

Computational Social Choice: The First Ten Years and Beyond

Haris Aziz¹, Felix Brandt², Edith Elkind³, and Piotr Skowron⁴

¹ Data61 and University of New South Wales, Australia

² Technische Universität München, Germany

³ University of Oxford, United Kingdom

⁴ Technische Universität Berlin, Germany

Abstract. Computational social choice is a research area at the intersection of computer science, mathematics, and economics that is concerned with aggregation of preferences of multiple agents. Typical applications include voting, resource allocation, and fair division. This chapter highlights six representative research areas in contemporary computational social choice: restricted preference domains, voting equilibria and iterative voting, multiwinner voting, probabilistic social choice, random assignment, and computer-aided theorem proving.

1 Introduction

Within the past few decades there has been a lively exchange of ideas between computer science, in particular artificial intelligence, algorithms and complexity theory, on the one hand, and economics, in particular game theory and social choice, on the other hand. This exchange goes in both directions, and is largely motivated by the emergence and the growing ubiquity of the Internet, which created a need for concepts concerning social interaction and cooperation provided by economics as well as for the algorithmic tools of computer science.

A recent example of this trend is the formation of the inter-disciplinary research area known as *computational social choice*, which combines ideas, models, and techniques from social choice theory with those of computer science. Social choice theory, which itself is already a multi-disciplinary area with contributions from economics, mathematics, political science, and philosophy, concerns the formal analysis and design of methods for aggregating the preferences of multiple agents. Typical applications include voting, resource allocation, and fair division. Computer science offers several powerful tools such as algorithm design, complexity theory, and communication complexity for analyzing such problems. At the same time, computer science has produced new application areas for social choice such as webpage ranking or collective decision-making in computational multi-agent systems.

In its most general form, social choice theory is concerned with a set of alternatives and a set of agents who possess binary preference relations (typically assumed to be complete and transitive) over the alternatives; a collection of

agents’ preference relations is called a preference profile. Problems of interest then include how to define (and find) a collective choice in form of a set of alternatives, a ranking of alternatives, or a lottery over alternatives that appropriately reflect the agents’ individual preferences. Collective outcomes and aggregation functions that return these outcomes are usually evaluated and compared by verifying whether they satisfy desirable properties, so-called axioms. Classic results in social choice—the most famous of which is certainly Arrow’s impossibility theorem [2]—have shown the incompatibility of certain sets of axioms, or characterized specific aggregation functions in terms of axioms they satisfy.

Most subareas of social choice (e.g., coalition formation, matching markets, and fair division) can be obtained as special cases of the general model described above by imposing structure on the set of alternatives and restricting the domain of preference relations accordingly. For example, in assignment problems, the goal is to find a fair and efficient assignment of objects to agents based on the agents’ preferences over objects. To cast an assignment problem as a social choice problem, we let the set of alternatives be the set of all possible allocations and postulate that agents are indifferent among all allocations in which they receive the same object. This conceptual insight is useful because it sometimes allows the transfer of positive results from superdomains to subdomains and that of negative results from subdomains to superdomains. However, most statements require a specific analysis of the domain in question and are often based on axioms that can only be meaningfully defined within this domain. Moreover, computational statements usually do not carry over from one domain to another due to the different representations.

Initial results in computational social choice focused on the computational complexity of aggregation functions that were proposed in the social choice literature. For example, it was shown that computing Kemeny’s rule, which returns collective consensus rankings and satisfies many desirable axioms, is NP-hard [23], and deciding whether a given alternative is on top of a consensus ranking is Θ_2^P -complete [83]. This and similar hardness results were followed by the analysis of heuristics, approximation algorithms, and fixed-parameter tractable algorithms for these problems. At the same time, new interesting computational problems concerning various ways of manipulating the election outcome were defined and investigated. Here, computational hardness is desirable and meant to serve as a shield against strategic behavior [see, e.g., 71, 69, 54, 70]. The contributions of contemporary computational social choice go far beyond purely algorithmic questions, and it has been claimed that computational social choice has revitalized the entire field of social choice theory. For instance, a recent result in computational social choice resolved a long-standing open problem in the cake cutting literature: Aziz and Mackenzie [6] proposed the first envy-free cake cutting protocol that requires a bounded number of queries and cuts. Even though the number of cuts is astronomically large, this result is surprising because experts believed that no such protocol exists.⁵

⁵ The number of cuts is upper bounded by $n^{n^{n^{n^n}}}$ where n is the number of agents.

Computational social choice is much too broad to be covered in its entirety in this chapter. We will therefore discuss a handful of new exciting research directions within computational social choice that have not been comprehensively addressed so far and that we consider to be particularly promising: restricted preference domains, voting equilibria and iterative voting, multiwinner voting, probabilistic social choice, random assignment, and computer-aided theorem proving.

2 Restricted Preference Domains

It is well known that when we aggregate the preferences of a group of agents by taking a majority vote over each pair of alternatives, we cannot ensure a rational outcome: the collective preference relation may fail to be transitive even if individual preferences are. This observation goes back to Condorcet (1785), and can be seen as the root cause for many impossibility results such as Arrow’s theorem [2].

Black [26] was the first to observe that this issue does not arise if the voters’ preferences are essentially one-dimensional: he defined the domain of *single-peaked preferences* and showed that for preference profiles that belong to this domain, the majority preference relation is necessarily transitive for an odd number of voters; this implies the existence of a Condorcet winner (an alternative that is preferred to every other alternative by a majority of voters). Informally, a preference profile is said to be single-peaked if the alternatives can be ordered on a line so that each voter has a favorite point (peak) on this line and his affinity for the alternatives declines as one moves away from the peak in either direction.

