

Single-Winner Voting with Alliances: Avoiding the Spoiler Effect

Grzegorz Pierczyński
AGH University, Poland
University of Warsaw, Poland
g.pierczynski@gmail.com

Stanisław Szufa
AGH University, Poland
CNRS, LAMSADE, Université Paris Dauphine-PSL, France
s.szufa@gmail.com

ABSTRACT

We study the setting of single-winner elections with ordinal preferences where candidates might be members of *alliances* (which may correspond to e.g., political parties, factions, or coalitions). However, we do not assume that candidates from the same alliance are necessarily adjacent in voters’ rankings. In such a case, every classical voting rule is vulnerable to the spoiler effect, i.e., the presence of a candidate may harm his or her alliance. We therefore introduce a new idea of *alliance-aware* voting rules which extend the classical ones. We show that our approach is superior both to using classical cloneproof voting rules and to running primaries within alliances before the election. We introduce several alliance-aware voting rules and show that they satisfy the most desirable standard properties of their classical counterparts as well as newly introduced axioms for the model with alliances which, e.g., exclude the possibility of the spoiler effect. Our rules have natural definitions and are simple enough to explain to be used in practice.

KEYWORDS

voting, single-winner elections, spoiler effect, alliances, primaries

ACM Reference Format:

Grzegorz Pierczyński and Stanisław Szufa. 2024. Single-Winner Voting with Alliances: Avoiding the Spoiler Effect. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024)*, Auckland, New Zealand, May 6 – 10, 2024, IFAAMAS, 9 pages.

1 INTRODUCTION

Imagine that there is a presidential election in the Republic of Social Choice. There are only two political parties, Party A and Party B. Party B is represented by only one candidate, Bob, while Party A has two leaders willing to candidate, Alice and Adam. Alice presents more center-wing views, and is therefore considered to have greater chances to win against Bob. On the other hand, Adam is much more popular within the electorate of Party A. In the end, both candidates decide to run in the election. The voters’ preferences are as follows:

46% of voters: Adam > Alice > Bob
5% of voters: Alice > Bob > Adam
43% of voters: Bob > Alice > Adam
6% of voters: Bob > Adam > Alice

Note that, apart from the clear supporters of Party A and Party B, there are also 5% of centrist voters, who, in general, like Party A more than Party B, yet they consider Adam to be too extreme, and hence prefer Bob to Adam, even though Bob is from Party B. On the other hand, there are 6% of Bob’s supporters who were impressed by Adam’s performance in the pre-election debate and consider him to be a better candidate than Alice.

Since the Republic of Social Choice uses Plurality in presidential elections, Bob is declared the winner. The supporters of Party A are very disappointed, especially since they know that they “could have won” if only Adam did not run in the election. Their main question is then how to avoid such situations in the future.

Some of them say that the solution is clear—there should have been only one running candidate from Party A instead of several ones. However, if Party A held a primary election, Adam would have won (no matter whether the voters taking part in the primaries represented the views of only the supporters of Party A or of the whole population). On the other hand, since Party A presents itself as a transparent and pro-democratic party, the backstage nomination of Alice was completely out of the question.

An alternative opinion is that the problem stemmed from using Plurality to determine the winner. “In our elections, we should use a more sophisticated rule that is independent of clones”, it is argued, “for example, STV or the Schulze’s method”. However, the proposed rules—as well as every other well-established ordinal cloneproof rule the supporters of Party A are aware of—would still elect Bob under the above voters’ preferences.

Studying the existing literature on social choice, the supporters of Party A could not find a fully satisfactory solution to their problem.

1.1 Our Contribution

The situation presented above is sometimes called the *spoiler effect* [4]—Adam not only did not win, but also prevented a member of his party from winning. We propose a new idea that solves this problem, *alliance-aware voting rules*, that take as input not only voters’ preferences, but also the information about the non-overlapping *alliances* (parties) of the candidates.

In Section 3, we propose several basic axioms that are desirable for alliance-aware voting rules. Our basic axioms, discussed in Section 3.1, prevent the spoiler effect among allies and ensure that it is not possible to affect the election result by certain sorts of manipulation (such as duplicating allies, or splitting a losing alliance into smaller ones). We argue that the rules based on the mechanics of Plurality and Maximin are best-suited for satisfying these axioms. Next, in Section 3.2 we present additional axioms that concern electing the best candidate within the winning alliance and argue that one could have two reasonable, yet contradictory, intuitions



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand. © 2024 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

here. Then in Section 4 we present four alliance-aware rules and show that they have good axiomatic properties (i.e., they satisfy our newly introduced notions and preserve the most critical properties of the classical variants of Plurality or Maximin they are based on).

In Section 5 we analyze our setting experimentally. We show that the standard variants of our alliance-aware rules often admit the spoiler effect among allies, even if we run primaries within each alliance before the voting. Further, we show that the additional axiomatic guarantees provided by our rules, fully eliminating this problem, do not require sacrificing much of social welfare.

Notably, all the four alliance-aware rules are polynomial-time computable and have simple, intuitive definitions. This fact makes them reasonable proposals for real-life elections. Additionally, in Section 6 we show that the extensions of Plurality can be implemented with simpler ballots than the full ordinal ones and discuss how our results can be applied to the model where the alliances are not fully disjoint.

1.2 Related Work

Our work is closely related to two widely studied topics in the literature on social choice—the spoiler effect and strategic candidacy.

The most important work regarding the spoiler effect is the famous work of Arrow [2], where the Arrow’s Impossibility Theorem is proved. This theorem states that every reasonable rule (satisfying the very weak conditions of unanimity and non-dictatorship) violates the *Independence of Irrelevant Alternatives* (IIA) axiom. It means that for every reasonable rule the removal of a non-winning candidate c may result in changing the result of the election—which means that the rule is vulnerable to the spoiler effect. This work has had tremendous impact on social choice, which manifested itself in the number of works that discuss, criticize, or propose weaker variants of this axiom (see, for example, the works of Bordes and Tideman [3], Maskin [16], Osborne [21], Ray [25] and Sen [29]).

Among such works, the one by Tideman [30] is the most relevant to our paper. Here the axiom of *independence of clones* was introduced. The idea is to focus only on avoiding the spoiler effect within groups of *clones* that intuitively correspond to candidates that are, in a certain formal sense, similar. We find this weakening natural and intuitive. However, in that work the similarity is deduced from the votes—the clones are required to be adjacent in all the voters’ rankings. Note that in our example, since a small fraction of voters does not rank Alice and Adam next to one another, this axiom does not provide any guarantees regarding vote-splitting. On the contrary, we assume that the information about the candidates’ similarity is given upfront.

