SIGecom Job Market Candidate Profiles 2025

This is the ninth annual collection of profiles of the junior faculty job market candidates of the SIGecom community. The forty candidates for 2025 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-candidates2025@acm.org.

–Vasilis Gkatzelis and Jason Hartline

Fig. 1. Generated using the research summaries of the candidates.

Contents

JERRY ANUNROJWONG (Homepage, CV, Google Scholar)

Thesis: Essays in Robust Auction Design and Electricity Markets ('25)

Advisor: Omar Besbes and Santiago R. Balseiro, Columbia University

Brief Biography: Jerry Anunrojwong is a Ph.D. candidate in the Decision, Risk, and Operations division at Columbia Business School, co-advised by Omar Besbes and Santiago R. Balseiro. He has been recognized with several honors, including being a finalist in the 2022 George Nicholson Student Paper Competition.

Research Summary: I am interested in *market design*, broadly construed. A common thread throughout my research is the critical role of participants' strategic behavior and their interplay with constraints on market structure and information access. In particular, my research is motivated by three themes in markets: (1) *robustness*, (2) *information design*, and (3) applications to *energy and sustainability*.

Energy. In [1], I study the cost of centralized versus decentralized batteries in energy markets. Privately-owned batteries aim to maximize profit, which may not align with the system efficiency goal of minimizing cost. Our model identifies three ways that batteries can distort the market. We also quantify the cost of distortions and calibrate our model with real data. In California (resp. Texas), we have a 15% (resp. 25%) difference in possible cost reductions. These strategic costs could be significant, but they pale in comparison to the cost reduction achieved by having enough batteries in the system. We also show that certain regulations can backfire, so market power mitigation protocols needs to be carefully designed.

Robustness. In [2], I consider a seller optimizing over randomized DSIC mechanisms to minimize the worst-case gap (regret or ratio) between mechanism revenue and the benchmark, where the only thing known about the value distribution of n buyers is that it is i.i.d. and its support is on $[a, b]$. I show that if a/b is below a threshold, second-price auctions (SPA) is optimal; if a/b is above another threshold, a new class of mechanisms I call pooling auctions (POOL) is optimal; if a/b is between the two thresholds, a randomization between SPA and POOL is optimal.

Information Design. In [3], I study the effectiveness of information design in reducing congestion in social services catering to users with varied levels of need. Each arriving user decides either to wait for the service by joining an unobservable FCFS queue, or to leave and get her outside option. I show that with enough heterogeneity in need, information design not only Pareto dominates full-info and no-info mechanisms, in some regimes it achieves the same welfare as the "first-best."

Representative Papers:

[1] Battery Operations in Electricity Markets: Strategic Behavior and Distortions (Under Review at *Management Science*), with S. Balseiro, O. Besbes, B. Xu.

[2] Robust Auction Design with Support Information (Minor Revision at *Management Science* & EC 2023), with S. Balseiro, O. Besbes.

[3] Information Design for Congested Social Services: Optimal Need-Based Persuasion (*Management Science* 2022 & EC 2020), with K. Iyer, V. Manshadi.

ESHWAR RAM ARUNACHALESWARAN (Homepage, CV, Google Scholar)

Thesis: Learning, Games and Fairness: Algorithms for Decision-Making in Complex Environments ('25)

Advisors: Sampath Kannan, Anindya De, University of Pennsylvania

Brief Biography: I'm a final year PhD student at the University of Pennsylvania, working on problems in the intersection of algorithmic game theory and online learning. In particular, I study (learning) algorithms as strategies for repeated games and the dynamics they induce. My research has been supported by a gift from AWS AI for research in Trustworthy AI. Currently, I'm visiting the Simons Institute, UC Berkeley, participating in the year long program on transformers and Language models, where I seek to find problems in the intersection of transformers and game theory.

Research Summary: My research focuses on the interactions between learning algorithms in strategic settings. I aim to understand the strategic responses elicited by these algorithms and the resulting dynamics. Central to my inquiry is the question: what are good learning algorithms for strategic interactions? My work entails evaluating learning algorithms based on how they shape time-dynamic, adaptive best-responses from the other player under varying payoff structures. The questions I address in my work have immense theoretical significance for understanding the field of online learning algorithms and how it intersects with dynamics in games.

Key contributions include defining and analyzing the notion of Pareto-optimality in learning settings [1], optimal commitment in repeated games against a single opponent [2]/ distributions of opponents (work in submission), studying the phenomenon of algorithmic collusion in competitive markets [4], and developing novel insights into the non-manipulability of learning algorithms in repeated games (ongoing work). Through these works, I have demonstrated the unique properties and strategic implications of various learning algorithms, such as no-swap-regret algorithms, and provided new tools for analyzing learning dynamics in games.

A central theme of my research is the use of games as a lens to study learning algorithms, often revealing deeper properties of the algorithms themselves. For example, our results highlight a fundamental distinction between FTRL and no-swap-regret algorithms, driven by the way FTRL algorithms can be exploited in repeated games due to their inherent memory. We also prove the asymptotic equivalence of all no-swap-regret algorithms in strategic environments, a particularly significant result in the light of recent breakthrough results about NSR algorithms.

Representative Papers:

[1] Pareto-Optimal Algorithms for Learning in Games (EC 24)
with Natalie Collina and Jon Schneider

[2] Efficient Stackelberg Strategies for Finitely Repeated Games (AAMAS 23)
with Natalie Collina and Michael Kearns

[3] Oracle Efficient Algorithms for Groupwise Regret (ICLR 24)
with Krishna Acharya, Sampath Kannan, Aaron Roth and Juba Ziani

[4] Algorithmic Collusion Without Threats (ITCS 25)
with Natalie Collina, Sampath Kannan, Aaron Roth, Juba Ziani

SANJAY CHANDLEKAR (Homepage, CV, Google Scholar)

Thesis: Autonomous Broker for Smart Grids: Auction Theory, Broker Design and Bidding Strategies (2024)

Advisor: Sujit Gujar, IIIT Hyderabad

Brief Biography: Sanjay is a final year PhD student in the Machine Learning Lab at IIIT Hyderabad, where he is advised by Professor Sujit Gujar. He has been working with TCS Research Labs since 2018 in the Data and Decision Science group. His research interests revolve around Game Theory, Auction Theory, Reinforcement Learning and Deep Learning. Sanjay holds a bachelor's degree in Computer Science from Nirma University and is a recipient of the Gold Medal of the program.

Research Summary: My research focus has been on optimizing bidding strategies in periodic double auctions (PDAs), specifically in the context of smart grids, aiming to improve smart grids' economic and functional efficiency from the distribution company's perspective. Smart grids operate across three key markets: wholesale, tariff, and balancing. Brokers must procure electricity via day-ahead PDAs from the wholesale market and sell it to customers through competitive tariff contracts that incentivize reduced energy consumption during peak times, thereby mitigating peak demands. Towards this, my work addresses challenges like minimizing procurement costs to buy electricity from the PDAs (most relevant for SIGecom), designing optimal tariff contracts to build customer portfolios to sell procured electricity and mitigating the recurring problem of peak demands in smart grids.

To minimize procurement costs, we design bidding strategies for day-ahead PDAs to help brokers procure electricity at the most economical prices. We leverage auction theory and game theory to analyze agent equilibrium behaviour in double auctions, including single-item, multi-item, single-shot[1, 3], and sequential auction[2, 3] settings. However, theoretical analysis either becomes intractable with the increase in the number of players or requires complete information assumption. To overcome the theoretical limitations and make the bidding strategies applicable to real-world scenarios, we implement these strategies in broker agents using reinforcement learning techniques[1, 3, 4]. We show that the learning-based strategies can effectively converge to theoretical equilibria and can be easily extended for more complex real-world scenarios. With the help of these bidding strategies along with the tariff module, our broker agent, VidyutVanika, became the winner of the International smart grid competition (PowerTAC) in 2021 and 2022.

Representative Papers:

[1] Multi-unit Double Auctions: Equilibrium Analysis and Bidding Strategy using DDPG in Smart-grids (AAMAS'22) with E. Subramanian, S. Bhat, P. Paruchuri, and S. Gujar

[2] Optimizing Prosumer Policies in Periodic Double Auctions Inspired by Equilibrium Analysis (IJCAI'24) with B. Manvi, and E. Subramanian

[3] Equilibrium Analysis and Strategic Bidding for Buyers in Multi-unit PDAs (Under Review in AIJ'24) with B. Manvi, E. Subramanian, and S. Gujar

[4] A Novel Bidding Strategy for PDAs using MCTS in Continuous Action Spaces (EMAS'24) with E. Subramanian

THÉO DELEMAZURE (Homepage, CV, Google Scholar)

Thesis: Small Change in Expressiveness, Big Change in Outcome Quality; Analysing Voting with Axioms and Data. ('25)

Advisor: Jérôme Lang and Dominik Peters (Paris Dauphine University)

Brief Biography: I am a fourth-year Ph.D. candidate at Paris Dauphine University - PSL. I hold a Master's degree in AI and a Diploma from école Normale Supérieure (ENS). During my studies, I completed internships at NYU, Nokia Bell Labs, and TU Berlin. My research focuses on voting theory and, more broadly, on social choice theory. I actively promote new ideas from the COMSOC community online, particularly related to electoral reforms.

Research Summary: My research focuses on voting theory. In one of my most recent works [1], we considered generalizations of the Instant Runoff Voting (IRV) rule to weak order preferences (i.e., people might have indifferences in their rankings). We defined the Approval-IRV rule, which generalizes IRV, and showed with axiomatic characterizations that this generalization is the only one that satisfies interesting normative properties. We also studied how this rule would behave using synthetic and real data.

Most of my works [1,2,3,5] follow the same structure: given a model with preferences (e.g., weak orders), and a question (e.g., how to generalize IRV?), what are the different solutions and how do we differentiate between them? The last part generally involves an *axiomatic analysis* of the solutions (often with impossibility and characterization results) and always an *experimental analysis*, in which we apply the different solutions to synthetic and real datasets.

To collect more datasets on which we can test the proposed solutions, I like to conduct experimental surveys in which we ask participants how they would have voted in specific elections (for instance, a presidential one) with alternative voting methods, such as approval voting or IRV. Another important goal of these surveys is to make participants more aware of the alternatives that exist for the voting systems we use. I was part of such a survey during the 2022 French presidential election and led one in 2024 for the European election in France. We designed a website from scratch for each survey and gathered several thousand responses. This also helped us gather interesting feedback from the participants about the different voting methods we suggested. These surveys are essential to my research, as I use at least one dataset collected from them in all of my projects.

Representative Papers:

[1] Generalizing Instant Runoff Voting to Allow Indifferences (EC 2024) with D. Peters

[2] Comparing Ways of Obtaining Candidate Orderings from Approval Ballots (IJCAI 2024) with C. Dong, D. Peters, M. Tydrichova

[3] Selecting the Most Conflicting Pair of Candidates (IJCAI 2024) with Ł. Janeczko, A. Kaczmarczyk and S. Szufa

[4] Independence of Irrelevant Alternatives under the Lens of Pairwise Distortions (AAAI 2024) with J. Lang and G. Pierczynski

[5] Approval With Runoff (IJCAI 2022) with J. Lang, J.-F. Laslier, R. Sanver

PETER DOE (Homepage, CV)

Thesis: Two-Sided Matching with Constraints: Theory-Driven Solutions ('25)

Advisors: Luciano Pomatto, Caltech; Federico Echenique, UC Berkeley

Brief Biography: I am a Ph.D. candidate in Social Sciences at the California Institute of Technology. I earned my B.B.A. from Baylor University in 2020 with majors in Mathematics, Economics, and Statistics. I have variety of interests at the intersection of computer science and economics, including matching, social choice, and algorithmic game theory.

Research Summary: I am a microeconomic theorist working on algorithmic market design, in particular two-sided matching. My research revolves around broadening classic matching models to incorporate matching activity outside of the regular market. I provide actionable market interventions grounded in theory.

In my job market paper [1], I propose a new solution concept for matching markets when some agents have already found partners. This initial match endows market participants with rights: every participant has the right to remain with her initial partner. Tension arises because an agent may wish to abandon his initial partner, but that requires his initial partner's approval. I propose a new equilibrium solution requiring that an agent can only object to a match if her initial partner agrees to the objection. I then present an algorithm that constructs such a match. My algorithm generalizes the Deferred Acceptance and Top Trading Cycles algorithms. I show my algorithm is immune to a variety of misreporting strategies, and it has applications to numerous markets including the resident-to-hospital match, college admissions, school choice, and labor markets.

