SIGecom Job Market Candidate Profiles 2021

This is the sixth annual collection of profiles of the junior faculty job market candidates of the SIGecom community. The twenty four candidates for 2021 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-candidates2021@acm.org.

–Vasilis Gkatzelis and Jason Hartline

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

Contents

JAMES P. BAILEY (Homepage, CV, Google Scholar)

Thesis: The Price of Deception in Social Choice

Advisor: Craig A. Tovey, Georgia Institute of Technology

Brief Biography: James Bailey is a visiting assistant professor at Texas A&M University in the Industrial and Systems Engineering Department. He was previously a postdoctoral researcher with Dr. Georgios Piliouras at the Singapore University of Technology and Design, and graduated from Georgia Tech's Algorithms, Combinatorics, and Optimization PhD program. His research focuses on building fundamental knowledge in online optimization, algorithmic game theory, and social choice by using multi-disciplinary approaches.

Research Summary: My research focuses on advancing the fundamental understanding of online learning in multi-agent systems using a multi-disciplinary understanding of complex systems. Unlike standard approaches that focus on regret and time-average convergence, I focus on characterizing the dynamics of agents' strategies. Using my new characterization, I improve the regret and convergence guarantees of gradient descent (GD) [2] and uncover a faster, decentralized algorithm for online optimization [1]. To build foundations for these results, I first prove that the continuous-time versions of follow-the-regularized-leader algorithms (FTRL), which includes GD, are Hamiltonian systems [3] – a well-understood system in physics. This connection opens a wide-range of tools and ideas to advance online optimization and algorithmic game theory.

I prove that FTRL in discrete-time is repelled by the set of Nash equilibria [4], contradicting the common belief that no-regret algorithms converge to Nash – a misnomer of time-average convergence. Moreover, I tightly characterize these dynamics to prove that GD in two-agent two-strategy zero-sum games achieves $O(\sqrt{T})$-regret when agents use arbitrary fixed step-sizes [2]. This matches the best-known bound using adaptive step-sizes and corrects the long-held belief that decaying step-sizes are needed for $o(T)$-regret in zero-sum games.

I also uncover a $O(1)$-regret decentralized algorithm for two-agent games by approximating continuous-time GD via Verlet integration [1] – an technique developed to approximate Hamiltonian systems in physics. The resulting approximation simply requires agents to take turns when updating in GD (alt-GD). I prove that an agent that uses alt-GD with arbitrary fixed-step size obtains $O(1)$-regret while using an arbitrary fixed step-size regardless of her/his opponent's actions. This is the first algorithm with $O(1)$-regret in general games and the first $O(1)$-regret algorithm in zero-sum games that makes no conditions on the actions of the opponent.

Representative Papers:

[1] Finite Regret and Cycles with Fixed Step-Size (COLT 2020) with G. Gidel and G. Piliouras

[2] Fast and Furious Learning in Zero-Sum Games: Vanishing Regret with Non-Vanishing Step Sizes (NeurIPS 2019) with G. Piliouras

[3] Multiagent Learning in Network Zero-Sum Games is a Hamiltonian System (AAMAS 2019 — nominated for best paper) with G. Piliouras

[4] MWU in Zero-Sum Games (EC 2018) with G. Piliouras

LUDWIG DIERKS (Homepage, CV, Google Scholar)

Thesis: Market Design for Cloud Computing

Advisor: Sven Seuken, University of Zurich

Brief Biography: Ludwig is a Ph.D. student advised by Prof. Sven Seuken in the Computation and Economics Research Group at the Department of Informatics of the University of Zurich. In 2018, Ludwig was a research intern at Microsoft in the Office of the Chief Economist. In 2015, he received an M.Sc. in Mathematics in Operations Research from TU Munich. His Master's thesis, under the supervision of Prof. Felix Brandt, was on "Cooperative Games with Payoff Restrictions." During his M.Sc., he visited the NCKU in Tainan, Taiwan (R.O.C.) for one year, under the supervision of Prof. Ruey-Lin Sheu.

Research Summary: In my research, I combine techniques from *game theory*, *artificial intelligence*, *data science* and *operations research* to understand markets and other complex systems with one or more *self-interested agents*. A special focus lies on the design of markets that are exposed to relatively unsophisticated users.

My PhD research was focused on the domain of *cloud computing*, though my current research is also branching off into a number of additional domains with their own unique challenges, such as *consumer software markets* and *adoption markets*.

One issue in cloud computing is that large amounts of capacity stand idle because they are reserved for purposes that do not require their continued use. A natural idea is to sell this idle capacity as preemptible capacity on a secondary market. In a paper with Sven Seuken [1], I focus on preemptible spot markets where users directly bid for capacity. We model the cloud provider's profit optimization problem by combining *queuing theory* and game theory to analyze the equilibria of the resulting queuing system. We show that, under a mild condition, a provider can simultaneously archive a profit increase and a Pareto improvement for the users by offering a secondary spot market.

Another challenge in cloud computing is that many modern cloud workload are characterized by scaling resource demands. A cloud provider has to continuously decide whether it is save to add additional workloads to a given compute cluster or if doing so would impact existing workloads' ability to scale. In joint work with Ian Kash and Sven Seuken [2], I formalize this "cluster admission problem" as a constrained partially observable Markov decision process (POMDP). As it is not feasible to optimally solve this POMDP, we relax it to design fast heuristic policies. Through simulations we show sizable performance gains for our policies over industry standards. We further propose the use of *variance-based pricing* to incentivize users to provide type classifications. In a paper with Sven Seuken [3], I analyze the competitive viability of variance-based pricing rules in a duopoly.

Representative Papers:

[1] Cloud Pricing: The Spot Market Strikes Back (EC'19; Management Science (forthcoming)) with S. Seuken

[2] On the cluster admission problem for cloud computing (NetEcon'19) with I. Kash and S. Seuken

[3] The Competitive Effects of Variance-based Pricing (IJCAI'20) with S. Seuken

ALON EDEN (Homepage, CV)

Thesis: Aspects of Mechanism Design: Correlation, Coordination, Competition and Pricing

Advisor: Michal Feldman and Amos Fiat, Tel Aviv University

Brief Biography: I am a postdoctoral fellow in the Harvard EconCS group under the supervision of Yiling Chen and David Parkes. I received my PhD in Computer Science from Tel Aviv University in 2019. 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: I am broadly interested in mechanism design and in algorithmic problems with economic motivation. I am fascinated by the following directions: (1) designing simple auctions with surprisingly strong guarantees, i.e. posted price mechanisms [3,4]; (2) designing auctions for settings where values can be correlated in various ways [1]; and (3) considering traditionally hard settings, and finding additional structure under which positive results can be attained [1,2].

In EC'16 [2], we present pricing schemes that achieve optimal welfare in many scenarios, including in the presence of matching markets and gross-substitutes valuations. In WINE'20, we prove an analogue of the famous Gul-Stacchetti Maximal-Domain Theorem for posted pricing schemes. In SODA'15 [1], we explore the power of dynamic pricing in online settings. We show that by using prices, one can get a near-optimal competitive ratio for a variety of well-known online problems. In ESA'18, we explore this framework for scheduling problems. In a recently submitted paper, we harness reinforcement learning methods to find good sequential pricing mechanisms for resource allocation problems.

In the Interdependent Values Setting (IDV), each agent has a private signal, and the valuation function of every agent depends on all signals. IDV generalizes the private correlated values model. In EC'18, we give the first approximately-optimal welfare and revenue maximization results that hold without single-crossing, a condition necessary to achieve full-efficiency. In EC'19 [1], we consider valuations that are submodular over signals, and obtain the first positive results for IDV in general combinatorial settings. In a recent submission, we obtain the first positive result for the PoA of simultaneous auctions in combinatorial settings with correlations.