So far only few papers studied the immunity of certain voting rules to the spoiler effect; a notable exception is the analysis of STV in the work of McCune and Wilson [17]. From the computational complexity perspective, the spoiler effect has been studied within so called *electoral control problems* [5, 10, 14, 18, 20].

The next group of related works shares with us the assumption that the information about candidates’ views is given, yet typically in the form of rankings over other candidates. Then the problem of the spoiler effect is studied from the strategic perspective (see the pioneering paper of Dutta et al. [6] who prove an analogous result to the Arrow’s one in this setting, and further follow-ups e.g.,

by Ehlers and Weymark [8], Lang et al. [13], Polukarov et al. [23] and Rodríguez-Álvarez [26, 27]).

Some of the most closely connected works to ours are the ones by Faliszewski et al. [11] and by Harrenstein et al. [12]. The first one studies the computational complexity of problems connected with selecting the optimal representative of a party in single-winner elections, while the second one characterizes Nash equilibria in a generalized Hotelling-Downs model.

2 PRELIMINARIES

An *election* is a tuple (C, V, \mathcal{A}) where $C = \{c_1, \dots, c_m\}$ is the set of m candidates, $V = \{v_1, \dots, v_n\}$ is the set of n voters, and $\mathcal{A} \subseteq 2^C$ denotes the set of *alliances* between candidates. Each voter $v_i \in V$ is associated with a strict linear order (ranking) $>_i$ over the candidates. For each pair of candidates $a, b \in C$, we let $\text{Pref}(a, b) = |\{v_i \in V : a >_i b\}|$ and say that a *wins head-to-head against* b if $\text{Pref}(a, b) > n/2$. We assume that $|\mathcal{A}| \geq 2$ and that every candidate belongs to exactly one alliance (possibly a single-element one), i.e., we have that $\bigcup_{A \in \mathcal{A}} A = C$ and for each pair of alliances $A_1, A_2 \in \mathcal{A}$ we have that $A_1 \cap A_2 = \emptyset$. For each candidate $c \in C$, by $A(c)$ we denote the alliance containing c . We say that candidates $a, b \in C$ are *allies* if they belong to the same alliance, i.e., $A(a) = A(b)$; otherwise, we say that a and b are *opponents*.

We consider two special cases of our model:

- *two-alliance* elections, where $|\mathcal{A}| = 2$,
- *no-ally* elections, where $|\mathcal{A}| = m$ (i.e., for every $A \in \mathcal{A}$ we have that $|A| = 1$).

Additionally, we say that a voter $v_i \in V$ is *alliance-consistent* if for each two allies a_1, a_2 there exists no candidate $b \notin A(a_1)$ such that $a_1 >_i b >_i a_2$.

A *voting rule* (or, a *rule*, for short) is a function that, for a given election, returns a winning candidate c .¹ We say that a voting rule f is *standard* if it does not use the information about alliances; i.e., for every two elections $E_1 = (C, V, \mathcal{A})$ and $E_2 = (C, V, \mathcal{A}')$ we have that $f(E_1) = f(E_2)$. If a voting rule is not standard, we say that it is *alliance-aware*.

Given an election E , we sometimes say that an alliance $A \in \mathcal{A}$ wins under f if a member of A wins under f , i.e., $f(E) \in A$. Analogously, an alliance $A \in \mathcal{A}$ loses under f if $f(E) \notin A$.

2.1 Standard Axioms of Voting

We consider the following well-known voting axioms, proposed in the literature for standard voting rules:

Majority consistency. A *majority winner* is a candidate who is ranked on top by a majority (more than $n/2$) of the voters.

We say that a voting rule is *majority-consistent* if it elects the majority winner whenever one exists.

Condorcet consistency. A *Condorcet winner* is a candidate who wins head-to-head with all the other candidates. We say that a voting rule is *Condorcet-consistent* if it elects the Condorcet winner whenever one exists.

The next axioms are based on the idea that “modifying” a considered election in some way should neither “harm” nor “help” certain

¹Typically, definitions of voting rules allow returning a set of tied winners. We assume that ties are broken lexicographically. The choice of a (deterministic) tie-breaking rule does not affect our results.

candidates. For a fixed voting rule f , we say that a candidate $c \in C$ is *harmed* by modifying election E into E' if $f(E) = c$ and $f(E') \neq c$. Similarly, we say that c is *helped* if $f(E) \neq c$ and $f(E') = c$.

Monotonicity. We say that a voting rule f is *monotone* if no candidate can be harmed by improving his or her position in rankings of some voters (without changing the relative order of the other candidates).

Independence of clones (cloneproofness). We say that a set of candidates $T \subseteq C$ are *clones* if there are no candidates $a, b \in T$, and $c \notin T$ such that for some voter $v_i \in V$ it holds $a >_i c >_i b$. Intuitively, candidates from T are adjacent in all voters' rankings. A voting rule f is *independent of clones* (or, *cloneproof*) if removing candidates from a set of clones T from the election neither harms nor helps candidates outside of T . Formally, it means that (1) if $f(E) \in T$ then $f(E') \in T$ and (2) if $f(E) \notin T$ then $f(E) = f(E')$.

2.2 Standard Voting Rules

In our work we consider the following well-established standard voting rules:

Positional scoring rules. Positional scoring rules elect the candidate maximizing the score, which can be characterized by a vector $\vec{s} = (\ell_1, \ell_2, \dots, \ell_m)$ such that $\ell_1 \geq \ell_2 \geq \dots \geq \ell_m$. A candidate c gets ℓ_i points from each voter v who ranks c at the i th position. The best-known positional scoring rule is Plurality, characterized by the vector $\vec{s} = (1, 0, \dots, 0)$. We further call the number of voters who rank a candidate c on top the *plurality score* of c .

Copeland. Copeland elects the candidate $c \in C$ who wins head-to-head against the greatest number of candidates.²

Maximin. Maximin elects the candidate $c \in C$ maximizing the *maximin score*, defined as $\min_{c' \in C} \text{Pref}(c, c')$.

STV. Single Transferable Vote proceeds recursively: In each round it eliminates the candidate with the smallest plurality score until there remains only one candidate.

Additionally, in the full version of our paper [22] we present the definition and analysis of the Schulze's rule [28].

