

Innovations in
ALGORITHMIC GAME THEORY
Jerusalem, May 22nd-26th, 2011



ORGANIZERS:

Michal Feldman
Noam Nisan



האוניברסיטה העברית בירושלים
The Hebrew University of Jerusalem

Innovations in ALGORITHMIC GAME THEORY

Jerusalem, May 22nd-26th, 2011

Partially supported by the Google Inter-university Center for Electronic Auctions and Markets
and Microsoft Israel R&D Center



Room 130, Feldman building, Edmund J. Safra, Givat Ram campus.

At the approximate age of 10 years, it seems that the field of Algorithmic Game Theory is maturing. IAGT seeks to sway its path for the next decade by focusing on research that carries strong conceptual messages and broadens the scope of the interface between economics, game theory and computer science. e.g., propose new concepts and techniques, introduce novel applications and draw attention to interesting connections between AGT's parent disciplines.

Time	Sunday May 22nd	Monday May 23rd	Tuesday May 24th	Wednesday May 25th	Thursday May 26th
09.00-09.30		Coffee	Coffee	Coffee	Coffee
09.30-11.10	Trip to Massada & Dead Sea	Michael O. Rabin Vincent Conitzer Albert Xin Jiang	Eva Tardos Lisa Fleischer Nati Linial	Sushil Bikhchandani Jing Chen Moshe Babaioff	Constantinos Daskalakis Shahar Dobzinski Jason Hartline
11.10-11.40	Leaving from IAS at 3:00	Coffee	Coffee	Coffee	Coffee
11.40-12.45		Paul Goldberg Moshe Tennenholtz	Asu Ozdaglar Avrim Blum	Michal Feldman Kevin Leyton-Brown Nicole Immorlica	Stefano Leonardi Silvio Micali
12.45-13.15		Lunch & Break	Lunch & Break		Lunch & Break
13.15-14.30				Lunch & Break	
14.30-16.10		Elchanan Mossel Abraham Neyman Nina Balcan	Noam Nisan Yiling Chen Vijay Vazirani	Tour of the old city of Jerusalem. Leaving from IAS at 14.30.	Anna Karlin Amos Fiat Liad Blumrosen
16.10-16.40		Coffee	Coffee		Coffee
16.40-17.45		Edith Elkind Sergiu Hart	Yishay Mansour Angelina Vidali Tim Roughgarden		Jeff Rosenschein Christos Papadimitriou
17.45-18.15					
18.15-19.00	Reception @ IAS	Poster Session			
19.00-20.00				Conference Banquet (After the tour) "Terasa" restaurant	

Sunday, May 22nd

03:00-14:00



Massada & The Dead Sea

Leaving from the main gate of campus.

Bring at least 2 liters of water, a hat and proper shoes for hiking.

Sunscreen may also be a good idea.

Packed breakfasts will be provided for all trip participants.

We will have lunch in Ein Gedi.

To dip in the dead sea - bring your swimsuits!

18:00-20:00



Reception

Feldman Building Lobby

Monday, May 23rd

Session Chair: Noam Nisan

09:00-09:30



Coffee

09:30-10:00

Michael O. Rabin

(The Hebrew University & Harvard University)

Cryptography and Solutions for Matching Problems

10:05-10:35

Vincent Conitzer

(Duke University)

Algorithms for Security Games

10:40-11:10

Albert Xin Jiang

(University of British Columbia)

Polynomial-time Computation of Exact Correlated Equilibrium in Compact Games

11:10-11:40



Coffee

11:40-12:10

Paul Goldberg

(University of Liverpool)

The Complexity of Homotopy Methods for Computing Nash Equilibria

12:15-12:45

Moshe Tennenholtz

(Technion & Microsoft Israel R&D Center)

Mechanisms for Multi-Level Marketing

12:45-14:30



Lunch & Break

Session Chair: Elias Koutsoupias

14:30-15:00

Elchanan Mossel

(UC Berkeley & Weizmann Institute)

If Not Agreeing to Disagree Then Agreeing on What?

15:05-15:35

Abraham Neyman

(The Hebrew University)

Open Problems in Repeated Games With Finite Automata

15:40-16:10

Nina Balcan

(Georgia Institute of Technology)

Learning Valuation Functions

16:10-16:40



Coffee

16:40-17:10

Edith Elkind

(Nanyang Technological University)

Ties Matter: Complexity of Voting Manipulation Revisited.

17:15-17:45

Sergiu Hart

(The Hebrew University)

Comparing and Measuring Risks

17:45-19:00



Poster Session

Tuesday, May 24th

Session Chair: Nicole Immorlica

09:00-09:30



Coffee

09:30-10:00

Eva Tardos

(Cornell University)

Network Formation in the Presence of Contagious Risk

10:05-10:35

Lisa Fleischer

(Dartmouth)

Competitive Strategies for Routing Flow Over Time.

10:40-11:10

Nati Linial

(The Hebrew University)

No Justified Complaints - Bottleneck-Based Fairness

11:10-11:40



Coffee

11:40-12:10

Asu Ozdaglar

(MIT)

Dynamics in Near-Potential Games

12:15-12:45

Avrim Blum

(Carnegie Mellon University)

The Price of Uncertainty: Safety Conditions for Multiagent Systems

12:45-14:30



Lunch & Break

Session Chair: Kevin Leyton-Brown

14:30-15:00

Noam Nisan

(The Hebrew University)

Non-Price Equilibria in Markets of Discrete Goods

15:05-15:35

Yiling Chen

(Harvard University)

Automated Market-Making via Online Convex Optimization

15:40-16:10

Vijay Vazirani

(Georgia Institute of Technology)

Extending General Equilibrium Theory to the Digital Economy

16:10-16:40



Coffee

16:40-17:10

Yishay Mansour

(Tel Aviv University)

Welfare and Profit Maximization with Production Costs

17:15-17:45

Angelina Vidali

(University of Vienna)