Representative Papers:

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

[2] The Competition Complexity of Auctions: A Bulow-Klemperer Result for Multi-Dimensional Bidders (EC'17) with M. Feldman, O. Friedler, I. Talgam-Cohen, and M. Weinberg

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

[4] Pricing Online Decisions: Beyond Auctions (SODA'15) with I. Cohen, A. Fiat, L. Jez

TOMER EZRA (Homepage, CV)

Thesis: Simple Mechanisms for Limited Information Settings

Advisor: Michal Feldman, Tel Aviv University

Brief Biography: Tomer is a fifth-year Ph.D. student at the Computer Science Department at Tel Aviv University, under the supervision of Michal Feldman, where he also obtained his MSc. Prior to his MSc, Tomer obtained his BSc at the Open University of Israel.

Research Summary: My main research interests are in Algorithmic Game Theory, Mechanism Design, and Optimal Stopping Theory. During my Ph.D., I study tradeoffs between simplicity and efficiency in various settings where the optimal mechanisms are poorly understood, intractable, or cannot be realistically implemented. In addition, I am also interested in fairness in mechanisms design, welfare guarantees of static and dynamic pricing, as well as optimal linear contracts with combinatorial actions' costs.

For example, in [1] we tackle the question raised by [Feige, STOC 06] of what is the best welfare approximation can a poly-communication protocol achieve for two subadditive buyers. A $\frac{1}{2}$-approximation can be guaranteed trivially by a truthful deterministic protocol with $O(1)$ communication. We show that outperforming this protocol requires exponential communication, settling this open problem. Moreover, we show that even randomized protocols cannot achieve better.

In [3], we study the effect of cognitive biases on the efficiency and stability of markets. Specifically, we study the *endowment effect*, coined by Nobel laureate Thaler, implying that consumers tend to inflate the value of items they own. Following up on an EC 2018 work, we provide a more flexible formulation of the endowment effect, using this formulation, we overcome the impossibility results of prior works, and provide results on the existence and efficiency of market equilibrium for generalized settings, as well as settings where the market designer can pre-pack items to indivisible bundles.

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

Representative Papers:

[1] Settling the Communication Complexity of Combinatorial Auctions with Two Subadditive Buyers (FOCS 2019) with M. Feldman, E. Neyman, I. Talgam-Cohen, and M. Weinberg

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

[3] A General Framework for Endowment Effects in Combinatorial Markets (EC 2020) with M. Feldman, and O. Friedler

YIDING FENG (Homepage, CV)

Thesis: Revelation Gap in Prior-independent Mechanism Design

Advisor: Jason Hartline, Northwestern University

Brief Biography: Yiding is a fifth-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 game theory, algorithm design and machine learning, with a focus on applications in market design and operations research. He was a visiting student at Stanford University (April–October 2019) hosted by Amin Saberi.

Research Summary: In algorithmic game theory, I design and analyze mechanisms that tradeoff good performance with other desirable features (e.g. simplicity, robustness). In [1], we study the performance gap between revelation (i.e. truth-telling) mechanisms (e.g. second-price auction) and non-revelation mechanisms (e.g. first-price auction, generalized second-price auction) in prior-independent mechanism design. This finding can be used as a theoretical evidence for the prevalence of non-truthful mechanisms in practice. In [2], we introduce an algorithmic framework which transforms mechanisms for linear utility agents to analogy mechanisms for general utility agents with provable performance guarantees. For reasons of tractability, mechanism design for linear utility agents is better understood than for general utility agents. Our work builds a bridge between those two settings and can be used as a black-box to extend results from the former to the latter.

Inspired by the emergence of various online markets, I work on online algorithm design problems for different platforms, i.e., revenue management and resource allocation in rental service, cloud computing service, and ride-sharing platform. In [3], we study how the ride-sharing platform can produce more efficient matching and pricing decisions by batching the requests. We introduce a two-stage stochastic matching model, and design online competitive algorithms for driver-weighted two-stage stochastic matching for maximizing supply efficiency, and two-stage joint matching and pricing for maximizing market efficiency.

My research interests also include the intersection of machine learning and economics, e.g., global convergence guarantees for gradient-based algorithms in the dynamic discrete choice model, a classic model in econometrics [4].

Representative Papers:

[1] An End-to-end Argument in Mechanism Design (FOCS 2018) with J. Hartline

[2] Optimal Auctions vs. Anonymous Pricing: Beyond Linear Utility (EC 2019) with J. Hartline, and Y. Li

[3] Two-stage Matching and Pricing with Applications to Ride Hailing (working paper) with R. Niazadeh, and A. Saberi

[4] Global Concavity and Optimization in a Class of Dynamic Discrete Choice Models (ICML 2020) with E. Khmelnitskaya, and D. Nekipelov

ZHE FENG (Homepage, CV, Google Scholar)

Thesis: Machine-Learning Aided Economic Design

Advisor: David C. Parkes, Harvard University

Brief Biography: Zhe Feng is a fifth year PhD candidate in Computer Science at Harvard University, where he is advised by David C. Parkes. He is a recipient of the 2019 Google PhD Fellowship in Algorithms, Optimizations and Markets. He spent two summers (2019, 2020) as a Research Intern at Google and one summer (2018) as a Core Data Science Intern at Facebook. His research broadly lies in the intersection between economics and computer science, and mainly focuses on understanding how to design better mechanisms through machine learning approaches.

Research Summary: Nowadays, online markets, for adverts, products, and services, grow larger and larger everyday. Market designers face challenges in designing a good market not only in regard to scalability, but also in seeking to use data to better understand the behavior of market participants. Motivated by these challenges, my research mainly focuses on two research questions, *economic design via machine learning* and *learning in online markets*.

Optimal economic design, especially auction design, has received a lot of attention in computer science in recent years. However, the design of optimal, dominant strategy incentive compatible (DSIC) mechanisms has remained largely open. Following the research line of data-driven, automated mechanism design, I have proposed three neural network architectures to approximately model optimal auctions [1]. RegretNet is a general neural network architecture, able to handle multi-item and multi-bidder settings. Simulations show that RegretNet can learn almost optimal auctions for essentially all settings for which there are known analytical solutions, and obtain novel mechanisms for settings in which the optimal mechanism is unknown. In addition, I have developed generalization results for both the expected revenue and expected regret of learned auction rules. RegretNet outputs an approximately-IC mechanism. In other work [2], I propose an approach to transform an approximately-IC mechanism into a fully BIC mechanism without loss of welfare and with only negligible loss in revenue. This work [1] has opened an area of economic design through deep learning and inspired a lot of follow-up work.

Another line of my research aims to understand the learning dynamics of strategic agents in online markets. In my work [3], I analyze online learning in auctions, where bidders need to learn their value for items and is only available when winning an auction. I develop algorithms with regret against the best fixed bid that converges exponentially faster, in terms of dependence on the action space, than what would follow from applying a generic bandit algorithm.

Representative Papers:

[1] Optimal Auctions through Deep Learning (ICML'19, Research Highlight in the Communications of the ACM, ArXiv) with P. Dütting, H. Narasimhan, D. C. Parkes, and S. S. Ravindranath

[2] Welfare-Preserving ε-BIC to BIC Transformation with Negligible Revenue Loss (ArXiv) with V. Conitzer, D. C. Parkes, and E. Sodomka

[3] Learning to Bid Without Knowing your Value (ACM-EC'18) with C. Podimata, V. Syrgkanis

MATHEUS VENTURYNE XAVIER FERREIRA (Homepage, CV)

Thesis: Decentralized Algorithms for Strategic Agents

Advisor: S. Matthew Weinberg, Princeton University

Brief Biography: Matheus is currently a Ph.D. candidate in computer science at Princeton University. Previously, Matheus received his bachelor's degree (2016) in computer engineering at Universidade Federal de Itajubá, Brazil, with a year abroad at the University of California, San Diego, via a BSMP fellowship (2014). His research area is broadly in algorithmic game theory and information security with applications in auction theory and cryptocurrencies. His research also encompasses policy and the incentives behind adopting better security practices that will motivate future decisions. He has served in the program committee for the Cryptoeconomic Systems journal (MIT Press), a peer mentor for first-year graduate students, a research mentor for four undergraduate research projects, and a peer educator. He is the recipient of a First-Year Fellowship in Engineering and a Dean's Grant at Princeton and is a LATinE fellow awarded by Purdue University.

