



<b>Jeff Decary</b> prediction markets, portfolio, market design, sports analytics	<b>12</b>
<b>Craig Fernandes</b> market design, revenue management, operations management	<b>13</b>
<b>Giannis Fikioris</b> learning in games, online learning, online fair allocation, auctions	<b>14</b>
<b>Maxwell Fishelson</b> online learning, regret minimization, calibration	<b>15</b>
<b>Abheek Ghosh</b> computational complexity, contest theory, learning dynamics, optimization	<b>16</b>
<b>Sumit Goel</b> mechanism design, game theory, contract theory, matching	<b>17</b>
<b>Andreas Haupt</b> market design, privacy, preference learning, online learning, competition	<b>18</b>
<b>Vishwa Prakash HV</b> fair division, algorithmic game theory, algorithm design, matching	<b>19</b>
<b>Stanisław Kaźmierowski</b> game theory, conflicts with multiple battlefields, equilibria computation	<b>20</b>
<b>Yoav Kolumbus</b> learning, AGT, operations, markets, auctions, dynamics, networks	<b>21</b>
<b>Pooja Kulkarni</b> data economics, online algorithms, submodular optimization, fair division	<b>22</b>
<b>Soonbong Lee</b> operations for social good, online and matching platforms, market design	<b>23</b>
<b>Ce Li</b> information design, mechanism design, human-AI interactions, ML	<b>24</b>
<b>Yuxin Liu</b> differential privacy, mechanism design, network science	<b>25</b>
<b>Gary Qiurui Ma</b> platform and gig economy, pricing, matching, market design	<b>26</b>
<b>Aniket Murhekar</b> incentives in machine learning, data economics, fair division, social choice	<b>27</b>

<b>Kalen Patton</b> online algorithms, resource allocation, load balancing, prophet inequalities	28
<b>Tomasz Ponitka</b> fair division, contract design	29
<b>Prasanna Ramakrishnan</b> computational social choice, approximation algorithms, game theory	30
<b>Kiran Rokade</b> game theory, networks, equilibrium computation, cooperation	31
<b>Shiri Ron</b> mechanism design, auction theory, communication complexity	32
<b>Shaul Rosner</b> game theory, congestion, contract design	33
<b>Šimon Schierreich</b> social choice, collective decision making, algorithmic game theory	34
<b>Daniel Schoepflin</b> practical mechanism design, bounded rationality, optimization	35
<b>Eklavya Sharma</b> data markets, fair allocation, social choice, approximation algorithms	36
<b>Suho Shin</b> mechanism design, delegation, online platforms, generative AI	37
<b>Ana-Andreea Stoica</b> machine learning, social networks, game theory, causal inference	38
<b>Yifan Wang</b> mechanism design, pricing, online algorithms	39
<b>Wentao Weng</b> machine learning, queueing theory, sustainable governance	40
<b>Nicholas Wu</b> dynamic games, mechanism design, information design	41
<b>Shengwei Xu</b> large language models, information elicitation, peer prediction	42
<b>Aviv Yaish</b> mechanism design, security, blockchain, decentralized finance	43

**Kunhe Yang**

incentive-aware machine learning, algorithmic game theory, AI alignment 44

**Konstantin Zabarnyi**

information design, cheap talk, mechanism design, robustness 45

**Wenxin Zhang**

dynamic resource allocation, online algorithms, generative AI efficiency 46

**Yichi Zhang**

information elicitation, evaluation with AI in the loop, multi-agent LLM 47

ALIREZA AMANIHAMEDANI ([Homepage](#), [CV](#))

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

**Advisor:** Ali Aouad, MIT

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

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

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

#### Representative Papers:

- [1] Adaptive Approximation Schemes for Matching Queues (STOC’25)  
with Ali Aouad, and Amin Saberi
- [2] Spatial Matching under Multihoming (major revision at Operations Research)  
with Ali Aouad, and Daniel Freund
- [3] Improved Approximations for Stationary Bipartite Matching: Beyond Probabilistic Independence (under review at Mathematics of Operations Research)  
with Ali Aouad, Tristan Pollner, and Amin Saberi

ABDELLAH AZNAG ([Homepage](#), [CV](#), [Google Scholar](#))

**Thesis:** Advances in Adaptive Data Collection ('26)

**Advisors:** Rachel Cummings, Adam N. Elmachtoub

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

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

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

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

#### Representative Papers:

- [1] An Active Learning Framework For Multi-group Mean Estimation (NeurIPS 2023 and major revision in Management Science) with A. N. Elmachtoub, and R. Cummings
- [2] Designing Lower Bounds for Active Learning in Multi-Armed Bandits (under review in Operations Research) with A. N. Elmachtoub, and R. Cummings
- [3] MNL-Bandit With Knapsacks: a Near-Optimal Algorithm (EC'21 and major revision in Operations Research) with V. Goyal, and N. Perivier
- [4] Calibrated experiments via Active Learning (working paper)
- [5] Low-rank adaptive active experiment design (working paper)

FEDERICO BOBBIO ([Homepage](#), [CV](#))

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

**Advisor:** Margarida Carvalho and Andrea Lodi, Université de Montréal

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

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

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

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

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

#### Representative Papers:

- [1] Capacity Planning in Stable Matching (Operations Research, 2025). With M. Carvalho, A. Lodi, I. Rios, and A. Torrico
- [2] Stable Matching with Contingent Priorities (EC 2025). With I. Rios, M. Carvalho, and A. Torrico.
- [3] Costly Measurements to Incentivize Spectrum Sharing (IEEE DySPAN 2025). With R. Berry, M. Honig, T. Nguyen, V. Subramanian, and R. Vohra.
- [4] Design and Implementation of a Heuristic-Enhanced Branch-and-Bound Solver for MILP (MIP Workshop 2022). With W. A. Silva, F. Caye, D. Liu, J. Pepin, C. Perreault-Lafleur, and W. St-Arnaud.
- [5] The First AI4TSP Competition: Learning to Solve Stochastic Routing Problems (Artificial Intelligence, 2023). With Y. Zhang, L. Bliet, and others.

