Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)

Strategic Candidacy Games with Lazy Candidates

Svetlana Obraztsova
Tel Aviv University
Tel Aviv, Israel

Edith Elkind
University of Oxford
Oxford, United Kingdom

Maria Polukarov
University of Southampton
Southampton, United Kingdom

Zinovi Rabinovich
Mobileye Vision Technologies Ltd.
Israel

Abstract

In strategic candidacy games, both voters and candidates have preferences over the set of candidates, and candidates may strategically withdraw from the election in order to manipulate the outcome according to their preferences. In this work, we extend the standard model of strategic candidacy games by observing that candidates may find it costly to run an electoral campaign and may therefore prefer to withdraw if their presence has no effect on the election outcome. We study the Nash equilibria and outcomes of natural best-response dynamics in the resulting class of games, both from a normative and from a computational perspective, and compare them with the Nash equilibria of the standard model.

1 Introduction

In his paper "Independence of clones as a criterion for voting rules" [1987] Nicolaus Tideman tells the following story:

> When I was 12 years old I was nominated to be treasurer of my class at school. A girl named Michelle was also nominated. I relished the prospect of being treasurer, so I made a quick calculation and nominated Michelle's best friend, Charlotte. In the ensuing election I received 13 votes, Michelle received 12, and Charlotte received 11, so I became treasurer.

Tideman uses this story to motivate the concept of *clones*—candidates that are so similar to each other that all voters rank them consecutively—and considers the problem of designing clone-proof voting rules. However, one can also look at this story from a somewhat different perspective: It is plausible that both Michelle and Charlotte preferred either of them becoming the treasurer to the outcome where both of them lost. Thus, either girl could have decided to withdraw her candidacy to obtain a more desirable election result.

Now, while a typical teenager may find it difficult to make a calculated choice in such situations (Tideman's younger self being an obvious exception), candidates in high-stakes elections often have strong preferences over possible election outcomes as well as strategic reasoning skills, and may therefore consider withdrawing from an election so as to change its winner to one they prefer to the current one. The resulting interaction among the candidates can be described as a non-cooperative game, where players are the candidates, each candidate can choose whether to participate or not, the voters (who are typically *not* considered to be strategic players—see, however, [Brill and Conitzer, 2015], where the model is extended to strategic voters) choose among the available candidates, and each player compares possible outcomes according to his ranking of the candidates. Such games are known as *strategic candidacy games*; they were introduced by Dutta, Le Breton and Jackson [2001; 2002], and have subsequently been studied by several other authors [Ehlers and Weymark, 2003; Eraslan and McLennan, 2004; Rodriguez-Alvarez, 2006a; 2006b; Lang *et al.*, 2013; Polukarov *et al.*, 2015].

In the model of Dutta et al., which was also adopted in subsequent work on strategic candidacy games, it is assumed that each player is indifferent between any two strategy profiles that result in the same election outcome. However, in practice candidates often have to take into account the monetary and reputational costs of running an electoral campaign. In particular, it is often natural to assume that candidates are "lazy", i.e., they prefer not to participate in the election if doing so has no impact on who becomes the winner. This type of secondary preferences has been recently considered in the context of analyzing the behavior of strategic *voters*, i.e., voters were assumed to abstain when they were not pivotal (see, e.g., [Desmedt and Elkind, 2010; Elkind *et al.*, 2014] and references therein). However, to the best of our knowledge, strategic candidacy games with lazy candidates have not been investigated in prior work.

The goal of this paper is to initiate the study of this class of games. For simplicity, we focus on the case where the voters make their choice among the available candidates using the Plurality rule. We consider Nash equilibria of such games as well as analyze two natural best-response dynamics for this setting: one where initially the set of active candidates is empty and the candidates join one by one, and one where initially all candidates are present, but then withdraw one by one. We relate the properties of strategic candidacy games with lazy candidates to those of "vanilla" strategic candidacy games, explore the role of Pareto dominance and Condorcet winners in this setting, and prove a number of computational complexity results. In particular, we show that checking whether a given strategic candidacy game with lazy candidates admits a Nash equilibrium is NP-complete.

2 Preliminaries and Model

We consider settings where there is a finite set of *potential candidates* $C = \{c_1, c_2, \dots, c_m\}$ and a finite set of *voters* $V = \{v_1, v_2, \dots, v_n\}$; we assume $C \cap V = \emptyset$. Given a set $A \subseteq C$, let $\mathcal{L}(A)$ denote the set of all linear orders over A. Each voter $v \in V$ is associated with a *preference order* $\succ_v \in \mathcal{L}(C)$; the list $P^V = (\succ_v)_{v \in V}$ is referred to as the *voters' preference profile*.

An election proceeds as follows. First, a subset of candidates $A \subseteq C$ announces that they will participate in the election; we refer to the candidates in A as the *actual candidates*. If $A = \emptyset$, a pre-determined candidate from C is declared to be the election winner. Otherwise, each voter $v \in V$ reports her preferences over the candidates in A; these preferences are obtained by restricting $\succ_v$ to A. We assume that all voters report their preferences sincerely. A *voting rule* takes the set A and the list of voters' preferences over A as input, and outputs a candidate $w \in A$; this candidate is called the *election winner*. In this paper, we will only consider the *Plurality rule*. Under this rule, the score $\mathrm{sc}(c)$ of a candidate $c \in A$ is equal to the number of voters who rank c first among the candidates in A; the winner is the candidate with the highest score, with ties broken according to a fixed priority order $\lhd$ over C (i.e., if $\mathrm{sc}(a) = \mathrm{sc}(b) > \mathrm{sc}(c)$ for every $c \in A \setminus \{a, b\}$ and $a \lhd b$, then a is declared to be the election winner). For consistency, we assume that, when $A = \emptyset$, the winner is the top candidate in $\lhd$; we denote this candidate by $c_\lhd$.

