

New Results on Equilibria in Strategic Candidacy

J erome Lang¹, Nicolas Maudet², and Maria Polukarov³

¹ LAMSADE, Universit  Paris-Dauphine, Paris, France
lang@lamsade.dauphine.fr

² LIP6, Universit  Pierre et Marie Curie, Paris, France
nicolas.maudet@lip6.fr

³ University of Southampton, United Kingdom
mp3@ecs.soton.ac.uk

Abstract. We consider a voting setting where candidates have preferences about the outcome of the election and are free to join or leave the election. The corresponding candidacy game, where candidates choose strategically to participate or not, has been studied in very few papers, mainly by Dutta et al. [5,6], who showed that no non-dictatorial voting procedure satisfying unanimity is candidacy-strategyproof, or equivalently, is such that the joint action where all candidates enter the election is always a pure strategy Nash equilibrium. They also showed that for voting trees, there are candidacy games with no pure strategy equilibria. However, no results were known about other voting rules. Here we prove several such results. Some are positive (a pure strategy Nash equilibrium is guaranteed for Copeland and the uncovered set, whichever is the number of candidates, and for all Condorcet-consistent rules, for 4 candidates). Some are negative, namely for plurality and maximin.

1 Introduction

The two main criteria for the evaluation of voting rules are their ability to resist various sorts of strategic behaviour and to adapt to changes in the environment. Many (if not most) papers in computational social choice deal with (at least) one of these issues. Typically, strategic behaviour is shown by the voters reporting insincere votes (*manipulation*); by a third party, usually the chair, acting on the set of voters or candidates (*control*), or on the votes (*bribery* and *lobbying*), or on the voting rule (*e.g.*, *agenda control*)¹; finally, it can arise among the candidates themselves, who may also have preferences about the outcome of the election. However, the latter case has received little attention in (computational) social choice comparing to the former two. One form thereof involves choosing optimal political platforms, but probably the simplest form comes from the very ability of candidates to *decide whether to run for the election or not*, which is the issue we address in this paper. The following table summarises this rough classification of strategic behaviour in voting, according to the identity of strategising agent(s) and also to another relevant dimension, namely what the strategic actions bear on—voters, votes or candidates (we omit the agenda to keep the table small).

¹ There are also some forms of strategic behaviour that are specific to multiwinner elections, such as gerrymandering (by the chair) or vote pairing (by the voters).

actions → agents ↓	<i>voters</i>	<i>votes</i>	<i>candidates</i>
<i>voters</i>	strategic participation	manipulation	-
<i>third party / chair</i>	voter control	bribery, lobbying	candidate control, cloning
<i>candidates</i>	-	-	<i>strategic candidacy</i>

Strategic candidacy does happen frequently in real-life elections, both in large-scale political elections and in small-scale, low-stake elections (*e.g.*, electing a chair in a research group). Throughout the paper we consider a finite set of *potential candidates*, which we simply call *candidates* when this is not ambiguous, and we make the following assumptions:

1. each candidate may choose to run or not for the election;
2. each candidate has a preference ranking over candidates;
3. each candidate ranks himself on top of his ranking;
4. the candidates' preferences are common knowledge among them;
5. the outcome of the election as a function of the set of candidates who choose to run is common knowledge among the candidates.

With the exception of 3, these assumptions were also made in the original model of Dutta et al. [5] which we discuss below. Assumption 2 amounts to saying that a candidate is interested only in the winner of the election² and has no indifferences or incomparabilities. Assumption 3 (considered as optional in [5]) is a natural domain restriction in most contexts. Assumptions 4 and 5 are common game-theoretic assumptions: note that we do not have to assume that the candidates know precisely how voters will vote, nor even the number of voters—they just have to know the choice function mapping every subset of candidates to a winner.

Existing work on strategic candidacy is rather scarce. It starts with [5] and [6], that formulate the strategic candidacy game and prove the following results (among others): (i) no non-dictatorial voting procedure satisfying unanimity is candidacy-strategyproof—or equivalently, is such that the joint action where all candidates enter the election is always a pure strategy Nash equilibrium; (ii) for the specific case of voting trees, there are candidacy games with no pure strategy Nash equilibria. These results are discussed further (together with simpler proofs) [7], and extended to voting correspondences [9,15] and to probabilistic voting rules [14].

Many questions remain unsolved. In particular, studying the solution concepts (such as Nash equilibria or strong equilibria) of a candidacy game would help predict the set of actual candidates and hence, the outcome of the vote, and therefore help design better elections. However, little is known about this: we only know that for any reasonable voting rule, there are some candidacy games for which the set of all candidates is not a Nash equilibrium, and that for voting trees, there exist a candidacy game with no pure strategy Nash equilibrium.

² In some contexts, candidates may have more refined preferences that bear for instance on the number of votes they get, how their score compares to that of other candidates etc. We do not consider these here.

