

Computational Social Dynamics

Martin Hoefer^{*1}, Sigal Oren^{*2}, Roger Wattenhofer^{*3}, and
Giovanna Varricchio^{†4}

- 1 Goethe-Universität – Frankfurt am Main, DE. mhoefer@em.uni-frankfurt.de
- 2 Ben Gurion University – Beer Sheva, IL. sigal3@gmail.com
- 3 ETH Zürich, CH. wattenhofer@ethz.ch
- 4 Goethe-Universität – Frankfurt am Main, DE. Varricchio@em.uni-frankfurt.de

Abstract

This report documents the program and outcomes of Dagstuhl Seminar 22452 “Computational Social Dynamics”. The seminar addressed social and dynamic problems in the field of algorithmic game theory, and their implications in numerous applications, such as fair division, financial networks, or behavioral game theory. We summarize organizational aspects of the seminar, the talk abstracts, and the problems that were discussed in the open problem sessions.

Seminar November 6–11, 2022 – <http://www.dagstuhl.de/22452>

2012 ACM Subject Classification Theory of computation → Algorithmic game theory and mechanism design; Information systems → Social networks; Theory of computation → Theory and algorithms for application domains; Theory of computation

Keywords and phrases algorithmic game theory, behavioral economics, fair division, financial networks, social networks

Digital Object Identifier 10.4230/DagRep.12.11.28

1 Executive Summary

Martin Hoefer (Goethe-Universität – Frankfurt am Main, DE)

Sigal Oren (Ben Gurion University – Beer Sheva, IL)

Roger Wattenhofer (ETH Zürich, CH)

License  Creative Commons BY 4.0 International license
© Martin Hoefer, Sigal Oren, and Roger Wattenhofer

Algorithmic techniques provide a powerful toolbox for understanding many phenomena arising in modern society. Often, these phenomena are related to dynamics (e.g., dynamic information spreading, or dynamics in social networks that result from agent interaction). A large part of the present social interaction on networks can be expressed using game-theoretic or microeconomic concepts, e.g., the dynamics of opinions in networks, pricing problems and viral marketing, network-based effects of opinions, group formation, cognitive bias, etc. These problems are rigorously analyzed in the area of algorithmic game theory, where researchers apply the algorithmic toolbox to analyze various social systems.

In this field, a number of applications have not received sufficient attention, which are recently becoming increasingly prominent. For example,

- issues of fairness and bias are central challenges in modern societies.
- behavioral economists are challenging standard assumptions that humans are maximizing expected utility. This change in perspective also poses new challenges for suitable models and algorithm design.

* Editor / Organizer

† Editorial Assistant / Collector



Except where otherwise noted, content of this report is licensed under a Creative Commons BY 4.0 International license

Computational Social Dynamics, *Dagstuhl Reports*, Vol. 12, Issue 11, pp. 28–44

Editors: Martin Hoefer, Sigal Oren, Roger Wattenhofer, and Giovanna Varricchio



Dagstuhl Reports

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

- an important trend in the modern economy are informational challenges, e.g., in problems involving recommendation, persuasion, delegation, or (smart) contract design.

The main aim of this seminar was to bring together a leading set of researchers to discuss these and other challenges, and to advance the state of the art in several new directions. The majority of participants were academics from computer science departments; some were from other disciplines such as economics, mathematics, or electrical engineering. All participants had strong interdisciplinary interests that typically span economics, game theory, and theoretical computer science.

The seminar started on Monday with an introductory session, in which participants introduced their name, affiliation, main research interest as well as a “crazy idea” or “provoking thought”. This session was very well-received by the participants, and it initiated discussions directly from the start. The subsequent program included four invited talks/tutorials of roughly one hour each, on several issues chosen by the organizers.

Monday Caragiannis talked about notions of fairness for resource allocation problems.

Tuesday Feldman presented recent innovations in algorithms for contract design.

Wednesday Plonsky discussed behavioral experiments and ideas to predict human behavior.

Thursday Wattenhofer surveyed issues in e-government, blockchains, and finance.

The contributed talks were solicited from the participants and lasted around 20-25 minutes each. In many of these talks there were lively discussions. They continued in the times from lunch to afternoon coffee, which were free for research and individual collaborative meetings.

There was a substantial set of presenters who focused on aspects of fair division. Rathi discussed epistemic notions of envy-freeness and their efficient computation. Bernadé presented results on online algorithms for fair division and guaranteeing notions of envy-freeness and Pareto-efficiency. Zick showed existence and efficient computation for several fairness criteria for matroid-rank valuations. Varricchio discussed algorithmic aspects of randomization and entitlements in fairness notions. Reiffenhäuser presented results on the quality of equilibria in fair division when strategic misreporting of preferences is allowed.

Another focus was the analysis of dynamics and stability in networks. Ventre discussed the clearing problem in financial networks with credit default swaps and showed FIXP-completeness results. Wilhelmi presented a game-theoretic model for clearing and results on equilibrium computation. Schmand considered a game-theoretic model for opinion formation in networks and convergence of natural best-response dynamics. Lenzner surveyed recent work on network creation games as well as network games that model segregation aspects.

