SIGecom Job Market Candidate Profiles 2019

Edited by VASILIS GKATZELIS and JASON HARTLINE

This is the fourth annual collection of profiles of the junior faculty job market candidates of the SIGecom community. The twenty five candidates for 2019 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-candidates2019@acm.org.

–Vasilis Gkatzelis and Jason Hartline

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

Contents

HEDYEH BEYHAGHI

Thesis. Matching Markets with Short Lists

Advisor. Éva Tardos, Cornell University

Brief Biography. Hedyeh Beyhaghi is a Ph.D. candidate at the Computer Science department of Cornell University advised by Éva Tardos. Her research lies in the intersection of computational aspects of economics and algorithms, with a focus on algorithmic mechanism design, simple mechanisms, and matching markets. Hedyeh was a long-term visitor at the Simons Institute, UC Berkeley for a semester on Economics and Computation in Fall 2015. She was an intern at Google-NYC in summer 2017, and a visiting student researcher at Princeton, hosted by Matt Weinberg in 2017-18.

Research Summary. I am broadly interested in algorithmic game theory and mechanism design. My research focuses on designing mechanisms for various markets and evaluating market outcomes with the goal of improving their efficiency. In auction design, for example, I focus on simple auctions with approximately optimal revenue. In matching markets, I analyze the loss of efficiency due to the market constraints in order to determine the conditions under which markets are more efficient. In cryptocurrencies, I analyze the effect of incentive issues that threaten the system and propose modifications to recover efficiency. Below I give two examples of my work in auction design and matching markets.

The problem of selling a single item to one of multiple independent but non-identical bidders occurs frequently when selling products via online auctions, and also in ad exchanges, where web publishers sell the ad slots on their webpages to advertisers. The widely used mechanisms for this setting are posted prices and the second price auction with reserves. We improve the best-known performance guarantees for these mechanisms. Motivated by our improved revenue bounds, we further study the problem of optimizing reserve prices in the second price auctions when the sorted order of personalized reserve prices among bidders is exogenous. We show that this problem can be solved polynomially. In addition, by analyzing a real auction dataset from Google's advertising exchange, we demonstrate the effectiveness of order-based pricing. These results also contribute to prophet inequalities and online algorithms.

In matching markets, sub-optimal outcomes can emerge as a result of constraints on the number of applications and interviews. We study the effects of such constraints in centralized matching markets, e.g. the National Residency Matching Program. We show when the number of interviews conducted by hospitals is limited, a market with a limit on the number of applications by residents achieves higher efficiency compared to a market with no limits; whereas with no limit on the number of interviews, matching size is increasing in the number of applications. Also, we find that a system of treating all applicants equally (where everyone has the same limit on the number of applications) is more efficient than allowing a small set to apply to one more/less position. By taking these non-intuitive phenomena into account, we are able to improve the market efficiency in spite of inevitable constraints.

Representative Papers.

[1] Improved Approximations for Free-Order Prophets and Second-Price Auctions (arXiv) with N. Golrezaei, R. Paes Leme, M. Pál, and B. Sivan

[2] Effect of Limited Number of Interviews on Matching Markets with É. Tardos

[3] Benchmark-Optimal Competition Complexity for Additive Buyers with S.M. Weinberg

GIANLUCA BRERO

Thesis. Machine Learning-powered Combinatorial Auctions

Advisor. Sven Seuken, University of Zurich

Brief Biography. Gianluca is Ph.D. candidate at the Department of Informatics of the University of Zurich advised by Sven Seuken. His research lies at the intersection of machine learning and market design, with a focus on designing market mechanisms that use machine learning internally to support market participants in complex settings. In summer 2016, he was a research intern at Microsoft Research-NYC. Gianluca holds a Bachelor Degree in Applied Mathematics from Turin Polytechnic and a double Master Degree in Mathematical Engineering from Turin Polytechnic and Milan Polytechnic. At Turin Polytechnic, he was awarded the Vallauri Prize as the top student in Mathematical Engineering graduated in 2014.

Research Summary. Market designers study how to set the "rules of a marketplace" such that market participants can coordinate effectively and voluntarily to achieve a good aggregate outcome. In complex domains, successful market designs should prevent participants from incurring prohibitively-high computational costs to coordinate on these outcomes. For example, in combinatorial domains, participants may have a specific value for any bundle of items they can obtain. As the number of items grows, it may quickly become unfeasible to understand how to allocate items efficiently.

In my research, I design combinatorial auctions (CAs) that use machine learning (ML) internally to deal with highly-complex valuations. The ML algorithms are meant to assist bidders in exploring their valuations and limit the amount of information they report while maintaining high allocative efficiency and incentivizing truthful reporting. These *ML-powered mechanisms* can be grouped into two main categories depending on their type of interaction with bidders.

The mechanisms of the first category interact with bidders via *value queries* where bidders report their values for obtaining bundles presented by the mechanism. In a AAAI'17 paper we introduced an elicitation paradigm where ML algorithms are trained on set of bundle-value pairs reported by each bidder to predict values bidders may have for each possible bundle. These predictions are used to determine a candidate efficient allocation maximizing the predicted total value to bidders. This allocation is then investigated via additional value queries. In a later IJCAI'18 paper, we showed how this elicitation paradigm can be integrated into incentive-aligned CA mechanisms that achieve high allocative efficiency in realistic spectrum auction settings with 98 items.

The mechanisms of the second category interact with bidders via *demand queries* where bidders report their favorite bundles at prices proposed by the mechanism. In a AAAI'18 paper we developed a probability model that allowed us to exploit data on bidder valuations to make inference over clearing prices. This model formed the basis of an auction process which alternates between refining estimates of bidder valuations and computing candidate clearing prices. Our auction proved to be highly competitive against standard auction schemes in terms of round of convergence, even when these schemes are optimally tuned to each valuation instance.

Representative Papers.

[1] Probably Approximately Efficient Combinatorial Auctions via Machine Learning (AAAI 2017) with B. Lubin and S. Seuken

[2] Combinatorial Auctions via Machine Learning-based Preference Elicitation (IJCAI 2018) with B. Lubin and S. Seuken

[3] A Bayesian Clearing Mechanism for Combinatorial Auctions (AAAI 2018) with S. Lahaie

NOAM BROWN

Thesis. Equilibrium Finding in Large Imperfect-Information Games

Advisor. Tuomas Sandholm, Carnegie Mellon University

Brief Biography. Noam Brown is a PhD student in computer science at Carnegie Mellon University advised by Tuomas Sandholm. His research combines reinforcement learning and game theory to develop AI systems capable of strategic reasoning in imperfect-information multi-agent interactions. He has applied this research to creating Libratus, the first AI to defeat top humans in no-limit poker. The breakthrough was published in Science and was a finalist for Science Magazine's 2017 Scientific Breakthrough of the Year. Noam received a NIPS Best Paper award in 2017 and an Allen Newell Award for Research Excellence. He is supported by an Open Philanthropy Project AI Fellowship and Tencent AI Lab Fellowship. Prior to starting a PhD, Noam worked at the Federal Reserve researching the effects of algorithm trading on financial markets. Before that, he developed algorithmic trading strategies commercially. Noam received a Bachelor's from Rutgers University in 2009.

Research Summary. From Deep Blue to AlphaGo, AI has made tremendous progress in the past few decades in the domain of perfect-information games. However, there has been historically little progress in dealing with hidden information in strategic settings because many of the implicit assumptions needed for perfect-information techniques to work no longer apply. For example, in imperfect-information games a subgame cannot be solved in isolation, and states may not have unique values in an equilibrium.

My research focuses on developing robust AIs capable of handling hidden information. I successfully applied this research toward developing *Libratus*, which in early 2017 became the first AI to defeat top humans in no-limit Texas hold'em poker, the primary benchmark for research on imperfect-information games and a key challenge problem for AI in general. While I have evaluated my techniques on variants of poker, none of my approaches are specific to poker. My research is now focused on extending these techniques to more complex real-world domains.

Recently, a number of researchers have attempted to apply deep learning to multi-agent interactions. For example, Facebook has set a goal of developing an AI capable of negotiation. However, the approaches typically taken do not have theoretical guarantees and are potentially exploitable by an opponent taking unusual actions. Being provably robust to adversarial attacks has always been a primary concern among imperfect-information AI researchers. Indeed, performance of AIs in the field are typically evaluated in terms of *exploitability*, which is how much more an AI would lose to a worst-case adversary relative to a perfect equilibrium AI. It is this focus on robustness that allowed *Libratus* to defeat top humans not just over a small sample, but over 120,000 hands of poker, even though the human opponents worked as a team to constantly probe the AI for any weakness.

This line of research is critical to safely bringing AI into the real world because most real-world strategic interactions involve some amount of hidden information, such as auctions, negotiations, cybersecurity, physical security, financial markets, business decisions, and military situations. For this reason, I believe equilibrium finding is the appropriate paradigm to approach strategic interactions. If a system is to be deployed at a global scale in the real world, it is essential that the system be robust to attacks and exploitative behavior.

Representative Papers.

[1] Superhuman AI for heads-up no-limit poker: Libratus beats top professionals (Science 2017) with T. Sandholm

[2] Safe and Nested Subgame Solving for Imperfect-Information Games (NIPS 2017) with T. Sandholm

[3] Strategy-Based Warm Starting for Regret Minimization in Games (AAAI 2016) with T. Sandholm

MITHUN CHAKRABORTY

Thesis. On The Aggregation Of Subjective Inputs From Multiple Sources

Advisor. Sanmay Das, Washington University in St. Louis

Brief Biography. I am currently a post-doctoral research fellow at the Department of Computer Science, National University of Singapore, working with Yair Zick. I obtained my PhD degree in Computer Science in May 2017 from the Department of Computer Science and Engineering at Washington University in St. Louis (WUSTL), where I was advised by Sanmay Das. I received a WU-CIRTL Practitioner Level recognition from the Teaching Center at WUSTL in May 2016, and the Department Chair Award for Outstanding Teaching in April 2016 for teaching Social Network Analysis (CSE316A) in Fall 2015. I was invited to give a talk on my research at the 20th Conference of the International Federation of Operational

Research Societies (IFORS) in July 2014. I earned my Bachelor's degree (with Honours) in Electronics and Telecommunication Engineering from Jadavpur University, Kolkata, India, in 2009.

