This article was downloaded by: [216.73.216.88] On: 09 November 2026, At: 10:41
Publisher: Institute for Operations Research and the Management Sciences (INFORMS)
INFORMS is located in Maryland, USA



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(6):v-ix. https://doi.org/10.1287/opre.2026.inthisissue.v74.n6

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

https://pubsonline.informs.org/journal/opre

OPERATIONS RESEARCH
Vol. 74, No. 6, November–December 2026, pp. v–ix
ISSN 0030-364X (print), ISSN 1526-5463 (online)

In This Issue

https://doi.org/10.1287/opre.2026.inthisissue.v74.n6

Copyright: © 2026 INFORMS

Downloaded from informs.org by [216.73.216.88] on 09 November 2026, at 10:41 . For personal use only, all rights reserved.

Will Hindsight Approach Work Well for Multi-Location Inventory Allocation?

Online retailers often decide where to place inventory before knowing which customers will arrive and in what order. A common solution is the hindsight or stochastic programming approach: allocate inventory as if future demand were known, then fulfill orders using a simple greedy rule, shipping from the cheapest available warehouse. In "Inventory Allocation under the Greedy Fulfillment Policy: The (Potential) Perils of the Hindsight Approach," Jasin, Liu, and Zhao show that this widely used method can be deceptively optimistic. The hindsight solution is equivalent to optimizing greedy fulfillment for the most favorable arrival sequence, rather than the sequence that unfolds. The authors characterize when this optimism is harmless—identical ordering costs, no warehouse capacity constraints, and independent demands—and show that violating any of these conditions can make the hindsight allocation arbitrarily worse than the true optimum under greedy fulfillment. Their results caution retailers to evaluate allocation heuristics against the fulfillment policy actually used.

Investing in Electricity Production Assets Under Price Ambiguity

How should investors set their capacity when market prices, which determine the revenues of their assets, are ambiguous? In "Capacity Investment Decisions in Equilibrium: A Distributionally Robust Approach," Abada and Ancel focus on European electricity markets and argue that the past distribution of prices is no longer trustworthy in crisis regimes such as those recently observed. Leveraging the distributionally robust optimization framework, the authors provide a novel and tractable formulation of the stochastic capacity expansion problem under price ambiguity in a power economy, leading to an equilibrium framework. They prove existence and show that their framework is compatible with the economic theory of pure and perfect competition. Through simulations, they quantify the detrimental impact of ambiguity on system welfare and installed capacity and estimate how certain contracts could help mitigate these effects.

Rethinking Artificial Intelligence-Assisted Priority Queuing: The Cost of Locking in Classifiers

When job types must be inferred from observable features, should artificial intelligence (AI) classifiers be locked before optimizing queue assignments, or should the full system be trained end to end? In "Feature-Driven Priority Queuing," Singh, Gurvich, and Van Mieghem show that the modular "type-first" approach—locking a type classifier and then, optimizing queue assignment based only on its output distribution—can be fundamentally suboptimal. The authors establish that type-first achieves optimality only when the locked classifier recovers the Bayes posterior almost everywhere, a condition rarely satisfied in high-dimensional or misspecified settings. In contrast, their "direct" approach jointly optimizes feature-to-queue mappings to minimize empirical waiting costs and converges to the theoretical optimum. Using 100,000 chest X-rays from the National Institutes of Health, they show that direct optimization substantially improves waiting-cost performance and reveal mechanisms, such as statistical pooling and strategic overprioritization of medium-cost types. The work speaks to the growing debate over whether AI systems should be locked or adaptable.

Blunting Gerrymandering with Multimember Districts

How much can election rules blunt the power of partisan mapmakers? In "Combating Gerrymandering with Ranked Choice Voting: An Experimental Analysis of Multimember Districts in the United States," Garg, Gurnee, Rothschild, and Shmoys study this question by combining computational redistricting with models of multiwinner voting. The authors compare single-member districts with multimember districts elected under rules such as single transferable vote (ranked choice voting), using algorithmically generated maps for the U.S. House of Representatives. Their results suggest that even modest three-member districts can help independent commissions reach proportional outcomes, sharply reducing the range of outcomes available to partisan gerrymanderers. The study also tests robustness under voter crossover and explores trade-offs involving geographic cohesion and within-party diversity. The paper points to election-system design as an operations research problem with direct democratic consequences.