Extending Characterizations of Truthful Mechanisms From Subdomains to Domains

17:50-18:20

Tim Roughgarden

(Stanford University)

Simple Auctions with Near-Optimal Equilibria

Wednesday, May 25th

Session Chair: Amos Fiat

09:00-09:30



Coffee

09:30-10:00

Sushil Bikhchandani

(UCLA)

Mechanism Design with Information Acquisition: Efficiency and Full Surplus Extraction

10:05-10:35

Jing Chen

(MIT)

Crowdsourced Bayesian Auctions

10:40-11:10

Moshe Babaioff

(Microsoft Research)

Peaches, Lemons, and Cookies: Designing Auction Markets with Dispersed Information

11:10-11:40



Coffee

11:40-12:10

Michal Feldman

(The Hebrew University)

Revenue Maximization in Probabilistic Single-Item Auctions via Signaling

12:15-12:45

Kevin Leyton-Brown

(University of British Columbia)

Dominant-Strategy Auction Design for Agents with Uncertain, Private Values

12:50-13:20

Nicole Immorlica

(Northwestern University)

Dueling Algorithms

13:20-14:30



Lunch & Break

14:30-19:00



Trip to the Old City

Leaving from the main campus gate.

19:00-20:30

Workshop Banquet

Dinner in the "Terasa" restaurant.

Thursday, May 26th

Session Chair: Anna Karlin

09:00-09:30



Coffee

09:30-10:00

Constantinos Daskalakis

(MIT)

On Optimal Multi-Dimensional Mechanism Design

10:05-10:35

Shahar Dobzinski

(Cornell University)

Bounding the Power of Truthfulness

10:40-11:10

Jason Hartline

(Northwestern University)

Truth or Envy?

11:10-11:40



Coffee

11:40-12:10

Stefano Leonardi

(Sapienza University of Rome)

Single Valued Combinatorial Auctions with Budgets

12:15-12:45

Silvio Micali

(MIT)

Collusive Dominant-Strategy Truthfulness

12:45-14:30



Lunch & Break

Session Chair: Michal Feldman

14:30-15:00

Anna Karlin

(University of Washington)

Selling in Exclusive Markets: Some Observations About Prior-Free Profit Maximization.

15:05-15:35

Amos Fiat

(Tel Aviv University)

Bribing a Guard and Related Issues

15:40-16:10

Liad Blumrosen

(The Hebrew University)

Only Valuable Experts Can Be Valued

16:10-16:40



Coffee

16:45-17:15

Jeff Rosenschein

(The Hebrew University)

Tight Bounds for Strategyproof Classification

17:20-17:50

Christos Papadimitriou

(UC Berkeley)

TBD

ABSTRACTS

Michael O. Rabin

(The Hebrew University & Harvard University)

Cryptography and Solutions for Matching Problems

We treat as an example stable matchings. There are n entities (say hospitals) and m candidates (say medical school graduates). Each entity ranks the candidates and each candidate ranks the entities. An administrator using a well known algorithm matches candidates to entities so that a certain stability condition is satisfied.

In certain situations it is desirable that the rankings/preferences remain secret, so that only the stable matching but no rankings are revealed by the administrator. We provide a surprising powerful method where the correctness, i.e. the stability of the announced matching, is proved to all participants without revealing any preferences.

The solution employs joint work with Y. Mansur, M. Muthukrishnan, and M. Yung.

Vincent Conitzer

(Duke University)

Algorithms for Security Games

I will talk about our work on algorithms for computing game-theoretic solutions (Stackelberg and Nash) in security games.

Joint work with Dmytro Korzhyk, Joshua Letchford, Kamesh Munagala, Ronald Parr (Duke), Manish Jain, Zhengyu Yin, Milind Tambe (USC), Christopher Kiekintveld (UT El Paso), Ondrej Vanek, Michal Pechoucek (Czech Technical University), and Tuomas Sandholm (CMU).

Albert Xin Jiang

(University of British Columbia)

Polynomial-time Computation of Exact Correlated Equilibrium in Compact Games

In a landmark paper, Papadimitriou and Roughgarden described a polynomial-time algorithm ("Ellipsoid Against Hope") for computing sample correlated equilibria of concisely-represented games. Recently, Stein, Parrilo and Ozdaglar showed that this algorithm can fail to find an exact correlated equilibrium, but can be easily modified to efficiently compute approximate correlated equilibria. Currently, it remains an open problem to determine whether the algorithm can be modified to compute an exact correlated equilibrium.

We show that it can, presenting a variant of the Ellipsoid Against Hope algorithm that guarantees the polynomial-time identification of exact correlated equilibrium. Also, our algorithm is the first to tractably compute correlated equilibria with polynomial-sized supports; such correlated equilibria are more natural solutions than the mixtures of product distributions produced previously, and have several advantages including requiring fewer bits to represent, being easier to sample from, and being easier to verify. Our algorithm differs from the original primarily in its use of a separation oracle that produces cuts corresponding to pure-strategy profiles. As a result, we no longer face the numerical precision issues encountered by the original approach, and both the resulting algorithm and its analysis are considerably simplified. Our new separation oracle can be understood as a derandomization of Papadimitriou and Roughgarden's original separation oracle via the method of conditional probabilities.

Joint work with Kevin Leyton-Brown

Paul Goldberg

(University of Liverpool)

The complexity of homotopy methods for computing Nash equilibria

Homotopy methods for solving games are based on the following approach. Given a game to be solved, construct a version where the payoffs have been changed so that there is an obvious Nash equilibrium. Then continuously adjust the payoffs towards those of the game of interest - the Nash equilibrium should change continuously, ending up at an equilibrium of that game. The Lemke-Howson algorithm is the best-known example of such a method. Surprisingly, the resulting equilibrium turns out to be PSPACE-complete to compute. I will give an overview of the general proof idea, and related open problems.