In a second paper [2] I investigate the distributional impact of agents preempting a centralized marketplace. The motivation behind this the empirical observation that many matching markets suffer from temporal unraveling, which is when matching agreements are made earlier and earlier in time. I show that less-desirable agents benefit from making earlier offers. I present a two-period model with two hospitals and a continuum of doctors. In the first period, doctors are uncertain about their preferences over hospitals, and hospitals can make exploding offers to the doctors. In the second period, doctors learn their preferences and a centralized clearinghouse coordinates the match between the hospitals and the (remaining) doctors. I show that the ex ante less popular hospital drives both hospitals to make exploding offers in the first period, which results in a match that is inefficient ex post. My result explains why less desirable programs are perceived as being harmed by the recent implementation of an "All-In" policy within the NRMP.

Representative Papers:

[1] Matching with a Status Quo: The Agreeable Core (poster at EC, 2024; working paper)

[2] Ranked-to-Match: The Effects of Early Matching in the NRMP (working paper)

SOROUSH EBADIAN (Homepage, CV, Google Scholar)

Thesis: Designing Fair and Socially-Aligned Decision-Making Mechanisms ('25)

Advisor: Nisarg Shah, University of Toronto

Brief Biography: Soroush is a PhD candidate in the Department of Computer Science at the University of Toronto. In fall 2023, he was a visitor at Harvard University, hosted by Ariel D. Procaccia. He holds an Ontario Graduate Fellowship (2023-24) and multiple departmental awards. Previously, he obtained a B.Sc. in Computer Engineering from Sharif University of Technology in 2020 and was a gold medalist in the Iranian National Olympiad in Informatics (2014).

Research Summary: My research lies at the intersection of computer science and economics, focusing on designing collective decision-making systems that are aligned with societal goals, ensuring fairness and economic efficiency.

Democratic systems. How can we best aggregate ranked preferences to select from a set of candidates? While ranked preferences often fail to capture voters' preference intensities, my work [1] resolves the long-standing open problem of identifying the optimal voting rule with minimum loss in welfare. I have also explored various elicitation formats and the simplicity and explainability of voting rules.

Sortition, an ancient democratic process, has been revived to select citizens' assemblies worldwide, often using uniform random selection. While this method ensures fairness, does it truly represent the population? In [2], we propose a new measure of representation, identifying cases where uniform random selection is fair and representative, and in others, designing algorithms that improve the trade-off.

Resource and task allocation. How can we fairly and efficiently divide goods or chores among people? For goods, the state-of-the-art method guarantees approximate envy-freeness and Pareto efficiency. For chores, my work [3] makes the first non-trivial progress to this problem when individuals classify tasks as easy or difficult. Our algorithm, conjectured to work for all instances, is now deployed to the not-for-profit website Spliddit.org. Moreover, in [4], we propose a new fairness notion, inspired by envy-freeness, that applies to a broad range of collective decision-making settings from voting to participatory budgeting and peer review.

Pluralistic AI alignment. My research also addresses aligning AI agents with multiple individuals having diverse preferences and goals [5]. By using insights from social choice and Markov decision processes, we develop methods for aggregating individual policies into a desirable collective policy.

Representative Papers:

[1] Optimized Distortion and Proportional Fairness in Voting (EC'22 and ACM TEAC) with A. Kahng, D. Peters, and N. Shah

[2] Is Sortition Both Representative and Fair? (NeurIPS'22) with G. Kehne, E. Micha, A. D. Procaccia, and N. Shah

[3] How to Fairly Allocate Easy and Difficult Chores (AAMAS'22) with D. Peters and N. Shah

[4] Harm Ratio: A Novel and Versatile Fairness Criterion (EAAMO'24) with R. Freeman and N. Shah

[5] Policy Aggregation (NeurIPS'24) with P. A. Alamdari and A. D. Procaccia

FRANCESCO FABBRI (Homepage, CV)

Thesis: Essays on Attention in Economics ('25)

Advisor: Pietro Ortoleva, Princeton University

Brief Biography: I am a PhD candidate in Economics at Princeton University. My research investigates theoretical problems related to information and uncertainty.

Research Summary: I am interested in addressing two broad questions: (i) the effect that information—whether exogenous or endogenously acquired— has on individuals' beliefs and subsequent choices; (ii) the design of decision-making models that rationalize behavioral anomalies, observed in the lab or theorized in existing models. I relate these questions to my papers, which I organize into three categories.

Rational Inattention. This framework models agents facing the trade-off between processing more information to improve decisions and saving on information costs. In [1], a producer sets the quality of its product before offering it for a fixed price to a consumer, who processes information about quality, incurring entropy costs. As the consumer becomes more attentive, quality rises, but trade frequency falls as high quality is produced less often. When attention costs vanish, only low quality is provided, causing market failure. [2] analyzes competitive firms' pricing strategies when consumers learn about prices only by paying attention to them. Two forces are at play: inattention may cause market failure, as firms charge high prices, but competition helps keep prices in check. When attention costs are sufficiently high, the effect of competition dominates, increasing trade and benefiting industries that obtain higher profits by competing rather than colluding.

Dynamic Games. The notion of Nash equilibrium is ubiquitous in the game theoretic analysis. However, establishing equilibrium existence has proven elusive in general settings, which are affected by dynamic considerations or where uncertainty entails uncountably infinite states. In [3], we relate equilibrium existence to players making imprecise observations about the history of the game and show that equilibrium exists in general games whenever any amount of idiosyncratic noise is included in players' observations.

Ambiguity aversion. This literature aims to accommodate the Ellsberg paradox and the related experimental evidence, which falsifies expected utility theory. In [4], I characterize the behavior of inattentive and ambiguity averse agents, showing it is consistent with violations of invariance under compression, a property that has been experimentally challenged. [5] axiomatically characterizes preferences that display less aversion to ambiguity as welfare improves.

Representative Papers:

[1] Attention Holdup (Job Market Paper)

[2] Competing to Commit: Markets with Rational Inattention (American Economic Review, 114, no.1, (2024):285-306), with C. Cusumano and F. Pieroth

[3] Stochastic Games with Noisy Informational Asymmetries (EC '24, Extended Abstract), with S. Moroni

[4] Rational Inattention with Ambiguity Aversion (Working Paper)

[5] Absolute and Relative Ambiguity Attitudes (Working Paper), with G. Principi and L. Stanca

ALIREZA FALLAH (Homepage, CV, Google Scholar)

Thesis: Algorithmic Interactions with Strategic Users: Incentives, Interplay, and Impact

Advisor: Asuman Ozdaglar, MIT

Brief Biography: Alireza Fallah is a postdoctoral researcher at UC Berkeley, hosted by Michael Jordan. In the summer of 2023, he obtained his Ph.D. in Electrical Engineering and Computer Science from MIT, where he worked with Asu Ozdaglar and Daron Acemoglu. He spent the fall of 2023 as the Gamelin Postdoctoral Fellow at the Simons Laufer Mathematical Sciences Institute (formerly MSRI), where he was a member of the Mathematics and Computer Science of Market and Mechanism Design program. He has received a number of awards and fellowships, including the honorable mention at the ACM SIGecom Doctoral Dissertation Award, the Ernst A. Guillemin MIT M.Sc. Thesis Award, the Apple Scholars in AI/ML Ph.D. Fellowship, the MathWorks Engineering Fellowship, and the Siebel Scholarship.

Research Summary: My research bridges economics, machine learning theory, and optimization to tackle the challenges arising from the interaction between machine learning (ML) algorithms and human behavior. Below are two key research directions I have explored:

• Data, the fuel that powers ML models and algorithms, is typically collected from users, which has raised various concerns, ranging from privacy to social welfare. My research addresses several key challenges that arise in this context, including: (i) understanding the impact of the emergence of data marketplaces and data monetization on user welfare and how regulations can improve it [1]; and (ii) designing mechanisms that provide both compensation and privacy guarantees to mitigate the issue of free-riding and incentivize data sharing [2].

• There is growing demand to embed societal values like privacy, fairness, and safety into ML models. While much research focuses on improving algorithms, the interaction between these models and the people they impact, particularly when individuals adjust their behavior in response to regulations, is underexplored. My research investigates how human incentives intersect with regulatory interventions, such as safety inspections in contract design [3] or fairness constraints in dynamic auction design [4], identifying when these interventions benefit users and when strategic behavioral shifts may cause unintended consequences.

Representative Papers:

[1] On Three-Layer Data Markets (Under review, shorter version accepted for oral presentation at the ICML Workshop on Agentic Markets), with A. Makhdoumi, A. Malekian, and M. I. Jordan

[2] Optimal and Differentially Private Data Acquisition: Central and Local Mechanisms (Operations Research 2023 & EC 2022), with A. Makhdoumi, A. Malekian, and A. Ozdaglar

[3] Contract Design with Safety Inspections (EC 2024), with M. I. Jordan.

[4] Fair Allocation in Dynamic Mechanism Design (NeurIPS 2024), with A. Ulichney and M. I. Jordan

KARL FEHRS (Homepage, CV)

Thesis: Optimization, Learning, and Fairness in Voting ('25)

Advisor: Ioannis Caragiannis, Aarhus University

Brief Biography: I am a Ph.D. student at the CS department of Aarhus University, Denmark, graduating in spring 2025. My research focus has been on topics at the interface of computer science and economics, usually motivated by questions in social choice theory. During my Ph.D. studies, I had the great honour to be hosted by Prof. Ariel Procaccia at Harvard University for a six month research stay. I hold a B.Sc. degree in CS from Goethe University Frankfurt, Germany, where I also took up my Master's studies in CS. I obtained my M.Sc. degree in CS from Aarhus University as part of my Ph.D. program. I also possess a German diploma in Business Law and work experience in the financial industry.

Research Summary: My research is primarily aimed at understanding decision making processes which are typically framed as voting problems. Specifically, this lead me to study approval-based committee scoring rules [1], metric multiwinner voting [2], and single winner voting under random utilities [3].

On the one hand, my work followed the well-established approach of viewing decision making processes as optimization problems which has motivated the notion of *distortion* in voting. Beyond the classic distortion setting, my research was directed at understanding the additional power of allowing a limited number of queries to the exact cardinal preferences of the agents [2,3]. We introduced a novel average-case notion of distortion in a random utility model [3]. In this model, we were able to show that as few as one query per agent can drastically improve the average distortion to an extend that is not achievable under the traditional worst-case notion of distortion.

In addition, my research employed established CS techniques from learning theory (PAC-learning [1]), complexity theory (parameterized complexity [1]), and the study of algorithms (clustering [2]). Most recently, I have been working in the sortition model of voting theory where my efforts are focused on leveraging insights from the study of clustering algorithms. I expect to make a contribution to the growing literature on the interplay of these two models (sortition/clustering) soon.

In my future research, I would like to further investigate distortion or similar measures of optimality under randomized preferences (e.g., utilities drawn from probability distributions, Mallow's rankings). Exploring this direction outside of the voting setting, e.g., for matching markets, interests me as well. I would also be happy to pursue other directions in computational social choice, e.g., the study of axiomatic properties such as proportionality in voting and beyond.

Representative Papers:

[1] The Complexity of Learning Approval-Based Multiwinner Voting Rules (AAAI-22) with I. Caragiannis

[2] Low-Distortion Clustering with Ordinal and Limited Cardinal Information (AAAI-24) with J. Burkhardt, I. Caragiannis, M. Russo, C. Schwiegelshohn, and S. Shyam

[3] Beyond the Worst Case: Distortion in Impartial Culture Electorates (WINE-24) with I. Caragiannis

SIMON FINSTER (Homepage, CV)

Thesis: Essays on Competitive and Strategic Bidding in Multi-Object Auctions ('22)

Advisor: Paul Klemperer, University of Oxford

Brief Biography: I'm a postdoctoral fellow at CREST (Paris) and the Inria group FairPlay, hosted by Patrick Loiseau and Bary Pradelski. Previously, I was an Associate Fellow at the Simons Laufer Mathematical Institute in Berkeley, CA. I completed my PhD in Economics at the University of Oxford, Nuffield College, a Masters at the Paris School of Economics (APE), and an undergraduate degree in Industrial Engineering at the KIT (Germany).

Research Summary: My main research agenda focuses on equity and fairness concerns in market design.

In my job market paper [1], we initiate the study of surplus equity in auctions for multiple items. We characterize the surplus-equitable mechanism and develop prior-free results on equity-preferred mixed pricing. Our results imply simple policy recommendations for electricity markets, e.g., the equity benefits of a tax discussed by the New Zealand Electricity Authority (2014), and for auctions of carbon emission permits, treasury bonds, and beyond.

I have also proposed and studied a framework for indivisible goods markets with preferences over how goods or services should be distributed among buyers [2]. In related work, we have explored the theoretical foundations of markets with budget-constrained buyers, with applications to the potentially exploitative behavior of digital monopolies [3].

I enjoy bringing market design to the field and the lab, aiming to create productive channels between theoretical and applied work. As such, we have devised pooled testing mechanisms for infectious diseases, in populations with heterogeneous social welfare weights [4]. I am also conducting a large-scale virtual lab experiment (>1100 participants) that sheds light on bidding behavior in auctions for substitutes. In my postdoc research group and beyond, I collaborate with computer scientists and mathematicians. In current work, we advance the understanding of strategic behavior in complex auctions using methods from machine learning.

