

Voting with Ties: Strong Impossibilities via SAT Solving

Felix Brandt
Technische Universität München
Munich, Germany
brandtf@in.tum.de

Christian Saile
Technische Universität München
Munich, Germany
saile@in.tum.de

Christian Stricker
Technische Universität München
Munich, Germany
stricker@in.tum.de

ABSTRACT

Voting rules allow groups of agents to aggregate their preferences in order to reach joint decisions. The Gibbard-Satterthwaite theorem, a seminal result in social choice theory, implies that, when agents have *strict* preferences, all anonymous, Pareto-optimal, and *single-valued* voting rules can be strategically manipulated. In this paper, we consider multi-agent voting when there can be ties in the preferences as well as in the outcomes. These assumptions are extremely natural—especially when there are large numbers of alternatives—and enable us to prove much stronger results than in the overly restrictive setting of strict preferences. In particular, we show that (i) all anonymous Pareto-optimal rules where ties are broken according to the preferences of a chairman or by means of even-chance lotteries are manipulable, and that (ii) all pairwise Pareto-optimal rules are manipulable, no matter how ties are broken. These results are proved by reducing the statements to finite—yet very large—problems, which are encoded as formulas in propositional logic and then shown to be unsatisfiable by a SAT solver. We also extracted human-readable proofs from minimal unsatisfiable cores of the formulas in question, which were in turn verified by an interactive higher-order theorem prover.

1 INTRODUCTION

Whenever a group of multiple agents aims at reaching a joint decision in a fair and principled way, they need to aggregate their individual preferences using a voting rule. Voting rules are studied in detail in social choice theory and are coming under increasing scrutiny from computer scientists who are interested in their computational properties or want to utilize them in computational multiagent systems [see, e.g., 12, 17].

One of the most prominent results in social choice theory, the Gibbard-Satterthwaite theorem, shows that, when agents have *strict* preferences, all anonymous, Pareto-optimal, and *single-valued* voting rules are susceptible to strategic manipulation [28, 45].¹ The restriction to single-valued rules has been identified as a major shortcoming of the theorem. For instance, Gärdenfors [24] claims that “[resoluteness] is a rather restrictive and unnatural assumption.” In a similar vein, Kelly [32] writes that “the Gibbard-Satterthwaite theorem [...] uses an assumption of singlevaluedness which is unreasonable” and Taylor [48] that “If there is a weakness to the Gibbard-Satterthwaite theorem, it is the assumption that winners are unique.” This sentiment is echoed by various other authors [see, e.g., 2–4, 15, 16, 20, 39].

¹The actual statement by Gibbard and Satterthwaite is somewhat stronger because it uses non-dictatorship instead of anonymity and non-imposition instead of Pareto-optimality.

The problem with single-valuedness is that the voting rule has to return a single alternative *based on the preferences only*. For example, if there are two alternatives, *a* and *b*, and two agents such that one prefers *a* and the other one *b*, there is no deterministic way of selecting a single alternative without violating basic fairness conditions such as anonymity and neutrality.

There already is a large number of results in the literature showing impossibility theorems for set-valued voting rules [see, e.g., 2, 4, 5, 15, 16, 24, 32, 35, 41, 44]. However, these results make relatively strong assumptions on the tie-breaking mechanisms and what the agents know about these mechanisms [15, 16, 24, 41, 44], or require strong additional assumptions [2, 4, 5, 24, 32, 35].² Some of the results above are derived by assuming that individual preferences do not admit ties [15, 16, 41, 44]. In many settings, especially when there are large numbers of alternatives, it is natural or even inevitable to allow agents to express indifferences between alternatives. Moreover, in many subdomains of social choice such as coalition formation, matching, or allocation, indifferences are inherently present in the agents’ preferences because agents are only concerned with their own coalition, matching partner, or allocation. As we show in this paper, the assumption of weak preferences allows us to prove impossibility theorems that only require surprisingly weak assumptions with respect to the tie-breaking mechanism.³

We have obtained these results using computer-aided theorem proving techniques that were pioneered by Tang and Lin [47] and have been successfully used to tackle other problems in social choice [see, e.g., 8, 9, 13, 14, 26, 47].⁴ The basic idea is to reduce the statement in question to a finite—yet very large—problem, which is encoded as a formula in propositional logic and then shown to be unsatisfiable by a SAT solver. We then extract a minimal unsatisfiable set of constraints from the formula and translate this back into a human-readable proof of the result. Despite great efforts to simplify the proof of our main result as much as possible, it remains rather complex as it argues about 21 different preference profiles. We therefore verified the proof using the interactive theorem prover Isabelle/HOL, which releases any need to verify our code for generating it. In contrast to previous papers in this stream of research, we are even able to give a *lower bound* on the proof complexity: no such proof is possible using less than 18 preference profiles. This can be considered as evidence that it is unlikely that the statement would have been proved without the help of computers, which underlines the potential of computer-aided theorem proving in social choice theory.

²Early results by Barberà [5] and Kelly [32], for example, required voting rules to be quasi-transitively rationalizable, a condition which is almost prohibitive on its own [see, e.g., 37].

³In fact, it has turned out that for strict preferences, positive results can be obtained under the same assumptions [see, e.g., 10, 19, 39].

⁴An overview of these techniques has been provided by Geist and Peters [27].

2 PRELIMINARIES

Let $A = \{a, b, \dots\}$ be a finite set of m alternatives and $N = \{1, \dots, n\}$ a finite set of agents. A (weak) preference relation is a complete, reflexive, and transitive binary relation on A . The preference relation of agent i is denoted by \succsim_i , the set of all preference relations by \mathcal{R} . We write $>_i$ for the strict part of \succsim_i , i.e., $x >_i y$ if $x \succsim_i y$ but not $y \succsim_i x$, and \sim_i for the indifference part of \succsim_i , i.e., $x \sim_i y$ if $x \succsim_i y$ and $y \succsim_i x$. A preference relation \succsim_i is called *strict* if it additionally is anti-symmetric, i.e., $x >_i y$ or $y >_i x$ for all distinct alternatives x, y . We will compactly represent a preference relation as a comma-separated list with all alternatives among which an agent is indifferent placed in a set. For example $x >_i y \sim_i z$ is represented by $\succsim_i: x, \{y, z\}$. A *preference profile* R is a function from a set of agents N to the set of preference relations \mathcal{R} . The set of all preference profiles is denoted by \mathcal{R}^N .

Our central objects of study are *social choice functions* (SCFs), i.e., voting rules that map a preference profile to a set of alternatives called the *choice set*. Formally, an SCF is a function

$$f: \mathcal{R}^N \rightarrow 2^A \setminus \emptyset.$$

Given a preference profile R , an alternative x *Pareto-dominates* another alternative y if $x \succsim_i y$ for all $i \in N$ and $x >_j y$ for some $j \in N$. An alternative is *Pareto-optimal* if it is not Pareto-dominated by some other alternative. The notion of Pareto-optimality can be used to define a simple SCF that returns the set of all Pareto-optimal alternatives. Formally,

$$PO(R) = \{x \in A: x \text{ is not Pareto dominated in } R\}.$$

An SCF f is said to be *Pareto-optimal* if $f(R) \subseteq PO(R)$ for all $R \in \mathcal{R}^N$.