Our considered alliance-aware voting rules will *extend* the above standard ones. Formally, an alliance-aware voting rule f extends a standard voting rule g , if for each no-ally election E it holds that $f(E) = g(E)$.

3 AXIOMS OF VOTING WITH ALLIANCES

We can clearly see that alliance-aware voting rules use more information than standard ones, which might potentially result in electing better candidates. In this section we present axioms that we consider natural for this setting.

3.1 Basic Axioms

Our overall intuition is that the additional information about alliances should be used to avoid vote-splitting between allies and *only* for that purpose. We formalize it via two following axioms:

²There are several variants of Copeland dealing differently with tied head-to-head comparisons (that is, head-to-head comparisons between candidates a and b where $\text{Pref}(a, b) = n/2$); the choice of a specific variant does not affect our results.

DEFINITION 1 (ALLY-NO-HARM). We say that a voting rule f satisfies *ally-no-harm criterion* if for every elections E, E' such that E' is obtained from E by removing a candidate c we have that

$$f(E) \notin A(c) \implies f(E') \notin A(c).$$

DEFINITION 2 (RESISTANCE TO ALLIANCE-SPLITTING). We say that a voting rule f is *resistant to alliance-splitting* if for every elections E, E' such that E' is obtained from E by splitting an alliance A into two alliances A_1, A_2 (so that $A_1 \cup A_2 = A$ and $A_1 \cap A_2 = \emptyset$) we have that

$$f(E) \notin A \implies f(E) = f(E').$$

Intuitively, the first axiom ensures that a candidate is never a spoiler for members of his or her alliance. The second axiom ensures that the information about alliances is not “abused” by an alliance-aware voting rule. The only consequence of splitting an alliance A into two alliances A_1, A_2 should be increasing the level of vote-splitting between members of A_1 and A_2 , which might result in decreasing their winning chances. Therefore, if no candidate from A wins anyway, splitting A should not affect the outcome of the rule at all. As a result, resistance to alliance-splitting intuitively rules out alliance-aware voting rules that are too artificial (whose definitions depend on, e.g. the number of alliances or their quantities) or that are too easily manipulable.

In order to fully avoid the spoiler effect, we also propose an analogue of the monotonicity axiom for the setting with alliances: improving the level of support of a candidate should never harm his or her alliance.

DEFINITION 3 (ALLIANCE MONOTONICITY). We say that a voting rule f is *alliance monotone* if for every elections E, E' such that E' is obtained from E by improving the position of some candidate c in a voter's ranking (without changing the relative order of other candidates), it holds that

$$f(E) \in A(c) \implies f(E') \in A(c).$$

Note that this axiom is incomparable to monotonicity—its guarantees are stronger for the case when $f(E) \in A(c) \setminus \{c\}$ and weaker for the case when $f(E) = c$.

If a rule satisfies ally-no-harm and alliance monotonicity, then an alliance never has an incentive to discourage its members from running or performing well in elections. However, for such rules another issue becomes critical: It should be never profitable for an alliance to *improve* its result by nominating several similar candidates instead of one as well. Otherwise, these rules would face the opposite problem of promoting bigger alliances, even if their size follows merely from duplicating candidates. We should require that the presence of a candidate may help his or her alliance only if it also enriches the actual offer for voters.

Formally, given an election, let us call two allies a_1, a_2 *similar* to each other if for each candidate b such that b is an opponent of a_1 and a_2 there is no voter v_i such that $a_1 >_i b >_i a_2$. Intuitively, this condition means that both candidates have the same support structure compared to candidates outside of $A(a_1)$.

Our concept of similar candidates is based on the same intuition as the idea of clones proposed by Tideman [30]. However, both definitions are actually incomparable, as we can see in Example 1:

EXAMPLE 1. Consider the following election with two voters:

$$v_1: \quad b > a_1 > a_2 > a_3$$

$$v_2: \quad a_2 > b > a_1 > a_3$$

Candidates a_1, a_2, a_3 are allies, while b is their opponent. Candidates b and a_1 are clones, yet they are not similar. On the other hand, candidates a_1 and a_3 are similar, but they are not clones.

Our next axiom can be viewed as an analogue of standard independence of clones, yet using the notion of similar candidates:

DEFINITION 4 (INDEPENDENCE OF SIMILAR ALLIES). We say that a voting rule f is independent of similar allies if for every elections E, E' such that E' is obtained from E by removing a candidate $c \in C$ similar to some other candidate $c' \in C$, we have that:

$$(1) \quad f(E) \in A(c') \implies f(E') \in A(c'),$$

$$(2) \quad f(E) \notin A(c') \implies f(E) = f(E').$$

Naturally, both axioms are incomparable. However, if we assume that every set of clones corresponds to a separate alliance³ then independence of similar allies implies independence of clones.

We believe that satisfying the four aforementioned axioms is crucial in the setting with alliances. Therefore, alliance-aware voting rules satisfying them will be called *basic*. Interestingly, at this point we can already exclude some certain types of alliance-aware voting rules—namely, those that extend non-monotone rules (such as STV), Copeland or scoring rules different than Plurality.

PROPOSITION 1. There exists no basic alliance-aware voting rule f that extends (1) a non-monotone standard rule, (2) the Copeland’s rule, or (3) a scoring rule different than Plurality.

PROOF. The extension of each non-monotone rule would violate alliance monotonicity for no-ally elections. For the Copeland’s rule, it is enough to consider a no-ally election E with three candidates a, b, c and three voters with preferences: $a > b > c, b > c > a$, and $c > a > b$. Then, under Copeland, every candidate is a tied winner. The election is symmetric, so without loss of generality we assume that a wins. Now, consider a modified election E' in which the candidate c' is added just below c in each ranking (the election remains no-ally). Now either b or c are elected by Copeland. Consider now election E'' obtained from E' by merging $A(c)$ and $A(c')$. From resistance to alliance-splitting, we obtain that the winner is c, c' , or b . However, now c and c' are similar, and by removing c' we go back to election E in which a is the winner, which violates independence of similar allies.

Now fix a scoring rule f , characterized by $\vec{s} = (\ell_1, \ell_2, \dots, \ell_m)$. Let $\ell_2 > \ell_m$ for $m > 2$ (otherwise, f is Plurality). Consider now a symmetric election E with two candidates a, b and two votes: $a > b$ and $b > a$. Now both candidates have equal score, so assume that a wins. Now add a candidate b' to each vote, just below b . Then the score of a is $\ell_1 + \ell_m$, while the score of b is $\ell_1 + \ell_2 > \ell_1 + \ell_m$, hence b wins. The remaining part of the proof (merging $A(b)$ and $A(b')$, and removing b') is the same as in the case of Copeland. \square

From now, we will therefore focus only on alliance-aware voting rules extending Plurality or Maximin.