Research Summary: Algorithmic Mechanism Design traditionally focuses on designing algorithms implemented by a central, trusted authority. Centralization minimizes inefficiencies, and the resulting mechanisms are simpler to analyze. However, with the increasing popularity of online auctions, transparency has become a first-order concern for users. Moreover, blockchains emerged as a decentralized economic system where the absence of central trust provides new business opportunities. In this direction, my thesis describes the challenges and offers new directions towards designing mechanisms with less trust on the mechanism designer.

My research proposes novel strategyproof auctions that are credible: where any deviation of the auctioneer from the promised auction is either detectable or unprofitable. Without additional assumptions, it was previously known that any credible, strategyproof optimal auction requires unbounded communication with the auctioneer. I showed how to instantiate communication efficient, credible, and strategyproof optimal auctions with mild cryptographic assumptions.

Blockchains and smart contracts are useful primitives for designing decentralized mechanisms. However, increasing network congestion undermines the robustness of any mechanism implemented on the blockchain, and block space is often allocated via a first-price auction that introduces inefficiencies and carries a high cognitive cost to agents. My research explores new block space auctions by leveraging miners' decentralization. On another front, my research shows that energy-efficient Proof-of-Stake (PoS) consensus can approximately preserve the incentives Proof-of-Work (PoW) provides to miners against withholding attacks when PoS has access to an unbiased and unpredictable source of randomness. I also show that, under similar assumptions, PoS does not always preserve the same incentives of PoW.

Representative Papers:

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

[2] Selling a Single Item with Negative Externalities (WWW 2019) with T. Chattopadhyay, N. Feamster, D. Y. Huang, and S. M. Weinberg

[3] Proof of Stake Mining Games with Perfect Randomness with S. M. Weinberg

KIRA GOLDNER (Homepage, CV, Google Scholar)

Thesis: Mechanism Design for a Complex World: Rethinking Standard Assumptions

Advisor: Anna Karlin, University of Washington

Brief Biography: Kira Goldner is both an NSF Mathematical Sciences and Data Science Institute postdoctoral fellow at Columbia University, hosted by Tim Roughgarden. She received her PhD in CS from the University of Washington under the advisement of Anna Karlin, during which she was supported by a 2017-19 Microsoft Research PhD Fellowship and a 2016-17 Google Anita Borg Scholarship. Kira co-founded Mechanism Design for Social Good (MD4SG), an interdisciplinary initiative working to improve access to opportunity for historically disadvantaged communities, which has grown into a community of over 2300 researchers globally. She has given tutorials on Open Problems in Mechanism Design for Social Good (WINE 2017) and Menu Complexity (EC 2018).

Research Summary: My research is in algorithmic mechanism design, where I primarily study multi-parameter settings, where it is unclear whether a problem is even tractable to solve. In addition, much of my recent work has been on working in mechanism design for social good, both in pursuing an understanding of the theoretical foundations and on solving domain-specific problems.

Within multi-parameter mechanism design, my work focuses on two primary themes. First, my work characterizes optimal mechanisms in *interdimensional* settings, a subclass we introduce of natural and ubiquitous multi-dimensional settings, e.g. where a seller offers items that vary in levels of quality [FGKK16, DGSSW20]. The complexity of maximizing revenue in these settings lies strictly between the "easy" case of single-parameter and the fully general "multi-dimensional" case of heterogenous items. Second, I design practical, *simple* yet near-optimal mechanisms whose guarantees are *robust* to various informational and behavioral assumptions, such as risk attitudes [CGMP18], interdependence [EFFG18, EFFGK19], complementarities [CDGM19], and lack of detailed information [GK16].

My recent work examines questions regarding access to opportunity and mitigation of social harms where incentives must be considered. Whether the designer is the government allocating licenses to limit carbon emissions [GIL20] or an employer contracting with health insurance providers to cover their patients [EGW20], it is still crucial to ensure that our systems give guarantees in the presence of strategic participants. Additionally, I am working on foundational questions for social good objectives, such as gains from trade, which measures the value that platforms (such as online labor markets) bring to their participants [BGG20, CGMZ20].

Representative Papers:

[1] Reducing Inefficiency in Carbon Auctions with Imperfect Competition (ITCS 2020) with N. Immorlica and B. Lucier

[2] Combinatorial Auctions with Interdependent Valuations: SOS to the Rescue (EC 2020: Best Paper with a Student Lead Author) with A. Eden, M. Feldman, A. Fiat, and A.R. Karlin

[3] Optimal Mechanism Design for Single-Minded Agents (EC 2020: Best Presentation by a Student or Postdoctoral Researcher) with N.R. Devanur, R.R. Saxena, A. Schvartzman, and S.M. Weinberg

YANNAI A. GONCZAROWSKI (Homepage, CV, Google Scholar)

Thesis: Aspects of Complexity and Simplicity in Economic Mechanisms

Advisors: Noam Nisan and Sergiu Hart, The Hebrew University of Jerusalem

Brief Biography: Yannai is a postdoc at MSR. He holds a PhD in Math, CS, and Rationality. He is also a professionally-trained opera singer, holding Master's and Bachelor's degrees in Classical Singing. Both as a lecturer and as a TA, he appeared repeatedly on the Outstanding Teachers list at the Hebrew U., on one year being ranked first among all lecturers and TAs in the CS department. His PhD thesis was recognized with several awards, including the Maschler Prize of the Israeli Game Theory Society and the ACM SIGecom Doctoral Dissertation Award. He is also the recipient of the ACM SIGecom Award for Best Presentation by a Student or Postdoc at EC'18 and of the 2020 INFORMS AMD Rothkopf Junior Researcher Paper Prize. Yannai was the (Virtual) Local Chair of EC'20, and his Erdős–Bacon Number is 5.

Research Summary: The explosion in online and computerized economic activity necessitates the study of economic mechanisms of unprecedented scale. How good can simple mechanisms be? How complex must optimal mechanisms be? And, most substantially, what are the precise trade-offs between simplicity and performance?

In papers in STOC'17 and STOC'18 [1], we study the *complexity of describing auctions*, bounding the menu-size of approximately optimal auctions. In a GEB'19 paper, we analyze the *communication complexity* of stable matching mechanisms.

In papers in STOC'17 and FOCS'18 [2], we study the *sample complexity of learning* up-to-ε optimal auctions, focusing on both computationally efficient results and information-theoretic results. We resolve standing open questions and show for the first time that up-to-ε optimal auctions can be learned from polynomially many samples from the buyers' type distributions even in multi-dimensional settings.

Fairness and inclusion call for mechanisms that are *simple to understand and participate in.* In papers in JET'18 [3] and EC'17, we characterize all "obviously strategyproof" (Li, 2018) mechanisms in canonical settings. In an EC'14 paper, I analyze what can be called the *complexity of manipulating* stable matching mechanisms.

On the applied side, we designed and deployed the Israeli Gap-Year Programs match [4]. This market has been run annually since 2018, and in 2020 matched 1,937 high-school seniors (out of a total of 6,118 candidates) to 42 different programs.

Representative Publications:

[1] Bounding the Menu-Size of Approximately Optimal Auctions via Optimal-Transport Duality (STOC 2018)

[2] The Sample Complexity of Up-to-ε Multi-Dimensional Revenue Maximization (FOCS 2018; minor revision requested at JACM) with S.M. Weinberg

[3] Stable Matching Mechanisms are Not Obviously Strategy-Proof (JET 2018) with I. Ashlagi

[4] Matching for the Israeli "Mechinot" Gap-Year Programs: Handling Rich Diversity Requirements (EC 2019; Best Paper at MATCH-UP 2019) with L. Kovalio, N. Nisan, and A. Romm

[5] Mathematical Logic through Python (textbook, forthcoming in Cambridge University Press) with N. Nisan

MAURICIO GONZALEZ-SOTO (Homepage, CV)

Thesis: Use and acquisition of causal relations for decision making under uncertainty using imperfect information games

Advisor: Luis Enrique Sucar and Hugo J. Escalante, Instituto Nacional de Astrofísica Óptica y Electrónica.