It is well-known that the Pareto dominance relation is transitive and that every Pareto-dominated alternative is Pareto-dominated by some alternative in $PO(R)$. This also implies that, for every $i \in N$, $PO(R)$ contains at least one alternative that is top-ranked by agent i because top-ranked alternatives can only be Pareto-dominated by other top-ranked alternatives. The previous observation allows us to define refinements of PO . For $i \in N$, PO_i returns all Pareto-optimal alternatives top-ranked by agent i .

$$PO_i(R) = \{x \in PO(A): \nexists y \in A \text{ with } y >_i x\}.$$

These SCFs are weakly dictatorial in the sense that they only return alternatives top-ranked by a pre-determined agent.

Two common symmetry conditions for SCFs are anonymity and neutrality. An SCF is *anonymous* if the choice set does not depend on the identities of the agents and *neutral* if it is symmetric with respect to alternatives. Formally, an SCF is anonymous if $f(R) = f(R')$ for all $R, R' \in \mathcal{R}^N$ and all bijections $\pi: N \rightarrow N$ such that $\succsim_i = \succsim'_{\pi(i)}$ for all $i \in N$. For a permutation π on A and a preference relation \succsim_i , $\pi(x) \succsim_i^\pi \pi(y)$ if and only if $x \succsim_i y$. An SCF f is *neutral* if $f(R^\pi) = \pi(f(R))$ for all $R \in \mathcal{R}^N$ and all permutations $\pi: A \rightarrow A$. PO is anonymous and neutral while PO_i clearly violates anonymity.

For a preference profile $R \in \mathcal{R}^N$, let

$$n_R(x, y) = |\{i \in N: x \succsim_i y\}|$$

be the number of agents who prefer x to y . The majority margin of x over y in R is denoted by $g_R(x, y)$ where

$$g_R(x, y) = n_R(x, y) - n_R(y, x).$$

An SCF f is *pairwise* if for all $R, R' \in \mathcal{R}^N$, $f(R) = f(R')$ whenever $g_R(x, y) = g_{R'}(x, y)$ for all alternatives $x, y \in A$. In other words, the choice set of a pairwise SCF only depends on the anonymized comparisons between pairs of alternatives. Since majority margins are invariant under permutations of agents, pairwise SCFs are anonymous. When ties are allowed, pairwise SCFs are slightly stronger than Fishburn's C2, which requires that the SCF only depends on n_R [23]. For example, PO satisfies C2, but violates pairwise SCFs. Nevertheless, there is a large number of attractive pairwise and Pareto-optimal SCFs [see, e.g., 21, 23]. Typical examples are Borda's rule, Kemeny's rule, the Simpson-Kramer rule (aka maximin), Schulze's rule, ranked pairs, or the essential set.

A very influential concept in social choice theory is that of a Condorcet winner, i.e., an alternative that is preferred to every other alternative by some majority of agents. Formally, an alternative x is a *Condorcet winner* in R if $g_R(x, y) > 0$ and a *Condorcet loser* if $g_R(x, y) < 0$ for all $y \in A \setminus \{x\}$, respectively. A *Condorcet extension* is an SCF that uniquely returns a Condorcet winner whenever one exists.

3 STRATEGYPROOFNESS OF SET-VALUED SOCIAL CHOICE FUNCTIONS

An important phenomenon in social choice is that agents misrepresent their preferences in order to obtain a more preferred outcome.

3.1 Tie-Breaking Mechanisms

When defining strategic manipulability for set-valued SCFs, one needs to specify how ties are broken and how much the agents know about the tie-breaking mechanism. In this paper, we will be concerned with the following tie-breaking mechanisms and corresponding epistemic assumptions.

- (1) Ties are broken arbitrarily; the agents do not have any knowledge about the tie-breaking mechanism.
- (2) Ties are broken by a chairman, i.e., there exists a strict tie-breaking ordering; this ordering is unknown to the agents.
- (3) Ties are broken by lottery and outcomes are compared based on expected utility; agents are neither aware of the probabilities nor of their concrete utility functions.⁵
- (4) Ties are broken by an even-chance (i.e., uniform) lottery and outcomes are compared based on expected utility; agents are unaware of their concrete utility functions.

Note that uncertainty on behalf of the agents in Assumption (1) and (2) is crucial. If the tie-breaking mechanism were deterministic and known to the agents, we are in the setting of single-valued SCFs and the Gibbard-Satterthwaite theorem applies. If the probabilities and concrete utility functions in Assumption (3) were known, a theorem by Hylland [29] implies that the only Pareto-optimal SCFs are randomizations over dictatorships.

⁵Utility functions are required to be consistent with ordinal preferences of the agents, i.e., if $x \succsim_i y$, then x has to yield at least as much utility as y .

3.2 Preference Extensions

As it turns out, each set of assumptions stated above can be modeled using so-called *preference extensions* which extend the agents' preferences over alternatives to preferences over sets of alternatives. The two preference extensions we consider in this paper are *Kelly's extension* [32] and *Fishburn's extension* [22, 25]. For all $X, Y \subseteq A$ and $\succsim_i \in \mathcal{R}$,

$$X \succsim_i^K Y \text{ iff } x \succsim_i y \text{ for all } x \in X, y \in Y, \text{ and} \quad (\text{Kelly})$$

$$X \succsim_i^F Y \text{ iff } X \setminus Y \succsim_i^K Y \text{ and } X \succsim_i^K Y \setminus X. \quad (\text{Fishburn})$$

The strict part of these relations will be denoted by $>_i^K$ and $>_i^F$, respectively. Note that

$$X >_i^K Y \text{ iff } X \succsim_i^K Y \text{ and there is } x \in X, y \in Y \text{ such that } x >_i y.$$

Similarly,

$$X >_i^F Y \text{ iff } X \succsim_i^F Y \text{ and there is } x \in X \setminus Y, y \in Y \\ \text{or } x \in X, y \in Y \setminus X \text{ such that } x >_i y.$$

It follows from the definitions that Fishburn's extension is a refinement of Kelly's extension, i.e.,

$$\succsim_i^K \subseteq \succsim_i^F \text{ and } >_i^K \subseteq >_i^F \text{ for every } \succsim_i \in \mathcal{R}.$$

3.3 Strategyproofness

With these extensions at hand, we can now formally define strategyproofness. An SCF f is *Kelly-manipulable* if there exist preference profiles $R, R' \in \mathcal{R}^N$, and an agent $i \in N$ such that $\succsim_j = \succsim_j'$ for all $j \neq i$ and $f(R') >_i^K f(R)$. f is said to satisfy *Kelly-strategyproofness* if it is not Kelly-manipulable. Fishburn-strategyproofness is defined analogously. The relationship between both preference extensions implies that Fishburn-strategyproofness is stronger than Kelly-strategyproofness.

