effect in the optimal portfolio has interesting implications. It generates a procyclical pattern in investor stock positions and time-varying risk aversion levels as well as a “buy high, sell low” market timing effect that may hurt HARA investors with low initial wealth.

Incorporating Fairness into Online Price Discrimination

In “Technical Note—Fairness-Aware Online Price Discrimination with Nonparametric Demand Models,” Chen, Lyu, Zhang, and Zhou explore how fairness can be integrated into dynamic pricing strategies. The authors propose a model that enforces price fairness constraints, ensuring that price differences between customer groups remain within a specified range. Their approach introduces a novel regret lower bound, which contrasts with the typical regret seen in traditional pricing algorithms. This shift underscores the added complexity of optimizing revenue while maintaining fairness. The study not only advances the understanding of fairness-aware dynamic pricing but also enriches the dynamic pricing literature by offering new lower-bound techniques. These insights may be useful for deriving lower bounds in other problems related to learning optimal prices under constraints. Their work contributes significantly to the growing need for ethical pricing practices in data-driven markets.

New Algorithms for Online Metric Matching with Stochastic Arrivals or Predictions

How do we match riders to drivers? Online metric matching offers a clean abstraction of this problem, where arriving riders are instantaneously matched to waiting drivers and the matching cost is measured by the pickup distance. Its challenge lies in making matching decisions without knowing the locations of future riders, for which the simple greedy approach yields suboptimal solutions. In “Online Metric Matching: Beyond the Worst Case,” Yang and Yu propose a novel algorithmic framework of designing algorithms for online metric matching when given access to additional information of riders’ locations in advance. They then apply this framework to derive new algorithms when the riders’ locations are independently sampled or when an untrusted prediction of riders’ locations is provided. In the former model, their algorithms achieve improved competitive ratio and regret guarantees for various settings. In the latter model, they present an algorithm whose performance smoothly depends on the prediction error while preserving the worst-case guarantee.

Assortment Optimization When All Products Have the Same Unit Margin

Should the seller display all available products in a setting with multiple products or only offer a subset? Is there a benefit of withholding some products? The seller does not want to display low-margin products due to cannibalization, but what if all products have the same margin? In “Performance of the Offer-Everything Policy,” Huh, Paat, and Queyranne argue that in an environment where the seller cannot adjust assortment based on customer types, it is best to offer all available products in an environment. The analysis is based on competitive ratio and expected revenue.

This research assures that a common practice of truthfully revealing product availability is a generally good policy to follow.

Ranking Inferences for General Multiway Comparisons

Ranking inference underpins critical decision making across diverse domains including, e.g., university ranking, journal ranking, academic paper review, voting, online recommendation and tournament competition. Beyond generating point estimates of ranks, these applications also demand confidence intervals for ranks through robust uncertainty quantification, in order to ensure reliable and informed decisions. However, existing approaches predominantly rely on classical models such as Bradley-Terry-Luce and Plackett-Luce, which assume homogeneous comparison structures and prove inadequate for complex real-world scenarios. In “Spectral Ranking Inferences Based on General Multiway Comparisons,” Fan, Lou, Wang, and Yu present a unified spectral ranking methodology for heterogeneous multiway comparisons, which simultaneously achieves statistical efficiency under minimal structural assumptions and computational scalability. The authors establish comprehensive ranking inference tools, grounded in the asymptotic normality theory and bootstrapping techniques, facilitating top-K selection, rank confidence interval construction, and hypothesis testing for cross-population or cross-period ranking consistency.

Can Customers Spot Price Discrimination?

Firms adjust prices to match supply with demand and increasingly use customer data to personalize pricing. Consequently, customers have limited visibility into price drivers and may worry about price discrimination based on their shopping behavior. For instance, a customer with high willingness to pay may receive a high price either because the product is popular or because the price has been personalized. In “Is Your Price Personalized? Alleviating Customer Concerns with Inventory Availability Information,” Aflaki and Zhang study whether observed prices are sufficiently informative to signal personalization in a dynamic setting. They find that price alone cannot resolve this uncertainty, which may lead to actions that harm both firms and customers. However, a simple binary signal disclosing whether inventory is low or abundant can, under certain conditions, ease customer concerns and benefit both the firm and its customers.

