

Convergence to Equilibria in Strategic Candidacy

Maria Polukarov
University of Southampton
United Kingdom

Svetlana Obraztsova
Tel Aviv University
Israel

Zinovi Rabinovich
Mobileye Vision Technologies Ltd.
Israel

Alexander Kruglyi
St.Petersburg State Polytechnical University
Russia

Nicholas R. Jennings
University of Southampton
United Kingdom

Abstract

We study equilibrium dynamics in candidacy games, in which candidates may strategically decide to enter the election or withdraw their candidacy, following their own preferences over possible outcomes. Focusing on games under Plurality, we extend the standard model to allow for situations where voters may refuse to return their votes to those candidates who had previously left the election, should they decide to run again. We show that if at the time when a candidate withdraws his candidacy, with some positive probability each voter takes this candidate out of his future consideration, the process converges with probability 1. This is in sharp contrast with the original model where the very existence of a Nash equilibrium is not guaranteed. We then consider the two extreme cases of this setting, where voters may block a withdrawn candidate with probabilities 0 or 1. In these scenarios, we study the complexity of reaching equilibria from a given initial point, converging to an equilibrium with a predetermined winner or to an equilibrium with a given set of running candidates. Except for one easy case, we show that these problems are NP-complete, even when the initial point is fixed to a natural—truthful—state where all potential candidates stand for election.

1 Introduction

The number of situations where people—and more recently, electronic agents—use voting mechanisms to make collective decisions, is hard to overestimate. Indeed, they get to vote in political elections on different levels, in selecting committees in professional and other organizations, choosing winners of various competitions, rating services and products they had consumed, scheduling meetings, allocating resources and planning joint actions, as well as expressing their opinion on all possible matters in surveys and polls.

Notoriously, most voting mechanisms (a.k.a. rules) are susceptible to various sorts of strategic behavior, shown either by voters misreporting their preferences (manipulation) or by a third party, typically the chair, trying to control the sets

of voters or candidates (voter/candidate control, cloning), influence the votes (bribery and lobbying) or affect the voting rule (agenda control). Finally, the candidates themselves may also have preferences about the outcome of the election and try to affect it by strategically choosing whether to stand for election or not. This latter issue of strategic candidacy we address in our work.

Most of the literature in computational social choice, though, focuses on strategic behaviors by voters and, in particular, on evaluating voting rules by their resistance to such behaviors, using computational complexity as a barrier to them (see, e.g., [Faliszewski *et al.*, 2010] for a survey of these works). Another natural approach is to analyse voting scenarios from a game-theoretic perspective, viewing strategic parties as players and examining possible stable outcomes of their interaction (i.e., Nash equilibria).

However, even though the first model for games with strategic voters dates back to the 1960's [Farquharson, 1969], this line of research only in recent decades has received serious attention in the algorithmic game theory and the social choice communities. A few works, in particular, consider Plurality voting games, characterizing their Nash equilibria (although under the restrictive assumption of single-peaked preferences) [Feddersen *et al.*, 1990], and examining their dominant strategy equilibria [Dhillon and Lockwood, 2004]. In [Messner and Polborn, 2007], the authors suggest a variation of a strong equilibrium and explore conditions for its existence and uniqueness. The most relevant to our work, however, is the paper by Meir *et al.* [Meir *et al.*, 2010] that studies equilibrium dynamics, based on myopic improving moves by single voters. The authors, in particular, demonstrate that convergence to equilibrium is guaranteed from any initial state, if players choose their best possible moves at every step and ties are broken lexicographically.

The literature on strategic candidacy is even more scarce. It starts with the works by Dutta *et al.* [Dutta *et al.*, 2001; 2002], who formulate the game and show that no reasonable (i.e., non-dictatorial and unanimous) voting rule can guarantee stability of the truthful¹ state where all candidates enter

¹Under the self-supporting candidate preferences assumption

the election. The authors also demonstrate examples of voting trees, where candidacy games may have no pure strategy Nash equilibrium. Going further in this direction, Lang et al. [Lang et al., 2013] prove more general results on the existence of equilibria in candidacy games, both positive and negative. Specifically, they show that in the case of 4 candidates a pure strategy equilibrium always exists for Condorcet-consistent rules; however, in the case of more than 4 candidates these rules may (Copeland) or may not (Maximin) admit such equilibria. Importantly, and particularly relevant to our paper, for Plurality voting with 4 or more candidates, there are candidacy games without pure equilibria.

Against this background, in this work we combine the study of equilibrium dynamics and strategic candidacy under Plurality. Indeed, the fact that pure strategy equilibria may not, in general, exist in these games, raises the question of whether one exists for a given preference profile and how (if at all) it can be reached dynamically from a given initial state. In practice, such dynamic processes often occur in online art, photo, literature or similar competitions, where each contestant has to present examples of their work, and can remove or replace any of them at any time by a given deadline. Other examples include sale campaigns where online shops choose items to be advertised on the main page and can replace or remove them at any time during the sales period, competitions among service providers for a public project, deciding on a team working plan, and so on.