Single-peaked preferences have received a considerable amount of attention from social choice researchers since Black’s pioneering work [e.g., 98]. More recently, it has been observed that restricting attention to such preferences can also simplify many problems in computational social choice. For instance, there are several voting rules that return a Condorcet winner whenever it exists and otherwise have to solve an NP-hard optimization problem. For any such rule computing the winner is easy if the number of voters is odd and their preferences are single-peaked: by Black’s result, we can simply return the Condorcet winner. For some of these rules, further effort yields polynomial-time algorithms for single-peaked profiles with an even number of voters [39]. The single-peaked ordering of the alternatives can also be used as a basis for a dynamic program; intuitively, one proceeds by computing a partial solution for each prefix of the alternative ordering. This approach leads, e.g., to a polynomial-time algorithm for computing the outcome of the Chamberlin–Courant multiwinner voting rule (see Section 4) for single-peaked profiles [24]. In recent years, many other computational social choice problems were shown to become easier for single-peaked preferences: examples include various forms of strategic behavior in elections [see, e.g., 73, 74] and preference elicitation [51]. Some of these easiness results extend to preference profiles that are nearly single-peaked, i.e., can be made

single-peaked by a small number of modifications (such as deleting a few voters or collapsing or swapping a few alternatives) [see, e.g., 55, 56, 74, 131].

In single-peaked profiles, alternatives can be positioned on a line in a way that respects the voters’ preferences. We can also obtain positive algorithmic results if alternatives can be mapped to a tree [134, 107] or a cycle [108]. Another approach is based on ordering voters rather than alternatives; the resulting domain of *single-crossing preferences* also admits efficient algorithms for a number of problems that are otherwise computationally hard [see, e.g., 122, 93].

An interesting class of algorithmic problems associated with restricted preference domains is to determine whether a given profile belongs to a particular domain. For single-peaked and single-crossing preferences, these problems admit efficient algorithms [58] as well as elegant characterizations in terms of forbidden substructures [21, 43]. By contrast, for trees, the complexity depends on whether we are satisfied with any tree, in which case there is an efficient algorithm [127], or whether we want to construct a tree that satisfies additional constraints, in which case the answer depends on the nature of the constraints [107]. A related question is whether voters and alternatives can be embedded into a d -dimensional space so that the preferences are driven by distances: for $d = 1$ the existence of such an embedding can be determined in polynomial time [58], but for $d > 1$ this problem is equivalent to the existential theory of reals (and thus, in particular, NP-hard) [106]. It is also hard to determine whether a preference profile is close to being single-peaked or single-crossing for many distance measures (with some notable exceptions) [56, 44, 68].

A more extensive survey of recent computational results for restricted domains is provided by Elkind et al. [65].

3 Voting Equilibria and Iterative Voting

In many voting scenarios, a voter or a group of voters can alter an election outcome to their benefit by misrepresenting their preferences; indeed, no reasonable voting rule is immune to this problem. As a consequence, understanding the complexity of finding a manipulative vote under various voting rules has been a prominent research topic in computational social choice since the inception of the area (see, e.g., the survey by Conitzer and Walsh [54]). However, the standard setting of voting manipulation assumes that only some of the voters are strategic, and the interests of all strategic voters are aligned.

When *all* voters act strategically, it is natural to assume that their behavior is governed by a game-theoretic solution concept, such as Nash equilibrium. However, it is not easy to identify an appropriate solution concept: voting games rarely admit dominant strategies, and they often have many Nash equilibria. For instance, under plurality voting with at least three voters, the situation where all voters vote for the same alternative is a Nash equilibrium, even if this alternative is universally hated, as no voter can unilaterally change the election outcome.

One can eliminate some of these equilibria by assuming that, in addition to preferences over alternatives, voters also have secondary preferences: e.g., they

may prefer not to lie unless a lie is clearly beneficial (such voters are called *truth-biased*), or they may prefer not to participate at all if their vote cannot influence the election outcome (such voters are called *lazy*). Either assumption eliminates many unrealistic Nash equilibria; the properties of the surviving equilibria and their computational complexity have been investigated by a number of authors [126, 101, 61, 62]. Another useful technique to get rid of many of the unintuitive outcomes is to focus on trembling-hand Nash equilibria [104].

An alternative approach is to move away from the assumption that all voters submit their ballots simultaneously. For instance, one can consider settings where voters submit their ballots one by one; the appropriate solution concept is then subgame-perfect Nash equilibrium [130, 57]. Alternatively, one can consider dynamic mechanisms, where voters take turns changing their ballots in response to the observed outcome, until no voter has an incentive to make a change: this line of work was initiated by Meir et al. [95], who focused on better/best-response dynamics of plurality voting, and has been subsequently extended to other voting rules [see, e.g., 114, 82, 102, 91]. Convergence and complexity of iterative voting depends on whether voters get to observe the full set of current ballots or just some aggregated information about the ballot profile [113, 67, 96], whether voters compute their best responses at each step, or may use other heuristics [82, 103], and whether voters exhibit secondary preferences, such as laziness or truth bias [111]; see the recent survey by Meir [94].

4 Multiwinner Voting

In multiwinner voting, the goal is to select not just a single winner, but a fixed-size set of winners (a *committee*). Multiwinner voting has a diverse set of applications, which include electing parliaments, shortlisting candidates for a job, selecting locations for public facilities, or deciding which products to advertise to customers. As a consequence, there is a wide variety of multiwinner voting rules, and the research challenge is to formulate desirable properties (axioms) for such rules so as to decide which rules are more suitable for each application, as well as to develop efficient algorithms for computing the outputs of such rules.

Multiwinner voting rules can be broadly classified according to their inputs; while most of the research so far focused on rules where voters have to rank the alternatives, there is also a substantial body of work on multiwinner rules that merely ask each voter to indicate which alternatives they approve.

For ranked ballots, an important class of multiwinner rules is that of *committee scoring rules* [64], which can be seen as analogues of the classic single-winner scoring rules. In more detail, under single-winner scoring rules, each voter assigns a certain number of points to each alternative, based on that alternative's position in her ranking. A typical example is the Borda rule: the Borda score assigned by voter v to alternative c is given by $m - j$, where j is the position of c in v 's ranking and m is the number of alternatives; the winner(s) are the alternatives with the highest total Borda score (summed over all voters). Similarly, under committee scoring rules, each voter assigns a score to each *committee*,

based on the set of positions of committee members in her preference order. For instance, the score that voter v assigns to a committee S can be the sum of Borda scores of all members of S , or the Borda score of v 's most preferred member of S ; under the former approach, the winning size- k committee consists of k alternatives with the highest Borda score, and under the latter approach we get the well-known Chamberlin–Courant rule [49].

