



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(3):iv–viii. <https://doi.org/10.1287/opre.2025.inthisissue.v73.n3>

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

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

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2025, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

In This Issue

<https://doi.org/10.1287/opre.2025.inthisissue.v73.n3>

Copyright: © 2025 INFORMS

Efficiently Computed Periodic Ordering Heuristics in Stochastic Distribution Systems

The difficulty of analyzing and optimizing the stochastic one-warehouse multiretailer problem under the (S, T) policy motivates the need to consider approximate but high-fidelity models that are easier to scrutinize. In “Heuristic (S, T) Solutions via an FPTAS for a One-Warehouse Multiretailer Problem,” Najy, Diabat, and Elbassioni show how to efficiently compute competitive solutions for such systems based on an approximate formulation of the inventory problem proposed by Chu and Shen (2010) and build on the assumption of power-of-two (POT) policies. The authors first devise a fully polynomial-time approximation scheme for the continuous relaxation of the model and then show how to round it to a POT solution within an improved approximation factor relative to the present literature. These solutions are shown via simulation to be highly competitive with optimal (S, T) solutions and are derived in a tiny fraction of the time needed by the current best (S, T) algorithms.

Optimizing the Path Towards Plastic-Free Oceans

Increasing ocean plastic pollution is irreversibly harming ecosystems and human activities. The Ocean Cleanup (TOC) is a Dutch nongovernmental organization dedicated to cleaning oceans from plastic pollution. Among its initiatives, TOC operates a plastic collection system in the Great Pacific garbage patch, an accumulation area twice the size of Texas and located between Hawaii and California. In “Optimizing the Path Towards Plastic-Free Oceans,” den Hertog, Pauphilet, Pham, Sainte-Rose, and Song design a routing algorithm that leverages data on weather conditions and plastic dispersion models to determine an optimal trajectory for their system, maximizing plastic collection. Unlike conventional routing problems, the trajectory directly impacts future plastic distribution, creating nonlinearities that make standard dynamic programming unsuitable. To address this, they devise a tailored algorithm based on a relaxation-induced search and a customized branch-and-bound scheme. Validated on one year of ocean data, their algorithm results in a 60% increase in plastic collection compared with current strategies.

Advancing Column Generation by a Novel Variable Fixing Method

In “DeLuxing: Deep Lagrangian Underestimate Fixing for Column-Generation-Based Exact Methods,” Yang introduces DeLuxing—an innovative variable-fixing technique that significantly advances column-generation-based exact methods for solving large-scale optimization problems, particularly vehicle routing problems (VRPs). DeLuxing leverages a novel linear programming formulation with a small subset of the enumerated variables, which is theoretically guaranteed to yield qualified dual solutions for computing Lagrangian underestimates. By eliminating over 75% of the unnecessary variables, DeLuxing drastically boosts computational efficiency, doubling the size of CMTVRPTW (capacitated multitrip vehicle routing problem with

time windows) instances that can be solved optimally. Additionally, this breakthrough accelerates top VRP solvers like RouteOpt by up to 71%. The core concept underpinning DeLuxing extends to broader contexts such as variable type relaxation and cutting plane addition, achieving an additional 25% speedup for difficult instances.

Time-Consistent Diversified Portfolio Allocations

Portfolio allocation problems typically involve diversification of risks, which can be achieved by risk budgeting strategies. Of additional importance is being time consistent—that is, ensuring that decisions about what to do in the future under certain outcomes remain optimal when those future outcomes are realized. In “Risk Budgeting Allocation for Dynamic Risk Measures,” Pesenti, Jaimungal, Saporito, and Targino develop time-consistent portfolio strategies, termed dynamic risk budgeting strategies, where the diversification is on the level of risk contribution of assets. The authors formalize the dynamic setting of risk budgeting strategies, recast dynamic risk budgeting strategies as solutions of sequences of convex optimization problems, and develop an efficient actor-critic deep reinforcement learning algorithm for estimating dynamic risk budgeting strategies. The algorithm and the dynamic risk budgeting strategies are illustrated on a complex simulation case study involving five assets and 12 time steps.

Statistical Arbitrage with Price Impact

A key challenge to turn price forecasts into profitable trading strategies is to mitigate the adverse impact of large trades. In “Trading with Concave Price Impact and Impact Decay—Theory and Evidence,” Hey, Mastromatteo, Muhle-Karbe, and Webster show that such statistical arbitrage problems can be solved in closed form for general nonparametric price forecasts and for price impact models that account for the nonlinearity in order sizes and gradual decay of impact observed empirically. The implications of these results are illustrated by an empirical case study, in which the price impact models are calibrated to proprietary metaorder data.