Moshe Tennenholtz

(Technion & Microsoft Israel R&D Center)

Mechanisms for Multi-Level Marketing

Multi-level marketing is a marketing approach that motivates its participants to promote a certain product among their friends. The popularity of this approach increases due to the accessibility of modern social networks, however, it existed in one form or the other long before the Internet age began (the infamous Pyramid scheme that dates back at least a century is in fact a special case of multi-level marketing). This paper lays foundations for the study of reward mechanisms in multi-level marketing within social networks. We provide a set of desired properties for such mechanisms and show that they are uniquely satisfied by geometric reward mechanisms. The resilience of mechanisms to false-name manipulations is also considered; while geometric reward mechanisms fail against such manipulations, we exhibit other mechanisms which are false-name-proof.

Joint work with Yuval Emek, Ron Karidi, and Aviv Zohar

Elchanan Mossel

(UC Berkeley & Weizmann Institute)

If Not Agreeing to Disagree Then Agreeing on What?

Beginning with results of Aumann and Geanakoplos & Polemarchakis extensive research has been devoted to establishing conditions under which repeated actions in a Bayesian setup lead to a common posterior. However, the analytic and computational difficulties involved in Bayesian game theory prevented thus-far from obtaining useful information regarding the value of the common posterior except in very simple cases. I will describe some recent work (with Omer Tamuz) presenting models where the Bayesian updates are computationally feasible and Bayesian learning is statistically consistent for large number of players on general connected networks followed by very recent developments (with Allan Sly and Omer Tamuz) showing that even for models where efficient Bayesian implementations are not known Bayesian learning is statistically consistent in a very general setup.

Abraham Neyman

(The Hebrew University)

Open Problems in Repeated Games with Finite Automata

The talk will state several open problems, and survey the related developments in the theory of repeated games with finite automata.

Nina Balcan

(Georgia Institute of Technology)

Learning Valuation Functions

A core element of microeconomics and game theory is that consumers have valuation functions over bundles of goods and that these valuation functions drive their purchases. In particular, the value given to a bundle need not be the sum of values on the individual items but rather can be a more complex function of how the items relate. Common valuation classes considered in the literature include OXS, submodular, and XOS valuations. Typically it is assumed that these valuations are known to the center or come from a known distribution. In this work we initiate the study of the approximate learnability of valuation classes in a distributional learning setting. We prove upper and lower bounds on the approximation guarantees achievable for learning over general data distributions by using only a polynomial number of examples.

Our work combines central issues in economics with central issues in optimization (submodular functions and matroids) with central issues in learning (learnability of natural but complex classes of functions in a distributional setting). Our analysis brings a twist on the usual learning theory models and uncovers some interesting structural and extremal properties of submodular functions, which are likely to be useful in other contexts as well.

Edith Elkind

(Nanyang Technological University)

Ties Matter: Complexity of Voting Manipulation Revisited

Computational complexity of voting manipulation is one of the most actively studied topics in computational social choice. Much of the existing work in this area, starting with the seminal paper of Bartholdi, Tovey, and Trick, assumes that ties are broken either in favor of the manipulator or adversarially to the manipulator. Under this assumption, most classic voting rules, such as Maximin, Copeland and all scoring rules, are easy to manipulate. In this talk, we will examine the role of this assumption in the existing easiness of manipulation results. We will show that if ties are broken by choosing a winner uniformly at random among all tied candidates, manipulation remains easy for all scoring rules, and, for a natural special class of voters' preferences, for the Maximin rule. However, it becomes hard for Maximin under general preferences as well as for the Copeland rule. We will also show hardness results for general polynomial-time computable tie-breaking rules.

Based on joint work with Svetlana Obraztsova (NTU) and Noam Hazon (CMU).

Sergiu Hart

(The Hebrew University)

Comparing and Measuring Risks

Stochastic dominance is a partial order on risky assets ("gambles") that is based on the uniform preference---of all decision-makers of an appropriate class---for one gamble over another. We modify this requirement, first, by taking into account the status quo (given by the current wealth) and the possibility of rejecting gambles, and second, by comparing rejections that are substantive (that is, uniform over wealth levels or over utilities). This yields two new stochastic orders: "wealth-uniform dominance" and "utility-uniform dominance". Unlike stochastic dominance, these two orders are complete: any two gambles can be compared. Moreover, they are equivalent to the orders induced by, respectively, the Aumann-Serrano (JPE 2008) index of riskiness and the Foster-Hart (JPE 2009) measure of riskiness.

<http://www.ma.huji.ac.il/hart/abs/risk-u.html>

<http://www.ma.huji.ac.il/hart/abs/risk.html>

Eva Tardos

(Cornell University)

Network Formation in the Presence of Contagious Risk

There are a number of domains where agents must collectively form a network in the face of the following trade-off: each agent receives benefits from the direct links it forms to others, but these links expose it to the risk of being hit by a cascading failure that might spread over multi-step paths. Financial contagion, epidemic disease, and the exposure of covert organizations to discovery are all settings in which such issues have been articulated. In this talk we formulate the problem in terms of strategic network formation, and provide asymptotically tight bounds on the welfare of both optimal and stable networks. Our analysis enables us to explore such issues as the trade-offs between clustered and anonymous market structures, and it exposes a fundamental sense in which very small amounts of "over-linking" in networks with contagious risk can have strong consequences for the welfare of the participants.

Joint work with Larry Blume, David Easley, Jon Kleinberg, and Robert Kleinberg

Lisa Fleischer

(Dartmouth)

Competitive Strategies for Routing Flow Over Time.

The study of routing games is motivated by the desire to understand the impact of individual user's decisions on network efficiency. To do this, prior work uses a simplified model of network flow where all flow exists simultaneously, and users route flow to either their maximum delay or their total delay. Both of these measures are surrogates for measuring how long it takes to get all of your traffic through the network over time.