An alternative approach is based on extending Condorcet's principle to the multiwinner setting. There are several ways to implement this idea: one can directly compare committees and ask for a committee that is preferred to every other committee by a majority of voters, or one can compare committees and individual alternatives, and require that each committee member is preferred to each non-member by a majority of voters [79, 112, 22], or, alternatively, that no 'large' group of voters prefers a non-member of the committee to each committee member [60, 20]. Yet another class of voting rules, which includes the popular single transferable vote rule, is based on iteratively adding alternatives to the committee and removing or reweighting the voters who approve these alternatives. The seemingly simpler setting of multiwinner voting with approval ballots also admits a variety of interesting voting rules; in fact, there are sophisticated approval-based multiwinner rules that date back to 19th century such as those due to Thiele and Phragmén. This area also has connections to the literature on apportionment [45].

Unfortunately, for many appealing multiwinner rules it is NP-hard to find a winning committee [90, 110, 92, 15, 123]. To circumvent these hardness results, researchers have developed approximation algorithms and used techniques of fixed-parameter tractability [see, e.g., 92, 24, 47, 121, 123]; also, it has been shown that many multiwinner voting rules become easier when voters' preferences belong to a restricted domain (see Section 2) [24, 122, 134, 59, 108, 105].

Computational social choice researchers have also contributed to understanding multiwinner voting rules from a normative perspective. For ranked ballots, Elkind et al. [64] put together a list of prominent multiwinner rules, formulated a number of desirable properties of such rules, and determined which of the rules in their list satisfied each property. Interestingly, they observed that an approximation algorithm for an NP-hard voting rule, when interpreted as a voting rule in its own right, may perform better according to these criteria than the rule it was meant to approximate; this is the case, for instance, for the Greedy Monroe rule of Skowron et al. [121]. For approval-based ballots, the *justified representation* axiom and its extensions have been used to explain some features of the important Proportional Approval Voting (PAV) rule [18, 117]. The class of committee scoring rules discussed in the beginning of this section was shown to admit an axiomatic characterization [124] that is reminiscent of Young's famous characterization of single-winner scoring rules [133]; the structure of this class and axiomatic properties of individual rules within this class are a subject of ongoing research [see, e.g., 123, 75, 76, 63, 120, 48]. We refer the reader to the recent survey by Faliszewski et al. [77] for further details on axiomatic and computational properties of multiwinner rules.

5 Probabilistic Social Choice

Randomization plays an important role in social choice theory. It is easily seen that deterministically picking a single winner is at variance with basic fairness principles (for example, when there are two alternatives and two voters such that each voter prefers a different alternative). In the past few years, there has been refreshed interest in voting rules that return probability distributions over alternatives (so-called lotteries) [e.g., 11, 13, 14, 34, 31, 32]. Often, the outcomes of these rules can also be interpreted as fractional shares of the alternatives.

Randomization may provide a way to circumvent classic impossibility results because the design space of probabilistic voting rules is much richer than that of deterministic ones. Since it is impractical to ask voters for their complete preferences over all lotteries, a common approach is to systematically extend preferences over alternatives to (possibly incomplete) preferences over lotteries via so-called preference extensions. There are various sensible preference extensions, which in turn lead to different generalizations of standard properties such as strategyproofness (no voter is better off by misrepresenting his preferences) and Pareto efficiency (no voter can be made better off without making another one worse off). A very influential preference extension is based on first-order stochastic dominance (*SD*). According to this extension, lottery p is preferred to lottery q if and only if, for every alternative x , the probability that p selects an alternative that is at least as good as x is greater or equal than the probability that q selects such an alternative. Equivalently, p is preferred to q if and only if, for every utility function consistent with the preferences over alternatives, p yields at least as much expected utility as q . A series of increasingly difficult theorems has recently culminated in a sweeping computer-aided impossibility, showing that there is no randomized rule that simultaneously satisfies *SD*-efficiency and weak *SD*-strategyproofness [32, see also Section 7].

Of course, every set-valued voting rule can be straightforwardly turned into a randomized rule by returning the uniform lottery over all winners. Randomization, however, allows for more elaborate rules that satisfy properties unmatched by deterministic rules. Particularly noteworthy in this context are *random serial dictatorship* (RSD) and *maximal lotteries* (ML).

RSD is defined by picking a sequence of the voters uniformly at random and then invoking *serial dictatorship* where voters proceed in a sequence, and each voter narrows down the set of alternatives by picking his most preferred alternatives among the ones selected by previous voters. RSD enjoys strong *SD*-strategyproofness, but violates *SD*-efficiency. It is often used in the domain of assignment where it is also referred to as *random priority* (see Section 6). While implementing RSD by uniformly selecting a sequence of agents and then running serial dictatorship is straightforward, it was shown that computing the resulting RSD probabilities is #P-complete [10], but fixed parameter tractable for parameters such as the number of voters or the number of alternatives [7].

ML is defined as the rule that returns all lotteries that are at least as good as any other lottery in a well-defined way. Maximal lotteries can thus be viewed as a probabilistic generalization of the notion of a Condorcet winner. However, in

contrast to deterministic Condorcet winners which often fail to exist, existence of maximal lotteries is guaranteed by the Minimax Theorem. ML satisfies a very strong notion of efficiency (stronger than *SD*-efficiency), but fails to be even weakly *SD*-strategyproof. Maximal lotteries are equivalent to mixed maximin strategies in a symmetric zero-sum game induced by the voters’ preferences and can be computed in polynomial time via linear programming. ML has been characterized as the only randomized voting rule that satisfies two fairly natural consistency conditions and it has been repeatedly recommended for practical use [see 33].⁶

Other computational work on probabilistic social choice deals with establishing hardness of manipulation via randomization [e.g., 53, 100, 129], approximating deterministic voting rules [e.g., 109, 25, 119], defining new randomized rules [e.g., 4, 8], and measuring the worst-case utilitarian performance of randomized voting rules [e.g., 1]. A more comprehensive overview of recent trends in probabilistic social choice is provided by Brandt [36].