In this paper, we go further in this direction and prove some positive as well as some negative results. We first consider the case of 4 candidates and show that a pure strategy Nash equilibrium always exists for Condorcet-consistent rules. Then we show that for Copeland and the uncovered set there is always an equilibrium in pure strategies, whichever is the number of candidates (although *strong* equilibria are not guaranteed to exist). On the negative side, we show that for plurality, for at least 4 candidates, and for maximin for at least 5 candidates, there are candidacy games without Nash equilibria.

Although it seems that strategic candidacy has not been considered yet in computational social choice, it is related to some questions that have received some attention in this community. First, the existence of strong equilibria is related to a stronger variant of candidate control (see the last paragraph of the conclusion). Other somewhat less related works that also consider a dynamic set of candidates are candidate cloning [8], possible winners with new candidates [3], and the unavailable candidate model [12].

The paper unfolds as follows. In Section 2 we define the strategic candidacy games and give a few preliminary results. In Section 3 we focus on the case of 4 candidates, whereas the case of 5 or more candidates is considered in Section 4. Finally, in Section 5 we discuss further issues, including the relation to candidate control.

2 Model and Preliminaries

In this section, we formally define the model of strategic candidacy and show that it induces a normal form game. We then present two simple results on the existence of Nash equilibria and strong equilibria in this setting.

2.1 Voting Rules

For completeness, we first define the common voting rules that we study in this paper.

There is a set of n voters electing from a set of m candidates. A single vote is a strict ordering of the candidates. A voting rule takes all the votes as input, and produces an outcome—a candidate, called *the winner* of the election. Although voting rules are usually defined for a fixed number of candidates, here we naturally extend the definition to an arbitrary number of candidates. All voting rules we consider in this work are *resolute*: we first define their irresolute version and assume that ties are broken up according to a fixed priority relation over candidates. Since voting rules are applied to varying sets of candidates, we assume that the tie-breaking rule is defined for the whole set of potential candidates, and projected to smaller sets of candidates; in other terms, if x has priority over y when all potential candidates run, this will still be the case for any set of candidates that contains x and y .

The *plurality* winner is the candidate that is ranked first by the largest number of voters. The *Borda* winner is the candidate who gets the highest Borda score: for each voter, a candidate c receives $q - 1$ points (where q is the number of candidates that are actually running) if it is ranked first by that voter, $q - 2$ if it is ranked second, and so on; the Borda score $B(c)$ of c is the total number of points he receives from all the voters.

Let $N(c, x)$ be the number of votes that rank c higher than x . The majority graph associated with a set of votes is the graph whose vertices are the candidates and containing an edge from x to y whenever $N(x, y) > \frac{n}{2}$ (when this holds we say that x

“beats” y). A candidate c is a *Condorcet winner* if x beats y for all $y \neq x$. A voting rule is *Condorcet-consistent* if it always elects a Condorcet winner when one exists.

The *maximin* rule chooses the candidate c for whom $\min_{x \in X \setminus \{c\}} N(c, x)$ is maximal. The Copeland⁰ (resp., Copeland¹) rule elects the candidate c maximising the number of candidates x such that $N(c, x) > \frac{n}{2}$ (resp., $N(c, x) \geq \frac{n}{2}$). The *uncovered set* (UC) rule selects the winner from the “uncovered set of candidates”: a candidate c belongs to the uncovered set if and only if, for any other candidate x , if x beats c then c beats some y that beats x .

2.2 Strategic Candidacy

There is a set $X = \{x_1, x_2, \dots, x_m\}$ of m potential candidates, and a set $V = \{1, 2, \dots, n\}$ of n voters. We assume that these sets of voters and candidates are disjoint. As is classical in social choice theory, each voter $i \in V$ has a *preference* relation P_i , over the different candidates—i.e., a strict order ranking the candidates. The combination $P = (P_1, P_2, \dots, P_n)$ of all the voters’ preferences defines their *preference profile*.

Furthermore, each candidate also has a strict preference ordering over the candidates. We naturally assume that the candidates’ preferences are *self-supported*—that is, the candidates rank themselves at the top of their ordering. Let $P^X = (P_c^X)_{c \in X}$ denote the candidates’ preference profile. Following P^X , the potential candidates may decide to enter an election or withdraw their candidacy. Thus, the voters will only express their preferences over a subset $Y \subseteq X$ of the candidates that will have chosen to participate in the election, and we denote by $P^{\downarrow Y}$ the restriction of P to Y . We assume that the voters are *sincere*.

Given a profile P of the voters’ preferences, a *voting rule* r defines a (single) winner among the *actual* candidates—i.e., given a subset $Y \subseteq X$ of candidates, it assigns to a (restricted) profile $P^{\downarrow Y}$ a member of Y . Each such voting rule r induces a natural *game form*, where the set of players is given by the set of potential candidates X , and the strategy set available to each player is $\{0, 1\}$ with 1 corresponding to entering the election and 0 standing for withdrawal of candidacy. A *state* s of the game is a vector of strategies $(s_c)_{c \in X}$, where $s_c \in \{0, 1\}$. For convenience, we use s_{-z} to denote $(s_c)_{c \in X \setminus \{z\}}$ —i.e., s reduced by the single entry of player z . Similarly, for a state s we use s_Z to denote the strategy choices of a coalition $Z \subseteq X$ and s_{-Z} for the complement, and we write $s = (s_Z, s_{-Z})$.

