Strategyproof Social Choice When Preferences and Outcomes May Contain Ties

Felix Brandt   Christian Saile   Christian Stricker
Technische Universität München

The Gibbard-Satterthwaite theorem implies that all anonymous, Pareto-optimal, and single-valued social choice functions can be strategically manipulated. In this paper, we investigate whether there exist social choice correspondences (SCCs), that satisfy these conditions under various assumptions about how single alternatives are eventually selected from the choice set. These assumptions include even-chance lotteries as well as resolute choice functions and linear tie-breaking orderings unknown to the agents. We show that *(i)* all anonymous Pareto-optimal SCCs where ties are broken according to some linear tie-breaking ordering or by means of even-chance lotteries are manipulable, and that *(ii)* all pairwise Pareto-optimal SCCs are manipulable for any deterministic tie-breaking rule. These results are proved by reducing the statements to finite—yet very large—formulas in propositional logic, which are then shown to be unsatisfiable by a computer.

1. Introduction

The Gibbard-Satterthwaite theorem has established that all non-dictatorial, non-imposing, and single-valued social choice functions (SCFs) are susceptible to strategic manipulation (Gibbard, 1973; Satterthwaite, 1975). While this sweeping impossibility has had a great effect on microeconomic theory at large, it has been repeatedly observed that the assumption of single-valuedness is somewhat restrictive, in particular in the context of voting.[1] For instance, Gärdenfors (1976) claims that "[resoluteness] is a rather restrictive and unnatural assumption." In a similar vein, Kelly (1977) writes that "the Gibbard-Satterthwaite theorem [. . .] uses an assumption of singlevaluedness which is unreasonable" and Taylor (2005) that "If there is a weakness to the Gibbard-Satterthwaite theorem, it is the assumption that winners are unique." This sentiment

[1]Gibbard himself acknowledged this limitation by proving a strengthening of his theorem that allows for randomized outcomes (Gibbard, 1977).

is echoed by various other authors (see, e.g., Barberà, 1977b; Duggan and Schwartz, 2000; Nehring, 2000; Barberà et al., 2001; Ching and Zhou, 2002). The problem with single-valuedness is not that the SCF has to return a single alternative, which is a well-motivated practical requirement, but that it has to select 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, which demand that all agents and all alternatives are treated equally. In the context of voting, these conditions are imperative because elections should be impartial.

Since anonymity is stronger than non-dictatorship and Pareto-optimality is stronger than non-imposition, the Gibbard-Satterthwaite theorem implies that no single-valued SCF can simultaneously satisfy anonymity, Pareto-optimality, and strategyproofness. In this paper, we investigate whether there exist set-valued SCFs, so-called *social choice correspondences (SCCs)*, that satisfy these properties under various assumptions on how ties are broken to eventually select a single alternative. Our main results are as follows.

1. There are no such SCCs when ties are broken according to some linear tie-breaking ordering unknown to the agents or by means of an even-chance lottery (Theorem 1).

2. There are no such pairwise SCCs for any deterministic tie-breaking rule unknown to the agents (Theorem 2).

In contrast to most existing results, our results crucially rely on the possibility of ties in the preferences, i.e., we require preferences to be complete and transitive, but not necessarily anti-symmetric. At the same time, we use much weaker notions of strategyproofness than related results. In the following, we first defend the assumption of weak preferences before pointing out in which sense our notions of strategyproofness are weaker than those used in existing impossibility theorems.

A large part of the social choice literature focusses on the special case of strict individual preferences. We believe that this assumption is mostly driven by mathematical convenience rather than by practical constraints. In most applications, indifferences are ubiquitous and sometimes even inevitable. A voter who strongly believes in open borders may find all nationalistic parties equally unacceptable, or an employee may be indifferent between all budget proposals that assign the same budget to his department. The case of indifferences is even more striking when social outcomes have the characteristics of private goods such as partitions of agents or assignments of objects to agents. In these settings, agents are likely to be indifferent between coalitions or assignments in which they are grouped with the same agents or in which they receive the same objects. Moreover, the sheer number of alternatives renders it impossible to come up with a strict ranking of all possible partitions, matchings, or assignments.

Independently of whether preferences are strict or weak, there are various ways of defining strategyproofness for SCCs. The tie-breaking assumptions mentioned above lead to so-called preference extensions that extend preferences over alternatives to incomplete preferences over sets. An incomplete preference extension allows for two different types of strategyproofness, one being stronger than the other. The first one, *strong* *strategyproofness*, counts every misrepresentation of an agent that changes the outcome to a set that is preferred or incomparable to the original choice set as a manipulation. The second one, *weak strategyproofness*, only counts a misrepresentation as a manipulation if the new set is comparable and preferred to the original choice set. The more pairs of sets are comparable, the less demanding becomes strong strategyproofness and the more demanding becomes weak strategyproofness. In the extreme case, where preferences over sets are complete, both notions coincide. The literature has produced several impossibility results for strong strategyproofness (e.g., Duggan and Schwartz, 2000; Barberà et al., 2001; Ching and Zhou, 2002; Benoît, 2002) and several possibility results for weak strategyproofness (e.g., Feldman, 1979; Nehring, 2000; Brandt, 2015; Brandt and Lederer, 2022). Our results show that, when allowing for weak preferences, the possibility results break down as there are far-reaching impossibility theorems that only require weak strategyproofness. Some SCCs become manipulable for weak preferences because a manipulator only needs to be strictly better off under *some* tie-breaking rule whereas he can be indifferent under all other tie-breaking rules. When weakening strategyproofness even further by demanding that a manipulator has to be strictly better off under *any* tie-breaking rule, the positive results for strict preferences are retained.

We have obtained our results using computer-aided theorem proving techniques that were pioneered by Tang and Lin (2009) and have been successfully used to tackle other problems in social choice (see, e.g., Geist and Endriss, 2011; Brandt and Geist, 2016; Brandt et al., 2017; Brandl et al., 2018, 2019). 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 so-called 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, which releases any need to verify our program for generating the proof. In contrast to previous papers, we are even able to give a *lower bound* on the proof complexity: no such proof is possible using less than 19 preference profiles. This can be considered as evidence that it is unlikely that the statement would have been proved without the help of computers, underlining the potential of computer-aided theorem proving in social choice theory.

The remainder of the paper is structured as follows. The setting and basic concepts from social choice theory are introduced in Section 2. Section 3 defines a rigorous model of strategyproofness when ties returned by the SCC are broken by preference-independent tie-breaking functions. This establishes a formal justification for considering the preference extensions attributed to Kelly and Fishburn. Our results and their consequences and limitations are presented in Section 4. Section 5 discusses the relationship between our theorems and related results from the literature. Finally, Section 6 concludes the paper with an overview of our results and a brief discussion of consequences for probabilistic social choice. Appendix A contains a complete human-readable proof of our main theorem and Appendix B provides details on the computer-aided methodology used to obtain the proof.

2. Preliminaries

Let A be a finite set of m alternatives and $N = \{1, \dots, n\}$ a finite set of agents. A *(weak) preference relation* is a complete and transitive binary relation on A. The preference relation of agent i is denoted by $\succcurlyeq_i$, the set of all preference relations by $\mathcal{R}$. We write $\succ_i$ for the strict part of $\succcurlyeq_i$, i.e., $x \succ_i y$ if $x \succcurlyeq_i y$ but not $y \succcurlyeq_i x$, and $\sim_i$ for the indifference part of $\succcurlyeq_i$, i.e., $x \sim_i y$ if $x \succcurlyeq_i y$ and $y \succcurlyeq_i x$. A preference relation $\succcurlyeq_i$ is called *strict* if it additionally is anti-symmetric, i.e., $x \succ_i y$ or $y \succ_i x$ for all distinct alternatives x, y. The set of all strict preference relations is denoted by $\mathcal{L}$. We will compactly represent a preference relation as a comma-separated list where all alternatives among which an agent is indifferent are written as a set. For example $x \succ_i y \sim_i z$ is represented by $\succcurlyeq_i\colon x, \{y, z\}$. 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 correspondences (SCCs)*, i.e., functions that map a preference profile to a set of alternatives called the *choice set*. Formally, an SCC is a function $f\colon \mathcal{R}^N \to \mathcal{C}$ where $\mathcal{C} = 2^A \setminus \emptyset$ denotes the non-empty subsets of A. Note that alternatives are assumed to be mutually exclusive, i.e., choice sets do not refer to committees of candidates, but rather to sets of equally capable candidates, from which a single candidate will eventually be selected.

Given a preference profile R, an alternative x *Pareto-dominates* another alternative y if $x \succcurlyeq_i y$ for all $i \in N$ and $x \succ_j y$ for some $j \in N$. An alternative is *Pareto-optimal* if it is not Pareto-dominated by some other alternative and $PO(R)$ denotes the SCC that returns the set of all Pareto-optimal alternatives. An SCC f is said to be *Pareto-optimal* if $f(R) \subseteq PO(R)$ for all $R \in \mathcal{R}^N$.

Another simple SCC with appealing strategic properties is *serial dictatorship*, which is based on a fixed, but arbitrary, ordering of the agents. First, the set of alternatives is restricted to the ones top-ranked by the first agent. Then, the next agent successively refines the set of alternatives to the set of most preferred alternatives from the remaining set. Formally, serial dictatorship returns $\max_{\succcurlyeq_n} \circ \dots \circ \max_{\succcurlyeq_1}(A)$, where $\max_{\succcurlyeq_i}(X)$ denotes the maximal elements of X according to the preference relation $\succcurlyeq_i$. Serial dictatorship satisfies Pareto-optimality and any reasonable form of strategyproofness, because choosing one's maximal elements is strategyproof for each agent, ruling out any possibility to manipulate. However, serial dictatorship is weakly dictatorial in the sense that it only returns alternatives top-ranked by a pre-determined agent.

