

Necessary and Sufficient Conditions for the Strategyproofness of Irresolute Social Choice Functions

Felix Brandt
Technische Universität München
85748 Garching bei München, Germany
brandtf@in.tum.de

Markus Brill
Technische Universität München
85748 Garching bei München, Germany
brill@in.tum.de

ABSTRACT

While the Gibbard-Satterthwaite theorem states that every non-dictatorial and resolute, i.e., single-valued, social choice function is manipulable, it was recently shown that a number of appealing irresolute Condorcet extensions are strategyproof according to Kelly’s preference extension. In this paper, we study whether these results carry over to stronger preference extensions due to Fishburn and Gärdenfors. For both preference extensions, we provide sufficient conditions for strategyproofness and identify social choice functions that satisfy these conditions, answering a question by Gärdenfors [15] in the affirmative. We also show that some more discriminatory social choice functions fail to satisfy necessary conditions for strategyproofness.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent Systems

General Terms

Theory

Keywords

Social Choice Theory, Strategyproofness, Preference Extensions

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 [17, 29]. However, the classic result by Gibbard and Satterthwaite only applies to *resolute*, i.e., single-valued, SCFs. This assumption has been criticized for being unnatural and unreasonable [15, 22]. As Taylor [34] puts it, “If there is a weakness to the Gibbard-Satterthwaite theorem, it is the assumption that winners are

unique.” For example, consider a situation with two agents and two alternatives such that each agent prefers a different alternative. The problem is not that a resolute SCF has to select a single alternative (which is a well-motivated practical requirement), but that it has to select a single alternative *based on the individual preferences alone* (see, e.g., [22]). As a consequence, the SCF has to be biased towards an alternative or a voter (or both). Resoluteness is therefore at variance with such elementary fairness notions as neutrality (symmetry among the alternatives) and anonymity (symmetry among the voters).

In order to remedy this shortcoming, Gibbard [18] 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 [19, 3]). 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 [36] (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 for *irresolute* SCFs with varying underlying notions of how to compare sets of alternatives with each other (e.g., [15, 1, 2, 22, 10, 5, 8, 28, 35]), many of which are surveyed by Taylor [34] and Barberà [4].

How preferences over sets of alternatives relate to or depend on preferences over individual alternatives is a fundamental issue that goes back to at least de Finetti [9] and Savage [30]. In the context of social choice the alternatives are usually interpreted as mutually exclusive candidates for a unique final choice. For instance, assume an agent prefers a to b , b to c , and—by transitivity— a to c . What can we reasonably deduce from this about his preferences over the subsets of $\{a, b, c\}$? It stands to reason to assume that he would strictly prefer $\{a\}$ to $\{b\}$, and $\{b\}$ to $\{c\}$. If a single alternative is eventually chosen from each choice set, it is safe to assume that he also prefers $\{a\}$ to $\{b, c\}$ (Kelly’s extension), but whether he prefers $\{a, b\}$ to $\{a, b, c\}$ already depends on (his knowledge about) the final decision process. In the case of a lottery over all pre-selected alternatives according to a known *a priori* probability distribution with full support, he would prefer $\{a, b\}$ to $\{a, b, c\}$ (Fishburn’s

ACM COPYRIGHT NOTICE. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Publications Dept., ACM, Inc., fax +1 (212) 869-0481, or permissions@acm.org. TARK 2011, July 12-14, 2011, Groningen, The Netherlands. Copyright ©2011 ACM. ISBN 978-1-4503-0707-9, \$10.00.

extension). This assumption is, however, not sufficient to separate $\{a, b\}$ and $\{a, c\}$. Based on a sure-thing principle which prescribes that alternatives present in both choice sets can be ignored, it would be natural to prefer the former to the latter (Gärdenfors' extension). Finally, whether the agent prefers $\{a, c\}$ to $\{b\}$ depends on his attitude towards risk: he might hope for his most-preferred alternative (leximax extension), fear that his worst alternative will be chosen (leximin extension), or maximize his expected utility.

In general, there are at least three interdependent reasons why it is important to get a proper conceptual hold and a formal understanding of how preferences over sets relate to preferences over individual alternatives.

Rationality constraints. The examples above show that depending on the situation that is being modeled, preferences over sets are subject to certain rationality constraints, even if the preferences over individual alternatives are not. Not taking this into account would obviously be detrimental to a proper understanding of the situation at hand.

Epistemic and informational considerations. In many applications preferences over all subsets may be unavailable, unknown, or at least harder to obtain than preferences over the individual alternatives. With a proper grasp of how set preferences relate to preferences over alternatives, however, one may still be able to extract important structural information about the set preferences. In a similar vein, agents may not be fully informed about the situation they are in, e.g., they may not know the kind of lottery by means of which final choices are selected from sets. The less the agents know about the selection procedure, the less may be assumed about the structural properties of their preferences over sets.

Succinct representations. Clearly, as the set of subsets grows exponentially in the number of alternatives, preferences over subsets become prohibitively large. Hence, explicit representation and straightforward elicitation are not feasible and the succinct representation of set preferences becomes inevitable. Preferences over individual alternatives are of linear size and are the most natural basis for any succinct representation. Even when preferences over sets are succinctly represented by more elaborate structures than just preferences over individual alternatives, having a firm conceptual grasp on how set preferences relate to preferences over single alternatives is of crucial importance.

Any function that yields a preference relation over subsets of alternatives when given a preference relation over individual alternatives is called a *preference extension* or *set extension*. How to extend preferences to subsets is a fundamental issue that pervades the mathematical social sciences and has numerous applications in a variety of its disciplines. One example given by Gärdenfors [16] is the following: "suppose one only has ordinal information about the welfare of the members of society. When is it possible to say that one group of people is better off than another group?"

