Report from Dagstuhl Seminar 13161

Interface of Computation, Game Theory, and Economics

Edited by

Sergiu Hart[1], Éva Tardos[2], and Bernhard von Stengel[3]

1 **Hebrew University, Jerusalem, IL,** hart@huji.ac.il
2 **Cornell University – Ithaca, US,** eva@cs.cornell.edu
3 **London School of Economics, GB,** stengel@nash.lse.ac.uk

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 13161 "Interface of Computation, Game Theory, and Economics". The workshop was strongly interdisciplinary, on the leading edge of current topics generally connected to algorithmic game theory: Mechanism design and auctions, interactions in networks, social models, and dynamics and equilibrium in games and markets. We summarize these topics, give the talk abstracts, and comment on experiences related to the organization of the workshop.

Seminar 14.–19. April, 2013 – www.dagstuhl.de/13161
1998 ACM Subject Classification J.4.1. Computer Applications – Economics
Keywords and phrases Algorithmic Game Theory, Economics, Internet, Nash Equilibrium, Mechanism Design, Auctions
Digital Object Identifier 10.4230/DagRep.3.4.69

1 Executive Summary

Sergiu Hart
Éva Tardos
Bernhard von Stengel

The aim of this seminar was to study research issues at the interface of computing, game theory and economics. It facilitated discussions among people working in different disciplines. The majority of participants were academics from computer science departments, and the others (about one third) from other disciplines such as economics or corporate research departments of Google or Microsoft. All have strong cross-disciplinary interests.

Economic transactions on the internet are of ever-increasing importance. In order to execute and support them algorithmically, it is important to understand the agents' incentives on one hand and computational constraints on the other hand. This is studied in approaches to mechanism design and auctions, which formed a large part of the topics of this workshop.

Theoretical and practical issues of *mechanism design* were topics of the following presentations: epistemic implementations with belief levels (Jing Chen), translating agent-provided inputs to optimization (Constantinos Daskalakis), reward schemes (Shahar Dobzinski), the difficulties of allocating more than one good (Sergiu Hart), advertisement exchanges (Vahab Mirrokni), mechanisms for the private supply of a public good (Rudolf Müller), truthfulness versus privacy (Aaron Roth), composing mechanisms (Vasilis Syrgkanis), and allocating indivisible objects (Rakesh V. Vohra).

Interface of Computation, Game Theory, and Economics, *Dagstuhl Reports*, Vol. 3, Issue 4, pp. 69–90
Editors: Sergiu Hart, Éva Tardos, and Bernhard von Stengel

Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Aspects of *auctions* concerned "expressiveness" about preferences (Paul Dütting), the approximate optimality of marginal revenue maximization (Jason D. Hartline), improving the design of online advertising auctions (Kevin Leyton-Brown), commitment (Katrina Ligett), inefficiency of multi-unit auctions (Vangelis Markakis), symmetric auctions (Mallesh Pai), interdependent values (Tim Roughgarden), and spectrum auctions (Ilya Segal).

Understanding the interconnectedness of complex economic systems requires models and theories for the underlying *network* structures and their dynamics. Networks were studied with respect to social segregation (Nicole Immorlica), practical market applications (Ramesh Johari), online creation (Thomas Kesselheim), competition (Brendan Lucier), and social contagion (Sigal Oren).

Social models, with bridges to mechanism design, were studied in presentations on division protocols (Simina Branzei), randomized social choice (Markus Brill), ranking methods (Gabrielle Demange), power changes in voting games (Edith Elkind), and incentives beyond selfishness (Guido Schäfer).

Achieving and computing an equilibrium in *dynamic models of large interactions* such as games and market models was studied for large aggregative games (Yakov Babichenko), new price updating in markets (Nikhil R. Devanur), payoff queries for games (Paul W. Goldberg), limit processes for evolutionary games (Bill Sandholm), and tournament competitions (Bernhard von Stengel).

The topics were chosen by the presenters, not by the organizers. The rather strong emphasis on mechanism design and auctions (which may have caused one single critical feedback comment on "too much groupthink") reflects the strong current interest in this area, in line with its economic importance, for example as the source of the riches of Google and other internet search engines.

2 Table of Contents

3 Overview of Talks

3.1 Best-Reply Dynamics in Large Aggregative Games

Yakov Babichenko (CalTech, US, babich@caltech.edu)

We consider small-influence aggregative games with a large number of players n. For this class of games we present a best-reply dynamic with the following two properties. First, the dynamic reaches Nash approximate equilibria fast (in at most $cn\log(n)$ steps for some constant $c > 0$). Second, Nash approximate equilibria are played by the dynamic with a limit frequency that is exponentially close to 1 (at least $1 - e^{-c'n}$ for some constant $c' > 0$).

3.2 Equilibria of Generalized Cut and Choose Protocols

Simina Branzei (Aarhus University, DK, simina@cs.au.dk)

Classic cake cutting protocols – which fairly allocate a divisible good among agents with heterogeneous preferences – are susceptible to manipulation. Do their strategic outcomes still guarantee fairness? We model the interaction among agents as a game and study its Nash equilibria. We show that each protocol in the novel class of generalized cut and choose protocols – which includes the most important discrete cake cutting protocols – is guaranteed to have an ε-equilibrium for all $\varepsilon > 0$. Moreover, we observe that the (approximate) equilibria of proportional protocols – which guarantee each of the n agents a $1/n$-fraction of the cake – must be (approximately) proportional. Finally, we design a generalized cut and choose protocol where all equilibrium outcomes satisfy the stronger fairness notion of envy-freeness.