Further talks addressed a number of different areas. Witkowski studied incentive-compatible forecasting competitions. Lavi considered incentives when several of these contests are available to participants. Klimm surveyed work on impartial selection, e.g., when agents have to decide on a representative member among themselves. Leyton-Brown posed new directions and open problems in the understanding of behavior in mobile gaming. Babichenko discussed aggregation mechanisms for anonymous information. Markakis presented a novel algorithm to compute an equilibrium in 2-player zero-sum games. Melnyk explained issues arising from Byzantine behavior in social choice. Last, but not least, Christodoulou gave a proof of the 23-year old Nisan-Ronen conjecture that every truthful mechanism for unrelated machine scheduling can only guarantee a linear approximation ratio.

The seminar was a big success. We believe it will stimulate new and very fruitful collaborations. We got laudatory feedback from many participants which is also reflected in the survey conducted by Dagstuhl.

We thank Giovanna Varricchio for serving as collector of abstracts and open problems.

2 Table of Contents

Executive Summary

<i>Martin Hoefer, Sigal Oren, and Roger Wattenhofer</i>	28
---	----

Overview of Talks


Regret-minimizing aggregation of anonymous information <i>Yakov Babichenko</i>	32
Dynamic fair division with partial information <i>Gerdus Benade</i>	32
Fairness in allocation problems (Tutorial) <i>Ioannis Caragiannis</i>	33
A proof of the Nisan-Ronen conjecture <i>Giorgos Christodoulou</i>	33
Algorithmic Contract Design (Tutorial) <i>Michal Feldman</i>	33
Impartial selection problems <i>Max Klimm</i>	33
From Monopoly to Competition: When Do Optimal Contests Prevail? <i>Ron Lavi</i>	34
Dynamics in Network Creation and Residential Segregation <i>Pascal Lenzner</i>	34
Mobile Games: An Exciting New Domain for Behavioral Game Theory <i>Kevin Leyton-Brown</i>	35
On new variants of multiplicative weights update and mirror prox methods for 0-sum games <i>Vangelis Markakis</i>	35
Byzantine Adversaries in Social Choice <i>Darya Melnyk</i>	36
Predictably Irrational? Can behavioral decision making research help predict human choice? (Tutorial) <i>Ori Plonsky</i>	36
New Fairness notions for Resource Allocation <i>Nidhi Rathi</i>	36
Allocating Indivisible Goods to Strategic Agents: Pure Nash Equilibria and Fairness <i>Rebecca Reiffenhäuser</i>	37
Asynchronous Opinion Dynamics in Social Networks <i>Daniel Schmand</i>	37
Best of Both Worlds: Agents with Entitlements <i>Giovanna Varricchio</i>	38
Systemic risk in financial networks – a computational perspective <i>Carmine Ventre</i>	39

Does Computation Change Society? (Tutorial)	
<i>Roger Wattenhofer</i>	39
Seniorities and Minimal Clearing in Financial Network Games	
<i>Lisa Wilhelmi</i>	40
Incentive-Compatible Forecasting Competitions	
<i>Jens Witkowski</i>	40
A Simple Vision for Fair Division	
<i>Yair Zick</i>	41
Open problems	
Condorcet outcome for m independent binary issues	
<i>Andrei Constantinescu</i>	41
Interval cover for spatial crowdsourcing – Improving approximation	
<i>Vangelis Markakis</i>	41
Complexity of Public Goods Games on Graph – Nash equilibrium existence (Solved)	
<i>Noam Nisan</i>	42
Optimizing over Serial Dictatorships – Lower bound for optimal matching	
<i>Nidhi Rathi</i>	42
Altruistic Hedonic Games (Short talk)	
<i>Jörg Rothe</i>	43
Tractability frontier of justice criteria under submodular valuations	
<i>Yair Zick</i>	43
Participants	44

3 Overview of Talks

3.1 Regret-minimizing aggregation of anonymous information

Yakov Babichenko (Technion – Haifa, IL)

License  Creative Commons BY 4.0 International license
© Yakov Babichenko

Joint work of Itai Arieli, Yakov Babichenko, Inbal Talgam-Cohen, Konstantin Zabarniy

We study a model in which a decision maker aggregates information obtained from several symmetric agents. Each agent provides the decision maker with a recommendation about a binary state of nature, where the state is drawn from a known prior distribution. While the decision maker knows the marginal distribution of each agent’s recommendation, the correlation between the recommendations is chosen adversarially. The decision maker’s goal is to choose an information aggregation function minimizing the regret – the difference between her own mistake probability when guessing the state of nature, and the mistake probability of a Bayesian decision maker knowing the correlation between the recommendations.

We provide a characterization of the minimal regret for any number of agents as the maximal Jensen gap of a convex function that captures the probability of a correct guess by a hypothetical Bayesian decision maker. For a large number of agents, we deduce that apart from some borderline cases, the unique optimal aggregation function is the random dictator rule that chooses an agent uniformly at random and adopts her recommendation.