A Data-Driven Early Warning System for Disease Outbreaks

Early detection of infectious disease outbreaks is essential for timely public health response, yet local case data are often sparse and noisy, making reliable monitoring difficult. In "Small-Area Estimation of Case Growths for Timely COVID-19 Outbreak Detection," She, Wang, Ayer, and Chhatwal propose a new statistical learning framework, that is, transfer learning random forests (TLRF), that improves the estimation of epidemic growth rates across small geographic regions.

When Screening and Vaccine Incentives Work Better Together

When vaccine mandates are not feasible, institutions may rely on vaccination rewards or routine screening of the unvaccinated. The study shows these tools can be more powerful when designed together. Exempting vaccinated individuals turns screening into a nonmonetary incentive while retaining its role in detecting infections. In "Behaviorally Informed Strategies for Infectious Disease Mitigation: Balancing Incentives and Screening," Tunc, Görgülü, Rabil, Bish, and Bish develop a behaviorally informed framework that accounts for vaccine hesitancy, community preferences, limited budgets, and testing capacity. Rather than committing to a rigid plan, the framework learns from surveys and vaccination data and adapts screening as community behavior unfolds. The analysis reveals that incentives and screening are not always substitutes: under some conditions, each strengthens the value of the

Downloaded from informs.org by [216.73.216.88] on 09 November 2026, at 10:41 . For personal use only, all rights reserved.

other. In a realistic college-campus case study using a detailed epidemic model, the adaptive strategy reduced infections and hospitalizations while adhering more closely to budget than incentive-only or screening-only approaches. The findings offer practical guidance for designing flexible mitigation plans for future outbreaks in universities and other close communities.

Testing Smarter, Not More: An Adaptive Approach

In "A Novel Adaptive Testing Scheme for Multidisease Testing," Bish, Bish, Bodur, and Görgülü develop an adaptive, cost-efficient testing scheme for lab-based multidisease testing, essential for diagnosis and screening, particularly for disease groups that share symptoms but require disease-specific treatments (e.g., respiratory, gastrointestinal, or sexually transmitted diseases). Their proposed adaptive strategy integrates multiplex assays and pooled testing in a novel way, utilizing modern automated testing technologies to build a customized follow-up test for each person, covering only the diseases flagged in that person's relevant pooled tests. Optimizing such adaptive designs is mathematically demanding, but in a case study of 18 respiratory diseases, based on data from the Centers for Disease Control and Prevention, their adaptive designs yield substantial cost savings (51%–61%) relative to current practice under forecast error and highly seasonal prevalences. By lowering the cost per subject, adaptive designs can expand testing capacity and support earlier detection under resource constraints.

Tensor and Matrix Methods for Multivariate Time Series, with Provable Guarantees

Real-world time series data—from financial markets to electricity demand—are often multivariate, noisy, and riddled with missing observations. In "On Multivariate Singular Spectrum Analysis: Tensor and Matrix Variants," by Agarwal, Alomar, and Shah introduce and rigorously analyze two extensions of singular spectrum analysis (SSA) for this challenging setting. Their matrix variant (mSSA) achieves imputation and forecasting error scaling as $1/\sqrt{\min(N,T)T}$, improving over both univariate SSA and standard matrix estimation methods that ignore temporal structure. A novel tensor variant (tSSA) further improves sample complexity in certain regimes, strengthening the link between time series analysis and tensor estimation. These results are made possible by a spatio-temporal factor model that accommodates a rich class of dynamics—harmonics, polynomials, and smooth periodic functions—which the authors extend further by introducing a "Hankel calculus" establishing closure under component-wise addition and multiplication. Empirically, mSSA matches state-of-the-art deep learning methods while significantly outperforming classical approaches such as vector autoregression.

Simplicity Meets Optimality in Online Resource Allocation

Online platforms frequently manage sequential allocation decisions where outcomes are uncertain, such as selecting which products to display to a user. While it is often assumed that complex "adaptive" algorithms—those reacting to real-time feedback like user choices—are necessary for optimal performance, Rajan Udwani's paper, "Optimality of Nonadaptive Algorithms in Online Submodular Welfare Maximization with Stochastic Outcomes," challenges this assumption.