3.3 On the Tradeoff between Economic Efficiency and Strategyproofness in Randomized Social Choice

Markus Brill (TU München, DE, brill@in.tum.de)

Two fundamental notions in microeconomic theory are efficiency – no agent can be made better off without making another one worse off – and strategyproofness – no agent can obtain a more preferred outcome by misrepresenting his preferences. When social outcomes are probability distributions (or lotteries) over alternatives, there are varying degrees of these notions depending on how preferences over alternatives are extended to preference over lotteries. We show that efficiency and strategyproofness are incompatible to some extent when preferences are defined using stochastic dominance (SD) and therefore introduce a natural weakening of SD based on Savage's sure-thing principle (ST). While random serial dictatorship is SD-strategyproof, it only satisfies ST-efficiency. Our main result is that strict maximal lotteries – an appealing class of social decision schemes due to Kreweras and Fishburn – satisfy SD-efficiency and ST-strategyproofness.

3.4 Epistemic Implementation

Jing Chen (IAS and Stony Brook, US, jingchen@csail.mit.edu)

In a setting of incomplete information, we model the hierarchy of the players' beliefs about each other's payoff types in a set-theoretic way. A player's beliefs can be totally arbitrary, and the beliefs of different players can be inconsistent with each other. In single-good auctions, for $k = 0, 1, \ldots$, we define a revenue benchmark G^k on the players' belief hierarchy. Intuitively, $G^k \geq v$ if and only if there exist at least two players "believing that there exists a player ..." (k times) valuing the good at least v. We construct an interim individually rational mechanism M that, without any clue about the players' beliefs and their rationality level, virtually guarantees revenue G^k whenever the players happen to be level-$(k+1)$ rational. We also separate the revenue achievable with level-k and level-$(k+1)$ rational players. For every $k \geq 0$, we show that no interim individually rational mechanism can virtually guarantee revenue G^k when the players' rationality level is k instead of $k+1$.

3.5 Reductions from Mechanism to Algorithm Design

Constantinos Daskalakis (MIT, US, costis@csail.mit.edu)

Algorithmic mechanism design centers around the following question: How much harder is optimizing an objective function over inputs that are furnished by rational agents compared to when the inputs are known? We provide a computationally efficient, black-box reduction from mechanism design (i.e. optimizing over rational inputs) to algorithm design (i.e. optimizing over known inputs) in general Bayesian settings. As an application of our reduction, we extend Myerson's celebrated auction to the multi-item setting.

3.6 A Ranking Method Based on Handicaps

Gabrielle Demange (Paris School of Economics, FR, demange@pse.ens.fr)

Ranking methods are a fundamental tool in many areas. Popular methods are based on the statements of "experts" and aggregate them in some way. As such, there is a variety of meaningful ranking methods, more or less adapted to the environment under consideration. We introduce and characterizes a new method, called the handicap-based method. The method assigns scores to the items and weights to the experts. Scores and weights form an equilibrium for a relationship based on the notion of handicaps. The method is, in a sense that we make precise, the counterpart of the counting method in environments that require intensity-invariance. Intensity-invariance is a desirable property when the intensity of the experts' statements has to be controlled. Otherwise, both the counting and handicap-based methods satisfy a property called homogeneity, which is a desirable property when cardinal statements matter, as is the case in many applications.

3.7 Tatonnement Beyond Gross Substitutes? Gradient Descent to the Rescue

Nikhil R. Devanur (Microsoft, Redmond, US, nikdev@microsoft.com)

Main reference Y.K. Cheung, R. Cole, N.R. Devanur, "Tatonnement beyond gross substitutes?: Gradient descent to the rescue," in Proc. of the 45th Annual ACM Symp. on Symposium on Theory of Computing (STOC'13), pp. 191–200, ACM, 2013.
URL http://dx.doi.org/10.1145/2488608.2488633

Tatonnement is a simple and natural rule for updating prices in exchange (Arrow-Debreu) markets. We define a class of markets for which tatonnement is equivalent to gradient descent. This is the class of markets for which there is a convex potential function whose gradient is always equal to the negative of the excess demand and we call it Convex Potential Function (CPF) markets. We show the following results. CPF markets contain the class of Eisenberg Gale (EG) markets, defined previously by Jain and Vazirani. The subclass of CPF markets for which the demand is a differentiable function contains exactly those markets whose demand function has a symmetric negative semi-definite Jacobian. We define a family of continuous versions of tatonnement based on gradient descent using a Bregman divergence. As we show, all processes in this family converge to an equilibrium for any CPF market. This is analogous to the classic result for markets satisfying the Weak Gross Substitutes property. A discrete version of tatonnement converges toward the equilibrium for the following markets of complementary goods; its convergence rate for these settings is analyzed using a common potential function. Fisher markets in which all buyers have Leontief utilities. The tatonnement process reduces the distance to the equilibrium, as measured by the potential function, to an ε fraction of its initial value in $O(1/\varepsilon)$ rounds of price updates. Fisher markets in which all buyers have complementary CES utilities. Here, the distance to the equilibrium is reduced to an ε fraction of its initial value in $O(\log(1/\varepsilon))$ rounds of price updates.

This shows that tatonnement converges for the entire range of Fisher markets when buyers have complementary CES utilities, in contrast to prior work, which could analyze only the substitutes range, together with a small portion of the complementary range.

3.8 Shared Resource Management via Reward Schemes

Shahar Dobzinski (Weizmann Institute, Rehovot, IL, dobzin@gmail.com)

Joint work of Dobzinski, Shahar; Ronen, Amir;

We study scenarios in which consumers have several options of using a shared resource (e.g., truck operators that can drive either in peak or off peak hours). Our goal is to design reward schemes that, in equilibrium, minimize the cost to society and the total sum of rewards. We introduce a simple reward scheme which does not require any knowledge of the private values of the consumers, yet its cost in equilibrium is always within a factor of $\sqrt{\alpha}$ of the cost of the optimal scheme that has complete knowledge of the consumers' valuations. Here, alpha is the ratio between the costs of the worst and best alternatives. We show the optimality of our scheme in various settings by providing lower bounds.

