



## Operations Research

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

### In This Issue

To cite this article:

(2022) In This Issue. Operations Research 70(3):iii-viii. <https://doi.org/10.1287/opre.2022.2297>

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

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

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

Copyright © 2022, 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.2022.2297>

Copyright: © 2022 INFORMS

### Introduction: Special Issue Honoring Kenneth Arrow

When our successors write about the first century of Operations Research, the name of Kenneth Arrow will be in lights. It is fitting that at this time, not long after his death, that we take a moment to consider Ken Arrow's academic legacy. In "Introduction: Special Issue Honoring Ken Arrow," Abbas and Bell highlight his contributions in the field of Operations Research and summarize the papers published in the special issue that speak to his legacy.

### Two Grades are Too Few in Voting!

In "Majority Judgment vs. Approval Voting," Balinski and Laraki show the two main criticisms by Approval Voting (AV) supporters have been that Majority Judgment (MJ) is not Condorcet'; and that it admits the no show" paradox. That MJ is not Condorcet-consistent is a good property shared with AV: the domination paradox shows majority rule may well err in an election between two. Whereas the no-show paradox is in theory possible with MJ it is as a practical matter impossible. Moreover, it is proven that MJ with three grades does not admit the no-show paradox. In contrast, AV suffers from serious drawbacks. With AV, voters cannot express their opinions adequately; experiments show that Approve is not the opposite of Disapprove; and AV does not admit the no-show paradox it admits the very closely allied no-show syndrome." Two grades are simply too few! The debate must concern three or more grades.

### How Does Model Uncertainty Affect the Demand for Learning?

In many situations, decisions under uncertainty are made with incomplete confidence about the probability law describing the environment (model uncertainty, or ambiguity), but with the opportunity to learn, at some cost, before choosing between actions. In "Optimal Learning Under Robustness and Time-Consistency," Epstein and Ji use a stylized model to derive the optimal demand for learning (or sampling) and how it is affected by model uncertainty. A presumption is that the decision-maker is conservative and seeks to make robust decisions. One surprising result is that, in a sequential hypothesis testing context, sensitivity analysis overstates the robustness value of sampling. In a different context where there exists an unambiguous (purely risky) action, but where learning is not feasible, Ellsberg has highlighted the intuition for choosing the risky option over ambiguous ones. In contrast, the authors demonstrate that where learning is feasible and optimal, then the risky action cannot be optimal.

### A Local Index of Absolute Risk Aversion in Rank-Dependent Utility

The well-known Pratt-Arrow approximation, developed independently by John W. Pratt and Kenneth Arrow, provides an insightful dissection of the risk premium under the expected

utility (EU) model. It is given by one-half the product of the variance of the risk and the local index of absolute risk aversion of the decision-maker. Quite surprisingly, despite many important developments on 'global' risk aversion in non-EU models, the 'local' approach to risk aversion has received little attention outside EU. By considering the first two dual moments — mean and maxiance — on equal footing with the first two primal moments — mean and variance —Eeckhoudt and Laeven develop in "Dual Moments and Risk Attitudes" a dissection of the risk premium under the popular rank-dependent utility (RDU) model. This yields a simple approximation to the risk premium and a local index of absolute risk aversion under the RDU model.

### Scenario-Focused Decision Analysis and Savage's Small Worlds

One of the strengths of decision analysis is that it can deal with most uncertainties; but, alas, not all. Sometimes uncertainties are too deep: i.e., within the time and data currently available, no agreement is possible between decision makers, experts, and stakeholders on their quantification as probabilities. The possible range of probabilities may be so great that any sensitivity study would show virtually all actions may be optimal. Recently, such cases have been approached by scenario-focused decision analyses in which the deep uncertainties are fixed at several 'interesting' values. These approaches are showing considerable potential, but there is a problem. The assumptions on which of decision analysis is based do not necessarily apply, because scenarios are not quite 'small worlds' in Savage's sense. In "Axiomatizing the Bayesian Paradigm in Parallel Small Worlds," French discusses the difficulty, offers a way forward and demonstrates some of the points within an example on nuclear energy strategy.