In this paper, we will be concerned with three of the most well-known preference extensions due to Kelly [22], Fishburn [14], and Gärdenfors [15]. On the one hand, we provide sufficient conditions for strategyproofness and identify social choice functions that satisfy these conditions. For example, we show that the top cycle is strategyproof according to Gär-

denfors' set extension, answering a question by Gärdenfors [15] in the affirmative. On the other hand, we propose necessary conditions for strategyproofness and show that some more discriminatory social choice functions such as the minimal covering set and the bipartisan set, which have recently been shown to be strategyproof according to Kelly's extension, fail to satisfy strategyproofness according to Fishburn's and Gärdenfors' extension. By means of a counter-example, we also show that Gärdenfors [15] *incorrectly* claimed that the SCF that returns the Condorcet winner when it exists and all Pareto-undominated alternatives otherwise is strategyproof according to Gärdenfors' extension.

2. PRELIMINARIES

In this section, we provide the terminology and notation required for our results.

2.1 Social Choice Functions

Let $N = \{1, \dots, n\}$ be a set of voters with preferences over a finite set A of alternatives. The preferences of voter $i \in N$ are represented by a complete and anti-symmetric *preference relation* $R_i \subseteq A \times A$.¹ The interpretation of $(a, b) \in R_i$, usually denoted by $a R_i b$, is that voter i values alternative a at least as much as alternative b . In accordance 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 $b R_i a$. As R_i is anti-symmetric, $a P_i b$ if and only if $a R_i b$ and $a \neq b$. The set of all preference relations over A will be denoted by $\mathcal{R}(A)$. The set of *preference profiles*, i.e., finite vectors of preference relations, is then given by $\mathcal{R}^*(A)$. The typical element of $\mathcal{R}^*(A)$ will be $R = (R_1, \dots, R_n)$.

The following notational convention will turn out to be useful. For a given preference profile R with $b R_i a$, $R_{i:(a,b)}$ denotes the preference profile

$$R_{i:(a,b)} = (R_1, \dots, R_{i-1}, R_i \setminus \{(b, a)\} \cup \{(a, b)\}, R_{i+1}, \dots, R_n).$$

That is, $R_{i:(a,b)}$ is identical to R except that alternative a is strengthened with respect to b within voter i 's preference relation.

Our central object of study are *social choice functions*, i.e., functions that map the individual preferences of the voters to a non-empty set of socially preferred alternatives.

DEFINITION 1. A *social choice function (SCF)* is a function $f : \mathcal{R}^*(A) \rightarrow 2^A \setminus \emptyset$.

An SCF f is said to be based on pairwise comparisons (or simply *pairwise*) if, for all preference profiles R and R' , $f(R) = f(R')$ whenever for all alternatives a, b ,

$$\begin{aligned} & |\{i \in N \mid a R_i b\}| - |\{i \in N \mid b R_i a\}| \\ &= |\{i \in N \mid a R'_i b\}| - |\{i \in N \mid b R'_i a\}|. \end{aligned}$$

In other words, the outcome of a pairwise SCF only depends on the comparisons between pairs of alternatives (see, e.g., [37, 38]).

¹For most of our results, we do not assume transitivity of preferences. In fact, Theorems 3 and 5 become stronger but are easier to prove for general—possibly intransitive—preferences. Theorems 4 and 6, on the other hand, become slightly weaker because there exist SCFs that are only manipulable if intransitive preferences are allowed. For all the manipulable SCFs in this paper, however, we show that they are manipulable even if transitive preferences are required.

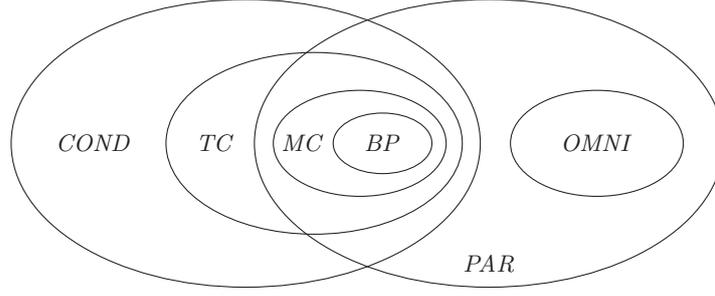


Figure 1: Set-theoretic relationships between the choice sets of the SCFs considered in this paper. The choice sets of SCFs that intersect in the diagram *always* intersect. The same is true for set-inclusions. Choice sets of SCFs that are disjoint in the diagram *may* have an empty intersection, i.e., there exist instances where the choice sets do not intersect.

For a given preference profile $R = (R_1, \dots, R_n)$, the *majority relation* $R_M \subseteq A \times A$ is defined by $a R_M b$ if and only if $|\{i \in N \mid a R_i b\}| \geq |\{i \in N \mid b R_i a\}|$. Let P_M denote the strict part of R_M . A *Condorcet winner* is an alternative a that is preferred to any other alternative by a strict majority of voters, i.e., $a P_M b$ for all alternatives $b \neq a$. An SCF is called a *Condorcet extension* if it uniquely selects the Condorcet winner whenever one exists.

We will now introduce the SCFs considered in this paper. With the exception of the Pareto rule and the omninomination rule, all of these SCFs are pairwise Condorcet extensions. Set-theoretical relationships between these SCFs are illustrated in Figure 1.

Pareto rule An alternative a is *Pareto-dominated* if there exists an alternative b such that $b P_i a$ for all voters $i \in N$. The Pareto rule PAR returns all alternatives that are *not* Pareto-dominated.

Omninomination rule The omninomination rule $OMNI$ returns all alternatives that are ranked first by at least one voter.

Condorcet rule The Condorcet rule $COND$ returns the Condorcet winner if it exists, and all alternatives otherwise.

Top Cycle Let R_M^* denote the transitive closure of the majority relation, i.e., $a R_M^* b$ if and only if there exists $k \in \mathbb{N}$ and $a_1, \dots, a_k \in A$ with $a_1 = a$ and $a_k = b$ such that $a_i R_M a_{i+1}$ for all $i < k$. The top cycle rule TC (also known as *weak closure maximality*, *GETCHA*, or the *Smith set*) returns the maximal elements of R_M^* , i.e., $TC(R) = \{a \in A \mid a R_M^* b \text{ for all } b \in A\}$ [20, 33, 31].

