Research Statement: Theory and Data for Better Decisions

Nicholas Mattei
IBM Research AI
nsmattei@gmail.com
www.nickmattei.net

My research broadly centers around the theory and practice of artificial intelligence (AI) and its applications; I am enticed by problems that require a blend of techniques to develop systems and algorithms that support decision making for autonomous agents and/or humans. ***My research vision is to use both theory and experiments to create novel algorithms, mechanisms, and systems that enable and support individual and group decision making.*** My research has advanced the state of the art in algorithms and complexity; game theory and computational social choice; data analytics, data mining, and machine learning; and preference reasoning.

One of the most enjoyable aspects of being a computer scientist is the myriad of settings where we can apply our mathematical and computational tools. Decisions are studied in numerous contexts, offering ample opportunities for interdisciplinary work. I have worked with a broad set of collaborators to find practical applications of my theoretical research in preferences, game theory, and social choice. ***This has led to projects, publications, patents, and relationships with researchers in a number of fields including computer science, economics, psychology, biology, education, and the humanities at a variety of universities as well as business units at IBM and Data61/NICTA.*** At heart I have an engineering mindset, I want to discover new theory and novel techniques in order to build systems that make the world a better, more interesting place by advancing the state of the art in intelligent systems.

My most enjoyable and successful projects are a blend of theoretical innovations, experimental analysis, and practical tool building. I feel that choice and preference research within computer science is walking the same road that Kagel and Roth [31] describe for experimental economics: evolving from theory, to simulated or re-purposed data, to full fledged laboratory and field experiments. This progression enabled a "conversation" to happen between experimental and theoreticians researchers that in turn significantly advanced the field. My work is focused on moving computer science research in decision making down the path described by Roth. In this research statement I summarize some of the main thrusts of my research and describe what I see as their future directions. For a comprehensive listing of my research activity please see www.nickmattei.net.

Social Choice and Preference Reasoning. The Internet enables computers and, by proxy, humans to communicate at distances and speeds previously unimaginable. Technology connects more decision makers (agents) into groups composed of human agents, computer agents, or a mix of the two. These groups of agents must make collective decisions subject to external and internal constraints and preferences in many important real-world settings including: selecting leaders, kidney exchanges, matching students to seats in schools, and allocating work or resources, e.g., [1, 26]. In all of these settings, self-interested agents submit their preferences to a centralized or de-centralized authority and outcomes are decided by a mechanism. Each mechanism for group decision making may or may not satisfy various important criteria, e.g., fairness and/or efficiency. The study of choice and mechanisms falls within the theory and AI related fields of reasoning, decision making, mechanism design, and computational social choice.

In collective decision making researchers use game theory [3] to study how and when agents can strategically misreport their preferences, i.e, when they can manipulate. The study of manipulation is about incentives and security: participants in an aggregation procedure should be incentivized to report the truth and/or be unwilling (computationally) or unable (axiomatically) to find a beneficial misreport. Current research in collective decision making makes strong assumptions about the ability or information provided to the agents, e.g., complete information, or focus purely on worst-case analysis, e.g., NP-hardness. This provides a limited view into many real-world settings; finding a beneficial misreporting is often easy given complete information and strict preferences. Researchers in economics have shown where the predictions of game theory are contradicted by data or experiment; giving rise to behavioral and experimental economics [31, 27]. Until recently there has not been a data driven research program in computers science that directly questions these strong assumptions in collective decision making [15, 26].

Establishing Empirical Research in Social Choice. During my PhD I was struck by the lack of empirical research in the computational social choice community. Coming directly from working as an aerospace research engineer at NASA, I wanted to know how these models worked in practice, on real-world data. To this end I performed empirical studies of voting systems, mining large sets of real-world preference data to create elections and group decision settings and found that many common assumptions about worst-case behavior, domain restrictions, and preferential paradoxes, did not occur in practice [34, 37]. With Prof. Toby Walsh I continued this work, expanding computational social choice research to include empirical testing. ***My work to promote and facilitate the use of data in computational social choice and preference reasoning has led to multiple publications and the establishment of a data repository at* PREFLIB.ORG *and* PREFLIB TOOLS *for working with that data, downloaded over thousands of times to date [36, 35].*** I have founded and organized the EXPLORE workshop series at the Autonomous Agents and Multiagent Systems (AAMAS) conference which has run continuously and with strong attendance since 2013. I will expand the quality and quantity of data and tools PREFLIB, and to use this data to drive research testing old assumptions and verifying new mechanisms.