Representative Papers:

[1] Equitable Pricing in Auctions (Working Paper, 2024). *Job Market Paper.* with P. Loiseau, S. Mauras, M. Molina, and B. Pradelski.

[2] Selling Multiple Complements with Packaging Costs. (Working Paper, 2024). *Young Economist Essay Award Finalist EARIE 2021.*

[3] Substitutes Markets with Budget Constraints: Solving for Competitive and Optimal Prices (WINE 2023). *R&R at Theoretical Economics.* with P. W. Goldberg, and E. Lock.

[4] Welfare-Maximizing Pooled Testing (EC 2023). *Exemplary paper award in applied modeling track.* with M. González Amador, E. Lock, F. Marmolejo-Cossío, E. Micha, and A. Procaccia.

MATTHIAS GREGER (Homepage, CV, Google Scholar)

Thesis: Collective Choice from the Probability Simplex with Application to Donor Coordination ('25)

Advisor: Felix Brandt, Technical University of Munich (TUM)

Brief Biography: I am a fourth-year PhD student at TUM. My research interests comprise various topics from social choice and game theory with a focus on *how to reach and justify fair outcomes.* Together with my co-authors, I received the Best Student Paper Award at WINE 2021. I hold a B.Sc. and M.Sc. in Mathematics from TUM. During my studies, I was supported by the German Academic Scholarship Foundation. At the moment, I am in the process of finishing my dissertation.

Research Summary: My research concerns the aggregation of preferences over a convex set of outcomes, e.g., the probability simplex ([2]), and is applicable but not limited to donor coordination. Specifically, I aim at (i) finding and computing fair and Pareto optimal outcomes, (ii) explaining decisions and singling out the contribution of each individual, and (iii) investigating the stability of outcomes. To achieve these goals, I apply methods from game theory, optimization, and dynamical systems.

Donor coordination deals with the problem of distributing donations from a set of agents among a set of public projects/goods. On the one hand, a central coordination mechanism should distribute the total budget in a Pareto optimal and fair way (or at least give such recommendations). On the other hand, agents want to choose the distributions of their individual donations depending on the distributions of others and the "nature" of the projects. In detail, projects might be interpreted as *substitutes* ([4]) or *complements* ([1,3]). For both models, we prove that maximizing a weighted product of the agents' utilities leads to Pareto optimal and fair outcomes.

In [4], we show that this mechanism incentivizes agents to contribute to public goods even when private goods are available. In [1,3], we prove that the corresponding mechanism is not only fair and Pareto optimal but also strategyproof. In addition, we address questions regarding computability and demonstrate how desired outcomes arise as limits of some natural proportional or best response dynamics for both models.

With my work, I hope to contribute to our general understanding of fairness and illuminate ways for a fair and efficient provision of public goods.

Representative Papers:

[1] Coordinating Charitable Donations (Working paper)
with F. Brandt, E. Segal-Halevi, and W. Suksompong

[2] Optimal Budget Aggregation with Single-Peaked Preferences (EC 2024)
with F. Brandt, E. Segal-Halevi, and W. Suksompong

[3] Balanced Donor Coordination (EC 2023)
with F. Brandt, E. Segal-Halevi, and W. Suksompong

[4] Funding Public Projects: A Case for the Nash Product Rule (Journal of Mathematical Economics 2022, WINE 2021)
with F. Brandl, F. Brandt, D. Peters, C. Stricker, and W. Suksompong

DANIEL HALPERN (Homepage, CV)

Thesis: Social Choice in the Modern Era: Navigating AI, Uncertainty, and Unprecedented Scale (2025)

Advisor: Ariel D. Procaccia, Harvard University

Brief Biography: I am a Ph.D. student in computer science at Harvard University, where I am funded by an NSF graduate research fellowship and a Siebel Scholarship. Before joining Harvard, I completed a Bachelor's of Science at the University of Toronto, where I worked with Nisarg Shah.

Research Summary: Broadly, my research applies tools from social choice theory and fair division to new contexts, often inspired by artificial intelligence. In these settings, classical frameworks frequently fall short, because traditional assumptions no longer hold. To address these challenges, I develop new theoretical models to design provably robust systems, which I empirically validate on real data, when possible. Below, I highlight three specific instances of this overarching agenda:

In [1], we consider the problem of fine-tuning a Large Language Model (LLM), improving its outputs using human preference data between different prompt answers. This is inherently a social choice problem, as we must aggregate heterogenous human preferences into a single output LLM. However, to facilitate training, the output here must be a reward function, which can assign a score to an arbitrary LLM output. This does not fit neatly into any existing social choice framework which typically output a single answer or set of answers. Nevertheless, variants of axioms from social choice can still apply. We show that current aggregation methods fail to satisfy fundamental social choice properties, and complement this with new aggregation rules to address these issues.

In [2], we take the perspective of an opinion aggregation website such as pol.is, which facilitates a large-scale discussion on complex issues, from how to regulate climate change to how to prioritize city funds. Participants express preferences on statements submitted by others, and then the platform generates a summary of the opinion space. At first glance, this again resembles a social choice problem: aggregate participant preferences over statements into a summary. However, due to the platform's scale, it is infeasible to ask each participant for preferences on all possible statements. Thus we must make do with partial data. We theoretically demonstrate the representation properties we can guarantee and validate these findings using real data from Polis discussions.

In [3], we explore a novel form of governance called Liquid Democracy, made possible by modern digital platforms. Through both theoretical models and real human experiments in classroom settings, we demonstrate its strong performance at improving collective decision-making.

Representative Papers:

[1] Axioms for AI Alignment (NeurIPS'24 Spotlight Presentation) with L. Ge, E. Micha, A.D. Procaccia, I. Shapira, Y. Vorobeychik, J. Wu

[2] Representation with Incomplete Votes (AAAI'23) with G. Kehne, A.D. Procaccia, J. Tucker-Foltz, and M. Wüthrich

[3] Tracking Truth in Liquid Democracy (*Management Science*, EC'23) with A. Berinsky, J.Y. Halpern, A. Jadbabaie, E. Mossel, A.D. Procaccia, and M. Revel

MEENA JAGADEESAN (Homepage, CV, Google Scholar)

Thesis: Machine Learning Ecosystems of Self-Interested Agents ('25)

Advisor: Michael I. Jordan and Jacob Steinhardt, UC Berkeley

Brief Biography: I am a 5th year PhD student in Computer Science at UC Berkeley, affiliated with the Berkeley AI Research Lab. My research is supported by an Open Philanthropy AI Fellowship and a P.D. Soros Fellowship. I've interned at Microsoft Research in the Economics & Computation Group (summers '23, '24).

Research Summary: Modern ML models (e.g., LLMs and recommender systems) interact with humans, companies, and other models within a broader ecosystem. However, interactions between these agents often lead to unintended ecosystem-level outcomes, including clickbait, safety violations, and market concentration.

My research takes an economic perspective of these multi-agent interactions, towards a vision of ML ecosystems operating as well-functioning markets. I view models, humans, and companies as self-interested agents optimizing their own objectives. I aim to characterize how multi-agent interactions shape ecosystem-level outcomes, and develop interventions to steer outcomes towards societal objectives.

In LLM ecosystems, the companies which train or fine-tune LLMs compete for user usage. My research demonstrates how this form of competition distorts model performance and market structure. Specifically, when companies fine-tune the same LLM, we show that training the LLM with more resources can reduce user welfare [1]. Furthermore, when companies train different LLMs, we show that new companies can enter the market with much less data than incumbents [4]. The underlying driver is that companies strategically train their models to attract users.

On recommendation platforms, the ML model used for recommendations facilitates competition between users. My research demonstrates how user incentives amplify the impact of details of the ML model. Specifically, for content recommendation, we characterize how the recommendation model's learned embeddings shape the supply of available content, due to creator incentives [2]. Moreover, for matching platforms with prices, where stability captures user incentives, we design bandit-based recommendation algorithms minimizing cumulative instability [3].

In other work, I also leverage an economic perspective of ML ecosystems to study human-AI interactions, competing platforms, AI policy, and algorithmic fairness.

Representative Papers:

[1] Improved Bayes Risk Can Yield Reduced Social Welfare Under Competition (NeurIPS 2023) with M. I. Jordan, J. Steinhardt, and N. Haghtalab

[2] Supply-Side Equilibria in Recommender Systems (NeurIPS 2023) with N. Garg and J. Steinhardt

[3] Learning Equilibria in Matching Markets from Bandit Feedback (Full version at Journal of the ACM; conference version at NeurIPS 2021) with A. Wei, Y. Wang, M. I. Jordan, and J. Steinhardt

[4] Safety vs. Performance: How Multi-Objective Learning Reduces Barriers to Market Entry (Under submission) with M. I. Jordan and J. Steinhardt

DEVANSH JALOTA (Homepage, CV, Google Scholar)

Thesis: Algorithm and Incentive Design for Sustainable Resource Allocation ('25)

Advisor: Marco Pavone & Yinyu Ye, Stanford University

Brief Biography: Devansh is a final-year Ph.D. candidate in Computational and Mathematical Engineering at Stanford University, where he is a Thomas C. Nelson Stanford Interdisciplinary Graduate Fellow. Prior to joining Stanford, Devansh received his B.Sc. in Civil and Environmental Engineering and B.A. in Applied Mathematics from UC Berkeley.

Research Summary: Devansh's research develops data-driven and online learning algorithms and incentive schemes to advance the science and practice of market design for sustainable and society-aware resource allocation. Blending ideas from operations, economics, and computer science, his research pushes the frontiers of resource allocation through both *foundational* and *application-driven* research.

On the foundational front, he introduces and studies models that incorporate *societal* and *practical* considerations, including equity, fairness, privacy, uncertainty, and security, into classical resource allocation problems and develops foundational tools and algorithms for decision-making in these more complex settings. For instance, his work has addressed the privacy and information availability concerns of traditional equilibrium pricing approaches that rely on complete information of user preferences in the context of electricity and Fisher markets [1, 4]. Moreover, his work has laid the foundations for accommodating complex constraints arising due to fairness and security concerns in classical equilibrium models and developed algorithms for computing equilibria in these more complex settings [2, 5].

On the applied front, he leverages theory into applications, tailoring AI and optimization-driven algorithms for domains spanning future mobility systems [3], artificial currency markets [1, 2], and electricity markets [4]. Notably, his work on future mobility systems develops methods to address emerging transportation equity and data privacy and uncertainty challenges, with a keen focus on addressing the inequity issues surrounding road congestion pricing [3].

Representative Papers:

[1] Stochastic Online Fisher Markets: Static Pricing Limits and Adaptive Enhancements (Operations Research (Forthcoming), WINE 2023) with Y. Ye

[2] Fisher Markets with Additional Linear Constraints: Equilibrium Properties and Efficient Distributed Algorithms (Games and Economic Behavior 2023) with M. Pavone, Q. Qi, and Y. Ye

[3] Balancing Fairness and Efficiency in Traffic Routing via Interpolated Traffic Assignment (Journal of Autonomous Agents and Multi-agent Systems 2023) with K. Solovey, M. Tsao, S. Zoepf, and M. Pavone

[4] Online Learning for Equilibrium Pricing in Markets under Incomplete Information (Major Revision at Operations Research) with H. Sun, N. Azizan

[5] When Simple is Near-Optimal in Security Games (Working Paper) with M. Ostrovsky, M. Pavone

ANAND KALVIT (Homepage, CV)

Thesis: Improved Asymptotics for Multi-armed Bandit Experiments under Optimism-based Policies: Theory and Applications (2023)

Advisor: Assaf Zeevi, Columbia University Graduate School of Business

Brief Biography: Anand Kalvit is a postdoctoral fellow at Stanford University's Immigration Policy Lab, where he is exploring incentive-aware online learning approaches for algorithmic refugee assignment. He completed his PhD in Decision, Risk, and Operations from Columbia University's Graduate School of Business and holds Bachelor's and Master's degrees in Electrical Engineering from IIT Mumbai, India. In Fall 2023, he was a research fellow at the Simons Laufer Mathematical Sciences Institute (formerly MSRI), Berkeley. His work has been recognized with spotlight papers at NeurIPS 2021 and the INFORMS RMP Conference 2022, along with finalist placements in multiple INFORMS competitions.

Research Summary: Anand's research focuses on sequential decision-making under uncertainty, blending theory and applications at the intersection of online learning, optimization, and mechanism design. A central theme in his current work is guiding agents toward beneficial actions in dynamic systems, with applications ranging from recommender systems to refugee allocation.

His doctoral thesis delves into the analysis of multi-armed bandit algorithms using the diffusion approximation framework and the design of optimal policies for complex settings such as dynamic marketplaces, aiming to address practical challenges and translate theoretical insights into real-world solutions.

Looking ahead, Anand aims to integrate and advance methods from machine learning, optimization, and economics into societally impactful domains such as AI for social good. His ongoing projects explore adaptive queueing and personalized recommendations in healthcare contexts, with the goal of developing robust frameworks for decision-making in non-stationary and resource-constrained environments.