Epidemic Forecasting on Networks: Bridging Local Samples with Global Outcomes

Forecasting how an epidemic will unfold, its trajectory and eventual size, usually presumes access to the full contact network. In practice, no one ever has the full mapping of networks. New research shows that full mapping is unnecessary. By observing only the local neighborhoods around a small number of individuals, the kind of partial data that can actually be collected during an epidemic, in "Epidemic Forecasting on Networks: Bridging Local Samples with Global Outcomes," Alimohammadi, Borgs, van der Hofstad, and Saberi prove that global outcomes, such as the eventual size of an epidemic, can be predicted with rigorous accuracy guarantees. The work formally connects what public health planners can observe locally—for example, connections revealed by contact tracing and similar methods—to how an epidemic behaves globally, giving them a principled way to turn limited, realistic data into reliable forecasts.

Knowing When to Hold Back: Smarter Messaging Keeps Customers

Marketers often assume that reaching customers more often produces better results. A new study suggests the opposite can be true. In "Doubly Adaptive Cascading Bandits with User Abandonment," Cao, Sun, and Shen examine how digital platforms can boost engagement without driving customers away through marketing fatigue. They develop a family of "doubly adaptive" learning algorithms that learn across users while also adjusting the messages sent to each individual in real time, weighing relevant content against the risk that one poorly targeted message prompts a customer to unsubscribe or delete an app. The authors prove their methods come close to the best possible performance and test them on real e-commerce data spanning tens of millions of interactions, where the approach retained far more users than standard benchmarks. The lesson is that sometimes sending fewer, better-timed messages is what keeps customers around.

Optimal Design of Default Donations

Many nonprofit organizations rely on suggested donation amounts to incentivize giving, yet there is little consensus on how to set those amounts. In "Optimal Design of Default Donations," Castro and Rodilitz develop a data-driven framework for designing optimal donation menus that maximize fundraising revenue while accounting for differences in donor preferences and responsiveness to suggestions. The authors show that commonly used heuristics can perform poorly and instead introduce a dynamic-programming algorithm that efficiently identifies revenue-maximizing donation menus. Their analysis also quantifies the value of donor information and demonstrates when personalized suggestions, larger menus, or a single targeted default are most effective. Applying the framework to data from a large fundraising experiment, the authors find that an optimally designed menu could substantially outperform the default amounts tested in practice. The results provide actionable guidance for nonprofits seeking to improve online fundraising through more effective choice architecture.

Going Beyond Single-Point Demand Forecasts in Real-Time Allocation

Many companies use machine learning to forecast future demand when making real-time allocation decisions, such as deciding how many airline seats, hotel rooms, or advertising opportunities to reserve for different customer groups. Most existing approaches rely on a single demand prediction, or "point estimate." In "Online Resource Allocation with Convex-Set Machine-Learned Advice," Golrezaei, Jaillet, and Zhou develop a new framework that goes beyond point estimates by representing forecasts as a convex uncertainty set—a range of possible demand outcomes for both high-reward and low-reward customers. This

Downloaded from informs.org by [216.73.216.88] on 09 November 2026, at 10:41 . For personal use only, all rights reserved.

richer form of machine-learned advice captures uncertainty, variability, and correlations in demand while also allowing realized demand to differ from the forecast. The authors develop adaptive online algorithms with provable theoretical guarantees, showing how to optimally balance performance when forecasts are accurate with robustness when forecasts are inaccurate or misleading. Numerical studies further demonstrate that the proposed approach outperforms methods based only on single-point predictions.

Online Rack Placement in Large-Scale Data Centers: Online Sampling Optimization and Deployment

Data centers have grown into major components of global supply chains. In "Online Rack Placement in Large-Scale Data Centers: Online Sampling Optimization and Deployment," by Baxi, Cummings, Jacquillat, Lo, McDonald, Mellou, Menache, and Molinaro develop, deploy, and assesse an optimization algorithm to improve how large-scale data centers place incoming server racks dynamically while balancing space, power, cooling, and reliability constraints. Poor placement decisions can leave valuable resources stranded, leading to high costs and reduced operating resilience. This paper formulates a large-scale online discrete optimization model and develops a new online sampling optimization (OSO) algorithm that anticipates future demand by repeatedly simulating future arrivals and reoptimizing decisions over time. Theoretical results provide performance guarantees, and computational results show its benefits against state-of-the-art reoptimization methods. The system was implemented as a decision-support tool and deployed across Microsoft's global fleet of data centers. Using postdeployment data, the paper shows that adoption of the tool reduced power stranding by one to three percentage points. At Microsoft's scale, these improvements translate into substantial financial savings and meaningful reductions in greenhouse gas emissions, demonstrating the real-world impact of optimization in cloud infrastructure management.

