Table of Contents

ACM SIGecom Exchanges, Vol. 23, No. 2, Jan 2026

Editors-in-Chief: Irene Lo and Sam Taggart

Communications Team: Yang Cai, Kira Goldner, and Jinzhao Wu

ACM Staff: Irene Frawley

Notice to Contributing Authors to SIG Newsletters

As a contributing author, you retain copyright to your article. ACM will refer all requests for republication directly to you.
By submitting your article for distribution in any newsletter of the ACM Special Interest Groups, you hereby grant to ACM the following non-exclusive, perpetual, worldwide rights:

—to publish your work online or in print on condition of acceptance by the editor
—to include the article in the ACM Digital Library and in any Digital Library-related services
—to allow users to make a personal copy of the article for noncommercial, educational, or research purposes
—to upload your video and other supplemental material to the ACM Digital Library, the ACM YouTube channel, and the SIG newsletter site

Furthermore, you affirm that:

—if third-party materials were used in your published work, supplemental material, or video, that you have the necessary permissions to use those third-party materials in your work

Editors' Introduction

NICK ARNOSTI
University of Minnesota
and
SAM TAGGART
Oberlin College

Happy and productive 2026 to the SIGecom community. We're thrilled to present this year's winter issue of the SIGecom Exchanges. Per tradition, we start the winter issue with the eleventh annual SIGecom job candidate profiles. Thanks to Vasilis Gkatzelis and Jason Hartline for preparing the profiles. Next, the letters section of the issue contains four contributions, each highlighting awarded work from the SIG across the last two years. Finally, the issue concludes two annotated reading lists.

The first letter is from Kei Ikegami, Atsuchi Iwasaki, Akira Matsushita, and Kyohei Okumura, whose paper "Evaluating the efficiency of regulation in matching markets with distributional disparities" won the award for best empirical paper at EC 2025. This work studies matching for medical residency with two twists: first, there are quotas restricting the number of residents assigned to hospitals and regions, and second, the work allows monetary subsidies. In the note, the authors give a wonderful overview of previous work on these topics, then sketch the main contributions of their EC paper.

Next is a letter from Bryan Wilder and Pim Welle, authors of the exemplary AI paper award at EC 2025 ("Learning Treatment Effects While Treating Those in Need"). This paper studies resource allocation problems such as those faced in provision of public services. The focus is two objectives: on the one hand, resources should be allocated to those for whom the effect will be greatest, and on the other, experimentation is necessary to infer who to target in the first place. The authors give a framework for systematically trading off these two objectives, and apply their approach on real data to get nearly the best of both worlds.

The third letter comes from Eshwar Ram Arunachaleswaran, Natalie Collina, Yishay Mansour, Mehryar Mohri, Balasubramanian Sivan, and Jon Schneider, the authors of the EC 2025 best paper "Swap Regret and Correlated Equilibria Beyond Normal-Form Games." The authors go beyond the content of this latter paper to discuss what they call "menus," a technique at the core of this and several related works. Menus give a clean distillation of outcomes achievable from play of an optimizer against a learning agent.

Our final letter is by Gabriele Farina. Gabriele won the SIGecom dissertation award in 2024, and we're delighted that he contributed a note surveying some of his recent work. The note is titled "Turning defense into offense in $O(\log 1/\varepsilon)$ steps: Efficient constructive proof of the minimax theorem," and presents another survey by way of a useful tool. In this case, the tool is an oracle-efficient method

---

Author's address: arnosti@umn.edu, staggart@oberlin.edu.

for solving zero-sum games over convex strategy spaces. Gabriele shows us several applications, including to equilibrium computation and approximate solution to variational inequalities.

We have two annotated reading lists in this issue. The first is from Tao Lin and Yiling Chen. Their reading list is a foray into the emerging literature on LLMs as economic agents. The end goal is to study their applications to information design and persuasion. Along the way, though, the authors cover related applications of LLMs to preference elicitation and decisionmaking.

The second list is from Rabanus Derr and Jessie Finocchiaro. They consider the multiple-class variant of calibrated prediction. Calibration with binary labels has received much recent attention at conferences like EC, with several recent breakthroughs (including many surveyed in the last issue of the Exchanges). The multi-class variant presents several new technical issues. The authors of this list give us a nice view into the diversity of approaches that have arisen in response.

This issue marks a transition in the editorial team for the Exchanges. Irene Lo is rotating out, after leading the newsletter through six excellent issues. We'd like to thank her for all she has done. Rotating onto the editorial team is Nick Arnosti. Nick is assistant professor in Industrial and Systems Engineering at the University of Minnesota. Thanks also go out as usual to communications chair Yang Cai, technical lead Jinzhao Wu, and social media chair Kira Goldner. Their help publishing this issue is greatly appreciated. Please continue to volunteer letters, surveys, annotated reading lists or position papers; your contributions make the Exchanges what it is. We hope you enjoy this issue.

SIGecom Job Market Candidate Profiles 2026

This is the eleventh annual collection of profiles of the junior faculty job market candidates of the SIGecom community. The forty one candidates for 2026 are listed alphabetically and indexed by research areas that define the interests of the community. The candidates can be contacted individually, or collectively via the moderated mailing list ecom-candidates2026@acm.org.

–Vasilis Gkatzelis and Jason Hartline

Fig. 1. Generated using the text from the candidate profiles.

Contents

ALIREZA AMANIHAMEDANI (Homepage, CV)

Thesis: Modern Markets: A Tale of Optimization, Incentives, and Learning ('26)

Advisor: Ali Aouad, MIT

Brief Biography: I am a final year PhD student in Management Science and Operations Research at London Business School, advised by Prof. Ali Aouad. I obtained my undergraduate degree in Computer Engineering from Sharif University of Technology, Iran. My research has been supported by the Google PhD Fellowship in Algorithms and Optimization (2025).

Research Summary: My research focuses on dynamic and data-driven optimization to design and analyze matching markets and digital platforms. The first main line of my research is on the optimization of dynamic matching markets. For instance, in ride-hailing, new drivers are continuously joining the platform while those already active may *abandon* if not matched promptly (e.g., to work on alternative platforms). Because maintaining "market thickness" in these dynamic systems may not be possible due to agents abandoning, designing efficient matching policies becomes challenging. A central theme of my research is to bridge online algorithms and stochastic control techniques to develop tractable yet provably good dynamic matching policies. Modeling these markets as queueing systems, in [1] I develop fully polynomial time approximation schemes for a broad range of markets such as organ allocation systems and ride-hailing platforms. These matching policies are crucially adaptive and leverage the real-time information of the market to make decisions. Another direction is to develop simpler policies, robust to broader market setups. In [3], I design such *proposal-based* matching policies for more general settings. Because the evolution of the queueing system is very intricate, prior work analyzed such algorithms through simplified approximations. In contrast, using tools from queueing theory, we developed a new framework that enables a more fine-grained analysis of the market, breaking the well-known $(1 - 1/e)$ approximation barrier—yielding the current best known algorithm for this class of problems.

The second line of my research studies another common feature of modern marketplaces which is the strategic behavior of both agents and platforms. In [2], I examine platforms' incentives under supplier multi-homing (i.e., participating across multiple platforms) and show that it can disincentivize platforms from adopting efficient dispatch algorithms, leading to highly inefficient equilibria. In ride-hailing, where drivers often multi-home, this *market failure* appears as platforms dispatching drivers located far from riders — a critical form of inefficiency. More recently, I am also studying how to leverage the advances in ML/AI to learn the behavior of market agents, enabling better predictions and improved market algorithms.

Representative Papers:

[1] Adaptive Approximation Schemes for Matching Queues (STOC'25)
with Ali Aouad, and Amin Saberi

[2] Spatial Matching under Multihoming (major revision at Operations Research)
with Ali Aouad, and Daniel Freund

[3] Improved Approximations for Stationary Bipartite Matching: Beyond Probabilistic Independence (under review at Mathematics of Operations Research)
with Ali Aouad, Tristan Pollner, and Amin Saberi

ABDELLAH AZNAG (Homepage, CV, Google Scholar)

Thesis: Advances in Adaptive Data Collection ('26)

Advisors: Rachel Cummings, Adam N. Elmachtoub

Brief Biography: I am a final year Ph.D. candidate at the Department of Industrial Engineering and Operations Research at Columbia University. I am advised by Prof. Rachel Cummings and Prof. Adam N. Elmachtoub. Prior to joining Columbia, I earned a BS and MS in Applied Mathematics from Ecole Polytechnique. I grew up in El Jadida, Morocco.

Research Summary: My research aims at developing the foundational principles to *strategically collect data within complex systems*. From dynamic pricing to clinical trials, the value of a final decision is inextricably linked to the quality of the data that informs it. Yet, data collection is often treated as a passive or ad-hoc process, which leads to critical oversight, biased conclusions, and significant opportunity costs. This reality motivates the central question of my work: How do we design policies that optimally balance the cost of acquiring new information against the value it provides for a decision? My work approaches this trade-off through two central themes. The first theme, *(1) Decision-aware Collection*, addresses settings where the value of data is directly tied to a well-defined downstream decision problem. The second theme, *(2) Reliability-aware Collection*, focuses on collecting data when an objective function is not yet determined, thereby building a complex system that is reliable.

My approach for tackling both of these themes is to view data collection through the lens of *information valuation*. The core challenge is not merely to design data collection policies, but to understand how these policies actively shape the flow of information, and how they reduce uncertainty. To do so, my research draws on tools from Statistical Learning, Dynamical Systems, and Information Geometry, to properly define and measure the value of information in complex settings.

Moving forward, my research will continue to advance a geometric approach to Active Learning, moving further beyond the conventional focus on sample complexity. A key focus will be studying other notions of reliability.

Representative Papers:

[1] An Active Learning Framework For Multi-group Mean Estimation (NeurIPS 2023 and major revision in Management Science)
with A. N. Elmachtoub, and R. Cummings

[2] Designing Lower Bounds for Active Learning in Multi-Armed Bandits (under review in Operations Research) with A. N. Elmachtoub, and R. Cummings

[3] MNL-Bandit With Knapsacks: a Near-Optimal Algorithm (EC'21 and major revision in Operations Research) with V. Goyal, and N. Perivier

[4] Calibrated experiments via Active Learning (working paper)

[5] Low-rank adaptive active experiment design (working paper)

FEDERICO BOBBIO (Homepage, CV)

Thesis: Dynamic Capacities and Priorities in Stable Matching ('24)

Advisor: Margarida Carvalho and Andrea Lodi, Université de Montréal

Brief Biography: I am a postdoc at Northwestern University, advised by Michael Honig, Randall Berry, Rakesh Vohra, Thanh Nguyen, and Vijay Subramanian.

Research Summary: My research centers on algorithmic mechanism design, crafting rules and incentives for fast, efficient, and fair decision-making. I design algorithms for markets with limited or no monetary transfers, where resource allocation relies on non-price instruments such as capacities, priorities, and inspections.

Mechanism design for education systems. Classical models assume fixed capacities and priorities, but real systems are flexible. I develop formulations and algorithms for stable matching with (i) capacity planning, e.g., extra spots [1] (runner-up Best Student Paper from CORS) and (ii) outcome-dependent priorities, e.g., sibling preferences in admissions [2]. For capacity planning, I provide MILP models, cutting-plane methods, and heuristics that yield transparent "policy dials" trading access vs merit. For contingent priorities, I introduce stability notions and give conditions for existence, algorithms for rank-optimal stable outcomes, and hardness results where guaranteed existence fails. These tools have been used to inform the Chilean school admission system.

Mechanisms under institutional frictions. In spectrum markets, payments to incumbents (e.g., radio astronomy) are inadmissible, and verification of interference is costly. I model a three-agent interaction incumbent-entrant-regulator, where the regulator uses access rules and selective inspections instead of prices. I show the optimal mechanism is deterministic and threshold-based with a knapsack structure: permit, deny, or inspect only on a targeted middle region to discipline misreporting. This yields auditable policies with welfare guarantees [3] (Policy Track, Best Paper).

Advancing the algorithmic backbone. Expressive non-price designs are useful only if they compute quickly and come with near-optimality certificates. I build learning-augmented branch-and-bound, large-neighborhood search, and cut-generation pipelines that solve large instances quickly. Work samples include a heuristic-enhanced MILP solver [4] (outstanding student submission) and a hybrid AI-and-optimization pipeline for stochastic routing [5] (first place, IJCAI'21).

Representative Papers:

[1] Capacity Planning in Stable Matching (Operations Research, 2025). With M. Carvalho, A. Lodi, I. Rios, and A. Torrico

[2] Stable Matching with Contingent Priorities (EC 2025). With I. Rios, M. Carvalho, and A. Torrico.

[3] Costly Measurements to Incentivize Spectrum Sharing (IEEE DySPAN 2025). With R. Berry, M. Honig, T. Nguyen, V. Subramanian, and R. Vohra.

[4] Design and Implementation of a Heuristic-Enhanced Branch-and-Bound Solver for MILP (MIP Workshop 2022). With W. A. Silva, F. Caye, D. Liu, J. Pepin, C. Perreault-Lafleur, and W. St-Arnaud.

[5] The First AI4TSP Competition: Learning to Solve Stochastic Routing Problems (Artificial Intelligence, 2023). With Y. Zhang, L. Bliek, and others.

ROBIN BOWERS (Homepage, CV)

Thesis: Beyond Pandora's Box: Algorithm and Mechanism Design with Costly Information Acquisition ('26)

Advisor: Bo Waggoner and Rafael Frongillo, CU Boulder

Brief Biography: I am a 5th year PhD student at the University of Colorado Boulder in Computer Science, where I also received my Master's degree. I earned my B.A. in Computer Science and Mathematics at Oberlin College. I founded the ongoing Algorithmic Fairness Reading Group at the University of Colorado Boulder, and co-organized the EC Gender Inclusion workshop in 2024 and 2025.

Research Summary: I am broadly interested in how algorithm and mechanism design incorporates hidden or expensive information. My work has largely focused on extending the Pandora's box model of costly information in various settings. This model of information acquisition is very natural in settings such as labor-intensive interviewing in job matching, but models often overlook these concerns.

Throughout my PhD, my work has incorporated Pandora's box models into various settings. We proved a price of anarchy result bounding the social welfare loss in a mechanism for matching with monetary transfers (as in, e.g., a job market) which allows for inspection [1]. This work blends techniques and concepts from auction design literature with Pandora's box-style approaches. Following up on our mechanism design problem, we proved that most intuitive algorithms for this problem fail when values may be negative, a critical regime for modeling certain settings [2]. These algorithmic results have deeper implications for the sophistication required of any matching mechanism attempting to coordinate inspecting agents.

My latest work extends the Pandora's box model to arbitrary sequences of information-gathering decisions. By combining prophet inequalities with a structural generalization of the Pandora's box problem, we provide algorithmic approximation results [3].

I have also worked on PTAS design for mechanisms, including problems such as unit-demand pricing [4].

My latest interests have been in more generally evaluating information in mechanism design settings, from peer prediction problems to using information structures to better aggregate reports from agents. I also have a longstanding interest in design of socially-responsible mechanisms and algorithms, particularly in the areas of social choice and matching.

Representative Papers:

[1] High Welfare Matching Markets via Descending Price (WINE '23) with B. Waggoner

[2] Matching with Nested and Bundled Pandora Boxes (WINE '24) with B. Waggoner

[3] Prophet Inequalities for Bandits, Cabinets, and DAGs (Working paper) with E. Lindgren, B. Waggoner

[4] Polynomial-Time Approximation Schemes via Utility Alignment: Unit-Demand Pricing and More (FOCS '25) with M. Garbea, E. Pountourakis, S. Taggart.

NATALIE COLLINA (Homepage, CV, Google Scholar)

Thesis: Learning and Strategic Interaction in Human–AI Systems ('26)

Advisors: Aaron Roth and Michael Kearns, University of Pennsylvania

Brief Biography: I am a 5th-year Ph.D. student at the University of Pennsylvania advised by Aaron Roth and Michael Kearns. In the summer of 2025, I interned at Microsoft Research, New England under the mentorship of Alex Slivkins. I have also collaborated with Jon Schneider at Google Research throughout my Ph.D. I am the recipient of several awards, including an IBM Ph.D. Fellowship and a joint Best Paper and Best Student Paper Award at EC 2025

Research Summary: My research studies learning and strategic interaction in human–AI systems. I use tools from online learning and game theory to understand how humans and AI agents learn from, adapt to, and influence one another. A central goal of my work is to develop a theory of human–AI collaboration—how cooperative and complementary behavior can emerge through repeated interaction, even when agents have distinct objectives [1] or partial information [2]. I also examine cases where these same learning dynamics give rise to algorithmic collusion, showing how AI agents can coordinate in ways that are hard to detect and fall outside our classical or legal understanding of collusion [3]. Together, these results provide a unified view of when learning promotes cooperation and when it enables strategic behavior that may be misaligned with social welfare.

A second, complementary line develops commitment as a design principle for learning algorithms, formalized through the notion of an algorithm's menu—the structured set of policies it commits to before interaction [4]. This framework, which was honored with the Best Paper Award at EC 2025, characterizes when such commitments are optimal and how they can preempt manipulation, encourage cooperation, or ensure robust outcomes in strategic environments [5]. Together, these directions aim to build a theoretical foundation for AI systems that behave predictably and beneficially in multi-agent settings.

Representative Papers:

[1] Emergent Alignment via Competition (in submission to ICLR 2026)
with S. Goel, A. Roth, E. Ryu, and M. Shi

[2] Tractable Agreement Protocols (STOC 2025)
with S. Goel, V. Gupta, and A. Roth

[3] Algorithmic Collusion Without Threats (ITCS 2025)
with E.R. Arunachaleswaran, S. Kannan, A. Roth, and J. Ziani

[4] Pareto-Optimal Algorithms for Learning in Games (EC 2024)
with E.R. Arunachaleswaran and J. Schneider

[5] Swap Regret and Correlated Equilibria Beyond Normal-Form Games (EC 2025)
with E.R. Arunachaleswaran, Y. Mansour, M. Mohri, J. Schneider, and B. Sivan

JEFF DECARY (Homepage, CV, Google Scholar)

Thesis: Winning Against All Odds: Combinatorial Optimization Approaches for Prediction Markets ('26)

Advisor: David Bergman, University of Connecticut

Brief Biography: Jeff Decary is a Ph.D. candidate in Business Administration (Operations & Information Management) at the University of Connecticut (UConn) School of Business, advised by David Bergman. His research lies at the intersection of optimization under uncertainty and market design for prediction markets, with applications in sports analytics. Prior to joining UConn, Jeff completed a Master's degree in applied mathematics at Polytechnique Montréal focusing on embedded neural networks, advised by Andrea Lodi.

Research Summary: Jeff's research centers on risk-sensitive combinatorial allocation, platform pricing, and econometric methods for causal inference. Methodologically, he integrates discrete, stochastic, and bilevel optimization, dynamic programming, and applied machine learning to design exact and simulation-based algorithms for high-stakes decisions. His work aims to deliver tools that help participants and platforms make better decisions in competitive environments.

Portfolio Optimization for Prediction Markets. A central question in his research is how traders make optimal decisions under uncertainty in prediction markets. In [1], he extends the classical Kelly criterion via a logic-based Benders decomposition for log-optimal portfolio allocation under combinatorial constraints, uncovering structural insights that enable scalable solution methods. In [2], he extends this perspective to contest-based markets (e.g., March Madness), developing simulation-based algorithms guided by a dynamic-programming structure that outperform experts in a real-world, high-stakes contest. Together, these works show how rigorous optimization uncovers structure and informs strategy in competitive markets.

Market Design for Prediction Markets. Another key question Jeff studies is how prediction-market platforms set prices, payout rules, and incentives to shape user behavior. In [1], he analyzes sportsbook parlay pricing and shows that small adjustments in payout odds can significantly shift portfolio allocations and platform risk exposure. In ongoing work [3], he models daily fantasy sports pricing as a bilevel optimization problem, capturing the interaction between platform pricing and participant strategy. Together, these studies highlight how market design choices influence competitiveness, fairness, and profitability in digital marketplaces.

Representative Papers:

[1] Log-Optimal Portfolio Construction for Binary Options with Combinatorial Constraints (under review at Management Science) with D. Bergman, and B. Zou

[2] The madness of multiple entries in march madness (Poster at EC, 2024; 2nd Round Revision at Production & Operations Management) with D. Bergman, C. Cardonha, J. Imbrogno, and A. Lodi

[3] Market-making for daily fantasy sports: Competitive pricing for sophistication (Working Paper) with D. Bergman, and R. Day

CRAIG FERNANDES (Homepage, CV)

Thesis: Operations Management for Innovative Markets ('26)

Advisor: Timothy Chan and Ningyuan Chen; University of Toronto

Brief Biography: Craig Fernandes is a final-year PhD candidate in Operations Research at the University of Toronto and a Vanier Scholar. His work has been published in Operations Research and Management Science and has been awarded 1st place at the CORS Student Paper Competition, the MIT SSAC Paper Competition and the INFORMS Case Competition. Craig interned with Amazon's SCOT team and was a visiting scholar at the Tuck School of Business at Dartmouth College.

Research Summary: My research examines innovative markets that are highly impactful yet often overlooked in traditional operations literature. Below, I outline three research themes that illustrate this work.

Analytical Market Design: This stream uses game theory to uncover hidden incentives in unique markets. In [1], I analyze income pools —contracts in superstar markets where individuals agree to share a portion of future earnings if they become highly successful — developing the first mathematical model to study their incentive structures. In [3], I examine academic conferences and design matching policies that encourage reviewers to exert effort, proposing an admission control policy that rejects papers from authors who performed poorly as reviewers. Notably, while writing this paper, NeurIPS, ICML, and CVPR independently implemented policies that mirror our proposed mechanism.

Data-Driven Market Design: This stream uses optimization and statistics to extract insights from messy data. In [4], we study how firms can set profit-maximizing prices from offline transaction data in business-to-business (B2B) markets, where final prices are determined through a quote-and-bargain process. Customers differ not only in willingness to pay but also in bargaining power, drawn jointly from an unknown distribution. Unlike prior work that assumes these distributions are known, we establish identifiability conditions and develop a data-driven algorithm that achieves near-optimal revenue.

Market Design in Sports: This stream uses Markov models to uncover hidden value in sports. In [2], I study the $30 billion NFL market, developing a dynamic programming and Markov-based value function that offers the first theoretical foundation for measuring how value is generated across plays, beyond what traditional box score metrics capture.

Representative Papers:

[1] Income Pools for Superstar Markets (Published, Management Science, 2025) with Timothy Chan and Ningyuan Chen

[2] Points Gained in Football (Published, Operations Research, 2021) with Timothy Chan and Martin Puterman

[3] Peer Review Market Design (ACM EC'25; Target: Operations Research) with James Siderius anad Raghav Singal

[4] Data-Driven B2B Pricing: Learning from Bargainers (Working paper) with Timothy Chan and Ningyuan Chen

GIANNIS FIKIORIS (Homepage, CV, Google Scholar)

Thesis: Learning in Games with State ('26)

Advisor: Éva Tardos, Cornell University

Brief Biography: Giannis Fikioris is a Ph.D. Candidate in the department of Computer Science at Cornell. He has been supported by the Google and NDSEG Fellowships, along with the Onassis and Gerondelis Scholarships. He has also interned twice in Google Research, at the Market Algorithms Team, advised by Mingfei Zhao and Yuan Deng. He has worked on learning in games with budgets and fairness in online resource allocation.

Research Summary: One of my main interests is learning in games, especially when the learning agents have long-term constraints, e.g., budget constraints. My goal in this area is to explore learning algorithms, or more generally, properties, that result in high utility or welfare. Specifically, I explore such properties in the following settings, from strongest to weakest (strong/weak loosely defined): (i) *adversarial setting*, where an agent is competing against other agents that are acting adversarially, (ii) *well-behaved settings*, where all the agents run a certain algorithm or algorithms that satisfy a certain property, (iii) *stochastic settings*, where an agent is competing against stationary competition. Below, I go over some of my past work that has examined these settings.

Liquid Welfare from high utility – type (ii) setting In [1], we proved that in Sequential Budgeted first-price Auctions, the resulting Liquid Welfare is high under the minimal assumption that agents have high competitive ratio. We were able to prove that under minimal conditions, aside from the resulting allocation having high utility for all players, their valuations for the items can be submodular and adversarially chosen.

Auctions with Spacing objectives – type (iii) setting In [2], we introduced a new model of sequential auctions, where the value of winning a round depends on the time since the last win. This is meant to capture settings like advertising, where the spacing between advertisements is important. Aside from modeling, we also simplify the offline problem of finding the optimal algorithm and offer an online learning algorithm with optimal $\tilde{O}(\sqrt{T})$ regret.

Online Resource Allocation – type (i) setting In a series of works (one of the most recent ones being [3]), we have examined the setting where a principal is repeatedly trying to fairly allocate a public good to agents without the use of money. We show that, under natural assumptions on the agent distributions, every agent can robustly (i.e., independent of how others behave) guarantee almost all of her ideal utility, i.e., her maximum utility under her nominal share of the resource.

Representative Papers:

[1] Liquid Welfare Guarantees for No-Regret Learning in Sequential Budgeted Auctions (EC'23 & Math of OR) with É. Tardos

[2] Learning in Budgeted Auctions with Spacing Objectives (EC'25 & Under Review at *JMLR*)
with R. Kleinberg, Y. Kolumbus, R. Kumar, Y. Mansour, and É. Tardos

[3] Beyond Worst-Case Online Allocation via Dynamic Max-min Fairness (EC'25 & Under Review at *Management Science*) with S. Banerjee and É. Tardos

MAXWELL FISHELSON (Homepage, CV, Google Scholar)

Thesis: Reliable Learning for Adversarial Environments ('26)

Advisor: Constantinos Daskalakis, Massachusetts Institute of Technology

Brief Biography: I'm a final year PhD student at the Massachusetts Institute of Technology, working on problems at the intersection of online learning and algorithmic game theory. In particular, I design algorithms for learning in games and for strong learning benchmarks like swap regret and calibration, which guarantee reliability and non-exploitability. My work has improved rates for trustworthy forecasting and extended these guarantees to high-dimensional action space domains. For the last year and a half, I have also been a student researcher at Google Research.

Research Summary: My research focuses on developing learning algorithms for adversarial environments, where standard machine learning methods can fail. I design and analyze algorithms that are reliable, trustworthy, and non-exploitable, even when facing strategic adversaries. My work seeks to understand what makes a good learning algorithm for these interactive settings, focusing on strong, distribution-free benchmarks like regret and calibration error.

My contributions in this area include establishing near-optimal, exponentially-improved regret bounds for fast learning in general games [1], and developing the first efficient algorithms for swap regret in high-dimensional action spaces [2]. My research has also advanced trustworthy forecasting by breaking a 27-year-old barrier in sequential calibration with an asymptotically faster algorithm [3], and introducing the first algorithm to break the exponential-dimension barrier for high-dimensional calibration over arbitrary convex sets [4].

The goal of my research is to meet the needs of modern decision-making tasks. These involve high-dimensional choices, such as tuning neural network weights, and a need for non-exploitability and reliability against strategic responses. My work aims to provide the foundational theoretical principles to meet these practical goals, inspired by this future role that AI will take on in society.

Representative Papers:

[1] Near-Optimal No-Regret Learning in General Games (NeurIPS 2021, Oral)
with Constantinos Daskalakis and Noah Golowich

[2] From External to Swap Regret 2.0: An Efficient Reduction for Large Action Spaces (STOC 2024, Invited to SICOMP Special Issue)
with Yuval Dagan, Constantinos Daskalakis, and Noah Golowich

[3] Breaking the $T^{2/3}$ Barrier for Sequential Calibration (STOC 2025, Invited to SICOMP Special Issue)
with Yuval Dagan, Constantinos Daskalakis, Noah Golowich, Robert Kleinberg, and Princewill Okoroafor

[4] High-Dimensional Calibration from Swap Regret (NeurIPS 2025, Oral)
with Noah Golowich, Mehryar Mohri, and Jon Schneider

ABHEEK GHOSH (Homepage, CV)

Thesis: Contests: Equilibrium Analysis, Design, and Learning ('24)

Advisor: Edith Elkind, Northwestern Univ.; Paul W. Goldberg, Univ. of Oxford

Brief Biography: I am currently a postdoctoral researcher supervised by Prof. Paul Goldberg, focusing primarily on fixed-point computation and TFNP complexity classes. During my PhD, advised by Profs. Edith Elkind and Paul Goldberg, I focused on contest theory and learning dynamics in games. I also completed an internship with Prof. Milind Tambe at Google Research, where I studied multi-armed bandits for resource allocation in social-good applications. Earlier, I was advised by Prof. Umang Bhaskar at TIFR Mumbai, where I worked on voting theory.

