

Voting by Axioms (Extended Abstract)*

Marie Christin Schmidtlein and Ulle Endriss

Institute for Logic, Language and Computation (ILLC), University of Amsterdam

Abstract


We develop an approach for collective decision making from first principles. In this approach, rather than using a—necessarily imperfect—voting rule to map any given scenario where individual agents report their preferences into a collective decision, we identify for every concrete such scenario the most appealing set of normative principles (known as axioms in social choice theory) that would entail a unique decision and then implement that decision. We analyse some of the fundamental properties of this new approach, from both an algorithmic and a normative point of view.

1 Introduction

There is a well-known mismatch between, on the one hand, seminal results in social choice theory—the principled study of decision making in groups—saying that it is essentially impossible to design an adequate rule for mapping the preferences of individuals into a collective decision [Arrow *et al.*, 2002; Brandt *et al.*, 2016] and, on the other hand, our everyday experience of “making things work”, often by using pragmatic methods—such as the infamous plurality rule—we know to be flawed. In fact, this pragmatic approach is not entirely without scientific justification. Results in behavioural social choice have shown that the problematic scenarios ultimately responsible for the mathematically enticing but otherwise discouraging findings of social choice theory are very rare in practice [Regenwetter *et al.*, 2006]. But too often the misguided take-away from this observed mismatch is to throw the baby out with the bathwater and to ignore the deep insights about sound and normatively grounded decision making provided by social choice theory altogether.

Instead, in recent work [Schmidtlein and Endriss, 2023] we have put forward an approach to collective decision making that is grounded in the axiomatic method of social choice theory [Plott, 1976; Thomson, 2001] but that accounts for the fact that it is impossible to design a “perfect” voting rule that will produce a suitable decision for every conceivable profile

of preferences reported by the members of a group. While in classical social choice theory axioms, i.e., formal renderings of normative principles, are used to motivate acceptable voting rules (that can then be applied in any concrete situation we might encounter), in our approach we take decisions from first principles—by appealing *directly* to axioms when proposing a decision in a given situation:



We call this approach “*voting by axioms*”. It is based on the fundamental idea that a set \mathcal{A} of axioms can be said to *force* a given outcome O for a given profile R of preferences if every voting rule F that satisfies all the axioms in \mathcal{A} would produce O when applied to R . When that is the case and when we find \mathcal{A} normatively appealing, then \mathcal{A} provides a perfect justification for choosing O . Of course, often this will not actually be the case. Indeed, the set \mathcal{A} might not force any outcome for R at all. So instead we work with an entire collection of axioms, ranked from most to least desirable. Then, if our favourite set of axioms does not force an outcome for the profile at hand, we can see whether the next best set might do so, and so forth. So we end up taking a decision that is suggested by the best possible set of normative principles available to us that actually speaks to the situation at hand.

While our approach, in principle, is relevant to any kind of decision making scenario, in practice it is most suited to high-stakes situations where a fairly small group of agents need to choose between a fairly small number of alternatives and where we require any decision taken to stand on sound normative grounds. Agents here could be human beings who are assisted by decision support technology implementing our approach, or they could be autonomous software agents acting on behalf of human stake-holders.

In methodological terms, our approach owes much to the development of the axiomatic method in social choice theory [Plott, 1976; Thomson, 2001], starting with the seminal work of Arrow [1951]. More specifically, our approach is inspired by work of Boixel and Endriss [2020] on ex-

*This is an extended abstract of a paper first published in the proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems [Schmidtlein and Endriss, 2023].

plainable decision making and has links to recent work in this area by several authors [Cailloux and Endriss, 2016; Belahcene *et al.*, 2019; Procaccia, 2019; Peters *et al.*, 2020; Boixel and de Haan, 2021; Boixel *et al.*, 2022; Nardi *et al.*, 2022; Suryanarayana *et al.*, 2022]. While contributions in that literature tend to focus on the task of generating human-readable explanations for *why* a given decision is forced, our concern here is more basic and we ask *whether* that decision is forced in the first place. Finally, there are connections to recent work on using SAT solvers to support the generation of proofs for impossibility theorems in social choice theory [Geist and Peters, 2017], an approach pioneered by Tang and Lin [2009], because one can use the same kind of encoding of axioms into propositional logic to develop practical implementations of our approach on top of a SAT solver.

2 The Approach

In this section we outline the approach of voting by axioms.