In such scenarios, one may care not only about converging to some equilibrium point, but reaching one where a particular candidate wins the election or a certain set of candidates still run. For instance, in many common scenarios where the candidates, once left, cannot enter the election again, the question of convergence becomes trivial as any sequence of improving moves is finite. However, reaching a state with predetermined characteristics appears computationally hard.

Interestingly, the two scenarios where the candidates are free to leave and enter the election any time, or only can leave once and never enter again, can be viewed as two extreme cases of a more general setting where voters may feel discouraged by candidates leaving the election and only return their votes to them, should they renew their candidacy, with some probability. It turns out that even assuming some arbitrarily small positive probability for each voter to refuse re-voting for a once withdrawn candidate, is sufficient to achieve convergence to equilibria with probability 1.

Contribution. The paper makes the following contributions:

1. We introduce a *dynamic candidacy game* model with *refusing voters*, where, with probability p_v , a voter v rejects a candidate who withdraws his candidacy. We show these games converge with probability 1, for any $p_v > 0$.
2. We define three decision problems, termed NE, WINNER and SET, which, given a profile of preferences and an initial state, decide whether there exists an improvement path leading to a Nash equilibrium, to an equilibrium with a predetermined winner or to one with a given set of running candidates, respectively.

(defined in Section 2).

3. For each of the three problems above, we consider its computational complexity in two variants, indexed 1 and 0, corresponding to refusing probabilities $p_v = 1$ and $p_v = 0, \forall v$. Except for NE1 where convergence is trivial, we show that these problems are NP-hard.
4. Finally, we establish NP-hardness of deciding the existence of an equilibrium with a predetermined winner in a static game, the problem that we term \exists WINNER. This is in contrast with games with strategic voters where this problem is trivially solvable.

Some proofs are omitted due to space limitations.

2 Model and preliminaries

We first recall some basic notions from voting theory and define candidacy games, following notation in [Lang et al., 2013]. We then describe the dynamic setting based on improvement moves by single candidates.

2.1 Candidacy games

There is a set of voters electing from a set of candidates. A single vote is a strict ranking of the candidates. A voting rule takes all the votes as input, and produces an outcome—a candidate that *wins* the election. Although voting rules are usually defined for a fixed number of candidates, for strategic candidacy settings the definition is naturally extended to an arbitrary finite number of candidates.

Formally, we have a set $C = \{c_1, c_2, \dots, c_{|C|}\}$ of *potential candidates*, and a set $V = \{v_1, v_2, \dots, v_{|V|}\}$ of *voters*. It is assumed that C and V are disjoint. Each voter $v \in V$ has a *preference* relation, $P_v \in \mathcal{L}(C)$, over the candidates, where for any finite set X , $\mathcal{L}(X)$ denotes the set of all strict linear orders on X . For any order $L \in \mathcal{L}(X)$ and $x, x' \in X$ we write $x \succ_L x'$ if L ranks x higher than x' . The combination $P^V = (P_v)_{v \in V}$ of all the voters' preferences defines their *preference profile*. Furthermore, each candidate $c \in C$ also has a preference ordering over the candidates, $P_c \in \mathcal{L}(C)$. It is often assumed that the candidates' preferences are *self-supporting*—that is, the candidates rank themselves at the top of their ordering.² Let $P^C = (P_c)_{c \in C}$ denote the candidates' preference profile, and $P = (P^V, P^C)$ represent the full profile of preferences of both the voters and the candidates.

Following P^C , the potential candidates may decide to enter the election or withdraw their candidacy. Thus, the voters will only need to express their preferences over a subset $A \subseteq C$ of *actual* candidates that will have chosen to participate in the election, and we denote by $P^{\downarrow A} \in \mathcal{L}(A)$ the restriction of P^V to A . Each voter v submits a *vote* (or *ballot*) $b_v^{\downarrow A} \in \mathcal{L}(A)$. For time being, we assume that the voters are *sincere*, that is, $b_v^{\downarrow A} = P_v^{\downarrow A}$.³ A voting profile $\mathbf{b}^{\downarrow A} = (b_v^{\downarrow A})_{v \in V}$ is a vector

²Even though our results do not rely on this assumption, we find it rather natural. For instance, our hardness proofs are valid, in particular, for the case where the initial state is fixed to be one with all potential candidates standing for election, which under self-supporting preferences corresponds to the truthful state.