More than just repurposing existing data, researchers in social choice should look to conduct experiments and gather data about how people really make decisions. To advance this line of work I have teamed up with psychologist Prof. Michel Regenwetter due to his work establishing behavioral social choice [48]. This has lead to an ongoing research collaboration with Prof. Regenwetter and his team on the intersection of social choice and preference modeling [45]. We have recently completed laboratory experiments and are finalizing our data analysis for a human subjects experiment to rigorously test CP-nets, a popular formalism that allows for compactly modeling conditional dependencies [18]. In parallel work on CP-nets we have created a novel software tool to generate CP-nets uniformly at random [6, 4] to facilitate the testing and verification of the data from the human subjects experiments. Provably generating data uniformly at random is a delicate process for CP-nets, requiring exact counting of combinatorial structures; many of the previous papers that claimed to be testing algorithms uniformly at random are provably incorrect and were using biased samples, which could cause misleading results. In a parallel line of work with colleagues from Kentucky and the University of Padova, we realized that CP-nets are restrictive on account of their rigid determinism; we wanted to make them more expressive. In a series of papers we generalized CP-nets, adding probabilities, to create PCP-nets [19]; providing applications, algorithms, and empirical results for working with this more general framework [21, 20]. ***This work led to multiple top venue publications at AAAI and AAMAS as well as a funded National Science Foundation (NSF) grant to rigorously test CP-nets with real people. We published a perspectives paper on real-world experiments in decision making [5] and are finalizing the analysis of the experimental results for publication. I will continue to work with collaborators outside computer science to rigorously test choice models in the real-world.***

Relaxing Strong Theoretical Assumptions in Collective Decision Making. Theoretical results on bribery and manipulation in social choice can be seen through both the lens of preventing bad behavior through computational barriers but also as a question of optimization. If we have some resources at our disposal, how do we maximize the support of an alternative, how can we best distribute our influence and resource in order to achieve consensus? If we assume agents have unlimited time, how much can they affect the final outcome or the welfare of other agents? Assuming complete and deterministic information makes traditional models extremely limited in their scope of understanding. Agents typically have uncertain, incomplete, or probabilistic information about their own preferences and the preferences of other agents; augmenting existing models with incomplete information, probabilistic uncertainty, and other noisy models of data provides more modeling power and extends the scope of questions that can be investigated. During my PhD I was first introduced to theoretical research in the area which assumed strict, deterministic preferences and certain actions. ***For voting and tournaments*** we expanded the base models in the literature to include uncertainty and probability distributions over preferences and manipulation actions, and adding budget for actions (e.g., for agents trying to buy other agents? vote for money). We showed that adding uncertainty does not always create harder instances [16, 41]. ***Combinatorial domains*** can express a decision problem as a sequence of problems with conditional dependencies. The canonical example is a group of friends deciding what appetizer, main course, desert, and wine should be served for a group meal as my choice of wine may depend on the main course. We generalized base models in the literature to this more complex domain and found that, for most of these problems, strategic agents can easily affect the outcome of the decision process [38, 39].