ROBIN BOWERS ([Homepage](#), [CV](#))

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

**Advisor:** Bo Waggoner and Rafael Frongillo, CU Boulder

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

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

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

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

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

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

#### **Representative Papers:**

- [1] High Welfare Matching Markets via Descending Price (WINE ’23)  
with B. Waggoner
- [2] Matching with Nested and Bundled Pandora Boxes (WINE ’24)  
with B. Waggoner
- [3] Prophet Inequalities for Bandits, Cabinets, and DAGs (Working paper)  
with E. Lindgren, B. Waggoner
- [4] Polynomial-Time Approximation Schemes via Utility Alignment: Unit-Demand Pricing and More (FOCS ’25) with M. Garbea, E. Pountourakis, S. Taggart.



NATALIE COLLINA ([Homepage](#), [CV](#), [Google Scholar](#))

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

**Advisors:** Aaron Roth and Michael Kearns, University of Pennsylvania

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

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

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

**Representative Papers:**

- [1] Emergent Alignment via Competition (in submission to ICLR 2026)  
with S. Goel, A. Roth, E. Ryu, and M. Shi
- [2] Tractable Agreement Protocols (STOC 2025)  
with S. Goel, V. Gupta, and A. Roth
- [3] Algorithmic Collusion Without Threats (ITCS 2025)  
with E.R. Arunachaleswaran, S. Kannan, A. Roth, and J. Ziani
- [4] Pareto-Optimal Algorithms for Learning in Games (EC 2024)  
with E.R. Arunachaleswaran and J. Schneider
- [5] Swap Regret and Correlated Equilibria Beyond Normal-Form Games (EC 2025)  
with E.R. Arunachaleswaran, Y. Mansour, M. Mohri, J. Schneider, and B. Sivan

JEFF DECARY ([Homepage](#), [CV](#), [Google Scholar](#))

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

**Advisor:** David Bergman, University of Connecticut

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

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

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

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

### Representative Papers:

- [1] Log-Optimal Portfolio Construction for Binary Options with Combinatorial Constraints (under review at Management Science) with D. Bergman, and B. Zou
- [2] The madness of multiple entries in march madness (Poster at EC, 2024; 2nd Round Revision at Production & Operations Management) with D. Bergman, C. Cardonha, J. Imbrogno, and A. Lodi
- [3] Market-making for daily fantasy sports: Competitive pricing for sophistication (Working Paper) with D. Bergman, and R. Day

CRAIG FERNANDES ([Homepage](#), [CV](#))

**Thesis:** Operations Management for Innovative Markets ('26)

**Advisor:** Timothy Chan and Ningyuan Chen; University of Toronto

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

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

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

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

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

### Representative Papers:

- [1] Income Pools for Superstar Markets (Published, Management Science, 2025)  
with Timothy Chan and Ningyuan Chen
- [2] Points Gained in Football (Published, Operations Research, 2021)  
with Timothy Chan and Martin Puterman
- [3] Peer Review Market Design (ACM EC'25; Target: Operations Research)  
with James Siderius and Raghav Singal
- [4] Data-Driven B2B Pricing: Learning from Bargainers (Working paper)  
with Timothy Chan and Ningyuan Chen

GIANNIS FIKIORIS ([Homepage](#), [CV](#), [Google Scholar](#))

**Thesis:** Learning in Games with State ('26)

**Advisor:** Éva Tardos, Cornell University

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

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

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

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

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

### Representative Papers:

- [1] Liquid Welfare Guarantees for No-Regret Learning in Sequential Budgeted Auctions (EC'23 & Math of OR) with É. Tardos
- [2] Learning in Budgeted Auctions with Spacing Objectives (EC'25 & Under Review at *JMLR*)  
with R. Kleinberg, Y. Kolumbus, R. Kumar, Y. Mansour, and É. Tardos
- [3] Beyond Worst-Case Online Allocation via Dynamic Max-min Fairness (EC'25 & Under Review at *Management Science*) with S. Banerjee and É. Tardos

MAXWELL FISHELSON ([Homepage](#), [CV](#), [Google Scholar](#))

**Thesis:** Reliable Learning for Adversarial Environments ('26)

**Advisor:** Constantinos Daskalakis, Massachusetts Institute of Technology

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

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

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

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

#### Representative Papers:

- [1] Near-Optimal No-Regret Learning in General Games (NeurIPS 2021, Oral)  
with Constantinos Daskalakis and Noah Golowich
- [2] From External to Swap Regret 2.0: An Efficient Reduction for Large Action Spaces (STOC 2024, Invited to SICOMP Special Issue)  
with Yuval Dagan, Constantinos Daskalakis, and Noah Golowich
- [3] Breaking the  $T^{2/3}$  Barrier for Sequential Calibration (STOC 2025, Invited to SICOMP Special Issue)  
with Yuval Dagan, Constantinos Daskalakis, Noah Golowich, Robert Kleinberg, and Princewill Okoroafor
- [4] High-Dimensional Calibration from Swap Regret (NeurIPS 2025, Oral)  
with Noah Golowich, Mehryar Mohri, and Jon Schneider

