

Preserving Condorcet Winners under Strategic Manipulation

Sirin Botan and Ulle Endriss

Institute for Logic, Language and Computation, University of Amsterdam
{sirin.botan, ulle.endriss}@uva.nl

Abstract

Condorcet extensions have long held a prominent place in social choice theory. A Condorcet extension will return the Condorcet winner as the unique winner whenever such an alternative exists. However, the definition of a Condorcet extension does not take into account possible manipulation by the voters. A profile where all agents vote truthfully may have a Condorcet winner, but this alternative may not end up in the set of winners if agents are acting strategically. Focusing on the class of tournament solutions, we show that many natural social choice functions in this class, such as the well-known Copeland and Slater rules, cannot guarantee the preservation of Condorcet winners when agents behave strategically. Our main result in this respect is an impossibility theorem that establishes that no tournament solution satisfying a very weak decisiveness requirement can provide such a guarantee. On the bright side, we identify several indecisive but otherwise attractive tournament solutions that do guarantee the preservation of Condorcet winners under strategic manipulation for a large class of preference extensions.

1 Introduction

The notion of strategyproofness has been studied across disciplines, including economics, game theory, and artificial intelligence (Arrow, Sen, and Suzumura 2010; Shoham and Leyton-Brown 2009; Brandt et al. 2016; Meir 2018). For multiagent systems, including those utilising voting to reach a collective decision, ensuring some level of strategyproofness is an important consideration (Amanatidis, Birmpas, and Markakis 2016; Hajaj et al. 2015).

By a seminal result in social choice theory, we know that every non-trivial resolute social choice function—or voting rule—is susceptible to strategic manipulation by voters (Gibbard 1973; Satterthwaite 1975). Our proposal in this paper is to define a more fine-grained strategyproofness axiom that dictates a specific kind of “undesirable” manipulation should not occur. We do not examine whether strategyproofness can be satisfied alongside a set of other desirable axioms, but whether the fact that strategyproofness fails can affect the “strength” of the other axiom(s). Our focus is on Condorcet extensions, and how failure of strategyproofness affects Condorcet consistency.

If an agent manipulates under a Condorcet extension—meaning she submits a preference other than her true one—it is possible that this will result in a reported profile without a Condorcet winner, despite the fact that the truthful profile has one. In particular, it is possible that a Condorcet extension in some cases fails to return the Condorcet winner of the truthful profile as the outcome. We examine exactly when the lack of strategyproofness affects whether we can trust that a Condorcet extension will give us all “true” Condorcet winners, even under the assumption that agents might vote strategically.

Strategyproofness on profiles with a Condorcet winner is also appealing from a practical viewpoint as real-world elections have a high probability of giving us a Condorcet winner (Gehrlein 2006; Gehrlein and Lepelley 2011).

Related work. There are various methods in the literature for circumventing the impossibility result due to Gibbard (1973) and Satterthwaite (1975). A well-employed strategy is to consider a restricted domain for the social choice function, where strategyproofness is guaranteed. Among these domain restrictions, the best-known is the single-peaked domain of Black (1948).¹ There are also many examples of positive results relative to more fine-grained axioms that focus on—and limit—the type of manipulation performed by the agent (Sato 2013; Bossert and Sprumont 2014; Kruger and Terzopoulou 2020). Similar positive results have been obtained by considering voters’ ignorance of others’ preference as an informational barrier (Conitzer, Walsh, and Xia 2011; Reijngoud and Endriss 2012; Osborne and Rubinstein 2003). Another approach has been to argue for the computational hardness of computing a successful manipulation strategy as a barrier to manipulation (Bartholdi, Tovey, and Trick 1989; Conitzer and Walsh 2016).

While the Gibbard-Satterthwaite Theorem pertains to resolute social choice functions, there are similar impossibility results for irresolute rules (Gärdenfors 1976; Kelly 1977; Barberà 1977; Duggan and Schwartz 2000; Sato 2008). These results differ in how they define manipulability, as they by necessity must make assumptions about agents’ preferences over sets of alternatives. Duggan and Schwartz (2000), for example, work with optimistic and pessimistic

¹For a more thorough treatment of single-peakedness and many other domain restrictions, see Gaertner (2001).

agents, while Gärdenfors (1976) defines manipulability relative to his eponymous preference extension.

Shifting focus away from resoluteness has also led to positive results regarding the strategyproofness of social choice functions. For example, Gärdenfors (1976) identifies two such strategyproof functions for the Gärdenfors extension and Brandt (2015) identifies social choice functions that are strategyproof under the Kelly preference extension.

The idea of preserving Condorcet winners has also been examined in the setting of probabilistic social choice. Hoang (2017) shows that *maximal lotteries* (Fishburn 1984) are strategyproof on profiles with Condorcet winners when based on the majority relation. Brandl, Brandt, and Stricker (2018) extend this result to all maximal lotteries.

Contribution. To distinguish between social choice functions that do not incentivise manipulation in profiles with a Condorcet winner and those that do, we introduce the notion of a *robust Condorcet extension*. We then show that no irresolute tournament solution that is *weakly resolute*—in the sense of returning a single winning alternative in at least one profile that does not have a Condorcet winner—can possibly be such a robust Condorcet extension. This class of weakly resolute rules includes the well-known social choice functions of Copeland (1951) and Slater (1961). Finally, we identify several attractive social choice functions that are robust Condorcet extensions (and thus fail weak resoluteness) for a large class of preferences. This includes, in particular, the minimal extending set (Brandt 2011) and its coarsenings.

Paper outline. The remainder of the paper is organised as follows. We introduce the framework and relevant literature in Section 2. In Section 3 we present an impossibility result for weakly resolute rules. In Section 4 we present a number of sufficient conditions for a Condorcet extension to be robust. We conclude in Section 5.

2 The Model

In this section we introduce the framework and notation we will be using throughout the paper. Much of the material up to Section 2.5 is familiar from social choice theory (Arrow, Sen, and Suzumura 2010; Brandt et al. 2016). In Section 2.5, we introduce the novel notion of a robust Condorcet extension, the central concept of this paper.

2.1 Preference Profiles

Let A be a finite set of *alternatives*, and $N = \{1, \dots, n\}$ a finite set of *agents*. A *preference profile* $\mathbf{P} = (\succ_1^{\mathbf{P}}, \dots, \succ_n^{\mathbf{P}})$ is a vector of strict linear orders over A , where $\succ_i^{\mathbf{P}}$ is the *preference relation* of agent i in the profile \mathbf{P} . For two profiles \mathbf{P} and \mathbf{P}' , and agent $i \in N$, we write $\mathbf{P} =_{-i} \mathbf{P}'$, and say they are i -variants, if $\succ_j^{\mathbf{P}} = \succ_j^{\mathbf{P}'}$ for all $j \in N \setminus \{i\}$. $\mathcal{L}(A)$ denotes the set of all linear orders over A , and $\mathcal{L}(A)^n$ denotes the set of all profiles for n agents.