Multilocation, Dynamic Staff Planning for a Healthcare System: Methodology and Application

In “Multilocation, Dynamic Staff Planning for a Healthcare System: Methodology and Application,” Rath, Rajaram, Hudson, and Mahajan develop and implement a robust optimization model for dynamically assigning anesthesiologists across multiple hospitals in a large health system. The model addresses uncertainty in surgical demand by incorporating a three-stage decision process: pre-surgical location assignments, midhorizon on-call deployment, and day-of realization of overtime or idleness. The authors formulate the problem as a multistage robust mixed-integer program and solve it efficiently using a novel nested column and constraint generation algorithm. The approach accounts for location-specific constraints and fairness in on-call duties using historical data to estimate demand uncertainty with a calibrated trade-off between optimality and robustness. Implemented at the University of Pittsburgh Medical Center, the model reduced annual staffing costs by 12% or roughly \$800,000 compared to existing practice. The framework is generalizable to other staffing problems in healthcare and provides operational insights on the value of

forecast accuracy, location flexibility, and equitable workload distribution.

New Algorithm Boosts Online Retail Decisions by Learning Product Positioning

In “Rate-Optimal Online Learning for Dynamic Assortment Selection with Positioning,” Luo, Sun, and Liu address a key challenge in online retail: product positioning. The authors propose a novel online learning framework called dynamic assortment selection with positioning (DAP). Unlike traditional models that focus solely on item selection, DAP also learns optimal product placement to maximize revenue. The researchers model customer choices using a multinomial logit framework, where item appeal depends on both intrinsic preference and display position. They demonstrate that ignoring position effects leads to suboptimal performance and introduce a new algorithm, TLR-UCB, which effectively incorporates adaptive position-dependent feedback through a geometric linear bandit structure and truncated linear regression techniques. Theoretical analysis confirms that TLR-UCB achieves optimal learning efficiency. To handle unknown position effects, they further develop EI-TLR, a two-stage policy that jointly estimates customer preferences and positioning impacts before applying a generalized TLR-UCB procedure. Extensive simulations show that both TLR-UCB and EI-TLR significantly outperform existing benchmarks, offering powerful tools for dynamic, data-driven assortment and layout optimization in online marketplaces.

Robustly Leveraging Predictions in Revenue Management

Machine learning algorithms are becoming increasingly powerful, but with that power comes greater complexity and opacity. As these models become more sophisticated, they become increasingly difficult to understand—and, crucially, harder to anticipate when and how they might fail. This makes it essential to incorporate their predictions in ways that remain robust to errors. In “Single-Leg Revenue Management with Advice,” Balseiro, Kroer, and Kumar develop an algorithm for the classical single-leg revenue management problem that robustly incorporates predictions. They uncover a fundamental tradeoff: Placing greater trust in predictive models can yield high performance when predictions are accurate but also makes algorithms vulnerable when predictions are off. The proposed algorithm achieves the optimal tradeoff between these goals, allowing decision makers to leverage machine learning predictions while guarding against their potential inaccuracies. By doing so, this work provides a principled approach to integrating powerful yet imperfect forecasts into real-world decision making.

Trading Prophets

The prophet inequality is a cornerstone of online decision making, comparing a sequential decision maker to a prophet who knows all outcomes in advance. In “Trading Prophets,” Correa, Cristi, Dütting, Hajiaghayi, Olkowski, and Schewior initiate the study of buy-and-sell prophet inequalities. Here, an online algorithm observes a sequence of prices, one after the other, to trade an item. At each time step, the algorithm can decide to buy and pay the current price if it does not already hold the item, or it can decide to sell and collect the current price as a reward if it holds the item. The authors identify settings where a constant-factor approximation to the all-knowing prophet benchmark can be achieved. Interestingly, these conditions differ from those required for standard prophet inequalities. Specifically, they show that no constant-factor inequality exists for arbitrary independent prices. In contrast, they

prove that a constant factor is achievable when independent prices arrive in a random order.

Smarter Supervision Lifts Rural Weavers’ Productivity

Frequent, predictable supervisor visits can be a powerful lever for boosting artisan productivity in distributed supply chains, according to a study conducted with Jaipur Rugs in India. In “A Data-Driven Approach to Improve Artisans’ Productivity in Distributed Supply Chains,” Singhvi, Singhvi, and Zhang analyze loom-level data and show that reducing the average gap between visits by just one day raises weaving rates by 8.5%—with more substantial gains on complex rugs and when visits follow consistent schedules. Building on these insights, they develop a routing and scheduling framework that targets those looms most in need of support. In a 25-week field implementation covering about 6,000 visits across 200 looms, sites assigned to optimized routes saw a 16.7% increase in weaving speed relative to controls, highlighting a practical, low-cost pathway to higher earnings for rural women weavers. The research suggests that data-driven supervision in other supply chains with a similar structure (e.g., smallholder agriculture) could boost productivity and earnings, offering an operational lever for poverty alleviation at scale.