ABHEEK GHOSH ([Homepage](#), [CV](#))

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

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

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

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

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

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

#### Representative Papers:

- [1] Computing Equilibrium Points of Electrostatic Potentials (in submission)  
with PW. Goldberg and A. Hollender
- [2] Fraud-Proof Revenue Division in Subscription Platforms (ICML 25)  
with TY. Neoh, N. Teh, and G. Tyrovolas
- [3] Best-Response Dynamics in Lottery Contests (EC 23)  
with PW. Goldberg

SUMIT GOEL ([Homepage](#), [CV](#), [Google Scholar](#))

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

**Advisor:** Federico Echenique, U.C. Berkeley

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

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

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

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

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

#### **Representative Papers:**

- [1] The Effect of Competition in Contests: A Unifying Approach (JMP)  
with A. Baranski
- [2] TTC Domains (WINE 2025)  
with Y. Tamura
- [3] Multi-Agent Contract Design with a Budget (EC 2024, R&R at GEB)  
with W. Hann-Caruthers
- [4] Stable Allocations in Discrete Exchange Economies (JET 2024)  
with F. Echenique, and S. Lee
- [5] Optimal Grading Contests (EC 2023, GEB 2025)



ANDREAS HAUPT ([Homepage](#), [CV](#), [Google Scholar](#))

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

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

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

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

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

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

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

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

#### Representative Papers:

- [1] Contextually Private Mechanisms (ACM EC '22, R&R at the AER)  
with Z. Hitzig
- [2] Risk Preferences of Learning Algorithms (Games and Economic Behavior '24)  
with A. Narayanan
- [3] Platform Preferencing and Price Competition I (Working Paper)  
with O. Hartzell
- [4] Certification Design for a Competitive Market (ACM EC '23)  
with B. Lucier and N. Immorlica



VISHWA PRAKASH HV ([Homepage](#), [CV](#), [Google Scholar](#))

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

**Advisor:** Prajakta Nimbhorkar, Chennai Mathematical Institute.

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

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

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

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

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

### Representative Papers:

- [1] EFX Exists for Three Types of Agents (EC 2025)  
with P. Ghosal, P. Nimbhorkar, and N. Varma
- [2] (Almost Full) EFX for Three (and More) Types of Agents. (AAAI 2025)  
with P. Ghosal, P. Nimbhorkar, and N. Varma
- [3] Almost and Approximate EFX for Few Types of Agents (working paper)  
with R. Mehta and P. Nimbhorkar
- [4] Fair and Efficient Completion of Indivisible Goods (AAAI 2025)  
with A. Igarashi, and R. Vaish
- [5] Fair and Efficient Allocation of Indivisible Mixed Manna (WINE 2025)  
with S. Barman, A. Sethia, and M. Suzuki
- [6] Weighted Proportional Allocations of Indivisible Goods and Chores: Insights via Matchings. (AAMAS 2024)  
with P. Nimbhorkar

STANISŁAW KAŻMIEROWSKI ([Homepage](#), [CV](#))

**Thesis:** Solving Succinct Games ('25)

**Advisor:** Marcin Dziubiński, University of Warsaw

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

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

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

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

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

#### **Representative Papers:**

- [1] Equilibria of the Colonel Blotto Games with Costs (AAAI 2025)
- [2] Computation of Nash Equilibria of Attack and Defense Games on Networks (SAGT 2023) with M. Dziubiński
- [3] Efficient Method for Finding Optimal Strategies in Chopstick Auctions with Uniform Objects Values (AAMAS 2024) with M. Dziubiński
- [4] An equilibrium analysis of the Arad-Rubinstein game (Journal of Economic Behavior & Organization) with C. Ewerhart

YOAV KOLUMBUS ([Homepage](#), [CV](#))

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

**Advisor:** Noam Nisan, The Hebrew University of Jerusalem

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

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

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

- (1) *System analysis:* Analyzing dynamics and trade-offs in strategic systems and markets with learning agents.
- (2) *Better learning and platforms:* Designing learning methods for improved performance in uncertain strategic settings and mechanisms suitable for learners.
- (3) *Interaction models:* Modeling interactions between strategic users and algorithmic agents under varied structures.

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

#### **Representative Papers:**

- [1] Markets with Heterogeneous Agents: Dynamics and Survival of Bayesian vs. No-Regret Learners (EC 2025) with D. Easley and É. Tardos
- [2] Auctions between Regret-Minimizing Agents (TheWebConf 2022) with N. Nisan
- [3] Contracting with a Learning Agent (NeurIPS 2024) with I. Talgam-Cohen, M. Weinberg, G. Guruganesh, M. Vlatakis, J. Wang, and J. Schneider

POOJA KULKARNI ([Homepage](#), [CV](#))

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

**Advisors:** Jugal Garg and Ruta Mehta, UIUC.

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

**Research Summary:** My research develops algorithms that make resource-allocation and market systems *stable and predictable*. Whether distributing disaster relief or selling digital goods such as API access to an LLM, stability can be ensured either from fairness or from equilibrium, where no individual or group benefits from deviating. My work advances these goals across three main lines:

*Online Fair Allocation.* In [1], I model disaster-relief allocation as an online discrete fair-division problem, where agents arrive sequentially. Such problems are notoriously hard—few positive results were previously known. My work introduced the ONLINEKTYPEFD model, inspired by learning-augmented algorithms, where agents belong to one of  $K$  known types. This yields the first non-trivial results for this setting. Specifically, we design deterministic algorithms for the Maximin Share (MMS) fairness notion. I am extending this line of work to randomized algorithms and other fairness notions.