For a profile \mathbf{P} , $\succ^{\mathbf{P}}$ (with asymmetric part $\succ^{\mathbf{P}}$) is the *majority relation* for \mathbf{P} and is defined such that $a \succ^{\mathbf{P}} a'$ if and only if $|\{i \in N \mid a \succ_i^{\mathbf{P}} a'\}| \geq |\{i \in N \mid a' \succ_i^{\mathbf{P}} a\}|$, for all $a, a' \in A$. An alternative $a \in A$ is a *Condorcet winner* in profile \mathbf{P} if it defeats every other alternative in a pairwise

majority contest, meaning $a \succ^{\mathbf{P}} a'$ for all $a' \in A \setminus \{a\}$. We define $\mathcal{D}_{\text{Condorcet}}$ as the set of all profiles with a Condorcet winner.

We say a relation \succ over A is *complete* if for all $a, b \in A$ it is the case that $a \succeq b$ or $b \succeq a$. A relation \succ is *connex* if $a \succ b$ or $b \succ a$ for all distinct $a, b \in A$. Note that the majority relation of any profile is complete, and any individual preference relation is connex.

2.2 Social Choice Functions

An irresolute *social choice function* (SCF) is a mapping from profiles to nonempty subsets of alternatives:

$$f : \mathcal{L}(A)^n \rightarrow 2^A \setminus \{\emptyset\}$$

To avoid having to break majority ties, we define SCFs for odd n . Of course we do not strictly speaking need odd n , but rather that profiles input to the function have a strict majority relation. A SCF f is a *Condorcet extension*, or is *Condorcet-consistent*, if it returns (only) the Condorcet winner whenever one exists.

For irresolute SCF, the size of the set of winning alternatives is an important consideration. As an example of a rule that returns quite large sets, take the rule that returns the Condorcet winner if one exists, and returns the whole set of alternatives otherwise. While this is clearly a Condorcet extension, it is a very *indecisive* rule, as it often results in many ties in the outcome.

A SCF is *resolute* if it always returns a singleton. In order to quantify how decisive an irresolute rule is, we define a weaker notion of resoluteness. We say f is *weakly resolute* if there exists a profile $\mathbf{P} \in \mathcal{L}(A)^n \setminus \mathcal{D}_{\text{Condorcet}}$ for which $|f(\mathbf{P})| = 1$. So, a rule is weakly resolute if it *sometimes* returns a singleton for a profile without a Condorcet winner. For some SCFs, we can directly compare how decisive they are relative to each other. A SCF f is a *refinement* of f' if for all profiles $\mathbf{P} \in \mathcal{L}(A)^n$ it is the case that $f(\mathbf{P}) \subseteq f'(\mathbf{P})$, meaning f always returns a subset of f' . If f is a refinement of f' , we say f' is a *coarsening* of f . If a rule is a refinement of another, it is clearly the more decisive of the two.

2.3 Tournaments and Tournament Solutions

A *tournament* \mathbf{T} is a pair $(S, \succ^{\mathbf{T}})$, where S is a set of nodes and $\succ^{\mathbf{T}}$ is an asymmetric and connex relation over S , which we call the *dominance relation* of the tournament. For a set S , we denote by $\mathcal{T}(S)$ all tournaments on S .

For a tournament $\mathbf{T} = (S, \succ^{\mathbf{T}})$, we say an alternative $a \in S$ *dominates* $a' \in S$ in the tournament \mathbf{T} if $a \succ^{\mathbf{T}} a'$. The *dominion* of a in \mathbf{T} is defined as $D_{\mathbf{T}}(a) = \{a' \in S \mid a \succ^{\mathbf{T}} a'\}$, the set of alternatives it dominates. The *dominators* of a in \mathbf{T} is defined as $\overline{D}_{\mathbf{T}}(a) = \{a' \in S \mid a' \succ^{\mathbf{T}} a\}$, the set of alternatives that dominate it. For $S' \subseteq S$, we define the restriction $\succ_{S'}^{\mathbf{T}} = \{(a, a') \in S' \times S' \mid a \succ^{\mathbf{T}} a'\}$, which is $\succ^{\mathbf{T}}$ restricted to the set S' . A *subtournament* of $\mathbf{T} = (S, \succ^{\mathbf{T}})$ is a tournament $(S', \succ_{S'}^{\mathbf{T}})$ where $S' \subseteq S$. Thus, a subtournament of \mathbf{T} is a subset of the nodes in \mathbf{T} , and the edges between those nodes.

We say $\pi : S \rightarrow S'$ is an isomorphism between two tournaments $\mathbf{T} = (S, \succ^{\mathbf{T}})$ and $\mathbf{T}' = (S', \succ^{\mathbf{T}'})$ if π is a bijection, and $a \succ^{\mathbf{T}} a' \Leftrightarrow \pi(a) \succ^{\mathbf{T}'} \pi(a')$ for all $(a, a') \in S \times S$.

We say a profile $\mathbf{P} \in \mathcal{L}(A)^n$ induces tournament $\mathbf{T} = (A, \succ^{\mathbf{T}})$ if $\succ^{\mathbf{P}} = \succ^{\mathbf{T}}$. So a profile induces a tournament if they range over the same alternatives, and the strict part of the majority relation is exactly the dominance relation of the tournament. Note that if a profile induces a tournament, the strict component of the majority relation of that profile must be connex. As tournaments do not speak about agents, we cannot directly talk about two tournaments being i -variants for some agent $i \in N$. Instead, we say two tournaments $\mathbf{T} = (A, \succ^{\mathbf{T}})$ and $\mathbf{T}' = (A, \succ^{\mathbf{T}'})$ are *single-agent variants*, and write $\mathbf{T} =_{-1} \mathbf{T}'$, if there exist a set of agents N and profiles $\mathbf{P}, \mathbf{P}' \in \mathcal{L}(A)^n$, for $n = |N|$, such that $\mathbf{P} =_{-i} \mathbf{P}'$ for some agent $i \in N$, and the profiles \mathbf{P} and \mathbf{P}' induce the tournaments \mathbf{T} and \mathbf{T}' , respectively.