Constant-Factor Policies for Consecutive-Stay Bookings

Railway operators and hotels must decide whether to accept bookings while preserving capacity for future customers. In "Constant-Factor Algorithms for Revenue Management with Consecutive Stays," Hu and Wu study this challenge in which requests span consecutive travel legs or nights. The authors develop polynomial-time policies for two settings: one in which a seller accepts or rejects a request and assigns an available resource and another in which customers choose from an assortment under the basic attraction model. When benchmarked against the optimal online policy, the algorithms deliver the first constant-factor revenue guarantees in these settings independent of the maximum stay length. For fixed request types (e.g., their stay intervals are fixed), they guarantee 63.2% of the optimal expected revenue for accept-or-reject decisions and 27.1% under customer choice; for random request types (e.g., their stay intervals are random), the guarantees are 39.9% and 17.1%, respectively. These constant-factor guarantees stand in sharp contrast to the prior nonconstant competitive ratios that are benchmarked against the offline optimum.

Increasing Parents' Engagement on EdTech Platforms: Evidence from the Field

Simple, well-designed digital nudges can meaningfully increase parents' participation in early childhood education, according to a large-scale study conducted with Rocket Learning in India. In "Increasing Parents' Engagement on EdTech Platforms: Evidence from the Field," Goyal, Okuno, Singhvi, and Singhvi examine peer- and self-comparison messages delivered through digital educational communities. In a randomized controlled trial involving nearly 154,000 parents, peer-comparison nudges increased the average number of parents active each day by 13.5%, whereas self-comparison nudges increased it by 8.0% relative to no nudge. Building on these insights, they develop a data-driven targeting policy that selects the nudge best suited to each parent. A second trial involving 122,000 parents shows that, whereas the targeting policy performs comparably to the best uniform policy on average, it identifies a substantial subgroup—8.8% of parents—for whom targeted nudging increases engagement by 5.4% over the uniform approach. The findings offer nonprofits and policymakers a scalable, low-cost tool for strengthening parental engagement, highlighting both the potential and limitations of personalized nudging.

Learning from Snapshots Is Not Enough: An Even-Driven Continuous-Time Reinforcement Learning Framework

Revenue management systems evolve continuously, but reinforcement learning often requires dividing time into a fixed grid. Fine grids improve accuracy but increase computation; coarse grids can sacrifice performance. In "Reinforcement Learning for Intensity Control: An Application to Choice-Based Network Revenue Management," Meng, Chen, and Gao address this tension with an event-driven continuous-time reinforcement learning framework. The key insight is that state jumps occur at customer arrival times and naturally partition each sample path, eliminating the need for up-front time discretization. Building on this structure, the authors extend policy evaluation to continuous time and develop event-driven actor-critic algorithms. A comprehensive numerical study shows their strong performance. In a bursty arrival environment, the proposed continuous-time approach achieves up to 16.64% higher revenue than a coarse-grid benchmark with comparable training time. The approach also handles a large-scale network revenue management problem with 100 resources and 200 products, and an extension to queue admission control further demonstrates the broad applicability of the framework.

AI-Enhanced Monte Carlo Methodology Achieves More Reliable Testing and Efficiency Certificate for Safety-Critical Systems

As artificial intelligence (AI) increasingly drives human-interacting intelligent physical systems such as self-driving vehicles, ensuring their safety against rare catastrophic events has also become a critical challenge. Traditional simulation-based testing methods are powerful in quantifying risks before real deployment. However, for black box systems that generically arise in AI-embedded applications, these methods can consistently underestimate the probabilities of rare failures and, dangerously, without diagnostically detectable signals. To address this perilous gap, the paper "Certifiable Deep Importance Sampling for Rare-Event Simulation of Black Box Systems," by Arief, Bai, Ding, He, Huang, Lam, and Zhao, introduces a framework based on an integration of neural networks into importance sampling, a variance reduction technique that has been found useful in rare-event simulation. This framework, which is called deep probabilistic accelerated evaluation (Deep-PrAE), uses deep neural network classifiers to learn rare-event geometries, which then calibrates importance samplers to estimate rare events with certifiable statistical guarantees. Moreover, these guarantees are designed to go beyond the conventional certifiable notions in rare-event simulation to effectively balance the applicability to black box settings

Downloaded from informs.org by [216.73.216.88] on 09 November 2026, at 10:41 . For personal use only, all rights reserved.

with the required sampling efficiency gain that pertains to rare-event computation. These findings demonstrate a vital step forward in the reliable and certifiable deployment of complex autonomous technologies.

