Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence

Group-Strategyproof Irresolute Social Choice Functions

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

Abstract

An important problem in voting is that agents may misrepresent their preferences in order to obtain a more preferred outcome. Unfortunately, this phenomenon has been shown to be inevitable in the case of resolute, i.e., single-valued, social choice functions. In this paper, we introduce a variant of Maskin-monotonicity that completely characterizes the class of pairwise irresolute social choice functions that are group-strategyproof according to Kelly's preference extension. The class is narrow but contains a number of appealing Condorcet extensions such as the *minimal covering set* and the *bipartisan set*, thereby answering a question raised independently by Barberà (1977) and Kelly (1977). These functions furthermore encourage participation and thus do not suffer from the no-show paradox (under Kelly's extension).

1 Introduction

One of the central results in social choice theory states that every non-trivial social choice function (SCF)—a function mapping individual preferences to a collective choice—is susceptible to strategic manipulation (Gibbard, 1973; Satterthwaite, 1975). However, the classic result by Gibbard and Satterthwaite only applies to *resolute*, i.e., single-valued, SCFs. The notion of a resolute SCF is rather restricted and artificial.[1] For example, consider a situation with two voters and two alternatives such that each voter prefers a different alternative. The problem is not that a resolute SCF has to pick a single alternative (which is a well-motivated practical requirement), but that it has to pick a single alternative *based on the individual preferences alone* (see, e.g., Kelly, 1977). As a consequence, resoluteness is at variance with elementary notions of fairness such as neutrality and anonymity.

In order to remedy this shortcoming, Gibbard (1977) went on to characterize the class of strategyproof *decision schemes*, i.e., aggregation functions that yield probability distributions over the set of alternatives rather than single alternatives (see also Barberà, 1979). This class consists of rather degenerate decision schemes and Gibbard's characterization is therefore commonly interpreted as another impossibility result. However, Gibbard's theorem rests on unusually strong assumptions with respect to the voters' preferences. In contrast to the traditional setup in social choice theory, which typically only involves ordinal preferences, his result relies on the axioms of von Neumann and Morgenstern (or an equivalent set of axioms) in order to compare lotteries over alternatives.

The gap between Gibbard and Satterthwaite's theorem for resolute SCFs and Gibbard's theorem for decision schemes has been filled by a number of impossibility results with varying underlying notions of how to compare sets of alternatives with each other (e.g., Gärdenfors, 1976; Barberà, 1977; Kelly, 1977; Duggan and Schwartz, 2000; Ching and Zhou, 2002; Sato, 2008), many of which are surveyed by Taylor (2005) and Barberà (2010). In this paper, we will be concerned with the weakest (and therefore least controversial) preference extension from alternatives to sets due to Kelly (1977). According to this definition, a set of alternatives is weakly preferred to another set of alternatives if all elements of the former are weakly preferred to all elements of the latter. Barberà (1977) and Kelly (1977) have shown independently that all non-trivial SCFs that are rationalizable via a quasi-transitive relation are manipulable in this model. However, as witnessed by various other (non-strategic) impossibility results that involve quasi-transitive rationalizability (e.g., Mas-Colell and Sonnenschein, 1972), it appears as if this property itself is unduly restrictive. As a consequence, Kelly (1977) concludes his paper by contemplating 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" and Barberà (1977) states that "whether a nonrationalizable collective choice rule exists which is not manipulable and always leads to nonempty choices for nonempty finite issues is an open question." Also referring to nonrationalizable choice functions, Kelly (1977) writes: "it is an open question how far nondictatorship can be strengthened in this sort of direction and still avoid impossibility results."