Instead of using these surrogates, we attempt a more direct study of how competition among users effects network efficiency by examining routing games in a flow-over-time model. We show that the network owner can reduce available capacity so that the competitive equilibrium in the reduced network is no worse than a small constant times the optimal solution in the original network using two natural measures of optimum: the time by which all flow reaches the destination, and the average amount of time it takes flow to reach the destination.

Joint work with Umang Bhaskar (Dartmouth) and Elliot Anshelevich (RPI)

Nati Linial

(The Hebrew University)

No Justified Complaints - Bottleneck-Based Fairness

Much work has been dedicated to fair sharing of resources. Our work has its origins in the management of large computational systems, where n users have access to m resources. Let r_{ij} be the amount of resource j that user i requests. The proportion between the different r_{ij} is fixed and reflects user i 's computational needs. Thus we should choose a real number $1 \geq x_i \geq 0$ for each i and give $x_i * r_{ij}$ units of resource j to user i .

Such an allotment x_1, \dots, x_n is feasible iff for every resource j we do not give away more than the available supply we have of it, which for simplicity and wlog we assume to be 1, i.e., $\sum_i x_i * r_{ij} \leq 1$. Resource j is considered a bottleneck if it is used up completely, i.e., $\sum_i x_i * r_{ij} = 1$.

In our model each user i has an entitlement $e_i > 0$ where the sum of e_i is 1. This number may represent user i 's share in investment in the system or the importance associated to his activity.

The allotment x_1, \dots, x_n is considered fair if for every user i there is a bottleneck resource j for which $x_i * r_{ij} \geq e_i$. In words "user i has no grounds for complaining, since there is a bottleneck resource on which he is receiving at least his entitlement".

Using ideas from the theory of ordinary differential equations we show that a fair allotment always exists. Many interesting questions remain open.

Joint work with Danny Dolev, Dror Feitelson, Joe Halpern, and Raz Kupferman

Asu Ozdaglar

(MIT)

Dynamics in Near-Potential Games

Potential games are a special class of games that allow for tractable dynamic analysis. Intuitively, games that are "close" to a potential game should share similar properties. In this talk, we formalize and develop this idea by quantifying to what extent the dynamic features of potential games extend to near-potential games. We first show that in an arbitrary finite game, the limiting behavior of better-response and best-response dynamics can be characterized by the approximate equilibrium set of a close potential game. Moreover, the size of this set is proportional to a closeness measure between the original game and the potential game. We then focus on logit response dynamics, which induce a Markov process on the set of strategy profiles of the game, and show that the stationary distribution of logit-response dynamics can be approximated using the potential function of a close potential game, and its stochastically stable strategy profiles can be identified as the approximate maximizers of this function. Our approach presents a systematic framework for studying convergence behavior of adaptive learning dynamics in finite strategic form games.

Joint work with Ozan Candogan and Pablo Parrilo

Avrim Blum

(Carnegie Mellon University)

The Price of Uncertainty: Safety Conditions for Multiagent Systems

Suppose that, through whatever means, a system has been able to reach a low-cost equilibrium state. What assurances can one provide that the state will not arbitrarily deteriorate? We in particular consider a model in which players in a potential game experience a small degree of uncertainty about their true costs, and therefore can be expected only to perform "nearly-best-response" actions to the current state. We consider a variety of potential games including fair cost-sharing games, set-cover games, routing games, and job-scheduling games. We show that in certain cases, even extremely small fluctuations in perceived costs can cause dynamics to spin out of control and move to states of much higher social cost, whereas in other cases these dynamics are much more stable even to large fluctuations. We also consider the resilience of these dynamics to a small number of Byzantine players about which no assumptions are made. We show that in certain cases even a single Byzantine player can cause best-response dynamics to transition to states of substantially higher cost, whereas in others these dynamics are much more resilient.

Joint work with Maria-Florina Balcan and Yishay Mansour

Noam Nisan

(The Hebrew University)

Non-Price Equilibria in Markets of Discrete Goods

We study markets of indivisible items in which price-based (Walrasian) equilibria often do not exist due to the discrete non-convex setting. Instead we consider Nash equilibria of the market viewed as a game, where players bid for items, and where the highest bidder on an item wins it and pays his bid. We first observe that pure Nash-equilibria of this game exactly correspond to price-based equilibria (and thus need not exist), but that mixed-Nash equilibria always do exist, and we analyze their structure in several simple cases where no price-based equilibrium exists. We also undertake an analysis of the welfare properties of these equilibria showing that while pure equilibria are always perfectly efficient ("first welfare theorem"), mixed equilibria need not be, and we provide upper and lower bounds on their amount of inefficiency.

Joint work with Avinatan Hassidim, Haim Kaplan and Yishay Mansour

Yiling Chen

(Harvard University)

Automated Market-Making via Online Convex Optimization

While computers have automated the operation of most financial markets, the underlying mechanism was designed for people to operate it. It is simple, not necessarily efficient, and has room for improvement. This work is an endeavor to design efficient automated market-making mechanisms that take into consideration of the logical relationships of securities.

We propose a general framework for the design of securities markets over combinatorial or infinite state spaces. The framework enables the design of computationally efficient markets tailored to an arbitrary, yet relatively small, space of securities with bounded payoff. We prove that any market satisfying a set of intuitive conditions must price securities via a convex cost function, which is constructed via conjugate duality. Rather than dealing with an exponentially large or infinite outcome space directly, our framework only requires optimization over a convex hull. By reducing the problem of automated market-making to convex optimization, where many efficient algorithms exist, we arrive at a range of new polynomial-time pricing mechanisms for various problems. We demonstrate the advantages of this framework with the design of some particular markets. We also show that by relaxing the convex hull we can gain computational tractability without compromising the market institution's bounded budget.