When Private Data Alone Drives Markets to Balance

In competitive markets, companies often lack access to their rivals’ sales, costs, and strategies. Can they still learn to make optimal decisions? In “Adaptive Learning in Uncertain and Sequential Competition,” Li and Mehrotra show that the answer is yes. Their research demonstrates that even without competitor data, firms can adaptively learn to make near-optimal choices using only their own operational information. More strikingly, when all players follow such self-driven learning, the entire market converges to a Nash equilibrium—the stable state predicted by economic theory—without explicit coordination. The study establishes theoretical guarantees for both convergence rates and regret performance and illustrates the framework in inventory management and dynamic pricing settings. These findings provide a foundation for data-driven decision making in competitive and uncertain environments and offer insights into how markets naturally self-organize.

Performance Analysis of Quantum Switches with Qubit Losses

Researchers have developed a novel model inspired by quantum switches to address the complexities of matching requests for entangled qubits in a discrete-time system. In “Matching Queues with Abandonments in Quantum Switches: Stability and Throughput Analysis,” Zubeldia, Jhunjhunwala, and Maguluri examine two types of arrivals: requests for entangled qubits between nodes and qubits supplied by nodes, which are subject to decoherence over time. Unlike classical queuing models, this system features server-less multiway matching and correlated abandonments, posing unique analytical challenges. By applying a max-weight policy, the researchers characterized the system’s stability using a two-time-scale fluid limit to account for qubit abandonments. They demonstrated that the max-weight policy is throughput optimal, outperforming nonidling policies under certain conditions. Intriguingly, the study revealed counterintuitive behavior: The longest request queue may grow temporarily, even in a stable system. These findings offer new insights into managing quantum-inspired systems with practical constraints, opening avenues for further research into quantum network optimization.

Privacy-Accuracy Trade-off from an Optimization Lens: Fix a Desired Level of Privacy, Then Maximize Accuracy

In differential privacy, the de facto standard for safeguarding individual information in data analysis, noise is added to statistics to limit the disclosure of sensitive information. Greater privacy requires more noise, creating a trade-off as the added noise reduces the accuracy of the resulting statistics. Historically, researchers have addressed this by restricting themselves to families of noise mechanisms that are sufficient for a predefined privacy level and proving their performance under specific conditions. In “Differential Privacy via Distributionally Robust Optimization,” Selvi, Liu, and Wiesemann propose a novel approach that guarantees optimal accuracy for any specified privacy level. They formulate the design of privacy mechanisms as an optimization problem that minimizes the expected loss associated with the random noise mechanism while encoding differential privacy as constraints. Through detailed analyses and by leveraging tools from distributionally robust optimization, the authors develop an efficient optimization algorithm and derive implementable solutions with provable guarantees to solve the problem within seconds.

Axiomatic Insights into Ambiguity Aversion: Binary Diversification and Exact Capacities

In decision sciences, ambiguity refers to uncertainty about event probabilities, and aversion to it is a well-studied phenomenon. In “Binary Diversification Characterizes Exact Capacities,” Hartmann and Kauffeldt consider the Choquet expected utility model, a framework for modeling ambiguity. They provide the first axiomatic characterization of exact capacities, solving an old problem. Their characterizing axiom, binary diversification, captures a specific level of ambiguity aversion—a preference for diversifications that lead to (at most) binary outcomes. The authors also propose an intuitive hierarchy of ambiguity aversion and illustrate that this hierarchy reflects increasingly strong levels of ambiguity aversion within the more general invariant biseparable model, whereas the hierarchy has limited bite in the Choquet model. They conclude by illustrating the implications for multiobjective shortest-path problems, demonstrating how the results can be applied.

