SIGecom Job Market Candidate Profiles 2016

Edited by SHADDIN DUGHMI, VASILIS GKATZELIS, and JASON HARTLINE

In the sixteen years since the inception of EC as the flagship conference of the SIGecom community, we have learned a lot about the efficiency (or lack thereof) of markets. We have studied markets for school choice, marriage, supply chains, advertising, etc. We understand the loss in welfare from decentralized versus centralized market mechanisms. We understand that information plays an important role in market efficiency and stability. Profiled herein are thirty candidates for the 2016 junior academic job market with research areas that span the interest of the SIGecom community. Though it may not be a complete list, we hope that it will help alleviate some of the informational inefficiencies in the most important market for the community, its junior faculty job market.

A table of contents lists the candidates in alphabetical order with keywords they provided. The candidate profiles are listed with brief biographies, research summaries, and three representative papers. Following the candidate profiles is an index of the candidates by keywords.

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

Contents

NICK ARNOSTI

Thesis. Frictions in Matching Markets

Advisor. Ramesh Johari, Stanford University

Brief Biography. I am a PhD candidate in the department of Management Science and Engineering at Stanford University, with a research focus on market design. I am drawn to this area because it offers the opportunity to develop deep and elegant theory that has clear practical applications. I have found that most of my projects start as conversations. I am insatiably curious, and have never found myself short of questions to ask or problems to work on. I believe that the best problems come from observing and understanding the operation of real-world marketplaces. Once I have selected a problem, the tools that I use to address it are drawn from my background in computer science, statistics, and probability. One technique that I have repeatedly found useful is to use simulations to identify interesting patterns and build intuition (though as a theoretician, I am not satisfied until I have leveraged this understanding into analytical proofs).

Research Summary. My research is in market design. I am particularly interested in studying the effects of frictions in matching markets. Frictions are barriers that market participants face when trying to find suitable partners. They take many forms, but two common ones are screening costs (incurred when evaluating potential partners) and application costs (incurred when communicating interest to these partners).

One recent paper of mine studies the outcomes of centralized matching markets when people learn their preferences through time-consuming interviews. In such markets, agents may waste interviews on others who are unlikely to be available to them, and may go unmatched despite the presence of an unrealized mutually agreeable match. I study the magnitude of these welfare losses, under various assumptions. In a school assignment setting, my work predicts a tradeoff between two tie-breaking procedures: one matches more students to their top choice, while the other matches more students overall. Thus, the choice of procedure should depend on how the district weighs these two goals.

Of course, most matching markets are not cleared centrally. One goal of mine is to better understand the costs and benefits of using a central clearinghouse when participants are initially uncertain about their preferences. Though the benefits are well-studied, the costs less understood. One cost is that participants in centralized procedures evaluate potential partners "in advance", and must choose between evaluating many partners (at high cost), and listing only a few (increasing the risk of going unmatched).

A theme in my research is a focus on simple mechanisms, rather than "optimal" ones. In work with Ramesh Johari and Yash Kanoria, I study congestion in decentralized matching markets. Rather than seeking a first best solution, we study the benefits that can accrue if the market operator merely limits the number of jobs for which each individual can apply.

Though simple interventions are appealing from a practical point of view, it is important to ask whether they are "adequate". One way to do this is to use the optimal procedure as a benchmark, and seek a simple procedure that performs well against this benchmark. This is the approach taken in my work with Marissa Beck and Paul Milgrom - we introduce a simple mechanism with several appealing properties, and then demonstrate that the gains of moving to any alternative mechanism are minimal.

Representative Papers.

[1] Short Lists in Centralized Clearinghouses (SSRN)
[2] Managing Congestion in Dynamic Matching Markets (SSRN)
with R. Johari and Y. Kanoria
[3] Adverse Selection and Auction Design for Internet Display Advertising (SSRN)
with M. Beck and P.R. Milgrom

MARKUS BRILL

Thesis. Set-valued Solution Concepts in Social Choice and Game Theory: Axiomatic and Computational Aspects

Advisor. Felix Brandt, Technische Universität München

Brief Biography. Markus Brill is a postdoc at the Department of Computer Science at Duke University. He received a Ph.D. degree in Computer Science (2012) and diploma (2008) and B.Sc. (2006) degrees in Mathematics from Technische Universität München (TUM), Germany. Markus is a graduate of the elite graduate program TopMath. For his Ph.D. thesis, he received the dissertation award of Bund der Freunde der TUM and an honorable mention for the Artificial Intelligence Dissertation Award sponsored by ECCAI. He is also a recipient of a Research Scholarship by ParisTech and currently holds a Feodor Lynen Research Fellowship, awarded by the Alexander von Humboldt Foundation.

Research Summary. I am fascinated by the formal analysis of scenarios in which multiple agents with possibly conflicting preferences interact. As such, my research interests lie at the intersection of computer science, theoretical economics, and mathematical social sciences. More precisely, my work focusses on axiomatic and computational aspects of game theory and social choice theory.

In my Ph.D. thesis, I have studied set-valued solution concepts. The first part of my thesis focusses on solution concepts for normal-form games that are based on varying notions of dominance [1]. The framework is very general and captures a number of concepts that have been proposed in the literature. In the second part of my thesis, I studied social choice functions and the complexity of the winner determination problem. I also studied the complexity of computing possible and necessary winners for partially specified instances and the axiomatic and asymptotic properties of tournament solutions [2].

My current research interests include various randomized solution concepts such as random serial dictatorship (RSD). My work has led to the identification of a fundamental tradeoff between strategyproofness and economic efficiency, and to results on the computational complexity of RSD [3]. I'm also working on finding axiomatically desirable and computationally feasible ways to randomized tiebreaking.

Moreover, I'm interested in mechanism design settings in which the power of the game designer is limited in the sense that she can only control part of the game. A natural question is whether one can "complete" an incompletely specified game in such a way that the resulting game has certain properties.

As a final example of ongoing work, let me mention an interdisciplinary research project whose goal is to arrive at numerical tradeoffs between different kinds of socially undesirable activities. For example, can we say that using one gallon of gasoline is just as bad for society as creating x bags of landfill trash? How would we arrive at a reasonable value of x? Such estimates would be useful to policy makers as well as well-meaning institutions and individuals. The vision is to create a system that can credibly arrive at numerical values for societal tradeoffs. A successful solution seems to require the application of techniques from a variety of research areas, such as game theory, social choice theory, mechanism design, prediction markets, etc.

Representative Papers.

[1] Computing Dominance-Based Solution Concepts (EC 2012)
with F. Brandt

[2] Minimal retentive sets in tournaments (SCW 2014)
with F. Brandt, F. Fischer, and P. Harrenstein.

[3] The computational complexity of random serial dictatorship (EL 2013)
with H. Aziz and F. Brandt

YUN KUEN CHEUNG

Thesis. Analyzing Tatonnement Dynamics in Economic Markets

Advisor. Richard Cole, New York University

Brief Biography. I am a postdoctoral scholar in the Computer Science Department at the University of Vienna, working with Monika Henzinger. In 2014, I received my PhD from the Courant Institute of Mathematical Sciences, New York University, where I was advised by Richard Cole. Before my PhD, I received an MPhil in Mathematics and a BSc in Mathematics and Physics, both from the Hong Kong University of Science and Technology (HKUST). My MPhil thesis was honored with the New World Mathematics Silver Award in 2010. I was a team member of the HKUST Programming team for two years, and was in the top four in two regional contests of the ACM ICPC. I was a bronze medallist in the 2004 International Mathematical Olympiad.

Research Summary. My research interests are in computational economics and algorithmic game theory. One of my main foci is the convergence analysis of the well-known tatonnement price dynamic. I am also working on mechanism design problems and game-theoretic aspects of markets.

Arguably, tatonnement is the most well-studied price dynamic in the theory of markets; it was introduced by Walras in 1874, along with the market equilibrium concept. Scarf showed tatonnement did not always converge to an equilibrium. Consequently, a fundamental problem is to identify broad classes of markets in which it does converge. Before my work, convergence was known for markets with goods that are substitutes, but little was known w.r.t. complementary goods. We are the first to show that tatonnement converges quickly in a number of interesting market classes with complementary goods, and more generally in some markets with both substitutes and complements.

One of the results mentioned above relies on the equivalence of tatonnement and coordinate descent for many markets. Motivated by asynchronous variants of tatonnement, we have been studying asynchronous coordinate descent (ACD), which has recently drawn attention in optimization theory. We designed a novel amortized technique to analyze ACD, with a general update rule that covers most known update rules, e.g. the round-robin rule.

Recently, I have worked on mechanism design problems with conflict-based negative externalities, e.g. in an ad auction, an advertiser's value for an ad slot drops when its rival has a better slot. We proposed a model for such externalities, and designed mechanisms with good approximation guarantees to the social welfare. One of our main results is to design a cone program (CP), which is a combination of a semi-definite program for independent set problems and the standard linear program for combinatorial auctions, and a rounding scheme for the CP, which achieves the best approximation ratio that one would expect.

I am now starting to work on game-theoretic aspects of markets. While truthful markets are efficient, individual agents can benefit by misreporting. Such manipulations are formulated by casting markets into games; thus analyses in terms of well-known effectiveness measures, e.g. the price of anarchy, are readily motivated. I am interested in how the amount of money an agent has and the similarity of her preferences to those of other agents influences both her power to manipulate and market efficiency.

Representative Papers.

[1] Tatonnement Beyond Gross Substitutes? Gradient Descent to the Rescue (STOC 2013) with R. Cole and N. Devanur

[2] Tatonnement in Ongoing Markets of Complementary Goods (EC 2012) with R. Cole and A. Rastogi

[3] Asynchronous Coordinate Descent and Tatonnement (submitted) with R. Cole

ILAN COHEN

Thesis. Online Algorithms and Game Theory

Advisor. Yossi Azar, Tel Aviv University

Brief Biography. Ilan Cohen is a PhD student at the Blavatnik School of computer science in Tel Aviv University under the supervision of Professor Yossi Azar. He holds an M.S. in computer science from Tel Aviv University and a B.S. cum laude in computer science from the Technion Institute in Haifa. His research involves online and approximation algorithms with game theoretical aspects. During the past three years, he has been a teaching assistant in the Algorithms course. Prior to his doctoral program, he worked as an algorithms developer and a programmer at IDF in the intelligence corps.

Research Summary. My research interests lie at the intersection of approximation algorithms, online algorithms and game theory. My work is divided into three parts. The first part adds game theoretical aspects to fundamental online problems. The second part involves oblivious algorithms that are motivated by designing prompt mechanisms for online bounded capacity auctions. The third part covers various subjects in online packing and covering problems.

