

Closing the Loop: Bringing Humans into Empirical Computational Social Choice and Preference Reasoning

Nicholas Mattei

Department of Computer Science, Tulane University
nsmattei@tulane.edu

Abstract

Research in both computational social choice and preference reasoning uses tools and techniques from computer science, generally algorithms and complexity analysis, to examine topics in group decision making. This has brought tremendous progress in the last decades, creating new avenues for research and results in areas including voting and resource allocation. I argue that of equal importance to the theoretical results are the impacts in research and development from the *empirical* part of the computer scientists toolkit: data, system building, and human interaction. I highlight work by myself and others to establish data driven, application driven research in the computational social choice and preference reasoning areas. Along the way, I highlight interesting application domains and important results from the community in driving this area to make concrete, real-world impact.

1 Introduction

The Internet enables computers and, by proxy, humans to both communicate and form groups at distances, speeds, and scales that recently lived only in the imagination. These groups of agents, be they human, computer, or a mix of the two, must make collective decisions subject to external and internal constraints and preferences in many important real-world settings. Group decisions have classically been the purview of research in *social choice*, which concerns itself with the study of how groups, where each member has their own preferences, make decisions that must be then shared by that group [Sen, 2018]. The field of *computational social choice (COMSOC)* adds to the standard model computational questions and modes of analysis including algorithms and complexity [Brandt *et al.*, 2016a] as well as incorporating work more broadly in artificial intelligence including preference reasoning and machine learning. [Domshlak *et al.*, 2011]. Taken together, this work falls mostly in multi-agent systems research [Shoham and Leyton-Brown, 2008; Wooldridge, 2009], where we study not only systems, but the interactions between many systems. This intersection of tools and techniques has led to broad and important results

in areas including resource allocation problems such as kidney exchange [Roth *et al.*, 2005; Dickerson *et al.*, 2014], school choice [Abdulkadiroğlu *et al.*, 2005], routing traffic [Kleinberg *et al.*, 1999], fair division of continuous resources [Aziz and Mackenzie, 2016], allocating radio spectra [Leyton-Brown *et al.*, 2017], fair assignment [Aziz, 2019; Lian *et al.*, 2018], and allocating security resources for patrol [Tambe, 2012]; as well as distributed decision tasks including (multi-winner) voting [Conitzer *et al.*, 2007; Faliszewski *et al.*, 2017], distributed democracy [Brill, 2018], and peer selection [Aziz *et al.*, 2019; Xu *et al.*, 2018].

In a traditional social choice setting we have a finite set of agents $N = \{1, \dots, n\}$ and a finite set of alternatives $A = \{1, \dots, m\}$. Each agent $i \in N$ has an order-able preference \succsim_i over the alternatives. In social choice we will focus on two important goals: *voting* and *allocation*. In voting, the preferences of all agents are combined to select either an element of or an ordering over the set of alternatives A to *share amongst the group*. In allocation, the items within A are to be distributed or *allocated* to the set of agents in N . The common thread among these settings is that we have a group of self-interested agents submit their preferences to a centralized or de-centralized authority and outcomes are decided by some mechanism. Each mechanism for group decision making may or may not satisfy various design goals such as fairness and/or efficiency as analyzed with formal models of agent behavior [Brandt *et al.*, 2016a; Shoham and Leyton-Brown, 2008]. In some instances these models of behavior provide a limited or (occasionally) unrealistic view into many real-world settings. Researchers in economics have shown where the predictions of game theory are contradicted by data or experiment; giving rise to behavioral and experimental economics [Kagel and Roth, 1995]. Consideration of the human element has led to significant impacts in the areas of markets and matching including school choice [Abdulkadiroğlu *et al.*, 2005] and kidney exchanges [Roth *et al.*, 2005]. However, until recently there has not been a data driven, systems oriented research program that directly questions these strong assumptions in COMSOC research.

My work focuses on:

1. closing the gap between these idealized models of agent behavior and reality by using empirical, data driven studies focusing on how humans really make decisions;
2. building concrete, real-world systems based on these

findings to encourage broad dissemination and availability of the research results; and

3. promoting empirical and data driven work in the community by building websites like www.preflib.org and running data focused workshops like the Exploring Beyond the Worst Case in Computational Social Choice (EXPLORE) and the Games, Agents, and Incentives Workshop (GAIW).

