

10th International Workshop on Computational Social Choice **COMSOC 2025**

September 17-19 • TU Wien • Austria



Organizing Committee

Péter Biró
Phuc Hung Hoang
Oliviero Nardi
Johannes Strasser

Jiehua Chen (chair)
Liana Khazaliya
Sofia Simola
Doris Brazda

Christian Hatschka
Merisa Mustajbašić
Manuel Sorge

Steering Committee

Omer Lev
William Zwicker

Reshef Meir

Marcus Pivato



W|W|T|F

Vienna Science
and Technology Fund

Artificial Intelligence Journal



Workshop Venue

TUtheSky, TU Wien, Vienna, Austria

Address: Getreidemarkt building of TU Wien, 11th floor
Getreidemarkt 9, 1060 Wien, Austria

COMSOC 2025 Program Committee

Péter Biró (PC co-chair)	Jiehua Chen (PC co-chair)	Stephane Airiau
Haris Aziz	Niclas Boehmer	Florian Brandl
Felix Brandt	Robert Brederick	Markus Brill
Martin Bullinger	Ioannis Caragiannis	Noam Hazon
Andreas Darmann	Ronald de Haan	Bart de Keijzer
Edith Elkind	Ulle Endriss	Piotr Faliszewski
Zack Fitzsimmons	Till Fluschnik	Rupert Freeman
Paul Goldberg	Judy Goldsmith	Umberto Grandi
Davide Grossi	Katarína Cechlárová	Sarit Kraus
Martin Lackner	Jérôme Lang	Kate Larson
Patrick Lederer	David Manlove	Nicholas Mattei
Nicolas Maudet	Reshef Meir	Neeldhara Misra
Shuichi Miyazaki	Hervé Moulin	Thanh Nguyen
Ali Ozkes	Dominik Peters	Jannik Peters
Marcus Pivato	Maria Polukarov	Clemens Puppe
Joerg Rothe	Ildikó Schlotter	Sanjukta Roy
Ulrike Schmidt-Kraepelin	Erel Segal-Halevi	Piotr Skowron
Warut Suksumpong	Stanisław Szufa	Nimrod Talmon
Zoi Terzopoulou	Magdalena Tydrichova	Tomasz Wąs
Anaëlle Wilczynski	Lirong Xia	Yongjie Yang
Yair Zick	William Zwicker	

Contents

1	COMSOC 2025 Program	1
2	Keynote Talk Kurt Mehlhorn	3
3	Keynote Talk Axel Ockenfels	4
4	Keynote Talk Moon Duchin	5
6	Oral Presentation Paper Abstracts	6
6.1	September 17, Morning Session	6
6.1	September 17, Afternoon Session	8
6.2	September 18, Morning Session	9
6.2	September 18, Afternoon Session	11
6.4	September 19, Morning Session	15
6.4	September 19, Afternoon Session	17
7	Poster Presentation Paper Abstracts	22
	Locations	39

Preface

Marking the 10th edition in the International Workshop Series on **COM**putational **SO**cial Choice, **COMSOC**, we are most excited to hold the COMSOC 2025 workshop on September 17–19, 2025, at TU Wien, in Vienna, Austria. Vienna’s intellectual heritage includes the Austrian School of Economics and its systematic study of individual choice and market behavior. Today, TU Wien continues this tradition of rigorous research through its outstanding work in logic, mathematics, and computer science.

COMSOC is an interdisciplinary research field at the interface of computer science, artificial intelligence, political and social sciences, and economics. It explores social choice and collective decision-making problems through the lenses of algorithms and complexity, discrete mathematics, combinatorial optimization, game theory, mechanism design, and economics. A key objective of the COMSOC workshop is to bring together research communities from these related areas, fostering cross-pollination of ideas and methodologies. Another important aim is to convey the excitement of recent research and emerging application areas, exposing participants to new techniques, novel theoretical frameworks, and real-world problems.

COMSOC is a rapidly growing research field. The workshop is non-archival, allowing authors to submit their work to other venues, which encourages broad participation and the sharing of cutting-edge research. This year, we received a record number of 102 submissions. From the 98 non-desk-rejected submissions, 39 submissions have not been published or are under submission anywhere, 28 are currently under submission to a conference or journal, and the remaining 31 have been published in a peer-reviewed conference or journal. The competition was exceptionally strong—papers required an average score of 7+ (accept) for oral presentation and 6+ (weak accept) for poster presentation. We ultimately accepted 33 papers for oral presentation and 37 for poster presentations. The accepted papers showcase the current state-of-the-art across COMSOC’s diverse research landscape, including axiomatic analysis of social choice methods, strategic behavior and manipulation in voting, multi-winner and committee elections, participatory budgeting mechanisms, fair resource allocation algorithms, preference aggregation procedures, matching under preferences, coalition formation, and experimental and empirical studies of related problems.

We thank our distinguished invited speakers—Moon Duchin, Kurt Mehlhorn, and Axel Ockenfels—whose expertise spans mathematical modeling of electoral systems, algorithmic game theory, and behavioral economics. We are grateful to our Program Committee and additional reviewers for their rigorous evaluation process, and to all authors for their contributions to advancing this vital field.

Additionally, we acknowledge our Organizing Committee’s dedicated efforts and thank TU Wien, the Vienna Science and Technology Fund (WWTF), Artificial Intelligence Journal (AIJ), and the European Association for Artificial Intelligence (EurAI) for their generous financial and logistical support.

This booklet contains abstracts of all invited keynote talks, and all accepted papers for oral and poster presentations, providing an overview of the research presented at COMSOC 2025. We hope it serves as both a valuable record of the current state of COMSOC research and a source of inspiration for future work in this rapidly evolving field.

COMSOC 2025 Organizing and Program Committee co-chairs:
Péter Biró and Jiehua Chen

COMSOC 2025 Program

<https://www.ac.tuwien.ac.at/comsoc2025/>

Main Venue: Getreidemarkt 9, 11th floor, 1060 Wien, Austria



Lunches

Time: 12:05-14:00
 Location: Mensa at Freihaus
 Wiedner Hauptstraße 8–10,
 2nd floor, 1040 Wien
 (see page 39)
 Transport: 10 minute walk from venue.

Parliament Tour

Date: Wednesday, September 17
 Time: 15:00-18:00
 Location: Dr.-Karl-Renner-Ring 3,
 1017 Wien (see page 40)
 Transport: 15-20 minute walk from
 venue.

Poster Session

Date: Wednesday, September 17
 Time: 18:30–21:00
 Location: Main venue

Group Photo

Date: Thursday, September 18
 Time: 17:30-18:00
 Location: Main venue

Workshop Dinner

Date: Thursday, September 18
 Time: 18:30-21:30
 Location: Stöckl im Park
 Prinz-Eugen-Straße 25,
 1030 Wien (see page 39)
 Website: <https://www.stoecklimpark.at>
 Transport: 25-30 minute walk from
 venue.

Wednesday, September 17

08:30–09:15 Registration and Welcoming

09:15–10:15 Keynote: **Kurt Mehlhorn**
Fair Division of Indivisible Goods

10:15–10:45 Coffee Break

10:45–11:05 Akrami, **Rathi**: *Epistemic EFX Allocations Exist for Monotone Valuations*

11:05–11:25 **Conati**, Niskanen, de Haan, Järvisalo:
Computing Efficient and Envy-Free Allocations under Dichotomous Preferences using SAT

11:25–11:45 **Cookson**, Ebadian, Shah: *Temporal Fair Division*

11:45–12:05 Bouveret, **Gilbert**, Lang, Méroué: *Constrained Serial Dictatorships can be Fair*

12:05–14:00 Lunch Break

14:00–14:20 Hornischer, **Terzopoulou**: *Learning How to Vote with Principles: Axiomatic Insights Into the Collective Decisions of Neural Networks*

14:20–14:40 Cailloux, **Hervouin**, Ozkes, Sanver: *Anonymous and neutral classification aggregation*

14:40–15:00 **Ravier**, Konieczny, Moretti, Viappiani:
From order lifting to social ranking: retrieving individual rankings from partial lifted preferences

15:00–18:00 Parliament Tour

18:30–21:00 Poster Session

Thursday, September 18		Friday, September 19	
08:45–09:45	Keynote: Axel Ockenfels <i>Behavioral Market Design</i>	08:45–09:45	Keynote: Moon Duchin <i>Beyond Sushi: Bridging from COMSOC to Practical Democracy</i>
09:45–10:15 Coffee Break		09:45–10:15 Coffee Break	
10:15–10:35	Dery : <i>Interactive and Iterative Peer Assessment: A Social Choice Approach to Eliciting and Aggregating Student Evaluations</i>	10:15–10:35	Brandt, Dong , Peters: <i>Condorcet-Consistent Choice Among Three Candidates</i>
10:35–10:55	Horn , Nüsken, Rothe, Seeger: <i>Skating System Unveiled: Exploring Preference Aggregation in Ballroom Tournaments</i>	10:35–10:55	Meir , Ghaime: <i>On Condorcet's Jury Theorem with Abstention</i>
10:55–11:15	Ghosh, Neoh, Teh, Tyrovolas : <i>Fraud-Proof Revenue Division on Subscription Platforms</i>	10:55–11:15	Ding , Yu: <i>An Axiomatic Characterization of the Minimax Voting Method</i>
11:15–11:25 Short Break		11:15–11:25 Short Break	
11:25–11:45	Delemazure , Peters, Freeman, Lang, Laslier: <i>Re-allocating Wasted Votes in Proportional Parliamentary Elections with Thresholds</i>	11:25–11:45	Berker , Casacuberta, Robinson, Ong, Conitzer, Elkind: <i>From Independence of Clones to Composition Consistency: A Hierarchy of Barriers to Strategic Nomination</i>
11:45–12:05	Boehmer , Fish, Procaccia: <i>Generative Social Choice: The Next Generation</i>	11:45–12:05	Schmidt-Kraepelin, Suksompong, Utke : <i>Discrete Budget Aggregation: Truthfulness and Proportionality</i>
12:05–14:00 Lunch Break		12:05–14:00 Lunch Break	
14:00–14:20	Prévotat , Terzopoulou, Zylbersztein: <i>Strategizing under Rule and Vote Uncertainty: An Experiment</i>	14:00–14:20	Dong, Frank , Peters, Suksompong: <i>Reconfiguring Proportional Committees</i>
14:20–14:40	Boehmer, Faliszewski, Janeczko, Kaczmarczyk, Lisowski, Pierczyński, Rey, Stolicki, Szufa , Wąs: <i>Guide to Numerical Experiments on Elections in COMSOC</i>	14:20–14:40	Kehne , Schmidt-Kraepelin, Sornat: <i>Robust Committee Voting, or The Other Side of Representation</i>
14:40–15:00	Faliszewski , Janeczko, Kaczmarczyk, Kurdziel, Pierczyński, Szufa: <i>Learning Real-Life Approval Elections</i>	14:40–15:00	Masařík, Pierczyński, Skowron : <i>A Generalised Theory of Proportionality in Collective Decision Making</i>
15:00–15:30 Coffee Break		15:00–15:30 Coffee Break	
15:30–15:50	Aziz, Lederer , Peters, Peters, Ritossa: <i>Committee Monotonicity and Proportional Representation for Ranked Preferences</i>	15:30–15:50	Papasotiropoulos, Pishbin , Skibski, Skowron, Wąs: <i>Method of Equal Shares with Bounded Overspending</i>
15:50–16:10	Bardal , Brill, McCune, Peters: <i>Proportionality in Practice: Quantifying Proportionality in Ordinal Elections</i>	15:50–16:10	Kraicz, Robinson, Elkind : <i>Streamlining Equal Shares</i>
16:10–16:30 Short Break		16:10–16:30	Baychkov , Brill, Utke: <i>Mixed Voting Rules for Participatory Budgeting</i>
16:30–16:50	Bullinger, Romen, Schlenga : <i>The Power of Matching for Online Fractional Hedonic Games</i>	16:30–16:50 Short Break	
16:50–17:10	Caragiannis , Roy: <i>Quantile Agent Utility and Implications to Randomized Social Choice</i>	16:50–17:10	Nguyen , Song: <i>Approximate Core of Participatory Budgeting via Lindahl Equilibrium</i>
17:10–17:30	Constantinescu , Wattenhofer: <i>Byzantine Game Theory: Sun Tzu's Boxes</i>	17:10–17:30	Faliszewski, Janeczko , Knop, Pokorný, Schierreich, Śluszniak, Sornat: <i>Participatory Budgeting Project Strength via Candidate Control</i>
17:30–18:00 Group Photo		17:30–19:00 Business Meeting	
18:30–22:00 Workshop Dinner			