Research Summary. My research interests span algorithmic economics, multi-agent systems, and machine learning.

A major focus of my work has been the aggregation of multiple subjective inputs, with applications to the design and analysis of prediction markets and more general market mechanisms. We established a formal connection between scoring rule-based prediction markets and opinion pools, a family of belief combination rules, under risk-averse trader models [2]. An IJCAI'16 paper proposes a new game-theoretic model of prediction market manipulation when participants can influence the forecast outcome, showing how the properties (e.g. truthfulness) of the Nash equilibrium depend on the probability of some agents being non-strategic. In an AAAI'15 paper, we proposed an extension of the logarithmic market scoring rule to a more practical setting (where traders specify both order sizes and acceptable prices), and empirically uncovered interesting properties of the resulting market ecosystem. I was also part of a team that conducted human-subject experiments to evaluate two different market making approaches ([3], AAAI'13). Some of my (joint) work on issues in information acquisition and aggregation with no regard to truth-telling incentives include a no-regret learning approach in multi-agent multi-armed bandits with costly communication (IJCAI'17) and an approximately Bayesian active learning algorithm for converging on a real-valued target from censored noisy signals (UAI'11).

More recently, I have become interested in the impact of diversity requirements on allocation mechanisms. A recent paper [1] looks at public housing allocation in Singapore where, to prevent de facto segregation, an upper bound is imposed on the number of flats in each estate block that can be purchased by applicants of each ethnic group, and analyzes both theoretically and empirically how such diversity constraints affect economic efficiency (welfare loss). I am currently working on more general related questions such as how to define diversity goals for various problems in computational social science (matching, committee selection, etc.) and how to extend concepts from the fair division literature (e.g. envy-freeness) to solution approaches in these domains.

My long-term plan is to not only pursue these strands individually but also explore how they can be tied together – in particular, how social desiderata (e.g. diversity and fairness) interact with individual economic incentives.

Representative Papers.

[1] Diversity Constraints in Public Housing Allocation (AAMAS 2018) with N. Benabbou, X. Ho, J. Sliwinski, and Y. Zick

[2] Market Scoring Rules Act As Opinion Pools For Risk-Averse Agents (NIPS 2015) with S. Das

[3] A Bayesian Market Maker (EC 2012) with A. Brahma, S. Das, A. Lavoie, and M. Magdon-Ismail

YU CHENG

Thesis. Computational Aspects of Optimal Information Revelation

Advisor. Shang-Hua Teng, University of Southern California

Brief Biography. Yu Cheng is currently a Postdoctoral Researcher in the Computer Science Department at Duke University, hosted by Vincent Conitzer, Rong Ge, Debmalya Panigrahi, and Kamesh Munagala. He obtained his Ph.D. in Computer Science from University of Southern California in 2017, advised by Shang-Hua Teng. He is interested broadly in theoretical computer science, with a focus on problems at the intersection of learning theory and game theory. These days, Yu is particularly interested in designing provably robust and scalable machine learning algorithms.

Research Summary. My research interests lie in the area of theoretical computer science, with a focus on learning theory and game theory. My research advances both algorithmic and complexity-theoretic understanding of many fundamental problems. Recently, I have been focusing on theoretical aspects of the challenges at different stages of building a machine learning system. I will give three examples below: the first example is on collecting data from strategic agents, the second example is on designing scalable algorithms that are provably robust against data contamination, and the third example is on how should machine learning systems reveal information to selfish agents.

In a recent paper, "When Samples Are Strategically Selected" with Vincent Conitzer and Hanrui Zhang, we study the classification problem when a machine learning system does not have direct access to the samples. Instead, it must rely on another party (a strategic agent) to provide samples. This agent may have different incentives, so we cannot trust that the samples are representative. The goal is to design a good classifier based on the provided samples, keeping in mind the agent's incentives. We introduce a theoretical framework for this problem and provide key structural and computational results.

In a COLT18 paper "Non-Convex Matrix Completion Against a Semi-Random Adversary" with Rong Ge, we study the matrix completion problem, where the goal is to recover a low-rank matrix after observing a random set of its entries. We consider a semi-random setting where an adversary may reveal additional entries to manipulate the output of the algorithm. We show that in this setting, existing non-convex algorithms may get stuck in bad local optima. We propose a nearly-linear time algorithm to fix this: after our preprocessing, the non-convex objective function has no bad local optima.

In an EC16 paper "Hardness Results for Signaling in Bayesian Zero-Sum and Network Routing Games" with Umang Bhaskar, Young Kun Ko, and Chaitanya Swamy, we study the problem of how to reveal information optimally. What is the best way for a machine learning system (e.g., Google Maps) to reveal information to selfish agents to induce a desirable outcome (e.g., reducing latency experienced by the drivers)? For Bayesian network routing games, we show that sometimes revealing partial information can be better than full revelation, and the problem of revealing information optimally is computationally hard in the worst-case.

Representative Papers.

[1] Non-Convex Matrix Completion Against a Semi-Random Adversary (COLT 2018) with R. Ge
[2] Hardness Results for Signaling in Bayesian Zero-Sum and Network Routing Games (EC 2016) with U. Bhaskar, Y. Ko, and C. Swamy
[3] Mixture Selection, Mechanism Design, and Signaling (FOCS 2015) with H.Y. Cheung, S. Dughmi, E. Emamjomeh-Zadeh, L. Han, and S.-H. Teng

ALON EDEN

Thesis. Harnessing Prices for Market Efficiency

Advisors. Michal Feldman and Amos Fiat, Tel-Aviv University

Brief Biography. Alon Eden is a PhD student in the Computer Science Department at Tel Aviv University, under the supervision of Michal Feldman and Amos Fiat, where he also obtained his MSc and BSc. He is interested in topics that combine economics and algorithms. He was a long term visitor during the Algorithms & Uncertainty semester at the Simons Institute during the Fall of 2016. His work also received the Best Paper Award at SAGT'17.

Research Summary. My main research interests are Mechanism Design, Online Algorithms and the combination of the two. I am generally interested in algorithmic problems with economic motivation.

During my PhD, I have studied pricing mechanisms for settings where agents arrive in an online fashion. The main advantage in using pricing mechanisms is that they do not require all agents to be present at the same time, and thus they can be implemented in a distributed manner. Furthermore, truthfulness is often immediate, as the agents are offered prices over different outcomes, and are not asked about their type. In [2], we show that prices are useful in market coordination, and can achieve optimal welfare in many scenarios, including in the presence of matching markets and gross substitutes valuations. In [1], we explore the power of dynamic pricing in online settings, where selfish agents seek to minimize the cost incurred to them by some service. We show that by using prices, one can get a near optimal competitive ratio for a variety of online problems, including metrical task systems, the k-server problem, and metric matching. In subsequent works, we explore this framework for scheduling problems.

In addition, I have studied multidimensional revenue maximization. We show that simple and widely used mechanisms can approximately maximize revenue in settings where an agent's valuation may exhibit complementarities. In a different paper, we extend the seminal result of Bulow and Klemperer to multidimensional settings. We show that given n additive bidders with values sampled from an identical product distribution over m items, adding $O(n + m)$ additional bidders and running the prior free VCG mechanism extracts more revenue than the optimal mechanism—an object which is not well understood, and computationally intractable. Lately, I have been intrigued by interdependent values settings, where agents' values are affected by the same parameters and can be arbitrarily correlated. Usually in these settings, a single-crossing condition is assumed on the valuations. In [3] we study what can be achieved without this assumption, and show that the best welfare we can hope for from selling a single item can be achieved by allocating to a random bidder. This motivates the approach of parameterizing how close the valuations are to being single-crossing. Using this parameterization, we manage to obtain positive results both for welfare and revenue maximization.

Representative Papers.

[1] Pricing Online Decisions: Beyond Auctions (SODA 2015)
with I.R. Cohen, A. Fiat, and L. Jez

[2] The Invisible Hand of Dynamic Market Pricing (EC 2016)
with M. Feldman, A. Fiat, and V. Cohen-Addad

[3] Interdependent Values without Single-Crossing (EC 2018)
with M. Feldman, A. Fiat and K. Goldner

KIRA GOLDNER

Thesis. Interdimensional Mechanism Design

Advisor. Anna Karlin, University of Washington

Brief Biography. Kira is a fifth-year PhD student at the University of Washington in Computer Science and Engineering, where she is advised by Anna Karlin. During her PhD, she co-organized the first two workshops on Mechanism Design for Social Good at EC 2017 and 2018, and the reading group by the same name. She has given tutorials on Mechanism Design for Social Good at WINE 2017 and on Menu Complexity at EC 2018. Kira spent two summers interning at Microsoft Research: with Nikhil Devanur in Redmond in 2017 and with Brendan Lucier in New England in 2018. She is a 2017-19 recipient of the Microsoft Research PhD Fellowship and was a 2016 recipient of the Google Anita Borg Scholarship. She graduated from Oberlin College in 2014 with a B.A. in mathematics.

Research Summary. My primary interest is in mechanism design. My work has focused on optimizing revenue in various settings, relaxing behavioral or informational assumptions, and tailoring mechanism design to specific social good scenarios.

