SIGecom Job Market Candidate Profiles 2023

This is the eighth annual collection of profiles of the junior faculty job market candidates of the SIGecom community. The thirty candidates for 2023 are listed alphabetically and indexed by research areas that define the interests of the community. The candidates can be contacted individually or via the moderated mailing list ecom-candidates2023@acm.org.

–Vasilis Gkatzelis and Jason Hartline

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

Contents

Samuel Altmann
empirical market design, auctions, econometrics **9**

Niclas Boehmer
social choice, fairness, matching, computational experiments **10**

Martin Bullinger
algorithms, game theory, computational complexity, social choice **11**

Jakub Černý
game theory, bounded rationality, equilibrium computation **12**

Bhaskar Ray Chaudhury
social choice theory, general equilibrium theory, fairness in ML **13**

Andrew Estornell
algorithmic fairness, strategic classification, adversarial ML, social choice 14

Tomer Ezra
online decision-making, combinatorial auctions, contracts, fair division 15

Alireza Fallah
machine learning theory, game theory and mechanism design, optimization 16

Yiding Feng
mechanism design, online algorithms, revenue management 17

Matheus Venturyne Xavier Ferreira
market design, security, blockchain, game theory 18

Maximilian Fichtl
equilibrium computation, online learning, optimization, auctions 19

Yotam Gafni
strategic machine learning, mechanism design, auctions, fair allocation 20

Sumit Goel
mechanism design, algorithmic game theory, social choice 21

Paul Gölz
democratic innovations, social good, citizens' assemblies, fair division 22

Anand Kalvit
data-driven decisions, online learning, adaptive experiments, platforms 23

Gregory Kehne
social choice, online algorithms, combinatorial optimization 24

Yoav Kolumbus
AGT, AI, online learning, dynamical systems, networks 25

Alexander Lam
mechanism design, facility location, proportional fairness, game theory 26

Majid Mahzoon
game theory, information economics, mechanism design 27

Evi Micha
algorithmic fairness, computational social choice, mechanism design 28

Divyarthi Mohan
mechanism design, opinion dynamics, strategic communication 29

Fedor Sandomirskiy
information economics, mechanism design, fair division algorithms **30**

Daniel Schoepflin
clock auctions, auction theory, obvious strategyproofness, fair division **31**

Sean Sinclair
data-driven decision making, reinforcement learning, fairness **32**

Clayton Thomas
explainable mechanism design, stable matching, market design **33**

Artem Tsikiridis
mechanism design, algorithmic game theory, auctions **34**

Jingyan Wang
fairness, machine learning, crowdsourcing **35**

Xintong Wang
market design, machine learning, game theory, agent-based simulation **36**

Hanrui Zhang
machine learning, mechanism design **37**

Junyao Zhao
optimization, mechanism design, beyond worst-case analysis, complexity **38**

SAMUEL ALTMANN (Homepage, CV)

Thesis: Essays in Empirical Market Design ('23)

Advisor: Ian Crawford, University of Oxford

Brief Biography: I am a DPhil Candidate in the Department of Economics at the University of Oxford, where I am advised by Ian Crawford. In 2018 I completed an MPhil in economics also at Oxford. Before that I graduated with a BSc in Philosophy and Economics from the LSE.

Research Summary: My research examines market design problems from an empirical perspective, with a particular focus on repeated auctions. My work considers problems in which the choice of mechanism depends on the features of agents' preferences. I use structural econometrics, guided by theory, to recover these preferences, before using simulations to evaluate alternate allocation mechanisms.

In [1] I study the allocation of food to food banks, and the allocation problem faced by Feeding America. Feeding America is a charity which distributes enough food to feed 130,000 people each day, across a national network of food banks. Under their current allocation system food banks bid in first-price auctions on truckloads of food using fake money. Previously, Feeding America used a simple queueing system which gave food banks little choice over the food they received. My paper examines how food banks benefit from using the more sophisticated allocation mechanism. The extent to which they benefit depends on the substitutability of different types of food, and the heterogeneity in demand both across food banks and across time. I estimate a dynamic auction model and simulate equilibrium allocations under several alternative mechanisms. I find that the vast majority of food banks are better off under the auction system than under any other system I consider. This is primarily driven by heterogeneity across time because the auction system allows food banks to signal when they need food the most.

In [2] I consider whether, and how, we can recover the values of forward-looking bidders from data on repeated multi-object auctions. I focus on repeated rounds of simultaneous first-price auctions, which are often used in government procurement. I prove that the distribution of values is non-parametrically identified and propose a computationally efficient procedure for estimating the model primitives. I apply the model to study the Michigan Department of Transport's procurement auctions for highway construction contracts. I find that auctioning contracts simultaneously leads to a more efficient allocation than auctioning contracts sequentially.

Representative Papers:

[1] Choice, Welfare, and Market Design: An Empirical Investigation of Feeding America's Choice System (job market paper, EC 2022)

[2] Identification and Estimation of a Dynamic Multi-Object Auction Model (working paper)

NICLAS BOEHMER (Homepage, CV, Research Summary)

Thesis: Bridging the Gap Between Theory and Practice by Using an Enriched Experimental Toolbox: Computational Social Choice in a Dynamic Environment ('23)

Advisor: Rolf Niedermeier (TU Berlin), Markus Brill (University of Warwick)

Brief Biography: I am a third-year PhD student at TU Berlin. Prior to that, I completed a BSc and MSc in Computer Science at RWTH Aachen and Oxford University, respectively finishing best in class. My paper [1] co-authored with Klaus Heeger received the Best Paper and Best Student Paper Award at WINE '20.

Research Summary: I co-authored 30 papers analyzing collective-decision making problems using computational experiments (see, e.g., [2,3]) and tools from complexity theory (see, e.g., [1,4]). A long-term research goal of mine is to strengthen the impact of computational social choice in real-world applications. However, I also look forward to tackling new types of problems from other (related) areas.

Experimental Toolbox. I contributed to enriching the toolbox for conducting computational experiments in social choice in two substantial ways. First, together with an undergraduate assistant, I collected and open-sourced a massive database of real-world preference data (consisting of 7582 elections divided into 25 datasets). Second, in a series of papers (see, e.g., [2]), I co-developed the "map of elections" framework for visualizing sets of elections. This framework has versatile use cases ranging from better understanding the relation between different methods to sample elections over composing a diverse dataset and analyzing the nature of real-world elections to visualizing non-aggregate results of experiments.

Computational Social Choice in a Dynamic Environment. Another major focus of my research lies on incorporating complex and dynamic facets of real-world social choice challenges into collective-decision making problems. First, motivated by changing preferences of agents, I studied the worst-case and average-case robustness of solutions and the problem of minimally adapting solutions to change. For instance, in [3], we found numerous real-world election winners that are very sensitive to random noise in the votes. Second, more broadly, I analyzed different ways of relaxing the classical paradigm of a single preference profile in the input for which a single solution needs to be found.

"Fair" Assignments. Lastly, in some of my works, I incorporated fairness or diversity constraints into different types of assignment problems, e.g., fairly assigning papers to reviewers or partitioning voters into voting districts (see, e.g., [4]).

Representative Papers:

[1] A fine-grained view on stable many-to-one matching problems with lower and upper quotas (WINE '20 and TEAC) with K. Heeger

[2] Expected frequency matrices of elections: Computation, geometry, and preference learning (NeurIPS '22) with R. Bredereck, E. Elkind, P. Faliszewski, and S. Szufa

[3] A quantitative and qualitative analysis of the robustness of (real-world) election winners (EAAMO '22) with R. Bredereck, P. Faliszewski, and R. Niedermeier

[4] The complexity of finding fair many-to-one matchings (ICALP '22) with T. Koana

MARTIN BULLINGER (Homepage, CV, Google Scholar)

Thesis: Desirable Outcomes in Coalition Formation ('23)

Advisor: Felix Brandt, Technical University of Munich

Brief Biography: Martin Bullinger is a Ph.D. candidate in his final year at the Technical University of Munich. He was visiting HPI, Potsdam for a research stay hosted by Tobias Friedrich and Pascal Lenzner, and he will be a visiting research scholar at UC Irvine hosted by Vijay Vazirani in spring. Martin was shortlisted for the best paper award at AAMAS 2021. He holds a B.Sc. and M.Sc. in Mathematics from TUM. During his studies, he received a scholarship by the German Academic Scholarship Foundation.

Research Summary: Martin's thesis focuses on coalition formation in the prominent framework of hedonic games. These consider the question of how to partition a set of agents into coalitions with respect to their preferences. The two main concerns for hedonic games are the representation of games and the definition of desirable outcomes. First, since a player can be part of an exponential number of different coalitions, a succinct representation of their preferences is required. To address this question, various game classes have been introduced during the last 20 years. Second, several concepts of stability and optimality have been defined to determine the desirability of partitions.

The first part of Martin's thesis considers classical solution concepts in classical games, resolving several open questions such as the complexity of contractual Nash stability in additively separable hedonic games [5], individual stability in symmetric fractional hedonic games [3], or strict popularity in roommate games [1]. Martin pioneers the study of Pareto optimality and popularity in hedonic games [1, 4].

The second part of the thesis concerns conceptual contributions. The consideration of majority-based stability concepts offers an appealing compromise between established single-agent stability concepts (AAAI 2022 and [5]). Moreover, the thesis contains the first in-depth study of dynamical stability obtained in a distributed way through myopic decisions of agents gaining novel insights in a broad variety of established settings ([3] and AAAI 2022).

Martin has broad research interests encompassing topics related to computational social choice, game theory, and computational complexity. Alongside his research on coalition formation, Martin has published work on Schelling segregation [2], manipulation in voting (JAIR, 2022), or network creation games (IJCAI, 2022).

Representative Papers:

[1] Finding and Recognizing Popular Coalition Structures (JAIR 2022, AAMAS 2020) with F. Brandt

[2] Welfare Guarantees in Schelling Segregation (JAIR 2021, AAAI 2021) with W. Suksompong and A. A. Voudouris

[3] Reaching Individually Stable Coalition Structures in Hedonic Games (AAAI 2021) with F. Brandt and A. Wilczynski

[4] Pareto-Optimality in Cardinal Hedonic Games (AAMAS 2020)

[5] Boundaries to Single-Agent Stability in Additively Separable Hedonic Games (MFCS 2022)

JAKUB ČERNÝ (Homepage, CV, Google Scholar)

Thesis: Coordination and Commitment in Boundedly Rational Interactions ('23)

Advisors: Bo An, NTU and Allan N. Zhang, A*STAR

Brief Biography: Jakub is a PhD candidate at NTU in Singapore, co-advised by Bo An and Allan N. Zhang. His research is supported by the Singapore international graduate award. During his PhD, Jakub interned for Avast/NortonLifeLock and is visiting UChicago in fall 2022, hosted by Haifeng Xu. Jakub is a recipient of A*STAR merit award and Cisco outstanding thesis award. He holds an MSc in applied mathematics from Charles University and an MSc *summa cum laude* in artificial intelligence from Czech Technical University. Prior to joining NTU, Jakub worked on a US Army Research Lab CRA project with teams from CMU and UTEP.