We say an element $a \in S$ is the *Condorcet winner* of the tournament $\mathbf{T} = (S, \succ^{\mathbf{T}})$ if $\overline{D}_{\mathbf{T}}(a) = \emptyset$. This corresponds to the notion of a Condorcet winner of a profile; if a tournament has a Condorcet winner, that alternative will be the Condorcet winner in any profile that induces this tournament. We write $\mathcal{T}_{\text{Condorcet}}$ to refer to the set of tournaments that have a Condorcet winner. A *tournament solution* is a mapping from tournaments to sets of alternatives, that does not distinguish between isomorphic tournaments:

$$F : \mathcal{T}(S) \rightarrow 2^S \setminus \{\emptyset\}$$

So $F(\mathbf{T}') = \{\pi(a) \mid a \in F(\mathbf{T})\}$ if π is an isomorphism between \mathbf{T} and \mathbf{T}' . For ease of reading, we sometimes write $F(\succ^{\mathbf{T}})$ to mean $F(S, \succ^{\mathbf{T}})$ when S is clear from context.

A SCF f is *equivalent* to a tournament solution F if $f(\mathbf{P}) = F(A, \succ^{\mathbf{P}})$ for all $\mathbf{P} \in \mathcal{L}(A)^n$. Note that the majority relation of this profile \mathbf{P} must be a strict order, as the SCF f is defined for odd n only. In a slight muddling of terminology, we will refer to SCFs that are equivalent to tournament solutions as *tournament-solution SCFs*.

Tournament solutions roughly correspond to Fishburn's C1 functions (Fishburn 1977), which require only the information in the majority graph to determine the winners. More precisely, tournament-solution SCFs correspond to *neutral* C1 functions,² as tournament solutions do not distinguish between isomorphic tournaments, and therefore do not favour any alternatives over others.

2.4 Extending Preferences

Because the rules we examine are irresolute—meaning they do not always return a single winner—we need to specify how agent preferences over alternatives are extended to preferences over *sets of alternatives*.

A preference extension e maps any given preference relation \succ over alternatives in A , to a relation \succeq^e (with strict part \succ^e) over sets of alternatives. We define two requirements for any preference extension e :

- (i) $x \succ y$ implies $\{x\} \succ^e \{y\}$ and,
- (ii) $X \succ^e Y$ implies that there exist some $x \in X$ and $y \in Y$ such that $x \succ y$ and $\{x, y\} \not\subseteq X \cap Y$.

²A SCF satisfies *neutrality* if for any profile \mathbf{P} and any permutation $\pi : A \rightarrow A$ it is the case that $f(\pi(\mathbf{P})) = \pi(f(\mathbf{P}))$.

The first requirement simply dictates that e stays faithful to the agent's preferences when comparing singleton sets. The second requires that e does not extend the preferences in a way that completely disagrees with an agent's preferences over alternatives. Our first requirement corresponds to the Extension Rule of Barberà, Bossert, and Pattanaik (2004), while our second requirement is similar in spirit to the l -extension of Kruger and Terzopoulou (2020). For an agent i with preference relation $\succ_i^{\mathbf{P}}$ in profile \mathbf{P} , we write $\succeq_i^{\mathbf{P}, e}$ to denote her preferences over sets of alternatives, extended according to e . We say an agent has e -preferences, if $\succeq_i^{\mathbf{P}, e}$ is her preference relation over sets of alternatives.

We say a preference extension e is *pessimistic* if $X \succ^e Y$ implies that there exists some $y \in Y$ such that $x \succeq y$ for all $x \in X$. Note that this defines a class of preference extensions. We highlight the *Gärdenfors extension* (Gärdenfors 1976) as a well-known example of a pessimistic preference extension. The Gärdenfors extension dictates that if one set is preferred over another, then any alternative added to the first set to reach the second is preferred to the alternatives in the initial set. Similarly, those alternatives removed from the initial set should be less preferred.

2.5 Robust Condorcet Extensions

We say an irresolute SCF f is *Condorcet-manipulable* by agent i in profile \mathbf{P} , under preference extension e , if there exists another profile $\mathbf{P}' =_{-i} \mathbf{P}$ such that $f(\mathbf{P}') \succ_i^{\mathbf{P}, e} f(\mathbf{P})$ and $\mathbf{P} \in \mathcal{D}_{\text{Condorcet}}$. We are now ready to present our central definition.

Definition 1. A SCF f is a *robust Condorcet extension* (or simply, is *robust*) under a preference extension e if f is *Condorcet-consistent* and not *Condorcet-manipulable* in any profile $\mathbf{P} \in \mathcal{D}_{\text{Condorcet}}$, by any $i \in N$ with e -preferences.

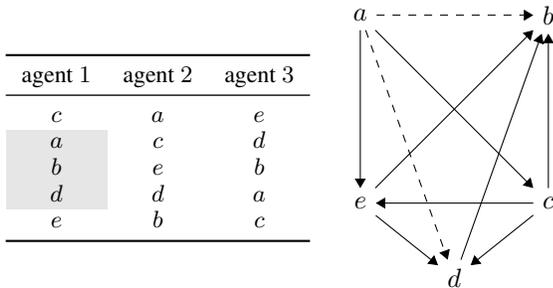
We sometimes write that a SCF is *robust* to mean that it is a robust Condorcet extension, as robustness is a property of Condorcet extensions. So a SCF is robust under a certain preference extension, if it is not Condorcet manipulable by any agent whose preferences over alternatives have been extended to sets of alternatives according to that extension.

While robustness is a weak strategyproofness requirement, it also speaks about how well a rule can preserve Condorcet winners. A robust Condorcet extension ensures that, if the truthful profile has a Condorcet winner, then it is a weakly dominant strategy for all agents to report their true preferences, thus ensuring that no Condorcet winner loses that designation because of strategic manipulation. A robust Condorcet extension therefore ensures that profiles with Condorcet winners are, in a sense, stable. We give an example of a Condorcet manipulation of the *Copeland* SCF,³ to demonstrate what failure of robustness looks like.

Example 1. Suppose the profile below, along with the corresponding majority graph, is the “truthful” profile, meaning

³The Copeland score of an alternative in a profile (for odd n) is the number of other alternatives it beats in a pairwise majority contest. The Copeland rule selects those alternatives with the highest Copeland scores (Copeland 1951).

all three agents have reported their true preferences. As alternative a is the Condorcet winner, Copeland will return a as the sole winner if all agents vote truthfully. Note, however, that agent 1 has the ability to reverse the dashed edges (a, b) and (a, d) in the majority graph, by moving b and d above a in her own preference order, while keeping their relative ranking as it is. Agent 1 also has an incentive to do so, as this would result in a majority graph where c —her top choice—is the only alternative with a single incoming edge.