Interdimensional Mechanism Design: Previous work in revenue maximization depicts a rigid dichotomy between two types of settings. On one hand, we have single-dimensional settings (e.g. single-item auctions), which are completely characterized by a simple closed-form theory; on the other, we have "multi-dimensional" settings (e.g. unit-demand or additive buyers), where optimal mechanisms are chaotic and often intractable. In "The FedEx Problem," we identify a space that is fundamentally different from both extremes, lying on the spectrum in between. We characterize the optimal mechanism for the "FedEx" setting, showing that it is simpler than optimal mechanisms for even the simplest multi-dimensional setting, yet far more complex than any single-dimensional setting. In subsequent work, my co-authors and I have shown that this class of problems contains many natural settings and continues to lie fundamentally in between single- and multi-dimensional with respect to menu-complexity, sample-complexity, and other metrics.

Behavioral and Informational Assumptions: My research also addresses mechanism design under more realistic assumptions about behavior or information. I have studied settings where (1) buyers are "uncertainty averse," i.e. will pay a premium for outcomes that occur with higher certainty; (2) buyers are "interdependent," i.e. have valuations that depend on the private information of others; (3) buyers have "proportional complementarities," i.e. have known correlations for the buyer's boost from getting multiple items; and (4) designers do not have information about the prior distributions of values or risk attitudes of buyers.

Mechanism Design for Social Good: Much of my ongoing work focuses on this theme. Often this involves addressing core theoretical mechanism design problems, but with differing objectives. For example: optimizing residual surplus (a likely objective for a single healthcare payer) and optimizing gains from trade (i.e. taking the perspective of a benevolent online platform). I am also interested in more domain-specific and applied problems, such as designing mechanisms for employers to contract with health insurance plans, and designing mechanisms to trade off carbon pollution with economic growth.

Representative Papers.

[1] The FedEx Problem (EC 2016)
with A. Fiat, A.R. Karlin, and E. Koutsoupias

[2] Revenue Maximization with an Uncertainty-Averse Buyer (SODA 2018)
with S. Chawla, J.B. Miller, and E. Pountourakis

[3] Interdependent Values without Single-Crossing (EC 2018)
with A. Eden, M. Feldman, and A. Fiat

YANNAI A. GONCZAROWSKI

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 PhD student in Math, CS, and Rationality at the Hebrew University of Jerusalem, where his advisors are Noam Nisan and Sergiu Hart. He holds an M.Sc. in Math (summa cum laude) and a B.Sc. in Math and CS (summa cum laude, Class Valedictorian) from the same institution. Yannai is also a professionally trained opera singer, holding a Master's as well as a Bachelor's degree in Classical Singing. Since the second year of his PhD, he has been a research intern (part-time) at Microsoft Research in Herzliya. Yannai started TAing while an undergrad, and has appeared repeatedly on the Outstanding Teachers list; he started lecturing on the first year of his PhD, and was ranked first among all CS lecturers and TAs that year. Yannai's PhD studies are supported by an Adams Fellowship of the Israel Academy of Sciences and Humanities. For his PhD thesis, he was awarded the 2018 Michael B. Maschler Prize of the Israeli Chapter of the Game Theory Society. His Erdős–Bacon Number is 5.

Research Summary. The recent explosion in online and computerized economic activity necessitates the study of economic mechanisms of unprecedented scale. For a wide range of mechanisms, from auctions to school-choice mechanisms, we are far from a satisfying answer to many fundamental questions: How good can simple mechanisms be? How complex must optimal mechanisms be? And, most substantially, what are the precise trade-offs between simplicity and quality? I am interested in such questions for various notions of complexity and simplicity.

In a STOC'17 paper and a STOC'18 paper, we study the complexity of describing auctions via the lens of their *menu-size*. As the revenue-optimal auction may have infinite menu-size (Daskalakis et al., 2013) even in simple settings, we study approximately optimal auctions and obtain asymptotic bounds on their menu-size. As we show, these results also have implications regarding the *communication complexity* of running such auctions. In a SODA'15 paper, we analyze the communication complexity of stable matching mechanisms.

Classic economic analysis traditionally assumes that the mechanism designer has access to a distribution from which participants' types are drawn. The *sample complexity* line of research initiated by Cole and Roughgarden (2014) at the interface of algorithmic game theory and learning theory relaxes this assumption to the more realistic one of having access to polynomially many samples from this fixed, yet unknown, type distribution; the goal is to learn from the samples an up-to-ε optimal mechanism for the unknown underlying distribution. In a STOC'17 paper and a FOCS'18 paper, we obtain computationally efficient results for single-parameter auctions, and information-theoretic results for multi-parameter auctions.

With the emergence of more and more online marketplaces, the profile of participants is transitioning from experts to laymen, necessitating mechanisms that are easy to understand and to participate in. In a JET'18 paper and an EC'17 paper, we characterize the set of *obviously strategyproof* (Li, 2018) mechanisms in popular settings. In an EC'14 paper, under certain scenarios I analyze what can be described as the complexity of manipulating stable matching mechanisms.

On the practical side, we redesigned the Israeli Pre-Military Academies Match. Outside mechanism design, my research includes analyzing certain congestion games via an analogy to a physical hydraulic system, and using epistemic logic to design optimal distributed protocols for canonical problems such as *consensus*.

Representative Papers.

[1] Bounding the Menu-Size of Approximately Optimal Auctions via Optimal-Transport Duality (STOC 2017)

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

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

RAGAVENDRAN GOPALAKRISHNAN

Thesis. Characterizing Distribution Rules for Cost Sharing Games

Advisor. Adam Wierman, California Institute of Technology

Brief Biography. Raga Gopalakrishnan is a postdoctoral associate at Cornell University, jointly hosted by Siddhartha Banerjee and Samitha Samaranayake. His Ph.D. is in Computer Science from Caltech, following which he held a postdoctoral position at the University of Colorado Boulder and Research Scientist/Manager roles at Xerox/Conduent (Transportation Services). His work concerns the design of service systems, especially the interplay between operational policies (e.g., routing, matching, staffing) and economic mechanisms (e.g., pricing, cost/revenue sharing, incentives), with a growing focus on applications to urban transportation. Through his industrial experience, he gained a hands-on perspective on leveraging theoretical foundations learned in academia to design effective real-world solutions, leading to an evolving emphasis on a practice-aware approach to operations research. His interdisciplinary work has appeared in diverse venues such as EC, AAMAS, IJCAI, CDC, Oper. Res., and Math. Oper. Res.

Research Summary. Urban mobility is undergoing a global transformation into a user-centric, service-oriented playing field, from a traditionally vehicle-centric framework. Technology advances that enable Mobility-as-a-Service and the sharing economy mold the underlying problem and design spaces along new dimensions. For example, single-payment multi-modal travel (involving more than one provider) is possible today, which not only introduces new real-time pricing and capacity/revenue sharing problems, but also connects economic objectives, e.g., social welfare and revenue, to operational measures, e.g., trip 'quality', in nontrivial ways. Consequently, I adapt traditional modelling frameworks/methodologies and analyze them to identify key insights to guide policy/decision-making in large, real-world systems.

My work on 'Service Systems with Strategic Servers' with Doroudi, Ward, and Wierman models and analyzes service systems when servers have discretion over service times, e.g., peer-reviewing, call centers, and ridehailing services, where optimal system policies turn out to be very different. In subsequent independent work, I consider the system configuration (one central queue vs. multiple parallel queues), expand the class of routing policies, and study the trade-off between system performance and servers' utilities.

My recent work on 'Sequential Individual Rationality and Fairness' with Mukherjee and Tulabandhula began at Xerox/Conduent from a project that I co-led, involving a multi-modal trip planner for a French transit network and piloting a local prototype in Bangalore. The focus is on QoS-aware pricing (by a commercial provider) or fair cost sharing (among peer-to-peer users) in shared service systems. Our solution, a dynamic utilitarian approach to model QoS and its impact, became part of the core cost sharing method of the ridesharing subsystem, and, together with my other sustainability-focused research on electric vehicles, earned me a *Scientific Excellence Award*.

My ongoing work with Hssaine, Banerjee, and Samaranayake aims to lay the groundwork for new fundamental research on overcoming economic inefficiency by tapping into potential operational cooperation (e.g., multi-modal trips in the urban mobility context), and characterizing environments in which such benefits are worth the overhead. Concurrently, I am leading the development of a data-driven, city-scale mobility simulator that can support various market models and economic control policies.

Representative Papers.

[1] Designing a Sustainable Marketplace for Modern Urban Mobility (working paper) with C. Hssaine, S. Banerjee, and S. Samaranayake

[2] Routing and Staffing When Servers Are Strategic (EC 2014, *Oper. Res.* 2016) with S. Doroudi, A.R. Ward, and A. Wierman

[3] Potential Games Are *Necessary* to Ensure Pure Nash Equilibria in Cost Sharing Games (EC 2013, *Math. Oper. Res.* 2014) with J.R. Marden and A. Wierman

ANILESH K. KRISHNASWAMY

Thesis. Democratic platforms: Engineering consensus and collaboration

Advisor. Ashish Goel, Stanford University

Brief Biography. Anilesh Krishnaswamy is a PhD student at Stanford University, advised by Prof. Ashish Goel. His research uses tools from game theory, social choice and algorithms to study both the theoretical and practical aspects of online platforms that enable democratic decision-making, social computing, and labor markets. 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 to decide the fate of millions of dollars of public funds. He was also a visiting researcher at Tel-Aviv University during May-June 2018. Prior to graduate studies at Stanford, he obtained a bachelor's degree in Electrical Engineering from the Indian Institute of Technology Madras.