*Submodular Fair Allocation.* Most real agents exhibit non-additive valuations like submodular valuations (capturing diminishing returns) yet fair-division theory largely assumes additive utilities. My work (e.g., [2]) extends notions like Nash Social Welfare and Maximin Share to submodular and fractionally subadditive valuations—relevant for settings such as course allocation. Since techniques for additive valuations rarely extend, my work has introduced new techniques such as *match-rematch* and *capped-welfare maximization*.

*Equilibrium in Data Economies.* Data is an economic asset. It is *non-rivalrous* i.e., can benefit multiple agents simultaneously. This fundamentally alters how exchange or trading economies with data will behave. To investigate this, in [3], I model data-sharing economies without payments, characterizing stable outcomes using the notion of *core* from cooperative game theory. In an ongoing work, I analyze data pricing, showing how non-rivalrousness fundamentally changes strategic pricing decisions and revenue compared to traditional goods.

#### Representative Papers:

- [1] Online Fair Division: Towards Ex-Post Constant MMS Guarantees (EC 2025) with R. Mehta, P. Shahkar
- [2] Approximating NSW under Submodular Valuations through (Un)Matchings (SODA 2020, TALG 2024) with J. Garg, R. Kulkarni
- [3] On the Existence and Complexity of Core Stable Data Exchanges (NeurIPS 2025) with J. Song, P. Shahkar, B. Chaudhury

SOONBONG LEE ([Homepage](#), [CV](#), [Google Scholar](#))

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

**Advisor:** Vahideh Manshadi (Yale School of Management)

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

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

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

#### Representative Papers:

- [1] Dynamic Matching with Post-Allocation Service and its Application to Refugee Resettlement (Accepted at *Management Science* & EC 2024) with K. Bansak, V. Manshadi, R. Niazadeh, and E. Paulson
- [2] Why the Rooney Rule Fumbles: Limitations of Interview-stage Diversity Interventions in Labor Markets (EC 2025 & Major Revision at *Operations Research*) with S. Farajollahzadeh, V. Manshadi, and F. Monachou
- [3] Who to Offer, and When: Redesigning Feeding America's Real-Time Donation Tool (Working Paper) with V. Manshadi, and D. Saban

CE LI ([Homepage](#), [CV](#))

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

**Advisor:** Bart Lipman, Department of Economics, Boston University

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

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

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

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

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

### Representative Papers:

- [1] Learning to Design Information (Job Market Paper), with Tao Lin
- [2] Information Design with Unknown Prior (ITCS 2025), with Tao Lin
- [3] From Best Responses to Learning: Investment Efficiency in Dynamic Environment (working paper), with Qianfan Zhang and Weiqiang Zheng

YUXIN LIU ([Homepage](#), [CV](#), [Google Scholar](#))

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

**Advisor:** M.Amin Rahimian, University of Pittsburgh

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

### Research Summary.

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

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

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

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

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

### Representative Papers:

- [1] Privacy-Aware Sequential Learning (Under Review at Operations Research) with M. A. Rahimian
- [2] Seeding with Differentially Private Network Information. (under submission) with M. A. Rahimian and F. Yu
- [3] Structural Dynamics of Harmful Content Dissemination on WhatsApp (ICWSM 2026) with M. A. Rahimian and K. Garimella



GARY QIURUI MA ([Homepage](#), [CV](#))

**Thesis:** Equilibrium Design in Online Marketplaces ('26)

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

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

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

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

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

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

#### Representative Papers:

- [1] Platform Equilibrium: Analyzing Social Welfare in Online Market Places (EC'24) with A. Eden, and D. Parkes
- [2] Disrupting Bipartite Trading Networks: Matching for Revenue Maximization (EC'24) with Y. Gonczarowski, L. D'Amico-Wong, and D. Parkes
- [3] Pricing with Tips in Three-Sided Delivery Platforms (working paper) with Y. Gonczarowski, and D. Parkes
- [4] Platform Behavior under Market Shocks: A Simulation Framework and Reinforcement-Learning Based Study (WWW'23) with X. Wang, A. Eden, C. Li, A. Trott, S. Zheng, and D. Parkes



ANIKET MURHEKAR ([Homepage](#), [CV](#), [Google Scholar](#))

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

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

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

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

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

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

#### Representative Papers:

- [1] Constant-Factor EFX Exists for Chores (STOC 2025)  
with J. Garg and J. Qin
- [2] On the Theoretical Foundations of Data Exchange Economies (EC 2025)  
with B. Chaudhury, J. Garg, and J. Song
- [3] You Get What You Give: Reciprocal Fair Federated Learning (ICML 2025)  
with B. Chaudhury, R. Mehta, and J. Song
- [4] Fair Federated Learning via the Proportional Veto Core (ICML 2024)  
with B. Chaudhury, Z. Yuan, B. Li, R. Mehta, and A.D. Procaccia

KALEN PATTON ([Homepage](#), [CV](#), [Google Scholar](#))

**Thesis:** Online Resource Allocation and Load Balancing Beyond  $\ell_p$ -Norms ('26)

**Advisor:** Sahil Singla, Georgia Institute of Technology

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

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

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

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

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

### Representative Papers:

- [1] Online Allocation with Concave Diminishing>Returns Objectives (SODA 2026)
- [2] Integral Algorithms for Online Set Cover and Load Balancing (FOCS 2025)  
with T. Kesselheim, M. Molinaro, and S. Singla
- [3] The Online Submodular Assignment Problem (FOCS 2024)  
with D. Hathcock, B. Jin, S. Sarkar, and M. Zlatin
- [4] Improved Mechanisms and Prophet Inequalities for Graphical Dependencies  
(EC 2024) with V. Livanos and S. Singla

TOMASZ PONITKA ([Homepage](#), [CV](#), [Google Scholar](#))

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

**Advisor:** Michal Feldman, Tel Aviv University

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

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

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

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

#### Representative Papers:

- [1] Breaking the Envy Cycle: Best-of-Both-Worlds Guarantees for Subadditive Valuations (EC 2024) with M. Feldman, S. Mavroug, and V. V. Narayan.
- [2] The Pseudo-Dimension of Contracts (EC 2025) with P. Dütting, M. Feldman, and E. Soumalias.