One of the most exciting things about computer science research is that we can both collaborate broadly with researchers in economics, psychology, biology and beyond and use our engineering skills to *create real systems* that can be completely automated or interact with humans at many different scales. I argue that a focus on data and real-world system building leads to the largest research impacts. This is an engineering mindset: to discover new theory and novel techniques to build systems that make the world a better place by advancing the state of the art in intelligent systems.

2 Empirical Computational Social Choice and Preference Reasoning

A key difference between social choice and *computational social choice* (COMSOC) is the focus on the concerns of a computer scientist: approximations, algorithms, and, I would argue, *systems*. The focus of my work has been to build systems and build theoretical models that help people and agents make decisions. In order to do this effectively we need to engage with the messy and sometimes downright irrational ways that choices and systems can be used in the real world – and integrate these aspects into preference reasoning and social choice research [Popova *et al.*, 2013; Allen *et al.*, 2015]. I argue that the most successful research projects are a blend of theoretical innovations, experimental analysis, and practical tool building. We should walk the same road that Kagel and Roth [1995] describe for experimental economics: evolving from theory, to simulated or re-purposed data, to full fledged laboratory and field experiments. This progression has enabled a “conversation” between experimental and theoretical researchers that, in turn, significantly advanced the field. In this section I detail some efforts to advance down this road.

Within the broader social choice field looking at voting and allocations, there has been work that analyzed large datasets to understand human behavior. Regenwetter *et al.* [2007] collected data from elections of professional societies, analyzing these elections for paradoxes and evidence of various proposed voter utility models. Likewise, the work of Plassmann and Tideman [2014] collected a number of historic elections in an effort to quantify the likelihood of various election paradoxes. In areas such as recommender systems and machine learning, using data to understand user decisions was already well established [Bache and Lichman, 2013; Kamishima, 2003]. Finally, in the broader matching community there is significant work on human behavior for market systems [Abdulkadiroğlu *et al.*, 2005; Roth *et al.*, 2005]. However, within COMSOC there had not been a broad, data driven program of research.

When I started my research, I wanted to see empirically

how these theories and models of agent behavior for, e.g., manipulation and voting paradoxes, manifested in the real world. I was struck by the small number of empirical research papers in the COMSOC community and wanted to expand the use of data [Walsh, 2011]. I performed empirical studies of voting systems, mining large sets of real-world preference data to create elections from user preferences, and found that many common assumptions about worst-case agent behavior, preference domain restrictions, and election paradoxes, did not occur in practice [Mattei, 2011; Mattei *et al.*, 2012]. To encourage more researchers to include experimental and empirical verification in their work, Toby Walsh and I established PREFLIB.ORG and PREFLIB TOOLS for working with real-world preference data [Mattei and Walsh, 2017; Mattei and Walsh, 2013]. To date we have gathered thousands of instances of data from real-world domains where agents express preferences in domains including voting, resource allocation, ranking, and rating. We unified this data into a common format, making it easy to access and reuse this data.

The use of PREFLIB and the broader mindset of empirical verification has expanded across the COMSOC community and beyond, feeding back into social choice and other domains. Examples include the area of multiwinner voting, which is typically a computationally hard problem, where Skowron *et al.* [2015] and others have used data and experiments to understand when certain multiwinner voting rules admit good approximations in practice or are computationally tractable. Brederbeck *et al.* [2019] use experiments on PREFLIB and simulated data to see when committees, selected using multiwinner voting rules, satisfy important properties such as justified representation. Finally, Szufa *et al.* [2020] use models generated from agent preferences to define notions of distances between elections, and then experimentally use optimization techniques to create novel multiwinner voting rules that satisfy good properties across a range of voting behavior. More than just enabling experiments, PREFLIB has provided a large library of data to use for looking at election paradoxes and inconsistencies. The data found in PREFLIB in conjunction with advanced automated proof finding techniques have led to a deeper understanding of the causes of voting paradoxes [Brandt *et al.*, 2016b].