Brief Biography: Mauricio received an Applied Mathematics Bachelors Degree as well as a Masters Degree in Data Science from the Instituto Tecnológico Autónomo de México (ITAM), where he got awarded a Special Mention for each of his theses. His Bachelors thesis was about a trayectorial decomposition of a certain type of stochastic process called a Continuous State Branching Process, and his Masters thesis was about Complex Networks and their use for the study of social networks.

Research Summary: Mauricio started his resarch by defining a *Causal Decision Problem* and studied a preliminar solution criterion for it. Using such criterion he was able to design a learning algorithm which found optimal actions in stochastic environments controlled by an unknown causal model with the same performance as the classical Q-learning, which was also able to learn a causal model of the environment. Then, Mauricio further extended his solution criterion to the case in which a decision maker does not know the causal model which controls her environment so she must use beliefs about possible causal models and, for each model in the support of her distribution, find the optimal action using an analogue of the preliminar solution, thus obtaining a causal version of Savage's Theorem. Later, Mauricio studied the extension of his decision-making results to the multi-agent setting which led him to define Causal Games, in which each players actions' are causally connected to consequences (outcomes) by an unknown causal model. Applying his previous results Mauricio defined the notion of a Causal Nash Equilibrium. He is now working on a random graph approach to causal structure discovery

Representative Papers:

[1] Causal Structure Learning using Random Graphs (Work in Progress: to be submitted to The 2020 NeurIPS Workshop on Causal Discovery and Causality-Inspired Machine Learning)
with Ivan Feliciano Avelino, L. Enrique Sucar and Hugo J. Escalante

[2] Von Neumann-Morgenstern and Savage Theorems for Causal Decision Making (Submitted to Artificial Intelligence Journal and oral presentation at New in ML Forum co-located with NeurIPS 2019) with L. Enrique Sucar and Hugo J. Escalante

[3] Causal Games and Causal Nash Equilibrium (1st Causal Reasoning Workshop at the Mexican International Conference on Artificial Intelligence 2019) with L. Enrique Sucar and Hugo J. Escalante

[4] Playing against Nature: causal discovery for decision making under uncertainty (Machine Learning for Causal Inference, Counterfactual Prediction and Autonomous Action (CausalML) Workshop at ICML 2018) with L. Enrique Sucar and Hugo J. Escalante

ALECK JOHNSEN (Homepage, CV)

Thesis: Information Design and Inference in Auctions

Advisor: Jason Hartline, Northwestern University

Brief Biography: Working in industry before graduate school, Aleck has accumulated 10 years of experience in trading, covering most roles: floor clerk and trader, algorithmic trader, algorithm designer, data scientist, and programmer; involving equities, fixed income, and options. He has a CFA charter and has done investor-side venture analysis of startups. During early graduate summers, Aleck wrote a C# e-file module (for transmitting to the IRS) for Forte International Tax (Evanston, IL). For 12 weeks (2016), Aleck visited Booking.com (Amsterdam) with advisor Jason Hartline and colleague Denis Nekipelov (Virginia), where they helped Booking's ranking team apply principles of auction theory. At Northwestern, Aleck taught EECS 214 Data Structures and Data Management. His undergrad degrees were in civil engineering; and mathematics; at University of Illinois (Urbana-Champaign).

Research Summary: Since visiting Booking.com, Aleck's research has focused on auction inversion techniques, covering both "online" implementation theory and "offline" inference. Towards implementation, [1] considers the design of a predictive interface (e.g., a dashboard) in order to get both agents wanting to best respond to predictions, and mechanism inference of values to run a desired social choice rule. An unpublished work for a first restricted setting solves a partial-differential-equation mapping from non-truthful agent bids back to private values such that by implementing a desired choice rule, the side-effects (e.g., payments) are consistent with the original agent reports. Towards inference, [2] shows how to infer agents' private values in auctions from the allocation rule and the prices they pay.

Aleck's research has also focused on robust mechanism design, including broadly in [3], which defines a new *counterfactual* regret concept to better capture informed-online-algorithm rationality in a principal-agent model; and fundamentally in [4], which identifies an equivalence between prior independent and prior free mechanism design via design of the latter's benchmark function, and applies the equivalence to a new solution for a canonical 2-agent auction problem. Further unpublished work applies Yao's Minimax Principal specifically in the prior independent setting which has similar structure to both tensor decomposition and to the economics topic of information design: if a correlated distribution (over n inputs) can be "decomposed" in two distinct ways – i.e., into distinct distributions over i.i.d. product distributions – then a lower bound on algorithm approximation follows as a consequence.

Representative Papers:

[1] Dashboard Mechanisms for Online Marketplaces (EC'19)
with J. Hartline, D. Nekipelov, and O. Zoeter

[2] Inference from Auction Prices (SODA'20)
with J. Hartline, D. Nekipelov, and Z. Wang

[3] Mechanisms for a No-regret Agent: Beyond the Common Prior (FOCS'20)
with M. Camara and J. Hartline

[4] Benchmark Design and Prior Independent Optimization (FOCS'20)
with J. Hartline and Y. Li

ANILESH K. KRISHNASWAMY (Homepage, CV)

Thesis: Robust and Fair Decision Making

Advisor: Vincent Conitzer and Kamesh Munagala, Duke University

Brief Biography: Anilesh Krishnaswamy is a Postdoctoral Associate in Computer Science at Duke University. His current research focuses on the game-theoretic and fairness aspects of machine learning, especially in settings with missing data and strategic behavior. He obtained his PhD from Stanford University, advised by Prof. Ashish Goel, with a thesis on the theoretical and practical aspects of online platforms that enable democratic decision-making. He is one of the designers of the Stanford Participatory Budgeting Platform (https://pbstanford.org/), a digital voting tool that is used by dozens of cities across the US.

Research Summary: Given the promise of learning algorithms in transforming societal decision making, it is imperative that we ensure that they are fair to all stakeholders involved. Practical mechanisms also need to be robust to mis-specified or under-specified model parameters, as well as missing data and strategic behavior. Therefore, it is important to work toward a broader theory of robustness that can simultaneously handle all the above concerns.

In machine learning, various group-based statistical notions of fairness have been proposed. However, in such approaches, no guarantees are provided for groups that the algorithm designer cannot foresee, and it is possible that the classifier performs relatively worse on some of them. To get around this issue, I'm interested in devising algorithms that achieve fairness guarantees for broad classes of groups (e.g. those with bounded VC dimension), and testing these methods on real datasets [1].

I am also interested in understanding how to quantify the fairness of social choice mechanisms, another pertinent question that has received little attention. To this end, I have worked on developing a novel method of quantifying fairness under the *metric distortion* framework (the agents and alternatives are assumed to lie in an unknown metric space, and the algorithm has access only to the rankings over the alternatives provided by the agents, with closer alternatives being ranked higher), using ideas from multi-objective optimization [2].

When machine learning techniques are used in applications such as credit approval and college admission, an agent may decide to strategically withhold some of her features (such as bad test scores) so as to be classified more favorably. This issue is compounded when some of the data is naturally missing. In recent work [3], I study this problem with the goal of designing algorithms that are robust to the strategic withholding of data. Testing these algorithms on real datasets leads to interesting observations on how different algorithms are useful in different scenarios based on the number of features, their discreteness, etc. [3].

Representative Papers:

[1] Fair for All: Proportional Fairness in Classification (Preprint)
with Y. Cheng, Z. Jiang, K. Munagala, and K. Wang

[2] Metric distortion of social choice rules: Lower bounds and fairness (EC 2017)
with A. Goel and K. Munagala

[3] Classification with Strategically Withheld Data (Preprint)
with H. Li, D. Rein, H. Zhang, and V. Conitzer

PHILIP LAZOS (Homepage, CV)

Thesis: Online Algorithms for Markets

Advisor: Elias Koutsoupias, University of Oxford

Brief Biography: Philip Lazos is a postdoctoral researcher at the Sapienza University of Rome, hosted by Stefano Leonardi. Prior to that, he was a postdoc at the University of Oxford, where he also completed his DPhil (more commonly referred to as PhD) in 2018, under the supervision of Elias Koutsoupias. He obtained his undergraduate degree in Electrical Engineering in 2015 at the National Technical University of Athens.