Candidacy games

Candidacy games are characterized by two important features: (1) candidates themselves also have preferences over the set C, and (2) each candidate can choose whether to run in the election.

Formally, each candidate $c \in C$ has two available actions: 1 (run) and 0 (abstain). Also, each candidate $c \in C$ is endowed with a preference order $\succ_c$ over C; the list $P^C = (\succ_c)_{c \in C}$ is referred to as the *candidates' preference profile*. We say that a candidate $c \in C$ has *self-supporting preferences* if $c \succ_c a$ for all $a \in C \setminus \{a\}$. While much of the work on candidacy games assumes that all candidates have self-supporting preferences, and, indeed, this assumption is reasonable in many real-life scenarios, there are settings where it is not satisfied. For example, a person may volunteer for a leadership position because her colleagues expect her to do so, but secretly hope not to be elected because of the associated workload. Therefore, instead of making this assumption in our work, we explicitly indicate which of our results hold for candidates with self-supporting preferences.

The tuple $\langle C, V, P^V, P^C, \lhd \rangle$ defines two related strategic games: the *vanilla strategic candidacy game (VSCG)* $\Gamma = \Gamma(C, V, P^V, P^C, \lhd)$ and the *lazy strategic candidacy game (LSCG)* $\Gamma^L = \Gamma^L(C, V, P^V, P^C, \lhd)$. In both games, the set of players is C and each player's set of actions is $\{0, 1\}$. We denote the action (strategy) of a player $c \in C$ by s_c; the vector $\mathbf{s} = (s_c)_{c \in C}$ is called the *strategy profile*. A strategy profile $\mathbf{s}$ defines the set of actual candidates $A(\mathbf{s}) = \{c \in C \mid s_c = 1\}$. Consequently, each strategy profile defines an outcome $w(\mathbf{s}) \in C$: $w(\mathbf{s})$ is simply the outcome of Plurality voting by voters in V over candidates in $A(\mathbf{s})$ when $A(\mathbf{s})$ is not empty (with ties broken according to $\lhd$), and $w(\mathbf{s}) = c_\lhd$ when $A(\mathbf{s}) = \emptyset$. We denote by $\mathrm{sc}(c, \mathbf{s})$ the Plurality score obtained by a candidate $c \in A(\mathbf{s})$ when voters in V vote over candidates in $A(\mathbf{s})$. The only difference between Γ and Γ^L is in the players' preferences over the states of the game: in Γ a player c prefers $\mathbf{s}$ to $\mathbf{t}$ (denoted by $\mathbf{s} >_c \mathbf{t}$) if and only if $w(\mathbf{s}) \succ_c w(\mathbf{t})$, whereas in Γ^L a player c prefers $\mathbf{s}$ to $\mathbf{t}$ (denoted by $\mathbf{s} >_c^L \mathbf{t}$) if and only if (i) $w(\mathbf{s}) \succ_c w(\mathbf{t})$ or (ii) $w(\mathbf{s}) = w(\mathbf{t})$, $s_c = 0$ and $t_c = 1$. That is, in Γ^L a candidate is willing to run if his presence can change the election outcome for the better; however, if his presence has no effect on the election outcome, he prefers not to participate. These definitions extend to voting rules other than Plurality in an obvious way.

We will be interested in Nash equilibria of VSCGs and LSCGs. Recall that a strategy profile $\mathbf{s}$ is said to be a *pure strategy Nash equilibrium* if no player in the game can profitably deviate. Formally, given a VSCG Γ (respectively, an LSCG Γ^L) described by a 5-tuple $\langle C, V, P^C, P^V, \lhd \rangle$, we say that a strategy profile $\mathbf{s}$ is a *pure strategy Nash equilibrium (PNE)* of Γ (respectively, Γ^L) if for every candidate $c \in C$ it is not the case that $\mathbf{t} >_c \mathbf{s}$ (respectively, $\mathbf{t} >_c^L \mathbf{s}$), where $\mathbf{t}$ is the strategy profile given by $t_a = s_a$ for $a \in C \setminus \{c\}$, $t_c = 1 - s_c$.

In what follows, we write $v : abc$ (respectively, $c : abc$) to indicate that the preference order of voter v (respectively, candidate c) is $a \succ b \succ c$. When voters' names are not important, we omit them, and write, e.g., (abc, abc, bca) to denote a profile where two voters rank the candidates as $a \succ b \succ c$ and one voter ranks the candidates as $b \succ c \succ a$. Also, we will sometimes identify a state $\mathbf{s}$ with its set of actual candidates $A(\mathbf{s})$.

3 Nash Equilibria

We start our investigation of lazy strategic candidacy games by exploring the properties of their PNE.

Properties of Nash Equilibria

In this section, we consider the role of common social choice concepts, such as Pareto dominance and Condorcet winners, in LSCGs. We will also compare these games to their vanilla counterparts.

Pareto Dominance Consider an LSCG Γ^L where candidate a *Pareto-dominates* candidate b in P^V, i.e., every voter in V prefers a to b. One would expect that b cannot be the election winner in a PNE of Γ^L. However, this turns out to depend on b's preferences and the tie-breaking order $\lhd$.

Proposition 1. *Let $\Gamma^L = \Gamma^L(C, V, P^C, P^V, \lhd)$. Suppose that $a \succ_v b$ for all $v \in V$. If (1) $a \succ_a b$ and (2) $a \lhd b$ then for* *every PNE* $\mathbf{s}$ *of* Γ^L *we have* $b \neq w(\mathbf{s})$. *However, if either of the conditions (1) and (2) is not satisfied,* b *can be a winner in a PNE of* Γ^L.