2.1 Social Choice Fundamentals

While the idea of voting by axioms can, in principle, be applied to any model of social choice [Arrow *et al.*, 2002], here we consider scenarios with a finite universe $N^* = \{1, \dots, n\}$ of potential *agents* and a finite set $X = \{1, \dots, m\}$ of *alternatives*, where at any given time the members of a some *electorate* $N \subseteq N^*$ each express their preferences by providing a full ranking of the elements of X , thereby giving rise to a *profile* of preferences, and we need to choose a nonempty *outcome* $O \subseteq X$ of intuitively “best” alternatives.

A *voting rule* F is a function mapping any such profile to an outcome. Well-known examples include the plurality rule, the Borda rule, and the Copeland rule [Zwicker, 2016].

A central concept in social choice theory are so-called *axioms*. These are normative principles describing properties we would expect to be satisfied by any good procedure taking decisions, such as a voting rule. Well-known examples include ANONYMITY, requiring that all agents are treated the same; PARETO, stipulating that dominated alternatives (i.e., alternatives y for which there is another alternative x that everyone expressing a preference prefers) are not chosen; CONDORCET, saying that whenever an alternative wins all pairwise majority contests only that alternative should be chosen; and REINFORCEMENT, demanding that, whenever the outcomes for two profiles with disjoint electorates have a nonempty intersection, that intersection will be the outcome when the union of both electorates report preferences.

The formal meaning of an axiom A can be fixed by referring to the set of all voting rules that satisfy it, its so-called *interpretation* $\mathbb{I}(A)$ [Boixel and Endriss, 2020]. Similarly, the interpretation $\mathbb{I}(\mathcal{A})$ of a set \mathcal{A} of axioms is the set of voting rules that satisfy all of the axioms in \mathcal{A} . Such a set is *nontrivial* if it is satisfied by at least one rule, i.e., if $\mathbb{I}(\mathcal{A}) \neq \emptyset$.

2.2 Simple Forcing

Suppose we are presented with a profile R and a set \mathcal{A} of axioms on which to base our decision which outcome to choose for R . Sometimes the axioms in \mathcal{A} might allow us to exclude certain alternatives from the outcome (e.g., when PARETO is

applicable), and sometimes \mathcal{A} might even fully determine—or *force*—a specific outcome O . Let us make this precise.

Definition 1. We say that a nontrivial axiom set \mathcal{A} *forces* an outcome O on a given profile R if every voting rule satisfying the axioms in \mathcal{A} would return that outcome:

$$F(R) = O \text{ for all } F \in \mathbb{I}(\mathcal{A}).$$

The case of trivial axiom sets is explicitly excluded from this definition, because a trivial axiom set would vacuously force every conceivable outcome O on every conceivable profile R .

The notion of forcing thus defined satisfies a number of simple structural properties that together demonstrate that it behaves as expected [Schmidlein and Endriss, 2023]:

- Adding an axiom to a set results in *more* profiles with a forced outcome: if $\mathcal{A} \subseteq \mathcal{A}'$ (and both are nontrivial) and if \mathcal{A} forces outcome O on profile R , then so does \mathcal{A}' .
- If an axiom set \mathcal{A} forces an outcome on *every* profile, then \mathcal{A} *characterises* some voting rule F : $\mathbb{I}(\mathcal{A}) = \{F\}$.

2.3 Ranked Forcing

We now generalise the fundamental idea of forcing by working with a ranking of several axiom sets rather than a single such set. A *ranked axiom corpus* is a pair $\langle \mathbb{A}, \succ \rangle$ consisting of a collection \mathbb{A} of axiom sets and a strict linear order \succ declared on this collection \mathbb{A} . We say that $\langle \mathbb{A}, \succ \rangle$ is *nontrivial* if every axiom set \mathcal{A} in \mathbb{A} is nontrivial.

Definition 2. We say that a nontrivial ranked axiom corpus $\langle \mathbb{A}, \succ \rangle$ *forces* an outcome O on a given profile R if at least one $\mathcal{A} \in \mathbb{A}$ forces some outcome on R and if O is the outcome forced by the top-ranked such axiom set in $\mathcal{A} \in \mathbb{A}$.

Thus, for any given profile R , we now look for the highest-ranked axiom set that actually forces an outcome on R .