### Severe Free Rider Problems with Investment in the Common Good

Many businesses and policy makers encounter a free rider problem when they have to decide when and how much to invest in the common good shared by other stakeholders. Although the extant theoretical literature on the management of the common good has shown a mild form of the free rider effect, in "Game of Variable Contribution to the Common Good under Uncertainty," Kwon shows that a simple game-theoretic model of investment in the common good can result in a war of attrition, which is a severe form of the free rider effect. As a result, each stakeholder would invest in the common good much more gradually than a socially optimal policy prescribes. The article also shows that if the stakeholders are asymmetric and the environment of the common good has randomness, this free rider effect can be significantly mitigated.

## When People Undertake Too Little Prevention

Prevention efforts, such as quitting smoking, flu vaccination, and exercising are of crucial importance in health policy, but people tend to undertake too little of them. The main reason is that most prevention efforts only reduce but do not completely eliminate the risk of poor health. This makes it harder for people to assess the benefits of prevention as they tend to misperceive and transform probabilities. In “When Risk Perception Gets in the Way: Probability Weighting and Underprevention,” Baillon, Bleichrodt, Emirmahmutoglu, Jaspersen, and Peter introduce psychological insights (probability weighting) in a model of optimal decision making and show that most people undertake too little prevention when the risk of poor health is between 10% and 80%. The authors discuss several policy measures to make people spend more on prevention.

## Managing Congestion When Acquiring Information

In “Search Under Accumulated Pressure” Alizamir, de Véricourt, and Sun explore how time pressure in the form of task accumulation might affect the information gathering process of a Decision Maker. Decision Makers often need to balance the benefit of improving their belief about the best choice against the cost of acquiring more information. In many instances, the cost to acquire further information depends on the size of the accumulated workload. The authors formulate this problem as a Partially Observable Markov Decision Process, and characterize the optimal control policy. The analysis reveals that the highest workload the Decision Maker is willing to sustain may increase in the middle of the information search. This runs counter to the more intuitive argument that as the search for information on a task progresses, the Decision Maker becomes less willing to sustain a high level of workload.

## The Design of the Procurement Process

Procurement strategies and techniques that reduce Cost of Goods Sold play a pivotal role in companies’ financial performance. While procurements in practice follow complicated procedures, most existing studies focus on contract negotiations or auctions, ignoring the implications of these procedures. In “The Strategic Benefit of Request for Proposal/Quotation,” Chu, Rong, and Zheng scrutinize the buyer’s solicitation and suppliers’ tendering process, i.e., request for proposal/quotation procedure (RFx). The RFx procedures almost always precede the contract negotiation even when the buyers have accurate cost estimates. In a complete information setting with two imperfectly substitutable suppliers, the research finds that the seemingly innocuous RFx stage (weakly) reduces the buyer’s purchasing cost. This is because the RFx procedure enables the buyer to endogenize the ensuing bargaining sequence and create supplier competition, as both suppliers would like to be the leader in the bargaining procedure.

## Sequential Decision Making Using Quantiles

The goal of a traditional Markov decision process (MDP) is to maximize the expectation of cumulative reward over a finite or infinite horizon. In many applications, however, a decision maker may be interested in optimizing a specific quantile of the cumulative reward. For example, a physician may want to determine the optimal drug regime for a risk-averse patient with the objective of maximizing the 0.10 quantile of the cumulative reward; this is the cumulative improvement in health that is expected to occur with at least 90% probability for the patient. In “Quantile Markov Decision Processes,” Li, Zhong, and Bran-deau provide analytic results to solve the quantile Markov

decision process (QMDP) problem. They develop an efficient dynamic programming procedure that finds the optimal QMDP value function for all states and quantiles in one pass. The algorithm also extends to the MDP problem with a conditional value-at-risk objective.

## Exact Penalization of Generalized Nash Equilibrium Problems

The Generalized Nash Equilibrium Problem (GNEP) extends the classical Nash Equilibrium Problem (NEP) by allowing individual players’ constraints, in addition to the payoff functions, to depend on rival players’ strategies. This extension adds considerably descriptive and explanatory power for the GNEP to model non-cooperative competition among multiple decision-making agents. Since its introduction in the 1950’s with the first model being the renowned Arrow-Debreu model of an abstract economy, the GNEP has risen beyond its mathematical-economic origin and become a common paradigm for various engineering disciplines. In “Exact Penalization of Generalized Nash Equilibrium Problems,” Ba and Pang study an exact penalization approach to the GNEP by penalizing the rival-dependent coupled constraints, thereby converting the GNEP into a NEP and facilitating the application of the theory and methods for the latter to address the former more challenging problem.