3.2 Dynamic fair division with partial information

Gerdus Benade (Boston University, US)


License  Creative Commons BY 4.0 International license
© Gerdus Benade

We consider the problem of fairly and efficiently allocating T indivisible items among n agents with additive preferences. Items arrive over a sequence of rounds, and must be allocated immediately and irrevocably before the next one arrives. When the agents’ valuations for the items are drawn from known distributions, it is possible to find allocations that are envy-free with high probability and Pareto efficient ex-post.

We study a partial-information setting, where values are drawn from unknown distributions and it is possible to elicit ordinal but not cardinal information. When a new item arrives, the algorithm can query each agent for the relative rank of this item with respect to a subset of the past items. We provide algorithms with strong simultaneous fairness and efficiency guarantees even with minimally expressive queries that ask for a comparison to a single previous item and show they are asymptotically optimal.

3.3 Fairness in allocation problems (Tutorial)


Ioannis Caragiannis (Aarhus University, DK)

License  Creative Commons BY 4.0 International license
© Ioannis Caragiannis

We will present variations of the problem of allocating indivisible items to agents with (mainly additive) valuations for the items. We will define basic fairness concepts such as proportionality and envy-freeness and discuss their basic properties. Next, we will introduce approximate versions of these concepts, such as envy-freeness up to some/any item (EF1/EFX) and maximin share fairness (MMS). We will present examples and many open problems.

3.4 A proof of the Nisan-Ronen conjecture

Giorgos Christodoulou (Aristotle University of Thessaloniki, GR)

License  Creative Commons BY 4.0 International license
© Giorgos Christodoulou

Noam Nisan and Amir Ronen conjectured that the best approximation ratio of deterministic truthful mechanisms for makespan-minimization for n unrelated machines is n . This work validates the conjecture.

3.5 Algorithmic Contract Design (Tutorial)

Michal Feldman (Tel Aviv University, IL)

License  Creative Commons BY 4.0 International license
© Michal Feldman
Joint work of Tomer Ezra, Paul Duetting, Michal Feldman, Thomas Kesselheim

Up until recently, Algorithmic Game Theory has mainly focused on the design of mechanisms that incentivize agents to truthfully report their private preferences. However, algorithms and incentives interact in many additional ways; the design of contracts being a prime example. While mechanism design deals with hidden preferences, contract design deals with hidden actions, and studies how best to incentivize agents to take costly actions, when their actions are hidden from the principal. With the transition of classic applications of contracts into computational platforms, algorithm design for such applications becomes timely and relevant.

In this talk, I will survey two papers on combinatorial contracts, which highlight different sources of complexity that arise in contract design.

3.6 Impartial selection problems

Max Klimm (TU Berlin, DE)

License  Creative Commons BY 4.0 International license
© Max Klimm

Impartial selection problems are concerned with situations where a group of agents selects a subset of the agents based on nominations from within the set. The fact that the agents act both as voters and as nominees may give rise to incentive issues when some agents may not be

willing to communicate their true opinion about who should be selected in order to influence their own chances of being selected. These issues may arise in a number of applications such as voting in committees, peer review in conferences where committee members also submit papers, and peer grading. One way to circumvent these incentive issues is to use impartial mechanisms that have the property that the selection probability of each agent is independent of its nominations. In this talk, I will give an overview on impartial selection mechanisms and point to some open problems in this area.

3.7 From Monopoly to Competition: When Do Optimal Contests Prevail?

Ron Lavi (University of Bath, GB)

License  Creative Commons BY 4.0 International license
© Ron Lavi

Joint work of Xiaotie Deng, Yotam Gafni, Ron Lavi, Tao Lin, Hongyi Ling

Main reference Xiaotie Deng, Yotam Gafni, Ron Lavi, Tao Lin, Hongyi Ling: “From Monopoly to Competition: Optimal Contests Prevail”, CoRR, Vol. abs/2107.13363, 2021.

URL <https://arxiv.org/abs/2107.13363>

We study competition among simultaneous heterogeneous contest designers in a general model that allows for a large space of contest design. Contestants choose in which contest to participate, and the goal of each contest designer is to maximize the contestants’ sum of efforts exerted in her contest. Our main result shows that, with symmetric contestants, optimal contests in the monopolistic setting (i.e., those that maximize the sum of efforts in a model with a single contest designer) form Pareto-optimal equilibria when contest designers compete. Under a natural assumption, monopolistic optimal contests are in fact dominant in the competitive case, and the equilibria that they form are unique. In many natural cases, they also maximize social welfare.

3.8 Dynamics in Network Creation and Residential Segregation

Pascal Lenzner (Hasso-Plattner-Institut, Universität Potsdam, DE)

License  Creative Commons BY 4.0 International license
© Pascal Lenzner

I will focus on dynamics in two different game-theoretic models, one focusing on network creation and the other one modelling location choice in a residential area.

Network creation games are motivated by the observation that the structure of many real-world networks like the Internet or (online) social networks is the outcome of a complex interaction of selfish agents. These settings can be modeled as strategic games and methods from Algorithmic Game Theory can be employed to rigorously analyze them. I will survey recent results on popular variants of network creation games with a focus on the most intriguing open problem: finding equilibria via suitable game dynamics.

