# SIGecom Job Market Candidate Profiles 2022

This is the seventh annual collection of profiles of the junior faculty job market candidates of the SIGecom community. The twenty four candidates for 2022 are listed alphabetically and indexed by research areas that define the interests of the community. The candidates can be contacted individually or via the moderated mailing list ecom-candidates2022@acm.org.

–Vasilis Gkatzelis and Jason Hartline


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

Contents

BEN ABRAMOWITZ (Homepage, Google Scholar)

**Thesis:** Deciding Who, What, Why, and How: Aggregating Preferences Over Agents, Alternatives, Axioms, and Rules

**Advisor:** Elliot Anshelevich, Rensselaer Polytechnic Institute

**Brief Biography:** Ben Abramowitz is a Computing Innovation Fellow (CRA/CCC) and postdoctoral researcher with Dr. Nicholas Mattei at Tulane University. Previously, Ben completed his PhD advised by Dr. Elliot Anshelevich at RPI. During that time, Ben held a visiting student position at the Weizmann Institute working with the group of Ehud Shapiro (2020), and an internship at IBM Research (2018). Ben's research interests continue to lie at the intersection of his three undergraduate majors: Math, Economics, and Computer Science.

**Research Summary:** I am primarily interested in the study of multi-agent systems, social choice, and the design of mechanisms/algorithms with economic and governmental applications. In general, the analytical tools and methods I most frequently use draw from game theory, graph theory, approximation algorithms, and complexity theory.

My research involves two complementary lines of work. The first involves the aggregation of agent preferences over other agents (e.g. network formation [1], coalition formation [1,4], proxy voting [2]) and/or generic alternatives (e.g. voting [3], project selection [4]). In these settings agents may be networked [1], lie in a common metric space [3], or select others to represent them [2]. For instance, in [1] we demonstrate that for a surprising broad class of problems we define involving networked agents, ordinal information derived from agent utilities is enough to closely approximate optimal social welfare. In [3] we prove that when agents and alternatives lie in a common metric space, a small amount of cardinal information added to ordinal preferences can greatly improve the best achievable approximations in voting (metric distortion). We also introduce the novel solution concept of *ideal candidate distortion*.

The second thread of my work involves grappling with agent preferences over axioms and rules (e.g. constitutional amendments [5], blockchain forking [4]). A central idea here is that the agents being modeled might care about any of the same things we as mechanism designers might. These problems lie somewhere between social choice and self-governance/self-organization.

**Representative Papers:**

[1] Utilitarians Without Utilities: Maximizing Social Welfare for Graph Problems Using Only Ordinal Preferences (AAAI 2018) with E. Anshelevich

[2] Flexible Representative Democracy: An Introduction with Binary Issues (IJCAI 2019) with N. Mattei

[3] Awareness of Voter Passion Greatly Improves the Distortion of Metric Social Choice (WINE 2019) with E. Anshelevich and W. Zhu

[4] Democratic Forking: Choosing Sides with Social Choice (ADT 2021) with E. Elkind, D. Grossi, E. Shapiro and N. Talmon

[5] How to Amend a Constitution? Model, Axioms, and Supermajority Rules (ADT 2021) with E. Shapiro and N. Talmon

GIANLUCA BRERO (Homepage, CV, Google Scholar)

**Thesis:** Machine Learning-powered Iterative Combinatorial Auctions

**Advisor:** Sven Seuken, University of Zurich

**Brief Biography:** Gianluca Brero is a postdoctoral researcher at Harvard University hosted by David Parkes. He is a recipient of the Early Postdoc-Mobility Fellowship (2020-2021) from the Swiss National Science Foundation. Gianluca received his Ph.D. in Computer Science from the University of Zurich where he was advised by Sven Seuken. He has also worked at Microsoft Research with Sébasien Lahaie. Gianluca holds a double Master's Degree in Mathematics from Turin Polytechnic and Milan Polytechnic and a Bachelor's Degree in Mathematics from Turin Polytechnic. At Turin Polytechnic, he was awarded the Vallauri Prize as the top student in Mathematics who graduated in 2014.

**Research Summary:** Despite great achievements in optimal economic design, many economic mechanisms used in real-world markets are still causing major inefficiencies. This is often the case for *indirect economic mechanisms*, which interact with participants multiple times and often perform suboptimally in the complex settings where they are applied. My goal is to leverage the most recent advances in machine learning to improve the design of indirect economic mechanisms.

I first focused on iterative combinatorial auctions, which are currently used to allocate resources in multi-billion dollar domains, like the sale of spectrum licenses. The key challenge when designing these auctions is that participants value exponentially many bundles of items, and bidding optimally may be infeasible. I introduce Bayesian approaches that leverage auction data to give bidding guidance via suitable ask prices. The resulting mechanisms dramatically reduce the number of auction rounds to clear the market. I also focus on auctions that use machine learning to directly identify the bundles on which participants should bid. In realistic settings, the resulting mechanism is demonstrated to outperform the Combinatorial Clock Auction, which is currently used in real-world spectrum auctions.

Despite their complexity, there is a considerable theoretical literature to guide the design of combinatorial auctions. This is not true when designing mechanisms for real-world settings such as the Uber or Amazon market platforms. The research I am pursuing during my postdoc makes use machine learning to derive optimal indirect mechanisms where there is little guidance from existing theory. In a first paper, I am working with collaborators to pioneer the use of reinforcement learning to design indirect economic mechanisms, focusing on settings where bidding is straightforward. We also use these insights to propose a reinforcement learning approach able to derive mechanism policies in a Stackelberg equilibrium with the induced agent behavior.

**Representative Papers:**

[1] Fast Iterative Combinatorial Auctions via Bayesian Learning (AAAI'19)
with S. Lahaie and S. Seuken

[2] Machine Learning-powered Iterative Combinatorial Auctions (arXiv)
with B. Lubin and S. Seuken

[3] Reinforcement Learning of Sequential Price Mechanisms (AAAI'21)
with A. Eden, M. Gerstgrasser, D. C. Parkes, and D. Rheingans-Yoo

MODIBO CAMARA (Homepage, CV)

**Thesis:** Complexities in Economic Theory

**Advisors:** Jason Hartline and Eddie Dekel, Northwestern University

**Brief Biography:** Modibo Camara is an economics Ph.D. candidate at Northwestern University and member of the Online Markets Lab. At Northwestern, he is co-advised by Jason Hartline (computer science) and Eddie Dekel (economics). He has worked as an intern at Microsoft Research, the Federal Reserve Board, and the Commodity Futures Trading Commission. Prior to graduate school, he earned a B.A. in economics and mathematics from the University of Pennsylvania.

**Research Summary:** I am an economic theorist focused on developing more credible models of behavior. In my work on bounded rationality, I try to understand what makes rational decision-making hard, using techniques developed in theoretical computer science and statistical learning. Then I explore the implications for predicted behavior and economic policy.

In my job market paper [1], I incorporate computational constraints into the theory of choice. I impose polynomial-time tractability as an axiom, and use the resulting framework to better understand common behavioral heuristics and violations of rationality. My dichotomy theorems show that choices satisfying rationality and tractability axioms correspond to forms of choice bracketing, a heuristic observed in lab experiments. Then I establish a choice trilemma: for many objective functions, choices can be rational and optimal, tractable and approximately optimal, or rational and tractable, but not all three. This suggests computationally-constrained agents may be better off violating standard rationality axioms.

In another line of work, I apply bounded rationality to mechanism design. Specifically, I study mechanism design where both the policymaker and agent are statistical learners. The goal is to relax traditional common knowledge assumptions and replace them with (i) a common history or (ii) a common dataset.

In [2], we consider a policymaker and single agent with a common history. The policymaker and agent interact repeatedly over time, with a hidden state of nature that is revealed after each period. Adapting ideas from adversarial online learning, we develop simple calibrated policies that ensure bounded regret, relative to the best-in-hindsight static policy. They require novel behavioral assumptions that capture concepts like "rationality" or "unpredictability" in a prior-free sense. This work was featured at the 2021 "Highlights Beyond EC" plenary session.