## On Solving a Class of Continuous Traffic Equilibrium Problems and Planning Facility Location Under Congestion

In “On Solving a Class of Continuous Traffic Equilibrium Problems and Planning Facility Location Under Congestion,” Wang, Ouhang, and She present methods to obtain analytical solutions to a class of continuous traffic equilibrium problems, where continuously distributed customers from a bounded 2-dimensional region travel along least-cost paths to reach one of several discretely located facilities for service. Closed-form solution can be obtained for linear cost-flux functions when the service region is simply connected and can be mapped into a certain regular shape through conformal mapping. These results shed light on interesting properties of traffic equilibrium in a continuous space, and also serves as the basis for service facility location design. Numerical and empirical examples are presented to illustrate application of the proposed traffic equilibrium and facility location design models.

## The Sufficiency of Local Information in Managing Supply Chains with Capacity Limits

Supply chains are distributed across multiple locations, and to effectively manage inventory in these channels requires knowledge of inventory levels at many sites. Such information is changing dynamically, so it is unrealistic to expect such information to be readily available, especially in channels with production capacity limits. In “Conveying Demand Information in Serial Supply Chains with Capacity Limits,” Kapuściński and Parker show that local information alone is sufficient to effectively manage inventory in capacity-limited supply chains. In supply chains with capacity limits, the retailer does not faithfully pass the demand information to her supplier but sends censored information in her order. Despite censoring, these orders contain sufficient information for the modified echelon base-stock policy, a full-information policy, to be mimicked. They provide evidence using numerical experiments and analytical bounds that the modified echelon base-stock policy performs superbly. The authors also describe information requirements used in supply chain literature and demonstrate that with an

incentive-compatible mechanism, similar to Lee and Whang (2000), local managers will follow the centralized inventory policy.

### Parallel Innovation Contests

Crowdsourcing using innovation contests has become a popular tool to source innovative solutions to various problems organizations face. With several innovation contests that run in parallel, two key questions are how the number of contests affects contest outcomes and whether solvers who can compete in these contests should focus their effort on a small group of contests. In “Parallel Innovation Contests” E. Körpeoğlu, C. G. Korpeoglu, and Hafalır address these key questions. The authors show that running up to a certain number of contests in parallel can benefit the overall outcome of these contests. Interestingly, more contests can be run in parallel if these contests seek disruptive innovation rather than incremental innovation. They also show that encouraging solvers to work on multiple contests in parallel can improve contest outcomes when these contests seek disruptive innovation. These findings can help guide organizations and crowdsourcing platforms that organize multiple contests in parallel.

### How Should Short-Haul Rail Be Used in Metropolitan Container Transportation?

When shipping ports are colocated with major population centers, the exclusive use of road transport for moving shipping containers across the metropolitan area is undesirable from both social and economic perspectives. Port shuttles, an integrated road and short-haul rail transport modality, are thereby gaining significant interest from governments and industry alike, especially in the Australian context. In “A Simultaneous Magnanti-Wong Method to Accelerate Benders Decomposition for the Metropolitan Container Transportation Problem,” Perrykkad, Ernst, and Krishnamoorthy explore the mathematics behind the optimal integration of road and port shuttle modalities for container transportation in metropolitan areas, including proofs of NP harness, a Benders decomposition, and an extensive computational study. Critically, to accelerate their Benders decomposition the authors develop the simultaneous Magnanti-Wong method: an extension of the classical Magnanti-Wong acceleration that preserves this problem’s important network substructure. In addition to the problem at hand, this technique shows promise more generally for Benders decompositions with special subproblem structure.