Representative Papers:

[1] Incentivized Exploration via Filtered Posterior Sampling (EC'24)
with Y. Gur, and A. Slivkins

[2] Complexity Analysis of a Countable-armed Bandit Problem (ALT'23)
with A. Zeevi

[3] Dynamic Learning in Large Matching Markets (NeurIPS'22)
with A. Zeevi

[4] Bandits with Dynamic Arm-acquisition Costs (Allerton'22)
with A. Zeevi

[5] A Closer Look at the Worst-case Behavior of Bandit Algorithms (NeurIPS'21)
with A. Zeevi

[6] From Finite to Countable-armed Bandits (NeurIPS'20)
with A. Zeevi

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 [3], 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 [1], 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 [2], 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 ongoing work (working single-author paper), I explore 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] Computation of Nash Equilibria of Attack and Defense Games on Networks (SAGT 2023) with M. Dziubiński

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

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

POOJA KULKARNI (Homepage, CV)

Thesis: Fair Allocation of Indivisibles Beyond Additive Valuations (2025)

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

Brief Biography: I am a final-year Computer Science Ph.D. student at the University of Illinois at Urbana-Champaign (UIUC). My research focuses on Fair Allocation of Resources across various settings, including discrete, continuous, offline, and online environments. I have published in top theory conferences such as SODA, ITCS, ICALP, and AI conferences like AAAI and AAMAS, with much of my work involving submodular and XOS maximization. Before my Ph.D., I completed my Master's at Indian Institute of Science (IISc) and Bachelor's at College of Engineering Pune (CoEP), earning a gold medal at both institutions. During my PhD, I have interned at Meta and was offered an internship at Google (declined).

Research Summary: Resource allocation, such as in food banks or ad-slot allocation, requires fairness to maintain societal harmony. Fairness has been studied extensively, leading to various notions. My research addresses two key questions: (1) Do fair allocations exist? (2) Can they be computed efficiently? I study these in both offline settings, where agents, goods, and preferences are known in advance, and online settings, where agents or goods arrive over time.

Offline In the offline setting, my research focuses on agents with submodular and XOS preference functions. I've studied fairness notions like Nash Social Welfare (NSW) for submodular [1] and XOS valuations [2], as well as Any Price Share (APS) for submodular valuations [3] and other subclasses beyond additive. My work involves techniques such as (1) Linear and concave programming, and (2) Combinatorial methods like greedy and local search, exposing me to concepts like continuous extensions of submodular functions, configuration LPs, and market-based fairness approximations. In an upcoming paper, we introduce a new dependent rounding scheme for submodular allocations.

Online We study a setting where goods are known in advance, but agents arrive over time. Fairness is challenging in offline settings and becomes more complex with unpredictable future demands. While comparing to a prophet often yields strong negative results, we can address this in two ways: (1) Using alternative benchmarks and (2) Leveraging predictive information. Given the advances in learning-augmented algorithms, my upcoming work goes for the second solution. We give positive results for MMS-fair share allocations using modest predictions about agent arrivals. My goal is to characterize the trade-off between the amount (and cost) of information learned and the achievable fairness approximation.

Representative Papers:

[1] Approximating Nash social welfare under Submodular Valuations through (Un)matchings (SODA, TALG) with Jugal Garg, Rucha Kulkarni

[2] Sublinear Approximation for Nash social welfare with XOS Valuations (ITCS) with Siddharth Barman, Anand Krishna, Shivika Narang

[3] $\frac{1}{2}$ Approximate MMS Allocation for Separable Piecewise Linear Concave Valuations (AAAI) with Chandra Chekuri, Rucha Kulkarni, Ruta Mehta

TAO LIN (Homepage, CV)

Thesis: Incentives and Learning in Information Design ('25)

Advisor: Yiling Chen, Harvard University

Brief Biography: Tao Lin is a 5th-year PhD student in Computer Science at Harvard University. He obtained a BSc in EECS from Peking University. During his PhD, Tao interned at ByteDance and Google, and received a Siebel Scholarship.

Research Summary: As machine learning algorithms increasingly shape real-world decision-making, the *strategic behavior* of participating agents – whether users or data providers – fundamentally impacts the algorithmic performance. The design of learning algorithms in strategic, dynamic multi-agent environments departs from the traditional machine learning paradigm that assumes exogenous, stationary data distributions. I investigate the complex interplay between incentives and learning in both theoretical economic models and real-world multi-agent systems, contributing to the community's common goal of building socially responsible AI systems.

- *Incentives and learning in mechanism design* [1]: I study the fundamental question of equilibrium convergence in repeated auctions with learning agents. Internet ad auctions provide a canonical example, where advertisers employ online learning algorithms to bid for ad slots. While convergence in truthful auctions is known, the dynamics in non-truthful auctions remained an open challenge. My work [1] provides the first complete characterization of when mean-based learning algorithms converge to Nash equilibrium in repeated first-price auctions and when they do not.

- *Incentives and learning in information design* [2, 3]: Traditional information design (Bayesian persuasion) assumes that agents can optimally process received information through Bayesian updating – an assumption that rarely holds in practice. My research [2] develops a novel framework where agents instead learn from experience using no-regret algorithms. Surprisingly, our result shows that the optimal utility of the sender remains nearly unchanged when facing learning agents, compared with the traditional model. Moreover, this finding generalizes to any general principal-agent problems including Stackelberg games and contract design.

- *Empirical study: recommender systems* [4]: During an internship at ByteDance, I conducted research on the impact of content creators' incentives on the polarization phenomenon in recommender systems. This work demonstrates how theoretical insights about strategic behavior can inform the design of deployed AI systems, advancing the goal of building healthy and sustainable online information ecosystems.

Representative Papers:

[1] Nash Convergence of Mean-Based Learning Algorithms in Auctions (WWW'22) with X. Deng, X. Hu, W. Zheng

[2] Generalized Principal-Agent Problem with a Learning Agent (working paper) with Y. Chen

[3] Multi-Sender Persuasion: A Computational Perspective (ICML'24) with S. Hossain, T. Wang, Y. Chen, DC. Parkes, and H. Xu

[4] User-Creator Feature Dynamics in Recommender Systems with Dual Influence (NeurIPS'24) with K. Jin, A. Estornell, X. Zhang, Y. Chen, Y. Liu

ANDREAS MAGGIORI (Homepage, CV, Google Scholar)

Thesis: Beyond Worst-Case Analysis, With or Without Predictions ('23)

Advisor: Ola Svensson, Rudiger Urbanke, EPFL

Brief Biography: I am a Postdoc at the Data Science Institute of Columbia University, working with Eric Balkanski and Will Ma. I earned my PhD from EPFL, and during that time, I interned twice at Google Research, hosted by Nikos Parotsidis and Ehsan Kazemi, respectively.

Research Summary: My research focuses on the field of *decision-making under uncertainty*–essentially, how to optimize a function when only partial information about its input is available. Traditional models of uncertainty assume that inputs are either adversarial (worst-case analysis) or drawn from a known distribution (stochastic analysis). Both models have limitations: worst-case analysis is often too pessimistic, ignoring useful information like typical instances and historical data, while stochastic analysis lacks robustness against noise, corruptions, outliers, and distribution shifts.

Learning-Augmented Algorithms ([1], [2], [3]): Hence, a central focus of my research is on intermediate models, particularly the area of learning-augmented algorithms. In this framework, we assume access to potentially imperfect predictions and seek to leverage these predictions without making any assumptions about their quality. Specifically, if the predictions are accurate, the algorithm should perform near optimally; if the predictions are poor, the algorithm must maintain robustness. This approach stems from the reality that machine learning models, while often accurate in practice, rarely offer worst-case guarantees.

However, relying on predictions can be a double-edged sword, as the algorithms' performance can be highly sensitive to bias present in those predictions.

Fairness [3], [4]: A recent theme in my work examines the fairness implications of optimization under uncertainty, especially in online decision-making. In [3], we prove that learning-augmented algorithms can be particularly vulnerable to small biases in predictions and lead to very unfair outcomes. To overcome this limitation, we design an algorithm that can use biased data to ameliorate its performance while making fair decisions.

Representative Papers:

[1] Learning Augmented Energy Minimization via Speed Scaling (Spotlight at NeurIPS 2020) with E. Bamas, L. Rohwedder, and O. Svensson

[2] The Primal-Dual Method for Learning Augmented Algorithms (Oral at NeurIPS 2020) with E. Bamas, and O. Svensson

[3] Fair Secretaries with Unfair Predictions (to appear at NeurIPS 2024) with E. Balkanski, and W. Ma

[4] Fair and Consistent Correlation Clustering (Under Submission, 2024) with E. Balkanski, and I. Chatzitheodorou

DIVYARTHI MOHAN (Homepage, CV)

Thesis: Simplicity and Optimality in Algorithmic Economics: Multi-Item Auctions and Social Learning ('21)

Advisor: S. Matthew Weinberg, Princeton University

Brief Biography: Divyarthi Mohan is a postdoctoral researcher at Boston University hosted by Prof. Kira Goldner. Previously, she was a postdoc at Tel Aviv University with Prof. Michal Feldman. She obtained her PhD in Computer Science at Princeton University in July 2021 advised by Prof. Matt Weinberg. Divya was awarded the class of 2021 Siebel Scholarship and the Simons-Berkeley Research Fellowship for Fall 2022. During her PhD, she received the School of Engineering and Applied Science's Award for Excellence in 2019 and the Department of Computer Science's Graduate Student Teaching Award in 2018.

Research Summary: My research addresses complex strategic behaviors and informational uncertainties in various settings of algorithmic economics, both by tackling fundamental problems in mechanism design and by developing/studying models for emerging applications and phenomena.

In particular, my expertise is in *multi-dimensional mechanism design* [1,4], where optimal mechanisms are often extremely complex, computationally intractable or even impossible without strong assumptions. My research tackles this through the algorithmic lens of approximation and designs simple, computationally efficient algorithms that are robust to strategic behavior.

My recent work explores settings with interdependencies in agents' values. The celebrated interdependent values model, awarded with the 2020 Nobel Prize in Economics, has been pivotal in studying auctions with more realistic representation of agent valuations, as they crucially rely on private information of others. I design simple algorithms and truthful mechanisms that guarantee a constant approximation to the optimal welfare, under additional informational challenges: namely, private valuation functions [1] or online arrival of agents [2].

In addition, my work in *social learning and strategic communication* investigates how and why (mis)information is propagated due to strategic interactions [3]. I develop and study simple models of communication with the goal of understanding prevalent social phenomena such as herding, polarization and echo chambers.

Representative Papers:

[1] Constant Approximation for Private Interdependent Valuations (FOCS 2023, Highlights Beyond EC 2024)
with A. Eden, M. Feldman, K. Goldner, and S. Mauras

[2] Optimal Stopping with Interdependent Values (EC 2024)
with S. Mauras and R. Reiffenhäuser

[3] Communication with Anecdotes (ITCS 2024)
with N. Haghtalab, N. Immorlica, B. Lucier, and M. Mobius

[4] Approximation Schemes for a Unit-Demand Buyer with Independent Items via Symmetries (FOCS 2019)
with P. Kothari, A. Schvartzman, S. Singla, S.M. Weinberg.

MATHIEU MOLINA (Homepage, CV, Google Scholar)

Thesis: Fairness and Sequential Decision-Making ('25)

Advisor: Vianney Perchet, ENSAE, Patrick Loiseau and Nicolas Gast, Inria

Brief Biography: Mathieu Molina is in the final year of his PhD in applied mathematics at the Inria FAIRPLAY team and ENSAE Paris, supervised by Patrick Loiseau, Vianney Perchet, and Nicolas Gast. He graduated from Mines ParisTech in 2021 and holds a master's degree in Artificial Intelligence, Systems, and Data from PSL University. His research interests lie at the intersection of sequential decision making, and fairness, in problems such as prophet inequalities, multi-armed bandits, auctions, and bipartite matching.

Research Summary: My research explores the impact of various constraints on decision-making in classical algorithmic settings, with a special emphasis on fairness and efficiency.

In my most recent work [1], I study a variant of an i.i.d. prophet inequality where a decision-maker competes with a prophet that selects the average of the two best items, instead of the maximum. We demonstrate that this modification sharply improves the competitive ratio from 0.745 to 0.966, and for the top ℓ items, the competitive ratio approaches 1 exponentially fast as ℓ grows. I am currently extending this work to analyze how the introduction of fairness penalties impacts the competitive ratio in this setting.

A key challenge in online decision-making with constraints is that these constraints may be uncertain and need to be learned over time. In my work [2], I study multi-armed bandits with covering constraints, where each arm must secure a minimum expected reward. We develop algorithms that optimally balance constraint violation and regret by choosing between pessimistic and optimistic constraint estimators. In [3] I work on an online allocation problem with a fairness penalty, but where the decision-maker lacks direct information about protected groups. In this setting, we allow the decision-maker to purchase data of varying quality and cost, and design an algorithm that balances fairness, efficiency, and data acquisition costs.