Research Summary: In my research, I attempt to go beyond the assumption of full rationality of players the traditional game theory is built upon. The implications of deviating from this assumption include the failure of standard theoretical properties, as well as increased computational complexity. For example, a (market) leader may no longer benefit from their commitment power, and the equilibria may not be computable in polynomial time, even in zero-sum games. This hints at the limited robustness of the traditional – entirely rational – solution concepts. My PhD project aims to characterize the consequences of deviating from full rationality domain agnostically and for equilibria beyond Nash. To this end, I restrict myself to the boundedly rational model of quantal response and study its integration into both one-shot or sequential games in three distinct settings.

In [1], we introduce a unilaterally rational concept of Quantal Nash equilibrium, describing a situation when an entirely rational player faces a cognitively restricted opponent. Among others, we prove one of its seemingly counterintuitive properties: a rational player may be at a loss when adopting a quantal Nash strategy instead of a Nash strategy against a quantal player. This motivates a generalization into a setting when the rational player has the ability to commit to a strategy beforehand [2,3]. These so-called Quantal Stackelberg equilibria are able to take advantage of the opponent's subrationality, yielding higher payoffs for the rational player over Nash equilibria. We identify conditions when the optimal strategies are polynomially approximable, yet finding them exactly remains NP-hard in general.

My recent work focuses on multiplayer games with signaling devices, giving rise to Quantal correlated equilibria [4]. Coordinating players via signals is beneficial for the designer, and the concept's hemicontinuity leads to efficient solvability.

Representative Papers:

[1] Complexity and Algorithms for Exploiting Quantal Opponents in Large Two-Player Games (AAAI'21) with D. Milec, V. Lisý, and B. An

[2] Dinkelbach-Type Algorithm for Computing Quantal Stackelberg Equilibrium (IJCAI'20) with V. Lisý, B. Bošanský, and B. An

[3] Computing Quantal Stackelberg Equilibrium in Extensive-Form Games (AAAI'21) with V. Lisý, B. Bošanský, and B. An

[4] Quantal Correlated Equilibrium in Normal Form Games (EC'22) with A. N. Zhang, and B. An

BHASKAR RAY CHAUDHURY (Homepage, CV)

Thesis: Finding Fair and Efficient Allocations ('21)

Advisor: Kurt Mehlhorn, Max Planck Institute for Informatics

Brief Biography: Bhaskar Ray Chaudhury is a Future Faculty Fellow Postdoctoral Research Scholar in the Computer Science department at UIUC. Prior to this, he was a PhD student at the Max Planck Institute for Informatics, supervised by Prof. Dr. Kurt Mehlhorn. He is broadly interested in economics and computation, with a particular emphasis on social choice theory, and its applications in relevant settings in Machine Learning. One of his results in social choice theory, won the Best Paper with a Student Lead Author Award and the Exemplary Paper in the Theory Track Award at ACM EC 2020.

Research Summary: The past two decades have witnessed an explosive integration of fairness and algorithms, i.e., (i) there is a large-scale use of algorithms to address fairness concerns, and (ii) with the advent of several algorithms and data driven procedures to make decisions of large societal impact, fairness has become an integral requirement of most algorithmic paradigms. My research (i) *algorithmically* resolves certain fundamental questions on fairness in *microeconomics* and *social choice theory* and (ii) is among the first to integrate classical fairness concepts from micro-economic theory to more modern avenues of fairness requirement like *fairness aware classification* and *federated learning*.

Fairness in Microeconomics and Social Choice Theory. Resource allocation and task distribution are two fundamental areas in Operations Research and Computer Science that demand fairness. Despite their agelessness and accessibility, these areas harbor several fundamental open problems on the existence and computation of fair allocations. Through a series of algorithmic works ([1-4]), I have resolved some of the important open problems in this area, and established connections to seemingly unrelated theories in computer science like *zero sum combinatorics*.

Integrating Fairness in Microeconomics in ML. Fairness aware classification and federated learning are two settings in ML that strongly resemble public decision making settings in social choice theory. My work ([5]) is among the first to adapt the notions of fairness and efficiency from microeconomics to the aforementioned settings, additionally providing these settings a strong axiomatic foundation.

Representative Papers:

[1] A Little Charity Guarantees Almost Envy-Freeness (SODA 2020, SICOMP 21) with T. Kavitha, K. Mehlhorn, and A. Sgouritsa

[2] EFX Exists for Three Agents (EC 2020, JACM 2022 (minor revision)) with J. Garg and K. Mehlhorn

[3] Polynomial Time Algorithms to Find an Approximate Competitive Equilibrium with Chores (SODA 2022, OR (minor revision)) with S. Boodaghians and R. Mehta

[4] Competitive Allocation of a Mixed Manna (SODA 2021, MOR 2022) with J. Garg, P. Mcglaughlin, and R. Mehta

[5] Fairness in Federated Learning via Core-Stability (NeurIPS 2022) with L. Li, M. Kang, B. Li, and R. Mehta

ANDREW ESTORNELL (Homepage, CV, Google Scholar)

Thesis: Incentives and Consequences of Fair Learning ('23)

Advisors: Yevgeniy Vorobeychik, Washington University in Saint Louis.
Sanmay Das, George Mason University.

Brief Biography: I am currently a fifth year PhD student studying Computer Science at Washington University in St Louis. My interests fall broadly within the fields of Adversarial Machine Learning, Optimization, Game Theory, Strategic Classification, and Algorithmic Fairness. Throughout my time at Washington University I have been fortunate enough to study a wide array of interesting topics within these fields. Looking forward I aim to help improve both the fairness and applicability of intelligent systems which are being deployed in the real world.

Research Summary: Recently my work has focused on Strategic Classification and Algorithmic Fairness with an emphasis on how these two fields relate to one another. Algorithmic Fairness, especially as it related to machine learning, is often studied under the assumption that those who interact with the machine learning model are truthful. However, these systems, especially those deployed in domains where people interact directly with the system such as loan applications, can create incentives for people to behave strategically in-order to achieve more preferential outcomes. Within these types of setting, I am interesting in the unique incentives and dynamics that arise from fair systems, in particular the ways in which strategic behavior can impact the fairness of group-fair machine learning systems [1]. More broadly speaking, I am also interested exploring both the ways such decision-making systems are vulnerable to manipulation in a wide variety of setting, e.g., election control [3], as well as the ways one can prevent manipulation of these systems, e.g., enforcing incentive compatibility in machine learning models [2]. The notions of fairness and strategic behavior are fundamentally intertwined as systems which we desire to be fair are precisely those for which decisions are consequential to humans. Thus, the consideration of strategic behavior is crucial to achieving fairness of decision making systems in practice.

Representative Papers:

[1] Unfairness Despite Awareness: Group-Fair Classification with Strategic Agents (Learning with Strategic Agents Workshop (LSA) and Strategic Machine Learning Workshop (StratML)) with S. Das, Y. Liu, Y. Vorobeychik,

[2] Incentivizing Truthfulness Through Audits in Strategic Classification (Conference on Artificial Intelligence (AAAI)) with S. Das, Y. Vorobeychik,

[3] Election Control by Manipulating Issue Significance (Conference on Uncertainty in Artificial Intelligence (UAI)) with S. Das, E. Elkind, Y. Vorobeychik.

TOMER EZRA (Homepage, CV, Google Scholar)

Thesis: Simple Mechanisms for Limited Information Settings ('22)

Ph.D. Advisor: Michal Feldman, Tel Aviv Univeristy
Postdoc Host: Stefano Leonardi, Sapienza University of Rome

Brief Biography: I am a first-year postdoc at Sapienza University, hosted by Stefano Leonardi. Previously, I was a Ph.D. student at Tel Aviv University and was advised by Michal Feldman, and was awarded the Deutsch Prize of Tel Aviv University for excellency.

Research Summary: I am primarily interested in Algorithmic Game Theory, Mechanism Design, and Optimal Stopping Theory. As part of my Ph.D. and postdoctoral studies, I investigated trade-offs between simplicity and efficiency in various settings where optimal mechanisms are poorly understood, intractable, or cannot be realistically implemented. Furthermore, I am interested in fair division of indivisible goods, welfare approximations, as well as contracts theory in combinatorial settings.

For example, in [1] we extend the principal-agent contract model to combinatorial settings. Our model involves a principal delegating a costly task to a strategic agent. The agent must select a set of actions among potential actions, each of which has a success probability associated with it. The principal's challenge then is to find a rewarding scheme that motivates the agent to increase success probability while keeping costs low. We show that under some assumptions on the structure of the success probability as a function of the set of actions, one can find the optimal rewarding scheme in polynomial time.

I also study problems related to optimal stopping theory. Previous work showed that the classical $\frac{1}{2}$-approximation for the prophet inequality extends to matching in bipartite graphs with 1-sided arrival. Does this guarantee extend to 2-sided arrival, or even better, to general graphs? In [2] we show that while the standard pricing approach does not extend naturally, a different approach can be used to achieve $\frac{1}{2}$-approximation even in the most general scenario. We do so by generalizing the technique of online contention resolution schemes into batched arrival scenarios.

Another domain that I work on is fair division of indivisible goods. For example, in [3] we consider the problem of fair allocation of indivisible goods to agents when the agents have unequal entitlements. Inspired by the fairness notion of the maximin share that is widely used in the equal entitlements case, we present a new share (the AnyPrice share) that is appropriate to settings with arbitrary entitlements. We then devise methods for proving (constructively) the existence of allocations that give all agents a good fraction of this share.

Representative Papers:

[1] Combinatorial Contracts (FOCS 2021)
with P. Dütting, M. Feldman, and T. Kesselheim

[2] Online Stochastic Max-Weight Matching: Prophet Inequality for Vertex and Edge Arrival Models (EC 2020, MOR 2021)
with M. Feldman, N. Gravin, and Z. Gavin Tang

[3] Fair-Share Allocations for Agents with Arbitrary Entitlements (EC 2021)
with M. Babaioff, and U. Feige

ALIREZA FALLAH (Homepage, CV, Google Scholar)

Thesis: Learning with Diverse Data: Adaptation, Collaboration and Conflict ('23)

Advisor: Asuman Ozdaglar, MIT

Brief Biography: Alireza Fallah is a Ph.D. candidate at the department of Electrical Engineering and Computer Science at the Massachusetts Institute of Technology, where he works with Prof. Asuman Ozdaglar and Prof. Daron Acemoglu. He has received a number of awards and fellowships, including the Ernst A. Guillemin Best M.Sc. Thesis Award, Apple Scholars in AI/ML Ph.D. fellowship, MathWorks Engineering Fellowship, and Siebel Scholarship. He has also worked as a research intern at Apple ML privacy team.