By the aforementioned result regarding the axiomatic characterisations of voting rules, a corpus \mathbb{A} including an axiom set \mathcal{A} that characterises a unique voting rule is a sufficient (but not a necessary) condition for $\langle \mathbb{A}, \succ \rangle$ forcing an outcome on *every* profile. So by placing a set \mathcal{A} that characterises some rule (which might not be ideal but offers a decent level of quality) at the bottom of the ranked axiom corpus, we can ensure that in every situation some outcome will be forced.

Example 1. Let $\langle \mathbb{A}, \succ \rangle$ be a corpus of the form $\{\text{CONDORCET}\} \succ \mathcal{A}$, where \mathcal{A} is some axiom set characterising the well-known Borda rule [Zwicker, 2016; Young, 1974]. It is easy to see that $\{\text{CONDORCET}\}$ forces an outcome on all profiles that have a *Condorcet winner*, i.e., an alternative that beats all others in pairwise majority contests, and that it does not force an outcome on any other profile. Since $\{\text{CONDORCET}\}$ is the top-ranked set in the corpus, on Condorcet profiles the corpus $\langle \mathbb{A}, \succ \rangle$ forces the singleton containing the Condorcet winner. On all other profiles $\langle \mathbb{A}, \succ \rangle$ forces the same outcome as would be returned by the Borda rule. Overall, we end up with the same form of decision making as has been proposed by Duncan Black back in 1958, who argued that we should choose the Condorcet winner when it exists and otherwise use the Borda rule [Black, 1958]. \triangle

It also is worth pointing out that, by placing an axiom set \mathcal{A} that characterises a unique rule F at the very top of the ranked axiom corpus, voting by axioms reduces to simply applying F to produce outcomes. In this sense our approach can be seen as generalising the classical approach of social choice theory, where we first use axioms to motivate voting rules and then apply those voting rules to concrete profiles.

We have not yet commented on the question where $\langle \mathbb{A}, \succ \rangle$ might come from. We might start out with a large set of candidate axioms and then use some or all of the (nontrivial) subsets of that set to populate \mathbb{A} . But supplying, from scratch, a complete and strict ranking over the sets in \mathbb{A} might be infeasible in practice. Instead, we might provide a ranking \triangleright on single axioms and then *lift* it to a ranking \succ on the sets in \mathbb{A} . There are myriad ways of how to lift an order on objects to an order on sets of objects [Barberà *et al.*, 2004], but one natural option would be to prefer small sets over large sets and to rank sets of the same cardinality lexicographically.

Another way of generating a ranking \succ on a collection \mathbb{A} of axiom sets would be to associate each axiom A with a cost $c(A)$. The cost of an axiom set then would be sum of the costs of its members and we could rank the axiom sets in \mathbb{A} from cheapest to most expensive. In this context, the cost $c(A)$ of an axiom A might reflect the effort of persuading someone to accept the normative principle underlying A .

3 Technical Results

In the full paper we establish a number of technical results regarding the new approach of voting by axioms, regarding both its algorithmic and its axiomatic features [Schmidtlein and Endriss, 2023]. In this section we sketch what these results entail but refer to the full paper for details and proofs.

3.1 Computational Intractability of Forcing

Voting by axioms clearly is a computationally demanding procedure. Given a ranked corpus $\langle \mathbb{A}, \succ \rangle$ of axiom sets and a profile R , we need to find the first axiom set in \mathbb{A} that forces an outcome on R . The operation at the core of the approach is that of checking whether one specific axiom set \mathcal{A} forces an outcome for a given profile R . A formal analysis of the computational complexity of this core problem thus can offer insights into inherent limitations of voting by axioms.

How difficult it is to determine whether \mathcal{A} forces an outcome on R will depend on the axioms involved; for some it might be straightforward and for others very difficult. To be able to offer a precise analysis of the computational difficulty of the problem, we need to be able to talk about the size of the input of the problem, i.e., the combined size of a given axiom set. To this end, we encode axioms using a simple logical language. We use propositional variables of the form $p_{R,x}$, where R is the name of a profile and x is the name of an alternative, to express that for profile R alternative x belongs to the outcome. The models of formulas in this logic are given by voting rules. For instance, rule F satisfies the formula $p_{R,x} \wedge \neg p_{R,y}$ if $x \in F(R)$ but $y \notin F(R)$.

Similar encodings have been used in the literature on computational social choice for a number of different purposes [e.g., Tang and Lin 2009; Cailloux and Endriss 2016].