Research Summary. The rise of online platforms has revolutionized traditionally small-group endeavors, such as informed decision-making and collaborative scientific research, by allowing them to operate in a distributed manner at much larger scales. My research focusses on the design of appropriate algorithmic and incentive mechanisms to fully leverage this development. In addition, the use of algorithms to make societal decisions necessitates a study of the *fairness* properties of such algorithms. In what follows, I will give a few examples of my work along these lines.

Direct democracy has gained a lot of traction lately, with exciting new paradigms such as participatory budgeting (PB), in which a local government body asks residents to vote on project proposals to decide how to allocate budgetary spending. In my work with the Stanford Crowdsourced Democracy Team, we introduced "Knapsack Voting", a novel voting method for aggregating voter preferences in this setting, and showed its advantages over methods currently used in practice. We also showed how our scheme works well in practice by implementing it on the digital voting platform we have deployed in many cities across the nation.

In an EC'14 paper, we study the problem of engineering large-scale collaborative research. The essential tension in such systems is the tendency of agents to withhold the sharing of partial progress in order to gain recognition for a larger breakthrough. We built a model that captures the incentives at play in such settings wherein a large project is divided into many smaller tasks with dependencies. We showed how allocating rewards to each of these smaller tasks in a locally fair manner can incentivize the immediate sharing of partial progress.

The utilitarian approach to social choice, in particular the problem of understanding the worst-case distortion of social choice rules, has gained a lot of attention recently. In an EC'17 paper, we looked at the *metric distortion* problem, wherein the agents and alternatives are assumed to lie in an unknown metric space. In addition to settling some open questions, we introduced a method of quantifying the "fairness" of social choice functions in this setting. We showed how constant factor approximations to a wide class of fairness measures can be achieved with very simple algorithms. I am also excited about ongoing work on mechanisms for achieving fair outcomes in other multi-agent problems such as bargaining and cake-cutting.

Representative Papers.

[1] Knapsack Voting for Participatory Budgeting (under review) with A. Goel, S. Sakshuwong, and T. Aitamurto

[2] Re-incentivizing discovery: Mechanisms for partial-progress sharing in research (EC 2014) with S. Banerjee and A. Goel

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

THANASIS LIANEAS

Thesis. Congestion Games: Stochastic Extensions and Techniques for Reducing the Price of Anarchy

Advisor. Dimitris Fotakis, Technical University of Athens

Brief Biography. Thanasis Lianeas is currently a Lecturer at the School of Electrical and Computer Engineering (ECE) of the National Technical University of Athens (NTUA). Before going back to Greece this year, he worked three years as a Post-doc at the University of Texas at Austin under the supervision of Prof. Evdokia Nikolova. He got his PhD from the School of ECE of the NTUA advised by Prof. Dimitris Fotakis and completed his undergraduate studies at the School of Applied Mathematical and Physical Sciences of NTUA. His research interests mainly lie in the Design and Analysis of Algorithms and in Algorithmic Game Theory, and he has taught or provided teaching assistance for classes like Algorithms (undergraduate and graduate level), Approximation Algorithms, Discrete Mathematics, Algorithmic Game Theory, Introduction to Computer Science and Introduction to Programming.

Research Summary. My research interests lie in the Design and Analysis of Algorithms and in Algorithmic Game Theory. Regarding Algorithmic Game Theory, I still see interesting open problems that lie in the area I have expertize on, i.e., Congestion Games, but also across the borders with other areas like Mechanism Design and Stochastic Optimization. Towards the side of Algorithms, I am interested in continuing my (recently started) research on algorithmic problems that arise in electricity networks including optimization problems with time-evolving costs, versions of which arise in electricity networks but have also stronger motivation, thus being of independent interest.

My previous research in Algorithmic Game Theory has been focused on Congestion Games, which provide a natural model for non-cooperative resource allocation in large-scale networks, e.g. transportation and communication networks. In an abstract description, Congestion Games (CGs) are games where selfish users compete over resources and perceive a cost for the resources they use, a cost that depends on the competition/congestion on the resources and each resource's "nature". My work deals with standard questions like computing, quantifying and improving equilibria of games that lie inside the classic framework but also in more generalized settings where uncertainty comes into play, with my research investigating how uncertainty and risk-aversion transform traditional models and methodologies (see e.g. the "Math of OR" and the "IJCAI '18" papers), or the resources themselves are operated by self-interested players yielding a more complex yet interesting bi-level game (see e.g. the "EC '18" paper).

Recently, I have started engaging with algorithmic problems that arise in electricity networks, partially tempted by the similarities of the physical laws that govern these networks to the equilibrium conditions in CGs. Interestingly, apart from practical ones, problems of great theoretical interest arise there, that lie under the umbrella of optimization with time-evolving costs, and, independently, have recently received the attention of researchers. Our first result in the area concerns the so called reconfiguration problem and appeared in the prestigious HICSS '18 conference, while, in preliminary work, taking advantage of the mentioned similarities with CGs, we also have results regarding the hardness of a Braess-type paradox that arises in such networks.

Representative Papers.

[1] Risk Averse Selfish Routing (Math of OR, to appear)
with N.E. Stier-Moses, and E. Nikolova

[2] Network Pricing: How to Induce Optimal Flows Under Strategic Link Operators (EC 2018) with J. Correa, C. Guzman, E. Nikolova, and M. Schroder

[3] When Does Diversity of User Preferences Improve Outcomes in Selfish Routing? (IJCAI 2018) with R. Cole and E. Nikolova

THODORIS LYKOURIS

Thesis. Adaptive decision making in a data-driven society

Advisor. Éva Tardos, Cornell University

Brief Biography. Thodoris Lykouris is a Ph.D. candidate at the Computer Science department of Cornell University, advised by Éva Tardos. His research spans across the areas of theoretical computer science, machine learning, and economics, with a particular emphasis on the interplay between data (learning theory) and incentives (game theory) for adaptive decision making. Thodoris is a recipient of the 2018 Google Ph.D. Fellowship on "Algorithms, Optimizations, and Markets", and a finalist in the 2017 INFORMS Applied Probability Society Best Student Paper Competition. He was a research intern at Microsoft Research Redmond (summer 2018), Google Research NYC (summer 2017), and Microsoft Research India (summer 2015). He was also a visiting student at TTI-Chicago (March-May 2018) hosted by Avrim Blum, and the Simons Institute (fall 2015, spring 2018). Thodoris received his Diploma from the EECS department of National Technical University of Athens.

Research Summary. The rise of online marketplaces has altered the nature of economic activity. Transactions are repeated rapidly; for example, in online advertising, a large volume of transactions with small individual value occurs every second. Marketplaces have access to vast pools of data from past transactions that can enhance the quality of the system. Some of this data is available to the participants who learn to adapt their behavior accordingly. My research aims to address the challenges and embrace the opportunities of such a strategic and data-driven era in marketplaces of economic significance.

An important question is how marketplaces and participants can adapt to past data under the presence of strategic interactions. For example, consider a recommendation system; participants provide feedback about their past experiences and the system can use this data to learn the quality of different restaurants, thereby making more informed recommendations. In a STOC'18 paper, we address a fundamental issue: the feedback may be fraudulent, e.g. due to fake reviews by competitors. Although classical stochastic bandit learning algorithms provide guidance on how to act in the face of uncertainty, they can be fooled even by minimally corrupted data. In contrast, we suggest a simple scheme that makes stochastic bandits robust to such strategic manipulations.

In many online marketplaces, platforms use past data for revenue management; most prominently, ridesharing systems use surge pricing to adapt to present demand and supply. This is an intricate task since surge prices in one location affect future supply in another. In an EC'17 paper, we design a queueing-theoretic approximation framework for such systems that returns point-to-point prices that compete well against the best adaptive prices despite being non-adaptive themselves. Our guarantee improves with the ratio of vehicles over locations and achieves asymptotic optimality.

Since strategic participants interact with each other repeatedly, it is crucial to understand the quality of the resulting outcomes. Learning has emerged as a behavioral assumption capturing the incentives of the participants without the unrealistic informational assumptions of equilibrium-based behavioral concepts. In a SODA'16 paper, we use the fact that learning can adapt to changes in the environment to prove that learning agents converge to efficient outcomes even when the population of the game is dynamically and abruptly evolving across time.

Representative Papers.

[1] Stochastic bandits robust to adversarial corruptions (STOC 2018) with V. Mirrokni and R. Paes Leme

[2] Pricing and optimization in shared vehicle systems: An approximation framework (EC 2017) with S. Banerjee and D. Freund

[3] Learning and efficiency in games with dynamic population (SODA 2016) with V. Syrgkanis and É. Tardos

HONGYAO MA

Thesis. Coordinating Behavior: A New Role for Mechanism Design

Advisor. David C. Parkes, Harvard University

Brief Biography. Hongyao Ma is a Ph.D. candidate in Computer Science in the Paulson School of Engineering and Applied Sciences at Harvard, advised by Prof. David C. Parkes. Her research is broadly situated at the interface between economics and computer science, and draws on concepts from multi-agent systems and game theory. She is particularly interested in designing incentive-aligned systems to bring people together in useful ways, in the presence of uncertainty, self-interest, and autonomy. Ma is a recipient of the Siebel Scholarship, Harvard SEAS fellowship, the UCLA-CSST scholarship, and the Derek Bok Center Certificate of Distinction in Teaching at Harvard. Ma received her M.S. in 2014 at Harvard, and B.E. in 2012 at Xi'an Jiaotong University, both in Electrical Engineering. She interned at Uber in the summer of 2017, visited Technion in the summer of 2016, worked as a research intern at AT&T Labs Research in the summer of 2013, and visited UCLA as a part of the CSST Program in 2011.