Additionally, I examine how fairness constraints impact the performance of established mechanisms. In my work on bipartite matching markets [4], I investigate the "price of fairness" in terms of utility loss. We show that under certain fairness constraints related to equality of opportunity, the worst-case utility loss is linear in the number of protected groups but independent of the size of the matching graph.

Representative Papers:

[1] Prophet Inequalities: Competing with the Top ℓ Items is Easy (SODA'25) with N. Gast, P. Loiseau, V. Perchet

[2] Multi-Armed Bandits with Guaranteed Revenue per Arm (AISTATS'24) with D. Baudry, N. Merlis, H. Richard, V. Perchet

[3] Trading-off price for data quality to achieve fair online allocation (NeurIPS'23) with N. Gast, P. Loiseau, V. Perchet

[4] The Price of Fairness in Bipartite Matching (Working paper) with R. Castera, F. Garrido-Lucero, S. Mauras, P. Loiseau, V. Perchet

PAOLA MOSCARIELLO (Homepage, CV)

Thesis: Redistricting with Endogenous Policies (2024)

Advisor: Leeat Yariv, Princeton University

Brief Biography: I am a Microeconomic theorist and a PhD candidate in the Economics Department at Princeton University. I specialize in applying tools from information design and optimal transport to topics in political economy and behavioral economics.

Research Summary: I am interested in a broad range of topics, including gerrymandering, committee decision making, school choice mechanisms, and experimental approaches to decision theory.

In my job market paper, "Redistricting with Endogenous Policies," I examine the interaction between partisan gerrymandering and the policy positions of candidates at the district level. I develop a model where a gerrymanderer partitions voters into equipopulous districts to maximize the expected number of districts won by one party. The key innovation is allowing candidates' positions to depend on the distribution of voters within a district, making voting behavior endogenous to the redistricting process itself. I solve the gerrymanderer's problem using tools from the optimal transport literature, mapping the redistricting problem to a Monge-Kantorovich transport problem. My findings show that optimal districts create a wedge between moderate and extreme opponents, encouraging the emergence of extreme candidates. By diluting the power of moderate voters, optimal redistricting generates a distribution of district representatives that has at least two modes, contributing to increased policy polarization.

My paper "Information Avoidance in School Choice," published in Games and Economic Behavior (2024), investigates how students' concerns about self-image can lead to strategic misreporting of preferences in school choice mechanisms.

I have several ongoing projects, coauthored with colleagues. For instance, my ongoing work, "Reputation in a Committee with Multiple Principals: The Case of the FOMC," examines how career concerns impact individual behavior and collective outcomes in the Federal Open Market Committee. Another example is a work titled "Caution in the Face of Complexity," which explores the interaction between complexity and ambiguity aversion in decision-making.

Representative Papers:

[1] Redistricting with Endogenous Policies (working paper)

[2] Information Avoidance in School Choice (Games and Economic Behavior, 2024)

[3] Caution in the Face of Complexity (work in progress)
with G. de Clippel, P. Ortoleva, and K. Rozen

[4] Reputation in a Commiette with Multiple Principals: The Case of the FOMC (work in progress)
with M. Iaryczower, and G. Lopez Moctezuma

ANIKET MURHEKAR (Homepage, CV, Google Scholar)

Thesis: Fairness, Efficiency, and Incentives in Allocation and ML Problems ('25)

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

Brief Biography: Aniket is a PhD candidate in Department of Computer Science at UIUC, where he is advised by Jugal Garg and Ruta Mehta. He has been a Research Intern at Google Research (2024) and Adobe Research (2022). He is the recepient of the Mavis Future Faculty Fellowship, the Siebel Scholarship, and the IIT Bombay Academic Prize. He holds a B.Tech. in CS from IIT Bombay.

Research Summary: My goal is to *develop algorithmic solutions with provable guarantees on fairness, efficiency, and incentives* for problems of societal and industrial importance. My work uses ideas from economics and social choice to quantify fairness and efficiency, and leverages techniques from algorithm design, economics, and game theory to design solutions. My research has two main directions:

(1) *Investigating fundamental questions regarding the existence and computation of fair and efficient allocations.* In discrete fair division, envy-freeness up to any item (EFX) is regarded as a central notion of fairness. The existence of EFX allocations is one of the most fundamental and enigmatic open problems in fair division. For allocating chores to agents with additive preferences, the existence of EFX allocations is open even for $n = 3$ agents, and the best result in terms of approximation was the existence of $O(n^2)$-EFX allocations. In recent work [1], we prove that 4-EFX allocations of chores always exist, thus showing the first constant-factor approximation of EFX. Another important open problem is the existence of allocations of chores that are both fair (EF1; a relaxation of EFX) and efficient (Pareto-optimal). We showed that such allocations exist and can be efficiently computed for certain structured instances, e.g., for $n = 3$ agents [2].

(2) *Addressing issues of fairness and incentives in machine learning systems*, such as federated learning (FL), by using ideas from game theory and social choice. FL allows agents with individual datasets to collaborate and train a joint model. However, differences in data distributions can lead to unfair and inefficient outcomes, and collusion among agents. Moreover, data-sharing costs may disincentivize agents from sharing their data, leading to free-riding. To address these issues, we used ideas from social choice theory to develop an FL protocol which returns a model that is fair, efficient, and robust to coalitions [3], and designed a mechanism inspired from public goods economics whose Nash equilibria incentivize data-sharing [4].

Representative Papers:

[1] Fair Division of Indivisible Chores via Earning Restricted Equilibria (*under submission*) with J. Garg and J. Qin

[2] Weighted EF1 and PO Allocations with Few Types of Agents or Chores (*IJCAI* '24) with J. Garg and J. Qin

[3] Fair Federated Learning via the Proportional Veto Core (*ICML* '24) with B.R. Chaudhury, Z. Yuan, B. Li, R. Mehta, and A.D. Procaccia

[4] Incentives in Federated Learning: Equilibria, Dynamics, and Mechanisms for Welfare Maximization (*NeurIPS* '23) with Z. Yuan, B.R. Chaudhury, B. Li, and R. Mehta

MARIOS PAPACHRISTOU (Homepage, CV, Google Scholar)

Thesis: From Contagion to Stability: Insights into Network Dynamics, Resilience and Stability (2025)

Advisor: Jon Kleinberg, Cornell University

Brief Biography: I am a fifth-year PhD Candidate at Cornell University advised by Jon Kleinberg. I work on the economics of neteworks exploring their roles within large-scale social and information systems, and understanding their wider societal implications. My research has been supported by an Onassis Scholarship, a LinkedIn Ph.D. Fellowship, a Cornell Fellowship, a grant from the A.G. Leventis Foundation, a grant from the Gerondelis Foundation. I have also spent time in industry, and particularly at Twitter, and Microsoft Research.

Research Summary: In my research, I leverage tools from probability, statistics, algorithms, economics, and machine learning to study how we can make complex network systems more resilient to cascading failures, study decentralized decision-making under privacy, develop statistical network models, and study whether human-like behavior emerges in complex systems simulated by agents powered by LLMs.

Specifically, I am interested in information diffusion and contagion, and how to remediate it. Examples include centralized decision-making algorithms for *remediating network contagion* [1, 2, 3], and decentralized learning and d*ecision-making algorithms to ensure resilience in the presence of privacy risks* [4, 5]. Moreover, I am interested in the *structural characteristics of networks that promote or stop contagion* [6], as well as how information diffuses in complex systems where the agents are black-box models simulated by LLMs and whether *LLMs can simulate human-like complex behavior* [7].

Representative Papers:

[1] Allocating Stimulus Checks in Times of Crisis (WWW 2022)
with J. Kleinberg

[2] Dynamic Interventions for Networked Contagions (WWW 2023)
with J. Kleinberg and S. Banerjee

[3] Optimal Resource Allocation for Remediating Networked Contagions (submitted R&R to Management Science, 2024) with J. Kleinberg and S. Banerjee

[4] Group Decision-Making among Privacy-Aware Agents (under review at Operations Research, 2024) with M. A. Rahimian

[5] Differentially Private Distributed Estimation and Learning (IISE Transactions, 2024) with M. A. Rahimian

[6] Core-periphery Models for Hypergraphs (KDD 2022)
with J. Kleinberg

[7] Network Formation and Dynamics among Multi-LLMs (working paper, 2024)
with Y. Yuan

MANEESHA PAPIREDDYGARI (Homepage, CV)

Thesis: Designing Markets for Information - A Generalized Approach. ('25)

Advisor: Bo Waggoner, University of Colorado, Boulder

Brief Biography: I am a fifth year PhD candidate at CU Boulder advised by Prof. Bo Waggoner and work closely with Prof. Rafael Frongillo. My research encompasses contract theory, prediction markets, designing Automated Market Makers (AMMs) and the economics of blockchain. During my PhD, I have been fortunate to be hosted as an intern by Prof. David Pennock and by Ethereum Foundation, and honoured for [1] to be selected for Highlights Beyond EC 2024. Prior to this, I completed my Masters in Economics at Delhi School of Economics and Bachelor's in Computer Science from IIIT-Hyderabad.

Research Summary: The explosion of interest in collecting data to train large-language models reinforces the need for eliciting more focused information from agents when appropriate, termed *information elicitation*. While its widespread adoption took a back seat due to regulations, my research looks into how seemingly unrelated tools can be used to accomplish elicitation. My research goal is to further develop fundamental insights into fields that surround information elicitation and develop robust theory on Automated Market Makers (AMMs).

Prediction markets, a well-studied field of EconCS, elicit predictions about a future event by enabling trading securities. Our work [1] lays a foundational bridge connecting them to Constant Function Market Makers (CFMMs), a prominent type of AMMs prevalent in the trillion-dollar trading landscape of Decentralized Finance (DeFi). This connection highlights a significant correlation between market-making axioms and desirable information-elicitation axioms. This connection also opens up a rich area for future research, as the literature in both fields can inter-operate and evolve together.

A key innovation in DeFi allows other agents, called Liquidity Providers (LP), to provide liquidity in the market to enable trades in exchange for trading fees. We leverage our equivalence result [1] to introduce a general LP framework to prediction markets in [2] and develop further insights into multidimensional fee.

In the classic principal-agent moral hazard problem, i.e. contract theory, the actions of agents are hidden from the principal. In [3] we show that contracts can be implemented via *Proper Scoring Rules* and this hidden action can be revealed without loss of generality. This is a helpful tool when the principal wishes to learn and better incentivize actions across time periods.

Representative Papers:

[1] An Axiomatic Characterization of CFMMs and Equivalence to Prediction Markets (ITCS 2024) with R. Frongillo and B. Waggoner

[2] A General Theory of Liquidity Provisioning for Automated Market Makers (Working Paper) with A. Bhaskara, and R. Frongillo

[3] Contracts with Information Acquisition, via Scoring Rules (EC 2022) with B. Waggoner

SIDDHARTH PRASAD (Homepage, CV, Google Scholar)

Thesis: Mechanism Design and Integer Programming in the Data Age ('25)

Advisors: Maria-Florina Balcan & Tuomas Sandholm, Carnegie Mellon University

Brief Biography: I am a final-year PhD student in Computer Science at Carnegie Mellon University. My research interests span artificial intelligence, mechanism and market design, machine learning, and operations research, and I have interned at Google Research where I worked on recommender systems with Craig Boutilier. My research has been recognized by a best poster honorable mention at the 2024 Mixed Integer Programming (MIP) workshop, an oral presentation at NeurIPS 2022, and a spotlight at NeurIPS 2021. I received a B.S. in Math and CS from Caltech.

Research Summary: The thesis of my research to date is that high performance—*e.g.,* revenue, social welfare, run-time, memory, *etc.*—in marketplaces can only be fully realized via a synergy of interdisciplinary approaches in mechanism design, integer programming, and machine learning. My goal is to improve computation and economic design for society's various markets. My research spans the full spectrum of new models/concepts, theory, and practical implementation/experiments.

Within mechanism design, I have designed algorithms, modeled new learning paradigms, and invented new mechanism classes for various settings including two-part tariffs, combinatorial auctions, shrinking markets, and general multidimensional mechanism design. A highlight here is the first framework for integrating side information into mechanisms to boost revenue while preserving efficiency and incentives in general multidimensional settings like combinatorial auctions [3].

Within integer programming, I have (1) developed a comprehensive generalization theory for data-driven cutting plane configuration [1, 2] and (2) made foundational (non-learning-based) contributions to the theory and practice of cutting planes [4]. Our generalization theory unveils new mathematical structure in the branch-and-cut algorithm and the canonical class of Gomory cuts [2] and is validated through experiments that show the impact of data-dependent parameter tuning. In a working paper [4] for which I was awarded *Best Poster Honorable Mention* at the 2024 MIP workshop, we propose a new technique for strengthening cover cuts—cutting planes that are critical to solvers like Gurobi—and fix an error in the definitive paper on this topic from 2000. We derive conditions when our new cuts define *facets* of the integer hull, which is the gold standard for cuts, and validate their practical use via experiments. Our cuts deliver strong numerical properties and are currently being tested within FICO Xpress and Cardinal Optimizer.