The outcome of a state s is $r(P^{\downarrow Y})$ where $c \in Y$ if and only if $s_c = 1$.³ Coupled with a profile P^X of the candidates’ preferences, this defines a normal form game $\Gamma = \langle X, P, r, P^X \rangle$ with m players. Here, player c prefers outcome $\Gamma(s)$ over outcome $\Gamma(s')$ if ordering P_c^X ranks $\Gamma(s)$ higher than $\Gamma(s')$.

2.3 Game-Theoretic Concepts

Having defined a normal form game, we can now apply standard game-theoretic solution concepts. Let $\Gamma = \langle X, P, r, P^X \rangle$ be a candidacy game, and let s be a state in Γ .

³ When clear from the context, we use vector s to also denote the set of candidates Y that corresponds to state s ; e.g., if $X = \{x_1, x_2, x_3\}$, we write $\{x_1, x_3\}$ and $(1, 0, 1)$ interchangeably.

We say that a coalition $Z \subseteq X$ has an *improving move* in s if there is s'_Z such that $\Gamma(s_{-Z}, s'_Z)$ is preferable over $\Gamma(s)$ by every player $z \in Z$. In particular, the improving move is *unilateral* if $|Z| = 1$. A (*pure strategy*) *Nash equilibrium (NE)* [13] is a state that has no unilateral improving moves. More generally, a state is a k -NE if no coalition with $|Z| \leq k$ has an improving move. A *strong equilibrium (SE)* ([1]) is a state that has no improving moves.

Example 1. Consider the game $\langle \{a, b, c, d\}, P, r, P^X \rangle$, where r is the Borda rule, and P and P^X are as follows⁴:

P	P^X
1 1 1 1 1 1 1	a b c d
b c c a d b a	a b c d
d d d c a c b	d a b a
a a b b c d c	b d a c
c b a d b a d	c c d b

The state (1,1,1,1) is not an NE: $abcd \mapsto c$, but $abc \mapsto a$, and d prefers a to c , so for d , leaving is an improving move. Now, (1,1,1,0) is an NE, as nobody has an improving move neither by joining (d prefers a over c), nor by leaving (obviously not a ; if b or c leaves then the winner is still a). It can be checked that this is also an SE.

2.4 Preliminary Results

Regardless of the number of voters and the voting rule, a straightforward observation is that a candidacy game with *three* candidates is guaranteed to possess an NE.⁵ This, however, is not true for SE.⁶ For any number of candidates, the following result holds.

Proposition 1. *Let $\Gamma = \langle X, P, r, P^X \rangle$ be a candidacy game where r is Condorcet-consistent. If P has a Condorcet winner c then for any $Y \subseteq X$,*

$$Y \text{ is a SE} \Leftrightarrow Y \text{ is an NE} \Leftrightarrow c \in Y.$$

Proof. Assume c is a Condorcet winner for P and let $Y \subseteq X$ such that $c \in Y$. Because r is Condorcet-consistent, and because c is a Condorcet winner for $P^{\downarrow Y}$, we have $r(P^{\downarrow Y}) = c$. Assume $Z = Z^+ \cup Z^-$ is a deviating coalition from Y , with Z^+ the candidates who join and Z^- the candidates who leave the election. Clearly, $c \notin Z$, as $c \in Y$ and c has no interest to leave. Therefore, c is still a Condorcet winner in $P^{\downarrow(Y \setminus Z^-) \cup Z^+}$,

⁴ In our examples, we assume a lexicographic tie-breaking. We also use the simplified notation $Y \mapsto x$ to denote that rule r applied to the subset of candidates $Y \subseteq X$ is x , and we omit curly brackets. The first row in P indicates the number of voters casting the different ballots.

⁵ This can be easily seen: Let $X = \{a, b, c\}$ and suppose w.l.o.g. that $abc \mapsto a$. If $\{a, b, c\}$ is not an NE, then either (1) $ab \mapsto b$ and c prefers b to a , or (2) $ac \mapsto c$ and b prefers c to a . Since b and c play symmetric roles, w.l.o.g., assume (1). Then $\{a, b\}$ is an NE.