Keynote Talk Kurt Mehlhorn

Fair Division of Indivisible Goods

Kurt Mehlhorn (Saarland University, Germany)

Wednesday, September 17, 09:15–10:15

Abstract. A set of indivisible goods, e.g., a car, a house, a toothbrush, . . . , has to be split among a set of agents in a *fair manner*. Each agent has its own valuation function for sets of goods. What constitutes a fair allocation? When does a fair allocation exist? If it exists, can we compute it efficiently? Can we approximate fair allocations? There are three main notions of fairness in the literature: Envy-based, Share-based, and maximum Nash Social Welfare. I will discuss all three notions and survey what is known.

Biography. Kurt Mehlhorn is Director Emeritus of the MPI for Informatics and Senior Professor of Computer Science at Saarland University. He headed the algorithms and complexity group at the MPI for Informatics. He is the author of six books and over 300 scientific publications, which include fundamental contributions to data structures, computational geometry, computer algebra, parallel computing, VLSI design, computational complexity, combinatorial optimization, graph algorithms, and fair allocation. He is also one of the people behind the LEDA software library.

Recently, he produced a series of video lectures “Ideas and Concepts of Computer Science” (in German) for a general audience. He supervised more than 90 students, many of whom have now faculty positions. He has received several prizes (Leibniz Award, EATCS Award, Zuse Medal, ACM Paris Kanellakis Theory and Practice Award, Erasmus Medal of the Academia Europaea) for his work. He holds Honorary Doctorate Degrees from Magdeburg, Aalto, Waterloo, Aarhus, Gothenburg and Patras universities and is an ACM Fellow. He is a member of the German Academy of Sciences Leopoldina, Academia Europaea, the German Academy of Science and Engineering acatech, the US Academy of Engineering, the Indian Academy of Engineering, and the US Academy of Science. From 2002 to 2008, he was vice president of the Max Planck Society. He is a co-founder of Algorithmic Solutions Software GmbH. He currently serves on the Research Advisory Board of Tata Consultancy, and the advisory board of Carnegie Mellon, Qatar.

Keynote Talk Axel Ockenfels

Behavioral Market Design

Axel Ockenfels (University of Cologne, Germany)

Thursday, September 18, 08:45–09:45

Abstract. Addressing pressing economic and social challenges – such as pandemics, climate change, and energy scarcity – requires changes in behavior. In this talk, I will use case studies, primarily from my own research, to illustrate how human behavior and bounded rationality influences the design of institutions aimed at aligning incentives and actions with overarching goals. Economic design research and behavioral science are often complementary, rather than substitutes, in promoting effective behavioral change.

Biography. Axel Ockenfels is Professor of Economics at the University of Cologne and Director at the Max Planck Institute for Research on Collective Goods in Bonn. He is also Director of the Cologne Laboratory of Economic Research (CLER), and Principal Investigator in the Center for Social and Economic Behavior (C-SEB) and the Excellence Cluster ECONtribute. He held visiting positions at Penn State, Harvard University, Stanford University, and UC San Diego.

Ockenfels' research focuses on behavioral economic engineering, which combines research in game theory and behavioral research to design markets, algorithms and competitive strategies.

Ockenfels is a member of the German National Academy of Sciences Leopoldina, the Berlin-Brandenburg Academy of Sciences and Humanities, the North Rhine-Westphalian Academy of Sciences, Humanities and the Arts, and of the National Academy of Science and Engineering acatech. He is also member of the Economists' Roundtable and of the Climate Economists' Roundtable at the Federal Chancellery (Bundeskanzleramt), of the Academic Advisory Board at the Federal Ministry for Economic Affairs and Climate Action (BMWK) and of the Specialist Group for Regulatory Issues at the Federal Network Agency (Bundesnetzagentur). He has received several prizes (Gottfried Wilhelm Leibniz Award, Gossen Prize of the German Economic Association, Zukunftspreis of the University of Cologne, and the Exeter Prize) for his work.

Keynote Talk Moon Duchin

Beyond Sushi: Bridging from COMSOC to practical democracy

Moon Duchin (Cornell University, NY USA)

Friday, September 19, 08:45–09:45

Abstract. In this talk, I'll start with a (qualified) success story for computational tools in democracy reform: A case study of U.S. redistricting, where graph algorithms have made surprising impacts on legal work around voting rights. Then I'll explain how perspectives from law and policy can help guide approaches to core questions in other areas of computational social choice, including definitions of key concepts like cohesion and proportionality. As one theme, I'll develop some descriptive statistics for preference profiles from political elections, which helps frame the challenges of building generative models of preference that better align with messy real-world behavior.

Biography. Moon Duchin is Professor of Mathematics at Cornell university. She is a member of the faculty at the Brooks School and the Department of Mathematics, affiliated with the Center for Data Science for Enterprise and Society. Duchin was hired as part of the Provost Office's Radical Collaboration initiative.

A prominent voice on fair redistricting, Duchin has developed mathematical models to analyze the potential and actual outcomes of changes to policy and voting districts. She has served as an expert in redistricting litigation in Wisconsin, North Carolina, Alabama, South Carolina, Pennsylvania, Texas, and Georgia. Recently, her work has turned to the study of alternative systems of election.

Her research has been recognized with a National Science Foundation Faculty Early Career Development Award (NSF CAREER), a Guggenheim Fellowship, and a Radcliffe Fellowship. She was a Sloan Professor at the Simons-Laufer Mathematical Sciences Research Institute (SLMath) for their recent program on Algorithms, Fairness, and Equity. She is a fellow of the American Mathematical Society.

Oral Presentation Paper Abstracts

Wednesday, September 17, 10:45–11:05

[77] **Epistemic EFX Allocations Exist for Monotone Valuations**

Authors. Hannaneh Akrami, **Nidhi Rathi**

Abstract. We study the fundamental problem of fairly dividing a set of indivisible items among agents with (general) monotone valuations. The notion of envy-freeness up to any item (EFX) is considered to be one of the most fascinating fairness concepts in this line of work. Unfortunately, despite significant efforts, existence of EFX allocations is a major open problem in fair division, thereby making the study of approximations and relaxations of EFX a natural line of research. Recently, Caragiannis et al. (2023) introduced a promising relaxation of EFX, called epistemic EFX (EEFX). An allocation is EEFX, if for every agent, it is possible to shuffle the items in the remaining bundles so that she becomes “EFX-satisfied”. Caragiannis et al. (2023) prove existence and polynomial-time computability of EEFX allocations for additive valuations. A natural question asks what happens when we consider valuations more general than additive? We address this important open question and answer it affirmatively by establishing the existence of EEFX allocations for an arbitrary number of agents with general monotone valuations. To the best of our knowledge, besides EF1, EEFX is the only known relaxation of EFX to have such strong existential guarantees. Furthermore, we complement our existential result by proving computational and information-theoretic lower bounds. We prove that even for an arbitrary number of (more than one) agents with identical submodular valuations, it is PLS-hard to compute EEFX allocations and it requires exponentially-many value queries to do so.

Wednesday, September 17, 11:05–11:25

[79] **Computing Efficient and Envy-Free Allocations under Dichotomous Preferences using SAT**

Authors. **Ari Conati**, Andreas Niskanen, Ronald de Haan, Matti Järvisalo

Abstract. We study the problems of computing envy-free Pareto-efficient allocations in the context of fair allocation and hedonic games under dichotomous preferences. We establish Σ_2^P -completeness of deciding the existence of envy-free Pareto-efficient allocations, refining earlier related results. We also develop iterative SAT-based exact algorithms for computing envy-free Pareto-efficient allocations, and extend the approach to computing minimum-envy Pareto-efficient allocations under different combinations of aggregation functions. We provide open-source implementations of the algorithms and show empirically that the approach scales to computing envy-free Pareto-efficient allocations up to hundreds of agents. This work has been published at the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025).

Wednesday, September 17, 11:25–11:45

[82] Temporal Fair Division

Authors. **Benjamin Cookson**, Soroush Ebadian, Nisarg Shah

Abstract. We study temporal fair division, whereby a set of agents are allocated a (possibly different) set of goods on each day for a period of days. We study this setting, as well as a number of its special cases formed by the restrictions to two agents, same goods on each day, identical preferences, or combinations thereof, and chart out the landscape of achieving two types of fairness guarantees simultaneously: fairness on each day (per day) and fairness over time (up to each day, or the weaker version, overall).

In the most general setting, we prove that there always exists an allocation that is stochastically-dominant envy-free up to one good (SD-EF1) per day and proportional up to one good (PROP1) overall, and when all the agents have identical preferences, we show that SD-EF1 per day and SD-EF1 overall can be guaranteed. For the case of two agents, we prove that SD-EF1 per day and EF1 up to each day can be guaranteed using an envy balancing technique. We provide counterexamples for other combinations that establish our results as among the best guarantees possible, but also leave open some tantalizing questions.

Wednesday, September 17, 11:45–12:05

[98] Constrained Serial Dictatorships can be Fair

Authors. Sylvain Bouveret, **Hugo Gilbert**, Jérôme Lang, Guillaume M  rou  

Abstract. When allocating indivisible items to agents, it is known that the only strategyproof mechanisms that satisfy a set of rather mild conditions are constrained serial dictatorships: given a fixed order over agents, at each step the designated agent chooses a given number of items (depending on her position in the sequence). Agents who come earlier in the sequence have a larger choice of items; however, this advantage can be compensated by a higher number of items received by those who come later. How to balance priority in the sequence and number of items received is a nontrivial question. We use a previous model, parameterized by a mapping from ranks to scores, a social welfare functional, and a distribution over preference profiles. For several meaningful choices of parameters, we show that the optimal sequence can be computed exactly in polynomial time or approximated using sampling. Our results hold for several probabilistic models on preference profiles, with an emphasis on the Plackett-Luce model. We conclude with experimental results showing how the optimal sequence is impacted by various parameters.