Based on joint work with Jacob Abernethy and Jennifer Wortman Vaughan

Vijay Vazirani

(Georgia Institute of Technology)

Extending General Equilibrium Theory to the Digital Economy

General Equilibrium Theory, the undisputed crown jewel of Mathematical Economics for over a century, gave a very satisfactory solution to the central problem of arriving at a principled method of pricing goods and services -- one that leads to an efficient allocation of scarce resources among alternative uses. The solution was based on Adam Smith's principle of maintaining parity between supply and demand, Walras' notion of equilibrium, and the Arrow-Debreu Theorem, which proved the existence of equilibrium in a very general model of the economy.

However, this solution, designed for conventional goods, is not applicable for digital goods -- once produced, a digital good can be reproduced at (essentially) zero cost, thus making its supply infinite. Considering the current size of the digital economy and its huge growth potential, it is imperative that we obtain an equally convincing theory for pricing of digital goods. In this talk, I will describe what appears to be a first step in this direction.

After taking into consideration idiosyncrasies of the digital realm, we give a market model that is appropriate for the digital setting, a notion of equilibrium for this model, and a proof of existence of equilibrium using Kakutani's fixed point theorem. For a special case, we also give a polynomial time algorithm and we show that a rational equilibrium always exists. Finally, I will outline a multitude of issues that still need to be addressed to obtain a theory for the digital economy that is on par with the original theory.

Based on the joint paper with Kamal Jain: <http://arxiv.org/abs/1007.4586>

Nicole Immorlica

(Northwestern University)

Dueling Algorithms

We revisit classic algorithmic search and optimization problems from the perspective of competition. Rather than a single optimizer minimizing expected cost, we consider a zero-sum game in which an optimization problem is presented to two players, whose only goal is to outperform the opponent. Such games are typically exponentially large zero-sum games, but they often have a rich combinatorial structure. We provide general techniques by which such structure can be leveraged to find minmax-optimal and approximate minmax-optimal strategies. We give examples of ranking, hiring, compression, and binary search duels, among others. We give bounds on how often one can beat the classic optimization algorithms in such duels.

Joint work with Adam Tauman Kalai, Brendan Lucier, Ankur Moitra, Andrew Postlewaite, and Moshe Tennenholtz.

Angelina Vidali

(University of Vienna)

Extending Characterizations of Truthful Mechanisms From Subdomains to Domains

The already extended literature in combinatorial auctions, public projects and scheduling demands a more systematic classification of the domains and a clear comparison of the results known. Connecting characterization results for different settings and providing a characterization proof using another characterization result as a black box without having to repeat a tediously similar proof is not only elegant and desirable, but also greatly enhances our intuition and provides a classification of different results and a unified and deeper understanding.

Characterizing the mechanisms for the domains of combinatorial auctions and scheduling unrelated machines are two outstanding problems in mechanism design. Since the scheduling domain is essentially the subdomain of combinatorial auctions with additive valuations, we consider whether one can extend a characterization of a subdomain to a domain. This is possible for two players (and for n -player mechanisms that satisfy S-monotonicity) if the only truthful mechanisms for the subdomain are the affine maximizers. Although this is not true for scheduling because besides the affine maximizers there are other truthful mechanisms (the threshold mechanisms), we still show that the truthful mechanisms that allocate all goods of practically any domain which is strictly superdomain of additive combinatorial auctions are only the affine maximizers.

Tim Roughgarden

(Stanford University)

Simple Auctions with Near-Optimal Equilibria

In practice, auction designers often exchange incentive guarantees for simplicity and tractability (e.g., in sponsored search or combinatorial auctions). How much welfare loss does this design choice cause? I'll talk about some techniques for answering this question.

Based largely on a SODA '11 paper with Kshipra Bhawalkar

Sushil Bikhchandani

(UCLA)

Mechanism Design with Information Acquisition: Efficiency and Full Surplus Extraction

Consider a mechanism design setting in which agents acquire costly information about an unknown, payoff-relevant state of nature. Information gathering is covert and the agents' information may be correlated. We investigate conditions under which (i) efficient implementation and (ii) full surplus extraction are Bayesian incentive compatible and interim individually rational.

Joint work with Ichiro Obara

Jing Chen

(MIT)

Crowdsourced Bayesian Auctions

We investigate the problem of generating revenue in auctions when the valuations of the players are jointly drawn from a (perhaps correlated) distribution D and "the players know each other better than the seller knows them." That is, when only the players have some information about D , so that the seller needs to elicit their help to sell the goods: he needs to crowdsource the auction to the players.

We work under a "very weak decentralized Bayesian assumption". Not only no information about D is common knowledge to the players, but no player even individually knows D . In fact, in our crowdsourced Bayesian assumption each player i only individually knows an arbitrary refinement of $D|t_i$, that is, D 's posterior based on his own type t_i .

Yet, we prove that such a weak assumption enables the seller to generate revenue that in expectation is close to the one he could generate by means of an optimal dominant-strategy-truthful mechanism if he himself knew D . Our results are "existential" for all auctions, and "constructive" for single-good ones.

Joint work with Pablo Azar and Silvio Micali

Moshe Babaioff

(Microsoft Research)

Peaches, Lemons, and Cookies: Designing Auction Markets with Dispersed Information

This paper studies the role of information asymmetries in second price, common value auctions. Motivated by information structures that arise commonly in applications such as online advertising, we seek to understand what types of information asymmetries lead to substantial reductions in revenue for the auctioneer. One application of our results concerns online advertising auctions in the presence of "cookies," which allow individual advertisers to recognize advertising opportunities for users who, for example, are customers of their websites. Cookies create substantial information asymmetries both *ex ante* and at the interim stage, when advertisers form their beliefs. The paper proceeds by first introducing a new refinement, which we call "tremble robust equilibrium" (TRE), which overcomes the problem of multiplicity of equilibria in many domains of interest.