We can fully describe a voting rule by taking a conjunction over all profiles and alternatives and including either the positive or negative literal, depending on whether or not the alternative is contained in the voting rule’s outcome for the profile in question. Thus, every conceivable axiom A can be expressed in our language by taking the disjunction over all the conjunctions corresponding to voting rules in $\mathbb{I}(A)$.

Example 2. The axiom PARETO can be encoded as follows:

$$\bigwedge_{y \in X} \bigwedge_{x \in X \setminus \{y\}} \bigwedge_{R: \forall i. (x,y) \in R_i} \neg p_{R,y}.$$

Here, the third conjunction operator is intended to range over all profiles R for which it is the case that every agent i who expresses a preference in R ranks x above y . Δ

Now we can define the *size* of an axiom A as the length of its logical encoding, and the *combined size* of an axiom set \mathcal{A} as the sum of the sizes of its members. With this definition in hand, we are now ready to state our complexity result.

Theorem 1. *The problem of deciding whether a given non-trivial axiom set \mathcal{A} forces an outcome for a given profile R is coNP-complete in the combined size of \mathcal{A} .*

So the problem is indeed computationally intractable. This might be unsurprising. In fact, the most interesting aspect of Theorem 1 is that the problem is *only* coNP-complete. For comparison, Boixel and de Haan [2021] have shown that a closely related problem, arising in the context of computing explanations for *why* a given set of axioms forces a given outcome, is hard for Σ_2^P , a complexity class that, under the usual complexity-theoretic assumptions, is located above coNP.

3.2 Practical Operationalisation via SAT Solving

The fact that the problem of deciding whether a given axiom set forces an outcome for a given profile is *only* coNP-complete suggests that it should be possible to tackle it with a SAT solver [Biere *et al.*, 2009], i.e., a tool for answering questions about the satisfiability of a given set of formulas of propositional logic. This opens up opportunities for practically feasible implementations of voting by axioms.

There, by now, is much precedent in the literature on computational social choice for the use of SAT solvers to reason about scenarios of collective decision making. Most such contributions have been concerned with offering computational support for proving impossibility theorems [Geist and Peters, 2017], but Nardi *et al.* [2022] also have used SAT solvers to design a practically viable algorithm for computing axiomatic justifications for election outcomes.

So let us briefly sketch how to reduce the problem of deciding whether a given axiom set \mathcal{A} forces a given outcome O on a given profile R to a query to a SAT solver. Recall that \mathcal{A} forcing O on R means that every voting rule F that satisfies \mathcal{A} returns O when applied to R . In other words, proposing any outcome other than O for R would be inconsistent with the axioms in \mathcal{A} for any well-formed voting rule F . So we can think of the task of proving that O is the right outcome as the task of proving that \mathcal{A} together with the assumptions that F is well-formed (in the sense of always returning a nonempty set of alternatives) and that O is *not* the outcome is logically inconsistent. Using our logical language, we can easily express

all of these requirements as formulas of propositional logic. We can then pass the conjunction of these formulas to a SAT solver and answer the original question of whether \mathcal{A} forces O on R in the affirmative in case the SAT solver should find that conjunction to be unsatisfiable.

3.3 Well-Behavedness of Intraprofile Axioms

Some axioms are particularly “well-behaved” in the context of voting by axioms. Specifically, this is the case for what Fishburn [1973] calls *intraprofile axioms*. Intraprofile axioms have a particularly simple structure in that they only speak about conditions on outcomes “one profile at a time”. In our logical encoding, they can be formulated as conjunctions of conditions that each relate to just one single profile.

Example 3. PARETO is an intraprofile axiom since, whenever the axiom requires alternative y to not be part of the outcome for a profile R due to y being dominated by x , this condition can be verified by considering R in isolation. On the other hand, ANONYMITY is not an intraprofile axiom, because each instance speaks about two different profiles. \triangle

As long as the axiom corpus we are working with is sufficiently rich, so as to guarantee that some outcome will be forced on any given profile, we can think of the procedure of voting by axioms as a voting rule in its own. Indeed, it amounts to a function mapping profiles to outcomes. In general, the axioms occurring within the corpus need not be axioms that are satisfied by that voting rule. But for intraprofile axioms we obtain a natural correspondence between axioms occurring in the corpus and axioms satisfying the rule.