As the Copeland winner is the alternative with the smallest number of incoming edges in the majority graph, c would be the lone Copeland winner if agent 1 misreports her preferences, meaning, Copeland incentivises a Condorcet-manipulation in this profile.

While the profile in Example 1 has a Condorcet winner, Copeland is not guaranteed to return this alternative as the winner (or even among them) unless we assume all agents vote truthfully. In the same profile, a robust Condorcet-extension would ensure no agent would have an incentive to misreport her preferences.

3 Impossibilities

We present a first impossibility result, showing there is no function that can meet our robustness requirement for all preference extensions.

Proposition 1. *No tournament-solution SCF is robust under all preference extensions.*

Proof. Let $A = \{a, b, c\}$, $N = \{1, 2, 3\}$,⁴ and let f be a Condorcet-consistent SCF equivalent to the tournament solution F . Suppose agent 1's preferences over sets of alternatives $\succeq_1^{P,e}$ are such that $X \succ_1^{P,e} Y$ if and only if one of the following holds:

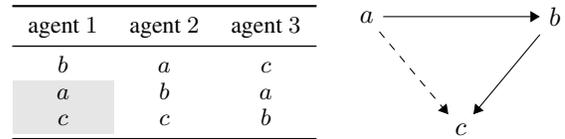
- $X = \{a, b, c\}$ and $Y = \{a\}$, or
- $X = \{x\}$ and $Y = \{y\}$ for some $x, y \in A$ s.t. $x \succ y$.

These preferences satisfy both our requirements for preference extensions, and are therefore a valid extension of \succ_1^P .

Let P be the profile shown below, with the induced tournament T on the right. As f is a Condorcet extension, $f(P) = \{a\}$. Let $P' =_{-1} P$, where $b \succ_1^{P'} c \succ_1^{P'} a$, meaning agent 1 reverses the edge (a, c) in the induced tournament by reversing the order of these alternatives in her ranking. The tournament T' , induced by P' , consists of

⁴While our proof uses three agents and three alternatives, we can always construct profiles with a larger number m of alternatives, where the majority relation is an m -cycle.

a 3-cycle.



As tournament solutions do not distinguish between isomorphic tournaments, $f(P') = F(T') = \{a, b, c\}$. As $\{a, b, c\} \succ_1^{P,e} \{a\}$, this would constitute a successful manipulation for agent 1, meaning f cannot be robust. \square

As no SCFs are robust for *all* preference extensions, we redirect our search to those that may be robust for *some* preference extension. We first recall a result by McGarvey (1953), which we will use to prove the main result of this section. We include the proof for the sake of completeness.

Theorem 1 (McGarvey, 1953). *Let A be a set of alternatives, and let \succeq be a complete relation over A . Then there is a profile $P \in \mathcal{L}(A)^n$ for some even n such that $\succeq^P = \succeq$, and if $a > b$, there are $\frac{n}{2} + 2$ agents ranking a over b in P .*

Proof. For a set of alternatives A and a relation \succeq (with strict component $>$) over A , the profile P is constructed for an even number of agents $N = \{i_{ab}, j_{ab} \mid (a, b) \in >\}$ as follows. For every pair of alternatives such that $a > b$, there are two voters i_{ab} and j_{ab} , with the following preferences:

$$a \succ_{i_{ab}}^P b \succ_{i_{ab}}^P x_1 \succ_{i_{ab}}^P \dots \succ_{i_{ab}}^P x_{|A|-2} \text{ and}$$

$$x_{|A|-2} \succ_{j_{ab}}^P \dots \succ_{j_{ab}}^P x_1 \succ_{j_{ab}}^P a \succ_{j_{ab}}^P b,$$

Here $\{x_1, \dots, x_{|A|-2}\} = A \setminus \{a, b\}$. For each agent in $N \setminus \{i_{ab}, j_{ab}\}$ who prefers a over b , there will be exactly one corresponding agent who prefers b over a , meaning in the profile P exactly $\frac{n}{2} + 2$ agents prefer a to b . As this holds for any pair of alternatives, it is clear that $\succeq^P = \succeq$. \square

Note that while our statement of McGarvey's Theorem is slightly stronger than the classical formulation, the proof and the profile constructed in the proof remain the same. We now show that weakly resolute rules fail robustness for all preference extensions, and further, that they are the only rules that do so. This strengthens an observation from Brandt, Brill, and Harrenstein (2016) stating that all weakly resolute rules are Kelly-manipulable.

Theorem 2. *A tournament-solution SCF is weakly resolute if and only if it fails robustness under all preference extensions.*

Proof. For the right-to-left direction we prove the contrapositive. That is, we suppose f is a tournament-solution SCF that fails weak resoluteness and show it must be robust under some preference extension. To see that this must be the case, note that any rule failing weak resoluteness never returns singletons outside $\mathcal{D}_{\text{Condorcet}}$. This means the preference extension ranging only over singletons would never (strictly)

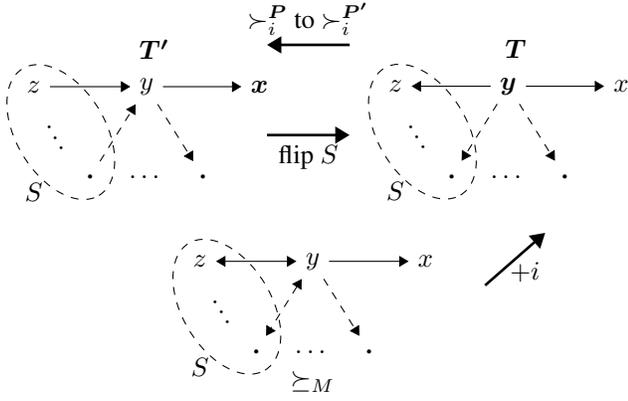


Figure 1: Tournaments T and T' —with winners marked in bold—and relation \succeq_M from the proof of Theorem 2. Ties are represented by bidirectional arrows.

favour a larger set over the singleton set with the Condorcet winner.⁵ As f will always return a set larger than a singleton outside the Condorcet domain, Condorcet-manipulation with these preferences is not possible, thereby making f robust under this preference extension.

For the left-to-right direction, let A be our set of alternatives. Suppose f is a weakly resolute tournament-solution SCF, equivalent to a tournament solution F . We show it is possible for an agent to manipulate f from a profile with a Condorcet winner under an arbitrary preference extension e , meaning f cannot be robust under any preference extension.