3.9 Expressiveness and Robustness of First-Price Position Auctions

Paul Dütting (EPFL Lausanne, CH, paul.duetting@epfl.ch)

It has been argued that increasing the expressiveness of an auction increases the quality of the outcomes it is able to support. Intuitively, more expressive auctions should allow agents to more accurately reveal their preferences in the presence of uncertainty. We study this issue in the context of a position auction in which valuations are one-dimensional but the designer is uncertain about the relative values of the positions. In this setting, efficient equilibria may fail to exist for simplified auctions that solicit only a single bid from each agent, but existence can be restored by increasing expressiveness. In particular, we show this to be the case for a generalized first-price (GFP) auction. In addition to the existence of an efficient Bayes-Nash equilibrium, the GFP auction is robust to varying assumptions about the information available to agents while second-price and VCG auctions are not. Technically, our main result is interesting because the Bayes-Nash equilibrium is constructed for a multi-dimensional bid space. The structure of the equilibrium bids moreover provides an intuitive explanation for why first-price payment rules may be able to support equilibria in a wider range of scenarios than second-price payment rules.

3.10 Dynamic Coalitional Games

Edith Elkind (Nanyang Technical University, Singapore, SG, eelkind@ntu.edu.sg)

Joint work of Elkind, Edith; Pasechnik, Dmitrii V.; Zick, Yair;
Main reference E. Elkind, D.V. Pasechnik, Y. Zick, "Dynamic weighted voting games," in Proc. of the Int'l Conf. on Autonomous Agents and Multi-Agent Systems (AAMAS'13), pp. 515-522, IFAAMAS, 2013.
URL http://dl.acm.org/citation.cfm?id=2485003

We define and study dynamic weighted voting games – weighted voting games where the weight of each player and the quota may change as a function of time. We investigate computational aspects of such games under the assumption that all weights and the quota are given by polynomials with integer coefficients. We focus on two types of algorithmic questions: computing a given solution concept at a particular point in time, and checking that a certain function of the game (e.g., the Shapley value of a given player or the value of the least core) remains within given bounds during a particular time interval. We provide pseudopolynomial algorithms for both types of problems, for a variety of solution concepts. We then use our results to investigate the changes in power distribution in the Council of the European Union over the next 50 years.

3.11 Payoff Queries

Paul W. Goldberg (University of Liverpool, UK, P.W.Goldberg@liverpool.ac.uk)

I give an informal introduction to "payoff query" algorithms, where an algorithm can specify a pure profile of a game (with initially unknown payoffs) and get told the payoffs for that pure profile. Given a class of games, and a solution concept, the challenge is to figure out the query complexity of solving an initially-unknown game from that class.

3.12 Two(!) Good To Be True

Sergiu Hart (Hebrew University, Jerusalem, IL, hart@huji.ac.il)

How to sell goods optimally? While the mechanism-design literature has solved this problem neatly when there is only one good, the multiple goods case turns out to be extremely difficult, mathematically and conceptually. Much of what is true for one good does not extend to multiple goods. We will try to explain the difficulties, show what can go wrong, and then present some universal approximation results. The talk is essentially self-contained; no background in mechanism design is necessary.

3.13 The Simple Economics of Approximately Optimal Auctions

Jason D. Hartline (Northwestern University, Evanston, US, hartline@eecs.northwestern.edu)

The intuition that profit is optimized by maximizing marginal revenue is a guiding principle in microeconomics. In the classical auction theory for agents with quasi-linear utility and single dimensional preferences, Bulow and Roberts (1989) show that the optimal auction of Myerson (1981) is in fact optimizing marginal revenue. In particular Myerson's virtual values are exactly the derivative of an appropriate revenue curve. We consider mechanism design in environments where the agents have multi-dimensional and non-linear preferences. Understanding good auctions for these environments is considered to be the main challenge in Bayesian optimal mechanism design. In these environments maximizing marginal revenue may not be optimal, and furthermore, there is sometimes no direct way to implementing the marginal revenue maximization mechanism. Our contributions are twofold: we give procedures for implementing marginal revenue maximization in general, and we show that marginal revenue maximization is approximately optimal. Our approximation factor smoothly degrades in a term that quantifies how far the environment is from an ideal one (i.e., where marginal revenue maximization is optimal).

3.14 An Analysis of One-Dimensional Schelling Segregation

Nicole Immorlica (Northwestern University, Evanston, US, nickle@eecs.northwestern.edu)

We analyze the Schelling model of segregation in which a society of n individuals live in a ring. Each individual is one of two races and is only satisfied with his location so long as at least half his $2w$ nearest neighbors are of the same race as him. In the dynamics, randomly-chosen unhappy individuals successively swap locations. We consider the average size of monochromatic neighborhoods in the final stable state. Our analysis is the first rigorous analysis of the Schelling dynamics. We note that, in contrast to prior approximate analyses, the final state is nearly integrated: the average size of monochromatic neighborhoods is independent of n and polynomial in w.

3.15 The Engineer as Economist: The Design of Online Market Platforms

Ramesh Johari (Stanford University, US, ramesh.johari@stanford.edu)

Markets are an ancient institution for matching those willing to provide a good or service with those who want it. Physical markets were typically slow to evolve, with simple institutions governing trade, and trading partners generally facing a daunting challenge in finding the "right" partner.

Online marketplaces dramatically disrupt this tradition. Such markets – from eBay, to Google's sponsored search auction, to online labor markets such as oDesk and TaskRabbit – can rapidly respond to evolving market trends, and "engineer" in fine grained ways the interactions of their participants with the platform. Further, the traditional difficulty of *finding* even one trading partner has been replaced with a new difficulty: how to narrow down a plethora of choices? Motivated by this new landscape, this talk will discuss some of the challenges that engineers face in designing and implementing emerging online marketplaces.