Second, we consider a special information structure, where only one bidder has access to superior information, and show that the seller's revenue in the unique TRE is equal to the expected value of the object conditional on the lowest possible signal, no matter how unlikely it is that this signal is realized. Thus, if cookies identify especially good users, revenue may not be affected much, but if cookies can (even occasionally) be used to identify very poor users, the revenue consequences are severe. In the third part of the paper, we study the case where multiple bidders may be informed, providing additional characterizations of the impact of information structure on revenue. Finally, we consider richer market designs that ensure greater revenue for the auctioneer, for example by auctioning the right to participate in the mechanism.

Joint work with Ittai Abraham, Susan Athey and Michael Grubb

Michal Feldman

(The Hebrew University)

Revenue Maximization in Probabilistic Single-Item Auctions via Signaling

Signaling is an important topic in the study of asymmetric information in economic settings. In particular, the transparency of information available to a seller in an auction setting is a question of major interest. We introduce the study of signaling when auctioning a probabilistic good whose actual instantiation is known to the auctioneer but not to the bidders. In particular, this can be used to model impressions selling in display advertising. Our study focuses on finding an optimal signaling scheme, captured by a segmentation of the possible goods into disjoint clusters. We show that while the problem is computationally hard, it possesses constant approximation for natural families of distribution functions over bidders valuations for the goods.

Joint work with Yuval Emek, Iftah Gamzu and Moshe Tennenholtz

Kevin Leyton-Brown

(University of British Columbia)

Dominant-Strategy Auction Design for Agents with Uncertain, Private Values

We study the problem of designing auctions for agents who incur a cost if they choose to learn about their own preferences. We reformulate the revelation principle for use with such deliberative agents. Then we characterize the set of single-good auctions giving rise to dominant strategies for deliberative agents whose values are independent and private. Interestingly, this set of dominant-strategy mechanisms is exactly the set of sequential posted-price auctions, a class of mechanisms that has received much recent attention.

This talk will describe joint work with David R.M. Thompson

Yishay Mansour

(Tel Aviv University)

Welfare and Profit Maximization with Production Costs

Combinatorial Auctions are a central problem in Algorithmic Mechanism Design: pricing and allocating goods to buyers with complex preferences in order to maximize some desired objective (e.g., social welfare, revenue, or profit). The problem has been well-studied in the case of limited supply (one copy of each item), and in the case of digital goods (the seller can produce additional copies at no cost). Yet the case of resources---think oil, labor, computing cycles, etc.---neither of these abstractions is just right: additional supplies of these resources can be found, but only at a cost (where the marginal cost is an increasing function).

In this work, we initiate the study of the algorithmic mechanism design problem of combinatorial pricing under increasing marginal cost. The goal is to sell these goods to buyers with unknown and arbitrary combinatorial valuation functions to maximize either the social welfare, or our own profit; specifically we focus on the setting of \emph{posted item prices} with buyers arriving online. We give algorithms that achieve constant factor approximations for a class of natural cost functions---linear, low-degree polynomial, logarithmic---and that give logarithmic approximations for all convex marginal cost functions (along with a necessary additive loss). We show that these bounds are essentially best possible for these settings.

This is a joint work with Avrim Blum, Anupam Gupta and Ankit Sharma

Constantinos Daskalakis

(MIT)

On Optimal Multi-Dimensional Mechanism Design

We solve the optimal multi-dimensional mechanism design problem when either the number of bidders is a constant or the number of items is a constant. In the first setting, we need that the values of each bidder for the items are i.i.d., but allow different distributions for each bidder. In the second setting, we allow the values of each bidder for the items to be arbitrarily correlated, but assume that the bidders are i.i.d. For all $\epsilon > 0$, we obtain an efficient additive ϵ -approximation, when the value distributions are bounded, or a multiplicative $(1-\epsilon)$ -approximation when the value distributions are unbounded, but satisfy the Monotone Hazard Rate condition. When there is a single bidder, we generalize these results to non identically distributed value distributions.

Based on joint works with Yang Cai and S. Matt Weinberg

Shahar Dobzinski

(Cornell University)

Bounding the Power of Truthfulness

It is widely believed that the power of computationally-efficient truthful mechanisms is severely limited. However, up until recently the community was missing the tools to prove this statement. In this talk we will discuss why previous attempts failed and some recent progress in proving bounds on the power of computationally-efficient truthful mechanisms. Specifically, we will demonstrate two techniques for proving impossibilities in two important settings: multi-unit auctions and combinatorial auctions with submodular valuations.

Jason Hartline

(Northwestern University)

Truth or Envy?

We consider (profit maximizing) mechanism design in general settings that include, e.g., position auctions (for selling advertisements on Internet search engines) and single-minded combinatorial auctions. We analyze optimal envy-free pricing in these settings and give economic justification for using optimal envy-free revenue as a benchmark for prior-free mechanism design and analysis. In addition to its economic justification, the envy-free revenue has a very simple structure and a strong connection to incentive compatibility (a.k.a., truthfulness) constraints in mechanism design.

As a first example of the connection between envy-free pricing and incentive compatible mechanism design, because the structures of optimal pricings and optimal mechanisms are similar, we give a mechanism design reduction from structurally rich environments including position auctions (and environments with a matroid structure) to multi-unit auction environments (i.e., auctioning k identical units to n unit-demand agents). For instance, via this reduction we are able to extend all prior-free digital good auctions to position auctions with a factor of two of loss in the approximation factor.

As a second example we extend a variant of the random sampling auction to get constant approximations for general downward closed (i.e., if we can serve a given set of agents, we can serve any subset) settings.

Stefano Leonardi

(Sapienza University of Rome)

Single Valued Combinatorial Auctions with Budgets

We consider budget constrained combinatorial auctions where each bidder has a private value for each of the items in some subset of the items and an overall budget constraint. Such auctions capture adword auctions, where advertisers offer a bid for those adwords that (hopefully) target their intended audience, and advertisers also have budgets. It is known that even if all items are identical and all budgets are public it is not possible to be truthful and efficient. Our main result is a novel auction that runs in polynomial time, is incentive compatible, and ensures Pareto-optimality. The auction is incentive compatible with respect to the private valuations whereas the budgets and the sets of interest are assumed to be public knowledge. This extends the result of Dobzinski, Lavi and Nisan (FOCS 2008) for auctions of multiple identical items with budgets to single-valued combinatorial auctions.