The first result of this paper is negative and shows that no Condorcet extension is strategyproof. The proof, how-

[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."

ever, crucially depends on strategic tie-breaking and hence does not work for strict preferences. We therefore turn to SCFs that cannot be manipulated by voters who misrepresent the *strict* part of their preference relation and show that all SCFs that satisfy a new variant of Maskin-monotonicity called *set-monotonicity* are strategyproof in this sense. Set-monotonicity requires the invariance of choice sets under the weakening of unchosen alternatives and is only satisfied by a handful of SCFs such as the *top cycle*, the *minimal covering set*, and the *bipartisan set*. Strategyproofness (according to Kelly's preference extension) thus draws a sharp line within the space of SCFs as many established SCFs (such as plurality, Borda's rule, and all weak Condorcet extensions like Llull's rule or Young's rule) are known to be manipulable for strict preferences (see, e.g., Taylor, 2005). We furthermore show that our characterization is complete for *pairwise* SCFs, i.e., SCFs whose outcome only depends on the comparisons between pairs of alternatives. Since set-monotonicity coincides with Maskin-monotonicity in the context of resolute SCFs, our characterization can thus be seen as a generalization of the Muller and Satterthwaite (1977) theorem within the setting of pairwise SCFs. We conclude the paper by pointing out that voters can never benefit from abstaining strategyproof pairwise SCFs. This does not hold for *resolute* Condorcet extensions, which is commonly known as the *no-show paradox* (Moulin, 1988).

Kelly's conservative preference extension has previously been primarily invoked in *impossibility* theorems because it is independent of the voters' attitude towards risk and the mechanism that eventually picks a single alternative from the choice set. Its interpretation in positive results, such as in this paper, is more debatable. Gärdenfors (1979) has shown that Kelly's extension is appropriate in a probabilistic context when voters are unaware of the lottery that will be used to pick the winning alternative. (Whether they are able to attach utilities to alternatives or not does not matter.) Alternatively, one can think of an independent chairman or a black-box that picks alternatives from choice sets in a way that prohibits a meaningful prior distribution. Whether these assumptions can reasonably be justified or such a device can actually be built is open to discussion. In particular, the study of distributed protocols or computational selection devices that justify Kelly's extension appears to be promising. Inspired by early work by Bartholdi, III et al. (1989), recent research in computer science investigated how to use computational hardness—namely NP-hardness—as a barrier against manipulation. However, NP-hardness is a worst-case measure and it would be much preferred if manipulation is hard on average. Recent negative results on the hardness of typical cases have cast doubt on this strand of research (see, e.g., Conitzer and Sandholm, 2006; Walsh, 2009; Isaksson et al., 2010), but more work remains to be done to settle the question completely. The current state of affairs is surveyed by Faliszewski and Procaccia (2010). If computational protocols or devices can be used to justify Kelly's extension by making "unpredictable" random selections, this might be an interesting alternative application of computational techniques to obtain strategyproofness.

2 Preliminaries

In this section, we provide the terminology and notation required for our results. We use the standard model of social choice functions with a variable agenda (see, e.g., Taylor, 2005).

2.1 Social Choice Functions

Let U be a universe of alternatives over which voters entertain preferences. The preferences of voter i are represented by a complete preference relation $R_i \subseteq U \times U$.[2] We have $a\ R_i\ b$ denote that voter i values alternative a at least as much as alternative b. In compliance with conventional notation, we write P_i for the strict part of R_i, i.e., $a\ P_i\ b$ if $a\ R_i\ b$ but not $b\ R_i\ a$. Similarly, I_i denotes i's indifference relation, i.e., $a\ I_i\ b$ if both $a\ R_i\ b$ and $b\ R_i\ a$. We denote the set of all preference relations over the universal set of alternatives U by $\mathcal{R}(U)$ and the set of *preference profiles*, i.e., finite vectors of preference relations, by $\mathcal{R}^*(U)$. The typical element of $\mathcal{R}^*(U)$ is $R = (R_1, \dots, R_n)$ and the typical set of voters is $N = \{1, \dots, n\}$. The set of *feasible sets* from which alternatives are to be chosen is the set of finite and non-empty subsets of U, denoted by $\mathcal{F}(U)$. Our central object of study are *social choice functions*, i.e., functions that map the individual preferences of the voters and a feasible set to a set of socially preferred alternatives.

Definition 1. *A* social choice function (SCF) *is a function* $f : \mathcal{R}^*(U) \times \mathcal{F}(U) \to \mathcal{F}(U)$ *such that* $f(R, A) \subseteq A$ *and* $f(R, A) = f(R', A)$ *for all feasible sets A and preference profiles* R, R' *such that* $R|_A = R'|_A$.