We first define two tournaments T and T' , which we will show are single-agent variants. As F is equivalent to a weakly resolute SCF, there is some tournament $T' = (A, \succ^{T'}) \in \mathcal{T}(A) \setminus \mathcal{T}_{\text{Condorcet}}$ such that $F(T') = \{x\}$ for an alternative $x \in A$. As x is not a Condorcet winner in T' , there must be some $y \in A$ such that $y \succ^{T'} x$, and by the same reasoning, there must be at least one alternative $z \in A$ such that $z \succ^{T'} y$. We conclude that the nodes $\{x, y, z\}$ and the edges $(y, x), (z, y)$ must be present in T' . For a visual representation, see Figure 1.

Let $S = \bar{D}_{T'}(y)$ be the dominators of y in T' . We define a second tournament $T = (A, \succ^T)$, where $y \succ^T a$ for all $a \in S$, and \succ^T agrees with $\succ^{T'}$ on all other pairs of alternatives. In other words, we simply reverse all incoming edges of y in T' to obtain T . Note that this makes y a Condorcet winner in T , meaning $F(T) = \{y\}$.

We now show that T and T' are single-agent variants. We start by constructing a profile P that induces T . To this end, consider a complete relation \succeq_M (with strict component \succ_M , and symmetric component \sim_M) over A , such that $\succeq_M = \succ^T \cup \{(a, y) \mid a \in S\}$. This means \succeq_M and \succ^T agree on all pairs of alternatives except those for which T and T' differ. In those cases, \succeq_M gives a tie between the

⁵Note that all relevant manipulations here would be between singleton outcomes, meaning we are taking advantage of strategyproofness of the Condorcet, or majority, rule (Campbell and Kelly 2003).

alternatives. By Theorem 1, we know there exists a profile $P^* = (\succ_1^{P^*}, \dots, \succ_n^{P^*})$ with majority relation \succeq_M . Further, we know that we can construct P^* with an even number of agents n , such that for any $a, a' \in A$, where $a \succ_M a'$, there are exactly $\frac{n}{2} + 2$ agents who prefer a to a' in P^* . We use P^* to construct the profile P . Let $P = (\succ_1^{P^*}, \dots, \succ_n^{P^*}, \succ_i^P)$, where $x \succ_i^P y \succ_i^P a$ for all $a \in A \setminus \{x, y\}$.

To see that P induces tournament T , note that for any pair of alternatives (a, a') , either

- (i) $a \succ_M a'$ —meaning $a \succ^T a'$, and $\frac{n}{2} + 2$ prefer a to a' in P^* , or
- (ii) $a \sim_M a'$ —meaning $a' = y$, and $a \in S$ (or vice versa).

If (i) is the case, a majority of agents in P will prefer a to a' regardless of agent i 's preferences; $\frac{n}{2} + 2$ agents still form a strict majority of $n + 1$ agents. If on the other hand (ii) is the case, we know from agent i 's preferences that $y \succ_i^P a$. As these alternatives were tied in P^* , adding agent i to the profile breaks these ties in favour of y , so a majority of agents in P will now prefer y to a . Thus the only differences between \succeq_M and \succeq^P relate to the pairs on which \succeq_M and \succeq^T differ. As the changes agree with \succ^T , this makes $\succ^T = \succ^P$, meaning P induces T . As $F(T) = \{y\}$, we can conclude that $f(P) = \{y\}$.

It now remains to construct a profile P' such that $P \dashv_i P'$ and P' induces T' . Let $P' = (\succ_1^{P'}, \dots, \succ_n^{P'}, \succ_i^{P'})$, and $x \succ_i^{P'} a \succ_i^{P'} y$, for all $a \in A \setminus \{x, y\}$, meaning agent i moves y to the bottom of their ranking. Clearly, P' is an i -variant of P . In the tournament induced by P' , it must be the case that the edges (a, y) for all $a \in S$ are present as the majority on these alternatives is dictated by agent i (and all other edges remain as they were in T). As these edges correspond exactly to those where T and T' differ, P' must induce T' , and as $F(T') = \{x\}$ we conclude $f(P') = \{x\}$.

Finally, let $\succeq_i^{P, e}$ be agent i 's true preferences over sets of alternatives, extended according to e . It is immediately clear, as both outcomes are singletons and $x \succ_i^P y$, that $f(P') \succ_i^{P, e} f(P)$. As P has a Condorcet winner, this constitutes a Condorcet-manipulation, meaning f cannot be robust under preference extension e . \square

We note that Theorem 2 applies to two of the most prominent Condorcet extensions—Copeland, and Slater.⁶

4 Robust Tournament Solutions

In this section, we present our robustness results for several tournament-solution SCF, and their coarsenings. Our results hold for all pessimistic extensions.

4.1 Relation to Kelly-Strategyproofness

While strategyproofness for say, Gärdenfors preferences, is not easily satisfied, there are several tournament-solution SCFs that have been shown to be strategyproof for the Kelly preference extension (Kelly 1977)—which we will refer to as k . For any two sets X and Y in $2^A \setminus \{\emptyset\}$, $X \succ^k Y$ if and

⁶The result also extends to Slater's weighted counterpart, the Kemeny rule (Kemeny 1959).

only if $x \succeq y$ for all $x \in X$ and all $y \in Y$, and there exists an $x \in X$ and a $y \in Y$ such that $x \succ y$. We say a SCF f is *Kelly-strategyproof* if no agent with Kelly preferences can successfully manipulate.

A SCF satisfying Gärdenfors-strategyproofness implies it also satisfies Kelly-strategyproofness, as the former must exclude more cases of manipulation to be satisfied. However, as robustness only requires taking into account comparisons where at least one singleton set is present, we can use strategyproofness results for Kelly preferences to show robustness for all pessimistic extensions, including Gärdenfors.

Proposition 2. *If a Condorcet-consistent SCF f is Kelly-strategyproof, then it is a robust Condorcet extension under any pessimistic preference extension.*

Proof. Suppose we have a rule f that is Kelly-strategyproof and let e be a pessimistic extension. That is, for any two profiles \mathbf{P} and \mathbf{P}' , and any agent $i \in N$, if $\mathbf{P}' =_{-i} \mathbf{P}$, then $f(\mathbf{P}') \not\prec_i^{P,k} f(\mathbf{P})$. Suppose \mathbf{P} has a Condorcet winner, meaning $f(\mathbf{P}) = \{a\}$ for some $a \in A$. Because $f(\mathbf{P}') \not\prec_i^{P,k} f(\mathbf{P})$, it cannot be the case that all elements of $f(\mathbf{P}')$ are preferred to a . So either $f(\mathbf{P}') = f(\mathbf{P})$ or there is some $a' \in f(\mathbf{P}')$ s.t. $a \succ_i^P a'$. It is then immediate from the definition of a pessimistic extension that $f(\mathbf{P}') \not\prec_i^{P,e} f(\mathbf{P})$. \square

We therefore get robustness under pessimistic preferences “for free” for Condorcet extensions known to be Kelly-strategyproof, such as the bipartisan set and the minimal covering set (Brandt 2015).⁷

4.2 Minimal Extending Set & Beyond

In this section we show that robustness for Condorcet extensions diverges from Kelly-strategyproofness, as we can find rules which satisfy the former while failing the latter. Before we present our positive robustness results, we need to define the Banks set and the minimal extending set.