Research Summary: My research is at the intersection of mechanism design, online algorithms and cryptocurrencies. I am interested in developing algorithms that cope with *uncertainty*, which can take many forms, such as incomplete information about future inputs, strategic behavior by selfish agents or unknown probability distributions. In many modern algorithmic settings the entities involved act rationally and try to improve their gain by strategically interacting with their environment, rather than simply cooperating as prescribed, which is the usual assumption in algorithm design. This is becoming an increasingly important issue in our world (e.g., managing the COVID-19 pandemic), and part of my research deals with these practical considerations as well.

A large part of my research deals with improving efficiency and robustness in markets. Starting by considering a more general market setting where agents can come and go (adversarially or at random), but can trade only one item, followed with improving the bounds known for many items assuming submodular valuations. The culmination of this line of research is [4], which improves on the best known guarantees despite using a single sample rather than prior distributions.

I have also worked in purely algorithmic settings, mostly in online algorithms and optimal stopping. My main focus has been on k-server based problems, such as a weak adversaries version or a combination with facility location. I have also extensively worked on secretary problems and prophet inequalities, leading to a recent development in the related Pandora's Box problem [1].

Finally much of my motivation stems from studying (and attempting to improve) the real-world effects of emerging technologies, such as exploring the world of blockchains and cryptocurrencies or automating group testing to contain the effects of the COVID-19 pandemic, for which we have received funding from the Global Challenges in Economics and Computation workshop.

Representative Papers:

[1] Pandora's Box Problem with Order Constraints (EC 2020), with Shant Boodaghians, Federico Fusco and Stefano Leonardi

[2] The Pareto Frontier of Inefficiency in Mechanism Design (WINE 2019), with Aris Filos-Ratsikas, Yiannis Giannakopoulos

[3] Fairness and Efficiency in DAG-based Cryptocurrencies (FC 2020), with Georgios Birmpas, Elias Koutsoupias and Francisco J. Marmolejo-Cossío

[4] Efficient Two-Sided Markets with Limited Information, *under submission*, with Paul Dütting, Federico Fusco, Stefano Leonardi and Rebecca Reiffenhäuser

BAR LIGHT (Homepage, CV)

Thesis: Market Design for Platforms and the Analysis of Large Games

Advisor: Gabriel Weintraub and Ramesh Johari, Stanford University

Brief Biography: I am a Ph.D. candidate at the Stanford Graduate School of Business under the supervision of Gabriel Weintraub and Ramesh Johari. I previously obtained an M.Sc. in Operations Research from Tel Aviv University under the supervision of Ehud Lehrer.

Research Summary: In recent years the amount of data collected by online platforms has increased massively. These data, together with the unique ability of online platforms to design their marketplaces, provide platforms with an unprecedented opportunity to make better *market design* choices. One branch of my research aims to provide tools and insights to make and understand these market design choices better. In my work with Ramesh Johari and Gabriel Weintraub [1] we study the following market design problem in a two-sided market: which sellers should the platform allow to participate and how much of its available information about participant sellers' quality should the platform share with buyers? Our main contribution is to provide conditions under which simple design choices commonly observed in practice, such as banning certain sellers from the platform while not distinguishing between participating sellers, maximize the platform's revenue.

Another branch of my research aims to provide theoretical tools that can be used by OR researchers and others in applied work. In particular, because online markets typically consist of many small buyers and sellers, I am interested in modeling, solving, and studying the properties of *large games*. In my work (to appear in *OR*) with Gabriel Weintraub [2], we study mean field equilibria (MFE). Our main contributions are finding conditions that ensure that the MFE is unique and applying our uniqueness result to dynamic oligopoly and macroeconomic models commonly used in previous literature. In two previous works, [3] and [4], I study the properties of large incomplete dynamic market models that are widely studied in the literature. The results in these papers shed light on the impact of market incompleteness on the properties of general equilibrium. I have also worked on dynamic optimization and applied probability. In my work [5] (to appear in *Math of OR*) I develop tools that enable deriving comparative statics results in complex dynamic environments. In my work [6] I improve the important Hoeffding's inequalities using information on the random variables' higher moments.

Representative Papers:

[1] Quality Selection in Two-Sided Markets: A Constrained Price Discrimination Approach (Working Paper) with R. Johari, and G. Weintraub

[2] Mean Field Equilibrium: Uniqueness, Existence, and Comparative Statics (Operations Research) with G. Weintraub

[3] A Bewley-Huggett Model with Many Consumption Goods (R&R in Operations Research - special issue honoring Kenneth Arrow)

[4] Uniqueness of Equilibrium in a Bewley-Aiyagari Model (Economic Theory)

[5] Stochastic Comparative Statics in Markov Decision Processes (Math of OR)

[6] Improving Hoeffding's Inequalities Using Higher Moments Information (arXiv)

JIAQI LU (Homepage, CV)

Thesis: Managing Stochastic Uncertainty in Dynamic Marketplaces

Advisor: Yash Kanoria, co-advised by Awi Federgruen, Columbia Business School

Brief Biography: Jiaqi is a Ph.D. candidate at Columbia Business School in the Decision, Risk, and Operations Division, where she is advised by Yash Kanoria and Awi Federgruen. She obtained her M.S. in Management Science and Engineering at Columbia University. Before that, she obtained B.E. in Industrial Engineering and B.A in English double major at Tsinghua University.

Research Summary: My research focuses on the modeling and optimization of operational challenges that arise in marketplaces as a result of uncertainty in various forms. Specific directions include the analysis and design of large matching markets, and the optimization of tactical decisions such as supply chain management and customer relationship management.

Research on the design of matching markets has focused primarily on assignments/mechanisms that achieve stability and market-clearing. However, search frictions and incomplete information may prevent the market from clearing in many contexts. For example, when a number of agents wait for each other to acquire information first, an "information deadlock" can occur. In [1] we identify a wide range of random sparse matching markets that suffer from information deadlock in the large market limit. We find that the occurrence of information deadlock crucially depends on the opportunity-to-agent ratio. We observe a phase transition from the deadlock regime to the deadlock-free regime when this ratio increases. The analysis develops tools inspired by methodology from statistical physics and allied areas.

For service firms, stochastic churn of customers is a central challenge. In [2], we show that the firm can increase customer lifetime value (CLV) by dynamically adjusting its service strategy based on the customer's current risk of churn: We study an online decision making problem to optimize CLV for a service (e.g., robo-advising) firm via a risky arm and a safe arm. The customer's churn probability at any time is a function of his recent-biased estimate of past experience. We fully characterize the firm's optimal policy, which reveals that the firm should be risk-averse when the customer is marginally satisfied and risk-seeking when the customer is marginally unsatisfied.

I also conduct research in [3] on how a firm should dynamically adjust its offshoring/onshoring sourcing decisions to a stochastically changing economy state that affects cost differential, demand, etc.. We demonstrate that volatility of the economy state may be exploited to the firm's benefit.

Representative Papers:

[1] When Does Competition and Costly Information Acquisition Lead to a Deadlock? (with manuscript) with N. Immorlica and Y. Kanoria

[2] Managing Customer Churn via Service Mode Control (Reject and Resubmit) with Y. Kanoria and I. Lobel

[3] Sourcing in an Increasingly Volatile World: Offshoring, Onshoring or Both? (in preparation) with A. Federgruen and Z. Liu

DIVYARTHI MOHAN (Homepage, CV)

Thesis: Simplicity and Optimality in Algorithmic Economics: Multi-Item Auctions and Opinion Dynamics in Social Networks

Advisor: S. Matthew Weinberg, Princeton University

Brief Biography: Divyarthi Mohan is a PhD candidate in Computer Science at Princeton University, advised by Prof. Matt Weinberg. She is a recipient of the 2021 Siebel Scholars Award. Divyarthi was also awarded the School of Engineering and Applied Science's Award for Excellence in 2019 and the Department of Computer Science's Graduate Student Teaching Award in 2018. During her PhD, Divyarthi was a summer research intern at Microsoft (2019) and Google (2020). Previously, she obtained her Master's in Theoretical Computer Science from The Institute of Mathematical Sciences, Chennai.