A *Condorcet winner* is an alternative a that, when compared with every other alternative b, is preferred by more voters, i.e., $|\{i \in N \mid a\ R_i\ b\}| > |\{i \in N \mid b\ R_i\ a\}|$ for all alternatives $b \neq a$. An SCF is called a *Condorcet extension* if it uniquely selects the Condorcet winner whenever one exists. The following notational convention will be very helpful throughout the paper. For a given preference relation R_i and alternatives a and b,

$$R_i^{(a,b)} = \{R_i\} \cup \{R_i \cup \{(a, b)\}\} \cup \{R_i \setminus \{(b, a)\} \cup \{(a, b)\}\}.$$

That is, $R_i^{(a,b)}$ is the set of all preference relations in which alternative a is weakly strengthened with respect to b.

A standard property of SCFs that is often considered is monotonicity. An SCF is monotonic if a chosen alternative remains in the choice set when it is strengthened in individual preference relations while leaving everything else unchanged.

Definition 2. *An SCF f is* monotonic *if for all feasible sets A, voters i, and preference profiles R and* R' *such that* $R_j = R'_j$ *for all* $j \neq i$, $a \in f(R, A)$ *and* $R'_i \in R_i^{(a,b)}$ *for some* $b \in A$ *implies* $a \in f(R', A)$.

Two other properties that will turn out to be useful in the context of this paper are the strong superset property and independence of unchosen alternatives. The strong superset

[2]Transitivity of individual preferences is not necessary for our results to hold. In fact, Theorem 3 is easier to prove for general—possibly intransitive—preferences. Theorem 4, on the other hand, would require a more cumbersome case analysis for transitive preferences.

property goes back to early work by Chernoff (1954) (see also Laslier, 1997) and requires that choice sets are invariant under the removal of unchosen alternatives.

Definition 3. *An SCF f satisfies the* strong superset property *(SSP) if for all feasible sets A, B and preference profiles R such that $f(R,A) \subseteq B \subseteq A$, $f(R,A) = f(R,B)$.*

Independence of unchosen alternatives was introduced by Laslier (1997) in the context of tournament solutions and requires that the choice set is invariant under modifications of the preference profile with respect to unchosen alternatives.

Definition 4. *An SCF f satisfies* independence of unchosen alternatives *(IUA) if for all feasible sets A and preference profiles R and R' such that $R_i|_{\{a,b\}} = R'_i|_{\{a,b\}}$ for all $a \in f(R,A)$, $b \in A$, and $i \in N$, $f(R,A) = f(R',A)$.*

2.2 Strategyproofness

An SCF is manipulable if one or more voters can misrepresent their preferences in order to obtain a more preferred outcome. Whether one choice set is preferred to another depends on how the preferences over individual alternatives are to be extended to sets of alternatives. In the absence of information about the mechanism that eventually picks a single alternative from any choice set, preferences over choice sets are typically obtained by the conservative extension $\widehat{R}_i$ (Barberà, 1977; Kelly, 1977), where for any pair of feasible sets A and B and preference relation R_i,

$$A \mathrel{\widehat{R}_i} B \text{ if and only if } a \mathrel{R_i} b \text{ for all } a \in A \text{ and } b \in B.$$

Clearly, in all but the simplest cases, $\widehat{R}_i$ is incomplete, i.e., many pairs of choice sets are incomparable. The strict part of $\widehat{R}_i$ is denoted by $\widehat{P}_i$, i.e., $A \mathrel{\widehat{P}_i} B$ if and only if $A \mathrel{\widehat{R}_i} B$ and $a \mathrel{P_i} b$ for at least one pair of $a \in A$ and $b \in B$.

Definition 5. *An SCF f is* $\widehat{R}$-manipulable by a group of voters *$G \subseteq N$ if there exists a feasible set A and preference profiles R, R' with $R_i = R'_i$ for all $i \notin G$ and $f(R,A) \neq f(R',A)$ such that*

$$f(R',A) \mathrel{\widehat{R}_i} f(R,A) \text{ for all } i \in G.$$

An SCF is $\widehat{R}$-strategyproof *if it is not $\widehat{R}$-manipulable by single voters, i.e., groups of size one. An SCF is* $\widehat{R}$-group-strategyproof *if it is not $\widehat{R}$-manipulable by any group of voters. $\widehat{P}$-strategyproofness and $\widehat{P}$-group-strategyproofness can be defined analogously.*

$\widehat{R}$-group-strategyproofness is particularly strong in the sense that none of the manipulating voters has to be *strictly* better off in the new preference profile. Obviously, every SCF that is $\widehat{R}$-group-strategyproof is also $\widehat{P}$-group-strategyproof.

