Strategic Abstention based on Preference Extensions: Positive Results and Computer-Generated Impossibilities

Florian Brandl Felix Brandt Christian Geist Johannes Hofbauer
Technische Universität München
85748 Garching bei München, Germany
{brandlfl,brandtf,geist,hofbauej}@in.tum.de

Abstract

Voting rules are powerful tools that allow multiple agents to aggregate their preferences in order to reach joint decisions. A common flaw of some voting rules, known as the *no-show paradox*, is that agents may obtain a more preferred outcome by abstaining from an election. We study strategic abstention for set-valued voting rules based on Kelly's and Fishburn's preference extensions. Our contribution is twofold. First, we show that, whenever there are at least five alternatives, *every* Pareto-optimal majoritarian voting rule suffers from the no-show paradox with respect to Fishburn's extension. This is achieved by reducing the statement 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 also provide a human-readable proof which we extracted from a minimal unsatisfiable core of the formula. Secondly, we prove that every voting rule that satisfies two natural conditions cannot be manipulated by strategic abstention with respect to Kelly's extension. We conclude by giving examples of well-known Pareto-optimal majoritarian voting rules that meet these requirements.

1 Introduction

Whenever a group of multiple agents aims at reaching a joint decision in a fair and satisfactory way, they need to aggregate their (possibly conflicting) preferences. 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 [Brandt *et al.*, 2013].

A common flaw of many such rules, first observed by Fishburn and Brams [1983], who called it the *no-show paradox*, is that agents may obtain a more preferred outcome by abstaining from an election. Following Moulin [1988], a voting rule is said to satisfy *participation* if it is immune to the no-show paradox. Moulin has shown that all resolute, i.e., single-valued, scoring rules (such as Borda's rule) satisfy participation while all resolute Condorcet extensions suffer from the no-show paradox. Condorcet extensions comprise a large class of voting rules that satisfy otherwise rather desirable properties.

In this paper, we study participation for *irresolute*, i.e., set-valued, *social choice functions (SCFs)*. A proper definition of participation for irresolute SCFs requires the specification of preferences over *sets* of alternatives. Rather than asking the agents to specify their preferences over all subsets (which would be bound to various rationality constraints and require exponential space), it is typically assumed that the preferences over single alternatives can be extended to preferences over sets. Of course, there are various ways how to extend preferences to sets (see, e.g., [Gärdenfors, 1979; Barberà *et al.*, 2004]), each of which leads to a different version of participation. A function that yields a (possibly incomplete) preference relation over subsets of alternatives when given a preference relation over single alternatives is called a *preference extension*. In this paper, we focus on two common preference extensions due to Kelly [1977] and Fishburn [1972], both of which arise under natural assumptions about the agents' knowledge of the tie-breaking mechanism that eventually picks a single alternative from the choice set (see, e.g., [Gärdenfors, 1979; Ching and Zhou, 2002; Sanver and Zwicker, 2012; Brandt and Brill, 2011; Brandt, 2015]). Kelly's extension, for example, can be motivated by assuming that the agents possess no information whatsoever about the tie-breaking mechanism. A common interpretation of Fishburn's extension, on the other hand, is that ties are broken according to the unknown preferences of a chairman. Since Fishburn's extension is a refinement of Kelly's extension it follows that Fishburn-participation is stronger than Kelly-participation. The idea pursued in this paper is to exploit the uncertainty of the agents about the tie-breaking mechanism in order to prevent strategic abstention. Our two main results are as follows.

- Whenever there are at least four alternatives, Pareto-optimality and Fishburn-participation are incompatible in the context of majoritarian SCFs. When there are at least five alternatives, this even holds for strict preferences.
- Every SCF that satisfies set-monotonicity and independence of indifferent voters satisfies Kelly-participation. Every set-monotonic majoritarian SCF satisfies Kelly-participation when preferences are strict.

The first result is obtained using computer-aided theorem proving techniques. In particular, we reduce the statement 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 also provide a human-readable proof for this result, which we extracted from a minimal unsatisfiable core of the SAT formula.

The conditions for the second result are easy to check and satisfied by a small number of well-studied SCFs, including Pareto-optimal majoritarian SCFs. In contrast to Moulin's negative result for resolute SCFs, there *are* appealing Condorcet extensions that satisfy Kelly-participation.

Our negative result holds even for strict preferences while our positive result holds even for weak preferences. The latter is somewhat surprising and stands in sharp contrast to the related finding that no Condorcet extension satisfies Kelly-*strategyproofness* when preferences are weak (recall that an SCF is strategyproof if no agent can obtain a more preferred outcome by misrepresenting his preferences) [Brandt, 2015].

