Abstracts

Contents

Vittorio Bilò
University of Salento, Italy

Achieving Envy-Freeness through Items Sale

We consider a fair division setting of allocating indivisible items to a set of agents. In order to cope with the well-known impossibility results related to the non-existence of envy-free allocations, we allow the option of selling some of the items so as to compensate envious agents with monetary rewards. In fact, this approach is not new in practice, as it is applied in some countries in inheritance or divorce cases. A drawback of this approach is that it may create a value loss, since the market value derived by selling an item can be less than the value perceived by the agents. Therefore, given the market values of all items, a natural goal is to identify which items to sell so as to arrive at an envy-free allocation, while at the same time maximizing the overall social welfare. Our work is focused on the algorithmic study of this problem, and we provide both positive and negative results on its approximability. When the agents have a commonly accepted value for each item, our results show a sharp separation between the cases of two or more agents. In particular, we establish a PTAS for two agents, and we complement this with a hardness result, that for three or more agents, the best approximation guarantee is provided by essentially selling all items. This hardness barrier, however, is relieved when the number of distinct item values is constant, as we provide an efficient algorithm for any number of agents. We also explore the generalization to heterogeneous valuations, where the hardness result continues to hold, and where we provide positive results for certain special cases.

Back to Contents Page →

Martin Bullinger
University of Oxford, UK

Stability in Random Hedonic Games

Partitioning a large group of employees into teams can prove difficult because unsatisfied employees may want to transfer to other teams. In this case, the team (coalition) formation is unstable and incentivizes deviation from the proposed structure. Such a coalition formation scenario can be modeled in the framework of hedonic games and a significant amount of research has been devoted to the study of stability in such games. Unfortunately, stable coalition structures are not guaranteed to exist in general and their practicality is further hindered by computational hardness barriers. We offer a new perspective on this matter by studying a random model of hedonic games. For three prominent stability concepts based on single-agent deviations, we provide a high probability analysis of stability in the large agent limit.

Our first main result is an efficient algorithm that outputs an individually and contractually Nash-stable partition with high probability. Our second main result is that the probability that a random game admits a Nash-stable partition tends to zero. Our approach resolves the two major downsides associated with individual stability and contractual Nash stability and reveals agents acting single-handedly are usually to blame for instabilities.

(Joint work with Sonja Kraiczy)

Back to Contents Page →

Ioannis Caragiannis
Aarhus University, Denmark

Two stories about distortion in social choice

The notion of distortion has received much attention in recent years by the computational social choice community. In general, distortion quantifies how the lack of complete information affects the quality of the social choice outcome. Ideally, a distortion of 1 means that the social choice outcome is the most efficient one.

In the talk, we will consider two related scenarios. The first one is inspired by voting under the impartial culture assumption. We assume that agents have random values for the alternatives, drawn from a probability distribution independently for every agent-alternative pair. We explore voting rules that use a limited number of queries per agent in addition to the agent's ordinal information. For simple distributions, we present rules that always select an alternative of social welfare that is only a constant factor away from the optimal social welfare (i.e., rules of constant distortion).

The second scenario is motivated by the practice of sortition. Here, we assume that agents correspond to points on a metric space. Our objective is to select, in a fair manner, a subset of the agents (corresponding to a citizens' assembly) so that for every election with alternatives from the same metric space, the most preferred alternative of the citizens' assembly has a social cost that is very close to that of the optimal alternative for the whole agent population. Our positive results indicate that assemblies of size logarithmic in the number of alternatives are sufficient to get constant distortion in this model.

The talk is based on two papers that are joint works with Karl Fehrs, and with Evi Micha and Jannik Peters, respectively.

Back to Contents Page →

Diptarka Chakraborty
National University of Singapore, Singapore

Fair Rank Aggregation

Aggregating multiple input rankings over a set of candidates to generate a consensus ranking is one of the fundamental ranking problems, having many applications in social choice theory, hiring, college admission, web search, and databases. However, the optimal consensus ranking might be biased against any individual candidate or candidates belonging to certain marginalized communities or groups. This has motivated studies of the rank aggregation problem from the fairness perspective. While finding a consensus ranking, the additional objective is to ensure fair representation of each group in the top positions of the final aggregated ranking. In this talk, we will discuss various algorithms to find such a fair ranking approximately.

Back to Contents Page →

Hau Chan
University of Nebraska-Lincoln, USA

Collective Decision-Making for Improving Accessibility to Facilities

Public facilities such as schools, hospitals, and transit stations are essential for many individuals; however, access to these facilities is not the same among different populations. Millions of individuals from low-income and underrepresented communities are unable to access various public facilities due to a lack of vehicles, limited finances, and inadequate infrastructure.