Online algorithms deal with making irrevocable decisions while handling a sequence of events. In our scenario, the events are strategic in nature and have a private cost function, and seek to maximize their utility, i.e. minimize their private cost incurred by making a decision plus the surcharge posted on the decision by our dynamic pricing scheme. An example of this is the "parking problem" where an online sequence of cars arrive in some metric space and need to park in a vacant parking spot. Online algorithms know the next car's destination and order it where to park, while in our setting the algorithm sets a surcharge for each parking place (without knowing the next car's destination) and defers the decision on where to park to the car itself. This scenario is natural for problems such as: k-server, online metric matching and metrical task systems. We achieve essentially the same approximation ratio (up to a constant) as the best known online algorithms for these problems.

A bounded capacity auction is a single-item periodic auction for bidders that arrive online, where the amount of participating bidders is bounded. The algorithm decides which agents will participate and the allocation and pricing rule. We show a reduction from a simple stochastic balls and bins game to this problem. Although the algorithm for the game is oblivious (i.e., it does not receive input), we devise a non-uniform randomized algorithm. We establish a lower bound of 1.5 and an upper bound of 1.55, which implies a 1.55 competitive ratio mechanism for this auction.

In online packing and covering problems we establish almost tight lower and upper bounds for packing multidimensional vectors into bins. In this work we give almost tight bounds on the number of bins where the competitive ratio depends on the number of dimensions and the ratio between the maximum coordinate to the bin size. Additionally, we have worked on online covering with convex objective functions, including application such as unrelated machine scheduling with startup costs.

Representative Papers.

[1] The Loss of Serving in the Dark (STOC 2013)
with Y. Azar and I. Gamzu

[2] Tight Bounds for Online Vector Bin Packing (STOC 2013)
with Y. Azar, S. Kamara, B. Shepherd

[3] Pricing Online Decisions: Beyond Auctions (SODA 2015)
with A. Eden, A. Fiat and L. Jez

JOHN P. DICKERSON

Thesis. A Unified Approach to Real-World Dynamic Matching and Barter Exchange

Advisor. Tuomas Sandholm, Carnegie Mellon University

Brief Biography. John is a Ph.D. candidate in the Computer Science Department at Carnegie Mellon University, where he works in the Electronic Marketplaces Lab with his advisor Tuomas Sandholm. John's research is at the intersection of computer science and economics, with a focus on solving practical economic problems using stochastic optimization. He has worked extensively on theoretical and empirical approaches to kidney exchange, where his work has set policy at the UNOS nationwide exchange; game-theoretic approaches to counter-terrorism, where his models have been deployed; and computational advertising through Optimized Markets, a CMU spin-off company. With Tuomas Sandholm, he created FutureMatch, a general framework for learning to match subject to human value judgments. FutureMatch won a 2014 HPCWire Supercomputing Award and now provides sensitivity analysis for matching policies at UNOS. He is the winner of a 2012–2015 NDSEG Fellowship and a 2015–2017 Facebook Fellowship.

Research Summary. The exchange of indivisible goods without money addresses a variety of constrained markets where a medium of exchange—such as money—is considered inappropriate. Participants are either matched directly with another participant or, in more complex domains, in barter cycles and chains with many other participants before exchanging their endowed goods. My thesis research addresses the design, analysis, and real-world fielding of dynamic matching markets and barter exchanges.

Specifically, I study competing dimensions found in both matching markets and barter exchange, such as uncertainty over the existence of possible trades, trade-offs between efficiency and fairness, and inherent market dynamism. For each individual dimension, I provide new theoretical insights as to the effect on market efficiency and match composition of clearing markets under models that explicitly consider those dimensions. I support each theoretical construct with new optimization models and techniques that focus on scalability and practical applicability. In the cases of uncertain trades and dynamic matching, where participants and potential trades arrive and depart over time, my algorithms perform substantially better than the status quo deterministic myopic matching algorithms used in practice, and also scale to larger instance sizes than prior methods. In the fairness case, I quantify the loss in system efficiency under a variety of equitable matching rules.

I address each dimension in "FutureMatch," a framework for learning to match in a general dynamic model. It takes as input a high-level objective decided on by experts, then automatically (i) learns based on data how to make this objective concrete and (ii) learns the "means" to accomplish its goal—a task that humans handle poorly. FutureMatch now provides sensitivity analysis for matching policies at the UNOS nationwide kidney exchange.

My research shows that taking a holistic approach to balancing efficiency and fairness can often practically circumvent negative theoretical results. The balance is struck computationally via extensive optimization of realistic stochastic models of markets. Yet, theory lends necessary intuition to modeling decisions and validity to optimization techniques. Moving forward, I will continue to produce new theoretical results and optimization algorithms in support of market clearing frameworks that adhere to reality, with an eye toward fielding these new mechanisms.

Representative Papers.

[1] Failure-Aware Kidney Exchange (EC 2013)
with A.D. Procaccia and T. Sandholm

[2] FutureMatch: Combining Human Value Judgments and Machine Learning to Match in Dynamic Environments (AAAI 2015) with T. Sandholm

[3] Ignorance is Almost Bliss: Near-Optimal Stochastic Matching With Few Queries (EC 2015) with A. Blum, N. Haghtalab, A.D. Procaccia, T. Sandholm, and A. Sharma

ARIS FILOS-RATSIKAS

Thesis. Social Welfare in Algorithmic Mechanism Design Without Money

Advisor. Peter Bro Miltersen, Aarhus University

Brief Biography. I am a PhD student at Aarhus University, Denmark, working on algorithmic mechanism design without money, specifically social welfare maximization in social choice and resource allocation problems, as well as fair division and markets, among others. My supervisor is Prof. Peter Bro Miltersen and I am expected to obtain my PhD degree in August 2015. I obtained my MSc degree in Computer Science from the University of Patras, Greece, under the supervision of Prof. Ioannis Caragiannis. The topic of my thesis was Algorithmic Mechanism design without money, with emphasis on kidney exchange problems. I obtained my undergraduate degree (a 5-year degree, equivalent to BSc and MSc) from the University of Patras, under the supervision of Prof. Ioannis Caragiannis. My undergraduate thesis was on algorithmic game theory, specifically congestion games and coordination mechanisms.

Research Summary. My main area of research is Algorithmic Game Theory and Algorithmic Mechanism Design. These fields lie in the intersection of Theoretical Computer Science and Economics and study the effects of strategic behavior of the participants in the performance of a system or an algorithm.

The main topic of my PhD thesis is approximately maximizing social welfare in general social choice and resource allocation problems without money, when the participants have unrestricted cardinal preferences over the set of outcomes. The goal is, for the different problems that I consider, to come up with truthful mechanisms with good approximation ratios or mechanisms with good Price of Anarchy guarantees. In the general social choice setting with unrestricted preferences, truthful deterministic mechanisms are severely limited by strong impossibility results. On the other hand, randomized mechanisms are possible and they do provide better approximation ratio guarantees. Comparing mechanisms and trying to come up with the best one is a topic of my thesis. A similar approach is adopted for the setting of one-sided matchings without money, where I study well-known mechanisms and prove their asymptotic optimality in terms of social welfare.

I have also been working on markets, studying both incentive properties and approximate fair solutions, for the fundamental Fisher market model, as well as markets with indivisibilities and subsets of Arrow-Debreu markets with quasilinear preferences. The subject of this work is to prove existence of pure Nash equilibria for these market models or come up with approximately clearing conditions that produce envy-free outcomes.

Another topic that I am interested in is fair division in general, in terms of incentives and economic properties as well as computational issues. I am interested in studying different well-known fair division protocols (such as the Adjusted Winner procedure) from a strategic point of view; proving equilibrium existence, Pareto-efficiency and social welfare guarantees in equilibrium. On the computational side, I have been working on fair division problems such as consensus halving or markets with indivisible items, trying to come up with hardness proofs and approximation algorithms for different goals.

I am also working on other subfields of algorithmic mechanism design, such as facility location problems, kidney exchange markets and structured assignment settings.

Representative Papers.

[1] Truthful Approximations to Range Voting (WINE 2014) with P.B. Miltersen

[2] Social Welfare in One-Sided Matchings: Random Priority and Beyond (SAGT 2014) with S.K.S. Frederiksen and J. Zhang

[3] The Adjusted Winner procedure: Characterizations and Equilibria (IJCAI 2015) with H. Aziz, S. Branzei, and S.K.S. Frederiksen

SAM GANZFRIED

Thesis. Computing Strong Game-Theoretic Strategies and Exploiting Suboptimal Opponents in Large Games

Advisor. Tuomas Sandholm, Carnegie Mellon University

Brief Biography. Sam received a PhD in computer science from Carnegie Mellon University in 2015 for his dissertation Computing Strong Game-Theoretic Strategies and Exploiting Suboptimal Opponents in Large Games and holds an A.B. in math from Harvard. His research interests include artificial intelligence, game theory, multiagent systems, multiagent learning, large-scale optimization, large-scale data analysis and analytics, and knowledge representation. He created two-player no-limit Texas hold em agent Tartanian7 that won the 2014 Annual Computer Poker Competition and Claudico that competed in the inaugural 2015 Brains vs. Artificial Intelligence competition against the strongest human specialists in the world for that poker variant: the humans won the latter by a margin that was statistically significant at the 90% level but not at the 95% level, and many exciting lessons were learned. He organized the AAAI Workshop on Computer Poker and Imperfect Information in 2014 and 2015.

Research Summary. Important problems in nearly all disciplines and on nearly all application domains involve multiple agents behaving strategically; for example, deploying officers to protect ports, determining optimal thresholds to protect against phishing attacks, and finding robust policies for diabetes management. Such problems are modeled under the framework of game theory. In many important games there is information that is private to only some agents and not available to other agents – for instance, in auctions each bidder may know his own valuation and only know the distribution from which other agents' valuations are drawn.

My research designs new approaches for strategic agents acting in large imperfect-information games. It includes novel algorithms, theoretical analysis, and large-scale implementation.

There are several major challenges that must be confronted when designing successful agents for large multiagent strategic environments. First, standard solution concepts such as Nash equilibrium lack theoretical justification in certain classes (e.g., games with more than two players). Second, computing these concepts is difficult in certain classes from a complexity-theoretic perspective. Third, computing these concepts is difficult in practice for many important games even for cases when they are well-motivated and polynomial-time algorithms exist (e.g., two-player zero-sum (competitive) games), due to enormous state spaces. And fourth, for all game classes, it is not clear if the goal should even be to compute a Nash equilibrium; one could achieve significantly higher payoff by learning to exploit opponents' mistakes. However, such exploitation must be done in a way that does not open oneself up to being exploited in turn by strong deceptive opponents.

While the approaches are domain independent, most of them have been motivated by and applied to the domain of poker. Poker has emerged as a major AI challenge problem. Poker is not simply a toy game; it is tremendously popular for humans, and online poker is a multi-billion dollar industry. For the past ten years, there has been a competition between the strongest computer poker agents held annually at the top AI conference. The version of two-player no-limit Texas hold 'em played has approximately 10^{165} states in its game tree. Several of the techniques I developed were utilized to create agents that won the 2014 competition and that competed against the strongest human specialists in 2015.