Participation is similar to, but logically independent from, strategyproofness. Manipulation by abstention is arguably a more severe problem than manipulation by misrepresentation for two reasons. First, agents might not be able to find a beneficial misrepresentation. It was shown in various papers that the corresponding computational problem can be intractable (see, e.g., [Faliszewski *et al.*, 2010]). Finding a successful manipulation by strategic abstention, on the other hand, is never harder than computing the outcome of the respective SCF. Secondly, one could argue that agents will not lie about their preferences because this is considered immoral (Borda famously exclaimed "my scheme is intended only for honest men"), while strategic abstention is deemed acceptable.[1]

2 Related Work

The problem of strategic abstention for irresolute SCFs has been addressed by Pérez [2001], Jimeno *et al.* [2009], and Brandt [2015]. Pérez [2001] examined the situation where an agent can cause his most preferred alternative to be excluded from the choice set when joining an electorate and showed that almost all Condorcet extensions suffer from this paradox. Jimeno *et al.* [2009], on the other hand, proved that manipulation by abstention is possible for most Condorcet extensions when agents compare sets according to an optimistic, pessimistic, or lexicographic extension. They mentioned the study of participation in the context of weak preferences and Fishburn's extension as interesting research directions for future work. Both of these questions are addressed in our paper. Brandt [2015] investigated strategyproofness for Kelly's extension and gave a simple argument connecting strategyproofness and participation. Brandl *et al.* [2015] studied participation for probabilistic social choice functions and, among other results, proposed functions where a participating agent is always *strictly* better off (unless he already obtains a most-preferred outcome). Abstention in slightly different contexts than the one studied in this paper has recently also caught the attention of computer scientists working on voting equilibria and campaigning (see, e.g., [Desmedt and Elkind, 2010; Baumeister *et al.*, 2012]).

[1]Alternatively, one could also argue that manipulation by misrepresentation is more critical because agents are tempted to act immorally, which is a valid, but different, concern.

The computer-aided techniques in this paper have been inspired by Tang and Lin [2009], who reduced well-known impossibility results for *resolute* SCFs—such as Arrow's theorem—to finite instances, which can then be checked by a SAT solver. Geist and Endriss [2011] extended this method to a fully-automatic search algorithm for impossibility theorems in the context of preference relations over sets of alternatives. More recently, Brandt and Geist [2014] proved both impossibility and possibility results regarding the strategyproofness of irresolute SCFs using this computer-aided approach. We strongly build on their methodology and extended it to cover the notion of participation, which—as we will see—requires a more advanced framework.

Also for other problems in economics the application of SAT solvers has proven to be quite effective. A prominent example is the ongoing work by Leyton-Brown [2014] in which SAT solvers are used for the development and execution of the FCC's upcoming reverse spectrum auction.

In some respect, our approach also bears some similarities to *automated mechanism design* (see, e.g., [Conitzer and Sandholm, 2002]), where desirable properties are encoded and mechanisms are computed to fit specific problem instances.

3 Preliminaries

Let A be a finite set of alternatives and $\mathbb{N}$ a countable set of agents of which we will consider finite subsets $N \subseteq \mathbb{N}$. Therefore, let $\mathcal{F}(\mathbb{N})$ denote the set of all finite and non-empty subsets of $\mathbb{N}$. 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 is denoted by $\mathcal{R}$. We write $\succ_i$ for the strict part of $\succsim_i$, i.e., $x \succ_i y$ 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., $x \succ_i y$ or $y \succ_i x$ 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 $x \succ_i y \sim_i z$ is represented by $\succsim_i \colon x, \{y, z\}$.

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 $\mathcal{R}^{\mathcal{F}(\mathbb{N})}$. For a preference profile $R \in \mathcal{R}^N$ and two agents $i \in N$, $j \in \mathbb{N}$, we define

$$R_{-i} = R \setminus \{(i, \succsim_i)\} \quad \text{and} \quad R_{+j} = R \cup \{(j, \succsim_j)\}.$$

The *majority relation* of R is denoted by $\succsim(R)$, where

$$x \succsim(R)\, y \text{ iff } |\{i \in N \colon x \succsim_i y\}| \geq |\{i \in N \colon y \succsim_i x\}|.$$

Its strict part is denoted by $\succ(R)$ and its indifference part by $\sim(R)$. An alternative x is a *Condorcet winner* in R if $x \succ(R)\, y$ for all $y \in A \setminus \{x\}$.

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

$$f \colon \mathcal{R}^{\mathcal{F}(\mathbb{N})} \to 2^A \setminus \emptyset.$$

Two minimal fairness conditions for SCFs are anonymity and neutrality. An SCF is *anonymous* if the outcome does not depend on the identities of the agents and *neutral* if it is symmetric with respect to alternatives. An SCF f is *majoritarian* (or a neutral C1 function) if it is neutral and for all $R, R' \in \mathcal{R}^{\mathcal{F}(\mathbb{N})}$, $f(R) = f(R')$ whenever $\succ(R) = \succ(R')$. Even the seemingly narrow class of majoritarian SCFs contains a variety of interesting functions (sometimes called *tournament solutions*). Examples include Copeland's rule, the top cycle, and the uncovered set (see, e.g., [Brandt *et al.*, 2015]). These functions usually also happen to be *Condorcet extensions*, i.e., SCFs that uniquely return a Condorcet winner whenever one exists.