### Data-driven Ordering in the Face of Fixed Setup Costs

When a new product has just been introduced or the economy has just entered a new phase, a firm is often at a loss as to what the underlying demand pattern has become let alone how best to respond to it. In “Dynamic Inventory Control with Fixed Setup Costs and Unknown Discrete Demand Distribution,” Davoodi, Katehakis, and Yang faced off this challenging problem by tailoring ordering decisions to empirical distributions formed out of past demand observations. In the presence of fixed setup costs, however, an (s,S) policy optimal in the conventional known-distribution setting would take many periods for its long-term benefit to be realized. Therefore, a good online policy has to balance between letting ordering decisions settle for long periods and adjusting them frequently to take advantage of newly available information. When properly balanced, such policies could indeed achieve tight bounds for the performance measure of regret.

### Assessing Distorted Information in Modern Warfare

Distorted information (misinformation and disinformation) has long been a part of warfare (see the writings of Sun-Tzu). However, the study of the ever-increasing use of distorted information in modern warfare has been rather limited. In “Misinformation and Disinformation in Modern Warfare,” Chang, Keblis, Li, Iakovou, and White model instances of today’s battlespace as a partially observable game with three agents, a leader and two followers, and examine the benefit to the leader (e.g., a military command) of modulating the communication of information between: (i) followers who are adversaries, and (ii) followers who are allies. Counter to intuition, the study shows that only under certain conditions it is optimal for the leader to degrade (enhance) the quality of the information communicated between adversarial (allied) followers. The developed methodology is applied to warfare instances encountered in the Battle of Mosul.

### Portfolio Selection: Trading off Expected Growth with Systemic Risk

How can we construct portfolios that perform well in the face of systemic events? The global financial crisis of 2007–2008 and the coronavirus disease 2019 pandemic have highlighted the importance of accounting for extreme form of risks. In “Systemic Risk-Driven Portfolio Selection,” Capponi and Rubtsov investigate the design of portfolios that trade off tail risk and expected growth of the investment. The authors show how two well-known risk measures, the value-at-risk and the conditional value-at-risk, can be used to construct portfolios that perform well in the face of systemic events. The paper uses U.S. stock data from the S&P500 Financials Index and Canadian stock data from the S&P/TSX Capped Financial Index, and it demonstrates that portfolios accounting for systemic risk attain higher risk-adjusted expected returns, compared with well-known benchmark portfolio criteria, during times of market downturn.

### Assortment Optimization with a Continuum of Products

In “Continuous Assortment Optimization with Logit Choice Probabilities and Incomplete Information,” by Peeters, den Boer, and Mandjes consider a novel formulation of the classical assortment optimization problem with multinomial logit demand and unknown model parameters. The novelty lies in the fact that the set of products is not finite but a continuum, motivated by the desire to understand the problem characteristics for many products, as well as by applications where products are characterized by a continuous quality variable. For settings with and without capacity constraints, the authors design data-driven decision policies and prove upper and lower bounds on the regret, which imply that these policies are asymptotically optimal.

### Using Menu Pricing to Engage Consumers with Heterogeneous Advertising Sensitivity on Media Platforms

Ad-supported media platforms often engage price-sensitive consumers by offering content for free while reaping revenue from advertisers. Many consumers are also sensitive to advertising, however, so platforms need to carefully balance their advertising quantities with the prices they offer to consumers. Several platforms use menu pricing to achieve this goal: they offer consumers a choice between a free-use/ad-supported option or a paid-use/no-ad option. In “Optimal Price/Advertising Menu for Two-Sided Media Platforms,” DeValve and Pekeć use

Lagrangian duality to establish the optimality of such menus in a model of consumers with heterogeneous advertising sensitivity, offering both theoretical justification for their use in practice as well as guidance on when and how to set positive prices. Moreover, the optimal menu characterization implies that platforms may have more of an incentive to offer menu pricing under competition (further explaining their prevalence in the competitive digital media market) and platforms primarily cater to consumers with low advertising sensitivity (suggestive of the proliferation of online advertising observed in practice).