Representative Papers.

[1] Endgame Solving in Large Imperfect-Information Games (AAMAS 2015) with T. Sandholm

[2] Safe Opponent Exploitation (TEAC 2015 and EC 2012) with T. Sandholm.

[3] Action Translation in Extensive-Form Games with Large Action Spaces: Axioms, Paradoxes, and the Pseudo-Harmonic Mapping (IJCAI 2013) with T. Sandholm

XI (ALICE) GAO

Thesis. Eliciting and Aggregating Truthful and Noisy Information

Advisor. Yiling Chen, Harvard University

Brief Biography. Xi (Alice) Gao is currently a postdoctoral research fellow in Computer Science at University of British Columbia, where she holds a prestigious Canadian NSERC Postdoctoral Fellowship. Alice's research tackles problems at the intersection of artificial intelligence, game theory, and crowdsourcing, using a mix of theoretical and experimental methods. Alice obtained her PhD in Computer Science from Harvard University in 2014. Her PhD dissertation received the 2014 Victor Lesser Distinguished Dissertation Runner-up Award and was also selected for Honorable Mention for the 2015 SIGecom Doctoral Dissertation Award. Her PhD research was supported by a Canadian NSERC Postgraduate Scholarship for Doctoral Students and she was named a 2014 Siebel Scholar. Previously, she earned her Bachelor's degree in Computer Science and Mathematics from University of British Columbia.

Research Summary. My research is in algorithmic game theory and broadly at the intersection of artificial intelligence, game theory, and crowdsourcing. I am driven by the desire to understand the strategic interactions of self-interested participants in complex systems and I aim to better design these systems to achieve desirable outcomes. In pursuing these goals, I draw insights from various disciplines such as artificial intelligence, game theory, statistics, etc. Moreover, I enjoy tackling problems using a mix of theoretical analyses and experimental studies.

My dissertation research focuses on developing and analyzing methods for eliciting and aggregating dispersed information. I have addressed a number of problems including eliciting truthful estimates of uncertain events using prediction markets, eliciting truthful evaluations of products and services using peer prediction methods, and ranking multiple alternatives by adaptively eliciting and aggregating noisy information. Currently, I am investigating ways of eliminating collusion in peer prediction mechanisms by using limited, costly access to ground truth provided by trusted evaluators.

Representative Papers.

[1] Trick or Treat: Putting Peer Prediction to the Test (EC 2014) with A. Mao, Y. Chen, and R.P. Adams

[2] Market Manipulation with Outside Incentives (JAAMAS 2015) with Y. Chen, R. Goldstein, and I.A. Kash

[3] Adaptive Polling for Information Aggregation (AAAI 2012)
with T. Pfeiffer, A. Mao, Y. Chen, and D.G. Rand

NIKOLAI GRAVIN

Thesis. Incentive Compatible Design of Reverse Auctions

Advisor. Dmitrii Pasechnik, University of Oxford

Brief Biography. Nick finished graduate school at Saint-Petersburg department of Steklov Mathematical Institute in Russia in 2010. At the same time he was a PhD student at the mathematical department of Nanyang Technological University in Singapore, which he finished in 2012. His research interests are twofold. In Mathematics he has been working in graph theory, convex and discrete geometry. In Theoretical Computer Science he is particularly interested in Algorithmic Mechanism Design and Equilibria computations. Nick is a recipient of a prestigious Microsoft Research Fellowship awarded to the top students in Asia.

Research Summary. My research lies in the areas of algorithmic mechanism design and game theory, with connections to on-line algorithms and learning theory. It involves the design and analysis of approximation algorithms for a variety of optimization problems. In my work, I often ask the following questions: Which metrics should be used to quantify performance and efficiency of an algorithm or mechanism in an economic setting?

Digital goods auctions. My work on digital goods auctions illustrates the importance of these questions. In this setting a monopolistic seller seeks to maximize profit from the sale of a single good available in unlimited supply. Digital goods with negligible costs for duplication and distribution such as pay-per-view television and downloadable audio files make a perfect example. [GHW 01] initiated the worst-case analysis of this problem in the mechanism design framework. This and many subsequent work study pricing mechanisms in the form of single-round, sealed-bid truthful auctions for selling digital good. Our characterization of the extremal distributions for the class of all monotone benchmarks provided a missing tool to derive tight worst-case results for a big family of meaningful benchmarks.

Simple Mechanisms for Complex Markets. This line of my work focuses on Bayesian framework allowing to circumvent undesirable computational hardness and performance gaps of the worst-case analysis. To study the large-scale, computer-aided combinatorial markets which are becoming a reality, with examples of FCC spectrum auction and internet-powered marketplaces like Ebay the CS community has generated a subfield of work on developing efficient algorithms and incentive compatible mechanism for combinatorial allocation problems. In the model of *combinatorial auction*, there is a large set M of m objects for sale, and n potential buyers. Each buyer has a private value function $v_i \colon 2^M \to \mathbb{R}_{\geq 0}$ mapping sets of objects to their associated values. The goal of the market designer is to decide how to allocate the objects among the buyers to maximize the social efficiency. In [3], [2] we study two simple and practical mechanisms in a Baysian setting: item-bidding simultaneous auctions that achieve nearly optimal social welfare at any Bayesian Nash equilibrium; posted price mechanism yielding the first Bayesian polytime constant-approximation truthful mechanism for submodular buyers.

Representative Papers.

[1] Combinatorial Auctions via Posted Prices (SODA 2015)
with M. Feldman and B. Lucier

[2] Optimal Competitive Auctions (STOC 2014)
with N. Chen and P. Lu

[3] Simultaneous Auctions are (Almost) Efficient (STOC 2013)
with M. Feldman, H. Fu, and B. Lucier

NIMA HAGHPANAH

Thesis. Optimal Multi-parameter Auction Design

Advisor. Jason Hartline, Northwestern University

Brief Biography. Nima Haghpanah is a postdoctoral associate at MIT CSAIL and Sloan School of Management, working with Costis Daskalakis and Itai Ashlagi. He obtained his PhD in 2014 from Northwestern University advised by Jason Hartline. His dissertation was on optimal multi-parameter auctions, and his research interests are mechanism design, pricing, algorithms, and optimization. Nima was awarded the Simon's Award for Graduate Students in Theoretical Computer Science 2013-2015, and Yahoo! Key Scientific Challenges in 2012.

Research Summary. Mathematical models are often designed to understand simples rules of nature. The design of real world systems, however, requires models encompassing high levels of detail, making them more complex and harder to analyze. I study these complex design problems, focusing on the design of mechanisms and markets, using tools and concepts from microeconomics analysis, algorithms, and optimization.

Bayesian mechanism design studies the design of systems to optimize revenue when the information is privately held by strategic agents. Myerson's (1981) seminal result studied this problem when there is a single item to sell. The result is celebrated because it identified optimal mechanisms that are practical and commonly used, and also proposed a universal analysis method based on virtual valuations. A fundamental assumption in Myerson's analysis was single-dimensionality: there is only one way to serve an agent, and the agent's utility is value minus payment. In my PhD thesis I studied a class of problems with wide variety of applications, in which there are multiple ways to serve an agent (e.g., the item may have different qualities or delivery times), and agents have general non-linear or multi-parameter preferences.

First, I extended the method of virtual values to analyze auctions in multi-parameter settings. I showed that virtual values that prove the optimality of a given auction exist, if the space of agent's possible preferences can be integrated in a consistent manner. Formalizing this intuition gives a system of differential equations, the solution to which would imply the optimality of the give mechanism. I applied this framework to recover conditions implying optimality of several natural mechanisms. Second, I showed that under general conditions, revenue optimization is governed by an intuitive and well-known principle: maximizing marginal revenue. Third, I showed that even without analytical understanding, optimal multi-agent mechanisms can be efficiently optimized.

In another line of research, I study mechanisms and pricing strategies in presence of network or allocation externalities. I showed that in single-parameter settings, the standard approach of reducing the problem to a worst-case optimization of virtual values can not be used to solve the problem. Instead, I reduced the problem to a novel average-case problem and used it to derive constant factor approximations to these problems.

Representative Papers.

[1] Reverse Mechanism Design, to appear (EC 2015)
with J. Hartline

[2] The Simple Economics of Approximately Optimal Auctions (FOCS 2013)
with S. Alaei, H. Fu, and J. Hartline

[3] Optimal Auctions with Positive Network Externalities (EC 2011)
with N. Immorlica, V. Mirrokni, and K. Munagala

DARRELL HOY

Thesis. Strategic Computation via Non-Truthful Mechanism Design

Advisor. Jason Hartline, Northwestern University

Brief Biography. Darrell is a Ph.D. candidate in the Theory and Economics group of the EECS department at Northwestern University, and has been in the Boston area visiting the Harvard EconCS group since 2014. Darrell received his undergraduate degree from Dartmouth College, where he greatly enjoyed a liberal arts education, racing for the cycling team as well as beginning to experiment with research. Between Dartmouth and Northwestern, Darrell worked in finance for Bridgewater Associates and launched a website to help cyclists find good roads to ride, sweetopenroads.com, and can be found searching for good roads and trails wherever they may be.

Research Summary. Sometimes the output of a computation affects the input. The way Google runs their page-ranking algorithm affects how site designers build websites. When Yelp asks for reviews, restaurants can change how they serve and incentivize customers to leave reviews. In such settings, a computation is more complicated than a collection of circuits: it involves all of the agents making strategic decisions in accordance with their own incentives rather than their instructions. I call these strategic computations, and I want to understand them. I'm working to push non-revelation mechanism design to be more useful as a general model of these strategic computations. I am particularly focused on developing analytical tools that are robust to a) the details of equilibrium, and b) robust to changes in the decision making frameworks of users, for instance, bidders who are risk-averse.

On the first front, my recent work with Sam Taggart and Jason Hartline [1] refined the smooth games and mechanisms approach for the objective of revenue in auctions. In a first-price auction with a few light regularity assumptions, we found implementing a reserve price is sufficient to eliminate the impact of agents with misaligned incentives, even in asymmetric settings where we do not have an analytical characterization of equilibrium.

Our current theories of mechanism design rely strongly on precise assumptions of the decision-making behavior of the agents: that they are risk-neutral and always choose their optimal action. I am working in this light to make our understanding of mechanisms more robust to the exact risk-attitudes and exact decision-making behavior of strategic agents. In work with Hu Fu and Jason Hartline [2], we showed that the first-price auction is approximately optimal when bidders exhibit a specific type of risk-attitude as well as when they are risk averse. I am very excited to continue pushing in these directions to broader notions of risk-aversion and other behavioral complications.

By pushing our understanding of non-revelation mechanism design to be more robust to the details of equilibrium and decision making behavior, I plan to push forward our understanding of how to think about and design general computations with strategic agents.