Next we introduce a very weak variable electorate condition which requires that a completely indifferent agent does not affect the outcome. An SCF f satisfies *independence of indifferent voters* (IIV) if

$$f(R) = f(R_{+i}) \text{ for all } R \in \mathcal{R}^{\mathcal{F}(\mathbb{N})},$$

where i is an agent who is indifferent between all alternatives, i.e., $x \sim_i y$ for all $x, y \in A$. It is easy to see that every majoritarian SCF satisfies anonymity, neutrality, and IIV.

We say that R' is an f-improvement over R if alternatives that are chosen by f in R are not weakened from R to R', i.e., for all $x \in f(R)$, $y \in A$, and $i \in N$, $x \succsim_i y$ implies $x \succsim'_i y$ and $y \succsim'_i x$ implies $y \succsim_i x$. An SCF f satisfies *set-monotonicity* if

$$f(R) = f(R') \text{ whenever } R' \text{ is an } f\text{-improvement over } R.$$

The two preference extensions we consider in this paper are Kelly's extension and Fishburn's extension. For all $X, Y \subseteq A$ and $\succsim_i \in \mathcal{R}$,

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

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

The strict part of these relations will be denoted by $\succ_i^K$ and $\succ_i^F$, respectively. It follows from the definitions that Fishburn's extension is a refinement of Kelly's extension, i.e., $\succsim_i^K \subseteq \succsim_i^F$ for every $\succsim_i \in \mathcal{R}$. In the interest of space, we refer to Section 1 (and the references therein) for justifications of these extensions.

With the preference extensions at hand, we can now formally define participation and strategyproofness. An SCF f is *Kelly-manipulable by strategic abstention* if there exists a preference profile $R \in \mathcal{R}^N$ with $N \in \mathcal{F}(\mathbb{N})$ and an agent $i \in N$ such that $f(R_{-i}) \succ_i^K f(R)$. 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 agents $j \neq i$ and $f(R') \succ_i^K f(R)$. f is said to satisfy *Kelly-participation* or *Kelly-strategyproofness* if it is not Kelly-manipulable by strategic abstention or Kelly-manipulable, respectively. Fishburn-participation and Fishburn-strategyproofness are defined analogously.

The following example illustrates the definitions of Kelly-participation and Fishburn-participation. Consider the preference profile R with six agents and four alternatives depicted below. The numbers on top of each column denote the identities of the agents with the respective preference relation.

1	2	3, 4	5, 6
c	d	a	b
d	b	c	a
b	a	d	c
a	c	b	d

R

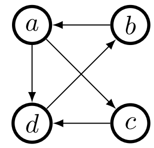

$\succ(R)$

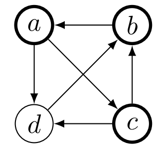

$\succ(R_{-6})$

The profile R induces the majority relation $\succsim(R)$ with its strict part $\succ(R)$. A well-studied majoritarian SCF is the *bipartisan set* [Laffond *et al.*, 1993; Dutta and Laslier, 1999]. The bipartisan set of R is $\{a, b, c, d\}$. If agent 6 leaves the electorate, we obtain the profile R_{-6}, which induces the majority relation $\succsim(R_{-6})$ whose bipartisan set is $\{a, b, c\}$. Observe that $\{a, b, c\} \succ_6^F \{a, b, c, d\}$, i.e., agent 6 can obtain a preferred outcome according to Fishburn's extension by abstaining from the election. Hence, the bipartisan set does not satisfy Fishburn-participation. However, $\{a, b, c\} \succ_6^K \{a, b, c, d\}$ does not hold and, hence, agent 6 cannot manipulate by abstaining according to Kelly's extension. In general, the bipartisan set satisfies Kelly-participation because it satisfies set-monotonicity and IIV (cf. Theorem 3).

We will relate participation and strategyproofness to *Pareto-optimality* in the subsequent sections. An alternative x is said to be *Pareto-dominated* (in $R \in \mathcal{R}^N$) by another alternative y if $y \succsim_i x$ for all $i \in N$ and there exists $j \in N$ such that $y \succ_j x$. Whenever there is no $y \in A$ that Pareto-dominates x, x is called *Pareto-optimal*. The *Pareto rule* (*PO*) is defined as the SCF that selects all Pareto-optimal alternatives.