It will turn out that SCFs that fail to be $\widehat{R}$-strategyproof can only be manipulated by breaking ties strategically, i.e., voters can obtain a more preferred outcome by only misrepresenting their *indifference* relation. In many settings, for instance when the choice infrastructure requires a strict ranking of the alternatives, this may be deemed acceptable. Accordingly, we obtain the following definition.

Definition 6. *An SCF is* strongly $\widehat{R}$-manipulable by a group of voters *$G \subseteq N$ if there exists a feasible set A and preference profiles R, R' with $R_i = R'_i$ for all $i \notin G$, $I_i \subseteq I'_i$ for all $i \in G$, and $f(R,A) \neq f(R',A)$ such that*

$$f(R',A) \mathrel{\widehat{R}_i} f(R,A) \text{ for all } i \in G.$$

An SCF is weakly $\widehat{R}$-group-strategyproof *if it is not strongly manipulable by any group of voters. Weak $\widehat{P}$-group-strategyproofness can be defined analogously.*

In other words, every strongly manipulable SCF admits a manipulation in which voters only misrepresent their strict preferences. Clearly, weak $\widehat{R}$-group-strategyproofness and $\widehat{R}$-group-strategyproofness coincide when voters have strict preferences.

Strategyproofness has been shown to be tightly connected to a strong version of monotonicity, which derives its name from a characterization of Nash-implementable SCFs due to Maskin (1999). An SCF satisfies Maskin-monotonicity if a chosen alternative remains in the choice set when it is (weakly) strengthened in individual preference relations and the relationships between other unrelated alternatives may be modified arbitrarily. Alternatively, it can be defined by requiring that an alternative remains in the choice set when weakening other alternatives.

Definition 7. *An SCF f satisfies* Maskin-monotonicity *if for all feasible sets A, voters i, and preference profiles R and R' such that $R_j = R'_j$ for all $j \neq i$, $x \in f(R,A)$ and $R'_i \in R_i^{(a,b)}$ for some $a \in A, b \in A \setminus \{x\}$ implies $x \in f(R',A)$.*

For strict individual preferences, Maskin-monotonicity precisely characterizes strategyproof resolute SCFs.

Theorem 1 (Muller and Satterthwaite, 1977). *If voters have strict preferences, a resolute SCF is group-strategyproof if and only if it satisfies Maskin-monotonicity.*

Unfortunately, as famously shown by Gibbard (1973) and Satterthwaite (1975), only trivial resolute SCFs satisfy Maskin-monotonicity. In Section 3.2, we will introduce a syntactically very similar condition that characterizes $\widehat{R}$-group-strategyproof irresolute SCFs.

3 Results

This section contains four results. First, we show that no Condorcet extension is $\widehat{P}$-strategyproof (Theorem 2). The proof of this claim, however, crucially depends on breaking ties strategically. We therefore study *weak* $\widehat{R}$-strategyproofness and obtain a much more positive characterization (Theorem 3) and show that the condition used for this characterization is necessary and sufficient in the case of pairwise SCFs (Theorem 4). Finally, we briefly examine the consequences of our results on strategic abstention (Proposition 2).