Representative Papers.

[1] Price of Anarchy for Auction Revenue (EC 2014)
with J. Hartline and S. Taggart

[2] Prior-independent auctions for risk-averse agents (EC 2013)
with H. Fu and J. Hartline

[3] A dynamic axiomatic approach to first-price auctions (EC 2013)
with K. Jain and C. Wilkens

MOHAMMAD REZA KHANI

Thesis. Revenue Efficient Mechanisms for Online Advertising

Advisor. Mohammad T. Hajiaghayi, University of Maryland

Brief Biography. Reza is a fourth year Ph.D. student in Department of Computer Science at University of Maryland working under supervision of Mohammad T. HajiAghayi. He got his M.Sc. degree from University of Alberta working with Mohammad R. Salavatipour. He did his undergraduate studies in computer engineering at Amirkabir University of Technology.

Research Summary. Generalized Second Price (GSP) auction (the current mechanism of choice in online advertising) has appealing properties when ads are simple (text based and identical in size). But GSP does not generalize to richer ad settings, whereas truthful mechanisms, such as VCG do. A straight switch from GSP to VCG either requires all bidders instantly bid truthfully or incurs significant revenue loss. We propose a transitional mechanism which encourages advertisers to update their bids to their valuations, while mitigating revenue loss. In settings where both GSP ads and truthful ads exist, it is easier to propose a payment function than an allocation function. We give a general framework for these settings to characterize payment functions which guarantee incentive compatibility of truthful ads, by requiring that the payment functions satisfy two properties.

Next, we study revenue monotonicity (revenue should go up as the number of bidders increases) of truthful mechanisms in online advertising. This natural property comes at the expense of social welfare - one can show that it is not possible to get truthfulness, revenue monotonicity, and optimal social welfare simultaneously. In light of this, we introduce the notion of Price of Revenue Monotonicity (PoRM) to capture the loss in social welfare of a revenue monotone mechanism. We design truthful and revenue monotone mechanisms for important online advertising auctions with small PoRM and prove a matching lower bound.

Finally, we study how to measure revenue of mechanisms in the prior free settings. One of the major drawbacks of the celebrated VCG auction is its low (or zero) revenue even when the agents have high values for the goods and a competitive outcome would have generated a significant revenue. A competitive outcome is one for which it is impossible for the seller and a subset of buyers to 'block' the auction by defecting and negotiating an outcome with higher payoffs for themselves. This corresponds to the well-known concept of core in cooperative game theory. We define a notion of core-competitive auctions. We say that an incentive-compatible auction is α-core-competitive if its revenue is at least $1/\alpha$ fraction of the minimum revenue of a core-outcome. We study designing core-competitive mechanisms for a famous online advertising scenario.

Representative Papers.

[1] Core-competitive Auctions (EC 2015)
with G. Goel and R. Paes Leme.

[2] Mechanism Design for Mixed Participants (to be submitted)
with Y. Bachrach, S. Ceppi, I.A. Kash, and P. Key

[3] Revenue Monotone Mechanisms for Online Advertising (WWW 2014)
with G. Goel

ELI MEIROM

Thesis. Games and Dynamics in Large Communication Networks

Advisor. Ariel Orda and Shie Mannor, Technion

Brief Biography. Eli Meirom received a B.Sc. degree in Math and Physics from the Technion, Israel (summa cum lauda), and an M.A in Physics (magna cum Lauda). He is currently pursuing a Ph.D degree in Electrical Engineering at the Technion, Israel. Previously, he held research positions in IBM Research and St. Jude Medical. He published papers in various fields, including quantum information, solid state physics, machine learning, game theory and computer networks. He was awarded the Applied Materials and Mel Berlin fellowships. His current research interests are in the intersection of machine learning and social networks, e.g., social network analysis and multi-agent dynamics on graphs.

Research Summary. The interplay between the network topology, and the network performance is critical in all networks. Additionally, a network must operate in secure and reliable fashion, in order to perform its function. We ask: How does the behavior, and performance requirements of the users, affect the network structure? What is the effect of reliability requirements on the network topology? In computer networks, can we design generic anti-malware measures that are based only on network structure and properties, rather than malware attributes?

Most of the studies in Network Formation Games (NFG) assume identical players, whereas the Internet is composed of many types of ASs, such as minor ISPs, CDNs, tier-1 ASs etc. We constructed a heterogeneous network formation games, and analyzed both is static and dynamic properties. Furthermore, game theoretic analysis is rarely confronted with real-world data. We took a step further, and considered real inter-AS topology data. Our model and its analysis resulted in some novel predictions regarding the evolution of the inter-AS topology.

In many NFG the resulting networks are very fragile. In a later work, we established a model that explicitly includes the agents' reliability requirements. We provide dynamical analysis of topological quantities, and explain the prevalence of some network motifs, i.e., sub-graphs that appear frequently in the Internet.

Attacks, from denial of service to state-driven cyberwarfare threaten to compromise our modern infrastructure. Anti-virus software can find the signature of known worm or virus. But how do we intercept malware spread in actionable time-scales, before we even know what is spreading?

When people contract a virus they might miss a work day, though of course many other factors could produce such behavior. Malware, similarly, produces slight deviations in system behavior, for example, a spike in network activity. Can we use indications of abnormality that are so weak that on their own they are statistically indistinguishable from noise, to make an accurate global diagnosis about a spreading contagion? We have shown that, given a map of nodes that experience suspicious behavior, we were able to identify a malware outbreak. We addressed this problem from a dynamic perspective as well, and showed that monitoring the dynamics of these weak signatures enables early detection, very shortly after the initial infiltration occurred.

Representative Papers.

[1] Network Formation Games with Heterogeneous Players and the Internet Structure (EC 2014) with S. Mannor and A. Orda

[2] Formation Games of Reliable Networks (INFOCOM 2015)
with S. Mannor and A. Orda

[3] Localized Epidemic Detection in Networks with Overwhelming Noise (SIGMETRICS 2015) with C. Milling, S. Shakkottai, C. Caramanis, S. Mannor, and A. Orda

ILAN NEHAMA

Thesis. Computational issues in Judgement Aggregation

Advisor. Noam Nisan, Hebrew University of Jerusalem

Brief Biography. Ilan received his B.A. in Math and B.Sc. in Computer Science at the Technion (Summa Cum Laude), and M.A. in Computer Science with specialization in Rationality at The Hebrew University under the supervision of Prof. Gil Kalai; Thesis: Implementing Social Choice Correspondences using k-Strong Nash Equilibrium (Summa Cum Laude, GPA: 97.77, 2/194). Ilan is a Ph.D. candidate at The Hebrew University (Benin School of Computer Science & Federmann Center for the Study of Rationality) and is expected to finish on 2016. During his graduate studies, he served as a lecturer in a Programming course, as well as a TA in several MA courses both in the Computer Science department - Mathematical Tools in CS, and the Economics department - Microeconomics A & Microeconomics B: Game Theory and Information Economics. Ilan's works are mainly in theoretical Game Theory, Social Choice, and Judgement Aggregation, and on computational aspects and the usage of methods from Computer Science in these fields.

Research Summary. I am interested in a Computer Science approach to questions in Social Choice and specifically to Judgement Aggregation. Judgement Aggregation (JA) investigates which procedures a group could or should use to form collective judgements on a given set of propositions or issues, based on the judgements of the group members. Judgement Aggregation is the subject of a growing body of works in economics, computer science, political science, philosophy, law, and other related disciplines. I find this field highly applicable to agent systems, voting protocols in a network and other frameworks in which one needs to aggregate a lot of opinions in a systematic way without letting the voters deliberate or without assuming a deliberation process.

I'm interested in shedding light on phenomena in Judgement Aggregation using approximation and perturbation viewpoints. That is, studying the way phenomena studied in the literature change when perturbing the classic strict properties. E.g., requiring an aggregation rule to satisfy a property with high probability (but not for sure), generalizing players' rationality constraints to being close to rational (bounded rationality). Dealing with probabilistic properties raises the question of choosing the 'right' underlying distribution or distributions family. It is clear that real-life distributions are not uniform, and indeed most of the current works that analyze rules using simulations, check non-uniform distributions. Nevertheless, there are very few analytical works dealing with such distributions. As part of the above perturbation paradigm, I study other natural distributions (e.g., Polya-Eggenberger models), and aim to extend works that study uniform distributions to analysis under other distributions families.

For example, I studied the perturbations of the 'Doctrinal Paradox' scenarios. In these scenarios, one looks for 'consistent' and 'independent' aggregation mechanism for a given agenda (=set of permissible opinions). In 'Approximately Classic Judgement Aggregation' I presented the relaxation where these two desired properties hold only with high probability, and showed that under uniform distribution, for conjunction and xor agendas, there is no non-trivial mechanisms that satisfy these perturbed constraints. In subsequent works, I show similar results for non-silencing agendas, which are most of the truth-functional agendas, and for non-uniform distributions, although still having independent representation.

Representative Papers.

[1] Approximately Classic Judgement Aggregation (AMAI 2013 & WINE 2011)

[2] Mechanism design on discrete lines and cycles (EC 2012)
with E. Dokow, M. Feldman, and R. Meir

[3] Complexity of Optimal Lobbying in Threshold Aggregation (AAMAS 2013 & ADT 2015)

SVETLANA OBRAZTSOVA

Thesis. Essays on the Complexity of Voting Manipulation

Advisor. Edith Elkind, University of Oxford

Brief Biography. Currently a postdoctoral fellow at Israeli Centre of Research Excellence (I-CORE) in Algorithms. My first postdoctoral appointment was with the CoreLab, NTUA, Greece. Prior to that, a dual, 4-year PhD program of the Steklov

Institute of Mathematics (St.Petersburg, Russia) and NTU (Singapore). My PhDs are in Mathematics and were received in June 2011 (from the Steklov Institute) and in October 2012 (from NTU). My work has been accepted and acknowledged by the research community. The collaboration with Prof. Elkind was nominated for the Best Paper Award at The 10th AAMAS Conference (4 nominees from 575 submissions). Same paper was among the three works selected to represent AAMAS-2011 at IJCAI-2011 (the Best Papers from Sister Conferences Track). Another joint work has received the Pragnesh Jay Modi Best Student Paper Award at AAMAS-2012, and also served as a representative work at AAAI-2012. My work on graph theory has won Google Europe Anita Borg Memorial Scholarship in 2008.

Research Summary. My research interests fall into the realm of Computational Economics or, more accurately, the Computational Social Choice.

Initially, during my PhD studies, I'd concentrated on the voting manipulation complexity, publishing several joint papers with Edith Elkind on the influence of the tie-breaking rule on the said complexity, and the complexity of optimal manipulations.