Representative Papers:

[1] Sample Complexity of Tree Search Configuration: Cutting Planes and Beyond (NeurIPS'21 Spotlight) with M.-F. Balcan, E. Vitercik & T. Sandholm.

[2] Structural Analysis of Branch-and-Cut and the Learnability of Gomory Mixed Integer Cuts (NeurIPS'22 Oral) with M.-F. Balcan, E. Vitercik & T. Sandholm.

[3] Bicriteria Multidimensional Mechanism Design with Side Information (NeurIPS'23) with M.-F. Balcan & T. Sandholm.

[4] New Sequence-Independent Lifting Techniques for Cutting Planes and When They Induce Facets (MIP'24) with E. Vitercik, M.-F. Balcan & T. Sandholm.

NIDHI RATHI (Homepage, CV, Google Scholar)

Thesis: Algorithmic and Hardness Results for Fundamental Fair-Division Problems (2021)

Advisor: Siddharth Barman and Mrinal K. Ghosh, Indian Institute of Science (IISc), Bangalore, India

Brief Biography: Nidhi Rathi is a Lise Meitner postdoctoral research fellow at Max Planck Institute for Informatics (MPI-INF), Saarbrücken, Germany, hosted by Danupon Nanongkai and Kurt Mehlhorn. Nidhi recieved her Ph.D. in Mathematics from IISc, Bangalore, India, advised by Siddharth Barman and Mrinal K. Ghosh. She has received a Commendation Certificate by the CS department, IISc, for her excellent Ph.D. thesis, and prestigious IBM Ph.D. Fellowship and Lise Meitner Postdoctoral Fellowship. Before joining MPI-INF, she was a postdoctoral research fellow at Aarhus University, Denmark, hosted by Ioannis Caragiannis.

Research Summary: My broad research interests lie in the design and analysis of algorithms with a focus on problems inspired by *computational social choice* and *algorithmic game theory*. The central focus of my research is fair division, which explores how to allocate a set of items among agents with varying preferences in a way that all parties consider as *fair*. While multiple hardness results exist for the problem of finding fair/efficient cake divisions (allocating a divisible resource), my work bypasses these computational barriers by [1] identifying the broad class of instances specified by a unifying property of monotone likelihood ratios for which polynomial-time algorithms exist for envy-freeness and various notions of economic efficiency, [2] developing an efficient algorithm with a multiplicative approximation factor of 1/2 (currently, the best known). It is often not possible to achieve fairness and efficiency together and distributions over (deterministic) allocations is a typical way of achieving the existence of such solutions. My work shows that that the above problem belongs to the complexity class of PPAD [3].

Whether EFX allocations exist is a major open problem in fair division of indivisible goods. In recent works [4,5], we propose a potent relaxation of EFX, namely, epistemic EFX, and show that it exists for any number of agents with monotone valuations and can be computed in polynomial time for additive valuations.

In general, I aim to explore different concepts of “fairness” in various theoretical problems in algorithmic design.

Representative Papers:

[1] Fair Cake Division Under Monotone Likelihood Ratios (EC’19 and MOR’22) with Siddharth Barman

[2] Fair and Efficient Cake Division with Connected Pieces (WINE’19) with Siddharth Barman, Eshwar Ram Arunachaleswaran and Rachitesh Kumar

[3] On the Complexity of Pareto-Optimal and Envy-Free Lotteries (AAMAS’24) with Ioannis Caragiannis and Kristoffer Arnsfelt Hansen

[4] New Fairness Concepts for Allocating Indivisible Items (IJCAI’23) with Ioannis Caragiannis, Jugal Garg, Eklavya Sharma, & Giovanna Varricchio

[5] Epistemic EFX Allocations Exist for Monotone Valuations (submitted to AAAI’24) with Hannaneh Akrami

ROJIN REZVAN (Homepage, CV, Google Scholar)

Thesis: Alternate Revenue Benchmarks: Approximation and Computation of simple vs. optimal in Multi-dimensional Bayesian settings (anticipated: 2025)

Advisor: Shuchi Chawla, University of Texas at Austin

Brief Biography: Rojin Rezvan is a final year PhD student at the University of Texas at Austin, advised by Shuchi Chawla. She received her masters degree from the University of Wisconsin-Madison. She is broadly interested in algorithmic game theory and mechanism design. More specifically, she has done research in multi-dimensional mechanism design, fairness in auctions and fair allocation. She is generally interested in the intersection of mechanism design and other fields such as fairness and decentralized systems.

Research Summary: One of the main focuses of my PhD is on the paradigm of "Simple vs. Optimal" in mechanism design for multi-dimensional settings. Multi-item mechanisms can have undesirable properties such as unbounded revenue, *lottery* options in the menu and super-additive pricing function. To circumvent these issues, there are two paths to take: 1) Make some assumptions, such as independence over item value distributions and the buyers' value functions, 2) Examine the validity of the benchmark. The approach we took in [1] and [2] was the latter.

Our proposal is to compare any *simple* mechanism we design to a more realistic benchmark, called "Buy-many". In this setting, it is assumed that each buyer can interact with the menu multiple times. This ensures that super-additive pricing will not happen. The main difference now is while optimal revenue may be unbounded, the gap between revenue of optimal simple mechanisms such as item pricing and optimal buy-many mechanisms is logarithmic in the number of items. In [1], we propose a structure necessary over the item values, with which we will get fine-grained results in terms of approximation and computation of the buy-many revenue via item pricing. In [2], we extend these results and definitions to multi-buyer setting.

I am also interested in algorithmic fairness. In [4], we ask: is it possible that certain allocation algorithms in ad auctions introduce unfairness to the allocations in addition to the data? The answer is yes: an algorithm that always allocates to the highest bidder, such as FPA, could potentially turn minor differences in bids to large differences in allocation. To circumvent the issue, we propose two different algorithms that ensure fairness, while losing a fraction of the optimal social welfare, or consequently revenue. Currently, I am working to extend this work to cases where the advertisers have budgets.

Representative Papers:

[1] Pricing Ordered Items (STOC 22) with S. Chawla, Y. Teng, C. Tzamos

[2] Buy-many Mechanisms for Many Unit-demand Buyers (WINE 23) with S. Chawla, Y. Teng, C. Tzamos

[3] Prophet Secretary Against the Online Optimal (EC 23) with P. Duetting, E. Gergatsouli, Y. Teng, and A. Tsigonias-Dimitriadis

[4] Individually Fair Auctions for Mutli-Slot Sponsored Search (Best student paper at FORC 22) with C. Chawla, N. Sauerberg

XIZHI TAN (Homepage, CV, Google Scholar)

Thesis: Learning-Augmented Mechanism Design ('25)

Advisor: Vasilis Gkatzelis, Drexel University

Brief Biography: Xizhi Tan is a fifth-year PhD student in Computer Science at Drexel University, advised by Prof. Vasilis Gkatzelis. She interned at Google Research during the summers of 2023 and 2024. Her work has received the Exemplary Theory Track Paper Award at EC 2024 as well as the Jay Modi Memorial Award. She was a finalist for the 2023 Meta Research PhD Fellowship.

Research Summary: Worst-case analysis has been the predominant method for mathematically evaluating algorithms in computer science. On the positive side, a worst-case guarantee provides a useful signal regarding the robustness of the algorithm. However, it can often lead to uninformative bounds or impossibility results that may not reflect the real obstacles that arise in practice. This is evident in the rapid progress of machine learning, which has produced highly effective algorithms, many lacking non-trivial worst-case guarantees.

Motivated by this discrepancy, a surge of recent research, known as "algorithms with predictions," aims to develop robust algorithms guided by machine-learned predictions that combine the *robustness* of worst-case guarantees with stronger performance when the predictions are *consistent* with the truth. Our work [1] applies this framework to economic systems and initiates the study of learning-augmented mechanism design in the presence of strategic agents.

We showcase the ubiquitous power of predictions in various societal and economic settings. In [1], we propose a mechanism for the strategic facility location problem that outputs the optimal solution when the prediction is correct, without sacrificing any worst-case guarantees. In the metric distortion problem, our voting rules reach the Pareto frontier of consistency and robustness [2]. Beyond social choice problems, predictions prove valuable in a variety of economic contexts. For example, [3] demonstrates how predictions help achieve improved revenue in online auction settings. In general combinatorial auction settings, our learning-augmented clock auction achieves nearly optimal welfare with correct predictions, which significantly improving upon the pessimistic $\log(n)$ worst-case bound of clock auctions, while still achieving the $\log(n)$ bound even if the prediction is arbitrarily bad [4].

Representative Papers:

[1] Learning-Augmented Mechanism Design: Leveraging Predictions for Facility Location (MOR '23, EC '22)
with P. Agrawal, E. Balkanski, V. Gkatzelis, and T. Ou

[2] Learning-Augmented Metric Distortion via (p, q)-Veto Core (EC '24)
with B. Berger, M. Feldman, and V. Gkatzelis

[3] Online Mechanism Design with Predictions (EC '24 Exemplary Theory Track Paper Award) with E. Balkanski, V. Gkatzelis, and C. Zhu

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

ANISH THILAGAR (Homepage, CV)

Thesis: Practical Guarantees in Forecasting Competitions: Accuracy, Efficiency, and Approximate Truthfulness ('25)

Advisor: Rafael Frongillo and Bo Waggoner, CU Boulder

Brief Biography: I am a 5th year PhD student in the CS Theory Group at CU Boulder. I was previously an undergraduate at Caltech where I double majored in Math and Computer Science and did research in machine learning, quatum computing, and DNA computing. In between, I spent 2 years as a Software Engineer at Google where I helped build and launch AutoML (now Vertex AI).

My main research interest is learning from strategic agents, but I am generally interested in theoretical machine learning, game theory, and mechanism design.

Research Summary: My work has primarily focused on designing and analyzing winner-take-all forecasting competitions with experts (Kaggle, Good Judgement Project, etc.). In these settings, forecasters submit predictions about future events to a mechanism that then chooses a single winner after the event outcomes are realized. It is well known that traditional mechanisms are not truthful and instead give players an incentive to extremize their predictions, so there is no guarantee that the experts report their beliefs or that the chosen winner is actually good.

We show that forecasters can instead have an incentive to behave the opposite way and misreport by hedging their beliefs. However, by analyzing the conditions that drive these contrasting behaviors, we are able to show that the traditional mechanism will be approximately truthful (strongly limiting how much any expert will mis-report) under practically achievable conditions [4].

While some truthful mechanisms for the general setting are known, we show that they all require a large number of events to guarantee the chosen winner is actually good. Instead, we present a class of mechanisms that are approximately truthful (strongly limiting how much any expert will mis-report) but require far fewer events to guarantee a good winner is chosen [1]. Then, we show that standard measures of correlation do not capture the notion that matters in this setting, and instead introduce a new metric that is able to do so. Furthermore, we show that our class of approximately truthful mechanisms are robust to those correlations, the first such guarantee in the literature [3].

Additionally, I have spent some time working on elicitation and loss function design. We found the first consistent polyhedral loss for top-k classification, and show that previously used losses solve other interesting problems [2].

Representative Papers:

[1] Efficient Competitions and Online Learning with Strategic Forecasters (EC'21) with R. Frongillo, R. Gomez, and B. Waggoner

[2] Consistent Polyhedral Surrogates for Top-k Classification and Variants (ICML'22) with J. Finocchiaro, R. Frongillo, and E. Goodwill

[3] Forecasting Competitions with Correlated Events (2023) with R. Frongillo, M. Lladser, and B. Waggoner

[4] A Strategic Analysis of Traditional Forecasting Competitions (2024) with M. Monroe, M. Hsu, and R. Frongillo

ARTEM TSIKIRIDIS (Homepage, CV)

Thesis: Design and Analysis of Auctions: Algorithms and Incentives ('23)

Advisor: Vangelis Markakis, Athens University of Economics and Business

Brief Biography: I am a postdoctoral researcher at Centrum Wiskunde & Informatica (CWI), hosted by Guido Schäfer. Previously, I completed my Ph.D. in Computer Science at Athens University of Economics and Business, where I was advised by Vangelis Markakis.

Research Summary: I am broadly interested in algorithmic game theory, mechanism design, and online algorithms. During my PhD, I focused on the design and analysis of auctions, specifically on protocols that retain provable performance guarantees while being implementable in real-life scenarios. An example is [1], where we studied core-selecting mechanisms, a formalism introduced by Ausubel and Milgrom (2002). Although these auctions are not generally truthful, they offer strong revenue guarantees and align bidder incentives (in a weaker sense). Our contributions include identifying core-selecting mechanisms suitable for practical implementation and designing a prior-free truthful mechanism competitive with the minimum revenue in the core. I've also worked on budget-feasible mechanisms and equilibria in non-truthful auctions.