Research Summary: My research focuses on developing mathematical tools and frameworks for understanding the challenges we face in deploying machine learning algorithms. These challenges vary from societal impacts of algorithms in their interactions with humans to the robustness certifications to perturbations, model misspecification, and changes in a dynamic environment. Hereunder, I mainly focus on my work on developing privacy mechanisms for data markets.

In [1], we consider a platform's problem of collecting data from privacy-sensitive users to estimate a parameter of interest. We formulate this question as a Bayesian-optimal mechanism design problem, in which an individual can share her data in exchange for a monetary reward, but at the same time has a private heterogeneous privacy cost which we quantify using differential privacy. We consider two popular differential privacy settings: central and local. In both settings, we establish minimax lower bounds for the estimation error and derive (near) optimal estimators for given heterogeneous privacy loss levels for users. Next, we pose the mechanism design problem as the optimal selection of an estimator and payments that elicit truthful reporting of users' privacy sensitivities. We further develop efficient algorithmic mechanisms to solve this problem in both privacy settings. Finally, in a follow-up work [2], we consider a model of utility for users in which their privacy loss is a heterogeneous combination of the local and central privacy losses.

In another work [3], we study a different aspect of the data market design: the optimal choice of architecture from both users' and the platform's points of view. From the users' point of view, we establish that the optimal privacy-preserving mechanism takes a simple "shuffling" form, whereby a user's data is revealed to the platform with some probability. Such a shuffling mechanism provides the best privacy guarantee to users for any given amount of learning about the underlying state. Furthermore, we show that when users gain from learning the underlying parameter increases, interestingly, their overall equilibrium utility could decrease.

Representative Papers:

[1] Optimal and Differentially Private Data Acquisition: Central and Local Mechanisms (Minor Revision, OR & EC 2022), with A. Makhdoumi, A. Malekian, and A. Ozdaglar

[2] Bridging Central and Local Differential Privacy in Data Acquisition Mechanisms (NeurIPS 2022), with A. Makhdoumi, A. Malekian, and A. Ozdaglar

[3] Privacy Costs Of Strategic Data Sharing: Implications Of Shuffling (working paper), with D. Acemoglu, A. Makhdoumi, A. Malekian, and A. Ozdaglar

YIDING FENG (Homepage, CV, Google Scholar)

Thesis: Revelation Gap in Prior-independent Mechanism Design ('21)

Advisor: Jason Hartline, Northwestern University

Brief Biography: Yiding Feng is a postdoctoral researcher at Microsoft Research New England, where he is a member of Economics and Computation group. He previously received his PhD from Department of Computer Science, Northwestern University in 2021 where his advisor was Jason D. Hartline. Before that, he received his BS degree from ACM Honors Class at Shanghai Jiao Tong University.

Research Summary: Thanks to the rapid growth of modern technology, online marketplaces have become an important component of today's economy. In many online marketplaces, the platforms are usually equipped with new algorithmic features and powers (e.g., repeated interaction with users, vast volumes of highly detailed user data), but they also confront additional technical challenges (e.g., the real-time aspect of applications, strategic behavior of users, necessity for fast algorithms, oligopolistic competition). The emergence of such power and challenge has led to new interest in computer science, operations research as well as economics, and enriched the fields with exciting problems.

My current research focuses on examining commonly used algorithms/mechanisms and developing new ones in various online marketplaces (e.g., advertising, ride hailing, cloud computing, vacation rental, video recommendation), taking into account potential concerns or features from *uncertainties* (due to the stochastic nature or real-time aspect of the application), *incentives* (due to the strategic behavior of users), *competition* (across platforms), or *presence of highly detailed user data*. As an interdisciplinary researcher, I approach these marketplaces with a fresh perspective as well as novel ideas from operations, economics, machine learning, and combinatorial optimization. Through collaboration with researchers in both academia and industry, I have primarily addressed problems in *online decision making (e.g., online resource allocation, online learning) [1], mechanism and information design* [2], and *optimization* [3]. While most of my works focus on theoretical research, I am also capable and have experience in data-driven research [3].

Representative Papers:

[1] Batching and Optimal Multi-stage Bipartite Allocations (R&R in Management Science, preliminary version in ITCS 2021) with R. Niazadeh

[2] Revelation Gap for Pricing from Samples (STOC 2021) with J. Hartline and Y. Li

[3] Controlling Epidemic Spread: Reducing Economic Losses with Targeted Closures (Management Science 2022) with J. Birge, and O. Candogan

MATHEUS VENTURYNE XAVIER FERREIRA (Homepage, CV, Google Scholar)

Thesis: Economics and Computation in Decentralized Systems ('22)

Advisor: S. Matthew Weinberg, Princeton University

Brief Biography: Matheus is currently a postdoc at Harvard University hosted by David C. Parkes. They obtained a Ph.D. (2022) and a MA (2018) in Computer Science from Princeton University under the supervision of Matt Weinberg. Previously, Matheus received a BS in Computer Engineering (2016) from the Federal University of Itajubá and spent a year as an exchange student at UC San Diego (2014). At Princeton, Matheus supervised four undergraduate research thesis (resulting in university awards and publications) and was a peer-mentor for first-year graduate students from diverse backgrounds (through GSP), and was an LGBTQIA peer mentor for Whitman College. Matheus is from Itabira, which is known as Brazilian's capital of poetry.

Research Summary: *How can one enforce commitment from online platforms holding a monopoly over transaction data?* In market design, one can enforce commitment via regulation. However, even with publicly available data, it is impossible to prove that a self-interested auctioneer is bidding on their auction with a fake identity. Not surprisingly, users can be skeptical of truthfully revealing their preferences when bidding on truthful auctions.

If one aims to design truthful revenue optimal auctions where the auctioneer cannot profit from cheating, a well-known impossibility result from economic theory states that ascending price auctions are the only alternative. Still, their high communication cost renders those impractical for many applications. My work overcomes this impossibility result by showing how cryptography allows us to design truthful (communication efficient) auctions where bidders need not trust auctioneers.

Decentralized blockchains are a promising tool for designing markets that require minimum commitment power from platforms. An auctioneer can implement an auction as a blockchain smart contract that provably executes the code faithfully. A leader selection mechanism chooses a random leader (i.e., block proposers) to aggregate bids into blocks, but must be unpredictable and unbiased to preserve the security and fairness of underlying applications. My work shows state-of-the-art Proof-of-Stake mechanisms can have their leader selection mechanism biased to a certain extent. Given such limitations, my work proposes a Proof-of-Stake leader selection framework that leverages existing public randomness (i.e., *NIST Randomness Beacon*) and can provide similar fairness guarantees of Proof-of-Work mechanisms but at negligible energy consumption.

Representative Papers:

[1] Credible, Truthful, and Two-Round (Optimal) Auctions via Cryptographic Commitments (EC 2020) with S. M. Weinberg.

[2] Proof of Stake Mining Games with Perfect Randomness (EC 2021) with S. M. Weinberg.

[3] Optimal Strategic Mining Against Cryptographic Self-Selection in Proof-of-Stake (EC 2022) with Y. L. S. Hahn; S. M. Weinberg; C. Yu.

MAXIMILIAN FICHTL (Homepage, CV)

Thesis: Algorithms for Computing Economic Equilibria ('23)

Advisor: Martin Bichler, Technical University of Munich

Brief Biography: Maximilian Fichtl is a computer science Ph.D. candidate at the Technical University of Munich. During the winter of 2021/22, he was a research intern at Meta Core Data Science in the area of economics & computation. He holds a M.Sc. in mathematics from the Technical University of Munich.

Research Summary: During my Ph.D. I have been working on computing Bayes-Nash equilibria in auction games and competitive equilibria in markets with indivisibilities. I mainly use methods from (discrete) convex analysis, linear and nonlinear optimization, and online learning.

In [1] we develop a new algorithm for computing competitive equilibrium prices in the strong-substitutes product-mix auction (SSPMA). The SSPMA is a sealed-bid auction format, allowing bidders to express arbitrary strong-substitutes preferences. While there exist generic algorithms for solving such markets, our algorithm takes advantage of the particular bid language the bidders use to report their preferences.

Articles [2] and [3] study the approximation of Bayes-Nash equilibria in auction games via methods from online learning. In [2], we model the bidders' strategy spaces as neural networks and use a pseudo-gradient method to approximate optimal strategies. In [3], we allow bidders to play mixed strategies and use a more "classical" discretization approach that enables us to apply the dual averaging algorithm and is easier to analyze theoretically. Experimental evidence suggests that both methods find close approximations to equilibria in a broad class of auctions.

In [4] we consider the problem of computing welfare-maximizing core outcomes in assignment markets with budget-constrained buyers. We propose an ascending auction that always terminates in a core outcome and – under additional assumptions – maximizes welfare. In sharp contrast to unconstrained assignment markets, we prove that maximizing welfare among core outcomes is NP-complete.

Representative Papers:

[1] Strong Substitutes: Structural Properties, and a New Algorithm for Competitive Equilibrium Prices. (Mathematical Programming, 2022)
with E. Baldwin, M. Bichler, and P. Klemperer

[2] Learning Equilibria in Symmetric Auction Games Using Artificial Neural Networks. (Nature Machine Intelligence, 2021)
with M. Bichler, S. Heidekrüger, N. Kohring, and P. Sutterer

[3] Computing Bayes Nash Equilibrium Strategies in Auction Games via Gradient Dynamics. (major revision, Operations Research)
with M. Oberlechner, and M. Bichler

[4] Core-Stability in Assignment Markets with Financially Constrained Buyers. (EC '22) with E. Batziou, and M. Bichler

[5] On the Expressiveness of Assignment Messages. (Economics Letters, 2021)

YOTAM GAFNI (Homepage, CV)

Thesis: False-name Attacks and Fair Allocation in Economic Mechanisms ('23)

Advisors: Ron Lavi and Moshe Tennenholtz, Technion

Brief Biography: Yotam Gafni is a direct PhD track in Operations Research at Technion, where he is advised by Ron Lavi and Moshe Tennenholtz. He received his BA in Math and Philosophy from the Hebrew University of Jerusalem.