Randomized Apportionment: A Fairer Distribution of Seats

The question of how to apportion the seats of the U.S. House of Representatives to states has fueled century-long political debates and sparked mathematical theory. Traditional deterministic methods, such as the Hamilton method or the currently used Huntington–Hill method, may result in paradoxes or substantially deviate from proportionality. In “In This Apportionment Lottery, the House Always Wins,” Gözl, Peters, and Procaccia propose a randomized approach that ensures each state receives its exact proportional share of seats in expectation and its proportional share, up to rounding, ex post. By incorporating randomization, the authors argue, the system can better adhere to the principle of proportional representation, minimizing the impact of small counting errors and ensuring fairness over time. In addition, their approach achieves house monotonicity, a property that prevents counterintuitive outcomes when the total number of seats changes. This is achieved through a novel cumulative rounding technique, a generalization of dependent rounding on bipartite graphs with potential applications beyond apportionment, including EU commission nominations and resource allocation.

Steering Customer Choices with Assortment Optimization

Businesses often seek to guide customer choices toward a common option to improve operational efficiency, enhance coordination, or optimize resource allocation. One example is attended home delivery platforms, which strategically select the set of delivery time slots offered to customers. By encouraging customers in the same residential area to choose overlapping slots, these platforms can reduce redundant trips and improve overall efficiency. In “Maximum Load Assortment Optimization: Approximation Algorithms and Adaptivity Gaps,” El Housni, Ibn Brahim, and Segev develop a general framework for optimizing assortment decisions to maximize the number of customers selecting the same option. Their work introduces algorithmic ideas that determine which options to present to customers in order to effectively influence their choices effectively. Beyond delivery scheduling, these algorithms apply to various settings where structuring choices to guide customer selections leads to better outcomes.

Breakthrough in Network Alignment: Chandeliers and Otter’s Constant

Network alignment or graph matching—figuring out how vertices across different networks correspond to each other—is a key challenge in many fields, from protecting online privacy to mapping biological data, improving computer vision, and even understanding languages. However, this problem falls into the class of notoriously difficult quadratic assignment problems, which are NP-hard to solve or approximate. Despite these challenges, researchers Mao, Wu, Xu, and Yu have made a major breakthrough. In “Random Graph Matching at Otter’s Threshold via Counting Chandeliers,” the authors introduce an innovative algorithm that can successfully match two random networks whenever the square of their edge correlation exceeds Otter’s constant (≈ 0.338). Their key innovation lies in counting chandeliers—specially designed tree-like structures—to identify corresponding vertices across the networks. The algorithm correctly matches nearly all vertices with high probability and even achieves perfect matching whenever the data allows. This is the first-ever polynomial-time algorithm capable of achieving perfect and near-perfect matching with an explicit constant correlation for both dense and sparse networks, bridging a long-standing gap between statistical limits and algorithmic performance.

Exponential Concentration in Stochastic Approximation

In “Exponential Concentration in Stochastic Approximation,” Law, Walton, and Yang analyze the convergence rates of stochastic approximation algorithms under nonvanishing gradient conditions. For these sharp (V -shaped) functions, standard Gaussian approximations fail, making tighter exponential concentration bounds more suitable. The authors prove that stochastic approximation algorithms, including Projected Stochastic Gradient Descent (PSGD), Kiefer-Wolfowitz, and Frank-Wolfe algorithms, exhibit exponential concentration near an optimum. A consequence is faster convergence rates, notably linear convergence, and $O(1/t)$ rates. The authors leverage techniques from Markov chain theory, specifically geometric ergodicity and Lyapunov drift analysis, thereby extending previous results by Hajek [Hitting-time and occupation-time bounds implied by drift analysis with applications. *Adv. Appl. Probab.* 14(3):502–525, 1982].

How to Efficiently Estimate Stability Measures: A Minimax Optimal Approach

In “Minimax Optimal Estimation of Stability Under Distribution Shift,” Namkoong, Ma, and Glynn address the challenge of

benchmarking the performance of decision policies and prediction models under distribution shift. Conventional risk measures and distributionally robust losses typically require specifying the magnitude of possible distribution shift—a quantity that is difficult to determine in practice. Instead, the authors consider a stability measure defined in terms of the acceptable level of performance degradation, which is more intuitive. To efficiently estimate this measure, the authors consider an estimator based on the dual formulation of the stability measure and show that this estimator is minimax optimal. Their results quantify the convergence rate of the estimator, which exhibits a fundamental phase shift behavior. In addition, they empirically observe that the stability measure reliably captures system performance under distribution shift in applications including queueing systems and healthcare prediction tasks.