3.16 Online Independent Set Beyond the Worst-Case: Secretaries, Prophets, and Periods

Thomas Kesselheim (Cornell University, US, kesselheim@cs.cornell.edu)

We investigate online algorithms for maximum (weight) independent set on graph classes with bounded inductive independence number like interval and disk graphs with applications to, e.g., task scheduling and spectrum allocation. In the online setting, it is assumed that nodes of an unknown graph arrive one by one over time. An online algorithm has to decide whether an arriving node should be included into the independent set. Unfortunately, this natural and practically relevant online problem cannot be studied in a meaningful way within a classical competitive analysis as the competitive ratio on worst-case input sequences is lower bounded by $\Omega(n)$. This devastating lower bound holds even for randomized algorithms on unweighted interval graphs and, hence, for the most restricted graph class under consideration.

As a worst-case analysis is pointless, we study online independent set in a stochastic analysis. Instead of focussing on a particular stochastic input model, we present a generic sampling approach that enables us to devise online algorithms achieving performance guarantees for a variety of input models. In particular, our analysis covers stochastic input models like the secretary model, in which an adversarial graph is presented in random order, and the prophet-inequality model, in which a randomly generated graph is presented in adversarial order. Our sampling approach bridges thus between stochastic input models of quite different nature. In addition, we show that the same performance guarantees can be obtained for a period-based input model that is inspired by practical admission control applications.

Our sampling approach yields an online algorithm for maximum independent set on interval and disk graphs with competitive ratio $O(1)$ with respect to all of the mentioned stochastic input models. More generally, for graph classes with inductive independence number ρ, the competitive ratio is $O(\rho^2)$. The approach can be extended towards maximum

weight independent set by losing only a factor of $O(\log n)$ in the competitive ratio with n denoting the (expected) number of nodes. This upper bound is complemented by a lower bound of $\Omega(\log n/\log^2 \log n)$ showing that our sampling approach achieves nearly the optimal competitive ratio in all of the considered models. Furthermore, we generalize our analysis to address several practically motivated extensions of the independent set problem, e.g., arrival and departure times of nodes or edge-weighted graphs capturing SINR-type interference conflicts in wireless networks.

3.17 Revenue Optimization in the Generalized Second-Price Auction

Kevin Leyton-Brown (University of British Columbia, Vancouver, CA, kevinlb@cs.ubc.ca)

We consider the question of how to optimize revenue in advertising auctions without departing from the generalized second-price. We consider several different GSP variants (including squashing and different types of reserve prices), and how to set their parameters optimally. Our main finding is that unweighted reserve prices (i.e., where each advertiser has the same per-click reserve price) are dramatically better than the quality-weighted reserve prices that have become common practice in the last few years. This result is extremely robust, arising from theoretical analysis as well as multiple computational experiments. Our work also includes one of the first studies of how squashing and reserve prices interact, and of how equilibrium selection affects the revenue of GSP when features such as reserves or squashing are applied.

3.18 Preplay Commitment in First-Price Auctions

Katrina Ligett (CalTech, Pasadena, US, katrina@caltech.edu)

We study a variation of the standard single-item sealed-bid first-price auction wherein all bidders know one another's valuations, and one bidder (the leader) publicly commits to a (possibly mixed) strategy before the others submit their bids. We formulate the auction as a two-stage Stackelberg game, and study the impact of commitment on the utilities of the bidders and the auctioneer. For the case where the leader's valuation is the highest or the second highest (including, e.g., when there are only two bidders), we characterize the commitment that maximizes the expected payoff of the leader. In this case, both the leader and the bidder with the highest valuation among the other bidders strictly benefit from the commitment—each obtains an expected payoff higher than that achieved at a Nash equilibrium of the standard first-price auction. For an important variant of our model where the leader's commitment is restricted to be a discrete random variable (and thus a credible commitment may be more practically implemented), we characterize the leader's optimal commitment as a solution to an optimization problem. There, we study the extent to which a discrete-valued commitment can approximate the maximum expected payoff achievable under committing to arbitrary mixed strategies.

3.19 A Model of Bertrand Price Competition in Networks

Brendan Lucier (Microsoft Research, Cambridge, US, brlucier@microsoft.com)

We study scenarios where multiple sellers of a homogeneous good compete on prices, where each seller can only sell to some subset of the buyers. Crucially, sellers cannot price-discriminate between buyers. We model the structure of the competition by a graph (or hyper-graph), with nodes representing the sellers and edges representing populations of buyers. We study equilibria in the game between the sellers, prove that they always exist, and present various structural, quantitative, and computational results about them. We also analyze the equilibria completely for a few cases. Many questions are left open.

3.20 On the Inefficiency of Standard Multi-unit Auction Formats

Vangelis Markakis (Athens University of Economics and Business, GR, markakis@gmail.com)

We study two standard multi-unit auction formats for allocating multiple units of a single good to multi-demand bidders. The first one is the Discriminatory Price Auction, which charges every winner his winning bids. The second is the Uniform Price Auction, which determines a uniform price to be paid per unit. Variants of both formats find applications ranging from the allocation of bonds to investors, to online sales over the internet, facilitated by popular online brokers.

For these multi-unit auction formats, we consider two bidding interfaces: (i) standard bidding, which is most prevalent in the scientific literature, and (ii) uniform bidding, which is the most widely used interface in practical applications. We evaluate the economic inefficiency of the two formats for both bidding interfaces, by means of upper and lower bounds on the Price of Anarchy for pure equilibria and mixed Bayes-Nash equilibria. Our results for bidders with submodular valuations improve upon bounds that have been obtained recently in [Markakis, Telelis, SAGT 2012] and [Syrgkanis, Tardos, STOC 2013]. Moreover, we also consider for the first time bidders with subadditive valuation functions and obtain constant upper bounds there as well.