⁶ Here is a counterexample (for which we thank an anonymous reviewer of the previous version of the paper). The selection rule is $abc \mapsto b$; $ab \mapsto a$; $ac \mapsto c$; $bc \mapsto c$; it can be easily implemented by the scoring rule with scoring vector (5, 4, 0) with 5 voters. Preferences of candidates are: $a : a \succ b \succ c$; $b : b \succ c \succ a$; $c : c \succ a \succ b$. The group deviations are: in $\{a, b, c\}$, c leaves; in $\{a, b\}$, b leaves and c joins; in $\{a, c\}$, b joins; in $\{b, c\}$, a joins; in $\{a\}$, c joins; in $\{b\}$, c joins; in $\{c\}$, a and b join.

which by the Condorcet-consistency of r implies that $r\left(P^{\downarrow(Y\setminus Z^-)\cup Z^+}\right) = c$, which contradicts the assumption that Z wants to deviate. We thus conclude that Y is an SE, and *a fortiori* an NE. Finally, let $Y \subseteq X$ such that $c \notin Y$. Then, Y is not an NE (and *a fortiori* not an SE), because c has an interest to join the election. \square

Now, if P has no Condorcet winner, the analysis becomes more complicated. We provide results for this more general case in the following sections. Interestingly, as we demonstrate, some Condorcet-consistent rules (e.g., Copeland and UC) do always possess a Nash equilibrium in this case, while some other (e.g., maximin) do not.

3 Four Candidates

With only 4 potential candidates, we exhibit a sharp contrast between Condorcet consistent rules, which all possess an NE, and scoring rules.

3.1 Scoring Rules

To study scoring rules, we make use of a very powerful result by Saari [16]. It states that for almost all scoring rules, any conceivable choice function can result from a voting profile. This means that our question boils down to checking whether a *choice function*, together with some coherent candidates' preferences, can be found such that no NE exists with 4 candidates. We solve this question by encoding the problem as an Integer Linear Program (ILP), the details of which can be found in Appendix. It turns out that such choice functions do exist: it then follows from Saari's result that counterexamples can be obtained for "most" scoring rules. We exhibit a profile for plurality.

Proposition 2. *For plurality and $m = 4$, there may be no NE.*

Proof. We exhibit a counterexample with 13 voters, whose preferences are contained in the left part of the table below. The top line indicates the number of voters with each particular profile. The right part of the table represents the preferences of the candidates.

3	1	1	1	1	1	1	2	2	a	b	c	d
d	d	d	a	a	a	b	b	c	a	b	c	d
c	b	a	b	c	d	c	a	b	b	a	d	a
a	c	b	c	b	b	d	c	d	c	c	a	b
b	a	c	d	d	c	a	d	a	d	d	b	c

\square

Similar constructions of profiles can thus be obtained for other scoring rules. However, Borda comes out as a very peculiar case [16] among scoring rules⁷. This is also verified for the case of strategic candidacy.

⁷ For a more detailed statement of this result, we point the reader to the work of Saari, in particular [17]. For the case of 4 candidates, families of scoring rules such that, when the scoring vector for 3 candidates is of the form $\langle w_1, w_2, 0 \rangle$, the vector for 4 candidates is of the form $\langle 3w_1, w_1 + 2w_2, 2w_2, 0 \rangle$ (for instance, $\langle \langle 3, 1, 0, 0 \rangle, \langle 1, 0, 0, 1, 0 \rangle \rangle$) are an exception in the sense that not all choice functions are implementable with them.

Proposition 3. *For Borda and $m = 4$, there is always an NE.*

We could check this by relying on the fact that Borda rule is represented by a weighted majority graph, and by adding the corresponding constraints into the ILP. The infeasibility of the resulting set of constraints shows that no instances without NE can be constructed. However, it takes only coalitions of pairs of agents to ruin this stability.

Proposition 4. *For Borda and $m = 4$, there may be no 2-NE.*

Proof. Consider the following game:

1	1	1	1		a	b	c	d
b	c	d	a		a	b	c	d
d	d	a	b		c	a	a	b
c	a	c	c		d	c	d	a
a	b	b	d		b	d	b	c


Here, only $s_1 = (0, 1, 1, 1)$ and $s_2 = (1, 1, 0, 1)$ are NE, with $bcd \mapsto b$, and $abd \mapsto d$. But from s_1 the coalition $\{a, c\}$ has an improving move to s_2 as they both prefer d over b . Now take s_2 : if b leaves and c joins, they reach $(1, 0, 1, 1)$, with $acd \mapsto c$ and both prefer c over d . □

3.2 Condorcet-Consistent Rules

We now turn our attention to Condorcet-consistent rules. It turns out that for *all of them*, the existence of an NE can be guaranteed.

Proposition 5. *For $m = 4$, if r is Condorcet-consistent, there always exists an NE.*

Proof. We start with a remark: although we do not assume that r is based on the majority graph, we nevertheless prove our result by considering all possible cases for the majority graphs (we get back to this point at the end of the proof). There are four graphs to consider (all others are obtained from these ones by symmetry).



For G_1 and G_2 , any subset of X containing the Condorcet winner is an NE (see Proposition 1). For G_3 , we note that a is a Condorcet loser. That is, $N(a, x) < N(x, a)$ for all $x \in \{b, c, d\}$. Note that in this case, there is no Condorcet winner in the reduced profile $P \downarrow \{b, c, d\}$ as this would imply the existence of a Condorcet winner in P (case G_1 or G_2). W.l.o.g., assume that b beats c , c beats d , and d beats b . W.l.o.g. again, assume that $bcd \mapsto b$. Then, $\{b, c\}$ is an NE. Indeed, in any set of just two candidates, none has an incentive to leave. Now, a or d have no incentive to join as this would not change the winner: in the former case, observe that b is the (unique) Condorcet winner in $P \downarrow \{a, b, c\}$, and the latter follows by our assumption. There is always an NE for G_3 .