More than just re-purposing existing data, researchers in social choice have conducted experiments to gather data about how people really make decisions and use this knowledge to feed back into the models used for theoretical research. Meir *et al.* [2020b] detail the design and implementation of an experimental system where choices of humans faced with group choice settings with access to uncertain poll information about an election are recorded. The goal of these experiments is to quantify how often and what types of strategic behavior people use in practice. Likewise, Mennle *et al.* [2015] ran a set of experiments in the assignment setting where agents could strategically misreport their preferences in order to gain more utility. The setting involved a computationally hard manipulation problem so most agents relied on simpler heuristics to find better, but not the best, strategic profiles. Finally, Scheuerman *et al.* [2019] have conducted a series of experiments on Mechanical Turk to understanding what factors of the decision environment including uncer-

tainty and payoff amounts affect when and how agents will act strategically in approval voting settings.

This use of empirical research has begun to feed back into the *theoretical research* of how agents make decisions when faced with specific group choice environments and mechanisms. Zou *et al.* [2015] use data from Doodle, arguably the most widely used system to make group decisions, and determined behavioral effects which led to different voting patterns such as the difference in approvals of public versus private polls. Zhao *et al.* [2018] make use of the OPERA system to devise a querying methodology for group decision making that minimizes the number of queries of individual users in order to come to a group decision quickly. Finally, Fairstein *et al.* [2019], present an analysis of the experimental data from Meir *et al.* [2020b] to propose a novel voter decision model which more accurately accounts for how agents make decisions in strategic voting environments.

3 Use Inspired Systems: Building Better Choices

In addition to a broader culture of experimentation with data, there has been a growing focus on COMSOC research that targets specific end users and systems. Perhaps some of the greatest successes in this area have been in the field of matching and resource allocation including kidney exchanges [Roth *et al.*, 2005], spectrum auctions [Leyton-Brown *et al.*, 2017], and security games [Tambe, 2012].

Since PREFLIB TOOLS a number of systems have been developed for general use. The Spliddit [Goldman and Procaccia, 2014] website allows any user to use some of the best algorithms in fair division research to split credit, rent, or even the cost of a cab. We have used similar ideas in order to help businesses in Australia divide the cost of on demand pickup and delivery problems, devising new heuristics and tools [Aziz *et al.*, 2016a]. For voting and group decision making the website OPERA [Chen *et al.*, 2020] provides a user friendly, interactive system to deploy a range of single and multiwinner voting rules for use in group decision making. Likewise, users interested in deploying results from randomized social choice can use the Pynx system [Brandt *et al.*, 2015]. All of these systems are highly polished, provide mobile and internet interfaces, and are used around the world.

Another significant application of research in COMSOC has been results in peer selection, conference reviewing, and other work assignment settings. The broad area of artificial intelligence is experiencing unprecedented growth in the number of papers written and submitted for review and new ideas are surfacing on how to handle this load [Shah *et al.*, 2018]. One idea for improving these systems is to create a market for the bids of agents on papers, as many papers are chronically under-demanded, leading to poor review quality [Meir *et al.*, 2020a]. Once papers have bids, there is still the assignment problem left to be solved. There have been a number of recent works that use real-world data to test novel objectives for fair and efficient assignments [Lian *et al.*, 2018; Aziz, 2019] as well as handle the discrepancy between scores assigned by the reviewers [Wang and Shah, 2019].

One idea to handle the load has been to enforce that sub-

mitters must also participate in the review process itself. This creates a peer selection problem, where agents evaluate one another and, based on these evaluations, a subset of agents are selected as winners. An obvious incentive problem arises in this setting: an agent may lie about their valuation for other agents in order to increase their chances of being selected. Peer selection has become an important topic in recent years for numerous applications: academic peer review including NSF grant reviewing; crowd-sourcing corporate or internal brainstorming sessions; and performing peer review for MOOCs. There are a number of novel strategy-proof mechanisms in this setting including partitioning methods [Kurokawa *et al.*, 2015], methods to handle complex interaction graphs between agents [Xu *et al.*, 2018], partitioning methods that re-balance themselves [Aziz *et al.*, 2016b; Aziz *et al.*, 2019], and methods that relax exactness to achieve better overall efficiency [Mattei *et al.*, 2020]. These results on assignment and selection will have broad impact not only for conference reviewing but hopefully the broader scientific enterprise and anywhere where work needs to be distributed and reviewed.