Based on Schelling’s seminal model for residential segregation, recently game-theoretic variants of it have become popular. In these models agents with different types strategically select a location on a given graph that models the residential area. Each agent’s utility depends on the agent type distribution in its respective neighborhood. Given this, studying the game dynamics of such models provides insights into why large metropolitan areas are segregated along various dimensions like ethnicity or household income. I will survey recent work on these dynamic aspects.

3.9 Mobile Games: An Exciting New Domain for Behavioral Game Theory

Kevin Leyton-Brown (University of British Columbia – Vancouver, CA)

License © Creative Commons BY 4.0 International license
© Kevin Leyton-Brown

Joint work of Kevin Leyton-Brown, Taylor Lundy, Narun Raman, Hu Fu

Two and a half billion people play mobile games worldwide, spending \$180.3B USD in 2021 alone. More than half of this revenue comes from free-to-play mobile games. These games make money through carefully optimized monetization strategies based on optional in-game purchases. Mobile games also often leverage social dynamics to recruit new players and convert existing players into paying customers. Because player behavior in games, while thoughtful, often deviates from standard definitions of rationality, mobile games constitute a rich laboratory for research in behavioural game theory.

In our first work on this topic we looked at the phenomenon of “skips” in which players pay real money to avoid game content. We leverage existing tools of mechanism design to answer questions like “Why would a player ever value skips?” and “Can expensive skips be good for player welfare?”. The talk will focus particularly on open questions from our ongoing work, which investigates elaborating both the behavioral model of player engagement (drawing on recent work from Kleinberg et. al.); incorporating the effects that virality and social networks can have on a game’s success; and studying the pervasive monetization scheme in which players are given a handful of “lives” and are forced to wait for a timer to reset these are depleted (or to buy more lives).

3.10 On new variants of multiplicative weights update and mirror prox methods for 0-sum games

Vangelis Markakis (Athens University of Economics and Business, GR)

License © Creative Commons BY 4.0 International license
© Vangelis Markakis

Main reference Michail Fasoulakis, Evangelos Markakis, Yannis Pantazis, Constantinos Varsos: “Forward Looking Best-Response Multiplicative Weights Update Methods for Bilinear Zero-sum Games”, in Proc. of the Proceedings of The 25th International Conference on Artificial Intelligence and Statistics, Proceedings of Machine Learning Research, Vol. 151, pp. 11096–11117, PMLR, 2022.

URL <https://proceedings.mlr.press/v151/fasoulakis22a.html>

We follow up on a recent line of works that attempt to establish last-iterate convergence results for iterative first-order methods in min-max optimization. Our work focuses on extra gradient learning algorithms for finding Nash equilibria in bilinear zero-sum games, motivated by related questions in online learning. A typical drawback of several first-order methods, including the standard versions of gradient descent and multiplicative weight updates, is that they converge only in an average sense, but not w.r.t. their last iterate. We propose a new algorithm, that can be considered as a variant of the Mirror Prox method (Nemirovski 2004), using a large learning rate parameter for the intermediate gradient step and a small learning rate in the final step of each iteration. Although counter-intuitive at first sight due to the irrationally large intermediate rate, we prove that this method attains last-iterate convergence. Furthermore, we perform experimental comparisons against other recently proposed algorithms (such as the optimistic variant of the multiplicative weights update method, by Daskalakis and Panageas (2019)), and show that our algorithm has significant practical potential since it offers substantial gains in terms of accelerated convergence.

3.11 Byzantine Adversaries in Social Choice

Darya Melnyk (Aalto University, FI)

License  Creative Commons BY 4.0 International license
© Darya Melnyk


Main reference Darya Melnyk, Yuyi Wang, Roger Wattenhofer: “Byzantine Preferential Voting”, in Proc. of the Web and Internet Economics, pp. 327–340, Springer International Publishing, 2018.

URL https://doi.org/10.1007/978-3-030-04612-5_22

In distributed computing, robust communication algorithms that handle a wide variety of adversaries, including network failures, node crashes, and even unpredictable behavior of the network participants, are being developed. In social choice theory, on the other hand, the participants are usually assumed to be selfish or biased towards some opinion, and they only influence the network with this goal in mind. In this talk, I will show how worst-case behavior of participants, the so-called Byzantine behavior, can influence the outcome of different voting protocols. To this end, I will talk about different ways of measuring how much the Byzantine party is able to influence the final result. These results can be used to design voting protocols that better represent the opinion of truthful voters.

3.12 Predictably Irrational? Can behavioral decision making research help predict human choice? (Tutorial)


Ori Plonsky (Technion – Haifa, IL)

License  Creative Commons BY 4.0 International license
© Ori Plonsky

Humans are said to be predictably irrational, yet accurate predictions of human choice behavior remain a major challenge. Behavioral science research often demonstrates contradictory deviations from rational choice, and predicting the direction of deviation is not easy. In this talk, I will review some of the classical findings and the most robust behavioral tendencies in human choice behavior. I will then demonstrate when and how understanding and accounting for them can aid the development of state-of-the-art predictive models.