The proof for G_4 is more complex and proceeds case by case. Since r is Condorcet-consistent, we have $acd \mapsto a$, $bcd \mapsto c$, $ab \mapsto b$, $ac \mapsto a$, $ad \mapsto a$, $bc \mapsto c$, $bd \mapsto d$ and $cd \mapsto c$. The sets of candidates for which r is undetermined are $abcd$, abc and abd .

We observe the following easy facts: (i) if $abcd \mapsto a$ then acd is an NE, (ii) if $abcd \mapsto c$ then bcd is an NE, (iii) if $abc \mapsto a$ then ac is an NE, (iv) if $abd \mapsto a$ then ad is an NE, (v) if $abc \mapsto c$ then bc is an NE. The only remaining cases are:

1. $abcd \mapsto b$, $abc \mapsto b$, $abd \mapsto b$.
2. $abcd \mapsto b$, $abc \mapsto b$, $abd \mapsto d$.
3. $abcd \mapsto d$, $abc \mapsto b$, $abd \mapsto b$.
4. $abcd \mapsto d$, $abc \mapsto b$, $abd \mapsto d$.

In cases 1 and 3, ab is an NE. In case 2, if a prefers b to c then abc is an NE, and if a prefers c to b , then bcd is an NE. In case 4, if a prefers c to d , then bcd is an NE; if b prefers a to d , then ad is an NE; finally, if a prefers d to c and b prefers d to a , then $abcd$ is an NE. To conclude, observe that the proof never uses the fact that two profiles having the same majority graph have the same winner.⁸ \square

The picture for 4 candidates shows a sharp contrast. On the one hand, the existence of choice functions shows that “almost all scoring rules” [16] may fail to have an NE. On the other hand, Condorcet-consistency alone suffices to guarantee the existence of an NE. (However, this criterion is not sufficient to guarantee stronger notion of stability: e.g., for Copeland, we could exhibit examples without any 2-NE.)

4 More Candidates

The first question which comes to mind is whether examples showing the absence of NE transfer to larger sets of candidates. They indeed do, under an extremely mild assumption. We say that a voting rule is *insensitive to bottom-ranked candidates* (IBC) if given any profile P over $X = \{x_1, \dots, x_m\}$, if P' is the profile over $X \cup \{x_{m+1}\}$ obtained by adding x_{m+1} at the bottom of every vote of P , then $r(P') = r(P)$. This property is extremely weak (much weaker than Pareto) and is satisfied by almost all voting rules studied in the literature (a noticeable exception being the veto rule).

Lemma 1. *For any voting rule r satisfying IBC, if there exists $\Gamma = \langle X, P, r, P^X \rangle$ with no NE, then there exists $\Gamma' = \langle X', P', r, P^Y \rangle$ with no NE, where $|X'| = |X| + 1$.*

Proof. Take Γ with no NE, with $X = \{x_1, \dots, x_m\}$. Let $X' = X \cup \{x_{m+1}\}$, P' the profile obtained from P by adding x_{m+1} at the bottom of every vote, and $P^{X'}$ be the candidate profile obtained by adding x_{m+1} at the bottom of every ranking of a candidate x_i , $i \leq m$, and whatever ranking for x_{m+1} . Let $Y \subseteq X$. Because Y is not an NE for Γ , some candidate $x_i \in X$ has an interest to leave or to join, therefore Y is not an NE either for Γ' . Now, consider $Y' = Y \cup \{x_{m+1}\}$. If $x_i \in X$ has an interest to leave (resp., join) Y , then because r satisfies IBC, the winner in $Y' \setminus \{x_i\}$ (resp., $Y' \cup \{x_i\}$) is the same as in $Y \setminus \{x_i\}$ (resp., $Y \cup \{x_i\}$), therefore $x_i \in X$ has an interest to leave (resp., join) Y' , therefore Y' is not an NE. \square

⁸ For instance, we may have two profiles P, P' both corresponding to G_4 , such that $r(P) = a$ and $r(P') = b$; the proof perfectly works in such a case.

Corollary 1. *For plurality and $m \geq 4$, there may be no NE.*

We now turn our attention to Condorcet-consistent rules, which all admit NE with 4 candidates. However, 5 candidates suffice to show that NE are not guaranteed any longer.

Proposition 6. *For maximin with $m = 5$, there may be no NE.*

Proof. The counterexample is given by the following pairwise comparison matrix, where the entry corresponding to row x and column y is equal to $N(x, y) - N(y, x)$. From Debord’s theorem [4] we know that there exists a profile with such a comparison matrix. The candidates’ preference profile is given on the right hand side. The tie-breaking priority is lexicographic.