At Data61/NICTA I continued to work on other fundamental aspects of social choice including tournament fixing, selecting multiple winners, and resource allocation. ***For tournaments***, a long standing open problem in scheduling and social choice is the complexity of determining a seeding, i.e., assignment of agents to leaves of the tournament tree, such that a particular agent can win the tournament. While this problem was known to be hard for certain probabilistic information models, we proved that it is a NP-hard problem, even if the designer knows exactly which teams will win and which will lose [7, 13]. Our result for the complete information model implied all the previous work on probabilistic information models, providing a more general result. ***Selecting a committee of multiple winners*** has its own set of challenges including ensuring a diverse committee and proper representation [25]. We closed multiple standing open problems in this area including the complexity of computing winners under proportional approval voting as well as generalizing the manipulation problem to this domain and settling the complexity of manipulation for repeated approval voting and satisfaction approval voting [8]. ***Resource allocation*** is a fundamental social choice task: given disjoint sets of agents and objects, assign the set of objects to the set of agents. In a series of papers we closed key theoretical questions surrounding the *probabilistic serial mechanism* [17] including the complexity of manipulation for the most general case as well as the existence and complexity of computing equilibria; these theoretical results were complemented by extensive experimental results, using preference distributions and real-world data from PrefLib, which showed that manipulation had minimal impact in terms of the price of anarchy and individual welfare loss [9, 10]. ***Along with many co-authors I have worked to advance the understanding of voting and decision making mechanisms, moving beyond the traditional assumptions of single decisions, single winners, strict orders, and perfect information. This work has led to multiple papers at top conferences including AAAI, IJCAI, and AAMAS as well as journal publications at JAIR and AIJ.*** For this research I co-authored a funded NSF proposal, won a best student paper award, received funding for extended international research visits, and worked with many fantastic and interesting people. I will push ahead with this thread of research with an eye to extend the scope of our theoretical understanding of individual and group decision making.

Novel Systems and Mechanisms for Individual and Group Decision Making. In addition to analyzing existing mechanisms, in partnership with government and industrial clients at Data61/NICTA and now IBM, my research has expanded into the areas of mechanism design [44] and decision support systems [46] with the explicit goal of creating mechanisms, algorithms, and systems with superior properties *in both theory and practice* for individual and collective decision making. My work in these areas has spanned a range specific settings including voting, peer selection, resource allocation, scheduling, and planning. During my PhD we developed a system that automatically generates natural language explanations for Markov decision processes (MDPs), a computational formalism for modeling and reasoning in probabilistic domains [47]. We applied our novel explanation methods to academic advising in a university; allowing us access to large datasets and domain experts. To validate our approach, we constructed and administered surveys, in consultation with faculty from psychology, to over 200 students and 15 domain experts. This robust study led to academic papers, models used at the International Planning Competition, and a set of best practices for advising undergraduate students [23, 24, 29, 40]. In addition we devised a family of ***novel voting procedures*** that is a generalization of scoring rules to include a tuning parameter, expressed as an order weighted average (OWA) [51] to reduce opportunities for manipulation. We studied the properties of these new rules in theory and leveraged PrefLib to show that on a variety of synthetic and real-world data, this tuning parameter worked as expected. Combining theory and data to deliver results with greater impact [28].

At Data61/NICTA I continued to work on developing novel systems and algorithms to support decision making in domains including peer selection, cost allocation, and scheduling. In the ***peer selection problem*** 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; crowdsourcing corporate or internal brainstorming sessions; and performing peer review for MOOCs. We devised a novel mechanism called *ExactDollarPartition* for this setting which is strategyproof and has better worst-case bounds than any other mechanism in the literature. We performed an extensive evaluation of all mechanisms in this domain, finding that our method is superior across a number of key metrics data derived from the target domain of NSF peer review [12, 14]. With an external client we leveraged theoretical results on ***cost sharing***, developed in cooperative game theory, and applied them to real-world routing and logistics problems. We proposed and deployed algorithms for novel approximations of the Shapley Value [50], a popular though computationally expensive metric, to allocate in large, industrial cost sharing problems where the underlying cost function is a solution to an NP-hard routing problem. This led to a successful integration with business team tools, allowing for greater analytic power than had been previous achievable [11]. In conjunction with the Disaster Recovery Group at Data61/NICTA and a number of students we developed a novel model and algorithms for what we called ***interdependent scheduling games***. After a natural disaster, multiple major stakeholders must coordinate repairs to interdependent networks in order to restore services, e.g., repairs to components of a gas network may need to be completed before the energy grid can be brought back online. We studied this setting theoretically and provided novel algorithms which we tested empirically on synthetic data [2].