3.13 New Fairness notions for Resource Allocation

Nidhi Rathi (Aarhus University, DK)

License  Creative Commons BY 4.0 International license
© Nidhi Rathi

Joint work of Ioannis Caragiannis, Jugal Garg, Nidhi Rathi, Eklavya Sharma, Giovanna Varricchio

We consider the fundamental problem in resource allocation setting that entails fairly dividing a set of discrete goods among agents. The notion of envy-freeness up to any good (EFX) is the most compelling notion of fairness in this line of work. We say an allocation is EFX when every agent (weakly) prefers her own bundle than any other agent j 's bundle after removing her least positively-valued item from j 's bundle. Despite significant efforts over the past few years, existence of EFX allocations is not known even for additive valuations. Therefore, there has been a lot of work focused on finding relaxations and approximations of EFX. In this talk, I will propose two natural relaxations (one comparison-based and one threshold-based) of fairness notions that is inspired from EFX.

The first notion is that of epistemic EFX (EEFX) which extends the definition of epistemic envy-freeness from [ABC+18]. An allocation is said to be EEFX if for every agent i , it is possible to shuffle/redistribute the goods of the other agents such that agent i does not envy any other agent up to any good (in this new allocation). Interestingly, we show that EEFX allocations are always guaranteed to exist for additive valuations, and can be found in polynomial-time. Next, we introduce a threshold-based fairness criteria of minimum EFX value (MXS) with agent thresholds defined using EFX allocations. The MXS threshold for agent i is defined to be the minimum value she receives in any allocation where she does not envy any other agent up to any good. Finally, we prove that it is NP-hard to compute the MXS threshold of any agent in a given fair-division instance. It is relevant to note that, despite this hardness constraint, we prove that an MXS allocation can always be computed in polynomial time for instances with additive valuations.

3.14 Allocating Indivisible Goods to Strategic Agents: Pure Nash Equilibria and Fairness

Rebecca Reiffenhäuser (University of Amsterdam, NL)

License © Creative Commons BY 4.0 International license
© Rebecca Reiffenhäuser

Main reference Georgios Amanatidis, Georgios Birmpas, Federico Fusco, Philip Lazos, Stefano Leonardi, Rebecca Reiffenhäuser: “Allocating Indivisible Goods to Strategic Agents: Pure Nash Equilibria and Fairness”, in Proc. of the Web and Internet Economics – 17th International Conference, WINE 2021, Potsdam, Germany, December 14-17, 2021, Proceedings, Lecture Notes in Computer Science, Vol. 13112, pp. 149–166, Springer, 2021.

URL https://doi.org/10.1007/978-3-030-94676-0_9

When allocating goods to agents (with additive valuations), both fairness and incentive compatibility are central goals. Unfortunately, each of those is often hard to achieve on its own, and aiming for both in the same (money-free) routine quickly leads to strong impossibility results. This work suggests a way around those, by abandoning the initial goal for a notion of the next-best thing: we show that the very natural round-robin algorithm for additive valuations actually produces EF1-fair, “truthful equilibria” – even though agents might have misreported their values.

3.15 Asynchronous Opinion Dynamics in Social Networks

Daniel Schmand (Universität Bremen, DE)


License © Creative Commons BY 4.0 International license
© Daniel Schmand

Opinion spreading in a society decides the fate of elections, the success of products, and the impact of political or social movements. The model by Hegselmann and Krause is a well-known theoretical model to study such opinion formation processes in social networks. In contrast to many other theoretical models, it does not converge towards a situation where all agents agree on the same opinion. Instead, it assumes that people find an opinion reasonable if and only if it is close to their own. The system converges towards a stable situation where agents sharing the same opinion form a cluster, and agents in different clusters do not influence each other.

We focus on the social variant of the Hegselmann-Krause model where agents are connected by a social network and their opinions evolve in an iterative process. When activated, an agent adopts the average of the opinions of its neighbors having a similar opinion. By this, the set of influencing neighbors of an agent may change over time. To the best of our knowledge, social Hegselmann-Krause systems with asynchronous opinion updates have only been studied with the complete graph as social network. We show that such opinion dynamics with random agent activation are guaranteed to converge for any social network. We provide an upper bound of $O(n|E|^2(\varepsilon/\delta)^2)$ on the expected number of opinion updates until convergence, where $|E|$ is the number of edges of the social network. For the complete social network we show a bound of $O(n^3(n^2 + (\varepsilon/\delta)^2))$ that represents a major improvement over the previously best upper bound of $O(n^9(\varepsilon/\delta)^2)$. Our bounds are complemented by simulations that indicate asymptotically matching lower bounds.

3.16 Best of Both Worlds: Agents with Entitlements

Giovanna Varricchio (Goethe-Universität – Frankfurt am Main, DE)

License  Creative Commons BY 4.0 International license
© Giovanna Varricchio

Joint work of Martin Hoefer, Marco Schmalhofer, Giovanna Varricchio
Main reference Martin Hoefer, Marco Schmalhofer, Giovanna Varricchio: “Best of Both Worlds: Agents with Entitlements”, CoRR, Vol. abs/2209.03908, 2022.
URL <https://doi.org/10.48550/arXiv.2209.03908>

Fair division of indivisible goods is a central challenge in artificial intelligence. For many prominent fairness criteria including envy-freeness (EF) or proportionality (PROP), no allocations satisfying these criteria might exist. Two popular remedies to this problem are randomization or relaxation of fairness concepts. A timely research direction is to combine the advantages of both, commonly referred to as Best of Both Worlds (BoBW).

