



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(3):vi-x. <https://doi.org/10.1287/opre.2026.inthisissue.v74.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 © 2026, INFORMS

Please scroll down for article—it is on subsequent pages



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

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

In This Issue

<https://doi.org/10.1287/opre.2026.inthisissue.v74.n3>

Copyright: © 2026 INFORMS

Studying How Inventory Placement Shapes Online Fulfillment Performance

A growing share of e-commerce operations involves deciding not only how to fulfill incoming orders, but also where to position inventory beforehand. In “Online Demand Fulfillment Problem with Initial Inventory Placement: A Regret Analysis,” Arlotto, Keskin, and Wei examine how these two decisions interact. The authors introduce a joint regret framework that evaluates the performance of a fulfillment policy together with the initial allocation of inventory across warehouses. Their analysis shows that probabilistic fulfillment inevitably accumulates regret that grows with the length of the planning horizon, regardless of how inventory is initially placed. By contrast, when combined with an appropriate offline inventory placement, the score-based approach achieves a regret bound that remains stable over time and scales only with the size of the system. The study offers insight into the role of initial placement in shaping fulfillment policy performance.

Optimal Electric Vehicle Charging Schedules

With the adoption of electric vehicles emerges a need for coordinated charging strategies to cater to the increased and often synchronized energy demand without violating physical infrastructure limits. Common charging strategies repeatedly solve a corresponding offline problem. In “Relating Electric Vehicle Charging to Speed Scaling with Job-Specific Speed Limits,” Winterschermann, Antoniadis, Gerards, Hoogsteen, and Hurink thoroughly analyze that offline problem and present the flow-based offline charging scheduler (FOCS), an algorithm that solves the problem to optimality. The authors continue to quantify the approximation ratio of repeatedly solving the offline problem throughout the day as opposed to solving the offline problem with a priori knowledge of future arrivals. Numerical experiments confirmed that the found worst-case approximation ratio of 4 is far from the achieved empirical ratio of 1.3 when evaluating the ratio for real-world electric vehicle charging instances. Experiments based on that same real-world data show that FOCS only takes seconds to schedule 400 EVs in 15-minute granularity, which is competitive with commercial solvers.

Axioms for Automated Market Makers: A Mathematical Framework in FinTech and Decentralized Finance

Automated Market Makers (AMMs) are smart contracts that have emerged as a cornerstone of Decentralized Finance (DeFi), replacing traditional intermediaries for trading digital assets. However, the underlying principles, risks, and economic properties of these novel markets require a rigorous mathematical foundation. In “Axioms for Automated Market Makers: A Mathematical Framework in Fintech and Decentralized Finance,” Bichuch and Feinstein construct such a foundation. The authors propose a set of intuitive axioms for the utility functions that govern AMMs. From this framework, they characterize the resulting swap and pricing

properties and define a novel measure of price impact to quantify and compare transaction costs across different AMM designs. The work also addresses practical issues by proposing a new fee structure that ensures indifference to transaction splitting and analyzing the critical risk of divergence loss for liquidity providers, offering a comprehensive guide for designing and analyzing these new financial technologies.

Designing Effective Procurement Mechanisms for Assembly Systems Under Supply Uncertainty

In “Procurement for Assembly Systems Under Supply Uncertainty: Optimal Mechanisms,” Bu, Dawande, and Janakiraman explore how manufacturers can effectively source components when suppliers face uncertain yields and possess private information about their production costs. Such challenges are common in industries such as electronics and automotive manufacturing, in which a shortage of a single part can halt the entire assembly line. The authors develop an optimal procurement mechanism that helps manufacturers decide how much to source from each supplier and how to structure payments contingent on uncertain delivery quantities when suppliers’ production efforts cannot be directly observed. Their analysis reveals a practical menu-based contract under which each supplier selects terms best suited to its situation and the manufacturer subsequently determines order quantities for each component. The study provides both theoretical insights and actionable guidance for managing complex supply relationships in uncertain environments.