Proof. Suppose that conditions (1) and (2) are satisfied, and suppose for the sake of contradiction that there is a PNE $\mathbf{s}$ with $w(\mathbf{s}) = b$. Note that $a \lhd b$ implies $b \neq c_\lhd$, and hence $t = \mathrm{sc}(b, \mathbf{s}) > 0$. Clearly, $a \notin A(\mathbf{s})$: otherwise no voter would rank b first. Consider the modified strategy profile $\mathbf{s}'$ with $s'_a = 1$, $s'_c = s_c$ for $c \in C \setminus \{a\}$. We have $\mathrm{sc}(a, \mathbf{s}') \geq t$: all voters who vote for b in $\mathbf{s}$ would switch to voting for a in $\mathbf{s}'$. On the other hand, since $b = w(\mathbf{s})$, for every candidate $c \in C \setminus \{a, b\}$ we have $\mathrm{sc}(c, \mathbf{s}') \leq \mathrm{sc}(c, \mathbf{s}) \leq t$, and $\mathrm{sc}(c, \mathbf{s}) = t$ implies $b \lhd c$. As $a \lhd b$, this means that $a = w(\mathbf{s}')$; since $a \succ_a b$, switching from $\mathbf{s}$ to $\mathbf{s}'$ is a profitable deviation for a.

It is easy to see that condition (1) is necessary: consider an election with $C = \{a, b\}$, a single voter who prefers a to b, and candidates' preferences given by $b \succ_a a$, $b \succ_b a$. Clearly, the resulting game has a PNE $\mathbf{s}$ with $A(\mathbf{s}) = \{b\}$.

To see that condition (2) is necessary, let $C = \{a, b, c, d\}$, and suppose that $b \lhd c \lhd a \lhd d$ and voters'/candidates' preferences are given by

$$P^V = (dcab, cdab, cdab, abcd, abcd),$$
$$P^C = (a : abcd, b : bacd, c : cabd, d : dbac).$$

Then the strategy profile $\mathbf{s}$ with $A(\mathbf{s}) = \{b, c, d\}$ is a PNE, even though all voters as well as candidate a prefer a to b. $\square$

Note that condition (1) is satisfied whenever all candidates have self-supporting preferences. Furthermore, the reader can check that Proposition 1 remains true for VSCGs.

Condorcet Winners An important notion in social choice is that of a *Condorcet winner*: this is a candidate that beats every other candidate in their pairwise election. Formally, given a profile P^V over a candidate set C and two candidates $a, b \in C$, let n_{ab} denote the number of voters in V who prefer a to b. A candidate $c \in C$ is said to be a *Condorcet winner* in (C, P^V) if $n_{ca} > n_{ac}$ for all $a \in C \setminus \{c\}$.

This concept turns out to be relevant for our analysis: if a candidate c is a Condorcet winner in (C, P^V), and c has self-supporting preferences, then the strategy profile $\mathbf{s}$ with $A(\mathbf{s}) = \{c\}$ is clearly a PNE of the respective game Γ^L. Indeed, if any other candidate decided to run, he would lose to c anyway, and c has no reason to withdraw.

In fact, we can obtain a full characterization of PNE with exactly one participating candidate, by modifying this concept so as to take into account the candidates' preferences and the tie-breaking order. Specifically, given a priority order $\lhd$ over C, we say that a candidate $c \in C$ is a $(P^C, \lhd)$-*Condorcet winner* in (C, P^V, P^C) if for each $a \in C \setminus \{c\}$ we have (i) $n_{ca} > n_{ac}$ or (ii) $n_{ca} = n_{ac}$ and $c \lhd a$ or (iii) $c \succ_a a$. Note that an election may have several $(P^C, \lhd)$-Condorcet winners; however, if the number of voters is odd and all candidates have self-supporting preferences, a candidate is a $(P^C, \lhd)$-Condorcet winner in (C, P^V, P^C) if and only if she is a Condorcet winner in (C, P^V).

We are now ready to state our characterization result.

Proposition 2. *Let* $\Gamma^L = \Gamma^L(C, V, P^C, P^V, \lhd)$. *A strategy profile* $\mathbf{s}^0$ *with* $A(\mathbf{s}^0) = \emptyset$ *is a PNE of* Γ^L *if and only if we have* $c_\lhd \succ_c c$ *for all* $c \in C \setminus \{c_\lhd\}$. *Further, for a given candidate* $c \in C$ *the strategy profile* $\mathbf{s}^c$ *with* $A(\mathbf{s}^c) = \{c\}$ *is a PNE of* Γ^L *if and only if* c *is a* $(P^C, \lhd)$*-Condorcet winner in* (C, P^V, P^C), $c \neq c_\lhd$, *and* $c \succ_c c_\lhd$.

Proof. We have $w(\mathbf{s}^0) = c_\lhd$. Thus, if every candidate $c \in C \setminus \{c_\lhd\}$ prefers $c_\lhd$ to herself, she has no reason to participate, and neither does $c_\lhd$. Conversely, if $c \succ_c c_\lhd$ for some $c \in C \setminus \{c_\lhd\}$, candidate c would prefer to join the election and win.

Now, consider a candidate c and the strategy profile $\mathbf{s}^c$. If c is a $(P^C, \lhd)$-Condorcet winner, no candidate that can change the election outcome by participating is willing to do so. Moreover, if $c \neq c_\lhd$ and $c \succ_c c_\lhd$, c prefers not to withdraw. Thus, under these conditions $\mathbf{s}^c$ is a PNE of Γ^L. Conversely, if $c = c_\lhd$ or $c_\lhd \succ_c c$, then c would prefer to withdraw, and if c is not a $(P^C, \lhd)$-Condorcet winner, there is a candidate a who beats c in their pairwise election and prefers herself to c; this candidate could then participate in the election and win. $\square$

Interestingly, the analysis in Proposition 2 shows that if $c_\lhd$ is a $(P^C, \lhd)$-Condorcet winner then there is no PNE where $c_\lhd$ is the unique candidate participating in the election; in contrast, if any other candidate is a $(P^C, \lhd)$-Condorcet winner and has self-supporting preferences, there is a PNE where this candidate is the unique participant. Elkind et al. [2014] observe that a similar phenomenon arises in equilibria of Plurality elections with lazy voters.