Recently, I have become interested in mechanism design in environments with predictions, also known as learning-augmented mechanism design. This beyond-worst-case analysis paradigm suggests augmenting algorithm inputs with predictions, leveraging potentially erroneous predictions to improve worst-case performance. The goal is to achieve strong performance when the prediction is perfect (consistency) while also providing guarantees when it's not (robustness). A recent line of work has proposed exploring strategic settings under this framework, a direction I find fascinating. In [2], for example, we study variants of the generalized assignment problem (GAP) in a setting without money. The main result is establishing the best possible consistency-robustness tradeoff for bipartite matching by designing a truthful mechanism that implements Gale-Shapley's deferred acceptance algorithm. Additionally, we design randomized mechanisms for more general GAP variants, achieving improved approximations compared to settings without predictions while maintaining a degree of robustness.

Another area I have worked on is stochastic optimization. In [3], we propose an extension of the Pandora's Box problem (Weitzman, 1979) that incorporates the notion of time in a general sense. For this NP-Hard problem, we provide an efficient constant-factor approximation to the optimal strategy of the decision maker.

Representative Papers:

[1] On Core-Selecting and Core-Competitive Mechanisms for Binary Single Parameter Auctions (WINE 2019) with E. Markakis

[2] To Trust or Not to Trust: Assignment Mechanisms with Predictions in the Private Graph Model (EC 2024) with R. Colini-Baldeschi, S. Klumper, and G. Schäfer

[3] Pandora's Box Problem Over Time (WINE 2024) with G. Amanatidis, F. Fusco and R. Reiffenhäuser

JAMIE TUCKER-FOLTZ (Homepage, CV)

Thesis: Algorithmic Institutional Fairness ('25)

Advisor: Ariel D. Procaccia, Harvard University

Brief Biography: I am a fifth year computer science PhD student at Harvard University. I earned my undergraduate degree in CS and mathematics from Amherst College, and a master's degree in CS from the University of Cambridge, where I studied on a Churchill Scholarship. At Harvard, I have been supported by both an NSF Graduate Research Fellowship and a Google PhD Fellowship.

Research Summary: I am broadly interested in applying ideas from theoretical computer science to improve political and economic institutions. I am particularly focused on algorithms for guaranteeing fairness in complex resource allocation tasks and democratic political processes. I work on a range of problems, either inspired by or directly applicable to real-world issues, drawing primarily upon methodologies from fair division, algorithmic game theory, and algorithmic techniques including combinatorial optimization, Markov chains, and computational geometry.

My first main research thread focuses on adapting benchmarks from the literature on fair division such as *proportionality* and *envy-freeness* beyond the realm of private goods to novel, high-stakes domains. For example, two of my recent papers have studied algorithms for political redistricting [2] and assigning school attendance zones [4], modeled as constrained fair division problems where the "players" are respectively political parties and demographic groups we wish to be fair toward. I have applied a similar axiomatic approach to the problem of randomized apportionment [3], refining the vast space of *ex ante proportional* algorithms by requiring that they are fair and predictable to arbitrary coalitions of parties.

My second research thread concerns graph algorithms for judging fairness in political redistricting, not via intrinsic fairness axioms but, rather, random sampling. If a given map gives some party far fewer congressional seats than a random map would, that is strong evidence it has been gerrymandered—this approach has been successfully used in high-profile litigation in the United States. There are numerous open technical questions regarding what we should mean by "random maps" and how we can efficiently sample them. My contributions include showing that the widely-studied *spanning tree distribution* favors geographically compact maps [5], and establishing the first polynomial-time algorithm to sample from it [1].

Representative Papers:

[1] Sampling Balanced Forests of Grids in Polynomial Time (STOC 2024)
with S. Cannon and W. Pegden

[2] You Can Have Your Cake and Redistrict It Too (EC 2023)
with G. Benadè and A.D. Procaccia

[3] Monotone Randomized Apportionment (EC 2024)
with J. Correa, P. Gölz, U. Schmidt-Kraepelin, and V. Verdugo

[4] School Redistricting: Wiping Unfairness Off the Map (SODA 2024)
with A.D. Procaccia and I. Robinson

[5] Compact Redistricting Plans Have Many Spanning Trees (SODA 2022)
with A.D. Procaccia

MARTIN VAETH (Homepage, CV)

Thesis: Essays in Information and Mechanism Design ('25)

Advisor: Roland Bénabou, Alessandro Lizzeri, and Fedor Sandomirskiy (Princeton University)

Brief Biography: I am currently in the sixth and final year of my PhD in Economics at Princeton University. Previously, I completed a MSc in Economics and Philosophy at the LSE and a BSc in Mathematics at Heidelberg University.

Research Summary: My research uses methods from information and mechanism design to tackle problems in bounded rationality and political economy.

My job market paper *Rational Voter Learning, Issue Alignment, and Polarization* [1] embeds costly information acquisition into a model of electoral competition to explain observed features of voter ideology. I model electoral competition between two parties when voters can learn about their political positions through flexibly acquiring costly information. Optimal learning creates *polarized* and *aligned* political preferences even when the true distribution of ideal points is unimodal and independent across policy issues. That is, voters' revealed positions are bimodally distributed and correlated across diverse issues such as taxation and abortion. When party positions are strategically chosen, voter and party polarization are mutually reinforcing, and both increase as information becomes less costly. This can explain the rise in polarization in recent decades in the US through the advance of information technologies like the internet.

My paper *Attention and Regret* [2] explores a connection between costly information acquisition/attention and emotions. The paper provides an evolutionary explanation for regret as an optimal self-control mechanism to motivate attention and thereby improve decision-making. The model endogenizes the optimal emotions as incentives for an agent who overweights the cost of attention, for example due to temptation or present bias. The optimal emotions turn out to follow the functional form of classical regret theory, which was proposed by Bell (1982) and Loomes and Sugden (1982) to explain behavioral anomalies. Further, the model advances regret theory by explaining why regret is stronger than rejoicing and why regret is stronger in simpler decision problems. Methodologically, the paper combines techniques from costly information acquisition with mechanism design.

In the paper *Imprecision Attenuates Updating* [3], I present a novel comparative statics result for Bayesian updating. Economists frequently model incomplete information through noisy normal signals about a normally distributed state. This signal structure can be used to explain behavioral inertia, as the posterior mean is compressed towards the prior mean, and this attenuation effect is stronger the less precise the signal. Despite the ubiquity of the normal-normal model, it was not known to what extent these properties generalize beyond normal distributions. I show that they generalize to all symmetric log-concave distributions.

Representative Papers:

[1] Rational Voter Learning, Issue Alignment, and Polarization (Under Review)

[2] Attention and Regret (R&R JPE)

[3] Imprecision Attenuates Updating (Working Paper)

GRIGORIS VELEGKAS (Homepage, CV, Google Scholar)

Thesis: Towards Addressing Challenges in Modern ML: Generalization and Responsible AI (2025)

Advisor: Amin Karbasi, Yale University

Brief Biography: I am a final-year PhD student in Computer Science at Yale University, advised by Amin Karbasi. I obtained my BSc and MSc in Electrical Engineering and Computer Science from NTUA. From May 2023 until October 2024 I was an intern at Google Research.

Research Summary: My research focuses on three main directions: **i)** understanding generalization properties of ML algorithms, **ii)** exploring responsible use of ML systems, and designing algorithms with provable replicability guarantees, **iii)** understanding the interaction between ML algorithms and mechanisms.

Modern ML systems are trained using neural networks with parameters vastly exceeding the amount of training data. Conventional theoretical analysis, based on the VC theory, suggests that such systems should suffer from *overfitting*. Thus, this worst-case analysis fails to explain their practical success. In [1] and [2] we derive *beyond* worst-case generalization guarantees, called *universal* rates, for two different learning tasks, interactive learning and regression. Under this notion of learnability, we derive much less pessimistic bounds than the VC theory suggests, that are more closely aligned with empirical observations in deep learning.

Another critical challenge is the replicability issue. The past decade has seen a *replicability crisis* across natural sciences, evident in ML. Many researchers struggle to replicate results from other studies and even their own. Thus, developing algorithms with *provable* replicability guarantees is crucial, ensuring that if executed twice on independent samples, results remain consistent. In [3], we develop replicable algorithms for learning large-margin halfspaces, a fundamental problem in learning theory, with *exponentially* smaller sample complexity than prior work.

Lastly, understanding how learning algorithms interact with economic mechanisms is equally important. Traditional analyses assume that entities interacting with mechanisms are *rational agents*. However, in applications like keyword auctions, *learning algorithms* bid on behalf of advertisers. Thus, it is crucial to study whether these analyses hold in their presence. In [4], we show that Myerson's celebrated auction is not optimal when bids are submitted by learning algorithms, proving that the revenue-optimal auction in this setting needs to be *randomized*.

Representative Papers:

[1] Universal Rates for Interactive Learning (NeurIPS 2022, Oral)
with S. Hanneke, A. Karbasi, and S. Moran

[2] Universal Rates for Regression: Separations between Cut-Off and Absolute Loss (COLT 2024) with I. Attias, S. Hanneke, A. Kalavasis, and A. Karbasi

[3] Replicable Learning of Large-Margin Halfspaces (ICML 2024, Spotlight)
with A. Kalavasis, A. Karbasi, K. G. Larsen, and F. Zhou

[4] Randomized Truthful Auctions with Learning Agents (NeurIPS 2024)
with G. Aggarwal, A. Gupta, and A. Perlroth

JEREMY VOLLEN (Homepage, CV, Google Scholar)

Thesis: Algorithms for Representation-based Fairness in Collective Decisions ('25)

Advisors: Haris Aziz and Toby Walsh, UNSW Sydney

Brief Biography: I am a Ph.D. candidate in my final year at UNSW Sydney. In Summer 2023, I visited Stanford University and was hosted by Ashish Goel. Previously, I received my Bachelor's in Computational Mathematics from Rice University.

Research Summary: My research focuses on the design of collective decision-making systems, with a particular focus on systems which guarantee fair outcomes. The central contributions of my work are two-fold: (1) novel and meaningful fairness definitions, which often take inspiration from social choice theory and fair division, and (2) efficient algorithms which advance the frontier of fair decision-making by provably guaranteeing the defined properties.

In even the simplest voting settings, deterministic algorithms encounter strong impossibilities when pursuing fair decisions. As seen by the universality of the coin toss, randomization is a natural approach to fairness in the face of these obstacles. This approach necessitates meaningful *ex-ante* fairness definitions. In the setting in which a collective must select k alternatives, we propose [1] a new definition which is stronger than all fairness properties known to admit efficient algorithms. We then use flow networks to design efficient algorithms which satisfy our definition in conjunction with other desiderata. One such algorithm computes lotteries over outcomes which additionally satisfy strong fairness guarantees *ex-post*. In [2], we extend this approach, known as "best-of-both-worlds fairness", to participatory budgeting (PB), a direct democratic process that allows participants to decide how to spend a central budget. In addition to designing fair algorithms, we develop a technique to implement arbitrary divisible outcomes by lotteries over discrete PB outcomes while minimizing the variance in budget spent ex-post.

My work also investigates contexts in which the use of techniques from computational social choice is relatively unexplored. One such work [3] studies the ubiquitous problem of centroid clustering. Motivated by scenarios in which data points correspond to individuals, we introduce properties which capture proportional representation in clustering outcomes. Our algorithms uphold the state-of-the-art with respect to existing fairness definitions while also providing proportional representation. In [4], we introduce a framework for PB-like processes for collectives that lack the institutional structure required to pool resources. We explore the extent to which systems can approximate welfare optimality while incentivizing participation.

Representative Papers:

[1] Maximum Flow is Fair: A Network Flow Approach to Committee Voting (EC 2024) with M. Suzuki

[2] Fair Lotteries for Participatory Budgeting (AAAI 2024) with H. Aziz, X. Lu, M. Suzuki, and T. Walsh

[3] Proportionally Representative Clustering (WINE 2024) with H. Aziz, B.E. Lee, and S. Morota Chu

[4] Coordinating Monetary Contributions in Participatory Budgeting (SAGT 2023) with H. Aziz, S. Gujar, M. Padala, and M. Suzuki

YONGZHAO WANG (Homepage, CV, Google Scholar)

Thesis: Multiagent Learning by Iterative Refinement of Game Models ('23)

Advisor: Michael Wellman, University of Michigan

Brief Biography: Yongzhao Wang is a postdoctoral fellow at the Alan Turing Institute in the United Kingdom hosted by Rahul Savani and Theodore Turocy. He is also an honorary research associate at the university of Liverpool, lecturing a second-year undergraduate course "Computer-Based Trading in Financial Markets" in Spring 2024 and Spring 2025. He received her Ph.D. in August 2023 from the CSE Department at the University of Michigan.