4 Computer-aided Theorem Proving

For some of our results, we are going to make use of the computer-aided proving methodology described by Brandt and Geist [2014]. The main idea is to prove statements by encoding a finite instance as a satisfiability problem, which can be solved by a computer using a SAT solver, and providing a (simple) reduction argument, which extends this result to arbitrary domain sizes. We extend their framework to also cater for indifferences in the majority relations, which is an important requirement for being able to deal with the notion of participation: if an agent with at least one strict preference abstains the election, the corresponding majority relation might already contain indifferences.

Note that the introduction of majority ties significantly increases the size of the search space (see Table 1), which makes any type of exhaustive search even less feasible. Apart from being able to treat such large search spaces, another major advantage of the computer-aided approach is that many similar conjectures and hypotheses (here, e.g., statements about other preference extensions) can be checked quickly using the same framework.

In the coming subsections, we are going to explain our extension and some core features of the computer-aided method; for details of the original approach, however, the reader is referred to Brandt and Geist [2014].

$\lvert A\rvert$	Brandt and Geist [2014]	This paper
3	49	**823,543**
4	50,625	$\sim \mathbf{2.5 \cdot 10^{49}}$
5	$\sim 7.9 \cdot 10^{17}$	$\sim \mathbf{9.4 \cdot 10^{867}}$
6	$\sim 5.8 \cdot 10^{100}$	$\sim \mathbf{6.8 \cdot 10^{38649}}$

Table 1: Number of different majoritarian SCFs. While Brandt and Geist [2014] could assume an odd number of agents with strict preferences, participation requires us to deal with variable electorates, and therefore weak majority relations.

4.1 Encoding Participation

At the core of the computer-aided approach lies an encoding of the problem to be solved as a SAT instance. For this, all axioms involved need to be stated in propositional logic. We take over the formalization of the optimized encoding by Brandt and Geist [2014], which contains the following relevant axioms: *functionality of the choice function*, the *orbit condition*, and *Pareto-optimality*. Pareto-optimality is encoded as being a refinement of the uncovered set. What remains is to encode the notion of participation. While this encoding turns out to be similar to the one of strategyproofness defined by Brandt and Geist [2014], it is more complex and not straightforward. In particular, it requires a novel condition that is equivalent to participation for majoritarian SCFs, which we are going to call *majority-participation*.

We are going to identify preference profiles with their corresponding majority relations (i.e., we write $f(\succsim)$ instead of $f(R)$). Moreover, the *inverse* of $\succsim$ is denoted by $\precsim$.

Definition 1. A majoritarian SCF f is *Fishburn-majority-manipulable by strategic abstention* if there exist majority relations $\succsim$, $\succsim'$ and a preference relation $\succsim_\mu \in \mathcal{R}$ such that $f(\succsim') \succ_\mu^F f(\succsim)$, with

$$\succ \cap \prec' \;=\; \emptyset, \tag{1}$$

$$(\succsim \setminus \succsim') \cup (\precsim' \setminus \precsim) \;\subseteq\; \succ_\mu \text{, and} \tag{2}$$

$$\sim \cap \sim' \;\subseteq\; \sim_\mu. \tag{3}$$

If the agents' preferences are required to be strict, it additionally has to hold that either $\succsim$ or $\succsim'$ is anti-symmetric. A majoritarian SCF f satisfies *Fishburn-majority-participation* if it is not Fishburn-majority-manipulable by strategic abstention.

Conditions (1) to (3) can intuitively be phrased as follows: (1) no strict relationship may be reversed between $\succ$ and $\succ'$, (2) $\succsim_\mu$ has to be in line with the changes from $\succsim$ to $\succsim'$, and (3) majority ties that occur in both majority relations must be reflected by an indifference in $\succsim_\mu$.

In the following lemma, we show that, for majoritarian SCFs, the condition of Fishburn-majority-manipulability corresponds to an abstaining agent with preferences $\succsim_\mu$ who thereby obtains a preferred outcome.[2]

[2]Note that both the definition of majority-participation and Lemma 1 are independent of a specific preference extension, and thus also applicable to, e.g., Kelly's extension.

Lemma 1. *A majoritarian SCF satisfies Fishburn-participation if and only if it satisfies Fishburn-majority-participation.*

Proof. Due to space restrictions we will only provide a short proof sketch here, the full proof is available from the authors upon request.