³In Section 3, we extend the model to scenarios where a voter—while still *not* being strategic—may feel discouraged, if his favored candidates leave the election, and would refuse to vote for them again, should they decide to re-enter the election.

of votes, one for each agent.

Given a set of actual candidates $A \subseteq C$, a *voting rule* $\mathcal{F} : \mathcal{L}(A)^{|V|} \rightarrow 2^A$ takes a voting profile as input, and produces an outcome—a nonempty subset of candidates, called *the winners* of election. Here we consider *resolute* voting rules $\mathcal{F} : \mathcal{L}(A)^{|V|} \rightarrow A$, which always return a single winner. That is, given their irresolute version, we assume that ties are broken according to a fixed *tie-breaking rule*. Specifically, we assume lexicographic tie-breaking—i.e., ties are broken according to some predetermined priority relation over the candidates. Since a voting rule is applied to varying sets of actual candidates, it is assumed 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 . In this paper, we particularly focus on *Plurality* voting rule, which decides the winner to be the candidate that is ranked first by most voters. Hence, we can simplify notation by restricting each ballot to specify only a single candidate—a voter’s *top choice* candidate among all the running candidates.

Each such voting setting induces a natural *game form*, where the set of players is given by the set of potential candidates C , 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 $(s_c)_{c \in C}$, where $s_c \in \{0, 1\}$. The outcome of a state s is $\mathcal{F}(\mathbf{b}^{\downarrow A})$ where $c \in A$ if and only if $s_c = 1$. Coupled with candidates’ preferences, this defines a normal form game with $|C|$ players, $\Gamma = \langle C, P, \mathcal{F} \rangle$, where $P = (P^V, P^C)$. Here, player c prefers outcome $\mathcal{F}(f)$ over outcome $\mathcal{F}(f')$ if P_c^C ranks $\mathcal{F}(f)$ higher than $\mathcal{F}(f')$.

2.2 Equilibrium dynamics

Having a normal form game defined, we can now apply standard game-theoretic solution concepts. Let $\Gamma = \langle C, P, \mathcal{F} \rangle$ be a candidacy game, and let s be a state in Γ . A player $c \in C$ has an *improving move* in s if there is s'_c such that c prefers $\Gamma(s_{-c}, s'_c)$ over $\Gamma(s)$. A (*pure strategy*) *Nash equilibrium* [Nash, 1951] is a state that has no such improving moves.

A *path* in $\{0, 1\}^{|C|}$ is a sequence $(s^0 \rightarrow s^1 \rightarrow \dots)$ of states such that for every $k \geq 1$ there exists a unique player, say candidate c , such that $s^k = (s'_c, s_{-c}^{k-1})$ for $s'_c \neq s_{-c}^{k-1}$ in $\{0, 1\}$. It is an *improvement path* if for all $k \geq 1$ it holds that $s^{k-1} \xrightarrow{c} s^k$ is an improvement move, where c is the unique deviator at step k . The setting of *dynamic candidacy* is based on myopic improvement dynamics as above: the candidates start by announcing some initial state, and then proceed and change their candidacy status in turns, one at a time, up until no one has an objection to the current outcome. We neither make assumptions on the initial profile s^0 , nor restrict the order, in which the players apply their moves, or the number of times for each candidate to change his status.

2.3 Reachable states

While it is known that a Nash equilibrium may not, in general, exist for candidacy games under Plurality [Lang *et al.*,

2013], the question remains of whether one is guaranteed for a given profile of preferences, and how it can be obtained. In particular, it is interesting to know whether an equilibrium state can be reached by a natural dynamic process based on improvement moves by single candidates as above. We call such a state *reachable*. What is even more important though, is to find out what candidates would still stand for election in the end of the process, or just who would be the final winner.

To this end, here we define and investigate the computational complexity of the following decision problems:

- NE. Given a candidacy game and its initial state s^0 , is there an equilibrium state, reachable from s^0 ?
- WINNER. Given an initial state s^0 of a candidacy game and a fixed candidate c , is there an equilibrium state, reachable from s^0 , in which c wins the election?
- SET. Given an initial state s^0 of a candidacy game and an equilibrium state s , is s reachable from s^0 ?

For each of these problems, we consider two of its variants depending on whether a candidate believes or not getting his previous votes back again, should he re-enter the election after having left it once. These are two extreme cases of a more general model, which we present next, where a voter refuses to re-vote for such a candidate with some known probability.

3 Refusing voters

Here we extend the dynamic candidacy setting to scenarios where withdrawals may cause the voters to ignore their once favorite candidates in the future. This is because the voters may either feel discouraged by and lose their trust and interest in the candidates who “let them down” by leaving the election, or simply stop updating their information about the withdrawn candidates and hence avoid making uninformed decisions. As we show, in this case convergence is guaranteed with probability 1.