3.1 Manipulation of Condorcet Extensions

We begin by showing that all Condorcet extensions are weakly $\widehat{P}$-manipulable, which strengthens previous results by Gärdenfors (1976) and Taylor (2005) who showed the same

2	…	2	1	…	1	1
$a_2, \ldots, a_m$	…	$a_1, \ldots, a_{m-1}$	$a_3, \ldots, a_m$	…	$a_1, \ldots, a_{m-2}$	$a_2, \ldots, a_{m-1}$
			a_1	…	a_{m-1}	a_m
a_1	…	a_m	a_2	…	a_m	a_1

Table 1: Preference profile R for $3m$ voters where $A = \{a_1, \ldots, a_m\}$. Voters are numbered from left to right.

statement for a weaker notion of manipulability and weak Condorcet extensions, respectively.[3]

Theorem 2. *Every Condorcet extension is $\widehat{P}$-manipulable when there are more than two alternatives.*

Proof. Let $A = \{a_1, \ldots, a_m\}$ with $m \geq 3$ and consider the preference profile R given in Table 1. For every alternative a_i, there are two voters who prefer every alternative to a_i and are otherwise indifferent. Moreover, there is one voter for every alternative a_i who prefers every alternative except a_{i+1} to a_i, ranks a_{i+1} below a_i, and is otherwise indifferent.

Since $f(R, A)$ yields a non-empty choice set, there has to be some $a_i \in f(R, A)$. Due to the symmetry of the preference profile, we may assume without loss of generality that $a_2 \in f(R, A)$. Now, let

$$R' = (R_1, R_2, R_3 \setminus \{(a_i, a_1) \mid 3 \leq i \leq m\}, R_4, \ldots, R_{3m}) \text{ and}$$
$$R'' = (R'_1, R'_2, R'_3, R'_4 \setminus \{(a_i, a_1) \mid 3 \leq i \leq m\}, R'_5, \ldots, R'_{3m}).$$

That is, R' is identical to R, except that voter 3 lifted a_1 on top and R'' is identical to R', except that voter 4 lifted a_1 on top. Observe that $f(R'', A) = \{a_1\}$ because a_1 is the Condorcet winner in R''.

In case that $a_2 \notin f(R', A)$, voter 3 can manipulate as follows. Suppose R is the true preference profile. Then, the least favorable alternative of voter 3 is chosen (possibly among other alternatives). He can misstate his preferences as in R' such that a_2 is not chosen. Since he is indifferent between all other alternatives, $f(R', A) \; \widehat{P}_3 \; f(R, A)$.

If $a_2 \in f(R', A)$, voter 4 can manipulate similarly. Suppose R' is the true preference profile. Again, the least favorable alternative of voter 4 is chosen. By misstating his preferences as in R'', he can assure that one of his preferred alternatives, namely a_1, is selected exclusively because it is the Condorcet winner in R''. Hence, $f(R'', A) \; \widehat{P}'_4 \; f(R', A)$. □

3.2 Weakly Group-Strategyproof SCFs

The previous statement showed that no Condorcet extension is $\widehat{P}$-strategyproof, let alone $\widehat{R}$-group-strategyproof. For our characterization of *weakly* $\widehat{R}$-group-strategyproof SCFs, we introduce the following variant of Maskin-monotonicity (cf. Definition 7).

Definition 8. *An SCF f satisfies* set-monotonicity *if for all feasible sets A, voters i, and preference profiles R and R' such that $R_j = R'_j$ for all $j \neq i$, $X = f(R, A)$ and $R'_i \in R_i^{(a,b)}$ for some $a \in A, b \in A \setminus X$ implies $X = f(R', A)$.*

In other words, an SCF satisfies set-monotonicity if the choice set is invariant under the weakening of unchosen alternatives.

Despite the similar appearance, set-monotonicity is logically independent of Maskin-monotonicity. However, set-monotonicity coincides with Maskin-monotonicity in the context of resolute SCFs and, in the presence of IUA, it is weaker than Maskin-monotonicity and stronger than monotonicity. The proof of the following proposition is omitted due to space restrictions.