³Since alliances are disjoint, this assumption implies that there are no nested groups of clones in an election. We will further address this limitation in Section 6.2.

3.2 Individual-Oriented Axioms


Observe that the axioms considered so far only say which alliances should or should not win in certain situations, yet they say nothing about which “precise” candidate from the winning alliance should win. Here we have only guarantees provided by standard axioms, such as the majority or Condorcet consistency. However, one could define here also additional *individual-oriented* axioms, attempting to formalize a notion of “a candidate who is individually strong” in the model with alliances. Consider the following election with three candidates $\{a_1, a_2, b\}$ and alliances $\{\{a_1, a_2\}, \{b\}\}$:

$$3 \text{ votes: } \quad a_1 > b > a_2$$

$$37 \text{ votes: } \quad b > a_1 > a_2$$

$$11 \text{ votes: } \quad b > a_2 > a_1$$

$$49 \text{ votes: } \quad a_2 > a_1 > b$$



Suppose that we use a basic rule f extending either Plurality or Maximin. Here, ally-no-harm requires that the winner should come from the alliance $\{a_1, a_2\}$. However, who should be the winner in this case? We could have two contradictory intuitions here:

INTUITION 1. The winner should be a_1 , since he or she won head-to-head against b , while a_2 lost. A victory against an opponent is more important to judge the quality of a candidate than a victory against an ally.

INTUITION 2. The winner should be a_2 , since he or she would be elected by both Plurality and Maximin if there were no alliances. The additional information about alliances should be used by f only to guarantee that the alliance $\{a_1, a_2\}$ wins, yet it should not affect the choice of the winner if the proper alliance wins anyway.

Let us first focus on Intuition 2. We formalize it as follows:

DEFINITION 5 (INDEPENDENT WINNER CONSISTENCY). Given election E and voting rule f , we say that candidate c is an independent winner under f , if $c = f(E')$, where E' is obtained from E by splitting $A(c)$ into $\{c\}$ and $A(c) \setminus \{c\}$.

We say that a voting rule f is independent-winner-consistent (or, IW-consistent, for short) if it elects the independent winner whenever one exist.

Intuitively, independent winner is a candidate c who would have won if he or she took part in the election as an individual candidate, not as a member of his or her alliance. Although in our paper we do not consider strategic behavior of candidates, we interpret this possibility as an evidence of the high voters’ support for c .

For each election and each basic alliance-aware voting rule, we can never have two independent winners. Indeed, suppose that for some election E and rule f there are two independent winners, a and b . First, from resistance to alliance-splitting we obtain that in the original election both $A(a)$ and $A(b)$ are winning, hence $A(a) = A(b)$. Second, consider the modified election E' in which both a and b are excluded from $A(a)$, and each of them forms a singleton alliance. We can obtain this election either by excluding a first and then b , or vice versa. Then, from resistance to alliance-splitting we obtain that both a and b should be winners in E' , a contradiction.

On the other hand, the independent winner may not exist, as we can see in the following example:

EXAMPLE 2. Consider the following election $E = (C, V, \mathcal{A})$ with $C = \{a_1, a_2, b_1, b_2\}$ and $\mathcal{A} = \{\{a_1, a_2\}, \{b_1, b_2\}\}$. Voters' preferences are the following:

$$\begin{array}{l} v_1: \quad b_2 > a_1 > b_1 > a_2 \\ v_2: \quad a_2 > b_2 > a_1 > b_1 \\ v_3: \quad b_1 > a_1 > a_2 > b_2 \end{array} \quad \begin{array}{ccc} & a_1 & \xrightarrow{2:1} & b_1 \\ & \downarrow 2:1 & \nearrow 2:1 & \uparrow 2:1 \\ & a_2 & \xrightarrow{2:1} & b_2 \end{array}$$

Consider a basic alliance-aware rule f . Let us assume that f elects a_1 in the above election and it elects b_1 after splitting alliance $\{a_1, a_2\}$ into $\{a_1\}$ and $\{a_2\}$. Then there is no independent winner. Indeed, a_1 and a_2 are not independent winners since they lose election after splitting the alliance $\{a_1, a_2\}$. On the other hand, if in the original election a_1 is the winner then, after splitting group $\{b_1, b_2\}$ into $\{b_1\}$ and $\{b_2\}$, a_1 is still the winner (from resistance to alliance-splitting). Hence, also b_1 and b_2 are not independent winners.

In such case, independent winner consistency does not specify who should win the election.

Let us now discuss Intuition 1. Here we propose the following definition:

DEFINITION 6 (SOLITARY WINNER CONSISTENCY). Given a 2-alliance election E and a voting rule f , we say that a candidate c is a solitary winner under f , if $c = f(E')$, where E' is obtained from E by removing all the candidates from $A(c)$, except for c .

We say that a voting rule f is solitary-winner-consistent (or, SW-consistent, for short) if for each election E such that the set of solitary winners W is nonempty, we have that $f(E) \in W$.

Intuitively, a solitary winner is a candidate c who would have won if he or she were chosen as the only representative of $A(c)$.

One could wonder why we require that an election E is two-alliance. In the full version of our paper [22] we discuss the definition of this axiom without this requirement. We show that then the definition would not be satisfiable by basic alliance-aware rules extending Plurality, Maximin or the Schulze's method. At the same time, we show that such a generalization in some cases seems to be too restrictive and therefore less appealing. Intuitively, in an election with more than two alliances, a candidate might be a solitary winner not because of his or her own high support, but because of the vote-splitting between the opponents.

On the contrary, Observation 1 and Observation 2 below illustrate that the implications of Definition 6 are very intuitive and natural. Note that, since we are interested only in basic alliance-aware rules (in particular, those resistant to alliance-splitting), the solitary winner axiom nonetheless provides certain guarantees also in elections with more than two alliances.

OBSERVATION 1. If a basic alliance-aware voting rule is majority-consistent then for each election E the set of solitary winners includes the subset of candidates who would be majority winners if all their allies were removed from E .

OBSERVATION 2. If a basic alliance-aware voting rule is Condorcet-consistent then for each election E the set of solitary winners is equal to the set of candidates who win head-to-head against every opponent.