PRASANNA RAMAKRISHNAN ([Homepage](#), [CV](#), [Google Scholar](#))

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

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

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

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

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

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

#### Representative Papers:

- [1] Breaking the Metric Voting Distortion Barrier (JACM 2024, SODA 2024)  
with M. Charikar, K. Wang, and H. Wu.
- [2] Six Candidates Suffice to Win a Voter Majority (STOC 2025)  
with M. Charikar, A. Lassota, A. Vetta, and K. Wang.
- [3] Approximately Dominating Sets in Elections (SODA 2026)  
with M. Charikar and K. Wang.
- [4] Distortion in Metric Matching with Ordinal Preferences (EC 2023)  
with N. Anari and M. Charikar.

KIRAN ROKADE ([Homepage](#), [CV](#), [Google Scholar](#))

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

**Advisor:** Francesca Parise, Cornell University

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

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

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

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

#### Representative Papers:

- [1] Graphon Games with Multiple Equilibria: Analysis and Computation (EC 2023, accepted at Mathematics of Operations Research) with F. Parise
- [2] Asymmetric Network Games:  $\alpha$ -Potential Function and Learning (ArXiv) with A. Jain, F. Parise, V. Krishnamurthy, E. Tardos
- [3] Learning in Time-Varying Monotone Network Games with Dynamic Populations (ArXiv, preliminary version in CDC 2023) with F. Al Taha (first author), F. Parise

SHIRI RON ([Homepage](#), [CV](#))

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

**Advisor:** Shahar Dobzinski, Weizmann Institute of Science

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

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

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

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

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

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

#### **Representative Papers:**

- [1] On the Hardness of Dominant Strategy Mechanism Design (STOC 2022)  
with Shahar Dobzinski and Jan Vondrák
- [2] Impossibilities for Obviously Strategy-Proof Mechanisms (SODA 2024)
- [3] On the Power of Randomization for Obviously Strategy-Proof Mechanisms (AAAI 2025) with Daniel Schoepflin

SHAUL ROSNER ([Homepage](#), [CV](#), [Google Scholar](#))

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

**Advisor:** Tami Tamir, Reichman University

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

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

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

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

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

#### **Representative Papers:**

- [1] Scheduling Games with Rank-Based Utilities (SAGT 2020, Games and Economic Behavior 2023) with T. Tamir
- [2] Cost-Sharing Games with Rank-Based Utilities (SAGT 2022, Theoretical Computer Science 2025) with T. Tamir
- [3] Nash Flows Over Time with Tolls (Accepted WINE 2025)  
with M. Schröder, and L. Vargas Koch
- [4] Bipartite Matching with Pair-Dependent Bounds (Under Review)  
with T. Tamir



ŠIMON SCHIERREICH ([Homepage](#), [CV](#), [Google Scholar](#))

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

**Advisor:** Dušan Knop, Czech Technical University in Prague

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

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

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

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

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

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

#### Representative Papers:

- [1] Evaluation of Project Performance in Participatory Budgeting (IJCAI '24)  
with N. Boehmer, P. Faliszewski, L. Janeczko, D. Peters, G. Pierczyński, P. Skowron, S. Szufa
- [2] Participatory Budgeting Project Strength via Candidate Control (IJCAI '25)  
with P. Faliszewski, L. Janeczko, D. Knop, J. Pokorný, M. Śluszniak, K. Sornat
- [3] The Complexity of Fair Division of Indivisible Items with Externalities (AAAI '24)  
with A. Deligkas, E. Eiben, V. Korchemna
- [4] Hedonic Diversity Games: A Complexity Picture With More Than Two Colors (AAAI '22 and AIJ) with R. Ganian, T. Hamm, D. Knop, O. Suchý
- [5] Host Community Respecting Refugee Housing (AAMAS '23)  
with D. Knop



DANIEL SCHOEPLIN ([Homepage](#), [CV](#), [Google Scholar](#))

**Thesis:** Designing and Analyzing Clock Auctions ('23)

**Advisor:** Vasilis Gkatzelis, Drexel University

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

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

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

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

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

#### Representative Papers:

- [1] Optimal Deterministic Clock Auctions and Beyond (ITCS '22)  
with G. Christodoulou and V. Gkatzelis
- [2] Bayesian and Randomized Clock Auctions (EC '22, OR '25)  
with M. Feldman, N. Gravin, and V. Gkatzelis
- [3] Deterministic Budget-Feasible Clock Auctions (SODA '22, OR (forthcoming))  
with E. Balkanski, P. Garimidi, V. Gkatzelis, and X. Tan
- [4] Clock Auctions Augmented with Unreliable Advice (SODA '25)  
with V. Gkatzelis and X. Tan
- [5] Multi-Parameter Mechanisms for Consumer Surplus Maximization (STOC '25)  
with T. Ezra and A. Shaulker

EKLAVYA SHARMA ([Homepage](#), [CV](#), [Google Scholar](#))

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

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

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

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

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

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

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

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

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

#### Representative Papers:

- [1] Improving Approximation Guarantees for Maximin Share (EC'24)  
with Hannaneh Akrami, Jugal Garg, Setareh Taki.
- [2] New Fairness Concepts for Allocating Indivisible Items (IJCAI'23)  
with Ioannis Caragiannis, Jugal Garg, Nidhi Rathi, Giovanna Varricchio.
- [3] Simplification and Improvement of MMS Approximation (IJCAI'23)  
with Hannaneh Akrami, Jugal Garg, Setareh Taki.
- [4] Exploring Relations among Fairness Notions in Discrete Fair Division (under submission, preprint on ArXiv), with Jugal Garg.