Wednesday, September 17, 14:00–14:20

[11] Learning How to Vote with Principles: Axiomatic Insights Into the Collective Decisions of Neural Networks

Authors. Levin Hornischer, **Zoi Terzopoulou**

Abstract. Can neural networks be applied in voting theory, while satisfying the need for transparency in collective decisions? We propose axiomatic deep voting: a framework to build and evaluate neural networks that aggregate preferences, using the well-established axiomatic method of voting theory. Our findings are: (1) Neural networks, despite being highly accurate, often fail to align with the core axioms of voting rules, revealing a disconnect between mimicking outcomes and reasoning. (2) Training with axiom-specific data does not enhance alignment with those axioms. (3) By solely optimizing axiom satisfaction, neural networks can synthesize new voting rules that often surpass and substantially differ from existing ones. This offers insights for both fields: For AI, important concepts like bias and value-alignment are studied in a mathematically rigorous way; for voting theory, new areas of the space of voting rules are explored.

Wednesday, September 17, 14:20–14:40

[36] Anonymous and neutral classification aggregation

Authors. Olivier Cailloux, **Matthieu Hervouin**, Ali I. Ozkes, Remzi Sanver

Abstract. Based on previous results in preference aggregation, we explore the possibility of defining anonymous and neutral aggregators in classification aggregation. We find a necessary condition on the number of individuals, objects, and categories for the existence of anonymous and neutral aggregators and propose such aggregators. We prove that this condition is tight for an equal number of individuals and objects. Experimental evidence suggests this tightness extends to cases with different numbers of individuals and objects, though this remains a conjecture requiring formal proof.

Wednesday, September 17, 14:40–15:00

[81] From order lifting to social ranking: retrieving individual rankings from partial lifted preferences

Authors. **Ariane Ravier**, Sébastien Konieczny, Stefano Moretti, Paolo Viappiani

Abstract. Several methods have been introduced in the literature to extend preferences over items from a population to preferences over the groups they may form - a problem known as the Order Lifting problem. The converse matter of deducing preferences over items from expressed preferences over the coalitions that may be formed within a population - a problem known as the Social Ranking Problem - has also been studied more recently.

In this paper, we investigate the links between these two problems: after an examination of the general case, we consider the impact of missing information, by studying which social ranking methods allow for the most accurate recovery of initial preferences over items, depending on the amount of missing information about coalitions. Finally, we consider the specific case in which preferences are only expressed over coalitions of same size, and examine the accuracy of social ranking methods when facing partial information about preferences over k -sized coalitions.

Thursday, September 18, 10:15–10:35

[5] Interactive and Iterative Peer Assessment: A Social Choice Approach to Eliciting and Aggregating Student Evaluations

Authors. **Lihi Dery**

Abstract. We present a peer assessment model grounded in computational social choice, aimed at mitigating grade inflation and strategic behavior in the evaluation of in-class project presentations. The model elicits individual evaluations through a structured rating process, augmented with pairwise comparison queries, to produce complete and tie-free rankings for each student. Aggregation proceeds in two stages: first, computing the median score for each project; second, applying a ranking-based voting rule to resolve ties. This hybrid approach combines the interpretability of score-based methods with the robustness of ordinal social choice mechanisms. Deployed in a university course, the model significantly reduced both grade inflation and cognitive load, demonstrating the value of integrating elicitation design with voting-based aggregation in peer assessment contexts.

Thursday, September 18, 10:35–10:55

[31] Skating System Unveiled: Exploring Preference Aggregation in Ballroom Tournaments

Authors. **Laryssa Horn**, Paul Nüsken, Jörg Rothe, Tessa Seeger

Abstract. The Skating System, which originated from the scrutineering system in dance sport tournaments, can be formulated as a voting system: We introduce and formalize the Skating System Single (SkS, for short), a new voting system embedded into the framework of computational social choice. Although SkS has similarities with Bucklin voting, it differs from it because it is subject to additional constraints when determining the election winners. Through an analysis of the axiomatic properties of SkS and of its vulnerability to manipulative and electoral control attacks, we compare SkS with Bucklin voting and provide insights into its potential strengths and weaknesses. In particular, we show that SkS satisfies nondictatorship as well as the majority criterion, positive responsiveness, monotonicity, and citizens' sovereignty but violates the Condorcet criterion, strong monotonicity, independence of clones, consistency, participation, resoluteness, and strategy-proofness. Further, we study manipulation—showing that the constructive coalitional weighted manipulation problem for SkS is NP-complete, while the destructive variant can be solved in polynomial time—and initiate the study of control.

Thursday, September 18, 10:55–11:15

[94] Fraud-Proof Revenue Division on Subscription Platforms

Authors. Abheek Ghosh, Tzeh Yuan Neoh, Nicholas Teh, **Giannis Tyrovolas**

Abstract. We study a model of subscription-based platforms where users pay a fixed fee for unlimited access to content, and creators receive a share of the revenue. Existing approaches to detecting fraud predominantly rely on machine learning methods, engaging in an ongoing arms race with bad actors. We explore revenue division mechanisms that inherently disincentivize manipulation. We formalize three types of manipulation-resistance axioms and examine which existing rules satisfy these. We show that a mechanism widely used by streaming platforms, not only fails to prevent fraud, but also makes detecting manipulation computationally intractable. We also introduce a novel rule, ScaledUserProp, that satisfies all three manipulation-resistance axioms. Finally, experiments with both real-world and synthetic streaming data support ScaledUserProp as a fairer alternative compared to existing rules.

Thursday, September 18, 11:25–11:45

[9] Reallocating Wasted Votes in Proportional Parliamentary Elections with Thresholds

Authors. **Théo Delemazure**, Dominik Peters, Rupert Freeman, Jérôme Lang, Jean-François Laslier

Abstract. In many proportional parliamentary elections, electoral thresholds (typically 3-5%) are used to promote stability and governability by preventing the election of parties with very small representation. However, these thresholds often result in a significant number of “wasted votes” cast for parties that fail to meet the threshold, which reduces representativeness. One proposal is to allow voters to specify replacement votes, by either indicating a second choice party or by ranking a subset of the parties, but there are several ways of deciding on the scores of the parties (and thus the composition of the parliament) given those votes. We introduce a formal model of party voting with thresholds, and compare a variety of party selection rules axiomatically, and experimentally using a dataset we collected during the 2024 European election in France. We identify three particularly attractive rules, called Direct Winners Only (DO), Single Transferable Vote (STV) and Greedy Plurality (GP).

Thursday, September 18, 11:45–12:05

[62] Generative Social Choice: The Next Generation

Authors. **Niclas Boehmer**, Sara Fish, Ariel D. Procaccia

Abstract. A key task in certain democratic processes is to produce a concise slate of statements that proportionally represents the full spectrum of user opinions. This task is similar to committee elections, but unlike traditional settings, the candidate set comprises all possible statements of varying lengths, and so it can only be accessed through specific queries. Combining social choice and large language models, prior work has approached this challenge through a framework of generative social choice. We extend the framework in two fundamental ways, providing theoretical guarantees even in the face of approximately optimal queries and a budget limit on the overall length of the slate. Using GPT-4o to implement queries, we showcase our approach on datasets related to city improvement measures and drug reviews, demonstrating its effectiveness in generating representative slates from unstructured user opinions.

Thursday, September 18, 14:00–14:20

[53] Strategizing under Rule and Vote Uncertainty: An Experiment

Authors. **Antoine Prévotat**, Zoi Terzopoulou, Adam Zylbersztein

Abstract. Individuals participate daily in collective decision-making. Sometimes they face a clear and explicit voting rule (such as during political elections), but often they are asked to express their preferences in situations where the voting rule is either unspecified or unknown (such as choosing a time slot for a work meeting or voting in an online contest). Understanding how personal preferences are translated into voting behavior is far from trivial. One key factor is the information individuals have about the votes of their peers, while another is the information they have about the applied voting rule. Although the first aspect has received substantial attention in the literature, the second remains largely unexplored. Our lab experiment introduces a setting in which participants do not know which voting rule will be used. We find that participants' beliefs about others' votes are systematically myopic, shaped by previous outcomes rather than rational expectations, though they become more accurate and confident over time as votes stabilize. Participants' votes are very rarely sincere, and remarkably, rule uncertainty does not discourage strategic behavior; rather, it appears to foster it in our experimental setting.

Thursday, September 18, 14:20–14:40

[78] Guide to Numerical Experiments on Elections in Computational Social Choice

Authors. Niclas Boehmer, Piotr Faliszewski, Łukasz Janeczko, Andrzej Kaczmarczyk, Grzegorz Lisowski, Grzegorz Pierczyński, Simon Rey, Dariusz Stoliczki, **Stanisław Szufa**, Tomasz Wąs

Abstract. We analyze how numerical experiments regarding elections were conducted within the computational social choice literature (focusing on papers published in the IJ-CAI, AAAI, and AAMAS conferences). We analyze the sizes of the studied elections and the methods used for generating preference data, thereby making previously hidden standards and practices explicit. In particular, we survey a number of statistical cultures for generating elections and their commonly used parameters.

Thursday, September 18, 14:40–15:00

[86] Learning Real-Life Approval Elections

Authors. **Piotr Faliszewski**, Łukasz Janeczko, Andrzej Kaczmarczyk, Marcin Kurdziel, Grzegorz Pierczyński, Stanisław Szufa

Abstract. We study the independent approval model (IAM) for approval elections, where each candidate has its own approval probability and is approved independently of the other ones. This model generalizes, e.g., the impartial culture, the Hamming noise model, and the resampling model. We propose algorithms for learning IAMs and their mixtures from data, using either maximum likelihood estimation or Bayesian learning. We then apply these algorithms to a large set of elections from the Pabulib database. In particular, we find that single-component models are rarely sufficient to capture the complexity of real-life data, whereas their mixtures perform well.

Thursday, September 18, 15:30–15:50

[17] Committee Monotonicity and Proportional Representation for Ranked Preferences

Authors. Haris Aziz, **Patrick Lederer**, Dominik Peters, Jannik Peters, Angus Ritossa

Abstract. We study committee voting rules under ranked preferences, which map the voters' preference relations to a subset of the candidates of predefined size. In this setting, the compatibility between proportional representation and committee monotonicity is a fundamental open problem that has been mentioned in several works. We address this research question by designing a new committee voting rule called the Solid Coalition Refinement (SCR) rule that simultaneously satisfies committee monotonicity and Dummett's PSC as well as one of its variants called inclusion PSC. This is the first rule known to satisfy both of these properties. Moreover, we show that this is effectively the best that we can hope for as other fairness notions adapted from approval voting are incompatible with committee monotonicity. We also prove that, for truncated preferences, the SCR rule still satisfies PSC and a property called independence of losing voter blocs, thereby refuting a conjecture of Graham-Squire et al. (2024). Finally, we discuss the consequences of our results in the context of rank aggregation.

Thursday, September 18, 15:50–16:10

[89] Proportionality in Practice: Quantifying Proportionality in Ordinal Elections

Authors. **Tuva Bardal**, Markus Brill, David McCune, Jannik Peters

Abstract. In real-world elections with ranked preferences, the Single Transferable Vote (STV) is the most widely used proportional voting method. STV is considered proportional because it satisfies an axiom known as proportionality for solid coalitions (PSC), requiring that large enough “solid coalitions” of voters are adequately represented. Using real-world data from local Scottish elections, we observe that solid coalitions of the required size rarely occur in practice. This observation challenges the importance of proportionality axioms and raises the question of how the proportionality of voting methods can be assessed beyond their axiomatic performance. We address these concerns by developing quantitative measures of proportionality. We apply these measures to evaluate the proportionality of voting rules on real-world election data. Besides STV, we consider the Single Non-Transferable Vote, the Expanding Approvals Rule, and Sequential Ranked-Choice Voting. We also study the effects of ballot truncation by artificially completing truncated ballots and comparing the proportionality of outcomes under complete and truncated ballots.