We remark that even if an instance (C, P^V, P^C) has a $(P^C, \lhd)$-Condorcet winner, it may have PNE where the $(P^C, \lhd)$-Condorcet winner does not participate.

Example 1. *Let* $C = \{a, b, c, d\}$, *and suppose that* $a \lhd b \lhd c \lhd d$ *and voters'/candidates' preferences are given by*

$$P^V = (bacd, cadb, dacb),$$
$$P^C = (a : abcd, b : bcda, c : cbda, d : dbca).$$

Candidate a *is the unique* $(P^C, \lhd)$*-Condorcet winner, yet the strategy profile* $\mathbf{s}$ *with* $A(\mathbf{s}) = \{b, c, d\}$ *is a PNE.*

Vanilla vs. Lazy Games It is not hard to verify that every PNE of a lazy strategic candidacy game is also a PNE of the respective vanilla strategic candidacy game.

Proposition 3. *Consider a tuple* $\langle C, V, P^C, P^V, \lhd \rangle$ *and let* $\Gamma = \Gamma(C, V, P^C, P^V, \lhd)$ *and* $\Gamma^L = \Gamma^L(C, V, P^C, P^V, \lhd)$. *Then every PNE of* Γ^L *is a PNE of* Γ.

Proof. Consider a strategy profile $\mathbf{s}$ for $\langle C, V, P^C, P^V, \lhd \rangle$ and a candidate $c \in C$. Suppose that $\mathbf{s}$ is a PNE of Γ^L with winner w. If $s_c = 1$, then, since $\mathbf{s}$ is a PNE of Γ^L, setting $s_c = 0$ would change the election outcome from w to some p with $w \succ_c p$; thus, in Γ player c would not want to change his action to 0 either. On the other hand, if $s_c = 0$, since $\mathbf{s}$ is a PNE of Γ^L, it has to be the case that setting $s_c = 1$ either does not change the election outcome at all, or changes it to one that c likes less than w. Hence c cannot profitably deviate in Γ either. $\square$

It is known that VSCGs under Plurality may have no PNE [Lang *et al.*, 2013]; Proposition 3 implies that this is also the case for LSCGs.

The converse of Proposition 3 is not true: a strategy profile that is a PNE of Γ may fail to be a PNE of Γ^L. An example is easy to construct: suppose, for instance, that $C = \{a, b\}$, all voters prefer a to b, and the candidates have self-supporting preferences. Then $\mathbf{s} = (1, 1)$ is a PNE of Γ, but in Γ^L candidate b prefers to abstain. A somewhat more complicated example shows that there are settings where Γ has a PNE, but Γ^L does not.

Example 2. *Let $C = \{a, b, c\}$, and suppose that we have*

$$P^V = (abc, bca, cab), \quad P^C = (a : abc, b : bca, c : cab).$$

Suppose also that $a \lhd b \lhd c$. Then $\Gamma(C, V, P^V, P^C, \lhd)$ has a PNE $\mathbf{s}$ with $A(\mathbf{s}) = \{a, b\}$, $w(\mathbf{s}) = a$. However, $\Gamma^L = \Gamma^L(C, V, P^V, P^C, \lhd)$ has no PNE. To see this, observe first that LSCGs have no PNE with exactly two participating candidates: indeed, the loser would prefer to withdraw. Moreover, as G^L has no $(P^C, \lhd)$-Condorcet winner, by Proposition 2 it has no PNE with one participating candidate, and since candidates' preferences are self-supporting, there is no PNE with zero participating candidates. It remains to observe that $(1, 1, 1)$ is not a PNE of G^L either, as candidate c would rather withdraw.

Complexity of PNE

Proposition 2 describes PNE where the number of participating candidates is zero or one. However, Example 1 illustrates that there can be PNE with three participants, and it is easy to extend the construction in that example to obtain PNE with any number of participants. This indicates that it may be nontrivial to identify all PNE of a given game, and provides motivation for studying the complexity of computing PNE in lazy strategic candidacy games.

To this end, we define the following decision problems:

- LAZYNE: Given a game $\Gamma^L = \Gamma^L(C, V, P^V, P^C, \lhd)$ and a strategy profile $\mathbf{s}$, decide whether $\mathbf{s}$ is a PNE of Γ^L.
- LAZY∃NE: Given a game $\Gamma^L = \Gamma^L(C, V, P^V, P^C, \lhd)$ decide whether Γ^L has a PNE.
- LAZY∃WNE: Given a game $\Gamma^L = \Gamma^L(C, V, P^V, P^C, \lhd)$ and a candidate $c \in C$, decide whether Γ^L has a PNE $\mathbf{s}$ with $w(\mathbf{s}) = c$.

It is easy to see that LAZYNE is in P: as each player only has two available actions, we only need to consider n possible deviations. In contrast, LAZY∃NE and LAZY∃WNE turn out to be computationally hard.

Theorem 1. *For lazy strategic candidacy games, the problems* LAZY∃NE *and* LAZY∃WNE *are* NP*-complete, even if all candidates have self-supporting preferences.*

Proof. Our observation that LAZYNE is polynomial-time solvable immediately implies that these problems are in NP.