SUHO SHIN ([Homepage](#), [CV](#))

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

**Advisor:** MohammadTaghi Hajiaghayi, University of Maryland

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

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

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

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

#### Representative Papers:

- [1] The Contest Behind the Feed: Optimal Contest for Recommender Systems (working paper, job market paper) with N. Golrezaei, M. Hajiaghayi
- [2] Delegated Choice with Combinatorial Constraints (EC 2025, under review at *Operations Research*) with K. Banihashem, M. Hajiaghayi, P. Krysta
- [3] Delegation with Costly Inspection (EC 2025) with M. Hajiaghayi, P. Krysta, M. Mahdavi
- [4] Gains-from-Trade in Bilateral Trade with a Broker (SODA 2025) with I. Hajiaghayi, M. Hajiaghayi, G. Peng
- [5] Tokenized Bandit for LLM Decoding and Alignment (ICML 2025) with M. Hajiaghayi, H. Xu, C. Yang

ANA-ANDREEA STOICA ([Homepage](#), [CV](#), [Google Scholar](#))

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

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

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

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

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

#### Representative Papers:

- [1] Policy Design in Long-Run Welfare Dynamics (ICLR'25)  
with Jiduan Wu, Rediet Abebe, and Moritz Hardt
- [2] Causal Inference from Competing Treatments (ICML'24)  
with Vivian Y. Nastl and Moritz Hardt
- [3] The Fairness-Quality Trade-off in Clustering (NeurIPS'24)  
with Rashida Hakim, Mihalis Yannakakis, and Christos Papadimitriou

YIFAN WANG ([Homepage](#), [CV](#))

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

**Advisor:** Sahil Singla, Georgia Tech

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

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

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

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

#### Representative Papers:

- [1] Single-Sample and Robust Online Resource Allocation (STOC'25)  
with Rohan Ghuge and Sahil Singla
- [2] Learning Optimal Posted Prices for a Unit-Demand Buyer (EC'25)  
with Yifeng Teng
- [3] Combinatorial Philosopher Inequalities (SODA'26)  
with Enze Sun and Zhihao Gavin Tang
- [4] Online Stochastic Matching with Unknown Arrival Order: Beating 0.5 against the Online Optimum (STOC'25) with Enze Sun and Zhihao Gavin Tang

WENTAO WENG ([Homepage](#), [CV](#))

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

**Advisor:** Daniel Freund, Thodoris Lykouris, MIT

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

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

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

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

#### Representative Papers:

- [1] Learning to Defer in Congested Systems: The AI-Human Interplay (Major Revision at Operations Research), with T. Lykouris.
- [2] Scheduling with Uncertain Holding Costs and its Application to Content Moderation (Major Revision at Management Science), with C. Gocmen, T. Lykouris, and D. Sinha.
- [3] The Transient Cost of Learning in Queueing Systems (NeurIPS'23, R&R at Operations Research), with D. Freund and T. Lykouris.
- [4] Efficient Decentralized Multi-agent Learning in Asymmetric Bipartite Queueing Systems (COLT'22, Operations Research'24), with D. Freund and T. Lykouris.
- [5] Regulating Wait-Driven Requests in Queues (EC'25), with D. Freund and D. Hausman.
- [6] The Dedicated Docket in U.S. Immigration Courts: An analysis of fairness and efficiency properties (EC'24, Major Revision at MSOM), with D. Freund.
- [7] Group fairness in dynamic refugee assignment (EC'23, Major Revision at Operations Research), with D. Freund, T. Lykouris, E. Paulson, and B. Sturt.

NICHOLAS WU ([Homepage](#), [CV](#))

**Thesis:** Essays in Microeconomic Theory ('26)

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

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

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

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

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

#### **Representative Papers:**

- [1] Solving Problems of Unknown Difficulty? (Job Market Paper)
- [2] From Doubt to Devotion: Trials and Learning-Based Pricing (accepted, Journal of Political Economy) with T. Gan
- [3] How Do Digital Advertising Auctions Impact Product Prices? (Review of Economic Studies 92 (4), 2330-2358) with D. Bergemann and A. Bonatti
- [4] Bidding with Budgets: Data-Driven Bid Algorithms in Digital Advertising (International Journal of Industrial Organization, 103172) with D. Bergemann and A. Bonatti



SHENGWEI XU ([Homepage](#), [CV](#))

**Thesis:** Aligned Information Elicitation for Text ('25)

**Advisor:** Grant Schoenebeck, University of Michigan

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

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

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

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

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

### Representative Papers:

- [1] Benchmarking LLMs’ Judgments with No Gold Standard. (ICLR 2025)  
with Y. Lu, Y. Zhang, Y. Kong, and G. Schoenebeck
- [2] Eliciting Informative Text Evaluations with Large Language Models. (EC’24)  
with Y. Lu, Y. Zhang, Y. Kong, and G. Schoenebeck
- [3] Stochastically Dominant Peer Prediction. (NeurIPS 2025)  
with Y. Zhang, D. Pennock, and G. Schoenebeck
- [4] Ad Insertion in LLM-generated Responses. (Under-review, presented at the Workshop on Frontiers of Online Advertising at EC’25)  
with Z. Chen, Z. Huang, G. Schoenebeck, and X. Deng