Minimal Covering Set A subset $C \subseteq A$ is called a *covering set* if for all alternatives $b \in A \setminus C$, there exists $a \in C$ such that $a P_M b$ and for all $c \in C \setminus \{a\}$, $b P_M c$ implies $a P_M c$ and $c P_M a$ implies $c P_M b$. Dutta [11] and Dutta and Laslier [12] have shown that there always exists a unique *minimal* covering set. The SCF MC returns exactly this set.

Bipartisan Set Consider the two-player zero-sum game in which the set of actions for both players is given by A and payoffs are defined as follows. If the first player chooses a and the second player chooses b , the payoff

for the first player is 1 if $a P_M b$, -1 if $b P_M a$, and 0 otherwise. The bipartisan set BP contains all alternatives that are played with positive probability in some Nash equilibrium of this game [23, 12].

Observe that PAR and $OMNI$ are only well-defined for transitive individual preferences.

2.2 Strategyproofness

An SCF is manipulable if one or more voters can misrepresent their preferences in order to obtain a more preferred choice set. While comparing choice sets is trivial for resolute SCFs, this is not the case for irresolute ones. 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 our investigation of strategyproof SCFs, we will consider the following three well-known set extensions due to Kelly [22], Fishburn [14],² and Gärdenfors [15]. Let R_i be a preference relation over A and $X, Y \subseteq A$ two non-empty subsets of A .

- $X R_i^K Y$ if and only if $x R_i y$ for all $x \in X$ and all $y \in Y$ [22]

One interpretation of this extension is that voters are unaware of the mechanism (e.g., a lottery) that will be used to pick the winning alternative [16].

- $X R_i^F Y$ if and only if $x R_i y$, $x R_i z$, and $y R_i z$ for all $x \in X \setminus Y$, $y \in X \cap Y$, and $z \in Y \setminus X$ [14]

One interpretation of this extension is that the winning alternative is picked by a lottery according to some underlying *a priori* distribution and that voters are unaware of this distribution [8]. Alternatively, one may assume the existence of a tie-breaker with linear, but unknown, preferences.

- $X R_i^G Y$ if and only if one of the following conditions is satisfied [15]:

- (i) $X \subset Y$ and $x R_i y$ for all $x \in X$ and $y \in Y \setminus X$
- (ii) $Y \subset X$ and $x R_i y$ for all $x \in X \setminus Y$ and $y \in Y$
- (iii) neither $X \subset Y$ nor $Y \subset X$ and $x R_i y$ for all $x \in X \setminus Y$ and $y \in Y \setminus X$

²Gärdenfors [16] attributed this extension to Fishburn because it is the weakest extension that satisfies a certain set of axioms proposed by Fishburn [14].

No interpretation in terms of lotteries is known for this set extension. Gärdenfors [15] motivates it by alluding to Savage’s sure-thing principle (when comparing two options, identical parts may be ignored). Unfortunately, the definition of this extension is somewhat “discontinuous,” which is also reflected in the hardly elegant characterization given in Theorem 5.

It is easy to see that these extensions form an inclusion hierarchy.

FACT 1. For all preference relations R_i and subsets $X, Y \subseteq A$,

$$X R_i^K Y \text{ implies } X R_i^F Y \text{ implies } X R_i^G Y.$$

For $\mathcal{E} \in \{K, F, G\}$, let $P_i^\mathcal{E}$ denote the strict part of $R_i^\mathcal{E}$. As R_i is anti-symmetric, so is $R_i^\mathcal{E}$. Therefore, we have $X P_i^\mathcal{E} Y$ if and only if $X R_i^\mathcal{E} Y$ and $X \neq Y$.

Based on these set extensions, we can now define three different notions of strategyproofness for irresolute SCFs. Note that, in contrast to some related papers, we interpret preference extensions as fully specified (incomplete) preference relations rather than minimal conditions on set preferences.

DEFINITION 2. Let $\mathcal{E} \in \{K, F, G\}$. An SCF f is $P^\mathcal{E}$ -manipulable by a group of voters $C \subseteq N$ if there exist preference profiles R and R' with $R_j = R'_j$ for all $j \notin C$ such that

$$f(R') P_i^\mathcal{E} f(R) \text{ for all } i \in C.$$

An SCF is $P^\mathcal{E}$ -strategyproof if it is not $P^\mathcal{E}$ -manipulable by single voters. An SCF is $P^\mathcal{E}$ -group-strategyproof if it is not $P^\mathcal{E}$ -manipulable by any group of voters.

Fact 1 implies that P^G -group-strategyproofness is stronger than P^F -group-strategyproofness, which in turn is stronger than P^K -group-strategyproofness.

3. RELATED WORK

Barberà [1] and Kelly [22] have shown independently that all non-trivial SCFs that are rationalizable via a quasi-transitive preference relation are P^K -manipulable. However, as witnessed by various other (non-strategic) impossibility results that involve quasi-transitive rationalizability (e.g., [24]), it appears as if this property itself is unduly restrictive. As a consequence, Kelly [22] 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.”

Strengthening earlier results by Gärdenfors [15] and Taylor [34], Brandt [6] showed that no Condorcet extension is P^K -strategyproof. The proof, however, crucially depends on strategic tie-breaking and hence does not work for anti-symmetric preferences. For this reason, only anti-symmetric preferences are considered in the present paper.

Brandt [6] also provided a sufficient condition for P^K -group-strategyproofness. *Set-monotonicity* can be seen as an irresolute variant of Maskin-monotonicity [25] and prescribes that the choice set is invariant under the weakening of unchosen alternatives.

DEFINITION 3. An SCF f satisfies *set-monotonicity* (*SET-MON*) if $f(R_{i:(a,b)}) = f(R)$ for all preference profiles R , voters i , and alternatives a, b with $b \notin f(R)$.

THEOREM 1 (Brandt [6]). *Every SCF that satisfies SET-MON is P^K -group-strategyproof.*

Set-monotonicity is a demanding condition, but a handful of SCFs such as the ones introduced in Section 2.1 are known to be set-monotonic. For the class of *pairwise* SCFs, this condition is also necessary, which shows that many well-known SCFs such as Borda’s rule, Copeland’s rule, Kemeny’s rule, the uncovered set, and the Banks set are not P^K -group-strategyproof.