To collectively improve the accessibility of individuals to facilities, the existing research literature has primarily focused on locating facilities to provide essential services subject to individuals' preferences on the locations of the facilities. Yet, creating and locating facilities can often be prohibitively expensive and time-consuming. Recognizing this challenge, our recent studies and various real-world initiatives aim to improve the accessibility of individuals to existing prelocated facilities collectively by strengthening existing infrastructures through structural modifications (e.g., constructing new roads, bridges, multi-use paths, or shuttle services).

In this talk, we will discuss our recent (theoretical and algorithmic) studies on modeling various structural modification strategies and designing (approximately optimal) strategyproof mechanisms to elicit (true) individual preferences on which communities to improve access and determine structural modification outcomes in order to improve facility accessibility.

Back to Contents Page →

Jiehua Chen
Vienna University of Technology, Austria

Fairness in Assignments with Congestion-Averse Agents: Concepts, Algorithms, and Complexity

The congested assignment problem is concerned with assigning agents to posts where agents care about both the posts and their congestion levels. Here, agents are averse to congestion, consistently preferring lower over higher congestion for the same resource. Such scenarios are prevalent across many domains, including traffic management and school choice, where fair resource allocation is crucial.

Congested assignment can be considered as a restricted variant of the Group Activity Selection problem, introduced by Darmann et al.

Additionally, it is related to many-to-one matching in matching under preferences.

In this talk, I will explore one ex-ante fairness concept, top-fairness, and two ex-post fairness concepts, envy-freeness and competitiveness.

The top-fairness and competitiveness concepts were recently introduced by Bogomolnaia and Moulin.

While a top-fair or envy-free assignment always exists and can be found easily, competitive assignments do not always exist. The talk will cover the following key points:

1. An efficient method to determine the existence of competitive or maximally competitive assignments for a given congestion profile.

2. Two optimization variants of congested assignments and their computational complexity: a) Finding a top-fair assignment that is envy-free b) Finding a top-fair assignment that is maximally competitive. Both variants are NP-hard, unfortunately.

3. Parameterized algorithms for these NP-hard problems.

Back to Contents Page →

Matthias Greger
Technical University of Munich, Germany

Preference Aggregation on the Probability Simplex

Social choice theory deals with the problem of aggregating individual preferences into a collective outcome. We consider a setting known as portioning where the set of possible outcomes consists of all lotteries over a finite set of alternatives and assume that each agent reports her favorite lottery. This talk investigates the compatibility of efficiency, strategyproofness, and fairness for various utility models (e.g., Leontief or norm-induced preferences) and covers impossibilities as well as suitable mechanisms.

Back to Contents Page →

Daniel Halpern
Harvard University, USA

Aggregating Preferences with Limited Queries

Social choice theory studies how to aggregate individual preferences into a collective decision for society. Traditionally, this assumes full access to each individual's complete set of preferences. However, modern online platforms promoting civic participation, such as pol.is, aim to solve social choice problems that do not fit neatly into this framework. These platforms aggregate complex preferences over a vast space of alternatives, rendering it infeasible to learn any individual's preferences completely. Instead, preferences are elicited by asking each user a simple query about a small subset of their preferences.

In this talk, I will present a simple model for analyzing what is possible in these scenarios, long with a variety of positive and negative results . It covers two recent papers:

EC'24 paper on ranked preferences: https://arxiv.org/abs/2402.11104
AAAI'23 paper on approval preferences: https://arxiv.org/abs/2211.15608

Contributions include:

- Positive algorithmic results: Efficient algorithms that produce representative outcomes with limited queries.

- Information-theoretic impossibilities: Fundamental limits on what can be learned, regardless of the number of queries.

- Query-complexity lower bounds: Situations where, even if it is possible in theory to achieve a desired outcome, an exponential number of queries may be required, making it practically infeasible.

Back to Contents Page →

Naoyuki Kamiyama
Kyushu University, Japan

Some recent results on super-stable matchings

Super-stability is one of the properties of a matching in the stable matching problem with ties. In this talk, I talk about some recent results on super-stable matchings. For example, I talk about the problem of modifying a given instance of the stable matching problem with ties in such a way that there exists a super-stable matching in the modified instance, and the problem of checking the existence of a super-stable matching in the setting where we are given generalized matroid constraints.

Back to Contents Page →

Alexander Lam
City University of Hong Kong, Hong Kong

Facility Location Games with Scaling Effects.

We take the classic facility location problem and consider a variation, in which each agent's individual cost function is equal to their distance from the facility multiplied by a scaling factor which is determined by the facility placement. In addition to the general class of continuous scaling functions, we also provide results for piecewise linear scaling functions which can effectively approximate or model the scaling of many real world scenarios. We focus on the objectives of total and maximum cost, describing the computation of the optimal solution. We then move to the approximate mechanism design setting, observing that the agents' preferences may no longer be single-peaked.