Proposition 1. *Maskin-monotonicity and IUA imply set-monotonicity. Set-monotonicity implies monotonicity and IUA.*

We are now ready to state the main result of this section.[4]

Theorem 3. *Every SCF that satisfies set-monotonicity is weakly $\widehat{R}$-group-strategyproof.*

Proof. Let f be an SCF that satisfies set-monotonicity and assume for contradiction that f is not weakly $\widehat{R}$-group-strategyproof. Then, there has to be a feasible set A, a group of voters $G \subseteq N$, and two preference profiles R and R' with $R_i = R'_i$ for all $i \notin G$ and $I_i \subseteq I'_i$ for all $i \in G$ such that $f(R', A) \neq f(R, A)$ and $f(R', A) \; \widehat{R}_i \; f(R, A)$ for all $i \in G$. We choose R and R' such that the size of the union of the symmetric differences of individual preferences $R \triangle R' = \bigcup_{i \in N}(R_i \setminus R'_i) \cup (R'_i \setminus R_i)$ is minimal, i.e., we look at a "smallest" counterexample in the sense that R and R' coincide as much as possible. Let $f(R, A) = X$ and $f(R', A) = Y$. We may assume $R \triangle R' \neq \emptyset$ as otherwise $R = R'$ and $X = Y$. Now, consider a pair of alternatives $a, b \in A$ such that, for some $i \in G$, $a \; P_i \; b$ and $b \; R'_i \; a$, i.e., voter i misrepresents his strict preference relation by strengthening b. The following case analysis will show that no such a and b exist, which implies that R and R' and consequently X and Y are identical, a

[3] A weak Condorcet winner is an alternative that is preferred by *at least* as many voters than any other alternative in pairwise comparisons. In contrast to Condorcet winners, weak Condorcet winners need not be unique. An SCF is called a *weak Condorcet extension* if it chooses the set of weak Condorcet winners whenever this set is non-empty. A large number of reasonable Condorcet extensions (including the minimal covering set and the bipartisan set) are not weak Condorcet extensions. Taylor (2005) calls the definition of weak Condorcet extensions "really quite strong" and refers to Condorcet extensions as "much more reasonable."

[4] Besides characterizing a class of SCFs that does not admit a strong manipulation, the proof of Theorem 3 shows something stronger about this class: In every manipulation where voters misrepresent strict preferences as well as indifferences, modifying the strict preferences is not necessary. The same outcome could have been obtained by only misrepresenting the indifference relation.

contradiction. To this end, let

$S = (R_1, \dots, R_{i-1}, R_i \setminus \{(a,b)\} \cup R'_i|_{\{a,b\}}, R_{i+1}, \dots, R_n)$ and
$S' = (R'_1, \dots, R'_{i-1}, R'_i \setminus \{\{(a,b)\}, \{b,a\}\} \cup R_i|_{\{a,b\}}, R'_{i+1}, \dots, R'_n)$.

In other words, S is identical to R, except that voter i's preferences over $\{a,b\}$ are as in R'. Similarly, S' is identical to R', except that voter i's preferences over $\{a,b\}$ are as in R.

Case 1 ($a \in X$ and $b \in Y$): $Y \widehat{R_i} X$ implies that $b R_i a$, a contradiction.

Case 2 ($a \notin X$): It follows from set-monotonicity that $f(S,A) = f(R,A) = X$. Thus, S and R' constitute a smaller counterexample since $|S \vartriangle R'| < |R \vartriangle R'|$.

Case 3 ($b \notin Y$): It follows from set-monotonicity that $f(S',A) = f(R',A) = Y$. Thus, R and S' constitute a smaller counterexample since $|R \vartriangle S'| < |R \vartriangle R'|$.

Hence, R and R' have to be identical, which concludes the proof. □

As mentioned before, when assuming that voters have strict preferences, weak strategyproofness and strategyproofness are equivalent. By showing that every monotonic SCF that satisfies SSP also satisfies set-monotonicity, we obtain the following useful corollary, the proof of which is omitted due to space restrictions.

Corollary 1. *Every monotonic SCF that satisfies SSP is weakly $\widehat{R}$-group-strategyproof.*

As mentioned in the introduction, there are few—but nevertheless quite attractive—SCFs that satisfy monotonicity and SSP, namely the *top cycle* (also known as *weak closure maximality*, *GETCHA*, or the *Smith set*), the *minimal covering set*, the *bipartisan set*, and their generalizations (see Bordes, 1976; Laslier, 1997; Dutta and Laslier, 1999). SSP and monotonicity do not completely characterize weak $\widehat{R}$-strategyproofness. SCFs that satisfy set-monotonicity but fail to satisfy SSP can easily be constructed.