AVIV YAISH ([Homepage](#), [CV](#), [Google Scholar](#))

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

**Advisor:** Aviv Zohar, The Hebrew University

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

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

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

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

#### Representative Papers:

- [1] Transaction Fee Mechanisms Robust to Welfare-Increasing Collusion (Minor Revision at Games and Economic Behavior, '25) with Y. Gafni
- [2] Barriers to Collusion-resistant Transaction Fee Mechanisms (EC'24) with Y. Gafni
- [3] Speculative Denial-of-Service Attacks in Ethereum (USENIXSEC'24) with K. Qin, L. Zhou, A. Zohar, and A. Gervais
- [4] Uncle Maker: (Time)Stamping Out the Competition in Ethereum (CCS'23 Distinguished Paper Award) with G. Stern, and A. Zohar
- [5] Blockchain Stretching & Squeezing: Manipulating Time for Your Best Interest (EC'22) with S. Tochner, and A. Zohar

KUNHE YANG ([Homepage](#), [CV](#), [Google Scholar](#))

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

**Advisor:** Nika Haghtalab, UC Berkeley

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

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

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

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

#### Representative Papers:

- [1] Distortion of AI Alignment: Does Preference Optimization Optimize for Preferences? (NeurIPS 25) with P. Gözl and N. Haghtalab
- [2] Platforms for Efficient and Incentive-Aware Collaboration (SODA 25) with N. Haghtalab and M. Qiao
- [3] Truthfulness of Calibration Measures (NeurIPS 24) with N. Haghtalab, M. Qiao, and E. Zhao
- [4] Calibrated Stackelberg Games: Learning Optimal Commitments Against Calibrated Agents (NeurIPS 23, Spotlight) with N. Haghtalab and C. Podimata

KONSTANTIN ZABARNYI ([Homepage](#), [CV](#), [Google Scholar](#))

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

**Advisors:** Dirk Bergemann, Yang Cai

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

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

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

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

#### Representative Papers:

- [1] Regret-Minimizing Bayesian Persuasion (EC'21 and Games and Economic Behavior) with Y. Babichenko, I. Talgam-Cohen, and H. Xu
- [2] A Random Dictator Is All You Need (EC'23 and American Economic Journal: Microeconomics) with I. Arieli, Y. Babichenko, and I. Talgam-Cohen
- [3] Algorithmic Cheap Talk (EC'24 and under review at Operations Research) with Y. Babichenko, I. Talgam-Cohen, and H. Xu
- [4] Persuasion Gains and Losses from Peer Communication (under review at Theoretical Economics) with T.T. Kerman, and A.P. Tenev

WENXIN ZHANG ([Homepage](#), [CV](#), [Google Scholar](#))

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

**Advisor:** Santiago Balseiro and Will Ma, Columbia University

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

### Research Summary:

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

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

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

### Representative Papers:

- [1] Dynamic Pricing for Reusable Resources: The Power of Two Prices (*Operations Research*, 2025) with S. Balseiro and W. Ma
- [2] Feature-Based Dynamic Matching (*Operations Research*, forthcoming; *EC*, 2023) with Y. Chen, Y. Kanoria, and A. Kumar
- [3] Distributed Load Balancing with Workload-Dependent Service Rates (*EC*, 2025; under review at *Operations Research*) with S. Balseiro, R. Kleinberg, V. Mirrokni, B. Sivan, and B. Wydrowski
- [4] Tail-Optimized Caching for LLM Inference (*NeurIPS*, 2025) with Y. Li, C. Moallemi, and T. Peng

YICHI ZHANG ([Homepage](#), [CV](#))

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

**Advisor:** Grant Schoenebeck, University of Michigan

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

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

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

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

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

#### Representative Papers:

- [1] Stochastically Dominant Peer Prediction. (NeurIPS'25)  
with S. Xu, D. Pennock, and G. Schoenebeck
- [2] Evaluating LLM-corrupted Crowdsourcing Data Without Ground Truth.  
(NeurIPS'25) with J. Pang, Z. Zhu, and Y. Liu
- [3] Mixture of Complementary Agents for Robust LLM Ensemble (Working paper)  
with K. Lu, Y. Zhang, J. Gao, L. Xia, and F. Yu
- [4] A System-Level Analysis of Conference Peer Review (EC'22, revision at OR)  
with F. Yu, G. Schoenebeck, and D. Kempe