At IBM I have continued to explore how we can use both theory and data to deliver better algorithms and systems for decision making across multiple domains. With the Australian Organ and Tissue Authority we have theoretically and empirically analyzed the efficiency and fairness tradeoffs of proposed ***online organ allocation mechanisms***. Though kidney exchanges [49, 22] have received generous attention in recent years, in many places deceased organ allocation still accounts for over 90% of transplants. With our partners we have analyzed both existing and proposed mechanisms for the online setting where both donors and patients arrive in an online manner [42, 43]. Using data from real-world conferences we were inspired to investigate novel algorithms for the ***conference paper assignment problem***. Expanding our earlier work with OWAs, we proposed novel algorithms which provide the conference organizer more latitude balance between utilitarian maximization and egalitarian fairness. We provide novel polynomial time algorithms for solving this problem and analyze the fairness and efficiency of these methods using real-world data from PrefLib [32]. As part of both internal development and client projects at IBM I have devised a number of novel modeling and decision making algorithms which are currently moving through the patent and publication pipeline. These research projects have given me a window into new domains of interest and I hope to publish these results soon. ***My work on designing novel mechanisms has led to multiple publications at top conferences and journals as well as two patents. We have provided theoretical analysis of these novel algorithms and leveraged the real-world data from PrefLib to demonstrate impact with data.*** The code developed for many of these projects is incorporated into PrefLib Tools, available free and open source. I will continue to use theory, informed by data, to develop novel mechanisms and systems for collective and individual decision making.

Research Directions and Future Work. Leveraging my work on PrefLib and the EXPLORE workshop series there is now more data used in computational social choice and more research using real-world data to study aggregation and decision mechanisms — I will continue to champion this direction. Using data and experiment to augment theoretical developments in individual and collective choice will enable more real-world impact. ***My future research will focus on measuring and identifying actual agent behavior, learned from real-world data, in aggregation and decision mechanisms; using this information to design, implement, and test new mechanisms that are more robust and more desirable.*** Leveraging PrefLib and the internet it is now possible learn and discover novel preference formalisms and domain restrictions from data. I plan to employ advances in machine learning along with best practices from experimental economics [31] and online platforms such as Mechanical Turk [33] accomplish this task. The design and development of tools to gather data and learn models of agent behavior in decision making tasks [30] is a rich and underexploited thread of research. With these novel models and additional data I can continue to create systems that support both humans and computers in the task of making collective decisions.

I have begun work with teams at IBM to create models, algorithms, and systems to support interactive group decision making. In many tasks, like hiring and proposal reviewing, agents iteratively narrow the list of winners to arrive at a final decision. Currently there are not good models for this sequential task. I will work to create these models and proposed efficient elicitation methods and that work in concert with novel, provably robust algorithms to more fully support the ways in which we make collective decisions in the real world. We have also begun work with information retrieval and knowledge extraction teams to create agents that provide information and preferences from structured and unstructured sources and participate in the collective decision making process. Developing these agents, which can augment augment human decision makers in a collaborative environment, is an exciting challenge.

There are opportunities for students at the undergraduate, masters, and PhD level to engage this research. I have numerous successful research collaborations with undergraduate and graduate students as well as summer interns at industrial labs. I have coauthored multiple funded grant proposals with experienced principal investigators, providing experience and a broad base for interdisciplinary work. Work in group decision making and decision support is popular with funding agencies and industry. Advancing preference handling and reasoning from a theoretical science to an experimental science is one of my long term goals. It is an important step for the field: deploying our algorithms and evaluating them in the wild through the use of human subjects experiments and real-world data. Harnessing the power of huge datasets and clever algorithms we can learn more about how humans reason and we can use computers to improve this reasoning process. My hope is that this will lead to better, more principled decision making and a better world.

References

[1] A. Abdulkadiroğlu, P. Pathak, and A. E. Roth. The New York City high school match. *American Economic Review*, 95(2):364–367, 2005.

[2] A. Abeliuk, H. Aziz, G. Berbeglia, S. Gaspers, P. Kalina, **N. Mattei**, D. Peters, P. Stursberg, P. Van Hentenryck, and T. Walsh. Interdependent scheduling games. In *Proc. 25th IJCAI*, pages 2–9, 2016.

[3] R. K. Ahuja, T. L. Magnanti, and J. B.Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.