Directly from the above observations, we note that in some elections there might be several solitary winners, while in some other ones there might be none.

Since solitary winner and independent winner consistency formalize Intuition 1 and Intuition 2 respectively, we can already note that no basic alliance-aware voting rule extending Plurality or Maximin can satisfy them both. In fact, this is true for every reasonable basic rule.

PROPOSITION 2. No deterministic basic alliance-aware rule that is majority-consistent for elections with two candidates, is both SW-consistent and IW-consistent.

PROOF. Consider a no-ally election E with three candidates $\{a, b, c\}$ and three voters with the following preferences: $a > b > c$, $b > c > a$, $c > a > b$. The election is symmetric, hence we can assume that a is the winner. Consider now a modified election E' obtained from E by merging $A(a)$ and $A(b)$. Now, from resistance to alliance-splitting, we know that alliance $\{a, b\}$ wins. It means that a is the independent winner. However, a loses head-to-head against c , while b wins, hence only b is a solitary winner. \square

In the next section we present basic alliance-aware voting rules satisfying the independent winner and solitary winner axioms. Apparently, all of them have simple and easy-to-explain definitions.

Which of them should be in this case preferred to be used in practice? The answer to this question depends on the context. For instance, if candidates can predict the results of the elections (for example, via polls), choose their alliances in a strategic way and bear no costs from starting individually in elections, we should definitely prefer IW-consistent rules. If this is not the case—for example, the membership of candidates in their alliances is fixed—we believe that SW-consistent rules lead to more intuitive results.

4 ALLIANCE-AWARE VOTING RULES

In this section we present four basic alliance-aware rules. Two of them extend Plurality and two extend Maximin. Orthogonally, two of them are solitary-winner-consistent, and the other two are independent-winner-consistent. Let us first present the key idea on which all the four rules are based.

Let $f \in \{\text{Plurality}, \text{Maximin}\}$ be a standard voting rule. Note that the definition of f is based on maximizing a specifically defined f -score—the plurality score of a candidate $c \in C$ is equal to the number of ballots ranking c on top, while the maximin score of c is equal to the number of supporters in his or her most difficult head-to-head comparison (i.e., $\min_{c' \in C} \text{Pref}(c, c')$). Both definitions of score have a very interesting property.

OBSERVATION 3. Consider an election E and an election E' obtained from E by removing some candidate c . Consider now a candidate $c' \neq c$. The plurality|maximin score of c' in E' is greater or equal to his or her score in E .

Intuitively, both in case of Plurality and Maximin, the presence of additional candidates can only decrease the score of c . We extend the idea of the f -score to the setting with alliances as follows:

DEFINITION 7 (ALLIANCE-AWARE f -SCORE). Let f be either Plurality or Maximin. Fix an election E and a candidate c . Let E' be the election obtained from E by removing all the candidates from $A(c) \setminus \{c\}$.

The alliance-aware f -score of c in E is defined as his or her standard f -score in E' .

According to the above definition, the alliance-aware plurality score is the number of votes in which no opponent of c is ranked higher than c . On the other hand, the alliance-aware maximin score of c is the worst-case number of supporters in his or her head-to-head comparisons against his or her opponents, i.e., $\min_{c' \in C \setminus A(c)} \text{Pref}(c, c')$.

4.1 Independent-Winner-Consistent Rules

Now we are ready to define the IW-consistent alliance-aware variant of Plurality and Maximin.

DEFINITION 8 (IW- f). Let $f \in \{\text{Plurality}, \text{Maximin}\}$. The IW- f rule has two rounds: The first one chooses the winning alliance and the second one chooses the winning candidate. In the first round, we compute alliance-aware f -scores for each candidate and we choose the alliance A whose member has the highest score. In the second round, only candidates from A can win, yet the other ones are not removed from the election. The winner is the candidate from A with the highest standard f -score.

Let us illustrate this definition via the following example:

EXAMPLE 3. Consider the election $E = (C, V, \mathcal{A})$, with six candidates $C = \{a_1, a_2, a_3, a_4, b, c\}$ and $\mathcal{A} = \{\{a_1, a_2, a_3, a_4\}, \{b\}, \{c\}\}$. Voters' preferences are the following:

30 votes: $b > a_2 > a_4 > a_3 > a_1 > c$
 35 votes: $a_1 > a_2 > a_4 > a_3 > b > c$
 5 votes: $a_2 > a_1 > a_4 > a_3 > b > c$
 30 votes: $a_3 > c > a_4 > b > a_2 > a_1$

Now let us compute IW-Plurality scores. For b , it is 30, for a_1 , a_2 and a_4 it is $35 + 5 = 40$, for a_3 it is $35 + 5 + 30 = 70$, for c it is 0. Candidate a_3 has the highest score, hence in the first round his or her alliance wins. In the second round, the winner is determined by standard Plurality, hence a_1 wins with score 35.

Let us now look at IW-Maximin. In the first round, we only care about head-to-head comparisons between opponents. The alliance-aware maximin score of a_1 and a_2 is $\min(40, 70) = 40$ (comparisons with b and c , respectively). For a_3 it is $\min(70, 100) = 70$, for a_4 it is $\min(70, 70) = 70$. Hence, it is clear that the alliance $A = \{a_1, a_2, a_3, a_4\}$ wins. In the second round, we additionally need to take into account the comparisons between candidates from A , possibly updating their score from the first round. Candidate a_2 wins against each ally: Against a_1 with vote proportion $65 : 35$, and against a_3 and a_4 with vote proportion $70 : 30$. Moreover, we can check that for a_1 , a_3 and a_4 this defeat is among the worst ones. Hence, the score of a_2 remains unchanged (40), the score of a_1 is 35 and the score of a_3 and a_4 is 30. Therefore, a_2 is the winner.

Let us now study the properties of these rules. We show that they both satisfy all the four basic axioms defined in Section 3.1, independent winner consistency, and the most important axiomatic properties of their standard analogues.

THEOREM 1. Let $f \in \{\text{Plurality}, \text{Maximin}\}$. Then IW- f is: (1) basic, (2) independent-winner-consistent, (3) majority-consistent, (4) monotone. Additionally, (5) IW-Maximin is Condorcet-consistent.