Formally, we assume that each time that a candidate c withdraws his candidacy from the election, each voter v decides to block this candidate with probability $p_v \in [0, 1]$, independently of other voters (unless he already blocked this candidate in previous steps).⁴ The only case, in which a voter v may return his vote to such a banned candidate c , is when the candidate for whom v has been currently voting decides to leave the election, and the voter also finds himself to have banned all the remaining candidates still standing for election, so he must reconsider and vote for one of them (that he prefers the most) again. We assume though that even in this case candidate c is formally considered as “banned” by voter v —that is, as soon as another candidate that had not been previously blocked by v enters the election, voter v moves his vote away from c .⁵ Importantly, the voters’ decisions are *not* strategic,

⁴Alternatively, one could consider a model where only those voters who currently support a given candidate, may decide to block him after his withdrawal. All our results hold for both these variants of the setting.

⁵Alternatively, one could assume that when voter v returns his vote to candidate c , he formally “unbans” him. All our results hold for both these variants of the setting.

that is, the voters always follow their true preferences, even if restricted to only a subset of available candidates.

We term a pair $(\Gamma, (p_v)_{v \in V})$, where Γ is a candidacy game and p_v , for $v \in V$, are probabilities as above, a *candidacy game with refusing voters*. Note that a state in this game, as well as a voting ballot and hence, a game outcome, are determined not only by the set of actual candidates, but also by the sets of banned candidates, one for each voter. That is, state S is a tuple (A, B) , where $A \subseteq C$ is a set of actual candidates and $B = (B_v)_{v \in V}$ where $\forall v \in V, B_v \subseteq C$ is a subset of potential candidates, which are banned by voter v . The corresponding ballot $b_v^{\downarrow A, B} \in \mathcal{L}(A)$ is then obtained by placing subset $A \setminus B_v$ in the top positions of the ballot, and subset $A \cap B_v$ —in the bottom positions (i.e., $\forall x \in A \setminus B_v$ and $\forall y \in A \cap B_v$ we have $x \succ_{b_v^{\downarrow A, B}} y$), while the internal order of candidates in each of the two subsets is determined by the true preference order P_v : i.e., $\forall x, x' \in A \setminus B_v, x \succ_{b_v^{\downarrow A, B}} x' \Leftrightarrow x \succ_{P_v} x'$ and $\forall y, y' \in A \cap B_v, y \succ_{b_v^{\downarrow A, B}} y' \Leftrightarrow y \succ_{P_v} y'$. The outcome of a state S is $\mathcal{F}(\mathbf{b}^{\downarrow A, B})$.

Next, we demonstrate that for any *positive* probabilities p_v , this game converges with probability 1 from any initial state. Our proof involves constructing a Markov chain that corresponds to a stochastic process based on the candidates' improvement moves and the voters' blocking actions, and showing that this chain is absorbing⁶.

Theorem 1. *Let $(\Gamma, (p_v)_{v \in V})$ be a candidacy game with refusing voters under Plurality. If $p_v > 0, \forall v \in V$, then with probability 1 any improvement path is finite.*

Proof. Consider a Markov chain over a set of states $S = \{S^1, S^2, \dots, S^{|S|}\}$ of the candidacy game with refusing voters. The process starts in one of these states and moves from one state to another at each step. Given a current state S^i , let p^{ij} denote the transition probability of moving from state S^i to state S^j at the next step; with probability p^{ii} the process remains in the same state S^i . Let us now determine these transition probabilities.

For each state S^i , let $C^i \subseteq C$ be the set of candidates who have an improving move from S^i . Note that every candidate c has exactly one possible move at each step, when the preference profile P and the vector of banned candidate sets $(B_v)_{v \in V}$ define whether this move is improving for c .

Assume first that C^i is non-empty. Then, one of the players in C^i will apply his improving move, and the process will move to another state S^j ; that is, $p^{ii} = 0$. To calculate p^{ij} , for each $c \in C^i$, let $p_c^i = \frac{1}{|C^i|}$ be the probability that c is randomly selected to move from state S^i ; we have $p_c^i > 0, \forall c \in C^i$, and $\sum_{c \in C^i} p_c^i = 1$. If $c \notin A^i$ then c enters the election at this step, and the process moves to state S^j where $A^j = A^i \cup \{c\}$ and $B_v^j = B_v^i$, for each $v \in V$. The corresponding transition probability is $p^{ij} = p_c^i$. Otherwise, if $c \in A^i$, then c withdraws his candidacy, and the process moves to state S^j where $A^j = A^i \setminus \{c\}$ and for each $v \in V$,

either $B_v^j = B_v^i$ or $B_v^j = B_v^i \cup \{c\}$. The corresponding transition probability is $p^{ij} = p_c^i \prod_{v \in V} p_v^{ij}$ where for each $v \in V$ the probabilities p_v^{ij} are given as follows. If c is banned by voter v in S^i , that is, if $c \in B_v^i$, then $p_v^{ij} = 1$ for $B_v^j = B_v^i$ (note that $B_v^i \cup \{c\} = B_v^i$ in this case). Otherwise, if c is not banned by voter v in S^i , then for $B_v^j = B_v^i$ we have $p_v^{ij} = 1 - p_v$, and for $B_v^j = B_v^i \cup \{c\}$ we have $p_v^{ij} = p_v$.^{7,8}