A tournament $\mathbf{T}' = (S', \succ^{\mathbf{T}'})$ is a *maximal transitive subtournament* of $\mathbf{T} = (S, \succ^{\mathbf{T}})$ if

- (i) \mathbf{T}' is a subtournament of \mathbf{T} ,
- (ii) $\succ^{\mathbf{T}'}$ is a transitive relation, and
- (iii) there is no other transitive subtournament $(S'', \succ^{\mathbf{T}''})$ of \mathbf{T} such that $S' \subset S''$.

We write $\hat{\mathbf{T}}$ to denote the set of all maximal transitive subtournaments of tournament \mathbf{T} , and $\text{top}(\succ)$ to denote the maximal element of the strict linear order \succ . Note that if a tournament \mathbf{T} has a Condorcet winner, it will be the maximal element of *all* maximal transitive subtournaments.⁸ The *Banks set* (Banks 1985) is the set of maximal elements of all maximal transitive subtournaments of a tournament:

$$\text{BA}(\mathbf{T}) = \{\text{top}(\succ_S^{\mathbf{T}}) \mid (S, \succ_S^{\mathbf{T}}) \in \hat{\mathbf{T}}\}.$$

⁷For a thorough treatment of Kelly-strategyproofness and functions which satisfy it, as well as definitions of the SCFs we have mentioned, we refer to Brandt, Brill, and Harrenstein (2016).

⁸As the existence of a Condorcet winner does not imply no cycles are present, there may indeed be several maximal transitive subtournaments.

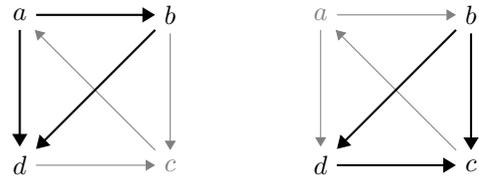
Because the Condorcet winner will top all maximal transitive subtournaments, the Banks set is a Condorcet extension.

A set $S \subseteq A$ is a *BA-stable set* of a tournament \mathbf{T} if $a \notin \text{BA}(S \cup \{a\}, \succ_{S \cup \{a\}}^{\mathbf{T}})$ for all $a \in A \setminus S$. A BA-stable set of a tournament \mathbf{T} is *minimal* if there is no BA-stable set S' of \mathbf{T} such that $S' \subset S$. The *minimal extending set* $\text{ME}(\mathbf{T})$ (Brandt 2011) of a tournament \mathbf{T} is the union of all minimal BA-stable sets of \mathbf{T} :

$$\text{ME}(\mathbf{T}) = \bigcup \{S \subseteq A \mid S \text{ is a minimal BA-stable set of } \mathbf{T}\}.$$

We give an example to shed some light on these definitions.

Example 2. *In the tournament \mathbf{T} below, the two maximal transitive subtournaments are indicated using darker edges. It is clear that the subtournaments are transitive, and they are both maximal; adding the last alternative will break transitivity. From examining these subtournaments, we can see that $\text{BA}(\mathbf{T}) = \{a, b\}$.*



The tournament has two minimal BA-stable sets: $\{a, b, d\}$ —because $c \notin \text{BA}(\mathbf{T})$, and $\{a, b, c\}$ —because $d \notin \text{BA}(\mathbf{T})$. ME will output the union of these sets: $\text{ME}(\mathbf{T}) = \{a, b, c, d\}$. Note that the set $\{b, c, d\}$ is not BA-stable, as $a \in \text{BA}(\mathbf{T})$.

ME is one of several tournament solutions that can be defined based on this notion of stability. The top cycle for example, is the union of all minimal CNL-stable sets (Brandt 2011), where CNL is the tournament solution returning the set of all Condorcet nonlosers—meaning alternatives with at least one outgoing edge. We will use BA and ME to refer both to the tournaments solutions, and the equivalent SCFs.

The minimal extending set is Kelly-manipulable. However, we show it is still robust under pessimistic preferences, and extend this result to all coarsenings of ME.

Theorem 3. *ME is a robust Condorcet extension under all pessimistic preference extensions.*

Proof. For a set of alternatives A , and a set of agents N , let $\mathbf{P} =_{-i} \mathbf{P}'$ be i -variant profiles for an agent $i \in N$. Let \mathbf{P} be such that $x \in A$ is the Condorcet winner in \mathbf{P} . Let $\mathbf{T} = (A, \succ^{\mathbf{T}})$ and $\mathbf{T}' = (A, \succ^{\mathbf{T}'})$ be the (single-agent variant) tournaments induced by \mathbf{P} and \mathbf{P}' , respectively.

We assume $\text{ME}(\mathbf{T}') \neq \text{ME}(\mathbf{T})$.⁹ Because of this, we know $\overline{D}_{\mathbf{T}'}(x)$ is nonempty, as the two outcomes cannot differ if x remains a Condorcet winner in \mathbf{T}' . As $\mathbf{P} =_{-i} \mathbf{P}'$, any changes going from \mathbf{T} to \mathbf{T}' must be counter to agent i 's preferences. This implies $x \succ_i^P a$ for all $a \in \overline{D}_{\mathbf{T}'}(x)$. So, all alternatives in $\overline{D}_{\mathbf{T}'}(x)$ are worse than x to agent i .

We want to show that there is some minimal BA-stable set S of \mathbf{T}' , such that $S \cap \overline{D}_{\mathbf{T}'}(x) \neq \emptyset$. This would guarantee the existence of an alternative $a \in \overline{D}_{\mathbf{T}'}(x)$ that is also

⁹If no such i -variants exist, robustness of the rule would immediately follow.

in $\text{ME}(\mathbf{T}')$, precluding agent i with pessimistic preferences from preferring this outcome to $\text{ME}(\mathbf{T})$.

So suppose for contradiction that S is a minimal BA-stable set of \mathbf{T}' such that $S \cap \overline{D}_{\mathbf{T}'}(x) = \emptyset$. The only way this can be the case is if $S \subseteq D_{\mathbf{T}'}(x) \cup \{x\}$. We consider two cases.