Thursday, September 18, 16:30–16:50

[27] The Power of Matching for Online Fractional Hedonic Games

Authors. **Martin Bullinger**, René Romen, Alexander Schlenga

Abstract. We study coalition formation in the framework of fractional hedonic games (FHGs). The objective is to maximize social welfare in an online model where agents arrive one by one and must be assigned to coalitions immediately and irrevocably. For general online FHGs, it is known that computing maximal matchings achieves the optimal competitive ratio, which is, however, unbounded for unbounded agent valuations.

We achieve a constant competitive ratio in two related settings while carving out further connections to matchings. If algorithms can dissolve coalitions, then the optimal competitive ratio of $1/(6 + 4\sqrt{2})$ is achieved by a matching-based algorithm. Moreover, we perform a tight analysis for the online matching setting under random arrival with an unknown number of agents. This entails a randomized $1/6$ -competitive algorithm for FHGs, while no algorithm can be better than $1/3$ -competitive.

Thursday, September 18, 16:50–17:10

[74] Quantile Agent Utility and Implications to Randomized Social Choice

Authors. **Ioannis Caragiannis**, Sanjukta Roy

Abstract. We initiate a novel direction in randomized social choice by proposing a new definition of agent utility for randomized outcomes. Each agent has a preference over all outcomes and a *quantile* parameter. Given a *lottery* over the outcomes, an agent gets utility from a particular *representative*, defined as the least preferred outcome that can be realized so that the probability that any worse-ranked outcome can be realized is at most the agent's quantile value. In contrast to other utility models that have been considered in randomized social choice (e.g., stochastic dominance, expected utility), our *quantile agent utility* compares two lotteries for an agent by just comparing the representatives, as is done for deterministic outcomes. We revisit questions in randomized social choice using the new utility definition. We study the compatibility of efficiency and strategyproofness for randomized voting rules, efficiency and fairness for randomized one-sided matching mechanisms, and efficiency, stability, and strategyproofness for lotteries over two-sided matchings. In contrast to well-known impossibilities in randomized social choice, we show that satisfying the above properties simultaneously can be possible.

Thursday, September 18, 17:10–17:30

[110] Byzantine Game Theory: Sun Tzu's Boxes

Authors. **Andrei Constantinescu**, Roger Wattenhofer

Abstract. We introduce the Byzantine Selection Problem, living at the intersection of game theory and fault-tolerant distributed computing. Here, an event organizer is presented with a group of n agents, and wants to select $l < n$ of them to form a team. For these purposes, each agent i self-reports a positive skill value v_i , and a team's value is the sum of its members' skill values. Ideally, the value of the team should be as large as possible, which can be easily achieved by selecting agents with the highest l skill values. However, an unknown subset of at most $t < n$ agents are byzantine and hence not to be trusted, rendering their true skill values as 0. In the spirit of the distributed computing literature, the identity of the byzantine agents is not random but instead chosen by an adversary aiming to minimize the value of the chosen team. Can we still select a team with good guarantees in this adversarial setting? As it turns out, deterministically, it remains optimal to select agents with the highest l values. Yet, if $t \geq l$, the adversary can choose to make all selected agents byzantine, leading to a team of value zero. To provide meaningful guarantees, one hence needs to allow for randomization, in which case the expected value of the selected team needs to be maximized, assuming again that the adversary plays to minimize it. For this case, we provide linear-time randomized algorithms that maximize the expected value of the selected team.

Friday, September 19, 10:15–10:35

[23] Condorcet-Consistent Choice Among Three Candidates

Authors. Felix Brandt, **Chris Dong**, Dominik Peters

Abstract. A voting rule is a Condorcet extension if it returns a candidate that beats every other candidate in pairwise majority comparisons whenever one exists. Condorcet extensions have faced criticism due to their susceptibility to variable-electorate paradoxes, especially the reinforcement paradox (Young and Levenglick, 1978) and the no-show paradox (Moulin, 1988). In this paper, we investigate the susceptibility of Condorcet extensions to these paradoxes for the case of exactly three candidates. For the reinforcement paradox, we establish that it must occur for every Condorcet extension when there are at least eight voters and demonstrate that certain refinements of maximin, a voting rule originally proposed by Condorcet (1785), are immune to this paradox when there are at most seven voters. For the no-show paradox, we prove that the only homogeneous Condorcet extensions immune to it are refinements of maximin. We also provide axiomatic characterizations of maximin and two of its refinements, Nanson’s rule and leximin, highlighting their suitability for three-candidate elections.

Friday, September 19, 10:35–11:05

[40] On Condorcet’s Jury Theorem with Abstention

Authors. **Reshef Meir**, Ganesh Ghalme

Abstract. The well-known Condorcet Jury Theorem states that, under majority rule, the better of two alternatives is chosen with probability approaching one as the population grows. We study an asymmetric setting where voters face varying participation costs and share a (possibly heuristic) belief about their pivotality.

In a costly voting setup where voters abstain if their participation cost is greater than their pivotality estimate, we identify a single property of the heuristic belief—weakly vanishing pivotality—that gives rise to multiple stable equilibria in which elections are nearly tied. In contrast, strongly vanishing pivotality (as in the standard Calculus of Voting model) yields a unique, trivial equilibrium where only zero-cost voters participate as the population grows. We then characterize when nontrivial equilibria satisfy a version of the Jury Theorem: below a sharp threshold, the majority-preferred candidate wins with probability approaching one; above it, both candidates either win with equal probability.

Friday, September 19, 11:05–11:15

[65] An Axiomatic Characterization of the Minimax Voting Method

Authors. **Yifeng Ding**, Xiaochen Yu

Abstract. Under the Minimax voting method, a candidate wins if her maximum margin of loss in pairwise majority comparisons is the smallest. It is one of the few Condorcet consistent voting methods guaranteeing that truthful voting does not make a voter’s unique favorite candidate lose. We show that two more axioms, one requiring that there can be only one winner if no two pairwise majority margins are equal and the other called Ordinal Margin Invariance stating that only the relative ordering of margins matters, suffice to precisely pin down Minimax on profiles without equal margins. We also axiomatize Minimax over all profiles. This is done by adding an axiom of continuity and either adding a further axiom of monotonicity or replacing the unique winner axiom above with Weak Positive Responsiveness, a weakening of Positive Responsiveness that features in May’s theorem for two-candidate voting.

Friday, September 19, 11:25–11:45

[91] From Independence of Clones to Composition Consistency: A Hierarchy of Barriers to Strategic Nomination

Authors. **Ratip Emin Berker**, Sílvia Casacuberta, Isaac Robinson, Christopher Ong, Vincent Conitzer, Edith Elkind

Abstract. We study two axioms for social choice functions that capture the impact of similar candidates: independence of clones (IoC) and composition consistency (CC). We clarify the relationship between these axioms by observing that CC is strictly more demanding than IoC, and investigate whether common voting rules that are known to be independent of clones (such as STV, Ranked Pairs, Schulze, and Split Cycle) are composition-consistent. While for most of these rules the answer is negative, we identify a variant of Ranked Pairs that satisfies CC. Further, we show how to efficiently modify any (neutral) social choice function so that it satisfies CC, while maintaining its other desirable properties. Our transformation relies on the hierarchical representation of clone structures via PQ-trees. We extend our analysis to social preference functions. Finally, we interpret IoC and CC as measures of robustness against strategic manipulation by candidates, with IoC corresponding to strategy-proofness and CC corresponding to obvious strategy-proofness.

Friday, September 19, 11:45–12:05

[30] Discrete Budget Aggregation: Truthfulness and Proportionality

Authors. Ulrike Schmidt-Kraepelin, Warut Suksompong, **Markus Utke**

Abstract. We study a budget aggregation setting where voters express their preferred allocation of a fixed budget over a set of alternatives, and a mechanism aggregates these preferences into a single output allocation. Motivated by scenarios in which the budget is not perfectly divisible, we depart from the prevailing literature by restricting the mechanism to output allocations that assign integral amounts. This seemingly minor deviation has significant implications for the existence of truthful mechanisms. Specifically, when voters can propose fractional allocations, we demonstrate that the Gibbard-Satterthwaite theorem can be extended to our setting. In contrast, when voters are restricted to integral ballots, we identify a class of truthful mechanisms by adapting moving-phantom mechanisms to our context. Moreover, we show that while a weak form of proportionality can be achieved alongside truthfulness, (stronger) proportionality notions derived from approval-based committee voting are incompatible with truthfulness.

Friday, September 19, 14:00–14:20

[39] Reconfiguring Proportional Committees

Authors. Chris Dong, **Fabian Frank**, Jannik Peters, Warut Suksompong

Abstract. An important desideratum in approval-based multiwinner voting is proportionality. We study the problem of reconfiguring proportional committees: given two proportional committees, is there a transition path that consists only of proportional committees, where each transition involves replacing one candidate with another candidate? We show that the set of committees satisfying the proportionality axiom of justified representation (JR) is not always connected, and it is PSPACE-complete to decide whether two such committees are connected. On the other hand, we prove that any two JR committees can be connected by committees satisfying a 2-approximation of JR. We also obtain similar results for the stronger axiom of extended justified representation (EJR). In addition, we demonstrate that the committees produced by several well-known voting rules are either connected or at least not isolated.

Friday, September 19, 14:20–14:40

[93] Robust Committee Voting, or The Other Side of Representation

Authors. **Gregory Kehne**, Ulrike Schmidt-Kraepelin, Krzysztof Sornat

Abstract. We study approval-based committee voting from a novel perspective. While extant work largely centers proportional representation of the voters, we shift our focus to the candidates while preserving proportionality. Intuitively, candidates supported by similar voter groups should receive comparable representation. Since deterministic voting rules cannot achieve this ideal, we develop randomized voting rules that satisfy ex-ante neutrality, monotonicity, and continuity, while maintaining strong ex-post proportionality guarantees.

Ensuring continuity of the candidate selection probabilities proves the most demanding of our ex-ante desiderata. We achieve it via voting rules that are algorithmically stable, a stronger guarantee which captures the continuity of the committee distribution under small changes. First, we introduce Softmax-GJCR, a randomized variant of the Greedy Justified Candidate Rule (GJCR) [Brill and Peters, 2023], which carefully leverages the slack in GJCR to satisfy our ex-ante properties. This polynomial-time algorithm satisfies EJR+ ex post, assures ex-ante monotonicity and neutrality, and provides $O(k^3/n)$ -stability (suppressing log factors). Building on our techniques for Softmax-GJCR, we also show that stronger stability guarantees can be attained by (i) allowing exponential runtime, (ii) relaxing EJR+ to an approximate α -EJR+, and (iii) relaxing EJR+ to JR.

We also demonstrate the utility of stable voting rules to other settings. In online dynamic committee voting, we show stable rules imply dynamic rules with low expected recourse, and illustrate this reduction for Softmax-GJCR. Our rules also satisfy a stronger form of stability that coincides with differential privacy, suggesting their applicability in privacy-sensitive domains.

Friday, September 19, 14:40–15:00

[101] A Generalised Theory of Proportionality in Collective Decision Making

Authors. Tomáš Masařík, Grzegorz Pierczyński, **Piotr Skowron**

Abstract. We consider a voting model, where a number of candidates need to be selected subject to certain feasibility constraints. The model generalizes 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 need 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 opinions 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 generalize 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.