3.21 Mechanism Design Problems in Ad Exchanges and Budget Constraints

Vahab Mirrokni (Google, New York, US, mirrokni@google.com)

I will give a survey of mechanism design problems motivated by ad exchanges and budget constraints in online advertising. For each problem, I will present preliminary/known results and pose open problems and research directions. Some topics that are discussed are auctions in the presence of intermediaries, optimal revenue-sharing double auctions, and pareto-optimal polyhedral clinching auctions.

3.22 Optimal Mechanism Design for the Private Supply of a Public Good

Rudolf Müller (Maastricht University, NL, r.muller@maastrichtuniversity.nl)

We study the problem of finding the profit-maximizing mechanism for a monopolistic provider of a single, non-excludable public good. This problem has been well studied for the case when agents' signals are independently distributed, but the literature is almost silent about the case of general joint distributions. Our model covers the most general setting, namely, we allow for correlation in the signal distribution as well as for informational externalities. We investigate the problem from an automated mechanism design perspective, meaning that we want to understand the algorithmic complexity of finding the optimal mechanism when we are given a finite set of signal profiles and their distribution.

We show that the optimal deterministic, ex-post incentive compatible, ex-post individual rational mechanism can be computed in polynomial time by reducing the problem to finding a maximal weight closure in a directed graph. Node weights in the graph correspond to conditional virtual values. When valuations are independent and independently distributed, the constructed mechanism is also optimal among all Bayes-Nash implementable and interim individual rational mechanisms. In contrast, for dependent valuations strictly higher profit can be achieved if one allows for interim individual rationality or Bayes-Nash implementability. By invoking techniques due to Cremer and McLean [1988], we show that optimal deterministic, interim individual rational, ex-post implementable mechanisms still can be found in polynomial time if the joint distribution of signals satisfies certain regularity conditions. Finally, we demonstrate that our techniques can be adapted for the excludable public good problem as well.

3.23 Selection and Influence in Cultural Dynamics

Sigal Oren (Cornell University, US, sigal@cs.cornell.edu)

One of the fundamental principles driving diversity or homogeneity in domains such as cultural differentiation, political affiliation, and product adoption is the tension between two forces: influence (the tendency of people to become similar to others they interact with) and selection (the tendency to be affected most by the behavior of others who are already similar). Influence tends to promote homogeneity within a society, while selection frequently causes fragmentation. When both forces are in effect simultaneously, it becomes an interesting question to analyze which societal outcomes should be expected.

In order to study the joint effects of these forces more formally, we analyze a natural model built upon active lines of work in political opinion formation, cultural diversity, and language evolution. Our model posits an arbitrary graph structure describing which "types" of people can influence one another: this captures effects based on the fact that people are only influenced by sufficiently similar interaction partners. In a generalization of the model, we introduce another graph structure describing which types of people even so much as come

in contact with each other. These restrictions on interaction patterns can significantly alter the dynamics of the process at the population level.

For the basic version of the model, in which all individuals come in contact with all others, we achieve an essentially complete characterization of (stable) equilibrium outcomes and prove convergence from all starting states. For the other extreme case, in which individuals only come in contact with others who have the potential to influence them, the underlying process is significantly more complicated; nevertheless we present an analysis for certain graph structures.

3.24 Symmetric Auctions

Mallesh Pai (University of Pennsylvania, US, mallesh@econ.upenn.edu)

Real-world auctions are often restricted to being symmetric (anonymous and nondiscriminatory) due to practical or legal constraints. We examine when this restriction prevents a seller from achieving his objectives. In an independent private value setting, we characterize the set of incentive compatible and individually rational outcomes that can be implemented via a symmetric auction. Our characterization shows that symmetric auctions can yield a large variety of discriminatory outcomes such as revenue maximization and affirmative action. We also characterize the set of implementable outcomes when individual rationality holds in an ex post rather than an interim sense. This additional requirement may prevent the seller from maximizing revenue.

3.25 Mechanism Design in Large Games and Differential Privacy

Aaron Roth (University of Pennsylvania, US, aaroth@cis.upenn.edu)

We study the design of mechanisms satisfying two desiderata – incentive compatibility and privacy. The first, requires that each agent should be incentivized to report her private information truthfully. The second, privacy, requires the mechanism not reveal "much" about any agent's type to other agents. We propose a notion of privacy we call Joint Differential Privacy. It is a variant of Differential Privacy, a robust notion of privacy used in the Theoretical Computer Science literature. We show by construction that such mechanisms, i.e. ones which are both incentive compatible and jointly differentially private exist when the game is "large", i.e., there are a large number of players, and any player's action affects any other's payoff by at most a small amount. Our mechanism adds carefully selected noise to no-regret algorithms similar to those studied in Foster-Vohra and Hart-Mas-Colell. It therefore implements an approximate correlated equilibrium of the full information game induced by players' reports.

3.26 Optimal Ex Post and Prior-Independent Auctions with Interdependent Values

Tim Roughgarden (Stanford University, US, tim@cs.stanford.edu)

We study optimal and approximately-optimal mechanism design questions in the interdependent values model, which generalizes the standard setting of independent and private values. We focus our attention on ex post incentive compatible and individually rational mechanisms, and develop an analog of Myerson's optimal auction theory that applies to many interdependent settings of interest. We demonstrate two applications for specific interdependent settings: First, a parallel result to the well-known optimality of the second-price auction with reserve for i.i.d. bidders, where the English auction replaces the second-price one. Second, we identify good prior-independent auctions – auctions with near-optimal expected revenue across a wide range of priors – for certain interdependent value settings.