Research Summary: In auction design, the celebrated VCG-mechanism is both incentive-compatible (IC) and maximizes social welfare, but is no longer IC when bidders may submit auction bids under multiple false identities. This was addressed in previous research mainly by designing alternative mechanisms. However, the conclusion was that efficiency must be sacrificed to regain IC. In [1] we take a Bayesian approach to characterize when, given some probabilistic assumptions on the bidders, VCG may still be IC, even with false identities. The basic intuition is that from the attacker's perspective, facing a stochastic environment, the attack could also go very badly. We find valuation classes and probability distributions where VCG is Bayesian IC. Our result is unique in achieving optimal (rather than approximate) efficiency and avoiding risk aversion (no overbidding) assumptions. In another paper for this setting, we define a novel notion of a loss averse agent. Loss aversion is a known behavioral phenomenon, where people prioritize not losing some baseline guarantee over gaining more in expectation. We formalize it as a robust solution concept, and our surprising result is that in this case VCG achieves optimal efficiency very generally. In [2], we explore the idea of analyzing mechanisms' response to false identities in the political sphere. We analyze power measures for parties' political power such as the Shapley-Shubik and the Banzhaf indices. We obtain multiple upper and lower bounds on how much parties can increase their relative power, given the strategic option to split into multiple entities. In [3], we consider false-name attacks on collaborative machine-learning frameworks. Collaborative machine-learning (federated learning being a notable example) is when machine-learning is done in a decentralized fashion. For example, this can be necessary because different stakeholders have different parts of the data required to train the model. We formalize protocols of sharing data over time, and show how using false-name attacks, an agent that is part of this system can learn the shared machine-learning model, while sharing fake data, resulting in the model used by others performing badly. Different protocols and algorithms have different guarantees to be false-name proof and these guarantees also depend on how many false-name identities an agent can generate, which relates to how strictly the platform verifies agents' identities.

Representative Papers:

[1] VCG under Sybil (False-name) Attacks — a Bayesian Analysis (AAAI 2020) with R. Lavi, M. Tennenholtz

[2] Worst-case Bounds on Power vs. Proportion in Weighted Voting Games with Application to False-name Manipulation (IJCAI 2021; JAIR) with R. Lavi, M. Tennenholtz

[3] Long-term Data Sharing under Exclusivity Attacks (EC 2022) with M. Tennenholtz

SUMIT GOEL (Homepage, CV, Google Scholar)

Thesis: Essays in Mechanism Design ('23)

Advisor: Federico Echenique, U.C. Berkeley

Brief Biography: Sumit Goel is a Ph.D. candidate in Economics (with a minor in Computer Science) at the California Institute of Technology advised by Federico Echenique. His research interests are in mechanism design and algorithmic game theory. Before graduate school, he completed a M.S. in Quantitative Economics from Indian Statistical Institute (Delhi) where he secured top rank in the program. He also holds a B. Tech. in Computer Science from Delhi Technological University.

Research Summary: My research interests lie broadly at the intersection of economics and computer science and in particular, I've worked on problems in mechanism design [1,2], market design [3], and contest theory [4].

In [1], we consider a facility location problem in two dimensions and show that for a large class of social cost functions (including p-norm and weighted utilitarian), the coordinate-wise median mechanism has the smallest worst-case approximation ratio among strategyproof mechanisms. We also find this ratio exactly for the minisum objective and obtain bounds for the general p-norm objective. The paper augments the previous literature that provides strong axiomatic foundations for the coordinate-wise median mechanism by demonstrating its quantitative optimality.

In [2], we consider a principal-agent project selection problem with asymmetric information. Under a natural partial verifiability constraint of no-overselling, we characterize the class of strategyproof mechanisms and find that the optimal mechanism has a simple cutoff structure. We further find that the principal gives more authority to an agent with more aligned preferences (ally principle).

In [3], we study exchange economies with indivisible goods and identify different sufficient conditions under which core allocations exist. We also propose an algorithm for finding stable allocations in these economies and show that the algorithm, when applied to a housing market, mimics the Top Trading Cycle (TTC) algorithm.

In [4], I study how competition influences the effort invested by agents in contests with incomplete information. In contrast to the complete information case, I find that a more competitive prize vector leads to a greater expected effort. I discuss the implications of this result for the design of grading contests. Assuming the value of a grade is determined by the information it reveals about the agents ability, I find that more informative grading schemes lead to more competitive prize vectors and thus induce higher effort in expectation.

Representative Papers:

[1] Optimality of the Coordinate-Wise Median Mechanism for Strategyproof Facility Location in Two Dimensions (SAGT 2022, Forthcoming at SCW) with W. Hann-Caruthers

[2] Project Selection with Partially Verifiable Information (WINE 2022) with W. Hann-Caruthers

[3] Stable Allocations in Discrete Economies (arXiv:2202.04706) with F. Echenique and S. Lee

[4] Prizes and Effort in Contests with Private Information (arXiv:2205.05207)

PAUL GÖLZ (Homepage, CV, Google Scholar)

Thesis: Social Choice for Social Good: Proposals for Democratic Innovation from Computer Science ('22)

Advisor: Ariel Procaccia, Carnegie Mellon University & Harvard University

Brief Biography: I'm a postdoc in Harvard's EconCS group and just received my Ph.D. in computer science from CMU. I co-founded and now co-lead the MD4SG working group on Civic Participation.

Research Summary: The question underlying my work is *how to coordinate individuals so that they can make decisions as a group?* To enable such collective action, I design effective, fair, and trustworthy decision-making processes, for which I draw on techniques from areas including computational social choice, fair division, optimization, mechanism design, smoothed analysis, and stochastic processes. My work approaches the design of decision-making processes in two domains: *(1) democratic innovations* and *(2) resource allocation*:

(1) Democratic Innovations. To turn the tide on a receding global democracy, a coalition of academics and practitioners are exploring *democratic innovations*: alternative institutions for decision-making that more directly involve constituents. Some of these innovations have met with resounding success in practice, and stand a good chances of becoming integral parts of our democracies.

My work contributes mathematical and algorithmic perspectives to the exploration of democratic innovations, often in collaboration with practitioners. My work on has particularly focused on *citizens' assemblies/sortition*, and one of the algorithms I designed is now routinely used to select citizens' assemblies in practice [1]. Other democratic innovations I have studied include *liquid democracy* and *apportionment* (e.g., [3]).

(2) Resource Allocation. In resource allocation problems, the group controls a set of scarce resources, and must jointly decide which individuals these resources should go to. I have studied a wide range of such settings: in *kidney exchange*, the resources are donor organs; in *refugee resettlement*, local capacity for hosting refugees (e.g., [2]); in *fair division*, the goods might be physical items in an estate; and in *fair machine learning*, slots in a graduate school program [4]. Each resource determines which properties a decision process must satisfy: the processes I design allocate resources where they maximize social welfare, satisfy fairness guarantees for individuals or demographic groups, or prevent incentives for manipulation. A major US resettlement agency now uses an algorithm from my work [2] to allocate refugees.

Representative Papers:

[1] Fair Algorithms for Selecting Citizens' Assemblies (*Nature* 2021) with B. Flanigan, A. Gupta, B. Hennig, and A.D. Procaccia

[2] Dynamic Placement in Refugee Resettlement (EC'21, minor revision @ *Operations Research*) with N. Ahani, A.D. Procaccia, A. Teytelboym, and A.C. Trapp.

[3] In This Apportionment Lottery, the House Always Wins (EC'22) with D. Peters and A.D. Procaccia

[4] Paradoxes in Fair Machine Learning (NeurIPS'19) with A. Kahng and A.D. Procaccia

ANAND KALVIT (Homepage, CV)

Thesis: Bandits and Concentration: Theory and Application to Matching Markets ('23)

Advisor: Assaf Zeevi, Columbia University

Brief Biography: I am a doctoral fellow at the Decision, Risk, and Operations division at the Graduate School of Business, Columbia University. Previously, I received my Bachelor's and Master's degrees in Electrical Engineering from IIT Bombay, India.

Research Summary: My research sits at the intersection of machine learning and operations management. Broadly, I am interested in sequential decision making under parameter uncertainty in modern operational settings. My work in this area has been anchored around: (i) developing smart and scalable algorithmic solutions for revenue management problems in complex economic environments (e.g., matching markets); and (ii) understanding the fundamental limits of achievable performance in stylized learning problems that can help guide the design of operationally efficient and scalable policies. Forthcoming paragraphs briefly elucidate these aspects.

In [3], we revisit the theoretical underpinnings of the canonical multi-armed bandit paradigm, and provide new results on the arm-sampling behavior of popular bandit policies leading to several important insights. Among these, it is shown that arm-sampling rates under the celebrated optimism-based UCB1 policy are asymptotically deterministic, regardless of the problem complexity. This discovery facilitates new sharp asymptotics for UCB1, and furthermore, lends itself to the first complete process-level characterization of the multi-armed bandit problem under UCB1 in the conventional *diffusion limit*. This, among other things, reveals profound distinctions between properties of UCB1 and the classical Thompson Sampling policy such as an *incomplete learning* phenomenon characteristic of the latter.

In my job market paper [1], we develop algorithm design principles for a prototypical matching problem faced by large centralized platforms where *jobs* must dynamically be matched to *workers* subject to uncertainty about arrivals, preferences, skills and population-level distributions thereof. Our work provides the first systematic treatment of the *choice overload* phenomenon pervading such marketplaces (and similar settings like server farms). We establish lower bounds on achievable performance and propose novel rate-optimal policies that adapt to aforementioned primitives via online learning; results have design implications more broadly.

Representative Papers:

[1] Dynamic Learning in Large Matching Markets (NeurIPS 2022) with A. Zeevi.

[2] Bandits with Dynamic Arm-acquisition Costs (Allerton Conference on Communication, Control, and Computing, 2022) with A. Zeevi.

[3] A Closer Look at the Worst-case Behavior of Multi-armed Bandit Algorithms (NeurIPS 2021) with A. Zeevi.

[4] From Finite to Countable-armed Bandits (NeurIPS 2020) with A. Zeevi.

GREGORY KEHNE (Homepage, CV, Google Scholar)

Thesis: Algorithms for Social Epistemology and Social Choice ('23)

Advisor: Ariel Procaccia, Harvard University

Brief Biography: Gregory Kehne is a Ph.D. student in the Computer Science Department at Harvard University, advised by Ariel Procaccia, and a recent recipient of the 2023 Siebel Scholarship. Prior to joining the Harvard EconCS group, he began his Ph.D. in the Algorithms, Combinatorics, and Optimization (ACO) program at Carnegie Mellon University. Greg holds an undergraduate degree in mathematics from Williams College, where his thesis received highest honors in 2016.

Research Summary: One of my research areas is in online algorithms. For many combinatorial optimization problems the online adversarial model admits strong information-theoretic lower bounds, and online algorithms are markedly outperformed by their offline counterparts. What aspects of the adversarial model entail this worse performance, and are they realistic? For covering integer programs, we find in [1] that it is both the unknown instance and the adversarial online order which make online covering hard. If the constraints of an unknown covering problem instead arrive in random order, then it is possible to essentially match the best offline guarantee. We are continuing to study this and other weakened online adversaries for related combinatorial optimization and network design problems.