4 Human Centered Choices and Preference Reasoning

An important topic that emerges when we focus on use inspired systems is understanding how, when, and where these systems are used. One aspect of this topic where I have been working is in helping our students think creatively about technological systems and their impacts [Burton *et al.*, 2018]. On the development side, economics and social choice have long had important things to say about the trade-offs sometimes necessary between efficiency and fairness including work focusing on the division of continuous resources such as land or water [Moulin, 2004] and on more discrete, indivisible settings such as goods and services [Thomson, 2011]. In political science and political economy we see extensive research into the fairness and equality of voting and other representation schemes [Brams, 2008] and more fundamentally in the areas of political economy having to do with justice and fair distribution of resources to individuals [Rawls, 1971].

As algorithms make more decisions in our daily lives, and the consequences of these decisions become more fraught [Corbett-Davies *et al.*, 2017], it is important to leverage the history of thinking about fairness from these diverse sources. Equity and other concerns, formalized as economic axioms have a long history in social choice both in allocation and voting [Sen, 2018] We hope to leverage this long history of study of the axioms, or properties, of the algorithms and aspects including *fairness* to improve our systems. As we build our systems we need to understand and embed the constraints of our society that can come from numerous places including morals, laws, business process, and other sources [Rossi and Mattei, 2019]. There has been extensive work in the COMSOC and preference reasoning communities to formalize new preference models for use in ethical decision making [Loreggia *et al.*, 2018] and even use aggregation rules to find compromises between the differing view of ethics that may be present in a community [Noothigattu *et al.*, 2017].

Kidney allocation and exchanges are perhaps one of the earliest, well known examples of use inspired systems where there has both been significant impacts on practice and significant attention paid to fairness and equity concerns [Roth *et al.*, 2005; Dickerson *et al.*, 2012]. There have been extensions looking at the trade-off between fairness and efficiency in these markets and work on using voting to align the outcomes of these algorithms to the values that human decision makers use in the United States [Freedman *et al.*, 2020]. We have also worked with the Australian Organ and Tissue Authority to theoretically and empirically analyze the efficiency and fairness trade-offs of their proposed online organ allocation mechanisms [Mattei *et al.*, 2017; Mattei *et al.*, 2018].

5 Conclusion

Work in multi-agent systems, COMSOC, and preference reasoning has and continues to have meaningful impact in the world in which we live. As we move forward in our research I argue that it is important to incorporate empirical, data driven understanding of human behavior in our models and focus on system building as a key part of research. With this comes a responsibility to do this work for the social good and to use our knowledge of decision making to inform how actual decisions are being made. There are many avenues to explore going forward including work in social networks, voter models, and online allocation that I have not touched on here. In sum, it is an exciting time to be a researcher in these areas as we can have broad, significant impact in the real world.

Acknowledgements

None of this would be possible without my coauthors, research partners, and mentors who have continually challenged me to do great work. I am especially grateful to Haris Aziz, John P. Dickerson, Judy Goldsmith, Omer Lev, and Elizabeth Mattei for comments on drafts.