[4] T. E. Allen, J. Goldsmith, H. E. Justice, **N. Mattei**, and K. Raines. Uniform random generation and dominance testing for CP-Nets. *JAIR*, 59: 771–813, 2017.

[5] T.E. Allen, M. Chen, J. Goldsmith, **N. Mattei**, A. Popova, M. Regenwetter, F. Rossi, and C. Zwilling. Beyond theory and data in preference modeling: Bringing humans into the loop. In *Proc. 4th ADT*, 2015.

[6] T.E. Allen, J. Goldsmith, H.E. Justice, **N. Mattei**, and K. Raines. Generating CP-nets uniformly at random. In *Proc. 30th AAAI*, 2016.

[7] H. Aziz, S. Gaspers, S. Mackenzie, **N. Mattei**, P. Stursberg, and T. Walsh. Fixing a balanced knockout tournament. In *Proc. 28th AAAI*, 2014.

[8] H. Aziz, S. Gaspers, J. Gudmundsson, S. Mackenzie, **N. Mattei**, and T. Walsh. Computational aspects of multi-winner approval voting. In *Proc. 14th AAMAS*, 2015.

[9] H. Aziz, S. Gaspers, S. Mackenzie, **N. Mattei**, N. Narodytska, and T. Walsh. Manipulating the probabilistic serial rule. In *Proc. 14th AAMAS*, 2015.

[10] H. Aziz, S. Gaspers, S. Mackenzie, **N. Mattei**, N. Narodytska, and T. Walsh. Equilibria under the probabilistic serial rule. In *Proc. 24th IJCAI*, 2015.

[11] 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.

[12] H. Aziz, O. Lev, **N. Mattei**, J. S. Rosenschein, and T. Walsh. Strategyproof peer selection: Mechanisms, analyses, and experiments. In *Proc. 30th AAAI*, 2016.

[13] H. Aziz, S. Gaspers, S. Mackenzie, **N. Mattei**, P. Stursberg, and T. Walsh. Fixing a balanced knockout and double elimination tournaments. *Accepted with Revisions at AI*, 2018.

[14] H. Aziz, O. Lev, **N. Mattei**, J. S. Rosenschein, and T. Walsh. Strategyproof peer selection using randomization, partitioning, and apportionment. *Accepted with Revisisons at AI*, 2018.

[15] J.J. Bartholdi, III, C.A. Tovey, and M.A. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3): 227–241, 1989.

[16] D. Binkele-Raible, G. Erdélyi, H. Fernau, J. Goldsmith, **N. Mattei**, and J. Rothe. The complexity of probabilistic lobbying. *Discrete Opt.*, 11 (1):1–21, 2014.

[17] A. Bogomolnaia and H. Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295–328, 2001.

[18] C. Boutilier, R.I. Brafman, C. Domshlak, H.H. Hoos, and D. Poole. CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *JAIR*, 21:135–191, 2004.

[19] C. Cornelio, J. Goldsmith, **N. Mattei**, F. Rossi, and K.B. Venable. Updates and uncertainty in CP-nets. In *Proc. 26th AUSAI*, 2013.

[20] C. Cornelio, U. Grandi, J. Goldsmith, **N. Mattei**, F. Rossi, and K.B. Venable. Voting with CP-nets using a probabilistic preference structure. In *Proc. 5th COMSOC*, 2014.

[21] C. Cornelio, U. Grandi, J. Goldsmith, **N. Mattei**, F. Rossi, and K.B. Venable. Reasoning with PCP-nets in a multi-agent context. In *Proc. 14th AAMAS*, 2015.

[22] 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.

[23] T. Dodson, **N. Mattei**, and J. Goldsmith. A natural language argumentation interface for explanation generation in Markov decision processes. In *Proc. 2nd ADT*, 2011.

[24] T. Dodson, **N. Mattei**, J. T. Guerin, and J. Goldsmith. An English-language argumentation interface for explanation generation with Markov decision processes in the domain of academic advising. *ACM: TiiS*, 3(3):1–30, 2013.

[25] E. Elkind, P., Faliszewski, P. Skowron, and A. Slinko. Properties of multiwinner voting rules. In *Proc. 13th AAMAS*, pages 53–60, 2014.