I also work on problems in social choice. The emerging popularity of direct democratic paradigms like sortition has heightened the need for a computational understanding of collective decision-making in new formats. Sortition entails randomly selecting a subset of a population in such a way that the subset is representative of the population, while providing equitable likelihood of selection to individuals. In [2] we study an established model of sortiton, wherein representative subsets satisfy intersecting demographic quotas, and we determine the extent to which representativeness guarantees, equity of individual selection probability, and the simplicity of the random process trade off. We will present an alternative model of sortition at NeurIPS '22, where the representativeness of a subset is instead a continuous objective. We explore the extent to which representativeness and selection equity trade off, and identify selection algorithms on the Pareto frontier.

There are also increasing efforts to transition traditional forms of decision-making and preference aggregation into the context of online platforms and communities. In response, other ongoing work extends approval-based multi-winner voting to larger scales and an online setting similar to that of [1]. In [3] we also advance the implicit utilitarian model of voting, deriving new upper and lower bounds on the optimal welfare regret of randomized voting rules.

Representative Papers:

[1] Random Order Set Cover is as Easy as Offline (FOCS 2021)
with A. Gupta and R. Levin

[2] Fair Sortition Made Transparent (NeurIPS 2021)
with B. Flanigan and A. Procaccia

[3] Worst-Case Voting When the Stakes are High (AAAI 2022)
with A. Kahng

YOAV KOLUMBUS (Homepage, CV, Google Scholar)

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

Advisor: Noam Nisan, The Hebrew University of Jerusalem

Brief Biography: Yoav Kolumbus is a PhD student in computer science at the Hebrew University of Jerusalem, where he is advised by Noam Nisan. He holds a M.Sc. and B.Sc. in Physics from the Hebrew University. Yoav is also a professionally trained musician and holds a Bachelor's degree in classical and Jazz performance on Double-bass from the Jerusalem Academy of Music and Dance.

Research Summary: Since the advent of the internet, and increasingly in recent years, many of our social and economic interactions have transformed to online systems which have significant impact on society. AGT has grown to be the central paradigm for analyzing and designing interactions in these online computational systems. However, with the rapidly increasing human activity invested in online interactions and the development of AI agents that interact with human users, many systems have evolved highly complex dynamics that are today poorly understood, and which often give rise to behaviors very different from those anticipated by their designers. The analysis and design of these systems require new tools and new concepts that will account for their hybrid interactions (combining humans and AI), their dynamic nature, their interconnectedness, and their scale. My research tackles questions of modeling and analysis in such systems, which include networked environments and interactions between strategic play and learning algorithms.

Consider for example the case of online auctions, which generate the bulk of revenues of internet giants like Google or Microsoft, and are, in a sense, the "economic engine" of the web. While economic theory provides us good understanding of traditional auctions, these ad-auctions typically occur at high frequencies which do not allow users (the advertisers) to "manually" bid in each auction. Instead, software "auto-bidding" agents are used to bid on the advertisers' behalf. These agents interact, learn from experience, and aim to optimize bids for their users.

In [1] we showed how dynamics in such systems can have very different outcomes from those anticipated by traditional auction theory, and how truthful auctions can lose their truthfulness property while non-truthful auctions can become truthful. In a companion paper [2], we formulated a general framework for analyzing interactions between users of learning agents in repeated games and studied emergent phenomena in several game classes.

Additional topics I have studied at the interface of AI, AGT, dynamics, and networks include mechanism design for investment networks [3], learning and prediction of human play in strategic interactions (IJCAI 2019), formal explanations of agent behavior in RL (NeurIPS 2022), and intervention policies for network contagion processes (Scientific Reports 2021).

Representative Papers:

[1] Auctions between Regret-Minimizing Agents (WWW 2022) with N. Nisan

[2] How and Why to Manipulate Your Own Agent: Modeling Games between Users of Learning Agents (NeurIPS 2022) with N. Nisan

[3] Optimal Collaterals in Multi-Enterprise Investment Networks (WWW 2022) with M. Babaioff and E. Winter

ALEXANDER LAM (Homepage, CV, Google Scholar)

Thesis: Strategic Behaviour and Proportional Fairness in Facility Location Games ('23)

Advisors: Toby Walsh and Haris Aziz, UNSW Sydney

Brief Biography: Alexander Lam is a final year computer science PhD candidate at UNSW Sydney, supervised by Haris Aziz and Toby Walsh. He was a research assistant at the City University of Hong Kong, working under Minming Li, and was a visiting research student at Hong Kong Polytechnic University, under the supervision of Bo Li. He has also interned with Data61. Previously, he received a Bachelor's Degree in Advanced Mathematics with First Class Honours from UNSW Sydney.

Research Summary: In the classic facility location problem, we are tasked with finding the optimal placement of a facility, such that its distance from the locations of a set of agents is minimized. This problem readily translates to many settings such as the geographical placement of schools or libraries, social choice with single-peaked preferences, and participatory budgeting. In my research, I aim to design strategyproof mechanisms which capture proportional fairness concerns for exogenous groups of agents at or near the same location. My research particularly draws from approximate mechanism design, game theory, and probability theory.

In [1], we formulate a hierarchy of proportional fairness axioms which provide, for each group of agents, a maximum distance guarantee which corresponds to the size of the group. We prove several characterization results of strategyproof mechanisms satisfying these various fairness axioms, and identify one of these mechanisms as the unique Nash equilibrium outcome for any mechanism satisfying natural fairness and monotonicity properties. Finally, we compute the welfare approximation ratios for certain strategyproof and proportionally fair mechanisms.

We also propose a proportional fairness axiom that is stronger than all of the main fairness notions discussed in [1] and cannot be satisfied by any strategyproof, deterministic mechanism. We therefore turn to randomized mechanisms in [2], and find a unique characterization of mechanisms satisfying this stronger proportional fairness axiom (in expectation), universal truthfulness and universal anonymity.

Finally, in [3], we apply our proportional fairness hierarchy in the context of the obnoxious facility location problem, in which agents instead wish to be located far away from the facility. We compute the price of fairness in both the deterministic and randomized setting, as well as the price of anarchy of welfare-optimal proportionally fair mechanisms in the former setting where strategyproofness is incompatible with proportional fairness.

Representative Papers:

[1] Strategyproof and Proportionally Fair Facility Location (WINE'22)
with H. Aziz, B. Lee and T. Walsh

[2] Random Rank: The One and Only Strategyproof and Proportionally Fair Randomized Facility Location Mechanism (NeurIPS'22)
with H. Aziz, M. Suzuki and T. Walsh

[3] Proportionally Fair Obnoxious Facility Location (Manuscript)
with H. Aziz, B. Li, F. Ramezani and T. Walsh

MAJID MAHZOON (Homepage, CV, Google Scholar)

Thesis: Essays on Network Economics ('23)

Advisors: Ali Shourideh and Maryam Saeedi, Carnegie Mellon University

Brief Biography: Majid is a PhD candidate in Economics at Tepper School of Business, Carnegie Mellon University. Previously, he completed his M.Sc. in Electrical and Computer Engineering at Carnegie Mellon University in 2017. He obtained his B.Sc. in Electrical Engineering from Sharif University of Technology, Iran in 2013.

Research Summary: My primary research is in the area of Game Theory and Information Economics, where I explore incentives for information transmission in organizations. I investigate how organizations' hierarchy configuration interacts with information transmission and the implications of information design for the pay-for-performance schemes used in organizations. I am also interested in the empirical study of networked markets motivated by Industrial Organization applications. I investigate these markets, the implications of different network setups for the market outcomes, and the regulations required to increase welfare and efficiency.

In my job market paper [1], I study intermediated communication in hierarchical organizations, such as a firm or the government, where information is transmitted through a chain of agents to a decision maker. I prove that we can focus on a simple class of equilibria where the only one who conceals information is the initial sender. When facing a binary decision, regardless of the number of intermediaries, at most four pivotal agents determine the amount of information communicated to the decision maker. In this case, my results underscore the importance of vice presidents in hierarchical organizations.

In a joint work with Ali Shourideh and Ariel Zetlin-Jones [2], we study moral hazard by combining the traditional principal-agent model with information design. We investigate the resulting contract when the agent controls principal's information about the output of agent's effort. This can be interpreted as an employee choosing the performance measure based on which to be paid. We fully characterize the optimal information structure from an agent's perspective in a general moral hazard setting, possibly with multiple dimensions for efforts and performance. When effort and performance are one-dimensional, under a general class of models, binary signals in the form of threshold signals are optimal.

On the empirical side, in a joint work with Ali Shourideh and Maryam Saeedi [3], we study networked markets using data on natural gas flows through US interstate pipelines. Our results illustrate the significant effect of mergers on prices even if the merging entities are not in the same physical market. Moreover, we identify network effects showing that a regional shock can have a significant effect on prices even in the farthest regions.

Representative Papers:

[1] Hierarchical Bayesian Persuasion: Importance of Vice Presidents (arXiv)

[2] Agent Optimal Learning in Moral Hazard (working paper) with A. Shourideh and A. Zetlin-Jones

[3] Pricing and Mergers in Complex Networks: The Case of Natural Gas Pipelines (work in progress) with A. Shourideh and M. Saeedi

EVI MICHA (Homepage, CV)

Thesis: Fair and Efficient Social Decision Making ('23)

Advisor: Nisarg Shah, University of Toronto

Brief Biography: Evi is a fifth-year Ph.D. student in Computer Science at the University of Toronto, advised by Nisarg Shah. She is also affiliated with the Vector Institute for Artificial Intelligence and the Schwartz Reisman Institute for Technology and Society. In Fall 2021, Evi was a visitor at Harvard University, hosted by Ariel D. Procaccia, and in Summer 2022, she was a visitor at UNSW Sydney, hosted by Haris Aziz. She was awarded Faculty of Arts and Science's General Motors Women in Science and Mathematics Award (2022).

Research Summary: My research lies at the intersection of Economics and Computation (EconCS), and spans areas such as *algorithmic fairness* and *computational social choice*. Broadly, one thread of my research discovers novel EconCS paradigms for understanding interactions between agents in multiagent systems, while another thread draws on the EconCS literature to craft and provably guarantee meaningful fairness definitions for domains such as voting, matching, and machine learning.

In the former thread, my research primarily focuses on collective decision-making by groups of users via voting. In [1], we study *sortition*, which is used worldwide to form citizens' assemblies by random selection of a panel from the population. We propose the first quantitative measure of representativeness of the panel, and characterize its tradeoff with a fairness property requiring each individual to have the same chance of being on the panel. My work [2] has also analyzed the efficiency of *liquid democracy*, a revolutionary democratic paradigm that combines the advantages of direct and representative democracies by allowing users to transitively delegate their voting power to other users whenever they wish to.