In [3], I again consider a policymaker and a single agent, but with a common i.i.d. dataset. The core idea is simple: if the available data convincingly demonstrates some fact about the world, the agent should believe that fact. Otherwise, her beliefs are left unspecified. I formalize this assumption and develop data-driven policies with good guarantees for the policymaker. I find that policies that are too complex (in a precise sense) can lead to unpredictable behavior by the agent.

**Representative Papers:**

[1] Computationally Tractable Choice (job market paper)
[2] Mechanisms for a No-Regret Agent: Beyond the Common Prior (FOCS 2020) with J. Hartline and A. Johnsen
[3] Mechanism Design with Common Data (working paper)

CHEN CHEN (Homepage, CV, Google Scholar)

**Thesis:** Performance Analysis in Large-Scale Stochastic Dynamic Programs

**Advisor:** Santiago Balseiro (Columbia) and David Brown (Duke)

**Brief Biography:** Chen is currently a postdoctoral scholar in the Operations Management area at Booth School of Business, University of Chicago, where he works with Prof. Ozan Candogan and Prof. Rad Niazadeh. He received his Ph.D. from Fuqua School of Business, Duke University, under the supervision of Prof. Santiago Balseiro and Prof. David Brown. His work has been awarded the first place of 2019 INFORMS Revenue Management and Pricing Student Paper Prize. He interned at Uber's Marketplace Optimization group in the summer of 2019, working on matching problems.

**Research Summary:** I am broadly interested in developing simple approximation policies – which are easy to compute, implement and interpret, and attain good performance guarantees – to improve the operations of modern marketplaces. Such policies are valued in practice and provide critical insights into the operational problems we study. On the practical side, I have followed this research guideline to address various operational challenges, including: (i) optimal design of experiments to collect information efficiently [1], (ii) dynamic pricing for ride-sharing platforms [2], and (iii) efficient algorithms for stochastic job scheduling [3].

Specifically, in [1] I develop near-optimal cluster/community-level experiments to minimize the variance of an unbiased Horvitz-Thompson estimator for the total market effect, against the worst-case value of the potential outcomes. The optimal community-level assignment is computationally expensive to solve, and can be difficult to implement due to the complicated correlation structure. Thus motivated, we develop a family of simple independent block randomization (IBR) experiments that are easy to compute and interpret, and can achieve much of the benefit from the optimal (correlated) randomized assignment. Specifically, the IBR experiments constitute a good approximation ratio guarantee for any problem instance, and are asymptotically optimal when the number of communities grows large.

In [2], I develop near-optimal dynamic pricing policies for relocating a limited number of resources over a large network of locations (e.g., as in ridesharing or other shared vehicle systems), based on Lagrangian relaxation of the capacity constraint at a subset of central "hub" locations. In [3], I study a stochastic scheduling problem with unrelated machines and use a novel information relaxation duality approach to show a simple static policy is asymptotically optimal in the regime of many jobs. Variations of this problem find broad applications in cloud computing, production systems, and healthcare system management among others.

**Representative Papers:**

[1] Near-Optimal Experimental Design for Networks: Independent Block Randomization (Working Paper) with O. Candogan and R. Niazadeh

[2] Dynamic Pricing of Relocating Resources in Large Networks (Management Science 2021) with S. Balseiro and D. Brown

[3] Static Routing in Stochastic Scheduling: Performance Guarantees and Asymptotic Optimality (Operations Research 2018) with S. Balseiro and D. Brown

MICHAEL CURRY (Homepage, CV, Google Scholar)

**Thesis:** Learning and Robustness With Applications To Mechanism Design

**Advisor:** John Dickerson & Tom Goldstein, University of Maryland

**Brief Biography:** Michael Curry is a PhD student at the University of Maryland. During the summer of 2020 he was a research intern in the G-RIPS program at UCLA IPAM, and during the summer of 2021 he was a research intern at Salesforce Research. His research lies at the intersection of mechanism design and machine learning, with a particular focus on the nascent set of tools provided by differential economics.

**Research Summary:** Much of my research has focused on differentiable economics. It's long been observed that mechanism design given samples from the valuation distribution is in essence a learning problem. And mechanisms are just functions, so why not represent them using the computational tools and rich function approximators that modern deep learning provides? This is the pitch of differentiable economics, a recent thread of work introduced in a number of papers, notably "Optimal Auctions Through Deep Learning" by Dütting et al. My work aims to extend this approach by adapting techniques from modern machine learning to improve the robustness and expressiveness of learned mechanisms.

In one thread of work, we apply tools from adversarially robust machine learning to mechanism design. We observe that strategyproofness of learned mechanisms is similar to adversarial robustness, and find that by applying techniques from adversarial robustness to modified auction network architectures, we can exactly certify the maximum violation of strategyproofness for any given valuation profile [1]. In ongoing work, we extend these results to provide global certificates over all valuations.

As a mentor to a number of extremely talented undergraduate students, I've also helped to extend the differentiable economics approach to incorporate notions of fairness – both formally defined and more recently, learned [2]. I've additionally made contributions to research on anticipating and preventing match failures in the kidney exchange system (which is now being applied as part of a pilot experiment) [3] and on adversarial robustness in non-mechanism-design settings.

**Representative Papers:**

[1] Certifying Strategyproof Auction Networks (NeurIPS 2020)
with M. Curry, P.-y. Chiang, T. Goldstein, and J. Dickerson

[2] PreferenceNet: Encoding Human Preferences in Auction Design with Deep Learning (NeurIPS 2021) with N. Peri, M. Curry, S. Dooley, and J. Dickerson

[3] Improving Policy-Constrained Kidney Exchange via Pre-Screening
(NeurIPS 2020) with D. McElfresh, M. Curry, T. Sandholm, and J. Dickerson

ALON EDEN (Homepage, CV, Google Scholar)

**Thesis:** Correlation, Coordination, Competition and Pricing

**Advisor:** Michal Feldman and Amos Fiat, Tel Aviv University

**Brief Biography:** I am a postdoc in Harvard's EconCS group hosted by Yiling Chen and David Parkes. I was awarded the Michael B. Maschler Prize of the Israeli Chapter of the Game Theory Society for the best PhD thesis of 2019. I was also awarded Best Paper at SAGT'17 and Best Paper with Lead Student Author at EC'19. I received honorable mention for the Best Presentation by a Student or Postdoctoral Researcher award at EC'19.

**Research Summary:** A process that started with the rise of the internet, has immensely sped-up during the COVID-19 pandemic — brick and mortar business gave way to an economy that is almost completely digital. This new economy gives rise to many challenges and opportunities. In my research, I tackle problems related to this new economy using algorithmic, game theoretic, and machine learning tools.

The Interdependent Values model, whose significance has been recognized by the 2020 Nobel prize in Economics, captures settings such as online marketplaces, where buyers only obtain partial information about the item being sold, and can benefit from information obtained by other bidders. In papers in EC'18, EC'19[2] and AAAI'21, we make big strides in devising practical mechanisms in this model that relax the rigid conditions, and bypass previous impossibilities.

In my research, I also advance the study of auction design in complex multidimensional settings. In EC'17 [3], we show that by augmenting the market with more buyers, which can be done in online marketplaces using ad campaigns, simple auctions outperform the revenue from the optimal design in the original market. In an OR'21 paper, we use parameterized analysis to devise simple mechanisms that give near-optimal revenue when considering items that have complements.

I also study online resource allocation problems, where strategic agents use the platform in an asynchronous manner. In SODA'15, EC'16 [4], and other papers, we show how simple pricing schemes can achieve optimal (or near optimal) performance in combinatorial markets, and other non-auction settings.

In a recent line of research, I use RL tools in order to devise simple and practical optimal mechanisms. In AAAI'21 [1], we use RL to devise optimal sequential prices mechanisms. In a working paper, we study other simple auction formats, where we use a novel RL framework to formalize the problem as finding an optimal leader strategy in a Stackelberg equilibrium.