Consequently, we characterize the conditions on scaling functions which ensure that agents have single-peaked preferences. Under these conditions, we find results on the total and maximum cost approximation ratios achievable by strategyproof and anonymous mechanisms.

Back to Contents Page →

David Manlove
University of Glasgow, UK

Couples can be tractable: New algorithms and hardness results for the Hospitals / Residents problem with Couples

The Hospitals / Residents problem with Couples (HRC) models the problem of allocating junior doctors to hospital posts, which has applications in centralised clearinghouses around the world. One such example is the National Resident Matching Program (NRMP) in the US, which handles over 40,000 applicants annually. In HRC, a solution is a stable matching or a report that none exists.

Informally, a stable matching ensures that no single doctor or couple, and no single hospital or to a pair of hospitals, would prefer to be matched to one another than to remain with their current assignment/s.

A key challenge is that an HRC instance need not admit a stable matching. We present a novel polynomial-time algorithm that can always find a near-feasible stable matching (adjusting the hospitals' capacities by at most 1) in an HRC instance where the couples' preferences are subresponsive (i.e., if one member switches to a better hospital, than the couple also improves) and subcomplete (i.e., each pair of hospitals that are individually acceptable to both members are jointly acceptable for the couple) by reducing it to an instance of the Stable Fixtures problem. We also present a polynomial-time algorithm for HRC in a sub-responsive, sub-complete instance that is a Dual Market, or where all couples are one of several possible types. We complement our algorithms with several hardness results. We show that HRC with subresponsive and sub-complete couples is NP-hard, even with other strong restrictions. We also show that HRC with a Dual Market is NP-hard under several simultaneous restrictions. Finally, we show that the problem of finding a matching with the minimum number of blocking pairs in HRC is not approximable within $m1-\varepsilon$, for any $\varepsilon>0$, where m is the total length of the hospitals' preference lists, unless P=NP, even if each couple applies to only one pair of hospitals. Our polynomial-time solvability results greatly expand the class of known tractable instances of HRC and provide additional evidence as to why long-standing entry-level labour markets that allow couples such as the NRMP remain successful to this day.

This is joint work with Gergely Csáji, Iain McBride and James Trimble. An extended abstract appears in the Proceedings of IJCAI 2024. The full paper can be found at https://arxiv.org/abs/2311.00405.

Back to Contents Page →

Vangelis Markakis
Athens University of Economics and Business, Greece

Winner Determination and Strategic Control in Conditional Approval Voting

The focus of this work is on multi-issue elections with preferential dependencies. We study a generalization of the classic Minisum approval voting rule, introduced by Barrot and Lang (2016), and referred to as Conditional Minisum (CMS), where the voters are allowed to submit conditional ballots and declare dependencies between different issues. Not surprisingly, the price we have to pay for this higher level of expressiveness is that we end up with a computationally hard rule. Motivated by this, we first focus on finding special cases that admit efficient algorithms for CMS. Our main result in this direction is that we identify the condition of bounded treewidth (of an appropriate graph, emerging from the provided ballots) as the necessary and sufficient condition for exact polynomial algorithms, under common complexity assumptions.

We then move to the design of approximation algorithms, where for the (still hard) case of binary issues, we identify restrictions, under which we provide the first multiplicative approximation algorithms for the problem. Finally, we also investigate the complexity of problems related to the strategic control of conditional approval elections by adding or deleting either voters or alternatives and we show that in most variants of these problems, CMS is computationally resistant against control. Overall, we conclude that CMS can be viewed as a solution with a satisfactory tradeoff between expressiveness and computational efficiency, when we have a limited number of dependencies among issues.

Back to Contents Page →

Jannik Peters
National University of Singapore, Singapore

Proportional Clustering and Social Choice

We will give an overview over the recently emerging field of proportional clustering. We will discuss different solution concepts including proportional fairness and individual fairness and discuss how these fairness notions are related. We further show how proportional multiwinner voting can be used to design proportional clustering algorithms and fairness notions. Finally, we give an application of proportional clustering in the setting of distortion and discuss how proportionality can relate to the decision-making quality of a partitioned panel.