Research Summary. I study how to design market-based systems for coordination problems, where society or the planner has an interest in people taking intended actions. Consider the problem of allocated resources going to waste: it is common for sport facilities to be fully booked, yet walking around the gym one sees many squash courts are in fact empty. Similar problems exist when asking consumers to reliably reduce electricity consumption if demand exceeds supply the following day. Consider ridesharing platforms like Uber and Lyft, where drivers are also customers. Drivers' strategic behavior, e.g. cherry-picking trips, declining trips to chase surge prices or wait for price hikes, undercut the platforms' reliability.

Common to many problems I study is a planning period followed by an action period, where people have uncertainty about the future at the time of planning, and will make decisions after uncertainty is resolved. I study the design of simple, indirect mechanisms based on novel penalty-bidding schemes, that achieve optimal utilization [3], or guarantee reliability. The maximum acceptable penalties provide good signals on people's reliability, and the actual penalties charged incentivize truthful reporting and shape downstream decisions. This work also lead to problems fundamental to mechanism design; for example, I have a tight characterization for the maximal non-quasi-linear utility domain where social choice mechanisms with good properties continue to exist [1].

In contrast to the above scenarios, penalizing drivers for declining dispatches in ridesharing is not effective, since penalties incentivize drivers to go offline, exacerbating the difficulties in estimating supply. Moreover, the need for strategic behavior is a symptom of a poorly designed system: sending drivers away from a sports stadium five minutes before a game ends at low prices is inefficient, and drivers are in fact acting to improve the system's efficiency. In recent work [2], for a model that captures the network flow imbalances and temporal variations in supply and demand, I propose the *Spatio-Temporal Pricing mechanism*, which uses information on future demand and supply, solves for the optimal driver flow and competitive equilibrium prices via linear programs, and is able to update the plan after any deviations. By dispatching and pricing trips correctly, the mechanism achieves envy-freeness and aligns incentives, without using penalties or time-extended contracts.

Representative Papers.

[1] Social Choice with Non Quasi-linear Utilities (EC 2018)
with R. Meir and D.C. Parkes

[2] Spatio-Temporal Pricing for Ridesharing Platforms (working paper, arXiv)
with F. Fang and D.C. Parkes

[3] Contingent Payment Mechanisms for Resource Utilization
(working paper, arXiv) with R. Meir, D.C. Parkes, and J. Zou

J. BENJAMIN MILLER

Thesis. Simple Mechanism Design with Robust Guarantees

Advisor. Shuchi Chawla, University of Wisconsin, Madison

Brief Biography. Benjamin Miller is a PhD candidate in the Department of Computer Sciences at the University of Wisconsin-Madison, advised by Shuchi Chawla. He completed his B.S. in Computer Science and Engineering Physics at Cornell University. His research focuses on simple mechanisms with robust guarantees against agent behavior. Benjamin was awarded a UW-Madison Cisco Systems Distinguished Graduate Fellowship in 2016-17 and 2017-18. In 2018, he was an intern in Facebook's Core Data Science group.

Research Summary. Most work in mechanism design assumes that agents' behavior can be captured by classical economic models, yet decades of behavioral research shows that these reliably fail to predict human behavior. This research also uncovers predictability which, perhaps surprisingly, can be tractably modeled. But mechanism design under these richer models remains largely unexplored. Below I highlight how my research begins to close that gap.

In the simplest auction setting—one item, one buyer—a deterministic price maximizes the seller's revenue, according to the classical risk-neutral model. In a SODA'18 paper, we study auction design under prospect theory, the predominant descriptive model of decision-making under uncertainty. We show that the optimal mechanism can use randomness to leverage a buyer's aversion to risk and extract more revenue than a deterministic mechanism. Nevertheless, we prove that a deterministic price recovers a significant fraction of the optimal revenue under a realistic bound on risk aversion. However, in a setting with repeated sales, we show that no constant approximation is possible unless the seller has detailed knowledge of the buyer's risk attitude.

In follow-up work [LMP], we investigate the same behavior model in settings with multiple items. In pathological instances, the model permits the seller to extract unbounded revenue, even without selling *any* item. The first challenge is thus to specify reasonable parameters in which to state our bounds. We study novel simple mechanisms which use limited randomness to exploit the buyer's risk attitude, and show bounds on the revenue relative to the revenue achievable under the classical model.

Finally, in an EC'16 paper, we study simple mechanisms for a seller with many goods and many buyers. Even with the classical model, the optimal auction may require randomized outcomes with sensitive dependence on buyers' bids. Furthermore, running the optimal auction—or even determining the participants' optimal bids—is often computationally intractable. We show that a much simpler mechanism which sequentially offers fixed prices (entry fee and item prices) earns a constant fraction of the optimal revenue. Our result belongs to a line of work on such simple mechanisms; ours is the first for multiple buyers with a broad class of valuations. Furthermore, in follow-up work (in preparation), we show that the mechanism's revenue is robust to prospect-theoretic behavior as studied above.

Representative Papers.

[1] Revenue Maximization with an Uncertainty-Averse Buyer (SODA 2018) with S. Chawla, K. Goldner, and E. Pountourakis

[2] Robust Mechanisms for a Prospect-Theoretic Buyer (in submission) with S. Liu and C.-A. Psomas

[3] Mechanism Design for Subadditive Agents via an Ex Ante Relaxation (EC 2016) with S. Chawla

RAD NIAZADEH

Thesis. Algorithms vs. Mechanisms: Mechanism Design for Complex Environments

Advisor. Robert Kleinberg, Cornell University

Brief Biography. Rad Niazadeh is a Motwani postdoctoral fellow at Stanford university, computer science department (theory group). During his postdoc, he is co-hosted by Tim Roughgarden, Amin Saberi and Moses Charikar. Prior to Stanford, he finished his PhD in computer science (minored in applied mathematics) from Cornell university, advised by Robert Kleinberg. His research interests lie broadly at the intersection of algorithm design, game theory and machine learning, with a focus on applications in market design and operations research. Rad was awarded the Motwani postdoctoral fellowship in 2017, Google PhD Fellowship (in market algorithms) in 2016, and Irwin Jacobs Fellowship in 2012. He was a long-term research scientist at Simons Institute (Berkeley) in Fall 2017. He was also an intern at Microsoft Research (Redmond) in summer 2016, Microsoft Research (New England) in summer 2015, and Yahoo! Research in Fall 2015.

Research Summary. Online marketplaces face various non-standard challenges stemming from intrinsic uncertainties (because of their stochastic or real-time natures), having to deal with incentives (because of human users), or existing non-convexities (because of complex structures or network effects). Broadly speaking, my research identifies, models, and tackles above challenges in various domains of market design, revenue management and machine learning. While current methods are broadly amplifying our understanding of many of these settings, I try to approach them from a fresh interdisciplinary perspective. I will describe some examples below.

When designing systems for human users, a triumph for the designer is to achieve Incentive Compatibility (IC). A striking question is then the following: Do running an IC mechanism require more computational resources than running an algorithm for the same problem? In [1], we resolve a five-year-old open question in this literature: *There is a polynomial time reduction from Bayesian IC mechanism design to Bayesian algorithm design for welfare maximization problems.* Unlike prior results, our reduction achieves exact incentive compatibility for problems with multidimensional and continuous type spaces.The key novel ingredient in our reduction is a generalization of the literature on Bernoulli factory in probability theory.

Inspired by spot market pricing in the cloud, I tackle the problem of dynamic pricing through the lens of adversarial online learning in a joint work with colleagues at Microsoft [2]. We reduce this problem (with extensions to multi-buyer auctions) to a pure learning framework termed as *multi-scale learning*, where the reward of each action is in a different range. By developing optimal no-regret algorithms for this learning framework using online mirror descent family, we obtain dynamic pricing/auctions with optimal convergence rates and multi-scale regret guarantees.

As more highlights, I have worked on open problems in mechanism design (e.g., see [3]) and non-convex optimization and machine learning. In my recent NIPS'18 paper with Roughgarden and Wang, we study the fundamental problem of maximizing a continuous non-monotone submodular function. Our main result is the *first* 2-approximation algorithm for this problem. This algorithm is optimal if polynomially many function queries are allowed, and has important applications in computer vision and revenue management under network effects.

Representative Papers.

[1] Bernoulli Factories and Black-Box Reductions in Mechanism Design (STOC 2017 & SIGecom Exchanges) with S. Dughmi, J. Hartline, and R. Kleinberg

[2] Online Auctions and Multi-scale Online Learning (JMLR 2018 & EC 2017) with S. Bubeck, N. Devanur, and Z. Huang

[3] Optimal Auctions vs. Anonymous Pricing (GEB 2018 & FOCS 2015) with S. Alaei, J. Hartline, M. Pountourakis, and Y. Yuan

EMMANOUIL POUNTOURAKIS

Thesis. Simple Mechanisms and Behavioral Agents: Towards a Theory of Realistic Mechanism Design

Advisor. Nicole Immorlica, Microsoft Research

Brief Biography. Emmanouil Pountourakis is a Post-doctoral fellow at University of Texas at Austin under the supervision of Evdokia Nikolova. During Spring 2018 he was a Fellow at the Real-Time Decision Making program at the Simons Institute. He received his PhD from Northwestern University in 2017 under the supervision of Nicole Immorlica. During his PhD he was a long-term visitor at Microsoft Research, completed a research internship at CWI Amsterdam, and was a student visitor at Hebrew University for the Special Semester in Algorithmic Game Theory in 2011. He holds an undergraduate and Masters degree in Computer Science from the University of Athens.

Emmanouil Pountourakis has a broad interest in algorithmic mechanism design. His current research focuses on revenue maximization in static and dynamic environments, and analyzing behavioral models of time-discounting and risk. Lately, his research focuses on using tools from algorithmic mechanism design to analyze energy market design.