To show that these problems are NP-hard, we provide a reduction from a restricted variant of the classic EXACT 3-COVER (X3C) problem. An instance of this problem is given by a set of ground elements $U = \{u_1, \dots, u_{3r}\}$ and a family $\mathcal{Z} = \{Z_1, \dots, Z_t\}$ of 3-element subsets of U; we assume that $Z_\ell = \{u_{i_\ell}, u_{j_\ell}, u_{k_\ell}\}$ for some $i_\ell, j_\ell, k_\ell \in \{1, \dots, 3r\}$. It is a "yes"-instance if there exists a subfamily $\widehat{\mathcal{Z}} \subset \mathcal{Z}$ such that $\cup_{Z \in \widehat{\mathcal{Z}}} Z = U$ and $Z_i \cap Z_j = \emptyset$ for all $Z_i, Z_j \in \widehat{\mathcal{Z}}$ with $i \neq j$; otherwise it is a "no"-instance. We additionally assume that $r \geq 2$ and each element of U is contained in exactly three distinct sets in $\mathcal{Z}$; note that this implies that $t = 3r$. This variant of X3C, which we will refer to as RX3C, has been shown to be NP-complete by Gonzalez [1984].

Given an instance $(U, \mathcal{Z})$ of RX3C with $|U| = 3r$, we set $q = 30r^2$ and construct an LSCG with the set of candidates $C = U \cup \mathcal{Z} \cup \{w_0, w_1, w_2\}$ and $n = 18r + 3r(q-1) + 3q - 3r$ voters. The preference profiles P^V and P^C are specified in Tables 1 and 2, respectively. In these tables we write U_i to denote the order $u_i \succ u_{i+1} \succ \dots \succ u_{3r} \succ u_1 \succ \dots \succ u_{i-1}$ over U; we write $U_i \setminus U'$, where $U' \subset U$, to refer to the order U_i with the candidates in U' removed. Also, we write $\mathcal{Z}_{-i}$ to refer to an arbitrary order over $\mathcal{Z} \setminus \{Z_i\}$; similarly, we write U and $\mathcal{Z}$ to refer to arbitrary orders over U and $\mathcal{Z}$, respectively. Finally, we set $w_1 \lhd U \lhd \mathcal{Z} \lhd w_2 \lhd w_0$. Note that the candidates' preferences are self-supporting. Let $\Gamma^L = \Gamma^L(C, V, P^V, P^C, \lhd)$.

We will now argue that if we have started with a "yes"-instance of RX3C then Γ^L has a PNE with winner w_1, and if we have started with a "no"-instance of RX3C then Γ^L has no PNE. This establishes NP-hardness of both LAZY∃NE and LAZY∃WNE.

Specifically, it can be checked that if a collection of subsets $\mathcal{Z}'$ provides an exact cover of U then the strategy vector $\mathbf{s}$ with $A(\mathbf{s}) = U \cup \{w_0, w_1, w_2\} \cup (\mathcal{Z} \setminus \mathcal{Z}')$ is a PNE with $w(\mathbf{s}) = w_1$.

For the converse direction we show that if Γ^L has a PNE $\mathbf{s}$, then $U \cup \{w_0, w_1, w_2\} \subseteq A(\mathbf{s})$ and $w(\mathbf{s}) = w_1$, and use this to conclude that candidates in $\mathcal{Z} \setminus A(\mathbf{s})$ provide an exact cover of U.

In more detail, observe first that this profile does not admit a $(P^C, \lhd)$-Condorcet winner: indeed, the candidates' preferences are self-supporting, for every candidate $c \in C \setminus U$ at least $3r(q-1) > n/2$ voters prefer any candidate in U to c, and for every candidate $u_i \in C$ at least $(3r-1)(q-1) > n/2$ voters prefer u_{i-1} to u_i (where we let $u_0 := u_{3r}$).

Now, suppose that Γ^L had a PNE $\mathbf{s}$. We will argue that $U \cup \{w_0, w_1, w_2\} \subseteq A(\mathbf{s})$ and $w(\mathbf{s}) = w_1$.

To show this, we will first prove that $A(\mathbf{s}) \cap U \neq \emptyset$. Indeed, if $A(\mathbf{s}) \cap U = \emptyset$, then if u_1 were to enter the election, he would receive $3r(q-1) > n/2$ votes and win. As u_1 has self-supporting preferences, this is a contradiction with $\mathbf{s}$ being a PNE of Γ^L.

Moreover, we have $U \subseteq A(\mathbf{s})$. Indeed, suppose that $U \setminus A(\mathbf{s}) \neq \emptyset$. Then there exists a candidate $u_i \in U \cap A(\mathbf{s})$ with $\mathrm{sc}(u_i, \mathbf{s}) \geq 2q-2$. As $A(\mathbf{s}) \cap U \neq \emptyset$, the score of each candidate in $\mathcal{Z} \cup \{w_0, w_1, w_2\}$ does not exceed $q + 18r < 2q-2$, so the winner is some candidate $u_j \in U$. Applying an inductive argument to candidates $u_{j-1}, \dots, u_1, u_{3r}, \dots, u_{j+1}$ (where $u_0 := u_{3r}$ and $u_{3r+1} := u_1$), we conclude that each of these candidates prefers not to participate in the election. Thus, $\mathrm{sc}(u_j, \mathbf{s}) \geq 3r(q-1)$, and therefore all candidates in $\mathcal{Z} \cup \{w_0, w_1, w_2\}$ prefer not to participate as well, so

Block 1 ($3r$ votes)				**Block 2** ($6r$ votes)				**Block 3** ($9r$ votes)				
Z_1	Z_2	$\dots$	Z_{3r}	Z_1	Z_1	$\dots$	Z_{3r}	Z_{3r}	$\dots$	Z_ℓ	Z_ℓ	Z_ℓ	$\dots$
w_1	w_1	$\dots$	w_1	w_2	w_2	$\dots$	w_2	w_2	$\dots$	u_{i_ℓ}	u_{j_ℓ}	u_{k_ℓ}	$\dots$
w_2	w_2	$\dots$	w_2	U_1	U_1	$\dots$	U_{3r}	U_{3r}	$\dots$	w_2	w_2	w_2	$\dots$
U_1	U_2	$\dots$	U_{3r}	$\mathcal{Z}_{-1}$	$\mathcal{Z}_{-1}$	$\dots$	$\mathcal{Z}_{-3r}$	$\mathcal{Z}_{-3r}$	$\dots$	$U_{i_\ell} \setminus \{u_{i_\ell}\}$	$U_{j_\ell} \setminus \{u_{j_\ell}\}$	$U_{k_\ell} \setminus \{u_{k_\ell}\}$	$\dots$
$\mathcal{Z}_{-1}$	$\mathcal{Z}_{-2}$	$\dots$	$\mathcal{Z}_{-3r}$	w_0	w_0	$\dots$	w_0	w_0	$\dots$	$\mathcal{Z}_{-\ell}$	$\mathcal{Z}_{-\ell}$	$\mathcal{Z}_{-\ell}$	$\dots$
w_0	w_0	$\dots$	w_0	w_1	w_1	$\dots$	w_1	w_1	$\dots$	w_0	w_0	w_0	$\dots$
									$\dots$	w_1	w_1	w_1	$\dots$