Research Summary: My recent research has focused on fixed-point computation problems. In an upcoming work, we study the complexity of computing a Nash equilibrium in normal-form games with a unique equilibrium. We show connections to the *unique end-of-line* (UEoL) problem. In another work, we study the complexity of computing equilibrium points of electrostatic potentials (Coulomb's law) [1]. We develop efficient algorithms to compute approximate equilibrium points with inverse-exponentially small error. For generalizations of the problem relevant to machine learning, we show connections to min–max optimization and the corresponding computational complexity classes. In another work, we resolve the complexity of computing KKT points of quadratic programs over a simplex.

Another area of my current research is contest theory. Contests are games in which agents compete for valuable rewards by exerting costly, irreversible effort. Classic examples include the all-pay auction and the Tullock contest. Contests model many important applications; for example, they arise naturally in proof-of-work and proof-of-stake blockchain protocols, and in competition among content creators on social-media platforms. In a sequence of papers, we study best-response, fictitious-play, and related learning dynamics in Tullock contests [3]. My two other lines of work in contests are equilibrium computation and price of anarchy in parallel Tullock and equal-sharing contests, and contest design to incentivize participation from weaker or under-represented agents in rank-order (all-pay-type) contests.

Other work includes fraud-proof revenue-division mechanisms for subscription platforms such as Spotify [2], near-optimal algorithms for restless multi-armed bandits and its applications to healthcare, information elicitation with credible agents, deliberative coalition formation, and distortion in voting mechanisms. I am broadly interested in combining tools from theoretical computer science and economics to solve practical problems. I plan to focus on computational complexity, particularly problems arising from dynamical systems and optimization, and learning in games, especially as applied to contests and auctions.

Representative Papers:

[1] Computing Equilibrium Points of Electrostatic Potentials (in submission)
with PW. Goldberg and A. Hollender

[2] Fraud-Proof Revenue Division in Subscription Platforms (ICML 25)
with TY. Neoh, N. Teh, and G. Tyrovolas

[3] Best-Response Dynamics in Lottery Contests (EC 23)
with PW. Goldberg

SUMIT GOEL (Homepage, CV, Google Scholar)

Thesis: Essays in Mechanism Design and Contest Theory ('23)

Advisor: Federico Echenique, U.C. Berkeley

Brief Biography: Sumit is a postdoc at NYU Abu Dhabi. He received his Ph.D. in Social Science (with a minor in CS) from Caltech. Before graduate school, he earned a Master's in Economics from Indian Statistical Institute (Delhi), and a Bachelor's in CS from Delhi Technological University.

Research Summary: I am a microeconomic theorist studying problems of economic design, with particular interests in the following areas.

A) Contest theory: Contests model environments where agents compete for valuable prizes by making costly investments. My research investigates how various design instruments affect investment behavior. [1] analyzes prize structures in rank-order contests, [3] studies design of target-based contracts, and [5] examines grading schemes when grades serve as signals of ability. Two additional papers explore the role of tie-breaking rules in Tullock-style contests, and feedback policies in dynamic all-pay auctions. These studies reveal two broad insights: under complete information, the choice of instrument is largely irrelevant, whereas under incomplete information, the winner-takes-all prize structure maximizes investment.

B) Allocation problems: Fair and efficient allocation of resources is a fundamental problem in economics and CS. My research seeks to overcome impossibility results by examining restricted preference domains and relaxing key axioms. [2] studies object reallocation and introduces the top-two condition as a useful richness criterion for identifying domains where TTC is uniquely desirable, or where alternative desirable mechanisms exist. [4] considers multiple objects, identifies conditions under which core allocations exist, and proposes a generalized TTC algorithm that finds an allocation in the stable set. A recent paper introduces a fairness notion of swap-bounded envy for multi-dimensional allocation problems and proposes a TTC+SD algorithm that yields fair and efficient allocation. A separate paper establishes an ordering of k-price auctions based on worst-case allocative efficiency.

C) Mechanism design without transfers: In the absence of transfers, the VCG mechanism is infeasible for implementing socially optimal outcomes. One paper quantifies the resulting welfare loss in a two-dimensional facility location problem, while another shows how partial verifiability can mitigate some of the losses in a principal-agent project selection problem.

Representative Papers:

[1] The Effect of Competition in Contests: A Unifying Approach (JMP)
with A. Baranski

[2] TTC Domains (WINE 2025)
with Y. Tamura

[3] Multi-Agent Contract Design with a Budget (EC 2024, R&R at GEB)
with W. Hann-Caruthers

[4] Stable Allocations in Discrete Exchange Economies (JET 2024)
with F. Echenique, and S. Lee

[5] Optimal Grading Contests (EC 2023, GEB 2025)

ANDREAS HAUPT (Homepage, CV, Google Scholar)

Thesis: The Economic Engineering of Personalized Experiences ('24)

Advisors: A. Bonatti, D. Hadfield-Menell, E. Maskin, D. Parkes

Brief Biography: Andreas Haupt is a Stanford HAI Postdoctoral Fellow jointly affiliated with Economics and Computer Science, advised by Erik Brynjolfsson and Sanmi Koyejo. He earned his PhD at MIT and master's degrees in Mathematics (2017) and Economics (2018) from the University of Bonn (with distinction). His policy experience includes work with the European Commission's Directorate-General for Competition and the U.S. Federal Trade Commission. He co-authors a forthcoming textbook on Machine Learning from Human Preferences and leads a doctoral course on the topic at Stanford.

Research Summary: Andreas designs and analyzes personalized platforms, with contributions to privacy-preserving mechanism design, learning from human preferences, and platform regulation.

Privacy. [1] develops a *need-to-know* notion of contextual privacy for extensive-form mechanism design: elicitation is justified only when needed to compute a social choice function. The paper characterizes which social choice functions admit such implementations and identifies auction formats (including ascending and descending implementations of the second-price auction choice rule) with improved privacy. Follow-up work studies statistical aspects of contextual privacy.

ML from Human Preferences. [2] analyzes online preference elicitation used in personalization. It shows that standard bandit algorithms favor—in a way the paper makes precise—actions with lower measurement noise or less heterogeneous responses and studies analogous behavior for gradient-based methods.

Platform Regulation. [3] studies platform preferencing on large marketplaces (e.g., Amazon). Because the platform is both market maker and seller, algorithmic curation can make in-house products disproportionately salient. The paper defines and measures preferencing, shows when it distorts price competition, and connects the results to current regulatory frameworks. Parallel work formalizes preferencing using counterfactual algorithmic fairness.

He also works in applied theory beyond platforms; see, e.g., [4] on voluntary carbon markets.

Representative Papers:

[1] Contextually Private Mechanisms (ACM EC '22, R&R at the AER)
with Z. Hitzig

[2] Risk Preferences of Learning Algorithms (Games and Economic Behavior '24)
with A. Narayanan

[3] Platform Preferencing and Price Competition I (Working Paper)
with O. Hartzell

[4] Certification Design for a Competitive Market (ACM EC '23)
with B. Lucier and N. Immorlica

VISHWA PRAKASH HV (Homepage, CV, Google Scholar)

Thesis: Existence Guarantees and Algorithms in Discrete Fair Division ('26)

Advisor: Prajakta Nimbhorkar, Chennai Mathematical Institute.

Brief Biography: Vishwa Prakash HV is a final-year Ph.D. candidate at the Chennai Mathematical Institute (CMI), under the supervision of Prof. Prajakta Nimbhorkar. He completed a Master's in Computer Science from CMI and a Bachelor's in Information Science & Engineering from Visvesvaraya Technological University. He was awarded the TCS Research Scholar Fellowship during his Ph.D. and the Cognizant Fellowship during his Master's.

Research Summary: I am broadly interested in *algorithmic game theory* and the study of combinatorial problems that arise in resource allocation. My doctoral research focuses on *fair division*, a core area of microeconomic theory concerned with fundamental questions: can a set of indivisible items be allocated among agents with diverse preferences in a *fair* manner? If so, for which notions of fairness, and can such allocations be computed efficiently?

A central theme of my work is the existence of allocations where no agent *envies* any strict subset of another's bundle—known as EFX allocations. In a sequence of papers [1–3], we give algorithms showing that such allocations exist for any number of agents when there are only a few *types* of agents. Together, these results support the conjecture that if EFX exists for k agents, then it also exists for any number of agents of k types.

In practical settings, an allocator may face partial constraints, such as preassigned goods (e.g., those specified in a will) or unavoidable tasks (e.g., breastfeeding). Work [4] examines the computational complexity of *completing* such partial allocations while ensuring standard fairness guarantees.

Ensuring *efficiency* alongside fairness is another key objective. In [5], this balance is explored in the context of *mixed manna*, establishing the existence of efficient and approximately envy-free allocations via fixed-point techniques. Finally, [6] provides structural characterizations of *weighted proportional allocations* under unequal entitlements, using tools from matching theory.

Representative Papers:

[1] EFX Exists for Three Types of Agents (EC 2025)
with P. Ghosal, P. Nimbhorkar, and N. Varma

[2] (Almost Full) EFX for Three (and More) Types of Agents. (AAAI 2025)
with P. Ghosal, P. Nimbhorkar, and N. Varma

[3] Almost and Approximate EFX for Few Types of Agents (working paper)
with R. Mehta and P. Nimbhorkar

[4] Fair and Efficient Completion of Indivisible Goods (AAAI 2025)
with A. Igarashi, and R. Vaish

[5] Fair and Efficient Allocation of Indivisible Mixed Manna (WINE 2025)
with S. Barman, A. Sethia, and M. Suzuki

[6] Weighted Proportional Allocations of Indivisible Goods and Chores: Insights via Matchings. (AAMAS 2024)
with P. Nimbhorkar

STANISŁAW KAŹMIEROWSKI (Homepage, CV)

Thesis: Solving Succinct Games ('25)

Advisor: Marcin Dziubiński, University of Warsaw

Brief Biography: I am a fourth-year PhD candidate at the University of Warsaw, Faculty of Mathematics, Informatics, and Mechanics, where I work on problems related to solving large games with succinct representation. During my PhD, I enjoyed a four-month-long internship at the Department of Economics of the University of Zurich, where I collaborated with Prof. Christian Ewerhart.

Research Summary: My research focuses on game theory, with a particular emphasis on the computation of equilibria in large games with succinct representations. I develop efficient algorithms to compute Nash equilibria in games with large, discrete strategy spaces, such as conflicts with multiple battlefields and network-based attack-defense games. A central challenge in these areas is the exponential growth in the number of strategies, where traditional methods often prove inefficient, and this is where my work seeks to innovate.

Beyond the computational aspect, I am also interested in the theoretical properties of equilibria. In our work on the Arad-Rubinstein game [4], we investigate how changing the tie-breaking rule affects the equilibrium set, revealing insights into strategic behavior, inefficiencies, and robustness.

To address the challenges posed by large games, I employ techniques such as strategy symmetrization, algorithmic optimization, and heuristic methods. For example, in article [2], we describe a network reduction operation that allows us to compute a Nash equilibrium in the Attack and Defense Game on Networks in polynomial time with respect to the number of nodes. In article [3], we present a polynomial-time algorithm for computing symmetrized payoffs in symmetric conflicts with multiple battlefields, reducing the game's size exponentially with a polynomial time cost. When combined with the Double Oracle Algorithm and a heuristic that leverages the model's structure, this method achieves a speedup of four orders of magnitude compared to classical approaches.

In my recent single-author paper[1], I explored a variant of the Colonel Blotto game that incorporates costs, demonstrating that it is strategically equivalent to a zero-sum Colonel Blotto game with one additional battlefield. This equivalence allows for the efficient computation of Nash equilibria in polynomial time with respect to the total number of battlefields and resources available to the players.

Representative Papers:

[1] Equilibria of the Colonel Blotto Games with Costs (AAAI 2025)

[2] Computation of Nash Equilibria of Attack and Defense Games on Networks (SAGT 2023) with M. Dziubiński

[3] Efficient Method for Finding Optimal Strategies in Chopstick Auctions with Uniform Objects Values (AAMAS 2024) with M. Dziubiński

[4] An equilibrium analysis of the Arad-Rubinstein game (Journal of Economic Behavior & Organization) with C. Ewerhart

YOAV KOLUMBUS (Homepage, CV)

Thesis: *Strategic Considerations and Learning in Complex Systems* ('23)

Advisor: Noam Nisan, The Hebrew University of Jerusalem

Brief Biography: I am an Assistant Research Professor of Computer Science and Economics at the Center for Data Science for Enterprise and Society at Cornell University, where I am fortunate to be mentored by David Easley, Robert Kleinberg, and Éva Tardos. Prior to joining Cornell, I completed my Ph.D. in Computer Science at the Hebrew University, advised by Noam Nisan. I also completed my B.Sc. and M.Sc. in Physics there, as well as a Bachelor's degree in Music at the Jerusalem Academy of Music, where I studied classical and jazz double bass.

Research Summary: Learning-based algorithmic agents and AI tools increasingly mediate major markets and online platforms, fundamentally reshaping their dynamics and economic outcomes. This shift raises foundational questions: How do existing mechanisms function when populated by such agents? What dynamics do they induce, and what outcomes emerge in equilibrium or through learning? How should human stakeholders — private users, advertisers, investors, and platform managers — interact with these tools to best achieve their objectives? And how should we redesign these systems and the learning agents themselves to promote efficiency, robustness, and other desirable outcomes in the AI era?

My research draws on machine learning, algorithmic game theory, microeconomics, and operations to study and design systems where learning agents and strategic players interact. I focus on understanding their dynamics and the incentives they induce, to develop mechanisms, policies, and algorithms that achieve desirable behaviors and outcomes. I approach this through three main lenses:

(1) *System analysis*: Analyzing dynamics and trade-offs in strategic systems and markets with learning agents.

(2) *Better learning and platforms*: Designing learning methods for improved performance in uncertain strategic settings and mechanisms suitable for learners.

(3) *Interaction models*: Modeling interactions between strategic users and algorithmic agents under varied structures.

My work integrates economic design, the algorithmic principles underlying learning, and the game-theoretic and network structures of interactions. It spans domains such as learning in markets [1], online advertising auctions [2], and contracting with learning agents [3], as well as decentralized routing, human–algorithm interaction, and foundational game-theoretic modeling for autonomous agents.

Representative Papers:

[1] Markets with Heterogeneous Agents: Dynamics and Survival of Bayesian vs. No-Regret Learners (EC 2025) with D. Easley and É. Tardos

[2] Auctions between Regret-Minimizing Agents (TheWebConf 2022)
with N. Nisan

[3] Contracting with a Learning Agent (NeurIPS 2024)
with I. Talgam-Cohen, M. Weinberg, G. Guruganesh, M. Vlatakis, J. Wang, and J. Schneider

POOJA KULKARNI (Homepage, CV)

Thesis: Fair Division: Addressing Complement-Free Valuations and Online Settings ('25)

Advisors: Jugal Garg and Ruta Mehta, UIUC.

Brief Biography: I am a Postdoctoral Researcher at Northwestern University, hosted by Prof. Samir Khuller. I received my Ph.D. in Computer Science from UIUC, advised by Jugal Garg and Ruta Mehta, and my masters and undergraduate degrees from IISc and CoEP, where I was the gold medalist at both institutions. I have industry and research experience through internships at Meta, Nvidia, and NTT Research. My research interests span Fair Allocation and Data Economics.

Research Summary: My research develops algorithms that make resource-allocation and market systems *stable and predictable*. Whether distributing disaster relief or selling digital goods such as API access to an LLM, stability can be ensured either from fairness or from equilibrium, where no individual or group benefits from deviating. My work advances these goals across three main lines:
Online Fair Allocation. In [1], I model disaster-relief allocation as an online discrete fair-division problem, where agents arrive sequentially. Such problems are notoriously hard—few positive results were previously known. My work introduced the OnlineKTypeFD model, inspired by learning-augmented algorithms, where agents belong to one of K known types. This yields the first non-trivial results for this setting. Specifically, we design deterministic algorithms for the Maximin Share (MMS) fairness notion. I am extending this line of work to randomized algorithms and other fairness notions.
Submodular Fair Allocation. Most real agents exhibit non-additive valuations like submodular valuations (capturing diminishing returns) yet fair-division theory largely assumes additive utilities. My work (e.g., [2]) extends notions like Nash Social Welfare and Maximin Share to submodular and fractionally subadditive valuations—relevant for settings such as course allocation. Since techniques for additive valuations rarely extend, my work has introduced new techniques such as *match–rematch* and *capped-welfare maximization*.
Equilibrium in Data Economies. Data is an economic asset. It is *non-rivalrous* i.e., can benefit multiple agents simultaneously. This fundamentally alters how exchange or trading economies with data will behave. To investigate this, in [3], I model data-sharing economies without payments, characterizing stable outcomes using the notion of *core* from cooperative game theory. In an ongoing work, I analyze data pricing, showing how non-rivalrousness fundamentally changes strategic pricing decisions and revenue compared to traditional goods.

Representative Papers:

[1] Online Fair Division: Towards Ex-Post Constant MMS Guarantees (EC 2025) with R. Mehta, P. Shahkar

[2] Approximating NSW under Submodular Valuations through (Un)Matchings (SODA 2020, TALG 2024) with J. Garg, R. Kulkarni

[3] On the Existence and Complexity of Core Stable Data Exchanges (NeurIPS 2025) with J. Song, P. Shahkar, B. Chaudhury

SOONBONG LEE (Homepage, CV, Google Scholar)

Thesis: Platform and Policy Design for Social Good: Modeling and Data-Driven Algorithmic Approaches ('26)

Advisor: Vahideh Manshadi (Yale School of Management)

Brief Biography: I am a fifth-year Ph.D. student in Operations at Yale School of Management, where I am fortunate to be advised by Prof. Vahideh Manshadi. My research interests center on platform and policy design for social good. My research has been recognized by several awards, including the MSOM Best Student Paper Prize and the Auctions and Market Design Rothkopf Prize.

Research Summary: My central research agenda is platform and policy design for social good. Drawing tools from optimization, data science, and economics, I develop data-driven solutions and analytic methods for resource allocation and matching problems arising in nonprofit and public sectors. In doing so, my goal is to amplify their social impact by improving operations and to generate insights for policymakers and managers seeking to address pressing societal challenges. I have worked on applications in refugee resettlement [1], diversity policies in labor markets [2], and food rescue [3].

Through my research, I aim to bridge theory and practice via two complementary streams. One thread focuses on developing data-driven solutions for tech-enabled nonprofits and public sector organizations. Much of this work has been grounded in close collaboration with partner organizations, including a major U.S. refugee resettlement agency [1] and Feeding America's online food rescue platform [3]. By combining data analysis with conversations with practitioners, I uncover operational challenges and design optimization algorithms that provide provable performance guarantees as well as practical appeals (e.g., strong empirical performance, computational efficiency, and interpretability), with the goal of proposing solutions that are deployable in practice. A complementary thread focuses on developing a general modeling framework to investigate policy questions related to societal challenges (e.g., diversity intervention in labor markets [2]), where access to fine-grained data is often limited. In these settings, I develop mathematical models that capture the core operational characteristics of the problem and conduct rigorous analysis to generate policy insights.

Representative Papers:

[1] Dynamic Matching with Post-Allocation Service and its Application to Refugee Resettlement (Accepted at *Management Science* & EC 2024) with K. Bansak, V. Manshadi, R. Niazadeh, and E. Paulson

[2] Why the Rooney Rule Fumbles: Limitations of Interview-stage Diversity Interventions in Labor Markets (EC 2025 & Major Revision at *Operations Research*) with S. Farajollahzadeh, V. Manshadi, and F. Monachou

[3] Who to Offer, and When: Redesigning Feeding America's Real-Time Donation Tool (Working Paper) with V. Manshadi, and D. Saban

CE LI (Homepage, CV)

Thesis: Essays in Economic Theory and Artificial Intelligence ('26)

Advisor: Bart Lipman, Department of Economics, Boston University

Brief Biography: Ce Li is an Economics PhD candidate at Boston University, advised by Professor Bart Lipman. She is the co-founder and co-organizer of the first ACM EC Workshop: *Information Economics* $\times$ *Large Language Models*. She obtained her S.M. in Health Data Science at Harvard University and finished her bachelor's studies in finance and statistics at Sun Yat-sen University.

Research Summary: My research develops theories at the AI-economics interface. I develop *microeconomic theories* (information design, mechanism design) and *learning algorithms* for *human-AI interactions* and *agentic decision-making* that are *efficient, trustworthy, and robust* to *profitable and societal concerns*, which provide insights for applications in *online platforms, LLMs, and AI agents*.

Information designers, such as large language models and online platforms, often do *not* know the *subjective* beliefs of their receivers or users. In my job market paper [1], we construct learning algorithms enabling the designer to learn the receiver's belief through repeated interactions. Our learning algorithms are *robust* to the receiver's *strategic manipulation* of the learning process of the designer. We study regret relative to two benchmarks to measure the performance of the learning algorithms. The static benchmark is T times the single-period optimum for the designer under a known belief. The *dynamic benchmark*, which is *stronger*, characterizes the *global dynamic optimality* for the designer under a known belief. Our learning algorithms allow the designer to achieve *no regret* against *both benchmarks* at fast speeds of $O(\log^2 T)$.

In [2], the information designer (e.g., platforms in a new market) does *not* know the prior belief of receivers (e.g., experienced users), which is the true distribution over the states. We design learning algorithms so that the designer learns the receiver's prior through repeated interactions. Our algorithms *superiorly* achieve no regret relative to optimality for the known prior at fast speeds: a *tight regret bound* $\Theta(\log T)$ in general and a *tight regret bound* $\Theta(\log\log T)$ for binary actions.

In [3], we consider an investing bidder who changes his value at a cost in truthful mechanisms. The bidder *may not best respond* but uses a *no-regret learning algorithm* to adapt his investment in a dynamic environment. An allocation algorithm's performance is measured by the approximation ratio between the induced welfare and optimal welfare benchmarks. We study how welfare guarantees from the allocation algorithms extend from static to dynamic settings. For the best-in-hindsight benchmark, the approximation ratios in the two settings coincide. For a time-varying benchmark, we characterize *tight* bounds on the ratio. Our work shows how welfare guarantees are maintained *robustly* when a bidder *cannot best respond* but *learns* his investment strategies in *uncertain dynamic environments*.

Representative Papers:

[1] Learning to Design Information (Job Market Paper), with Tao Lin

[2] Information Design with Unknown Prior (ITCS 2025), with Tao Lin

[3] From Best Responses to Learning: Investment Efficiency in Dynamic Environment (working paper), with Qianfan Zhang and Weiqiang Zheng

YUXIN LIU (Homepage, CV, Google Scholar)

Thesis: Learning, Sharing, and Spreading under Privacy Constraints ('26)

Advisor: M.Amin Rahimian, University of Pittsburgh

Brief Biography: I am a third-year PhD student in Operations Research at the Department of Industrial Engineering, University of Pittsburgh. Prior to Pitt, I received my B.S. in Economics from Jilin University and my M.E. in System Engineering from Tianjin University. My research passion lies at the intersection of differential privacy, mechanism design, network science, and information diffusion.

Research Summary.

My research integrates probability, economics, and computer science to study how to balance information utility and privacy in algorithmic decision-making. As data-driven platforms increasingly rely on personal information, society faces a fundamental tension: while sharing data enhances learning and efficiency, it also increases privacy risks. I aim to design mechanisms that protect individual privacy without sacrificing performance, challenging the common belief that privacy inevitably hinders learning or welfare.

My first line of work develops the theoretical foundations of privacy-preserving sequential learning. In this setting, agents act in sequence, observing predecessors' actions but not private signals. I show that with well-designed privacy mechanisms—such as smooth randomized response—learning can, paradoxically, accelerate even under stronger privacy. *Privacy-Aware Sequential Learning* [1] establishes new results on asymptotic efficiency under differential privacy and provides insights for organizations balancing confidentiality and adaptivity.

A second line focuses on privacy-aware optimization and influence in networks. *Seeding with Differentially Private Network Information* [2] studies how to design near-optimal interventions when network data are privatized, bridging differential privacy and operations research to show that effective interventions remain possible under noise.

Finally, Structural Dynamics of Harmful Content Dissemination on WhatsApp [3] examines information diffusion and harmful content on encrypted platforms. Using network reconstruction and causal modeling, I uncover how misinformation and propaganda spread through large-scale, partially observable cascades.

Together, these projects build a unified agenda at the intersection of algorithmic privacy, information design, and social learning, advancing both the science and governance of privacy toward systems that achieve win–win outcomes where privacy is preserved and performance enhanced.

Representative Papers:

[1] Privacy-Aware Sequential Learning (Under Review at Operations Research) with M. A. Rahimian

[2] Seeding with Differentially Private Network Information. (under submission) with M. A. Rahimian and F. Yu

[3] Structural Dynamics of Harmful Content Dissemination on WhatsApp (ICWSM 2026) with M. A. Rahimian and K. Garimella

GARY QIURUI MA (Homepage, CV)

Thesis: Equilibrium Design in Online Marketplaces ('26)

Advisor: David C. Parkes, Yannai A. Gonczarowski, Harvard University

Brief Biography: I am a fifth year Computer Science PhD student at Harvard University, and previously obtained my undergraduate degree in CS from HKUST. During my PhD, I interned at Microsoft Research New England.

Research Summary: Online platforms such as Amazon and DoorDash connect multiple sides of a market—buyers, sellers, and couriers—to facilitate trade. My research examines how these platforms use pricing, matching, and other design tools to balance supply and demand, while maximizing revenue or social welfare.

Pricing. During the COVID lockdown, delivery platforms raised transaction fees charged to restaurants—sometimes reaching thirty percent. [1] studies how a platform's revenue-maximizing transaction fee affects social welfare. In the model, sellers and buyers form a bipartite network, where sellers can pay a fraction of their competitive-equilibrium price to join the platform and access all buyers. The sellers' participation decisions constitute a game, and we show that under mild regulations of the transaction fee, equilibrium welfare can be lower bounded by a fraction of the social optimum.

Matching. Beyond pricing, platforms can shape market outcomes through strategically matching sellers to buyers. For example, the European Commission sued Amazon for favoring sellers that use its fulfillment services in product recommendations to buyers. [2] studies the platform's computational problem of matching sellers to buyers, where the platform needs to balance revenue gains from a trade against the externalities the trade create. The study further establishes lower bounds on social welfare when the platform matches strategically to maximize its own revenue.

Other Perspectives. Platforms need to account for a range of other factors when balancing supply and demand. [3] examines one such factor—tipping. In delivery platforms like DoorDash, tips are specified by buyers and visible to couriers prior to delivery. This visibility causes couriers to favor high-tip orders, leaving low-tip orders unserved. [3] studies how platforms set prices and courier compensations to clear the market while incorporating these optional tips. Other factors, such as market shocks, taxation policies, and fee caps, are explored in [4] through simulation, which studies the platform's optimal responses under different regulatory environments.