Research Summary. My research interests lie in the intersection of theoretical computer science, economics, and game theory. The main focus of my work has been in algorithmic mechanism design, the theory of algorithm design in strategic environments. Particularly, I am interested in analyzing the performance of simple mechanisms that are prevalent in practical applications, as well as analyzing settings where the agent behavior deviates from the standard model. My contributions are summarized below.

Simple Auctions for Revenue-Maximization: There is a rich literature characterizing the revenue-maximizing auctions, however optimal mechanisms are commonly observed in practice. One such example is revenue maximization using simple auctions for the single item setting. We provided a tight analysis of the approximation ratio of anonymous pricing to the standard upper bound benchmark of ex-ante relaxation. It implied tightened upper bounds for the approximation of optimal auction from four to $e \approx 2.718$ for both anonymous posted pricing and second price auction with anonymous reserve. In another project we study the dynamic environment of selling a single item over several periods. We show that the optimal discriminatory mechanism obtains arbitrarily low revenue in a perfect Bayesian equilibrium. Surprisingly, restricting the feasible space of the mechanism to only anonymous pricing schemes allows good revenue guarantees.

In addition, I am interested in studying models beyond the standard rational agent model. For example, one of my works studies optimal contract design with a present-biased agent. When making a decision, the present-biased agent overestimates her present utility. The contract designer seeks to exploit that behavior to maximize their own utility. We characterize the optimal protocol and introduce regulations that reduce the exploitative power of the designer. Furthermore, I am interested in settings where the agent is not risk-neutral. In a recent work we study risk-robust auctions under a risk-averse agent in static and dynamic settings. In another project, we characterize the optimal mechanism for selling a single item to a risk-seeking agent.

Finally, my current research focuses on mechanism design for electricity market design. Particularly, I am studying the Market-clearing mechanism of the deregulated energy market as well as analyzing energy distribution systems with prosumers: consumers that in addition may produce energy.

Representative Papers.

[1] Repeated Sales with Multiple Strategic Buyers (EC 2017)
with N. Immorlica, B. Lucier, and S. Taggart

[2] Procrastination with Variable Present Bias (EC 2016)
with N. Gravin, N. Immorlica, and B. Lucier

[3] Optimal Auctions vs Anonymous Pricing (FOCS 2015)
with S. Alaei, J. Hartline, R. Niazadeh, and Y. Yuan

CHRISTOS-ALEXANDROS PSOMAS

Thesis. Algorithmic Mechanism Design in Dynamic Environments

Advisor. Christos Papadimitriou, University of California, Berkeley

Brief Biography. Christos-Alexandros (Alex) Psomas is a postdoctoral researcher in the Computer Science Department at Carnegie Mellon University, hosted by Ariel Procaccia. He received his PhD in 2017 from the Department of Electrical Engineering and Computer Sciences at UC Berkeley, under the supervision of Christos Papadimitriou. He is broadly interested in algorithmic economics, including topics such as algorithmic mechanism design, computational social choice, fair division and auction theory. During his PhD he spent two summers as a research intern at the International Computer Science Institute with Eric Friedman, a summer at Microsoft Research Redmond with Nikhil Devanur, and a summer as an instructor for the Discrete Mathematics and Probability Theory course at UC Berkeley. Prior to joining UC Berkeley, Alex received a MS in Logic, Algorithms and Computation from the University of Athens and a BS in Computer Science from Athens University of Economics and Business.

Research Summary. My research focuses on the intersection of Theoretical Computer Science and Economics. I apply tools and insights from computer science to study problems in a variety of economic environments. Some highlights of my research include the following.

Revenue maximization in dynamic environments. In mechanism design we typically study one-shot auctions. However, multi-stage environments are more common in the real world. My work in this space begins with the observation that restricting to static mechanisms results in significant losses in revenue. Why is it then that dynamic mechanisms are not more prevalent in practice? Surprisingly, we prove [PPPR'16] that the problem of maximizing revenue from selling two items across two rounds, arguably the simplest meaningful dynamic mechanism design problem imaginable, is computationally intractable. On the other hand, we prove [LP'18] that introducing additional competition to the market and running a simple static mechanism in each round is better, in terms of revenue, than running the optimal dynamic auction. Combined, these results provide a partial answer to the above question.

Fair division in dynamic environments. Consider the following problem faced by a food bank. Food donations arrive, and must be delivered to nonprofit organizations such as food pantries and soup kitchens. Items are often perishable, which is why allocation decisions must be made quickly, and donated items are typically leftovers, leading to lack of information about items that will arrive in the future. In this setting, how should we make allocation decisions in a way that is fair to the donation recipients? In [BKPP'18] we study such a dynamic fair division problem, where indivisible goods arrive and must be allocated immediately and irrevocably to one of n agents. Our goal is to minimize the maximum envy between any two agents, after all the goods have been allocated. We give a polynomial-time, deterministic and asymptotically optimal algorithm with vanishing envy, i.e. the maximum envy divided by the number of items T goes to zero as T goes to infinity. On the applied side, we've been working with a food bank called called 412 Food Rescue in order to develop an algorithm that would automate their allocation decisions in a way that is ethical, fair and efficient.

Representative Papers.

[1] On the Complexity of Dynamic Mechanism Design (SODA 2016)
with C. Papadimitriou, G. Pierrakos, and A. Rubinstein

[2] On the Competition Complexity of Dynamic Mechanism Design (SODA 2018)
with S. Liu.

[3] How to Make Envy Vanish Over Time (EC 2018)
with G. Benade, A. Kazachkov, and A. Procaccia

GORAN RADANOVIC

Thesis. Elicitation and Aggregation of Crowd Information

Advisor. Boi Faltings, Swiss Federal Institute of Technology in Lausanne

Brief Biography. Goran Radanovic is a postdoctoral researcher at Harvard University, where he works with Prof. David C. Parkes on topics related to value-aligned artificial intelligence, social-facing computing, and human-AI collaboration. His research particularly focuses on incentive mechanism design and reinforcement learning with humans. Goran Radanovic received his Ph.D. in Computer Science from the Swiss Federal Institute of Technology in Lausanne (EPFL) in 2016, under the supervision of Prof. Boi Faltings. He is a recipient of Early Postdoc.Mobility Fellowship (2016-2018) from Swiss National Science Foundation, and was awarded with EPFL Ph.D. Distinction for an outstanding dissertation in 2017.

Research Summary. I am interested in understanding the effect of belief misalignment in multi-agent systems with humans. Belief misalignment can arise, e.g., due to human behavioral biases or inherent subjectivity of prior information. Much of my research focuses on two domains: data elicitation and human-AI collaboration.

When gathering data from distributed agents, one has to account for the cost the agents suffer from data mining, and compensate them for it, e.g., using incentive mechanisms. If the accuracy of the gathered data cannot be directly verified, incentive mechanisms have to be based on examining consistency between reported values. This implies that a proper incentive mechanism has to take into account how agents reason about each other, which is reflected through their beliefs that are not necessarily known by the mechanism. We develop and test novel game-theoretic incentive mechanisms that provide proper rewards for accurate reporting of complex (non-binary) information without requiring access to agents' beliefs. Our results generalize the existing class of *Peer-Prediction* mechanisms and demonstrate boundaries on what is possible in the considered elicitation settings. To this end, we show the importance of acquiring an additional report from each agent when agents have unconstrained beliefs, usefulness of partially truthful incentives in eliciting accurate aggregates, and a possibility of utilizing a decomposable micro-task elicitation structure to obtain stronger incentive properties.

Within the human-AI collaboration domain, my research focuses on designing AI policies that account for human imperfections. These imperfections are incorporated into the belief system of a human agent and can model behavioral biases that preclude a human agent from deriving optimal policies. For example, consider an intervention game where an AI agent controls an object of interest (e.g., a self-driving car), but a human agent can make costly interventions to take control. Knowing that humans overestimate low-probability events, the AI agent can avoid costly interventions by adjusting its policy to the risk-aversion of the human agent. We model the scenario using *Multi-View Decision Processes* — a reinforcement learning framework that we show to be (weakly) more general than stochastic games. We further derive the basic properties of the framework and develop algorithms for computing the agents' policies.

Representative Papers.

[1] A Robust Bayesian Truth Serum for Non-Binary Signals (AAAI 2013) with B. Faltings

[2] Incentives for Effort in Crowdsourcing Using The Peer Truth Serum (ACM TIST 2016) with B. Faltings and R. Jurca

[3] Multi-View Decision Processes: The Helper-AI Problem (NIPS 2017) with C. Dimitrakakis, D.C. Parkes, and P. Tylkin

AMIN RAHIMIAN

Thesis. Learning and Decision Making in Groups

Advisor. Ali Jadbabaie, Massachusetts Institute of Technology

Brief Biography. In 2017, I received a postdoctoral fellowship from MIT Institute for Data, Systems, and Society to conduct research at the intersection of social and computational sciences, where I am currently working with Dean Eckles and Elchanan Mossel. Previously, I received my AM in Statistics from the Wharton School, MS in Systems Engineering, and PhD in Electrical and Systems Engineering all from the University of Pennsylvania. I was a finalist in 2015 Facebook Fellowship Competition, as well as 2016 ACC Best Student Paper Competition. In 2016 I also received an outstanding poster award at the Stochastic Networks Conference, as well as the Graduate Award for Best Teaching Assistant at the Doctoral Level form the Penn ESE department. My research interests are at the intersection of network science, statistics, control and decision theory, with applications to social and economic networks.

Research Summary. The majority of my work is devoted to challenges of design, analysis, control, and decision making in networks. Understanding that which paradigms best describe and predict aggregate behavior in networks is important both for understanding whether information transmission is efficient, and for thinking through social and organizational policy designs that rely on information dissemination.