We now connect both strategyproofness notions to the tie-breaking assumptions given in Section 3.1 [see 18, 25, 43, for more details]. It is fairly easy to see that Kelly-strategyproofness is equivalent to strategyproofness under Assumptions (1) and (3).⁶ Moreover, strategyproofness under Assumption (2) is equivalent to Fishburn-strategyproofness and strategyproofness under Assumption (4) corresponds to a strategyproofness notion stronger than Fishburn-strategyproofness.⁷ Hence, any impossibility for Fishburn-strategyproofness also implies an impossibility for tie-breaking by even-chance lotteries.

The following example (due to Feldman [19]) illustrates the definitions of Kelly-strategyproofness and Fishburn-strategyproofness. To this end, consider the preference profile R .

$$\succsim_1: a, \{b, c\} \quad \succsim_2: \{b, c\}, a$$

Clearly, $PO(R) = \{a, b, c\}$. Now assume that Agent 1 changes his preferences to \succsim_1' resulting in preference profile R' .

$$\succsim_1': a, b, c \quad \succsim_2': \{b, c\}, a$$

⁶Note that a Kelly-manipulable SCF can be *potentially* manipulated (once ties have been broken). It is possible to define an even weaker notion of strategyproofness where the SCF can be definitely manipulated. See Remark 4 in Section 5.2

⁷Another justification for Fishburn's extension using *a priori* probability distributions was proposed by Ching and Zhou [15].

Alternative c is Pareto-dominated by alternative b in R' and $PO(R') = \{a, b\}$. This does not constitute a Kelly-manipulation because $X = \{a, b\}$ is not preferred to $Y = \{a, b, c\}$ according to Kelly's extension applied to the preference relation \succsim_1 (note that $a >_1 b, a \in Y, b \in X$). In fact X and Y are incomparable according to Kelly's extension. This is in line with Assumption (1) because there could be a tie-breaking mechanism that always selects b from X and a from Y . The picture looks different for Fishburn's extension, however, as $X >_1^F Y$. To see that this concurs with Assumption (2), consider a chairman with preferences c, a, b . This chairman will select a from X and c from Y and $a >_i c$. For all other strict preference relations of the chairman, Agent 1 will be indifferent between the eventually chosen alternatives. Clearly, X is also preferred to Y according to Assumption (4): for all utility functions consistent with \succsim_1 , the expected utility for an even-chance lottery between a and b exceeds that of an even-chance lottery between all three lotteries.

The example shows that PO is Fishburn-manipulable. By contrast, as first shown by Feldman [19], PO does satisfy Kelly-strategyproofness. Since Feldman proves this statement by making reference to stronger strategyproofness notions, we give a self-contained proof below.

THEOREM 3.1. *PO is Kelly-strategyproof.*

PROOF. Assume for contradiction that there are two preference profiles R and R' , and an agent $i \in N$ such that $\succsim_j = \succsim_j'$ for all $j \neq i$ and $PO(R') >_i^K PO(R)$. Recall that PO_i is non-empty for all $i \in N$, i.e., PO returns at least one top-ranked alternative of each agent. This implies that a manipulation from R to R' would only be possible if there is some $x \in PO(R) \setminus PO(R')$ and there is no $x' \in PO(R)$ with $x' \sim_i x$. Since $x \notin PO(R')$, there has to be some $y \in PO(R')$ such that y Pareto-dominates x in R' . Moreover, y does not Pareto-dominate x in R . This implies that $x >_i y$. Since $x \in PO(R)$ and $y \in PO(R')$, it is impossible that $PO(R') >_i^K PO(R)$. \square

PO_i constitutes a family of weakly dictatorial SCFs satisfying Pareto-optimality and any reasonable form of strategyproofness, because agent i gets a subset of his most preferred alternatives, ruling out any possibility to manipulate and all other agents having no influence at all.

4 COMPUTER-AIDED THEOREM PROVING

Our results are obtained using the computer-aided proving methodology described by Brandt and Geist [13]. First, we provide a reduction argument in Lemma 5.1, which allows us to prove a statement for general domain sizes by restricting ourselves to a finite number of agents and alternatives. However, in these restricted domains, the number of anonymous SCFs is huge (see Table 1), which renders any type of exhaustive search infeasible. For our main results, we require three agents and three alternatives, which already admits about $3.3 \cdot 10^{384}$ possible anonymous SCFs.⁸ Thus, heuristic search algorithms as provided by state-of-the-art SAT solvers are required. Apart from allowing us to deal with enormous search spaces, the computer-aided approach has the major advantage that related

⁸For comparison, this search space exceeds that of Theorem 3 by Brandt and Geist [13] and lies in between Theorems 1 and 2 by Brandt et al. [9].

conjectures and hypotheses, e.g., statements including additional axioms, can be checked quickly using the same framework.

m	n	Preference profiles	SCFs
3	3	455	$\sim 3.3 \cdot 10^{384}$
3	4	1,820	$\sim 1.2 \cdot 10^{1,538}$
4	3	73,150	$\sim 1.2 \cdot 10^{86,031}$
4	4	1,426,425	$\sim 9.4 \cdot 10^{1,677,605}$

Table 1: Number of different profiles and possible SCFs when assuming anonymity.

4.1 SAT-Solving and Proof Extraction

Basically, the core of the computer-aided approach is the aforementioned encoding of the problems to be solved as a SAT instances in *conjunctive normal form* (CNF). For this, all axioms involved need to be stated in propositional logic. All variables are of the form $c_{R,X}$ with a preference profile R and a set of alternatives $X \subseteq A$. The semantics of these variables are that $c_{R,X}$ if and only if $f(R) = X$, i.e., the SCF f selects the set of alternatives X as the choice set for the preference profile R .

Although an encoding with variables $c_{R,x}$ for single alternatives x rather than choice sets would require less variable symbols, it would significantly increase the complexity of the clauses for some axioms, especially for strategyproofness. Due to the fact that strategyproofness clauses outnumber all other clauses combined, we chose the former encoding with more variables but much easier clauses.

First, we ensure that the variables $c_{R,X}$ indeed model a function rather than an arbitrary relation, i.e., for each preference profile R , there is exactly one choice set X such that the variable $c_{R,X}$ is set to true. We split this into *choice set existence*,

$$\left(\forall R \in \mathcal{R}^N \right) (\exists X \subseteq A) c_{R,X} \equiv \bigwedge_{R \in \mathcal{R}^N} \bigvee_{X \subseteq A} c_{R,X},$$

and *uniqueness*,

$$\begin{aligned} \left(\forall R \in \mathcal{R}^N \right) \left((\forall Y, Z \subseteq A) Y \neq Z \rightarrow \neg(c_{R,Y} \wedge c_{R,Z}) \right) \\ \equiv \bigwedge_{R \in \mathcal{R}^N} \bigwedge_{Y \neq Z} (\neg c_{R,Y} \vee \neg c_{R,Z}). \end{aligned}$$

By contrast to these rather elaborate axioms, the formalization of Pareto-optimality can be easily written without logical disjunctions as

$$\begin{aligned} \left(\forall R \in \mathcal{R}^N \right) (\forall x \notin PO(R)) x \notin f(R) \\ \equiv \bigwedge_{R \in \mathcal{R}^N} \bigwedge_{x \notin PO(R)} \bigwedge_{X \ni x} \neg c_{R,X}. \end{aligned}$$