[26] P. Faliszewski and A. D. Procaccia. AI's war on manipulation: Are we winning? *AI Magazine*, 31(4):53–64, 2010.

[27] J. K. Goeree and C. A. Holt. Ten little treasures of game theory and ten intuitive contradictions. *American Economic Review*, pages 1402–1422, 2001.

[28] J. Goldsmith, J. Lang, **N. Mattei**, and P. Perny. Voting with rank dependent scoring rules. In *Proc. 28th AAAI*, 2014.

[29] J. Guerin, J. Hanna, L. Knouse, **N. Mattei**, and J. Goldsmith. The academic advising planning domain. In *3rd International Planning Competition Workshop at the 22nd International Conference on Automated Planning and Scheduling (ICAPS)*, 2012.

[30] J. S. Hartford, J. R. Wright, and K. Leyton-Brown. Deep learning for predicting human strategic behavior. In *Advances in Neural Information Processing Systems (NIPS)*, pages 2424–2432, 2016.

[31] J. H. Kagel and A. E Roth. *The Handbook of Experimental Economics*. Princeton University, 1995.

[32] 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. 32nd AAAI*, 2018.

[33] W. Mason and S. Suri. Conducting behavioral research on Amazon's Mechanical Turk. *Behavior Research Methods*, 44(1):1–23, 2012.

[34] **N. Mattei**. Empirical evaluation of voting rules with strictly ordered preference data. In *Proc. 2nd ADT*, 2011.

[35] **N. Mattei** and T. Walsh. PrefLib: A library for preferences, HTTP://WWW.PREFLIB.ORG. In *Proc. 3rd ADT*, 2013.

[36] **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.

[37] **N. Mattei**, J. Forshee, and J. Goldsmith. An empirical study of voting rules and manipulation with large datasets. In *Proc. 4th COMSOC*, 2012.

[38] **N. Mattei**, M. S. Pini, F. Rossi, and K. B. Venable. Bribery in voting over combinatorial domains is easy. In *Proc. 11th AAMAS*, 2012.

[39] **N. Mattei**, M. S. Pini, F. Rossi, and K. B. Venable. Bribery in voting with CP-nets. *AMAI*, 68(1–3):135–160, 2013.

[40] **N. Mattei**, T. Dodson, J. T. Guerin, J. Goldsmith, and J. M. Mazur. Lessons learned from development of a software tool to support academic advising. In *American Society for Engineering Education (ASEE) Zone 1 Conference*, 2014.

[41] **N. Mattei**, J. Goldsmith, A. Klapper, and M. Mundhenk. On the complexity of bribery and manipulation in tournaments with uncertain information. *J. of Applied Logic*, 13(4):557–581, 2015.

[42] **N. Mattei**, A. Saffidine, and T. Walsh. Mechanisms for online organ matching. In *Proc. 26th IJCAI*, 2017.

[43] **N. Mattei**, A. Saffidine, and T. Walsh. Fairness in deceased organ matching. Manuscript, 2017.

[44] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press Cambridge, 2007.

[45] A. Popova, M. Regenwetter, and **N. Mattei**. A behavioral perspective on social choice. *AMAI*, 68(1–3):135–160, 2013.

[46] D. J. Power and R. Sharda. Decision Support Systems. In S. Nof, editor, *Springer Handbook of Automation*, pages 1539–1548. Springer, 2009.

[47] M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 2014.

[48] M. Regenwetter, B. Grofman, A. A. J. Marley, and I. M. Tsetlin. *Behavioral Social Choice: Probabilistic Models, Statistical Inference, and Applications*. Cambridge Univ. Press, 2006.

[49] A. E. Roth, T. Sönmez, and M. U. Ünver. Pairwise kidney exchange. *Journal of Economic theory*, 125(2):151–188, 2005.

[50] L. S. Shapley. A value for n-person games. In H.W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games*, volume 2 of *Annals of Mathematical Studies*. Princeton University Press, 1953.

[51] R.R. Yager. On ordered weighted averaging aggregation operators in multicriteria decisionmaking. *IEEE Transactions on Systems, Man and Cybernetics*, 18(1):183–190, 1988.