In my postdoc work (with Dean Eckles, Elchanan Mossel, and Subhabrata Sen), we study how interventions that change network structure can substantially affect the spread of new ideas, products, and conventions over a social network. For simple contagion models borrowed from epidemic spread, highly clustered networks slow spread compared with more random networks, such that interventions that randomly rewire edges would increase spread. However, for other contagion models that require multiple exposures before adoption (i.e. complex contagions), such as those motivated by myopic best-reply in games with strategic complements, recent work has argued for the opposite conclusion: highly clustered, rather than random, networks facilitate spread. Here we show that slight modifications of prior analyses, by allowing a small $O(1/\sqrt{n})$ probability of simple (single exposure) adoptions, which make them more realistic, reverse this result. We verify our theoretical findings using simulations on networks of Chinese and Ugandan villages.

In my PhD work (with Jan Hazla, Elchanan Mossel, and Ali Jadbabaie), we address the computations that Bayesian agents undertake in an opinion exchange model over a network. The agents act repeatedly on private information and take myopic actions that maximize their expected utility according to a fully rational posterior. We show that such computations are NP-hard for two natural utility functions, including the case where agents reveal their posteriors. Our results are robust in the sense that they show NP-hardness of distinguishing (and therefore also approximating) between posteriors that are concentrated on two distinct states of the world. We also provide an algorithm to compute agents' actions using iterated elimination of infeasible signals and show that if the network is transitive, the algorithm can be modified to run in polynomial time.

Representative Papers.

[1] Long ties accelerate realistic complex contagions (Job Market Paper)
with D. Eckles, E. Mossel, and S. Sen

[2] Bayesian Decision Making in Groups is Hard (SSRN)
with J. Hazla, A. Jadbabaie, and E. Mossel

ASSAF ROMM

Thesis. Essays in Market Design

Advisors. Alvin Roth and Drew Fudenberg, Harvard University

Brief Biography. Following a BSc and MA from the Hebrew University, I've completed my PhD in Economics from Harvard at 2015, and joined the Hebrew University Department of Economics and the Center for Study of Rationality as an Assistant Professor. I've received an ISF grant and a BSF grant (two out of the three leading grants in Israel). Starting from 2017, I also work at Microsoft Research, Israel, which has been a great experience. This has also revealed the number of people in CS whose research interests are very similar to mine. I am in the process of relocating due to my wife's career, and figured it may be a good time to also look at positions in CS departments. Market design is a multi-disciplinary field, and I, in particular, am a multi-disciplinary person, and due to my background I am very comfortable working in either environment.

Research Summary. I'm generally interested in AGT, but more specifically in market design. My papers are always motivated by an economic question, and guided by economic intuitions, but often use tools that are traditionally associated with CS and OR.

My research has both theoretical and applied aspects. In the past 5 years I have redesigned several markets in Israel (with collaborators), and those often led to interesting research directions. Working on school choice led to a paper on tie-breaking methods and another one on fairness properties of different algorithms. The Israeli Medical Internship Match was followed by a study on couples in assignment problems. The Israeli Psychologists Master's Match raised questions about how people react to strategy-proof mechanisms, and about properties of the core (the set of stable matchings) in matching markets with contracts. The recent pre-military academies match was mostly focused on designing good algorithms when stable matchings could not be found, and the on-going lawyers match leads to many directions that are related to signaling mechanisms in general.

Other work, inspired by real markets but not directly related to any specific one, involves studying two-sided markets with heterogeneous goods, showing an approximate law of one price. This specific work utilizes tools from both matching theory and graph theory to expand a classic model of economic markets (presented by Shapley and Shubik in the 70's) and study likely properties of such markets.

Representative Papers.

[1] Redesigning the Israeli Psychology Master's Match (EC 2017, AER P&P 2017) with A. Hassidim and R.I. Shorrer

[2] An approximate law of one price in random assignment games (EC 2015) with A. Hassidim

[3] Assigning more students to their top choices: a tiebreaking rule comparison (EC 2015) with I. Ashlagi and A. Nikzad

KARTHIK ABINAV SANKARARAMAN

Thesis. Randomized Algorithms for Computational Economics and Machine Learning

Advisor. Aravind Srinivasan, University of Maryland, College Park

Brief Biography. I am a PhD student at the University of Maryland, College Park advised by Aravind Srinivasan. My research focuses on understanding the power of randomized algorithms to tackle problems in matching markets, online learning theory and algorithmic statistics. I am particularly excited about designing and understanding algorithms in *challenging* environments (*e.g.,* noisy, adversarial, changing distributions, etc.). In all of these cases deterministic algorithms perform very poorly while clever randomized algorithms can achieve surprisingly good performance. During my time as a PhD student, I have interned at Adobe, IBM, and Microsoft, working with world-class researchers from a variety of fields, all of whom have shaped my research philosophy, bag of techniques and world-view of important problems.

Research Summary. My current research focus is on the foundations and applications of algorithms to problems in computational economics and data science. The first problem we study comes from online matching markets. A unifying model to study allocations in modern markets is online bipartite matching. However, in many of these applications, we have additional constraints that prevent us from using vanilla bipartite matching models and hence calls for newer models and algorithms. A core challenge is to develop algorithms that are robust to work in the online setting and models which incorporate the availability of historical data. In a series of works, including the AAAI 18 paper, along with my co-authors, we consider many variants of matching markets including crowdsourcing, recommender systems, and ride-sharing and give provable algorithms. Along the way, we develop algorithms for natural foundational questions (e.g., the SODA 18 paper) related to allocation problems. Finally, we also evaluate our algorithms on many real-world datasets for these domains. Next, we study the dynamic pricing problem and the more general multi-armed bandits with knapsacks. We study problems in the stochastic and the more challenging adversarial setting. The core foundational contribution is algorithms that work for a very general problem with near-optimal guarantees. From an applied side, we show many applications including dynamic pricing, dynamic assortment, and network revenue optimization. The AIStats 18 paper considers the combinatorial semi-bandit problem with knapsacks where one is selling multiple goods simultaneously, under the assumption that the user valuation for these goods is drawn i.i.d. from an unknown distribution. A sale of good leads to consumption of some shared resources and the goal is to learn the optimal prices maximizing the total revenue under a hard constraint on the total available resources. In a working paper, we generalize to a setting where the user valuation can be adversarial; we present a general algorithm and prove that its optimal. Finally, we study algorithmic foundations of causal inference, a tool used commonly in econometrics. We consider the linear structural equation model (SEM) of causality and study the robustness of these algorithms to noise (via its condition number) from both theoretical and experimental perspectives. We characterize for a large subclass of SEM, the behavior of the condition number and validate it on a popular sociology dataset.

Representative Papers.

[1] Allocation Problems in Ride-Sharing Platforms (AAAI 18)
with J. Dickerson, A. Srinivasan, and P. Xu

[2] Algorithms to Approximate Column-Sparse Packing Programs (SODA 18)
with B. Brubach, A. Srinivasan, and P. Xu

[3] (Adversarial) Bandits with Knapsacks (in preparation)
with N. Immorlica, R. Schapire, and A. Slivkins

YIXIN TAO

Thesis. Market efficiency, dynamics, and optimization

Advisor. Richard Cole, New York University

Brief Biography. Yixin Tao is a Ph.D. student in the Computer Science Department, Courant Institute of Mathematical Sciences, New York University, supervised by Richard Cole. Before coming to NYU, he received my BS degree in Computer Science from ACM Honor Class at Shanghai Jiao Tong University.

Research Summary. My research focuses on Algorithmic Game Theory (AGT) and Optimization. In AGT, my work mainly concerns market efficiency and market dynamics; and in Optimization, my work addresses asynchronous implementations of coordinate descent.

Market Efficiency: In the standard economic model, many classes of markets have efficient equilibria, assuming agents are non-strategic, i.e. they declare their true demands when offered goods at particular prices. An important question is how much the equilibria degrade in the face of strategic behavior, i.e. what is the Price of Anarchy (PoA) of the market viewed as a mechanism? We focus on two types of market, Walrasian equilibria for indivisible goods and Fisher Markets for divisible goods. Our main result is that, given suitable assumptions, in the large market setting, the PoA of the Nash Equilibrium (NE), including pure, mixed and Bayes NE, tends to 1 as the market size increases.

Market Dynamics: A major goal in AGT is to justify equilibrium concepts from an algorithmic and complexity perspective. One appealing approach is to identify natural distributed algorithms that converge quickly to an equilibrium. In the Fisher Market setting, we established new convergence results for two generalizations of Proportional Response when buyers have CES utility functions. The starting points are new convex and convex-concave formulations of such markets. The two generalizations correspond to suitable mirror descent algorithms applied to these formulations. Our results follow from new notions of strong Bregman convexity and of strong Bregman convex-concave functions, and associated linear rates of convergence.

Asynchronous Optimization: Coordinate descent is a core tool in machine learning and elsewhere. Large problem instances are common. To help solve them, two orthogonal approaches are known: acceleration and parallelism.One important issue in parallel implementations is whether the different processors are using up-to-date information. Asynchronous updating avoids the need for up-to-date data, reducing and potentially eliminating the need for waiting. In our work, we give a comprehensive analysis of Asynchronous Stochastic Accelerated Coordinate Descent. We show: A linear speedup for strongly convex functions so long as the parallelism is not too large; a substantial, albeit sublinear, speedup for strongly convex functions for larger parallelism; a substantial, albeit sublinear, speedup for convex function.

Representative Papers.