Representative Papers:

[1] Platform Equilibrium: Analyzing Social Welfare in Online Market Places (EC'24) with A. Eden, and D. Parkes

[2] Disrupting Bipartite Trading Networks: Matching for Revenue Maximization (EC'24) with Y. Gonczarowski, L. D'Amico-Wong, and D. Parkes

[3] Pricing with Tips in Three-Sided Delivery Platforms (working paper) with Y. Gonczarowski, and D. Parkes

[4] Platform Behavior under Market Shocks: A Simulation Framework and Reinforcement-Learning Based Study (WWW'23) with X. Wang, A. Eden, C. Li, A. Trott, S. Zheng, and D. Parkes

ANIKET MURHEKAR (Homepage, CV, Google Scholar)

Thesis: Algorithms and Solution Concepts for Allocation and Collaboration ('25)

Advisors: Jugal Garg and Ruta Mehta, University of Illinois at Urbana-Champaign

Brief Biography: I am a postdoctoral researcher at Northwestern University, hosted by Edith Elkind. I completed my Ph.D. in Computer Science from UIUC, during which I interned at Google Research (2024) and Adobe Research (2022). I am a recipient of the Simons-Berkeley Research Fellowship (Spring 2026), the Mavis Future Faculty Fellowship, the Siebel Scholarship, and the IIT Bombay Academic Prize. I earned a B.Tech. in Computer Science from IIT Bombay.

Research Summary: I am broadly interested in using ideas from economics, game theory, and social choice theory to develop algorithms and solution concepts that promote incentive alignment, fairness, and efficiency in modern economies. My recent work focuses on allocation and collaboration problems involving entities such as *chores* and *data*, which differ from traditional goods in their nature, their complexity, and the extent to which they have been studied.

Chores are economic bads that impose *dis*utility on their consumers, and their allocation poses challenges that are fundamentally different—and often more difficult—than those encountered in the allocation of goods. In discrete fair division, *envy-freeness up to any item* (EFX) is a central notion of fairness whose existence remains a fundamental and enigmatic open question. For chore allocation under additive preferences, existence of EFX allocations is open even for $n = 3$ agents, and the best result guaranteed a weak approximation of $O(n^2)$-EFX. My recent work [1] provides the first *constant*-factor approximation to EFX by introducing a novel market-based framework for chore allocation called the earning-restricted equilibrium, thereby significantly advancing our understanding of the problem.

Data has emerged as a fundamental economic resource in the AI era that differs from traditional goods due to its non-rival and replicable nature. This motivates a principled study of the *economic foundations of data-sharing frameworks*—including data exchange platforms [2], federated learning (FL) [3,4], and data marketplaces. My work on data exchange economies [2] establishes the existence of exchanges that satisfy core-stability and reciprocity, two benchmark notions of fairness and stability. In federated learning (FL), where clients jointly train models while retaining private datasets, my works develop mechanisms and protocols to mitigate free-riding and incentivize participation while ensuring fairness and efficiency, using ideas from mechanism design [3] and social choice theory [4].

Representative Papers:

[1] Constant-Factor EFX Exists for Chores (STOC 2025)
with J. Garg and J. Qin

[2] On the Theoretical Foundations of Data Exchange Economies (EC 2025)
with B. Chaudhury, J. Garg, and J. Song

[3] You Get What You Give: Reciprocal Fair Federated Learning (ICML 2025)
with B. Chaudhury, R. Mehta, and J. Song

[4] Fair Federated Learning via the Proportional Veto Core (ICML 2024)
with B. Chaudhury, Z. Yuan, B. Li, R. Mehta, and A.D. Procaccia

KALEN PATTON (Homepage, CV, Google Scholar)

Thesis: Online Resource Allocation and Load Balancing Beyond ℓ_p-Norms ('26)

Advisor: Sahil Singla, Georgia Institute of Technology

Brief Biography: I am a 5th year PhD student in the Algorithms, Combinatorics, and Optimization program at Georgia Tech. Prior to my PhD, I completed my undergraduate study at Georgia Tech with a BS in Mathematics and a BS in Computer Science. I am primarily interested in analysis of online algorithms, and my work has been published at top venues including EC, SODA, and FOCS.

Research Summary: My research focuses on developing online algorithms for allocating a sequence of arriving items to agents in order to optimize some objective. Such problems fall into the broad field of *optimization under uncertainty*, as each item must be allocated immediately without knowledge of future arrivals.

Online allocation is a fundamental challenge in many real-world settings, such as in allocating requests to servers in cloud computing, or in allocating ad space in online advertising markets. In such settings, the objectives we seek to optimize may be complex, taking into account problem-specific constraints, utility metrics, and notions of fairness. However, classical results only obtain guarantees for special cases of highly structured objectives. I am interested in extending these algorithms to account for a wider variety of objective functions, and to determine what structural properties of these objectives are needed for online algorithms to be effective. My research examines these questions for both maximization (resource allocation) and minimize (load balancing) problems.

In *online resource allocation* settings, I am particularly interested in online submodular welfare maximization, online bipartite matching, and prophet inequalities. My recent work on these problems [1, 3] has involved developing new primal-dual tools to obtain optimal competitive bounds for various integral and fractional online resource allocation settings. In stochastic settings, my work [4] has also studied to what extent we can design algorithms when we have correlated prior distributions.

In *online load balancing* settings, I am interested in extending poly-logarithmic competitive algorithms from standard makespan-minimization (i.e. minimizing the ℓ_∞ norm of machine loads) to more general classes of convex objectives. In recent work [2], my collaborators and I developed new algorithmic techniques for online load balancing problems beyond traditional "relax-and-round" methods, allowing us to show new competitive bounds by avoiding barriers from large integrality gaps.

.

Representative Papers:

[1] Online Allocation with Concave Diminishing-Returns Objectives (SODA 2026)

[2] Integral Algorithms for Online Set Cover and Load Balancing (FOCS 2025) with T. Kesselheim, M. Molinaro, and S. Singla

[3] The Online Submodular Assignment Problem (FOCS 2024) with D. Hathcock, B. Jin, S. Sarkar, and M. Zlatin

[4] Improved Mechanisms and Prophet Inequalities for Graphical Dependencies (EC 2024) with V. Livanos and S. Singla

TOMASZ PONITKA (Homepage, CV, Google Scholar)

Thesis: Algorithms in Conflict: Optimal Tradeoffs in Fair Division and Contract Design ('26)

Advisor: Michal Feldman, Tel Aviv University

Brief Biography: I am a fourth-year Ph.D. student at Tel Aviv University, advised by Michal Feldman. During my Ph.D., I spent nine months at Sapienza University of Rome, hosted by Stefano Leonardi, and one month at Stanford University, hosted by Aviad Rubinstein. I completed my undergraduate studies at the University of Oxford, where I worked with Elias Koutsoupias.

Research Summary: I am broadly interested in designing algorithms that reconcile conflicting interests of economic agents in complex combinatorial environments. More specifically, my research studies two classic problems in economics, fair division and contract design, through the lens of theoretical computer science. The core contribution of my work is a set of new positive and negative bounds that characterize the Pareto frontier between conflicting objectives in these settings, such as fairness versus efficiency and expressiveness versus representation error.

Fair Division. An example of my contribution to fair division appears in [1]. In this work, we study how to allocate indivisible resources among agents using randomized lotteries. We show that it is possible to design a lottery that achieves a "best-of-both-worlds" guarantee: it is approximately envy-free in expectation (ex-ante 1/2-EF) and approximately envy-free up to any good in every realized outcome (ex-post 1/2-EFX). Notably, we provide the first algorithm achieving such guarantees for subadditive valuations and complement this with impossibility results that partially characterize the Pareto frontier between ex-ante and ex-post fairness. Our approach builds on a careful randomization of the classic Envy Cycles procedure, derived through a novel connection to Markov chain decomposition.

Contract Design. An example of my work on contract design is in [2]. Here, we study how to incentivize an agent to exert costly effort when the agent's type is unknown and must be learned from data. We introduce the concept of pseudo-dimension from statistical learning theory into this setting to design sample-efficient algorithms. Our main results establish nearly tight tradeoffs between pseudo-dimension, sample complexity, and representation error across key contract classes, establishing nearly optimal upper and lower bounds. We further extend our analysis to combinatorial settings, where the agent can choose among subsets of actions. This extension uncovers a new connection between the sample complexity of learning near-optimal linear contracts and the number of critical values previously used to bound the time complexity of computing such contracts.

Representative Papers:

[1] Breaking the Envy Cycle: Best-of-Both-Worlds Guarantees for Subadditive Valuations (EC 2024) with M. Feldman, S. Mauras, and V. V. Narayan.

[2] The Pseudo-Dimension of Contracts (EC 2025)
with P. Dütting, M. Feldman, and E. Soumalias.

PRASANNA RAMAKRISHNAN (Homepage, CV, Google Scholar)

Thesis: The Possibility of Approximately Optimal Social Choice ('26)

Advisors: Moses Charikar and Li-Yang Tan, Stanford University

Brief Biography: I am a PhD student in computer science at Stanford University. I graduated with a BS in mathematics and an MS in computer science from Stanford in 2020. My research is in theoretical computer science, with a focus on computational social choice.

Research Summary: My work uses the lens of theoretical computer science to study how groups of individuals can make collective decisions. Decisions like these, broadly categorized as *elections*, pervade a diverse range of settings, from choosing political leaders to selecting new hires, declaring winners of competitions, picking a restaurant with friends, and fine-tuning language models using human preferences.

From Condorcet's paradox (1785) to Arrow's impossibility theorem (1950), centuries of classical social choice theory have carried a largely pessimistic message: many natural criteria for voting rules are either unattainable or mutually incompatible. On the other hand, theoretical computer science has pragmatic ways of addressing its own kind of impossibilities: when finding an optimal solution is intractable, *relax* the problem and find an *approximately optimal* solution instead. By borrowing the philosophy of approximation algorithms, and harnessing its rich toolkit of probabilistic, game-theoretic, and combinatorial techniques, my work seeks to develop a more optimistic theory of social choice.

In research recognized with a Best Paper Award at SODA 2024 [1], we designed a randomized voting rule with provably better approximation guarantees than any deterministic rule in the popular *metric distortion* model, breaking a longstanding barrier. More recently, we studied how Condorcet's paradox changes in multi-winner elections. In 2011, Elkind, Lang, and Saffidine asked if it is always possible to select a small set of winners that are *collectively* preferred over every candidate by a majority of voters. While one might expect the size of the set to need to grow with the number of candidates, we showed that in *every* election just six candidates suffice [2].

Representative Papers:

[1] Breaking the Metric Voting Distortion Barrier (JACM 2024, SODA 2024) with M. Charikar, K. Wang, and H. Wu.

[2] Six Candidates Suffice to Win a Voter Majority (STOC 2025) with M. Charikar, A. Lassota, A. Vetta, and K. Wang.

[3] Approximately Dominating Sets in Elections (SODA 2026) with M. Charikar and K. Wang.

[4] Distortion in Metric Matching with Ordinal Preferences (EC 2023) with N. Anari and M. Charikar.

KIRAN ROKADE (Homepage, CV, Google Scholar)

Thesis: Large Network Games: Equilibrium Analysis, Computation and Learning ('26)

Advisor: Francesca Parise, Cornell University

Brief Biography: Kiran Rokade is a PhD candidate at Cornell working at the intersection of game theory, networks and dynamical systems theory, planning to defend his thesis by May 2026. Prior to this, he obtained his MS in Electrical Engineering from the Indian Institute of Technology (IIT) Madras where he worked on networked dynamical systems and distributed optimization algorithms. He is on the job market for a postdoctoral position starting Fall 2026.

Research Summary: Systems in which a large number of decision makers interact via a network are ubiquitous, e.g., social networks and networks of firms, modeled by network games. While equilibrium outcomes of small network games are well-understood, the analysis and algorithms often do not scale well with the number of players in large network games, e.g., dynamic pricing by ∼2 million sellers on Amazon, opinion formation by ∼3 billion users on Facebook. In my PhD, I work on developing tractable methods for equilibrium analysis, computation and independent learning in large network games.

Graphon games are infinite population limits of network games. In [1], we characterize the relationship between (possibly infinite) equilibria of graphon games and associated large network games. Using this relationship, we develop computationally efficient algorithms to compute approximate Nash equilibria of large network games via commonly-observed network structures, such as low rank and community blocks.

Players often adapt their strategies dynamically via interactions with their neighbors in the network. In [2], we show that network games in which the network is "nearly symmetric" are α-potential games, where α depends on the amount of asymmetry in the network. Based on α, we derive two algorithms for players to learn approximate Nash equilibria. We show that α scales well with the network size for several network models. In [3], we assume that players update their strategies using gradient-based learning, where a new network is sampled at each time step from an underlying network. We show that the sequence of learned strategies converges to an approximate Nash equilibrium of a game played over the underlying network, and that players' regret over time is sublinear, under a monotonicity assumption.

Representative Papers:

[1] Graphon Games with Multiple Equilibria: Analysis and Computation (EC 2023, accepted at Mathematics of Operations Research) with F. Parise

[2] Asymmetric Network Games: α-Potential Function and Learning (ArXiv) with A. Jain, F. Parise, V. Krishnamurthy, E. Tardos

[3] Learning in Time-Varying Monotone Network Games with Dynamic Populations (ArXiv, preliminary version in CDC 2023) with F. Al Taha (first author), F. Parise

SHIRI RON (Homepage, CV)

Thesis: Mechanisms for Strategic Agents: An Exploration of Incentive Compatibility Notions ('25)

Advisor: Shahar Dobzinski, Weizmann Institute of Science

Brief Biography: I am currently a postdoctoral researcher at Tel Aviv University, hosted by Michal Feldman and Inbal Talgam-Cohen. In the summer of 2025, I obtained my PhD in Computer Science at the Weizmann Institute of Science, advised by Shahar Dobzinski. During my studies, I have interned at Microsoft Research with Moshe Babaioff and have also worked on the 5G spectrum auction under the guidance of Liad Blumrosen. I am fortunate to be a recipient of a number of fellowships and awards, including the Azrieli Fellowship, Maschler Prize and the Dimitris N. Chorafas Award.

Research Summary: My research investigates the communication complexity of welfare maximization in combinatorial auctions. For this investigation, we take the lens of approximation algorithms since exact efficiency is communication-wise infeasible for most valuation classes.

Our goal is to design auctions that are incentive-compatible, approximately optimal and feasible from a communication complexity perspective. Despite significant interest, it is unknown whether such auctions exist. It is conjectured that the answer to this question is negative, but we lack the techniques to prove it.

So far, most of the effort has been put into understanding the power of ex-post incentive compatible mechanisms, i.e. mechanisms that admit strategies that form a Nash equilibrium. The notion of dominant-strategy mechanisms, i.e., mechanisms that admit dominant strategies, has received less attention. In [1], we distinguish ex-post incentive compatible mechanisms and dominant-strategy mechanisms by showing that for gross substitutes, which are the frontier for which welfare maximization is "easy" communication-wise, exact efficiency combined with implementation in dominant strategies necessitates exponential communication.

We then consider the notion of obvious strategy-proofness. We show that even for additive or unit demand bidders, where welfare can be maximized by dominant strategy mechanisms with polynomial communication, approximate efficiency and obvious strategy-proofness are incompatible for deterministic mechanisms [2]. In fact, even randomized obviously strategy-proof mechanisms cannot extract more than $\frac{7}{8}$ of the optimal welfare in expectation [3].

By investigating this hierarchy of incentive compatibility notions and characterizing what is achievable at each level, we clarify connections among them and inform the design of mechanisms with more reliable behavioral guarantees.

Representative Papers:

[1] On the Hardness of Dominant Strategy Mechanism Design (STOC 2022) with Shahar Dobzinski and Jan Vondrák

[2] Impossibilities for Obviously Strategy-Proof Mechanisms (SODA 2024)

[3] On the Power of Randomization for Obviously Strategy-Proof Mechanisms (AAAI 2025) with Daniel Schoepflin

SHAUL ROSNER (Homepage, CV, Google Scholar)

Thesis: Models of Congestion with Non-Standard Utilities ('24)

Advisor: Tami Tamir, Reichman University

Brief Biography: Shaul Rosner is a postdoctoral researcher at Tel Aviv University hosted by Prof. Michal Feldman and Prof. Inbal Talgam Cohen. He obtained his PhD in Computer Science at Reichman University in 2024 advised by Prof. Tami Tamir.

Research Summary: During my PhD, I focused on analysis of the stability and efficiency of models which depend on user congestion. Congestion on a resource can have a significant effect on the quality of a service received by a user. This impact can be both negative (traffic congestion, server load), and positive (shared cost). In particular, my work focuses on alternatives to the standard goals of cost minimization/maximization, which studies in behavioral science suggest do not necessarily fit users' behavior in real-life applications. My work aims to narrow this gap, by developing and analyzing new models describing these behaviors.

An example is [1,2], where we considered players which value overtaking competitors over minimizing their own cost. For both negative congestion ([1]) and positive congestion ([2]), this creates models which are less stable than their standard cost minimization variants. However, we showed interesting non-trivial subclasses for which an equilibrium can be computed efficiently. For these models, I studied the existence of Nash equilibrium profiles and algorithmic approaches for finding and approximating them, hardness of computation and approximation of finding such profiles, and equilibrium inefficiency. For computationally hard problems, I also considered ways to restrict the model to create classes of computationally feasible instances.

Beyond congestion and contract models, I have also studied related algorithmic problems, including maximal matching [4] and Nash flows over time [3]. These projects explore strategic and combinatorial structures, highlighting the themes of equilibrium existence and computation that appear in my main line of research.

More recently, I have started working on problems in the space of algorithmic contract design. In particular, I am studying the impact of non-discrete actions on the approximability of optimal contracts. I am also looking at modifications of existing principal-agent models, in order to better pinpoint the inefficiencies of linear contracts when compared to general contracts.

Representative Papers:

[1] Scheduling Games with Rank-Based Utilities (SAGT 2020, Games and Economic Behavior 2023) with T. Tamir

[2] Cost-Sharing Games with Rank-Based Utilities (SAGT 2022, Theoretical Computer Science 2025) with T. Tamir

[3] Nash Flows Over Time with Tolls (Accepted WINE 2025) with M. Schröder, and L. Vargas Koch

[4] Bipartite Matching with Pair-Dependent Bounds (Under Review) with T. Tamir

ŠIMON SCHIERREICH (Homepage, CV, Google Scholar)

Thesis: Multivariate Complexity and Structural Restrictions in Computational Social Choice ('25)

Advisor: Dušan Knop, Czech Technical University in Prague

Brief Biography: I am a postdoctoral fellow at AGH University of Kraków, working with Piotr Faliszewski. I earned my PhD in theoretical computer science from Czech Technical University in Prague and was a Fulbright Fellow at Penn State with Hadi Hosseini. My research explores algorithmic aspects of collective decision making, including voting, fair division, and coalition formation.

Research Summary: My research lies at the intersection of computer science and economics. I am primarily interested in algorithmic problems in computational social choice and algorithmic game theory–that is, in the design and analysis of collective decision-making processes. I focus on developing algorithms with theoretically guaranteed properties such as fairness and efficiency, while ensuring these algorithms scale well with increasing amounts of data.

Participatory Budgeting. In [1,2], we study *participatory budgeting* and propose a robust framework that enhances the interpretability of election outcomes, thereby promoting greater acceptance and satisfaction among participants.

Fair Division. The goal here is to find a fair allocation of items among agents with individual preferences. In [3], we study the model of *fair division with externalities*, where agents derive utility not only from their own items but may also consider which items not allocated to them are received by others.

Coalition Formation. In coalition formation, the objective is to partition a set of agents into subgroups—called coalitions—such that agents are satisfied with their assigned coalition. In [4], we study *hedonic diversity games*, in which agents are divided into several types and their preferences depend on the composition of their coalition with respect to these types.

Matching in Networks. In [5], we introduce a novel model of *refugee housing*, addressing scenarios where a community must allocate refugees to available homes while respecting the preferences of both refugees and local inhabitants. We identify conditions under which such stable housings exist and develop algorithms to compute them efficiently.

Representative Papers:

[1] Evaluation of Project Performance in Participatory Budgeting (IJCAI '24) with N. Boehmer, P. Faliszewski, Ł. Janeczko, D. Peters, G. Pierczyński, P. Skowron, S. Szufa

[2] Participatory Budgeting Project Strength via Candidate Control (IJCAI '25) with P. Faliszewski, Ł Janeczko, D. Knop, J. Pokorný, M. Słuszniak, K. Sornat

[3] The Complexity of Fair Division of Indivisible Items with Externalities (AAAI '24) with A. Deligkas, E. Eiben, V. Korchemna

[4] Hedonic Diversity Games: A Complexity Picture With More Than Two Colors (AAAI '22 and AIJ) with R. Ganian, T. Hamm, D. Knop, O. Suchý

[5] Host Community Respecting Refugee Housing (AAMAS '23) with D. Knop

DANIEL SCHOEPFLIN (Homepage, CV, Google Scholar)

Thesis: Designing and Analyzing Clock Auctions ('23)

Advisor: Vasilis Gkatzelis, Drexel University

Brief Biography: I am a postdoctoral fellow at Rutgers University - DIMACS. Previously, I was a postdoctoral fellow at the Simons Laufer Mathematical Sciences Institute in the Fall 2023 Mathematics and Computer Science of Market and Mechanism Design program. I obtained my Ph.D. in Computer Science in August 2023 at Drexel University where I was advised by Vasilis Gkatzelis. My thesis was awarded Drexel University's university-wide 2023 Outstanding Dissertation Award.

Research Summary: While I have broad interests in algorithmic game theory and combinatorial optimization, my research centers on *practical mechanism design*.

As an example, my thesis studies the important class of mechanisms known as *(deferred acceptance) clock auctions*. Milgrom and Segal proposed clock auctions as a practical alternative to the well-known sealed-bid auction format. They argued that clock auctions are particularly well-suited for practical application, but, prior to my thesis, relatively little was known about the performance of clock auctions.

In [1], we provide a deterministic, prior-free clock auction for *arbitrary* downward-closed feasibility constraints which *achieves the best possible* welfare guarantees. In [2], we initiate the study of clock auctions with *Bayesian prior information* and show how this information (or randomization) can allow clock auctions to achieve *exponentially improved* approximation guarantees. In [3], we consider a reverse auction setting and examine the well-studied budget feasible mechanism design problem from the perspective of clock auctions. We design a clock auction for this problem which is the *first* deterministic truthful mechanism *of any kind* to give an $O(1)$-approximation to the optimum when the auctioneer has a submodular valuation, resolving one of the most important open problems in this literature.

During my postdoctoral studies, I have continued to study clock auctions but have also expanded my research in other directions. In [4], we study *learning augmented* clock auctions. Our clock auction achieves near-optimal welfare when equipped with an accurate prediction, but still retains the best possible approximation guarantees when the prediction is arbitrarily bad. In [5], we study mechanism design for *consumer surplus* (i.e., utility), providing a mechanism framework giving optimal approximation guarantees for a variety of multi-parameter auction settings.

Representative Papers:

[1] Optimal Deterministic Clock Auctions and Beyond (ITCS '22)
with G. Christodoulou and V. Gkatzelis

[2] Bayesian and Randomized Clock Auctions (EC '22, OR '25)
with M. Feldman, N. Gravin, and V. Gkatzelis

[3] Deterministic Budget-Feasible Clock Auctions (SODA '22, OR (forthcoming))
with E. Balkanski, P. Garimidi, V. Gkatzelis, and X. Tan

[4] Clock Auctions Augmented with Unreliable Advice (SODA '25)
with V. Gkatzelis and X. Tan

[5] Multi-Parameter Mechanisms for Consumer Surplus Maximization (STOC '25)
with T. Ezra and A. Shaulker

EKLAVYA SHARMA (Homepage, CV, Google Scholar)

Thesis: Allocation Problems in Fair Division and Data Markets ('26)

Advisor: Jugal Garg, University of Illinois at Urbana-Champaign (UIUC)

Brief Biography: I am a PhD candidate at UIUC. I completed my Masters from the Indian Institute of Science (IISc), and my Bachelors from BITS Pilani. At UIUC, I received the Mavis Future Faculty Fellowship and the Brainin Fellowship. At IISc, I received the best masters thesis award.

Research Summary: My research is about allocating resources among agents with diverse preferences. Specifically, I have worked on fair item allocation, and recently, the design of data marketplaces.

Economists and computer scientists have developed insightful theories on the equilibrium and dynamics of markets. However, many of these theories are inapplicable when the goods being sold are *data*, because unlike traditional goods, data can be replicated at near-zero cost. Hence, my research aims to better understand the economics of data. In ongoing work, we study how to price m datasets to maximize total revenue, given the valuations of n buyers. We show that, unlike for traditional goods, competitive equilibria do not always maximize revenue, and we give alternative methods for optimally pricing data.

Research on the fair allocation of indivisible items aims to address two key questions: (i) How should we formally define *fairness*? (ii) Do fair allocations always exist, and can they be computed efficiently? My research contributes to both fronts.

EFX is regarded as one of the strongest notions of fairness, but whether EFX allocations exist has been an open problem since 2016. In [2], we propose a slight relaxation of EFX, called EEFX, and show that EEFX allocations always exist and can be efficiently computed for additive valuations. This is arguably the strongest relaxation of EFX for which a positive result is known.

MMS is another strong fairness notion. MMS was shown to be infeasible, and so, its multiplicative approximations were studied. In [3], we improved the best-known approximation factor, and at the same time, significantly simplified the analysis. Moreover, we devised a computer-aided search technique to identify bottlenecks in current techniques. This insight was used in subsequent work by others that further improved the approximation factor. In [1], we analyzed known algorithms from a different perspective to get better *randomized* guarantees for MMS.

In addition to the directions above, I have worked on Nash equilibria for repeated games, distortion of voting rules, and algorithms for geometric bin packing.

Representative Papers:

[1] Improving Approximation Guarantees for Maximin Share (EC'24)
with Hannaneh Akrami, Jugal Garg, Setareh Taki.

[2] New Fairness Concepts for Allocating Indivisible Items (IJCAI'23)
with Ioannis Caragiannis, Jugal Garg, Nidhi Rathi, Giovanna Varricchio.

[3] Simplification and Improvement of MMS Approximation (IJCAI'23)
with Hannaneh Akrami, Jugal Garg, Setareh Taki.

[4] Exploring Relations among Fairness Notions in Discrete Fair Division (under submission, preprint on ArXiv), with Jugal Garg.

SUHO SHIN (Homepage, CV)

Thesis: Delegation: From Classic Auctions to the Modern Digital Economy ('26)

Advisor: MohammadTaghi Hajiaghayi, University of Maryland

Brief Biography: I'm a fourth-year Ph.D. student in Computer Science at the University of Maryland, College Park. Previously, I was a software and machine learning engineer at Coupang and LINE. I completed my M.S. in Electrical Engineering and B.S. in Mathematics at KAIST and was a gold medalist in the Korean Mathematical Olympiad in 2008.

Research Summary: I am interested in mechanism design and market design, broadly construed. A core thread throughout my research is the development of the algorithmic foundations of delegated decision-making under uncertainty, a theme at the intersection of computer science, operations research, and economics. Across classical markets, modern digital platforms, and AI systems, decision-makers increasingly rely on autonomous or strategic agents to act on their behalf. For instance, non-experts delegate decision-making to experts, content platforms delegate high-quality content production to creators, and individuals now largely delegate information-seeking to large language models. A critical challenge in such settings is that delegatees may have ulterior motives, misaligned incentives, or incomplete information, which can lead to undesirable outcomes for the delegator. How can delegation be structured so that the resulting system remains efficient, fair, and economically sustainable?

I study how a principal—such as a regulator, platform, or algorithm designer—can design mechanisms, learning procedures, and information structures that align the incentives of self-interested or information-restricted agents. I investigate this question in various real-world scenarios: delegation of content production in online platforms [1], delegation of choices [2,3], delegation in auctions [4], and delegation in information-seeking process [5]. These results characterize how much the principal's utility decreases compared to the first-best ideal scenario and what efficient algorithms or mechanisms can achieve better outcomes.

My ultimate goal is to build a coherent framework of delegation problems that unifies incentive design and decision making process for human-AI systems, online marketplaces, and various organizations.