Finally, if $C^i = \emptyset$ then the process stays in S^i with probability $p^{ii} = 1$ (i.e., state S^i is absorbing). The other transition probabilities are zeroes. Observe that these probabilities do not depend upon which states the chain was in before the current state S^i , so the Markov property does indeed hold.

We now turn to show that this chain is absorbing. That is, it has at least one absorbing state, and it is possible to reach such a state from every state in the chain. To this end, from any initial state S^0 , we construct a path $(S^0 \rightarrow S^1 \rightarrow \dots \rightarrow S^t)$ with an absorbing terminal state S^t and positive transition probabilities. At each step $i = 1, \dots, t$ on this path, if there exists a candidate c in A^{i-1} who wants to withdraw his candidacy, let S^i be the state with $A^i = A^{i-1} \setminus \{c\}$ where as many voters as possible have banned c after his withdrawal. If no candidate wants to leave the election in S^{i-1} then choose any candidate $c \in C^{i-1} \setminus A^{i-1}$ who would like to enter the election and let S^i be the state with $A^i = A^{i-1} \cup \{c\}$ and $B_v^i = B_v^{i-1}$ for each $v \in V$. From above, the transition probability $p^{i-1, i}$ is positive (unless there are no improving moves from state S^{i-1} , in which case we are done). Now, note that any time that a candidate enters the election, he gets some votes (otherwise, he cannot change the outcome and this is not an improving move). However, by our path definition, if/when he withdraws his candidacy, all the voters block him, so he cannot get any more votes should he enter the election again and so will stay aside. Indeed, such a candidate could only hope to get votes from those voters who have banned all the candidates by this step, but, as we mentioned before, the only case, in which a voter may return his vote to a banned candidate, is when another candidate whom this voter has been currently supporting (who is also his last unbanned candidate) decides to leave the election, and the voter has to move his vote to those currently standing for election. However, the candidate that re-enters the election cannot get his vote. Thus, each candidate can enter the election at most once (and only if he was not running in the initial state), and will never re-enter again since having left it (which also can happen at most once).⁹ Since we only have $|C|$ potential

⁷If for some voter v the withdrawing candidate c was his last unbanned candidate, v must return his vote to candidate c' whom he prefers the most in A^j . In the setting where v unbans c' in such a scenario, the process moves to state S^j with $B_v^j = B_v^i \setminus \{c'\}$ or $B_v^j = B_v^i \cup \{c\} \setminus \{c'\}$, with the same transition probabilities (that is, $p_v^{ij} = 1 - p_v$ and $p_v^{ij} = p_v$, respectively).

⁸In the case where only those voters who support candidate c in S^i can block him after his withdrawal, the probability p_v^{ij} also depends on whether v votes for c or not.

⁹In the case where only those voters who support a given candidate in the current state can block him after his withdrawal, a candidate can re-enter the election at most $|V|$ times on our path. Indeed, by our path definition, if/when he withdraws his candidacy, all the voters who can block him will do so, and there is always at least

⁶That is, it has at least one absorbing state (which transits to itself with probability 1), and it is possible to reach some absorbing state from every state in the chain.

candidates, after at most $2|C|$ steps¹⁰ (each with a positive probability), the path will reach its terminal state S^t where no candidate wishes to join or leave the election, and hence, $p^{tt} = 1$ (that is, S^t is absorbing).

Finally, knowing that in an absorbing Markov chain, the probability for the process to be absorbed is 1 (see e.g., [John G. Kemeny and Snell, 1976]), completes our proof. \square

In the following section, we consider the two extremes of Plurality candidacy game with refusing voters ($\Gamma, (p_v)_{v \in V}$), where $p_v = 0$ or $p_v = 1, \forall v \in V$. These cases correspond to two natural instances of the game; the former coincides with the original model where voters follow no other considerations but their preference orders; in the latter, since all the voters block each withdrawn candidate, no one has incentives to ever renew their candidacy. For these two cases, we study the computational complexity of NE, WINNER and SET.