[1] An Analysis of Asynchronous Stochastic Accelerated Coordinate Descent (in submission) with R. Cole

[2] Dynamics of Distributed Updating in Fisher Markets (EC 2018) with Y. Cheung and R. Cole

[3] Large Market Games with Near Optimal Efficiency (EC 2016) with R. Cole

ALAN TSANG

Thesis. Strategic Voting and Social Networks

Advisor. Kate Larson, University of Waterloo

Brief Biography. Alan Tsang is Postdoctoral Researcher at the National University of Singapore, hosted by Yair Zick as part of the Data-Driven Strategic Collaboration Group. He obtained his PhD at the Cheriton School of Computer Science at the University of Waterloo, under the supervision of Kate Larson. He was supported by the David R. Cheriton Graduate Scholarship, and the Ontario Graduate Scholarship. Prior to that, he completed his M.Math (specializing in Graph Theory) and his B.Math (Bioinformatics option) at the University of Waterloo. Alan also has experience working in various industry position via the co-operative education program during his bachelors degree, as well as working in a start-up after his graduation. In addition, Alan has extensive classroom experience as tutorial leader for weekly review and QA sessions, as well as lecturer for the Introduction to AI course (offered to both upper-year undergraduates and graduate students).

Research Summary. My primary research interest in exploring decision making in social networks. I use a variety of computational and statistical tools in my approach, and leverage game theoretic concepts in understanding the rationale behind emergent behaviors. My past research sought inspiration from and collaborations with a variety of fields, ranging from psychology, political science, to forestry and fire management. I continue to seek opportunities interdisciplinary collaboration.

Social networks model the flow of information through human interactions. In [1], I examine examine how information and opinions propagate through networks where agents are skeptical. In this "opinion dynamics" domain, each agent has an opinion that evolves over time based on interactions with her neighbors, growing closer over time. My contribution introduces the concept of skepticism, where agents distrust neighbors who hold opinions different from their own, a psychological phenomenon called confirmation bias. We show that even in networks with high levels of skepticism, agents are able to come to consensus. However, in homophilic networks – ones where agents preferentially connect to those similar to themselves – agents may fail to converge to consensus. Presently, I am examining novel ways of modeling opinions that are defined as intervals in the opinion space, where an opinion specifies a range of feasible values; this is ongoing work with collaborators from the Nanyang Technological University and the University of Kentucky.

I am also interested in examining how agents make rational decisions based on information received through social networks. In particular, in [2,3] I use game theoretic concepts to model when voters exercise strategic voting in social networks. My model suggests one reason why relatively few voters vote strategically in real world elections: network homophily. When people surround themselves with like-minded individuals, their perception of the world can become skewed. Voters who favor an unpopular candidate mistakenly believe there is sufficient support for their candidate, and fail to see the opportunity to vote strategically instead. Our model empirically demonstrates the emergence and ramifications of this so-called "Echo Chamber Effect". A current line of research looks to examine how human participants vote in such a social network context, and whether this effect can be replicated in laboratory conditions.

Representative Papers.

[1] Opinion Dynamics of Skeptical Agents (AAMAS 2014)
with K. Larson

[2] The Echo Chamber: Strategic Voting and Homophily in Social Networks (AAMAS 2016) with K. Larson

[3] Boundedly Rational Voters in Large(r) Networks (AAMAS 2018)
with A. Salehi-Abari and K. Larson

PAN XU

Thesis. Online and Offline Matching Algorithms for E-Commerce

Advisors. John Dickerson and Aravind Srinivasan, University of Maryland, College Park

Brief Biography. Pan is a Ph.D. student in the Department of Computer Science at the University of Maryland (UMD), College Park, co-advised by Dr. John Dickerson and Dr. Aravind Srinivasan. Pan's research interests broadly span the intersection of Algorithms, Operations Research, and Artificial Intelligence. He currently focuses on algorithm design for offline and online matching models and their applications in various real matching markets, including rideshare, data center scheduling, and labor markets. Previously, he worked intensively on high-dimensional data indexing and bilevel linear program optimization problems. He has received several fellowships and awards, including a Miller Graduate Fellowship (2009-2012, Iowa State University), a Research Excellence Award (2013, Iowa State University), an Ann G. Wylie Dissertation Fellowship (2018-2019, UMD), and an Outstanding Graduate Assistant Award (2018, UMD).

Research Summary. Stochastic matching models capture noisy real-world applications such as labor markets, online dating, and organ allocation, where uncertainty exists for the outcome of each potential pairing of two participants in the system. In the offline setting, we are given a general graph where each edge e has a non-negative weight and a *known* existence probability, while each vertex v has a patience $t_v \in Z_+$, representing a "budget" on potential failures to match (e.g., an upper bound on the number of job interviews or in-person dates). Our task is to construct a maximum expected-weight matching by probing each edge sequentially subject to both the matching and patience constraints. The Online Stochastic Matching (OSM) is primarily inspired by the Internet advertising business. Consider a bipartite graph $G = (U, V, E)$ where U (the offline advertisers) is known *a priori*, while V (the online impressions/keywords) arrive sequentially in a random manner; upon each arrival of a vertex $v \in V$, a central clearinghouse needs to make an *irrevocable and immediate* decision to match or reject. The task is to design an online matching policy such that the size of the final matching is maximized. My work (APPROX, 2015; Algorithmica, 2017; SODA 2018) studied the offline stochastic matching and its generalization to k-column sparse packing programs, while (ESA, 2016; AAMAS, 2017) considered OSM and its variants under known arrival distributions. More details can be found from my website: https://sites.google.com/site/panxupi.

To complement my theoretical work, I also apply online matching models in other domains. For "crowdsourcing human resources" marketplaces (matching online workers to offline tasks), we proposed a Multi-Budgeted OSM model, where we assume each match will consume multiple offline resources (AAMAS, 2017). For ridesharing platforms (matching online users to offline drivers), we crafted a variant of OSM, called *Online Matching with (offline) Reusable Resources*, to capture the fact that most drivers will be available again after completing a trip (AAAI, 2018). For online food-ordering platforms like Grubhub, we developed a new model, called *Online Task Assignment with Two-Sided Arrival*, to address the new challenge that both workers and tasks can join the system dynamically (AAMAS, 2018).

Representative Papers.

[1] Assigning Tasks to Workers based on Historical Data: Online Task Assignment with Two-sided Arrivals (AAMAS 2018) with J. Dickerson, K.A. Sankararaman, and A. Srinivasan

[2] Allocation Problems in Ride-Sharing Platforms: Online Matching with Offline Reusable Resources (AAAI 2018) with J. Dickerson, K.A. Sankararaman, and A. Srinivasan

[3] Algorithms to Approximate Column-Sparse Packing Programs (SODA 2018) with B. Brubach, K.A. Sankararaman, and A. Srinivasan

JUBA ZIANI

Thesis. Data for markets; markets for data

Advisors. Katrina Ligett and Adam Wierman, California Institute of Technology

Brief Biography. I am currently a PhD student in the Computing and Mathematical Sciences Department at Caltech, where I am advised by Katrina Ligett and Adam Wierman. My research lies at the intersection of Computer Science and Economics, with a focus on markets for data, data privacy, and fairness. I was a visiting student at Microsoft Research New England, hosted by Vasilis Syrgkanis, in Spring 2017, a research intern at Microsoft Research New England with Brendan Lucier in Summer 2017, and a visiting student at University of Pennsylvania, hosted by Aaron Roth, in Summer 2018. I received my B.Sc in 2011 and M.Sc in Engineering in 2013 from École Supérieure d'Électricité, France, as well a M.Sc in Operations Research from Columbia University in 2012.

Research Summary. I work on many aspects of large datasets, with an emphasis on strategic considerations and societal concerns.

One aspect of my research is to study mechanism design for markets for data. I have worked on designing mechanisms to efficiently buy and aggregate data from strategic agents; specifically, in an EC'18 paper with Y. Chen, N. Immorlica, B. Lucier, and V. Syrgkanis, we study the question of how to optimally collect data from and compensate strategic agents with private costs correlated to their data, in order to optimize the accuracy of a statistical estimate of interest under a budget constraint. I am also interested in privacy concerns in markets for data; in particular, I have an ITCS'15 paper with R. Cummings, K. Ligett, A. Roth and Z. S. Wu in which we design a mechanism that aggregates data from agents and gives differential privacy guarantees to the participants.

The introduction of information in standard market settings can also have surprising effects. I am currently working on a paper that asks the question of how the presence of a third-party data provider or source of outside information affects mechanism design. Y. Cai, F. Echenique, H. Fu, K. Ligett, A. Wierman and I consider a setting in which the bidders may have incomplete information about the type of the item that they are bidding on, and the third-party data provides additional information to the bidders about the item type. We show that in such a setting, simple mechanisms cannot achieve a constant fraction of the revenue of the optimal auction.

Large datasets also offer an opportunity to better understand strategic behavior. I have worked on two papers that address the question of how observations on strategic agents' behavior in complex settings can be used to recover their utilities. I have studied this question in a manuscript with V. Chandrasekaran and K. Ligett, and contributed to a paper with V. Syrgkanis and E. Tamer that is currently in the revise and resubmit phase at Econometrica.

Recently, I have also begun to study additional societal concerns that arise from the use of personal data. In a manuscript with N. Immorlica and K. Ligett, we study unfairness in university admissions caused by unequal access to strategic signaling. In a manuscript with S. Kannan and A. Roth, we study whether and – when possible – how a school can guarantee that similar students from different populations will be treated equivalently down the job market pipeline.

Representative Papers.

[1] Optimal Data Acquisition for Statistical Estimation (EC 2018)
with Y. Chen, N. Immorlica, B. Lucier, and V. Syrgkanis

[2] Accuracy for Sale: Aggregating Data with a Variance Constraint (ITCS 2015)
with R. Cummings, K. Ligett, A. Roth, and Z.S. Wu

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

Index