In general, we show that for every preference profile R that allows for a Fishburn-manipulation by abstention by agent μ, the two majority graphs $\succsim(R)$ and $\succsim(R_{-\mu})$ together with $\succsim_\mu$ satisfy all required conditions. In return, whenever we have two majority relations $\succsim$, $\succsim'$ and a preference relation $\succsim_\mu$ with the properties stated in Definition 1, we can assign integer weights to all pairs of alternatives and, by Debord [1987], use these to determine a preference profile R' that induces the majority relation $\succsim'$. Together with $R = R'_{+\mu}$ we obtain $\succsim(R') = \succsim'$, $\succsim(R) = \succsim$ and thus $f(R') \succ_\mu^F f(R)$ for majoritarian SCFs f. □

Fishburn-majority-participation can then be encoded in propositional logic (with variables $f_{\succsim,X}$ representing $f(\succsim) = X$) as the following simple transformation shows:

$$\neg\left(f(\succsim') \succ_\mu^F f(\succsim)\right) \equiv \bigwedge_{Y \succ_\mu^F X} (\neg f_{\succsim',X} \vee \neg f_{\succsim,Y})$$

for all majority relations $\succsim$, $\succsim'$ and preference relations $\succsim_\mu$ satisfying conditions (1) to (3).

4.2 Proof Extraction

A very interesting feature of the approach by Brandt and Geist [2014] is the possibility to extract human-readable proofs from an unsatisfiability result by the SAT solver. This is done by computing a minimal unsatisfiable set (MUS), an inclusion-minimal set of clauses that is still unsatisfiable.[3] This MUS can then, assisted by our encoder/decoder program, be read and transformed into a standard human-readable proof. Different proofs can be found by varying the MUS extractor or by encoding the problem for different subdomains, such as neighborhoods of a set of profiles or randomly sampled subdomains, respectively.

5 Results and Discussion

In general, participation and strategyproofness are not logically related. However, extending an observation by Brandt [2015], it can be shown that strategyproofness implies participation under certain conditions. The proof of this statement is omitted due to space restrictions.

Lemma 2. *Consider an arbitrary preference extension. Every SCF that satisfies IIV and strategyproofness satisfies participation. When preferences are strict, every majoritarian SCF that satisfies strategyproofness satisfies participation.*

As a consequence of Lemma 2, some positive results for Kelly-strategyproofness and Fishburn-strategyproofness carry over to participation. We will complement these results by impossibility theorems for Fishburn-participation and a positive result for Kelly-participation, which specifically does not hold for Kelly-strategyproofness.

[3]We used PICOMUS, which is part of the PICOSAT distribution [Biere, 2008].

5.1 Fishburn-participation

It turns out that Pareto-optimality is incompatible with Fishburn-participation in majoritarian social choice. The corresponding Theorems 1 and 2 and their proofs were obtained using the computer-aided method laid out in Section 4. In order to simplify the original proofs, which were found by the computer, we first state a lemma, which offers further insights into the possible choices of majoritarian SCFs that satisfy Fishburn-participation and Pareto-optimality.

To state Lemma 3 we introduce some additional notation: an alternative x (McKelvey) covers an alternative y if x is at least as good as y compared to every other alternative. Formally, x *covers* y if $x \succ y$ and, for all $z \in A$, both $y \succsim z$ implies $x \succsim z$, and $z \succsim x$ implies $z \succsim y$. The *uncovered set* of $\succsim$, denoted $UC(\succsim)$, is the set of all alternatives that are not covered by any other alternative. By definition, UC is a majoritarian SCF.

Brandt *et al.* [2014] have shown that every majoritarian and Pareto-optimal SCF selects a subset of the (McKelvey) uncovered set. We show that an SCF that additionally satisfies Fishburn-participation furthermore only depends on the majority relation between alternatives in the uncovered set.

Lemma 3. *Let f be a majoritarian and Pareto-optimal SCF that satisfies Fishburn-participation. Let R, R' be preferences profiles such that $\succsim(R)|_{UC(\succsim(R))} = \succsim(R')|_{UC(\succsim(R'))}$. Then $f(R) = f(R') \subseteq UC(\succsim(R))$.*

The proof of Lemma 3 is omitted due to space restrictions.[4]

Now, let us turn to our impossibility theorems. The computer found these impossibilities even without using Lemma 3. However, the formalization of the lemma allowed the SAT solver to find smaller proofs and makes the human-readable proofs more intuitive.

Theorem 1. *There is no majoritarian and Pareto-optimal SCF that satisfies Fishburn-participation if $|A| \geq 4$.*

Proof. Let f be a majoritarian and Pareto-optimal SCF satisfying Fishburn-participation. We first prove the statement for $A = \{a, b, c, d\}$ and reason about the outcome of f for some specific majority relations. Throughout this proof, we are going to make extensive use of Lemma 1, which allows us to apply Fishburn-majority-participation instead of regular Fishburn-participation. Intuitively, the proof strategy is to alter the majority relations $\succsim$, $\succsim'$, and $\succsim''$ as depicted below by letting varying agents *join* some underlying electorate, which will exclude certain choices of f (by an application of Fishburn-majority-participation), until we reach a contradiction. In the figures of the strict part of the majority relations we highlight alternatives that have to be chosen by f with a thick border.