Friday, September 19, 15:30–15:50

[50] Method of Equal Shares with Bounded Overspending

Authors. Georgios Papasotiropoulos, **Seyedeh Zeinab Pishbin**, Oskar Skibski, Piotr Skowron, Tomasz Wąs

Abstract. Pure proportional voting rules can sometimes lead to highly suboptimal outcomes. We introduce the Method of Equal Shares with Bounded Overspending (BOS Equal Shares), a robust variant of the Method of Equal Shares that balances proportionality and efficiency. BOS Equal Shares addresses inefficiencies implied by strict proportionality, yet still provides fairness guarantees, similar to the original Equal Shares. Our extensive empirical analysis shows excellent performance of BOS Equal Shares across several metrics. In the course of the analysis, we also study a fractional variant of the Method of Equal Shares.

Friday, September 19, 15:50–16:10

[96] Streamlining Equal Shares

Authors. Sonja Kraiczy, Isaac Robinson, **Edith Elkind**

Abstract. Participatory budgeting (PB) is a form of citizen participation that allows citizens to decide how public funds are spent. Through an election, citizens express their preferences on various projects (spending proposals). A voting mechanism then determines which projects will be approved. The Method of Equal Shares (MES) is the state of the art algorithm for a proportional, voting based approach to participatory budgeting and has been implemented in cities across Poland and Switzerland. A significant drawback of MES is that it is not exhaustive meaning that it often leaves a portion of the budget unspent that could be used to fund additional projects. To address this, in practice the algorithm is combined with a completion heuristic - most often the ADD-ONE heuristic which artificially increases the budget until a heuristically chosen threshold. This heuristic is computationally inefficient and will become computationally impractical if PB is employed on a larger scale. We propose the more efficient ADD-OPT heuristic for Exact Equal Shares (EES), a variation of MES that is known to retain many of its desirable properties. We solve the problem of identifying the next budget for which the outcome for EES changes in $O(mn)$ time for cardinal utilities and $O(m^2n)$ time for uniform utilities, where m is the number of projects and n is the number of voters. Our solution to this problem inspires the efficient ADD-OPT heuristic which bypasses the need to search through each intermediary budget. We perform comprehensive experiments on real-world PB instances from Pabulib and show that completed EES outcomes usually match the proportion of budget spent by completed MES outcomes. Furthermore, the ADD-OPT heuristic matches the proportion of budget spent by ADD-ONE for EES.

Friday, September 19, 16:10–16:30

[87] Mixed Voting Rules for Participatory Budgeting

Authors. **Anton Baychkov**, Markus Brill, Markus Utke

Abstract. Designing and analyzing voting rules for Participatory Budgeting (PB) elections is an active research area in computational social choice. Many PB voting rules aim to optimize a specific objective. For instance, the ubiquitous Greedy rule attempts to maximize utilitarian welfare, while the Method of Equal Shares (MES) aims to achieve proportional representation of voter preferences. However, it is often desirable to achieve good outcomes on multiple objectives rather than a close-to-perfect outcome for one. Inspired by mixed-member systems that are often used for parliamentary elections, we introduce mixed voting rules for PB. These are composed of a sequence of two or more rules that can each spend some fraction of the overall budget in order to add projects to the set selected by earlier rules. We develop a theoretical framework for formulating and analyzing mixed PB voting rules, and explore how existing rules can be adapted to this framework. We particularly focus on MES and its potential to address imbalances in representation created by earlier rules. We propose different ways to adjust MES voter budgets based on how satisfied voters are with previously chosen projects, and examine how well the resulting rules approximate well-known proportionality axioms such as EJR+. We complement our theoretical results with an empirical analysis of real-world PB elections, investigating how mixed rules perform compared to their constituent rules.

Friday, September 19, 16:50–17:10

[84] Approximate Core of Participatory Budgeting via Lindahl Equilibrium

Authors. Thanh Nguyen, Haoyu Song

Abstract. Participatory budgeting is a democratic process in which citizens decide how to allocate public funds among proposed projects. In practice, participants typically submit ordinal preferences—such as rankings or approvals—rather than numerical utilities. A central fairness concept in this setting is the proportional core, which ensures that no group of agents can reallocate their proportional share of the budget to strictly improve their outcomes. However, the core may be empty under general ordinal preferences, motivating the study of approximate core solutions that relax this requirement while preserving its fairness spirit.

We improve the best-known approximation bound from 32 to 6.24 under general monotone ordinal preferences. In structured domains such as committee selection with approval ballots or ranking preferences, we further refine our approach to achieve stronger guarantees of 5.15 and 5.10, respectively. Our main innovation is the introduction of the Lindahl Equilibrium with Ordinal preference (LEO), a novel continuous relaxation inspired by the classical Lindahl equilibrium. This framework bridges traditional economic insights with the discrete, ordinal setting of social choice and participatory budgeting, and has potential applicability in a broad range of settings.

Friday, September 19, 17:10–17:30

[88] Participatory Budgeting Project Strength via Candidate Control

Authors. Piotr Faliszewski, Łukasz Janeczko, Dušan Knop, Jan Pokorný, Šimon Schierreich, Mateusz Ślusznia, Krzysztof Sornat

Abstract. We study the complexity of candidate control in participatory budgeting elections. The goal of constructive candidate control is to ensure that a given candidate wins by either adding or deleting candidates from the election (in the destructive setting, the goal is to prevent a given candidate from winning). We show that such control problems are NP-hard to solve for many participatory budgeting voting rules, including Phragmén and Equal-Shares, but there are natural cases with polynomial-time algorithms. We also argue that control by deleting candidates is a useful tool for assessing the performance (or, strength) of initially losing projects, and we support this view with experiments on real-life PB instances.

Poster Presentation Paper Abstracts

[3] Delivering Fairly in the Gig Economy

Authors. Hadi Hosseini, **Šimon Schierreich**

Abstract. Distributing services, goods, and tasks in the gig economy heavily relies upon on-demand workers (aka agents), leading to new challenges varying from logistics optimization to the ethical treatment of gig workers. We focus on fair and efficient distribution of delivery tasks—placed on the vertices of a graph—among a fixed set of agents. We consider the fairness notion of minimax share (MMS), which aims to minimize the maximum (submodular) cost among agents and is particularly appealing in applications without monetary transfers. We propose a novel efficiency notion—namely non-wastefulness—that is desirable in a wide range of scenarios and, more importantly, does not suffer from computational barriers. Specifically, given a distribution of tasks, we can, in polynomial time, i) verify whether the distribution is non-wasteful and ii) turn it into an equivalent non-wasteful distribution. Moreover, we investigate several fixed-parameter tractable and polynomial-time algorithms and paint a complete picture of the (parameterized) complexity of finding fair and efficient distributions of tasks with respect to both the structure of the topology and natural input restrictions. Finally, we highlight how our findings shed light on computational aspects of other well-studied fairness notions, such as envy-freeness and its relaxations.

[6] On voting rules satisfying false-name-proofness and participation

Authors. Agustín G. Bonifacio, **Federico Fioravanti**

Abstract. We consider voting rules in settings where voters' identities are difficult to verify. Voters can manipulate the process by casting multiple votes under different identities or abstaining from voting. Immunities to such manipulations are called *false-name-proofness* and *participation*, respectively. For the universal domain of (strict) preferences, these properties together imply *anonymity* and are incompatible with *neutrality*. For the domain of preferences defined over all subsets of a given set of objects, both of these properties cannot be met by *onto* and *object neutral* rules that also satisfy the *tops-only* criterion. However, when preferences over subsets of objects are restricted to be separable, all these properties can be satisfied. Furthermore, the domain of separable preferences is maximal for these properties.

[7] Why Instant-Runoff Voting Is So Resilient to Coalitional Manipulation: Phase Transitions in the Perturbed Culture

Authors. **François Durand**

Abstract. Previous studies have shown that Instant-Runoff Voting (IRV) is highly resistant to coalitional manipulation (CM), though the theoretical reasons for this remain unclear. To address this gap, we analyze the susceptibility to CM of three major voting rules—Plurality, Two-Round System, and IRV—within the Perturbed Culture model. Our findings reveal that each rule undergoes a phase transition at a critical value θ_c of the concentration of preferences: the probability of CM for large electorates converges exponentially fast to 1 below θ_c and to 0 above θ_c . We introduce the Super Condorcet Winner (SCW), showing that its presence is a key factor of IRV's resistance to coalitional manipulation, both theoretically and empirically. Notably, we use this notion to prove that for IRV, $\theta_c = 0$, making it resistant to CM with even minimal preference concentration.

[8] Diverse Committees with Incomplete or Inaccurate Approval Ballots

Authors. Feline Lindeboom, Martijn Brehm, **Daide Grossi**, Pradeep Murukannaiah

Abstract. We study diversity in approval-based committee elections with incomplete or inaccurate information. We define diversity according to the Maximum Coverage problem, which is known to be NP-complete, with a best attainable polynomial time approximation ratio of $1 - 1/e$. In the incomplete information setting, voters vote only on a small portion of the candidates, and we prove that getting arbitrarily close to the optimal approximation ratio w.h.p. requires $\Omega(m^2)$ non-adaptive queries, where m is the number of candidates. This motivates studying adaptive querying algorithms, that can adapt their querying strategy to information obtained from previous query outcomes. In that setting, we lower this bound to only $\Omega(m)$ queries. We propose a greedy algorithm to match this lower bound up to log-factors. We prove the same $\tilde{O}(m)$ bound for the generalized problem of Max Cover over a matroid constraint, using a local search algorithm. Specifying a matroid of valid committees lets us implement extra structural requirements on the committee, like quota. In the inaccurate information setting, voters' responses are corrupted with a small probability. We prove $\tilde{O}(nm)$ queries are required to attain a $(1 - 1/e)$ -approximation with high probability, where n is the number of voters. While the proven bounds show that all our algorithms are viable asymptotically, they also show that some of them would still require large numbers of queries in instances of practical relevance. Using real data from Polis as well as synthetic data, we observe that our algorithms perform well also on smaller instances, both with incomplete and inaccurate information.

[10] Multiwinner Voting with Interval Preferences under Incomplete Information

Authors. **Drew Springham**, Edith Elkind, Maria Polukarov, Bart de Keijzer

Abstract. In multiwinner approval elections with many candidates, voters may struggle to determine their preferences over the entire slate of candidates. It is therefore of interest to explore which (if any) fairness guarantees can be provided under reduced communication. In this paper, we consider voters with one-dimensional preferences: voters and candidates are associated with points on the real line, and each voter's approval set forms an interval of the real line. We put forward a probabilistic preference model, where the voter set consists of sigma different groups; each group is associated with a distribution over an interval of the real line, so that each voter draws the endpoints of her approval interval from the distribution associated with her group. We present an algorithm for computing committees that provide *Proportional Justified Representation+* (PJR+), which proceeds by querying voters' preferences, and show that, in expectation, it suffices to make $O(\log(\sigma \cdot k))$ queries per voter, where k is the desired committee size.

[13] Sociotropic Behavior in Voting

Authors. Jan Maly, **Oliviero Nardi**, Zoi Terzopoulou

Abstract. This paper proposes a formal model of sociotropic voting, where voters' ballots are based on the preferences of their peers in a group rather than their individual preferences. This behavior, prevalent in social settings and elections, can paradoxically lead to inferior outcomes for the group. We assume that voters have perfect knowledge of others' preferences and use an internal aggregation rule to construct their ballots, which may be distinct from the external voting rule. We find that sociotropic voting can alter outcomes in ordinal single-winner and in approval-based multi-winner elections, but not in single-winner elections under the classical approval voting rule. Our theoretical and experimental results show that sociotropic voting can be harmful when external voting rules are designed for consensus and fairness, such as Condorcet consistent rules. However, it can enhance equity in settings using simple voting rules like plurality.