6 Random Assignment

Random assignment is concerned with the probabilistic assignment of m objects to n agents. Each agent specifies transitive and complete preferences over the objects, and the goal is to allocate the objects among the agents in a fair, efficient, and strategyproof manner. When fairly assigning indivisible objects, randomization is necessary in order to satisfy agents with identical preferences. Possible applications include assigning dormitories to students, jobs to applicants, rooms to housemates, processor time slots to jobs, parking spaces to employees, offices to workers, kidneys to patients, etc.

For simplicity, it is typically assumed that $m = n$, that each agent has demand for exactly one object (unit demand), and that individual preferences are strict. A *random assignment* is a probability distribution over deterministic assignments and can be represented by a matrix that specifies, for each agent and each object, the probability with which the agent receives the object. The matrix is bistochastic, which means that all row sums and all column sums are equal to 1. A *random assignment rule* maps each preference profile to a random assignment. It is assumed that agents are only concerned about their individual random assignment, given by the corresponding row of the bistochastic matrix. In order to reason about the axiomatic properties of random assignments and random assignment rules, preferences over objects can be extended to preferences over lotteries just as described in Section 5. It is then possible to define efficiency, strategyproofness, and envy-freeness (no agent prefers another agent’s random assignment) based on the *SD* preference extension. Bogomolnaia and Moulin [29] have shown that no random assignment rule satisfies *SD*-efficiency, strong *SD*-strategyproofness and equal treatments of equals (agents with identical preferences receive identical individual random assignments), even when individual

⁶ ML is one of several voting rules used by the online voting tool Pnyx (<https://pnyx.dss.in.tum.de>).

preferences are strict. The tradeoff among these properties is the subject of ongoing research [e.g., 5, 97].

RSD, as described in Section 5, has a particularly natural interpretation in random assignment: a sequence of agents is picked uniformly at random and then one agent after another picks his most preferred of the remaining objects. RSD inherits strategyproofness from the more general voting domain, but still violates *SD*-efficiency and only satisfies a weak notion of *SD*-envy-freeness. The computational properties of RSD mentioned in Section 5 also hold within the domain of assignment [10, 7, 116].

A well-studied alternative to RSD is the probabilistic serial (PS) rule, which is *SD*-efficient, strongly *SD*-envy-free, and weakly *SD*-strategyproof (as long as individual preferences are strict).⁷ Under *PS*, agents ‘eat’ the most preferred available object at an equal rate until all objects are consumed. When a most preferred object is completely consumed, agents eat their next most preferred object that is still available. The fraction of any object consumed by an agent is the probability of the agent receiving that object. There have been a number of appealing axiomatic characterizations of PS using *SD*-efficiency, *SD*-envy-freeness, and additional properties [28, 128, 27]. Furthermore, PS has been extended in a number of ways [e.g., 86, 132, 3, 46, 118]. In particular, there is a natural extension to the more general case of multi-unit demand [29, 89, 84]. Just as in the case of ties in the preferences, weak *SD*-strategyproofness breaks down when allowing multi-unit demand. However, it has been shown that the problem of manipulating PS to increase one’s expected utility is NP-hard [16].

Maximal lotteries, as described in Section 5, are known as *popular* random assignments within the domain of assignment [87, 12]. In contrast to the voting domain, however, popular random assignments are rarely unique and popularity can be seen as a property of random assignment rules rather than a rule by itself. In this sense, popularity is stronger than *SD*-efficiency and it is violated by both RSD and PS. Moreover, popularity has been shown to be incompatible with each of weak *SD*-strategyproofness and weak *SD*-envy-freeness [42]. Kavitha et al. [87] have shown that popular random assignments can be computed in polynomial time via linear programming.

Apart from examining existing rules, the structure and computational complexity of efficiency notions constitutes an interesting research area [14, 17]. There is a close connection between probabilistic assignment of indivisible objects and deterministic allocation of divisible objects [see, e.g., 9]. Other recent work has focused on theoretically and experimentally analyzing the performance of random assignment rules [e.g., 78, 85, 19] and extending the model to allow for other richer features such as incorporating side constraints [46], priorities of objects over agents [88], endowments [3], or optional participation [35].

⁷ No rule satisfies these conditions when ties are allowed in the agents’ preferences [86].

7 Computer-aided Theorem Proving

Due to its rigorous axiomatic foundation and its emphasis on impossibility results, social choice theory is particularly well-suited for computer-aided theorem proving techniques. Apart from work that is directed towards formalizing and verifying existing results [see, e.g., 99, 81], a number of recent papers have proved new theorems with the help of computers [37, 30, 41, 32]. This branch of research was initiated by Tang and Lin [125], who reduced well-known impossibility results such as Arrow’s theorem to finite instances, which could then be checked by SAT solvers.

In more detail, the approach for these proofs usually goes along the following lines:⁸ First, it is manually proven that if there exists a voting rule that satisfies a given set of axioms for $m + 1$ alternatives and $n + 1$ voters, then we can also find a voting rule that satisfies the same set of axioms for m alternatives and n voters. The contrapositive of this statement can serve as an induction step for impossibility theorems: If there is *no* voting rule satisfying some axioms for fixed m and n , then there is no such rule for any larger m and n either. Thus, it suffices to prove the impossibility for fixed—and ideally small— m and n . Checking whether there *exists* a voting rule that satisfies certain axioms even for small m and n can be very difficult and is obviously a much more complex task than checking whether a *given* voting rule satisfies the axioms. Exhaustive search is infeasible because the number of possible voting rules is prohibitively large. These problems are therefore typically tackled using general problem solvers such as SAT (propositional satisfiability), SMT (satisfiability modulo theories), ASP (answer set programming), or MIP (mixed integer programming). In most cases, the axioms are encoded as a propositional formula and a SAT solver is asked whether this formula has a satisfying assignment. If it does, the satisfying assignment can be translated back to a concrete voting rule that satisfies the given axioms. If the formula is unsatisfiable, no such voting rule exists. Many SAT instances are initially computationally infeasible and can only be solved after leveraging insights into the axioms and finding a restricted domain of preference profiles sufficient for the impossibility. A common criticism of computer-aided proving methods is that the verdict of the computer usually stands without human-readable proof. Fortunately, when relying on SAT solving, this criticism can be addressed by extracting a human-readable proof from an inclusion-minimal unsatisfiable set of clauses returned by the SAT solver. This approach, pioneered by Brandt and Geist [37], has been successfully applied in several recent papers [30, 32, 41].