4 Complexity of reaching equilibria

We start by showing that the WINNER problem is NP-hard in both cases. Our proofs involve reducing from Exact 3-Set Cover (X3C) and Restricted Exact 3-Set Cover (RX3C). We note that our reductions hold, in particular, for the special case where the initial state is truthful. In addition, in case of WINNER0, the winner in an equilibrium state that corresponds to a solution of the reduced problem, is unique. The latter then implies the computational hardness of NE0 and of the problem of deciding the very existence of a Nash equilibrium with a predetermined winner, regardless of dynamic processes, which we denote \exists WINNER.

For completeness, we first define X3C and RX3C:

- X3C. Given a set $U = \{u_1, \dots, u_{3m}\}$ and a family $Z = \{z_1, \dots, z_n\}$ of triples $z_j = \{u_{j_1}, u_{j_2}, u_{j_3}\} \subseteq U, j = 1, \dots, n$, is there a subfamily Z' of Z such that every element in U is contained in exactly one triple of Z' ?
- RX3C. Same as X3C, with the additional restriction that each element of U appears exactly in three triples.

We are now ready to state our results.

Theorem 2. WINNER1 is NP-complete.

Proof. First, observe that the problem is in NP. Indeed, having a state s coupled with a path ($s^0 \rightarrow \dots \rightarrow s$), it takes polynomial time to check whether s is an equilibrium state with a given winner and whether ($s^0 \rightarrow \dots \rightarrow s$) is an improvement path that leads from the initial state s_0 to s . Note that each valid improvement path may only contain withdrawals, and so is of polynomial length.

To show hardness, we reduce from X3C with $n \geq 3m$, by constructing an instance of WINNER1 as follows. Let $C = Z \cup Q \cup U \cup D \cup \{w_0, w_1, w\}$ be the set of candidates where Z and Q each contains n elements as the number of triples in X3C, U is a set of $3m$ candidates corresponding to the ground set in X3C, D is a (large) set of dummy players,

one such voter. That is, after $|V|$ entrances, the candidate gets to the point where he is banned by all the voters.

¹⁰ $2|V||C|$ steps, in the case with only supporting voters being able to block a candidate.

and w_0, w_1, w are single distinguished candidates. The candidates' preferences are in Table 1.

There is a set V of voters, divided into 7 blocks. As can be seen from the voters' preference profile in Table 2, Blocks 1 and 4 each contain n voters, Block 2 has $2n$ voters, and Block 3 has $3n$. Furthermore, there are $3m(f - 1)$ voters in Block 5, $n(f - 1) + 2f - 3m$ voters in Block 6, and f voters in Block 7, where f is a large constant (it is sufficient to have $f > 7(n^3 + m)$). For any subset $X \subseteq C$, by X cycle we denote a fragment of preference lists where members of X appear in the same cycling order, with the starting point being alternated—that is, for $X = \{x_1, x_2, \dots, x_k\}$, we have $(x_1x_2 \dots x_k), (x_2 \dots x_kx_1)$, and so on.

The tie-breaking rule prioritizes w over all other candidates, who can be ordered arbitrarily between them.

Let the initial state be $s^0 = (1, 1, \dots, 1)$ where all the candidates run. We show that if X3C has a solution, then there exists an improvement path from s^0 to an equilibrium state s where the winner is w . Otherwise, there is no such reachable equilibrium state. At s^0 , the winner is w_0 with f points. Note that since candidates never have an incentive to re-enter the election after having left it once (as $p = 1$), the score of the winner cannot decrease along an improvement path, as leaving candidates only give their points to remaining candidates and never take them back.

Consider now the first n steps of the process. We show that only candidates $z_j \in Z$ can move at these steps. Indeed, look at the first leaving candidate who is not in Z . It cannot be any dummy player $d_i \in D$, as there are no dummies in the top of the preference profile, and only those in the top for at least some of the voters can pass their points to remaining candidates and change the outcome. Now, anyone of the players $q_j \in Q, u_i \in U, w_0, w_1$ or w , can only give a point to a dummy player (or to w_1 in Blocks 1-4, in case they have been previously “opened” (i.e., got to the top of the preference list) by some z_j that withdrawn his candidacy). But all of them prefer dummy players the least, so they wouldn't make a move (in fact, the dummies even have no chances to win the election, so there's no point to do it for them at all.) As for w_1 , he is also of low priority for q_j 's and u_i 's—he only beats w_0 , but the latter stops being a winner after the very first improving move, so in the later moves changing the winner to w_1 cannot be beneficial for q_j 's and u_i . Hence, only $z_j \in Z$ can move at first n steps, making their favourite q_j 's or u_i 's win the election.