References

- [Abdulkadiroğlu *et al.*, 2005] A. Abdulkadiroğlu, P. A. Pathak, and A. E. Roth. The new york city high school match. *American Economic Review*, 95(2):364–367, 2005.
- [Allen *et al.*, 2015] T.E. Allen, M. Chen, J. Goldsmith, N. Mattei, A. Popova, M. Regenwetter, F. Rossi, and C. Zwillig. Beyond theory and data in preference modeling: Bringing humans into the loop. In *Proc. 4th ADT*, 2015.
- [Aziz and Mackenzie, 2016] H. Aziz and S. Mackenzie. A discrete and bounded envy-free cake cutting protocol for any number of agents. In *Proc. 57th FOCS*, pages 416–427, 2016.
- [Aziz *et al.*, 2016a] H. Aziz, C. Cahan, C. Gretton, P. Kilby, N. Mattei, and T. Walsh. A study of proxies for shapley allocations of transport costs. *JAIR*, 56:573–611, 2016.
- [Aziz *et al.*, 2016b] H. Aziz, O. Lev, N. Mattei, J. S. Rosenschein, and T. Walsh. Strategyproof peer selection: Mechanisms, analyses, and experiments. In *Proc. 30th AAAI*, 2016.
- [Aziz *et al.*, 2019] H. Aziz, O. Lev, N. Mattei, J. S. Rosenschein, and T. Walsh. Strategyproof peer selection using randomization, partitioning, and apportionment. *Artif. Intell.*, 275:295–309, 2019.
- [Aziz, 2019] H. Aziz. Developments in multi-agent fair allocation. *CoRR*, abs/1911.09852, 2019.
- [Bache and Lichman, 2013] K. Bache and M. Lichman. UCI Machine Learning Repository, 2013. University of California, Irvine, School of Information and Computer Sciences.
- [Brams, 2008] Steven J Brams. Mathematics and democracy: Designing better voting and fair division procedures. *Mathematical and Computer Modelling*, 48(9-10):1666–1670, 2008.
- [Brandt *et al.*, 2015] F. Brandt, G. Chabin, and C. Geist. Pnyx: a powerful and user-friendly tool for preference aggregation. In *Proc. 14th AAMAS*, pages 1915–1916, 2015.
- [Brandt *et al.*, 2016a] F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- [Brandt *et al.*, 2016b] F. Brandt, C. Geist, and M. Strobel. Analyzing the practical relevance of voting paradoxes via ehrhart theory, computer simulations, and empirical data. In *Proc. 16th AAMAS*, pages 385–393, 2016.
- [Bredereck *et al.*, 2019] R. Bredereck, P. Faliszewski, A. Kaczmarczyk, and R. Niedermeier. An experimental view on committees providing justified representation. In *Proc. 28th IJCAI*, pages 109–115. AAAI Press, 2019.
- [Brill, 2018] M. Brill. Interactive democracy. In *Proc. 17th AAMAS*, pages 1183–1187, 2018.
- [Burton *et al.*, 2018] E. Burton, J. Goldsmith, and N. Mattei. How to teach computer ethics through science fiction. *Commun. ACM*, 61(8):54–64, 2018.
- [Chen *et al.*, 2020] Y. Chen, J. Qian, J. Wang, L. Xia, and G. Zahavi. OPRA: An open-source online preference reporting and aggregation system, 2020.
- [Conitzer *et al.*, 2007] V. Conitzer, T. Sandholm, and J. Lang. When are elections with few candidates hard to manipulate? *JACM*, 54(3):14, 2007.
- [Corbett-Davies *et al.*, 2017] S. Corbett-Davies, E. Pierson, A. Feller, S. Goel, and A. Huq. Algorithmic decision making and the cost of fairness. In *Proc. 23rd SIGKDD*, pages 797–806. ACM, 2017.
- [Dickerson *et al.*, 2012] J. P. Dickerson, A. D. Procaccia, and T. Sandholm. Optimizing kidney exchange with transplant chains: Theory and reality. In *Proc. 11th AAMAS*, pages 711–718, 2012.
- [Dickerson *et al.*, 2014] J. P. Dickerson, A. D. Procaccia, and T. Sandholm. Price of fairness in kidney exchange. In *Proc. 13th AAMAS*, pages 1013–1020, 2014.
- [Domshlak *et al.*, 2011] C. Domshlak, E. Hüllermeier, S. Kaci, and H. Prade. Preferences in AI: An overview. *AI*, 175(7):1037–1052, 2011.
- [Fairstein *et al.*, 2019] R. Fairstein, A. Lauz, R. Meir, and K. Gal. Modeling people’s voting behavior with poll information. In *Proc. 18th AAMAS*, 2019.
- [Faliszewski *et al.*, 2017] P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Multiwinner voting: A new challenge for social choice theory. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 2. 2017.
- [Freedman *et al.*, 2020] R. Freedman, J. S. Borg, W. Sinnott-Armstrong, J. P. Dickerson, and V. Conitzer. Adapting a kidney exchange algorithm to align with human values. *AI*, page 103261, 2020.