Two representative results in this branch of research are an improved computer-aided proof of Moulin’s *No-Show Paradox* and an impossibility for randomized voting rules mentioned in Section 5. The first proof requires only 12 voters (instead of Moulin’s 25) and this bound is furthermore shown to be tight [41]. The computer proof (unexpectedly) exploits certain automorphisms in preference profiles, which makes the proof easier to verify and arguably more

⁸ This section focuses on voting, but all techniques can be similarly applied to other social choice domains such as assignment or coalition formation.

elegant than Moulin’s proof. The second result shows the incompatibility of *SD*-efficiency and weak *SD*-strategyproofness and strengthens a number of previous impossibilities [32]. Since working with lotteries requires real-valued arithmetic (rather than only Boolean logic), the statement was obtained via an SMT solver. The resulting proof is rather complex and difficult to verify for humans. It was therefore translated back into a proof in higher-order logic, which was in turn verified via the interactive theorem prover Isabelle/HOL.⁹

An important benefit of the described approach is its universality and flexibility. As soon as a problem has been formalized, it is straightforward to adapt individual axioms or alter the encoding so that related problems can be solved, too. For a more comprehensive account of computed-aided theorem proving in social choice theory, the reader is referred to the survey by Geist and Peters [80].

8 Further Reading

There are various excellent sources that extensively cover the existing literature, most notably the recently released *Handbook of Computational Social Choice* [40] and an upcoming book on trends in computational social choice [66]. Further overviews and introductions were provided by Rothe [115], Brandt et al. [38], Conitzer [52], Faliszewski and Procaccia [69], Faliszewski et al. [72], and Chevaleyre et al. [50].

Acknowledgements

Haris Aziz is supported by a Julius Career Award. Felix Brandt is supported by the DFG under grant BR 2312/11-1. Edith Elkind and Piotr Skowron are supported by the ERC under grant 639945 (ACCORD). Piotr Skowron is also supported by a Humboldt Research Fellowship for Postdoctoral Researchers. The authors thank Florian Brandl and Dominik Peters for helpful feedback.

References

1. E. Anshelevich and J. Postl. Randomized social choice functions under metric preferences. In *Proc. of 25th IJCAI*, pages 46–59, 2016.
2. K. J. Arrow. *Social Choice and Individual Values*. New Haven: Cowles Foundation, 1st edition, 1951.
3. S. Athanassoglou and J. Sethuraman. House allocation with fractional endowments. *International Journal of Game Theory*, 40(3):481–513, 2011.
4. H. Aziz. Maximal Recursive Rule: A New Social Decision Scheme. In *Proc. of 23rd IJCAI*, pages 34–40, 2013.
5. H. Aziz and Y. Kasajima. Impossibilities for probabilistic assignment. *Social Choice and Welfare*, 2017. Forthcoming.

⁹ Unfortunately, the theorems considered in this context are much too complex to be proven directly by higher-order theorem provers.

6. H. Aziz and S. Mackenzie. A discrete and bounded envy-free cake cutting protocol for any number of agents. In *Proc. of 57th FOCS*, pages 416–427, 2016.
7. H. Aziz and J. Mestre. Parametrized algorithms for random serial dictatorship. *Mathematical Social Sciences*, 72:1–6, 2014.
8. H. Aziz and P. Stursberg. A generalization of probabilistic serial to randomized social choice. In *Proc. of 28th AAAI*, pages 559–565, 2014.
9. H. Aziz and C. Ye. Cake cutting algorithms for piecewise constant and piecewise uniform valuations. In *Proc. of 10th WINE*, pages 1–14, 2014.
10. H. Aziz, F. Brandt, and M. Brill. The computational complexity of random serial dictatorship. *Economics Letters*, 121(3):341–345, 2013.
11. H. Aziz, F. Brandt, and M. Brill. On the tradeoff between economic efficiency and strategyproofness in randomized social choice. In *Proc. of 12th AAMAS*, pages 455–462, 2013.
12. H. Aziz, F. Brandt, and P. Stursberg. On popular random assignments. In *Proc. of 6th SAGT*, pages 183–194. Springer, 2013.
13. H. Aziz, F. Brandl, and F. Brandt. On the incompatibility of efficiency and strategyproofness in randomized social choice. In *Proc. of 28th AAAI*, pages 545–551, 2014.
14. H. Aziz, F. Brandl, and F. Brandt. Universal Pareto dominance and welfare for plausible utility functions. *Journal of Mathematical Economics*, 60:123–133, 2015.
15. H. Aziz, S. Gaspers, J. Gudmundsson, S. Mackenzie, N. Mattei, and T. Walsh. Computational aspects of multi-winner approval voting. In *Proc. of 14th AAMAS*, pages 107–115, 2015.
16. H. Aziz, S. Gaspers, S. Mackenzie, N. Mattei, N. Narodytska, and T. Walsh. Manipulating the probabilistic serial rule. In *Proc. of 14th AAMAS*, pages 1451–1459, 2015.
17. H. Aziz, S. Mackenzie, L. Xia, and C. Ye. Ex post efficiency of random assignments. In *Proc. of 14th AAMAS*, pages 1639–1640, 2015.
18. H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 2016.
19. H. Aziz, J. Chen, A. Filos-Ratsikas, S. Mackenzie, and N. Mattei. Egalitarianism of random assignment mechanisms. In *Proc. of 15th AAMAS*, 2016.
20. H. Aziz, E. Elkind, P. Faliszewski, M. Lackner, and P. Skowron. The Condorcet principle for multiwinner elections: From shortlisting to proportionality. 2017. Working paper.
21. M. A. Ballester and G. Haeringer. A characterization of the single-peaked domain. *Social Choice and Welfare*, 36(2):305–322, 2011.
22. S. Barberà and D. Coelho. How to choose a non-controversial list with k names. *Social Choice and Welfare*, 31(1):79–96, 2008.
23. 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.