Joint work with Amos Fiat, Jared Saia and Piotr Sankowski

Silvio Micali

(MIT)

Collusive Dominant-Strategy Truthfulness

Fifty years ago, Vickrey published his famous mechanism for auctioning a single good in limited supply. The main property of Vickrey's mechanism is efficiency in dominant strategies. In absence of collusion, this is a wonderful efficiency guarantee. We note, however, that collusion is far from rare in auctions, and if some colluders exist and have some wrong beliefs, then the Vickrey mechanism dramatically loses its efficiency. Accordingly, we put forward a new mechanism that guarantees efficiency by providing a richer set of strategies and ensuring that it is dominant for every player to reveal truthfully not only his own valuation, but also with whom he is colluding, if he is indeed colluding with someone else. (Our result meaningfully bypasses prior impossibility proofs.)

Joint work with Jing Chen

Anna Karlin

(University of Washington)

Selling in Exclusive Markets: Some Observations About Prior-Free Profit Maximization.

We study the problem of designing truthful mechanisms to maximize a seller profit without prior knowledge of buyers' valuations. We look at a very simple setting where the seller can sell to any number of buyers in one of two exclusive markets. We show that even in this simple setting, a strong benchmark proposed by Hartline and Roughgarden (STOC' 08) is unattainable - that is, no truthful in expectation mechanism can be constant competitive to this benchmark. On the other hand, competitiveness to a symmetrized version of this benchmark is attained by a very simple mechanism.

Joint work with Thach Nguyen and Yuval Peres

Amos Fiat

(Tel Aviv University)

Bribing a Guard and Related Issues

We describe how to give a truthful in expectation mechanism that, with high probability, can implement almost anything (with several restrictions). This can be viewed as a significant strengthening of the Abrue-Matsushima mechanisms. Whereas A-M mechanisms require that the private type of the agent be known to all agents and use iterative elimination of dominated strategies as a solution concept, our mechanism is (almost) dominant strategy truthful and assumes nothing about the state of knowledge of the other agents.

Joint work with Elias Koutsoupias and Angelina Vidali

Liad Blumrosen

(The Hebrew University)

Only Valuable Experts Can Be Valued

Suppose a principal Alice wishes to reduce her uncertainty regarding some future payoff. Consider a self-proclaimed expert Bob that may either be an informed expert knowing an exact (or approximate) distribution of a future random outcome that may affect Alice's utility, or an uninformed expert who knows nothing more than Alice does. Alice would like to hire Bob and solicit his signal. Her goal is to incentivize an informed expert to accept the contract and reveal his knowledge while deterring an uninformed expert from accepting the contract altogether.

Following a powerful negative result by Olszewski and Sandroni, we reexamine the notion of an expert and conclude that knowing some hidden variable (i.e., the description of the aforementioned distribution), does not make Bob an expert, or at least not a valuable expert. The premise of our paper is that if Alice only tries to incentivize experts which are valuable to her decision making then she can indeed screen them from uninformed experts. On a more technical level, we consider the case where Bob's signal about the distribution of a future event cannot be an arbitrary distribution but rather comes from some subset P of all possible distributions. We give rather tight conditions on P (which relate to its convexity), under which screening is possible. We formalize our intuition that if these conditions are not met then an expert is not guaranteed to be valuable. We give natural and arguably useful scenarios where indeed such a restriction on the distribution arise.

Joint work with Moshe Babaioff, Nicolas S. Lambert and Omer Reingold

Jeff Rosenschein

(The Hebrew University)

Tight Bounds for Strategyproof Classification

Strategyproof (SP) classification considers situations in which a decision-maker must classify a set of input points with binary labels, minimizing expected error. Labels of input points are reported by self-interested agents, who may lie so as to obtain a classifier more closely matching their own labels. These lies would create a bias in the data, and thus motivate the design of truthful mechanisms that discourage false reporting.

I'll present an overview of our work on strategyproof classification, and answer some questions left open by previous research on the topic, in particular regarding the best approximation ratio (in terms of social welfare) that an SP mechanism can guarantee for n agents. Our primary result is a lower bound of $3-2/n$ on the approximation ratio of SP mechanisms under the shared inputs assumption; this shows that the previously known upper bound (for uniform weights) is tight. The proof relies on a result from Social Choice theory, showing that any SP mechanism must select a dictator at random, according to some fixed distribution. We then show how different randomizations can improve the best known mechanism when agents are weighted, matching the lower bound with a tight upper bound. These results contribute both to a better understanding of the limits of SP classification, as well as to the development of similar tools in other, related domains such as SP facility location and judgment aggregation.

This is joint work with Reshef Meir, Shaul Almagor, and Assaf Michaely

Christos Papadimitriou

(UC Berkeley)