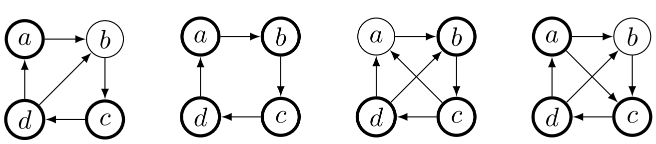

$\succ$ $\qquad$ $\succ_{+1}$ $\qquad$ $\succ_{+2}$ $\qquad$ $\succ_{+3}$

[4]Lemma 3 can be strengthened in various respects. It also holds for the *iterated* uncovered set, all preference extensions satisfying some mild conditions, and probabilisitic social choice functions.

First consider $\succsim$. Adding an agent with preferences $\succsim_1\colon \{a, b, c\}, d$ possibly yields $\succsim_{+1}$ where, due to symmetry, $f(\succsim_{+1}) = \{a, b, c, d\}$. As f satisfies Fishburn-participation, nothing that is strictly preferred to $\{a, b, c, d\}$ according to $\succsim_1^F$ can be chosen in $\succsim$. Thus, $d \in f(\succsim)$. Adding another agent with $\succsim_1\colon b, \{a, c, d\}$ may also lead to $\succsim_{+1}$. Hence $f(\succsim) \neq \{b, d\}, \{a, b, d\}, \{b, c, d\}$.

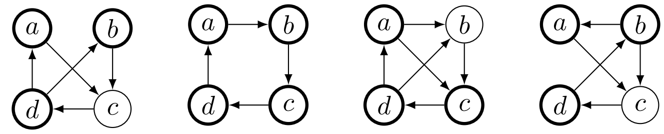

$\succ'$ $\qquad$ $\succ'_{+1}$ $\qquad$ $\succ'_{+2}$ $\qquad$ $\succ'_{+3}$

In a similar fashion, we obtain—step by step and using the majority relations depicted—that $f(\succsim) = \{a, c, d\}$, $f(\succsim') = \{a, b, d\}$, and finally $f(\succsim'') = \emptyset$, a contradiction. The details of these cases have to be omitted due to space constraints and are available from the authors upon request.

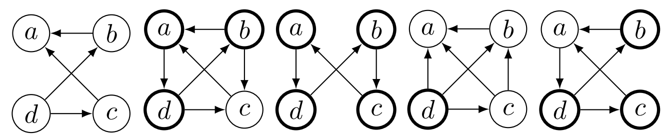

$\succ''$ $\qquad$ $\succ''_{+1}$ $\qquad$ $\succ''_{+2}$ $\qquad$ $\succ''_{+3}$ $\qquad$ $\succ''_{+4}$

Now let $|A| \geq 5$. It follows from Lemma 3 that the choice of f does not depend on covered alternatives. Hence the statement follows by extending the majority relations depicted above to A such that all alternatives but a, b, c, d are covered. $\square$

We could verify with our computer-aided approach that this impossibility still holds for strict preferences when there are at least 5 alternatives.

Theorem 2. *There is no majoritarian and Pareto-optimal SCF that satisfies Fishburn-participation if $|A| \geq 5$, even when preferences are strict.*

The shortest proof of Theorem 2 that we were able to extract from our computer-aided approach still uses 124 different instances of manipulation by abstention. The proof of Theorem 1, by comparison, consists of 10 such instances. As a consequence, the complete proof of Theorem 2 has to be omitted here; a computer-generated version, however, is available from the authors upon request. Theorems 1 and 2 are both tight in the sense that, whenever there are less than four or five alternatives, respectively, there exists an SCF that satisfies the desired properties.

An interesting question is whether these impossibilities also extend to other preference extensions. Given the computer-aided approach, this can be easily checked by simply replacing the preference extension in the SAT encoder. For instance, it turns out that the impossibility of Theorem 1 still holds if we consider a coarsening of Fishburn's extension which can only compare sets that are contained in each other. Kelly's extension on the other hand does not lead to an impossibility for $|A| \leq 5$, which will be confirmed more generally in the next section.

Extension	Property	Strict preferences	Weak preferences
Kelly	Participation	$\oplus$ All set-monotonic SCFs that satisfy IIV (Theorem 3)	
	Strategy-proofness	$\oplus$ All set-monotonic SCFs[b]	$\ominus$ No Condorcet extension[b]
Fishburn	Participation	$\ominus$ No majoritarian & Pareto-optimal SCF ($\lvert A\rvert \geq 5$) (Theorem 2)	$\ominus$ No majoritarian & Pareto-optimal SCF ($\lvert A\rvert \geq 4$) (Theorem 1)
		$\oplus$ Few undiscriminating SCFs, e.g., $COND^c$, TC^g, and PO^e (Lemma 2), and all scoring rules[f]	
	Strategy-proofness	$\ominus$ No majoritarian & Pareto-optimal SCF ($\lvert A\rvert \geq 5$)[d]	$\ominus$ No pairwise & Pareto-optimal SCF ($\lvert A\rvert \geq 4$)[a]
		$\oplus$ Few undiscriminating SCFs, e.g., $COND^c$, TC^g, and PO^e	