Case 1: Suppose $x \notin S$. As x dominates all alternatives in $D_{\mathbf{T}'}(x)$, it will dominate all alternatives in S , as $S \subseteq D_{\mathbf{T}'}(x)$. This means x is a Condorcet winner in the tournament $(S \cup \{x\}, \succ_{S \cup \{x\}}^{\mathbf{T}'})$, and thus, $x \in \text{BA}(\succ_{S \cup \{x\}}^{\mathbf{T}'})$. So S cannot be a BA-stable set, contradicting our assumption that it is a minimal one.

Case 2: Suppose instead $x \in S$. To reach our contradiction, we want to show there exists an alternative $a \in \overline{D}_{\mathbf{T}'}(x)$ such that $a \in \text{BA}(\succ_{S \cup \{a\}}^{\mathbf{T}'})$ —which would imply S is not BA-stable. We use an algorithm proposed by Hudry (2004) to find such an alternative $a \in \text{BA}(\succ_{S \cup \{a\}}^{\mathbf{T}'})$. We start at step 1 with a transitive subtournament of $(S \cup \{a\}, \succ_{S \cup \{a\}}^{\mathbf{T}'})$. Let $\mathbf{S}_1 = (\{x, a\}, \succ_{\{x, a\}}^{\mathbf{T}'})$, for some $a \in \overline{D}_{\mathbf{T}'}(x)$. We label all remaining elements of S —which are all elements of $D_{\mathbf{T}'}(x)$ —in any order from 2 to $|S|$. At step k , we look at the alternative labelled k , and add it to the tournament \mathbf{S}_{k-1} to create \mathbf{S}_k , if it does not break transitivity to do so. As a dominates x , and x dominates all $a' \in D_{\mathbf{T}'}(x)$, adding any alternative outside the dominion of a will break transitivity, as it will create a 3-cycle. Thus, at any step, an alternative $a' \in D_{\mathbf{T}'}(x)$ will only be added to the tournament if $a \succ^{\mathbf{T}'} a'$. When the algorithm terminates after iterating through all alternatives, we will be left with a subtournament $\mathbf{S}_{|S|}$ of $(S \cup \{a\}, \succ_{S \cup \{a\}}^{\mathbf{T}'})$. It is easy to see the resulting tournament will be transitive, and it will indeed be a maximal transitive subtournament of $(S \cup \{a\}, \succ_{S \cup \{a\}}^{\mathbf{T}'})$, as no further alternatives can be added to the tournament without breaking transitivity. Importantly, the maximal element of the resulting subtournament will be a , meaning $a \in \text{BA}(\succ_{S \cup \{a\}}^{\mathbf{T}'})$. Thus, S cannot be an BA-stable set, which contradicts our assumption that it is a minimal one.

As we have shown that no subset of $D_{\mathbf{T}'}(x) \cup \{x\}$ can be a BA-stable set of \mathbf{T}' , any minimal BA-stable set must contain at least one element of $\overline{D}_{\mathbf{T}'}(x)$, meaning it cannot be the case that $\text{ME}(\mathbf{T}') \succ_i^{P, g} \text{ME}(\mathbf{T})$ when e is a pessimistic preference extension. \square

In terms of decisiveness, ME is among the more decisive tournament solutions that fail weak resoluteness, as it is a refinement of several prominent tournament solutions, including the top cycle and the Banks set (Brandt, Harrenstein, and Seedig 2017). We now show that all coarsenings of a robust SCF inherit the robustness property.

Proposition 3. *If a Condorcet-consistent SCF f is robust under pessimistic preferences, then all Condorcet-consistent coarsenings of f are robust under pessimistic preferences.*

Proof. Let f be a SCF that is robust under pessimistic preferences, and let f' be a Condorcet-consistent coarsening of

f . Let \mathbf{P} be a profile with a Condorcet winner a . Note that $f(\mathbf{P}) = f'(\mathbf{P}) = \{a\}$ as both are Condorcet extensions.

Suppose \mathbf{P}' is an i -variant of \mathbf{P} for some agent $i \in N$. Because f is robust under pessimistic preferences, either (i) there must be some $a' \in f(\mathbf{P}')$ such that $a \succ_i^P a'$, or (ii) $f(\mathbf{P}) = f(\mathbf{P}')$.

If (i) is the case, then as $f(\mathbf{P}') \subseteq f'(\mathbf{P}')$, a' is also an element of $f'(\mathbf{P}')$. As $f'(\mathbf{P}) = \{a\}$, we know $a' \in f'(\mathbf{P}') \setminus f'(\mathbf{P})$, meaning by definition of a pessimistic extension that it cannot be the case that $f'(\mathbf{P}') \succ_i^{P, e} f'(\mathbf{P})$ for any pessimistic extension e .

If (ii) is the case, we know a must also be the Condorcet winner in \mathbf{P}' as f cannot satisfy weak resoluteness if it is robust under any preference extension, and therefore does not return singletons outside the Condorcet domain. Since f' is also a Condorcet extension, we know $f'(\mathbf{P}') = \{a\}$, meaning $f'(\mathbf{P}') \not\succeq_i^{P, e} f'(\mathbf{P})$. \square

Corollary 1. *Condorcet-consistent coarsenings of ME are robust under pessimistic preferences.*

Corollary 1 follows from Proposition 3 and Theorem 3, and it establishes the robustness of the Banks set. Note that Corollary 1 is not restricted to tournament-solution SCFs, but holds for all Condorcet-consistent SCFs.

5 Conclusion

We have introduced the strategyproofness-related notion of a robust Condorcet extension. We have argued that Condorcet extensions that are robust are preferable to those that are not, as we can trust that they will return true Condorcet winners when they exist. We have introduced an axiom—weak resoluteness—and shown that no weakly resolute tournament solution can be a robust Condorcet extension. Finally, we have shown that the minimal extending set is a robust Condorcet extension under all pessimistic preferences, and have extended this result to all coarsenings of ME.

We have argued that in lieu of searching for fully strategyproof rules, a fruitful endeavour is to explore immunity against more specific manipulations that may interact with, and compromise, other desirable properties satisfied by manipulable social choice functions. We have scratched the surface in this paper, but have limited our exploration to robustness of irresolute rules in general, and tournament solutions in particular. These are, of course, only a small class of all Condorcet extensions, and it remains to be seen if similar results can be obtained for other classes. Finally, there are many other areas of social choice theory open to similar notions of strategyproofness. For example, while proportionality and strategyproofness cannot be satisfied by the same multiwinner approval voting rule (Peters 2018; Kluyving et al. 2020), one might want to ask to what degree proportionality is affected by the failure of strategyproofness.