	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>
<i>a</i>	0	-3	3	-1	1
<i>b</i>	3	0	-3	3	1
<i>c</i>	-3	3	0	-1	-1
<i>d</i>	1	-3	1	0	-5
<i>e</i>	-1	-1	1	5	0

<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>
<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>
<i>c</i>	<i>e</i>	<i>d</i>	<i>a</i>	<i>b</i>
<i>b</i>	<i>c</i>	<i>a</i>	<i>c</i>	<i>a</i>
<i>e</i>	<i>a</i>	<i>e</i>	<i>b</i>	<i>d</i>
<i>d</i>	<i>d</i>	<i>b</i>	<i>e</i>	<i>c</i>

The proof goes by exhibiting all cases. For each subset we indicate the deviation (the winner being shown using bold font): $\mathbf{ad} \rightarrow \mathbf{abd} \rightarrow \mathbf{abcd} \rightarrow \mathbf{abcde} \rightarrow \mathbf{bcde} \rightarrow \mathbf{cde} \rightarrow \mathbf{acde} \rightarrow \mathbf{abcde}$; $\mathbf{ab} \rightarrow \mathbf{abc} \rightarrow \mathbf{abce} \rightarrow \mathbf{bce} \rightarrow \mathbf{ce} \rightarrow \mathbf{ace} \rightarrow \mathbf{abce}$; $\mathbf{abde} \rightarrow \mathbf{bce}$; $\mathbf{ae} \rightarrow \mathbf{abe} \rightarrow \mathbf{abce}$; $\mathbf{ac} \rightarrow \mathbf{acd} \rightarrow \mathbf{abcd}$; $\mathbf{bd} \rightarrow \mathbf{bcd} \rightarrow \mathbf{abcd}$; $\mathbf{bde} \rightarrow \mathbf{bcde}$; $\mathbf{bc} \rightarrow \mathbf{abc}$; $\mathbf{be} \rightarrow \mathbf{bce}$; $\mathbf{cd} \rightarrow \mathbf{cde}$; $\mathbf{ce} \rightarrow \mathbf{ace}$; $\mathbf{de} \rightarrow \mathbf{ade}$. □

Corollary 2. *For maximin and $m \geq 5$, there may be no NE.*

This negative result does not extend to all Condorcet-consistent rules. In particular, next we show the existence of NE for Copeland and the uncovered set (UC), under deterministic tie-breaking, for any number of candidates.

Proposition 7. *For Copeland⁰, with any number of candidates, there is always an NE.*

Proof. Let P be a profile and \rightarrow_P its associated majority graph. Let $C(x, P)$ be the number of candidates $y \neq x$ such that $x \rightarrow_P y$. Let $COP^0(P)$ be the set of the Copeland⁰ cowinners for P , i.e., the set of candidates maximising $C(\cdot, P)$, and $Cop^0(P) = c$ the Copeland⁰ winner—the highest-priority candidate in $COP^0(P)$. Consider $Dom(c) = \{c\} \cup \{y | c \rightarrow_P y\}$. Note that $C(c, P \downarrow Dom(c)) = |Dom(c)| - 1 = q = C(c, P)$. Also, since any $y \in Dom(c)$ is beaten by c , we have $C(y, P \downarrow Dom(c)) \leq q - 1$.

We claim that $Dom(c)$ is an NE. Note that c is a Condorcet winner in the restriction of P to $Dom(c)$, and *a fortiori*, in the restriction of P to any subset of $Dom(c)$. Hence, c is the Copeland⁰ winner in $Dom(c)$ and any of its subsets, and no candidate in $Dom(c)$ has an incentive to leave.

Now, assume there is a candidate $z \in X \setminus Dom(c)$ such that $Cop^0(P \downarrow Dom(c) \cup \{z\}) \neq c$. Note that $c \not\rightarrow_P z$ as z does not belong to

$Dom(c)$; so, $C(c, P \downarrow_{Dom(c) \cup \{z\}}) = q$. For any $y \in Dom(c)$ we have $C(y, P \downarrow_{Dom(c) \cup \{z\}}) \leq (q - 1) + 1 = q = C(c, P \downarrow_{Dom(c) \cup \{z\}})$. If $C(y, P \downarrow_{Dom(c) \cup \{z\}}) < C(c, P \downarrow_{Dom(c) \cup \{z\}})$, then y is not the Copeland⁰ winner in $P \downarrow_{Dom(c) \cup \{z\}}$. If $C(y, P \downarrow_{Dom(c) \cup \{z\}}) = C(c, P \downarrow_{Dom(c) \cup \{z\}})$, then $C(y, P) \geq C(c, P)$. That is, either $c \notin COP^0(P)$, a contradiction, or both y, c are in $COP^0(P)$. In that case, the tie-breaking priority ensures that $Cop^0(P \downarrow_{Dom(c) \cup \{z\}}) \neq y$.