3.27 Large Deviations and Stochastic Stability in the Large Population Limit

Bill Sandholm (University of Wisconsin, Madison, US, whs@ssc.wisc.edu)

In this talk I will give an overview about new tools for equilibrium selection in games in the framework of stochastic evolutionary game theory. In this project we investigate stochastic stability theory from a new angle. We elaborate on the precise role of the parameters in this game theoretic model, which are the population size and the level of noise in the agents' updating decisions. Stochastic stability theory is concerned with understanding the long-run properties of the stochastic evolutionary game dynamics when these parameters are taken to their respective limits separately, or simultaneously. For each possible way of taking limits, we present the appropriate technique to understand the long-run of the game dynamics. This requires a novel and interesting combination of various mathematical techniques, such as large deviations theory and optimal control. We also discuss the computational problem of stochastic stability in simple game settings.

3.28 Altruism and Spite in Games

Guido Schäfer (Guido Schäfer, CWI, Amsterdam, NL, G.Schaefer@cwi.nl)

In most game-theoretical studies it is assumed that the decision makers base their decisions on purely selfish grounds. This assumption is in stark contrast with a large body of research in experimental economics and the social sciences, which suggest that decision makers are often motivated by other-regarding preferences such as altruism, spite or fairness. Very little attention has been given to the analysis of the impact of such alternative behaviors. In this talk, we review some recent advances in the study of the inefficiency of equilibria when players are (partially) altruistic or spiteful and highlight a few counter-intuitive results.

3.29 Heuristic Auctions and U.S. Spectrum Repurposing

Ilya Segal (Stanford University, US, ilya.segal@stanford.edu)

We examine a novel class of procurement auctions for single-minded bidders, in which the auctioneer selects a set of bids to be accepted subject to complicated feasibility constraints that preclude optimization-based winner determination. (This setting is inspired by the U.S. Federal Communication Commission's problem of buying out a subset of broadcast TV licenses to clear spectrum for broadband use while retuning the remaining stations into the remaining TV spectrum subject to interference constraints and a possible budget constraint.) Instead, we propose a class of computationally feasible "greedy deferred-acceptance heuristic" auctions for calculating both a feasible set of winning bids and "threshold" payments which induce strategy-proof bidding. The calculation iteratively rejects the "highest-scoring" bid that could still be feasibly rejected, with a bidder's score based on its bid value and possibly on the bids already rejected. (The latter dependence could be used to ensure that the total "threshold payments" satisfy the auctioneer's budget constraint, or that the winners do not get excessive payments.) This class of "deferred acceptance" heuristic auctions differs from the previously studied "deferred rejection/instance acceptance" heuristic auctions. In particular, we show that deferred-rejection heuristic auctions with threshold payments: (1) are equivalent to clock auctions with descending bidder-specific prices in which bidders who haven't quit are acquired at their final clock prices; (2) are (weakly) group strategy-proof, and so are the corresponding clock auctions for any information disclosure policy; (3) are outcome-equivalent to their paid-as-bid counterparts with the same allocation rule under full information: A paid-as-bid heuristic auction has a Nash equilibrium with the same outcome as its threshold-auction counterpart, which is a unique outcome surviving iterated deletion of weakly dominated strategies. In contrast, the Vickrey auction generally fails properties (1)–(3), except when bidders are substitutes, in which it can be implemented as a heuristic or clock auction.

3.30 Composable and Efficient Mechanisms

Vasilis Syrgkanis (Cornell University, US, vasilis@cs.cornell.edu)

Online markets require simple, and well-designed systems that work well even if users participate in multiple ones in parallel. Traditional mechanism design has considered mechanisms only in isolation, and the mechanisms it proposes tend to be complex and impractical. In contrast, players typically participate in various mechanisms that are run by different principals (e.g. different sellers on eBay or different ad-exchange platforms) and coordinating them to run a single combined mechanism is infeasible or impractical. Even the simple and elegant Vickrey auction loses some of its appeal when not in isolation: when the overall value of each player is a complex function of the outcomes of different Vickrey auctions, the global mechanism is no longer truthful.

We initiate the study of efficient mechanism design with guaranteed good properties even when players participate in multiple different mechanisms simultaneously or sequentially. We

define the class of smooth mechanisms, related to smooth games defined by Roughgarden, that can be thought of as mechanisms that generate approximately market clearing prices. We show that smooth mechanisms result in high quality outcome in equilibrium both in the full information setting and in the Bayesian setting with uncertainty about participants, as well as in learning outcomes. Our main result is to show that such mechanisms compose well: smoothness locally at each mechanism implies efficiency globally.

3.31 Rounding and the Allocation of Indivisible Objects

Rakesh V. Vohra (Northwestern University, Evanston, US, r-vohra@kellogg.northwestern.edu)

The problem of allocating indivisible objects arises in the allocation courses, spectrum licenses, landing slots at airports and assigning students to schools. We propose a technique for making such allocations that is based on rounding a fractional allocation. Under the assumption that no agent wants to consume more than k items, the rounding technique can be interpreted as giving agents lotteries over approximately feasible integral allocations that preserve the ex-ante efficiency and fairness properties of the initial fractional allocation. The integral allocations are only approximately feasible in the sense that upto $k-1$ more units than the available supply of any good is allocated.

3.32 Equilibria in the Challenge Tournament

Bernhard von Stengel (London School of Economics, stengel@nash.lse.ac.uk)

Arad and Rubinstein (2013) describe a challenge tournament where n players have a binary choice and play a round-robin tournament where they score against each other randomly for a stake that depends on their choices. The player with the highest total score wins, with ties resolved randomly. They conjecture that for $n > 3$ the only equilibrium is that all players take the riskier choice. We propose an elementary proof of this conjecture based on a dominance argument.