Two common symmetry conditions for SCCs are anonymity and neutrality. An SCC 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 SCC is anonymous if $f(R) = f(R')$ for all $R, R' \in \mathcal{R}^N$ and all permutations $\pi\colon N \to N$ such that $\succcurlyeq_i = \succcurlyeq'_{\pi(i)}$ for all $i \in N$. For a permutation $\pi\colon A \to A$ and a preference relation $\succcurlyeq_i$, we define $\succcurlyeq_i^\pi$ as the preference relation where alternatives are renamed according to π, i.e., $\pi(x) \succcurlyeq_i^\pi \pi(y)$ if and only if $x \succcurlyeq_i y$. For a preference profile $R \in \mathcal{R}^N$, R^π denotes the preference profile that maps each $i \in N$ to $\succcurlyeq_i^\pi$. Similarly, π is extended to sets of alternatives $A \in \mathcal{C}$ by letting $\pi(A) = \{\pi(x)\colon x \in A\}$. An SCC f is *neutral* if $f(R^\pi) = \pi(f(R))$ for all $R \in \mathcal{R}^N$ and all permutations $\pi\colon A \to A$. PO is anonymous and neutral while serial dictatorship clearly violates anonymity (while satisfying neutrality).

For a preference profile $R \in \mathcal{R}^N$, let $n_R(x,y) = |\{i \in N : x \succcurlyeq_i y\}|$ be the number of agents who prefer x to y. The majority margin of x over y in R is denoted by $g_R(x,y) = n_R(x,y) - n_R(y,x)$. An SCC 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 SCC only depends on the anonymized comparisons between pairs of alternatives (see, e.g., Young, 1974; Zwicker, 1991). Since majority margins are invariant under permutations of agents, pairwise SCCs are anonymous.[2] When ties are allowed, pairwiseness is slightly stronger than Fishburn's C2, which requires that the SCC only depends on n_R (Fishburn, 1977). This is due to the fact that a pair of opposed preferences affects the majority margin exactly like indifferences do. Hence, PO satisfies C2, but violates pairwiseness because you cannot tell whether x Pareto-dominates y by only looking at $g_R(x,y)$.

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 $g_R(x,y) > 0$ for all $y \in A \setminus \{x\}$. A *Condorcet extension* is an SCC that uniquely returns a Condorcet winner whenever one exists. Almost all Condorcet extensions considered in the literature also happen to be pairwise.

3. A Formal Model of Manipulation Under Tie-Breaking

When defining strategic manipulability for set-valued SCCs, one needs to specify how ties are broken and what the agents know about the tie-breaking mechanism. To this end, we introduce tie-breaking functions g that map choice sets to single alternatives. This provides a clean separation between preference-based selection (by means of a set-valued SCC f) and non-preference-based selection (by means of a tie-breaking function g). While agents are fully informed about f, they only have incomplete information about g. We will distinguish between two different types of tie-breaking functions: deterministic ones (such as a fixed linear tie-breaking ordering) and randomized ones (such as returning an even-chance lottery over selected alternatives).

3.1. Deterministic Tie-Breaking

Let us first consider deterministic tie-breaking functions $g : \mathcal{C} \to A$. If g is known by all agents, this model is equivalent to that of single-valued SCFs and the Gibbard-Satterthwaite theorem applies. We therefore make the weaker—yet reasonable—assumption that agents are unaware of the concrete tie-breaking function g, but only know that g belongs to a certain *class* of functions G. Based on this uncertainty, we define the following strong notion of strategy manipulability (which in turn results in a weak notion of strategyproofness).

[2]Note that, in contrast to other papers, we do not require pairwise SCCs to be neutral (cf. Brandl et al., 2015; Aziz et al., 2018).

An SCC f is *G-manipulable* if there exist preference profiles $R, R' \in \mathcal{R}^N$, and an agent $i \in N$ with $\succcurlyeq_j = \succcurlyeq'_j$ for all $j \neq i$ such that

$$g(f(R')) \succcurlyeq_i g(f(R)) \text{ for all } g \in G \text{ and}$$
$$g(f(R')) \succ_i g(f(R)) \text{ for some } g \in G.$$

f is *G-strategyproof* if it is not G-manipulable.

A very weak notion of strategyproofness is obtained when quantifying over *all* possible tie-breaking functions G_{all} where

$$G_{\text{all}} = \left\{ g \in A^{\mathcal{C}} \colon g(X) \in X \text{ for all } X \in \mathcal{C} \right\}.$$

A natural subset of G_{all} is given by all tie-breaking functions that break ties according to a linear tie-breaking ordering $\geq$ on A (such as the alphabetical ordering). G_{lin} consists of all tie-breaking functions that consistently return the alternative in the choice set that is ranked highest according to some fixed tie-breaking ordering. Formally,

$$G_{\text{lin}} = \{ g \in G_{\text{all}} \colon \text{there is } \geq \in \mathcal{L} \text{ such that for all } X \in \mathcal{C} \text{ and } y \in X,\ g(X) \geq y \}.$$

If $\geq$ is known by the agents, this setting is equivalent to that of single-valued SCFs. By quantifying over all possible linear tie-breaking orderings, we obtain G_{lin}-strategyproofness, which is stronger than G_{all}-strategyproofness. This notion can, for example, be motivated by the existence of a chairman, whose preferences are unknown and who eventually picks a single alternative from the choice set.

3.2. Randomized Tie-Breaking

Another common way to break ties is to utilize randomization. A randomized tie-breaking function is any function $g : \mathcal{C} \to \Delta(A)$ where $\Delta(A)$ is the set of all probability distributions (or lotteries) over alternatives in A. Lotteries are typically compared by assuming the existence of a utility function $u : A \to \mathbb{R}$ that assigns numeric values to alternatives and that is consistent with the agent's ordinal preferences. A utility function u is consistent with a preference relation $\succsim$ if, for all $x, y \in A$, $x \succsim y$ if and only if $u(x) \geq u(y)$. The set of all utility functions consistent with $\succsim$ will be denoted by $U(\succsim)$. With slight abuse of notation, we extend any utility function u to lotteries via *expected* utility. For some class of randomized tie-breaking functions G^{Δ}, an SCC f is *G^{Δ}-manipulable* if there exist preference profiles $R, R' \in \mathcal{R}^N$, and an agent $i \in N$ with $\succcurlyeq_j = \succcurlyeq'_j$ for all $j \neq i$ such that

$$u(g(f(R')) > u(g(f(R))) \text{ for all } g \in G^{\Delta} \text{ and all } u \in U(\succsim_i).$$

G^{Δ}_{all}-strategyproofness is obtained when quantifying over *all* lotteries that randomize among alternatives in the choice set, i.e.,

$$G^{\Delta}_{\text{all}} = \left\{ g \in \Delta(A)^{\mathcal{C}} \colon \text{supp}(g(X)) = X \text{ for all } X \in \mathcal{C} \right\}.$$

A very natural special case of randomized tie-breaking is to randomize uniformly among the alternatives, i.e., G^Δ consists of the unique function *uni*, which maps every choice set to an even-chance lottery over its alternatives. A more general set of tie-breaking functions is obtained by fixing an *a priori* weight function $w : A \to \mathbb{R}_{>0}$ and then assign probabilities to alternatives in the choice set that are proportional to their weights.

$$G^\Delta_{\text{pro}} = \left\{ g \in G^\Delta_{\text{all}} \colon \text{there is } w \in (\mathbb{R}_{>0})^A \text{ such that for all } X \in \mathcal{C},\, g(X) \cong w|_X \right\}.$$

If the weight function assigns the same weight to all alternatives, this corresponds to an even-chance lottery. In Section 6, we discuss generalizations of this model where f directly maps to the set of all lotteries $\Delta(A)$.

3.3. Preference Extensions

As it turns out, each variant of G-manipulation introduced in the previous section can be modeled using so-called *preference extensions* which extend the agents' preferences over alternatives to preferences over sets of alternatives. The three preference extensions considered in this paper are *Kelly's extension*, *Fishburn's extension*, and the *even-chance extension* (Kelly, 1977; Fishburn, 1972; Gärdenfors, 1979). For all $X, Y \in \mathcal{C}$ and $\succcurlyeq_i \in \mathcal{R}$,

$$X \succcurlyeq_i^K Y \text{ iff } x \succcurlyeq_i y \text{ for all } x \in X,\, y \in Y, \tag{Kelly}$$

$$X \succcurlyeq_i^F Y \text{ iff } X \backslash Y \succcurlyeq_i^K Y \text{ and } X \succcurlyeq_i^K Y \backslash X, \text{ and} \tag{Fishburn}$$

$$X \succcurlyeq_i^E Y \text{ iff } |\{x \in X \colon x \succcurlyeq_i z\}|/|X| \geq |\{y \in Y \colon y \succcurlyeq_i z\}|/|Y| \text{ for all } z \in X \cup Y. \tag{Even-chance}$$

The strict parts of these extensions require that at least one of the relations in the definition holds strictly and will be denoted by $\succ_i^K$, $\succ_i^F$, and $\succ_i^E$, respectively. It follows from the definitions that the even-chance extension is a refinement of Fishburn's extension, which in turn is a refinement of Kelly's extension. This inclusion relationship also holds for the strict parts of the three relations. For every $\succcurlyeq_i \in \mathcal{R}$,

$$\succcurlyeq_i^K \subseteq \succcurlyeq_i^F \subseteq \succcurlyeq_i^E \text{ and } \succ_i^K \subseteq \succ_i^F \subseteq \succ_i^E.$$

Even when $m = 3$ (which is sufficient for our impossibility theorems), the even-chance extension is incomplete because the sets $\{b\}$, $\{a, c\}$, and $\{a, b, c\}$ are incomparable with respect to $\succcurlyeq_i \colon a, b, c$. Furthermore, for three alternatives, the only difference between the even-chance extension and Fishburn's extension is that $\{a, b\} \succ_i^E \{a, c\}$ and $\{a, c\} \succ_i^E \{b, c\}$ (while these pairs of sets are incomparable according to Fishburn's extension).