In the latter thread, my research focuses on developing a theory of algorithmic fairness for diverse applications. In [3], we borrow the idea of Nash welfare from economics to design low-regret fair algorithms for a multi-agent version of the classical *multi-armed bandit* problem, in which pulling an arm generates different stochastic rewards to different agents. Nash welfare serves as a middle-ground solution between utilitarian and egalitarian welfare. Another work of mine [4] borrows the idea of proportional fairness from economics to design algorithms for *clustering* that achieve a constant approximation to the fairness guarantee. This guarantee demands that if k cluster centers are placed given n data points in a metric space, no group of at least n/k points admit a new cluster center that they all prefer.

Representative Papers:

[1] Is Sortition Both Representative and Fair? (NeurIPS 2022)
with S. Ebadian, G. Kehne, A. D. Procaccia, and N. Shah.

[2] A Contribution to the Critique of Liquid Democracy (IJCAI 2019)
with I. Caragiannis

[3] Fair Algorithms for Multi-Agent Multi-Armed Bandits (NeurIPS 2021)
with S. Hossain, and N. Shah

[4] Proportionally Fair Clustering Revisited (ICALP 2020)
with N. Shah

DIVYARTHI MOHAN (Homepage, CV, Google Scholar)

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

Advisor: Matt Weinberg, Princeton University

Brief Biography: I am a Postdoctoral Fellow at Tel Aviv University's EconCS lab hosted by Michal Feldman. I received my PhD in Computer Science from Princeton University in 2021. I was awarded the Simons-Berkeley Research Fellowship for Fall 2022, and the class of 2021 Siebel Scholarship. During my Phd, I 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: The main focus of my research is in mechanism design beyond single-dimensional settings, where optimal mechanisms are often complex and computationally intractable, and my work tackles this through the lens of approximations. In [1], I study the fundamental problem of selling n heterogeneous items to a unit-demand buyer and provide a mechanism that gets *near-optimal* revenue in time quasipolynomial in n. Notably this the *first sub-exponential approximation scheme* for a general class of valuations. In [2], I consider *rich ad auctions* where the advertisers have a multi-dimensional type and the platform faces a knapsack constraint. While the VCG mechanism can obtain the optimal social welfare truthfully, it is computationally inefficient and impractical. I design a simple, greedy-like, polytime truthful mechanism that achieves a 3-approximation to the optimal welfare. More recently, in [3] I study mechanism design for the public projects problem when the agents have *interdependent* valuations.

In addition, I study how information/misinformation propagates due to strategic interactions, using simple and natural models that help understand social phenomena such as information cascades, echo chambers, and homophily. My work [4], considers a communication game between two agents, where the sender shares one of the signals she observed (an anecdote) instead of sending arbitrary messages. I study the equilibrium effects induced by the sender's temptation to influence the receiver with *biased anecdotes*, and investigate how this can give rise to homophily. In [BIMW20], I studied the effects of network structure on information aggregation when agents interact by changing their opinion to match the majority of their neighbours. For example, I show that an information cascade is avoided when the network is a tree formed through preferential attachment.

Representative Papers:

[1] Approximation Schemes for a Unit-Demand Buyer with Independent Items via Symmetries (FOCS 19)
with P. Kothari, A. Schvartzman, S. Singla, and S.M. Weinberg

[2] Simple Mechanisms for Welfare Maximization in Rich Advertising Auctions (NeurIPS 22) with G. Aggarwal, K. Bhawalkar, A. Mehta, and A. Psomas

[3] Interdependent Public Projects (SODA 23)
with A. Cohen, M. Feldman, and I. Talgam-Cohen

[4] Communicating with Anecdotes
with N. Haghtalab, N. Immorlica, B. Lucier, and M. Mobius

FEDOR SANDOMIRSKIY (Homepage, CV, Google Scholar)

Thesis: The Value of Information in Repeated Games ('14)

Advisor: Victor Domansky, Ernst Presman

Brief Biography: I am a postdoc at Caltech. Before, I was a postdoc at the Technion game theory group, a member of the Mechanism Design for Data Science group, and worked at the Game Theory Lab of HSE University. I've served on the program committee of EC 2020-22. My research has received EC 2020 best paper award and has been published by top economic & OR journals and conferences.

Research Summary: I am a microeconomic theorist specializing in *information economics* and *economic design*. My research brings new methodological tools to understand interactions among strategic agents, and improve the outcomes of these interactions by proposing new mechanisms. It blends microeconomic insights with ideas from algorithmic game theory and relies on the interplay of probability, convexity, and functional analysis. I am especially interested in the following topics.

A) *Information economics.* Information economics studies how agents' behavior is affected by the information they have access to. My key interest is in methods for studying problems where different agents acquire information from *different sources.* My main contribution to this field [1] describes agents' belief distributions that can emerge when there are multiple sources of information.

B) *Mechanisms and algorithms for fair and efficient resource allocation.* The classical economic design suffers from impossibility results. My key interest is in methods escaping the impossibilities. My main contribution is to fair allocation of resources. My research has highlighted a surprising dissimilarity between the classic setting with goods and problems with *bads* (chores) and inspired economic & CS literature on bads. My papers describe the pseudo-market approach for bads [2], develop a novel practical sharing-minimization approach [3], and provide the first results on the robust use of statistical data in fair division [4].

C) *Common patterns behind* A), B), *and beyond.* Seemingly unrelated problems of economic theory are connected to majorization, optimal transport, maximal flows, mathematical tomography, and measures with given marginals. My papers uncover and exploit these connections. My current goal is to understand the connections themselves. My recent WPs link multi-item multi-bidder auctions and multi-agent persuasion to optimal transportation. Ongoing projects explore consumer preference aggregation and establish a direct connection between information and Bayesian mechanism design through their mutual link to majorization.

Representative Papers:

[1] Feasible Joint Posterior Beliefs (JPE 2021, EC 2020 best paper award)
with I. Arieli, Y. Babichenko, and O. Tamuz

[2] Competitive Division of a Mixed Manna (Econometrica 2017)
with A. Bogomolnaia, H. Moulin, and E. Yanovskaya

[3] Fair Division with Minimal Sharing (Operations Research 2022)
with E. Segal-Halevi

[4] Online Fair Division Problem (Management Science 2022)
with A. Bogomolnaia and H. Moulin

DANIEL SCHOEPFLIN (Homepage, CV, Google Scholar)

Thesis: The Design and Analysis of Deferred Acceptance Clock Auctions ('23)

Advisor: Vasilis Gkatzelis, Drexel University

Brief Biography: Daniel Schoepflin is a Ph.D. candidate in Computer Science at Drexel University, where he is advised by Vasilis Gkatzelis. Previously, he obtained his undergraduate degree in Electrical Engineering at Drexel University.

Research Summary: My research centers around practical mechanism design, in particular, *(deferred acceptance) clock auctions*. Milgrom and Segal proposed clock auctions for the 2017 FCC spectrum reallocation auctions and argued that they are an ideal solution for many settings, in part, due to the unique list of practical properties which they satisfy. These properties include: (i) *obvious strategyproofness* (OSP), a very strong incentive property which loosely means that participants can "trivially" verify their optimal truth-telling strategy and which has been shown empirically to elicit truth-telling behavior from agents more readily than non-OSP truthful mechanisms; (ii) *(weak) group-strategyproofness*, meaning that clock auctions are robust to coalitions of bidders; (iii) *unconditional winner privacy*, which ensures that winners reveal the minimal information required to the auctioneer to "certify" that they should win; and (iv) *credibility*, meaning that it is in the best interest of revenue seeking auctioneers to follow the stated rules of the auction thereby not requiring participants to trust the auctioneer. Despite this compelling list of properties, relatively little was known about the performance of clock auctions. To address this, my thesis designs clock auctions for forward and backward auction settings and analyzes their performance from the perspective of approximation.

In [1], we provide a deterministic prior-free clock auction for *any* setting with an arbitrary downward closed feasibility constraint which obtains optimal social welfare approximation. In [2], we initiate the study of *Bayesian* clock auctions and propose the class of *single-price clock auctions* (SPCA), which closely resemble the well-known class of (sequential) posted-price mechanisms (SPM) - another class of simple clock auctions. While both SPCA and SPM offer a single price to each bidder, SPM must guarantee service if a bidder accepts her offer whereas SPCA makes service decisions after all bidders have responded. We show that this "deferred acceptance" allows for exponentially better approximation. We also show how one can gradually complicate the clock auction format and retain the same improved guarantees with gradually less prior information. In [3], we consider a reverse auction setting and the well-known "budget-feasible mechanism design problem". We design a deterministic clock auction which is the first deterministic truthful $O(1)$-approximate mechanism when the auctioneer has a submodular valuation function, thus resolving one of the most important open problems in this line of literature.

Representative Papers:

[1] Optimal Deterministic Clock Auctions (ITCS 2022)
with G. Christodoulou and V. Gkatzelis

[2] Bayesian and Randomized Clock Auctions (EC 2022)
with M. Feldman, V. Gkatzelis, and N. Gravin

[3] Deterministic Budget-feasible Clock Auctions (SODA 2022)
with E. Balkanski, P. Garimidi, V. Gkatzelis, and X. Tan

SEAN SINCLAIR (Homepage, CV, Google Scholar)

Thesis: Adaptivity, Structure, and Objectives in Sequential Decision Making ('23)

Advisor: Christina Lee Yu and Siddhartha Banerjee, Cornell University

Brief Biography: Sean Sinclair is a final-year Ph.D. student in the School of Operations Research and Information Engineering at Cornell University, where he is co-advised by Professors Christina Yu and Sid Banerjee. He was selected for the 2022 Future Leaders Summit at the Michigan Institute for Data Science. In 2020 and 2022 he was a visitor at the Simons Institute for the programs on the Theory of Reinforcement Learning and Data-Driven Decision Processes. In 2021 he was an intern at Microsoft Research working with Adith Swaminathan and Ching-An Chen on using reinforcement learning for cloud computing services. Prior to graduate school, he earned his BS in Mathematics and Computer Science from McGill University and served as a teacher in Ghana with the Peace Corps.

Research Summary: My research focuses on developing *algorithms for data-driven sequential decision making* for societal applications. For example, Feeding America collects truckloads of food from manufacturers and grocery stores, and distributes them to regional food banks under demand uncertainty. A typical solution approach either: (*i*) builds a simulator capturing the uncertainty, or (*ii*) solves an offline planning problem over predicted uncertainty sets. Reinforcement learning (RL) provides a framework built on MDPs, giving a way to design algorithms which jointly learn the model and policy. However, in practice it is difficult to define a ground truth objective, and algorithms requires modifications to scale for real-world systems.