Generative AI Learns to Sample Extreme Events

Many rarest events often carry highest stakes, such as a market crash, a supply-chain rupture, or an extreme clinical outcome. Today's leading generative AI models (e.g., diffusion models) have shown excessive power to simulate high-dimensional distributions, but they are incapable of capturing rare events in the heavy tails of a distributions. In "Learning to Simulate from Heavy-Tailed Distribution via Diffusion Model," Liu, Zhu, Jia, He, and Zheng show that this mismatch is structural rather than incidental: standard diffusion models systematically fail to capture tail behavior, both when learning from heavy-tailed data and when generating from it. In *Learning to Simulate from Heavy-Tailed Distribution via Diffusion Model*, the authors pinpoint why standard diffusion models break down on heavy-tailed targets and propose a Student-t-based heavy-tailed diffusion (SHD) framework that fixes both ends of the pipeline. Across synthetic Pareto data, vector autoregressive systems, queueing networks, bike-sharing demand, and stock equity returns, SHD beats Gaussian-noise baselines, particularly in the tails where operations research decisions hinge.

Optimizing Warehouse Space: Algorithmic Advances in Inventory Staggering

Since the mid-1960s, the inventory staggering problem has challenged supply chain experts and logisticians aiming to minimize peak storage requirements across multiple periodic replenishment policies. Despite its practical prevalence in warehousing, aerospace, and production planning, foundational computational questions have remained largely unresolved. In "New Approximation Guarantees for the Inventory Staggering Problem," Alon and Segev discuss devising a host of algorithmic techniques and analytical ideas—some being entirely novel and some leveraging well-studied concepts in combinatorics and number theory—for improving upon existing approximation guarantees for the inventory staggering problem.

Exact Matching from Multiple Graphs

Networks often describe the same underlying objects from different viewpoints: users on social platforms, proteins across species, or maps built by collaborating robots. But when labels on the nodes of the network are hidden or scrambled, how much correlation among the networks is enough to recover the true correspondence? In "Aligning Multiple Inhomogeneous Random Graphs: Fundamental Limits of Exact Recovery," Ameen and Hajek study this question for multiple heterogeneous random graphs, where different pairs of nodes may have different connection probabilities. The paper introduces a simple principle: first obtain reliable partial pairwise matchings, then combine them through transitive closure. The authors identify conditions under which this procedure exactly matches all nodes, and show that for homogeneous Erdős–Rényi graphs these conditions are information-theoretically sharp. A key insight is that multiple graphs can succeed even when no pair of graphs alone contains enough information for exact matching. The work also develops related results on k-cores of inhomogeneous random graphs.

Convergence of Optimistic Learning in Games and the Role of Forgetfulness

Online learning algorithms solve games by repeatedly updating players' strategies. A natural hope is that the latest strategy improves at a predictable rate. In "Separation of Nonergodic Uniform Convergence Rates for Regularized Learning in Games," Cai, Farina, Grand-Clement, Kroer, Lee, Luo, and Zheng show that this intuition can fail for optimistic follow-the-regularized-leader methods, including the widely used optimistic multiplicative weights algorithm. In two-player zero-sum games, the authors separate three notions of nonergodic performance: last-iterate, random-iterate, and best-iterate convergence. They prove that no instance-independent last-iterate rate exists for the broad algorithmic family and establish strong lower bounds for random iterates. Yet, a useful positive result remains; in 2×2 games, optimistic multiplicative weights achieve a uniform best-iterate rate. The analysis traces the slowdown to a lack of "forgetfulness"; accumulated past losses can keep the dynamics moving away from equilibrium long after they reach its neighborhood. The results suggest that practitioners should distinguish carefully between the latest strategy, a randomly selected strategy, and the best observed strategy when evaluating learning dynamics.

Nonsmooth Learning Algorithms Behave Differently from Smooth Ones

Many learning algorithms, including Q-learning, update their estimates through noisy recursive rules. When these updates use a constant step size and involve nonsmooth operators, their long-run behavior can be difficult to characterize: the iterates do not converge to a single point but instead settle into a stationary distribution, and classical tools for smooth dynamics may fail. In "Prelimit Coupling and Steady-State Convergence of Constant-Step-Size Nonsmooth Contractive Stochastic Approximation," Zhang, Huo, Chen, and Xie develop a prelimit coupling framework for analyzing such nonsmooth stochastic approximation procedures. They establish Wasserstein convergence to a unique stationary distribution, characterize the limiting steady-state distribution as the step size vanishes, and show that nonsmoothness can create a bias of order of the square root of the step size, unlike the bias linear in the step size that is typical of smooth dynamics. The results also justify Richardson–Romberg extrapolation as a tool for bias reduction, with applications to Q-learning.