As we have just mentioned, after the first such step, w_0 loses and can't become a winner ever again, but he keeps its f points, so a new winner must receive at least f points. Now look at step m , where both w_1 and w reach f points. In the next $n - m$ steps (if there are such steps) w will be getting 1 additional point at each step, and w_1 will be receiving 2 points. Hence, w will never become a winner, unless the process stops at step m where w wins the election with f points (by the tie-breaking with w_1). Note though, that at each of the first m steps, a withdrawn candidate z_j was giving a point to some candidate u_i , initially having the score of $f - 1$. That is, if w is the winner after step m , then no u_i has received more than one additional point—i.e., exactly m of them have been opened exactly once, and we have an

Z block			Q block			U block			D block			w_0	w_1	w
z_1	\dots	z_n	q_1	\dots	q_n	u_1	\dots	u_{3m}	d_1	\dots	$d_{ D }$	w_0	w_1	w
q_1	\dots	q_n	w			w			w			w_1	w_1	w_1
w			U			$U \setminus \{u_i\}$			w_1			U	U	U
U			Z			Z			U			Z	Z	Z
$Z \setminus \{z_j\}$			$Q \setminus \{q_j\}$			Q			Z			Q	Q	Q
$Q \setminus \{q_j\}$			w_1			w_1			Q			w	w_0	w_0
w_1			w_0			w_0			w_0			D	D	D
w_0			D			D			$D \setminus \{d_i\}$					

Table 1: WINNER1. Candidates' preferences.

Block 1				Block 2				Block 3				Block 4							
z_1	z_2	\dots	z_n	z_1	z_1	z_2	\dots	z_1	z_1	z_1	\dots	z_n	z_n	z_n	z_n	z_1	z_2	\dots	z_n
w	w	\dots	w	w_1	w_1	w_1	\dots	w_1	w_1	w_1	\dots	w_1	w_1	w_1	w_1	q_1	q_2	\dots	q_n
w_1	w_1	\dots	w_1					w_1	w_1	w_1	\dots	w_1	w_1	w_1	w_1	w_1	w_1	\dots	w_1
D cycle				D cycle				D cycle				D cycle				D cycle			
$Z \setminus \{z_j\}$ cycle				$Z \setminus \{z_j\}$ cycle				$Z \setminus \{z_j\}$ cycle				$Z \setminus \{z_j\}$ cycle				$Z \setminus \{z_j\}$ cycle			
U cycle				U cycle				$U \setminus \{u_{j_i}\}$ cycle				U cycle				U cycle			
Q cycle				Q cycle				Q cycle				Q cycle				$Q \setminus \{q_j\}$ cycle			
w_0	w_0	\dots	w_0	w_0	w_0	w_0	\dots	w_0	w_0	w_0	\dots	w_0	w_0	w_0	w_0	w_0	w_0	\dots	w_0
w	w	\dots	w	w	w	w	\dots	w	w	w	\dots	w	w	w	w	w	w	\dots	w

Block 5				Block 6				Block 7	
$u_1 \dots u_1$	\dots	$u_{3m} \dots u_{3m}$		$q_1 \dots q_1$	\dots	$q_n \dots q_n$	$w_1 \dots w_1$	$w \dots w$	$w_0 \dots w_0$
$f-1$		$f-1$		$f-1$		$f-1$	$f-2m$	$f-m$	f
				D cycle		D cycle			D cycle
				Z cycle		Z cycle			Z cycle
				$U \setminus \{u_i\}$ cycle		U cycle			U cycle
				Q cycle					Q cycle
$w_1 \dots w_1$	\dots	$w_1 \dots w_1$		$w_1 \dots w_1$	\dots	$w_1 \dots w_1$	$w_1 \dots w_1$		$w_1 \dots w_1$
$w_0 \dots w_0$	\dots	$w_0 \dots w_0$		$w_0 \dots w_0$	\dots	$w_0 \dots w_0$	$w_0 \dots w_0$		$w_0 \dots w_0$
$w \dots w$	\dots	$w \dots w$		$w \dots w$	\dots	$w \dots w$	$w \dots w$		$w \dots w$

Table 2: WINNER1. Voters' preferences.

exact cover in X3C. The reverse direction is trivial: take a cover, and let the corresponding z_j 's leave one after another, in descending order. \square

We now turn to the case with $p_v = 0, \forall v \in V$. Since now the candidates may not only leave, but also re-enter the election, our proof requires a much more involved hardness reduction. Specifically, we reduce from RX3C, building on the counterexample for the existence of equilibria under Plurality presented in [Lang *et al.*, 2013], using it as a sub-block in our constructed preference profile, to help us lead the process into a cycle when the reduced problem has no solution.

Theorem 3. WINNER0 is NP-hard.

Note that both Theorem 2 and Theorem 3 use the truthful state $s^0 = (1, 1, \dots, 1)$ as initial point, so their results hold, in particular, for this important special case. Also, from the proof of Theorem 3 we derive another useful observation.

Lemma 1. There is an instance of WINNER0 with a fixed candidate, w , who wins the election in any equilibrium state if and only if there is an exact cover of the reduced RX3C.