Pre-hedging

In “Pre-hedging,” Muhle-Karbe and Oomen study a dealer that pre-hedges an anticipated potential trade and analyses how this affects the client’s overall execution outcome. It shows that pre-hedging can benefit both parties: Improved risk management over an extended horizon then enables the dealer to charge reduced spreads that more than offset any adverse impact the pre-hedging activity has on the execution price. However, when a dealer pre-hedges too aggressively, this can be detrimental to the client. Timing uncertainty of the potential trade is an effective control held by the client to mitigate any counterproductive pre-hedging. Our results are robust to a setting where competing dealers simultaneously pre-hedge.

Multiasset Optimal Execution via Deep Learning for High-Dimensional Continuous-Time Stochastic Control

In “Deep Learning for High-Dimensional Continuous-Time Stochastic Optimal Control Without Explicit Solution,” Dupret and Hainaut introduce the generalized policy iteration physics-informed neural network, a novel deep learning algorithm for solving high-dimensional continuous-time stochastic optimal control problems even when the optimal control does not admit explicit solution. The method combines physics-informed neural networks with an actor-critic structure based on generalized

policy iteration and uses separate networks to approximate both the value function and the multidimensional optimal control. This approach provides a global approximation of the solution across time and space, enabling fast online evaluation. Theoretical guarantees on convergence and optimality are provided, whereas its accuracy and efficacy are empirically validated through two important numerical examples from operations research. Thereby, the authors generalize the Almgren–Chriss framework arising from optimal execution in finance by allowing both temporary and permanent price impacts to be fully nonlinear and by considering a multidimensional setting with multiple cointegrated assets.

From Discrete to Continuum: Risk Preference Heterogeneity in Traffic Equilibrium

How do travelers with heterogeneous risk preferences choose routes in congested networks with uncertain travel times? Conventional models characterize risk preference heterogeneity using a small number of discrete user classes. In “Mean-Risk Traffic Assignment Under the Continuously Distributed Risk-Aversion Factor,” Xu, Li, Xie, Chen, and Liu develop a new traffic equilibrium framework that directly accommodates a continuous distribution of risk aversion. The resulting model generalizes the classical Wardrop equilibrium to settings with both stochastic travel times and continuous preference heterogeneity, providing a more refined representation of real-world routing behavior. This study demonstrates that the continuous approach outperforms the discrete approach by not only eliminating the inherent discretization bias—that can mislead infrastructure investment decisions—but also delivers superior computational performance, requiring less run time and memory. More broadly, the framework could be potentially applied to a wide range of biobjective congestion routing games where user preferences are continuously distributed.

The Price of Attention: Ranking Products for Maximum Revenue

How should an online retailer rank products when customers have limited attention spans? In “Revenue Maximization and Learning in Product Ranking,” Chen, Li, and Yang tackle this classic problem by extending the well-known cascade model to account for two crucial, real-world factors: customers view only a random number of items, and the firm’s goal is to maximize revenue, not just clicks. This creates a difficult trade-off between ranking popular, low-price items and riskier, high-price ones. The authors propose the “Best- x ” algorithm, an efficient method for finding a near-optimal ranking. They prove it guarantees a revenue of at least $1/e$ (approximately 37%) of that achievable by a clairvoyant who knows each customer’s attention span in advance. For cases where product attractiveness and attention distributions are unknown, the authors also devise the RankUCB online learning algorithm, which learns personalized rankings from customer interactions and achieves near-optimal performance over time.