My work bridges algorithmic techniques in reinforcement learning to an operations management perspective with an emphasis on models, data uncertainty, and objectives. Recent theoretical contributions include instance-specific optimal regret guarantees for nonparametric RL [2], Pareto-optimal fair resource allocation for mobile food pantries [1], and data-efficient algorithms for cloud computing allocations [3]. Complement to this, I also design open-source code instrumentation and methodology to empirically analyze the multi-criteria performance of algorithms on these problems [4].

Representative Papers:

[1] Sequential Fair Allocation: Achieving the Optimal Envy-Efficiency Tradeoff Curve (under review at *Operations Research*)
with G. Jain, S. Banerjee, and C. L. Yu

[2] Adaptive Discretization for Online Reinforcement Learning (under review at *Operations Research*) with S. Banerjee, and C. L. Yu

[3] Hindsight Learning in MDPs with Exogenous Inputs (working paper)
with F. Frujeri, C.-A. Cheng, and A. Swaminathan

[4] ORSuite: Benchmarking Suite for Sequential Operations Models (ACM PER) with C. Archer, S. Banerjee, M. Cortez, C. Rucker, M. Solberg, Q. Xie and C. Yu.

CLAYTON THOMAS (Homepage, CV)

Thesis: Explainable Mechanism Design ('23)

Advisor: Matt Weinberg, Princeton University

Brief Biography: Clayton Thomas is a final-year PhD candidate at Princeton University. His work strives towards new ways of understanding and explaining mechanisms. He has received a 2021 Princeton School of Engineering Award for Excellence, a 2023 Siebel Scholar award, and a 2022-2023 Wallace Memorial Fellowship (awarded to three students across the Princeton School of Engineering). He has co-organized a 2021 Princeton CS Student Theory Day, sub-reviewed for numerous CS theory conferences, spoken in and organized several reading groups, co-organized the 2022 SIGecom seminar series, and he was the EC 2022 website chair/student general chair.

Research Summary: Everybody loves strategyproof mechanisms. However, in real world settings such as school choice, their desirable properties crucially hinge on agents *understanding* their strategyproofness, and unfortunately, empirical papers show that real-world agents attempt to strategically manipulate these mechanisms. Moreover, previous theoretical frameworks such as obvious strategyproofness (OSP) can get very little traction towards explaining school choice mechanisms (as my paper [3] helps to show), and gives no guidance on how to proceed when an OSP implementation of a mechanism does not exist.

Nevertheless, in [1] we provide a new, simple description of deferred acceptance (DA, the most common strategyproof matching rule in practice) which makes strategyproofness easy to see. Our general framework, which we investigate both theoretically and empirically, gives a format for describing *any* strategyproof mechanism to one player at a time in a way that exposes strategyproofness, even when no OSP implementation exists. Complementing our positive results, we study the complexity of descriptions in this framework, and precisely quantify how and why descriptions of DA in our framework must differ substantially from traditional descriptions. In ongoing follow-up work, we examine how our framework interacts with other approaches toward explaining matching mechanisms, such as publishing admission cutoffs after the matching is calculated.

Where traditional computer science studies the ease of running an algorithm on a computer, my thesis studies the ease of explaining an algorithm to a human. I wish to continue this and other deeply interdisciplinary work between economics and computation, from classical questions such as those we address in [2,3,4] to new directions in explainability and interpretability as in [1].

Representative Papers:

[1] Strategyproofness-Exposing Mechanism Descriptions (Working paper) with Y. Gonczarowski and O. Heffetz

[2] Exponential Communication Separations Between Notions of Selfishness (STOC'21) with A. Rubinstein, R. R. Saxena, S. M. Weinberg, and J. Zhao.

[3] Classification of Priorities Such that DA is OSP-Implementable (EC'21)

[4] The Short-Side Advantage in Random Matching Markets (SOSA'22) with L. Cai

ARTEM TSIKIRIDIS (Homepage, CV, Google Scholar)

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

Advisor: Vangelis Markakis, Athens University of Economics and Business

Brief Biography: I am a final year Ph.D. student in Computer Science at Athens University of Economics and Business, advised by Vangelis Markakis. I received my B.Sc. in Management Science in 2015 and my M.Sc. in Computer Science in 2017, both from AUEB. I am interested in algorithmic game theory and problems on the interface of microeconomics and theoretical computer science. For Fall 2022, I am an intern at CWI, hosted by Guido Schäfer. Moreover, I have interned at CERN as an undergrad. I was awarded the Best Paper Award at SAGT'17.

Research Summary: I am interested in the study of combinatorial auctions from an interdisciplinary perspective. My research is focused on both the design of new protocols and the analysis of existing ones. My objective is to highlight design principles that auction mechanisms should follow so that they retain provable performance guarantees and, at the same time, are implementable in real-life scenarios.

A central theme of my research is the notion of a core-selecting mechanism (due to Ausubel & Milgrom). Such mechanisms have been known to possess good revenue guarantees and some of their variants have been used in practice. Despite their popularity, these auctions are generally non-truthful. As a result, current research has focused either on identifying core-selecting mechanisms with minimal incentives to deviate from truth-telling or on proposing truthful mechanisms whose revenue is competitive against core outcomes. In [1], we contributed to both of these directions. We started with a comparative statics analysis on the core polytope, which we utilized in two ways. Firstly, we proposed a truthful mechanism that is $O(\log n)$-competitive against the minimum revenue in the core for downward-closed single-parameter domains. Secondly, we studied the existence of non-decreasing payment rules, meaning that the payment of each bidder is a non-decreasing function of her bid. This property had been advocated by the related literature but it was an open question if there exist non-decreasing mechanisms that have minimal incentives for bidders to lie. We described such a class of rules.

Another theme of my research is the study of multi-unit auctions. Despite the fact that the VCG mechanism has important theoretical advantages under the 'core' lens, conceptually simpler non-truthful auction protocols are used in practice, such as the uniform price auction and the discriminatory auction. In [2], we settled the price of anarchy of pure Nash equilibria of the uniform price auction. Then, we focused in [3] on improving the known PoA bounds for broader classes of equilibria of the discriminatory auction, by providing a partial characterization of inefficiency in this format and by using variational calculus to construct inefficient instances.

Representative Papers:

[1] On Core-Selecting and Core-Competitive Mechanisms for Binary Single Parameter Auctions (WINE'19) with E. Markakis

[2] Tight Welfare Guarantees for Pure Nash Equilibria of the Uniform Price Auction (ToCS, SAGT'17-Best Paper) with G. Birmpas, E. Markakis, and O. Telelis

[3] Towards a Characterization of Worst Case Equilibria in the Discriminatory Price Auction (WINE'21) with E. Markakis and A. Sgouritsa

JINGYAN WANG (Homepage, CV, Google Scholar)

Thesis: Understanding and Mitigating Biases in Evaluation ('21)

Advisor: Nihar Shah, Carnegie Mellon University

Brief Biography: Jingyan Wang is a Ronald J. and Carol T. Beerman President's postdoctoral fellow in the H. Milton Stewart School of Industrial and Systems Engineering at Georgia Institute of Technology. Her research interests lie in analyzing and improving evaluation problems such as admissions and hiring, with tools from statistics and machine learning. She is the recipient of the Best Student Paper Award at AAMAS 2019. She received her Ph.D. from the School of Computer Science at Carnegie Mellon University in 2021, and her B.S. in Electrical Engineering and Computer Sciences from the University of California, Berkeley in 2015.

Research Summary: Evaluation – to estimate the quality of items or people – arises in many applications such as admissions, hiring, healthcare and law. These are high-stakes decision-making problems, where people desire the outcome to be equitable. However, the complexity of such evaluation tasks imposes fundamental challenges in obtaining accurate and fair results. The goal of my research is to *design, understand and improve these evaluation systems*. I use technical tools from statistics and computer science to develop theoretical insights, and work on policies to make practical impact. My research focuses on four components: *people, algorithms, designs* and *policies*.

I analyze the bounded rationality of *people*, by identifying human bias including miscalibration [1], personal experience [2] and ordering effect [3]. I propose intuitive models that do not rely on restrictive parametric assumptions, and develop algorithms that provably correct these biases.

I develop *algorithms* that integrate considerations of *accuracy, fairness* and *uncertainty*. In improving fairness, for the problem of learning from pairwise comparisons, I propose a simple and effective modification to the maximum likelihood estimator. In reducing the uncertainty involved in recruitment strategies, I formulate a combinatorial optimization problem, and present various algorithms that achieve desirable guarantees.

I study *design* decisions on how to engage people in evaluation. I examine different allocation schemes that assign tasks to evaluators in a distributed fashion, and characterize the tradeoffs between these schemes using a mixture of theoretical and experimental methods.

Finally, beyond technical work, I work with conferences and institutions to implement real *policy* improvements for bias mitigation in practice.

Representative Papers:

[1] Your 2 is My 1, Your 3 is My 9: Handling Arbitrary Miscalibrations in Ratings (AAMAS 2019) with N. B. Shah

[2] Debiasing Evaluations That Are Biased by Evaluations (AAAI 2021) with I. Stelmakh, Y. Wei, N. B. Shah

[3] Modeling and Correcting Bias in Sequential Evaluation (arXiv preprint 2022) with A. Pananjady

XINTONG WANG (Homepage, CV, Google Scholar)

Thesis: Computational Modeling and Design of Financial Markets: Towards Manipulation-Resistant and Expressive Markets ('21)

Advisor: Michael Wellman, University of Michigan

Brief Biography: Xintong Wang is a postdoctoral fellow in the Harvard EconCS group hosted by David Parkes. She received her Ph.D. in April 2021 from the CSE Department at the University of Michigan. During her Ph.D., she worked as a research intern at Microsoft Research NYC (mentored by David Pennock) and J.P. Morgan AI Research (mentored by Tucker Balch). Prior to Michigan, she received her Bachelor's degree from Washington University in St. Louis in 2015.

Research Summary: I work on developing principled computational approaches to model complex agent behaviors and to design better market-based algorithmic systems, drawing on tools from machine learning, game theory, and optimization.

One line of my research has focused on understanding and mitigating *market manipulation* in financial markets, where actions are taken to mislead others, so as to induce outcomes that increase one's own profit at the expense of market efficiency. We developed a multi-agent, causal model of *spoofing* (a common form of manipulation), demonstrating its effectiveness in a game-theoretic equilibrium, where other agents act rationally to learn from market data and trade. We further designed agent strategies that are robust to spoofing and market mechanisms that render spoofing unprofitable through the targeted hiding of market information [1].