However, accepting manipulation attempts as a given, I've quickly moved to the study of resulting Nash Equilibria (NEs), and general stable states. First, I've contributed to the study of biased voters, as means to reduce the set of NEs. E.g., I've investigated the effects of the truth and lazy biases. Second, I took interest in iterative voting, which lead to several results on its stability conditions and stable state characteristics.Taking the lead on the iterative voting processes, I've also researched Iterative Candidacy Games.

These results have proved of a great interest to the community in general, and a well grouped set of them was organised into the "Voting and Candidacy Games" Tutorial at AAMAS 2015.

There are several research questions that I continue to address, both building on my previous success and establishing new research directions. The first natural direction is to further exploit the idea of rewarding truthfulness, extending it to other voting rules. It is also possible to additionally enrich the truth-biased voting model by introducing the concept of "lie degree", based on the distance between the submitted vote and the truthful preference. Thus making the current model a binary sub-case. A variety of distance measures can be used here, which suggests a rich ground for new ideas and publications.

Another research direction in my immediate plans is to further develop the concept of candidates' game. In spite of recent contributions, including my own, it remains one of the least studied and, yet, most intriguing research prospects in Social Choice. Current results are only partial and specialised, and are in dire need for generalisation across voting rules and mutual voter-candidate choice dynamics.

Finally, a budding research direction is the link between voting manipulations and preference elicitation procedures, as expressed by the dynamics of these these processes. Inevitably with great success, the first set of results was presented with my co-authors to limited audiences, building up to their wider acceptance at IJCAI-2015.

Representative Papers.

[1] Ties matter: complexity of voting manipulation revisited (AAMAS 2011)

with E. Elkind and N. Hazon

[2] Optimal manipulation of voting rules (AAMAS 2012)
with E. Elkind

[3] Plurality voting with truth-biased agents (SAGT 2013)
with E. Markakis and D. Thompson

JOEL OREN

Thesis. Multi-Winner Social Choice: Algorithmic and Game-Theoretic Approaches

Advisor. Allan Borodin and Craig Boutilier, University of Toronto

Brief Biography. Joel Oren is a Ph.D student at the Department of Computer Science, University of Toronto, Canada. He is supervised by Allan Borodin and Craig Boutilier, and will be graduating in August, 2015. He received his M.Sc in Computer Science, at the University Toronto under supervision of Allan Borodin. He received his B.Sc in Computer Science with honors from Ben Gurion University, Israel, where he was supervised by Avraham Melkman. He received the Ontario Graduate Scholarship award in 2010, 2011, and 2013. He has also received the Raymond Reiter award by the department of Computer Science, University of Toronto, for the year 2015.

Research Summary. I am interested in the study of influence diffusion in social networks. In particular, I focus on the game theoretic study of games and mechanisms for competing diffusion processes. Moreover, I take an active part in research on the design of efficient and parallelizable algorithms for large scale social networks, for addressing problems such as cascade detection, and fitting of influence diffusion model parameters. My goal is to reason about the spread processes that are captured by real-world datasets.

During my PhD studies, my research had two recurring themes. The first major theme of research is the study of stochastic processes in social networks. I performed research in approximation algorithms and mechanism design for competitive influence diffusion. I am presently engaged in applying current algorithmic paradigms for parallel computing (e.g., MapReduce) for data-mining tasks related to social networks, such as influence estimation and maximization, structural analysis, etc. With my collaborators, I have taken part in the design of highly-parallelizable algorithms for estimating influence diffusion processes in massive social networks, with provable performance guarantees.

The second area that I have contributed to deals with applying probabilistic models to decision-making problems at the core of computational social choice. With my collaborators, I have studied slate optimization and efficient preference elicitation, while crucially building on belief models of agent preferences. We have revisited existing problems in these regimes through a probabilistic lens, and offered a rigorous study, which offered great insight, that goes beyond existing worst-case models. Our work includes an analysis of the capabilities of top-k voting for preference elicitation, and online algorithms for slate optimization.

Representative Papers.

[1] Strategyproof Mechanisms for Competitive Influence in Networks (WWW 2013)
with A. Borodin, M. Braverman and B. Lucier

[2] Influence at Scale: Distributed Computation of Complex Contagion in Networks (KDD 2015) with B. Lucier and Y. Singer

[3] Efficient Voting via the Top-k Elicitation Scheme: a Probabilistic Approach (EC 2014) with Y. Filmus

EMMANOUIL POUNTOURAKIS

Thesis. Simple Mechanisms in Static and Dynamic Settings

Advisor. Nicole Immorlica, Microsoft Research

Brief Biography. Emmanouil Pountourakis is currently a fifth year PhD student in the Department of Electrical Engineering and Computer Science at Northwestern University, advised by Nicole Immorlica. Since 2014 he has been a long term visitor at Microsoft Research, New England. He holds an undergraduate and masters degree in Computer Science from the University of Athens. During Summer 2012 he completed a research internship at CWI, Amsterdam. He was a student visitor in the Institute of Advanced Studies at Hebrew University for the Special Semester in Algorithmic Game Theory in 2011. Emmanouil Pountourakis has a broad interest in algorithmic mechanism design. His current research focuses on revenue maximization in static and dynamic environments. In the past he has worked on a variety of topics including cost-sharing, matching, and mechanism design without money.

Research Summary. Most of my research lies in algorithmic mechanism design and lately focuses on revenue maximization in static and dynamic settings. I am particularly interested in the study of anonymous pricing. For several years this widely studied mechanism was known to have a revenue approximation ratio within [2,4] for selling a single item. My recent work gives the first improvement of this gap to $[2.23, e]$. Furthermore, I am currently studying anonymous pricing in the dynamic environment of repeated sales of a single item. This setting gives rise to various strategic behaviors with negative implications for the revenue. Surprisingly, anonymous pricing suffers less than its discriminatory counterpart and may provide higher revenue. My on-going work further investigates this phenomenon. Also, I am interested in the interaction of revenue maximization and different behavioral models. My recent work studies optimal contract design with a present-biased agent. When making a decision, the present-biased agent overestimates her present utility by a multiplicative factor. The contract designer exploits this behavior to maximize his revenue. My work introduces regulations and studies optimal contracts under them in the hopes of reducing the exploitative power of the mechanism.

Earlier problems I've worked on include cost-sharing, matching, and mechanism design without money. My contributions are briefly outlined below. A central challenge in cost sharing is the design of group-strategyproof mechanisms, that is mechanisms that are resilient to group manipulation. My work gives a complete characterization of group-strategyproof cost sharing mechanisms and studies their performance in terms of budget balance, that is what fraction of the cost can be covered using payments assigned to the agents. In another work I generalize the stable-marriage problem. Agents have limited information about their peers illustrated by a social graph. This relaxes the stability constraint and gives rise to the optimization problem of finding a maximum size stable matching. My work gives an approximation algorithm with matching lower bound. Finally, I studied the problem of task allocation without payments using the assumption of binding reports: the agent incurs a cost which is the maximum of the cost they reported and their actual cost. My work studies the problem of choosing a base of a matroid, e.g. spanning tree, and designs optimal truthful approximate mechanisms under this assumption.

Representative Papers.

[1] Mechanisms for Hiring a Matroid Base without Money (SAGT 2014) with G. Schfer

[2] Socially Stable Matchings in the Hospitals/Residents Problem (WADS 2013) with G. Askalidis, N. Immorlica, A. Kwanashie, and D. Manlove

[3] On Budget-Balanced Group-Strategyproof Cost-Sharing Mechanisms (WINE 2012) with N. Immorlica

BAHARAK RASTEGARI

Thesis. Stability in Markets with Power Asymmetry

Advisors. Kevin Leyton-Brown and Anne Condon, University of British Columbia

Brief Biography. Baharak Rastegari is a research associate at the School of Computing Science, University of Glasgow (UofG). She is currently working on an EPSRC project titled “Efficient Algorithms for Mechanism Design Without Monetary Transfer”, which is a joint project between the universities of Glasgow, Liverpool, and Oxford. Her main areas of interest are Game Theory and Bioinformatics. For the past five years she has been focused on solving problems concerning matching markets under preferences; before that, she worked on revenue properties of combinatorial auctions and the prediction of RNA secondary structures. She received her Ph.D. from the University of British Columbia (UBC), Canada, in 2013. She holds an M.Sc. in Computer Science (UBC 2004) and a B.Sc. in Computer Engineering (Sharif University of Technology, Iran, 2002). She enjoys teaching and has given several lectures in various courses, including Foundations of Multiagent Systems (UBC) and Algorithmics (UofG).

Research Summary. I enjoy solving mathematical problems and in particular designing efficient algorithms for various problems, or proving that none exists. My particular focus has been on two-sided matching markets, where agents have preferences over one another; these preferences might be partially known, and might not necessarily be strict. The goal is to compute matchings of the agents to one another that are optimal with respect to the given preferences. Optimality can refer to classical concepts such as stability (which ensures that no two agents have an incentive to form an arrangement outside of the matching) or Pareto optimality (which guarantees that no coalition of agents can improve without harming someone else). Applications arise in entry-level labor markets such as the allocation of junior doctors to hospitals.

In a paper in EC 2013 we considered two-sided matching markets with incomplete information — that is, agents' own preference profiles are only partially known and true preferences can be learned via so-called interviews. The goal was to identify a centralized interview policy, i.e. an algorithm that adaptively schedules interviews in order to produce a matching that is stable with respect to agents' true preferences, and that is furthermore optimal for one given side of the market. We showed that an interview-minimizing policy can be computed in exponential time and we gave evidence showing that it is likely that no polynomial-time computable policy exists unless P=NP. Additionally, we provided a polynomial time algorithm that identifies an interview-minimizing policy for a restricted setting inspired by real world applications.

In a paper in EC 2014 we studied truthful mechanisms for finding large Pareto optimal matchings in two-sided matching markets with one-sided preferences (the so-called "house allocation problem"). We provided a natural and explicit extension of the classical Random Serial Dictatorship Mechanism to the case where preferences may include ties. Consequently we obtained a universally truthful randomized mechanism for finding a Pareto optimal matching and showed that it achieves an approximation ratio of $\frac{e}{e-1}$.

In earlier work, appearing at AAAI 2007 and SODA 2009, we studied combinatorial auctions involving the so-called revenue monotonicity property (which guarantees that a seller's revenue weakly increases as the number of bidders grows), investigating the existence of truthful mechanisms for this restriction.

Representative Papers.

[1] Size versus Truthfulness in the House Allocation Problem (EC 2014) with P. Krysta, D. Manlove, and J. Zhang

[2] Two-sided Matching with Partial Information (EC 2013) with A. Condon, N. Immorlica, and K. Leyton-Brown

[3] Revenue Monotonicity in Deterministic, Dominant-Strategy Combinatorial Auctions (AIJ 2011) with A. Condon and K. Leyton-Brown

NISARG SHAH

Thesis. Economic Foundations for Practical Social Computing

Advisor. Ariel D. Procaccia, Carnegie Mellon University

Brief Biography. Nisarg Shah is a Ph.D. candidate in the Computer Science Department at Carnegie Mellon University, advised by Ariel Procaccia. His broad research agenda in algorithmic economics includes topics such as computational social choice, fair division, game theory (both cooperative and noncooperative), and prediction markets. He focuses on designing theoretically grounded methods that have practical implications. Shah is the winner of the 2013-2014 Hima and Jive Graduate Fellowship and the 2014-2015 Facebook Fellowship.