Optimizing Assortment and Inventory Under Stockout-Based Substitution

In the realm of retail, a common challenge arises when retailers must decide the set of substitutable products to offer and how much inventory to stock at the beginning of a selling season when customers choose and substitute among the products that are available to them. This problem is challenging, as the demand for each product is shaped by the inventories of all products and the substitution behavior. In “Technical Note—Leveraging the Degree of Dynamic Substitution in Assortment and Inventory Planning,” Zhang, Ma, and Topaloglu focus on the stockout-based substitution and develop a novel rounding scheme based on a fluid relaxation. The authors provide a rigorous theoretical analysis of the method, improving earlier results. In particular, for the multinomial logit model, they achieve an optimality gap that

does not depend on the number of products. This result has practical implications for online retailers with vast product varieties, ensuring reliable performance, even when both the demand and the number of products are large.

Revenue Management with Dependent Demands

In the revenue management literature, a standard approach to model the demand is based on dividing the selling horizon into a number of small time periods such that there is at most one customer arrival at each time period and that the customer arrivals at different time periods are independent of each other. Under this demand model, if the mean demand for a product increases with rate T , then the standard deviation of the demand must increase with rate square root of T . In other words, large demand volume and large demand variability cannot coexist. In “Technical Note—Revenue Management with Calendar-Aware and Dependent Demands: Asymptotically Tight Fluid Approximations,” Li, Rusmevichientong, and Topaloglu consider demand models that can simultaneously incorporate large demand volume and large demand variability while also allowing dependence between the demands over nonoverlapping time intervals. Under such demand models, they give efficiently computable policies that admit performance guarantees.

Algorithms for Dynamic Resource Allocation and Novel Drivers of Performance

In “Dynamic Resource Allocation: Algorithmic Design Principles and Spectrum of Achievable Performances,” Besbes, Kanoria, and Kumar consider a broad class of dynamic resource allocation problems and study the performance of practical algorithms. In particular, they focus on the interplay between the distribution of request types and achievable performance, given the broad set of configurations that can be encountered in practical settings. Although prior literature studied either a small number of request types or a continuum of types with no gaps, the present work allows for a large class of type distributions. Using initially the prototypical multisecretary problem to explore fundamental performance limits as a function of type distribution properties, the authors develop a new algorithmic property “conservativeness with respect to gaps” that guarantees near-optimal performance. In turn, they introduce a practically motivated, simulation-based algorithm and establish its near-optimal performance, not only for multisecretary problems, but also for general dynamic resource allocation problems.

Pricing and Assortment Optimization over a Network of Consumers

In “Personalized Pricing and Assortment Optimization Under Consumer Choice Models with Local Network Effects,” Xie and Wang introduce a consumer choice model where each consumer’s utility is affected by their neighbors’ purchase probabilities in a network. The authors first characterize the choice probabilities and then consider the associated personalized assortment optimization problem. Although this problem is NP-hard, we show that for star networks, the optimal assortment to the central consumer cannot be strictly larger than that without network effects; and the optimal assortment to each peripheral consumer must be a revenue-ordered assortment. Then, the authors propose a novel idea in which the sellers are allowed to offer “randomized assortments” to each node in the network by viewing each node as a consumer group. The authors show that randomized assortments may further increase the revenue, and for general network setting, under certain conditions, the optimal assortment must be a combination of two adjacent revenue-ordered assortments.

Finally, the authors study the associated optimal pricing problem and develop structural properties and efficient algorithms.

Rethinking Informational Assumptions in Mechanism Design

Mechanism design has proven to be very effective is soliciting buyers’ incentives and finding the optimal way of allocating scarce resources. A common assumption in classic mechanism design, however, is that buyers have full information about their values. Motivated by application in e-commerce, Alaei, Makhdoui, and Malekian, in “Revenue Maximization Under Unknown Private Values with Nonobligatory Inspection,” consider settings without this assumption and find the optimal selling mechanism for them.