In the context of *prediction market* and *financial options market* where agent preferences/beliefs are diverse and of complex structures (e.g., hierarchical, combinatorial), I have designed mechanisms that can both incorporate rich agent information in obtaining better information aggregation or economic efficiency, while still achieving computational efficiency [2, 3]. In [3], we further proposed a novel instrument, *combinatorial financial option*, which enables the elicitation of future correlations among multiple random variables, by generalizing standard financial options to support trading any linear combination of underlying assets.

In a recent line of research, we use deep RL to model the fee-setting and matching behavior of an economic platform (e.g., UberEats) under market shocks (e.g., the pandemic) that disrupt the traditional, offline market [4]. By further imposing different interventions (e.g., fee caps and taxation) to the platform, we aim to evaluate and inform the design of regulatory policies to align platform incentives (e.g., revenue) with broader social considerations (e.g., seller diversity, sustainability).

Representative Papers:

[1] Spoofing the Limit Order Book: A Strategic Agent-Based Analysis (Games'21) with C. Hoang, Y. Vorobeychik, and M. Wellman

[2] Log-time Prediction Markets for Interval Securities (AAMAS'21) with M. Dudík, D. Pennock, and D. Rothschild

[3] Designing a Combinatorial Financial Options Market (EC'21) with D. Pennock, N. Devanur, D. Rothschild, B. Tao, and M. Wellman

[4] Platform Behavior under Market Shocks: A Reinforcement-Learning Based Study (Manuscript) with G. Ma, A. Edon, C. Li, A. Trott, S. Zheng, D. Parkes

HANRUI ZHANG (Homepage, CV)

Thesis: Designing and Analyzing Machine Learning Algorithms in the Presence of Strategic Manipulation ('23)

Advisor: Vincent Conitzer, Carnegie Mellon University

Brief Biography: Hanrui Zhang is a PhD student at Carnegie Mellon University, advised by Vincent Conitzer. He interned twice at Google, where he worked closely with the OMEGA Market Algorithms Team headed by Vahab Mirrokni. He was named a finalist for the 2021 Facebook Fellowship, and won the Best Student Paper Award at the ESA 2020 conference.

Research Summary: Machine learning algorithms now play a major role in all kinds of decision-making scenarios, such as college admissions, credit approval, and resume screening. When the stakes are high, self-interested agents — about whom decisions are being made — are increasingly tempted to manipulate the machine learning algorithm, in order to better fulfill their own goals, which are generally different from the decision maker's. The main focus of my research is on *designing and analyzing machine learning algorithms that are robust against strategic manipulation.* My research sets the foundations for several key problems in machine learning in the presence of strategic behavior:

(i) Empirical risk minimization and generalization in classification problems [1]: Traditional wisdom suggests that a classifier trained on historical observations (i.e., an empirical risk minimizer) usually also works well on future data points to be classified. Is this still true in the presence of strategic manipulation?

(ii) Distinguishing distributions with samples [2, 3]: Due to various constraints, often we have to judge the quality of a data point based on a few samples (e.g., screening job candidates based on a few representative papers). How should we calibrate our judgment when these samples are strategically selected or transformed?

(iii) Planning in Markov decision processes [4, 5]: Dynamic decision-making problems (traditionally modeled using Markov decision processes) can be solved efficiently when the decision maker always has complete and reliable information about the state of the world, as well as full control over which actions to take. What happens when the state of the world is reported by a strategic agent, or when a self-interested agent may interfere with the actions taken?

Representative Papers:

[1] Incentive-Aware PAC Learning (AAAI 2021) with V. Conitzer

[2] When Samples Are Strategically Selected (ICML 2019) with Y. Cheng and V. Conitzer

[3] Distinguishing Distributions When Samples Are Strategically Transformed (NeurIPS 2019) with Y. Cheng and V. Conitzer

[4] Automated Dynamic Mechanism Design (NeurIPS 2021) with V. Conitzer

[5] Efficient Algorithms for Planning with Participation Constraints (EC 2022) with Y. Cheng and V. Conitzer

JUNYAO ZHAO (Homepage, CV)

Thesis: Submodular Maximization and Mechanism Design: Beyond Worst-Case Analysis and Optimal Communication Complexity ('23)

Advisor: Aviad Rubinstein, Stanford University

Brief Biography: Junyao is a fifth-year PhD student in the theory group of the CS Department at Stanford University. He is broadly interested in TCS, game theory, and machine learning. During his PhD, he strives to (i) design algorithms and mechanisms with better beyond worst-case performance for submodular optimization and auctions and (ii) understand the extra communication cost caused by enforcing incentive compatibility constraints in algorithmic mechanism design. In 2021 & 2022 summer, he was a research intern working on contract theory and autobidding in the Market Algorithms team at Google Mountain View. Before PhD, he obtained his MSc degree with distinction in CS from ETH Zürich.

Research Summary: *Beyond worst-case analysis.* We revisit the classic cardinality-constrained submodular maximization problem from a novel beyond-worst-case perspective. Specifically, We study a semi-adversarial setting where the submodular function is adversarially chosen, but the cardinality constraint satisfies some beyond-worst-case assumptions: (i) One assumption we consider is that the constraint is sampled from an arbitrary distribution, and in this case we show that the classic greedy algorithm is still optimal for monotone submodular functions and enjoys better-than-worst-case performance for realistic distributions [1]. (ii) Another assumption we consider is that the cardinality constraint is significantly smaller than the number of the elements which we choose from, and under this assumption, we give a 0.539-approximation algorithm for any non-negative submodular function [2], successfully circumventing all the existing hardness results in the worst-case setting.

Communication Complexity. In algorithmic mechanism design, we design mechanisms that satisfy certain incentive compatibility constraints. Does incentive compatible (IC) mechanisms require a higher communication complexity than non-incentive-compatible (non-IC) algorithms to achieve the same performance? (i) In [3], we establish optimal randomized communication complexity for revenue-maximizing auctions – Our randomized IC auction protocol is infinitely more efficient than any deterministic IC protocol, and our tight lower bound shows that IC protocols require exponentially higher communication cost than non-IC protocols. (ii) Moreover, in [4], we show that stronger IC constraints can incur exponentially larger communication cost than weaker IC constraints.

Representative Papers:

[1] Budget-Smoothed Analysis for Submodular Maximization (ITCS 2022) with A. Rubinstein

[2] Maximizing Non-Monotone Submodular Functions over Small Subsets: Beyond 1/2-Approximation (ICALP 2022) with A. Rubinstein

[3] The Randomized Communication Complexity of Randomized Auctions (STOC 2021) with A. Rubinstein

[4] Exponential Communication Separations between Notions of Selfishness (STOC 2021) with A. Rubinstein, R.R. Saxena, C. Thomas, and S.M. Weinberg

Index

combinatorial auctions
Tomer Ezra, 15
general equilibrium theory
Bhaskar Ray Chaudhury, 13

adaptive experiments
Anand Kalvit, 23
adversarial ML
Andrew Estornell, 14
agent-based simulation
Xintong Wang, 36
algorithmic fairness
Andrew Estornell, 14
Evi Micha, 28
algorithmic game theory
Sumit Goel, 21
Yoav Kolumbus, 25
Artem Tsikiridis, 34
algorithms
Martin Bullinger, 11
artificial intelligence
Yoav Kolumbus, 25
auctions
Samuel Altmann, 9
Maximilian Fichtl, 19
Yotam Gafni, 20
Daniel Schoepflin, 31
Artem Tsikiridis, 34

beyond worst-case analysis
Junyao Zhao, 38
blockchain
Matheus V. X. Ferreira, 18
bounded rationality
Jakub Černý, 12

citizens' assemblies
Paul Gölz, 22
clock auctions
Daniel Schoepflin, 31
combinatorial optimization
Gregory Kehne, 24
complexity
Junyao Zhao, 38
computational complexity
Martin Bullinger, 11
computational experiments
Niclas Boehmer, 10
computational social choice
Evi Micha, 28
contracts
Tomer Ezra, 15
crowdsourcing
Jingyan Wang, 35

data-driven decision making
Anand Kalvit, 23
Sean Sinclair, 32
democratic innovations
Paul Gölz, 22
dynamical systems
Yoav Kolumbus, 25

econometrics
Samuel Altmann, 9
empirical market design
Samuel Altmann, 9
equilibrium computation
Jakub Černý, 12
Maximilian Fichtl, 19
explainable mechanism design
Clayton Thomas, 33

facility location
Alexander Lam, 26
fair allocation
Yotam Gafni, 20
fair division
Tomer Ezra, 15
Paul Gölz, 22
Fedor Sandomirskiy, 30
Daniel Schoepflin, 31
fairness
Niclas Boehmer, 10
Sean Sinclair, 32
Jingyan Wang, 35
fairness in ML
Bhaskar Ray Chaudhury, 13

game theory
Jakub Černý, 12

Martin Bullinger, 11
Matheus V. X. Ferreira, 18
Alexander Lam, 26
Majid Mahzoon, 27
Xintong Wang, 36
game theory and mechanism design
Alireza Fallah, 16

information economics
Majid Mahzoon, 27
Fedor Sandomirskiy, 30

machine learning
Jingyan Wang, 35
Xintong Wang, 36
Hanrui Zhang, 37
machine learning theory
Alireza Fallah, 16
market design
Matheus V. X. Ferreira, 18
Clayton Thomas, 33
Xintong Wang, 36
matching
Niclas Boehmer, 10
mechanism design
Yiding Feng, 17
Yotam Gafni, 20
Sumit Goel, 21
Alexander Lam, 26
Majid Mahzoon, 27
Evi Micha, 28
Divyarthi Mohan, 29
Fedor Sandomirskiy, 30
Artem Tsikiridis, 34
Hanrui Zhang, 37
Junyao Zhao, 38

networks
Yoav Kolumbus, 25

obvious strategyproofness
Daniel Schoepflin, 31
online algorithms
Yiding Feng, 17
Gregory Kehne, 24
online decision-making
Tomer Ezra, 15
online learning
Maximilian Fichtl, 19
Anand Kalvit, 23
Yoav Kolumbus, 25
opinion dynamics
Divyarthi Mohan, 29
optimization
Alireza Fallah, 16
Maximilian Fichtl, 19
Junyao Zhao, 38

platforms
Anand Kalvit, 23
proportional fairness
Alexander Lam, 26

reinforcement learning
Sean Sinclair, 32
revenue management
Yiding Feng, 17

security
Matheus V. X. Ferreira, 18
social choice
Niclas Boehmer, 10
Martin Bullinger, 11
Bhaskar Ray Chaudhury, 13
Andrew Estornell, 14
Sumit Goel, 21
Gregory Kehne, 24
social good
Paul Gölz, 22
stable matching
Clayton Thomas, 33
strategic classification
Andrew Estornell, 14
strategic communication
Divyarthi Mohan, 29
strategic machine learning
Yotam Gafni, 20