THEOREM 2 (Brandt [6]). *Every pairwise SCF that is P^K -group-strategyproof satisfies SET-MON.*

Strategyproofness according to Kelly’s extension thus draws a sharp line within the space of SCFs as almost all established non-pairwise SCFs (such as plurality and all weak Condorcet extensions like Young’s rule) are also known to be P^K -manipulable (see, e.g., [34]).

The state of affairs for Gärdenfors’ and Fishburn’s extensions is less clear. Gärdenfors [15] has shown that *COND* and *OMNI* are P^G -group-strategyproof. In an attempt to extend this result to more discriminatory SCFs, he also claimed that $COND \cap PAR$, which returns the Condorcet winner if it exists and all Pareto-undominated alternatives otherwise, is P^G -strategyproof. However, we show that this is not the case (Proposition 2). Gärdenfors concludes that “we have not been able to find any more decisive function which is stable [strategyproof] and satisfies minimal requirements on democratic decision functions.” We show that *TC* is such a function (Corollary 1).

Apart from a theorem by Ching and Zhou [8], which uses an unusually strong definition of strategyproofness, we are not aware of any characterization result using Fishburn’s extension. Feldman [13] has shown that the Pareto rule is P^F -strategyproof and Sanver and Zwicker [27] have shown that the same is true for *TC*.

4. RESULTS

This section contains our results. All omitted proofs can be found in the appendix.

4.1 Necessary and Sufficient Conditions for Group-Strategyproofness

We first introduce a new property that requires that modifying preferences between chosen alternatives may only result in smaller choice sets. Set-monotonicity entails a condition called *independence of unchosen alternatives*, which states that the choice set is invariant under modifications of the preferences between unchosen alternatives. Accordingly, the new property will be called *exclusive independence of chosen alternatives*, where “exclusive” refers to the requirement that unchosen alternatives remain unchosen.

DEFINITION 4. An SCF f satisfies *exclusive independence of chosen alternatives* (*EICA*) if $f(R') \subseteq f(R)$ for all pairs of preference profiles R and R' that differ only on alternatives in $f(R)$, i.e., $R_i|_{\{a,b\}} = R'_i|_{\{a,b\}}$ for all $i \in N$ and all alternatives a, b with $b \notin f(R)$.

It turns out that, together with SET-MON, this new property is sufficient for an SCF to be group-strategyproof according to Fishburn’s preference extension.

THEOREM 3. *Every SCF that satisfies SET-MON and EICA is P^F -group-strategyproof.*

For *pairwise* SCFs, the following weakening of EICA can be shown to be necessary for group-strategyproofness according to Fishburn’s extension. It prescribes that modify-

ing preferences among chosen alternatives does not result in a choice set that is a strict superset of the original choice set.

DEFINITION 5. An SCF f satisfies *weak EICA* if $f(R) \not\subseteq f(R')$ for all pairs of preference profiles R and R' that differ only on alternatives in $f(R)$.

THEOREM 4. *Every pairwise SCF that is P^F -group-strategy-proof satisfies SET-MON and weak EICA.*

We now turn to P^G -group-strategyproofness. When comparing two sets, P^G differs from P^F only in the case when neither set is contained in the other. The following definition captures exactly this case.

DEFINITION 6. An SCF f satisfies the *symmetric difference property (SDP)* if either $f(R) \subseteq f(R')$ or $f(R') \subseteq f(R)$ for all pairs of preference profiles R and R' such that $R_i|_{\{a,b\}} = R'_i|_{\{a,b\}}$ for all $i \in N$ and all alternatives a, b with $a \in f(R) \setminus f(R')$ and $b \in f(R') \setminus f(R)$.

THEOREM 5. *Every SCF that satisfies SET-MON, EICA, and SDP is P^G -group-strategyproof.*

As was the case for Fishburn's extension, a set of necessary conditions for pairwise SCFs can be obtained by replacing EICA with weak EICA.

THEOREM 6. *Every pairwise SCF that is P^G -group-strategy-proof satisfies SET-MON, weak EICA, and SDP.*

4.2 Consequences

We are now ready to study the strategyproofness of the SCFs defined in Section 2. It can be checked that *COND* and *TC* satisfy SET-MON, EICA, and SDP and thus, by Theorem 5, are P^G -group-strategyproof.

COROLLARY 1. *COND and TC are P^G -group-strategyproof.*

OMNI, *PAR*, and $COND \cap PAR$ satisfy SET-MON and EICA, but not SDP.

COROLLARY 2. *OMNI, PAR, and $COND \cap PAR$ are P^F -group-strategyproof.*

As *OMNI*, *PAR*, and $COND \cap PAR$ are not pairwise, the fact that they violate SDP does *not* imply that they are P^G -manipulable. In fact, it turns out that *OMNI* is strategyproof according to Gärdenfors' extension, while *PAR* and $COND \cap PAR$ are not.

PROPOSITION 1. *OMNI is P^G -group-strategyproof.*

PROPOSITION 2. *PAR and $COND \cap PAR$ are P^G -manipulable.*

Proof. Consider the following profile $R = (R_1, R_2, R_3, R_4)$.

R_1	R_2	R_3	R_4
c	c	a	a
d	d	b	b
b	a	c	c
a	b	d	d

It is easily verified that $PAR(R) = \{a, b, c\}$. Now let $R' = (R'_1, R_2, R_3, R_4)$ where $d R'_1 c R'_1 a R'_1 b$. Obviously, $PAR(R') = \{a, c, d\}$ and $\{a, c, d\} P_1^G \{a, b, c\}$ because $d R_1 b$. I.e., the first voter can obtain a preferable choice set by misrepresenting his preferences. As neither R nor R' has a Condorcet winner, the same holds for $COND \cap PAR$. \square

Finally, we show that *MC* and *BP* violate weak EICA, which implies that both rules are manipulable according to Fishburn's extension.

COROLLARY 3. *MC and BP are P^F -manipulable.*

Proof. By Theorem 4 and the fact that both *MC* and *BP* are pairwise, it suffices to show that *MC* and *BP* violate weak EICA. To this end, consider the following profile $R = (R_1, R_2, R_3, R_4, R_5)$ and the corresponding majority graph representing P_M .