Representative Papers:

[1] The Contest Behind the Feed: Optimal Contest for Recommender Systems (working paper, job market paper) with N. Golrezaei, M. Hajiaghayi

[2] Delegated Choice with Combinatorial Constraints (EC 2025, under review at *Operations Research*) with K. Banihashem, M. Hajiaghayi, P. Krysta

[3] Delegation with Costly Inspection (EC 2025) with M. Hajiaghayi, P. Krysta, M. Mahdavi

[4] Gains-from-Trade in Bilateral Trade with a Broker (SODA 2025) with I. Hajiaghayi, M. Hajiaghayi, G. Peng

[5] Tokenized Bandit for LLM Decoding and Alignment (ICML 2025) with M. Hajiaghayi, H. Xu, C. Yang

ANA-ANDREEA STOICA (Homepage, CV, Google Scholar)

Thesis: Algorithmic Design for Social Networks: Inequality, Bias, and Diversity ('22)

Advisor: Augustin Chaintreau, Columbia University; Moritz Hardt, Max Planck Institute for Intelligent Systems

Brief Biography: I am a Research Group Leader in the Social Foundations of Computation Department at the Max Planck Institute for Intelligent Systems in Tübingen, Germany, advised by Moritz Hardt. I completed my Ph.D. in Computer Science at Columbia University in 2022, advised by Augustin Chaintreau. I spent Fall 2022 as a Simons Fellow at the Simons Institute for Theory of Computing in Berkeley, attending the Graph Limits and Processes on Networks program.

Research Summary: My research investigates how AI can add value to online platforms, markets, and social systems: I develop methods that efficiently allocate resources, balance social and optimization objectives, and minimize negative externalities. To solve such tasks, I design algorithms with proven performance guarantees that can effectively solve unsupervised problems at scale. Additionally, I develop statistical methods to handle complex social data influenced by competition and interference, using insights from causal inference and mechanism design. My goal is to understand and improve how platform design affects people's welfare and access to opportunities in the long run. During my PhD in Computer Science and my current postdoctoral appointment, I have contributed novel methodology and theoretical foundations for diagnosing and mitigating inequality in algorithms deployed in social settings.

In my recent work, I combine causal inference tools and auction design to provide statistical guarantees of experiments in which the users' attention is split among multiple inference tasks, constructing provably optimal estimators under equilibrium [2]. Furthermore, I study mechanisms for targeting people with interventions by leveraging an economic model of income fluctuations over time to show that targeting the most vulnerable individuals can be long-term more efficient in improving a population's social welfare than targeting those who benefit immediately [1]. My research also studies trade-offs between objectives in resource allocation systems: I worked on novel algorithms that trace the Pareto front between generally NP-hard optimization objectives (such as clustering) and fairness objectives with strong approximation guarantees for a wide range of objective functions [3].

Representative Papers:

[1] Policy Design in Long-Run Welfare Dynamics (ICLR'25)
with Jiduan Wu, Rediet Abebe, and Moritz Hardt

[2] Causal Inference from Competing Treatments (ICML'24)
with Vivian Y. Nastl and Moritz Hardt

[3] The Fairness-Quality Trade-off in Clustering (NeurIPS'24)
with Rashida Hakim, Mihalis Yannakakis, and Christos Papadimitriou

YIFAN WANG (Homepage, CV)

Thesis: Learning to Price: From Samples to Queries ('26)

Advisor: Sahil Singla, Georgia Tech

Brief Biography: I am a 5th-year Ph.D. student studying in Computer Science at Georgia Institute of Technology, advised by Prof. Sahil Singla. Before that, I received my Bachelor's degree in Computer Science from Tsinghua University. During my Ph.D., I have interned at Google Research (summer 2025).

Research Summary: I am broadly interested in algorithmic game theory, mechanism design, and online algorithms. In particular, my research analyzes models in computational economics from new perspectives, focusing on two main directions:

(1) *Learning Bayesian mechanisms with limited information.* To study the learnability of Bayesian mechanisms, a common approach is to analyze their sample complexity. However, standard formulations assume the seller has access to a large amount of samples, which is often impractical. My research addresses this limitation by restricting the seller's information. In [1], we design a $(1-\epsilon)$-competitive mechanism for the standard online resource allocation problem with large item supply, even when only single sample is provided. In [2], we study the *query complexity* of the posted pricing mechanism for a unit-demand buyer, that is, the number of auctions a seller needs to run to learn a near-optimal pricing mechanism under the constraint that the seller can only observe the buyer's consumption behavior in each auction, which reflects the most realistic feedback available in practice.

(2) *Online models beyond competitive analysis.* Many allocation and auction models in computational economics require the algorithm to make online decisions, where information about the future is revealed over time. Traditionally, the performance of an algorithm is measured by comparing to the optimal offline solution, leading to the framework of competitive analysis. However, this framework is often overly pessimistic, motivating the need to move beyond worst-case guarantees. My research studies the online combinatorial allocation model with stochastic agents, which is known to be tight 0.5-competitive when agents' valuations are XOS. In [3], we show that by switching the benchmark from comparing to the optimal offline solution to the weaker (inefficient) optimal online algorithm, there exists a poly-time algorithm that achieves a $0.5 + \Omega(1)$ approximation when agents are submodular. Furthermore, when agents are unit-demand, we show in [4] that a $0.5 + \Omega(1)$ approximation is achievable even when the arrival order of agents is unknown.

Representative Papers:

[1] Single-Sample and Robust Online Resource Allocation (STOC'25) with Rohan Ghuge and Sahil Singla

[2] Learning Optimal Posted Prices for a Unit-Demand Buyer (EC'25) with Yifeng Teng

[3] Combinatorial Philosopher Inequalities (SODA'26) with Enze Sun and Zhihao Gavin Tang

[4] Online Stochastic Matching with Unknown Arrival Order: Beating 0.5 against the Online Optimum (STOC'25) with Enze Sun and Zhihao Gavin Tang

WENTAO WENG (Homepage, CV)

Thesis: Optimizing Human Decision-Making for Safer Societies ('26)

Advisor: Daniel Freund, Thodoris Lykouris, MIT

Brief Biography: I am a PhD candidate in the department of Electrical Engineering and Computer Science at MIT, with my B.E. in Computer Science from Tsinghua University.

Research Summary: My work optimizes the efficiency of human decision-making for safer societies. Combining tools from machine learning and queueing theory, I design data-driven algorithms and provide practical insights for content moderation and the adjudication of legal proceedings.

My first line of research studies *AI-powered queueing control* to optimize the assignment and prioritization of content-moderation tasks for human reviewers. Building on my experience with a production-scale content moderation system during my internship at Meta, my work uncovers and rectifies fundamental deficiencies of status quo queueing controls, which (a) fail to collect sufficient AI training data [1] and ignore the uncertainty in content virality [2], and (b) rely on fully specified system parameters [3] and centralized decision-making [4], which are impractical as social media platforms employ thousands of outsourced human reviewers.

My second line of research studies *efficient and fair* resource allocation in the adjudication process of humanitarian immigration. Through engagement with domain experts, I study how different queueing designs developed to reduce asylum backlogs may compromise due process for asylum seekers. Specifically, my work characterizes fairness issues created by the use of Last-In-First-Out in asylum interviews [5] and the use of fast tracks in immigration courts [6]. Methodologically, my research draws on and challenges the traditional literature on rational queueing to capture a counterintuitive *benefit of waiting* in asylum systems. I have also developed matching algorithms that safeguard group fairness in AI-based optimization, which support GeoMatch, a nonprofit refugee resettlement platform [7].

Representative Papers:

[1] Learning to Defer in Congested Systems: The AI-Human Interplay (Major Revision at Operations Research), with T. Lykouris.

[2] Scheduling with Uncertain Holding Costs and its Application to Content Moderation (Major Revision at Management Science), with C. Gocmen, T. Lykouris, and D. Sinha.

[3] The Transient Cost of Learning in Queueing Systems (NeurIPS'23, R&R at Operations Research), with D. Freund and T. Lykouris.

[4] Efficient Decentralized Multi-agent Learning in Asymmetric Bipartite Queueing Systems (COLT'22, Operations Research'24), with D. Freund and T. Lykouris.

[5] Regulating Wait-Driven Requests in Queues (EC'25), with D. Freund and D. Hausman.

[6] The Dedicated Docket in U.S. Immigration Courts: An analysis of fairness and efficiency properties (EC'24, Major Revision at MSOM), with D. Freund.

[7] Group fairness in dynamic refugee assignment (EC'23, Major Revision at Operations Research), with D. Freund, T. Lykouris, E. Paulson, and B. Sturt.

NICHOLAS WU (Homepage, CV)

Thesis: Essays in Microeconomic Theory ('26)

Advisor: Dirk Bergemann (Yale University), Johannes Hörner (Toulouse School of Economics)

Brief Biography: I am a PhD candidate in economics at Yale University. My research in microeconomic theory explores how incentives and information shape behavior in dynamic and strategic environments. Prior to my PhD, I obtained my B.S. and M.Eng. from MIT in computer science.

Research Summary: My research focuses on understanding how information and incentives shape behavior in dynamic and strategic environments, with applications to the digital economy and technological innovation.

In many important settings, actions are taken without full knowledge of their consequences, and those actions shape the information that can be learned. I study these settings using tools from dynamic games and mechanism design. In [1], I examine how a problem-solver allocates effort when facing uncertainty about which approach to take and how quickly a solution should arrive. Learning induces the optimal policy to alternate between exploring new approaches and revisiting previously discarded ones. I apply this framework to startup investment, providing a novel rationale for frontloaded incentive contracts. In [2], we examine how a privately informed seller can overcome buyer skepticism by exploiting the possibility of learning with dynamic pricing. We consider a dynamic trade environment where a seller has private information about product match quality, which affects the buyer's private consumption experience. We show that equilibrium mechanisms take the form of free/discounted trials or dynamic tiered pricing—prevalent features in digital markets—and that access to consumer data can reduce sellers' revenue.

In related work, I use mechanism and information design tools to study digital advertising and platform competition. [3] presents a model where platforms can leverage data advantages in auction design. We show the platform-optimal mechanism is a managed campaign that conditions on-platform prices on off-platform prices, attaining vertical integration profits while increasing off-platform prices and decreasing consumer surplus relative to data-augmented auctions. In [4], we analyze equilibrium properties of auto-bidding algorithms and derive the revenue-maximizing algorithm, providing conditions where an equilibrium with auto-bidding algorithms generates more or less revenue compared to manual bidding.

Representative Papers:

[1] Solving Problems of Unknown Difficulty? (Job Market Paper)

[2] From Doubt to Devotion: Trials and Learning-Based Pricing (accepted, Journal of Political Economy) with T. Gan

[3] How Do Digital Advertising Auctions Impact Product Prices? (Review of Economic Studies 92 (4), 2330-2358) with D. Bergemann and A. Bonatti

[4] Bidding with Budgets: Data-Driven Bid Algorithms in Digital Advertising (International Journal of Industrial Organization, 103172)
with D. Bergemann and A. Bonatti

SHENGWEI XU (Homepage, CV)

Thesis: Aligned Information Elicitation for Text ('25)

Advisor: Grant Schoenebeck, University of Michigan

Brief Biography: Shengwei is a 5th-year PhD student in Information Science at the University of Michigan, advised by Prof. Grant Schoenebeck. He obtained a BSc in Computer Science from Peking University. During his PhD, Shengwei interned at Google.

Research Summary: I study incentives and strategic behavior in interactions among humans, AI agents, and algorithms, leveraging algorithmic game theory, information elicitation, and large language models (LLMs) to build trustworthy and sustainable AI ecosystems. This agenda builds on my previous work in algorithmic game theory, with a focus on peer prediction.

Information elicitation and evaluation mechanisms for text.[1,2] In the strategic setting, if an evaluator can be "gamed", high scores may reflect surface-level tricks rather than genuine quality. To address this issue, We introduce GEM (Generative Estimator for Mutual Information) that captures semantic informativeness of the textual report, building on information theory and peer prediction. In our empirical study on ICLR review datasets, GEM differentiates between LLM-generated fictitious reviews and genuine human reviews, and demonstrates improved accuracy and robustness against manipulations than baselines, including LLM-as-a-Judge.

Building on GEM, we develop GRE-bench (Generating Review Evaluation Benchmark), where LLMs are tasked with writing reviews for recently published papers, ensuring that evaluation data is not leaked in model pretraining corpora. This addresses data contamination and helps establish fairer and more reliable benchmarks by continuously curating fresh, unseen tasks.

Ad Insertion in LLM-generated Responses.[4] Given the cost of training and deploying LLMs, a natural approach of monetization is to integrate ads into LLM responses, similar to current search engines. We develop a framework for ad auctions within LLM chatbots. Through user studies and surveys, we identify people's ethical expectations about LLM advertising. Informed by these insights, our framework integrates these legal and moral constraints into the incentive-compatible auction mechanism: sponsored content must be clearly disclosed; ads should be contextually coherent to protect user experience; and the framework mitigates hallucination-caused misinformation by design.

Representative Papers:

[1] Benchmarking LLMs' Judgments with No Gold Standard. (ICLR 2025)
with Y. Lu, Y. Zhang, Y. Kong, and G. Schoenebeck

[2] Eliciting Informative Text Evaluations with Large Language Models. (EC'24)
with Y. Lu, Y. Zhang, Y. Kong, and G. Schoenebeck

[3] Stochastically Dominant Peer Prediction. (NeurIPS 2025)
with Y. Zhang, D. Pennock, and G. Schoenebeck

[4] Ad Insertion in LLM-generated Responses. (Under-review, presented at the Workshop on Frontiers of Online Advertising at EC'25)
with Z. Chen, Z. Huang, G. Schoenebeck, and X. Deng

AVIV YAISH (Homepage, CV, Google Scholar)

Thesis: Intelligent Economic Agents, Cryptocurrencies & Distributed Ledgers ('24)

Advisor: Aviv Zohar, The Hebrew University

Brief Biography: Aviv is a postdoc at Yale University, where he makes and breaks distributed systems by bridging economic theory and practice. His approach is driven by a philosophy of constructive deconstruction: pushing systems to their limits is key to making them robust. Aviv's work has been recognized across several fields: security (*CCS Distinguished Paper award*), economics (*CBER Best Paper award*), and industry (*three prizes* from various organizations). He earned his Ph.D. in Computer Science from the Hebrew University (HUJI), where he was the sole lecturer for large-scale courses and won a teaching award. During his studies, he served as a consultant at Matter Labs. His honors include the *AIANI* and *Jabotinsky* fellowships, and inclusion in *HUJI's top 10 CS teaching staff of '20*, *CBER's Top PhD Graduates of '23-'24*, and *HUJI's 40 Under 40 of '25* lists.

Research Summary: Financial systems are increasingly decentralized, demanding a fusion of distributed systems security and economic theory. A key enabler of this change, blockchain technology, promises more private and egalitarian economic mechanisms, built by facilitating consensus between pseudonymous actors. However, these systems' theoretical security may mask significant real-world risks.

My work bridges this gap between theory and practice. In [4], we devise and analyze the first practical attack on a major consensus protocol and mitigations for it, both theoretically and experimentally. By uncovering evidence of a similar attack being performed on Ethereum, the 2$^{\text{nd}}$ most popular cryptocurrency, we resolve a decade-old pursuit for evidence of attacks on major systems. In [5], we show that decentralized financial applications are susceptible to manipulations of underlying consensus protocols, and derive system parameters robust to such deviations.

To allocate transactions to blocks, cryptocurrencies use auction-esque transaction fee mechanisms (TFMs). Roughgarden (JACM'24) asks whether there is a TFM that is incentive compatible for users and system operators, and is also resistant to collusion between both. We resolve this question in the negative in [1,2] for deterministic TFMs by characterizing collusion-resistant designs, and show limits on the efficiency of randomized protocols. We ground TFM theory in practice by devising attacks against popular designs in [3], and also present safer designs.

Representative Papers:

[1] Transaction Fee Mechanisms Robust to Welfare-Increasing Collusion (Minor Revision at Games and Economic Behavior, '25) with Y. Gafni

[2] Barriers to Collusion-resistant Transaction Fee Mechanisms (EC'24) with Y. Gafni

[3] Speculative Denial-of-Service Attacks in Ethereum (USENIXSEC'24) with K. Qin, L. Zhou, A. Zohar, and A. Gervais

[4] Uncle Maker: (Time)Stamping Out the Competition in Ethereum (CCS'23 Distinguished Paper Award) with G. Stern, and A. Zohar

[5] Blockchain Stretching & Squeezing: Manipulating Time for Your Best Interest (EC'22) with S. Tochner, and A. Zohar

KUNHE YANG (Homepage, CV, Google Scholar)

Thesis: Designing and Evaluating AI Systems in Strategic and Agentic Environments ('26)

Advisor: Nika Haghtalab, UC Berkeley

Brief Biography: I am a fifth-year PhD student in Computer Science at the UC Berkeley, advised by Prof. Nika Haghtalab. During my Ph.D., I have had the pleasure of visiting the TTIC in Summer 2022 and Harvard University in Summer 2025, hosted respectively by Prof. Avrim Blum and Prof. Ariel Procaccia. I was a finalist for the 2023 Meta Research PhD Fellowship in the Economics and Computation track. I earned my bachelor's degree from Yao Class at Tsinghua University.

Research Summary: As AI systems operate in environments with humans and other AI agents, their behavior and performance are shaped by the incentives and strategic behaviors that arise within their interactions. My research develops theoretical foundations for designing and evaluating AI systems in such strategic environments, combining tools from learning theory, game theory, and economics to ensure robustness, efficiency, and alignment with human and societal objectives.

On the learning side, my research focuses on learning human-centric policies from behavioral or preference feedback that is diverse, sparse, and strategically shaped by incentives and information asymmetries. In the context of AI alignment, we introduce the notion of *distortion*, adapted from social choice theory, to capture pluralistic AI alignment from heterogenous feedback [1]. This notion reveals the failure modes of alignment methods such as RLHF, which can systematically misrepresent the average human utilities of diverse users, and offers guidance for designing more robust optimization approaches grounded in equilibrium concepts from game theory. In principal–agent settings, we propose the *Calibrated Stackelberg Games* framework [4], which models agents as using calibrated forecasts to make data-driven responses under uncertainty about the principal's strategy. We design algorithms for achieving calibration on the agent side and near-optimal learning for the principal. I also design *multi-agent platforms* that enable efficient and stable collaboration among strategic agents via algorithmic information design [2].

On the evaluation side, I focus on designing *incentive-aware evaluation metrics* that are robust to strategic manipulations. In [3], we initiate the study on the *truthfulness of calibration*, a property that requires a calibration measure to automatically incentivize the predictor to incorporate their most accurate information and report their true beliefs.

Representative Papers:

[1] Distortion of AI Alignment: Does Preference Optimization Optimize for Preferences? (NeurIPS 25) with P. Gölz and N. Haghtalab

[2] Platforms for Efficient and Incentive-Aware Collaboration (SODA 25) with N. Haghtalab and M. Qiao

[3] Truthfulness of Calibration Measures (NeurIPS 24) with N. Haghtalab, M. Qiao, and E. Zhao

[4] Calibrated Stackelberg Games: Learning Optimal Commitments Against Calibrated Agents (NeurIPS 23, Spotlight) with N. Haghtalab and C. Podimata

KONSTANTIN ZABARNYI (Homepage, CV, Google Scholar)

Thesis: Information Design in the Twenty-First Century: A Computer Science Perspective ('24)

Advisors: Dirk Bergemann, Yang Cai

Brief Biography: I am a postdoc at the Center for Algorithms, Data, and Market Design at Yale. I earned my PhD in computer science at the Technion – Israel Institute of Technology, advised by Yakov Babichenko and Inbal Talgam-Cohen. My PhD thesis won the SIGecom Doctoral Dissertation Award. During my PhD years, I was awarded the PBC (VATAT) Scholarship in Data Science, the Student Research Prize for Cross-PI Collaboration in Data Science Funded by PBC (VATAT), and an Excellent Teaching Assistant Award from the Technion. Earlier, I obtained my BSc in computer science and math at the Technion (summa cum laude).

Research Summary: I study the intersection of computer science, operations research, and theoretical economics. A significant share of my research is analyzing strategic information revelation, including the models of Bayesian persuasion and cheap talk, through an algorithmic lens. Both models involve a *sender* strategically revealing information to *receivers*. However, while in Bayesian persuasion the sender can commit to an information disclosure policy, in cheap talk the sender has no commitment power. I also study simple, approximately optimal mechanism design.

I aim to derive more realistic versions of existing models. In the context of Bayesian persuasion, I weaken stringent assumptions on the sender's power. In particular, I analyze intricate communication structures inspired by real-world applications. In [4], we study scenarios in which, counterintuitively, increased communication between receivers hurts them and benefits the sender. In [3], we are the first to systematically study computational aspects of cheap talk, thus analyzing stripping the sender away from her commitment power through an algorithmic lens. In [1], we consider another relaxation of the sender's power – namely, assuming that the sender only has partial knowledge of the receiver's utility (i.e., we analyze a robust Bayesian persuasion framework); we show that a qualitative knowledge of the receiver's utility suffices to guarantee the sender a surprisingly low Savage regret.

In general, I study a range of robust frameworks and the interplay between robustness paradigms. In [2], we show that for robust aggregation of binary recommendations, the three common robustness paradigms – maximin, regret and approximation ratio – are equivalent when the number of recommendations is large. Discovering further settings with such equivalence is an interesting open question.

Representative Papers:

[1] Regret-Minimizing Bayesian Persuasion (EC'21 and Games and Economic Behavior) with Y. Babichenko, I. Talgam-Cohen, and H. Xu

[2] A Random Dictator Is All You Need (EC'23 and American Economic Journal: Microeconomics) with I. Arieli, Y. Babichenko, and I. Talgam-Cohen

[3] Algorithmic Cheap Talk (EC'24 and under review at Operations Research) with Y. Babichenko, I. Talgam-Cohen, and H. Xu

[4] Persuasion Gains and Losses from Peer Communication (under review at Theoretical Economics) with T.T. Kerman, and A.P. Tenev

WENXIN ZHANG (Homepage, CV, Google Scholar)

Thesis: Dynamic Resource Allocation for Large-Scale Service Systems ('26)

Advisor: Santiago Balseiro and Will Ma, Columbia University

Brief Biography: I am a PhD candidate in the Decision, Risk, and Operations Division at Columbia Business School. I hold a B.E. in Industrial Engineering from Tsinghua University in 2021, and interned at Google Research in the summers of 2024 and 2025. I have been recognized with several honors, including being a finalist in the 2025 Applied Probability Society Best Student Paper Competition.

Research Summary:

My research focuses on dynamic resource allocation: I build application-grounded models and design deployable algorithms with provable guarantees. This agenda spans classic operations research and computer science settings (pricing, matching) and emerging challenges in artificial intelligence (AI) serving. Looking ahead, my agenda is to make AI serving more efficient through better resource allocation.

New frameworks for dynamic pricing and matching. For reusable resources with general service time distributions, an optimal online policy must track each service unit's usage history, leading to an exponential state space. We overcome this curse of dimensionality with a stock-dependent policy that only tracks the number of busy units, which can significantly outperform the best static policy both in theory and in practice [1]. For matching markets with effectively infinite (feature-rich, continuous) demand/supply types, canonical approaches break. We model them as feature vectors drawn from some known distributions and propose a model-predictive control heuristic that is provably near-optimal [2].

Making AI serving more efficient. In [3], we study how to route requests across geographically distributed AI data centers. To capture the complexity of GenAI inference such as request batching, we abstract each data center with a workload-dependent service rate. Based on this, we show that routing based on the highest marginal service rate is asymptotically optimal for both throughput and latency. This work was conducted while I interned at Google Research, and now informs the design of inference serving for a major AI company. In [4], we study how to optimize the tail latency in conversational AI such as chatbots. We propose a new stochastic model of stateful conversations that captures temporal locality and design a simple modification of the Least-Recently-Used (LRU) policy that is provably optimal for our tail metric.

Representative Papers:

[1] Dynamic Pricing for Reusable Resources: The Power of Two Prices (*Operations Research*, 2025) with S. Balseiro and W. Ma

[2] Feature-Based Dynamic Matching (*Operations Research*, forthcoming; *EC*, 2023) with Y. Chen, Y. Kanoria, and A. Kumar

[3] Distributed Load Balancing with Workload-Dependent Service Rates (*EC*, 2025; under review at *Operations Research*) with S. Balseiro, R. Kleinberg, V. Mirrokni, B. Sivan, and B. Wydrowski

[4] Tail-Optimized Caching for LLM Inference (*NeurIPS*, 2025) with Y. Li, C. Moallemi, and T. Peng

YICHI ZHANG (Homepage, CV)

Thesis: Incentivizing Effort and Honesty for High-quality Information ('24)

Advisor: Grant Schoenebeck, University of Michigan

Brief Biography: I'm a postdoctoral associate at the Center for Discrete Mathematics and Theoretical Computer Science (DIMACS), Rutgers University, hosted by David Pennock and Lirong Xia. Before my current position, I received my Ph.D. from the School of Information, University of Michigan. I earned my B.S. from Shanghai Jiao Tong University, China.

Research Summary: My research aims to mitigate *misalignment* in decision and AI policies, where the achieved outcomes diverge from the intended outcomes. Misalignment often stems from issues of data quality, driven by two main challenges: (1) the divergent incentives of data producers and consumers, and (2) uncertainty about the reliability of data producers. To mitigate misalignment, I develop theory-grounded, manipulation-resistant *automatic evaluators* that can incentivize genuine human feedback, assess data quality, and steer AI training.

The first challenge concerns incentives: data producers seek higher pay for less effort, which conflicts with the requester's goal of high-quality data. Since ground-truth verification is rarely scalable, my research develops *peer prediction* mechanisms that reward high-effort human feedback relative to peers. In one representative work [1], I design mechanisms whose score distribution under high-effort reporting *first-order stochastically dominates* that under any low-effort manipulations, guaranteeing incentive alignment under any monotone reward scheme.

The idea of peer prediction can be adapted to design evaluators to benchmark AI judges and assess human contributions. This addresses the second challenge, data evaluation, which is increasingly important as modern AI participates in both data generation and evaluation. In [2], I examine how and when peer prediction can measure the informativeness of data producers, even when an unknown fraction rely on LLMs to generate superficial responses. Leveraging these insights, I further investigate how to select complementary AI agents to form better collaboration [3].

A closely related application is peer review, where authors aim to maximize acceptance, while conferences aim to maximize quality. When review policies overlook strategic author behavior, especially under noisy peer reviews, misaligned objectives translate into misaligned decisions. My research investigates the system-level consequences of this misalignment and how the design of review policies can mitigate it [4].

Representative Papers:

[1] Stochastically Dominant Peer Prediction. (NeurIPS'25) with S. Xu, D. Pennock, and G. Schoenebeck

[2] Evaluating LLM-corrupted Crowdsourcing Data Without Ground Truth. (NeurIPS'25) with J. Pang, Z. Zhu, and Y. Liu

[3] Mixture of Complementary Agents for Robust LLM Ensemble (Working paper) with K. Lu, Y. Zhang, J. Gao, L. Xia, and F. Yu

[4] A System-Level Analysis of Conference Peer Review (EC'22, revision at OR) with F. Yu, G. Schoenebeck, and D. Kempe

Index

Regulating Matching Markets with Distributional Constraints

KEI IKEGAMI
University of Tokyo
and
ATSUSHI IWASAKI
University of Electro-Communications
and
AKIRA MATSUSHITA
Kyoto University
and
KYOHEI OKUMURA
University of Wisconsin–Madison

Distributional constraints arise naturally in many matching markets, requiring the number of matches of specific types to satisfy predetermined bounds. This article reviews recent developments in the design and analysis of matching markets under such constraints. We discuss existing theoretical and empirical approaches. We then describe the results of [Ikegami et al. 2025], which develops a new framework for matching markets with distributional constraints and applies it to the Japan Residency Matching Program. The analysis illustrates how data can be used to evaluate regulatory instruments and to construct subsidy schemes that implement constrained-efficient outcomes.