Remarkably, the robustness of the minimal covering set and the bipartisan set with respect to strategic manipulation also extends to *agenda manipulation*. The strong superset property precisely states that an SCF is resistant to adding and deleting losing alternatives. Moreover, both SCFs are composition-consistent, i.e., they are strongly resistant to the introduction of clones (Laffond et al., 1996).[5] Scoring rules like plurality and Borda's rule are prone to both types of agenda manipulation as well as to strategic manipulation.

3.3 Weakly Group-Strategyproof Pairwise SCFs

In this section, we identify a natural and well-known class of SCFs for which the characterization given in the previous section is complete. An SCF f is said to be based on pairwise comparisons (or simply *pairwise*) if, for all preference profiles R, R' and feasible sets A, $f(R,A) = f(R',A)$ if and only if for all $a, b \in A$,

$$|\{i \in N \mid a\ P_i\ b\}| - |\{i \in N \mid b\ P_i\ a\}| = |\{i \in N \mid a\ P'_i\ b\}| - |\{i \in N \mid b\ P'_i\ a\}|.$$

[5] In addition to these attractive properties, the minimal covering set and the bipartisan set can be computed efficiently (Brandt and Fischer, 2008).

In other words, the outcome of a pairwise SCF only depends on the comparisons between pairs of alternatives (see, e.g., Young, 1974). The class of pairwise SCFs is quite natural and contains a large number of well-known voting rules such as Kemeny's rule, Borda's rule, maximin, ranked pairs, and all rules based on simple majorities (e.g., Copeland's rule, the Slater set, the top cycle, the uncovered set, the Banks set, the minimal covering set, and the bipartisan set). We now show that set-monotonicity is *necessary* for the $\widehat{R}$-strategyproofness of pairwise SCFs.

Theorem 4. *Every weakly $\widehat{R}$-strategyproof pairwise SCF satisfies set-monotonicity.*

Proof. We need to show that every pairwise SCF that fails to satisfy set-monotonicity is strongly $\widehat{R}$-manipulable. Suppose SCF f does not satisfy set-monotonicity. Then, there exists a feasible set A, a preference profile R, a voter i, and two alternatives $a, b \in A$ with $b\ R_i\ a$ and $b \notin f(R,A) = X$ such that $f(R',A) = Y \neq X$ where R' is a preference profile such that $R_j = R'_j$ for all $j \neq i$ and $R'_i = R_i^{(a,b)} \setminus R_i$.

First, define R_{n+1} and R'_{n+1} by letting

$$R_{n+1} = (U \times U) \setminus \{(a,b)\} \cup (R_i \cap \{(a,b)\}) \text{ and}$$
$$R'_{n+1} = (U \times U) \setminus \{(b,a)\} \cup (R'_i \cap \{(b,a)\}).$$

We now define two preference profiles with $n+1$ voters where voter i is indifferent between a and b and the crucial change in preference between a and b has been moved to voter $n+1$. Let

$$S = (R_1, \dots, R_{i-1}, R_i \cup \{(a,b)\}, R_{i+1}, \dots, R_n, R_{n+1}) \text{ and}$$
$$S' = (R'_1, \dots, R'_{i-1}, R'_i \cup \{(b,a)\}, R'_{i+1}, \dots, R'_n, R'_{n+1}).$$

It follows from the definition of pairwise SCFs that $f(S,A) = f(R,A) = X$ and $f(S',A) = f(R',A) = Y$. If $b\ P_i\ a$, we have $Y\ \widehat{R}_{n+1}\ X$ and f can be manipulated by voter $n+1$ at preference profile S by misstating his strict preference $b\ P_{n+1}\ a$ as $b\ I'_{n+1}\ a$. If, on the other hand, $b\ I_i\ a$, we have $X\ \widehat{R}'_{n+1}\ Y$ and f can be manipulated by voter $n+1$ at preference profile S' (by misstating his strict preference $a\ P'_{n+1}\ b$ as $a\ I_{n+1}\ b$). Hence, f is strongly $\widehat{R}$-manipulable. □

Proposition 1 entails the following useful corollary of Theorem 4.