Platform Disintermediation: Information Effects and Pricing Remedies

In “Platform Disintermediation: Information Effects and Pricing Remedies,” Sekar and Siddiq analyze how platforms use pricing and informational levers to combat disintermediation, where sellers transact off-platform to avoid commission fees despite losing payment protection. The authors find that a platform facing a high threat of disintermediation may optimally raise its commission rates in a high-information environment, prioritizing revenue maximization from remaining secure transactions rather than preventing all off-platform transactions. Furthermore, they show that

increasing sellers’ switching costs may actually reduce platform revenue by enabling sellers to pass the cost onto on-platform buyers via higher prices. The analysis also demonstrates that a platform can benefit from providing a partially informative signal about buyer riskiness, as full information reduces the value of the platform’s protection and strengthens the disintermediation incentive. Finally, whereas access-based pricing (upfront fees) prevents disintermediation, it may yield less revenue than commissions when seller quality is highly heterogeneous.

From Loot Boxes to Better Design: Pricing Randomized Products

Randomized rewards—often described as “mystery packs” or “blind-box” mechanics—are now a familiar feature in many video games. Yet choosing a loot box’s price and the odds of each possible drop is not just a design decision; it is a challenging optimization problem. In “Algorithms for Loot Box Design,” Han, Ryan, and Tong develop an algorithmic framework for designing loot boxes by selecting a purchase price and item drop probabilities to maximize expected revenue under player choice behavior. They show that the general problem is computationally hard, but also identify economically motivated restrictions on player utilities that make the design problem tractable. When the number of items is fixed, the authors provide an exact polynomial-time algorithm under one class of utility structures and efficient approximation algorithms with provable guarantees under another. The analysis also links loot-box design to classic pricing ideas, offering guidance on how to translate item-level values and rarity into transparent, well-performing randomized reward systems.

From Big Data to Real-Time Decisions: Online Tensor Inference

Modern digital platforms, from e-commerce and online advertising to mobile health, generate massive streams of high-dimensional data that must be analyzed in real time. In “Online Tensor Inference,” Wen, Sun, and Zhang develop a new statistical framework that enables both efficient learning and rigorous inference for streaming tensor data. The authors propose an online low-rank tensor estimation method based on stochastic gradient descent that processes observations sequentially without storing historical data, overcoming the memory and scalability limitations of traditional offline approaches. Beyond estimation, the authors introduce a novel online debiasing technique that delivers valid confidence intervals and hypothesis tests on the fly without data splitting. Theoretical results establish near-minimax-optimal convergence rates and asymptotic normality for general linear functionals of tensors. Together, these advances provide a principled foundation for real-time, statistically grounded decision making in fast-changing, data-rich environments.

Operationalizing Semipersonalized Pricing

How can modern firms leverage feature information to set prices in way that is both profitable and practical? In “Optimal Feature-Based Market Segmentation and Pricing,” Cui and Hamilton address this question by analyzing feature-based market segmentation and pricing (FBMSP), a semipersonalized approach to pricing where firms use customer characteristics to group buyers and set segment-specific prices. Although businesses often rely on heuristic “segment-then-price” methods, in their article the authors show that under realistic statistical assumptions, the jointly optimal segmentation and pricing policy can be computed efficiently. Further, using structural results about the optimal FBMSP, the authors prove that semipersonalized pricing quickly converges to the performance of fully personalized pricing, motivating its use in practice. Finally, in a case study on U.S. home

mortgage data, they apply their method and show it significantly outperforms traditional heuristics, achieving near-maximal revenue with only a few segments. This research offers both practical tools and theoretical insights for firms navigating the balance between personalization and implementability in pricing.

Unifying the Regret Spectrum in Data-Driven Newsvendor

The data-driven newsvendor problem seeks to optimize inventory decisions using samples from an unknown demand distribution. Although this problem has attracted significant attention, previous studies have typically analyzed specific distribution classes or regret definitions in isolation. In “Survey of Data-driven Newsvendor: Unified Analysis and Spectrum of Achievable Regrets,” Chen and Ma present a unified analysis that synthesizes these settings and simplifies existing proofs. The study utilizes a notion of clustered distributions defined via the cumulative distribution function (CDF). This approach demonstrates that the achievable regret covers the entire spectrum of convergence rates between $1/\sqrt{n}$ and $1/n$. Beyond the theoretical unification, the authors show through simulations that this CDF-based notion accurately predicts the empirical regret and captures how the difficulty of the problem evolves with sample size. This work provides insights into understanding the value of data in the newsvendor problem and, more broadly, decision making under uncertainty.