[14] An Incentive-Compatible Utilitarian Voting Procedure for Permanent Citizens' Assemblies

Authors. **Rajarshi Ghosh**, Marcus Pivato

Abstract. We consider a committee of voters randomly drawn from a larger population facing an infinite sequence of voting decisions, akin to a citizen jury. We propose a new voting mechanism for such juries where each voter has a privately known von Neumann-Morgenstern (vNM) utility function over social alternatives in each decision, and is asked to report a real-valued 'valuation' for each alternative of a decision. We further impose a probability of being removed from the committee for the next decision dependent on the report of a voter. If a voter is removed, then they are replaced by some non-committee member from the larger population. We show that when the voters' discount factor is large enough, imposing a probability equal to a scalar multiple of the Vickrey-Clarke-Groves tax leads to truthful revelation by the voters and consequently utilitarian efficient outcomes at a Bayesian Nash equilibrium.

[15] Budget Division Via Continuous Thiele's Rules

Authors. Jonathan Wagner, Reshef Meir

Abstract. We introduce the class of Continuous Thiele's Rules that generalize the familiar Thiele's rules of multi-winner voting to divisible Participatory Budgeting. Each rule in that class maximizes $\sum_i f(\pi_i)$ where π_i is agent i 's satisfaction and f could be any twice differentiable, increasing and concave real function. Based on a single parameter we call the 'Inequality Aversion' of f (elsewhere known as "Relative Risk Aversion"), we derive bounds on the egalitarian and utilitarian welfare loss, and the approximation of individual and group Fair Share. This leads to a quantifiable, continuous presentation of their inevitable trade-offs. In particular, we show that many of these rules satisfy Individual Fair Share, and that the Nash Product Rule satisfies Average Fair Share in our setting.

[20] Metric Distortion in Peer Selection*Authors.* Javier Cembrano, **Golnoosh Shahkarami**

Abstract. In the metric distortion problem, a set of voters and candidates lie in a common metric space, and a committee of k candidates is to be elected. The goal is to select a committee with a small social cost, defined as an increasing function of the distances between voters and selected candidates, but a voting rule only has access to voters' ordinal preferences. The distortion of a rule is then defined as the worst-case ratio between the social cost of the selected set and the optimal set, over all possible preferences and consistent distances.

We initiate the study of metric distortion when voters and candidates coincide, which arises naturally in peer selection, and provide tight results for various social cost functions on the line metric. We consider both utilitarian and egalitarian social cost, given by the sum and maximum of the individual social costs, respectively. For utilitarian social cost, we show that the simple voting rule that selects the k middle agents achieves a distortion that varies between 1 and 2 as k varies between 1 and n when the cost of an individual is the sum of their distances to all selected candidates (additive aggregation). When the cost of an individual is their distance to their q th closest candidate (q -cost), we provide positive results for $q = k = 2$ but mostly show that negative results for general elections carry over to our restricted setting: No constant distortion is possible when $q \leq k/2$ and no distortion better than $3/2$ is possible for $q \geq k/2 + 1$. For egalitarian social cost, a rule that selects extreme agents achieves the best-possible distortion of 2 for additive cost and q -cost with $q > k/3$, whereas no bounded distortion is possible for $q \leq k/3$. Our results suggest that having a common set of voters and candidates allows for better constants compared to the general setting, but cases in which no constant is possible in general remain hard under this restriction.

[21] Voter Priming Campaigns: Strategies, Equilibria, and Algorithms*Authors.* Jonathan Shaki, **Yonatan Aumann**, Sarit Kraus

Abstract. Issue salience is a major determinant in voters' decisions. Candidates and political parties campaign to shift salience to their advantage - a process termed priming. We study the dynamics, strategies and equilibria of campaign spending for voter priming in multi-issue multi-party settings. We consider both parliamentary elections, where parties aim to maximize their share of votes, and various settings for presidential elections, where the winner takes all. For parliamentary elections, we show that pure equilibrium spending always exists and can be computed in time linear in the number of voters. For two parties and all settings, a spending equilibrium exists such that each party invests only in a single issue, and an equilibrium can be computed in time that is polynomial in the number of issues and linear in the number of voters. We also show that in most presidential settings no equilibrium exists. Additional properties of optimal campaign strategies are also studied.

[29] Solving Four Open Problems about Core Stability in Altruistic Hedonic Games*Authors.* Jörg Rothe and Ildikó Schlotter

Abstract. Hedonic games—at the interface of cooperative game theory and computational social choice—are coalition formation games in which the players have preferences over the coalitions they can join. Kerkmann et al. [15] introduced altruistic hedonic games where the players’ utilities depend not only on their own but also on their friends’ valuations of coalitions. The complexity of the verification problem for core stability has remained open in four variants of altruistic hedonic games: namely, for the variants with average- and minimum-based “equal-treatment” and “altruistic-treatment” preferences. We solve these four open questions by proving the corresponding problems coNP-complete; our reductions rely on rather intricate gadgets in the related networks of friends.

[34] Probability of a Condorcet Winner for Large Electorates: An Analytic Combinatorics Approach*Authors.* François Durand, Emma Caizergues, Élie de Panafieu, Marc Noy, Vlado Ravelomanana

Abstract. We study the probability that a given candidate is an $\vec{\alpha}$ -winner, i.e., a candidate preferred to each other candidate j by a fraction α_j of the voters. This extends the classical notion of Condorcet winner, which corresponds to the case $\vec{\alpha} = (1/2, \dots, 1/2)$. Our analysis is conducted under the general assumption that voters have independent preferences, illustrated through applications to well-known models such as Impartial Culture and the Mallows model. While previous works use probabilistic arguments to derive the limiting probability as the number of voters tends to infinity, we employ techniques from the field of analytic combinatorics to compute convergence rates and provide a method for obtaining higher-order terms in the asymptotic expansion. In particular, we establish that the probability of a given candidate being the Condorcet winner in Impartial Culture is $a_0 + a_{1,n}n^{-1/2} + O(n^{-1})$, where we explicitly provide the values of the constant a_0 and the coefficient $a_{1,n}$, which depends solely on the parity of the number of voters n . Along the way, we derive technical results in multivariate analytic combinatorics that may be of independent interest.

[35] Scaling Fair and Efficient Course Allocation: From Theory to Practice

Authors. **Paula Navarrete Diaz**, Cyrus Cousins, George Bissias, Yair Zick

Abstract. Universities regularly face the challenging task of assigning classes to thousands of students while considering their preferences, along with course schedules and capacities. Ensuring the effectiveness and fairness of course allocation mechanisms is crucial to guaranteeing student satisfaction and optimizing resource utilization. We approach this problem from an economic perspective, using formal justice criteria to evaluate different algorithmic frameworks. To evaluate our frameworks, we conduct a large scale survey of university students at University of Massachusetts Amherst, collecting over 1000 student preferences. This is, to our knowledge, the largest publicly available dataset of student preferences. We develop software for generating synthetic student preferences over courses, and implement four allocation algorithms: the serial dictatorship algorithm used by University of Massachusetts Amherst; Round Robin; an Integer Linear Program; and the Yankee Swap algorithm. We propose improvements to the Yankee Swap framework to handle scenarios with item multiplicities. Through experimentation with the Fall 2024 Computer Science course schedule at University of Massachusetts Amherst, we evaluate each algorithm's performance relative to standard justice criteria, providing insights into fair course allocation in large university settings.

[45] Frankenmandering: Repeated Social Graph Gerrymandering

Authors. Sahil Agarwal, Svetlana Obraztsova, Zinovi Rabinovich, **Alan Tsang**

Abstract. Opinion influence by means of manipulating a social graph structure, as well as election manipulation by means of artificial grouping (gerrymandering), are well-known fields of application and investigation in computational social choice and machine learning. However, while studied separately, in real life, both of these manipulations (social graph alteration and gerrymandering) occur simultaneously. In this paper, we offer the first model of such a simultaneous process, which takes the form of repeated gerrymandering with an underlying social graph for opinion diffusion. We term this process "Frankenmandering", and provide the first steps in its analysis: examples of principal feasibility and the impact of the underlying social graph.

[48] Stable Matching under Matroid Rank Valuations

Authors. Alon Eden, **Vignesh Viswanathan**, Yair Zick

Abstract. We study a two-sided matching model where one side of the market (hospitals) has combinatorial preferences over the other side (doctors). Specifically, we consider the setting where hospitals have matroid rank valuations over the doctors, and doctors have either ordinal or cardinal unit-demand valuations over the hospitals. While this setting has been extensively studied in the context of one-sided markets, it remains unexplored in the context of two-sided markets.

When doctors have ordinal preferences over hospitals, we present simple sequential allocation algorithms that guarantee stability, strategyproofness for doctors, and approximate strategyproofness for hospitals. When doctors have cardinal utilities over hospitals, we present an algorithm that finds a stable allocation maximizing doctor welfare; subject to that, we show how one can maximize either the hospital utilitarian or hospital Nash welfare. Moreover, we show that it is NP-hard to compute stable allocations that approximately maximize hospital Nash welfare.

[49] Cycles in Liquid Democracy: A Game-Theoretic Justification

Authors. Markus Brill, Rachael Colley, Anne-Marie George, **Grzegorz Lisowski**, Georgios Papasotiropoulos, Ulrike Schmidt-Kraepelin

Abstract. Liquid democracy is a voting scheme in which individuals either vote directly or delegate their vote to others. A common critique is that delegation cycles can occur, seemingly resulting in unused voting power. Yet, practitioners argue that delegation cycles are not only unproblematic but are even intentionally formed by participants. This divergent view stems from differing interpretations of delegations: in practice, delegations serve as backup options that can be overridden at any time by direct voting, whereas the literature often treats voting and delegating as mutually exclusive. Bringing theory closer to reality, we introduce a probabilistic model that captures strategic behavior under uncertainty. Within this model, we study the existence and structure of Nash equilibria, revealing that delegation cycles naturally emerge. We further examine the quality of equilibria via a Price-of-Anarchy approach. To complement these findings, we perform computational experiments using best-response dynamics.

[52] Pareto-efficiency of ordinal multiwinner voting rules

Authors. Jean Lainé, **Jérôme Lang**, Ipek Özkal-Sanver, Remzi Sanver

Abstract. We investigate the Pareto efficiency of ordinal multiwinner voting rules. Defining the Pareto-optimality of a committee requires relating the voters' rankings over individual candidates to their preferences over committees. We consider two well-known extension principles that extend rankings over candidates to preferences over committees: responsive and lexicographic. As the responsive extension outputs partial orders, we consider two associated Pareto-optimality notions: a committee is possibly (respectively, necessary) Pareto-optimal if it is Pareto-optimal for some (respectively, every) completion of these partial orders. We define several Pareto-efficiency notions for multiwinner rules, depending on whether some (respectively, all) committees in the output are Pareto-optimal for one of the latter notions. We review what we believe to be a complete list of ordinal multiwinner rules that have been studied in the literature, and identify which Pareto-efficiency notions they satisfy. We find that, somewhat surprisingly, these rules show a huge diversity: some satisfy the strongest notion, some do not even satisfy the weakest one, with many other rules at various intermediate levels.