Categories and Subject Descriptors: J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms: Economics, Design, Human factors

Additional Key Words and Phrases: Matching with Distributional Constraints, Discrete Choice, Taxation, Subsidy

1. INTRODUCTION

Many matching markets are subject to distributional requirements that limit how matches can be allocated across groups or locations. These requirements are commonly implemented through rules that restrict the number of matches of particular types. Examples include affirmative action policies in college admissions, gender quotas in electoral systems, and regional caps in the Japan Residency Matching Program, which limits placements in urban hospitals to maintain adequate staffing in rural areas.

Cap-based policies are widely used as they are straightforward to implement. However, a cap is a blunt instrument that may prevent high-surplus matches. Monetary interventions, such as taxes and subsidies, offer a natural yet comparatively

Authors' addresses: ikegamikei@e.u-tokyo.ac.jp, atsushi.iwasaki@uec.ac.jp, amatsushita@i.kyoto-u.ac.jp, kyohei.okumura@wisc.edu

underexplored alternative in this context. Their potential advantage lies in their ability to account for the intensity of preferences: a well-calibrated targeted subsidy could, in principle, influence marginal participants without generating large welfare losses. Whether this theoretical potential translates into meaningful efficiency gains in practice, and under what conditions, remains a central open question for both theory and policy.

This article first reviews recent theoretical and empirical work on matching markets with distributional constraints. It then describes the framework developed in [Ikegami et al. 2025], which provides a unified approach to analyzing cap-based and monetary interventions in matching markets with distributional constraints. Applied to the Japan Residency Matching Program, the framework uses aggregate match data to quantify the effects of alternative policies and to construct subsidy schemes that implement constrained-efficient outcomes, defined as outcomes that maximize social surplus subject to the constraints.

2. EXISTING WORK ON MATCHING WITH DISTRIBUTIONAL CONSTRAINTS

2.1 Theoretical Work

Distributional objectives, particularly floor constraints, arise in many matching markets beyond the Japanese medical residency system. Rural doctor shortages have been documented in the United States [Fogarty et al. 2025], India [Alcoba 2009], Australia [Nambiar and Bavas 2010], and Korea [Chae 2025]. Similar distributional concerns also appear in other settings, such as teacher assignments to public schools in France with minimum staffing regulations [Terrier 2014] and the assignment of newly graduated cadets to U.S. military branches subject to minimum staffing requirements [Fragiadakis and Troyan 2017].

A large literature in matching theory studies mechanisms for addressing distributional imbalances, primarily in non-transferable utility (NTU) environments where agents cannot endogenously adjust transfers as part of the matching process. Early contributions typically model distributional concerns using capacity or upper-bound constraints. In school choice, [Abdulkadiroğlu and Sönmez 2003] introduce type-specific reserves that protect access for particular student types by reserving seats ex ante, without imposing ex post minimum assignment requirements. In a related vein, [Kamada and Kojima 2015; 2018] analyze the Japanese residency market by imposing regional upper bounds on matches with urban hospitals and proposing a flexible deferred acceptance algorithm that allocates limited urban capacity in response to residents' demand. Their approach encourages rural matches indirectly through urban ceilings, but it does not explicitly model lower bounds and therefore does not guarantee that minimum staffing constraints are satisfied ex post.

A distinct strand of the literature instead studies floor constraints, which impose minimum matching requirements ex post. With hard floor constraints, a matching must satisfy prescribed minimums even when demand is insufficient. [Ehlers et al. 2014] first formalize this setting and show an incompatibility between feasibility—simultaneously satisfying upper and lower bounds—and stability, defined by fairness and non-wastefulness. Motivated by this impossibility, they introduce soft floor constraints that may be violated when necessary. Building on this insight, [Fragiadakis et al. 2015] design mechanisms that offer alternative tradeoffs between fairness and non-wastefulness, while [Tomoeda 2018] provides sufficient conditions on hospital preferences under which feasible and stable matchings with floor constraints exist.

Other work focuses on mechanism design under hard floor constraints. [Goto et al. 2016] propose a strategy-proof mechanism that is non-wasteful and weakly Pareto efficient, though not necessarily stable. [Fragiadakis and Troyan 2017] allow for wasteful matchings and design a strategy-proof mechanism that endogenously adjusts ceiling constraints to achieve fairness while satisfying floor requirements, showing that doctors unanimously prefer this mechanism to one with fixed ceilings. [Akin 2021] instead weaken fairness to guarantee existence: their algorithm first runs deferred acceptance without floors and then applies a serial dictatorship to satisfy remaining minimums, yielding a strategy-proof outcome.

The literature also differs in the structure of constraints it considers. Many papers impose institution-level constraints, specifying lower and upper bounds independently for each school. By contrast, other work, including [Kamada and Kojima 2015; 2018] and [Ikegami et al. 2025], studies *regional constraints* that span multiple institutions and restrict the total number of matches across a group of schools.

Compared to the NTU model, relatively few papers study constraints and policy interventions in transferable utility (TU) matching models. [Kojima et al. 2020] and [Jalota et al. 2025] analyze the existence of equilibria under various constraints. Regarding the design of optimal taxes and subsidies, [Yokote 2020] studies a many-to-one TU matching framework [Kelso and Crawford 1982] with *interval constraints*, which impose lower and upper bounds on the number of matches at each hospital.[1] [Ikegami et al. 2025] instead studies the design of taxes and subsidies under regional constraints and extends the model to incorporate unobserved heterogeneity and to enable empirical analysis using aggregate-level matching data.

2.2 Empirical Work

A central empirical challenge is to assess how distributional constraints perform in practice, including whether they achieve their intended objectives and what welfare consequences they generate. When multiple policy instruments can be used to satisfy the same constraints, an additional question is how they compare in terms of social welfare. More generally, even in the absence of a theoretically dominant policy, it is natural to ask whether data can be used to design interventions that perform well in a given market. Addressing these issues requires empirical analysis beyond purely theoretical considerations. Yet, relative to the theoretical literature, empirical work on matching markets with constraints remains limited. We briefly review the related empirical work.

[Agarwal 2015] provides a set of policy analyses of the U.S. residency matching market. The paper applies an NTU matching framework to construct an empirical model of the medical match, which is estimated using actual matching outcomes and the rank-order lists submitted to the DA-like algorithm used in practice.

Building on the preference estimates from [Agarwal 2015], [Agarwal 2017] studies

[1][Yokote 2020] also develops a general discrete-optimization result based on hierarchical affine constraints. While the paper applies this result only to interval constraints, it may potentially be useful for studying taxes and subsidies under other types of constraints, including those considered in [Ikegami et al. 2025].

the effects of price- and quantity-based regulations aimed at addressing geographic imbalances in the supply of medical residents. The analysis suggests that neither higher wages (price regulation) nor capacity adjustments across regions (quantity regulation) substantially increase the number of residents placed in rural locations. However, both policies affect the composition of matches by altering the distribution of resident quality across regions and improving the quality of residents assigned to rural areas. The author argues that price regulation screens residents who are relatively more willing to work in rural locations, generating welfare gains for residents that may outweigh the associated fiscal costs. The paper also analyzes a restricted class of counterfactual policies and evaluates their effects, taking policy interventions as exogenously specified. The counterfactual policies considered there are not derived from an explicit theoretical benchmark.

Many of empirical market design studies examine the effect of altering the allocation algorithm. Examples include the design of waiting lists in kidney exchange and public housing allocation [Agarwal et al. 2021; Waldinger 2021], as well as the comparison between the deferred acceptance and Boston mechanisms in school choice [Abdulkadiroglu et al. 2011]. Closer to the literature on matching with constraints, [Combe et al. 2022] study the French teacher assignment market, with a particular focus on assignments to rural schools. Across these studies, policy interventions act through changes to the matching algorithm rather than through monetary instruments.

3. TU MODEL OF MATCHING WITH REGIONAL CONSTRAINTS

In [Ikegami et al. 2025], we develop a framework to design and analyze policies in a matching market with regional constraints. Our framework accommodates both *taxation policies*, which levy taxes or subsidies on specific matches, and *cap-based policies*, which impose quantity constraints on the number of available positions. In the proposed framework, there exists a taxation policy that induces the constrained-efficient matching.

3.1 Baseline Model

We consider a one-to-one, two-sided matching market with doctors $i \in I$ on one side and job slots $j \in J$ on the other, and there are finitely many regions. The job slots are owned by hospitals, and each hospital belongs to one region. A policymaker faces *regional constraints*. The regional constraints specify lower and upper bounds on the number of matches realized in each region.

Agents form a stable outcome à la [Shapley and Shubik 1971]. Without policy intervention, the realized matching may not meet the regional constraints. One class of interventions available to the policymaker is *taxation policies*, which alter the distribution of the joint surplus among agents to satisfy the regional constraints. When a doctor i and a slot j are matched, they generate an *(individual-level) gross joint surplus* $\Phi_{ij} \in \mathbb{R}$. The tax $w_z \in \mathbb{R}$ is imposed on each match (i, j) in region z, with negative taxes being interpreted as subsidies. With taxation policy $w = (w_z)_{z \in Z}$, each matched pair divides the *net joint surplus* $\Phi_{ij} - w_{z(j)}$ instead of the gross joint surplus. A matching market is characterized by a tuple (I, J, Φ, w), and the stable outcome under a taxation policy is defined as follows:

Definition 3.1 *Stable outcome*. Given (I, J, Φ, w), a profile $(d, (u, v))$ of feasible matching $d = (d_{ij})_{i,j}$ and equilibrium payoff profiles $u = (u_i)_i$ and $v = (v_j)_j$ forms a *stable outcome* if it satisfies:[2]

(1) Individual rationality: For all $i \in I$, $u_i \geq \Phi_{i,j_0}$, with equality if i is unmatched. For all $j \in J, v_j \geq \Phi_{i_0,j}$, with equality if j is unmatched.

(2) No blocking pairs: For all $i \in I$ and $j \in J$, $u_i + v_j \geq \Phi_{ij} - w_{z(j)}$, with equality if $d_{ij} = 1$.

The policymaker may also employ *cap-based policies*, which restrict the set of positions offered by hospitals to induce a desired allocation of applicants. Formally, a *cap-based policy* is specified by a subset $J' \subseteq J$, representing the positions that remain available after the policy is imposed. Given such a policy, let $\Phi' := (\Phi_{ij})_{i\in I,\, j\in J'}$ denote the restriction of the surplus matrix to doctors and the remaining slots. Under a cap-based policy J', agents form a stable outcome in the induced matching market $(I, J', \Phi', 0)$. Such quantity-based interventions are motivated by the objective of redirecting applicants toward understaffed regions by limiting capacity in high-demand areas. We assume that no taxation is applied when a cap-based policy is in place, that is, $w \equiv 0$.[3]

3.2 Results for the Baseline Model

A matching d is *constrained-efficient* if it maximizes total surplus $\sum_{i,j} d_{ij}\Phi_{ij}$ subject to the regional constraints. In the baseline model, we can show that the policymaker can compute an *optimal taxation policy* w^* that implements the constrained-efficient matching as an equilibrium outcome *if she knows the joint surplus generated by each pair.*

The optimal taxation policy, characterized by Lagrange multipliers associated with the regional constraints, can be efficiently computed via linear programming. Moreover, even if the policymaker can impose pair-specific taxes, the constrained-efficient allocation can be implemented with a uniform tax within each region. This structure substantially simplifies policy design and execution.

The constrained-efficient allocation serves as a welfare benchmark. In particular, the optimal taxation policy yields a weakly higher surplus than any cap-based policy satisfying the same regional constraints, providing a quantitative basis for evaluating the welfare losses of alternative policies.

4. IMPLEMENTING OPTIMAL TAXATION POLICY

To implement the optimal taxation policy derived in the previous section, the policymaker needs to know the preferences of market participants. We can show that the taxation policy is implementable using past match data under certain structural assumptions. We illustrate how to apply the proposed method using newly collected data on the Japan Residency Matching Program.

[2] $d_{ij} = 1$ if i and j are matched; $d_{ij} = 0$, otherwise. i_0 and j_0 denote the outside options.

[3] In our subsequent analysis incorporating unobserved heterogeneity, we assume that positions within a given region are removed uniformly at random under a cap-based policy.

4.1 Model with Unobserved Heterogeneity

Let X represent the finite set of observable characteristics, or *types*, of doctors. Each doctor $i \in I$ has a type $x(i) \in X$. Similarly, let Y represent the finite set of observable characteristics of job slots, with each slot $j \in J$ having a type $y(j) \in Y$. Agents with the same type are indistinguishable to the policymaker, but there can be *unobservable heterogeneity*: doctors of the same type x or job slots of the same type y may generate different joint surpluses when matched. We assume each job slot type $y \in Y$ belongs to a unique region. Let μ_{xy} denote the number of matches between type-x doctors and type-y job slots. We call $\mu = (\mu_{xy})_{x \in X, y \in Y}$ an *aggregate-level matching.*

Types $y \in Y$ and regions $z \in Z$ can be interpreted in various ways. For example, in the context of the Japan Residency Matching Program, a type y corresponds to a hospital, and a region z may correspond to a district (e.g., a prefecture). In other contexts, a type could represent a subcategory of occupation (e.g., registered nurse, physician assistant), and a region could represent a broader occupational category (e.g., healthcare).

4.2 Implementing Optimal Taxation using Data

Suppose that the policymaker observes historical aggregate match outcomes $(\mu_{xy})_{xy}$. Under the identification results of [Galichon and Salanié 2021], the *aggregate joint surplus* Φ_{xy}—the analogue of individual-level surplus Φ_{ij} in the presence of unobserved heterogeneity—can be recovered under certain structural assumptions. Embedding our model in this framework, we show that the optimal taxation policy can be computed from the identified joint surplus by solving a convex optimization problem, the aggregate counterpart of the linear program in the baseline model.

To apply the results to the data and compute the optimal taxes and subsidies, we estimate the joint surplus Φ_{xy} parametrically, decomposing it into doctors' and hospitals' systematic utilities. The specification uses observable attributes of medical schools x and hospitals y, including wages, quality measures, geographic distance, and past match frequencies. Wage information is especially important, as it allows us to express surplus—and, in counterfactual exercises, taxes and subsidies—in monetary units. Because [Galichon and Salanié 2021] does not address the use of transfer data in empirical analysis, we refer to [Ikegami et al. 2025] for a detailed description of the empirical model and the estimation procedure.

We use the estimates to conduct counterfactual simulations that compare the current regulatory regime with two alternative policies. The first scenario, *Artificial Caps* (AC), replicates the allocation under the cap-based policy currently implemented in the Japan Residency Matching Program. The second, *No Caps* (NC), provides an unconstrained welfare benchmark by removing all caps and reinstating all residency positions eliminated between 2017 and 2019. The third, *Optimal Subsidy* (OS), constructs the constrained-optimal benchmark: caps are removed, and subsidies are chosen to maximize total surplus subject to distributional constraints requiring designated rural regions to receive at least the number of residents assigned under AC.

Comparing welfare across these three scenarios isolates the sources of efficiency loss. By construction, NC attains the highest surplus and AC the lowest. The difference between NC and OS reflects the welfare cost of the distributional constraints themselves, while the difference between OS and AC measures the loss due to the choice of policy instrument. The simulations indicate that the former loss is small (approximately 6 million JPY per month), whereas the latter is large (exceeding 2,600 million JPY per month).[4] Thus, the cap-based policy has limited effectiveness in redirecting residents during the 2017–2019 period, and comparable distributional goals can be achieved at substantially lower welfare cost using a targeted subsidy scheme.

5. CONCLUSION

We conclude by highlighting several directions for future research. First, the framework of [Ikegami et al. 2025] abstracts from the possibility that the policymaker can create new positions. Extending the analysis to optimal capacity planning problems that allow for both expansions and contractions would be valuable. Second, practical policy design is often constrained by budget balance, which motivates the study of optimal policies under explicit budget constraints. Finally, the implementability result in [Ikegami et al. 2025]—the computability of optimal taxation policies using aggregate match data—relies on structural assumptions introduced in [Galichon and Salanié 2021]. Evaluating the restrictiveness of these assumptions, examining how they may be relaxed with richer data, and studying optimal policy design under ambiguity about agents' preferences are important avenues for future research.

REFERENCES

Abdulkadiroglu, A., Che, Y.-K., and Yasuda, Y. 2011. Resolving conflicting preferences in school choice: The "boston mechanism" reconsidered. *American Economic Review 101,* 1, 399–410.

Abdulkadiroğlu, A. and Sönmez, T. 2003. School choice: A mechanism design approach. *American Economic Review 93,* 3, 729–747.

Agarwal, N. 2015. An Empirical Model of the Medical Match. *American Economic Review 105,* 7, 1939–1978.

Agarwal, N. 2017. Policy analysis in matching markets. *American Economic Review 107,* 5, 246–50.

Agarwal, N., Ashlagi, I., Rees, M. A., Somaini, P., and Waldinger, D. 2021. Equilibrium allocations under alternative waitlist designs: Evidence from deceased donor kidneys. *Econometrica 89,* 1, 37–76.

Akin, S. 2021. Matching with floor constraints. *Theoretical Economics 16,* 3, 911–942.

Alcoba, N. 2009. Hitting the road to alleviate india's rural doctor shortage. *Canadian Medical Association Journal 180,* 10, E34–E36.

Chae, H.-S. 2025. Rural hospitals are short of doctors and not even high pay will attract them. Accessed: 2025-12-21.

Combe, J., Tercieux, O., and Terrier, C. 2022. The design of teacher assignment: Theory and evidence. *Review of Economic Studies 89,* 6, 3154–3222.

Ehlers, L., Hafalir, I. E., Yenmez, M. B., and Yildirim, M. A. 2014. School choice with controlled choice constraints: Hard bounds versus soft bounds. *Journal of Economic Theory 153*, 648–683.

Fogarty, C. T., Byun, H., and Huffstetler, A. N. 2025. Family physician workforce trends: The toll on rural communities. *The Annals of Family Medicine 23,* 6, 535–538.

[4]The estimated total surplus under NC is approximately 65,000 million JPY per month.

Fragiadakis, D., Iwasaki, A., Troyan, P., Ueda, S., and Yokoo, M. 2015. Strategyproof matching with minimum quotas. *ACM Transactions on Economics and Computation 4*.

Fragiadakis, D. and Troyan, P. 2017. Improving matching under hard distributional constraints. *Theoretical Economics 12,* 2, 863–908.

Galichon, A. and Salanié, B. 2021. Cupid's Invisible Hand: Social Surplus and Identification in Matching Models. *The Review of Economic Studies 89,* 5, 2600–2629.

Goto, M., Iwasaki, A., Kawasaki, Y., Kurata, R., Yasuda, Y., and Yokoo, M. 2016. Strategyproof matching with regional minimum and maximum quotas. *Artificial Intelligence 235*, 40–57.

Ikegami, K., Iwasaki, A., Matsushita, A., and Okumura, K. 2025. Evaluating the efficiency of regulation in matching markets with distributional disparities. In *Proceedings of the 26th ACM Conference on Economics and Computation (EC'25)*.

Jalota, D., Ostrovsky, M., and Pavone, M. 2025. Matching with transfers under distributional constraints. *Games and Economic Behavior 152*, 313–332.

Kamada, Y. and Kojima, F. 2015. Efficient matching under distributional constraints: Theory and applications. *American Economic Review 105,* 1, 67–99.

Kamada, Y. and Kojima, F. 2018. Stability and strategy-proofness for matching with constraints: A necessary and sufficient condition. *Theoretical Economics 13,* 2, 761–793.

Kelso, A. and Crawford, V. 1982. Job matching, coalition formation, and gross substitutes. *Econometrica 50,* 6, 1483–1504.

Kojima, F., Sun, N., and Yu, N. N. 2020. Job matching under constraints. *American Economic Review 110,* 9, 2935–47.

Nambiar, M. and Bavas, J. 2010. Rudd plan won't fix rural doctor shortage: Rda. Accessed: 2025-12-21.

Shapley, L. S. and Shubik, M. 1971. The assignment game I: The core. *International Journal of Game Theory 1,* 1, 111–130.

Terrier, C. 2014. Matching practices for secondary public school teachers – france. Matching in Practice Country Profile No. 20. Accessed: 2025-12-21.

Tomoeda, K. 2018. Finding a stable matching under type-specific minimum quotas. *Journal of Economic Theory 176*, 81–117.

Waldinger, D. 2021. Targeting in-kind transfers through market design: A revealed preference analysis of public housing allocation. *American Economic Review 111,* 8, 2660–96.

Yokote, K. 2020. On optimal taxes and subsidies: a discrete saddle-point theorem with application to job matching under constraints. *Journal of Mechanism and Institution Design 5,* 1, 37–77.

Balancing learning and targeting in predictive allocation

BRYAN WILDER
Carnegie Mellon University
and
PIM WELLE
Allegheny County Department of Human Services

1. INTRODUCTION

This letter provides an overview of our recent work on "Learning Treatment Effects While Treating Those in Need" (published at the 2025 ACM Conference on Economics and Computation) as well as a more general perspective on design goals for algorithmic systems that are used to allocate limited resources in policy settings. Our motivation is the kind of algorithms that are used widely at present to prioritize candidates for various kinds of social interventions: public housing assistance, drop-out prevention programs in education, unconditional cash transfers in development, or a variety of other social services. By far the most common way of constructing such systems is the lens of *predictive allocation*: the algorithm designer identifies an outcome that the program seeks to alter (long-term homelessness, dropping out of school, etc) and constructs a predictive model for that outcome [Vaithianathan and Kithulgoda 2020; Aiken et al. 2022; Pan et al. 2017; Toros and Flaming 2017]. Candidates are ranked by predictions of risk so that, e.g., limited spots in a housing program might be offered to those at greatest predicted risk of long-term homelessness.

We argue that algorithm designers should take a more expansive view of the design goals that such systems aim to optimize. Our work focuses specifically on the tradeoff between predictive targeting and causal learning: if resources are allocated strictly according to predicted risk, it becomes difficult to tease out to what extent the program improves outcomes for its beneficiaries. Learning such causal effects requires experimentation, where otherwise similar candidates sometimes receive different allocation decisions. However, randomized experimentation would depart from the goal of strictly aligning allocations with predictions of risk. The fear of denying services to higher-need applicants is one of the main reasons that randomized controlled trials are relatively uncommon in many policy settings. However, high-quality evidence about program effectiveness is absolutely critical in order to improve services over time.

Our work asks how stark the tradeoffs really are. We construct a multiobjecitve

Authors' addresses: `bwilder@cmu.edu`, `Paul.Welle@alleghenycounty.us`

framework for optimizing a policy for allocating a limited resource. The goal is to optimally trade off between the goals of (1) allocating resources to individuals that have been identified as having higher levels of risk or need; and, (2) introducing enough random variation to accurately estimate the average treatment effect of the intervention. Our method allows policymakers to explore the Pareto frontier between these goals. We apply this framework to data on human service programs in Allegheny County, Pennsylvania and find that the tradeoffs between these goals are empirically quite tractable: with careful design, policymakers can get most of one without giving up too much of the other.

More broadly, we suggest that algorithm designers should incorporate other goals (like causal evaluation) explicitly into the design of allocative systems. Framing allocation as a prediction problem naturally suggests that the best policy is one that aligns with the most accurate predictive model. Prediction may often be a reasonable way to construct a proxy for an individual's need for a service. However, it is still only a proxy (and one which we have shown in other work may not always align well with other goals, like benefit from an intervention [Sharma and Wilder 2025]). It may be very worthwhile to give up a bit of predictive performance in order to ensure that algorithmic systems reflect the much broader array of challenges for decision making in social systems.

2. A MULTIOBJECTIVE FRAMEWORK FOR LEARNING AND TARGETING

The specific setting studied in our paper concerns the allocation of a single limited resource. The policymaker chooses a function $p : \mathcal{X} \to [0, 1]$ which maps an individuals' features $X \in \mathcal{X}$ to the probability with which they are allocated treatment. The policymaker specifies upfront a function $u(X)$ which specifies their preferences for which individuals should be allocated the resource; we refer to u as their *targeting utility function*. In our running example, u is the output of a predictive model for some adverse outcome: the policymaker prefers that the resource be allocated to individuals with higher levels of predicted risk. However, our algorithmic framework is agnostic as to how u is constructed. Individuals have potential outcomes $(Y(0), Y(1))$, where outcome $Y(1)$ is realized if they receive the intervention and $Y(0)$ is realized if they do not. We assume that $(X, Y(0), Y(1))$ are drawn iid from some joint distribution. The average treatment effect of the intervention is given by $\tau = \mathbb{E}[Y(1) - Y(0)]$. Our goal is to design an allocation p which will enable accurate estimates of τ while still scoring highly according to u. Allocations are subject to a budget constraint that at most a b fraction of the population can receive the intervention, i.e., that $\mathbb{E}[p(X)] \leq b$.

To formalize the goal of accurately estimating treatment effects, we turn to the semiparametric efficiency bound for the average treatment effect [Hahn 1998]. Suppose that we construct an estimate $\hat{\tau}$ for the ATE using any of the many methods developed in the causal inference literature. We would like to minimize the expected error in the estimate, $\mathbb{E}[(\hat{\tau} - \tau)^2]$. Since standard estimators are unbiased, the entirety of the mean-squared error in an estimate is given by the variance of the estimator. This is precisely the quantity that the efficiency bound gives: the smallest possible variance for any estimator.

We show that, under many circumstances, the efficiency bound is itself upper bounded by a function of the form

$$\mathbb{E}\left[\frac{a_1(X)}{p(X)} + \frac{a_0(X)}{1-p(X)}\right]$$

for a pair of functions $a_1(X)$ and $a_0(X)$. Intuitively, the performance of estimators for the treatment effect degrades rapidly as $p(X)$ gets closer to 0 or 1, i.e., as allocations get closer to deterministic. The efficiency bound is exactly given by taking a_0 and a_1 to be the conditional outcome variances, $a_0(X) = \mathrm{Var}(Y(0)|X)$ and $a_1(X) = \mathrm{Var}(Y(1)|X)$ [Hahn 1998]. Since the variances are typically unknown before collecting experimental data, we show how structural properties, side information, or assumptions about the distribution can be used to construct an upper bound on the worst-case variance. The particular values of a_0 and a_1 for the optimization problem then arise from such choices (e.g., stipulating that outcomes are binary, or using historical data to learn about $Y(0)$).

One possible allocation policy is to spend the entire budget on the b-fraction of the population with the highest value of $u(X)$. This corresponds to the predictive allocation strategy often employed in practice. However, since $p(X)$ is either 0 or 1 for every X, estimating treatment effects is impossible. At the other end of the spectrum is a completely uniform allocation policy, $p(X) = b$ for all X, which corresponds to a randomized controlled trial. This maximizes the power to estimate treatment effects (formally, minimizes the efficiency bound when a_1 and a_0 are constant) but entirely sacrifices the targeting goal.

While these extremes are the strategies typically employed in practice (predictive allocation most frequently, randomized trials rarely), our paper proposes an entire spectrum of allocation policies in between. Formally, these are derived as solutions to a family of multiobjective optimization problems that balance the targeting and causal learning goals. In particular, we aim to find allocation probabilities p which solve an optimization problem of the form

$$\begin{aligned}
&\min_p \mathbb{E}\left[\frac{a_1(X)}{p(X)} + \frac{a_0(X)}{1-p(X)}\right] \\
&\mathbb{E}[p(X)u(X)] \geq c \\
&\mathbb{E}[p(X)] \leq b \\
&p(X) \in [\gamma, 1-\gamma] \quad \forall X \in \mathcal{X}.
\end{aligned} \tag{1}$$

where the parameter c is chosen by the policymaker to control the extent to which allocation should align with their targeting utility function u. As discussed in our work, this core formulation is also easily extensible to include constraints enforcing additional goals. The last constraint exerts limited influence on the actual solution since optimal policies will never have $p(X)$ at 0 or 1 by virtue of the objective function. However, the parameter γ measuring closeness to deterministic solutions appears in the exact statement of our sample complexity bounds, with more samples required when the designer wishes to optimize over closer-to-deterministic policies. A complete discussion of this topic can be found in our full paper.