### **A New Learning Algorithm, with Applications to Inventory Management**

A fundamental yet notoriously difficult problem in operations management is the periodic inventory control problem under positive lead time and lost sales. More recently, there has been interest in the problem setting where the demand distribution is not known a priori and must be learned from the observations made during the decision-making process. In “Learning in Structured MDPs with Convex Cost Functions: Improved Regret Bounds for Inventory Management,” Agrawal and Jia present a reinforcement learning algorithm that uses the observed outcomes of past decisions to implicitly learn the underlying dynamics and adaptively improve the decision-making strategy over time. They show that, compared with the best base-stock policy, their algorithm achieves an optimal regret bound in terms of the time horizon and scales linearly with the lead time of the inventory ordering process. Furthermore, they demonstrate that their approach is not restricted to the inventory problem and can be applied in an almost black box manner to more general reinforcement learning problems with convex cost functions.

### **Platforms: How Much and What Should They Control?**

Platforms exert control in many ways: from determining which buyers are shown to which sellers to directly controlling allocation of aggregate supply across different consumer markets. How should platforms exercise control to increase social welfare? In “Transparency and Control in Platforms for Networked Markets,” Pang, Lin, Fu, Kleeman, Bitar, and Wierman examine the trade-offs that emerge between efficiency loss, transparency, and control across different platform designs. They show that open access platforms incentivize increased production toward perfectly competitive levels and limit efficiency loss, whereas controlled allocation designs can lead to supply-side withholding, resulting in unbounded efficiency loss. Balancing transparency and control, discriminatory access designs strictly improve upon the worst-case efficiency loss of both open access and controlled allocation platforms. Besides providing insights into the design and operation of two-sided platforms, this paper contributes to the growing theory of Cournot competition in networked markets.

### **Designing Incentives to Promote Treatment Adherence**

Premature cessation of antibiotic therapy is common and can severely compromise health outcomes, potentially leading to worsening health, disease transmission, and antibiotic resistance. In “Design of Incentive Programs for Optimal Medication Adherence in the Presence of Observable Consumption,” Suen, Negoescu, and Goh investigate the problem of designing a schedule of incentive payments to induce socially optimal treatment adherence levels in settings in which treatment adherence can be observed but patient preferences for treatment adherence

are heterogeneous and unobservable to a health provider. The novel elements of this problem stem from its institutional features: there is a single incentive schedule applied to all patients, incentive payments must be increasing in patients’ adherence, and patients cannot be a priori prohibited from any given levels of adherence. The authors develop models to design optimal incentives incorporating these features and conduct a numerical study of the tuberculosis epidemic in India.

### **Selecting Planning Objectives for Radiation Therapy: Integrating Inverse Optimization and Feature Selection**

The quality of radiation therapy treatment plans and the efficiency of the planning process are heavily affected by the choice of planning objectives. Although simple objectives enable efficient treatment planning, the resulting treatment quality might not be clinically acceptable; complex objectives can generate high-quality treatment, yet the planning process becomes computationally prohibitive. In “Objective Selection for Cancer Treatment: An Inverse Optimization Approach,” by integrating inverse optimization and feature selection techniques, Ajayi, Lee, and Schaefer propose a novel objective selection method that uses historical radiation therapy treatment data to infer a set of planning objectives that are tractable and parsimonious yet clinically effective. Although the objective selection problem is a large-scale bilevel mixed-integer program, the authors propose various solution approaches inspired by feature selection greedy algorithms and patient-specific anatomical characteristics.

### **Understanding How Transportation Bottlenecks Affect Commodity Prices**

The impact of transportation constraints in commodity markets has become increasingly relevant as many markets experience demand and supply growth that continues to outpace the growth in transportation infrastructure. In “Spatial Price Integration in Commodity Markets with Capacitated Transportation Networks,” Birge, Chan, Pavlin, and Zhu examine the relationship between the spatial distribution of commodity prices and the underlying transportation network that supports the flow of these commodities. The authors show that under mild assumptions, the prices between all locations must be bounded when there are no bottlenecks in the network. Conversely, a bottleneck can cause different locations to incur a congestion surcharge that pushes prices out of these bounds. The authors propose a time series analysis technique using mixed integer optimization that estimates the value of these surcharges from commodity prices and then apply the technique to study how gasoline prices changed after a series of well-documented supply chain disruptions in the Southeastern U.S. gasoline market.