Research Summary. I am interested in the general problem of how to use inputs from multiple agents for computing a social outcome; examples include political elections, crowdsourcing, and multi-agent resource allocation. My thesis research investigates real-world social computing settings in which monetary exchange is prohibited, and uses theoretical insights to design well-founded solutions.

For example, in computational fair division, my latest project [1], in collaboration with a large California school district, deals with the design and implementation of a method for fairly allocating classrooms to charter schools. On the theoretical level, we show that our approach is provably fair and provides worst-case optimization guarantees. And from the practical viewpoint, a scalable implementation of our mechanism requires a number of innovations; its deployment is an ongoing, intricate project. I have also done significant work on the fair allocation of computational resources in clusters (see, e.g., [2]). In particular, we were among the first to study dynamic fair division in a setting where agents arrive over time, thereby pushing the conceptual limits of fair division theory itself.

As another example, my work in computational social choice focuses on settings where an objective ground truth exists, and the input votes provide noisy estimates of this ground truth. In a sequence of papers (see, e.g., [3]), we design social choice methods that provide accuracy guarantees with respect to wide families of possible noise distributions, or even with respect to worst-case noise. My latest project handles correlated noise that arises from an underlying social network. In an ongoing collaboration with Facebook, we aim to design and deploy more efficient rules when the social network structure is known.

Going beyond these specific examples, I am excited about the broader potential of algorithmic economics to make a real-world impact. This potential is evidenced by widely deployed game-theoretic algorithms for protecting critical infrastructure sites (an area that I have contributed to); and by the popularity of the fair division website Spliddit.org, which I am helping to develop. As the field continues to mature, and its theoretical foundations become firmer, I am certain that more opportunities will arise for applying algorithmic economics for societal good.

Representative Papers.

[1] Leximin Allocations in the Real World (EC 2015)
with A.D. Procaccia and D. Kurokawa

[2] Beyond Dominant Resource Fairness: Extensions, Limitations, and Indivisibilities (ACM TEAC and EC 2012) with D. Parkes and A.D. Procaccia

[3] When Do Noisy Votes Reveal the Truth? (EC 2013)
with I. Caragiannis and A.D. Procaccia

OR SHEFFET

Thesis. Beyond Worst-Case Analysis in Privacy and Clustering: Exploiting Explicit and Implicit Assumptions

Advisor. Avrim Blum, Carnegie Mellon University

Brief Biography. I'm a postdoctoral fellow at the Center for Research on Computation and Society at the School for Engineering and Applied Sciences at Harvard University, under the supervision of Prof. Salil Vadhan. Before joining Harvard, I was a Research Fellow at the Theoretical Foundations of Big Data program at the Simons Institute for the Theory of Computing in UC Berkeley. I completed my PhD in computer science from Carnegie Mellon University, advised by Prof. Avrim Blum. I got my M.Sc in computer science from the Weizmann Institute of Science, where I was advised by Prof. Oded Goldreich. I have a B.Sc in CS and Math from the Hebrew University in Jerusalem, Israel, where I worked with Prof. Nati Linial as part of the Amirim honors program.

Research Summary. My interests lie in many fields within computer science that touch on, and benefit from, rigorous mathematical theory. Projects I have worked on span areas such as algorithm design, machine learning and clustering, ranking and voting, algorithmic game theory and social network analysis. Complementary to this range of interests, my focus in recent years has been on the notion of differential privacy — a powerful, rigorous mathematical guarantee of privacy.

The main goal of my work is to design new differentially private data analysis techniques in the above-mentioned fields. In the coming years I will further pursue the study of the back-and-forth connections between differential privacy and these fields, as well as aim to establish new connections with other fields of big data analysis. One particular direction I am actively pursuing nowadays is the ability do statistical inference with differential privacy. Apparently, existing techniques in statistics – such as sampling from a posterior or regularizations – preserve privacy for the right choice of parameters.

The fact that privacy concerns are rooted in economic incentives is well-known. (To illustrate, think of the next two questions: Will you let me read your emails? Will you let me read your emails for a million dollars?) It is therefore natural to study the motivation for privacy from a game theoretic perspective, of selfish utility-maximizing agents. A recent work of mine, which I am currently continuing, studies the behavior of rational agents under concrete privacy concerns, where we show that some privacy concerns lead to agents behaving at equilibrium in a way that is differentially private, while in a different setting agents' behavior are diametrically different.

I am currently a member of Harvard's Privacy Tools project, aimed at implementing differentially private techniques in order to release information about real datasets. As part of my involvement with the project I have interacted with researchers from very different fields, like statisticians, social scientists and even lawyers. Though they approach data and think of data-driven tools in a very different way than in CS, I found this collaboration to be very rewarding and I plan on continuing such collaborations in the future. This is one of the goals I have set for myself: to promote the use of differentially private tools and to assist in the diffusion of differential privacy from CS to other scientific disciplines.

Representative Papers.

[1] Privacy Games (WINE 2014)
with Y. Chen and S. Vadhan

[2] The Johnson-Lindenstrauss Transform Itself Preserves Differential Privacy (FOCS 2012) with J. Blocki, A. Blum, and A. Datta

[3] Optimal Choice Functions: A Utilitarian View (EC 2012)
with C. Boutilier, I. Caragiannis, S. Haber, T. Lu, and A.D. Procaccia

VASILIS SYRGKANIS

Thesis. Efficiency of Mechanisms in Complex Markets

Advisor. Eva Tardos, Cornell University

Brief Biography. Vasilis Syrgkanis is a Postdoc researcher at Microsoft Research, NYC. He received his PhD in 2014, from the Computer Science Department of Cornell University under the supervision of Prof. Eva Tardos. His research interests include algorithms, game theory, auction theory, mechanism design, crowdsourcing, econometrics, online learning theory and computational complexity. He is the recipient of the Simons Graduate Fellowship in Theoretical Computer Science 2012-2014 and his research has received the best paper award at the ACM Conference on Economics and Computation. During his PhD he spent three summers as a research intern at Microsoft Research. Prior to his PhD he completed his undergraduate in Electrical Engineering and Computer Science at the National Technical University of Athens.

Research Summary. My research addresses the design and analysis of complex electronic marketplaces. It lies at the intersection of computer science and economics and more specifically in the areas of algorithms, game theory, mechanism design, econometrics and online learning theory, addressing optimization problems in the presence of incentives. I am interested in developing theoretical tools for analyzing and designing online markets, focusing on their distinct characteristics and their large scale nature. Some of the key topics I have worked on are:

Analysis and Design of Distributed Mechanisms. How efficient is a market composed of simple, distributed mechanisms for allocating resources and how should we design these local mechanisms in a way that global market efficiency is guaranteed? Most of mechanism design has focused on the design of centralized mechanisms that run in isolation. In "Composable and Efficient Mechanisms", we tackle the problem of designing distributed mechanisms and give an essential local property that each mechanism should satisfy for the market to achieve global efficiency, even under learning behavior and incomplete information. We show that this property is satisfied by several simple mechanisms, many of which are currently used in practice. Our work unifies a large set of results in the recent literature of characterizing the efficiency of simple and distributed mechanisms, including my work on the efficiency of sequential auctions, and has been subsequently applied and generalized in several settings. In "Bayesian Games and the Smoothness Framework" I also provide a more general approach for quantifying the efficiency in any incomplete information game.

Algorithmic Game Theory and Data Science. How can approaches from algorithmic game theory impact traditional econometrics and how can we use data to inform our theorems? I have explored two directions in this area. In "Econometrics for Learning Agents", we use an online learning theory approach to model strategic behavior in repeated game theoretic environments and based on that propose an econometric theory for inferring private parameters of participants. In "Robust Data-Driven Efficiency Guarantees in Auctions", we propose an approach for incorporating observed data of strategic behavior to infer efficiency guarantees in games that are better than the worst-case theoretical guarantees.

Representative Papers.

[1] Composable and Efficient Mechanisms (STOC 2013)
with E. Tardos

[2] Econometrics for Learning Agents (EC 2015)
with D. Nekipelov and E. Tardos

[3] Bayesian Incentive-Compatible Bandit Exploration (EC 2015)
with Y. Mansour and A. Slivkins

BO TANG

Thesis. On Optimization Problems in Auction Design

Advisor. Paul Goldberg, University of Oxford

Brief Biography. Bo is PhD student in Economics and Computation Group at the Department of Computer Science, University of Liverpool, supervised by Prof. Xiaotie Deng, Prof. Paul Goldberg and Dr. Giorgos Christodoulou. Before coming to Liverpool, he got a bachelor degree from Shanghai Jiao Tong University in China. He was a Research Intern at Microsoft Research Asia and a Research Assistant at Nanyang Technological University and Columbia University.

Research Summary. Bo investigated the effect of agents' manipulation in economic market in the following three projects: Price-taking vs Strategic manipulation, PoA of Simple Auctions Auctions and Auction Design with a Revenue Target.

Price-taking vs Strategic manipulation: The interplay of demand and price has been modeled as market equilibrium in economics. It have been shown that this classical model is not robust under strategic playing. Bo studied how much utility the buyer can gain by manipulation. For a general class of valuations, he proved this improvement is bounded by his utility as a price-taker and diminishes when the market grows larger and also provided several sufficient conditions for that price-taking approaches best strategic behavior. As a corollary, when the buyers are even allowed to form coalitions, the pricing-taking behavior is also a good approximation to the best response.

Simultaneous Simple Auctions Auctions: These simple auctions were applied by eBay to sell miscellaneous items via running single-item bidding auctions on each item simultaneously. Nevertheless, these succinct auctions disregard the interdependence between items like substitution and complementation. Thus, in such auctions the resulting allocation might be inefficient, that is, it doesn't allocate items to the buyers who want them most. Bo examined this efficiency loss in such item-bidding auctions and provided a tight lower bound for a class of valuations by constructing an inefficient Nash equilibrium. This result closes the gap of the efficiency bound for simultaneous simple auctions which has been studied in a series of literature from ICALP 2008 to STOC 2013.

Auction Design with a Revenue Target: The common objective for an auctioneer is to maximize the expected revenue raised from this auction. Actually, the objective can be generalized as a particular function of revenue when the seller is of special types like risk-averse, risk-seeking and goal-oriented. For instance, an auctioneer in debt would like to maximize the probability to earn a target revenue and pay off his debt. In contrast to the results for expected revenue, Bo showed the NP-hardness of computing the optimal auction with these objective functions even in single-parameter settings. On the positive side, polynomial-time algorithms can be developed to find the optimal mechanism for special cases based on a novel characterization of optimal auctions with a revenue target.