Block 4 ($3r(q-1)+3q-3r$ votes)						
u_1	u_2	$\dots$	u_{3r}	w_2	w_1	w_0
$U_1 \setminus \{u_1\}$	$U_2 \setminus \{u_2\}$	$\dots$	$U_{3r} \setminus \{u_{3r}\}$	U	U	U
$\mathcal{Z}$	$\mathcal{Z}$	$\dots$	$\mathcal{Z}$	$\mathcal{Z}$	$\mathcal{Z}$	$\mathcal{Z}$
w_2	w_2	$\dots$	w_2	w_0	w_2	w_2
w_0	w_0	$\dots$	w_0	w_1	w_0	w_1
w_1	w_1	$\dots$	w_1	$(q-2r)$ copies	$(q-r)$ copies	q copies
$(q-1)$ copies	$(q-1)$ copies		$(q-1)$ copies			

Table 1: Proof of Theorem 1: voters' preferences

$A(\mathbf{s}) = \{u_j\}$. But this is a contradiction with Proposition 2, as u_j is not a $(P^C, \lhd)$-Condorcet winner. The contradiction shows that $U \subseteq A(\mathbf{s})$.

We can now conclude that $\{w_0, w_1, w_2\} \subset A(\mathbf{s})$: if this was not the case, u_1 would get at least $2q - 2r - 1$ point and win, in which case u_{3r} would prefer to withdraw. It follows that $w_1 = w(\mathbf{s})$: otherwise a candidate $u_i \in U \setminus \{w(\mathbf{s}), u_{3r}\}$ would prefer to withdraw, as this would make u_{i+1} the winner, and u_i prefers u_{i+1} to all candidates other than himself and w_1.

Now that we have established that $w_1 = w(\mathbf{s})$, we can observe that exactly $2r$ candidates from $\mathcal{Z}$ participate in the election: if $|A(\mathbf{s}) \cap \mathcal{Z}| = 2r + x$ for $x > 0$ then $\mathrm{sc}(w_1, \mathbf{s}) < q = \mathrm{sc}(w_0, \mathbf{s})$, and if $|A(\mathbf{s}) \cap \mathcal{Z}| = 2r - x$ for $x > 0$ then $\mathrm{sc}(w_2, \mathbf{s}) = q - 2r + 2(r + x) > q - r + (r + x) = \mathrm{sc}(w_1, \mathbf{s})$.

Let $\widehat{\mathcal{Z}} = \mathcal{Z} \setminus A(\mathbf{s})$. We have shown that $|\widehat{\mathcal{Z}}| = r$. We will now argue that $\widehat{\mathcal{Z}}$ provides an exact cover of U. Indeed, suppose that this is not the case. As $|\widehat{\mathcal{Z}}| = r$, it follows that there exists an element $u_i \in U$ that appears in two distinct sets $Z_j, Z_k \in \widehat{\mathcal{Z}}$. But then we have $\mathrm{sc}(u_i, \mathbf{s}) \geq q + 1 > q = \mathrm{sc}(w_1, \mathbf{s})$, a contradiction with $w_1 = w(\mathbf{s})$. $\square$

$Z_\ell, \ell = 1, \dots, 3r$	$u_i, i = 1, \dots, 3r$	w_1	w_2	w_0
Z_ℓ	u_i	w_1	w_2	w_0
u_{i_ℓ}	w_1	w_2	w_1	w_1
u_{j_ℓ}	$U_i \setminus \{u_i\}$	U	U	U
u_{k_ℓ}	$\mathcal{Z}$	$\mathcal{Z}$	$\mathcal{Z}$	$\mathcal{Z}$
w_1	w_2	w_0	w_0	w_2
$U_\ell \setminus \{u_{i_\ell}, u_{j_\ell}, u_{k_\ell}\}$	w_0			
$\mathcal{Z}_{-\ell}$				
w_2				
w_0				

Table 2: Proof of Theorem 1: candidates' preferences

It is natural to ask whether the hardness results established in Theorem 1 still hold if the number of voters or the number of candidates is small. It turns out that both LAZY∃NE and LAZY∃WNE become easier under these constraints.

In more detail, it is immediate that both LAZY∃NE and LAZY∃WNE admit an algorithm whose running time is $2^m \mathrm{poly}(n, m)$, where $n = |V|$ and $m = |C|$: we can simply enumerate all strategy profiles, and, for each of them, check whether it is a PNE and, in case of LAZY∃WNE, compute its winner. In fact, this procedure works not just for Plurality, but for any voting rule with a polynomial-time winner determination procedure. To obtain an algorithm that performs well when the number of voters n is small, we observe that (1) in any PNE of a Plurality-based LSCG the number of candidates with a positive score does not exceed n, and (2) in any PNE $\mathbf{s}$ we have $\mathrm{sc}(c, \mathbf{s}) > 0$ for all $c \in A(\mathbf{s})$. Therefore, it suffices to consider strategy profiles $\mathbf{s}$ with $|A(\mathbf{s})| \leq n$; the number of such profiles does not exceed m^n. The following proposition summarizes these observations in the language of fixed-parameter complexity (see, e.g., Niedermeier [2006]).