Similar to the choice set uniqueness axiom, strategyproofness from some preference extension \mathcal{E} can be encoded by

$$\begin{aligned} \left(\forall R \in \mathcal{R}^N \right) (\forall \succ_i \in R) (\forall \succ'_i \in \mathcal{R}) \neg \left(f(R_{i \rightarrow \succ'_i}) \succ_i^{\mathcal{E}} f(R) \right) \\ \equiv \bigwedge_{R \in \mathcal{R}^N} \bigwedge_{\succ_i \in R} \bigwedge_{\succ'_i \in \mathcal{R}} \bigwedge_{Y \succ_i^{\mathcal{E}} X} \left(\neg c_{R_{i \rightarrow \succ'_i}, Y} \vee \neg c_{R, X} \right), \end{aligned}$$

with $R_{i \rightarrow \succ'_i}$ denoting the preference profile R where agent i 's preference relation is replaced with \succ'_i .

After encoding the axioms using a Java program, satisfiability of the SAT instance is checked with the LINGELING solver family by Biere [7]. If an instance turns out to be unsatisfiable, we extract a *minimal unsatisfiable core* (also called a *minimal unsatisfiable set* (MUS)), a feature which is offered by a range of SAT solvers. A subset of clauses is an *unsatisfiable core* if it is already unsatisfiable by itself. If removing any clause from the unsatisfiable core renders it satisfiable, it is called minimal. However, although an MUS is inclusion-minimal, it is not necessarily a *smallest* unsatisfiable core, i.e., a core with a minimal number of clauses or variables. In particular, a MUS does not have to be unique.

Especially with regard to proof extraction later on, we intend to find a smallest minimal unsatisfiable set (SMUS). This can theoretically be done with the software tool MARCO by Liffiton et al. [33], as it offers an option to search for an SMUS.⁹ Although MARCO does not terminate in any reasonable amount of time for our problem sizes,¹⁰ it returns the smallest MUS found *so far*. Thus, it yields increasingly better approximations of an SMUS over time.

Additionally, we aim at minimizing the number of required preference profiles instead of the number of all clauses of the CNF formula. One of the reasons behind this is that strategyproofness is responsible for most of the clauses in our SAT instances, resulting in MARCO spending most of the runtime on optimizing the size of the MUS concerning the number of applications of strategyproofness only, instead of rather concentrating on the number of different preference profiles involved in it. We realized this optimization objective by using *group-oriented* CNF formulas and declaring clauses of the choice set existence axioms as *interesting* groups and all the remaining clauses as a single *don't care* group. This technique significantly increases the performance of our search for a (group-oriented) SMUS. See Liffiton and Sakallah [34] for more details on group-oriented SAT solving.

Moreover, with the group-oriented approach, we can now also give lower bounds for the number of profiles needed in such impossibility proofs. We achieve this with the tool FORQES by Ignatiev et al. [30], as it supports a restricted version of group-oriented SAT solving, namely the specification of *don't care* clauses.¹¹ In contrast to MARCO, it does not compute or return approximations of an SMUS during its runtime, but rather iteratively rules out the existence of a MUS of a given size starting with the trivial size of just one clause (and finally returns an SMUS if not aborted before).

After finding a sufficiently small MUS, a proof trace can be extracted from the MUS with the help of certain SAT solvers like PicoSAT by Biere [6]. If this yields a reasonably sized proof trace, we can directly create a pen-and-paper proof by going through its main steps and translating the clauses back to the preference profile level. For this we use a dictionary containing the correspondences between SAT variables and preference profile/choice set combinations. We used this method, e.g., for Theorem 5.4.

⁹The option `--smus` for searching an SMUS is available up to version 1.1 only.

¹⁰The decision problem corresponding to finding an SMUS, i.e., whether there exists an MUS of size less than or equal to k is Σ_2^P -complete [30, 42].

¹¹The *don't care* clauses have to be specified in a separate CNF file and called via the `-n` option.

4.2 Formal Verification

If the computer-generated proof exceeds a reasonable size, it becomes a tedious and error-prone task for humans to translate the output of the SAT solver to a human-readable proof and thereby checking correctness. Simply accepting the black-box-like output of the SAT solver as a proof is not sufficient, as one has (i) to trust the correctness of the SAT solver and (ii) to rely on the correctness of the Java code that generates the CNF formula in the first place. The first concern is less problematic and is addressed by using a verified SAT solver [36]. However, more importantly, there is no guarantee that the Java program meets its specification. Even a verified SAT solver may produce an overall unsound proof due to a bug in the Java code for encoding the axioms. To tackle this issue, we make use of the interactive theorem prover Isabelle/HOL [see, e.g., 40] to produce a machine-verified proof. The main application of the generic proof assistant Isabelle is the formalization of mathematical proofs and formal verification. Building on the framework introduced by Brandl et al. [8], the set of preference profiles and conditions obtained from the MUS is translated to Isabelle and the user interactively develops the proof. This approach entirely removes the dependence on the unverified Java program and we obtain an independent Isabelle proof that can be even checked manually step by step. Trustworthiness of Isabelle is considerably high as it is widely used for verification tasks.¹²

5 RESULTS

We start by showing that for impossibility results using Pareto-optimality and strategyproofness, it suffices to prove that the axioms are incompatible for some fixed number of alternatives and agents.

LEMMA 5.1. *Let f be an anonymous SCF f that satisfies Pareto-optimality and strategyproofness for A and N . Then there is an anonymous SCF f' that satisfies these axioms for $A' \subseteq A$ and $N' \subseteq N$.*

PROOF. We define an embedding φ of preference profiles $R' = (\succsim'_1, \dots, \succsim'_{n'})$ over N' and A' into preference profiles R over N and A by means of extending the existing preferences with $D = A \setminus A'$ as new bottom-ranked, hence Pareto-dominated alternatives and adding indifferent agents:

$$\varphi(R') = R$$

$$\text{with } \succsim_i = \begin{cases} \succsim'_i \cup (A \times D) & \text{if } i \in N', \\ A \times A & \text{otherwise.} \end{cases}$$

Now let $f'(R') = f(\varphi(R'))$. f' is anonymous since f is anonymous and agents in N only differ by their preferences over A' . Pareto-optimality of f' holds because f is Pareto-optimal and $PO(R) = PO(R')$. Finally, f' is strategyproof because f is strategyproof and the choice sets of f' under the two profiles R' and $(R')_{i \mapsto \succsim'_i}$ are equal to the choice sets of f under the two (extended) profiles R and $R_{i \mapsto \succsim_i}$, respectively. \square

It is easily seen that Lemma 5.1 also holds for neutral SCFs.

¹²Using higher-order proof assistants such as Isabelle/HOL to prove these theorems in the first place is currently completely out of reach.