### **Efficient Fair Division with Minimal Sharing**

When assets are to be divided among several partners, for example, a partnership split, fair division theory can be used to determine a fair allocation. The applicability of existing approaches is limited as they either treat assets as divisible resources that end up being shared among participants or deal with indivisible objects providing only approximate fairness. In practice, sharing is often possible but undesirable, and approximate fairness is not adequate, particularly for highly valuable assets. In “Efficient Fair Division with Minimal Sharing,” Sandomirskiy and Segal-Halevi introduce a novel approach offering a middle ground: The number of shared objects is minimized while maintaining exact fairness and economic efficiency. This minimization can be conducted in polynomial time for generic instances if the number of agents or objects is fixed. Experiments

on real data demonstrate a substantial improvement over current methods.

### Pricing and Optimization in Shared Vehicle Systems: An Approximation Framework

The optimal management of shared vehicle systems, such as bike-, scooter-, car-, or ride-sharing, is more challenging compared with traditional resource allocation settings because of the presence of spatial externalities—changes in the demand/supply at any location affect future supply throughout the system within short timescales. These externalities are well captured by steady-state Markovian models, which are therefore widely used to analyze such systems. However, using Markovian models to design pricing and other control policies is computationally difficult because the resulting optimization problems are high dimensional and nonconvex. In “Pricing and Optimization in Shared Vehicle Systems: An Approximation Framework,” Banerjee, Freund, and Lykouris design a framework that provides near-optimal policies, for a range of possible controls, that are based on applying the possible controls to achieve spatial balance on average. The optimality gap of these policies improves as the ratio between supply and the number of locations increases and asymptotically goes to zero.

### Designing Dashboards for Valid Inference from Online A/B Tests

In “Always Valid Inference: Continuous Monitoring of A/B Tests,” Johari, Koomen, Pekelis, and Walsh study the design of dashboards for analysis of randomized controlled experiments, known as A/B tests in online platforms. A typical application of A/B testing is to compare a new version of a web page to an existing version, or a new version of a feature to an existing feature. A/B tests are typically analyzed via dashboards that display standard frequentist inferential quantities (i.e., p-values and confidence intervals). However, a significant practical concern is that A/B tests are typically \*continuously monitored\* by experimenters, leading to invalid inference. Leveraging insights from sequential hypothesis testing, the authors report on a novel approach to computing “always valid” p-values and confidence intervals that are robust to continuous monitoring of A/B tests by experimenters. Their results suggest that always valid measures also admit more efficient detection of true effects, and naturally combine with common procedures for control of the error rates in multiple hypothesis testing. In addition to the methodological contribution, the paper describes the successful use of such measures in a large-scale commercial A/B testing platform.

### Primal-Dual Analysis of Distributionally Robust Linear-Discrete Optimization with Marginals

In optimization problems, decisions are often made in the face of uncertainty that might arise in the form of random costs or benefits. In “Distributionally Robust Linear and Discrete Optimization with Marginals,” Chen, Ma, Natarajan, Simchi-Levi, and Yan study a robust bound of linear and discrete optimization problems in which the objective coefficients are random and the set of admissible joint distributions is assumed to be specified only up to the marginals. They provide a primal-dual formulation for this problem, and in the process, unify existing results with new results. They establish NP-hardness of computing the bound for general polytopes and identify two sufficient conditions—one based on a dual formulation, and one based on sublattices that provide a class of polytopes where the robust bounds are efficiently computable.

### A New Method for Dynamic Learning and Doing

For a large class of learning-and-doing problems, two processes are intertwined in the analysis: a forward process that updates the decision maker’s belief or estimate of the unknown parameter, and a backward process that computes the expected future values. The mainstream literature focuses on the former process. In contrast, in “Dynamic Learning and Decision Making via Basis Weight Vectors,” Zhang proposes a new method based on pure backward induction on the continuation values created by feasible continuation policies. When the unknown parameter is a continuous variable, the method represents each continuation-value function by a vector of weights placed on a set of basis functions. The weight vectors that are potentially useful for the optimal solution can be found backward in time exactly (for very small problems) or approximately (for larger problems). A simulation study demonstrates that an approximation algorithm based on this method outperforms some popular algorithms in the linear contextual bandit literature when the learning horizon is short.