Combinatorial Disjunctive Constraints: New Techniques for Piecewise Linear Relaxations of Nonlinear Functions

Many optimization problems in fields such as chemical engineering, robotics, and marketing involve nonlinear functions, which are computationally challenging to solve. A common approach to approximating these complex nonlinear functions is through piecewise linear functions. In “Building Formulations for Piecewise Linear Relaxations of Nonlinear Functions,” Lyu, Hicks, and Huchette introduce new mixed-integer programming formulations for piecewise linear relaxations of nonlinear functions using combinatorial disjunctive constraints and independent branching techniques. In particular, the authors demonstrate that these formulations significantly reduce computational time compared to other relaxation approaches while providing strong dual bounds for the original problems.

On the Sparsity of Optimal Linear Decision Rules for a Class of Robust Optimization Problems with Box Uncertainty Sets

One of the major reasons for the popularity of robust optimization is that these problems are often amenable to efficient approximations in operational planning problems where decisions must be made over time under uncertainty. The most common approximation is based on restricting the control policies to “linear decision rules”, that is, linear functions of the information revealed up in the past. In “On the Sparsity of Optimal Linear Decision Rules for a Class of Robust Optimization Problems with Box Uncertainty Sets,” Lu and Sturt prove for a class of robust inventory management problems that there always exists an optimal linear decision rule in which the number of nonzero parameters in the linear decision rule grows linearly in the number of time periods. This is the first result to prove that optimal linear decision rules are sparse in a widely studied class of robust optimization problems with many time periods. Harnessing this sparsity guarantee, the authors develop a novel reformulation technique and active set algorithm for computing optimal linear decision rules that yield a 32× speedup over state-of-the-art linear programming solvers in numerical experiments on production–inventory problems with hundreds of time periods.

Ambiguities in Average Speed of Answer Targets

Staffing problems are often formulated as satisfization problems, in which the cost of servers is minimized subject to quality of service constraints. These constraints are intended to capture customers’ disutility from waiting or, at least, its structure. In “Technical Note—What’s in a Constraint? On the Ambiguity of Standard Delay Targets,” Soh and Gurvich show that such targets—especially the popular average speed of answer—are ambiguous: they give rise to multiple optimal solutions (prioritization policies), each consistent with different assumptions about how customers value their time. By choosing, among all optimal solutions, the one that minimizes a weighted index of diversion (a generalization of variance for the multiclass queue), a service provider can ensure that its ASA-based staffing and prioritization decisions align with a convex model of customer delay disutility. Nonambiguity can also be enforced by restricting attention to fixed queue ratio priority policies.

Bayesian Process Control with Smart Sensors

Control charts are practical tools for fault detection and recovery. However, traditional control charts rely on random samples collected from a production process at fixed time intervals, causing late detection if sampling intervals are too long or excessive sampling if the intervals are too short. In “Event-Triggered Bayesian Control Chart,” Abbou and Makis develop a novel control chart leveraging real-time data from smart sensors to jointly decide when to collect samples and when to stop the production process, leading to quick fault detection and recovery using few samples. Applying optimal stopping theory and dynamic programming analysis, the authors establish the average-cost optimality of their control chart and propose an efficient procedure for computing the optimal sampling and stopping thresholds. Through an empirical study, the control chart is shown to achieve substantial cost savings compared to benchmarks. Furthermore, thanks to its event-triggering mechanism, the proposed control chart requires little data communication from sensors, which is crucial from an energy-efficiency perspective.

Data-Driven Policy Optimization with Robust Prescriptive Analytics

In “Robust Actionable Prescriptive Analytics,” Chen, Sim, Zhang, Zhao, and Zhou present a significant advancement in prescriptive analytics. The authors propose a novel robust prescriptive analytics framework that bridges data-driven decision making and actionable policy optimization. Unlike traditional approaches that follow a “predict, then optimize” methodology, this framework directly maps side information to optimized decisions, ensuring both interpretability and implementability. Leveraging a robust satisficing approach, the model effectively mitigates overfitting to empirical data while maintaining computational tractability. The authors also introduce tree-based static and affine policies for enhanced interpretability, and they demonstrate the framework’s practical value through a portfolio optimization case study. This innovative approach provides a powerful tool for decision makers seeking robust, data-driven policies across various operational contexts.