TBD

Workshop Participants

Itai Ashlagi	iashlagi@mit.edu	Sefi Naor	naor@cs.technion.ac.il
Yossi Azar	azar@tau.ac.dot.il	Ilan Nehama	ilan_n@cs.huji.ac.il
Moshe Babaioff	moshe@microsoft.com	Abraham Neyman	aneyman@math.huji.ac.il
Maria Florina Balcan	ninamf@cc.gatech.edu	Thach Nguyen	nthach@gmail.com
Yonatan Bazak	hubh@cs.huji.ac.il	Noam Nisan	noam@cs.huji.ac.il
Gil Ben-Zvi	benzvi@gmail.com	Raz Nissim	raznissim@gmail.com
Oren Ben-Zwi	orenb@technion.ac.il	Svetlana Olonetsky	svetlana.olonetsky@gmail.com
Kshipra Bhawalkar	kshipra@stanford.edu	Sigal Oren	sigal@cs.cornell.edu
Florian Biermann	florianb@mssc.huji.ac.il	Asu Ozdaglar	asuman@mit.edu
Sushil Bikhchandani	sbikhcha@anderson.ucla.edu	Renato Paes-Leme	renatoppl@cs.cornell.edu
Avrim Blum	avrim@cs.cmu.edu	Christos Papadimitriou	christos@cs.berkeley.edu
Liad Blumrosen	blumrosen@huji.ac.il	Katia Papakonstantinopoulou	katia@di.uoa.gr
Ronen Brafman	brafman@cs.bgu.ac.il	Ron Peretz	ronprtz@gmail.com
Michael Brautbar	brautbar@cis.upenn.edu	Georgios Piliouras	georgios.piliouras@gmail.com
Elisa Celis	ecelis@u.washington.edu	Maria Polukarov	mp3@ecs.soton.ac.uk
Jing Chen	jingchen@csail.mit.edu	Emmanouil Pountourakis	malgor21@gmail.com
Yiling Chen	yiling@eecs.harvard.edu	George Rabanca	grabanca@gmail.com
Liron Cohen	lirchi@gmail.com	Michael O. Rabin	rabin@seas.harvard.edu
Vincent Conitzer	conitzer@cs.duke.edu	Zinovi Rabinovich	zinovi.rabinovich@gmail.com
Endre Csóka	csokaendre@gmail.com	Tom Ron	rontom@gmail.com
Constantinos Daskalakis	costis@csail.mit.edu	Amir Ronen	amirro@il.ibm.com
Shahar Dobzinski	dobzin@gmail.com	Jeffrey S. Rosenschein	jeff@cs.huji.ac.il
Yair Dombb	yairbiu@gmail.com	Tim Roughgarden	tim@cs.stanford.edu
Edith Elkind	eedith@gmail.com	Dmitrijs Rutko	dim_rut@inbox.lv
Uriel Feige	uriel.feige@weizmann.ac.il	Roy Schwartz	schwartz@cs.technion.ac.il
Michal Feldman	mfeldman@huji.ac.il	Ella Segev	ellasgv@bgu.ac.il
Moran Feldman	moranfe@cs.technion.ac.il	Shai Shalev-Shwartz	shais@cs.huji.ac.il
Amos Fiat	fiat@tau.ac.il	Mordechai Shalom	cmshalom@cs.technion.ac.il
Lisa Fleischer	lkf@cs.dartmouth.edu	Or Sheffet	osheffet@cs.cmu.edu
Kobi Gal	kobig@bgu.ac.il	Yaron Singer	yaron@cs.berkeley.edu
Ratnik Gandhi	ratnik@tifr.res.in	Alexander Skopalik	ASkopalik@ntu.edu.sg
Jugal Garg	jugal@cse.iitb.ac.in	Rann Smorodinsky	rann@ie.technion.ac.il
Paul W. Goldberg	P.W.Goldberg@liverpool.ac.uk	Lampros Stavrogiannis	ls8g09@ecs.soton.ac.uk
Sharon Goldberg	goldbe@cs.bu.edu	Vasilis Syrgkanis	vasilis@cs.cornell.edu
Tali Gutman	gutmant@gmail.com	Omer Tamuz	omertamuz@gmail.com
Sergiu Hart	hart@huji.ac.il	Eva Tardos	eva@cs.cornell.edu
Jason Hartline	hartline@eecs.northwestern.edu	Moshe Tennenholtz	moshet@ie.technion.ac.il
Avinatan Hassidim	avinatah@gmail.com	David Thompson	daveth@cs.ubc.ca
Liat Hayun	liat.hayun@gmail.com	Shai Vardi	shaivardi@gmail.com
Pavel Hubacek	hubacek@cs.au.dk	Paul Varkey	pvarkey@cs.uic.edu
Nicole Immorlica	nicimm@gmail.com	Vijay Vazirani	vazirani@cc.gatech.edu
Albert Xin Jiang	jiang@cs.ubc.ca	Angelina Vidali	angvid@gmail.com
Haim Kaplan	haimk@post.tau.ac.il	Ilya Volkovich	ilyav@cs.technion.ac.il
Anna Karlin	karlin@cs.washington.edu	Ariella Voloshin	ariella.voloshin@gmail.com
Erez Karpas	karpase@technion.ac.il	Yoav Wilf	yoavwilf@gmail.com
Elias Koutsoupias	elias@di.uoa.gr	Shmuel Zaks	zaks@cs.technion.ac.il
Ron Lavi	ronlavi@ie.technion.ac.il	Roie Zivan	zivan.roie@gmail.com
Daniel Lehmann	lehmann@cs.huji.ac.il	Aviv Zohar	avivz@microsoft.com
Dariusz Leniowski	d.leniowski@mimuw.edu.pl	Michael Zuckerman	michael.zuckerman@mail.huji.ac.il
Stefano Leonardi	leon@dis.uniroma1.it		
Omer Lev	omerl@cs.huji.ac.il		
Kevin Leyton-Brown	jiang@cs.ubc.ca		
Edo Liberty	edo.liberty@gmail.com		
Katrina Ligett	katrina@cs.cornell.edu		
Nathan Linial	nati@cs.huji.ac.il		
Yishay Mansour	mansour@tau.ac.il		
Avishay Maya	avishaym@gmail.com		
Ruta Mehta	ruta@cse.iitb.ac.in		
Reshef Meir	reshef24@gmail.com		
Silvio Micalli	silvio@csail.mit.edu		
Dov Monderer	dov@ie.technion.ac.il		
Jamie Morgenstern	jamiemmt@cs.cmu.edu		
Elchanan Mossel	mossel@stat.berkeley.edu		
Daniil Musatov	musatych@gmail.com		