PROOF. Regarding (1): First, directly from the definition of f -score and Observation 3 we obtain that removing a candidate from a losing alliance does not affect the score of his or her allies at all, while the score of his or her opponents may only increase. Hence, ally-no-harm is satisfied. Second, splitting a losing alliance A into two alliances A_1, A_2 makes the f -score of every candidate $c \in A_1$ be counted as if he or she was the only candidate from A_1 , not from A —which together with Observation 3 implies that his or her alliance-aware f -score decreases. On the other hand, the alliance-aware f -score of the candidates beyond A remains unaffected. Hence, resistance to alliance-splitting is satisfied. Third, consider two similar candidates a_1, a_2 and their opponent b . Now, the subset of votes in which a_1 is ranked higher than b is the same as the subset of votes in which a_2 is ranked higher than b . It means that after removing a_2 from the election, both the alliance-aware plurality score and the alliance-aware maximin score of b remain unchanged. This fact, together with ally-no-harm, implies independence of similar allies. Fourth, alliance monotonicity follows from the fact that improving a position of a candidate c in a voter's ranking can never decrease neither the alliance-aware f -score of c , nor the alliance-aware f -score of any ally of c , nor the standard f -score of c . (2) follows from the fact that if a candidate c has the highest alliance-aware f -score among all the candidates in the scenario where only he or she runs as an individual candidate then she would also have the highest standard f -score where all the candidates run as individual ones. (3) and (5) follow from the fact that if a candidate c has standard f -score $s > n/2$, then no opponent has greater alliance-aware f -score than $n - s < n/2$. Then, $A(c)$ would be elected in the first round and c would win in the second one since f is majority-consistent (Condorcet-consistent for $f = \text{Maximin}$). For (4) the argument is the same as for alliance monotonicity. \square

4.2 Solitary-Winner-Consistent Rules

Let us now present the extensions of Plurality and Maximin that satisfy solitary winner consistency. We will again present the joint definition for both rules.

DEFINITION 9 (SW- f). Let $f \in \{\text{Plurality}, \text{Maximin}\}$. The SW- f rule has two rounds. In the first round, we compute alliance-aware f -scores for each candidate, and we choose the set of candidates T that obtain higher score than $n/2$. Then in the second round, we eliminate all the candidates from $C \setminus T$ and elect the candidate from T with the highest standard f -score. If $T = \emptyset$, the rule terminates after the first round, returning the candidate with the highest f -score.

Let us illustrate this definition via the same example as before:

EXAMPLE 4. Consider the same election as in Example 3. Let us begin with SW-Plurality. Alliance-aware plurality scores are the same as before, yet now we directly elect a_3 , since he or she is the only candidate with greater score than $n/2$.

For SW-Maximin, the scores in the first round are the same as in Example 3, yet now only a_3 and a_4 advance to the second round, as their score is greater than $n/2$. In the runoff, we compare only a_3 and a_4 head-to-head and declare a_4 the winner (with vote proportion $70 : 30$).

Similarly as in the case of IW-rules, SW-rules preserve the most important axiomatic properties of their standard analogues.

	IW-Plur.	SW-Plur.	IW-Maxi.	SW-Maxi.
Ally-no-harm	✓	✓	✓	✓
Res. to alliance-split.	✓	✓	✓	✓
Ind. of similar allies	✓	✓	✓	✓
Alliance monotonicity	✓	✓	✓	✓
Monotonicity	✓	✓	✓	✓
Majority consistency	✓	✓	✓	✓
Condorcet consistency	✗	✗	✓	✓
IW consistency	✓	✗	✓	✗
SW consistency	✗	✓	✗	✓

Table 1: The axiomatic comparison of alliance-aware rules.

THEOREM 2. *Let $f \in \{\text{Plurality}, \text{Maximin}\}$. Then SW- f is: (1) basic, (2) solitary-winner-consistent, (3) majority-consistent, (4) monotone. Additionally, (5) SW-Maximin is Condorcet-consistent.*

PROOF. For (1) the argumentation is the same as in case of Theorem 1. (2) follows from the characterization of solitary winners presented in Observation 1 and Observation 2. Additionally, for SW-Plurality we note that if $T = \emptyset$ and solitary winner consistency (possibly together with resistance to alliance-splitting) requires electing c , he or she needs to have the highest alliance-aware plurality score. It holds because after removing the allies of c from the election and merging the alliances of his or her opponents, the alliance-aware score of c is the same, and the opponents' score could only increase. (3) and (5) follow from the fact that the majority winner (Condorcet winner, for $f = \text{Maximin}$) would always advance to the second round and win there. (4) follows from the fact that improving a position of a candidate c in a voter's ranking may only improve the alliance-aware f -score of c , does not affect the alliance-aware f -score of his or her allies (hence, does not change the subset of candidates advancing to the second round), and increases the standard f -score of c in the second round. \square

To conclude this section, in Table 1 we present an overview of the axioms satisfied by our IW- and SW- voting rules.

5 EXPERIMENTAL ANALYSIS

In this section, we conduct two experiments, which we treat as sanity checks. In the first one we verify how often primary elections within alliances fail to select an optimal representative. In the second one, we compare alliance-aware rules with the standard ones from the utilitarian perspective.

Since we lack the real-life election data with complete rankings and alliance affiliations, we focus on synthetically generated elections. Below we define two statistical cultures, one based on impartial culture, and the other one based on Euclidean model [9, 19].

Impartial Culture (IC). Under the impartial culture model, both preference orders of the voters and alliance affiliations of the candidates are sampled uniformly at random.

Euclidean. In a d -dimensional Euclidean model, each candidate, voter, and alliance is assigned an ideal point, drawn independently from the uniform distribution on $(0, 1)^d$. Then, the alliance of a given candidate is the closest one (using ℓ_2

distance) to to him or her. Finally, the ranking of each voter is obtained by ordering candidates according to the increasing ℓ_2 distance from the voter's ideal point.

Note that we omit here several other well-known statistical cultures, e.g., the Pólya-Eggenberger urn model [7] or the Mallows model [15]. However, it seems unclear how to properly adapt them to the setting with alliances. For cultures different than impartial culture, choosing alliance membership uniformly at random does not seem appealing, since we expect that the candidates from the same alliance should be somehow "closer" to each other. On the other hand, for cultures different than the Euclidean one, there is no natural notion of "distance" between candidates. We therefore leave the study of other models for future research.

5.1 Non-optimal Primary Winner

In the following experiment we study standard Plurality and Maximin with the following modification: Only one candidate from each alliance takes part in the main election. We assume that representatives of all alliances are chosen via primaries. The primaries may be *joint*, where we select primary winners using all the ballots, or *disjoint*, where for each alliance A we select its representative using only the ballots of supporters of A , i.e., voters who ranked a candidate from A on top.