This is a strictly convex optimization problem in p, subject to linear constraints. The challenge is that this optimization problem is over the space of policies, i.e., functions $p : \mathcal{X} \to [0, 1]$. In other policy learning settings, a typical approach is to optimize over a parameterized class of policies (e.g., linear models, decision trees, etc.), with sample complexity bounds for optimal policy learning then depending on a measure of the complexity of this class like the VC dimension [Athey and Wager 2021; Chernozhukov et al. 2019; Swaminathan and Joachims 2015].

In our work, we show that solutions to the optimization problem above have a special structure. Effectively, the function p can be parameterized very concisely, just via a handful of parameters that correspond to the dual variables. This viewpoint enables efficient optimization and learning. To formalize this idea, we take the dual of the population optimization problem. Let the constraints be collectively denoted as functions g_j, $j = 1...J$. The dual is

$$\max_{\lambda \geq 0} \min_{p \in [\gamma, 1-\gamma]^{\mathcal{X}}} \mathbb{E}\left[\frac{a_1(X)}{p(X)} + \frac{a_0(X)}{1-p(X)} + \sum_{j=1}^{J} \lambda_j (g(p(X), X) - c_j)\right]$$

where λ_j is the dual variable associated constraint j. The dual has the attractive property that the inner objective is separable across X, allowing us to push the min inside the expectation:

$$\max_{\lambda \geq 0} \mathbb{E}\left[\min_{p(X) \in [\gamma, 1-\gamma]} \frac{a_1(X)}{p(X)} + \frac{a_0(X)}{1-p(X)} + \sum_{j=1}^{J} \lambda_j (g(p(X), X) - c_j)\right].$$

The strategy proposed in our work is to estimate the dual parameters using a sample of X's drawn from the distribution of interest by solving the *sample* problem

$$\max_{\lambda \geq 0} \frac{1}{n} \sum_{i=1}^{n} \left[\min_{p(X_i) \in [\gamma, 1-\gamma]} \frac{a_1(X_i)}{p(X_i)} + \frac{a_0(X_i)}{1-p(X_i)} + \sum_{j=1}^{J} \lambda_j (g(p(X_i), X_i) - c_j)\right]. \quad (2)$$

This is a strictly convex optimization problem in n variables that can easily be solved using standard methods. Then, to compute the optimal allocation probability $p(X)$ for any given individual, we solving the inner minimization problem at the estimated dual parameters, i.e., we compute

$$\widehat{p}(X) = \text{argmin}_{p(X) \in [\gamma, 1-\gamma]} \frac{a_1(X)}{p(X)} + \frac{a_0(X)}{1-p(X)} + \sum_{j=1}^{J} \widehat{\lambda}_j (g(p(X), X) - c_j) \quad (3)$$

separately for each individual and randomize them to be treated with probability $\widehat{p}(X)$.

Theoretically, we establish that this strategy enjoys strong sample complexity guarantees. Intuitively, the class of policies that we optimize over has only J parameters, one per constraint (of which there were two in the core formulation discussed above). We show that the number of samples from the covariate distribution required in order to produce an approximately optimal policy scales at parametric, root-n rates, with the complexity of the policy class represented by a function of J:

Proposition 1 *[informal]: In order to obtain policy $\hat{p}$ that is within ϵ of both optimality and feasibility for Problem 1, it suffices to have $n = O\left(\frac{J^3}{\epsilon^2}\right)$ samples.*

Empirically, we simulate the policies output by our method on historical data

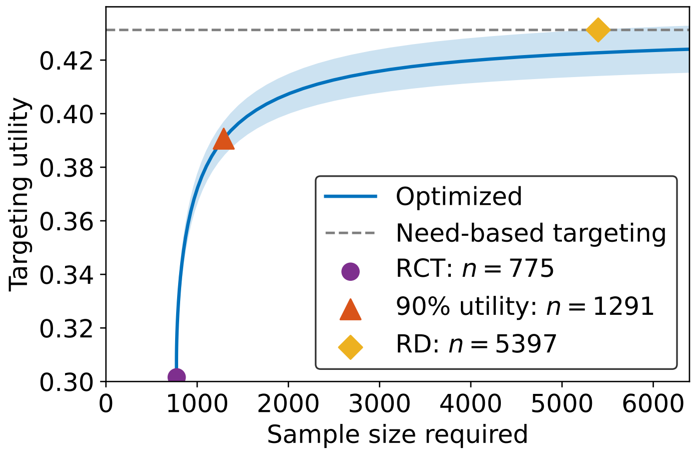

Fig. 1. Tradeoff between predictive targeting and causal learning on historical data from the Allegheny County Human Services Department. The horizontal axis shows the number of samples that would be required to power an estimate of the average treatment effect. The vertical axis shows the fraction of individuals who would reenter jail in the next year that are targeted for the program. "Optimized" refers to the Pareto frontier of designs output by our method while "need-based targeting" refers to allocating entirely to individuals with the highested predicted risk. "'RCT" gives the performance of a randomized controlled trial that uniformly allocates the treatment. "90% utility" labels the point on the Pareto curve output by our method that achieves 90% of the targeting utility of need-based targeting. Finally, "RD" shows the number of samples required to power a regression discontinuity estimate of the local treatment effect.

from the Allegheny County Human Services Department. In this letter, we reproduce one example result, which uses administrative data on the outcomes of formerly incarcerated people reentering the community. We mimic standard practice by fitting a machine learning model to predict the probability that a given individual will reenter jail within the next month. Our evaluation models a hypothetical program which offers an intervention to avert jail reentry. The policymaker would like to offer this program preferentially to individuals with a higher probability of reentering jail and we characterize the tradeoff between this targeting goal and learning causal effects. To provide a concrete interpretation for power to learn causal effects, we rescale asymptotic variance into the number of samples that would be required to estimate of the average treatment effect with 5% type-1 error and 80% power.

Our main empirical result is that the tradeoff between these goals is substantially mitigated by optimal policy design. Figure 1 shows the tradeoff in the form the Pareto curve output by our method, obtained by varying the value of the targeting utility constraint c. At one end of the spectrum we have a randomized controlled trial that allocates the intervention uniformly, without consideration for risk. As the value of c increases, the allocation converges towards thresholding on the predicted risk. However, the shape of the curve in between is highly concave, indicating that the policymaker can obtain most of one goal without giving up too much of the other. For example, the allocation policies output by our method achieve 90% of the maximum possible performance at targeting high-risk individuals while requiring less than twice the number of samples in order to estimate treatment effects as the ideal randomized trial. Similar findings surface across a number of other programmatic settings and simulations. At the far end of the spectrum, we compare to the number of samples that would be required to power an estimate of a local average treatment effect via a regression discontinuity design if the policymaker instead used pure predictive allocation. The number of samples is much larger, several times that required by our method at the elbow of the curve. This indicates that the price (in terms of causal learning) required to eliminate the last bit of randomization is very high, imposing a sample size which is infeasible for many programmatic settings.

3. DISCUSSION AND CONCLUSION

We have shown how the targeting interventions to people at higher levels of present need can be reconciled with the goal of introducing randomization to learn whether these interventions produce real benefits. There is still research to be done in order to bring these ideas to the operational reality of public services. For example, in this work, we focused on designing policies that adhere to a set of constraints in-expectation on a fixed distribution. In reality, allocations are subject to a variety of other complications in implementation, like hard budget constraints, dynamic arrivals of applicants and resources, drift in covariate distributions, and so on. In subsequent work, we showed how techniques from combinatorial optimization allow us to enforce hard budget constraints (which requires dependent assignments) while still providing statistical valid estimates of treatment effects [Yamin et al. 2025]. In fact, we find that doing so actually provides substantial variance reduction benefits and may be useful in other experimental design settings as well. However, a variety of more complex settings remain open.

Broadly though, the takeaway remains: by integrating causal learning as a goal alongside prediction in designing allocation policies, it is often possible to get much of the best of both worlds in practice. Our hope is that these kinds of techniques enable much more widespread use of rigorous program evaluation strategies by lowering the cost (in terms of foregone targeting) to running randomized trials. While prediction is one way of encoding goals for allocation, we believe that many settings will benefit at least as much from better, faster, and more widely available evidence about "what works". Is is tempting for algorithm designers to focus on the details of allocation, trying to optimize the individual-level match between people and potential interventions. But this focus should not crowd out our ability to learn and improve interventions over time, ultimately improving the quality of services available to everyone.

4. ACKNOWLEDGMENTS

This material is based upon work supported by the AI Research Institutes Program funded by the National Science Foundation under AI Institute for Societal Decision Making (AI-SDM), Award No. 2229881, and the AI2050 program at Schmidt Sciences (Grant #G-22-64474).

REFERENCES

Aiken, E., Bellue, S., Karlan, D., Udry, C., and Blumenstock, J. E. 2022. Machine learning and phone data can improve targeting of humanitarian aid. *Nature 603,* 7903, 864–870.

Athey, S. and Wager, S. 2021. Policy learning with observational data. *Econometrica 89,* 1, 133–161.

Chernozhukov, V., Demirer, M., Lewis, G., and Syrgkanis, V. 2019. Semi-parametric efficient policy learning with continuous actions. *Advances in Neural Information Processing Systems 32*.

Hahn, J. 1998. On the role of the propensity score in efficient semiparametric estimation of average treatment effects. *Econometrica*, 315–331.

Pan, I., Nolan, L. B., Brown, R. R., Khan, R., van der Boor, P., Harris, D. G., and Ghani, R. 2017. Machine learning for social services: a study of prenatal case management in illinois. *American journal of public health 107,* 6, 938–944.

Sharma, V. and Wilder, B. 2025. Comparing targeting strategies for maximizing social welfare with limited resources. In *The Thirteenth International Conference on Learning Representations*.

Swaminathan, A. and Joachims, T. 2015. Batch learning from logged bandit feedback through counterfactual risk minimization. *The Journal of Machine Learning Research 16,* 1, 1731–1755.

Toros, H. and Flaming, D. 2017. Prioritizing which homeless people get housing using predictive algorithms. *Available at SSRN 2960410*.

Vaithianathan, R. and Kithulgoda, C. I. 2020. Using predictive risk modeling to prioritize services for people experiencing homelessness in allegheny county: Methodology update.

Yamin, K., Kennedy, E., and Wilder, B. 2025. Dependent randomized rounding for budget constrained experimental design. In *The 41st Conference on Uncertainty in Artificial Intelligence*.

Menus: A Framework for Learning Against Strategic Opponents

ESHWAR RAM ARUNACHALESWARAN
SESCO Enterprises
and
NATALIE COLLINA
University of Pennsylvania
and
YISHAY MANSOUR
Google Research
and
MEHRYAR MOHRI
Google Research
and
BALASUBRAMANIAN SIVAN
Google Research
and
JON SCHNEIDER
Google Research

Algorithms increasingly mediate repeated strategic interactions in marketplaces, from automated pricing to auction bidding. When one party commits to a learning algorithm, the other party can respond strategically over time by steering the algorithm's internal state toward a favorable long-run outcome. This note surveys a line of work that studies this "learning-as-commitment" perspective via a geometric object we call a *menu*: the convex set of long-run outcomes an opponent can induce against a fixed learning rule. Menus provide a common language for (i) comparing learning algorithms against strategic opponents, (ii) optimizing over learning rules under uncertainty about opponent objectives, and (iii) characterizing when an opponent can manipulate learning dynamics beyond what they could achieve with a static strategy. Using this machinery, we converge upon no-swap-regret algorithms as an "optimal" commitment strategy for robust learning against a strategic opponent. We also identify principled generalizations of no-swap-regret beyond normal-form games that preserve the same strategic guarantees while remaining computationally tractable.

Categories and Subject Descriptors: [**Theory of Computation**]: Algorithmic Game Theory and Mechanism Design

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Online Learning, Swap Regret, Strategizing with Algorithms

Authors' addresses: `eshwarram.arunachaleswaran@gmail.com`, `ncollina@seas.upenn.edu`

1. INTRODUCTION

Algorithms are increasingly used to make repeated decisions in strategic environments. A canonical example is automated pricing: two sellers may repeatedly compete to sell an identical item, with the lower posted price winning the sale (a Bertrand-style competition). In principle, sellers could adjust prices manually, checking in periodically to respond to market conditions. In practice, many rely on automated pricing tools that update based on observed competitor prices. Related examples include automated bidders in auction markets and algorithmic agents in platform-mediated interactions.

When a firm deploys a learning algorithm, it is not merely choosing a sequence of actions. It is committing to a *policy* that maps the evolving interaction history to future behavior. This changes the strategic problem faced by the opposing party: a sophisticated opponent can choose time-varying actions designed to exploit the learning rule itself. For instance, if a pricing algorithm undercuts competitors but "resets" when prices dip too low (a common heuristic), an informed competitor can deliberately trigger the reset and then exploit the resulting price increase. In such interactions, the opponent can steer the learner's dynamics to ensure their own long-term success. This raises a basic design question:

Design problem. If you are deploying a learning algorithm in a repeated, general-sum game, how should you design it to be robust to an opponent who responds strategically over time?

Setup We consider a repeated two-player normal-form game, where in each round t, the first player, called the learner, select a distribution $x_t \in \Delta^m$ over m pure actions, and the second player, called the optimizer, selects a distribution $y_t \in \Delta^n$ over n pure actions. The players receive bilinear utilities $u_L(x_t, y_t)$ and $u_O(x_t, y_t)$–in other words, their utilities are only a function of the probability that each move pair occurs, a linear function applied to $x_t \otimes y_t$. The learner's strategy is an algorithm $\mathcal{A}$ that maps the history of play so far $H_t = (x_1, y_1), (x_2, y_2) \cdots (x_{t-1}, y_{t-1})$ to an action x_t for the next round. The optimizer selects a sequence of actions[1] $y_{1:T}$ strategically in order to maximize their own long-term utility. We let $\mathcal{A}(y_{1:T})$ denote the move pair distribution induced by the optimizer playing $y_{1:T}$ against $\mathcal{A}$.

Definition 1.1 *Optimizer best-response.* We consider an optimizer who, when faced with the algorithm $\mathcal{A}$, plays the sequence

$$\mathsf{BR}(\mathcal{A}, u_O) = \arg\max_{y_{1:T} \in \mathcal{Y}^T} u_O(\mathcal{A}(y_{1:T}))$$

Here we omit details on tie-breaking[2].

What is a good algorithm for a learner to deploy in such an environment? A natural starting point is the standard learning-theoretic solution concept of regret

[1] In fact, the optimizer may choose to play their own adaptive strategy instead of a fixed sequence; however, in this note we primarily discuss learner algorithms which do not have correlated randomness, against which there is no need for the optimizer to play adaptively. Thus, for simplicity of notation, we will write the optimizer's strategy as a sequence rather than an algorithm.

[2] A careful formulation of tie-breaking is needed to ensure continuity of the *learner's* outcomes, see [Arunachaleswaran et al. 2024] for a rigorous treatment

minimization. The (external) regret of the learner's realized play is

$$\mathrm{Regret}_T \;:=\; \max_{x\in\mathcal{X}} \sum_{t=1}^{T} u_L(x,y_t) \;-\; \sum_{t=1}^{T} u_L(x_t,y_t),$$

and we say the algorithm is *no-regret* if $\mathrm{Regret}_T = o(T)$ (equivalently, average regret vanishes).

A stronger notion is *swap regret*, which compares to the best *action remapping* $\phi : \mathcal{X} \to \mathcal{X}$ applied to the learner's realized actions:

$$\mathrm{SwapRegret}_T \;:=\; \max_{\phi:\mathcal{X}\to\mathcal{X}} \sum_{t=1}^{T} u_L(\phi(x_t),y_t) \;-\; \sum_{t=1}^{T} u_L(x_t,y_t),$$

and we say the algorithm has *no swap regret* if $\mathrm{SwapRegret}_T = o(T)$.

These guarantees may seem strong, but they are achievable: there are online algorithms with vanishing regret against *any* (even adaptive) sequence of opponent actions. In a repeated *zero-sum* game, playing a no-regret algorithm is essentially the right strategy: a strategic opponent can always hold the learner to its minimax value.

However, many market interactions are *general-sum*. Here the opponent is not trying to minimize the learner's utility, they are trying to maximize their own. In this setting, regret guarantees can become a poor proxy for strategic robustness. A vivid illustration appears in [Braverman et al. 2018], which shows that in a repeated auction with a mean-based, no-regret buyer, a strategic seller can extract essentially full surplus. The message is that the strategic guarantees of many no-regret algorithms are too weak when an opponent optimizes *against the algorithm.*

So what should replace regret as the organizing principle? Across a series of papers [Arunachaleswaran et al. 2024; 2025; Arunachaleswaran et al. 2025], we study this question under varying information regimes. A common technical tool is a geometric abstraction that makes the commitment aspect explicit: *menus of algorithms.*

2. MENUS OF ALGORITHMS.

The central theme of this letter is a view of algorithms in terms of the geometric set of all outcomes they can induce, an object which we call an algorithm's "menu".

Definition 2.1 *Finite-time menu.* For a horizon dependent algorithm $\mathcal{A}^T$, we define its menu $\mathcal{M}(\mathcal{A}^T)$ as the convex hull of all vectors of the form $\frac{1}{T}\sum_t x_t \otimes y_t$ for all possible optimizer sequences $y_1, y_2 \ldots y_T$ and the corresponding induced sequences $x_1, x_2..x_T$ of the algorithm, i.e. $x_t = \mathcal{A}^T(H_t)$.

We refer to such vectors as Correlated Strategy Profiles or CSPs. CSPS have a clean interpretation, at least in normal-form games: they are a distribution over the learner's and optimizer's joint actions.

Definition 2.2 *Correlated Strategy Profile (CSP).* A CSP is a vector of the form $\frac{1}{T}\sum_t x_t \otimes y_t$.

If a learning algorithm $\mathcal{A}$, which we define as a composition of horizon dependent algorithms $\mathcal{A}^1, \mathcal{A}^2 \ldots$, satisfies the property that its menus $\mathcal{M}(\mathcal{A}^1), \mathcal{M}(\mathcal{A}^2) \ldots$

converge under the Hausdorff metric to some set $\mathcal{M}$, we define its menu $\mathcal{M}(\mathcal{A})$ to be the limit set[3] $\mathcal{M}$ (see [Arunachaleswaran et al. 2024] for a rigorous treatment).

Definition 2.3 *Menu.* For a consistent algorithm $\mathcal{A} = \mathcal{A}^1, \mathcal{A}^2 \ldots$, we define its menu $\mathcal{M}(\mathcal{A})$ as $\lim_{T\to\infty} \mathcal{M}(\mathcal{A}^T)$

A set $\mathcal{M} \subseteq \Delta^{mn}$ is said to be a menu if and only if there exists a consistent algorithm $\mathcal{A}$ such that $\mathcal{M}$ is the menu of $\mathcal{A}$. To elucidate this idea, we now provide a few examples of algorithms and their corresponding menus.

2.1 Example Menus

Warm-up example 1. Consider the learning algorithm $\mathcal{A}^*$ that ignores the interaction history and plays x^* on every round:

$$\mathcal{A}^*_t(h_{t-1}) := x^* \quad \text{for all } t \geq 1 \text{ and all histories } h_{t-1}.$$

What outcomes can be induced against $\mathcal{A}^*$? As the learning algorithm plays x^* each day, and as the optimizer can play any sequence of their actions, the menu of $\mathcal{A}^*$ precisely **the convex set of CSPs of the form $x^* \otimes y$ for $y \in \mathcal{Y}$**.

Warm-up example 2. Next, let us consider an adaptive example: the learning algorithm $\hat{\mathcal{A}}$, which plays x^* as long as the opponent has never played $\hat{y}$, and plays $\hat{x}$ otherwise:

$$\mathcal{A}_t(\hat{h_{t-1}}) \;:=\; \begin{cases} \hat{x} & \text{if } \hat{y} \notin \{h_{t-1}\}, \\ x^* & \text{otherwise,} \end{cases}$$

As the optimizer could play $\hat{y}$ each day, the menu of $\hat{\mathcal{A}}$ includes the CSP $\hat{x} \otimes \hat{y}$. Furthermore, at any point the optimizer may play some action $y \neq \hat{y}$; at this point, $\hat{\mathcal{A}}$ begins behaving exactly as $\mathcal{A}^*$. Thus, its menu is **the convex hull of $\hat{x} \otimes \hat{y}$ and CSPs of the form $x^* \otimes y$ for $y \in \mathcal{Y}$.** [4]

2.2 Characterization of Menus

Given these examples, a natural question to ask is what subsets of Δ^{mn} are true menus–i.e., induced as the menu of some algorithm.

THEOREM 2.4 [ARUNACHALESWARAN ET AL. 2024]. *A closed, convex subset $\mathcal{M} \subseteq \Delta_{mn}$ is an asymptotic menu iff for every $y \in \Delta_n$, there exists a $x \in \Delta_m$ such that $x \otimes y \in \mathcal{M}$.*

The necessity of the condition is easy to argue, since every menu must contain a point corresponding to when the optimizer picks the same action in each round. The sufficiency of the condition is proved via Blackwell Approachability.

[3]If this property is satisfied, we call $\mathcal{A}$ a *consistent* algorithm; in [Arunachaleswaran et al. 2024] we show that many focal learning algorithms, such as no-swap regret algorithms, are consistent.
[4]It is possible for an optimizer to induce a move pair distribution of the form $\hat{x} \otimes y \neq \hat{y}$ in the single day after they defect from $\hat{y}$ and before the learner can react. However, the contribution of this day to the CSP will disappear in the limit, and thus this CSP is not contained in the menu.

2.3 Menus as a tool for learning in strategic settings

The concept of a menu allows us to recast the game between the learner and optimizer in geometric terms[5]:

—The learner picks an algorithm $\mathcal{A}$ and offers the menu $\mathcal{M}(\mathcal{A})$ to the optimizer.

—The optimizer picks an outcome CSP $\varphi \in \mathcal{M}(\mathcal{A})$ that maximizes its utility with ties broken in favor of the learner. The utilties of both the learner and optimizer are linear functions of CSPs in $\mathcal{M}(\mathcal{A})$. Specifically, the learner gets utility:

$$V_L(\mathcal{M}(\mathcal{A}), u_O) = \max\{u_L(\varphi) | \varphi \in \arg\max_{\phi \in \mathcal{M}(\mathcal{A})} u_O(\phi)\}$$

Importantly, given the menu $\mathcal{M}$ of an algorithm, it is easy to see what outcome a strategic optimizer with utility u_O would induce: it will be the CSP of $\mathcal{M}$ in the extreme direction of u_O.

It is also possible to recast the no-regret and no-swap-regret properties purely in terms of menus - to aid with this we define two special menus. We say that the CSP φ is *no-regret* if it satisfies the no-regret constraint

$$\sum_{i \in [m]} \sum_{j \in [n]} \varphi_{ij} u_L(i,j) \geq \max_{j^* \in [n]} \sum_{i \in [m]} \sum_{j \in [n]} \varphi_{ij} u_L(i,j^*). \tag{1}$$

Similarly, say that the CSP φ is *no-swap-regret* if, for each $j \in [n]$, it satisfies

$$\sum_{i \in [m]} \varphi_{ij} u_L(i,j) \geq \max_{j^* \in [n]} \sum_{i \in [m]} \varphi_{ij} u_L(i,j^*). \tag{2}$$

For a fixed u_L, we will define the *no-regret menu* $\mathcal{M}_{\mathrm{NR}}$ to be the convex hull of all no-regret CSPs, and the *no-swap-regret menu* $\mathcal{M}_{\mathrm{NSR}}$ to be the convex hull of all no-swap-regret CSPs. Menus provide a clean geometric statement of the no-regret and no-swap-regret properties. These algorithmic properties (restrictions on transcripts) translate to geometric containment of menus within corresponding polytopes (restrictions on outcomes).

Theorem 2.5 [Arunachaleswaran et al. 2024]. *A learning algorithm $\mathcal{A}$ is no-regret iff $\mathcal{M}(\mathcal{A}) \subseteq \mathcal{M}_{NR}$. Likewise, $\mathcal{A}$ is no-swap-regret iff $\mathcal{M}(\mathcal{A}) \subseteq \mathcal{M}_{NSR}$.*

We provide a visualization of these containment properties in Figure 1.

2.4 Results about Menus

Among no-regret algorithms, it is well known that different algorithms can induce distinctly different outcomes: for instance, multiplicative weights does not satisfy the no-swap-regret property, while other constructions (e.g., [Blum and Mansour 2007]) do satisfy this strictly stronger guarantee. However, it was not a priori

[5]While this game is a proxy for the actual finite horizon game between the learner and the optimizer, as long as the optimizer picks the best utility point in $\mathcal{M}(\mathcal{A}^T)$ and tie-breaks in favor of the learner, it is not hard to show that the sequence of utilities obtained by the learner converges in the limit to the utility $V_L(\mathcal{M}(\mathcal{A}), u_O)$ achieved by the learner in this platonic menu version of the game ([Arunachaleswaran et al. 2024]).

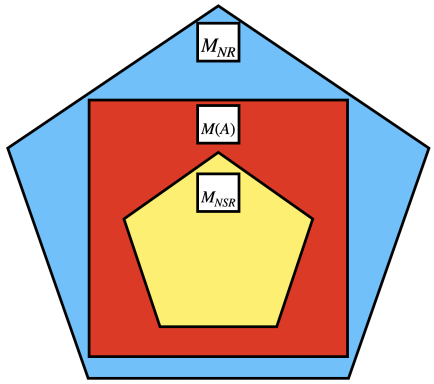

Fig. 1. Different no-regret menus in relation to each other. Here, A is any no-regret algorithm and $M(A)$ is its menu. $M(A)$ must be contained within M_{NR}, and must contain M_{NSR}.

clear whether algorithmically distinct NSR algorithms—such as [Blum and Mansour 2007] and [Dagan et al. 2024]—admit the same set of achievable outcomes, or whether their differing update rules lead to fundamentally different strategic guarantees. The menu framework resolves this question.

Theorem 2.6 [Arunachaleswaran et al. 2024]. *If $\mathcal{A}$ is a no-swap-regret algorithm, then $\mathcal{M}(\mathcal{A}) = \mathcal{M}_{NSR}$. Further, the no-swap-regret menu $\mathcal{M}_{NSR}$ is the convex hull of all CSPs of the form $x \otimes y$, with $x \in \Delta_m$ and $y \in BR_L(x)$.*

This result both collapses the menu of all no-swap-regret algorithms to the same menu and simultaneously provides a vertex definition of this polytope, to go along with its hyperplane based definition. We discuss the implications of this result -

(1) First, all no-swap-regret algorithms are *asymptotically equivalent*, in the sense that regardless of which no-swap-regret algorithm you run, any asymptotic strategy profile you converge to under one algorithm, you could also converge to under another algorithm (for appropriate play of the other player). This is true even when the no-swap-regret algorithms appear qualitatively quite different in terms of the strategies they choose (compare e.g. the fixed-point based algorithm of [Blum and Mansour 2007] with the more recent algorithms of [Dagan et al. 2024] and [Peng and Rubinstein 2024]).
(2) In particular, there is no notion of regret that is meaningfully *stronger* than no-swap-regret for learning in (standard, normal-form) games. That is, there is no regret-guarantee you can feasibly insist on that would rule out some points of the no-swap-regret menu while remaining no-regret in the standard sense. In other words, the no-swap-regret menu is *minimal* among all no-regret menus: every no-regret menu contains $\mathcal{M}_{\text{NSR}}$, and no asymptotic menu (whether it is no-regret or not) is a subset of $\mathcal{M}_{\text{NSR}}$.
(3) Finally, these claims are *not* generally true for external regret. There are different no-regret algorithms with very different asymptotic menus (as a concrete example, $\mathcal{M}_{\text{NR}}$ and $\mathcal{M}_{\text{NSR}}$ are often different, and they are both asymptotic

menus of some learning algorithm by Theorem 2.4).