With these extensions at hand, we can formally define strategyproofness of SCCs without making reference to tie-breaking functions. An SCC f is *Kelly-manipulable* if there exist preference profiles $R, R' \in \mathcal{R}^N$, and an agent $i \in N$ such that $\succcurlyeq_j = \succcurlyeq_j'$ for all $j \neq i$ and $f(R') \succ_i^K f(R)$. f is said to satisfy *Kelly-strategyproofness* if it is not Kelly-manipulable. Fishburn-strategyproofness and even-chance-strategyproofness are defined analogously. The relationship between the preference extensions implies that even-chance-strategyproofness is stronger than Fishburn-strategyproofness and that

Fishburn-strategyproofness is stronger than Kelly-strategyproofness. The connection between both strategyproofness notions and the tie-breaking assumptions given in Sections 3.1 and 3.2 is as follows. For any SCC f, we have that

(i) f is Kelly-strategyproof iff it is G_{all}-strategyproof iff it is G_{all}^{Δ}-strategyproof,

(ii) f is Fishburn-strategyproof iff it is G_{lin}-strategyproof iff it is G_{pro}^{Δ}-strategyproof, and

(iii) f is even-chance-strategyproof iff it is $\{uni\}$-strategyproof.

The equivalences between Kelly-strategyproofness and G_{all}-strategyproofness and G_{all}^{Δ}-strategyproofness are relatively straightforward (see, e.g., Erdamar and Sanver (2009, Theorem 3.1) for the former and Gärdenfors (1979, Prop. 5) for the latter). The equivalence of Fishburn-strategyproofness and G_{lin}-strategyproofness is essentially due to the fact that for all $g \in G_{\text{lin}}$ and $X, Y \in \mathcal{C}$, $g(X), g(Y) \in X \cap Y$ implies $g(X) = g(Y)$ (see, e.g., Erdamar and Sanver, 2009, Theorem 3.4). The equivalence between Fishburn-strategyproofness and G_{pro}^{Δ}-strategyproofness was shown by Ching and Zhou (2002, Lemma 1). Even-chance-strategyproofness and $\{uni\}$-strategyproof are equivalent because of standard stochastic dominance arguments (see, e.g., Gärdenfors, 1979).

The more the agents know about the tie-breaking mechanism, the stronger the corresponding notion of strategyproofness. In this sense, Kelly-strategyproofness is the weakest possible notion of strategyproofness because it assumes that agents do not know anything about tie-breaking (except that ties will eventually be broken).

4. Results

We are now ready to state our results. We start by showing that the Pareto rule satisfies Kelly-strategyproofness but violates Fishburn-strategyproofness. This observation will lead to the main theorem, showing that *every* anonymous and Pareto-optimal SCC is Fishburn-manipulable. We then prove that this impossibility remains intact when weakening Fishburn-strategyproofness to Kelly-strategyproofness and restricting attention to the broad class of pairwise SCCs.

4.1. Manipulation of the Pareto Rule

In order to illustrate the definitions of Kelly-strategyproofness and Fishburn-strategyproofness, consider the Pareto rule PO and the following preference profile R.

$$\succcurlyeq_1 : a, \{b, c\} \qquad\qquad \succcurlyeq_2 : \{b, c\}, a$$

Clearly, $PO(R) = \{a, b, c\}$. Now assume that Agent 1 changes his preferences to $\succcurlyeq_1'$ resulting in preference profile R'.

$$\succcurlyeq_1' : a, b, c \qquad\qquad \succcurlyeq_2' : \{b, c\}, a$$

Alternative c is Pareto-dominated by alternative b in R' and $PO(R') = \{a, b\}$. This does not constitute a Kelly-manipulation because

$$\{a, b\} \not\succ_1^K \{a, b, c\}$$

In fact, these sets are incomparable according to $\succ_1^K$. This is in line with our observations from Section 3 because there could be a deterministic tie-breaking function that selects b from $\{a, b\}$ and a from $\{a, b, c\}$. The picture looks different for Fishburn's extension, however, as

$$\{a, b\} \succ_1^F \{a, b, c\}.$$

To see that this concurs with the equivalence of Fishburn's extension and linear tie-breaking orderings, consider the tie-breaking ordering $\geq$ with $c \geq a \geq b$. According to this ordering, a will be selected from $\{a, b\}$ and c from $\{a, b, c\}$. Since $a \succ_1 c$ and for all other tie-breaking orderings, Agent 1 is indifferent between the eventually chosen alternatives, we have a Fishburn-manipulation. $\{a, b\}$ is also preferred to $\{a, b, c\}$ when ties are broken by even-chance lotteries: for all utility functions consistent with $\succcurlyeq_1$, the expected utility for an even-chance lottery between a and b exceeds that of an even-chance lottery between all three alternatives.

The example shows that PO is Fishburn-manipulable (and consequently also even-chance-manipulable). By contrast, as first shown by Feldman (1979), PO does satisfy Kelly-strategyproofness. Since Feldman proves this statement by making reference to stronger strategyproofness notions, we give a self-contained proof below.

Proposition 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 $\succcurlyeq_j = \succcurlyeq'_j$ for all $j \neq i$ and $PO(R') \succ_i^K PO(R)$. 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. This also implies that PO contains at least one top-ranked alternative from every agent because top-ranked alternatives can only be Pareto-dominated by other top-ranked alternatives. Hence, $PO(R')$ contains only top-ranked alternatives of agent i while $PO(R)$ contains at least one alternative that is not top-ranked by agent i. This means that there is some $x \in PO(R) \setminus PO(R')$ and there is no $x' \in PO(R) \setminus \{x\}$ 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 $x \succ_i y$. Since $x \in PO(R)$ and $y \in PO(R')$, it is impossible that $PO(R') \succ_i^K PO(R)$. □

Remark 1 (Refinements of PO). PO is not the most discriminating Kelly-strategyproof SCC. For example, using a proof similar to that of Proposition 1, it can be shown that the SCC that returns all Pareto-optimal alternatives that are top-ranked by at least one agent is Kelly-strategyproof as well.

Remark 2 (Group-strategyproofness). Proposition 1 and Remark 1 also hold when replacing strategyproofness with *group-strategyproofness* where a group of agents can manipulate such that all of them are strictly better off (see, also, Bandyopadhyay, 1983b; Umezawa, 2009).

4.2. Fishburn-strategyproofness

The example given in the previous section shows that PO is Fishburn-manipulable. Our main theorem significantly strengthens this observation by showing that *every anonymous refinement* of PO is Fishburn-manipulable.

Theorem 1. *There is no anonymous SCC that satisfies Pareto-optimality and Fishburn-strategyproofness for $m \geq 3$ and $n \geq 3$.*

The proofs of this and the next theorem are obtained using the computer-aided proving methodology described by Brandt and Geist (2016). The idea is to first manually prove a reduction argument (Lemma 1), which essentially establishes that an impossibility result using Pareto-optimality and strategyproofness can already be shown by proving the incompatibility of these properties for some fixed number of alternatives and agents.

Lemma 1. *Let f be an anonymous SCC that satisfies Pareto-optimality and strategyproofness for A and N. Then there is an anonymous SCC f' that satisfies these axioms for any $A' \subseteq A$ and $N' \subseteq N$.*

Proof. We define an embedding φ of preference profiles $R' = (\succcurlyeq'_1, \dots, \succcurlyeq'_{n'})$ 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: $\varphi(R') = R$ with

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

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 $(R')_{i \mapsto \succcurlyeq'_i}$ are equal to the choice sets of f under the two (extended) profiles R and $R_{i \mapsto \succcurlyeq_i}$, respectively. □

It is easily seen that Lemma 1 also holds for neutral SCCs.

We have used a computer program to generate a proof by contradiction for three agents and three alternatives, which essentially boils down to an extensive case analysis. Similar proofs were found by humans to show impossibility theorems in the context of random assignment and probabilistic social choice (e.g., Bogomolnaia and Moulin, 2001; Bogomolnaia et al., 2005; Brandl et al., 2016; Chang and Chun, 2017; Nesterov, 2017; Aziz et al., 2018; Chun and Yun, 2020). Despite the finiteness of the domain we consider, the number of anonymous SCCs is still enormous (see Table 1), which renders any type of exhaustive search infeasible. There are already about $3.3 \cdot 10^{384}$ possible anonymous SCCs when $m = n = 3$. We therefore formulated the existence of an SCC with the desired properties as a formula in propositional logic and then consulted a state-of-the-art SAT solver, which uses heuristic search algorithms to decide the satisfiability of such formulas. Apart from enabling us to deal with enormous search spaces, the computer-aided approach has the major advantage that related conjectures and hypotheses, e.g.,

m	n	Preference profiles	SCCs	Pareto-optimal SCCs
3	3	455	$\sim 3.3 \cdot 10^{384}$	$\sim 5.0 \cdot 10^{179}$
3	4	1,820	$\sim 1.2 \cdot 10^{1{,}538}$	$\sim 2.8 \cdot 10^{933}$
4	3	73,150	$\sim 1.2 \cdot 10^{86{,}031}$	$\sim 2.2 \cdot 10^{42{,}914}$

Table 1: Number of different profiles and Pareto-optimal SCCs when assuming anonymity.

statements including weaker axioms, can be checked quickly using the same framework. This is reflected in many of the technical remarks given after the theorems.

Starting from the initial, unsatisfiable SAT instance, we used Marco to find a small (group-oriented) MUS, which utilizes the 21 profiles (out of the 455 possible anonymous preference profiles when $m = n = 3$) given in Table 3 of the Appendix. 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 19 profiles with Forqes. This lower bound also implies that for any domain (with $m = n = 3$) that consists of at most 18 preference profiles, there exists some anonymous SCC that satisfies Pareto-optimality and Fishburn-strategyproofness. With the help of PicoSAT, we extracted a proof out of the MUS which is divided into eight main proof steps.[3] This proof has been verified by Isabelle and proof replication data is publicly available (Brandt et al., 2018). The complete human-readable proof of Theorem 1 is given in Appendix A. A more detailed description of the proof methodology is given in Appendix B.