Now we verify how frequently we witness choosing a non-optimal primary winner, i.e., a situation where there exists an alliance which is losing, but which could have won if another candidate was selected as the primary winner.

In Table 2 we present the results. Each entry is an average over 1000 elections with 10 candidates, 101 voters and 2 alliances. (additional results with 8|10|12 candidates, 2|3 alliances, and including the Schulze's method are available in the full version of our paper [22]). First of all, for disjoint primaries we observe a much higher probability of witnessing a non-optimal primary winner. As expected, the probability of witnessing a non-optimal primary winner is smaller under Maximin, than under Plurality—but still not negligible. The exception are 1D elections, but note that Maximin under joint primaries is Condorcet-consistent and in such elections (given the odd number of voters) the Condorcet winner always exists.

This simulation supports our claim that primaries are indeed not a good solution. Naturally, the exact probability of witnessing a non-optimal primary winner depends on the specific model. However, the very fact is that this value is significant for at least some models shows that this is an actual problem.

5.2 Cost of Using Alliance-Aware Rules

One might argue that using the alliance-aware voting rule, we can significantly decrease the social welfare. In the following experiment, we show that is not the case. Specifically, we compare the social welfare obtained by standard and alliance-aware rules.

Given an election $E = (C, V, \mathcal{A})$, let $c_w \in C$ be the winner of that election. We define the social welfare of candidate c_w as $\sum_{v \in V} (m - 1 - \text{pos}_v(c_w))$, where $\text{pos}_v(c_w)$ denotes the position of candidate c_w in vote v . In other words, the social welfare of a given candidate is equal to his or her Borda score. The social welfare of an election E given voting rule f , is defined as the social welfare of the candidate winning under rule f .

	Rule	IC	Euc 3D	Euc 2D	Euc 1D
spoilers	Plurality (Joint Prim.)	45.6%	10.9%	14.5%	0.3%
	Plurality (Disjoint Prim.)	49.3%	24.2%	37.5%	55.6%
	Maximin (Joint Prim.)	35.3%	1.7%	1.6%	0.0%
	Maximin (Disjoint Prim.)	40.2%	11.3%	17.7%	67.1%
average social welfare	Plurality	480.8	564.0	523.8	474.1
	Plurality (Joint Prim.)	490.3	616.6	595.6	580.5
	Plurality (Disjoint Prim.)	486.5	588.9	551.6	525.8
	IW-Plurality	478.9	568.1	532.1	493.6
	SW-Plurality	483.9	626.1	607.9	580.8
	Maximin	498.5	643.8	631.1	581.6
	Maximin (Joint Prim.)	497.0	643.8	631.1	581.6
	Maximin (Disjoint Prim.)	492.0	628.7	605.0	515.0
	IW-Maximin	497.9	643.9	631.0	581.6
	SW-Maximin	496.4	643.8	631.1	581.6

Table 2: Probability of non-optimal primary winner occurrence under Plurality and Maximin and the average social welfare of our rules compared to their classical analogues with and without primaries. Setup: 1000 elections (per culture) with 10 candidates, 101 voters, and 2 alliances.

In Table 2 we present the results. Fortunately, we see that the social welfare of alliance-aware variants of Maximin rule, provide the same social welfare as standard Maximin. For Plurality, the social welfare of alliance-aware variants is (in most cases) not only not smaller, but even larger than that of standard Plurality. Note that the Plurality with primaries is also providing larger social welfare than the standard version.

6 SUMMARY AND DISCUSSION

In this paper, we have proposed an idea of *alliance-aware voting rules* as a proposal to eliminate the spoiler effect between candidates from the same alliance. We have introduced first axioms that are desirable in this model as well as four alliance-aware voting rules with good axiomatic properties. Our rules completely eliminate the need for alliances to run primary elections, which, as we have shown, often may return candidates that are not optimal if the goal is to win the whole election. We have argued that they provide similar social welfare to the voters as the standard rules.

Let us complete our paper by discussing how our methods could be implemented in real-life elections and proposing one interesting extension of our model.

6.1 Implementation Details

The key factor that we find crucial about our rules is that they are easy to explain and understand. This is in contrast to many rules that are independent of clones (like Ranked Pairs proposed by Tideman [30] or the Schulze’s method). However, one could argue that for IW-Plurality and SW-Plurality, there is still one issue which makes them more complicated (and therefore, less appealing) than standard Plurality: Namely, the complexity of voters’ ballots. For standard Plurality it is enough for every voter to provide only his or her top choice, while our rules, as defined in Section 4, seem to require full rankings.

However, it is actually not the case. For each voter v_i , let us denote by T_i the largest set of candidates from the same alliance A , such that v_i ranks no candidate outside of A higher than a member of T_i . In words, T_i is the set of candidates who form the longest single-alliance prefix of v_i ’s ranking. Now we can see that for IW-Plurality, we in fact only need a voter v_i to provide his or her top choice together with set T_i (without even the need to order it). Indeed, the remaining information is not used by the rule neither in the first round, nor in the second one. This corresponds to a mixture of an approval and a plurality ballot.

For SW-Plurality, the information about the v_i ’s preferences over candidates beyond T_i might be used in the second round. However, we can implement this method in a similar way as it is usually the case for the Plurality with runoff: The two rounds can be implemented as two separate votings. In the first round, each voter casts an approval ballot, approving the candidates from T_i . Then candidates who gain approvals from more than 50% of voters, advance to the runoff (if there is no such candidate, the candidate with the greatest number of approvals wins). In the second round, the voters may vote using standard plurality ballots.

6.2 Extension: Nested Alliances

In our work we have assumed for clarity that alliances do not overlap. However, our axioms and voting rules can be naturally extended to the setting where set \mathcal{A} forms a *laminar family*, i.e., for each two alliances we have that either one of them is a subset of another, or they are disjoint. This may correspond to the situation in which we have e.g., political parties, their coalitions, and factions within them. Intuitively, now candidates a and b might be at the same time allies with respect to one alliance, and opponents with respect to some other (inner) one. The general idea behind our axioms remains the same, yet now their definitions need to include the fact that a candidate c might be a member of multiple alliances.

Our voting rules, instead of two rounds, now could possibly have at most $k + 1$ ones, where k is the maximal nesting depth of an alliance; intuitively, after each round we recount the alliance-aware scores of candidates, skipping the alliances that are supersets of the set of potential election winners. We present the details of our adapted definitions in the full version of our paper [22].