**Representative Papers:**

[1] Reinforcement Learning of Sequential Price Mechanisms (AAA'21)
with G. Brero, M. Gerstgrasser, D. Parkes, and D. Rheingans-Yoo

[2] Combinatorial Auctions with Interdependent Valuations: SOS to the Rescue (EC'19) with M. Feldman, A. Fiat, K. Goldner, and A. Karlin

[3] The Competition Complexity of Auctions (EC'17)
with M. Feldman, O. Friedler, I. Talgam-Cohen, and M. Weinberg

[4] The Invisible Hand of Dynamic Market Pricing (EC'16)
with V. Cohen-Addad, M. Feldman, A. Fiat

JESSIE FINOCCHIARO (Homepage, CV, Google Scholar)

**Thesis:** Designing Convex Surrogate Loss Functions for General Prediction Tasks

**Advisor:** Dr. Rafael Frongillo, CU Boulder

**Brief Biography:** Jessie is a final-year PhD student at the University of Colorado Boulder working with Dr. Rafael Frongillo, and is a 2019 NSF Graduate Research Fellow. Her research focuses on the design of proper scoring rules for general prediction tasks, and is interested in understanding how the choice of scoring rule affects algorithmic decision-making and recommendations. She holds a MS in Computer Science from CU Boulder and BS in Mathematics and Computer Science from Florida Southern College.

**Research Summary:** Algorithms are ubiquitously used to inform human decision-making, often designed to predict by minimizing the average error over a dataset by assigning a loss by comparing predictions to observed outcomes. Depending on the loss function, the prediction minimizing the loss should be some summary statistic, or *property* of the underlying data distribution, such as the median, variance, or expected value. In practice, researchers most often start with a loss function and later observe what property corresponds to minimizing that loss; conversely, I study property elicitation, which starts with a property to learn, and designs loss functions whose minimizers correspond to the property value. In [3], we show that the condition of statistical consistency implies indirect property elicitation, demonstrating the use of elicitation as a handy tool to construct lower bounds on the dimension of consistent convex surrogates (e.g., [2]).

In [1], we present a framework for studying the design of piecewise linear and convex surrogate loss functions that correspond to discrete prediction tasks, whose original target function is computationally hard to optimize. One example of an embedding in action is the use of hinge loss as a surrogate for 0-1 loss, though the framework extends to general prediction tasks such as structured prediction, top-$k$ classification, and ordered partitions.

Finally, I am interested in the impact of target prediction tasks and constraints on algorithmic decision-making. In [4], we compare some modelling assumptions, constraints, and objectives commonly found in the mechanism design and machine learning literatures, drawing some of the strength and weaknesses of each, along with lessons transferred across the two fields.

**Representative Papers:**

[1] An Embedding Framework for Consistent Polyhedral Surrogates (NeurIPS 2019) with R. Frongillo and B. Waggoner

[2] Embedding Dimension of Polyhedral Losses (COLT 2020) with R. Frongillo and B. Waggoner

[3] Unifying Lower Bounds on Prediction Dimension of Consistent Convex Surrogates (NeurIPS 2021) with R. Frongillo and B. Waggoner

[4] Bridging Machine Learning and Mechanism Design towards Algorithmic Fairness (FAccT 2021) with R. Maio, F. Monachou, G. Patro, M. Raghavan, A. Stoica, and S. Tsirtsis

YUAN GAO (Homepage, CV)

**Thesis:** Optimization Models and Methods for Equilibrium Computation and Mechanism Design

**Advisor:** Christian Kroer, Columbia University

**Brief Biography:** Yuan is a 5th-year PhD student in Operations Research at Columbia University. Previously, he obtained a Bachelor's Degree in Applied Mathematics from National University of Singapore.

**Research Summary:** I develop optimization models and methods for equilibrium computation and mechanism design while contributing to general optimization algorithms and analysis.

In [1], I studied optimization-based methods for computing market equilibria (ME) under different practically used utility functions. Using proportionality of ME, I showed that their equilibrium-capturing convex programs (ECCP), after simple reformulations, satisfy a weakened "relaxed convexity" (Proximal-PŁ) condition. Hence, the proximal gradient method applied to them achieves linear convergence. This is the first linear convergence result for computing market equilibria via first-order methods. Furthermore, I showed that proximal gradient with an economical linesearch scheme applied to general Proximal-PŁ problems achieves linear convergence as well. For quasilinear utilities, I proposed a new Shmyrev-type ECCP, for which mirror descent achieves $O(1/T)$ last-iterate convergence. This yields a new form of *Proportional Response* dynamics.

In [2], I extended the concepts of linear Fisher markets and ME to a finite measurable item space and proposed infinite-dimensional Eisenberg-Gale-type (EG) convex programs. Then, I showed that their solutions correspond to ME and that various structural properties of ME in the finite case also hold in the infinite-dimensional case. When buyers have piecewise linear valuations on a closed interval, I showed that the EG convex program exhibits a finite-dimensional convex conic reformulation, which can be solved efficiently using off-the-shelf interior-point optimization software. I also showed that it can be solved using the ellipsoid method, which gives the first polynomial-time cake-cutting algorithm for piecewise linear valuations. Finally, I extended most of the results to the case of quasilinear utilities.

In [3], I developed an online allocation and pricing mechanism by applying dual averaging on a reformulated ECCP. It is simple, scalable, and interpretable as first-price auctions with pacing. I showed that, under stochastic item arrivals, the buyers' pacing multipliers, utilities and expenditures converge to their respective equilibrium quantities of an underlying static market. It is also the first algorithm that achieves these guarantees in a stochastic setting. As such, the realized allocations of the mechanism satisfy desirable fairness and efficiency asymptotically.

**Representative Papers:**

[1] First-Order Methods for Large-Scale Market Equilibrium Computation (NeurIPS 2020) with C. Kroer

[2] Infinite-Dimensional Fisher Markets and Tractable Fair Division (minor revision, Operations Research; AAAI 2021) with C. Kroer

[3] Online Market Equilibrium with Application to Fair Division (Accepted, NeurIPS 2021) with C. Kroer and A. Peysakhovich

CHAMSI HSSAINE (Homepage, CV, Google Scholar)

**Thesis:** People-Centric Operations of Societal Systems

**Advisor:** Sid Banerjee, Cornell University

**Brief Biography:** Chamsi Hssaine is a final-year Ph.D. student in the School of Operations Research and Information Engineering at Cornell University, where she is advised by Professor Sid Banerjee. She was selected for the 2020 Rising Stars in EECS workshop at UC Berkeley, as well as the 2020 Rising Scholars conference at the Stanford Graduate School of Business. In 2019, she was a visitor at the Simons Institute for the program on Online and Matching-Based Market Design. She graduated *magna cum laude* from Princeton University in 2016, with a B.S. in Operations Research and Financial Engineering.

**Research Summary:** My research centers around *algorithm and incentive design for smart societal systems*. Combining rigorous theory, data-driven approaches, and large-scale optimization, my work develops tools for the *people-centric* design and operations of these systems. In particular, a unique and unifying thread across all my research is a focus on incorporating more realistic models of behavior under incentives, and better understanding the effect of policy decisions on stakeholders.

One of my areas of focus is developing theoretical foundations of how data-driven decision-making impacts *competition* [1] and *collaboration* [4] in multi-agent environments. In [1], we tackle the problem of characterizing market outcomes when competing firms deploy popular online learning algorithms. We introduce a broad class of games of price competition that subsumes many well-validated behavioral models and show that, in these games, gradient-based learning dynamics may converge to outcomes in which firms can experience unbounded losses in revenue compared to the best price equilibrium. To address this concern, we propose a novel learning algorithm which not only provably avoids convergence to such bad outcomes, but also successfully converges to this best equilibrium in experiments.

Much of my work has also been concerned with how ride-hailing platforms interact with the broader ecosystem in which they operate. In [2], we consider the problem of designing *fair* profit-maximizing compensation schemes for gig economy workers, when they have limited information on the underlying algorithms generating their wages. In [3], we investigate how ride-hailing services can be leveraged to create *demand-reactive* transit networks at scale, to create more equitable transportation systems that balance *access* and *efficiency*. Toward this goal, we develop fast approximation algorithms with provable guarantees.