Representative Papers.

[1] Envy-Free Sponsored Search Auctions with Budgets (IJCAI 2015) with J. Zhang

[2] Pricing Ad Slots with Consecutive Multi-unit Demand (SAGT 2013) with P. Goldberg, X. Deng, Y. Sun and J. Zhang

[3] The Simulated Greedy Algorithm for Several Submodular Matroid Secretary Problems (STACS 2013) with T. Ma and Y. Wang

PANOS TOULIS

Thesis. Causal Inference under Network or Strategic Interference

Advisor. David C. Parkes (co-advised by Edoardo M. Airoldi and Donald B. Rubin), Harvard University

Brief Biography. I obtained my B.Sc. in Electrical Engineering from the Aristotle University (Greece) in 2005. Between 2006-2009 I worked on applications of intelligent agent systems, and between 2008-2009 I worked in the UAE on the creation of the first Arabic-speaking humanoid robot. In 2009, I moved to the U.S. and obtained my M.Sc. in CS at Harvard, and in 2011 I joined the Statistics Ph.D. program. At Harvard, I work on projects at the intersection of game theory, causal inference, and large data analysis through stochastic gradient descent. I received the 2015 Arthur P. Dempster prize for my work in implicit stochastic gradient descent, and the 2013 Thomas R. Ten Have award for my work in causal inference with interference. In 2012, I helped the Obama For America analytics team to do experimental design in voter mobilization on Facebook. My work has been supported by the 2015 LinkedIn EGC award, the 2012 Google US/Canada Ph.D. Fellowship in Statistics, and the Hellenic Harvard Foundation.

Research Summary. In social and economic contexts, there is an abundance of algorithms on how to mobilize voters over social media, or set prices in online ad auctions, or do viral marketing for a new product. However there is a shortage of methods to empirically evaluate them. In the context of statistical experimentation, the algorithms are considered to be treatments applied on certain units of analysis, e.g., voters, auctions, or customers. A fundamental problem in the evaluation of such treatments is interference. My research has focused on causal inference of treatment effects under three different forms of interference.

In social network interference, units affect each other through a pre-existing social network, e.g., friends affecting each other in their voting behavior. A key challenge is to use better randomizations of treatment on networks, and to use statistical models to disentangle the interference spillover effect from the primary effect of the treatment itself.

In strategic interference, units affect each other through their strategic actions; for example, in an online ad auction advertisers adjust their bids in response to new prices, thus affecting the competition. A key challenge is then to adjust causal inference to estimate long-term effects, i.e., effects that would be observed if we waited long enough until a new equilibrium was reached in response to new prices.

In experimental strategic interference, a new problem arises when the treatments to be evaluated are themselves self-interested agents, each having a strategic choice of what version of treatment to apply. In an experimental evaluation, the experimenter wants to know how an agent would perform if it adopted its natural behavior, defined as the choice of treatment version that the agent would make if there was no competition. However, agents can game the experiment by adopting different behaviors than their natural behavior, e.g., by applying more risky treatment versions. The goal is therefore to design an incentive-compatible experiment where agents will choose to adopt their natural behavior.

Parallel to causal inference, I have also been interested in incentive problems of mechanisms operating on random graphs in the context of kidney exchanges, and in estimation problems with large data sets using implicit stochastic approximations.

Representative Papers.

[1] Estimation of Causal Peer Influence Effects (ICML 2015) with E. Kao

[2] Incentive-Compatible Experimental Design (EC 2015) with D.C. Parkes, E. Pfeffer, and J. Zou

[3] Design and Analysis of Multi-Hospital Kidney Exchange Mechanisms Using Random Graphs (GEB 2015) with D.C. Parkes

DANIEL URIELI

Thesis. Learning Agents for Sustainable Energy

Advisor. Peter Stone, University of Texas at Austin

Brief Biography. Daniel Urieli is a PhD candidate (graduating in 2015) in The Department of Computer Science at The University of Texas at Austin. Daniel works with Professor Peter Stone on designing autonomous learning agents for sustainable energy problems. As a part of his research, Daniel designed a state-of-the-art, smart-grid energy trading agent that won several research competitions in 2013 and 2015. Before that, Daniel designed a learning agent for smart HVAC thermostat control, which is a part of a pending U.S. patent application by UT Austin. Previously, Daniel was a main contributor to the UT Austin Villa team, which won first place at the international RoboCup competitions in 2011 and 2012, in the 3D simulation league. Before joining UT Austin, Daniel completed a dual major B.Sc. in mathematics and computer science, and an M.Sc. (summa-cum-laude) in computer science at Tel Aviv University, and developed software for microprocessor power-delivery optimization at Intel.

Research Summary. The vision of a smart electricity grid is central to the efforts of moving society to a sustainable energy consumption. The main goals of the smart grid include (1) integration of intermittent, renewable energy sources, (2) reducing the peak electricity demand, and (3) automated energy efficiency. A main milestone for achieving these goals is "customer participation in power markets through demand-side-management". Demand-side management refers to adapting customer demand to supply conditions. Our research advances towards this milestone by designing state-of-the-art autonomous learning agents for energy trading and for energy efficiency.

In the context of goals (1)-(2) we designed TacTex, an autonomous energy trading agent that won several Power Trading Agent Competitions (Power TAC). The goal of Power TAC is to test novel energy market structures in simulation. This is important due to the high cost of failure in the real-world (like California-2001). In Power TAC, autonomous brokers compete for making profits in a realistic simulation of future smart-grid energy markets. Such brokers must be able to continually (1) learn (2) predict (3) plan (4) adapt in uncertain conditions. Our research gives insights regarding (1) computational techniques that are required for designing a successful broker (2) the overall impact of such autonomous brokers on the economy. We formalized the energy trading problem as two interdependent (intractable) utility-maximization problems. TacTex approximates their solutions by combining online reinforcement learning with efficient model-based optimization. Using TacTex, we investigated a widely-proposed method for demand-side management called Time-Of-Use tariffs (TOU), achieved state-of-the-art performance, and pointed out challenges and impacts of using TOU in competitive markets.

In the context of goal (3), we developed a thermostat-controlling agent that learns and adapts using advanced reinforcement learning. Our agent saves 7%-15% of the yearly energy consumption of a heat-pump HVAC system while maintaining occupants' comfort unchanged compared with the widely-used strategy, as observed in simulated experiments using a realistic simulator developed by the U.S. Department of Energy. Since HVAC systems are among the largest energy consumers, such savings can have a significant impact the total electricity demand. Our agent is a part of a pending U.S. patent application by The University of Texas at Austin.

Representative Papers.

[1] Autonomous Electricity Trading using Time-Of-Use Tariffs in a Competitive Market (under review) with P. Stone

[2] TacTex'13: A Champion Adaptive Power Trading Agent (AAAI 2014) with P. Stone

[3] A Learning Agent for Heat-Pump Thermostat Control (AAMAS 2013) with P. Stone

ANGELINA VIDALI

Thesis. Game-theoretic Analysis of Networks: Designing Mechanisms for Scheduling

Advisor. Elias Koutsoupias, University of Oxford

Brief Biography. Angelina Vidali is a Postdoctoral Researcher at Pierre and Marie Curie University-LIP6. She received her PhD from the Department of Informatics of the University of Athens (Greece), advised by Elias Koutsoupias. She also held Postdoctoral Researcher Positions at the Max Planck Institute for Informatics (Germany), at the University of Vienna (Austria) and at Duke University (USA). At Duke she organized an interdisciplinary seminar series (Departments of Economics, Computer Science and Fuqua School of Business) sponsored by Yahoo. Her research and studies have been supported by grants from the Vienna Science and Technology Fund, the Alexander von Humboldt foundation, the Alexandros Onassis foundation, the University of Athens, the Greek State Scholarship Foundation and the Greek Secretariat for Research and Technology.

Research Summary. My research lies in the intersection of computer science and economics; a timely, new, exciting research area with unexplored, well-motivated research directions, facing more and more challenges as electronic markets, cloud computing and crowdsourcing gain in market share and as markets and processes get reshaped by social networks. Computer science with its methodology and novel approaches addresses new questions and sheds new light on fundamental problems in economics. The internet enables us to run auctions, crowdsourcing contests and to compute tasks using cloud computing. The bidders/players/machines are not physically present but connected through the internet and linked through a social network structure, making these auctions easily accessible to a broader public and a part of our everyday life. MSN, Google, Yahoo and eBay need to design auctions for new settings such as sponsored search auctions, display ads, digital goods and pricing of cloud computing services, that will maximize their revenue but also guarantee customer satisfaction. As new markets emerge we need to build realistic new models and analyze them, while classic results from economics improve our intuition and provide us a solid background. This emerging new area has a lot more to contribute to computer science and economics in the coming years.

Representative Papers.

[1] A Characterization of n-Player Strongly Monotone Scheduling Mechanisms (IJCAI 2015) with A. Kovacs

[2] Mechanism Design for Scheduling with Uncertain Execution Time (AAAI 2014) with V. Conitzer

[3] A $1 + \phi$ Lower Bound for Truthful Scheduling Mechanisms (MFCS 2007) with E. Koutsoupias

ELAINE WAH

Thesis. Computational Models of Algorithmic Trading in Financial Markets

Advisor. Michael Wellman, University of Michigan

Brief Biography. Elaine Wah is currently a PhD candidate in Computer Science & Engineering at the University of Michigan. Her research interests lie at the intersection of finance and artificial intelligence, specifically in applying computational methods to study algorithmic trading in financial markets. Her dissertation work employs agent-based modeling and simulation to capture current market structure and to investigate the impact of algorithmic trading on market participants. She is a recipient of an NSF IGERT Fellowship and a Rackham Predoctoral Fellowship, and she received the Pragnesh Jay Modi Best Student Paper Award at AAMAS 2015. She has interned previously in the Division of Economic and Risk Analysis at the U.S. Securities and Exchange Commission, and she is spending summer 2015 as a Research Intern at Microsoft Research NYC. Prior to Michigan, she completed a BS in Electrical Engineering at the University of Illinois at Urbana-Champaign and an MS in Computer Science at UCLA.

Research Summary. Algorithmic trading, the use of quantitative algorithms to automate the submission of orders, is responsible for the majority of trading activity in today's financial markets. To better understand the societal implications of such trading, I construct computational agent-based models comprised of investors and algorithmic traders. I examine two overlapping types of algorithmic traders: high-frequency traders (of recent Flash Boys fame) who exploit speed advantages for profit, and market makers who facilitate trade and supply liquidity by simultaneously maintaining offers to buy and sell. I employ simulation and empirical game-theoretic analysis to study trader behavior in equilibrium, that is, when all traders best respond to their environment and other agents' strategies. I focus on the impact of algorithmic trading on allocative efficiency, or overall gains from trade.