Theorem 2. *For any nontrivial ranked axiom corpus $\langle \mathbb{A}, \succ \rangle$ and any intraprofile axiom A , we can find a voting rule $F \in \mathbb{I}(A)$ that agrees with the outcomes forced by $\langle \mathbb{A}, \succ \rangle$ on every profile R for which A belongs to the top-ranked axiom set in \mathbb{A} that forces an outcome on R .*

Thus, for any intraprofile axiom A , the procedure of voting by axioms across all profiles where A belongs to the top-ranked axiom set forcing an outcome coincides with the behaviour of a voting rule that satisfies that axiom A . The same is not true for arbitrary axioms. Indeed, if an axiom requires the outcome of one profile to be dependent on the outcome of another profile, this dependency might not be preserved when two distinct axiom sets in our corpus end up forcing outcomes on these profiles. In such a situation, the two forced decisions together might not be consistent with the axiom. Theorem 2 holds because this problem does not occur for intraprofile axioms, as they do not allow for such dependencies.

3.4 Characterisation of Induced Rules

Once again, if the ranked axiom corpus $\langle \mathbb{A}, \succ \rangle$ is sufficiently rich to force an outcome on every possible profile, we can think of the operation of ranked forcing as defining a voting rule F . We refer to F as the voting rule *induced* by $\langle \mathbb{A}, \succ \rangle$.

We already pointed out in Section 2.3 that, if a voting rule F is characterised by an axiom set \mathcal{A} and we place \mathcal{A} at the top of a ranked axiom corpus, then running the procedure of voting by axioms is equivalent to applying F . We conclude by presenting a result that further refines this simple insight by establishing that, if characterising axioms are placed high

enough in the ranking of a corpus, then the induced voting rule will be the characterised rule itself. In particular, the induced voting rule will satisfy the characterising axioms.

Theorem 3. *For every axiom set \mathcal{A} that uniquely characterises a voting rule F , and for every nontrivial ranked axiom corpus $\langle \mathbb{A}, \succ \rangle$ for which \mathcal{A} is a subset of the top-ranked axiom set in \mathbb{A} for every profile R , it is the case that the characterised rule F coincides with the rule induced by $\langle \mathbb{A}, \succ \rangle$.*

4 Conclusion

We introduced a novel approach to collective decision making from first principles. Instead of using a—necessarily imperfect—voting rule, we proposed to use axioms to determine and justify outcomes in voting scenarios. By using a collection of multiple axiom sets, this approach allows us to involve many (even mutually inconsistent) axioms in the decision process. At the same time it must be noted that this method is only ever as good as the axioms it uses. The decisions taken will be appropriate only if the axioms are.

We suggested one way of implementing the framework to compute outcomes based on forcing, namely by encoding axioms in a propositional logic and using a SAT solver to determine forced outcomes. Nonetheless, the complexity result we obtained indicate that realising voting by axioms will be a challenging task in practice. We further explained how our approach can be seen as an extension of the classical approach of social choice theory by showing that the procedure of voting by axioms can sometimes be represented by a voting rule. We highlighted two cases in which our procedure enjoys particularly attractive properties.

Future work should be dedicated to making it easier to use the procedure of voting by axioms in practice. One aspect of this would be to develop a formal language for encoding axioms that is more compact and that lends itself more easily to presenting axioms to users in human-readable form. Some steps in this direction have been taken by Boixel and de Haan [2021], and more broadly in the literature on modelling social choice scenarios in mathematical logic [Endriss, 2011]. Another aspect would be to develop heuristics leading to faster algorithms to determine forcing, for instance, along the lines of recent work by Nardi *et al.* [2022]. Related to this point, it also would be interesting to study special classes of axioms, such as the algebraic axioms of Kaminski [2004], for which forcing is easier to determine. Finally, more research is needed on how to support users to construct a corpus of axioms and, specifically, a ranking of sets of axioms drawn from that corpus. The two approaches we sketched, one favouring smaller axiom sets and one assigning costs to axioms, represent just two of many possible directions to explore.

References

- [Arrow *et al.*, 2002] Kenneth J. Arrow, Amartya K. Sen, and Kotaro Suzumura, editors. *Handbook of Social Choice and Welfare*, volume 1. Elsevier, 2002.
- [Arrow, 1951] Kenneth J. Arrow. *Social Choice and Individual Values*. John Wiley & Sons, 1951.