Hence, $Cop^0(P \downarrow_{Dom(c) \cup \{z\}}) = z$. By $Cop^0(P) = c$ we have $C(z, P \downarrow_{Dom(c) \cup \{z\}}) \leq C(z, P) \leq C(c, P) \leq q$; therefore, $C(z, P \downarrow_{Dom(c) \cup \{z\}}) = q$, and the tie-breaking priority favours z over c . But then, $C(z, P) = C(c, P)$, i.e., both c and z are in $COP^0(P)$, and the tie-breaking priority ensures that $Cop^0(P \downarrow_{Dom(c) \cup \{z\}}) \neq z$, a contradiction. Therefore, the Copeland⁰ winner in $P \downarrow_{Dom(c) \cup \{z\}}$ must be c , which implies that z has no incentive to join $Dom(c)$. \square

Note that if the number of voters is odd, we do not have to care about head-to-head ties. In this case, all Copeland ^{α} rules, where each agent in a head-to-head election gets $0 \geq \alpha \geq 1$ points in the case of a tie (Copeland⁰ being a special cases), are equivalent, and the result above holds. However, if the number of voters is even, this is not necessarily the case. Thus, in particular, for Copeland^{0.5} (more often referred to as Copeland), $Dom(c)$ is generally no more an NE, and we do not know whether the existence of an NE is guaranteed or not.

Proposition 8. *For UC, with any number of candidates, and an odd number of voters, there is always an NE.*

Proof. Let c be the UC winner in P , i.e., the highest-priority candidate in $UC(P)$. Consider (again) $Dom(c) = \{c\} \cup \{y | c \rightarrow_P y\}$. We claim that $Dom(c)$ is an NE.

Since c is a Condorcet winner in the restriction of P to $Dom(c)$, and *a fortiori*, in the restriction of P to any subset of $Dom(c)$, it is the UC winner in $Dom(c)$ and in any of its subsets, and no candidate in $Dom(c)$ wants to leave.

Now, let $z \in X \setminus Dom(c)$. Since $z \notin Dom(c)$, we have $c \not\rightarrow_P z$ and hence, $z \rightarrow_P c$, as n is odd. Since $x \in UC(P)$, there must be $y \in Dom(c)$ such that $y \rightarrow_P z$. This implies that $x \in UC(P \downarrow_{Dom(c) \cup \{z\}})$, which, due to tie-breaking priority, yields that c is the UC winner in $P \downarrow_{Dom(c) \cup \{z\}}$. Thus, z has no incentive to join $Dom(c)$. \square

Note that the proofs of Propositions 7 and 8 also show that for Copeland⁰ and UC, there always exists an NE in which the winning candidate is the winner in the full profile (with all candidates present)⁹.

5 Conclusions

In this work, we further explored the landscape of strategic candidacy in elections and obtained several positive results (for Condorcet-consistent rules with 4 candidates; for two versions of Copeland, as well as for the uncovered set, with any number of candidates) and several negative results (for plurality and maximin). Many cases remain open, especially Borda with more than 4 candidates.

⁹ We thank Edith Elkind for this remark.

Another line for further research is the study of the set of states that can be reached by some improvement path (e.g., best or better response dynamics) starting, say, from the set or all potential candidates. In some cases, even when the existence of NE is guaranteed (e.g., for Copeland), we could already come up with examples showing that no equilibrium point is reachable by a sequence of better responses. But other types of dynamics can also be considered.

Finally, there is an interesting connection between strategic candidacy and control by deleting or adding candidates [2,11], as well as *multimode* control [10] where the chair is allowed both to delete *and* to add some candidates. Strategic candidacy relates to a slightly more demanding notion of control, which we can call *consenting control*, in which candidates have to agree to be added or removed. For instance, s is an SE if there is no consenting destructive control by removing+adding candidates against the current winner $r(X_s)$. Not only this notion is of independent interest, but also, complexity results for control may allow to derive complexity results for the problem of deciding the existence of NE or SE in a strategic candidacy game.

Acknowledgements. We would like to thank Michel Le Breton, Vincent Merlin, Edith Elkind and the anonymous reviewers for helpful discussions.

References

1. Aumann, R.: Acceptable points in general cooperative n-person games. In: Contributions to the Theory of Games IV. Annals of Mathematics Study, vol. 40, pp. 287–324 (1959)
2. Bartholdi, J., Tovey, C., Trick, M.: How hard is it to control an election? Social Choice and Welfare 16(8-9), 27–40 (1992)
3. Chevaleyre, Y., Lang, J., Maudet, N., Monnot, J., Xia, L.: New candidates welcome! possible winners with respect to the addition of new candidates. Mathematical Social Sciences 64(1), 74–88 (2012)
4. Debord, B.: Caractérisation des matrices de préférences nettes et méthodes d’agrégation associées. Mathématiques et Sciences Humaines 97, 5–17 (1987)
5. Dutta, B., Breton, M.L., Jackson, M.O.: Strategic candidacy and voting procedures. Econometrica 69, 1013–1037 (2001)
6. Dutta, B., Breton, M.L., Jackson, M.O.: Voting by successive elimination and strategic candidacy in committees. Journal of Economic Theory 103, 190–218 (2002)
7. Ehlers, L., Weymark, J.A.: Candidate stability and nonbinary social choice. Economic Theory 22(2), 233–243 (2003)
8. Elkind, E., Faliszewski, P., Slinko, A.M.: Cloning in elections: Finding the possible winners. J. Artif. Intell. Res (JAIR) 42, 529–573 (2011)
9. Eraslan, H., McLennan, A.: Strategic candidacy for multivalued voting procedures. Journal of Economic Theory 117(1), 29–54 (2004)
10. Faliszewski, P., Hemaspaandra, E., Hemaspaandra, L.: Multimode control attacks on elections. JAIR 40, 305–351 (2011)
11. Hemaspaandra, E., Hemaspaandra, L., Rothe, J.: Anyone but him: The complexity of precluding an alternative. Artificial Intelligence 171(5-6), 255–285 (2007)
12. Lu, T., Boutillier, C.: The unavailable candidate model: a decision-theoretic view of social choice. In: ACM Conference on Electronic Commerce 2010, pp. 263–274 (2010)
13. Nash, J.: Non-cooperative games. Annals of Mathematics 54(2), 286–295 (1951)