Profile	Agent 1	Agent 2	Agent 3
1	$a \sim c > b$	$a \sim c > b$	$b > c > a$
2	$c > a \sim b$	$b > c > a$	$c > b > a$
3	$a \sim c > b$	$b > c > a$	$c > b > a$
4	$a \sim c > b$	$a > b > c$	$b > c > a$
5	$c > a \sim b$	$a > b > c$	$b > c > a$
6	$b > a \sim c$	$c > a \sim b$	$b > c > a$
7	$a \sim c > b$	$b > a \sim c$	$b > c > a$
8	$b \sim c > a$	$a > b \sim c$	$a > c > b$
9	$b \sim c > a$	$b > a \sim c$	$a > b > c$
10	$c > a \sim b$	$a > b > c$	$c > b > a$
11	$a \sim c > b$	$a > b > c$	$c > b > a$
12	$c > a \sim b$	$b > a > c$	$c > b > a$
13	$a \sim c > b$	$b > a > c$	$c > b > a$
14	$a > b \sim c$	$c > a \sim b$	$a > c > b$
15	$b \sim c > a$	$a > b \sim c$	$a > b > c$
16	$a \sim b > c$	$c > a \sim b$	$a > b > c$
17	$a > b \sim c$	$a > b > c$	$b > c > a$
18	$a \sim c > b$	$b > a \sim c$	$b > a > c$
19	$a > b \sim c$	$c > a \sim b$	$a > b > c$
20	$b > a \sim c$	$a > b > c$	$b > a > c$
21	$b \sim c > a$	$a > b > c$	$b > c > a$

Table 2: The 21 profiles needed for the proof of Theorem 5.2.

5.1 Fishburn-strategyproofness

Recall from Section 3.3 that PO is Fishburn-manipulable. Our main theorem is a much more general statement showing that every anonymous and Pareto-optimal SCF is Fishburn-manipulable.

THEOREM 5.2. *There is no anonymous SCF that satisfies Pareto-optimality and Fishburn-strategyproofness for $m \geq 3$ and $n \geq 3$.*

The full proof of Theorem 5.2 is omitted due to space constraints. Instead, we provide some information on its size and structure and prove a weaker version of Theorem 5.2 for neutral SCFs (Corollary 5.3). Starting from the initial, unsatisfiable SAT instance, we used MARCO to find a small (group-oriented) MUS, which utilizes the 21 profiles listed in Table 2. Although the MUS does not guarantee that this is the minimal number of profiles needed, no significantly easier proof of this form exists, because we were able to compute a lower bound of 18 profiles with FORQES. With the help of PICO SAT, we extracted a proof out of the MUS which is divided in 22 main proof steps.¹³ The first step in the raw proof trace looks as follows.

```
111 -4 0 1 2 3 5 9 10 11 12 13 18 20 21 22 23 26 27
      28 29 30 31 32 33 34 45 46 47 48 56 57 63 64
      65 66 67 68 69 70 71 72 73 74 75 76 77 96 98
      99 101 102 103 104 106 107 0
```

with 111 designating the clause ID of this step (also counting all clauses from the MUS). The numbers after the first zero indicate the IDs of the clauses (corresponding to profiles, manipulation instances and previous proof steps) that have to be applied. The literal in front of the first zero indicates the result—here that the

¹³We applied Pareto-optimality constraints manually before these proof steps to make the proof as compact as possible.

$f(\cdot)$	$\{a\}$	$\{b\}$	$\{a, b\}$	$\{c\}$	$\{a, c\}$	$\{b, c\}$	$\{a, b, c\}$
R^1		P		Q		O	
R^2		U		L		V	
R^3		S		T		R	
R^4	A	H	I	E	N	M	J
R^5	?	B	C	F	?	G	D
R^{10}	?	?	?	K	?	?	?

Table 3: Main steps in the proof of Theorem 5.2.

Step	Sequence of profiles
A	4, 1, 3, 2, 12, 10, 11, 13, 18, 20, 9, 21, 5
B, C, D	5, 16, 19, 14, 8, 15, 17, 4, 21, 9, 20, 18, 7
E	4, 21, 9, 20, 18, 7
F, G	5, 21, 9, 20, 18, 7, 6
H, I, J	4, 5, 17
K	10, 5
L	2, 10, 12
M, O, P	1, 4, 5
R, S, U, V	2, 3, 1

Table 4: Step details in the proof of Theorem 5.2.

SAT variable 4 is false. The second zero indicates the end of this proof step.

Translating the raw output into a human-readable format, it turns out that only profiles R^1 to R^5 , and R^{10} are directly affected in the results of these main proof steps. The rest of the profiles are needed in intermediate steps. Table 3 shows the sequence of the proof steps indicated by letters in lexicographic ordering. A red entry designates a proof step that shows that the profile in this row cannot be assigned the choice set corresponding to its column. Accordingly, when all possible choice sets except one are ruled out for a given profile, the remaining set must be chosen, as indicated in green. Observe that alternative a is Pareto-dominated in the first three profiles. Thus, the choice sets $\{a\}$, $\{a, b\}$, $\{a, c\}$, and $\{a, b, c\}$ are not allowed for these profiles. The light orange area highlights the final contradiction that no choice set can be assigned to profile R^2 . Each main proof step consists of a separate proof by contradiction using the profiles listed in Table 4 with the dependencies between the steps shown in Figure 1. Exemplary, proof step A, which shows $f(R^4) \neq \{a\}$, involves 40 applications of Fishburn-strategyproofness involving profiles 4, 1, 3, 2, 12, 10, 11, 13, 18, 20, 9, 21, and 5 in this ordering and is a direct prerequisite for the steps B, C, D and N. The very first step assumes for contradiction that $f(R^4) = \{a\}$ and yields that $f(R^1) = \{c\}$, because a is Pareto-dominated by c in R^1 and $b \in f(R^1)$ would contradict Fishburn-strategyproofness.

We now give an excerpt of our Isabelle proof. In contrast to the encoding for SAT-solving, where there is only one type of variables, the formalization in Isabelle employs variables for alternatives, agents, preference profiles, and social choice functions and allows for a straightforward formulation of the result in higher-order logic. According to Table 2, we instantiate preference profiles $R1$ and $R4$




Figure 1: Dependencies of proof steps in Theorem 5.2

with fixed variables a, b, c for alternatives and $A1, A2, A3$ for agents, respectively.