In the case of equally entitled agents and additive valuations, a lottery that is simultaneously ex-ante EF and ex-post envy-free up to one good (EF1) is known to exist. With our work, we broaden the picture by focusing on agents with different entitlements. Our main result is a lottery that is ex-ante weighted envy-free (WEF), as well as ex-post weighted proportional up to one good (WPROP1) and weighted transfer envy-free up to one good (WEF(1, 1)). In addition, we show that this result is tight – ex-ante WEF is incompatible with any stronger ex-post WEF relaxation.

We also try to extend positive results to more expressive valuation functions; unfortunately, our techniques partially apply in more general settings. Extending BoBW results is an interesting and challenging problem, and developing new tools is a fundamental next step in this direction.

3.17 Systemic risk in financial networks – a computational perspective

Carmine Ventre (King's College London, GB)

License © Creative Commons BY 4.0 International license
© Carmine Ventre

Joint work of Stavros Ioannidis, Bart de Keijzer, Carmine Ventre

We are given a network modelling assets and liabilities of the financial institutions in the system. We study the following basic question: Can we efficiently compute the exposure rate of each bank to defaults in the system?

The answer is yes if there are no financial derivatives (i.e., conditional obligations) in the network. When we introduce derivatives, specifically Credit Default Swaps (CDS), the problem is complete for the complexity class PPAD if one is content with "almost" (that is, weakly approximate) solutions that could grossly under- and over-estimate each bank's exposure.

What about solutions where the rate is precise to say 1%? We prove that computing these strong approximations up to any given precision is complete for the class FIXP, capturing hardness due to numerical aspects (in particular, the irrationality of the actual solution) in addition to the combinatorial issues modelled by PPAD.

We also study the relationship between the network structure and the (ir)rationality of the solution, the robustness of our findings to different payment rules of insolvent banks and the computational complexity of questions motivated by the needs of regulators. Overall, our results support a ban of the purely speculative naked CDS and uncover a connection between FIXP and efficient algorithms in the real computational model of Blum-Shub-Smale.

3.18 Does Computation Change Society? (Tutorial)

Roger Wattenhofer (ETH Zürich, CH)

License © Creative Commons BY 4.0 International license
© Roger Wattenhofer

In our Dagstuhl Seminar, we talk about better mechanisms to organize problems of society in a better (for instance fairer) way. We believe that by now computation should have had a huge impact on society, governments, and markets. However, many government mechanisms are still organized the same way as 100 years ago, pretty much ignoring all technical progress. In this talk, we discuss a few examples of computation in society. In the first part, we discuss systems that can help making voting decisions, or the benefits of electronic voting systems. We even fantasize about having a democracy without the need of politicians (made scalable by artificial intelligence). In the second part of the talk, we discuss some aspects of markets, in particular decentralized exchanges.

3.19 Seniorities and Minimal Clearing in Financial Network Games

Lisa Wilhelmi (Goethe-Universität – Frankfurt am Main, DE)

License © Creative Commons BY 4.0 International license
© Lisa Wilhelmi

Joint work of Martin Hoefer, Lisa Wilhelmi

Main reference Martin Hoefer, Lisa Wilhelmi: “Seniorities and Minimal Clearing in Financial Network Games”, in Proc. of the Algorithmic Game Theory – 15th International Symposium, SAGT 2022, Colchester, UK, September 12-15, 2022, Proceedings, Lecture Notes in Computer Science, Vol. 13584, pp. 187–204, Springer, 2022.

URL https://doi.org/10.1007/978-3-031-15714-1_11

Financial network games model payment incentives in the context of networked liabilities. In this paper, we advance the understanding of incentives in financial networks in two important directions: minimal clearing (arising, e.g., as a result of sequential execution of payments) and seniorities (i.e., priorities over debt contracts).

We distinguish between priorities that are chosen endogenously or exogenously. For endogenous priorities and standard (maximal) clearing, the games exhibit a coalitional form of weak acyclicity. A strong equilibrium exists and can be reached after a polynomial number of deviations. Moreover, there is a strong equilibrium that is optimal for a wide variety of social welfare functions. In contrast, for minimal clearing there are games in which no optimal strategy profile exists, even for standard utilitarian social welfare. Perhaps surprisingly, a strong equilibrium still exists and, for a wide range of strategies, can be reached after a polynomial number of deviations. In contrast, for exogenous priorities, equilibria can be absent and equilibrium existence is NP-hard to decide, for both minimal and maximal clearing.