R_1	R_2	R_3	R_4	R_5
d	c	b	e	d
e	b	c	a	c
a	a	e	b	a
b	e	a	d	b
c	d	d	c	e

It can be checked that $MC(R) = BP(R) = \{a, b, c\}$. Define $R' = R_{1:(c,b)}$, i.e., the first voter strengthens c with respect to b . Observe that P_M and P'_M disagree on the pair $\{b, c\}$, and that $MC(R') = BP(R') = \{a, b, c, d, e\}$. Thus, both *MC* and *BP* violate weak EICA and the first voter can manipulate because $\{a, b, c, d, e\} P_1^F \{a, b, c\}$. \square

The same example shows that the tournament equilibrium set [32] and the minimal extending set [7], both of which are only defined for an odd number of voters and conjectured to be P^K -group-strategyproof, are P^F -manipulable.

5. CONCLUSION

In this paper, we investigated the effect of various preference extensions on the manipulability of irresolute SCFs. We proposed necessary and sufficient conditions for strategyproofness according to Fishburn's and Gärdenfors' set extensions and used these conditions to illuminate the strategyproofness of a number of well-known SCFs. Our results are summarized in Table 1. As mentioned in Section 3, some of these results were already known or—in the case of P^F -strategyproofness of the top cycle—have been discovered independently by other authors. In contrast to the papers by Gärdenfors [15], Feldman [13], and Sanver and Zwicker [27], which more or less focus on particular SCFs, our axiomatic approach yields unified proofs of most of the statements in the table.³

Many interesting open problems remain. For example, it is not known whether there exists a Pareto-optimal pairwise SCF that is strategyproof according to Gärdenfors' extension. Recently, the study of the manipulation of irresolute SCFs by other means than untruthfully representing one's preferences—e.g., by abstaining the election [26, 21]—has been initiated. For the set extensions considered in this paper it is unknown which SCFs can be manipulated by abstention. It would be desirable to also obtain characterizations of these classes of SCFs and, more generally, to improve our understanding of the interplay between both types of manipulation. For instance, it is not difficult to show that the negative results in Corollary 3 also extend to manipulation by abstention.

³The results in the leftmost column of Table 1 are due to Brandt [6] and are included for the sake of completeness.

Table 1: Summary of results.

	P^K -strategyproof	P^F -strategyproof	P^G -strategyproof
<i>OMNI</i>	✓	✓	✓ ^a
<i>COND</i>	✓	✓	✓ ^a
<i>TC</i>	✓	✓ ^b	✓
<i>PAR</i>	✓	✓ ^c	–
<i>COND</i> \cap <i>PAR</i>	✓	✓	–
<i>MC</i>	✓	–	–
<i>BP</i>	✓	–	–

^aGärdenfors [15]

^bSanver and Zwicker [27]

^cFeldman [13]

Another interesting related question concerns the epistemic foundations of the above extensions. Most of the literature in social choice theory focusses on well-studied economic models where agents have full knowledge of a random selection process, which is often assumed to be a lottery with uniform probabilities. The study of more intricate distributed protocols or computational selection devices that justify certain set extensions appears to be very promising. For instance, Kelly’s set extension could be justified by a distributed protocol for “unpredictable” random selections that do not permit a meaningful prior distribution.

6. ACKNOWLEDGMENTS

This material is based on work supported by the Deutsche Forschungsgemeinschaft under grants BR 2312/6-1 and BR 2312/7-1. We gratefully acknowledge the support of the TUM Graduate School’s Faculty Graduate Center CeDoSIA at Technische Universität München, Germany.

REFERENCES

- [1] S. Barberà. Manipulation of social decision functions. *Journal of Economic Theory*, 15(2):266–278, 1977.
- [2] S. Barberà. The manipulation of social choice mechanisms that do not leave “too much” to chance. *Econometrica*, 45(7):1573–1588, 1977.
- [3] S. Barberà. A note on group strategy-proof decision schemes. *Econometrica*, 47(3):637–640, 1979.
- [4] 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.
- [5] S. Barberà, B. Dutta, and A. Sen. Strategy-proof social choice correspondences. *Journal of Economic Theory*, 101(2):374–394, 2001.
- [6] F. Brandt. Group-strategyproof irresolute social choice functions. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, 2011. Forthcoming.
- [7] F. Brandt. Minimal stable sets in tournaments. *Journal of Economic Theory*, 2011. Forthcoming.
- [8] S. Ching and L. Zhou. Multi-valued strategy-proof social choice rules. *Social Choice and Welfare*, 19:569–580, 2002.
- [9] B. de Finetti. La prévision: ses lois logiques, ses sources subjectives. *Annales de l’institut Henri Poincaré*, 7(1):1–68, 1937.
- [10] J. Duggan and T. Schwartz. Strategic manipulability without resoluteness or shared beliefs: Gibbard-Satterthwaite generalized. *Social Choice and Welfare*, 17(1):85–93, 2000.
- [11] B. Dutta. Covering sets and a new Condorcet choice correspondence. *Journal of Economic Theory*, 44:63–80, 1988.
- [12] B. Dutta and J.-F. Laslier. Comparison functions and choice correspondences. *Social Choice and Welfare*, 16(4):513–532, 1999.
- [13] A. Feldman. Manipulation and the Pareto rule. *Journal of Economic Theory*, 21:473–482, 1979.
- [14] P. C. Fishburn. Even-chance lotteries in social choice theory. *Theory and Decision*, 3:18–40, 1972.
- [15] P. Gärdenfors. Manipulation of social choice functions. *Journal of Economic Theory*, 13(2):217–228, 1976.
- [16] P. Gärdenfors. On definitions of manipulation of social choice functions. In J. J. Laffont, editor, *Aggregation and Revelation of Preferences*. North-Holland, 1979.
- [17] A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41:587–602, 1973.
- [18] A. Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45(3):665–681, 1977.
- [19] A. Gibbard. Straightforwardness of game forms with lotteries as outcomes. *Econometrica*, 46(3):595–614, 1978.
- [20] I. J. Good. A note on Condorcet sets. *Public Choice*, 10:97–101, 1971.
- [21] 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.
- [22] J. Kelly. Strategy-proofness and social choice functions without single-valuedness. *Econometrica*, 45(2):439–446, 1977.
- [23] G. Laffond, J.-F. Laslier, and M. Le Breton. The bipartisan set of a tournament game. *Games and Economic Behavior*, 5:182–201, 1993.
- [24] A. Mas-Colell and H. Sonnenschein. General possibility theorems for group decisions. *Review of Economic Studies*, 39(2):185–192, 1972.