3. APPLICATIONS OF THE MENU FRAMEWORK

The rest of this note is a tour of three papers that develop the menu framework in complementary directions. One direction asks what it means for an algorithm to be *globally good* against all possible optimizers, and whether standard learning algorithms are actually optimal in that sense. A second direction asks how to choose an algorithm when you have a known *distribution* over opponent types. And a third direction revisits the classical swap-regret and correlated equilibrium story, and asks: outside of normal-form games, what does it mean for a learning algorithm to be *non-manipulable*?

3.1 Pareto-optimal algorithms: when is one learning algorithm strictly better than another?

If you must commit to a learning algorithm in a general-sum repeated game, but you do not know the utility of the optimizer, what does it mean for this commitment to be "optimal"? In [Arunachaleswaran et al. 2024], we formalize a dominance relation between learning algorithms that is explicitly strategic.

We say an algorithm $\mathcal{A}$ is *Pareto-dominated* if there exists another algorithm $\mathcal{A}'$ that achieves at least as high long-run utility as $\mathcal{A}$ against *every* opponent payoff function, and strictly higher utility for at least one opponent type. An algorithm is *Pareto-optimal* if it is not Pareto-dominated. This definition is intentionally weak: it does not insist on optimality against each opponent type, only that the algorithm is not uniformly outperformed across all types. Nonetheless, we show in [Arunachaleswaran et al. 2024] that many standard no-regret algorithms, such as Multiplicative Weights and Follow-the-Perturbed-Leader, fail even this lenient benchmark. This leads to the question of whether there is any algorithm that is both no-regret and Pareto-optimal.

Our main positive result is that the stronger condition of *no-swap-regret* is sufficient for Pareto-optimality. We prove this result by redefining Pareto-optimality in terms of menus and characterizing it via geometric properties, which we then show that no-regret menus satisfy. One key property is inclusion-minimality; a menu $\mathcal{M}$ is *inclusion-minimal* if there is no valid menu whichx is strictly contained within $\mathcal{M}$.

Theorem 3.1 Informal. *Any algorithm whose menu is inclusion-minimal and contains the maximum learner utility CSP is Pareto-optimal.*

3.2 Unknown opponents: choosing an algorithm under a distribution over optimizer types

The Pareto-optimality benchmark is most natural when the optimizer's utility is completely unknown. In the other extreme, when the learner has full knowledge of the optimizer's utility, this becomes a Stackelberg game in the strategy space of algorithms, and approximately optimal solutions are known for this problem [Collina et al. 2023]. But in many applications, one has partial information: a model of competitor behavior in pricing, historical data about user tradeoffs on a platform, or a prior over agent "types." This motivates a Bayesian version of the algorithm-design problem.

In [Arunachaleswaran et al. 2025], the learner commits to an algorithm, the opponent payoff function is drawn from a distribution $\mathcal{D}$, and the opponent then best-responds in algorithm space. The learner's objective is to maximize expected long-run utility given knowledge of $\mathcal{D}$. It is not clear how much the learner can infer about the optimizer's realized type; as the optimizer is non-myopic, it might be in their best interest to obfuscate. Additionally, the learner must balance their attempt to learn, to the extent possible, versus exploiting their information.

Naively, this problem may initially seem computationally prohibitive: the space of history-dependent algorithms is enormous, and computing a best response to a fixed algorithm can be NP-hard (making it tricky to systematically predict or evaluate the best response against a candidate commitment). Menus again make the problem tractable in natural regimes. Each opponent type induces a linear objective over the learner's menu, and a distribution over opponent types therefore induces a distribution over linear objectives. The learner's problem becomes: choose an algorithm whose menu scores well in expectation under these objectives.

We optimize implicitly over the set of all menus by focusing only on what outcomes the menu incentivizes for each optimizer. This yields efficient methods for computing approximately optimal commitment algorithms under $\mathcal{D}$ (namely, polynomial time under constant game size or bounded support assumptions). Importantly, our characterization of menus constrains this search so as to guarantee that any optimized menu can be implemented by an actual learning algorithm.

We also study a robustness-motivated restriction in which the learner insists on no-regret behavior while still optimizing expected performance under $\mathcal{D}$. Here, the menu viewpoint again clarifies the answer: the right algorithmic class is no swap regret, which both preserves the relevant strategic robustness and admits optimization of the induced menu.

3.3 Swap regret and correlated equilibria beyond normal-form games: starting from non-manipulability

In normal-form repeated games, no swap regret has a clean strategic interpretation: it is tightly connected to correlated equilibrium and captures a sense in which the opponent cannot benefit from dynamically steering the learner. This motivates taking "absence of dynamic manipulation" as a primitive desideratum.

Non-manipulability. Fix a learner algorithm $\mathcal{A}$ and an opponent utility function u_O. Let $V_T^{\text{dyn}}(\mathcal{A}, u_O)$ be the maximum expected total utility the opponent can achieve over T rounds using any history-dependent strategy against $\mathcal{A}$. Let $V_T^{\text{stat}}(\mathcal{A}, u_O)$ be the maximum expected total utility when the opponent is restricted to a *static* mixed strategy (the same distribution each round). We call $\mathcal{A}$ *non-manipulable* if for every u_O,

$$V_T^{\text{dyn}}(\mathcal{A}, u_O) - V_T^{\text{stat}}(\mathcal{A}, u_O) \;=\; o(T).$$

Informally: regardless of the opponent's objective, they cannot extract linear-in-T advantage by playing dynamically rather than statically.

Beyond Normal-Form Games Many strategic environments are not well-modeled as repeated normal-form games. Returning to pricing: payoffs may depend on per-period costs, demand states, or other exogenous factors, leading to repeated

Bayesian or more structured games. It was not clear what the "right" generalization of swap regret was in these richer settings; decompositions of action distributions are no longer unique, leading to many possible ways to "swap" the same mixed action. Thus, it was an open question whether there existed an efficient algorithm which was non-manipulable and no-regret. Prior algorithms that guaranteed non-manipulability and no-regret in these games were not known to be computationally tractable. Meanwhile, more tractable regret notions did not actually rule out manipulability. [Arunachaleswaran et al. 2025] studies a broad class called *Polytope games* (which generalizes both Bayesian games and Extensive-form games) and uses non-manipulability as the starting point for defining an appropriate notion of swap regret.

In menu language, non-manipulability has a particularly crisp geometric meaning: beneficial manipulation manifests as an extreme point that cannot be generated by any static product distribution. Equivalently, $\mathcal{A}$ is non-manipulable if and only if every extreme point of its menu is a product distribution. The no-regret condition further constraints what sort of product distributions the menu must contain. While the question of "what should be swapped" does not have a natural translation beyond normal-form games, the menu interpretation of non-manipulability operates identically.

We introduce a new notion of swap regret, *profile swap regret*, which is measured by the distance from a CSP to a specific "non-manipulable" menu composed of product distribution extreme points. Profile swap regret is efficiently achievable (our work gives algorithms with $O(\sqrt{T})$ profile swap regret) and guarantees non-manipulability.

4. CLOSING THOUGHTS

Against non-myopic opponents, a learning algorithm effectively serves as a commitment, shaping the set of long-run outcomes the opponent can induce. Menus provide an explicit geometric characterization of these outcomes and a unifying lens for comparing algorithms (via Pareto dominance), selecting algorithms under priors (Bayesian design), and identifying stability notions that prevent manipulation in richer game models (profile swap regret).

While this work focuses on the two-player setting, the dynamics of multiple interacting learning algorithms and strategic agents remains an open frontier, where menus may help define appropriate notions of equilibria. Moreover, in complex strategy spaces, such as high-dimensional auctions or pricing environments, understanding which menus admit computationally efficient optimization becomes a critical challenge. Ultimately, as algorithms increasingly govern economic interactions, viewing them not merely as adaptive procedures but as geometric commitments that induce constrained outcome sets will be essential for designing robust and transparent marketplaces.

REFERENCES

Arunachaleswaran, E. R., Collina, N., Mansour, Y., Mohri, M., Schneider, J., and Sivan, B. 2025. Swap regret and correlated equilibria beyond normal-form games. In *Proceedings of the 26th ACM Conference on Economics and Computation*. 130–157.

Arunachaleswaran, E. R., Collina, N., and Schneider, J. 2024. Pareto-optimal algorithms for learning in games. In *Proceedings of the 25th ACM Conference on Economics and Computation*. 490–510.

Arunachaleswaran, E. R., Collina, N., and Schneider, J. 2025. Learning to play against unknown opponents. In *Proceedings of the 26th ACM Conference on Economics and Computation*. 478–504.

Blum, A. and Mansour, Y. 2007. From external to internal regret. *Journal of Machine Learning Research 8,* 6.

Braverman, M., Mao, J., Schneider, J., and Weinberg, M. 2018. Selling to a no-regret buyer. In *Proceedings of the 2018 ACM Conference on Economics and Computation*. 523–538.

Collina, N., Arunachaleswaran, E. R., and Kearns, M. 2023. Efficient stackelberg strategies for finitely repeated games. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*. 643–651.

Dagan, Y., Daskalakis, C., Fishelson, M., and Golowich, N. 2024. From external to swap regret 2.0: An efficient reduction for large action spaces. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*. 1216–1222.

Peng, B. and Rubinstein, A. 2024. Fast swap regret minimization and applications to approximate correlated equilibria. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*. 1223–1234.

Turning defense into offense in $O(\log 1/\epsilon)$ steps: Efficient constructive proof of the minimax theorem

GABRIELE FARINA
Massachusetts Institute of Technology

Von Neumann's minimax theorem asserts that the ability to defend against any opponent strategy implies the existence of an offensive strategy that guarantees the same value. This note revisits that symmetry from a constructive, oracle-based point of view. Given access to a *defense oracle* that, for any opponent strategy x, returns a response y guaranteeing payoff at least v, we ask how efficiently one can compute an *offense* strategy $y^\star$ that guarantees value at least $v - \epsilon$ against all x. A classical construction via no-regret learning yields such a $y^\star$ after $O((1/\epsilon)^2)$ calls to the defense oracle. In this note, I describe a different construction that uses only $O(\log(1/\epsilon))$ calls to the oracle (up to polynomial factors in the dimension and encoding size). I then illustrate this primitive through three applications: computing Φ-equilibria in convex and extensive-form games beyond polynomial type, computing expected solutions to variational inequalities, and computing expected fixed points of possibly discontinuous maps.

Categories and Subject Descriptors: Mathematical optimization [**Theory of computation**]: Design and analysis of algorithms

General Terms: Algorithms; Theory

Additional Key Words and Phrases: Minimax theorem; equilibrium computation; variational inequalities

1. INTRODUCTION

Suppose we are playing a two-player zero-sum game, and strategies are allowed to be mixed: each (mixed) strategy is represented as a point in a convex and compact set $\mathcal{X}$ or $\mathcal{Y}$. If for any mixed strategy $x \in \mathcal{X}$ that my opponent could pick, I know how to construct a counterstrategy $y = h(x) \in \mathcal{Y}$ that secures me a certain score v, then von Neumann's minimax theorem states that I must have a strategy $y^\star$ that secures me the same score v no matter what my opponent could pick.

This symmetry between defense and offense is anything but obvious. In early work on game theory in the 1920s, Borel was able to prove special cases of this symmetry for small symmetric zero-sum games, but wrongly conjectured that the property could not hold beyond 3×3 games (Borel, 1921; Weinstein, 2022). He was proven wrong in 1928, when von Neumann proved his minimax theorem (v. Neumann, 1928), marking a turning point for modern game theory, and much more.

Since von Neumann's original proof, a repertoire of conceptually different proofs of the minimax theorem have been proposed. Many of them proceed by way of heavy machinery, including fixed-point theorems from topology, compactness arguments in functional analysis, and separation results in convex geometry (Borwein, 2016, §1.1). However, those arguments are typically nonconstructive, in the sense that they prove the *existence* of a good offense policy, but shed little light on how one could compute it starting from knowledge of how to play defense, that is, from the function h.

Author email: `gfarina@mit.edu`

This note concerns *constructive* proofs of the minimax theorem:

> *Given access to an oracle h that constructs good defense policies, how can we construct a good offense policy?*

Crucially, we will not require that $\mathcal{Y}$ be efficiently representable: we will simply require that h can be queried. This will enable constructing good offense policies even if $\mathcal{Y}$ is too complex to optimize over, as will be the case in the three applications mentioned in Sections 5–7.

As we recall in Section 3, a classical (and beautiful) such constructive proof is possible via a connection to regret minimization. However, that construction is quite slow: to find an offense strategy that gives value at least $v - \epsilon$ starting from defense strategies of value v, the construction requires evaluating $O((1/\epsilon)^2)$ defensive strategies.

In contrast, I will present a different construction that only requires $\log(1/\epsilon)$ calls to h. This rate improves exponentially over the regret-based one, and shows that the cost of computing near-optimal offense strategies can grow polynomially in the number of accuracy *bits*, rather than the inverse of the accuracy itself. As a consequence, under standard assumptions, such as polyhedrality of the strategy sets, this improvement leads to *exact* computation of offensive strategy in polynomial time. The construction extends the ellipsoid-against-hope algorithm of Papadimitriou and Roughgarden (2008) for computing correlated equilibria in succinct games, generalizing it beyond correlated equilibria, beyond polyhedrality, and beyond the need for a polynomial number of vertices, and is presented in a recent paper that Charis Pipis and I wrote for NeurIPS (Farina and Pipis, 2024).

We found this construction to be a helpful computational primitive for a number of problems that bear a connection with the minimax theorem, including applications in equilibrium computation, fixed points, and variational inequalities, as briefly recounted in Sections 5–7.

2. SETUP AND NOTATION

Let $\mathcal{X}$ and $\mathcal{Y}$ be convex and compact domains for which we have oracle access (*e.g.*, an efficient separation oracle, or a projection oracle), and $u : \mathcal{X} \times \mathcal{Y} \to [0, 1]$ be a bilinear utility function, that is, $u(x, y) = x^\top Ay$ for some matrix A. The x-player will want to minimize u, while the y-player will want to maximize it.

Von Neumann's minimax theorem states that

$$\min_{x\in\mathcal{X}} \max_{y\in\mathcal{Y}} u(x, y) = \max_{y\in\mathcal{Y}} \min_{x\in\mathcal{X}} u(x, y).$$

The x-player is our opponent, while we are the y-player.

For any strategy $x \in \mathcal{X}$ of the opponent, the value $\max_{y\in\mathcal{Y}} u(x, y)$ is the best score that we can secure for ourselves by responding to x. Hence, the statement $\min_{x\in\mathcal{X}} \max_{y\in\mathcal{Y}} u(x, y) \geq v$ means that no matter what strategy x the opponent picks, we can always respond with a strategy that guarantees us a score of at least v. In other words, the statement $\min_{x\in\mathcal{X}} \max_{y\in\mathcal{Y}} u(x, y) \geq v$ is equivalent to the existence of a function $h : \mathcal{X} \to \mathcal{Y}$ such that for all $x \in \mathcal{X}$, $u(x, h(x)) \geq v$. Such a function h is what we will call a *defense oracle*.

On the other hand, the statement $\max_{y\in\mathcal{Y}} \min_{x\in\mathcal{X}} u(x, y) \geq v$ means that there is a strategy $y^\star \in \mathcal{Y}$ for us, such that no matter how the opponent responds, our score will be at least v. Von Neumann's minimax theorem states the rather unintuitive fact that if a defense oracle h exists, then such a strategy $y^\star$ must also exist.

In the following we will be interested in the following question: *given access to a defense oracle h, how can we construct a strategy $y^\star$ that guarantees us a score of at least $v - \epsilon$ against any opponent's strategy, in $O(\log(1/\epsilon))$ time (ignoring polynomial factors in the dimension and diameter of the sets $\mathcal{X}$ and $\mathcal{Y}$)?*

3. WARMUP: A SLOW CONSTRUCTION USING REGRET MINIMIZATION

A classical, neat, and deceptively short constructive proof of the minimax theorem can be derived from the mere fact that no-regret learning algorithms exist.

The idea is very simple. Suppose we are a y-player concerned with constructing an *offense* strategy $y^\star$ for ourselves, but so far have only figured out how to play *defense* via a defense oracle h. We do not know how our opponent, the x-player, will behave, so we'd better do some training. We will do so by simulating a plausible, fictitious x-player who plays according to a no-regret learning algorithm. In each round t, the x-player selects a strategy $x_t \in \mathcal{X}$, and we respond by playing the defense strategy $y_t = h(x_t)$, inducing a loss of Ay_t to the x-player.

After T rounds of this simulation, the regret incurred by the x-player is by definition

$$\mathrm{Reg}_T = \max_{\hat{x}\in\mathcal{X}} \sum_{t=1}^{T} (x_t^\top A y_t - \hat{x}^\top A y_t).$$

If the simulated x-player is no-regret, then by definition the regret is sublinear in the time horizon T. For concreteness, several algorithms are known to guarantee regret on the order of $\sqrt{T}$, ignoring polynomial factors in the dimension and diameter of the strategy set $\mathcal{X}$.

Dividing by T, and plugging in the bound $\mathrm{Reg}_T \le \sqrt{T}$ to get an estimate, we find that

$$\max_{\hat{x}\in\mathcal{X}} \frac{1}{T}\sum_{t=1}^{T} \hat{x}^\top A y_t \ge \frac{1}{T}\sum_{t=1}^{T} x_t^\top A y_t - \frac{1}{\sqrt{T}}.$$

The value of $x_t^\top A y_t$ is at least v for all t, since y_t is a defense strategy by construction. Furthermore, the average $(1/T)\sum_{t=1}^{T} \hat{x}^\top A y_t$ is equal to $\hat{x}^\top A \bar{y}$, where $\bar{y} = (1/T)\sum_{t=1}^{T} y_t$ is the average of the defense strategies played. Hence, we have shown that for any $\hat{x} \in \mathcal{X}$,

$$\hat{x}^\top A \bar{y} \ge v - \frac{1}{\sqrt{T}}.$$

In other words, the average strategy $\bar{y}$ is an offense strategy that guarantees us a score of at least $v - 1/\sqrt{T}$ against any opponent's strategy. Choosing $T = 1/\epsilon^2$, we find that $\bar{y}$ is an offense strategy that guarantees us a score of at least $v - \epsilon$.

In summary, the mere existence of regret minimization (for which many readily implementable algorithms exist) gives rise to a constructive proof of the minimax theorem. However, this construction is quite slow, requiring on the order of $1/\epsilon^2$ calls to the defense oracle h to construct an offense strategy that guarantees us a score of at least $v - \epsilon$.

4. TURNING DEFENSE INTO OFFENSE IN $O(\log(1/\epsilon))$ STEPS

In this section, we present a different construction that turns a defense oracle h into an offense strategy $y^\star$ in only $O(\log(1/\epsilon))$ calls to the defense oracle h. We presented this construction in Farina and Pipis (2024). It borrows several important ideas from the ellipsoid against hope algorithm for computing correlated equilibria in succinct games (Papadimitriou and Roughgarden, 2008), but generalizes them into a general framework disconnected from correlated equilibria. Furthermore, unlike Papadimitriou and Roughgarden (2008), our construction does not require polyhedral sets with a polynomial number of vertices, and as we will show later applies to games of non-polynomial type as well.

A main ingredient of the construction is the following. Consider the set of "unbeatable" strategies for the opponent, defined as

$$\Omega \coloneqq \{x \in \mathcal{X} : x^\top A y < v \quad \forall y \in \mathcal{Y}\}.$$

Being defined by linear inequalities, this set is convex. More importantly, this set is *empty*, since by assumption we have access to an oracle h that for any $x \in \mathcal{X}$ constructs a strategy $y = h(x)$ that guarantees us value at least v against x. For any strategy $x \in \mathcal{X}$, we can certify that $x \notin \Omega$ by considering the linear constraint corresponding to $y = h(x)$, which is violated, and hence we have a separation oracle for the set Ω.

Step 1: Reduction of constraints via the ellipsoid method. Ignoring details about numerical precision and bit representations, having access to a separation oracle for Ω allows us to run the ellipsoid method to certify that Ω is near empty, in the sense of having arbitrarily low volume $\leq \epsilon$, in a number of steps polynomial in the dimension and diameter of $\mathcal{X}$ and logarithmic in $1/\epsilon$. Each step of the ellipsoid method requires a call to the separation oracle for the ellipsoid center x, which in turn requires a call to the defense oracle h if $x \in \mathcal{X}$, or a call to the separation oracle for $\mathcal{X}$ if $x \notin \mathcal{X}$.

One can think of the ellipsoid method as a method for selecting, out of the infinitely many that define Ω, a polynomial subset of constraints that already certify (near) emptiness of the set. Under suitable assumptions on the bit representation of the problem, this can be easily also turned into a polynomial certificate of emptiness.

Step 2: An application of Farkas lemma. Once the set of constraints has been appropriately sparsified over the course of $T = O(\log(1/\epsilon))$ steps, we are left with a set of strategies

$$\tilde{\Omega} \coloneqq \{x \in \mathcal{X} : x^\top A y_t < v \quad \forall t = 1, \ldots, T\}$$

defined by a *finite* subsystem of linear inequalities that is already empty (or near-empty).

The Farkas lemma fully characterizes what it means for a set of linear inequalities to be infeasible: the only way that can happen is if there is a nonnegative, nontrivial linear combination of the inequalities that gives rise to a contradiction. In our case, this means that there must exist nonnegative multipliers $\lambda_1, \ldots, \lambda_T \geq 0$, not all zero, such that

$$\sum_{t=1}^{T} \lambda_t x^\top A y_t \geq v \sum_{t=1}^{T} \lambda_t \quad \forall x \in \mathcal{X}.$$

Dividing both sides by $\sum_{t=1}^{T} \lambda_t$ (which is positive by nontriviality), we find that there must exist a convex combination of the strategies y_t (which were obtained by calling the oracle h on the feasible centers of the ellipsoids constructed during the ellipsoid method) that guarantees us value at least v against any opponent's strategy $x \in \mathcal{X}$.

To find such multipliers λ_t, it suffices to maximize the concave function

$$d : \Delta^T \to \mathbb{R}, \qquad d(\lambda) \coloneqq \min_{x \in \mathcal{X}} \sum_{t=1}^{T} \lambda_t x^\top A y_t$$

over the T-dimensional simplex Δ^T. This is typically straightforward to carry out efficiently using standard convex optimization methods.

Compared to the regret minimization construction of Section 3, the construction in this section does not use uniform averaging, but at the same time it requires simulating defending against only $T = O(\log(1/\epsilon))$ strategies of the opponent before an offensive strategy can be extracted.

Remark 4.1. In the above, the convexity and compactness of $\mathcal{X}$ were needed to run the ellipsoid method on Ω. However, only convexity (and not compactness) was used for $\mathcal{Y}$. In other words, the same argument yields a constructive proof of the slightly stronger statement that

$$\min_{x \in \mathcal{X}} \sup_{y \in \mathcal{Y}} u(x, y) = \sup_{y \in \mathcal{Y}} \min_{x \in \mathcal{X}} u(x, y)$$

whenever $\mathcal{X}$ is convex and compact, $\mathcal{Y}$ is convex, and u is bilinear. This will come in handy in our applications, where $\mathcal{Y}$ will often be the space of distributions over a convex compact set.

5. APPLICATION 1: EFFICIENT COMPUTATION OF EQUILIBRIA IN CONVEX GAMES

We can combine the previous construction with an observation by Hart and Schmeidler (1989) to produce algorithms that compute exact Φ-equilibria in multiplayer games, beyond the polynomial type requirement of the original paper of Papadimitriou and Roughgarden (2008).

To fix ideas, let's suppose we would like to compute an exact coarse correlated equilibrium in a convex game. Each player i has a convex and compact strategy set $\mathcal{X}_i$, and a utility function $u_i : \mathcal{X}_1 \times \cdots \times \mathcal{X}_n \to [0, 1]$ that is multilinear. A coarse correlated equilibrium is a joint distribution over strategy profiles such that no player can gain in expectation by unilaterally deviating to any fixed strategy. Formally, a distribution y over $\mathcal{X}_1 \times \cdots \times \mathcal{X}_n$ is a coarse correlated equilibrium if for all players i and all strategies $x_i \in \mathcal{X}_i$,

$$\mathbb{E}_{z \sim y}[u_i(z)] \geq \mathbb{E}_{z \sim y}[u_i(x_i, z_{-i})].$$

We can think of coarse correlated equilibria as a minimax strategy for the game between a *mediator* who picks y, and an *opponent* who picks a player i and a deviation strategy $x_i \in \mathcal{X}_i$ to check whether the equilibrium conditions are satisfied. It is direct to check that the preconditions (Remark 4.1) for the construction of

Section 4 hold. To construct an optimal strategy for the mediator player y we can then start from determining whether y has a defense oracle against any x.

> In the case of coarse correlated equilibria, a defense oracle is quite simple: if the mediator knows the strategy of the opponent, that is, the distribution over pairs (i, x_i), they can respond by picking a product distribution y whose moments match the marginals of the opponent's distribution. This trivially guarantees that no player can gain by deviating to any fixed strategy, since the expected utility of each player under y is equal to the expected utility when the opponent picks their deviation.

The construction of Section 4 then allows us to turn this defense oracle into an offense strategy (distribution over strategy profiles) that works no matter what the opponent picks, and that is a coarse correlated equilibrium by definition.

Since each product distribution returned by the defense oracle admits a succinct representation (via its marginals), all while allowing d to be maximized efficiently, this construction shows that coarse correlated equilibria can be computed in polynomial time well beyond games of polynomial type. In particular, for extensive-form games, which were explicitly excluded by Papadimitriou and Roughgarden (2008), it shows that a coarse correlated equilibrium can be computed in time polynomial in the game tree size.

Finally, the same idea applies more generally to more complicated notions of Φ-equilibria, as long as a defense oracle can be constructed efficiently (usually via computing a suitable fixed point of the deviation functions). For example, this was a key ingredient in Daskalakis et al. (2025). Furthermore, this framework recovers the algorithm for extensive-form correlated equilibria of Huang and von Stengel (2008) as a particular instantiation.

6. APPLICATION 2: VARIATIONAL INEQUALITIES

Another application that my coauthors and I found useful concerns variational inequalities. We studied this in a joint paper with Zhang, Anagnostides, Tewolde, Berker, Farina, Conitzer, and Sandholm (2025a).

Given a convex and compact set $\mathcal{Z} \subseteq \mathbb{R}^d$ and a function $F : \mathcal{Z} \to \mathbb{R}^d$, the variational inequality problem asks to find a point $z^\star \in \mathcal{Z}$ such that

$$F(z^\star)^\top (x - z^\star) \geq 0 \quad \forall x \in \mathcal{Z}.$$

When F is the gradient of a function f, the variational inequality captures the first-order optimality conditions for the constrained optimization problem $\min_{z \in \mathcal{Z}} f(z)$. When $\mathcal{Z} = \Delta^m \times \Delta^n$ and F is the linear map

$$F(z_1, z_2) = -\begin{pmatrix} 0 & A \\ B^\top & 0 \end{pmatrix} \begin{pmatrix} z_1 \\ z_2 \end{pmatrix},$$

the variational inequality captures the Nash equilibria of the bimatrix game with payoff matrices $A, B \in \mathbb{R}^{m \times n}$. Given the hardness of approximating Nash equilibria, the previous example immediately shows that, even for *linear* maps F, approximating a solution to the variational inequality is computationally intractable.

However, we can compute a distribution $y \in \Delta(\mathcal{Z})$ over $\mathcal{Z}$ such that

$$\mathbb{E}_{z\sim y}[F(z)^\top (x - z)] \geq -\epsilon \quad \forall x \in \mathcal{Z} \tag{1}$$

in poly$(d, \log(1/\epsilon))$ time, using the construction of Section 4.[1] In other words, as long as we are willing to seek a *distribution* over points in $\mathcal{Z}$ (rather than a single point), we can bypass the hardness results tied to variational inequalities.

To do so, observe that (1) captures a game between the y-player and the x-player, with bilinear utility $(x, y) \mapsto \mathbb{E}_{z\sim y}[F(z)^\top (x - z)]$.