24. N. Betzler, A. Slinko, and J. Uhlmann. On the computation of fully proportional representation. *Journal of AI Research*, 47:475–519, 2013.
25. E. Birrell and R. Pass. Approximately strategy-proof voting. In *Proc. of 22nd IJCAI*, pages 67–72, 2011.
26. D. Black. On the rationale of group decision-making. *Journal of Political Economy*, 56(1):23–34, 1948.
27. A. Bogomolnaia. Random assignment: Redefining the serial rule. *Journal of Economic Theory*, 158:308–318, 2015.
28. A. Bogomolnaia and E. J. Heo. Probabilistic assignment of objects: Characterizing the serial rule. *Journal of Economic Theory*, 147:2072–2082, 2012.
29. A. Bogomolnaia and H. Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295–328, 2001.
30. F. Brandl, F. Brandt, C. Geist, and J. Hofbauer. Strategic abstention based on preference extensions: Positive results and computer-generated impossibilities. In *Proc. of 24th IJCAI*, pages 18–24, 2015.
31. F. Brandl, F. Brandt, and J. Hofbauer. Incentives for participation and abstention in probabilistic social choice. In *Proc. of 14th AAMAS*, pages 1411–1419, 2015.
32. F. Brandl, F. Brandt, and C. Geist. Proving the incompatibility of efficiency and strategyproofness via SMT solving. In *Proc. of 25th IJCAI*, pages 116–122, 2016.
33. F. Brandl, F. Brandt, and H. G. Seedig. Consistent probabilistic social choice. *Econometrica*, 84(5):1839–1880, 2016.
34. F. Brandl, F. Brandt, and W. Suksompong. The impossibility of extending random dictatorship to weak preferences. *Economics Letters*, 141:44–47, 2016.
35. F. Brandl, F. Brandt, and J. Hofbauer. Random assignment with optional participation. In *Proc. of 16th AAMAS*, 2017. Forthcoming.
36. F. Brandt. Rolling the dice: Recent results in probabilistic social choice. In U. Endriss, editor, *Trends in Computational Social Choice*. AI Access, 2017.
37. F. Brandt and C. Geist. Finding strategyproof social choice functions via SAT solving. *Journal of Artificial Intelligence Research*, 55:565–602, 2016.
38. F. Brandt, V. Conitzer, and U. Endriss. Computational social choice. In G. Weiß, editor, *Multiagent Systems*, pages 213–283. MIT Press, 2nd edition, 2013.
39. F. Brandt, M. Brill, E. Hemaspaandra, and L. Hemaspaandra. Bypassing combinatorial protections: Polynomial-time algorithms for single-peaked electorates. *Journal of AI Research*, 53:439–496, 2015.
40. F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
41. F. Brandt, C. Geist, and D. Peters. Optimal bounds for the no-show paradox via SAT solving. *Mathematical Social Sciences*, 2017. Forthcoming.

42. F. Brandt, J. Hofbauer, and M. Suderland. Majority graphs of assignment problems and properties of popular random assignments. In *Proc. of 16th AAMAS*, 2017. Forthcoming.
43. R. Bredereck, J. Chen, and G. J. Woeginger. A characterization of the single-crossing domain. *Social Choice and Welfare*, 41(1):989–998, 2013.
44. R. Bredereck, J. Chen, and G. Woeginger. Are there any nicely structured preference profiles nearby? *Mathematical Social Sciences*, 79:61–73, 2016.
45. M. Brill, J. F. Laslier, and P. Skowron. Multiwinner approval rules as apportionment methods. In *Proc. of 31st AAAI*, 2017. Forthcoming.
46. E. Budish, Y.-K. Che, F. Kojima, and P. Milgrom. Designing random allocation mechanisms: Theory and applications. *American Economic Review*, 103(2):585–623, 2013.
47. J. Byrka and K. Sornat. PTAS for minimax approval voting. In *Proc. of 10th WINE*, pages 203–217, 2014.
48. I. Caragiannis, S. Nath, A. D. Procaccia, and N. Shah. Subset selection via implicit utilitarian voting. In *Proc. of 25th IJCAI*, 2016.
49. J. R. Chamberlin and P. N. Courant. Representative deliberations and representative decisions: Proportional representation and the Borda rule. *The American Political Science Review*, 77(3):718–733, 1983.
50. Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. A short introduction to computational social choice. In *Proc. of 33rd SOFSEM*, pages 51–69. Springer, 2007.
51. V. Conitzer. Eliciting single-peaked preferences using comparison queries. *Journal of AI Research*, 35:161–191, 2009.
52. V. Conitzer. Making decisions based on the preferences of multiple agents. *Communications of the ACM*, 53(3):84–94, 2010.
53. V. Conitzer and T. Sandholm. Universal voting protocol tweaks to make manipulation hard. In *Proc. of 18th IJCAI*, pages 781–788, 2003.
54. V. Conitzer and T. Walsh. Barriers to manipulation in voting. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
55. D. Cornaz, L. Galand, and O. Spaajaard. Bounded single-peaked width and proportional representation. In *Proc. of 20th ECAI*, pages 270–275, 2012.
56. D. Cornaz, L. Galand, and O. Spaajaard. Kemeny elections with bounded single-peaked or single-crossing width. In *Proc. of 23th IJCAI*, pages 76–82, 2013.
57. Y. Desmedt and E. Elkind. Equilibria of plurality voting with abstentions. In *Proc. of 11th ACM-EC*, pages 347–356. ACM Press, 2010.
58. J. Doignon and J. Falmagne. A polynomial time algorithm for unidimensional unfolding representations. *Journal of Algorithms*, 16(2):218–233, 1994.
59. E. Elkind and M. Lackner. Structure in dichotomous preferences. In *Proc. of 24th IJCAI*, pages 2019–2025, 2015.
60. E. Elkind, J. Lang, and A. Saffidine. Condorcet winning sets. *Social Choice and Welfare*, 44(3):493–517, 2015.