Among other possible extensions of our setting we could mention e.g., the model where alliances overlap arbitrarily or the multi-winner model. Besides, all the topics studied in the literature for classical single-winner voting rules (e.g., distortion [1, 24], resistance to manipulations, etc.) are interesting for alliance-aware rules as well. We leave these open problems for future research.

ACKNOWLEDGEMENTS

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 101002854). Additionally, Grzegorz Pierczyński was supported by Poland’s National Science Center grant no. 2022/45/N/ST6/00271.



REFERENCES

- [1] E. Anshelevich, A. Filos-Ratsikas, N. Shah, and A. A. Voudouris. 2021. Distortion in Social Choice Problems: The First 15 Years and Beyond. In *International Joint Conference on Artificial Intelligence*. <https://api.semanticscholar.org/CorpusID:232075911>
- [2] K. J. Arrow. 1950. A difficulty in the concept of social welfare. *Journal of political economy* 58, 4 (1950), 328–346.
- [3] G. Bordes and T. N. Tideman. 1991. Independence of irrelevant alternatives in the theory of voting. *Theory and Decision* 30 (1991), 163–186.
- [4] C. Börgers. 2010. *Mathematics of social choice: voting, compensation, and division*. SIAM.
- [5] Y. Chevaleyre, J. Lang, N. Maudet, and J. Monnot. 2010. Possible winners when new candidates are added: The case of scoring rules. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 24. 762–767.
- [6] B. Dutta, M. O. Jackson, and M. Le Breton. 2001. Strategic candidacy and voting procedures. *Econometrica* 69, 4 (2001), 1013–1037.
- [7] F. Eggenberger and G. Pólya. 1923. Über die statistik verketteter vorgänge. *ZAMM-Journal of Applied Mathematics and Mechanics/Zeitschrift für Angewandte Mathematik und Mechanik* 3, 4 (1923), 279–289.
- [8] L. Ehlers and J. A. Weymark. 2003. Candidate stability and nonbinary social choice. *Economic Theory* 22 (2003), 233–243.
- [9] J. M. Enelow and M. J. Hinich. 1984. *The Spatial Theory of Voting: An Introduction*. Cambridge University Press, Cambridge, UK.
- [10] G. Erdélyi, M. Neveling, C. Reger, J. Rothe, Y. Yang, and R. Zorn. 2021. Towards completing the puzzle: complexity of control by replacing, adding, and deleting candidates or voters. *Autonomous Agents and Multi-Agent Systems* 35 (2021), 1–48.
- [11] P. Faliszewski, L. Gourvès, J. Lang, J. Lesca, and J. Monnot. 2016. How hard is it for a party to nominate an election winner?. In *Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016*. 257–263.
- [12] Paul Harrenstein, Grzegorz Lisowski, Ramanujan Sridharan, and Paolo Turrini. 2021. A Hotelling-Downs Framework for Party Nominees. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*. 593–601.
- [13] J. Lang, N. Maudet, and M. Polukarov. 2013. New results on equilibria in strategic candidacy. In *Algorithmic Game Theory: 6th International Symposium, SAGT 2013, Aachen, Germany, October 21–23, 2013. Proceedings* 6. Springer, 13–25.
- [14] H. Liu, H. Feng, D. Zhu, and J. Luan. 2009. Parameterized computational complexity of control problems in voting systems. *Theoretical Computer Science* 410, 27–29 (2009), 2746–2753.
- [15] C. L. Mallows. 1957. Non-Null Ranking Models. I. *Biometrika* 44, 1-2 (June 1957), 114–130.
- [16] E. Maskin. 2020. A modified version of Arrow’s IIA condition. *Soc. Choice Welf.* 54, 2-3 (2020), 203–209.
- [17] D. McCune and J. Wilson. 2023. Ranked-choice voting and the spoiler effect. *Public Choice* (2023), 1–32.
- [18] R. Meir, A. D. Procaccia, J. S. Rosenschein, and A. Zohar. 2008. Complexity of strategic behavior in multi-winner elections. *Journal of Artificial Intelligence Research* 33 (2008), 149–178.
- [19] S. Merrill. 1984. A Comparison of Efficiency of Multicandidate Electoral Systems. *American Journal of Political Science* 28 (1984), 23–48.
- [20] M. Neveling, J. Rothe, and R. Zorn. 2020. The complexity of controlling Condorcet, fallback, and k-veto elections by replacing candidates or voters. In *Computer Science—Theory and Applications: 15th International Computer Science Symposium in Russia, CSR 2020, Yekaterinburg, Russia, June 29–July 3, 2020, Proceedings* 15. Springer, 314–327.
- [21] D. K. Osborne. 1976. Irrelevant Alternatives and Social Welfare. *Econometrica* 44, 5 (1976), 1001–1015.
- [22] G. Pierczyński and S. Szufa. 2024. Single-Winner Voting with Alliances: Avoiding the Spoiler Effect. [arXiv:2401.16399 \[cs.GT\]](https://arxiv.org/abs/2401.16399)
- [23] M. Polukarov, S. Obratsova, Z. Rabinovich, A. Kruglyi, and N. R. Jennings. 2015. Convergence to equilibria in strategic candidacy. In *Proceedings of the 24th International Conference on Artificial Intelligence*. 624–630.
- [24] A. D. Procaccia and J. S. Rosenschein. 2006. The distortion of cardinal preferences in voting. In *International Workshop on Cooperative Information Agents*. Springer, 317–331.
- [25] P. Ray. 1973. Independence of irrelevant alternatives. *Econometrica: Journal of the Econometric Society* (1973), 987–991.
- [26] C. Rodríguez-Álvarez. 2006. Candidate stability and probabilistic voting procedures. *Economic Theory* 27 (2006), 657–677.
- [27] C. Rodríguez-Álvarez. 2006. Candidate stability and voting correspondences. *Social Choice and Welfare* 27, 3 (2006), 545–570.
- [28] M. Schulze. 2011. A new monotonic, clone-independent, reversal symmetric, and condorcet-consistent single-winner election method. *Social choice and Welfare* 36 (2011), 267–303.
- [29] A. Sen. 1970. *Collective Choice and Social Welfare*. Harvard University Press.
- [30] T. N. Tideman. 1987. Independence of clones as a criterion for voting rules. *Social Choice and Welfare* 4 (1987), 185–206.