Table 2: Overview of results for participation and strategyproofness with respect to Kelly's and Fishburn's extension and strict/weak preferences. The symbol $\oplus$ marks sufficient conditions while $\ominus$ marks impossibility results. [a]: follows from a result by Aziz *et al.* [2014] about probabilistic social choice functions, [b]: Brandt [2015], [c]: Brandt and Brill [2011], [d]: Brandt and Geist [2014], [e]: Feldman [1979], [f]: Moulin [1988], [g]: Sanver and Zwicker [2012]

5.2 Kelly-participation

Theorems 1 and 2 are sweeping impossibilities within the domain of majoritarian SCFs. For Kelly's extension, we obtain a much more positive result that covers attractive majoritarian and non-majoritarian SCFs. Brandt [2015] has shown that set-monotonicity implies Kelly-strategyproofness for strict preferences and that no Condorcet extension is Kelly-strategyproof when preferences are weak. We prove that set-monotonicity (and the very mild IIV axiom) imply Kelly-participation even for weak preferences. We have thus found natural examples of SCFs that violate Kelly-strategyproofness but satisfy Kelly-participation.[5]

Theorem 3. *Let f be an SCF that satisfies IIV and set-monotonicity. Then f satisfies Kelly-participation.*

Proof. Let f be an SCF that satisfies IIV and set-monotonicity. Assume for contradiction that f does not satisfy Kelly-participation. Hence there exist a preference profile R and an agent i such that $f(R_{-i}) \succ_i^K f(R)$. Let $X = f(R)$, $Y = f(R_{-i})$, and $Z = A \setminus (X \cup Y)$. By definition of $Y \succ_i^K X$ we have that $x \sim_i y$ for all $x, y \in X \cap Y$.

We define a new preference relation $\succsim_i'$ in which all alternatives in Y are tied for the first place, followed by all alternatives in $X \setminus Y$ as they are ordered in $\succsim_i$, and all remaining alternatives in one indifference class at the bottom of the ranking. Formally,

$$\succsim_{i'} = (Y \times A) \cup \succsim_i|_X \cup (A \times Z).$$

Let i'' be an agent who is indifferent between all alternatives, i.e., $x \sim_{i''} y$ for all $x, y \in A$. Since f satisfies IIV we have that $f(R_{-i+i''}) = f(R_{-i})$.

By definition, $R_{-i+i'}$ is an f-improvement over both R and $R_{-i+i''}$. Hence, set-monotonicity implies that $f(R_{-i+i'}) = f(R)$ and $f(R_{-i+i'}) = f(R_{-i+i''})$. In summary, we obtain

$$f(R_{-i+i'}) = f(R_{-i+i''}) = f(R_{-i}) \succ_i^K f(R) = f(R_{-i+i'}),$$

which is a contradiction. □

[5] It is easily seen that the proof of Theorem 3 straightforwardly extends to *group*-participation, i.e., no group of agents can obtain a unanimously more preferred outcome by abstaining.

Two rather undiscriminating SCFs that satisfy both IIV and set-monotonicity are the *Pareto rule* and the *omninomination rule* (which returns all alternatives that are ranked first by at least one agent). Majoritarian SCFs satisfy IIV by definition and there are several appealing majoritarian SCFs that satisfy set-monotonicity, among those for instance the *top cycle*, the *minimal covering set*, and the *bipartisan set* (see, e.g., [Brandt, 2015; Brandt *et al.*, 2015]). These majoritarian SCFs are sometimes criticized for not being discriminating enough. The computer-aided approach described in this paper can be used to find more discriminating SCFs that still satisfy Kelly-participation. We thus found a refinement of the bipartisan set that, for $|A| = 5$, selects only 1.43 alternatives on average, and satisfies Kelly-participation. For comparison, the bipartisan set (the smallest previously known majoritarian SCF satisfying Kelly-participation) yields 2.68 alternatives on average.

6 Conclusion

Previous results have indicated a conflict between strategic non-manipulability and Condorcet-consistency [Moulin, 1988; Pérez, 2001; Jimeno *et al.*, 2009; Brandt, 2015]. For example, Moulin [1988] has shown that no resolute Condorcet extension satisfies participation and Brandt [2015] has shown that no irresolute Condorcet extension satisfies Kelly-strategyproofness. Theorem 3 addresses an intermediate question and finds that—perhaps surprisingly—there *are* attractive Condorcet extensions that satisfy Kelly-participation, even when preferences are weak. On the other hand, we have presented elaborate computer-generated impossibilities (Theorems 1 and 2), which show that these encouraging results break down once preferences are extended by the more refined Fishburn extension. These findings improve our understanding of which behavioral assumptions allow for aggregation functions that are immune to strategic abstention.