Tackling High-Dimensional Tensor Clustering

In “Jointly Modeling and Clustering Tensors in High Dimensions,” Cai, Zhang, and Sun address the challenge of jointly modeling and clustering tensors by introducing a high-dimensional tensor mixture model with heterogeneous covariances. The proposed mixture model exploits the intrinsic structures of tensor data. The authors develop a computationally efficient high-dimensional expectation conditional maximization (HECM) algorithm and show that the HECM iterates, with an appropriate initialization, converge geometrically to a neighborhood that is within statistical precision of the true parameter. The theoretical analysis is nontrivial because of the dual nonconvexity arising from both the expectation maximization-type estimation and the nonconvex objective function in the M step. They also study the convergence rate of the algorithm when the number of clusters is overspecified and when the signal-to-noise ratio diminishes with sample size. The efficacy of the proposed method is demonstrated through numerical experiments and a real-world medical data application.

Is it Better Not to Diversify?

Diversification is generally regarded as an efficient tool to reduce portfolio risks. In “Technical Note—An Unexpected Stochastic Dominance: Pareto Distributions, Dependence, and Diversification,” Chen, Embrechts, and Wang show that the weighted average of independent and identically distributed (i.i.d.) Pareto random variables with infinite mean is larger than one such random variable in the sense of first-order stochastic dominance, and thus diversification is, surprisingly, worse than no diversification. The relation implies superadditivity of value-at-risk, a regulatory risk measure used in the finance and insurance sectors. The obtained relation also holds under some form of negative dependence.

Modeling and Algorithmic Design for Data-Driven Minimax Optimization

The realm of data-driven optimization has garnered substantial attention in recent years. However, there remains a paucity of research addressing challenges posed by intricate data-driven constraints. In “Data-Driven Minimax Optimization with Expectation Constraints,” Yang, Li, and Lan concentrate on the realm of data-driven minimax optimization while considering expectation constraints. To grapple with these complex models, the authors introduce a novel class of optimal primal-dual algorithms. Their work showcases the practical efficacy of these algorithms through the resolution of real-world problems, including data-driven robust pricing and the maximization of the area under the ROC curve (AUC) while adhering to fairness constraints.

Solving Large-Scale Linear Programs by Randomly Sampling Columns

Large-scale linear programs (LPs) play a pivotal role in various applications, including the classic cutting-stock problem and the vehicle routing problem. The standard solution approach for these LPs, namely, column generation (CG), is often computationally intractable because of the NP-hard nature of the corresponding subproblem in many applications. In “Column-Randomized Linear Programs: Performance Guarantees and Applications,” Akchen and Mišić introduce a randomized method that involves first sampling a set of columns from the original LP and subsequently solving an LP composed of the sampled columns, termed the column-randomized LP. The authors analyze the optimality gap of the column-randomized LPs and establish conditions under which the gap is small. Empirical findings demonstrate the effectiveness of the column-randomized LP approach, showcasing its advantages over the CG approach in two practical applications: the cutting-stock problem and nonparametric choice model estimation.

Optimal Strategies for Wordle and its Variants

Motivated by the recent, explosive popularity of Wordle and a suite of suboptimal attempts to solve the game, Bertsimas and Paskov, in “An Exact Solution to Wordle,” propose a new and scalable framework to exactly solve Wordle and any of its variants. Namely through modeling the game as a finite-state Markov decision process and deriving several theoretical and computational speed-ups, the authors outline an algorithm to solve the game. Ultimately, they find that “SALET” is the best starting word and that the best strategy solves the game in 3.42 moves in expectation. Results and analysis are also extended for recent variants and settings, such as Wordle Hard Mode or optimizing against a worst-case situation.

Data-Driven Compositional Optimization in Misspecified Regimes

With a manifold growth in the scale and intricacy of systems, the challenges of parametric misspecification become pronounced. These concerns are further exacerbated in compositional settings, which emerge in problems complicated by modeling risk and robustness. In “Data-Driven Compositional Optimization in Misspecified Regimes,” Yang, Fang, and Shanbhag consider the resolution of compositional stochastic optimization problems, plagued by parametric misspecification. In considering settings where such misspecification may be resolved via a parallel learning process, the authors develop schemes that can contend with diverse forms of risk, dynamics, and nonconvexity. They provide asymptotic and rate guarantees for unaccelerated and accelerated schemes for convex, strongly convex, and nonconvex problems in a two-level regime with extensions to the multilevel setting. Surprisingly, the nonasymptotic rate guarantees show no degradation from the rate statements obtained in a correctly specified regime and the schemes achieve optimal (or near-optimal) sample complexities for general T -level strongly convex and nonconvex compositional problems.