$R1 = A1: [a, c], b \quad A2: [a, c], b \quad A3: b, c, a$
 $R4 = A1: [a, c], b \quad A2: a, b, c \quad A3: b, c, a$

Next we derive the restrictions Pareto-optimality imposes on the outcome of the social choice function scf . While it is obvious that Pareto-optimality implies a must not be in $scf R1$, the formal proof requires this to be derived from the elementary definition of Pareto-optimality.

```
lemma : "a ∉ scf R1"
  by (rule pareto_efficiency)
     (simp_all add: eval_pareto R1.eval)
```

In a similar vein, we proceed with the strategyproofness condition. We prove that agent $A3$ in preference profile $R4$ must not prefer the outcome of $scf R1$ to the outcome of $scf R4$.

```
lemma : "¬ scf R1 >[Fishb(R4 A3)] scf R4"
  by (intro strategyproof'[where j = A3])
     (simp_all add: R1.eval R4.eval)
```

This example showcases the interactive nature of Isabelle. The user gives commands and Isabelle's automation completes the proof, checks that no cases are unintentionally omitted and derives the equivalent statement in a quantifier-free formula.

$$b \in scf R1 \wedge a \in scf R4 \vee c \in scf R1 \wedge$$

$$(a \in scf R4 \vee b \notin scf R1 \wedge b \in scf R4) \vee$$

$$b \in scf R1 \wedge b \notin scf R4 \wedge a \in scf R4 \vee$$

$$c \in scf R1 \wedge c \notin scf R4 \wedge (a \in scf R4 \vee$$

$$b \in scf R4) \vee (c \in scf R4 \rightarrow b \in scf R4 \vee$$

$$b \notin scf R1) \wedge (c \in scf R1 \rightarrow c \in scf R1 \vee$$

$$b \notin scf R1)$$

In order to verify Theorem 5.2, we apply these steps to the preference profiles and manipulation instances indicated by the MUS. The generated quantifier-free formulas are gathered and the conjunction thereof is proven to be unsatisfiable. Replication

data for the Isabelle proof can be found in Harvard Dataverse under <https://dataverse.harvard.edu/privateurl.xhtml?token=02959985-6dbe-4901-8e52-14805bbd5ba4>.

For the reader's benefit we now give a full, human-readable proof of a significantly weaker version of Theorem 5.2 which additionally requires neutrality. This proof is based on only three preference profiles (rather than 21) and requires only five strategyproofness applications (rather than 89).

COROLLARY 5.3. *There is no neutral and anonymous SCF that satisfies Pareto-optimality and Fishburn-strategyproofness for $m \geq 3$ and $n \geq 2$.*

PROOF. Let $N = \{1, 2\}$ and $A = \{a, b, c\}$ and assume for contradiction that f is a neutral and anonymous SCF that satisfies Pareto-optimality and Fishburn-strategyproofness. First, consider preference profile R^1 .

$$\succsim_1^1: a, b, c \quad \succsim_2^1: b, a, c \quad .$$

By anonymity and neutrality, $a \in f(R^1)$ if and only if $b \in f(R^1)$. Together with c being Pareto-dominated by b , this implies $f(R^1) = \{a, b\}$. This already determines the choice set for the following preference profile R^2 .

$$\succsim_1^2: a, b, c \quad \succsim_2^2: \{b, c\}, a$$

Both $f(R^2) = \{a\}$ and $f(R^2) = \{b\}$ would allow for manipulations since the second agent prefers $\{a, b\}$ to $\{a\}$ in R^2 and $\{b\}$ to $\{a, b\}$ in R^1 . Furthermore, $c \notin f(R^2)$ since alternative c is Pareto-dominated by b , hence $f(R^2) = \{a, b\}$. Lastly, we consider preference profile R^3 .

$$\succsim_1^3: a, \{b, c\} \quad \succsim_2^3: \{b, c\}, a$$

By anonymity and neutrality, $b \in f(R^3)$ if and only if $c \in f(R^3)$. However, if $\{b, c\} \subseteq f(R^3)$, then the first agent can deviate from \succsim_1^3 to \succsim_1^2 . This only leaves $f(R^3) = \{a\}$, which allows the first agent to deviate from \succsim_1^2 to \succsim_1^3 , a contradiction. \square

Remark 1. Using a more complicated proof, it can be shown that Corollary 5.3 even holds for a weakening of Fishburn-strategyproofness where choice sets can only be compared when they are disjoint or contained in each other. We are unable to prove the same for Theorem 5.2, even when $m = 4$ and $n = 3$.

Remark 2. The axioms of Theorem 5.2 are independent of each other. PO satisfies all axioms except Fishburn-strategyproofness, PO_i for some $i \in N$ satisfies all axioms except anonymity, and the trivial SCF which always returns all alternatives satisfies all axioms except Pareto-optimality. Also, the bounds used in the theorem ($m \geq 3$ and $n \geq 3$) are tight, as confirmed by the SAT solver.

Remark 3. The main result by Brandt and Geist [13] shows that all Pareto-optimal *majoritarian* SCFs are Fishburn-manipulable when $m \geq 5$ and $n \geq 7$. This result is implied by Theorem 5.2, except that the result by Brandt and Geist even holds for strict preferences.

Remark 4. When assuming strict preferences, there are various Fishburn-strategyproof SCFs satisfying Pareto-optimality, e.g., the omninomination rule [11]. Theorem 5.2 shows that these SCFs cannot be extended to weak preferences without giving up one of these desirable properties.

5.2 Kelly-strategyproofness

It is not possible to replace Fishburn-strategyproofness with Kelly-strategyproofness in Theorem 5.2 because PO is Kelly-strategyproof. We therefore focus on pairwise SCFs when dealing with Kelly-strategyproofness and directly present a human-readable proof of the following impossibility.

THEOREM 5.4. *There is no pairwise SCF that satisfies Pareto-optimality and Kelly-strategyproofness for $m \geq 3$ and $n \geq 3$.*

PROOF. Let $N = \{1, 2, 3\}$ and $A = \{a, b, c\}$ and assume for contradiction that there is a pairwise SCF f that satisfies Pareto-optimality and Kelly-strategyproofness. If not stated otherwise, the absolute values of the majority margins in the following applications of pairwise SCFs are always one. First, consider the classic Condorcet profile R^1 .

$$\succsim_1^1: a, b, c \quad \succsim_2^1: c, a, b \quad \succsim_3^1: b, c, a.$$

Due to the symmetry of the profile, we may assume without loss of generality that $b \in f(R^1)$. Now consider R^2 .

$$\succsim_1^2: a, b, c \quad \succsim_2^2: c, a, b \quad \succsim_3^2: b, \{a, c\},$$

R^2 and R^1 only differ in the third agent's preferences. By Fishburn-strategyproofness, $b \in f(R^2)$, as otherwise Agent 3 could obtain a preferred choice set by changing his preferences from \succsim_3^2 to \succsim_3^1 . Now consider R^3 , which has the same majority margins as R^2 .

$$\succsim_1^3: a, b, c \quad \succsim_2^3: \{a, c\}, b \quad \succsim_3^3: b, c, a,$$

Since $g_{R^2} = g_{R^3}$, $b \in f(R^3)$. Now consider R^4 .

$$\succsim_1^4: a, b, c \quad \succsim_2^4: a, c, b \quad \succsim_3^4: b, c, a,$$

R^4 differs from R^3 by the second agent's preferences \succsim_2^4 . It follows that $b \in f(R^4)$. Otherwise, the second agent could misrepresent his preferences \succsim_2^3 by \succsim_2^4 and obtain the Kelly-preferred choice set $f(R^4)$ without b . Finally, consider R^5 .

$$\succsim_1^5: a, b, c \quad \succsim_2^5: \{a, b, c\} \quad \succsim_3^5: \{a, b, c\}.$$

Since $g_{R^5} = g_{R^4}$, $b \in f(R^5)$ holds as well. However, b is Pareto-dominated by a in R^5 , yielding a contradiction. \square

Remark 1. The conjunction of pairwise SCFs and Pareto-optimality implies that Condorcet losers should not be chosen whenever the pairwise majority relation is transitive. We have shown that Theorem 5.4 also holds when pairwise SCFs are replaced with this weaker, but technical, condition and $n \geq 4$. Note that this technical assumption is weaker than requiring the SCF to be a Condorcet extension or to never return Condorcet losers.¹⁴ Interestingly, the smallest MUS we found for this statement consists of 29 profiles.