- [25] E. Maskin. Nash equilibrium and welfare optimality. *Review of Economic Studies*, 66(26):23–38, 1999.
- [26] J. Pérez. The strong no show paradoxes are a common flaw in Condorcet voting correspondences. *Social Choice and Welfare*, 18(3):601–616, 2001.
- [27] M. R. Sanver and W. S. Zwicker. Monotonicity properties and their adaption to irresolute social choice rules. Unpublished Manuscript, 2010.
- [28] S. Sato. On strategy-proof social choice correspondences. *Social Choice and Welfare*, 31:331–343, 2008.
- [29] 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.
- [30] L. J. Savage. *The Foundations of Statistics*. Wiley Publications in Statistics. Wiley and Sons, 1954.
- [31] T. Schwartz. *The Logic of Collective Choice*. Columbia University Press, 1986.
- [32] T. Schwartz. Cyclic tournaments and cooperative majority voting: A solution. *Social Choice and Welfare*, 7: 19–29, 1990.
- [33] J. H. Smith. Aggregation of preferences with variable electorate. *Econometrica*, 41(6):1027–1041, 1973.
- [34] A. D. Taylor. *Social Choice and the Mathematics of Manipulation*. Cambridge University Press, 2005.
- [35] M. Umezawa. Coalitionally strategy-proof social choice correspondences and the Pareto rule. *Social Choice and Welfare*, 33:151–158, 2009.
- [36] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 2nd edition, 1947.
- [37] H. P. Young. An axiomatization of Borda’s rule. *Journal of Economic Theory*, 9:43–52, 1974.
- [38] W. S. Zwicker. The voter’s paradox, spin, and the Borda count. *Mathematical Social Sciences*, 22:187–227, 1991.

APPENDIX

A. OMITTED PROOFS

A.1 Notation

For a preference relation R_i , let R_i^- denote the preference relation where all preferences are reversed, i.e., $a R_i^- b$ if and only if $b R_i a$.

Furthermore, define the *distance* $\delta(R_i, R'_i)$ between two preference relations R_i and R'_i as the number of (unordered) pairs of alternatives on which they disagree. As preference relations are complete and anti-symmetric, $\delta(R_i, R'_i) = |R_i \setminus R'_i| = |R'_i \setminus R_i|$. The distance between two preference profiles R and R' of length n is defined as $\delta(R, R') = \sum_{i=1}^n \delta(R_i, R'_i)$.

A.2 Proof of Theorem 3

We first need a lemma, which states that EICA together with SET-MON implies the following: if only preferences between chosen alternatives are modified and some alternatives leave the choice set, then at least one of them was weakened with respect to an alternative that remains chosen.

LEMMA 1. *Let f be an SCF that satisfies SET-MON and EICA and consider a pair of profiles R, R' that differ only on alternatives in $f(R)$. If $f(R') \subset f(R)$, then there exist $i \in N$, $x \in f(R) \setminus f(R')$ and $y \in f(R')$ such that $x R_i y$ and $y R'_i x$.*

Proof. Assume for contradiction that $f(R') \subset f(R)$ and $R' \setminus R = \bigcup_{i \in N} (R'_i \setminus R_i)$ does not contain a pair (y, x) with $y \in f(R')$ and $x \in f(R) \setminus f(R')$. Then each pair $(y, x) \in R' \setminus R$ belongs to exactly one of the following two classes.

CLASS 1. $y, x \in f(R')$

CLASS 2. $y \in f(R) \setminus f(R')$, $x \in A$

We now start with preference profile R' and change the preferences in $R' \setminus R$ one after the other to arrive at profile R . We first change the preferences for all pairs (y, x) from Class 1 and denote the resulting profile by R'' . As R' and R'' differ only on alternatives in $f(R')$, EICA implies that $f(R'') \subseteq f(R')$. We then change the preferences for all pairs (y, x) from Class 2. By definition, the resulting profile is R . As $f(R'') \subseteq f(R')$, $y \notin f(R')$ implies $y \notin f(R'')$. Thus, in this second step, only alternatives $y \notin f(R'')$ are weakened and SET-MON implies that $f(R) = f(R'')$. But $f(R) = f(R'') \subseteq f(R')$ contradicts the assumption that $f(R')$ is a strict subset of $f(R)$. \square

THEOREM 3. *Every SCF that satisfies SET-MON and EICA is P^F -group-strategyproof.*

Proof. Let f be an SCF that satisfies SET-MON and EICA and assume for contradiction that f is not P^F -group-strategyproof. Then, there have to be a group of voters $C \subseteq N$ and two preference profiles R and R' with $R_j = R'_j$ for all $j \notin C$ such that $f(R') P_i^F f(R)$ for all $i \in C$. We choose R and R' such that $\delta(R, R')$ is minimal, i.e., we look at a “smallest” counter-example in the sense that R and R' coincide as much as possible. Let $f(R) = X$ and $f(R') = Y$. We may assume $\delta(R, R') > 0$ as otherwise $R = R'$ and $X = Y$. Now, consider a pair of alternatives $a, b \in A$ such that, for some $i \in C$, $a R_i b$ and $b R'_i a$, i.e., voter i misrepresents his preference relation by strengthening b . The following argument will show that no such a and b exist, which implies that R and R' and consequently X and Y are identical, a contradiction. We need the following two claims.