61. E. Elkind, E. Markakis, S. Obraztsova, and P. Skowron. Equilibria of plurality voting: Lazy and truth-biased voters. In *Proc. of 8th SAGT*, pages 110–122, 2015.
62. E. Elkind, E. Markakis, S. Obraztsova, and P. Skowron. Complexity of finding equilibria of plurality voting under structured preferences. In *Proc. of 15th AAMAS*, pages 394–401, 2016.
63. E. Elkind, P. Faliszewski, J. Laslier, P. Skowron, A. Slinko, and N. Talmon. What do multiwinner voting rules do? An experiment over the two-dimensional euclidean domain. In *Proc. of 31st AAAI*, 2017. Forthcoming.
64. E. Elkind, P. Faliszewski, P. Skowron, and A. Slinko. Properties of multiwinner voting rules. *Social Choice and Welfare*, 2017. Forthcoming.
65. E. Elkind, M. Lackner, and D. Peters. Structured preferences. In U. Endriss, editor, *Trends in Computational Social Choice*. AI Access, 2017.
66. U. Endriss, editor. *Trends in Computational Social Choice*. AI Access, 2017.
67. U. Endriss, S. Obraztsova, M. Polukarov, and J. S. Rosenschein. Strategic voting with incomplete information. In *Proc. of 25th IJCAI*, pages 236–242, 2016.
68. G. Erdélyi, M. Lackner, and A. Pfandler. Computational aspects of nearly single-peaked electorates. *Journal of AI Research*, 2017. Forthcoming.
69. P. Faliszewski and A. D. Procaccia. AI’s war on manipulation: Are we winning? *AI Magazine*, 31(4):53–64, 2010.
70. P. Faliszewski and J. Rothe. Control and bribery in voting. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
71. P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. A richer understanding of the complexity of election systems. In S. Ravi and S. Shukla, editors, *Fundamental Problems in Computing: Essays in Honor of Professor Daniel J. Rosenkrantz*. Springer, 2009.
72. P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. Using complexity to protect elections. *Communications of the ACM*, 53(11):74–82, 2010.
73. P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. The shield that never was: Societies with single-peaked preferences are more open to manipulation and control. *Information and Computation*, 209(2): 89–107, 2011.
74. P. Faliszewski, E. Hemaspaandra, and L. A. Hemaspaandra. The complexity of manipulative attacks in nearly single-peaked electorates. *Artificial Intelligence*, 207:69–99, 2014.
75. P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Committee scoring rules: Axiomatic classification and hierarchy. In *Proc. of 25th IJCAI*, pages 250–256, 2016.
76. P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Multiwinner analogues of the plurality rule: Axiomatic and algorithmic views. In *Proc. of 30th AAAI*, pages 482–488, 2016.
77. 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*. AI Access, 2017.

78. A. Filos-Ratsikas, S. K. S. Frederiksen, and J. Zhang. Social welfare in one-sided matchings: Random priority and beyond. In *Proc. of 7th International Symposium on Algorithmic Game Theory (SAGT)*, pages 1–12, 2014.
79. W. V. Gehrlein. The Condorcet criterion and committee selection. *Mathematical Social Sciences*, 10(3):199–209, 1985.
80. C. Geist and D. Peters. Computer-aided methods for social choice theory. In U. Endriss, editor, *Trends in Computational Social Choice*. AI Access, 2017.
81. U. Grandi and U. Endriss. First-order logic formalisation of impossibility theorems in preference aggregation. *Journal of Philosophical Logic*, 42(4):595–618, 2013.
82. U. Grandi, A. Loreggia, F. Rossi, K. B. Venable, and T. Walsh. Restricted manipulation in iterative voting: Condorcet efficiency and Borda score. In *Proc. of 3rd ADT*, pages 181–192, 2013.
83. E. Hemaspaandra, H. Spakowski, and J. Vogel. The complexity of Kemeny elections. *Theoretical Computer Science*, 349(3):382–391, 2005.
84. E. J. Heo. Probabilistic assignment with multiple demands: A generalization of the serial rule and its characterization. *Journal of Mathematical Economics*, 54:40–47, 2014.
85. H. Hosseini, K. Larson, and R. Cohen. Investigating the characteristics of one-sided matching mechanisms. In *Proc. of 15th AAMAS*, pages 1443–1444, 2016.
86. A.-K. Katta and J. Sethuraman. A solution to the random assignment problem on the full preference domain. *Journal of Economic Theory*, 131(1):231–250, 2006.
87. T. Kavitha, J. Mestre, and M. Nasre. Popular mixed matchings. *Theoretical Computer Science*, 412(24):2679–2690, 2011.
88. O. Kesten and U. Unver. A theory of school choice lotteries. *Theoretical Economics*, pages 543–595, 2015.
89. F. Kojima. Random assignment of multiple indivisible objects. *Mathematical Social Sciences*, 57(1):134–142, 2009.
90. R. LeGrand, E. Markakis, and A. Mehta. Some results on approximating the minimax solution in approval voting. In *Proc. of 6th AAMAS*, pages 1193–1195, 2007.
91. O. Lev and J. S. Rosenschein. Convergence of iterative scoring rules. *Journal of AI Research*, 57:573–591, 2016.
92. T. Lu and C. Boutilier. Budgeted social choice: From consensus to personalized decision making. In *Proc. of 22nd IJCAI*, pages 280–286, 2011.
93. K. Magiera and P. Faliszewski. How hard is control in single-crossing elections? In *Proc. of 21st ECAI*, pages 579–584, 2014.
94. R. Meir. Iterative voting. In U. Endriss, editor, *Trends in Computational Social Choice*. AI Access, 2017.
95. R. Meir, M. Polukarov, J. S. Rosenschein, and N. R. Jennings. Convergence to equilibria in plurality voting. In *Proc. of 24th AAAI*, pages 823–828, 2010.
96. R. Meir, O. Lev, and J. S. Rosenschein. A local-dominance theory of voting equilibria. In *Proc. of 15th ACM-EC*, pages 313–330, 2014.