**Representative Papers:**

[1] Pseudo-Competitive Games and Algorithmic Pricing (under review at *Management Science*) with S. Banerjee and V. Kamble

[2] Earning sans Learning: Noisy Decision-Making and Labor Supply on Gig Economy Platforms (under review at *Management Science*) with D. Freund

[3] Real-Time Approximate Routing for Smart Transit Systems (SIGMETRICS 2021) with S. Banerjee, N. Perivier and S. Samaranayake

[4] Information Signal Design for Incentivizing Team Formation (WINE 2018) with S. Banerjee

RAVI JAGADEESAN (Homepage, CV, Google Scholar)

**Thesis:** Market Design for Matching and Auctions

**Advisor:** Scott Kominers, Harvard Business School

**Brief Biography:** Ravi Jagadeesan is a postdoctoral scholar in the Department of Economics at Stanford University, hosted by Al Roth. His research interests are in market design, public economics, and statistics. He received a Ph.D. in Business Economics from Harvard University in 2020, where he was supported by a NSF Graduate Research Fellowship. Prior to starting his Ph.D., he completed his undergraduate work at Harvard in 2018, with an A.B. in mathematics and an A.M. in statistics.

**Research Summary:** My research in market design focuses on the design of market-clearing mechanisms for matching markets and auctions. My work seeks to incorporate important features of real-world settings that matter for market design.

One strand of work studies interconnected markets that suffer from transactional frictions such as transaction taxes and commissions. In [1], we show how to incorporate transactional frictions into a model of matching in trading networks, and characterize when competitive equilibria in this setting. We also a provide cooperative foundation for competitive equilibrium pricing in the presence of frictions that applies when agents evaluate potential deviations myopically.

A second strand of work studies markets with indivisibilities in which budget constraints or other financial constraints limit how much participants can spend. Such constraints are important in markets such as housing markets and spectrum auctions, but make utility non-quasilinear and generally cause the gross substitutability condition to fail. In [2], we develop a duality method for analyzing markets for indivisible goods with non-quasilinear utility. We show that the structure of agents' substitution effects fundamentally determines whether competitive equilibria exist. Using this insight, we introduce a weakening of the gross substitutes condition, called the *net substitutes* condition, that is compatible with financial constraints and still ensures the existence of competitive equilibria. In [3], we apply our duality method to study matching makets with budget constraints. We show that under the net substitutes condition, stable outcomes are guaranteed to exist, but the deferred acceptance algorithm may fail to find a stable outcome.

I also have research interests in statistics. For example, in [4], we develop experimental designs for causal inference in settings in which exposing agents to treatments affects their neighbors in a network.

**Representative Papers:**

[1] Trading Networks with Frictions (EC 2018; *Econometrica* 2019)
with T. Fleiner, Z. Jankó, and A. Teytelboym

[2] The Equilibrium Existence Duality: Equilibrium with Indivisibilities and Income Effects (EC 2020)
with E. Baldwin, O. Edhan, P. Klemperer, and A. Teytelboym

[3] Matching and Prices (EC 2021) with A. Teytelboym

[4] Designs for Estimating the Treatment Effect in Networks with Interference (*Annals of Statistics* 2020) with N. S. Pillai and A. Volfovsky

SÜLEYMAN KERIMOV (Homepage, CV, Google Scholar)

**Thesis:** Essays on Dynamic Matching and Market Design

**Advisors:** Itai Ashlagi, Stanford University; Itai Gurvich, Northwestern University

**Brief Biography:** I am a final year Ph.D. student in Operations Research at Stanford University, where I am advised by Itai Ashlagi and Itai Gurvich. I received my B.S. in Mathematics from Bilkent University in 2016.

**Research Summary:** I am interested in market design, matching and applied probability. My research is motivated by frictions that arise in various marketplaces, including kidney exchange. My thesis has focused on challenges that originate due to stochasticity and liquidity in matching markets, and their effects on the ability to apply simple policies that can achieve efficient outcomes.

In [1], we study how to optimally match agents in centralized dynamic matching markets, where agents arrive stochastically and match values are heterogeneous. In such dynamic markets, there is an inherent trade-off between short- and long-term allocative efficiency. Delaying actions to accumulate inventory creates a positive externality from forming future matches that generate high value. This delay, however, inevitably compromises short-term value. We find that when the market exhibits even a small imbalance in terms of agents' arrival rates, there exists a simple periodic clearing policy with a carefully chosen period length that nearly maximizes total match value simultaneously at all times; the policy is hindsight optimal and the trade-off is essentially moot. We also establish that acting greedily is suboptimal. As a follow-up, we restrict our attention to two-way matching networks in [2], and we show that suitably designed greedy policies are hindsight optimal. Both papers also reveal an intimate connection between queueing theory and matching theory. We characterize a parameter called the *general position gap*, which is analogous to *network utilization* parameter in queueing theory that dictates a lower bound on stationary queue-lengths. Similarly, the general position gap dictates a lower bound on regret in matching networks, and we establish that this lower bound is achievable under the proposed periodic clearing and greedy policies.

In [3], we consider economies without monetary transfers, where scrip systems serve an alternative to sustain cooperation, improve efficiency and mitigate free riding. We find that even with minimal liquidity in the market, in the sense that only few agents are available to provide service, cooperation can be sustained by balancing service provisions among agents. This suggests that scrip systems can lead to efficient outcomes in kidney exchange platforms by sustaining cooperation between hospitals, where compatibility (liquidity) in such markets is sparse.

**Representative Papers:**

[1] Dynamic Matching: Characterizing and Achieving Constant Regret (Submitted to Management Science; SSRN) with I. Ashlagi, I. Gurvich

[2] On the Optimality of Greedy Policies in Dynamic Matching (Submitted to Operations Research; SSRN) with I. Ashlagi, I. Gurvich

[3] Scrip Systems with Minimal Availability (Working paper; WINE 2019) with I. Ashlagi

HANNAH LI (Homepage, CV, Google Scholar)

**Thesis:** Decision-Making Using Platform Experiments: What Goes Wrong Under Interference

**Advisor:** Ramesh Johari and Gabriel Weintraub, Stanford University

**Brief Biography:** Hannah is a PhD Candidate at Stanford University, where she is advised by Ramesh Johari and Gabriel Weintraub. She is part of the Operations Research group in MS&E and the Society and Algorithms Lab. Her research uses math modeling, optimization, and causal inference in order to analyze and design data science methodology for marketplace platforms. She has worked with several platforms, including Airbnb, Common App, Opendoor, and Vinted. Before coming to Stanford, she graduated from Pomona College with a degree in mathematics.

**Research Summary:** I develop generalizable statistical methods for online platforms. I put these theoretically grounded advances into practice by engaging with industry partners. This requires a close collaboration with my partners from the initial problem formulation to the final implementation. I demonstrate the impact of this approach using applications in experimental design and education settings.

My thesis focuses on experimental design in two-sided marketplaces in the presence of interference. These platforms often use experiments to test new interventions and decide whether to implement platform-wide. However, treatment effect estimates in these experiments can be biased, because users in the market can affect or "interfere" with the outcomes of others. In [1], we study commonly used experiment types analytically and show that bias depends on the supply and demand imbalance in the market. We introduce and study a novel experiment design based on two-sided randomization (TSR), which leverages the connection to market imbalance to reduce bias. In follow-up work, we investigate the variance [2] and the bias that interference creates in standard error estimates. Finally, in ongoing work, we evaluate the impact of these experimental biases on the accuracy of the resulting decisions. Throughout this process, I worked with industry partners to evaluate the magnitude of these effects in practice as well as run live TSR experiments.

In another line of work, I utilize operations research techniques to design accessible education systems, both through theoretical methods and industry collaboration. In [3] we develop a theoretical model to study the often-debated decision of whether to drop standardized testing requirements in school admissions and, more broadly, the design of fair selection processes in capacity constrained systems. On the implementation side, I am working with the Common App to build data-driven solutions for quantifying and improving the equity of higher education.