- [Goldman and Procaccia, 2014] J. Goldman and A. D. Procaccia. Spliddit: Unleashing fair division algorithms. 13(2):41–46, 2014.
- [Kagel and Roth, 1995] J. H. Kagel and A. E. Roth. *The Handbook of Experimental Economics*. Princeton University, 1995.
- [Kamishima, 2003] T. Kamishima. Nantonac collaborative filtering: recommendation based on order responses. In *Proc. 9th SIGKDD*, pages 583–588, 2003.
- [Kleinberg *et al.*, 1999] Jon Kleinberg, Yuval Rabani, and Éva Tardos. Fairness in routing and load balancing. In *40th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 568–578. IEEE, 1999.
- [Kurokawa *et al.*, 2015] D. Kurokawa, O. Lev, J. Morgenstern, and A. D. Procaccia. Impartial peer review. In *Proc. 23rd IJCAI*, 2015.
- [Leyton-Brown *et al.*, 2017] K. Leyton-Brown, P. Milgrom, and I. Segal. Economics and computer science of a radio spectrum reallocation. *Proc. Natl. Acad. Sci. USA*, 114(28):7202–7209, 2017.
- [Lian *et al.*, 2018] J. W. Lian, N. Mattei, R. Noble, and T. Walsh. The conference paper assignment problem: Using order weighted averages to assign indivisible goods. In *Proc. 33rd AAAI*, 2018.
- [Loreggia *et al.*, 2018] A. Loreggia, N. Mattei, F. Rossi, and K. B. Venable. Preferences and ethical principles in decision making. In *Proc. 1st AIES*, 2018.
- [Mattei and Walsh, 2013] N. Mattei and T. Walsh. Preflib: A library for preferences, [HTTP://WWW.PREFLIB.ORG](http://www.preflib.org). In *Proc. 3rd ADT*, 2013.
- [Mattei and Walsh, 2017] N. Mattei and T. Walsh. A PREFLIB.ORG Retrospective: Lessons Learned and New Directions. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 15, pages 289–309. AI Access Foundation, 2017.
- [Mattei *et al.*, 2012] N. Mattei, J. Forshee, and J. Goldsmith. An empirical study of voting rules and manipulation with large datasets. In *Proc. 4th COMSOC*, 2012.
- [Mattei *et al.*, 2017] N. Mattei, A. Saffidine, and T. Walsh. Mechanisms for online organ matching. In *Proc. 26th IJCAI*, 2017.
- [Mattei *et al.*, 2018] N. Mattei, A. Saffidine, and T. Walsh. Fairness in deceased organ matching. In *Proc. 1st AIES*, 2018.
- [Mattei *et al.*, 2020] N. Mattei, P. Turrini, and S. Zhydkov. Peer-Nomination: Relaxing exactness for increased accuracy in peer selection. *CoRR*, abs/2004.14939, 2020.
- [Mattei, 2011] N. Mattei. Empirical evaluation of voting rules with strictly ordered preference data. In *Proc. 2nd ADT*, 2011.
- [Meir *et al.*, 2020a] R. Meir, J. Lang, J. Lesca, N. Kaminski, and N. Mattei. A market-inspired bidding scheme for peer review paper assignment. In *Proc. 2nd GAIW Workshop at AAMAS*, 2020.
- [Meir *et al.*, 2020b] Reshef Meir, Kobi Gal, and Maor Tal. Strategic voting in the lab: compromise and leader bias behavior. *Auton. Agents Multi Agent Syst.*, 34(1):31, 2020.
- [Mennle *et al.*, 2015] T. Mennle, M. Weiss, B. Philipp, and S. Seuken. The power of local manipulation strategies in assignment mechanisms. In *Proc. 24th IJCAI*, 2015.
- [Moulin, 2004] Hervé Moulin. *Fair division and collective welfare*. MIT press, 2004.
- [Noothigattu *et al.*, 2017] R. Noothigattu, S. Gaikwad, E. Awad, S. Dsouza, I. Rahwan, P. Ravikumar, and A. D. Procaccia. A voting-based system for ethical decision making. In *Proc. 32nd AAAI*, 2017.