Acknowledgements

We thank the reviewers for their extensive feedback.

References

- Amanatidis, G.; Birmpas, G.; and Markakis, E. 2016. On Truthful Mechanisms for Maximin Share Allocations. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI-2016)*.
- Arrow, K. J.; Sen, A.; and Suzumura, K. 2010. *Handbook of Social Choice and Welfare*, volume 2. Elsevier.
- Banks, J. S. 1985. Sophisticated Voting Outcomes and Agenda Control. *Social Choice and Welfare* 1(4): 295–306.
- Barberà, S. 1977. Manipulation of Social Decision Functions. *Journal of Economic Theory* 15(2): 266–278.
- Barberà, S.; Bossert, W.; and Pattanaik, P. K. 2004. Ranking Sets of Objects. In *Handbook of Utility Theory*, 893–977. Springer.
- Bartholdi, J. J.; Tovey, C. A.; and Trick, M. A. 1989. The Computational Difficulty of Manipulating an Election. *Social Choice and Welfare* 6(3): 227–241.
- Black, D. 1948. On the Rationale of Group Decision-making. *Journal of Political Economy* 56(1): 23–34.
- Bossert, W.; and Sprumont, Y. 2014. Strategy-proof Preference Aggregation: Possibilities and Characterizations. *Games and Economic Behavior* 85: 109–126.
- Brandl, F.; Brandt, F.; and Stricker, C. 2018. An Analytical and Experimental Comparison of Maximal Lottery Schemes. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI-2018)*.
- Brandt, F. 2011. Minimal Stable Sets in Tournaments. *Journal of Economic Theory* 146(4): 1481–1499.
- Brandt, F. 2015. Set-Monotonicity Implies Kelly-Strategyproofness. *Social Choice and Welfare* 45(4): 793–804.
- Brandt, F.; Brill, M.; and Harrenstein, P. 2016. Tournament Solutions. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*, chapter 3. Cambridge University Press.
- Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D. 2016. *Handbook of Computational Social Choice*. Cambridge University Press.
- Brandt, F.; Harrenstein, P.; and Seedig, H. G. 2017. Minimal Extending Sets in Tournaments. *Mathematical Social Sciences* 87: 55–63.
- Campbell, D. E.; and Kelly, J. S. 2003. A strategy-proofness characterization of majority rule. *Economic Theory* 22(3): 557–568.
- Conitzer, V.; and Walsh, T. 2016. Barriers to Manipulation in Voting. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*, chapter 6. Cambridge University Press.
- Conitzer, V.; Walsh, T.; and Xia, L. 2011. Dominating Manipulations in Voting with Partial Information. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI-2011)*.
- Copeland, A. H. 1951. A “Reasonable” Social Welfare Function. University of Michigan Seminar on Applications of Mathematics to the Social Sciences.
- Duggan, J.; and Schwartz, T. 2000. Strategic Manipulability Without Resoluteness or Shared Beliefs: Gibbard-Satterthwaite Generalized. *Social Choice and Welfare* 17(1): 85–93.
- Fishburn, P. C. 1977. Condorcet Social Choice Functions. *SIAM Journal on Applied Mathematics* 33(3): 469–489.
- Fishburn, P. C. 1984. Probabilistic Social Choice Based on Simple Voting Comparisons. *The Review of Economic Studies* 51(4): 683–692.
- Gaertner, W. 2001. *Domain Conditions in Social Choice Theory*. Cambridge University Press.
- Gärdenfors, P. 1976. Manipulation of Social Choice Functions. *Journal of Economic Theory* 13(2): 217–228.
- Gehrlein, W. V. 2006. *Condorcet’s Paradox*. Springer.
- Gehrlein, W. V.; and Lepelley, D. 2011. *Voting Paradoxes and Group Coherence*. Springer.
- Gibbard, A. 1973. Manipulation of Voting Schemes: A General Result. *Econometrica* 41(4): 587–601.
- Hajaj, C.; Dickerson, J. P.; Hassidim, A.; Sandholm, T.; and Sarne, D. 2015. Strategy-Proof and Efficient Kidney Exchange Using a Credit Mechanism. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI-2015)*.
- Hoang, L. N. 2017. Strategy-proofness of the Randomized Condorcet Voting System. *Social Choice and Welfare* 48(3): 679–701.
- Hudry, O. 2004. A Note on “Banks Winners in Tournaments are Difficult to Recognize” by G.J. Woeginger. *Social Choice and Welfare* 23(1): 113–114.
- Kelly, J. S. 1977. Strategy-proofness and Social Choice Functions without Singlevaluedness. *Econometrica* 45(2): 439–446.
- Kemeny, J. G. 1959. Mathematics without Numbers. *Daedalus* 88(4): 577–591.
- Kluiving, B.; de Vries, A.; Vrijbergen, P.; Boixel, A.; and Endriss, U. 2020. Analysing Irresolute Multiwinner Voting Rules with Approval Ballots via SAT Solving. In *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI-2020)*.
- Kruger, J.; and Terzopoulou, Z. 2020. Strategic Manipulation with Incomplete Preferences: Possibilities and Impossibilities for Positional Scoring Rules. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2020)*.
- McGarvey, D. C. 1953. A Theorem on the Construction of Voting Paradoxes. *Econometrica* 21(4): 608–610.
- Meir, R. 2018. Strategic voting. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 13(1): 1–167.

- Osborne, M. J.; and Rubinstein, A. 2003. Sampling Equilibrium, with an Application to Strategic Voting. *Games and Economic Behavior* 45(2): 434–441.
- Peters, D. 2018. Proportionality and Strategyproofness in Multiwinner Elections. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems, (AAMAS-2018)*.
- Reijngoud, A.; and Endriss, U. 2012. Voter Response to Iterated Poll Information. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2012)*.
- Sato, S. 2008. On Strategy-proof Social Choice Correspondences. *Social Choice and Welfare* 31(2): 331–343.
- Sato, S. 2013. A Sufficient Condition for the Equivalence of Strategy-proofness and Nonmanipulability by Preferences Adjacent to the Sincere one. *Journal of Economic Theory* 148(1): 259–278.
- Satterthwaite, M. A. 1975. Strategy-proofness and Arrow's Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions. *Journal of Economic Theory* 10(2): 187–217.
- Shoham, Y.; and Leyton-Brown, K. 2009. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press.
- Slater, P. 1961. Inconsistencies in a Schedule of Paired Comparisons. *Biometrika* 48(3–4): 303–312.