Looking Beyond Mean Delay in Queue Scheduling

A central aim of queueing theory is to design scheduling policies that reduce delays. Much of literature focuses on minimizing *mean* delay, but, in practice, it can be more important to reduce the delay *tail*—that is, the probability of especially large delays. There is a subtle interplay between these objectives. For instance, the SRPT policy (shortest remaining processing time) prioritizes short jobs over long jobs, which optimizes mean delay. Sometimes, SRPT also has optimal delay tail, but other times, its delay tail is *pessimal*. Thus, good mean delay need not imply good delay tail.

In “When Does the Gittins Policy Have Asymptotically Optimal Response Time Tail in the $M/G/1?$,” Scully and van Kreveld explore the interplay between mean delay and delay tail for systems with *unknown job durations*. The article studies the *Gittins policy*, which optimizes mean delay in this setting. Curiously, its delay tail can be optimal, pessimal, or in between, depending on the job duration distribution. Fortunately, the worst case can be avoided: the article shows how to modify Gittins to avoid pessimal delay tail while maintaining near-optimal mean delay.

Improving Operating Room, Surgery, and Anesthesiologist Scheduling Using Stochastic and Robust Optimization

Efficient planning and scheduling of operating room (OR) activities is crucial for managing costs and delivering high-quality surgical care. However, this task is extremely complex for several reasons. First, it requires coordinating multiple resources, such as ORs and anesthesiologists. Second, in addition to limited OR capacity and time, there is a significant shortage of anesthesiologists required to perform surgeries. Third, surgery durations are uncertain and difficult to predict. Ignoring such uncertainty may lead to substantial overtime, idling, and surgery delays, among other schedule deficiencies. Thus, hospital managers could benefit greatly from advanced methodologies to improve OR utilization, surgical care, and quality as well as to minimize OR operational costs. In “Stochastic Optimization Approaches for an Operating Room and Anesthesiologist Scheduling Problem,” Tsang, Shehadeh, Curtis, Hochman, and Brentjens propose computationally tractable stochastic programming and distributionally robust optimization methodologies for an integrated allocation, assignment, sequencing, and scheduling problem under uncertainty involving multiple ORs, anesthesiologists, and surgery types. Using real-world surgery data and a case study from a health system in New York, they conduct extensive experiments demonstrating the computational efficiency of the proposed methodologies, allowing for their implementation in practice. Moreover, they show the negative consequences of adopting the existing non-integrated approaches and provide valuable practical insights.

New Robust Optimization Models Tackle Ambiguity in High-Risk Decision Making

Decision making under uncertainty involves both ambiguity and risk. In “Robust CARA Optimization,” Chen and Sim have developed innovative optimization models designed for ambiguity-averse decision makers whose risk preference is consistent with constant absolute risk aversion (CARA). The research delves into maximizing the worst-case expected exponential utility amid uncertainties from independent factors with ambiguous marginals. To enhance computational feasibility, the authors developed a series of approximations: starting with tractable concave functions in affinely perturbed cases, advancing to concave piecewise affinely perturbed scenarios, and culminating in novel multi-deflected linear decision rules for adaptive optimization. This comprehensive framework extends to a multi-period consumption model, ultimately forming an exponential conic optimization problem efficiently solvable with existing solvers. Practical applications demonstrated in project and multiperiod inventory management highlight the models’ potential to surpass existing stochastic optimization methods, especially under high risk aversion.

A New Practical Algorithm Enables Optimal Preconditioning

A classic and important question in optimization and numerical methods is how to find diagonal preconditioners to maximally

reduce the condition number of any matrix with full rank. Until recently, few practical methods could handle large sparse matrices. In “Optimal Diagonal Preconditioning,” Qu, Gao, Hinder, Ye, and Zhou show that this problem can be modeled using quasiconvex optimization and semidefinite programming. Leveraging these insights, they develop algorithms to efficiently find optimal diagonal preconditioners for large sparse systems. They find that although heuristic diagonal preconditioners are popular in practice, their performance at reducing condition numbers could have a significant gap to optimal diagonal preconditioners. This work provides theoretical foundation for future works on optimal preconditioning as well as practical implementations that could be used to build more sophisticated software for optimal preconditioning at scale. A main advantage of the framework in this paper is its potential to be scaled up to handle even larger matrices, which is an exciting direction for numerical optimization.