Corollary 2. *Every weakly $\widehat{R}$-strategyproof pairwise SCF satisfies IUA.*

As a consequence, a large number of Condorcet extensions (e.g., Copeland's rule, the uncovered set, the Banks set, and the Slater set) are not weakly $\widehat{R}$-group-strategyproof because they are known to fail IUA (Laslier, 1997).

Theorem 3 and Theorem 4 establish that set-monotonicity is necessary and sufficient for the weak $\widehat{R}$-group-strategyproofness of pairwise SCFs, which can be seen as a generalization of Theorem 1 to irresolute SCFs within the setting of pairwise SCFs.

Theorem 5. *A pairwise SCF is weakly $\widehat{R}$-group-strategyproof if and only if it satisfies set-monotonicity.*

3.4 Weak Strategyproofness and Participation

Brams and Fishburn (1983) introduced a particularly natural variant of strategic manipulation where voters obtain a more preferred outcome by abstaining the election. An SCF is said to satisfy *participation* if voters are never better off by abstaining. A common criticism of Condorcet extensions is that they do not satisfy participation and thus suffer from the so-called *no-show paradox* (Moulin, 1988). However, Moulin's proof strongly relies on resoluteness. Irresolute Condorcet extensions that satisfy $\widehat{P}$-participation, which is defined in analogy to $\widehat{P}$-strategyproofness, do exist and, in the case of pairwise SCFs, there is a close connection between (weak) strategyproofness and participation as shown by the following proposition, the proof of which is omitted due to space restrictions.

Proposition 2. *Every weakly $\widehat{P}$-strategyproof pairwise SCF satisfies $\widehat{P}$-participation.*

Consequently, according to Kelly's preference extension, all SCFs satisfying set-monotonicity are immune to strategic abstention, which adds to the appeal of this compelling class of functions.

Acknowledgments

I am grateful to Markus Brill and Paul Harrenstein for helpful discussions. This material is based on work supported by the Deutsche Forschungsgemeinschaft under grants BR 2312/3-3 and BR 2312/7-1.

References

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

S. Barberà. A note on group strategy-proof decision schemes. *Econometrica*, 47(3):637–640, 1979.

S. Barberà. Strategy-proof social choice. In K. J. Arrow, A. K. Sen, and K. Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 2, chapter 25, pages 731–832. Elsevier, 2010.

J. Bartholdi, III, C. A. Tovey, and M. A. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.

G. Bordes. Consistency, rationality and collective choice. *Review of Economic Studies*, 43(3):451–457, 1976.

S. Brams and P. Fishburn. Paradoxes of preferential voting. *Mathematical Magazine*, 56(5):207–214, 1983.

F. Brandt and F. Fischer. Computing the minimal covering set. *Mathematical Social Sciences*, 56(2):254–268, 2008.

H. Chernoff. Rational selection of decision functions. *Econometrica*, 22:422–443, 1954.

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

V. Conitzer and T. Sandholm. Nonexistence of voting rules that are usually hard to manipulate. In *Proc. of 21st AAAI Conference*, pages 627–634. AAAI Press, 2006.

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.

P. Faliszewski and A. D. Procaccia. AI's war on manipulation: Are we winning? *AI Magazine*, 31(4):53–64, 2010.

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.

A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41:587–602, 1973.

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

M. Isaksson, G. Kindler, and E. Mossel. The geometry of manipulation: A quantitative proof of the gibbard-satterthwaite theorem. In *Proc. of 51st FOCS Symposium*, pages 319–328, 2010.

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

G. Laffond, J. Lainé, and J.-F. Laslier. Composition-consistent tournament solutions and social choice functions. *Social Choice and Welfare*, 13:75–93, 1996.

J.-F. Laslier. *Tournament Solutions and Majority Voting*. Springer, 1997.

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

E. Maskin. Nash equilibrium and welfare optimality. *Review of Economic Studies*, 66(26):23–38, 1999.

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

E. Muller and M. A. Satterthwaite. The equivalence of strong positive association and strategy-proofness. *Journal of Economic Theory*, 14(2):412–418, 1977.

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

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

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

T. Walsh. Where are the really hard manipulation problems? The phase transition in manipulating the veto rule. In *Proc. of 21st IJCAI*, pages 324–329, 2009.

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