Optimizing Prices for Viral Demand and Scarcity

Dynamic pricing models often assume that inventory and demand scale proportionally, but this “fluid” view breaks down when products go viral. In “Regime-Dependent Approximations for the Single-Item Dynamic Pricing Problem,” Abdallah and Reed investigate market extremes, specifically the “large market regime,” where inventory is scarce relative to surging demand. Their analysis reveals critical pitfalls in common heuristic approaches. The authors demonstrate that intuitive “price high and wait” policies are ineffective and, remarkably, that fluid static policies are not even first-order optimal in this context. Instead, they establish that a dynamic run-out-rate policy is essential to achieve both first- and second-order asymptotic optimality. Leveraging extreme value theory, this research provides a robust framework for managing severe supply-demand imbalances, ensuring that firms can effectively capture value in inventory-constrained environments.

Pricing with a Ripple Effect

How should firms price when each sale sparks the next? In “Dynamic Pricing Under Self-Exciting Arrival Processes,” Yuan, Du, and Hu study pricing decisions in markets where customer purchases actively stimulate future demand through word-of-mouth and social influence. Using a self-exciting (Hawkes) process to model the demand process, the authors show that optimal prices depend on both time and an “excitement level” summarizing accumulated customer influence. A key insight is that optimal prices may rise or fall with demand momentum, depending on whether the market is in a growth or saturation phase. The authors further demonstrate that simple, easy-to-implement deterministic (nonstationary) pricing heuristics can perform nearly as well as fully dynamic policies at large demand volumes. These results provide actionable guidance for firms operating in social media-driven markets, where this hour’s customers shape the next hour’s demand, and highlight the importance of explicitly accounting for the ripple effect of purchases in pricing design.

Balancing Risk and Robustness in Dynamic Decision Making

Many real systems, such as networks, finance, and safety-critical autonomy, must hedge against rare but costly events. Risk-sensitive control formalizes this idea by optimizing an exponential cost objective that prioritizes reliability over just average performance. Classical dynamic programming methods such as value iteration and policy iteration are well-understood in this risk-sensitive setting. However, modified policy iteration (MPI), which combines the strengths of both through partial policy evaluation, has lacked any theoretical understanding. In “On the Convergence of Modified Policy Iteration in Risk-Sensitive Exponential Cost Markov Decision Processes,” Murthy, Moharrami, and Srikanth address this gap. The authors analyze MPI for risk-sensitive Markov decision processes governed by a multiplicative Bellman equation, develops normalization and contraction tools suited to this setting, and proves both convergence and finite-time guarantees. The results provide a principled foundation for algorithms that combine computational efficiency with robustness, supporting the development of reinforcement learning methods that emphasize long-term reliability.

Bayesian Learning for Data-Driven Stochastic Optimal Control

Stochastic optimal control (SOC) provides a principled approach to dynamic decision making under uncertainty. It models the transition of the system state with a dynamic equation driven by randomness, with the assumption that the distribution of randomness is known. However, in practical problems, modeling randomness frequently depends on data or observations, which introduces uncertainty regarding the distribution of randomness. To address the data-driven SOC problem, in “Episodic Bayesian Optimal Control with Unknown Randomness Distributions,” Shapiro, Zhou, Lin, and Wang propose a new approach that incorporates Bayesian learning with optimal control in an episodic manner. They show the convergence and convergence rate results for their approach. They also develop an efficient computational method for a class of problems that have convex cost functions and linear state dynamics.