A Network View of Bilevel Optimization

Many decision problems are inherently hierarchical; a leader acts first while anticipating that a follower will respond optimally. Determining the leader's best decision is notoriously difficult in areas such as infrastructure design, transportation, and security, especially when decisions are discrete. In "Network Relaxations for Combinatorial Bilevel Optimization Under Linear Interactions," Lozano, Bergman, and Cire discuss results that demonstrate the practical value of exploiting value function structure and network approximations to enhance the tractability of bilevel optimization.

Finite-Time Guarantees for Scheduling

Scheduling underpins modern infrastructure, including data centers, wireless networks, service systems, and graphics processing unit clusters. In these environments, bursty, nonstationary demand and stringent service objectives call for rigorous finite-time performance guarantees. In "Finite-Time Minimax Bounds and an Optimal Lyapunov Policy in Queueing Control," Liu, Tan, and Xu develop a minimax framework for analyzing finite-time scheduling performance. The framework quantitatively characterizes how the expected total queue length scales with system capacity and variability in arrivals and departures, providing a basis for evaluating and comparing scheduling policies. The authors also introduce LyapOpt, which accounts for both first- and second-order effects in service allocation, whereas

MaxWeight optimizes only first-order drift. Under a specific structured condition, LyapOpt achieves optimal finite-time performance while retaining stability guarantees. In contrast, MaxWeight's finite-time performance scales suboptimally with system parameters. These results clarify the limitations of classical drift-based scheduling and motivate new queueing-control methods with rigorous finite-time guarantees.

When "Optimal" Advice Gets Ignored

Algorithms can recommend decisions that look optimal on paper but conflict with how people assess the situation, making those recommendations less likely to be followed. In "Conformal Inverse Optimization for Adherence-Aware Prescriptive Analytics," Chan, Delage, and Lin address this challenge by using past decisions to learn how people evaluate different options. Because those decisions may reflect diverse preferences, noisy observations, or imperfect models, the authors do not rely on a single estimated preference profile. Instead, they construct an uncertainty set around the inferred parameters and use robust optimization to generate recommendations that perform well both objectively and from the decision maker's perspective. Their method comes with statistical and performance guarantees and outperforms standard inverse optimization in numerical experiments. In a simulated Toronto food-delivery study, it substantially improves estimated courier adherence, maintaining comparable delivery times, showing how better-aligned recommendations can benefit both users and organizations.

A New Class of Overbooking Problems in Container Shipping Industry

Container carriers lose billions of dollars each year when shippers reserve slots but fail to show up with the corresponding cargo. In "Prophet Inequalities for a New Class of Overbooking Problems in Container Shipping," Feldman, Huang, and Kouvelis study an online booking problem under a deposit-based mechanism designed to better align incentives and mitigate the costs of overbooking. The authors derive near-optimal online policies for two demand settings: one in which shippers' no-show behavior is coupled through a spot market, and another in which no-shows are independent. In the coupled setting, a randomized threshold policy guarantees at least 81.9% of clairvoyant profit. In the independent setting, a simple greedy policy is shown to be asymptotically 0.5-optimal as the number of slots tends to infinity. Extensive simulations further demonstrate strong and robust practical performance.

Incentivizing Effort While Learning the Market Condition

When a firm hires an agent to generate customers in a new market, neither party initially knows whether the market is favorable. In "Dynamic Contract Design with Learning," Wang, Liang, and Sun study the interaction between learning and moral hazard: effort is costly and unobservable, so a missing arrival may reflect an unfavorable market, bad luck, or shirking. The firm's goal is to design a dynamic contract governing arrival-based payments and termination, thereby motivating effort while learning whether to continue. Yet shirking can make the principal more pessimistic than the agent and create scope to manipulate future payments. The authors prove that an optimal contract inducing full effort until termination exists, but computing it appears prohibitively complex. They develop two simple online contracts with payments and termination thresholds fixed in advance. Both induce full effort and achieve $O(\ln T)$ regret over T periods against a first-best, full-information benchmark. A matching $\Omega(\ln T)$ lower bound proves rate optimality. The results extend to continuous time and differing prior beliefs.

Downloaded from informs.org by [216.73.216.88] on 09 November 2026, at 10:41 . For personal use only, all rights reserved.