Research Summary: My primary interests are understanding simplicity in algorithmic economics and when simple mechanisms produce *good outcomes*. My research focuses on Multi-item Mechanisms Design, and in this line of work we developed the first sub-exponential *approximation scheme*.

Designing mechanisms to maximize revenue is a fundamental problem in mathematical economics and has various real world applications like online ad auctions, spectrum auctions, and pricing bundles of goods. Unfortunately, optimal auctions for selling multiple items are extremely complex and are intractable by numerous formal measures. Our work [1] shows that simple mechanisms can achieve almost optimal revenue, in the single unit-demand buyer and n independent items setting. We approached the trade-offs of simplicity formally through the lens of *computational complexity* and *(symmetric) menu complexity*. Our main result provides a mechanism that gets a $(1-\epsilon)$-approximation to the optimal revenue in time quasi-polynomial in n. In ongoing work (with the Market Algorithms team at Google), we are studying auction design for ads with rich size constraints, which we model as *1.5-dimensional*. Our goal is to design simple auctions for revenue and welfare.

My other research interest is Social Learning and understanding the effects of various simple dynamics. In [2], we studied how information aggregates in networks through simple majority dynamics. We consider a model where agents make binary decisions, and each agent gets a private signal denoting which decision is better. To make an informed decision they interact with their neighbors. More recently in [3], we study social learning between two Bayesian agents (a sender and a receiver) with limited/structured communication. In our model, to reflect real social exchanges, agents simply share a signal they observed instead of sending arbitrary messages.

Representative Papers:

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

[2] Asynchronous Majority Dynamics in Preferential Attachment Trees (ICALP 2020) with M. Bahrani, N. Immorlica, and S.M. Weinberg

[3] Understanding Anecdotal Communication (In Preparation)
with N. Haghtalab, N. Immorlica, B. Lucier, and M. Mobius

DOMINIK PETERS (Homepage, CV, Google Scholar)

Thesis: Fair Division of the Commons

Advisor: Edith Elkind, University of Oxford

Brief Biography: I am a postdoc at Harvard University (and previously Carnegie Mellon University) with Ariel Procaccia. I received my DPhil in 2019 at the University of Oxford, advised by Edith Elkind. My dissertation won the Victor Lesser Distinguished Dissertation Award and the EurAI Distinguished Dissertation Award.

Research Summary: I work on computational social choice. Recently, I have focussed on budgeting problems: a group needs to decide how to use common resources, by voting over which uses are best. Many cities run *Participatory Budgeting* (PB) elections, where residents vote over project proposals to be funded by part of the government budget. I'm particularly interested in ensuring that these budgeting decisions are *proportional*, in the sense that subpopulations with shared interests are adequately represented in the funding decision. Most real-world PB elections use variants of welfare-maximization, and such systems can leave large groups of voters without any representation. In recent work with Piotr Skowron [1], we have formalized the notion of proportionality as it applies to PB; as it turns out, there appear to be two distinct and incompatible interpretations of the concept. We argue that one of them is better, and have designed a promising voting rule satisfying it. These results apply to cases where projects are discrete and need a specific amount of money like in most PB instances. In other contexts, projects can be flexibly scaled, motivating a divisible or fractional model. We have shown that such models also admit attractive proportional voting rules [2, 3].

Another research interest is proving new impossibility theorems in voting theory, using a computer-aided method: encode axioms in a logical language and then use a solver (e.g., a SAT solver) to check whether some voting rule satisfies them. Amazingly, it is often possible to extract human-readable impossibility proofs from this approach, using minimal unsatisfiable sets. For example, we can show that no PB voting rule is strategyproof [4]. The method can also sharpen an impossibility due to Moulin: we can show that every Condorcet voting rule suffers from the *no-show paradox* when there are ≥ 12 voters, but that the paradox can be avoided for ≤ 11 voters [5]. Moulin proved the impossibility for 25 voters, but the computer-generated proof for 12 voters is arguably more elegant than the human one.

Representative Papers:

[1] Proportionality and the Limits of Welfarism (EC 2020)
with P. Skowron

[2] Truthful Aggregation of Budget Proposals (EC 2019)
with R. Freeman, D. M. Pennock, J. Wortman Vaughan

[3] Portioning using Ordinal Preferences: Fairness and Efficiency (IJCAI 2019)
with S. Airiau, H. Aziz, I. Caragiannis, J. Kruger, J. Lang

[4] Proportionality and Strategyproofness in Multiwinner Elections (AAMAS 2018)

[5] Optimal Bounds for the No-Show Paradox via SAT Solving (Mathematical Social Sciences 2017, and AAMAS 2016) with F. Brandt, C. Geist

PENGYU QIAN (Homepage, CV, Google Scholar)

Thesis: Online Decision-Making in Networked Marketplaces

Advisor: Yash Kanoria, Columbia University

Brief Biography: Pengyu Qian is a Ph.D. student in Graduate School of Business at Columbia University, where he is advised by Yash Kanoria. He received B.S. in Mathematics from Peking University in 2015.

Research Summary: My research studies networked marketplaces with an emphasis on online decision-making in such marketplaces, using tools from applied probability and modern optimization. I am interested in foundational theoretical models motivated by problems in revenue management and pricing, and matching markets. My research emphasizes algorithms and mechanisms that not only have good theoretical guarantees, but also are simple, robust, and hence practical for real-world systems.

Matching platforms must dynamically manage and balance the "inventory" of (heterogeneous) agents on the two sides of the market, in a stochastic environment. For instance, in shared transportation (e.g., ride-hailing) systems, vehicles circulate in the network, agents are (primarily) distinguished by geographical location, and a platform must efficiently manage the geographical distribution of available vehicles using available control levers (admission/pricing/dispatch control). Even in this "simple" setting, prior state-of-the-art policies required perfect knowledge of future demand arrival rates (else their performance is far from optimal) and were therefore of limited practicality. In [1,2], we introduce simple and practical control policies that make decisions based only on the current geographical distribution of vehicles. We show that these policies are near-optimal, even without knowing demand arrival rates. Notably, [2] introduces a novel "Mirror Backpressure" (MBP) control methodology which significantly generalizes the celebrated backpressure policy (a "congestion-aware greedy" approach) for control of queueing networks. MBP is highly flexible, comes with provable guarantees, and may be broadly useful.

In another line of work, I am interested in improving our understanding of equilibria in marketplaces [3,4], towards designing better marketplaces. In [3], we study the waiting-list mechanism, which is widely used in object allocation problems such as the assignment of public housing. We show allocative efficiency of this mechanism in settings with an arbitrary number of agent/item types. Our analysis overcomes the curse of dimensionality by establishing a novel connection with stochastic gradient descent, and may be of independent interest.

Representative Papers:

[1] Dynamic Assignment Control of a Closed Queueing Network under Complete Resource Pooling (SIGMETRICS'18) with S. Banerjee, Y. Kanoria

[2] Blind Dynamic Resource Allocation in Closed Networks via Mirror Backpressure (EC'20) with Y. Kanoria

[3] Price Discovery and Efficiency in Waiting Lists: A Connection to Stochastic Gradient Descent (EC'20) with I. Ashlagi, J. Leshno, A. Saberi

[4] Which Random Matching Markets Exhibit a Stark Effect of Competition? (SODA'21) with Y. Kanoria, S. Min

SCOTT RODILITZ (Homepage, CV)

Thesis: Managing Growth and Efficiency on Crowdsourced Online Platforms

Advisor: Vahideh Manshadi, Yale University

Brief Biography: Scott is a fifth-year PhD student at the Yale School of Management, where he is advised by Vahideh Manshadi. He is the third-place winner of the Rothkopf Junior Researcher Prize (awarded by the INFORMS Auction & Market Design Section) and a two-time finalist for the INFORMS Public Sector OR Best Paper Award (in 2019 and 2020). In 2019, he was a visitor at the Simons Institute for the Theory of Computing for the program on Online and Matching-Based Market Design, and in that summer he worked as a research science intern at Lyft. Before Yale, he graduated from Williams College in 2013 with a BA in Mathematics and Economics and taught for three years at Kingswood Oxford School.