¹⁴Brandt [10, Theorem 2] has also shown that every Condorcet extension is Kelly-manipulable. While his proof needs $3m$ agents, we require Pareto-optimality for our reduction argument.

Remark 2. The axioms of Theorem 5.4 are independent of each other. Borda’s rule satisfies all axioms except Kelly-strategyproofness, *PO* satisfies all axioms except pairwise, and the trivial SCF which always returns all alternatives satisfies all axioms except Pareto-optimality. Also, the bounds used in the theorem ($m \geq 3$ and $n \geq 3$) are tight, as confirmed by the SAT solver.

Remark 3. Theorem 5.4 implies Theorem 4 by Aziz et al. [1], who use a stronger notion of strategyproofness and furthermore require $m, n \geq 4$. Brandl et al. [9, Table 2] mention a consequence of this theorem for Fishburn-strategyproof SCFs. Interestingly, this consequence follows from both Theorem 5.2 and Theorem 5.4.

Remark 4. When assuming strict preferences, there are attractive pairwise Kelly-strategyproof SCFs satisfying Pareto-optimality, e.g., the essential set [10]. Theorem 5.4 shows that these SCFs cannot be extended to weak preferences without giving up one of these desirable properties. The same is true if we instead define Kelly’s extension by requiring that X is preferred to Y if and only if every alternative in X is *strictly* preferred to every alternatives in Y [10, Remark 6].

6 CONCLUSION AND DISCUSSION

We investigated the existence of anonymous, Pareto-optimal, and strategyproof SCFs when there may be ties in the preferences as well as in the outcomes. Our main results are as follows.

- (1) There are no such SCFs when ties are broken according to the preferences of a chairman or by means of an even-chance lottery (Theorem 5.2).
- (2) There are no such pairwise SCFs, no matter how ties are broken (Theorem 5.4).

The computer-aided proof of Theorem 5.2 is rather complex and we have shown that no significantly easier proof exists. The interpretation using even-chance lotteries can be used to relate this statement to a recent significant result in probabilistic social choice [8]. Brandl et al. [8] have shown that there is no anonymous, neutral, and strategyproof SCF that satisfies a strengthening of Pareto-optimality called *SD*-efficiency, when ties are broken by *arbitrary* lotteries (known to the agents). Our result is weaker in that it only allows for even-chance lotteries, but it is stronger in that it only requires Pareto-optimality (rather than *SD*-efficiency) and dispenses with neutrality. Even-chance lotteries are the most natural—and sometimes the only acceptable—form of randomized tie-breaking [see, e.g., 22]. Moreover, it may be very difficult to implement non-uniform lotteries in practice. Neutrality seems like an appealing fairness criterion, but can be overly restrictive in some settings [see, e.g., 46, 49]. In fact, many voting rules used in the real world such as supermajority rules are not neutral. It was technically very challenging and required the aid of computers to prove Theorem 5.2 without the assumption of neutrality. Mossel and Rácz [38] faced similar difficulties when generalizing the quantitative Gibbard-Satterthwaite theorem by Isaksson et al. [31] to non-neutral SCFs.

Due to our weak assumptions about tie-breaking mechanisms, our results use much weaker notions of strategyproofness than those by Duggan and Schwartz [16], Ching and Zhou [15], Rodríguez-Álvarez [41], and Sato [44]. In part, this is possible because we allow for weak preferences.

Our results are tight in the sense that omitting any of the axioms, reducing the number of agents or alternatives, or requiring strict preferences immediately allows for positive results. This underlines the adequacy of impossibility results to improve our understanding of what *can* be achieved and to guide practitioners looking for attractive SCFs.

There are few opportunities to strengthen our results even further. One is to check whether Theorem 5.2 holds for weak Pareto-optimality (where an alternative is Pareto-dominated if *all* agents are strictly better off).¹⁵ Another one is to try to replace anonymity with weaker conditions such as non-dictatorship conditions for weak preferences. This would, however, require a new and more complicated reduction argument.

REFERENCES

- [1] H. Aziz, F. Brandl, and F. Brandt. 2014. On the Incompatibility of Efficiency and Strategyproofness in Randomized Social Choice. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*. AAAI Press, 545–551.
- [2] T. Bandyopadhyay. 1983. Manipulation of non-imposed, non-oligarchic, non-binary group decision rules. *Economics Letters* 11, 1–2 (1983), 69–73.
- [3] T. Bandyopadhyay. 1983. Multi-Valued Decision Rules and Coalitional Non-Manipulability. *Economics Letters* 13, 1 (1983), 37–44.
- [4] S. Barberà. 1977. The Manipulation of Social Choice Mechanisms That Do Not Leave “Too Much” to Chance. *Econometrica* 45, 7 (1977), 1573–1588.
- [5] S. Barberà. 1977. Manipulation of Social Decision Functions. *Journal of Economic Theory* 15, 2 (1977), 266–278.
- [6] A. Biere. 2008. PicoSAT Essentials. *Journal on Satisfiability, Boolean Modeling and Computation (JSAT)* 4 (2008), 75–79.
- [7] A. Biere. 2013. Lingeling, Plingeling and Treengeling entering the SAT competition 2013. In *Proceedings of the SAT Competition 2013*. 51–52.
- [8] F. Brandl, F. Brandt, M. Eberl, and C. Geist. 2017. Proving the Incompatibility of Efficiency and Strategyproofness via SMT Solving. *J. ACM* (2017). Forthcoming. Preliminary results appeared at IJCAI-2016.
- [9] F. Brandl, F. Brandt, C. Geist, and J. Hofbauer. 2015. Strategic Abstention based on Preference Extensions: Positive Results and Computer-Generated Impossibilities. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*. AAAI Press, 18–24.
- [10] F. Brandt. 2015. Set-Monotonicity Implies Kelly-Strategyproofness. *Social Choice and Welfare* 45, 4 (2015), 793–804. Preliminary results appeared at IJCAI-2011.
- [11] F. Brandt and M. Brill. 2011. Necessary and Sufficient Conditions for the Strategyproofness of Irresolute Social Choice Functions. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*. ACM Press, 136–142.
- [12] F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia (Eds.). 2016. *Handbook of Computational Social Choice*. Cambridge University Press.
- [13] F. Brandt and C. Geist. 2016. Finding Strategyproof Social Choice Functions via SAT Solving. *Journal of Artificial Intelligence Research* 55 (2016), 565–602. Preliminary results appeared at AAMAS-2014.
- [14] F. Brandt, C. Geist, and D. Peters. 2017. Optimal Bounds for the No-Show Paradox via SAT Solving. *Mathematical Social Sciences* 90 (2017), 18–27. Special Issue in Honor of Hervé Moulin. Preliminary results appeared at AAMAS-2016.
- [15] S. Ching and L. Zhou. 2002. Multi-valued strategy-proof social choice rules. *Social Choice and Welfare* 19, 3 (2002), 569–580.
- [16] J. Duggan and T. Schwartz. 2000. Strategic Manipulability without Resoluteness or Shared Beliefs: Gibbard-Satterthwaite Generalized. *Social Choice and Welfare* 17, 1 (2000), 85–93.
- [17] U. Endriss (Ed.). 2017. *Trends in Computational Social Choice*. AI Access.
- [18] B. Erdamar and M. R. Sanver. 2009. Choosers as extension axioms. *Theory and Decision* 67, 4 (2009), 375–384.
- [19] A. Feldman. 1979. Manipulation and the Pareto Rule. *Journal of Economic Theory* 21 (1979), 473–482.
- [20] A. Feldman. 1979. Nonmanipulable multi-valued social choice decision functions. *Public Choice* 34 (1979), 177–188.
- [21] F. Fischer, O. Hudry, and R. Niedermeier. 2016. Weighted Tournament Solutions. In *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia (Eds.). Cambridge University Press, Chapter 4.