Remark 3 (Independence of axioms). The axioms of Theorem 1 are independent of each other. PO satisfies all axioms except Fishburn-strategyproofness, serial dictatorship satisfies all axioms except anonymity, and the trivial SCC 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 4 (Majoritarian SCCs). Brandt and Geist (2016, Theorem 3) have shown that all Pareto-optimal *majoritarian* SCCs are Fishburn-manipulable when $m \geq 5$ and $n \geq 7$. This result is weaker than Theorem 1, except that it even holds for strict preferences.

Remark 5 (Strict preferences). When assuming strict preferences, there are Fishburn-strategyproof SCCs satisfying Pareto-optimality, e.g., PO or the SCC that returns all top-ranked alternatives (Feldman, 1979; Brandt and Brill, 2011). Theorem 1 shows that these SCCs cannot be extended to weak preferences without giving up one of these desirable properties. The same is true if we instead define Fishburn's extension using strict preferences, i.e., X is preferred to Y if and only if $x \succ_i y$ for all $x \in X \backslash Y$,

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

$y \in Y$ and all $x \in X$, $y \in Y \backslash X$. This preference extension can be obtained via the framework introduced in Section 3 by changing the definition of G-manipulability such that for all $g \in G$, $g(f(R')) \neq g(f(R))$ implies $g(f(R')) \succ_i g(f(R))$.

Remark 6 (Weak Pareto-optimality). *Weak Pareto-optimality* requires that an alternative y should not be selected whenever there is another alternative x such that $x \succ_i y$ for all $i \in N$. Theorem 1 does not hold when replacing Pareto-optimality with weak Pareto-optimality because the SCC that returns all top-ranked alternatives satisfies weak Pareto-optimality and Fishburn-strategyproofness. Note that the SCC that returns all weakly Pareto-optimal alternatives violates Fishburn-strategyproofness. This can be seen by replacing the second agent's preferences in the example given in Section 4.1 with $\succcurlyeq_2 : b, \{a, c\}$.

For the reader's benefit, we now give a simple human-readable proof of a significantly weaker version of Theorem 1 which additionally assumes neutrality.[4] This proof is based on only three preference profiles (rather than 21) and requires only five strategyproofness applications (rather than 89).

Corollary 1. *There is no neutral and anonymous SCC 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 SCC that satisfies Pareto-optimality and Fishburn-strategyproofness. First, consider preference profile R^1.

$$\succcurlyeq_1^1 : a, b, c \qquad\qquad \succcurlyeq_2^1 : b, a, c$$

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 both a and b), this implies $f(R^1) = \{a, b\}$. This already determines the choice set for the following preference profile R^2.

$$\succcurlyeq_1^2 : a, b, c \qquad\qquad \succcurlyeq_2^2 : \{b, c\}, a$$

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.

$$\succcurlyeq_1^3 : a, \{b, c\} \qquad\qquad \succcurlyeq_2^3 : \{b, c\}, a$$

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 $\succcurlyeq_1^3$ to $\succcurlyeq_1^2$. This only leaves $f(R^3) = \{a\}$, which allows the first agent to deviate from $\succcurlyeq_1^2$ to $\succcurlyeq_1^3$, a contradiction. □

[4] Neutrality may seem like an innocuous fairness criterion, but it can be overly restrictive in some settings (see, e.g., Sen, 1970, Section 6.1 (Critique of Anonymity and Neutrality)). In fact, many SCCs used in the real world such as supermajority rules are not neutral.

Remark 7 (Weakening Fishburn-strategyproofness). Using a more complicated proof, it can be shown that Corollary 1 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 1, even when $m = 4$ and $n = 3$ or when $m = 3$ and $n = 4$.

Since even-chance-strategyproofness is stronger than Fishburn-strategyproofness, we get the following immediate corollary of Theorem 1.

Corollary 2. *There is no anonymous SCC that satisfies Pareto-optimality and even-chance-strategyproofness for $m \geq 3$ and $n \geq 3$.*

Proving Corollary 2 is not significantly easier than proving Theorem 1. Forqes provides a lower bound of 15 preference profiles.

4.3. Kelly-strategyproofness

It is not possible to replace Fishburn-strategyproofness with Kelly-strategyproofness in Theorem 1 because PO is Kelly-strategyproof (Proposition 1). We therefore focus on the class of pairwise SCCs (see Section 1), which excludes PO.

Pairwise SCCs constitute a rich and well-studied subclass of SCCs, which contains all *majoritarian* SCCs (i.e., SCCs that only depend on the *sign* of g_R such as Copeland's rule, the top cycle, and the uncovered set). There is a large number of attractive pairwise and Pareto-optimal SCCs (see, e.g., Fishburn, 1977; Fischer et al., 2016). Typical examples are Borda's rule, Kemeny's rule, the Simpson-Kramer rule (aka maximin), Nanson's rule, Schulze's rule, ranked pairs, and the essential set. Virtually all Condorcet extensions considered in the literature are pairwise SCCs.

Theorem 2. *There is no pairwise SCC 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 SCC f that satisfies Pareto-optimality and Kelly-strategyproofness. If not stated otherwise, the absolute values of the majority margins in the following applications of pairwiseness are always 1. First, consider the classic Condorcet profile R^1.

$$\succcurlyeq_1^1 \colon a, b, c \qquad \succcurlyeq_2^1 \colon c, a, b \qquad \succcurlyeq_3^1 \colon b, c, a$$

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

$$\succcurlyeq_1^2 \colon a, b, c \qquad \succcurlyeq_2^2 \colon c, a, b \qquad \succcurlyeq_3^2 \colon b, \{a, c\}$$

R^2 and R^1 only differ in the third agent's preferences. By Kelly-strategyproofness, $b \in f(R^2)$, as otherwise Agent 3 could obtain a preferred choice set by changing his preferences from $\succcurlyeq_3^2$ to $\succcurlyeq_3^1$. Now consider R^3, which has the same majority margins as R^2.

$$\succcurlyeq_1^3 \colon a, b, c \qquad \succcurlyeq_2^3 \colon \{a, c\}, b \qquad \succcurlyeq_3^3 \colon b, c, a$$

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

$$\succcurlyeq_1^4 \colon a, b, c \qquad \succcurlyeq_2^4 \colon a, c, b \qquad \succcurlyeq_3^4 \colon b, c, a$$

R^4 differs from R^3 by the second agent's preferences $\succcurlyeq_2^4$. It follows that $b \in f\big(R^4\big)$. Otherwise, the second agent could misrepresent his preferences $\succcurlyeq_2^3$ by $\succcurlyeq_2^4$ and obtain the Kelly-preferred choice set $f\big(R^4\big)$ without b. Finally, consider R^5.

$$\succcurlyeq_1^5 \colon a, b, c \qquad \succcurlyeq_2^5 \colon \{a, b, c\} \qquad \succcurlyeq_3^5 \colon \{a, b, c\}$$

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

The original proof of Theorem 2 found by the SAT solver consisted of nine preference profiles, and we used FORQES to verify that no proof with less than nine profiles exists. The given proof only argues about five profiles because the first step ("without loss of generality") implicitly makes reference to profiles that are not spelled out explicitly.

Remark 8 (Independence of axioms). The axioms of Theorem 2 are independent of each other. Borda's rule satisfies all axioms except Kelly-strategyproofness, PO satisfies all axioms except pairwiseness, and the trivial SCC 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 9 (Condorcet winners). The conjunction of pairwiseness and Pareto-optimality implies that Condorcet winners should be chosen whenever the pairwise majority relation is transitive and its margins have absolute value 1. We have shown that Theorem 2 also holds when pairwiseness and Pareto-optimality are replaced with this weaker, but technical, condition and $m = 3$ and $n = 4$. Interestingly, the SMUS we found for this statement also consists of nine profiles. Kelly-strategyproof Condorcet extensions have been further explored by Brandt et al. (2022, Remark 12).

Remark 10 (Condorcet losers). For $m = 3$ and $n = 4$, we found a proof consisting of 29 profiles that shows the incompatibility of Kelly-strategyproofness, Pareto-optimality, and the condition that the choice set should not contain a *Condorcet loser*, i.e., an alternative x such that $g_R(x, y) < 0$ for all $y \in A \setminus \{x\}$. However, the former condition does not allow for an induction step (even when paired with Pareto-optimality). This has been rectified using a manual proof by Brandt et al. (2022, Theorem 3).

Remark 11 (BD-strategyproofness). Theorem 2 implies Theorem 6 by Aziz et al. (2018), who use a stronger notion of strategyproofness in the context of probabilistic social choice and furthermore require $m, n \geq 4$.

Remark 12 (Strict preferences). When assuming strict preferences, there are attractive pairwise Kelly-strategyproof SCCs satisfying Pareto-optimality, e.g., the uncovered set, the minimal covering set, and the essential set (Brandt, 2015). Theorem 2 shows that these SCCs 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 using strict preferences, i.e., X is preferred to Y if and only if $x \succ_i y$ for all $x \in X$, $y \in Y$ (Brandt, 2015, Remark 6). This preference extension can be obtained via the framework introduced in Section 3 either by defining G-manipulability using $g(f(R')) \succ_i g(f(R))$ for all $g \in G$ or by changing the definition of G^{Δ}_{all} to $G^{\Delta}_{\text{all}} = \{g \in \Delta(A)^{\mathcal{C}} \colon \text{supp}(g(X)) \subseteq X \text{ for all } X \in \mathcal{C}\}$.

Remark 13 (Weak Pareto-optimality). Theorem 2 does not hold when replacing Pareto-optimality with weak Pareto-optimality (see Remark 6). The SCC that returns all weakly Pareto-optimal alternatives satisfies pairwiseness, weak Pareto-optimality, and Kelly-strategyproofness.

Remark 14 (Dichotomous preferences). When preferences are dichotomous (i.e., each preference relation admits at most two indifference classes), both impossibilities do not hold because approval voting satisfies all desired conditions. PO, however, still violates Fishburn-strategyproofness, which can be seen by considering the following two profiles R and R'.