On Sinkhorn’s Algorithm and Choice Modeling

Choice modeling is an important topic that underlies a wide range of applications involving human decision making, and it traces its roots to the 1920s. Matrix balancing has an equally long history and wide applicability (e.g., in transportation and mobility networks). Recently, its celebrated Sinkhorn’s algorithm has been instrumental in the efficient approximation of optimal transport distances. However, the two topics have largely developed independently. In “On Sinkhorn’s Algorithm and Choice Modeling,” Qu, Galichon, Gao, and Ugander establish extensive connections between a class of Luce choice models and a common matrix-balancing problem. They leverage these connections to resolve open problems on the convergence of Sinkhorn’s algorithm for nonnegative matrices, characterizing its global linear convergence rate in terms of the algebraic connectivity and deriving the sharp asymptotic rate. The connections established in this paper between two seemingly unrelated topics help the transmission of ideas and lead to further interesting results.

Two Prices Unlock Big Gains for Reusable Resources

How much sophistication is needed to price reusable resources, like hotel rooms and cloud computing, when usage durations are *not memoryless*? *Surprisingly little*. In “Dynamic Pricing for Reusable Resources: The Power of Two Prices,” Balseiro, Ma, and

Zhang propose a class of dynamic *stock-dependent* policies that achieve significant improvements over static pricing by only looking at how many units are busy and ignoring how long they have been busy. Using an “insensitivity” property of loss networks, they show that optimizing within this policy class can be formulated as a tractable convex optimization problem. Better yet, the performance loss of the optimal stock-dependent policy can be achieved by a simple two-price policy: charge a *high* price when inventory falls below a threshold and a *low* price otherwise. Extensions to multiple resources and customer classes, together with extensive simulations, confirm that “just a little” dynamicity can go a long way.

Strategic Servers with Individual Preferences in Heavy Traffic

Agents in service systems are known to strategically adjust their service speeds to maximize idle time. “Many-Server Queueing Systems with Heterogeneous Strategic Servers in Heavy Traffic” by Büke, dos Reis, and Platonov is the first work to rigorously examine strategic heterogeneous servers in queueing systems. The authors address this technically challenging problem using a novel asymptotic framework that provides valuable insights into equilibrium behavior. The analysis goes beyond existing literature by incorporating a remarkably general class of utility functions and routing policies, allowing for variation in agents’ marginal utility of idleness. One of the key results shows that the celebrated square-root staffing rule retains its optimality—even in the presence of strategic behavior—when agents are sufficiently averse to low levels of idleness.

Online Planning in Nonstationary Environments

In “Online Planning in Nonstationary Environments,” Cheung and Lyu develop a new computational approach to tackle a central challenge in online planning: balancing real-time decisions with long-term performance in dynamic systems. Traditional solutions often rely on intractable dynamic programming, but the new method delivers near-optimal performance efficiently, even in nonstationary environments. The framework addresses a broad class of planning problems with concave objectives and convex constraints, applicable to real-world scenarios like coupon assignment, order fulfillment, and resource allocation. The study considers two decision-making settings: one with unlimited access to data samples via simulations and another with finite sample data. By leveraging gradient-based insights from offline simulations, the researchers propose an offline-to-online framework that performs well in both settings. Notably, in the sampling setting, the approach improves as more data or longer planning horizons become available. Numerical tests on coupon assignment and supply chain management problems show significant gains over existing methods.

Fairness in Online Selection Problems

Two of the most studied models in online decision making are the secretary problem and the prophet inequality problem. Both capture the challenge of making irrevocable choices under uncertainty. But, what happens when candidates come from different groups and fairness enters the picture? In “Fairness and Bias in Online Selection,” Correa, Cristi, Dütting, and Norouzi-Fard introduce and analyze multicolor variants of these problems. In these models, each candidate belongs to a “color,” and comparisons are only meaningful within the same color. This captures real-world situations where crossgroup rankings are unreliable or biased—for instance, when evaluating students from different schools or job applicants from diverse backgrounds. For the

multicolor secretary problem, the authors characterize the optimal online algorithm. In contrast to the offline optimum—which always selects from the most promising group—the optimal online algorithm is inherently fairer. For the multicolor prophet inequality, the authors design algorithms that enforce target selection probabilities across groups, ensuring equitable treatment.