Proposition 4. *The problems* LAZY∃NE *and* LAZY∃WNE *are in* FPT *with respect to the number of candidates m and in* XP *with respect to the number of voters n.*

4 Best-response Dynamics

In practice, knowing that a game has a PNE does not necessarily mean that players will reach a stable outcome, as it may be difficult for computationally bounded players to find it. Therefore, it is important to understand which states of the game can be achieved by players that optimize their behavior in an iterative fashion. To this end, we will now consider *myopic dynamics* of LSCGs, i.e., sequences of states where each state is obtained from the previous one by allowing a single player to deviate in a way that is beneficial for her.

It is natural to assume that once a candidate withdraws from the election, he cannot rejoin it; indeed, this is typically the case in real-life elections. Under this assumption, if all candidates were present in the initial state, then at each subsequent step one candidate withdraws from the election until none of the remaining candidates has an incentive to do so; we refer to this type of dynamics as W-dynamics. We will also consider the complementary setting where initially the set of candidates is empty, and the candidates join the election one by one and are not allowed to leave; we refer to this type of dynamics as J-dynamics. Under both types of dynamics each candidate can change his action at most once, and therefore any such dynamic process converges in at most m steps, where m is the number of candidates. Note, however, that the final state need not be a PNE: e.g., under W-dynamics it may be the case that a candidate who has withdrawn at an earlier point would prefer to rejoin, but is prevented from doing so by the "no rejoining" rule. Thus, the set of states reachable by such dynamics can be seen as a distinct solution concept for LSCGs.

We will now discuss some properties of the W-dynamics and J-dynamics. For simplicity, we assume that the number of voters is odd and the candidates' preferences are self-supporting (and hence a candidate is a $(P^C, \lhd)$-Condorcet winner if and only if she is a Condorcet winner); however, our results extend to the general case.

Consider a game $\Gamma^L = \Gamma^L(C, V, P^C, P^V, \lhd)$. We have argued that if (C, P^V) has a Condorcet winner $c \neq c_\lhd$ then Γ^L has a PNE $\mathbf{s}$ with $A(\mathbf{s}) = \{c\}$. We will now investigate whether such a PNE can be achieved by W-dynamics or J-dynamics.

Proposition 5. *Let $\Gamma^L = \Gamma^L(C, V, P^C, P^V, \lhd)$, and suppose that $c \neq c_\lhd$ in a Condorcet winner in (C, P^V). Then there is a J-dynamics that leads to the state $\mathbf{s}$ with $A(\mathbf{s}) = \{c\}$. However, there may be no W-dynamics that terminates in a state $\mathbf{s}'$ with $c \in A(\mathbf{s}')$, and there may exist a J-dynamics that terminates in a state $\mathbf{s}''$ with $c \notin A(\mathbf{s}'')$.*

Proof. The J-dynamics where agent c joins at the first step terminates as soon as c joins, since the resulting state is a PNE. For W-dynamics, consider the profile in Example 1: here the Condorcet winner a is the only candidate that has an incentive to leave from the original state $(1, 1, 1, 1)$, so he is not present in any terminal state of the W-dynamics. Further, the reader can verify that $\emptyset \to \{d\} \to \{c, d\} \to \{b, c, d\}$ is a J-dynamics for this profile that results in a stable state where a does not participate. □

It is also interesting to ask whether our dynamics can terminate in the state where all candidates are present. Clearly, for W-dynamics this is only possible if this state is a PNE. In contrast, the following example shows that J-dynamics can reach this state even if it is not a PNE.

Example 3. *Let $C = \{a, b, c, d\}$, and suppose that $a \lhd b \lhd c \lhd d$ and voters'/candidates' preferences are given by*

$$P^V = (abcd, bcda, cdab, dabc),$$
$$P^C = (a : abcd, b : bcda, c : cdab, d : dabc).$$

Observe that $\{d\} \to \{c, d\} \to \{b, c, d\} \to \{a, b, c, d\}$ is a J-dynamics for this profile, and its final state is not a PNE.

Proposition 5 and Example 3 further illustrate that the set of terminal states of W- and J-dynamics may differ from the set of PNE of a given LSCG.

Finally, we can show that deciding whether either of our dynamics can terminate in a state with a given winner/set of actual candidates is NP-hard; the proof is obtained by modifying the construction in the proof of Theorem 1 and is omitted due to space constraints.

Theorem 2. *Given an LSCG game Γ^L, a candidate c and a state $\mathbf{s}$ of Γ^L, the following problems are NP-complete:*

- *deciding whether J-dynamics may converge to $\mathbf{s}$;*
- *deciding whether W-dynamics may converge to $\mathbf{s}$;*
- *deciding whether J-dynamics may converge to a state $\mathbf{s}'$ with $w(\mathbf{s}') = c$;*
- *deciding whether W-dynamics may converge to a state $\mathbf{s}'$ with $w(\mathbf{s}') = c$.*

5 Related Work and Discussion

Early work on strategic candidacy games [Dutta *et al.*, 2001; 2002; Ehlers and Weymark, 2003; Eraslan and McLennan, 2004; Rodriguez-Alvarez, 2006b; 2006a] focused on voting rules other than Plurality. Vanilla strategic candidate games under Plurality are briefly considered by Lang et al. [2013] and then in more detail by Polukarov et al. [2015]. In particular, Lang et al. argue that any VSCG with three candidates and self-supporting preferences has a PNE, and provide an example of a Plurality-based VSCG with four candidates and no PNE. The first of these results stands in contrast with Example 2, where we construct a three-candidate LSCG with no PNE. Polukarov et al. primarily focus on dynamics of VSCGs. However, their model is different from the one considered in Section 4, in that they allow candidates to rejoin an election after having withdrawn. Their results imply that the problems of checking whether a given VSCG has a PNE, or, more narrowly, a PNE with a specific winner (i.e., the vanilla analogues of LAZY∃NE and LAZY∃WNE) are NP-complete.