3.20 Incentive-Compatible Forecasting Competitions

Jens Witkowski (Frankfurt School of Finance & Management, DE)

License © Creative Commons BY 4.0 International license
© Jens Witkowski

Joint work of Jens Witkowski, Rupert Freeman, Jennifer Wortman Vaughan, David Pennock, Andreas Krause

We initiate the study of incentive-compatible forecasting competitions in which multiple forecasters make predictions about one or more events and compete for a single prize. We have two objectives: (1) to incentivize forecasters to report truthfully and (2) to award the prize to the most accurate forecaster. Proper scoring rules incentivize truthful reporting if all forecasters are paid according to their scores. However, incentives become distorted if only the best-scoring forecaster wins a prize, since forecasters can often increase their probability of having the highest score by reporting more extreme beliefs. In this paper, we introduce two novel forecasting competition mechanisms. Our first mechanism is incentive compatible and guaranteed to select the most accurate forecaster with probability higher than any other forecaster. Moreover, we show that in the standard single-event, two-forecaster setting and under mild technical conditions, no other incentive-compatible mechanism selects the most accurate forecaster with higher probability. Our second mechanism is incentive compatible when forecasters’ beliefs are such that information about one event does not lead to belief updates on other events, and it selects the best forecaster with probability approaching 1 as the number of events grows. Our notion of incentive compatibility is more general than previous definitions of dominant strategy incentive compatibility in that it allows for reports to be correlated with the event outcomes. Moreover, our mechanisms are easy to implement and can be generalized to the related problems of outputting a ranking over forecasters and hiring a forecaster with high accuracy on future events.

3.21 A Simple Vision for Fair Division

Yair Zick (University of Massachusetts – Amherst, US)

License © Creative Commons BY 4.0 International license
© Yair Zick

Main reference Vignesh Viswanathan, Yair Zick: “A General Framework for Fair Allocation with Matroid Rank Valuations”, CoRR, Vol. abs/2208.07311, 2022.

URL <https://doi.org/10.48550/arXiv.2208.07311>

We study the problem of fairly allocating a set of indivisible goods among agents with matroid rank valuations. We present a simple framework that efficiently computes any fairness objective that satisfies some mild assumptions. Along with maximizing a fairness objective, the framework is guaranteed to run in polynomial time, maximize utilitarian social welfare and ensure strategyproofness. We show how our framework can be used to achieve four different fairness objectives: (a) Prioritized Lorenz dominance, (b) Maxmin fairness, (c) Weighted leximin, and (d) Max weighted Nash welfare. In particular, our framework provides the first polynomial time algorithms to compute weighted leximin and max weighted Nash welfare allocations for matroid rank valuations.

4 Open problems

4.1 Condorcet outcome for m independent binary issues

Andrei Constantinescu (ETH Zürich, CH)

License © Creative Commons BY 4.0 International license
© Andrei Constantinescu

An assembly of n voters wants to decide on m independent binary issues. The preferences of each voter and potential outcomes of the decision-making process are represented by m -bit vectors. In a direct race between two potential outcomes, a voter will vote for the outcome closer in Hamming distance to their vector, abstaining if indifferent.

Problem: Find if there is a Condorcet outcome; i.e. one not defeated by any other outcome in a direct race.

4.2 Interval cover for spatial crowdsourcing – Improving approximation

Vangelis Markakis (Athens University of Economics and Business, GR)

License © Creative Commons BY 4.0 International license
© Vangelis Markakis

Consider a set of tasks, say $\{1, \dots, m\}$, that are ordered in a line, and a set of available workers $\{1, \dots, n\}$. Each worker is able to work on a subinterval T_i of tasks for a cost c_i . A worker, if selected, can contribute with a value q_i to the tasks in the interval T_i , and each task j has a demand of d_j . The goal is to find an interval cover that completes every task and is of minimum cost. The optimization problem is known to be NP-hard but not APX-hard. Furthermore, the best-known approximation is 3.

Question 1: What’s the best approximation for the minimization problem?

Consider the strategic variant of the problem where costs are private information of workers. Denoted by Δ the maximum number of workers that can contribute to a task, the best-known approximation for a strategyproof mechanism is of Δ .

Question 2: What’s the best truthful approximation for the minimization problem?

4.3 Complexity of Public Goods Games on Graph – Nash equilibrium existence (Solved)

Noam Nisan (*The Hebrew University of Jerusalem, IL*)

License © Creative Commons BY 4.0 International license
© Noam Nisan

Main reference Matan Gilboa, Noam Nisan: “Complexity of Public Goods Games on Graphs”, CoRR, Vol. abs/2207.04238, 2022.

URL <https://doi.org/10.48550/arXiv.2207.04238>

This is a specific problem that was left open in the recent paper “Complexity of Public Goods Games on Graphs” by M. Gilboa and N. Nisan, SAGT 2022.

The game is played on a graph $G = (V, E)$ where each node is an agent that needs to choose among two actions $\{1, 0\}$. The action “1” is a best reply of a node if and only if at most two of its neighbours in the graph play 1. A pure Nash equilibrium of the game is an assignment $a : V \rightarrow \{0, 1\}$ such that for every vertex v we have:

$$a(v) = 1 \text{ iff } |\{u \mid (v, u) \in E \text{ and } a(u) = 1\}| \leq 2.$$

Question: Is it true that for every G there exists a pure Nash Equilibrium? Can you find one in polynomial time?