**Representative Papers:**

[1] Experimental Design in Two-Sided Platforms: An Analysis of Bias (Forthcoming in *Management Science*, EC'20)
with R. Johari, I. Liskovich, and G. Weintraub

[2] Interference, Bias, and Variance in Two-sided Marketplace Experimentation: Guidance for Platforms (EC'21 Workshop on Design of Online Platforms)
with G. Zhao, R. Johari, and G. Weintraub

[3] Dropping Standardized Testing for Admissions (FAccT'20 and EAAMO '21, under review in *Management Science*) with N. Garg and F. Monachou

YINGKAI LI (Homepage, CV)

**Thesis:** Simple Mechanisms for Trading Information

**Advisor:** Jason Hartline, Northwestern University

**Brief Biography:** Yingkai is a final-year PhD candidate at Northwestern University in Computer Science, where he is advised by Jason Hartline. His research interests lie broadly at the intersection of computer science and economics, with a focus on mechanism design, information design and repeated games. He was a summer intern at MSR in year 2020 and 2021 hosted by Brendan Lucier and Alex Slivkins.

**Research Summary:** There is a drastic growth in the online markets, which brings new opportunities and new challenges for designing good mechanisms. Specifically, it is generally computational hard to compute the optimal mechanisms. Thus simple and practical mechanisms are widely adopted, and my research focuses on providing the theoretical foundations for those mechanisms, e.g., showing their approximate optimality and robustness.

For multi-agent Bayesian auctions, in [5] and [FHL-21], we show that simple mechanisms such as sequential posted pricings are approximately optimal for wide classes of non-linear utilities. In [4], we show that a simple random markup mechanism is the prior-independent optimal mechanism, i.e., the mechanism that minimizes the approximation ratio when the seller is ignorant of the prior distribution.

Recently, I have focused on the topic of trading information. In [2], we provide characterizations on the optimal reward scheme for incentivizing the agent to exert effort in acquiring information. Moreover, we design a novel approximately optimal reward scheme when there are multiple tasks for the agent, and show that it is crucial to link the incentives of the agent across different task. In [1], we consider the revenue maximization problem for a principal selling data to an agent with private preference and endogenous information. We fully characterize the optimal mechanism and show that pricing for revealing full information reaches half of the optimal revenue.

I also share interests in topics related to repeated games. For example, in [3], we characterize the equilibrium behavior of the patient player in reputation games, and in [LP-21], we consider a time misspecification model in repeated games and identify novel cyclic behaviors in the actions sequentially chosen by the agent.

**Representative Papers:**

[1] Selling Data to an Agent with Endogenous Information (Arxiv)

[2] Optimization of Scoring Rules (Arxiv, Best poster in EC 20)
with J. Hartline, L. Shan, and Y. Wu

[3] Equilibrium Behaviors in Repeated Games (JET 21)
with H. Pei

[4] Benchmark Design and Prior-independent Optimization (FOCS 20)
with J. Hartline and A. Johnsen

[5] Optimal Auctions vs. Anonymous Pricing: Beyond Linear Utility (EC 19)
with Y. Feng and J. Hartline

SIMON MAURAS (Homepage, CV, Google Scholar)

**Thesis:** Analysis of Random Models for Stable Matchings

**Advisor:** Claire Mathieu and Hugo Gimbert, CNRS, France

**Brief Biography:** Simon Mauras is currently a Ph.D. candidate in Computer Science at Université de Paris, advised by Claire Mathieu and Hugo Gimbert. In 2022, he will be joining the Economics and Computation Laboratory at Tel Aviv University as a Postdoctoral researcher, hosted by Michal Feldman. In 2019, he was a visitor at Simons Institute for the Theory of Computing for the program on Online and Matching-Based Market Design. In 2018, he obtained a M.Sc. in Computer Science from École Normale Supérieure de Lyon in France. Since 2013, he volunteers to select and train French high-school students for the International Olympiads in Informatics.

**Research Summary:** In my research, I study problems at the intersection between algorithms, probability and game theory. More specifically, I worked on rank aggregation [1], stable matchings [2,3], data-structures and online optimization.

In my PhD thesis, I look at random models of two-sided matching markets, where agents draw their preferences from a known distribution. A recent line of works show that stable matchings are essentially unique, either because agents have short preference lists [IM15], or because the market is unbalanced [AKL17]. This phenomenon is referred to as core-convergence, and implies incentive compatibility. In the first part of my thesis, addressing the question "who can manipulate?", I explore an alternative explanation for core-convergence, based on the fact that agents have correlated preferences [3]. In complement, I look at the incomplete information game where applicants must behave strategically because of an upper quota on the number of applications they can submit. In the second part, addressing the question "who gets what?", I study the output of Gale and Shapley's deferred acceptance procedure. Under certain input preference distributions, I show that the output distribution does not depend on which side proposes [2], and I give exact and approximate formula for the probability that two agents will be matched.

Beyond theoretical results, my contribution to several research projects consists in experimental results obtained via computer simulations. In a paper published in PLOS Computational Biology [4], we quantify the relative efficiency of multiple telecommuting strategies on the probability of a COVID outbreak in a workplace, simulating the propagation of a disease in the graph of contacts between employees.

**Representative Papers:**

[1] How to Aggregate Top-Lists: Approximation Algorithms via Scores and Average Ranks (SODA'20) with C. Mathieu

[2] Two-Sided Random Matching Markets: Ex-Ante Equivalence of the Deferred Acceptance Procedures (EC'20; forthcoming in TEAC)

[3] Two-Sided Matching Markets with Strongly Correlated Preferences (FCT'21) with H. Gimbert and C. Mathieu

[4] Mitigating COVID-19 outbreaks in workplaces and schools by hybrid telecommuting (PLOS Computational Biology) with V. Cohen-Addad, G. Duboc, M. Dupré la Tour, P. Frasca, C. Mathieu, L. Opatowski and L. Viennot

FAIDRA MONACHOU (Homepage, CV)

**Thesis:** Discrimination, Diversity, and Information in Selection Problems

**Advisor:** Itai Ashlagi, Stanford University

**Brief Biography:** Faidra Monachou is a final-year Ph.D. candidate in Management Science and Engineering at Stanford University. She is interested in market and information design, with a particular focus on the interplay between policy design and discrimination in education and labor. Faidra's research has been supported by various scholarships and fellowships from Stanford Data Science, Stanford HAI, Google, and other organizations. She won the Best Paper with a Student Presenter Award at ACM EAAMO'21. She co-chaired the MD4SG'20 workshop and co-organizes the Stanford Data Science for Social Good program.

**Research Summary:** My research lies at the intersection of *operations* and *social sciences*. I study policy and market design questions through an *informational lens*.

A major line of my research is concerned with understanding the role that information plays in *discrimination*, especially in admissions in education [1] and online labor markets [3]. Despite the large empirical literature on discrimination, our theoretical understanding is limited. Thus, I develop a theoretical framework to study how a decision-maker concerned with both merit and diversity, selects candidates under imperfect information, limited capacity, and legal constraints. In [1], motivated by recent decisions to drop standardized testing in admissions, we apply this framework to study how information differences lead to disparities across equally skilled groups. Furthermore, we quantify the trade-off between information and access in test-free and test-based policies with and without affirmative action. In follow-up work, I extend the model in [1] to study how *privilege* differences lead to *intra-* and *inter-group* disparities. We establish that the direction of discrimination at the observable level differs from the unobservable level. We also evaluate common affirmative action policies and design an optimal policy under legal constraints.

I also study fundamental questions on *information design in markets* [2, 4]. In [2], using majorization techniques, we identify the welfare-maximizing mechanism in a continuum dynamic market with objects and agents with private types, supermodular utilities, and quasilinear payoffs in their waiting cost. We show that it can be implemented in two ways: (i) via a system of monotone disjoint queues; (ii) using information design, via a single FCFS queue with deferrals paired with an information disclosure policy that pools adjacent object types.