$\succsim_1 \colon \{b, c\}, a$	$\succsim_2 \colon a, \{b, c\}$	$\succsim_3 \colon a, \{b, c\}$
$\succsim_1' \colon \{b, c\}, a$	$\succsim_2' \colon a, \{b, c\}$	$\succsim_3' \colon \{a, b\}, c$

It is easily verified that $PO(R) = \{a, b, c\}$, $PO(R') = \{a, b\}$, and $\{a, b\} \succsim_3^F \{a, b, c\}$.

5. Discussion of Related Work

There is already a large body of literature dealing with impossibility theorems for strategyproof SCCs and a comparison is in order. From a bird's eye view, existing results require rather restrictive additional assumptions or utilize stronger notions of strategyproofness, but they already hold for the domain of strict individual preferences.

Early results by Barberà (1977a) and Kelly (1977) using Kelly-strategyproofness (or even weaker notions) required SCCs to be *quasi-transitively rationalizable*, a condition which is almost prohibitive on its own (see, e.g., Mas-Colell and Sonnenschein, 1972).[5] MacIntyre and Pattanaik (1981) and Bandyopadhyay (1983a) use similar—albeit slightly weaker—rationalizability conditions such as *minimal binariness* or *quasi-binariness* while Brandt (2015), improving on a result by Gärdenfors (1976) for a strengthening of Fishburn-strategyproofness, uses Condorcet-consistency. Barberà (1977b) restricts attention to *positively responsive* SCCs, which are almost always single-valued. Only very few commonly considered SCCs satisfy positive responsiveness, most notably Borda's rule and some of its variations.

Results with less restrictive assumptions typically require significantly stronger notions of strategyproofness. For example, while Ching and Zhou (2002) also use Fishburn's preference extension like we do in Theorem 1, they require that the original

[5]This is acknowledged by Kelly (1977) who writes that "one plausible interpretation of such a theorem is that, rather than demonstrating the impossibility of reasonable strategy-proof social choice functions, it is part of a critique of the regularity [rationalizability] conditions."

outcome is Fishburn-comparable to the choice set for *any* misrepresentation of preferences. This results in *strong* Fishburn-strategyproofness (see Section 1), which is stronger than Fishburn-strategyproofness and any other form of strategyproofness mentioned in this section. Perhaps the best known result in this area is due to Duggan and Schwartz (2000), who have shown that any non-imposing and non-dictatorial SCC that satisfies a condition they refer to as *residual resoluteness* can be manipulated by an optimist or by a pessimist. The corresponding notion of strategyproofness is stronger than Fishburn-strategyproofness (when preferences are strict) and stronger than Kelly-strategyproofness (without imposing any restrictions on preferences). Their remaining conditions—except residual resoluteness—are weaker than ours: non-imposition is weaker than Pareto-optimality and non-dictatorship is weaker than anonymity. Residual resoluteness is a rather technical condition that is difficult to justify (see, e.g., Rodríguez-Álvarez, 2007; Sato, 2008), and which is, for example, violated by the Pareto rule. To the best of our knowledge, the only impossibility using Fishburn-strategyproofness is a theorem by Brandt and Geist (2016), which only holds for the restricted class of majoritarian SCCs, which form a subset of pairwise SCCs.

The models considered by Barberà et al. (2001), Benoît (2002), Özyurt and Sanver (2009), and Sato (2014) differ from the ones considered so far in that they assume that agents submit *complete* preference relations over sets, subject to certain rationality constraints. Hence, their results can be interpreted as results about single-valued SCFs where the set of alternatives is defined as the set of all non-empty subsets of some set of candidates and the domain of admissible preferences is restricted. Barberà et al. (2001) prove a remarkable analogue of the Gibbard-Satterthwaite theorem for the domain of all preference relations over sets that concur with Fishburn's extension. However, Barberà et al.'s notion of strategyproofness (as well as that of Özyurt and Sanver) is stronger than Fishburn-strategyproofness because it rests on the assumption that all sets are comparable. The strategyproofness notions used by Benoît and Sato are weaker than those of Duggan and Schwartz (2000) and Barberà et al. (2001), but incomparable to both Fishburn- and Kelly-strategyproofness. Their results rely on the comparatively strong assumption that an agent who prefers a to b to c may prefer the set $\{b\}$ to the set $\{a, c\}$ whereas we always deem these sets to be incomparable.[6] Benoît and Sato also employ additional technical conditions called *near unanimity* and *non-decisiveness*, respectively, which are reminiscent of Duggan and Schwartz's residual resoluteness and which are controversial when there are only few agents. For example, near unanimity is violated by Borda's rule for three agents and three alternatives (Sato, 2014).

In contrast to our theorems, the results by Duggan and Schwartz (2000), Barberà et al. (2001), Ching and Zhou (2002), Benoît (2002), Rodríguez-Álvarez (2007), Sato (2008, 2014), and Brandt and Geist (2016) all hold even when preferences are strict.

[6] Within the domain of strict preferences, our proof of Theorem 1 does not make any assumptions on the comparability of sets that is not also made by Benoît (2002) and Sato (2014).

6. Conclusion

We investigated the existence of anonymous, Pareto-optimal, and strategyproof SCCs when there may be ties in the preferences as well as in the outcomes. Our main results are negative and, together with existing positive results, sharpen the boundary of strategyproof social choice. The computer-aided proof of Theorem 1 is quite involved and we have shown that no significantly easier proof exists.

Model	G-Manipulation	Pref. extension	Result
	$\lvert G\rvert = 1$	none	Gibbard (1973) and Satterthwaite (1975)
$\mathcal{R}^N \xrightarrow{f} \mathcal{C} \xrightarrow{g} A$	$G \subseteq G_{\text{lin}}$	Fishburn	Theorem 1
	$G \subseteq G_{\text{all}}$	Kelly	Theorem 2 (pairwise SCCs)
	$G = \{uni\}$	even-chance	Corollary 2
$\mathcal{R}^N \xrightarrow{f} \mathcal{C} \xrightarrow{g} \Delta(A)$	$G \subseteq G^{\Delta}_{\text{pro}}$	Fishburn	Theorem 1
	$G \subseteq G^{\Delta}_{\text{all}}$	Kelly	Theorem 2 (pairwise SCCs)
$\mathcal{R}^N \xrightarrow{f} \Delta(A)$	n/a	stochastic dom.	Brandl et al. (2018) (neutral, *SD*-efficient SDSs)
$U^N \xrightarrow{f} \Delta(A)$	n/a	expected utility	Hylland (1980) (*ex ante* efficient cardinal SDSs)

Table 2: Summary of main results plus two results from probabilistic social choice for comparison. Each line corresponds to an impossibility theorem using anonymity, Pareto-optimality, and strategyproofness.

Our results and their interpretation based on the tie-breaking rules introduced in Section 3 are summarized in Table 2. For comparison, the table also contains the Gibbard-Satterthwaite theorem and two related results from probabilistic social choice, which studies *social decision schemes (SDSs)*, i.e., functions that directly map to the set of all lotteries. When agents submit expected utility functions (or, equivalently, complete preferences over lotteries that adhere to the von Neumann-Morgenstern axioms), one speaks of *cardinal SDSs*. Hylland (1980) has proven that no cardinal SDS satisfies non-dictatorship, strategyproofness, and a strengthening of Pareto-optimality called *ex ante* efficiency. Brandl et al. (2018) showed that this impossibility still holds when agents only submit preferences over alternatives, strategyproofness is weakened to *SD*-strategyproofness, *ex ante* efficiency is weakened to *SD*-efficiency, and non-dictatorship is strengthened to anonymity and neutrality. When comparing this result to Corollary 2, it turns out that Corollary 2 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., Fishburn, 1972). Moreover, it may be difficult to implement non-uniform lotteries in practice.

Both Kelly- and Fishburn-strategyproofness can be translated to the probabilistic setting by applying them to the support of lotteries. For strict preferences, both notions are much weaker than *SD*-strategyproofness. Similarly, SDSs can be translated to the set-valued setting by only considering the support of the resulting lotteries. Thus, Gibbard (1977)'s *random dictatorship*, for example, translates to the *omninomination rule*, which returns all alternatives that are top-ranked by some agent.

It follows from Gibbard's characterization that this rule is Fishburn-strategyproof for strict preferences. Theorem 1 implies that the omninomination rule cannot be extended to weak preferences without giving up Fishburn-strategyproofness or Pareto-optimality. When preferences are weak, Fishburn-strategyproofness is incomparable to *SD*-strategyproofness. In fact, random serial dictatorship (a natural extension of random dictatorship to weak preferences) is *SD*-strategyproof and Pareto-optimal, but violates Fishburn-strategyproofness.[7] Kelly-strategyproofness, on the other hand, remains weaker than *SD*-strategyproofness even when preferences are weak. Hence, Theorem 2 implies that no pairwise SDS can be *SD*-strategyproof and Pareto-optimal at the same time.

Our definitions of Pareto-optimality and strategyproofness crucially rely on the assumption that indifferences are due to equal valuations rather than the absence of comparisons. In the latter case, weak Pareto-optimality (see Remarks 6 and 13), and weaker notions of strategyproofness (see Remarks 5 and 12) are appropriate. Any of these weakenings renders our results moot. The same is true when requiring that individual preferences are strict. For example, the *essential set* (Dutta and Laslier, 1999; Laslier, 2000) and its probabilistic counterpart *maximal lotteries* (Fishburn, 1984; Brandl et al., 2022) satisfy Kelly-strategyproofness when preferences are strict and a weakening of Kelly-strategyproofness when preferences are weak (see Remark 12 and Aziz et al. (2018)).

Acknowledgments

This material is based on work supported by the Deutsche Forschungsgemeinschaft under grants BR 2312/11-1, BR 2312/11-2, and BR 2312/12-1. The authors thank Florian Brandl, Patrick Lederer, and Dominik Peters for helpful discussions and Manuel Eberl for Isabelle support. Preliminary results of this paper were presented at the 5th German Day on Computational Game Theory (Augsburg, February 2018), at the 7th International Workshop on Computational Social Choice (Troy NY, June 2018), and at the 17th International Conference on Autonomous Agents and Multiagent Systems (Stockholm, July 2018).