Optimizing School Seat Allocation to Improve Access and Fairness

A growing shortage of public school seats in Chile has left thousands of students unassigned each year. In “Capacity Planning in Stable Matching,” Bobbio, Carvalho, Lodi, Rios, and Torrico develop a novel framework that jointly determines where to expand school capacities and computes a student-optimal stable assignment in the enlarged market. The authors develop exact and heuristic methods that make this theoretically complex problem tractable in practice. Using rich administrative data from the Chilean school choice system, the framework demonstrates how adding a limited number of seats can trigger improvement chains benefiting multiple students, also revealing diminishing marginal returns to capacity expansion. Beyond the Chilean context, the framework provides a versatile toolkit that can be adapted to other constrained allocation problems, offering a rigorous foundation for data-driven policy design in education and beyond.

Entropy-Regularized Wasserstein Distributionally Robust Optimization

Uncertainty in data poses a central challenge in operations research. Distributionally robust optimization (DRO) offers a principled framework for addressing this challenge by producing solutions resilient to distributional variations. Among various DRO approaches, the Wasserstein DRO has received significant attention though its computational efficiency relies on stringent assumptions, and its worst case distributions are typically discrete. In “Sinkhorn Distributionally Robust Optimization,” Wang, Gao, and Xie leverage the Sinkhorn distance—an entropy-regularized variant of the Wasserstein distance—to more realistically model uncertainty, enhancing computational efficiency. The authors establish a strong duality reformulation and propose a first order stochastic mirror descent algorithm with provable complexity guarantees for general loss functions. Unlike Wasserstein DRO, Sinkhorn DRO yields continuous worst case distributions, offering a more flexible representation of practical uncertainties. Extensive experiments in the newsvendor problem, portfolio optimization, and adversarial classification demonstrate its superior performance in both out-of-sample performance and efficiency.

Integrated Conditional Estimation-Optimization

Many real-world optimization problems involve uncertain parameters with probability distributions that can be estimated using contextual feature information. In contrast to the standard approach of first estimating the distribution of uncertain parameters and then optimizing the objective based on the estimation, in “Integrated Conditional Estimation-Optimization” Qi, Grigas, and Shen propose an integrated conditional estimation-optimization (ICEO) framework that estimates the underlying conditional distribution of the random parameter while considering the structure of the optimization problem. The authors directly model the relationship between the conditional distribution of the random parameter and the contextual features and then estimate the probabilistic model with an objective that aligns with the downstream optimization problem. They show that our ICEO approach is asymptotically consistent under moderate regularity conditions and further provide finite performance guarantees.

Computationally, performing estimation with the ICEO approach is a nonconvex and often nondifferentiable optimization problem. The authors propose a general methodology for approximating the potentially nondifferentiable mapping from estimated conditional distribution to the optimal decision by a differentiable function, which greatly improves the performance of gradient-based algorithms. Numerical experiments demonstrate the empirical success of our approach in different situations, including with limited data samples and model mismatches.

Low-Rank Approaches to Large-Scale DNN Optimization

Doubly nonnegative (DNN) relaxations are a powerful tool for approximating large-scale mixed-binary quadratic programs, but their size—often involving millions of constraints—makes them difficult to solve. In “A Low-Rank Augmented Lagrangian Method for Doubly Nonnegative Relaxations of Mixed-Binary Quadratic Programs,” Hou, Tang, and Toh propose RiNNAL, a Riemannian augmented Lagrangian method that exploits the low-rank structure often present in optimal solutions. By applying a low-rank decomposition, RiNNAL reformulates most quadratic constraints into simpler affine ones, reducing problem complexity and mitigating issues caused by violations of Slater’s condition. A key innovation is showing that the required metric projection onto a certain algebraic variety, although nonconvex in form, can be solved as a convex optimization problem under mild regularity conditions, enforced through constraint reformulation. RiNNAL is versatile, handling not only DNN relaxations but also general semidefinite programs with polyhedral constraints. Extensive experiments on challenging benchmarks confirm its efficiency and robustness, making RiNNAL a promising solver for large-scale optimization problems.