14. Rodriguez-Alvarez, C.: Candidate stability and probabilistic voting procedures. *Economic Theory* 27(3), 657–677 (2006)
15. Rodriguez-Alvarez, C.: Candidate stability and voting correspondences. *Social Choice and Welfare* 27(3), 545–570 (2006)
16. Saari, D.: A dictionary of voting paradoxes. *Journal of Economic Theory* 48 (1989)
17. Saari, D.: Election results and a partial ordering for positional ordering. In: Schofield, N.J. (ed.) *Collective Decision-Making: Social Choice and Political Economy*, pp. 93–110. Kluwer (1996)

Appendix: ILP Formulation

Let S be the set of ($2^{|X|}$) states, and $A(s)$ be the set of agents candidating in state s .

Choice functions without NE. We introduce a binary variable $w_{s,i}$, indicating that agent i wins in state s . We add constraints enforcing that a winner in each state s is unique:

$$\forall i \in X, \forall s \in S : w_{s,i} \in \{0, 1\} \quad (1)$$

$$\forall s \in S : \sum_{i \in X} w_{s,i} = 1 \quad (2)$$

$$\forall s \in S, \forall i \in X \notin A(s) : w_{s,i} = 0 \quad (3)$$

Now we denote by $D(s)$ the set of *possible deviations* from state s (states where a single agent's candidacy differs from s). We denote by $a(s, t)$ an agent potentially deviating from s to t . Binary variables $d_{s,t}$ indicate a deviation from s to t . In each state, there must be at least one deviation, otherwise this state is an NE:

$$\forall s \in S, \forall t \in S : d_{s,t} \in \{0, 1\} \quad (4)$$

$$\forall s \in S : \sum_{t \in D(s)} d_{s,t} \geq 1 \quad (5)$$

Now, we introduce constraints related to the preferences of the candidates. For this purpose, we introduce a binary variable $p_{i,j,k}$, indicating that agent i prefers candidate j over candidate k . If there is indeed a deviation from s to t , the deviating agent must prefer the winner of the new state over the winner of the previous state:

$$\forall s \in S, \forall t \in D(s), \forall i \in X, \forall j \in X : w_{s,i} + w_{t,j} + d_{s,t} - p_{a(s,t),j,i} \leq 2 \quad (6)$$

Finally we ensure that the preferences are irreflexive and transitive¹⁰, and respect the constraint of being self-supported:

$$\forall i \in X, \forall j \in X : p_{i,j,j} = 0 \quad (7)$$

$$\forall a \in X, \forall i \in X \forall j \in X, \forall k \in X : p_{a,i,j} + p_{a,j,k} - p_{a,i,k} \leq 1 \quad (8)$$

$$\forall i \in X, \forall j \in X : p_{i,i,j} = 1 \quad (9)$$

¹⁰ Notice that this ILP does not necessarily contain complete preferences: the program only needs to check those preference relations that correspond to possible deviations. Any linear extension of these (partial) preferences gives an instance with complete preferences.

Constraints for Borda. We introduce a new integer variable $N_{i,j}$ to represent the number of voters preferring i over j in the weighted tournament. We first make sure that the values of $N_{i,j}$ are coherent throughout the weighted tournament:

$$\forall i \in X, \forall j \in X, \forall k \in X, \forall l \in X : N_{i,j} + N_{j,i} = N_{k,l} + N_{l,k} \quad (10)$$

In each state, when agent i wins, we must make sure that his total amount of points is the highest among all the agents in this state (note that i can simply tie with those agents that i is prioritised over by the tie-breaking; we omit this for the sake of readability):

$$\begin{aligned} & \forall s \in S, \forall i \in A(s), \forall j \in A(s) \setminus \{i\} : \\ & (1 - w_{s,i}) \times M + \sum_{j \in A(s) \setminus \{i\}} N_{i,j} > \sum_{j \in A(s) \setminus \{k\}} N_{k,j} \end{aligned} \quad (11)$$

Here M is an arbitrary large value, used to relax the constraint when $w_{s,i}$ is 0.