Back to Contents Page →

Ulrike Schmidt-Kraepelin
TU Eindhoven, Netherlands

Monotone Randomized Apportionment

Apportionment is the act of distributing the seats of a legislature among political parties (or states) in proportion to their vote shares (or populations). A famous impossibility by Balinski and Young (2001) shows that no apportionment method can be proportional up to one seat (quota) while also responding monotonically to changes in the votes (population monotonicity). Grimmett (2004) proposed to overcome this impossibility by randomizing the apportionment, which can achieve quota as well as perfect proportionality and monotonicity — at least in terms of the expected number of seats awarded to each party. Still, the correlations between the seats awarded to different parties may exhibit bizarre non-monotonicities. When parties or voters care about joint events, such as whether a coalition of parties reaches a majority, these non-monotonicities can cause paradoxes, including incentives for strategic voting. We propose monotonicity axioms ruling out these paradoxes, and study which of them can be satisfied jointly with Grimmett's axioms. Essentially, we require that, if a set of parties all receive more votes, the probability of those parties jointly receiving more seats should increase. Our work draws on a rich literature on unequal probability sampling in statistics (studied as dependent randomized rounding in computer science). Our main result shows that a sampling scheme due to Sampford (1967) satisfies Grimmett's axioms and a notion of higher-order correlation monotonicity.

This talk is based on joint work with José Correa, Paul Gölz, Jamie Tucker-Foltz, and Victor Verdugo.

Back to Contents Page →

Piotr Skowron
University of Warsaw, Poland

Generalised Theory of Proportionality in Collective Decision Making

We consider a voting model, where a number of candidates need to be selected subject to certain feasibility constraints. The model generalises committee elections (where there is a single constraint on the number of candidates that need to be selected), various elections with diversity constraints, the model of public decisions (where decisions needs to be taken on a number of independent issues), and the model of collective scheduling. A critical property of voting is that it should be fair -- not only to individuals but also to groups of voters with similar opinions on the subject of the vote; in other words, the outcome of an election should proportionally reflect the voters' preferences. We formulate axioms of proportionality in this general model. Our axioms do not require predefining groups of voters; to the contrary, we ensure that the opinion of every subset of voters whose preferences are cohesive-enough are taken into account to the extent that is proportional to the size of the subset. Our axioms generalise the strongest known satisfiable axioms for the more specific models. We explain how to adapt two prominent committee election rules, Proportional Approval Voting (PAV) and Phragmén Sequential Rule, as well as the concept of stable-priceability to our general model. The two rules satisfy our proportionality axioms if and only if the feasibility constraints are matroids.

Back to Contents Page →

Zhaohong Sun
Kyushu University, Japan

Stable Matching in Practice: Daycare Matching Markets in Japan

---

We study a practical two-sided matching problem of allocating children to daycare centers, which has significant social implications. We are cooperating with several municipalities in Japan and our goal is to devise a reliable clearing algorithm to deal with the problem. In this presentation, I will describe the design of our new algorithm that minimizes the number of unmatched children while ensuring stability. We evaluate our algorithm using real-life data sets, and experimental results demonstrate that our algorithm surpasses the commercial software that currently dominates the market in terms of both the number of matched children and the number of blocking coalitions (measuring stability). Our findings have been reported to local governments, and one city has just adopted our proposed algorithm instead of the existing solution. Moreover, our model and algorithm have broader applicability to other important matching markets, such as hospital-doctor matching with couples and school choice with siblings.

Back to Contents Page →

Yu Yokoi
Tokyo Institute of Technology, Japan

Popular Matching under Matroid and Optimality Constraints

Popularity of matching is an adaptation of the concept of a weak Condorcet winner to the setting of matching under preferences. This talk presents recent algorithmic results on popular matching problems under various constraints, such as matroid constraints and size constraints. The framework also includes the popular arborescence problem, which has a connection to liquid democracy.

Back to Contents Page →

William Zwicker
Union College, USA

The Median Procedure – a Universal Aggregation Rule?

The median procedure MP (Barthélemy and Monjardet, 1981) aggregates a sequence of binary relations from some input class I into a single relation (with ties allowed) in some output class O. Varying the choice of I and O gives rise to a remarkable range of known rules as special cases of MP, including:

(1) Plurality voting,
(2) Approval voting,
(3) Kemeny voting,
(4) Borda voting (with outcome a winner),
(5) Mirkin aggregation of equivalence relations (a form of cluster analysis),

but not

(6) Borda voting (with outcome a ranking),
(7) The Mean Rule (Duddy and Piggins),
(8) j,k-Kemeny (a version of Kemeny for weak orders), or
(9) Any of the known Condorcet extensions: Copeland, minimax, etc.

MP is usually defined by choosing the relation in O "closest" (using a form of Hamming distance) to the inputs. But an alternative formulation using inner (aka, "dot") product and orthogonal decomposition is better equipped to explain how and why computational complexity varies among the rules listed above, and why rules (6), (7) and (8) – but not (9) – also arise from MP when an extra projection step is inserted. This formulation suggests that rules (1) - (8) all aggregate information in essentially the same way, but differ with regard to which dimensions of information are taken into account.

*Union College Mathematics Department and Murat Sertel Center for Advanced Economic

Back to Contents Page →