97. T. Mennle and S. Seuken. An axiomatic approach to characterizing and relaxing strategyproofness of one-sided matching mechanisms. In *Proc. of 15th ACM-EC*, pages 37–38, 2014.
98. H. Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4):437–455, 1980.
99. T. Nipkow. Social choice theory in HOL: Arrow and Gibbard-Satterthwaite. *Journal of Automated Reasoning*, 43:289–304, 2009.
100. S. Obraztsova and E. Elkind. On the complexity of voting manipulation under randomized tie-breaking. In *Proc. of 22nd IJCAI*, pages 319–324, 2011.
101. S. Obraztsova, E. Markakis, and D. R. M. Thomson. Plurality voting with truth-biased agents. In *Proc. of 6th SAGT*, pages 26–37, 2014.
102. S. Obraztsova, O. Lev, E. Markakis, Z. Rabinovich, and J. S. Rosenschein. Beyond plurality: Truth-bias in binary scoring rules. In *Proc. of 14th AAMAS*, pages 1733–1734, 2015.
103. S. Obraztsova, E. Markakis, M. Polukarov, Z. Rabinovich, and N. R. Jennings. On the convergence of iterative voting: How restrictive should restricted dynamics be? In *Proc. of 29th AAAI*, pages 993–999, 2015.
104. S. Obraztsova, Z. Rabinovich, E. Elkind, M. Polukarov, and N. R. Jennings. Trembling hand equilibria of plurality voting. In *Proc. of 25th IJCAI*, pages 440–446, 2016.
105. D. Peters. Single-peakedness and total unimodularity: Efficiently solve voting problems without even trying. Technical report, arXiv.org, 2017.
106. D. Peters. Recognising multidimensional Euclidean preferences. In *Proc. of 31st AAAI*, 2017.
107. D. Peters and E. Elkind. Preferences single-peaked on nice trees. In *Proc. of 30th AAAI*, pages 594–600, 2016.
108. D. Peters and M. Lackner. Preferences single-peaked on a circle. In *Proc. of 31st AAAI*, 2017. Forthcoming.
109. A. D. Procaccia. Can approximation circumvent Gibbard-Satterthwaite? In *Proc. of 24th AAAI*, pages 836–841, 2010.
110. A. D. Procaccia, J. S. Rosenschein, and A. Zohar. On the complexity of achieving proportional representation. *Social Choice and Welfare*, 30:353–362, 2008.
111. Z. Rabinovich, S. Obraztsova, O. Lev, E. Markakis, and J. S. Rosenschein. Analysis of equilibria in iterative voting schemes. In *Proc. of 29th AAAI*, pages 1007–1013, 2015.
112. T. C. Ratliff. Some startling inconsistencies when electing committees. *Social Choice and Welfare*, 21(3):433–454, 2003.
113. A. Reijngoud and U. Endriss. Voter response to iterated poll information. In *Proc. of 11th AAMAS*, pages 635–644, 2012.
114. R. Reyhani and M. C. Wilson. Best reply dynamics for scoring rules. In *Proc. of 20th ECAI*, pages 672–677, 2012.
115. J. Rothe, editor. *Economics and Computation: An Introduction to Algorithmic Game Theory, Computational Social Choice, and Fair Division*. Springer, 2015.

116. D. Saban and J. Sethuraman. The complexity of computing the random priority allocation matrix. *Mathematics of Operations Research*, 40(4):1005–1014, 2015.
117. L. Sánchez-Fernández, E. Elkind, M. Lackner, N. Fernández, J. A. Fisteus, P. Basanta Val, and P. Skowron. Proportional justified representation. In *Proc. of 31st AAAI*, 2017.
118. L. J. Schulman and V. V. Vazirani. Allocation of divisible goods under lexicographic preferences. In *Proc. of 35th FSTTCS*, pages 543–559, 2015.
119. T. C. Service and J. A. Adams. Strategyproof approximations of distance rationalizable voting rules. In *Proc. of 11th AAMAS*, pages 569–576, 2012.
120. P. Skowron. What do we elect committees for? A voting committee model for multi-winner rules. In *Proc. of 24th IJCAI*, pages 1141–1148, 2015.
121. P. Skowron, P. Faliszewski, and A. Slinko. Achieving fully proportional representation: Approximability result. *Artificial Intelligence*, 222:67–103, 2015.
122. P. Skowron, L. Yu, P. Faliszewski, and E. Elkind. The complexity of fully proportional representation for single-crossing electorates. *Theoretical Computer Science*, 569:43–57, 2015.
123. P. Skowron, P. Faliszewski, and J. Lang. Finding a collective set of items: From proportional multirepresentation to group recommendation. *Artificial Intelligence*, 241:191–216, 2016.
124. P. Skowron, P. Faliszewski, and A. Slinko. Axiomatic characterization of committee scoring rules. In *Proc. of 6th COMSOC*, 2016.
125. P. Tang and F. Lin. Computer-aided proofs of Arrow’s and other impossibility theorems. *Artificial Intelligence*, 173(11):1041–1053, 2009.
126. D. R. M. Thomson, O. Lev, K. Leyton-Brown, and J. Rosenschein. Empirical analysis of plurality election equilibria. In *Proc. of 12th AAMAS*, pages 391–398, 2013.
127. M. A. Trick. Recognizing single-peaked preferences on a tree. *Mathematical Social Sciences*, 17(3):329–334, 1989.
128. M. U. Ünver, O. Kesten, M. Kurino, T. Hashimoto, and D. Hirata. Two axiomatic approaches to the probabilistic serial mechanism. *Theoretical Economics*, 9:253–277, 2014.
129. T. Walsh and L. Xia. Lot-based voting rules. In *Proc. of 11th AAMAS*, pages 603–610, 2012.
130. L. Xia and V. Conitzer. Stackelberg voting games: Computational aspects and paradoxes. In *Proc. of 24th AAAI*, pages 921–926, 2010.
131. Y. Yang and J. Guo. The control complexity of r -approval: from the single-peaked case to the general case. In *Proc. of 13th AAMAS*, pages 621–628, 2014.
132. O. Yilmaz. Random assignment under weak preferences. *Games and Economic Behavior*, 66(1):546–558, 2009.
133. H. P. Young. Social choice scoring functions. *SIAM Journal on Applied Mathematics*, 28(4):824–838, 1975.
134. L. Yu, H. Chan, and E. Elkind. Multiwinner elections under preferences that are single-peaked on a tree. In *Proc. of 23rd IJCAI*, pages 425–431, 2013.