[54] Selecting Interlacing Committees

Authors. Chris Dong, **Martin Bullinger**, Tomasz Wąs, Larry Birnbaum, and Edith Elkind

Abstract. Polarization is a major concern for a well-functioning society. Often, mass polarization of a society is driven by polarizing political representation, even when the latter is easily preventable. The existing computational social choice methods for the task of committee selection are not designed to address this issue. We enrich the standard approach to committee selection by defining two quantitative measures that evaluate how well a given committee interconnects the voters. Maximizing these measures aims at avoiding polarizing committees. While the corresponding maximization problems are NP-complete in general, we obtain efficient algorithms for profiles in the voter-candidate interval domain. Moreover, we analyze the compatibility of our goals with other representation objectives, such as excellence, diversity, and proportionality. We identify trade-offs between approximation guarantees, and describe algorithms that achieve simultaneous constant-factor approximations.

[55] Temporal Fair Division of Indivisible Items

Authors. Edith Elkind, Alexander Lam, Mohamad Latifian, Tzeh Yuan Neoh, **Nicholas Teh**

Abstract. We study a fair division model where indivisible items arrive sequentially, and must be allocated immediately and irrevocably. Previous work on online fair division has shown impossibility results in achieving approximate envy-freeness under these constraints. In contrast, we consider an informed setting where the algorithm has complete knowledge of future items, and aim to ensure that the cumulative allocation at each round satisfies approximate envy-freeness—which we define as temporal envy-freeness up to one item (TEF1). We focus on settings where items can be exclusively goods or exclusively chores. For goods, while TEF1 allocations may not always exist, we identify several special cases where they do—two agents, two item types, generalized binary valuations, unimodal preferences—and provide polynomial-time algorithms for these cases. We also prove that determining the existence of a TEF1 allocation is NP-hard. For chores, we establish analogous results for the special cases, but present a slightly weaker intractability result. We also establish the incompatibility between TEF1 and Pareto-optimality, with the implication that it is intractable to find a TEF1 allocation that maximizes any p-mean welfare, even for two agents.

[58] Fair Allocation with Initial Utilities

Authors. Niclas Boehmer and **Luca Kreisel**

Abstract. The problem of allocating indivisible resources to agents arises in a wide range of domains, including treatment distribution and social support programs. An important goal in algorithm design for these problems is fairness, where the focus in previous work has been on ensuring that the computed allocation provides equal treatment to everyone. However, this perspective disregards that agents may start from unequal initial positions, which is crucial to consider in settings where fairness is understood as equality of outcome. In such settings, the goal is to create an equal final outcome for everyone by leveling initial inequalities through the allocated resources. To close this gap, focusing on agents with additive utilities, we extend the classic model by assigning each agent an initial utility and study the existence and computational complexity of several new fairness notions following the principle of equality of outcome. Among others, we show that complete allocations satisfying a direct analog of envy-freeness up to one resource (EF1) may fail to exist and are computationally hard to find, forming a contrast to the classic setting without initial utilities. As our main contribution, we propose a new, always satisfiable fairness notion, called minimum-EF1-init and design a polynomial-time algorithm based on an extended round-robin procedure to compute complete allocations satisfying this notion.

[60] Proportionality in Thumbs Up and Down Voting

Authors. Sonja Kraiczy, **Georgios Papasotiropoulos**, Grzegorz Pierczyński, Piotr Skowron

Abstract. In the classic committee election setting each voter approves a subset of candidates and the goal is to select k winners based on these preferences. A central focus of recent research in the area has been to achieve proportional representation, as formalized by the family of Justified Representation (JR) axioms. In this work, we explore notions of proportionality in a more expressive setting that allows voters to downvote candidates—a common feature on online polling platforms and beyond. We propose two conceptually distinct interpretations of down votes, resulting in different perspectives of proportionality. In the first, preventing the election of disapproved candidates is as important to voters as electing approved ones. In the second, approvals and disapprovals are treated separately, with each receiving its own fairness guarantees. For each approach, we introduce suitable axioms capturing proportionality and examine their satisfiability by appropriate variants of Phragmén’s rule, Proportional Approval Voting rule (PAV) and the Method of Equal Shares (MES).

[61] Proportional Selection in Networks

Authors. Georgios Papasotiropoulos, Oskar Skibski, Piotr Skowron, **Tomasz Wąs**

Abstract. We address the problem of selecting k representative nodes from a network, aiming to achieve two objectives: identifying the most influential nodes and ensuring the selection proportionally reflects the network’s diversity. We propose two approaches to accomplish this, analyze them theoretically, and demonstrate their effectiveness through a series of experiments.

[63] Connected Equitable Cake Division via Sperner's Lemma

Authors. Umang Bhaskar, **A. R. Sricharan**, Rohit Vaish

Abstract. We study the problem of fair cake-cutting where each agent receives a connected piece of the cake. A division of the cake is deemed fair if it is equitable, which means that all agents derive the same value from their assigned piece. Prior work has established the existence of a connected equitable division for agents with nonnegative valuations using various techniques. We provide a simple proof of this result using Sperner's lemma. Our proof extends known existence results for connected equitable divisions to significantly more general classes of valuations, including nonnegative valuations with externalities, as well as several interesting subclasses of general (possibly negative) valuations.

[64] Core-Stable Kidney Exchange via Altruistic Donors

Authors. **Gergely Csáji**, Thanh Nguyen

Abstract. Kidney exchange programs—connecting hospitals in the U.S. or countries in Europe—aim to improve efficiency by pooling donors and patients on a centralized platform. Core-stable allocations are essential for sustaining cooperation and preventing defections. However, when the core is empty, incentive problems emerge—hospitals or countries may withhold easily matched pairs for internal use, reducing the overall efficiency and scope of the exchange. We propose a solution to restore core stability by supplementing the platform with additional altruistic donors. Although the worst-case requirement for the number of such donors may be large, we show that in realistic settings, only a few are sufficient. We analyze two models of the compatibility graph: one based on random graphs, and the other on compatibility types. When only pairwise exchanges are allowed, the number of required altruists is bounded by the maximum number of independent odd cycles—disjoint odd cycles with no edges between them. This parameter grows logarithmically with market size in the random graph model, and is at most one-third the number of compatibility types in the type-based model. When small exchange cycles are allowed, it suffices for each participating organization to receive a number of altruists proportional to the number of compatibility types. Finally, simulations show that far fewer altruists are needed in practice than theory suggests.

[67] Collusion in the polluted river problem

Authors. **Mikel Álvarez-Mozos**, René van den Brink, Martí Jané-Ballarín

Abstract. We study the effects of collusion in the polluted river problem. In this model, agents along a river network have to share the costs of cleaning the entire river. We define a collusion neutrality property which requires that, if all agents upstream of a given point agree to collude and act as one, then the cost to be paid by this new entity equals the sum of the costs originally allocated to the colluding agents. This property is related to the Unlimited Territorial Integrity doctrine used to settle international disputes over water streams flowing through several territories. We introduce a new cost sharing method, called the Upstream Recursive Equal Sharing (URES) method, and axiomatically characterize it using our new collusion neutrality axiom. We also provide a new characterization for the well-known Local Responsibility Sharing (LRS) method that uses the axiom and differs from the URES method only in one property that explains how costs are shared at the source of the river. Finally, we derive and characterize a family of collusion neutral methods that contains the LRS method as an extreme and the URES method as an intermediate method. All members of the family can be computed by means of a recursive algorithm.

[75] Computing Efficient Envy-Free Partial Allocations of Indivisible Goods

Authors. **Bin Sun**, Robert Brederick, Junjie Luo, Andrzej Kaczmarczyk

Abstract. Envy-freeness is one of the most prominent fairness concepts in the allocation of indivisible goods. Even though trivial envy-free allocations always exist, rich literature shows this is not true when one additionally requires some efficiency concept (e.g., completeness, Pareto-efficiency, or social welfare maximization). In fact, in such case even deciding the existence of an efficient envy-free allocation is notoriously computationally hard. In this paper, we explore the limits of efficient computability by relaxing standard efficiency concepts and analyzing how this impacts the computational complexity of the respective problems. Specifically, we allow partial allocations (where not all goods are allocated) and impose only very mild efficiency constraints, such as ensuring each agent receives a bundle with positive utility. Surprisingly, even such seemingly weak efficiency requirements lead to a diverse computational complexity landscape. We identify several polynomial-time solvable or fixed-parameter tractable cases for binary utilities, yet we also find NP-hardness in very restricted scenarios involving ternary utilities.

[80] Whoever Said Money Won't Solve All Your Problems? Weighted Envy-free Allocation with Subsidy

Authors. Noga Klein Elmalem, Haris Aziz, **Rica Gonen**, Xin Huang, Kei Kimura, Indrajit Saha, Erel Segal-Halevi, Zhaohong Sun, Mashbat Suzuki, Makoto Yokoo

Abstract. We explore solutions for fairly allocating indivisible items among agents assigned weights representing their entitlements. Our fairness goal is weighted-envy-freeness (WEF), where each agent deems their allocated portion relative to their entitlement at least as favorable as any other's relative to their own. Often, achieving WEF necessitates monetary transfers, which can be modeled as third-party subsidies. The goal is to attain WEF with bounded subsidies.

Previous work relied on characterizations of unweighted envy-freeness (EF), that fail in the weighted setting. This makes our new setting challenging. We present polynomial-time algorithms that compute WEF allocations with a guaranteed upper bound on total subsidy for monotone valuations and various subclasses thereof.

We also present an efficient algorithm to compute a fair allocation of items and money, when the budget is not enough to make the allocation WEF. This algorithm is new even for the unweighted setting.

[83] Fairly Stable Two-Sided Matching with Indifferences

Authors. **Benjamin Cookson**, Nisarg Shah

Abstract. Stability has been a foundational criterion for two-sided matching. When agents on one side have weak preferences involving indifferences, the seminal work of Kesten and Ünver [2015] proposes the Fractional Deferred Acceptance (FDA) algorithm for computing a fractional matching that satisfies (ex ante) stability along with a fairness criterion that ensures no discrimination among (equally-preferred) agents on one side.

We show that their algorithm can actually fail to terminate, refuting their claim of (polynomial-time) termination. Using substantially new algorithmic ideas, we develop an algorithm, Doubly-Fractional Deferred Acceptance Via Strongly Connected Components (DFDA-SCC), which can handle agents on both sides exhibiting indifferences and, in polynomial time, compute a fractional matching satisfying ex ante stability and no ex ante discrimination among agents on both sides.

[90] An Impossibility Result for Strongly Group-Strategyproof Multi-Winner Approval-Based Voting

Authors. **Iannis Caragiannis**, Rob LeGrand, Evangelos Markakis, Emmanouil Pountourakis

Abstract. Multi-winner approval-based voting has received considerable attention recently, as an election format. A voting rule in this setting takes as input ballots in which each agent approves a subset of the available alternatives and outputs a committee of alternatives of a given size k . We consider the scenario when a coalition of agents can act strategically and alter their ballots so that the new outcome is strictly better for some coalition member and at least as good for anyone else in the coalition. Voting rules that are robust against this strategic behaviour are called strongly group-strategyproof. We prove that, for $k \in \{1, 2, \dots, m-2\}$, strongly group-strategyproof multi-winner approval-based voting rules which furthermore satisfy the minimum efficiency requirement of unanimity do not exist, where m is the number of available alternatives. Our proof builds a connection to single-winner voting with ranking-based ballots and exploits the infamous Gibbard-Satterthwaite theorem to reach the desired impossibility result. Our result has implications for paradigmatic problems from the area of approximate mechanism design without money and indicates that strongly group-strategyproof mechanisms for minimax approval voting, variants of facility location, and classification can only have an unbounded approximation ratio.