**Representative Papers:**

[1] Dropping Standardized Testing for Admissions Trades Off Information and Access (FAccT 2021, EAAMO 2021 - *Best Paper with a Student Presenter Award*), with N. Garg and H. Li

[2] Optimal Dynamic Allocation: Simplicity through Information Design (EC 2021), with I. Ashlagi and A. Nikzad

[3] Discrimination in Online Markets: Effects of Social Bias on Learning from Reviews and Policy Design (NeurIPS 2019), with I. Ashlagi

[4] Counterbalancing Learning and Strategic Incentives in Allocation Markets (NeurIPS 2021), with J. Kang, M. Koren, and I. Ashlagi

VISHNU NARAYAN (Homepage, CV)

**Thesis:** Combinatorial Auctions and Allocations

**Advisor:** Adrian Vetta, McGill University

**Brief Biography:** Vishnu Narayan is a Ph.D. candidate at McGill University advised by Adrian Vetta. For Fall 2021, he is a Fellow of the Harvard SEAS where he is hosted by Ariel Procaccia and is working on fair division and computational social choice. Before his PhD, Vishnu completed an M.Math. in Combinatorics and Optimization at UWaterloo (Canada). He has a B.Eng. from R.V. College of Engineering (India) where he received the Best Outgoing Student Award. He has also received the Harold H. Helm Fellowship and Murata Family Fellowship from the McGill Faculty of Science, and won a Best Paper Award at SAGT 2019. Alongside his central thesis research, Vishnu is also broadly interested in CS theory and has published research in algorithms, combinatorics, optimization, and game theory. He is very enthusiastic about teaching courses in these areas and has a Teaching Assistant Award from the McGill School of Computer Science. Outside of academia, he enjoys competitive programming and strategy games.

**Research Summary:** I am drawn to problems that tie my previous expertise in combinatorics and optimization together with the Econ-CS domain. One theme of my doctoral work concerns sequential auctions, which are perhaps the most natural method by which to sell a set of items. The interplay between the potentially arbitrary combinatorial valuation functions of the agents and the simplicity of the mechanism in which they bid leads to equilibria that are notoriously hard to analyze. We showed that the declining price anomaly – an observation that is empirically supported by dozens of papers on real-world sequential auctions – does not always occur in the subgame-perfect equilibria of these auctions. This work won the best paper award at SAGT 2019 [1]. In other work, we analyzed the risk-free strategy (ie. the safety strategy) and corresponding equilibria of an agent in these auctions.

A second theme of my Ph.D. research concerns fair division. Our paper, that I presented at EC'20 [2], looks at achieving envy-freeness in indivisible-item instances through the use of an extra divisible good (i.e. a subsidy, or, equivalently, transfer payments between the agents). We gave a tight upper bound on the amount of subsidy sufficient to eliminate envy, resolving two conjectures. We've since expanded this work to study the extent to which one can concurrently obtain envy-freeness and high welfare through the use of transfers [3]. I hope to complete a comprehensive analysis of the inclusion of transfers in envy-free fair division. Moving forward, I am focused on resolving the plethora of open problems in combinatorial auctions, fair division and social choice, which continue to grow in number and importance.

**Representative Papers:**

[1] The Declining Price Anomaly is not Universal in Multi-Buyer Sequential Auctions (but almost is) (SAGT 2019) with E. Prebet and A. Vetta

[2] One Dollar Each Eliminates Envy (EC 2020)
with J. Brustle, J. Dippel, M. Suzuki, and A. Vetta

[3] Two Birds With One Stone: Fairness and Welfare via Transfers (SAGT 2021) with M. Suzuki and A. Vetta

CHARA PODIMATA (Homepage, CV, Google Scholar)

**Thesis:** Incentive-Aware Machine Learning

**Advisor:** Yiling Chen, Harvard

**Brief Biography:** Chara is a PhD candidate in the EconCS group at Harvard, where she is advised by Yiling Chen. Her research is supported by a Microsoft Dissertation Grant and a Siebel Scholarship. During her PhD, she interned twice for MSR NYC (mentored by Jennifer Wortman Vaughan and Aleksandrs Slivkins) and once for Google Research NYC (mentored by Renato Paes Leme). She has given tutorials related to strategic learning at EC20 and FAccT21. Outside of research, she spends her time adventuring with her pup, Terra.

**Research Summary:** My research addresses questions related to strategic behavior in Machine Learning (ML). My overarching goal is to build a theory of incentives in ML, where the deployed algorithms have provable guarantees even against strategic agents. Towards this goal, I design new ML algorithms, that are robust to strategic noise, rather than the traditional notions of stochastic or adversarial noise. Drawing intuition from Mechanism Design, strategic robustness can take the form of either *Incentive-Compatibility (IC)* or *Incentive-Awareness (IA)*.

IC learning algorithms guarantee that it is in the agents' best interest to report their data truthfully. In this direction, my work has provided strong IC results for the fundamental tasks of high-dimensional linear regression [1] and online prediction with expert advice [2]. A key theme in both is that the results are bridging between seemingly unrelated problems; the Ham Sandwich Theorem from Computational Geometry in [1] and the wagering mechanisms from Mechanism Design in [2].

Despite the various encouraging results, IC remains a relatively hard goal. Instead, the more relaxed desideratum for some settings is IA, where the algorithm adapts to the incentives of the agents it faces. One of the most studied IA ML problems is "strategic classification", for which in [3] we model a broad class of utility and loss functions for the agents and the learner. Our main result is a novel adaptive discretization algorithm, whose regret scales with the power of the strategic agent. All these game-theoretic formulations assume that agents act according to a specific behavioral model, and existing IA algorithms can have poor performance if the realized behavior of some agents is not the one assumed (i.e., irrationalities). In [4], we design an ML algorithm whose performance degrades gracefully with the number of such irrational agents. This extends the classical Ulam's game (binary search with lies) to a contextual version and unknown number of lies.

**Representative Papers:**

[1] Strategyproof Linear Regression in High Dimensions (EC18, best-paper finalist) with Y. Chen, A. Procaccia, and N. Shah

[2] No-Regret and Incentive-Compatible Online Learning (ICML20) with R. Freeman, D. Pennock, and J. Vaughan

[3] Learning Strategy-Aware Linear Classifiers (NeurIPS20) with Y. Chen, and Y. Liu

[4] Contextual Search in the Presence of Irrational Agents (STOC21) with A. Krishnamurthy, T. Lykouris, and R. Schapire

ARIEL SCHVARTZMAN COHENCA (Homepage, CV )

**Thesis:** Circumventing Lower Bounds in Mechanism and Tournament Design

**Advisor:** S. Matthew Weinberg, Princeton University

**Brief Biography:** Ariel Schvartzman Cohenca is a Postdoctoral Associate at DIMACS supervised by David Pennock. Ariel's work focuses in understanding the trade-off between optimality and simplicity in the design of multi-dimensional auctions. He was awarded the Department of Computer Science's Graduate Student Teaching Award in 2017, and the School of Engineering and Applied Science's Award for Excellence in 2018. Ariel obtained his PhD in June 2020 from Princeton University, under the supervision of S. Matthew Weinberg. He obtained his B.S. in Mathematics with Computer Science from MIT in 2015.

**Research Summary:**

Optimal mechanism design beyond single-item settings remains a central question at the intersection of economics and computer science. The problem is intricate for a number of reasons: the mechanisms may be bizarre, computationally hard to find or simply too complex to present to a bidder. The community's focus, thus, has shifted from to asking, for instance, how complex must a mechanism be in order to extract 99% of the optimal revenue? My work joins that of others in quantifying this trade-off explicitly. Our results suggest that significantly simpler mechanisms can compete with optimal ones if the seller is willing to lose 1% of the optimal revenue [1]. Another way to circumvent some of the impossibility results in optimal mechanism design is to reduce the dimensionality of the problem by restricting the relationship between the items for sale. Our work extends the so-called FedEx problem to the case of Single-Minded buyers and shows that even in this simpler setting, optimal mechanisms may remain prohibitively large to write down ([2], Saxena et al. SODA 2018).

In settings where buyers have correlated valuations simple (or even finite) mechanisms have no hope of competing with optimal ones, even approximately. In light of this, we begin the study of beyond-worst case approximations for correlated bidders via the smoothed-analysis framework. Our results she light on the properties that make correlated distributions inapproximable [3].

Finally, I am also interested in mechanism design for tournaments: how should a tournament designer pick a reasonable winner from a set of teams? We show that simple tournament formats are optimal among all fair ones that dissuade collusion (Schneider et al., ITCS 2017, Schvartzman et al., ITCS 2020).