CLAIM 1. $b \in Y$

In order to prove this claim, suppose that $b \notin Y$. It follows from SET-MON that $f(R'_{i:(a,b)}) = f(R') = Y$. Thus, R and $R'_{i:(a,b)}$ constitute a smaller counter-example since $\delta(R, R'_{i:(a,b)}) = \delta(R, R') - 1$. This is a contradiction because $\delta(R, R')$ was assumed to be minimal.

CLAIM 2. $a \in Y$

The following case distinction shows that Claim 2 holds. Suppose $a \notin Y$. If $a \notin X$ either, SET-MON implies that $f(R_{i:(b,a)}) = f(R) = X$. Thus, $R_{i:(b,a)}$ and R' constitute a smaller counter-example since $\delta(R_{i:(b,a)}, R') = \delta(R, R') - 1$. On the other hand, if $a \in X \setminus Y$, $b \in Y$ and $a R_i b$ contradict the assumption that $Y P_i^F X$.

We thus have $\{a, b\} \subseteq Y$ for every pair (a, b) such that some voter $i \in C$ misrepresents his preference between a and b . In particular, this means that R and R' differ only

on alternatives in $Y = f(R')$. Therefore, Lemma 1 implies⁴ that either $X = Y$ or $X \subset Y$ and there exist $y \in Y \setminus X$ and $x \in X$ such that $y R'_i x$ and $x R_i y$. Both cases contradict the assumption that $Y P_i^F X$.

Hence, we have shown that no such R and R' exist, which concludes the proof. \square

Note that the preferences of voter i in the profile $R_{i:(b,a)}$ might not be transitive. Therefore, one has to be careful when applying the preceding proof to *PAR* and *OMNI*, as those SCFs are only defined for transitive preferences. One can however generalize the definition of both SCFs to intransitive preference profiles in such a way that all arguments in the proof remain valid.⁵

A.3 Proof of Theorem 4

THEOREM 4. *Every pairwise SCF that is P^F -group-strategy-proof satisfies SET-MON and weak EICA.*

Proof. We need to show that every pairwise SCF that violates either SET-MON or weak EICA is P^F -manipulable.

First, let f be a pairwise SCF that violates SET-MON.⁶ Then, there exist a preference profile $R = (R_1, \dots, R_n)$, a voter i , and two alternatives a, b with $b R_i a$ and $b \notin f(R) = X$ such that $f(R_{i:(a,b)}) = Y \neq X$.

Let R_{n+1} be a preference relation such that $b R_{n+1} a$ and $Y P_{n+1}^F X$ (such a relation exists because $b \notin X$) and let $R_{n+2} = R_{n+1}^-$. Let S denote the preference profile $S = (R_1, \dots, R_n, R_{n+1}, R_{n+2})$. It follows from the definition of pairwise SCFs that $f(S) = f(R) = X$ and $f(S_{n+1:(a,b)}) = f(R_{i:(a,b)}) = Y$.

As $Y P_{n+1}^F X$, we have that f can be manipulated by voter $n+1$ at preference profile S by misstating his preference $b R_{n+1} a$ as $a R_{n+1} b$. Hence, f is P^F -manipulable.

Second, let f be a pairwise SCF that violates weak EICA. Then, there exist two preference profiles $R = (R_1, \dots, R_n)$ and $R' = (R'_1, \dots, R'_n)$ that differ only on alternatives in $f(R)$, such that $f(R) = X \subset Y = f(R')$. Let $C \subseteq N$ be the group of voters that have different preferences in R and R' , i.e., $C = \{i \in N \mid R_i \neq R'_i\}$. Without loss of generality, we can assume that $C = \{1, \dots, c\}$, where $c = |C|$. For all $i \in \{1, \dots, c\}$, let R_{n+i} be a preference relation such that $Y P_i^F X$ and $R_i \setminus R'_i \subseteq R_{n+i}$ (such a preference relation exists because $X \subset Y$ and $R_i \setminus R'_i \subseteq X \times X$) and let $R_{n+c+i} = R_{n+i}^-$. Furthermore, for all $i \in \{1, \dots, c\}$, let $R'_{n+i} = R_{n+i} \setminus R_i \cup R'_i$. I.e., R'_{n+i} differs from R_{n+i} on exactly the same pairs of alternatives as R'_i differs from R_i .

⁴Observe that we apply Lemma 1 with $f(R) = Y$ and $f(R') = X$.

⁵To see this, define *OMNI*(R) to contain all those alternatives a for which there exists a voter i with $a R_i b$ for all $b \neq a$. The definition of *PAR* can remain unchanged. Generalized in this way, the choice set of either function may be empty for intransitive preferences. It can however easily be shown that *PAR* and *OMNI* still satisfy SET-MON in the case of non-empty choice sets. As the sets X and Y used in the proof of Theorem 3 are non-empty, the latter condition is then sufficient for the argument in the proof to go through.

⁶Brandt [6] has shown that this implies P^K -manipulability in a setting where ties are allowed.

Consider the preference profiles

$$S = (R_1, \dots, R_n, R_{n+1}, \dots, R_{n+c}, R_{n+c+1}, \dots, R_{n+2c}) \text{ and} \\ S' = (R_1, \dots, R_n, R'_{n+1}, \dots, R'_{n+c}, R_{n+c+1}, \dots, R_{n+2c}).$$

It follows from the definition of pairwise SCFs that $f(S) = f(R) = X$ and $f(S') = f(R') = Y$. As $Y P_{n+i}^F X$ for all $i \in \{1, \dots, c\}$, we have that f can be manipulated by the group $\{n+1, \dots, n+c\}$ at preference profile S by misstating their preferences R_{n+i} as R'_{n+i} . Hence, f is P^F -manipulable. \square

A.4 Proof of Theorem 5

THEOREM 5. *Every SCF that satisfies SET-MON, EICA, and SDP is P^G -group-strategyproof.*

Proof. Let f be an SCF that satisfies SET-MON, EICA, and SDP, and assume for contradiction that f is not P^G -group-strategyproof. Then, there have to be a group of voters $C \subseteq N$ and two preference profiles R and R' with $R_j = R'_j$ for all $j \notin C$ such that $f(R') P_i^G f(R)$ for all $i \in C$. Choose R and R' such that $\delta(R, R')$ is minimal and let $X = f(R)$ and $Y = f(R')$.