I also investigate the potential for a frequent call market, in which orders are matched to trade at discrete periodic intervals rather than continuously, to mitigate the latency advantages of high-frequency traders. Frequent call markets have been proposed as a market design solution to the latency arms race perpetuated by high-frequency traders in continuous markets, but the path to widespread adoption of these call markets is unclear. I demonstrate that switching to a frequent call market eliminates the advantage of speed and promotes efficiency, and I formulate a game of strategic market choice to characterize the market conditions under which fast and slow traders choose to trade in a frequent call market versus a continuous double auction.

Representative Papers.

[1] Latency Arbitrage, Market Fragmentation, and Efficiency: A Two-Market Model (EC 2013) with M.P. Wellman

[2] Welfare Effects of Market Making in Continuous Double Auctions (AAMAS 2015) with M.P. Wellman

[3] Strategic Market Choice: Frequent Call Markets vs. Continuous Double Auctions for Fast & Slow Traders (AMMA 2015) with D.R. Hurd and M.P. Wellman

MATT WEINBERG

Thesis. Algorithms for Strategic Agents

Advisor. Costis Daskalakis, MIT

Brief Biography. Matt is currently a postdoc in the Computer Science department at Princeton University, hosted by Mark Braverman. From 2010-2014 he was a PhD student with Costis Daskalakis in Computer Science at MIT. Prior to that, Matt completed his B.A. in Math at Cornell University, where he worked with Bobby Kleinberg. During his time at MIT, Matt spent a summer interning with Microsoft Research New England, mentored two high school students in AGT research through the MIT PRIMES program, and supervised an undergraduate research project (UROP). He also spent the summers from 2009-2011 doing math and crypto research with the Department of Defense and the Institute for Defense Analyses. Before college, Matt grew up in Baltimore, MD. Outside of work, he spends the majority of his time training, teaching, and competing in taekwondo.

Research Summary. My research focuses largely on Algorithmic Mechanism Design, and has also made contributions to optimal stopping theory and convex optimization. At a high level, I like to study fundamental problems in Algorithmic Mechanism Design and distill from them, to the extent possible, purely algorithmic questions. Resolving such questions then develops new tools for these fundamental problems of study in a way that also contributes to more classical areas of Theoretical Computer Science.

One example of my work in this direction culminated in my thesis, and addresses the following question: In traditional algorithm design, some input is given and some output is desired. How much (computationally) harder is it to solve the same problem when the input is held instead by strategic agents with their own preferences over potential outputs? This broad question captures, for instance, the problem of optimal mechanism design from a computational perspective. My thesis provides a generic reduction from solving any optimization problem on strategic input to solving a perturbed version of that same problem when the input is directly given. In other words, we have shown how to answer questions in mechanism design by solving purely algorithmic problems.

My work also addresses mechanism design from other angles, such as understanding the quality of simple versus optimal auctions. Surprisingly, even in settings where the optimal auction is prohibitively complex or computationally intractable, we are able to show that very simple auctions can still perform quite well. My research in this area also develops new prophet inequalities and other online algorithms. Recently, I've also become interested in tackling from a mechanism design perspective more applied problems where strategic interaction is involved, such as peer grading in MOOCs.

Representative Papers.

[1] Understanding Incentives: Mechanism Design Becomes Algorithm Design (FOCS 2013) with Y. Cai and C. Daskalakis

[2] Matroid Prophet Inequalities (STOC 2012) with R. Kleinberg

[3] A Simple and Approximately Optimal Mechanism for an Additive Buyer (FOCS 2014) with M. Babaioff, N. Immorlica, and B. Lucier

JAMES R. WRIGHT

Thesis. Behavioral Game Theory: Predictive Models and Mechanisms

Advisor. Kevin Leyton-Brown, University of British Columbia

Brief Biography. James Wright is a Ph.D. candidate in computer science at the University of British Columbia, advised by Kevin Leyton-Brown. He holds an M.Sc. from the University of British Columbia (2010) and a B.Sc. from Simon Fraser University (2000). He studies problems at the intersection of behavioral game theory and computer science, with a focus on applying both machine learning techniques and models derived from experimental and behavioral economics to the prediction of human behavior in strategic settings. He also studies the implications of behavioral game theoretic models on multiagent systems and mechanisms. James's expected graduation date is June 2016.

Research Summary. A wealth of experimental evidence demonstrates that human behavior in strategic situations is often poorly predicted by classical economic models. Behavioral game theory studies deviations of human behavior from the standard assumptions, and provides many models of these deviations. These models typically focus on explaining a single anomaly. Although understanding individual anomalies is valuable, the resulting models are not always well-suited to predicting how people will behave in generic settings, which limits their application to questions of interest in algorithmic game theory, such as "What is the optimal mechanism for implementing a particular objective?".

I am interested applying machine learning techniques to construct behavioral game theoretic models that have high predictive accuracy, and in applying these models to problems in algorithmic game theory. As an example of the first direction, I previously analyzed and evaluated behavioral models in simultaneous-move games, eventually identifying a specific class of models (iterative models) as the state of the art. I then proposed and evaluated an extension that improves the prediction performance of any iterative model by better incorporating the behavior of nonstrategic agents.

Despite growing interest in behavioral game theory over the past decade, many important questions about its application to areas such as mechanism design remain open. For example, foundational analytic techniques such as the revelation principle may not be straightforwardly applicable under some classes of behavioral model. One direction of my current research aims to determine to which classes of behavioral model do such principles apply, and how to handle the cases where they don't apply.

Representative Papers.

[1] Beyond Equilibrium: Predicting Human Behavior in Normal-Form Games (AAAI 2010) with K. Leyton-Brown

[2] Behavioral Game-Theoretic Models: A Bayesian Framework for Parameter Analysis (AAMAS 2012) with K. Leyton-Brown

[3] Level-0 Meta-Models for Predicting Human Behavior in Games (EC 2014) with K. Leyton-Brown

JIE ZHANG

Thesis. Incentive Ratio and Market Equilibrium

Advisor. Xiaotie Deng, City University of Hong Kong

Brief Biography. From 2008.08 to 2011.07 Jie was a PhD student in City University of Hong Kong, advised by Xiaotie Deng. During the last six months of the PhD he was in Harvard University as a visiting student, hosted by Yiling Chen. From 2011.10 to 2014.03 he was a postdoc at Aarhus University, working under Peter Bro Miltersen. After that he moved to University of Oxford for another postdoc (research associate in the UK), working with Elias Koutsoupias.

Research Summary. My research mainly focuses on Algorithmic Game Theory. It analyzes strategic behaviors of rational agents, and designs efficient algorithms and mechanisms in multi-agent environments under incentive constraints, as well as equilibrium analysis.

Identifying, understanding and modeling agents' incentives and strategic behavior in cooperative and competitive environments is considered to be one of the most important subjects in the study of Internet markets. The analysis of agents' behavior aids mechanism designers to better achieve their objectives, such as information aggregation, revenue maximization, social welfare optimization, and so on. My objective is to design efficient algorithms and mechanisms to align the incentives of the agents with that of society, by employing algorithmic game theory methodology.

I have worked mostly, but not only, on the following topics: game theoretical analysis of market equilibrium; Prediction markets; mechanism design; fixed-point models and complexity of computing Nash equilibrium; fair division in resource allocation.

Representative Papers.

[1] What You Jointly Know Determines How You Act Strategic Interactions in Prediction Markets (EC 2013) with X.A. Gao and Y. Chen

[2] How Profitable are Strategic Behaviors in a Market? (ESA 2011) with N. Chen and X. Deng

[3] Externalities in Cake Cutting (IJCAI 2013) with S. Brnzei and A.D. Procaccia

YAIR ZICK

Thesis. Arbitration, Fairness and Stability: Revenue Division in Collaborative Settings

Advisor. Edith Elkind, University of Oxford

Brief Biography. Yair Zick is a postdoctoral research fellow in the computer science department of Carnegie Mellon University, hosted by Anupam Datta and Ariel D. Procaccia. His research interests span cooperative game theory, computational social choice and their applications to domains such as security, privacy, machine learning and education. He completed his PhD at Nanyang Technological University under the supervision of Edith Elkind where he was funded by the Singapore A*STAR scholarship. As a graduate student, he has coauthored nine publications (seven of which as a main author), all appearing in top AI conferences. His first paper, "Arbitrators in Overlapping Coalition Formation Games" received the AAMAS 2011 Pragnesh Jay Modi best student paper award; his dissertation "Arbitration, Fairness and Stability: Revenue Division in Collaborative Settings", has received the 2014 IFAAMAS Victor Lesser distinguished dissertation award.

Research Summary. As a graduate student, I mostly worked on cooperative game theory, and coauthored several papers on cooperative games and computational social choice. My thesis mainly focused on overlapping coalition formation (OCF). In OCF games, each player possesses some divisible resource (say, time or processing power), and may contribute a fractional amount of it to joint tasks with other players. These fractional coalitions generate revenue, which must be divided among participants. An outcome (a division into coalitions plus a division of revenue) of an OCF game is stable if no subset of agents can deviate —reallocate resources and revenue such that all of its members are strictly better off. The key observation here is that when agents deviate, they may still be invested in projects involving non-deviators. For example, if an agent receives payoffs from several projects but would like to withdraw only from one of them, the profitability of deviation strongly depends on how non-deviators react.

I studied OCF games with arbitration functions. These functions describe the way non-deviators react to deviation; their lenience governs the stability of OCF games. The structure of arbitration functions has far-reaching implications on stability, as well as the computational problem of finding such outcomes.

Upon graduation, I joined Carnegie Mellon University as a postdoctoral research fellow. During my time here, I expanded my research interests to include machine learning, privacy, security, and fair division. I am currently involved in several exciting projects. We are exploring an interesting link between game theory and causality in machine learning environments. In addition, we apply PAC learning techniques to cooperative games and to fair allocation of indivisible goods. Our results in an ongoing rent division project have been implemented on the `Spliddit.org` website, with more results on the way! In a foray to the field of AI and education, we implement machine learning models in order to elicit various student metrics from course data.

I am passionate about applying game theoretic notions to other fields, and I always enjoy learning about new fields and techniques. Furthermore, I am keenly interested in empirical analysis of game theoretic solution concepts: what outcomes are considered fair by people? Do their notions of fairness coincide with our formal definitions?

Representative Papers.

[1] Influence in Classification via Cooperative Game Theory (IJCAI 2015) with A. Datta, A. Datta, and A.D. Procaccia

[2] Learning Cooperative Games (IJCAI 2015) with M.F. Balcan and A.D. Procaccia

[3] Arbitration and Stability in Cooperative Games with Overlapping Coalitions (JAIR 2014) with E. Markakis and E. Elkind

Index