Solution: During the Dagstuhl Seminar, M. Klimm came up with an elegant solution by showing this game is equivalent to a congestion game, which implies that best-response dynamics converge to a pure Nash equilibrium in polynomial time. This solution was subsequently developed by M. Klimm and M.J. Stahlberg into their recent manuscript “Complexity of equilibria in binary public goods games on undirected graphs” (see <https://arxiv.org/pdf/2301.11849.pdf>).

Worth of mentioning is the independent and concurrent work “A Characterization of Complexity in Public Goods Games” by M. Gilboa (see <https://arxiv.org/pdf/2301.11580.pdf>).

4.4 Optimizing over Serial Dictatorships – Lower bound for optimal matching

Nidhi Rathi (*Aarhus University, DK*)

License © Creative Commons BY 4.0 International license
© Nidhi Rathi

Main reference Ioannis Caragiannis, Nidhi Rathi: “Optimizing over Serial Dictatorships”, CoRR, Vol. abs/2202.08097, 2022.

URL <https://arxiv.org/abs/2202.08097>

Consider a bipartite graph with n agents on one side and items on another. We access the weights via making following queries. On a query of the form (agent i , an ordered sequence S of agents), the oracle outputs the highest-valued item available for agent i , when she gets

to pick after the sequence of agents S . Our recent work proves that there always exists an action sequence that results into a maximum weight matching. We also prove that we can find an action sequence with $O(n^2)$ queries that produces a 2-approximate matching.

Question: Can we improve the lower bound on the number of queries or the approximation factor to the optimal matching?

4.5 Altruistic Hedonic Games (Short talk)

Jörg Rothe (*Heinrich-Heine-Universität Düsseldorf, DE*)

License  Creative Commons BY 4.0 International license
© Jörg Rothe

Hedonic games are coalition formation games in which players have preferences over the coalitions they can join. While previous models of representing hedonic games were based upon selfish players only, in a JAIR 2022 paper we have proposed a novel model for them that takes into account not only the players' own preferences but also their friends' preferences. In this talk, we briefly introduce our model, survey some known results, and outline open questions.

4.6 Tractability frontier of justice criteria under submodular valuations

Yair Zick (*University of Massachusetts – Amherst, US*)

License  Creative Commons BY 4.0 International license
© Yair Zick

When agents have binary submodular valuation functions over items, we can efficiently compute allocations that satisfy several justice criteria (leximin, EF-X, max. Nash welfare, (approximate) MMS, and more).

What happens when we move beyond binary valuations?

Recently, we were able to prove that when agents have submodular valuations with bivalued marginal gains (say, a and b), it is possible to still efficiently compute allocations that satisfy weaker notions of justice, when a divides b .

It is also known that for general submodular valuations, it is impossible to efficiently compute even relatively mild justice criteria (e.g. finding an allocation that maximizes the minimum agent welfare, also known as the Santa Claus problem), let alone stronger ones like leximin allocations.

Where is the tractability frontier? Can we efficiently compute just outcomes with trinary valuations (under some assumptions)? What happens when we assume that all marginal gains follow a certain distribution? Understanding these domains will allow us to move towards better algorithmic frameworks for fair allocation.

Participants

- Yakov Babichenko
Technion – Haifa, IL
- Gerdus Benadè
Boston University, US
- Ioannis Caragiannis
Aarhus University, DK
- Giorgos Christodoulou
Aristotle University of
Thessaloniki, GR
- Andrei Constantinescu
ETH Zürich, CH
- Michal Feldman
Tel Aviv University, IL
- Tobias Harks
Universität Passau, DE
- Martin Hofer
Goethe-Universität –
Frankfurt am Main, DE
- Max Klimm
TU Berlin, DE
- Ron Lavi
University of Bath, GB
- Pascal Lenzner
Hasso-Plattner-Institut,
Universität Potsdam, DE
- Kevin Leyton-Brown
University of British Columbia –
Vancouver, CA
- Vangelis Markakis
Athens University of Economics
and Business, GR
- Darya Melnyk
Aalto University, FI
- Noam Nisan
The Hebrew University of
Jerusalem, IL
- Sigal Oren
Ben Gurion University –
Beer Sheva, IL
- Ori Plonsky
Technion – Haifa, IL
- Maria Polukarov
King's College London, GB
- Nidhi Rathi
Aarhus University, DK
- Rebecca Reiffenhäuser
University of Amsterdam, NL
- Jörg Rothe
Heinrich-Heine-Universität
Düsseldorf, DE
- Daniel Schmand
Universität Bremen, DE
- Giovanna Varricchio
Goethe-Universität –
Frankfurt am Main, DE
- Carmine Ventre
King's College London, GB
- Roger Wattenhofer
ETH Zürich, CH
- Lisa Wilhelmi
Goethe-Universität – Frankfurt
am Main, DE
- Jens Witkowski
Frankfurt School of Finance &
Management, DE
- Yair Zick
University of Massachusetts –
Amherst, US