[95] AI-Generated Compromises for Coalition Formation

Authors. Eyal Briman, Ehud Shapiro, **Nimrod Talmon**

Abstract. The challenge of finding compromises between agent proposals is fundamental to AI sub-fields such as argumentation, mediation and negotiation. Building on this tradition, Elkind et al. introduced a process for coalition formation that seeks majority-supported proposals preferable to the status quo, using a metric space where each agent has an ideal point. The crucial step in this iterative process involves identifying compromise proposals around which agent coalitions can unite. How to effectively find such compromise proposals, however, remains an open question. We address this gap by formalizing a holistic model that encompasses agent bounded rationality and uncertainty and developing AI models to generate such compromise proposals. We focus on the domain of collaboratively writing text documents – e.g., to enable the democratic creation of a community constitution. We apply NLP (Natural Language Processing) techniques and utilize LLMs (Large Language Models) to create a semantic metric space for text and develop algorithms to suggest suitable compromise points. To evaluate the effectiveness of our algorithms, we simulate various coalition formation processes and demonstrate the potential of AI to facilitate large-scale democratic text editing, such as collaboratively drafting a constitution—an area where traditional tools are limited.

[97] Maximizing Index Diversity in Committee Elections

Authors. **Paula Böhm**, Robert Bredereck, Till Fluschnik

Abstract. We introduce two models of multiwinner elections with approval preferences and labelled candidates that take the committee’s diversity into account. One model aims to find a committee with maximal diversity given a scoring function (e.g. of a scoring-based voting rule) and a lower bound for the score to be respected. The second model seeks to maximize the diversity given a minimal satisfaction for each agent to be respected. To measure the diversity of a committee, we use multiple diversity indices used in ecology and introduce one new index. We define (desirable) properties of diversity indices, test the indices considered against these properties, and characterize the new index. We analyze the computational complexity of computing a committee for both models with most of the indices considered and scoring functions of well-known voting rules, and investigate the influence of weakening the score or satisfaction constraints on the diversity empirically.

[99] It’s Not All Black and White: Degree of Truthfulness for Risk-Avoiding Agents

Authors. Eden Hartman, **Erel Segal-Halevi**, Biaoshuai Tao

Abstract. The classic notion of *truthfulness* requires that no agent has a profitable manipulation — an untruthful report that, for *some* combination of reports of the other agents, increases her utility. This strong notion implicitly assumes that the manipulating agent either knows what all other agents are going to report, or is willing to take the risk and act as-if she knows their reports.

Without knowledge of the others’ reports, most manipulations are *risky* – they might decrease the manipulator’s utility for some other combinations of reports by the other agents. Accordingly, a recent paper (Bu, Song and Tao, “On the existence of truthful fair cake cutting mechanisms”, *Artificial Intelligence* 319 (2023), 103904) suggests a relaxed notion, which we refer to as *risk-avoiding truthfulness (RAT)*, which requires only that no agent can gain from a *safe* manipulation — one that is sometimes beneficial and never harmful. Truthfulness and RAT are two extremes: the former considers manipulators with complete knowledge of others, whereas the latter considers manipulators with no knowledge at all. In reality, agents often know about some — but not all — of the other agents. This paper introduces the *RAT-degree* of a mechanism, defined as the smallest number of agents whose reports, if known, may allow another agent to safely manipulate, or n if there is no such number. This notion interpolates between classic truthfulness (degree n) and RAT (degree at least 1): a mechanism with a higher RAT-degree is harder to manipulate safely.

To illustrate the generality and applicability of this concept, we analyze the RAT-degree of prominent mechanisms across various social choice settings, including auctions, indivisible goods allocations, cake-cutting, voting, and two-sided matching.

[106] On Iterative Voting Outcomes in Plurality Elections

Authors. Vincent Mousseau, Henri Surugue, **Magdaléna Tydrichova**, Anaëlle Wilczynski

Abstract. This article deals with iterative voting under the plurality rule, where voters can strategically perform sequential deviations. Most works in iterative voting focus on convergence properties or evaluate the quality of the resulting outcome. However, the iterative winner depends on the sequence of voters' deviations. We propose to analyze to what extent this impacts the outcome of iterative voting by adopting a qualitative, quantitative and computational approach. In particular, we introduce the notions of possible and necessary iterative winners. We first study the extreme configuration for the existence of a necessary winner, where no voter has an incentive to deviate from her truthful ballot. We show that this phenomenon occurs with high probability under impartial cultures. Then, we explore the computational complexity of determining possible and necessary iterative winners, proving that the two problems fall in different complexity classes. Finally, we investigate the election of the Condorcet winner as an iterative winner and theoretically prove that the Condorcet efficiency of plurality is increased by considering its iterative voting version.

[107] Agreement among Voting Rules under Single-Peaked Preference Distributions

Authors. Vincent Mousseau, **Henri Surugue**, Anaëlle Wilczynski

Abstract. Many different voting rules have been proposed in the literature and they can select very different alternatives. This naturally raises the question of whether this diversity in outcomes often occurs. Previous works have shown that the probability that voting rules agree on the same outcome is generally quite low under impartial culture. In this article, we use a similar probabilistic approach on single-peaked cultures, which are more structured and typically more realistic than impartial culture. We provide conditions for voting rules to agree under standard single-peaked cultures, and show that the probability of agreement between rather large families of voting rules is much higher under such cultures, with fast convergence of this probability with respect to the number of voters. We finally provide some insights on other structured preference distributions, observing that many exhibit similar convergence in agreement, including the Mallows' distribution. Our study reveals a tendency of several well-known voting cultures to bias the outcome of voting rules, which is worth knowing before conducting experiments on synthetic data.

[108] An $O(\log \log n)$ -Approximation for Submodular Facility Location

Authors. Fateme Abbasi, Marek Adamczyk, Miguel Bosch-Calvo, Jarosław Byrka, Fabrizio Grandoni, **Krzysztof Sornat**, Antoine Tinguely

Abstract. In the Submodular Facility Location problem (SFL) we are given a collection of n clients and m facilities in a metric space. A feasible solution consists of an assignment of each client to some facility. For each client, one has to pay the distance to the associated facility. Furthermore, for each facility f to which we assign the subset of clients S^f , one has to pay the opening cost $g(S^f)$, where $g(\cdot)$ is a monotone submodular function with $g(\emptyset) = 0$.

SFL captures practical scenarios where the cost of opening a facility is a (non-linear, still “tractable”) function of the set of served clients. For example, each client might have different types of needs, and satisfying such needs might have a submodular impact on the opening cost (regardless of the facility location). SFL can model participatory budgeting within a spatial voting framework, where clients are voters positioned in a metric space, and facilities correspond to projects considered for implementation.

SFL is APX-hard since it includes the classical (metric uncapacitated) Facility Location problem (with uniform facility costs) as a special case. Svitkina and Tardos [SODA’06] gave the current-best $O(\log n)$ approximation algorithm for SFL. The same authors pose the open problem whether SFL admits a constant approximation and provide such an approximation for a very restricted special case of the problem.

We make some progress towards the solution of the above open problem by presenting an $O(\log \log n)$ approximation. Our approach is rather flexible and can be easily extended to generalizations and variants of SFL. In more detail, we achieve the same approximation factor for the natural generalizations of SFL where the opening cost of each facility f is of the form $p_f + g(S^f)$ or $w_f \cdot g(S^f)$, where $p_f, w_f \geq 0$ are input values.

We also obtain an improved approximation algorithm for the related Universal Stochastic Facility Location problem. In this problem one is given a classical (metric) facility location instance and has to a priori assign each client to some facility. Then a subset of active clients is sampled from some given distribution, and one has to pay (a posteriori) only the connection and opening costs induced by the active clients. The expected opening cost of each facility f can be modelled with a submodular function of the set of clients assigned to f .

[109] Reducing Leximin Fairness to Utilitarian Optimization

Authors. Eden Hartman, Yonatan Aumann, Avinatan Hassidim, **Erel Segal-Halevi**

Abstract. Two prominent objectives in social choice are *utilitarian* - maximizing the sum of agents’ utilities, and *leximin* - maximizing the smallest agent’s utility, then the second-smallest, etc. Utilitarianism is typically computationally easier to attain but is generally viewed as less fair. This paper presents a general reduction scheme that, given a utilitarian solver, produces a distribution over states (deterministic outcomes) that is leximin in expectation. Importantly, the scheme is robust in the sense that, given an *approximate* utilitarian solver, it produces a lottery that is approximately-leximin (in expectation) - with the same approximation factor. We apply our scheme to several social choice problems: stochastic allocations of indivisible goods, giveaway lotteries, and fair lotteries for participatory budgeting.

[112] Condorcet Winners and Anscombe's Paradox Under Weighted Binary Voting


Authors. Carmel Baharav, **Andrei Constantinescu**, Roger Wattenhofer

Abstract. We consider voting on multiple independent binary issues. In addition, a weighting vector for each voter defines how important they consider each issue. The most natural way to aggregate the votes into a single unified proposal is issue-wise majority (IWM): taking a majority opinion for each issue. However, in a scenario known as Ostrogorski's Paradox, an IWM proposal may not be a Condorcet winner, or it may even fail to garner majority support in a special case known as Anscombe's Paradox.

We show that it is co-NP-hard to determine whether there exists a Condorcet-winning proposal even without weights. In contrast, we prove that the single-switch condition provides an Ostrogorski-free voting domain under identical weighting vectors. We show that verifying the condition can be achieved in linear time and no-instances admit short, efficiently computable proofs in the form of forbidden substructures. On the way, we give the simplest linear-time test for the voter/candidate-extremal-interval condition in approval voting and the simplest and most efficient algorithm for recognizing single-crossing preferences in ordinal voting.

We then tackle Anscombe's Paradox. Under identical weight vectors, we can guarantee a majority-supported proposal agreeing with IWM on strictly more than half of the overall weight, while with two distinct weight vectors, such proposals can get arbitrarily far from IWM. The severity of such examples is controlled by the maximum average topic weight $\tilde{w}_{max}$: a simple bound derived from a partition-based approach is tight on a large portion of the range $\tilde{w}_{max} \in (0, 1)$. Finally, we extend Wagner's rule to the weighted setting: an average majority across topics of at least $3/4$'s precludes Anscombe's paradox from occurring.


Main Venue & Mensa Map



Main Venue (TUtheSky): Getreidemarkt 9, 11th floor, 1060 Wien


Mensa (Campus Freihaus): Wiedner Hauptstraße 8-10, 2nd floor, 1040 Wien

Dinner Map (TUtheSky → Stöckl im Park), September 18th, 18:30-21:30



Location: Stöckl im Park, Prinz-Eugen-Straße 25, 1030 Wien
Website: <https://www.stoecklimpark.at>


Parliament Tour (TUtheSky → Parliament)



Parliament Tour

Date and time: Wednesday, September 17, 16:00-18:00
Address: Dr.-Karl-Renner-Ring 3, 1017 Wien
Transport: 15–20 minute walk from the main venue.
Route: See map above or

<https://maps.app.goo.gl/k4Zs7xbbTrmgNeKYA>



COMSOC 2025 Website

<https://www.ac.tuwien.ac.at/comsoc2025/>