- [Barberà *et al.*, 2004] Salvador Barberà, Walter Bossert, and Prasanta K. Pattanaik. Ranking sets of objects. In Salvador Barberà, Peter J. Hammond, and Christian Seidl, editors, *Handbook of Utility Theory*, pages 893–977. Springer, 2004.
- [Belahcene *et al.*, 2019] Khaled Belahcene, Christophe Labreuche, Nicolas Maudet, Vincent Mousseau, and Wassila Ouerdane. Comparing options with argument schemes powered by cancellation. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI-2019)*, pages 1537–1543, 2019.
- [Biere *et al.*, 2009] Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh, editors. *Handbook of Satisfiability*. Frontiers in Artificial Intelligence and Applications. IOS Press, 2009.
- [Black, 1958] Duncan Black. *The Theory of Committees and Elections*. Cambridge University Press, 1958.
- [Boixel and de Haan, 2021] Arthur Boixel and Ronald de Haan. On the complexity of finding justifications for collective decisions. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI-2021)*, pages 5194–5201, 2021.
- [Boixel and Endriss, 2020] Arthur Boixel and Ulle Endriss. Automated justification of collective decisions via constraint solving. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2020)*, pages 168–176, 2020.
- [Boixel *et al.*, 2022] Arthur Boixel, Ulle Endriss, and Ronald de Haan. A calculus for computing structured justifications for election outcomes. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI-2022)*, pages 4859–4866, 2022.
- [Brandt *et al.*, 2016] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- [Cailloux and Endriss, 2016] Olivier Cailloux and Ulle Endriss. Arguing about voting rules. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2016)*, pages 287–295, 2016.
- [Endriss, 2011] Ulle Endriss. Logic and social choice theory. In Amitabha Gupta and Johan van Benthem, editors, *Logic and Philosophy Today*, volume 2, pages 333–377. College Publications, 2011.
- [Fishburn, 1973] Peter C. Fishburn. *The Theory of Social Choice*. Princeton University Press, 1973.
- [Geist and Peters, 2017] Christian Geist and Dominik Peters. Computer-aided methods for social choice theory. In Ulle Endriss, editor, *Trends in Computational Social Choice*, pages 249–267. AI Access, 2017.
- [Kaminski, 2004] Marek Kaminski. Social choice and information: The informational structure of uniqueness theorems in axiomatic social theories. *Mathematical Social Sciences*, 48:121–138, 2004.
- [Nardi *et al.*, 2022] Oliviero Nardi, Arthur Boixel, and Ulle Endriss. A graph-based algorithm for the automated justification of collective decisions. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2022)*, pages 935–943, 2022.
- [Peters *et al.*, 2020] Dominik Peters, Ariel D. Procaccia, Alexandros Psomas, and Zixin Zhou. Explainable voting. In *Proceedings of the 34th International Conference on Neural Information Processing Systems (NeurIPS-2020)*, volume 33, pages 1525–1534, 2020.
- [Plott, 1976] Charles R. Plott. Axiomatic social choice theory: An overview and interpretation. *American Journal of Political Science*, 20(3):511–596, 1976.
- [Procaccia, 2019] Ariel D. Procaccia. Axioms should explain solutions. In Jean-François Laslier, Hervé Moulin, M. Remzi Sanver, and William S. Zwicker, editors, *The Future of Economic Design*, pages 195–199. Springer, 2019.
- [Regenwetter *et al.*, 2006] Michel Regenwetter, Bernard Grofman, Ilia Tsetlin, and Anthony A. J. Marley. *Behavioral Social Choice: Probabilistic Models, Statistical Inference, and Applications*. Cambridge University Press, 2006.
- [Schmidtlein and Endriss, 2023] Marie Christin Schmidtlein and Ulle Endriss. Voting by axioms. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2023)*, pages 2067–2075, 2023.
- [Suryanarayana *et al.*, 2022] Sharadhi Alape Suryanarayana, David Sarne, and Sarit Kraus. Justifying social-choice mechanism outcome for improving participant satisfaction. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2022)*, pages 1246–1255, 2022.
- [Tang and Lin, 2009] Pingzhong Tang and Fangzhen Lin. Computer-aided proofs of Arrow’s and other impossibility theorems. *Artificial Intelligence*, 173(11):1041–1053, 2009.
- [Thomson, 2001] William Thomson. On the axiomatic method and its recent applications to game theory and resource allocation. *Social Choice and Welfare*, 18(2):327–386, 2001.
- [Young, 1974] H. Peyton Young. An axiomatization of Borda’s rule. *Journal of Economic Theory*, 9(1):43–52, 1974.
- [Zwicker, 2016] William S. Zwicker. Introduction to the theory of voting. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, pages 23–56. Cambridge University Press, 2016.