- [Plassmann and Tideman, 2014] F. Plassmann and T. N. Tideman. How frequently do different voting rules encounter voting paradoxes in three-candidate elections? *Social Choice and Welfare*, 42(1):31–75, 2014.
- [Popova *et al.*, 2013] A. Popova, M. Regenwetter, and N. Mattei. A behavioral perspective on social choice. *AMAI*, 68(1–3):135–160, 2013.
- [Rawls, 1971] J. Rawls. *A Theory of Justice*. Harvard University Press, 1971.
- [Regenwetter *et al.*, 2007] M. Regenwetter, M.-H. R. Ho, and I. Tsetlin. Sophisticated approval voting, ignorance priors, and plurality heuristics: A behavioral social choice analysis in a Thurstonian framework. *Psychological Review*, 114(4):994–1014, 2007.
- [Rossi and Mattei, 2019] F. Rossi and N. Mattei. Building ethically bounded AI. In *Proc. 33rd AAAI*, 2019.
- [Roth *et al.*, 2005] A. E. Roth, T. Sönmez, and M. U. Ünver. Pairwise kidney exchange. *Journal of Economic theory*, 125(2):151–188, 2005.
- [Scheuerman *et al.*, 2019] J. Scheuerman, J. L. Harman, N. Mattei, and K. B. Venable. Heuristic strategies in uncertain approval voting environments. *arXiv preprint arXiv:1912.00011*, 2019.
- [Sen, 2018] Amartya Sen. *Collective Choice and Social Welfare*. Harvard University Press, 2018.
- [Shah *et al.*, 2018] N. B. Shah, B. Tabibian, K. Muandet, I. Guyon, and U. Von Luxburg. Design and analysis of the NIPS 2016 review process. *JMLR*, 19(1):1913–1946, 2018.
- [Shoham and Leyton-Brown, 2008] Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Game-theoretic, and Logical Foundations*. Cambridge University Press, 2008.
- [Skowron *et al.*, 2015] P. Skowron, P. Faliszewski, and A. Slinko. Achieving fully proportional representation: Approximability results. *AI*, 222:67–103, 2015.
- [Szufa *et al.*, 2020] S. Szufa, P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Drawing a map of elections in the space of statistical cultures. In *Proc. 19th AAMAS*, pages 1341–1349, 2020.
- [Tambe, 2012] M. Tambe. *Security and Game Theory - Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press, 2012.
- [Thomson, 2011] W. Thomson. Fair allocation rules. In *Handbook of Social Choice and Welfare*, volume 2, pages 393–506. Elsevier, 2011.
- [Walsh, 2011] T. Walsh. Where are the hard manipulation problems? *JAIR*, 42:1–29, 2011.
- [Wang and Shah, 2019] J. Wang and N. B. Shah. Your 2 is my 1, your 3 is my 9: Handling arbitrary miscalibrations in ratings. In *Proc. 18th AAMAS*, pages 864–872, 2019.
- [Wooldridge, 2009] M. J. Wooldridge. *An Introduction to Multi-Agent Systems, Second Edition*. Wiley, 2009.
- [Xu *et al.*, 2018] Y. Xu, H. Zhao, X. Shi, and N. B. Shah. On strategyproof conference peer review. *arXiv preprint arXiv:1806.06266*, 2018.
- [Zhao *et al.*, 2018] Z. Zhao, H. Li, J. Wang, J. O. Kephart, N. Mattei, H. Su, and L. Xia. A cost-effective framework for preference elicitation and aggregation. In *Proc. 34th UAI*, pages 446–456, 2018.
- [Zou *et al.*, 2015] J. Zou, R. Meir, and D. Parkes. Strategic voting behavior in Doodle polls. In *Proc. 18th CSCW*, pages 464–472, 2015.