4 Further Participants and Session Chairs

In addition to the speakers listed above, the following researchers participated in this Dagstuhl workshop, often co-authors of the given presentations:

- Giorgos Christodoulou (University of Liverpool, GB)
- Michal Feldman (Hebrew University, Jerusalem, IL) – session chair Thursday afternoon
- Felix Fischer (University of Cambridge, GB)
- Yannai A. Gonczarowski (Hebrew University, Jerusalem, IL)
- Penélope Hernández (University of Valencia, ES) – session chair Friday morning
- Martin Höfer (MPI für Informatik – Saarbrücken, DE)

- Max Klimm (TU Berlin, DE)
- Elias Koutsoupias (University of Oxford, GB) – session chair Tuesday afternoon
- Jeffrey MacKie-Mason (University of Michigan, US) – session chair Tuesday morning
- Hervé Moulin (Rice University – Houston, US) – session chair Wednesday morning
- Dimitrii V. Pasechnik (Nanyang Technical University, Singapore, SG)
- Rahul Savani (University of Liverpool, GB)
- Michael Schapira (Hebrew University, Jerusalem, IL)
- Éva Tardos (Cornell University, US) – session chair Monday afternoon
- Berthold Vöcking (RWTH Aachen, DE) – session chair Thursday morning
- Jens Witkowski (Universität Freiburg, DE)

5 Open Problems

Two open problem sessions were held on Tuesday and Thursday afternoon after the afternoon talks. Of the discussed research questions, the following is published here (the open problem is stated at the end of the following exposition; the topic was not presented in a talk but in the open problem session).

5.1 Fair Division of a Max-Flow

Hervé Moulin (Rice University, Houston, US, moulin@rice.edu)

The problem of fairly dividing a max-flow has a long history, going back to the work of Megiddo, Brown, Hall and Vohra, and others. In all applications there is a distribution network that is capacity constrained (capacities may be on nodes or edges), and agents need to be allocated different amounts of the commodity based on objective or subjective needs. While standard models of flows in networks are typically concerned with the optimization of an exogenous objective, there is a substantial subset of that literature where the goal is to not only maximize the quantity distributed, but also to ensure that the distribution is equitable. As an example, Brown (O.R., 1979) discussed a sharing problem motivated by the equitable distribution of coal during a prolonged coal strike. Other settings where similar problems come up is in distributing aid or relief during a catastrophic event (e.g., food during a famine).

The classical papers of Megiddo and Brown call a max-flow fair if it equalizes the allocation of shares between the relevant agents (located on some nodes or edges); they typically use the lexicographic ordering to select the fairest max-flow, and design efficient algorithms to find it. Crucially, they take an ethically neutral view of the structure of the graph: if we can give equal shares of the flow to all agents, we should do it, ignoring the differences in their connectivity altogether. This is appropriate in contexts where agents should not be held responsible for their connections/capacities, such as the distribution of relief supplies mentioned above. It is not appropriate when agents should be held responsible for their connections in the network, so that a "better connected", "more central", agent ought to receive a greater share of the resources. Think of the distribution of a workload between contractors with different abilities to perform certain jobs.

The formal property of "consistency" captures the influence of the network structure on the final division of the flow: each agent has a claim only on the subflows that he is connected to, and the division of a given subflow relies on the residual claims of the relevant agents (total claim minus shares in other subflows; as in the bargaining model of Kleinberg and Tardos). Consistency is a key component of several other axiomatic fair division models.

Jay Sethuraman and I have developed a fairly complete model for the case of bipartite graphs but extending the approach to a general flowgraph raises several conceptual difficulties that I would like to discuss with the workshop participants.

6 Organizational Issues

The following comments about the organization of the workshop and lessons learnt may be of general interest.

6.1 Invitations

By far the most time-consuming organizational issue was whom to invite and when. In the end, the seminar had 48 participants, which were well accommodated (very few junior people sharing rooms) while a second smaller seminar took place in parallel at Schloss Dagstuhl.

The selection of invitees was difficult because of the broad scope of the workshop across economics and computer science.

In order to explain the workshop to economists, Sergiu Hart as the organizer closest to economics sent a separate email to them explaining the "by invitation only" and funding model. In addition, we asked for a subjective probability of attendance in order to estimate the number of participants. Both were very helpful.

One suggestion we have is the possibility to "*overbook*" the seminar even if confirmation numbers are relatively high (within small limits, of course). The number of people confirming after the first round of invitations was higher than expected. This left almost no room for a second round of invitations until a few months before the start of the workshop when people had to decide (and, as expected, some could not come after all). At this point (about six months before the meeting), it would have been good if we had been allowed to invite one or two senior researchers that we had not invited in the first round or which were suggested to us by colleagues, because in the end we did have the space and at worst could have asked more people to share rooms. For one person the eventual invitation came too late. However, the high caliber of those who had confirmed continued to make the workshop very attractive. The overall coordination with the Dagstuhl service team (Annette Beyer) was excellent.

In the end, we were happy with the selection that we could make: If in doubt, preferring an early-career researcher to a more established one. In addition, PhD students of the organizers could fill last-minute open places (in the feedback for the workshop, one participant suggested not to allow this because of a perceived restriction of the diversity of the group, but we think this is one of the privileges the organizers should enjoy, and because it allows to use free spaces with flexibility).

Moreover, the suggested "affirmative action" of Dagstuhl was successful in the sense that we had a relative large number of female participants (10 out of 48). In addition, three of them had a baby or small child looked after by the nanny of the Center, so families were well accommodated. Dagstuhl is exemplary here!

6.2 During the Workshop

We wanted to have a workshop with sufficient free time for informal talks and collaboration. For that purpose, we wrote about two months before the workshop an email to the effect of "if you really want to give a talk, please send us your suggested topic". This was deliberately not very encouraging and resulted in 12 suggested topics, not enough to fill the workshop.

We separately encouraged those whom we knew as particularly good speakers to propose a topic.