These results together with the analysis in our paper indicate that VSCGs and LSCGs are broadly similar: in both, Condorcet winners can win in PNE under mild additional assumptions, and computational problems related to PNE of both types of games are NP-hard. However, Example 2 points to a crucial difference: laziness may destroy stability. Intuitively, the impact of laziness can be quite substantial, as illustrated by the following simple probabilistic argument.

Consider a strategic candidacy game with $C = \{a, b, c\}$ and $P^C = (a : abc, b : bca, c : cab)$. We will argue that if voters' preferences are drawn uniformly at random, then with probability at least $1/2$ the profile $\mathbf{s} = (1, 1, 1)$ is a PNE of the respective VSCG; however, it is not a PNE of the respective LSCG. Indeed, assume without loss of generality that $w(\mathbf{s}) = a$. Then a and c cannot benefit from withdrawing, and b only wants to withdraw if the number of bca-voters is larger that the number of bac-voters. As voters who rank b first are equally likely to be bac-voters or bca-voters, the first claim follows. For the second claim, note that b prefers to withdraw from $\mathbf{s}$.

Extending the analysis in the example above to estimate the frequency of profiles that admit PNE (or, more narrowly, have $\mathbf{s} = (1, \ldots, 1)$ as PNE), both for vanilla and for lazy strategic candidacy games, is an interesting direction for future work; however, Theorem 1 and the hardness results of Polukarov et al. indicate that this may be a difficult problem.

In this context, it may be interesting to note that, in contrast, Plurality voting games with lazy voters and lexicographic tie-breaking are known to have a very simple structure: in PNE of such games at most one voter casts his vote, and all other voters abstain. On the other hand, when voters are not lazy, there are PNE with many voters [Elkind *et al.*, 2014]. Thus, the similarity between VSCGs and LSCGs should not be taken for granted.

6 Conclusions and Future Work

We have introduced strategic candidacy games with lazy candidates and analyzed the properties of their PNE. In contrast with much of the previous work, we explicitly examined the role of the assumption that candidates' preferences are self-supporting, as well as the impact of the tie-breaking order: for instance, Proposition 1 indicates that we need to take these into account when asking whether Pareto-dominated candidates can appear in PNE of LSCGs.

As argued in Section 5, it would be interesting to obtain quantitative results regarding frequency of profiles with PNE in LSCGs and VSCGs; in particular, it is important to know whether the profile $\mathbf{s}$ with $A(\mathbf{s}) = C$ is considerably less likely to be stable in LSCGs compared to VSCGs.

Further, our model implicitly assumes that for each candidate the campaign costs are small: a candidate is willing to participate even if all that it accomplishes is changing the outcome from his 8th most preferred candidate to his 7th most preferred candidate. Alternatively, we could explicitly model both participation costs and utilities from each outcome, and allow for the possibility that some of the changes in the election outcomes are not significant enough to justify participation. While this model is richer than the one we have considered, for large enough costs, the associated computational problems may become simpler, as in all but a handful of situations most of the candidates would prefer to withdraw.

Acknowledgments

This work was partially supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement number 337122, RFFI 14-01-00156-a, and STSM Grant from the COST Action IC1205.

References

[Brill and Conitzer, 2015] M. Brill and V. Conitzer. Strategic voting and strategic candidacy. In *AAAI'15*, pages 819–826, 2015.

[Desmedt and Elkind, 2010] Y. Desmedt and E. Elkind. Equilibria of plurality voting with abstentions. In *ACM EC'10*, pages 347–356, 2010.

[Dutta *et al.*, 2001] B. Dutta, M. L. Breton, and M. O. Jackson. Strategic candidacy and voting procedures. *Econometrica*, 69:1013–1037, 2001.

[Dutta *et al.*, 2002] B. Dutta, M. L. Breton, and M. O. Jackson. Voting by successive elimination and strategic candidacy in committees. *Journal of Economic Theory*, 103:190–218, 2002.

[Ehlers and Weymark, 2003] L. Ehlers and J. A. Weymark. Candidate stability and nonbinary social choice. *Economic Theory*, 22(2):233–243, 2003.

[Elkind *et al.*, 2014] E. Elkind, E. Markakis, S. Obraztsova, and P. Skowron. Equilibria of plurality voting: Lazy and truth-biased voters. Arxiv abs/1409.4132, 2014.

[Eraslan and McLennan, 2004] H. Eraslan and A. McLennan. Strategic candidacy for multivalued voting procedures. *Journal of Economic Theory*, 117(1):29–54, 2004.

[Gonzalez, 1984] T. F. Gonzalez. Clustering to minimize the maximum intercluster distace. *Theoretical Computer Science*, 38:293–306, 1984.

[Lang *et al.*, 2013] J. Lang, N. Maudet, and M. Polukarov. New results on equilibria in strategic candidacy. In *SAGT'15*, pages 13–25, 2013.

[Niedermeier, 2006] R. Niedermeier. *Invitation to Fixed-Parameter Algorithms*. Oxford University Press, 2006.

[Polukarov *et al.*, 2015] M. Polukarov, S. Obraztsova, Z. Rabinovich, A. Kruglyi, and N. Jennings. Convergence to equilibria in strategic candidacy. In *IJCAI'15*, 2015.

[Rodriguez-Alvarez, 2006a] C. Rodriguez-Alvarez. Candidate stability and probabilistic voting procedures. *Economic Theory*, 27(3):657–677, 2006.

[Rodriguez-Alvarez, 2006b] C. Rodriguez-Alvarez. Candidate stability and voting correspondences. *Social Choice and Welfare*, 27(3):545–570, 2006.

[Tideman, 1987] N. Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4(3):185–206, 1987.