References

H. Aziz, F. Brandl, F. Brandt, and M. Brill. On the tradeoff between efficiency and strategyproofness. *Games and Economic Behavior*, 110:1–18, 2018.

T. Bandyopadhyay. Manipulation of non-imposed, non-oligarchic, non-binary group decision rules. *Economics Letters*, 11(1–2):69–73, 1983a.

T. Bandyopadhyay. Coalitional manipulation and the Pareto rule. *Journal of Economic Theory*, 29(2):359–363, 1983b.

[7]However, random serial dictatorship violates *SD*-efficiency, because any anonymous, neutral, and *SD*-strategyproof SDS does (Brandl et al., 2018).

S. Barberà. Manipulation of social decision functions. *Journal of Economic Theory*, 15 (2):266–278, 1977a.

S. Barberà. The manipulation of social choice mechanisms that do not leave "too much" to chance. *Econometrica*, 45(7):1573–1588, 1977b.

S. Barberà, B. Dutta, and A. Sen. Strategy-proof social choice correspondences. *Journal of Economic Theory*, 101(2):374–394, 2001.

J.-P. Benoît. Strategic manipulation in voting games when lotteries and ties are permitted. *Journal of Economic Theory*, 102(2):421–436, 2002.

A. Biere. PicoSAT essentials. *Journal on Satisfiability, Boolean Modeling and Computation*, 4:75–79, 2008.

A. Biere. Lingeling, Plingeling and Treengeling entering the SAT competition 2013. In *Proceedings of the SAT Competition 2013*, pages 51–52, 2013.

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

A. Bogomolnaia, H. Moulin, and R. Stong. Collective choice under dichotomous preferences. *Journal of Economic Theory*, 122(2):165–184, 2005.

F. Brandl, F. Brandt, and J. Hofbauer. Incentives for participation and abstention in probabilistic social choice. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1411–1419, 2015.

F. Brandl, F. Brandt, and W. Suksompong. The impossibility of extending random dictatorship to weak preferences. *Economics Letters*, 141:44–47, 2016.

F. Brandl, F. Brandt, M. Eberl, and C. Geist. Proving the incompatibility of efficiency and strategyproofness via SMT solving. *Journal of the ACM*, 65(2):1–28, 2018.

F. Brandl, F. Brandt, C. Geist, and J. Hofbauer. Strategic abstention based on preference extensions: Positive results and computer-generated impossibilities. *Journal of Artificial Intelligence Research*, 66:1031–1056, 2019.

F. Brandl, F. Brandt, and C. Stricker. An analytical and experimental comparison of maximal lottery schemes. *Social Choice and Welfare*, 58(1):5–38, 2022.

F. Brandt. Set-monotonicity implies Kelly-strategyproofness. *Social Choice and Welfare*, 45(4):793–804, 2015.

F. Brandt and M. Brill. 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)*, pages 136–142, 2011.

F. Brandt and C. Geist. Finding strategyproof social choice functions via SAT solving. *Journal of Artificial Intelligence Research*, 55:565–602, 2016.

F. Brandt and P. Lederer. Characterizing the top cycle via strategyproofness. 2022. Working paper.

F. Brandt, C. Geist, and D. Peters. Optimal bounds for the no-show paradox via SAT solving. *Mathematical Social Sciences*, 90:18–27, 2017. Special Issue in Honor of Hervé Moulin.

F. Brandt, M. Eberl, C. Saile, and C. Stricker. The incompatibility of Fishburn-strategyproofness and Pareto-efficiency. *Archive of Formal Proofs*, 2018. ISSN 2150-914x. URL `http://isa-afp.org/entries/Fishburn_Impossibility.html`.

F. Brandt, M. Bullinger, and P. Lederer. On the indecisiveness of Kelly-strategyproof social choice functions. *Journal of Artificial Intelligence Research*, 2022. Forthcoming.

H.-I. Chang and Y. Chun. Probabilistic assignment of indivisible objects when agents have the same preferences except the ordinal ranking of one object. *Mathematical Social Sciences*, 90:80–82, 2017.

S. Ching and L. Zhou. Multi-valued strategy-proof social choice rules. *Social Choice and Welfare*, 19(3):569–580, 2002.

Y. Chun and K. Yun. Upper-contour strategy-proofness in the probabilistic assignment problem. *Social Choice and Welfare*, 54(4):667–687, 2020.

J. Duggan and T. Schwartz. Strategic manipulability without resoluteness or shared beliefs: Gibbard-Satterthwaite generalized. *Social Choice and Welfare*, 17(1):85–93, 2000.

B. Dutta and J.-F. Laslier. Comparison functions and choice correspondences. *Social Choice and Welfare*, 16(4):513–532, 1999.

B. Erdamar and M. R. Sanver. Choosers as extension axioms. *Theory and Decision*, 67 (4):375–384, 2009.

A. Feldman. Manipulation and the Pareto rule. *Journal of Economic Theory*, 21:473–482, 1979.

F. Fischer, O. Hudry, and R. Niedermeier. Weighted tournament solutions. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 4. Cambridge University Press, 2016.

P. C. Fishburn. Even-chance lotteries in social choice theory. *Theory and Decision*, 3 (1):18–40, 1972.

P. C. Fishburn. Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, 33(3):469–489, 1977.

P. C. Fishburn. Probabilistic social choice based on simple voting comparisons. *Review of Economic Studies*, 51(4):683–692, 1984.

P. Gärdenfors. Manipulation of social choice functions. *Journal of Economic Theory*, 13 (2):217–228, 1976.

P. Gärdenfors. On definitions of manipulation of social choice functions. In J. J. Laffont, editor, *Aggregation and Revelation of Preferences*. North-Holland, 1979.

C. Geist and U. Endriss. Automated search for impossibility theorems in social choice theory: Ranking sets of objects. *Journal of Artificial Intelligence Research*, 40:143–174, 2011.

C. Geist and D. Peters. Computer-aided methods for social choice theory. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 13, pages 249–267. AI Access, 2017.

A. Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4): 587–601, 1973.

A. Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45 (3):665–681, 1977.

A. Hylland. Strategyproofness of voting procedures with lotteries as outcomes and infinite sets of strategies. Mimeo, 1980.

A. Ignatiev, A. Previti, M. Liffiton, and J. Marques-Silva. Smallest MUS extraction with minimal hitting set dualization. In *Proceedings of the 21st International Conference on Principles and Practice of Constraint Programming*, volume 9255 of *Lecture Notes in Computer Science (LNCS)*, pages 173–182. Springer-Verlag, 2015.

J. S. Kelly. Strategy-proofness and social choice functions without single-valuedness. *Econometrica*, 45(2):439–446, 1977.

J.-F. Laslier. Aggregation of preferences with a variable set of alternatives. *Social Choice and Welfare*, 17(2):269–282, 2000.

M. H. Liffiton and K. A. Sakallah. Algorithms for computing minimal unsatisfiable subsets of constraints. *Journal of Automated Reasoning*, 40(1):1–33, 2008.

M. H. Liffiton, A. Previti, A. Malik, and J. Marques-Silva. Fast, flexible MUS enumeration. *Constraints*, 21(2):223–250, 2016.

I. MacIntyre and P. K. Pattanaik. Strategic voting under minimally binary group decision functions. *Journal of Economic Theory*, 25(3):338–352, 1981.

F. Marić. Formal verification of a modern SAT solver by shallow embedding into Isabelle/HOL. *Theoretical Computer Science*, 411(50):4333 – 4356, 2010.

A. Mas-Colell and H. Sonnenschein. General possibility theorems for group decisions. *Review of Economic Studies*, 39(2):185–192, 1972.

K. Nehring. Monotonicity implies generalized strategy-proofness for correspondences. *Social Choice and Welfare*, 17(2):367–375, 2000.

A. Nesterov. Fairness and efficiency in strategy-proof object allocation mechanisms. *Journal of Economic Theory*, 170:145–168, 2017.

T. Nipkow, L. C. Paulson, and M. Wenzel. *Isabelle/HOL – A Proof Assistant for Higher-Order Logic*, volume 2283 of *Lecture Notes in Computer Science (LNCS)*. Springer-Verlag, 2002.

S. Özyurt and M. R. Sanver. A general impossibility result on strategy-proof social choice hyperfunctions. *Games and Economic Behavior*, 66(2):880–892, 2009. Special Section in Honor of David Gale.

C. Rodríguez-Álvarez. On the manipulation of social choice correspondences. *Social Choice and Welfare*, 29(2):175–199, 2007.

S. Sato. On strategy-proof social choice correspondences. *Social Choice and Welfare*, 31: 331–343, 2008.

S. Sato. A fundamental structure of strategy-proof social choice correspondences with restricted preferences over alternatives. *Social Choice and Welfare*, 42(4):831–851, 2014.

M. A. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, 1975.

A. K. Sen. *Collective Choice and Social Welfare*. North-Holland, 1970.

P. Tang and F. Lin. Computer-aided proofs of Arrow's and other impossibility theorems. *Artificial Intelligence*, 173(11):1041–1053, 2009.

A. D. Taylor. *Social Choice and the Mathematics of Manipulation*. Cambridge University Press, 2005.

M. Umezawa. Coalitionally strategy-proof social choice correspondences and the Pareto rule. *Social Choice and Welfare*, 33:151–158, 2009.

H. P. Young. An axiomatization of Borda's rule. *Journal of Economic Theory*, 9(1): 43–52, 1974.

W. S. Zwicker. The voter's paradox, spin, and the Borda count. *Mathematical Social Sciences*, 22(3):187–227, 1991.

APPENDIX

A. Proof of Theorem 1

Proof. Let f be a Fishburn-strategyproof and Pareto-optimal SCC. A recurrent step in the proof is to argue that, for two preference profiles $R, R' \in \mathcal{R}^N$ and some agent $i \in N$ such that $\succcurlyeq_j = \succcurlyeq'_j$ for all $j \neq i$, $f(R) = \{x\}$ implies that $f(R') = \{y\}$ for some $y \in A$. In particular, it can be shown that the following four implications hold when $f(R) = \{x\}$ and $m = 3$.