As P^G coincides with P^F on all pairs where one set is contained in the other set, and, by Theorem 3, f is P^F -group-strategyproof, we can conclude that neither $X \subseteq Y$ nor $Y \subseteq X$. Thus, SDP implies that there exist pairs (x, y) with $x \in X \setminus Y$ and $y \in Y \setminus X$ such that some voters have modified their preference between x and y , i.e., $(x, y) \in (R_i \setminus R'_i) \cup (R'_i \setminus R_i)$ for some $i \in C$. Each such pair (x, y) thus belongs to at least one of the following two classes:

CLASS 1. $(x, y) \in R_i \setminus R'_i$ for some $i \in C$

CLASS 2. $(x, y) \in R'_i \setminus R_i$ for some $i \in C$

We go on to show that Class 1 contains at least one pair. Assume for contradiction that all pairs belong to Class 2 and let $(x, y) \in R'_i \setminus R_i$ be one of these pairs. As $x \notin Y$, SET-MON implies that $f(R'_{i:(y,x)}) = f(R') = Y$. As $f(R'_{i:(y,x)}) P_i^G f(R)$ for all $i \in C$ and $\delta(R, R'_{i:(y,x)}) = \delta(R, R') - 1$, R and $R'_{i:(y,x)}$ constitute a smaller counterexample, contradicting the minimality of $\delta(R, R')$.

Thus, there is at least one pair (x, y) that belongs to Class 1, i.e., a pair (x, y) with $x \in X \setminus Y$ and $y \in Y \setminus X$ such that $x R_i y$ for some voter $i \in C$. But this contradicts the assumption that $Y P_i^G X$ for all $i \in C$, and completes the proof. \square

A.5 Proof of Theorem 6

THEOREM 6. *Every pairwise SCF that is P^G -group-strategy-proof satisfies SET-MON, weak EICA, and SDP.*

Proof. By Theorem 4 and the fact that P^G -group-strategy-proofness implies P^F -group-strategyproofness, it remains to be shown that every pairwise SCF that violates SDP is P^G -manipulable. Suppose that f is pairwise and violates SDP. Then, there exist two preference profiles $R = (R_1, \dots, R_n)$ and $R' = (R'_1, \dots, R'_n)$ such that $X = f(R)$ and $Y = f(R')$ are not contained in one another, and $R_i|_{\{x,y\}} = R'_i|_{\{x,y\}}$ for all $i \in N$ and all alternatives x, y with $x \in X \setminus Y$ and $y \in Y \setminus X$.

The proof now works analogously as the proof of Theorem 4. For each voter i with $R_i \neq R'_i$ we have two new voters R_{n+i} and R_{n+c+i} such that $Y P_{n+i}^G X$, $R_i \setminus R'_i \subseteq R_{n+i}$, and $R_{n+c+i} = R_{n+i}^-$. By letting $R'_{n+i} = R_{n+i} \setminus R_i \cup R'_i$ and

defining S and S' as in the proof of Theorem 4, we can show that f is P^G -manipulable. \square

A.6 Proof of Corollary 1

COROLLARY 1. *COND and TC are P^G -group-strategy-proof.*

Proof. By Theorem 5, it is sufficient to show that *COND* and *TC* satisfy *SET-MON*, *EICA*, and *SDP*.

If a preference profile R does not have a Condorcet winner, *COND* trivially satisfies the three properties because all alternatives are chosen. If R has a Condorcet winner, *EICA* is again trivial and *SET-MON* and *SDP* are straightforward.

The (easy) fact that *TC* satisfies *SET-MON* was shown by Brandt [6]. To see that *TC* satisfies *EICA*, consider $b \notin TC(R)$. By definition of *TC*, $b R_M^* a$ for no $a \in TC(R)$. As R and R' differ only on alternatives in $TC(R)$, it follows that $b R_M^* a$ for no $a \in TC(R)$, and thus $a \notin TC(R')$.

Finally, to see that *TC* satisfies *SDP*, observe that $a P_M b$ for all $a \in TC(R)$ and $b \notin TC(R)$. Thus, if $x \in TC(R) \setminus TC(R')$ and $y \in TC(R') \setminus TC(R)$, we have $x P_M y$ and $y P_M' x$. This implies that at least one voter has modified his preference between x and y . \square

A.7 Proof of Corollary 2

COROLLARY 2. *COND \cap PAR and PAR are P^F -group-strategyproof.*

Proof. By Theorem 3, it is sufficient to show that *COND \cap PAR* and *PAR* satisfy *SET-MON* and *EICA*.

SET-MON holds because a Pareto-dominated alternative remains Pareto-dominated when it is weakened. For *EICA*, observe that transitivity of Pareto-dominance implies that each Pareto-dominated alternative is dominated by a Pareto-undominated one. Therefore, if $a \notin PAR(R)$ and R and R' differ only on alternatives in $PAR(R)$, then $a \notin PAR(R')$ because a is still Pareto-dominated by the same alternative. \square

A.8 Proof of Proposition 1

PROPOSITION 1. *OMNI is P^G -group-strategyproof.*

Proof. Assume for contradiction that *OMNI* is not P^G -group-strategyproof. Then, there have to be a group of voters $C \subseteq N$ and two preference profiles R and R' with $R_j = R'_j$ for all $j \notin C$ such that $OMNI(R') P_i^G OMNI(R)$ for all $i \in C$. Denote $X = OMNI(R)$ and $Y = OMNI(R')$.

As P^G coincides with P^F on all pairs where one set is contained in the other set, and, by Theorem 3, *OMNI* is P^F -group-strategyproof, we can conclude that neither $X \subseteq Y$ nor $Y \subseteq X$. Choose $x \in X \setminus Y$ and $y \in Y \setminus X$ arbitrarily. On the one hand, $x \in X \setminus Y$ implies the existence of a voter $i \in C$ with $x R_i a$ for all $a \neq x$. On the other hand, $Y P_i^G X$ implies $y R_i x$, a contradiction. \square