An overview of the main results of this paper and how they relate to other related results is given in Table 2.

References

[Aziz *et al.*, 2014] H. Aziz, F. Brandl, and F. Brandt. On the incompatibility of efficiency and strategyproofness in randomized social choice. In *Proc. of 28th AAAI Conference*, pages 545–551. AAAI Press, 2014.

[Barberà *et al.*, 2004] S. Barberà, W. Bossert, and P. K. Pattanaik. Ranking sets of objects. In S. Barberà, P. J. Hammond, and C. Seidl, editors, *Handbook of Utility Theory*, volume II, chapter 17, pages 893–977. Kluwer Academic Publishers, 2004.

[Baumeister *et al.*, 2012] D. Baumeister, P. Faliszewski, J. Lang, and J. Rothe. Campaigns for lazy voters: truncated ballots. In *Proc. of 11th AAMAS Conference*, pages 577–584. IFAAMAS, 2012.

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

[Brandl *et al.*, 2015] F. Brandl, F. Brandt, and J. Hofbauer. Incentives for participation and abstention in probabilistic social choice. In *Proc. of 14th AAMAS Conference*. IFAAMAS, 2015. Forthcoming.

[Brandt and Brill, 2011] F. Brandt and M. Brill. Necessary and sufficient conditions for the strategyproofness of irresolute social choice functions. In *Proc. of 13th TARK Conference*, pages 136–142. ACM Press, 2011.

[Brandt and Geist, 2014] F. Brandt and C. Geist. Finding strategyproof social choice functions via SAT solving. In *Proc. of 13th AAMAS Conference*, pages 1193–1200. IFAAMAS, 2014.

[Brandt *et al.*, 2013] F. Brandt, V. Conitzer, and U. Endriss. Computational social choice. In G. Weiß, editor, *Multiagent Systems*, chapter 6, pages 213–283. MIT Press, 2nd edition, 2013.

[Brandt *et al.*, 2014] F. Brandt, C. Geist, and P. Harrenstein. A note on the McKelvey uncovered set and Pareto optimality. 2014. Working Paper.

[Brandt *et al.*, 2015] F. Brandt, M. Brill, and P. Harrenstein. Tournament solutions. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 3. Cambridge University Press, 2015. Forthcoming.

[Brandt, 2015] F. Brandt. Set-monotonicity implies Kelly-strategyproofness. *Social Choice and Welfare*, 2015. Forthcoming.

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

[Conitzer and Sandholm, 2002] V. Conitzer and T. Sandholm. Complexity of mechanism design. In *Proc. of 18th UAI Conference*, pages 103–110, 2002.

[Debord, 1987] B. Debord. Caractérisation des matrices des préférences nettes et méthodes d'agrégation associées. *Mathématiques et sciences humaines*, 97:5–17, 1987.

[Desmedt and Elkind, 2010] Y. Desmedt and E. Elkind. Equilibria of plurality voting with abstentions. In *Proc. of 11th ACM-EC Conference*, pages 347–356. ACM Press, 2010.

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

[Faliszewski *et al.*, 2010] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. Using complexity to protect elections. *Communications of the ACM*, 53(11):74–82, 2010.

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

[Fishburn and Brams, 1983] P. C. Fishburn and S. J. Brams. Paradoxes of preferential voting. *Mathematics Magazine*, 56(4):207–214, 1983.

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

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

[Geist and Endriss, 2011] 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.

[Jimeno *et al.*, 2009] J. L. Jimeno, J. Pérez, and E. García. An extension of the Moulin No Show Paradox for voting correspondences. *Social Choice and Welfare*, 33(3):343–459, 2009.

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

[Laffond *et al.*, 1993] G. Laffond, J.-F. Laslier, and M. Le Breton. The bipartisan set of a tournament game. *Games and Economic Behavior*, 5(1):182–201, 1993.

[Leyton-Brown, 2014] K. Leyton-Brown. Pragmatic algorithmic game theory. *Keynote address at the ACM-EC Conference*, June 2014.

[Moulin, 1988] H. Moulin. Condorcet's principle implies the no show paradox. *Journal of Economic Theory*, 45:53–64, 1988.

[Pérez, 2001] J. Pérez. The strong no show paradoxes are a common flaw in Condorcet voting correspondences. *Social Choice and Welfare*, 18(3):601–616, 2001.

[Sanver and Zwicker, 2012] M. R. Sanver and W. S. Zwicker. Monotonicity properties and their adaption to irresolute social choice rules. *Social Choice and Welfare*, 39(2–3):371–398, 2012.

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