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 Divison 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	
08:45–09:45	Keynote: **Axel Ockenfels** *Behavioral Market Design*
09:45–10:15 Coffee Break	
10:15–10:35	**Dery**: *Interactive and Iterative Peer Assessment: A Social Choice Apprach to Eliciting and Aggregating Student Evaluations*
10:35–10:55	Horn, Nüsken, Rothe, **Seeger**: *Skating System Unveiled: Exploring Preference Aggregation in Ballroom Tournaments*
10:55–11:15	Ghosh, Neoh, Teh, **Tyrovolas**: *Fraud-Proof Revenue Division on Subscription Platforms*
11:15–11:25 Short Break	
11:25–11:45	**Delemazure**, Peters, Freeman, Lang, Laslier: *Reallocating Wasted Votes in Proportional Parliamentary Elections with Thresholds*
11:45–12:05	**Boehmer**, Fish, Procaccia: *Generative Social Choice: The Next Generation*
12:05-14:00 Lunch Break	
14:00–14:20	**Prévotat**, Terzopoulou, Zylbersztejn: *Strategizing under Rule and Vote Uncertainty: An Experiment*
14:20–14:40	Boehmer, Faliszewski, Janeczko, Kaczmarczyk, Lisowski, Pierczyński, Rey, Stolicki, **Szufa**, Wąs: *Guide to Numerical Experiments on Elections in COMSOC*
14:40–15:00	**Faliszewski**, Janeczko, Kaczmarczyk, Kurdziel, Pierczyński, Szufa: *Learning Real-Life Approval Elections*
15:00–15:30 Coffee Break	
15:30–15:50	Aziz, **Lederer**, Peters, Peters, Ritossa: *Committee Monotonicity and Proportional Representation for Ranked Preferences*
15:50–16:10	**Bardal**, Brill, McCune, Peters: *Proportionality in Practice: Quantifying Proportionality in Ordinal Elections*
16:10–16:30 Short Break	
16:30–16:50	Bullinger, Romen, **Schlenga**: *The Power of Matching for Online Fractional Hedonic Games*
16:50–17:10	**Caragiannis**, Roy: *Quantile Agent Utility and Implications to Randomized Social Choice*
17:10–17:30	**Constantinescu**, Wattenhofer: *Byzantine Game Theory: Sun Tzu's Boxes*
17:30–18:00 Group Photo	
18:30–22:00 Workshop Dinner	

Friday, September 19	
08:45–09:45	Keynote: **Moon Duchin** *Beyond Sushi: Bridging from COMSOC to Practical Democracy*
09:45–10:15 Coffee Break	
10:15–10:35	Brandt, **Dong**, Peters: *Condorcet-Consistent Choice Among Three Candidates*
10:35–10:55	**Meir**, Ghaime: *On Condorcet's Jury Theorem with Abstention*
10:55–11:15	**Ding**, Yu: *An Axiomatic Characterization of the Minimax Voting Method*
11:15–11:25 Short Break	
11:25–11:45	**Berker**, Casacuberta, Robinson, Ong, Conitzer, Elkind: *From Independence of Clones to Composition Consistency: A Hierarchy of Barriers to Strategic Nomination*
11:45–12:05	Schmidt-Kraepelin, Suksompong, **Utke**: *Discrete Budget Aggregation: Truthfulness and Proportionality*
12:05–14:00 Lunch Break	
14:00–14:20	Dong, **Frank**, Peters, Suksompong: *Reconfiguring Proportional Committees*
14:20–14:40	**Kehne**, Schmidt-Kraepelin, Sornat: *Robust Committee Voting, or The Other Side of Representation*
14:40–15:00	Masařík, Pierczyński, **Skowron**: *A Generalised Theory of Proportionality in Collective Decision Making*
15:00–15:30 Coffee Break	
15:30–15:50	Papasotiropoulos, **Pishbin**, Skibski, Skowron, Wąs: *Method of Equal Shares with Bounded Overspending*
15:50–16:10	Kraiczy, Robinson, **Elkind**: *Streamlining Equal Shares*
16:10–16:30	**Baychkov**, Brill, Utke: *Mixed Voting Rules for Participatory Budgeting*
16:30–16:50 Short Break	
16:50–17:10	**Nguyen**, Song: *Approximate Core of Participatory Budgeting via Lindahl Equilibrium*
17:10–17:30	Faliszewski, **Janeczko**, Knop, Pokorný, Schierreich, Słuszniak, Sornat: *Participatory Budgeting Project Strength via Candidate Control*
17:30–19:00 Business Meeting	

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 Méroué

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 Zylbersztejn

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 Stolicki, **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 IJCAI, 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 $t >= l$, 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-word 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 spend 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 Słuszniak, 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.