A second, more short-term email resulted in over 20 further suggestions, so that we had to make a choice but in the end accommodated nearly everyone.

We allocated 32 talks (with the abstracts above) of 30 minutes each to the following slots (morning + afternoon): 1+5 (Monday), 5+3 (Tuesday), 5+0 (Wednesday), 5+3 (Thursday), 5+0 (Friday), plus 2 open problem sessions at the end of the afternoons of Tuesday and Thursday.

Monday morning started with a "lightning talk" session where every speaker sent in advance one or two PDF slides about themselves and introduced themselves for at least one minute but no longer than two minutes. This worked very well and was considered a success. One suggestion is a large *countdown clock* (maybe in software) that would have helped to put the speakers more at ease that they use their time properly, because a talk time of at most two minutes is unfamiliar to most people. Only the very junior participants talked about themselves not long enough (i.e., less than one minute). The length restriction allowed the lightning talks to proceed in two sessions of about one hour each; five minutes per talk would have been too much.

In fact, the lightning talk already provided information about the quality of a speaker. One suggestion was to let participants *vote* on the permitted time of a talk for each speaker after the lightning talks, which could work as follows: Standard slot speaking time is 20 minutes, but, say, 10 slots of 30 minutes are also available, and everybody would vote for their most preferred speakers for those longer slots. In that way, winning a longer slot time would be an honour, rather than a shorter time an embarrassment. (The full vote count would not be made public, only the winners.)

Because of the lightning talks, Monday had one scientific talk in the morning and five afternoon talks. Every other day had five morning talks, and the only other afternoon talks were three each (plus an open problems session) on Tuesday and Thursday from 16:00, so that a large part of the afternoon and every evening was free for personal collaboration. This was highly appreciated.

In retrospect, we could have possibly been more selective and dictatorial in reducing or lengthening individual talk times, or inviting longer tutorials (apart from the suggested vote after the lightning talks). On the other hand, the equal talk length allowed us to make these decisions at the workshop with ease (with a Google Docs spreadsheet for joint editing). Only the Monday schedule was decided shortly in advance of the meeting. At the workshop, we could also discuss the content of some presentations with those participants who wanted advice, which was helpful. In addition, the topics followed in relatively random order, which gave each day variety and allowed everyone to listen regularly to a topic close to their interests.

All talk topics and most abstracts were sent to us ahead of time, but only about half of them were uploaded as system-wide materials. As suggested in the feedback for the workshop, we should have made this *mandatory* for being allowed to give a talk, for the benefit of the participants. On the other hand, people attended most talks, despite of limited advance information.

There was no need for working groups or parallel sessions because the group was largely interested in and listening to all topics.

The free Wednesday afternoon was used for a local hike, and we were very lucky to have good weather throughout the workshop. Most took part and liked it. A small number of participants would have preferred a more exciting sightseeing tour of, for example, the Völklinger Hütte, but the long bus journey and the fact that a few had already seen it made us decide against it. The site office was very helpful with advice and local maps.

The quality of talks was very high, in line with our expectations and selection of the participants.

Participants

- Yakov Babichenko
CalTech, US
- Simina Branzei
Aarhus University, DK
- Markus Brill
TU München, DE
- Jing Chen
Institute of Advanced Study – Princeton, US
- Giorgos Christodoulou
University of Liverpool, GB
- Constantinos Daskalakis
MIT, US
- Gabrielle Demange
Paris School of Economics, FR
- Nikhil R. Devanur
Microsoft – Redmond, US
- Shahar Dobzinski
Weizmann Inst. – Rehovot, IL
- Paul Dütting
EPFL – Lausanne, CH
- Edith Elkind
Nanyang TU – Singapore, SG
- Michal Feldman
The Hebrew University of Jerusalem, IL
- Felix Fischer
University of Cambridge, GB
- Paul W. Goldberg
University of Liverpool, GB
- Yannai A. Gonczarowski
The Hebrew University of Jerusalem, IL
- Sergiu Hart
The Hebrew University of Jerusalem, IL
- Jason D. Hartline
Northwestern University – Evanston, US
- Penelope Hernandez Rojas
University of Valencia, ES
- Martin Hoefer
MPI für Informatik – Saarbrücken, DE
- Nicole Immorlica
Northwestern University – Evanston, US
- Ramesh Johari
Stanford University, US
- Thomas Kesselheim
Cornell University, US
- Max Klimm
TU Berlin, DE
- Elias Koutsoupias
University of Oxford, GB
- Kevin Leyton-Brown
University of British Columbia – Vancouver, CA
- Katrina Ligett
CalTech – Pasadena, US
- Brendan Lucier
Microsoft Research New England – Cambridge, US
- Jeffrey MacKie-Mason
University of Michigan, US
- Vangelis Markakis
Athens University of Economics and Business, GR
- Vahab Mirrokni
Google – New York, US
- Hervé Moulin
Rice University – Houston, US
- Rudolf Müller
Maastricht University, NL
- Sigal Oren
Cornell University, US
- Mallesh Pai
University of Pennsylvania, US
- Dimitrii V. Pasechnik
Nanyang TU – Singapore, SG
- Aaron Roth
University of Pennsylvania, US
- Tim Roughgarden
Stanford University, US
- William H. Sandholm
University of Wisconsin – Madison, US
- Rahul Savani
University of Liverpool, GB
- Guido Schäfer
CWI – Amsterdam, NL
- Michael Schapira
The Hebrew University of Jerusalem, IL
- Ilya R. Segal
Stanford University, US
- Vasilis Syrgkanis
Cornell University, US
- Éva Tardos
Cornell University, US
- Berthold Vöcking
RWTH Aachen, DE
- Rakesh V. Vohra
Northwestern University – Evanston, US
- Bernhard von Stengel
London School of Economics, GB
- Jens Witkowski
Universität Freiburg, DE