Research Summary: In my research, I shed light on how public and civic sector organizations can improve the design of their online platforms, with a particular focus on the utilization of *crowdsourcing* [1] and *platform-based markets* [2,3].

My research combines insights from data with rigorous modeling frameworks to provide actionable recommendations designed to address unique features of the public and civic sectors. Informed by my experience as a research science intern at Lyft, I employ interpretable machine learning approaches to unlock the (often untapped) potential of data outside the private sector. I then incorporate that data into models of the complex stochastic processes which govern crowdsourcing and platform-based markets. Though grounded by practical considerations, these models lead to broad theoretical insights in addition to specific platform design recommendations which draw on techniques from the analysis of online algorithms.

As one example, in [2] we address a fundamental design question for notification systems through the lens of Food Rescue U.S. (FRUS), a leading online food recovery platform. In order to ensure the successful transportation of perishable donated food, FRUS sends last-minute notifications to a subset of its volunteers. Though notifying more volunteers increases the likelihood that someone responds, it also contributes to *notification fatigue*, a common issue for crowdsourcing platforms. We model notification decisions as a generalization of online bipartite matching, and leveraging historical data, we design a policy which provides strong theoretical guarantees as well as substantial improvement over the current practice at FRUS. Our framework and data-driven approach can also be utilized in private sector applications where tensions arise between maximizing an immediate payoff and limiting the negative impact of excessive notifications (e.g. customizing online notification systems commonly used for marketing and social network engagement).

Representative Papers:

[1] Diffusion in Random Networks: Impact of Degree Distribution (*Operations Research* 2020) with V. Manshadi and S. Misra

[2] Online Policies for Efficient Volunteer Crowdsourcing (EC 2020, under review at *Management Science*) with V. Manshadi

[3] Commitment on Volunteer Crowdsourcing Platforms: Implications for Growth and Engagement (under review at *Manufacturing & Service Operations Management*) with I. Lo, V. Manshadi, and A. Shameli

MASHBAT SUZUKI (Homepage, CV, Google Scholar)

Thesis: Fair Division , Social Choice

Advisor: Adrian Vetta, McGill University

Brief Biography: Mashbat Suzuki is a final year PhD candidate at McGill University, where he is supervised by Adrian Vetta. He recieved B.A and M.Sc in Mathematics at McGill University focusing on manifold learning. His research lie in economics and computation, specifically areas of fair division, resource allocation and social choice.

Research Summary: A major part of my thesis is focused on studying fair division of indivisible items. When items are indivisible, envy-free allocation is often impossible to achieve. As a result, there has been much focus on weaker notions of fairness such as EF-1 and maximin share, etc. In our work [1], we show that exact envy-free allocation is possible if we include a small amount of subsidy payments to each agent. Specifically, we show that there exist an allocation and small amount of payments (only depending on the number of agents) such that each agent prefers its bundle plus payment over the bundle plus payment of every other agent. Furthermore, we show that such allocation and payments can be computed efficiently. Our work settle a conjecture of Helpern and Shaw, en route we show interesting connections between fair division of divisible goods, and indivisible goods.

Building upon this work, in [3] we study the problem of simultaneously achieving envy freeness and efficiency when we are allowed to introduce payments to each agent. This is a natural extension of our work in [1], where we focused mainly on fairness. We show that amount of payments/transfers required to achieve both envy freeness and high social welfare is large. Furthermore, we give an algorithm that simultaneously achieves desired social welfare and envy properties while the total payment is almost as small as possible.

Another major topic of my research focuses on voting in social groups. Specifically, we study the evolution of a social group when admission to the group is determined via consensus or unanimity voting. In [2], we study the setting where in each time period, the candidates appear uniformly over a metric space. We prove, on a unit ball the expected cardinality of the social group after T time periods is $\Theta(T^{\frac{1}{8}})$, while on the unit square it is least $\Omega(\ln T)$ but at most $O(\ln T \cdot \ln \ln T)$. As a result, our work demonstrates that the dynamics of such voting process depend heavily on the geometry of the underlying metric space. We also provide analytic tools to study problems related to the spatial theory of voting in higher dimensions.

Representative Papers:

[1] One Dollar Each Eliminates Envy (EC 2020) with J. Brustle, J. Dippel, V. Narayan, A. Vetta

[2] How Many Freemasons Are There? The Consensus Voting Mechanism in Metric Spaces (SAGT 2020) with A. Vetta

[3] Simultaneously Achieving Fairness and Efficiency via Transfer Payments (Preprint) with V. Narayan, A. Vetta

YIFENG TENG (Homepage, CV)

Thesis: The Power of Simple Pricings in Revenue and Welfare Maximization

Advisor: Shuchi Chawla, University of Wisconsin-Madison

Brief Biography: Yifeng Teng is a PhD candidate in Department of Computer Sciences at University of Wisconsin-Madison, advised by Prof. Shuchi Chawla. He is broadly interested in topics in theoretical computer science, especially algorithmic mechanism design and online algorithms. During his doctorate studies, he spent two summers as an intern with sell-side ads team and market algorithms team in Google Research New York. Prior to graduate studies at UW-Madison, he obtained a bachelor's degree from Institute of Interdisciplinary Information Sciences (Yao Class) at Tsinghua University, advised by Prof. Pingzhong Tang.

Research Summary:

Revenue maximization of general multi-item mechanisms has been long-time considered as a hard problem. In particular, it is well known that the revenue gap between optimal randomized mechanisms and simple mechanisms can be even unbounded. In work [1] [2], we advocate studying a new revenue benchmark, namely the optimal revenue achievable by "buy-many" mechanisms, that is much better behaved. In a buy-many mechanism, the buyer is allowed to interact with the mechanism multiple times instead of just once. We show that the optimal buy-many revenue for any n item setting is at most $O(\log n)$ times the revenue achievable by item pricing. Furthermore, a finite menu-size mechanism suffices to achieve $(1-\epsilon)$ fraction of optimal buy-many revenue. All such results hold for the most general setting possible, for any correlated distribution over arbitrary valuation function.

Another part of my research focuses on online welfare maximization. In work [3], we study several online resource allocation problems motivated by the online cloud market. We obtain asymptotically tight approximations to social welfare via posted pricing mechanisms, even compared to arbitrary online algorithms that do not need to satisfy incentive constraints. We also get a nearly tight supply vs. competitive ratio tradeoff when the seller has multiple copies of each item.

Finally, I have also worked on a broader class of online algorithms.In work [4], we study the Pandora's Box problem, which captures optimization problems with stochastic input where the algorithm can obtain instantiations of input variables at some cost. While all previous work on this class of problems assumes that different random variables in the input are distributed independently, we provide the first approximation algorithms for Pandora's Box-type problems with correlations.

Representative Papers:

[1] Buy-Many Mechanisms Are Not Much Better Than Item Pricing (EC'19) with S. Chawla and C. Tzamos

[2] Menu-size Complexity and Revenue Continuity of Buy-many Mechanisms (EC'20) with S. Chawla and C. Tzamos

[3] Pricing for Online Resource Allocation: Intervals and Paths (SODA'19) with S. Chawla and B. Miller

[4] Pandora's Box with Correlations: Learning and Approximation (FOCS'20) with S. Chawla, E. Gergatsouli, C. Tzamos and R. Zhang

ELLEN VITERCIK (Homepage, CV, Google Scholar)

Thesis: Automated Algorithm and Mechanism Configuration

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

Brief Biography: Ellen Vitercik is a computer science PhD student at Carnegie Mellon University (CMU). Her research interests are artificial intelligence, machine learning theory, algorithm design, and mechanism design. Among other honors, Ellen received the IBM PhD Fellowship, the NSF Graduate Research Fellowship, and a fellowship from CMU's Center for Machine Learning and Health. At EC 2019, she won the Exemplary Artificial Intelligence Track Paper Award and the Best Presentation by a Student or Postdoctoral Researcher Award. In 2015, she graduated *Junior Phi Beta Kappa* and *Summa Cum Laude* from Columbia.