**Representative Papers:**


[1] Approximation Schemes for a Buyer with Independent Items via Symmetries (FOCS 2019) with P. Kothari, D. Mohan, S. Singla, and S. M. Weinberg

[2] Optimal Mechanism Design for Single-Minded Agents (EC 2020) with N. Devanur, K. Goldner, R. R. Saxena, and S. M. Weinberg

[3] Smoothed Analysis of Multi-Item Auctions with Correlated Values (EC 2019) with A. Psomas, and S. M. Weinberg

ANA-ANDREEA STOICA (Homepage, CV)

**Thesis:** Diversity and Inequality in Social Networks

**Advisor:** Augustin Chaintreau, Columbia University

**Brief Biography:** Ana-Andreea Stoica is a Ph.D. candidate at Columbia University, advised by Augustin Chaintreau. Ana holds a Bachelor's degree in Mathematics from Princeton University (2016). Her work focuses on mathematical models, data analysis, and inequality in social networks. Ana is an awardee of the 2019 J.P. Morgan Ph.D. Fellowship and has interned at Microsoft Research NYC in 2019. Since 2019, she has been co-organizing the Mechanism Design for Social Good (MD4SG) initiative. Ana served as a Program Co-Chair for the inaugural ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization in 2021, which has grown out of the MD4SG initiative.

**Research Summary:** My research develops theories about robust models to support decision-makers in bridging learning objectives with human incentives, particularly in settings where there are underlying social and information networks. A main theme in my research is to leverage algorithmic and computational techniques to shed light on the root drivers of bias in automated decision-making and to bring algorithmic, mechanism design, and machine learning insights to bear on effective interventions to mitigate such bias. Towards this end, I draw on insights from graph theoretical models of interaction and machine learning methods to understand complex structures and incentives in networks.

In [1], we redefine an optimization problem for social influence maximization, entailing the algorithmic selection of nodes chosen due to their advantageous position in a network. We find analytical conditions in which algorithms that are aware of the network structure can be more efficient and mitigate inequality within a population by selecting the most promising individuals in a more efficient way. This project builds on my previous work [3] that uses network models in proving theoretical conditions for establishing inequality bounds in recommendation algorithms. In [2], we leverage voting mechanisms that embed users' geography, constraints, and preferences, and use them to split people in districts or communities that fairly represent their voices. Our theoretical analysis proves the computational complexity of this problem, while providing algorithms that improve demographic diversity in segregated clusters. Inspired by these, I aim to continue to work on algorithms that learn individuals' position and relations in order to cluster them based on their preferences and constraints, with a focus on feature-aware design. My most recent lines of work focus on redesigning clustering algorithms as well as multi-objective optimization for the purpose of resource allocation.

**Representative Papers:**

[1] Seeding Network Influence in Biased Networks and the Benefits of Diversity (TheWebConf'20) with J. X. Han and A. Chaintreau

[2] Minimizing Margin of Victory for Fair Political and Educational Districting (AAMAS'20) with A. Chakraborty, P. Dey, and K. P. Gummadi

[3] Algorithmic Glass Ceiling in Social Networks: The effects of social recommendations on network diversity (TheWebConf'18) with C. Riederer and A. Chaintreau

WEI TANG (Homepage, CV)

**Thesis:** Examining the Interplay Between Humans and Algorithm Design

**Advisor:** Chien-Ju Ho, Washington University in St. Louis

**Brief Biography:** Wei Tang is a PhD candidate in the Department of Computer Science & Engineering at the Washington University in St. Louis, advised by Chien-Ju Ho. His research interests are in online learning, algorithmic economics, optimization, and behavioral experiments, with a focus on developing theoretically rigorous, empirically grounded frameworks to understand and design human-centered algorithms. He received the B.E. degree from Tianjin University in 2017.

**Research Summary:** The goal of my research is to develop theoretically rigorous, empirically grounded frameworks to understand and design algorithmic systems that integrate humans as a component in the design process. This *human-centered* focus is interdisciplinary in nature, and my work draws on ideas from online learning, algorithmic economics, optimization, and behavioral social science.

One line of my research has focused on how human behavior impacts efficient algorithm design [1, 2]. For example, I studied how to include human *biased* behavior in online learning frameworks [1]. I explored two natural behavior models where I showed that one allows an efficient learning algorithm while the other makes the efficient learning infeasible. The results demonstrate the importance of understanding human behavior in algorithm design. Another line of my research focuses on how algorithms impact human welfare. For example, I explored long-term impact of actions, which often came up when human well-being is involved, informed by the consequential decisions [3]. To formulate the problem, I generalized the multi-armed bandit by introducing the *impact functions* to encode the dependency between the action history and the arm rewards. This dependency structure makes the problem substantially more challenging where applying standard bandit algorithms leads to linear regret. I then proved that, under mild conditions, efficient algorithms with theoretical guarantees for this problem are possible.

In addition to theoretical analysis, I conducted behavioral experiments to understand how humans respond to information in practice in an information design game [4]. The experimental results show that human behavior significantly deviates from standard model, which assumes Bayesian rational. I showed that an alternative model (discrete choice model coupled with probability weighting) better aligns with human's real behavior. My vision in this line of work is to strengthen the bond between theoretical and empirical analyses of human-centered algorithm design.

**Representative Papers:**

[1] Bandit Learning with Biased Human Feedback (AAMAS'19)
with CJ. Ho

[2] Linear Models are Robust Optimal Under Strategic Behavior (AISTATS'21)
with CJ. Ho, and Y. Liu

[3] Bandit Learning with Delayed Impact of Actions (NeurIPS'21)
with CJ. Ho, and Y. Liu

[4] On the Bayesian Rational Assumption in Information Design (HCOMP'21)
with CJ. Ho

DAVID WAJC (Homepage, CV, Google Scholar)

**Thesis:** Matching Theory Under Uncertainty

**Advisor:** Bernhard Haeupler, Carnegie Mellon University & ETH Zurich

**Brief Biography:** David is the 2020 Motwani postdoctoral fellow in theoretical computer science at Stanford University, where he is hosted by Amin Saberi. Prior to that, he completed his PhD in computer science at CMU, as part of the interdisciplinary Algorithms, Combinatorics and Optimization program. He has gone on numerous extended academic visits over the years, and is looking forward to the day where academic travel is once again the norm, following these uncertain times. Fittingly, his research focuses on algorithms under uncertainty.

**Research Summary:** The proliferation of user-facing mobile and web-based apps has made online problems more prominent than ever. In many such applications algorithms must match agents (e.g., riders and passengers) immediately and irrevocably upon arrival of agents or matching opportunities, which are a priori unknown. How well can algorithms do in the face of such uncertainty? Much of my research addresses this question.

In a seminal work in the early 90's, Karp et al. proved the sub-optimality of the greedy algorithm for online matching—at least in bipartite graphs under one-sided vertex arrivals—and asked whether their results extend to more general settings. In a FOCS'19 paper [1], we resolved this decades-old question, showcasing the importance of granularity of sequential decision making: greedy is sub-optimal for general vertex arrivals, but optimal for edge arrivals.

In an EC'21 paper [3], we studied Online Bayesian Selection (OBS) in two-sided matching markets. Deviating from most prior work on OBS, rather than ask how well we can approximate a "prophetic" offline algorithm, we ask how well we can approximate the optimal *online* algorithm. Surprisingly, we prove that it is PSPACE-hard to approximate this algorithm within some factor $\alpha < 1$. On the other hand, we also design polytime online algorithms that approximate the optimal online algorithm within a factor better than the best-possible prophet inequality.

Changing data is also central to the *dynamic algorithms* literature, which studies how quickly solutions can be updated after changes to the data. For the well-studied dynamic matching problem, the fastest known algorithms were all randomized, and crucially assumed non-adaptively generated data. Unfortunately, this assumption does not hold in user-facing applications. Naturally, a central question in the dynamic algorithms literature is whether or not randomization is still useful without this assumption. In a single-authored paper at STOC'20 [2], I answered this question for dynamic matching, presenting the first randomized algorithms for this problem which work in the face of adaptively-changing data.