### Decomposition Branching for Mixed Integer Programming

Applications of mixed integer programming can be found in many industries, such as transportation, healthcare, energy, and finance, and their economic impact is significant. It is also well known that mixed integer programs (MIPs) can be very difficult to solve. Their challenge continues to stimulate research in the design and implementation of efficient and effective techniques that can better solve them. In “Decomposition Branching for Mixed Integer Programming,” Yildiz, Boland, and Savelsbergh introduce a novel and powerful approach for solving certain classes of mixed integer programs (MIPs): decomposition branching. Two seminal and widely used techniques for solving MIPs, branch-and-bound and decomposition, form its foundation. Computational experiments with instances of a weighted set covering problem and a regionalized p-median facility location problem with assignment range constraints demonstrate its efficacy: it explores far fewer nodes and can be orders of magnitude faster than a commercial solver and an automatic Dantzig-Wolfe approach.

### Structural and Algorithmic Results for Matchings, Beyond Stability

Since the seminal work of Gale and Shapley, stable assignments have received widespread attention for their mathematical elegance and broad applicability. However, in applications such as the school choice problem, in which public schools are often perceived as commodities and only students’ welfare matters, enforcing stability implies a loss of efficiency for the students. In “Legal Assignments and Fast EADAM with Consent via Classical Theory of Stable Matchings,” Faenza and Zhang study two extensions of the traditional model—legal assignments and efficiency adjusted deferred acceptance mechanism (EADAM)—that strive to regain this loss in efficiency. The authors establish a tight connection between legal and stable assignments, which allows them to use critical structural tools of stable matchings, such as the concept of rotations, to design provably fast algorithms for (1) optimizing linear functions over the set of legal assignments and (2) finding the outcome of EADAM. These algorithmic results greatly improve the applicability of both extensions as witnessed by a complexity analysis and experimental results.

### Subsampling to Enhance Efficiency in Input Uncertainty Quantification

Quantifying the impact of input estimation errors in data-driven stochastic simulation often encounters substantial computational challenges due to the entanglement of Monte Carlo and input data noises. In “Subsampling to Enhance Efficiency in Input Uncertainty Quantification,” Lam and Qian propose a subsampling framework to bypass this computational bottleneck, by leveraging the form of the output variance and its estimation error in terms of data size and sampling effort. Compared with standard subsampling in the literature, our motivation is distinctly to reduce the sampling complexity of the two-layer bootstrap required in simulation uncertainty quantification. Compared with standard bootstraps, our subsampling approach provably and experimentally leads to more accurate variance and confidence interval estimations under the same amount of simulation budget.

### Dual Bounds of Sparse Principal Component Analysis

Sparse principal component analysis (PCA) is a widely used dimensionality reduction tool in machine learning and statistics. Compared with PCA, sparse PCA enhances the interpretability by incorporating a sparsity constraint. However, unlike PCA, conventional heuristics for sparse PCA cannot guarantee the qualities of obtained primal feasible solutions via associated dual bounds in a tractable fashion without underlying statistical assumptions. In “Using  $\ell_1$ -Relaxation and Integer Programming

to Obtain Dual Bounds for Sparse PCA,” Dey, Mazumder, and Wang present a convex integer programming (IP) framework of sparse PCA to derive dual bounds. They show the worst-case results on the quality of the dual bounds provided by the convex IP. Moreover, the authors empirically illustrate that the proposed convex IP framework outperforms existing sparse PCA methods of finding dual bounds.

### System Optimization and Competitive Equilibrium with Risk-Averse Agents

According to the classical welfare theorems of economics, markets with competing agents operating under well-known assumptions will deliver socially optimal outcomes. In “Dynamic Risked Equilibrium,” Ferris and Philpott show how these theorems can be extended to a dynamic stochastic setting in which perfectly competitive agents optimize using dynamic nested coherent risk measures. Random events are assumed to have finite probability distributions and are modeled using a scenario tree. The key ingredient in this theory is the assumption of a complete market for trading risk in each state of the world that enables the different risk measures of the agents to be integrated into a system risk measure that is used in the optimization of risk-adjusted expected social welfare. The correspondence between social optimization and competitive partial equilibrium enables regulators to monitor the performance of incomplete, imperfectly competitive markets compared with socially optimal benchmarks, and identify interventions that can be used to improve them.