(i) If agent i top-ranks x, possibly with other Pareto-dominated alternatives, in R', then $f(R') = \{x\}$.

(ii) If agent i does not top-rank x in R and some alternative is Pareto-dominated both in R and R', then $f(R') = \{x\}$.

(iii) If agent i's unique least-preferred alternative in R is x, then $f(R') = \{x\}$.

(iv) If y Pareto-dominates x in R' and agent i top-ranks exactly x and y in R', then $f(R') = \{y\}$.

All four implications already hold for Kelly's extension and are easily proved. Implications *(i)* and *(iv)* make use of strategyproofness from R' to R: if agent i ranks alternative x top in R' and if $f(R) = \{x\}$, then strategyproofness implies that $f(R')$ must consist only of top ranked alternatives of agent i in R'. Otherwise, agent i can get a strictly preferred outcome by changing his preferences from R' to R.
For implication *(ii)*, observe that there is some alternative y in the top indifference class of agent i in R, while, by assumption, alternative x is not top ranked. Since $f(R) = \{x\}$, strategyproofness with respect to agent i from R to R' implies that neither $f(R') = \{x, y\}$ nor $f(R') = \{y\}$. By Pareto-optimality, we have that the remaining alternative $z \notin f(R')$, therefore we have that $f(R') = \{x\}$.
Lastly, implication *(iii)* is an immediate consequence of strategyproofness from R to R': Any outcome other than the singleton consisting of the last ranked alternative would be strictly preferred by agent i.

Now, for the proof of the theorem, let $N = \{1, 2, 3\}$ and $A = \{a, b, c\}$ and consider the preference profiles given in Table 3. We determine restrictions on the outcome of f imposed by Pareto-optimality and Fishburn-strategyproofness until we arrive at a contradiction. We will use the four implications from above by referring to their numbers and instantiating with the appropriate preference profiles. To improve readability, it will not be explicitly mentioned that Pareto-dominated alternatives, which are marked in gray, must not be contained in any choice set. Anonymity is used implicitly by identifying preference profiles with *sets* of individual preference relations.

Interestingly, Kelly-strategyproofness suffices for the entire proof except for two implications involving the possibility that $f(R^2) = \{a, b, c\}$ in the proof of Claim (2). We will explicitly mention that Fishburn's extension is used in these cases. It turns out that

Profile	Preference relations		
R^1	$\{a, c\}, b$	a, b, c	b, c, a
R^2	$c, \{a, b\}$	a, b, c	b, c, a
R^3	$c, \{a, b\}$	a, b, c	c, b, a
R^4	$c, \{a, b\}$	b, c, a	c, b, a
R^5	$\{a, c\}, b$	$\{a, c\}, b$	b, c, a
R^6	$b, \{a, c\}$	$c, \{a, b\}$	b, c, a
R^7	$\{a, c\}, b$	$b, \{a, c\}$	b, c, a
R^8	$\{b, c\}, a$	$a, \{b, c\}$	a, c, b
R^9	$\{b, c\}, a$	$b, \{a, c\}$	a, b, c
R^{10}	$\{a, c\}, b$	b, c, a	c, b, a
R^{11}	$\{a, c\}, b$	a, b, c	c, b, a
R^{12}	$c, \{a, b\}$	b, a, c	c, b, a
R^{13}	$\{a, c\}, b$	b, a, c	c, b, a
R^{14}	$a, \{b, c\}$	$c, \{a, b\}$	a, c, b
R^{15}	$\{b, c\}, a$	$a, \{b, c\}$	a, b, c
R^{16}	$\{a, b\}, c$	$c, \{a, b\}$	a, b, c
R^{17}	$a, \{b, c\}$	a, b, c	b, c, a
R^{18}	$\{a, c\}, b$	$b, \{a, c\}$	b, a, c
R^{19}	$a, \{b, c\}$	$c, \{a, b\}$	a, b, c
R^{20}	$b, \{a, c\}$	a, b, c	b, a, c
R^{21}	$\{b, c\}, a$	a, b, c	b, c, a

Table 3: Anonymous preference profiles for three agents used in the proof of Theorem 1. Pareto-dominated alternatives are marked in gray.

only profiles R^1 to R^4 are directly affected in the results of the main proof steps, which can be broken down into eight claims. The rest of the profiles are needed in intermediate steps.

(1) $f(R^1) \neq \{a\}$

Assume for contradiction that $f(R^1) = \{a\}$. The following chain of implications shows that this entails $f(R^{12}) = \{c\}$. For convenience, we restate the preference profiles involved in these implications on the right-hand side. Preference relations that changed from one profile to another are highlighted in gray.

	$f(R^1) = \{a\}$	a, b, c	$\{a, c\}, b$	b, c, a
$\overset{(iv)}{\Longrightarrow}$	$f(R^5) = \{c\}$	$\{a, c\}, b$	$\{a, c\}, b$	b, c, a
$\overset{(i)}{\Longrightarrow}$	$f(R^{10}) = \{c\}$	c, b, a	$\{a, c\}, b$	b, c, a
$\overset{(i)}{\Longrightarrow}$	$f(R^4) = \{c\}$	c, b, a	$c, \{a, b\}$	b, c, a
$\overset{(ii)}{\Longrightarrow}$	$f(R^{12}) = \{c\}$	c, b, a	$c, \{a, b\}$	b, a, c

From $f(R^{12}) = \{c\}$ we can infer that $f(R^{13}) \subseteq \{a, c\}$ because otherwise there is a manipulation from R^{13} to R^{12}.

	$f(R^{12}) = \{c\}$	$c, \{a, b\}$	b, a, c	c, b, a
$\Rightarrow$	$f(R^{13}) \subseteq \{a, c\}$	$\{a, c\}, b$	b, a, c	c, b, a

Further, $f(R^1) = \{a\}$ implies that $f(R^{11}) = \{a\}$.

	$f(R^1) = \{a\}$	b, c, a	$\{a, c\}, b$	a, b, c
$\overset{(iii)}{\Longrightarrow}$	$f(R^{11}) = \{a\}$	c, b, a	$\{a, c\}, b$	a, b, c

Now $f(R^{13}) \subseteq \{a, c\}$ and $f(R^{11}) = \{a\}$ imply $f(R^{13}) = \{a\}$, since both $f(R^{13}) = \{c\}$ and $f(R^{13}) = \{a, c\}$ would violate strategyproofness from R^{13} to R^{11}.

	$f(R^{11}) = \{a\}$	a, b, c	$\{a, c\}, b$	c, b, a
$\Rightarrow$	$f(R^{13}) = \{a\}$	b, a, c	$\{a, c\}, b$	c, b, a

From $f(R^{13}) = \{a\}$ we deduce that $f(R^2) = \{a\}$ using the following chain of implications.

	$f(R^{13}) = \{a\}$	c, b, a	$\{a, c\}, b$	b, a, c
$\overset{(iii)}{\Longrightarrow}$	$f(R^{18}) = \{a\}$	$b, \{a, c\}$	$\{a, c\}, b$	b, a, c
$\overset{(i)}{\Longrightarrow}$	$f(R^{20}) = \{a\}$	$b, \{a, c\}$	a, b, c	b, a, c
$\overset{(ii)}{\Longrightarrow}$	$f(R^9) = \{a\}$	$b, \{a, c\}$	a, b, c	$\{b, c\}, a$
$\overset{(ii)}{\Longrightarrow}$	$f(R^{21}) = \{a\}$	b, c, a	a, b, c	$\{b, c\}, a$
$\overset{(iii)}{\Longrightarrow}$	$f(R^2) = \{a\}$	b, c, a	a, b, c	$c, \{a, b\}$

Also, $f(R^{12}) = \{c\}$ implies $f(R^3) = \{c\}$.

	$f(R^{12}) = \{c\}$	b, a, c	$c, \{a, b\}$	c, b, a
$\overset{(iii)}{\Longrightarrow}$	$f(R^3) = \{c\}$	a, b, c	$c, \{a, b\}$	c, b, a

However, $f(R^3) = \{c\}$ and $f(R^2) = \{a\}$ violate strategyproofness from R^2 to R^3.

$f(R^2) = \{a\}$	b, c, a	$c, \{a, b\}$	a, b, c
$f(R^3) = \{c\}$	c, b, a	$c, \{a, b\}$	a, b, c

Hence, the assumption that $f(R^1) = \{a\}$ was incorrect.

(2) $f(R^2) \notin \{\{b\}, \{a, b\}, \{a, b, c\}\}$

Assume for contradiction that (2) is false. This implies $f(R^{16}) = \{a\}$, since $f(R^{16}) = \{a, c\}$ and $f(R^{16}) = \{c\}$ would violate Fishburn-strategyproofness from R^{16} to R^2.

$f(R^2) \in \{\{b\}, \{a, b\}, \{a, b, c\}\}$	b, c, a	$c, \{a, b\}$	a, b, c
$\Rightarrow f(R^{16}) = \{a\}$	$\{a, b\}, c$	$c, \{a, b\}$	a, b, c

Now $f(R^{16}) = \{a\}$ implies $f(R^{17}) = \{a\}$.

$f(R^{16}) = \{a\}$	$\{a, b\}, c$	a, b, c	$c, \{a, b\}$
$\overset{(i)}{\Longrightarrow} f(R^{19}) = \{a\}$	$a, \{b, c\}$	a, b, c	$c, \{a, b\}$
$\overset{(i)}{\Longrightarrow} f(R^{14}) = \{a\}$	$a, \{b, c\}$	a, c, b	$c, \{a, b\}$
$\overset{(ii)}{\Longrightarrow} f(R^{8}) = \{a\}$	$a, \{b, c\}$	a, c, b	$\{b, c\}, a$
$\overset{(i)}{\Longrightarrow} f(R^{15}) = \{a\}$	$a, \{b, c\}$	a, b, c	$\{b, c\}, a$
$\overset{(iii)}{\Longrightarrow} f(R^{17}) = \{a\}$	$a, \{b, c\}$	a, b, c	b, c, a

From (1) we know that $f(R^1) \neq \{a\}$, therefore $f(R^{17}) = \{a\}$ yields either $f(R^1) = \{c\}$ or $f(R^1) = \{a, c\}$ by strategyproofness from R^1 to R^{17}.