Research Summary: Algorithms often have tunable parameters that significantly impact runtime and solution quality. Domain experts typically hand-tune these parameters, a highly time-consuming and error-prone task. I study how to use machine learning and data from a particular application domain to configure algorithm parameters. Algorithm configuration has enormous real-world potential and has been studied for over two decades, but almost all existing research is purely experimental; there is surprisingly little known from a theoretical perspective. In light of this, I study how to select algorithm configurations with provably high performance on problems from the application at hand. My research has provided the first learning-theoretic guarantees for algorithm configuration in diverse domains, including integer programming [1], clustering, and computational biology.

My research on mechanism design is integral to this broader agenda. A mechanism is a special type of algorithm which helps sets of agents come to collective decisions. Mechanisms often have tunable parameters—such as reserves—that impact revenue and social welfare. In an EC'18 paper [2], we study how to learn mechanism parameter settings with high revenue in expectation over the distribution defining the buyers' values, given only samples from the distribution. We provide a general theory that applies to a diverse array of multi-item pricing, auction, and lottery mechanisms. We prove new guarantees for mechanism classes not previously studied in the literature on sample-based mechanism design, and otherwise match or improve over the best-known guarantees for many classes.

In an EC'19 paper [3] which won the *Exemplary Artificial Intelligence Track Paper Award* and the *Best Presentation by a Student or Postdoctoral Researcher Award*, we study a related statistical question about mechanism design. Given a manipulable mechanism, can we use samples from the distribution over buyers' values to estimate how much the buyers will be willing to lie about their values? Though the mechanism is not incentive compatible, we show how to estimate the extent to which the mechanism is *approximately* incentive compatible.

Representative Papers:

[1] Learning to Branch (ICML 2018) with M.-F. Balcan, T. Dick, and T. Sandholm

[2] A General Theory of Sample Complexity for Multi-Item Profit Maximization (EC 2018) with M.-F. Balcan and T. Sandholm

[3] Estimating Approximate Incentive Compatibility (EC 2019) with M.-F. Balcan and T. Sandholm

MINGFEI ZHAO (Homepage, CV, Google Scholar)

Thesis: Simple Auctions in Bayesian Mechanism Design

Advisor: Yang Cai, Yale University

Brief Biography: Mingfei is a Ph.D. candidate in Department of Computer Science at Yale University, where he is advised by Yang Cai. Before coming to Yale, he received a M.Sc. in Computer Science and spent two years as a Ph.D. student at McGill University, advised by Yang Cai. Before that, he was an undergraduate student in Institute for Interdisciplinary Information Sciences (Yao Class) at Tsinghua University. He did an internship at Google Mountain View with Saeed Alaei in 2018. He is a recipient of 2017 Richard H. Tomlinson Doctoral Fellowship. His research interests broadly lie in mechanism design and algorithmic game theory, and mainly aim to design simple mechanisms in auctions and two-sided markets.

Research Summary:

Designing simple and approximately revenue-optimal mechanisms in multi dimensional combinatorial auctions has been extensively studied in recent years, as the optimal mechanisms in such settings are usually complex and randomized. In [1], we unify and improve all previous results, as well as generalize the results to broader cases. In particular, we provide two classes of simple mechanisms, where the best achieves a constant factor of the optimal revenue for multiple buyers with XOS valuations. And the approximation ratio is a logarithm in the number of items when the buyers are subadditive. In [2], we study a generalization of the above setting, where the seller may have a randomized production cost on items and aims to design truthful auctions that maximize her profit, i.e. revenue minus cost. We introduce a new class of mechanisms called permit-selling, and prove a constant factor approximation of the optimal profit for multiple matroid-rank buyers.

Another line of my work focuses on designing truthful and budget-balanced mechanisms that approximately maximizes the Gains from Trade (GFT, expected gain in the social welfare induced by the mechanism) in two-sided markets. This line of research is motivated by stock exchange and online ad exchange platforms. While the impossibility result by Myerson and Satterthwaite shows that no mechanisms can achieve full GFT and the optimal mechanism is rather complex even in bilateral trading (1 buyer, 1 seller, 1 item), we provide simple mechanisms that achieve half of the optimal GFT for any single-dimensional two-sided market [3]. In a recent work [4], we present the first worst-case approximation guarantee for GFT in multi-dimensional two-sided markets.

Representative Papers:

[1] Simple Mechanisms for Subadditive Buyers via Duality (STOC 2017, SIGecom Exchange Apr. 2019) with Y. Cai

[2] Simple Mechanisms for Profit Maximization in Multi-item Auctions (EC 2019) with Y. Cai

[3] Approximating Gains from Trade in Two-sided Markets via Simple Mechanism (EC 2017, preparing for OR) with J. Brustle, Y. Cai and F. Wu

[4] On Multi-Dimensional Gains from Trade Maximization (SODA 2021) with Y. Cai, K. Goldner, and S. Ma

JUBA ZIANI (Homepage, CV, Google Scholar)

Thesis: Data: Implications for Markets and for Society

Advisors: Katrina Ligett, Hebrew University; Adam Wierman, Caltech.

Brief Biography: I am a Warren Center Postdoctoral Fellow at the University of Pennsylvania, hosted by Sampath Kannan, Michael Kearns, Aaron Roth, and Rakesh Vohra. Prior to this, I was a PhD student in the Computing and Mathematical Sciences Department at Caltech, where I was advised by Katrina Ligett and Adam Wierman and where I received the Bhansali Family Doctoral prize for best dissertation in Computer Science. During my PhD, I visited and then interned at Microsoft Research New England in 2017, visited UPenn in Summer 2018, and participated in the Privacy semester at Simons in Spring 2019.

Research Summary: I study the challenges that arise from transactions and interactions involving data, often in the presence of strategic behavior. In particular, my research focuses on data markets, differential privacy, fairness in machine learning, and strategic considerations in machine learning.

One major focus of my work is building theoretical foundations for data markets. In [1], we study the optimization and mechanism design problems faced by an analyst who must acquire and aggregate costly data from strategic agents, under a budget constraint. We consider agents whose costs are correlated with their data, as is often the case with private or sensitive information.

Promoting access to data of course introduces societal concerns. The first major concern I study is data privacy, for which I use techniques from differential privacy to protect the privacy of individuals whose sensitive data is used in computations—see my website for more information. The second central concern that I aim to address is to ensure that data about agents is not used to discriminate against them. One of my main contributions to the fairness literature is [2]. There, we study a college education and job market pipeline in which a university can use affirmative action as a way to influence the Bayesian posteriors of potential employers down the pipeline, rather than simply as a way to correct for past disparities; we characterize which statistical notions of fairness are achievable in this setting.

Additionally, I am interested in how strategic considerations affect and inform agents' data and in how data influences agents' behavior. In [3], we examine how the presence of a strategic third-party data provider affects mechanism design in settings with information asymmetries, such as ad auctions. We show that traditional simple mechanisms cannot achieve good revenue guarantees, which contrasts with a recent line of work that advocates for simple designs in incomplete information settings. In more recent work, I have been examining how a learner should interact with agents who strategically alter their data in machine learning settings.

Representative Papers:

[1] Optimal Data Acquisition for Statistical Estimation (EC'18)
with Y. Chen, N. Immorlica, B. Lucier, and V. Syrgkanis

[2] Downstream Effects of Affirmative Action (FAT* 2019)
with S. Kannan, and A. Roth

[3] Third-party Data Providers Ruin Simple Mechanisms (SIGMETRICS 2020)
with Y. Cai, F. Echenique, H. Fu, K. Ligett, A. Wierman

Index