Research Summary: My broad and long-term research interests lie in Artificial Intelligence (AI) with a focus on game theory, reinforcement learning, and large language models in the study of multiagent systems. My work has focused on developing scalable and robust automated game-theoretic analysis frameworks for large complex multiagent systems, characterized by a large number of strategies and players, incomplete information across players, or sequential decision-making processes. Specifically, I have explored the following areas:

(1) *Learning in large and complex multiagent systems*: As the multiagent systems become large and complex, traditional game-theoretic analysis often struggles to handle. I integrated various learning methods (e.g., deep RL) in AI with game theory, leveraging AI's capabilities to improve the effectiveness of game-theoretic analysis in tackling these large, complex scenarios.

(2) *Large language models (LLMs) for multiagent systems*: Driven by the success of LLMs, the potential applications of LLMs in combination with game theory has become a promising new area of research. I investigated this combination from two directions: (1) integrating game-theoretic methods and software as modules within LLMs for automated game-theoretic analysis from natural language and (2) embedding LLMs within existing game-theoretic frameworks, such as simulating human behavior for the study of behavioral game theory.

(3) *Applications in cybersecurity and financial markets*: Multiagent systems are universal in the real world from entertainment games like poker to financial markets with millions of traders and cybersecurity scenarios involving attackers and defenders. In the financial domain, I applied game-theoretic analysis to understand various financial operations, market making, and market manipulation. In cybersecurity, I collaborated with domain experts to construct realistic game models and developed automated defense systems using AI and game theory.

Representative Papers:

[1] Market Making with Learned Beta Policies (ICAIF-24) with Rahul Savani, Anri Gu, Chris Mascioli, Theodore Turocy, and Michael Wellman.

[2] A Strategic Analysis of Prepayments in Financial Credit Networks (IJCAI-24) with Hao Zhou, Konstantinos Varsos, Nicholas Bishop, Rahul Savani, Anisoara Calinescu, and Michael Wooldridge.

[3] Evaluating Strategy Exploration in Empirical Game-Theoretic Analysis (AAMAS-22) with Michael Wellman.

MITCHELL WATT (Homepage, CV, Google Scholar)

Thesis: Designing Price Mechanisms for Large Markets

Advisor: Paul Milgrom, Stanford University

Brief Biography: I am a Ph.D. candidate in economics at Stanford University. I hold a Masters in Public Policy from Harvard Kennedy School and a Bachelor of Science in mathematics from the University of Queensland. Outside of academia, I have worked as a consultant at Auctionomics on online display advertising auctions and was a policy adviser and speechwriter for an Australian parliamentarian.

Research Summary: I study microeconomic theory and market design, with a focus on questions relevant to public policy and regulation.

My job market paper [1] and companion paper [2] study the design of in-kind subsidies for redistribution. In [1], we characterize the optimal subsidy mechanism when consumers can "top up" subsidized allocations in a private market. The ability to top up requires the planner to offer subsidies increasing in the consumer's demand for the good, limiting the scope of redistribution via subsidies. When the social planner seeks to redistribute to consumers with lower demand, subsidies are optimal only if lump-sum transfers are unavailable and the cost of public funds is lower than the average weight the planner assigns to consumer surplus, leading to subsidies for consumption up to a maximum level. When the social planner seeks to redistribute to consumers with higher demand, the social planner may prefer in-kind subsidies to lump-sum transfers, providing discounts for consumption beyond a minimum level. The optimal mechanisms have features of food stamps and fare capping programs observed in practice. In [2], we study the case in which consumers participate in either a private market or a subsidized program, but not both. This widens the scope of redistribution for the planner compared to the case with topping up, and the optimal mechanism has three components: a public option, nonlinear subsidies, and laissez-faire consumption.

In [3] and [4], I study Walrasian economies, relaxing in each case one of the standard assumptions of the classical model: convex preferences (in [3]) and price-taking (in [4]). In [3], we introduce Markup equilbrium, an extension of Walrasian equilibrium that adds a markup to the prices that consumers pay to ensure existence even in nonconvex quasilinear economies. Markup equilibria are resource-feasible, incur no budget deficit, and require little more communication and computation than the Walrasian equilibrium. In [4], I study the rate of convergence of price-taking incentives in the Walrasian model, showing that the price impact of misreports is inversely proportional to the number of agents (with high probability) when the expected demand correspondence is strongly monotone.

Representative Papers:

[1] In-Kind Subsidies with Topping Up (Job Market Paper) with Z. Y. Kang

[2] Optimal In-Kind Redistribution with Z. Y. Kang

[3] A Walrasian Mechanism with Markups for Nonconvex Economies (EC 2022, Revise and Resubmit at *Review of Economic Studies*) with P. Milgrom

[4] Strong Monotonicity and Perturbation-Proofness of Walrasian Equilibrium (Best Paper by Young Researcher at Econometric Society Australasian Meeting 2023)

JIBANG WU (Homepage, CV, Google Scholar)

Thesis: Strategic Alignment for AI Systems ('25)

Advisor: Haifeng Xu, University of Chicago

Brief Biography: Jibang is currently a final year PhD student in Computer Science at University of Chicago advised by Prof. Haifeng Xu. His research interest lies at the interface between game theory, learning theory and optimization, with the primary focus on modeling and solving multi-agent decision-making problems under complex, unknown environment. His work has recently received the Stigler center PhD dissertation award.

Research Summary: While intelligent systems are becoming more advanced and influential, their design often overlooks critical incentive structures within their operating environments, risking unintended and potentially harmful consequences. My research aims to advance the design principles and approaches of intelligent systems towards *strategic alignment*, a concept centered on aligning the interests of all stakeholders to achieve mutually beneficial outcomes. Examining the theoretical foundations of machine learning and algorithmic economics, my research branches into two key components of intelligent systems:

- Decision Alignment: *Learning for Strategic Decision-Making.* The outcomes of data-driven decisions can be subject to the strategic responses from stakeholders under conflicting interest and asymmetric information. My work [1,2] models the strategic interactions in the multi-agent decision-making processes and adopts the online learning framework to analyze the adaptive decision optimization problems in a complex, unknown environment.

- Feedback Alignment: *Learning from Strategic Data Sources.* The data empowering machine learning systems can be strategically withheld or manipulated by the data providers. My work [3,4] models data providers' incentives on the learning outcomes and adopts a mechanism design perspective to analyze how different design of statistical methods (e.g., ranking, classification or calibration) or monetary incentives could induce more desirable equilibrium outcomes.

Building upon the two threads above, my ongoing research agenda is focusing on strategic alignment problems in the emerging realm of generative AI. In particular, I am interested in developing practical techniques to 1) build *incentive-aware AI agents with strategic intelligence and rationalizable behaviors*; 2) align the economic incentives of users, model developers and data providers *for more sustainable AI ecosystems.*

Representative Papers:

[1] Markov Persuasion Process and Its Efficient Reinforcement Learning (EC 22)
with Z. Zhang, Z. Feng, Z. Wang, Z. Yang, M. I. Jordan, H. Xu

[2] Robust Stackelberg Equilibria (EC 23)
with J. Gan, M. Han, H. Xu

[3] Auctioning with Strategically Reticent Bidders (WINE 24)
with A. Badanidiyuru, H. Xu

[4] An Isotonic Mechanism for Overlapping Ownership (ongoing)
with H. Xu, Y. Guo, W. Su

BRIAN HU ZHANG (Homepage, CV, Google Scholar)

Thesis: New Solution Concepts, Algorithms, and Applications for Extensive-Form Games: Learning, Correlation, Communication, and Common Knowledge ('25)

Advisor: Tuomas Sandholm, Carnegie Mellon University (CMU)

Brief Biography: I am a final-year PhD student in Computer Science at CMU. I have been honored to receive the *inaugural* CMU Hans J. Berliner Graduate Fellowship in AI (awarded to one student annually) for my research.

Research Summary: My research focuses on algorithms and solution concepts for solving large imperfect-information games. I have developed *new state-of-the-art and provably optimal* algorithms for many zero-sum and general-sum games. Examples include the *first solution concept and efficient algorithm* for hidden-role or social deduction games [1], the *first provably-optimal parameterized algorithm* for optimal (*e.g.,* welfare-maximizing) extensive-form correlated equilibria [2], and the *first efficient no-regret learning algorithms* for achieving the strongest robustness notions known to be efficiently achievable [3].

A foundation that enabled these results (and more!) stems from intrinsic *connections* that I have discovered among problems previously treated as unrelated. For example, the new no-regret algorithms above were made possible in part by exploiting new insights from the seemingly-unrelated problem of generalized mechanism design [3]. These connections are part of a *new framework* I have developed that 1) unifies under a single umbrella a growing range of problems such as principal-agent problems (including mechanism, information, and contract design) and optimal extensive-form correlated equilibria, and 2) *reduces them all to zero-sum games* (two-player and team games, respectively). This allows them to be solved with the rich zero-sum game toolbox—including ML techniques—and motivates further advances to the state of the art in zero-sum games [5].

My research has also led to *practical breakthroughs.* I have created the *first superhuman AI for dark chess* [4], the most complex turn-based game in which superhuman AI has been achieved. Further, my new algorithm and solution concept for hidden-role games led to the *first exact solutions to several variants of the popular game Avalon* [1], the most complex hidden-role game ever solved.

Representative Papers:

[1] Hidden-Role Games: Equilibrium Concepts and Computation (EC'24) with L. Carminati, G. Farina, N. Gatti, & T. Sandholm

[2] Optimal Correlated Equilibria in General-Sum Extensive-Form Games: Fixed-Parameter Algorithms, Hardness, and Two-Sided Column-Generation (EC'22; accepted to appear in Math of OR) with G. Farina, A. Celli, & T. Sandholm

[3] Efficient Φ-Regret Minimization with Low-Degree Swap Deviations in Extensive-Form Games (NeurIPS'24) with I. Anagnostides, G. Farina, & T. Sandholm

[4] Superhuman Performance in Dark Chess via General-Purpose Search Techniques in Imperfect-Information Games (Working paper'24) with T. Sandholm

[5] Computing Optimal Equilibria and Mechanisms via Learning in Zero-Sum Extensive-Form Games (NeurIPS'23) with G. Farina, I. Anagnostides, F. Cacciamani, S. McAleer, A. Haupt, A. Celli, N. Gatti, V. Conitzer, & T. Sandholm

JIAYU (KAMESSI) ZHAO (Homepage, CV)

Thesis: Incentivizing Flexibility in Platform Operations ('25)

Advisor: Daniel Freund, MIT

Brief Biography: I am a final year PhD student at the Operations Research Center at MIT, where I am advised by Prof. Daniel Freund. Prior to my PhD, I graduated Summa Cum Laude from Columbia University in 2020 with a B.S. degree in Operations Research.

Research Summary: My research studies how two-sided service platforms, via market and algorithm designs, can incentivize agents' flexibility to enhance operational efficiency. Such incentives facilitate the matching between supply and demand sides of the market by easing the heterogeneity in space (e.g., Lyft's relocation incentives for drivers), time (Uber's 'wait and save' discount for riders), among others. Motivated by the increasing uses of such flexibility incentives today, my research studies the design of flexibility in platforms, using a combination of tools from stochastic decision-making and game theory.

The first part of my research agenda studies the market design questions around flexible operations. While flexibility incentives are common on both the demand (e.g., 'wait and save' feature at Uber) and the supply side (Ride streak bonuses at Uber) of platforms, they have been treated *in isolation* in the literature and in practice. By modeling how these incentives influence the likelihood of compatibility between agents and the resulting matching size, my work [1] is the first to investigate the management of two-sided flexibility in platforms. Aside from this horizontal interplay of flexibility incentives on different market sides, my research also investigates the vertical supply chain implications of ride-hailing platforms' flexibility decisions [2]. When dual-sourcing autonomous vehicles (AVs) and flexible human drivers with self-scheduling capacity, platforms (e.g., Uber's operations in Phoenix) make dispatch prioritization decisions to fulfill demand through a hybrid fleet, which affects the incentives of AV suppliers and human drivers. I study how potential incentive misalignment can hinder successful AV deployments and provide contracting solutions to overcome them.

The second aspect of my research focuses on algorithms that provide *better customization and timing* to harness flexibility. For instance, booking platforms can adjust their admission control decisions in real-time by considering customers' heterogeneous probabilities of being no-shows (i.e., not requiring service) and their compensation requirements for overbooking. I study an online resource allocation problem that allows overbooking in [3] and propose a policy that improves the additive profit loss guarantee (compared to a clairvoyant) in T periods from $\Omega(\sqrt{T})$ in the literature to $\mathcal{O}(1)$ in our paper.

Representative Papers:

[1] Two-sided flexibility in platforms. (MIT ORC Best Student Paper Award) with D. Freund, and S. Martin

[2] On the supply of autonomous vehicles in platforms (EC'24) with D. Freund, and I. Lobel

[3] Overbooking with bounded loss. (EC'21, Mathematics of Operations Research) with D. Freund

Index