$f(R^{17}) = \{a\}$	$a, \{b, c\}$	a, b, c	b, c, a
$\Rightarrow f(R^1) \in \{\{c\}, \{a, c\}\}$	$\{a, c\}, b$	a, b, c	b, c, a

However, $f(R^1) \in \{\{c\}, \{a, c\}\}$ contradicts the assumption that $f(R^2)$ is either $\{b\}$, $\{a, b\}$, or $\{a, b, c\}$. In each of the cases, Fishburn-strategyproofness from R^2 to R^1 is violated.

$f(R^2) \in \{\{b\}, \{a, b\}, \{a, b, c\}\}$	$c, \{a, b\}$	a, b, c	b, c, a
$f(R^1) \in \{\{c\}, \{a, c\}\}$	$\{a, c\}, b$	a, b, c	b, c, a

(3) $f(R^1) \neq \{c\}$

Assume for contradiction that $f(R^1) = \{c\}$. The following chain of implications shows that this implies $f(R^7) = \{b\}$.

	$f(R^1) = \{c\}$	$\{a, c\}, b$	b, c, a	a, b, c
$\overset{(iv)}{\Longrightarrow}$	$f(R^{21}) = \{b\}$	$\{b, c\}, a$	b, c, a	a, b, c
$\overset{(i)}{\Longrightarrow}$	$f(R^9) = \{b\}$	$\{b, c\}, a$	$b, \{a, c\}$	a, b, c
$\overset{(i)}{\Longrightarrow}$	$f(R^{20}) = \{b\}$	b, a, c	$b, \{a, c\}$	a, b, c
$\overset{(ii)}{\Longrightarrow}$	$f(R^{18}) = \{b\}$	b, a, c	$b, \{a, c\}$	$\{a, c\}, b$
$\overset{(i)}{\Longrightarrow}$	$f(R^7) = \{b\}$	b, c, a	$b, \{a, c\}$	$\{a, c\}, b$

However, $f(R^7) = \{b\}$ and $f(R^1) = \{c\}$ contradicts strategyproofness from R^1 to R^7.

$f(R^1) = \{c\}$	a, b, c	$\{a, c\}, b$	b, c, a
$f(R^7) = \{b\}$	$b, \{a, c\}$	$\{a, c\}, b$	b, c, a

(4) $f(R^2) \notin \{\{c\}, \{b, c\}\}$

Assume for contradiction that (4) does not hold. This implies $f(R^{21}) = \{b\}$, since $a \in f(R^{21})$ would violate strategyproofness from R^{21} to R^2.

	$f(R^2) \in \{\{c\}, \{b, c\}\}$	$c, \{a, b\}$	a, b, c	b, c, a
$\Rightarrow$	$f(R^{21}) = \{b\}$	$\{b, c\}, a$	a, b, c	b, c, a

Now $f(R^{21}) = \{b\}$ implies $f(R^6) = \{b\}$:

	$f(R^{21}) = \{b\}$	b, c, a	$\{b, c\}, a$	a, b, c
$\overset{(i)}{\Longrightarrow}$	$f(R^9) = \{b\}$	$b, \{a, c\}$	$\{b, c\}, a$	a, b, c
$\overset{(i)}{\Longrightarrow}$	$f(R^{20}) = \{b\}$	$b, \{a, c\}$	b, a, c	a, b, c
$\overset{(ii)}{\Longrightarrow}$	$f(R^{18}) = \{b\}$	$b, \{a, c\}$	b, a, c	$\{a, c\}, b$
$\overset{(i)}{\Longrightarrow}$	$f(R^7) = \{b\}$	$b, \{a, c\}$	b, c, a	$\{a, c\}, b$
$\overset{(iii)}{\Longrightarrow}$	$f(R^6) = \{b\}$	$b, \{a, c\}$	b, c, a	$c, \{a, b\}$

However, $f(R^6) = \{b\}$ and the assumption that $f(R^2)$ is either $\{c\}$ or $\{b, c\}$ contradicts strategyproofness from R^2 to R^6 in either case.

	$f(R^2) \in \{\{c\}, \{b, c\}\}$	a, b, c	$c, \{a, b\}$	b, c, a
$\Rightarrow$	$f(R^6) = \{b\}$	$b, \{a, c\}$	$c, \{a, b\}$	b, c, a

(5) $\underline{b \notin f(R^1)}$

Assume for contradiction that (5) does not hold. From (2) and (4) we deduce that either $f(R^2) = \{a\}$ or $f(R^2) = \{a, c\}$. This contradicts strategyproofness from R^1 to R^2.

$f(R^1) \ni b$	$\{a,c\}, b$	a, b, c	b, c, a
$f(R^2) \in \{\{a\}, \{a,c\}\}$	$c, \{a,b\}$	a, b, c	b, c, a

(6) $\underline{f(R^3) \neq \{c\}}$

Assume for contradiction that $f(R^3) = \{c\}$. From (2) and (4) we have again that either $f(R^2) = \{a\}$ or $f(R^2) = \{a, c\}$. In both cases, strategyproofness from R^2 to R^3 is violated.

$f(R^2) \in \{\{a\}, \{a,c\}\}$	b, c, a	$c, \{a,b\}$	a, b, c
$f(R^3) = \{c\}$	c, b, a	$c, \{a,b\}$	a, b, c

(7) $\underline{f(R^4) \neq \{c\}}$

Assume for contradiction that $f(R^4) = \{c\}$. This implies $f(R^3) = \{c\}$, a contradiction to (6).

	$f(R^4) = \{c\}$	b, c, a	$c, \{a,b\}$	c, b, a
$\overset{(ii)}{\Longrightarrow}$	$f(R^{12}) = \{c\}$	b, a, c	$c, \{a,b\}$	c, b, a
$\overset{(iii)}{\Longrightarrow}$	$f(R^3) = \{c\}$	a, b, c	$c, \{a,b\}$	c, b, a

(8) $\underline{f(R^4) = \{c\}}$

From (1), (3), and (5) we know that $f(R^1) = \{a, c\}$. This implies $f(R^5) = \{c\}$, as otherwise strategyproofness from R^5 to R^1 would be violated.

	$f(R^1) = \{a,c\}$	a, b, c	$\{a,c\}, b$	b, c, a
$\Rightarrow$	$f(R^5) = \{c\}$	$\{a,c\}, b$	$\{a,c\}, b$	b, c, a

Now $f(R^5) = \{c\}$ implies $f(R^4) = \{c\}$.

	$f(R^5) = \{c\}$	$\{a,c\}, b$	$\{a,c\}, b$	b, c, a
$\overset{(i)}{\Longrightarrow}$	$f(R^{10}) = \{c\}$	c, b, a	$\{a,c\}, b$	b, c, a
$\overset{(i)}{\Longrightarrow}$	$f(R^4) = \{c\}$	c, b, a	$c, \{a,b\}$	b, c, a

Since $f(R^4) = \{c\}$ and $f(R^4) \neq \{c\}$ is a contradiction, (8) and (7) conclude the proof, showing that no Fishburn-strategyproof and Pareto-optimal SCC exists. $\square$

B. Computer-Aided Theorem Proving

Basically, the core of the computer-aided approach is the encoding of the problems to be solved as 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 $c_{R,X}$ where R is a preference profile and $X \subseteq A$ a set of alternatives. The semantics of these variables are that $c_{R,X}$ is true if and only if $f(R) = X$, i.e., the SCC 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*,

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

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 for some preference extension $\mathcal{E}$ can be encoded as

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

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

After encoding the axioms using a Java program, satisfiability of the SAT instance is checked with the LINGELING solver family by Biere (2013). 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. Any unsatisfiable subset of clauses is an *unsatisfiable core*. 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, neither an MUS nor a smallest unsatisfiable core has to be unique.

Especially with regard to proof extraction later on, we aimed at finding a smallest minimal unsatisfiable set (SMUS), for which we used the software tool Marco by Liffiton et al. (2016). Rather than merely minimizing the number of *clauses* of the CNF formula, we are interested in proofs that minimize the number of required *preference profiles*. 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 (2008) for more details on group-oriented SAT solving).

Moreover, using the group-oriented approach, we can now also give lower bounds for the number of profiles needed in impossibility proofs. The number of profiles seems to be a reasonable measure of proof complexity, even though it is, of course, very well possible that proofs using more profiles turn out to be "easier," e.g., by requiring fewer case distinctions. We achieve the lower bound with the tool Forqes by Ignatiev et al. (2015), as it supports a restricted version of group-oriented SAT solving, namely the specification of *don't care* clauses. 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 an MUS of a given size starting with the trivial size of just one clause (and finally returns an SMUS if not aborted prematurely).

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 (2008). 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 statements about preference profiles. For this, we use a dictionary containing the correspondences between SAT variables and tuples consisting of preference profiles and choice sets.

If computer-generated proofs exceed a certain 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 program that generates the CNF formula. The first concern is less problematic and is addressed by using a verified SAT solver (Marić, 2010). 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 program for encoding the axioms. To tackle this issue, we make use of the interactive theorem prover Isabelle (see, e.g., Nipkow et al., 2002) 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. (2018), the set of preference profiles and conditions obtained from the MUS is translated to higher-order logic 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 even be checked manually step by step. Trustworthiness of Isabelle is considerably high as it is widely used for verification tasks.[8] For more background on computer-aided theorem proving in the context of social choice theory, the reader is referred to Geist and Peters (2017).

[8]Using higher-order proof assistants such as Isabelle to prove these theorems in the first place is currently completely out of reach due to performance limitations.