## Index

- algorithm design
  - Vishwa Prakash HV, 19
- optimization
  - Daniel Schoepflin, 35
- active learning
  - Abdellah Aznag, 8
- AI alignment
  - Kunhe Yang, 44
- algorithmic game theory
  - Vishwa Prakash HV, 19
  - Yoav Kolumbus, 21
  - Šimon Schierreich, 34
  - Kunhe Yang, 44
- algorithmic mechanism design
  - Federico Bobbio, 9
- approximation algorithms
  - Prasanna Ramakrishnan, 30
  - Eklavya Sharma, 36
- auctions
  - Giannis Fikioris, 14
  - Yoav Kolumbus, 21
  - Shiri Ron, 32
- bounded rationality
  - Daniel Schoepflin, 35
- calibration
  - Maxwell Fishelson, 15
- causal inference
  - Ana-Andreea Stoica, 38
- cheap talk
  - Konstantin Zabarnyi, 45
- collective decision making
  - Šimon Schierreich, 34
- communication complexity
  - Shiri Ron, 32
- computational complexity
  - Abheek Ghosh, 16
- computational social choice
  - Prasanna Ramakrishnan, 30
- conflicts with multiple battlefields
  - Stanisław Kaźmierowski, 20
- congestion
  - Shaul Rosner, 33
- contest theory
  - Abheek Ghosh, 16
- contract design
  - Sumit Goel, 17
  - Tomasz Ponitka, 29
  - Shaul Rosner, 33
- cooperation
  - Kiran Rokade, 31
- data economics
  - Pooja Kulkarni, 22
  - Aniket Murhekar, 27
- data markets
  - Eklavya Sharma, 36
- data-driven & online algorithms
  - Alireza AmaniHamedani, 7
- delegation
  - Suho Shin, 37
- differential privacy
  - Yuxin Liu, 25
- dynamic games
  - Nicholas Wu, 41
- dynamic resource allocation
  - Wenxin Zhang, 46
- dynamics
  - Yoav Kolumbus, 21
- equilibrium computation
  - Stanisław Kaźmierowski, 20
  - Kiran Rokade, 31
- evaluation with AI in the loop
  - Yichi Zhang, 47
- fair division
  - Vishwa Prakash HV, 19
  - Pooja Kulkarni, 22
  - Aniket Murhekar, 27
  - Tomasz Ponitka, 29
  - Eklavya Sharma, 36
- fairness
  - Robin Bowers, 10
- game theory
  - Sumit Goel, 17
  - Stanisław Kaźmierowski, 20
  - Prasanna Ramakrishnan, 30

- Kiran Rokade, 31
- Shaul Rosner, 33
- Ana-Andreea Stoica, 38
- generative AI
  - Suho Shin, 37
- generative AI efficiency
  - Wenxin Zhang, 46
- human-AI collaboration
  - Natalie Collina, 11
- human-AI interaction
  - Ce Li, 24
- incentive-aware machine learning
  - Kunhe Yang, 44
- incentives in machine learning
  - Aniket Murhekar, 27
- information design
  - Ce Li, 24
  - Nicholas Wu, 41
  - Konstantin Zabarnyi, 45
- information elicitation
  - Shengwei Xu, 42
  - Yichi Zhang, 47
- information geometry
  - Abdellah Aznag, 8
- large language models
  - Shengwei Xu, 42
- learning
  - Yoav Kolumbus, 21
- learning dynamics
  - Abheek Ghosh, 16
- learning in games
  - Giannis Fikioris, 14
- load balancing
  - Kalen Patton, 28
- machine learning
  - Ce Li, 24
  - Ana-Andreea Stoica, 38
  - Wentao Weng, 40
- market design
  - Alireza AmaniHamedani, 7
  - Jeff Decary, 12
  - Craig Fernandes, 13
  - Soonbong Lee, 23
- Gary Qiurui Ma, 26
- markets
  - Yoav Kolumbus, 21
- matching
  - Alireza AmaniHamedani, 7
  - Federico Bobbio, 9
  - Robin Bowers, 10
  - Sumit Goel, 17
  - Vishwa Prakash HV, 19
  - Gary Qiurui Ma, 26
- mechanism design
  - Robin Bowers, 10
  - Sumit Goel, 17
  - Ce Li, 24
  - Yuxin Liu, 25
  - Shiri Ron, 32
  - Suho Shin, 37
  - Yifan Wang, 39
  - Nicholas Wu, 41
  - Konstantin Zabarnyi, 45
- multi-agent LLM
  - Yichi Zhang, 47
- multi-armed bandits
  - Abdellah Aznag, 8
- networks
  - Yoav Kolumbus, 21
  - Yuxin Liu, 25
  - Kiran Rokade, 31
- online algorithms
  - Pooja Kulkarni, 22
  - Kalen Patton, 28
  - Yifan Wang, 39
  - Wenxin Zhang, 46
- online and matching platforms
  - Soonbong Lee, 23
- online fair allocation
  - Giannis Fikioris, 14
- online learning
  - Natalie Collina, 11
  - Giannis Fikioris, 14
  - Maxwell Fishelson, 15
- online platforms
  - Suho Shin, 37
- operations
  - Yoav Kolumbus, 21

- operations for social good
  - Soonbong Lee, 23
- operations management
  - Craig Fernandes, 13
- optimization
  - Federico Bobbio, 9
  - Abheek Ghosh, 16
- Pandora’s box
  - Robin Bowers, 10
- peer prediction
  - Shengwei Xu, 42
- platform and gig economy
  - Gary Qiurui Ma, 26
- portfolio
  - Jeff Decary, 12
- practical mechanism design
  - Daniel Schoepflin, 35
- prediction markets
  - Jeff Decary, 12
- pricing
  - Gary Qiurui Ma, 26
  - Yifan Wang, 39
- prophet inequalities
  - Kalen Patton, 28
- queueing
  - Alireza AmaniHamedani, 7
  - Wentao Weng, 40
- regret minimization
  - Maxwell Fishelson, 15
- resource allocation
  - Kalen Patton, 28
- revenue management
  - Craig Fernandes, 13
- robustness
  - Konstantin Zabarnyi, 45
- social choice
  - Aniket Murhekar, 27
  - Šimon Schierreich, 34
  - Eklavya Sharma, 36
- social networks
  - Ana-Andreea Stoica, 38
- spectrum markets
  - Federico Bobbio, 9
- sports analytics
  - Jeff Decary, 12
- submodular optimization
  - Pooja Kulkarni, 22
- sustainable governance
  - Wentao Weng, 40
- trustworthy AI
  - Natalie Collina, 11
- value of information
  - Robin Bowers, 10