Designing Fair and Incentive-Compatible Matching Systems

Many service systems face the challenge of dynamically matching customers who have heterogeneous preferences with an available pool of service providers in a fair and equitable manner. The first come, first served (FCFS) service discipline offers a straightforward and fair approach to allocate arriving customers to available servers and manage congestion. However, FCFS may limit the ability of service providers to reach an efficient allocation based on the individual preferences of customers and the characteristics of servers. In “Designing Service Menus for Bipartite Queuing Systems,” Caldentey, Hillas, and Gupta develop a queueing framework to address this dynamic matching problem. Their proposed methodology enables service providers to offer a menu of different service classes, each served by a distinct pool of servers. By thoughtfully designing this menu, a service provider can encourage self-interested customers to select service classes in a way that optimizes the trade-off between quality of matching and waiting time delays.

Black-Box Quantile Optimization via Finite-Difference-Based Gradient Approximation

Risk management necessitates consideration of metrics such as quantiles to supplement conventional mean performance measures. In “Quantile Optimization via Multiple-Timescale Local Search for Black-Box Functions,” Hu, Song, and Fu consider the problem where the goal is to optimize the quantile of a black-box output. They introduce two new iterative multitimescale stochastic approximation algorithms utilizing finite-difference-based gradient estimators. The first algorithm requires $2d + 1$ samples of the black-box function per iteration, where d is the number of decision variables. The second employs a simultaneous-perturbation-based gradient estimator that uses only three samples per iteration, irrespective of the number of decision variables. The authors prove strong local convergence for both algorithms and analyze their finite-time convergence rates through a novel fixed-point argument. These algorithms perform well across a varied set of benchmark problems.

A New Optimization Methodology Helps Allocate Vaccines to Combat a Pandemic

During the COVID-19 pandemic, the extraordinary speed of vaccine developments quickly gave rise to a critical operational problem: How to distribute a scarce stockpile of vaccines across communities? In “Branch-and-Price for Prescriptive Contagion Analytics,” Jacquillat, Li, Rame, and Wang address this question by presenting a methodology to allocate shared resources across subpopulations governed by contagion dynamics. This problem combines the difficulties of mixed-integer nonconvex optimization and

those of optimization with constraints governed by ordinary differential equations. By combining novel column generation, approximate dynamic programming, and branch-and-bound elements, the authors’ methodology can solve large and otherwise-intractable instances, outperforming state-of-the-art benchmarks. From a practical standpoint, their approach can significantly enhance the effectiveness of a vaccination campaign. Ultimately, the results from the paper outline vaccine allocation—beyond vaccine development and vaccine manufacturing—as a critical lever to mitigate the impact of a pandemic on public health.

Building Sparse and Temporally or Spatially Linked Regression Models

The framework of slowly varying regression under sparsity addresses many problems in machine learning where the underlying model is sparse and varies slowly and sparsely. This includes problems with temporally or spatially varying structures. For example, in the temporal case, the factors important in predicting the energy consumption in a building can vary depending on the hour of the day or the period of the year. In the spatial case, the factors that affect house prices can differ by neighborhood. In “Slowly Varying Regression Under Sparsity,” Bertsimas, Digalakis, Jr., Li, and Lami rigorously formulate such models and propose new theories, algorithms, and software for solving them. The authors thoroughly evaluate their framework on synthetic and real-world case studies and make their implementations available open-source to facilitate adoption by practitioners.

Handling Heterogeneous Machines in Malleable Scheduling

Parallelization is an important and widespread technique to speed up the completion of time-critical tasks, not only in high-speed computing, but also in operations planning in production and logistics. A fundamental model in this context is that of *malleable jobs*, each of which can be assigned to a subset of machine for parallel processing. In “Assigning and Scheduling Generalized Malleable Jobs Under Subadditive or Submodular Processing Speeds,” Fotakis, Matuschke, and Papadigenopoulos go beyond the by now well-understood identical-machine setting in malleable scheduling and develop algorithmic approaches for scheduling malleable jobs under various discrete concavity assumptions on the joint processing speeds of the assigned (possibly very heterogeneous) machines. They show that under these assumptions, the task of finding a *schedule* of small makespan can be reduced to that of finding an *assignment* with small maximum machine load. For this latter problem, numerous efficient approximation algorithms are derived and their practical performance explored in a computational experiments. These results indicate that the computational challenges posed by parallelization in heterogeneous environments can indeed be overcome, enabling the optimization of heavily parallelized schedules in the aforementioned applications.