**Representative Papers:**

[1] Online Matching with General Arrivals (FOCS'19)
with B. Gamlath, M. Kapralov, A. Maggiori, O. Svensson

[2] Rounding Dynamic Matchings Against an Adaptive Adversary (STOC'20)
Single authored

[3] Near-Optimum Online Ad Allocation for Targeted Advertising (EC'21)
with C. Papadimitriou, T. Pollner and A. Saberi

KANGNING WANG (Homepage, CV, Google Scholar)

**Thesis:** Approximately Fair and Efficient Economic Solutions

**Advisor:** Kamesh Munagala, Duke University

**Brief Biography:** Kangning is a fifth-year Ph.D. student in Computer Science at Duke University, advised by Kamesh Munagala. He spent two summers as a research intern at Google Research: hosted by Jieming Mao and Renato Paes Leme in 2020, and by Aranyak Mehta in 2021. His work received Best Paper Award at WINE 2018. He got his bachelor's degree from Yao Class, Tsinghua University.

**Research Summary:** In many economic scenarios, the ideal solutions are inaccessible, due to selfishness of agents, lack of information, computational hardness, over-idealization of the benchmarks, etc. My research focuses on providing theoretically guaranteed compromises in these situations.

In *multi-winner elections*, the *core* is a classical notion whose general idea has existed for more than a century. In a core solution, no group of voters can deviate to another proportionally-sized set of candidates while benefiting every voter in the group. It is a strong notion of proportionality and fairness, stating that every group of voters should be well represented. Unfortunately, core solutions may not exist in many settings. In a line of work, we demonstrate how to bypass the impossibility results via scaling down the budgets [1] or utilities [SODA 2022] of deviating voters, or via randomization [EC 2019].

The *metric distortion* problem for social choice rules has recently gained significant attention: voters and candidates lie in a metric space, and a central designer needs to select a candidate to minimize the sum of its distances to each voter. However, the designer only knows the ordinal preferences of the voters instead of the exact cardinal metric. In [2], we improved the efficiency guarantee by proposing a novel social choice rule.

In *bilateral trade*, the celebrated result of Myerson and Satterthwaite states that no budget-balanced mechanism between two self-interested agents can be efficient for the society. In [3], we provide the first constant-approximation mechanism to the first-best gains-from-trade, a natural benchmark measuring efficiency. This settles a long-standing open question mentioned in multiple previous works. In another work with J. Mao and R. Paes Leme, we demonstrate how efficiency can be achieved if the two agents can credibly exchange information by engaging in a bargaining process prior to trading.

**Representative Papers:**

[1] Approximately Stable Committee Selection (STOC 2020)
with Z. Jiang and K. Munagala

[2] Improved Metric Distortion for Deterministic Social Choice Rules (EC 2019)
with K. Munagala

[3] Approximately Efficient Bilateral Trade (Manuscript)
with Y. Deng, J. Mao, and B. Sivan

FANG-YI YU (Homepage, CV)

**Thesis:** Dynamics on Social Networks

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

**Brief Biography:** I am a postdoctoral fellow in the Harvard EconCS group under the supervision of Yiling Chen. I obtained my Ph.D. degree in Computer Science in August 2019 from the Computer Science and Engineering Division at the University of Michigan. I received the B.S. degree in Electrical Engineering with double major in Mathematics from the National Taiwan University in 2013.

**Research Summary:** I am interested in the theoretical aspect of multi-agent systems. My work focuses on understanding how these multi-agent systems interact with information and explores how these systems can solve social problems, including learning, cooperation, and communication.

My research often explores these emergence phenomena, where the systems can collectively exhibit properties that individual agents do not have, and design robust interventions for social impacts. I will highlight examples of my work in two settings: 1) Information elicitation, where each agent's information is correlated and is used to acquire and aggregate their information [1]; 2) Social and economic networks, where each agent's actions only directly impact those around it [2,3]. My recent and forthcoming work also focuses on 3) dynamics of learning with multiple agents where each agent's decision affects data they aim to predict and each other.

The goal of information elicitation is to design systems that elicit or aggregate information from people. When an agent's information is private or subjective, peer prediction mechanisms exploit the interdependence in agents' signals to incentivize agents to report their private signal truthfully even when the reports cannot be directly verified. Using robust and variational statistics, I design several robust mechanisms that learn collective structure in the multi-agent system through interaction among agents, and use those structural properties to offset noisy, strategic, or even adversarial behavior. [1]

One line of my work studies the long-term behavior of dynamical systems with applications, including contagions and opinion formation on social networks, local search algorithms, and equilibria of no-regret learners. In [2,3], I study a large family of stochastic processes containing stochastic gradient descent with a uniform step size. We show that those systems can escape the non-attracting fixed point and converge to an attracting fixed point in $\Theta(n \log n)$. This result improves previous analysis of stochastic gradient descent escaping saddle points and provides new insight into evolutionary stable strategies in evolutionary game theory.

**Representative Papers:**

[1] Learning and Strongly Truthful Multi-Task Peer Prediction: A Variational Approach (ITCS '21) with G. Schoenebeck

[2] Escaping Saddle Points in Constant Dimensional Spaces: An Agent-based Modeling Perspective (EC '20) with G. Schoenebeck

[3] Consensus of Interacting Particle Systems on Erdos-Renyi Graphs (SODA 18) with G. Schoenebeck

MANOLIS ZAMPETAKIS (Homepage, CV)

**Thesis:** Statistics in High Dimensions without IID Samples: Truncated Statistics and Minimax Optimization

**Advisor:** Constantinos Daskalakis, Massachusetts Institute of Technology (MIT)

**Brief Biography:** Manolis Zampetakis is a postdoctoral researcher at the Foundations of Data Science Institute (FODSI) in UC Berkeley, working with Prof. Michael Jordan. He received his Ph.D. at MIT, advised by Prof. Constantinos Daskalakis. His research interests include theoretical machine learning, statistics, optimization, mechanism design, and complexity theory. He is a recipient of the 2018 Google PhD Fellowship on "Algorithms, Optimizations, and Markets". His Ph.D. dissertation was awarded the 2020 ACM SIGecom Doctoral Dissertation Award.

**Research Summary:** The design of modern ML systems commonly makes idealized assumptions about their input, their goal, and the environment that they operate in. My research is focused on the theoretical understanding of the opportunities and limitations arising for ML when these assumptions are dropped.

Specifically, some of the most pervasive assumptions of ML systems are: (1) we are given access to data that adequately captures the application environment, (2) there exists a single objective function to be optimized, and (3) the data provider is not affected by the outcome chosen by the ML system. As part of my work, I designed and analyzed methods that handle the absence of (1)-(3) and I provided fundamental lower bounds on the effectiveness of ML systems in these cases.

(1) Truncated Samples. Truncation occurs when data falling outside of a subset of the population are not observable. Such phenomena have been known to affect data analysis in a counterintuitive way, as in *Berkson's paradox*. In a recent line of work, we provide the first provably efficient methods to accomplish basic statistical tasks *out of exclusively censored data* (see [1] and follow-up works).

(2) Multiple Objectives. A basic step towards ML applicable to complex environments is to incorporate the existence of multiple, many times opposing, objective functions. The most fundamental problem in this direction is *min-max optimization*. In our work, we provide the first exponential separation in the query & time complexity required to solve min-max optimization problems vs single-objective ones [2]. We complement our intractability results by providing novel techniques to design algorithms that at least achieve asymptotic convergence (see Ph.D. Thesis).

(3) Strategic Agents. The output of ML systems can be used to make decisions that in turn will affect the data. Ignoring this feedback loop leads to fallacious decisions and provides the wrong incentives for adversarial behavior in the collection of data. My work on Mechanism Design (see, e.g., [3]) provides tools for decision making in the presence of such feedback loops.

**Representative Papers:**

[1] Efficient Statistics, in High Dimensions, from Truncated Samples (FOCS 2018) with C. Daskalakis, T. Gouleakis, and C. Tzamos

[2] The Complexity of Constrained Min-Max Optimization (STOC 2021) with C. Daskalakis, and S. Skoulakis

[3] Robust Learning of Optimal Auctions (NeurIPS 2021 *Spotlight*) with W. Guo, and M. Jordan

## Index