¹⁵Theorem 5.4 does not hold for this weaker notion of Pareto-optimality because the SCF that returns all weakly Pareto-optimal alternatives is pairwise and Kelly-strategyproof. There are simple examples showing that this SCF fails to be Fishburn-strategyproof.

- [22] P. C. Fishburn. 1972. Even-chance lotteries in social choice theory. *Theory and Decision* 3, 1 (1972), 18–40.
- [23] P. C. Fishburn. 1977. Condorcet Social Choice Functions. *SIAM J. Appl. Math.* 33, 3 (1977), 469–489.
- [24] P. Gärdenfors. 1976. Manipulation of Social Choice Functions. *Journal of Economic Theory* 13, 2 (1976), 217–228.
- [25] P. Gärdenfors. 1979. On definitions of manipulation of social choice functions. In *Aggregation and Revelation of Preferences*, J. J. Laffont (Ed.). North-Holland.
- [26] C. Geist and U. Endriss. 2011. Automated Search for Impossibility Theorems in Social Choice Theory: Ranking Sets of Objects. *Journal of Artificial Intelligence Research* 40 (2011), 143–174.
- [27] C. Geist and D. Peters. 2017. Computer-aided Methods for Social Choice Theory. In *Trends in Computational Social Choice*, U. Endriss (Ed.). Chapter 13.
- [28] A. Gibbard. 1973. Manipulation of Voting Schemes: A General Result. *Econometrica* 41, 4 (1973), 587–601.
- [29] A. Hylland. 1980. Strategyproofness of Voting Procedures with Lotteries as Outcomes and Infinite Sets of Strategies. (1980). Mimeo.
- [30] A. Ignatiev, A. Previti, M. Liffiton, and J. Marques-Silva. 2015. Smallest MUS Extraction with Minimal Hitting Set Dualization. In *Proceedings of the 21st International Conference on Principles and Practice of Constraint Programming*, Vol. 9255. Springer-Verlag New York, Inc., New York, NY, USA, 173–182.
- [31] M. Isaksson, G. Kindler, and E. Mossel. 2012. The Geometry of Manipulation: A Quantitative Proof of the Gibbard-Satterthwaite Theorem. *Combinatorica* 32 (2012), 221–250. Preliminary results appeared at FOCS-2010.
- [32] J. S. Kelly. 1977. Strategy-Proofness and Social Choice Functions Without Single-Valuedness. *Econometrica* 45, 2 (1977), 439–446.
- [33] M. H. Liffiton, A. Previti, A. Malik, and J. Marques-Silva. 2016. Fast, flexible MUS enumeration. *Constraints* 21, 2 (2016), 223–250.
- [34] M. H. Liffiton and K. A. Sakallah. 2008. Algorithms for Computing Minimal Unsatisfiable Subsets of Constraints. *Journal of Automated Reasoning* 40, 1 (2008), 1–33.
- [35] I. MacIntyre and P. K. Pattanaik. 1981. Strategic voting under minimally binary group decision functions. *Journal of Economic Theory* 25, 3 (1981), 338–352.
- [36] F. Marić. 2010. Formal verification of a modern SAT solver by shallow embedding into Isabelle/HOL. *Theoretical Computer Science* 411, 50 (2010), 4333 – 4356.
- [37] A. Mas-Colell and H. Sonnenschein. 1972. General Possibility Theorems for Group Decisions. *Review of Economic Studies* 39, 2 (1972), 185–192.
- [38] E. Mossel and M. Z. Rácz. 2015. A quantitative Gibbard-Satterthwaite theorem without neutrality. *Combinatorica* 35, 3 (2015), 317–387. Preliminary results appeared at STOC-2012.
- [39] K. Nehring. 2000. Monotonicity implies generalized strategy-proofness for correspondences. *Social Choice and Welfare* 17, 2 (2000), 367–375.
- [40] T. Nipkow, L. C. Paulson, and M. Wenzel. 2002. *Isabelle/HOL – A Proof Assistant for Higher-Order Logic*. Lecture Notes in Computer Science (LNCS), Vol. 2283. Springer-Verlag.
- [41] C. Rodríguez-Álvarez. 2007. On the manipulation of social choice correspondences. *Social Choice and Welfare* 29, 2 (2007), 175–199.
- [42] V. Ryvchin and O. Strichman. 2011. Faster Extraction of High-Level Minimal Unsatisfiable Cores. In *Proceedings of the 14th International Conference on Theory and Applications of Satisfiability Testing - SAT 2011 (Lecture Notes in Computer Science (LNCS))*, Karem A. Sakallah and Laurent Simon (Eds.), Vol. 6695. Springer, 174–187.
- [43] M. R. Sanver and W. S. Zwicker. 2012. Monotonicity properties and their adaption to irresolute social choice rules. *Social Choice and Welfare* 39, 2–3 (2012), 371–398.
- [44] S. Sato. 2008. On strategy-proof social choice correspondences. *Social Choice and Welfare* 31 (2008), 331–343.
- [45] M. A. Satterthwaite. 1975. Strategy-Proofness and Arrow’s Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions. *Journal of Economic Theory* 10, 2 (1975), 187–217.
- [46] A. K. Sen. 1970. *Collective Choice and Social Welfare*. North-Holland.
- [47] P. Tang and F. Lin. 2009. Computer-aided proofs of Arrow’s and other impossibility theorems. *Artificial Intelligence* 173, 11 (2009), 1041–1053. Preliminary results appeared at AAAI-2008.
- [48] A. D. Taylor. 2005. *Social Choice and the Mathematics of Manipulation*. Cambridge University Press.
- [49] W. S. Zwicker. 2008. Consistency without neutrality in voting rules: When is a vote an average? *Mathematical and Computer Modelling* 48, 9 (2008), 1357 – 1373.