A defense oracle of value 0 for the y-player is easy to construct. If we know the point $x \in \mathcal{Z}$ that the opponent picks, we could respond by selecting the distribution $y = h(x)$ that puts all the mass on the point x, which trivially satisfies

$$\mathbb{E}_{z\sim y}[F(z)^\top (x - z)] = F(x)^\top (x - x) = 0 \geq 0.$$

Hence, we can use our general construction from Section 4 to efficiently upgrade the defense oracle into an "offensive" point $y \in \Delta(\mathcal{Z})$ that works for all x, that is, a solution to the expected variational inequality problem (1).

In a recent paper (Zhang et al., 2025a), we further strengthen this construction to show that we can efficiently compute a stronger notion of expected solution: find a distribution $y \in \Delta(\mathcal{Z})$ over $\mathcal{Z}$ such that

$$\mathbb{E}_{z\sim y}[F(z)^\top (Q(z) - z)] \geq -\epsilon \quad \forall \text{ affine endomorphisms } Q : \mathcal{Z} \to \mathcal{Z}.$$

The special case in which Q is taken to be a constant map recovers (1). It can be shown easily that for suitably constructed operators F, these notions recover coarse-correlated equilibria and refinements of correlated equilibria in concave games.

For general operators F (i.e., not necessarily representing utility functions in games), it is not understood what these concepts recover, but to my knowledge, they are the strongest *tractable* relaxation of the notion of solution for general variational inequalities that we can compute today.

7. APPLICATION 3: EXPECTED FIXED POINTS

As a third application of this construction, consider the problem of computing a fixed point of a function $f : \mathcal{Z} \to \mathcal{Z}$. When f is continuous and $\mathcal{Z} \subseteq \mathbb{R}^d$ is convex and compact, Brouwer's fixed point theorem guarantees the existence of a fixed point $z^\star \in \mathcal{Z}$ such that $f(z^\star) = z^\star$. However, computing such a fixed point is in general computationally intractable, given its intimate connection to Nash equilibria. However, in a recent paper with Zhang, Anagnostides, Tewolde, Berker, Farina, Conitzer, and Sandholm (2025b), we show that we can compute a distribution $y \in \Delta(\mathcal{Z})$ over $\mathcal{Z}$ such that

$$\|\mathbb{E}_{z\sim y}[f(z) - z]\|_2 \leq \epsilon \tag{2}$$

in $O(\log(1/\epsilon))$ steps, by reduction to our construction of Section 4.

[1]Problems of the form (1) are called *outgoing minimax problems* by Foster and Hart (2021).

We can interpret (2) as a game. On one side, we, the y-player, seek to find a distribution with the property above. To keep us honest, our opponent, the x-player, can select a direction x from the unit Euclidean ball $\mathbb{B}^d$ and win the game if they can prove that y does not induce (in expectation) a fixed point in the direction x, that is, if $\mathbb{E}_{z\sim y}[x^\top(f(z)-z)] < 0$. In other words, we are seeking to solve the game

$$\arg\max_{y\in\Delta(\mathcal{Z})}\min_{x\in\mathbb{B}^d}\mathbb{E}_{z\sim y}[x^\top(f(z)-z)],$$

whose value we know is 0. Once again, we are within the hypotheses of the main construction (see also Remark 4.1).

To solve this game, we can again think in terms of defense and offense. If we knew the direction x that the opponent would pick, we could respond by picking as response the distribution $y = h(x)$ defined as

$$h(x) = \text{distribution putting all mass on} \arg\max_{z\in\mathcal{Z}} x^\top z.$$

This is a valid defense oracle: it guarantees $\mathbb{E}_{z\sim y}[x^\top(f(z)-z)] \le 0$ given that $x^\top z \ge x^\top f(z)$ for the maximizer z in the support of y by construction.

We can then use the construction of Section 4 to turn this defense oracle into a distribution $y \in \Delta(\mathcal{Z})$, that works no matter the direction x picked by the opponent. In particular, a convex combination of the defense strategies constructed during the process can be constructed in time polynomial in the dimension and $\log(1/\epsilon)$, and yields an offense strategy $y^\star$ such that

$$-\epsilon \le \min_{x\in\mathbb{B}^d}\mathbb{E}_{z\sim y^\star}[x^\top(f(z)-z)] = -\|\mathbb{E}_{z\sim y^\star}[f(z)-z]\|_2,$$

that is,

$$\|\mathbb{E}_{z\sim y^\star}[f(z)-z]\|_2 \le \epsilon.$$

As a side note, this construction works even if f is not continuous. We found it to be a fundamental building block for sidestepping the computational equivalence barrier of Hazan and Kale (2007) when constructing Φ-regret minimizers for nonlinear deviations Φ (Zhang et al., 2025b).

8. ACKNOWLEDGMENTS

I am grateful to Samuel Sokota, Brian Hu Zhang, Juan Carlos Perdomo, Jon Schneider, Martino Bernasconi, Charis Pipis, and Samuel Taggart for their useful feedback on various drafts of this note.

REFERENCES

Borel, E. 1921. La théorie du jeu et les équations intégralesa noyau symétrique. *Comptes rendus de l'Académie des Sciences 173,* 1304-1308, 58.

Borwein, J. M. 2016. A very complicated proof of the minimax theorem. *Minimax Theory and its Applications 1,* 1, 21–27.

Daskalakis, C., Farina, G., Fishelson, M., Pipis, C., and Schneider, J. 2025. Efficient learning and computation of linear correlated equilibrium in general convex games. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*. 542–553.

Farina, G. and Pipis, C. 2024. Polynomial-time computation of exact Φ-equilibria in polyhedral games. *Advances in Neural Information Processing Systems (NeurIPS) 37*, 4207–4234.

Foster, D. P. and Hart, S. 2021. Forecast hedging and calibration. *Journal of Political Economy 129,* 12, 3447–3490.

Hart, S. and Schmeidler, D. 1989. Existence of correlated equilibria. *Mathematics of Operations Research 14,* 1, 18–25.

Hazan, E. and Kale, S. 2007. Computational equivalence of fixed points and no regret algorithms, and convergence to equilibria. *Advances in Neural Information Processing Systems 20*.

Huang, W. and von Stengel, B. 2008. Computing an extensive-form correlated equilibrium in polynomial time. In *International Workshop on Internet and Network Economics (WINE)*. 506–513.

Papadimitriou, C. H. and Roughgarden, T. 2008. Computing correlated equilibria in multi-player games. *Journal of the ACM (JACM) 55,* 3, 1–29.

v. Neumann, J. 1928. Zur theorie der gesellschaftsspiele. *Mathematische annalen 100,* 1, 295–320.

Weinstein, J. 2022. On a mistranslation of a mistake about minimax.

Zhang, B. H., Anagnostides, I., Tewolde, E., Berker, R. E., Farina, G., Conitzer, V., and Sandholm, T. 2025a. Expected variational inequalities. In *International Conference on Machine Learning (ICML)*.

Zhang, B. H., Anagnostides, I., Tewolde, E., Berker, R. E., Farina, G., Conitzer, V., and Sandholm, T. 2025b. Learning and computation of Φ-equilibria at the frontier of tractability. In *ACM Conference on Economics and Computation (ACM EC)*. 1105–1105.

Information Design with Large Language Models: An Annotated Reading List

TAO LIN
The Chinese University of Hong Kong, Shenzhen
and
SAFWAN HOSSAIN
Harvard University
and
YILING CHEN
Harvard University

This annotated reading list surveys emerging works at the intersection of information design and large language models (LLMs). While classical information design theory studies signaling in abstract mathematical models, real-world communication often occurs in natural language. We highlight papers that use LLMs to elicit and communicate information through natural-language messages, cover techniques used in "information design + LLM" research such as language-space optimization and LLM proxies, and discuss papers on LLM persuasion. We aim to illustrate the potential of LLMs to bridge the gap between the theory and practice of information design.

Categories and Subject Descriptors: F [**Theory of Computation**]: Theory of computing—*Algorithmic mechanism design*

General Terms: Economics; Design; Human factors; Languages; Theory

Additional Key Words and Phrases: Persuasion, Signaling, Decision-making, Generative AI, LLM

1. INTRODUCTION

Information design asks a deceptively simple question: How should a player strategically select and present information to shape the belief and behavior of others? Traditionally, this question has been studied in stylized models, where the transmitted information is modeled as an abstract random variable (called "signal") correlated with the state of the world. However, real-world communication often occurs in *natural language*, which has long been an unwieldy object for formal analysis. The recent breakthroughs in generative AI, particularly Large Language Models (LLMs), have the potential to bridge this gap.

LLMs offer numerous new opportunities to information design. First, by engaging users in natural conversations, an LLM can elicit users' private preferences that are difficult to elicit through traditional methods (e.g., surveys) and one-shot interactions. Second, LLMs can serve as proxies for human decision-makers who receive information in natural language. Instead of assuming ideal Bayesian decision-makers, we can study how an LLM – prompted to simulate humans or directly make decisions on behalf of humans – responds to linguistic messages, and

Authors' addresses: `lintao@cuhk.edu.cn`, `shossain@g.harvard.edu`, `yiling@seas.harvard.edu`

use that as a testbed for designing communication policies. Third, LLMs offer a way to search over the vast space of natural language and interaction protocols, rather than over the abstract space of information structures (signaling schemes), enabling the design of information strategies that are both theoretically grounded and practically implementable.

This non-exhaustive annotated reading list surveys emerging works at this interface between information design and LLMs. It highlights papers that use generative AI to elicit and communicate information to influence decision-makers (1)(2)(3)(4). Some papers not directly related to information design are included because they provide a toolbox for "information design + LLM" research, such as language optimization techniques (5) and LLM proxies for human decision-makers (6)(7). As information design is closely related to persuasion, some empirical papers on "LLM + persuasion" are also discussed (8)(9)(10). The selected papers aim to give SIGecom readers an entry point to this rapidly developing area, emphasize conceptual connections between classical information design theory and modern generative AI technologies.

(1) "Information Design With Large Language Models" (arXiv, 2025).
Duetting, P., Hossain, S., Lin, T., Leme, R. P., Ravindranath, S. S., Xu, H., and Zuo, S.

This work defines an information design (Bayesian persuasion) problem with framing effect, where the framing of a signal shapes the receiver's prior belief in a non-Bayesian way, while the signal further updates the receiver's belief via Bayes rule. Besides theoretical characterization of the optimal joint design of framing and signaling, the paper uses LLMs to simulate the framing-to-belief mapping of the receiver and to optimize framing using a hill-climbing-based prompt optimization method.

(2) "Verbalized Bayesian Persuasion" (arXiv, 2025).
Li, W., Lin, Y., Wang, X., Jin, B., Zha, H., and Wang, B.

This work formulates Bayesian persuasion as a verbalized game, representing states and signals as texts, and signaling schemes (stochastic mappings from states to signals) as "writing styles" that control LLM's outputs. Equilibrium computation and prompt optimization techniques (e.g, OPRO (5)) are used to jointly optimize the signaling scheme and the receiver's responding strategy in language space, obtaining outcomes that are comparable to the theoretical Bayesian persuasion outcomes.

(3) "AI Realtor: Towards Grounded Persuasive Language Generation for Automated Copywriting" (arXiv, 2025).
Wu, J., Yang, C., Wu, Y., Mahns, S., Wang, C., Zhu, H., Fang, F., and Xu, H.

This work uses LLM to generate textual product descriptions (for houses) to highlight certain features of the product while adhering to the true product attributes. The set of highlighted features is personalized to the buyer, whose preference is elicited by pre-participation survey as well as LLM-assisted inference from revealed preference during interactions. Human-subject experiments demonstrate both attractiveness and validity of LLM-generated descriptions.

(4) "Algorithmic Persuasion Through Simulation" (arXiv, 2023).
Harris, K., Immorlica, N., Lucier, B., and Slivkins, A.

This theoretical work studies a Bayesian persuasion problem where the sender infers the receiver's private type by querying an oracle that simulates the receiver's behavior given signals. Such an oracle can be implemented by LLM (see (6)(7) for justification). The optimal joint querying and signaling strategy is characterized.

(5) "Large Language Models as Optimizers" (ICLR, 2024).
Yang, C., Wang, X., Lu, Y., Liu, H., Le, Q. V., Zhou, D., and Chen, X.

This influential work from Google DeepMind proposes a prompt optimization method, OPRO (Optimization by PROmpting), which treats an LLM as a gradient-free optimizer that, given the problem description in natural language, iteratively improves prompts based on previously tried prompts and their evaluation scores. OPRO and similar methods have been used to optimize linguistic strategies in previous works on information design with LLMs (1)(2).

(6) "Large Language Models as Simulated Economic Agents: What Can We Learn from Homo Silicus?" (EC, 2024).
Filippas, A., Horton, J. J., and Manning, B. S.

This work prompts LLMs with different "personas" and performs classical behavioral economics and social science experiments on those LLMs. By comparing LLM and human responses, this work demonstrates LLMs' ability to simulate human populations with diverse preferences, thereby providing new opportunities and drastically reducing the cost of human-subject experiments. Such an LLM proxy idea has been used by existing works on information design with LLM (1)(4).

(7) "LLM-Powered Preference Elicitation in Combinatorial Assignment" (arXiv, 2025).
Soumalias, E., Jiang, Y., Zhu, K., Curry, M., Seuken, S., and Parkes, D. C.

A more recent work using LLMs as proxies for humans, specifically for preference elicitation in combinatorial assignment. Students describe their preferences for course assignments to LLMs in natural language. Assignment algorithms then make queries to LLMs instead of students, alleviating students' burden.

(8) "The Persuasive Power of Large Language Models" (AAAI Conference on Web and Social Media, 2024).
Breum, S. M., Egdal, D. V., Gram Mortensen, V., Møller, A. G., and Aiello, L. M.

Information design is closely related to persuasion. This is a foundational work on the empirical evaluation of the persuasion ability of LLMs. It conducts systematic experiments where human subjects are exposed to arguments generated by either humans or LLMs, and measures how these arguments influence participants' beliefs or decisions. LLM-generated arguments are demonstrated to be comparable and sometimes more persuasive than human-written ones.

(9) "Persuasion with Large Language Models: a Survey" (arXiv, 2024).
Rogiers, A., Noels, S., Buyl, M., and Bie, T. D.

An empirical survey of how LLMs persuade humans in real-world domains such as politics, marketing, public health, and e-commerce. It identifies key factors for persuasive effectiveness, including personalization, interactivity, prompt design, model scale, and disclosure of AI authorship. It also surveys experiment paradigms, ethical considerations, and regulatory landscape.

(10) "The Earth is Flat because...: Investigating LLMs' Belief towards Misinformation via Persuasive Conversation" (ACL, 2024).
Xu, R., Lin, B., Yang, S., Zhang, T., Shi, W., Zhang, T., Fang, Z., Xu, W., and Qiu, H.

Complementing works on LLMs' ability to persuade humans (8)(9), this ACL 2024 Outstanding Paper flips the lens to examine LLMs' susceptibility to being persuaded by misinformation. The authors curated a dataset, "Farm (Fact-to-Misinform)", consisting of factual questions paired with wrong answers and arguments. They presented such misinformation to LLMs in multi-turn dialogues, finding that LLMs' initial correct beliefs can be manipulated over turns.

REFERENCES

Breum, S. M., Egdal, D. V., Gram Mortensen, V., Møller, A. G., and Aiello, L. M. 2024. The Persuasive Power of Large Language Models. *Proceedings of the International AAAI Conference on Web and Social Media 18*, 152–163.

Duetting, P., Hossain, S., Lin, T., Leme, R. P., Ravindranath, S. S., Xu, H., and Zuo, S. 2025. Information Design With Large Language Models. *CoRR abs/2509.25565*.

Filippas, A., Horton, J. J., and Manning, B. S. 2024. Large Language Models as Simulated Economic Agents: What Can We Learn from Homo Silicus? In *Proceedings of the 25th ACM Conference on Economics and Computation*. ACM, New Haven CT USA, 614–615.

Harris, K., Immorlica, N., Lucier, B., and Slivkins, A. 2023. Algorithmic Persuasion Through Simulation. *CoRR abs/2311.18138*.

Li, W., Lin, Y., Wang, X., Jin, B., Zha, H., and Wang, B. 2025. Verbalized Bayesian Persuasion. *CoRR abs/2502.01587*.

Rogiers, A., Noels, S., Buyl, M., and Bie, T. D. 2024. Persuasion with Large Language Models: a Survey. *CoRR abs/2411.06837*.

Soumalias, E., Jiang, Y., Zhu, K., Curry, M., Seuken, S., and Parkes, D. C. 2025. LLM-Powered Preference Elicitation in Combinatorial Assignment. *CoRR abs/2502.10308*.

Wu, J., Yang, C., Wu, Y., Mahns, S., Wang, C., Zhu, H., Fang, F., and Xu, H. 2025. AI Realtor: Towards Grounded Persuasive Language Generation for Automated Copywriting. *CoRR abs/2502.16810*.

Xu, R., Lin, B., Yang, S., Zhang, T., Shi, W., Zhang, T., Fang, Z., Xu, W., and Qiu, H. 2024. The Earth is Flat because...: Investigating LLMs' Belief towards Misinformation via Persuasive Conversation. In *Proceedings of the 62nd Annual Meeting of the Association for Computational*

Linguistics (Volume 1: Long Papers). Association for Computational Linguistics, Bangkok, Thailand, 16259–16303.

Yang, C., Wang, X., Lu, Y., Liu, H., Le, Q. V., Zhou, D., and Chen, X. 2024. Large Language Models as Optimizers. In *The Twelfth International Conference on Learning Representations*.

SIGEcom Exchanges Annotated Reading List: Multiclass Calibration

RABANUS DERR
University of Tübingen, Tübingen AI Center
and
JESSIE FINOCCHIARO
Boston College

ML model evaluation often takes one of two main approaches: *risk minimization*, associated with "high accuracy" or *calibration*, meaning that predictions are "trustworthy" and can be interpreted from a probabilistic lens. There is an extensive line of work which has studied the relationship between risk minimization and calibration, mostly focusing on the binary outcome setting. Even in the binary setting, there are a variety of proposed calibration metrics which non-trivially interact. In the multiclass label setting, the choices to be made are even more complex and particularly there are different semantics for different notions. Here, we briefly present an annotated reading list reviewing some of the proposed definitions and their relationships.

Introduction

The classical understanding of calibration is that a prediction is calibrated if among the days on which the probability of rain was forecasted is p, the average number of rainy days is p, e.g., [DeGroot and Fienberg 1983]. This is formalized by saying a predictor $f : \mathcal{X} \to [0, 1]$ is *calibrated* if

$$\Pr[Y = 1 | f(X) = p] \approx p \qquad \forall p \in \mathbf{im}(f) \ . \tag{1}$$

In the binary setting, Equation (1) satisfies two nice desiderata of trustworthiness:

a. *self-referential*: predicting p means the true probability of the positive label is p, and
b. it *precisely estimates the loss* incurred by a decision maker using the prediction.

The first of the two desiderata is relatively self-explanatory. The second requires more context in terms of a decision maker. We understand a decision maker as an agent equipped with a loss function mapping from actions and outcomes to scores. The decision maker can *precisely estimate the loss*, if the given prediction allows the decision-maker to precisely compute the expected loss for a taken action in comparison to the actual incurred loss for the same taken action. It is a matter of some computations to show that Equation (1) fulfills this requirement, if the decision maker orients its action only based on the prediction, e.g., [Zhao et al. 2021].

Vanilla calibration (Equation (1)) as provided above gives trustworthiness desiderata when the prediction task is binary, e.g., rain or no rain. If the set of considered possible outcomes grows, e.g., rain, sun, cloudy, i.e., $\mathcal{Y}$ is non-binary but finite, then

Authors' addresses: `rabanus.derr@uni-tuebingen.de`, `finocch@bc.edu`

the formal definitions of calibration, and their achieved trustworthiness desiderata need reconsideration.

A naïve extension of Equation (1) following [Kull et al. 2019] to,

$$\Pr[Y = y | f(X) = p] \approx p_y \qquad \forall y \in \mathcal{Y}, \forall p \in \mathbf{im}(f) \ , \tag{2}$$

where $f : \mathcal{X} \to \Delta(\mathcal{Y})$ and p_y denotes the y-component of the probability vector p, is problematic. The sample complexity of calibration grows in the number of conditional probabilities considered (cf., the "for all"-quantifier over $\mathbf{im}(f)$). Hence, Equation (2) called *full distribution calibration* results in an exponential blowup in sample complexity. Scholarship essentially suggests two ways around the problem.

One solution is to consider only a relevant subset of conditions usually defined by the downstream decision makers. That is, calibration is aimed to be achieved only around decision boundaries of the decision makers. This requirement can be further weakened by focusing only on the loss estimates instead of the action recommendations made through the predictions; this notion is called *decision calibration*.

One alternative proposal is that calibration is generalized and achieved for relevant summary statistics, such as the mean or class-wise distributions. The variety of definitions for multiclass calibration have been mainly proposed with a focus on computationally constructing predictors with small statistical complexity. This final semantic generalization of calibration is called *property calibration*.

Interestingly, recent work shows that the semantic notions of decision calibration and property calibration have a strict separation in the multiclass setting. In the following reading list, we try to reflect the approaches to multiclass calibration. The list is not meant to be exhaustive, but rather should demonstrate differing notions.

Reading List

Distribution calibration

(1) Kull et al. [2019] propose a "natively multiclass calibration" method. In doing so, they offer a clean introduction of a natural extension of Equation (1) to multiple classes Equation (2), albeit with exponential computational and sample complexity. They further relate the full calibration extension to multiple classes to other suggestions made in literature which demand for class-wise calibration, respectively best-class calibration.

(2) Gopalan et al. [2024] discuss the fragile relationship between definitions of calibration in multiclass settings and the need to balance (a) sample complexity, (b) computational complexity, and (c) robustness of calibration notions. To this end, they propose the metric of *smooth projected calibration error* for multiclass settings and analyze the sample and computational complexities of attaining a predictor with low calibration error in this sense. Their work focuses on distributional predictors which might be used for binary subset selection problems as the downstream decision.

Decision calibration

(3) Zhao et al. [2021] propose a definition of calibration for multiclass settings that is motivated by the usefulness of predictions for downstream decision-

making. Motivated by *loss outcome indistinguishability*, they say a predictor is *decision calibrated* if a loss-minimizing decision-maker attains near-optimal loss by trusting a model's probabilistic predictions. Importantly, they require the action space of the decision-maker to be polynomially bounded in the number of classes, which stands in contrast to distribution calibration, where action spaces are not considered, and distance is simply measured from a predicted to observed distribution.

(4) Fröhlich and Williamson [2024] study the evaluation of imprecise forecasts. That means that forecasts are not single probability distributions, but sets of probability distributions. The paper focuses on loss functions and calibration as evaluation metrics. Even though imprecise forecasting is a rather exotic topic, they are a perfect ground to study the meaning of evaluations. In particular, the authors argue that a distinction between the goal of trustworthy uncertainty estimates and the goal of recommending favorable actions is required for imprecise forecasts, but not for precise ones.

(5) Noarov et al. [2025] study the computation of sequential predictions which fulfill a polynomial number of unbiasedness conditions. In particular, the authors can guarantee sample efficient predictions which are calibrated in a multiclass setting. This is achieved by putting focus on decision relevant conditions, i.e., the unbiasedness conditions can be defined through action policies by the decision maker.

Property calibration

(6) Jung et al. [2021] provide computational methods to predict such that "moment multicalibration" is met. Multicalibration is the extension of calibration as in Equation (1) to simultaneously hold on a set of subgroups $\mathcal{G} \subseteq 2^{\mathcal{X}}$. Moment calibration refers to the understanding of Equation (1) as a moment matching task. That is, the expected value of the output should be equal to the prediction, conditioned on the prediction. The authors extend this moment matching idea beyond the first order moment to higher order moments, including the variance.

(7) Noarov and Roth [2023] generalize the notion of moment multicalibration [Jung et al. 2021] into Γ-multicalibration for continuous, real-valued properties.[1] In this work, the authors propose a definition of Γ-multicalibration, which intuitively suggests that, conditioned on a model f predicting a property value r (e.g., predicting *the mean* is r), the property value should be approximately r. They characterize the set of "calibratable" properties Γ and present batch and online algorithms to Γ-multicalibrate a given predictor $f : \mathcal{X} \to \mathbb{R}$ for a set of labels $\mathcal{Y} \subseteq \mathbb{R}$ for continuous, real-valued Γ.

(8) Gneiting and Resin [2023] take a statistical perspective on forecast evaluation and model diagnostics. The aim of the paper is to develop calibration for real-valued forecasts. In particular, the authors suggest the notion of T-calibration using the concept of identifiability for properties[2]. Their definition is, up to

[1]Properties are functions $\Gamma\colon \Delta_{\mathcal{Y}} \to \mathcal{R}$ mapping distributions over labels to descriptive statistics, such as the mean $\Gamma(p) = \mathbb{E}_{Y\sim p}[Y]$, or mode $\Gamma(p) = \arg\max_y p_y$.

[2]An identifiable property is a property $\Gamma\colon \Delta_{\mathcal{Y}} \to \mathcal{R}$ such that there exists a function $\nu\colon \mathcal{Y}\times\mathcal{R} \to \mathcal{R}$ with $\Gamma(p) = \gamma \Leftrightarrow \mathbb{E}_{Y\sim p}[\nu(Y,\gamma)] = 0$. For instance, the mean or the median.

minor details, equivalent to Γ-calibration from [Noarov and Roth 2023]. The developments within this paper and [Noarov and Roth 2023] have, even though strongly related, happened independently.

Relationships between semantic notions

(9) Derr et al. [2025] examine the works above (among others), and proposes the semantic clusters of distribution calibration, property calibration, and decision calibration to characterize the differences and relationships between the semantic notions. In the binary setting, Derr et al. [2025] shows the semantic notions are equivalent, but establishes that decision calibration and property calibration are strictly separate in multiclass settings.

REFERENCES

DeGroot, M. H. and Fienberg, S. E. 1983. The comparison and evaluation of forecasters. *Journal of the Royal Statistical Society: Series D (The Statistician) 32,* 1-2, 12–22.

Derr, R., Finocchiaro, J., and Williamson, R. C. 2025. Three types of calibration with properties and their semantic and formal relationships.

Fröhlich, C. and Williamson, R. C. 2024. Scoring rules and calibration for imprecise probabilities.

Gneiting, T. and Resin, J. 2023. Regression diagnostics meets forecast evaluation: conditional calibration, reliability diagrams, and coefficient of determination. *Electronic Journal of Statistics 17,* 2, 3226 – 3286.

Gopalan, P., Hu, L., and Rothblum, G. N. 2024. On computationally efficient multi-class calibration. In *Proceedings of Thirty Seventh Conference on Learning Theory*, S. Agrawal and A. Roth, Eds. Proceedings of Machine Learning Research, vol. 247. PMLR, 1983–2026.

Jung, C., Lee, C., Pai, M., Roth, A., and Vohra, R. 2021. Moment multicalibration for uncertainty estimation. In *Proceedings of Thirty Fourth Conference on Learning Theory*, M. Belkin and S. Kpotufe, Eds. Proceedings of Machine Learning Research, vol. 134. PMLR, 2634–2678.

Kull, M., Perello Nieto, M., Kängsepp, M., Silva Filho, T., Song, H., and Flach, P. 2019. Beyond temperature scaling: Obtaining well-calibrated multi-class probabilities with dirichlet calibration. In *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Eds. Vol. 32. Curran Associates, Inc.

Noarov, G., Ramalingam, R., Roth, A., and Xie, S. 2025. High-dimensional prediction for sequential decision making. In *Forty-second International Conference on Machine Learning*.

Noarov, G. and Roth, A. 2023. The statistical scope of multicalibration. In *Proceedings of the 40th International Conference on Machine Learning*, A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett, Eds. Proceedings of Machine Learning Research, vol. 202. PMLR, 26283–26310.

Zhao, S., Kim, M., Sahoo, R., Ma, T., and Ermon, S. 2021. Calibrating predictions to decisions: A novel approach to multi-class calibration. In *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds. Vol. 34. Curran Associates, Inc., 22313–22324.