Integrating Descriptive and Prescriptive Analytics: Optimized Dimensionality Reduction

Most moment-based distributionally robust optimization problems can be reformulated as semidefinite programming (SDP) problems, which can be solved in polynomial time. However, solving high-dimensional SDPs is often time-consuming. Existing approximation methods typically reduce the dimensionality of random parameters before solving the approximated SDPs. This sequential approach relies solely on statistical information to reduce the high-dimensional uncertainty space, which may not yield the best approximation performance. In “Optimized Dimensionality Reduction for Moment-Based Distributionally Robust Optimization,” Jiang, Cheng, Pan, and Shen introduce an optimized dimensionality reduction approach that integrates the dimensionality reduction of random parameters with subsequent optimization problems. This integration enables two outer approximations and one inner approximation of the original problem, all represented as low-dimensional SDPs that can be solved efficiently, providing two lower bounds and one upper bound, respectively. More importantly, these approximations can theoretically achieve the optimal value of the original high-dimensional SDPs, resulting in a zero gap.

Efficient Funding Mechanisms for Innovative Projects

In “Funding the Real Deal: Dynamic Moral Hazard with Adverse Selection,” Tian, Zhang, Sun, and Duenyas develop a model to help funders design optimal contracts for research projects when agents have both private information and hidden actions. The study addresses a common challenge in innovation financing: The agent (such as a researcher or developer) knows more about project costs and effort levels than the principal (the funder). Using a dynamic contracting framework where breakthroughs arrive randomly, the authors show that optimal incentives can be structured through simple time-dependent payment schemes. Depending on the agent’s cost type, the principal may offer either a *linear contract* with rewards that decline steadily over time, or a *one-switch contract* that introduces a single drop in rewards at a key moment. The framework provides a fast, interpretable solution method and offers insights for effectively funding R&D, climate, and pharmaceutical innovation under uncertainty.

Nature Doesn’t Play Dice, It Plays to Win

Decision making under uncertainty can be brittle, often failing when real-world data deviates from training assumptions. In “Nash Equilibria, Regularization, and Computation in Optimal Transport-Based Distributionally Robust Optimization,” Shafiee, Aolaritei, Dörfler, and Kuhn frame this problem as a game between a decision maker and an adversary, nature, who strategically corrupts the data distribution to create a worst case scenario with the cost of these changes defined by optimal transport theory. The authors establish conditions under which a stable outcome, a Nash equilibrium, exists and provide efficient methods to compute it. A key insight is that nature’s optimal strategy corresponds to generating remarkably deceptive adversarial examples; in an image classification task, this strategy can transform an image of an “8” into a convincing “3.” This work provides a powerful framework for developing more reliable models by understanding and countering worst case data perturbations.

Divergence Geometry Meets Covariance Estimation

How can we systematically improve high-dimensional covariance estimates without leaning on fragile Gaussian or other distributional assumptions? Yue, Rychener, Kuhn, and Nguyen develop a unifying geometric framework for *covariance estimation* that answers this question in “A Geometric Unification of Distributionally Robust Covariance Estimators: Shrinking the Spectrum by Inflating the Ambiguity Set.” The authors model covariance estimation as a distributionally robust optimization problem that minimizes the worst-case Frobenius loss over an ambiguity set defined by a divergence between covariance matrices. Under certain structural conditions, every such model yields a nonlinear shrinkage estimator that is orthogonally invariant, shrinks eigenvalues toward zero, and provably improves conditioning while remaining efficiently computable. The framework covers a wide range of divergences—such as Kullback-Leibler, Fisher-Rao, and Wasserstein—and leads to explicit new estimators with finite-sample and consistency guarantees. Numerical experiments in portfolio optimization and classification demonstrate performance competitive with state-of-the-art shrinkage methods.