Adaptive Lagrangian Policies for Inventory Control

Consider the problem of managing inventory in a multiwarehouse, multistore setting where each store can be periodically replenished from potentially a set of nearby warehouses. This is a practically relevant problem, and its optimal policy is not easy to compute because of the curse of dimensionality. Existing works in the literature have proposed a policy based on Lagrangian relaxation of the original stochastic problem. This vanilla policy applies the control parameters derived from the Lagrangian model and has been shown to perform well in some cases. In “Adaptive Lagrangian Policies for a Multiwarehouse, Multistore Inventory System with Lost Sales,” Chao, Jasin, and Miao go one step farther

and develop adaptive Lagrangian policies that update (at certain update times) the control parameters of the vanilla policy based on the historical demand realizations. The authors analytically show that their proposed adjustment significantly improves the performance of the vanilla policy.

Balancing Speed and Value in On-Demand Matching Platforms

In “Matching Impatient and Heterogeneous Demand and Supply,” Aveklouris, DeValve, Stock, and Ward consider a fundamental trade-off faced by many platforms (e.g., Uber/Lyft) that match supply (e.g., drivers) and demand (e.g., riders) dynamically over time: making matches quickly capitalizes on the value of current supply and demand in the system, whereas waiting may enable better matches at the risk of losing impatient customers. They show that this trade-off can be balanced by waiting a short amount of time before making matches: long enough to gather enough agents to make valuable matches but not so long that impatient agents are likely to leave. Intuitively, this balance depends on how long agents are willing to wait, on average, but the authors show that it also depends on the full distribution of the willingness to wait (i.e., not only mean, but also variance and higher moments play a role). Thus, approaches that only take into account the mean willingness to wait may perform quite poorly. Further, the authors develop an algorithm to rank matching priorities in order to achieve an optimized trade-off between speed and value of matches.

Second-Price Auctions Are Robustly Optimal

Market designers often only have partial information about the environment and prefer simple mechanisms that are robust to the underlying uncertainty. In “On the Robustness of Second-Price Auctions in Prior-Independent Mechanism Design,” Anunrojwong, Balseiro, and Besbes study the fundamental problem of selling an item to n buyers, when only an upper bound on values is known and the seller minimizes worst-case regret. They establish that the same mechanism is robustly optimal across a wide range of environments for distributions of values: i.i.d., mixtures of i.i.d., exchangeable, and affiliated (the last two capturing positive dependence). Moreover, this robust mechanism is a second-price auction with random reserve price depending on n . Without positive dependence, the problem reduces to the one-buyer case. This result supports the wide use of second-price auctions in

practice and allows them to quantify the robust value of competition in auctions.

Feasible Online Learning with Gradient Feedback

Online gradient descent (OGD) is well-known to be doubly optimal under strong convexity or monotonicity assumptions: (1) in the single-agent setting, it achieves an optimal regret of $\Theta(\log T)$ for strongly convex cost functions, and (2) in the multiagent setting of strongly monotone games with each agent employing OGD we obtain last-iterate convergence of the joint action to a unique Nash equilibrium at an optimal rate of $\Theta(\frac{1}{T})$. Whereas these finite-time guarantees highlight its merits, OGD has the drawback that it requires knowing the strong convexity/monotonicity parameters. In “Adaptive, Doubly Optimal No-Regret Learning in Strongly Monotone and Exp-Concave Games with Gradient Feedback,” Jordan, Lin, and Zhou design a fully adaptive OGD algorithm, AdaOGD, that does not require a priori knowledge of these parameters. In the single-agent setting, the algorithm achieves $O(\log^2(T))$ regret under strong convexity, which is optimal up to a log factor. Further, if each agent employs AdaOGD in strongly monotone games, the joint action converges in a last-iterate sense to a unique Nash equilibrium at a rate of $O(\frac{\log^3 T}{T})$, again optimal up to log factors. The algorithms are illustrated in a learning version of the classic newsvendor problem, in which, because of lost sales, only (noisy) gradient feedback can be observed. The results immediately yield the first feasible and near-optimal algorithm for both the single-retailer and multiretailer settings.

Tight Guarantees for Multiunit Prophet Inequalities and Online Stochastic Knapsack

In “Tight Guarantees for Multiunit Prophet Inequalities and Online Stochastic Knapsack,” Jiang, Ma, and Zhang analyze the multiunit prophet inequalities and online knapsack problems. They formulate the problems as the online contention resolution scheme problems. Then, they develop linear programming formulations to represent the online contention resolution scheme problems. They characterize the optimal solution of the linear programming and derive the corresponding optimal policies. For the multiunit prophet inequalities, they are able to obtain the tight ratios, which improved over the previous ones. For the online stochastic knapsack problem, their ratio is also optimal and the best up to date. They finally generalize their results to the multiresource setting and expand the applicability of their approach.