Now, let \exists WINNER denote the following decision problem: Given a preference profile P and a fixed candidate c , is there an equilibrium state s , in which c wins the election? The following Theorem 4 is then a direct corollary of Lemma 1.

Theorem 4. \exists WINNER and NE0 are NP-hard.

Note that for NE, this result is in contrast with the case of $p_v = 1$ where a stable state is easy to reach.

Theorem 5. NE1 always returns “yes”. Moreover, a stable state is reachable in linear time.

Finally, we modify our reductions for the WINNER problem to show NP-hardness of SET.

Theorem 6. SET0 and SET1 are NP-hard.

5 Discussion and Future Work

In this paper, we initiate the study of equilibrium dynamics in candidacy games. While such dynamic processes have recently been in the focus of active research in the context of strategic voting, the case where candidates behave strategically remained unexplored so far. Naturally, we first focus on the simple Plurality rule.

Remarkably, solution sets for voting and candidacy games possess very different properties, which also imply differences between their corresponding dynamic processes. Thus, voting games have multiple equilibria, some of which are highly undesirable (e.g., where all the voters select the same—and least preferred by all—candidate). However, restricting the set of equilibria to only those, reachable dynamically from the truthful state, appears useful in excluding such bad equilibria [Branzei *et al.*, 2013]. In contrast, for candidacy games the very existence of an equilibrium is not guaranteed under Plurality rule, thus implying the need of studying the existence and reachability of equilibria for given profiles of preferences. Also, in this context, besides seeking a state with a particular winner, it is sensible to look for a state with a certain set of running candidates. The main difference

though, is in that equilibrium dynamics here have no aim of eliminating equilibria but merely finding them.

Importantly, adding dynamics to searching for equilibria, also has different impact on the complexity of the problem. Thus, finding any stable state is easy in voting games, while checking that a given equilibrium is reachable is NP-hard; on the other hand, for truth-biased voters (who always vote truthfully unless can change the outcome in their favour by deviating), the complexity of these two problems gets reversed [Rabinovich *et al.*, 2014]. In contrast, as this paper shows, in candidacy games the problems of finding a stable state with a given winner, or reaching such a state dynamically, are both NP-hard.

On the other hand, notice that converging to a state with a given set of actual candidates, which is NP-hard in the original case with no refusing voters, becomes polynomial time solvable when voters refuse to return their votes to withdrawn candidates with probability 1. The question is then whether the problem can be efficiently solved with high probability for any refusing probability $p_v > 0$. In other words, (when) is it possible to circumvent this computational hardness in the model with refusing voters.

We thus believe that our work makes a first step in several exciting directions. First, we hope that studying dynamic processes can shed light on the properties of candidacy games under Condorcet-consistent rules, which all admit equilibria for 4 candidates, but split as the number of candidates grows. Second, given the hardness results we presented here, one may seek opportunities for getting positive results in terms of computational complexity, e.g., by restricting the space of preference profiles to single-crossed or single-peaked domains. Furthermore, the model with refusing voters should be further extended (e.g., to scenarios where voters can block a candidate for only a period of time or unblock them with some probability at each step). Finally, it is interesting to investigate equilibrium dynamics in the setting with both candidates and voters being strategic [Brill and Conitzer, 2015 forthcoming]. While there are multiple equilibrium states in this case, the question of whether they (and which of them) are reachable remains open.

Acknowledgments

This work has been cofinanced by the UK Research Council for project ORCHID—grant EP/I011587/1, the COST Action IC1205 on Computational Social Choice, and RFFI 14-01-00156-a. Svetlana Obratsova also gratefully acknowledges funding from the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013)—ERC grant agreement number 337122.

References

- [Branzei *et al.*, 2013] Simina Branzei, Ionanis Caragiannis, Jamie Morgenstern, and Ariel D. Procaccia. How bad is selfish voting? In *AAAI*, 2013.
- [Brill and Conitzer, 2015 forthcoming] M. Brill and V. Conitzer. Strategic voting and strategic candidacy. In *AAAI*, 2015, forthcoming.
- [Dhillon and Lockwood, 2004] Amrita Dhillon and Ben Lockwood. When are plurality rule voting games dominance-solvable? *Games and Economic Behavior*, 46(1):55–75, 2004.
- [Dutta *et al.*, 2001] B. Dutta, M. Le Breton, and M. O. Jackson. Strategic candidacy and voting procedures. *Econometrica*, 69:1013–1037, 2001.
- [Dutta *et al.*, 2002] B. Dutta, M. Le Breton, and M. O. Jackson. Voting by successive elimination and strategic candidacy in committees. *Journal of Economic Theory*, 103:190–218, 2002.
- [Faliszewski *et al.*, 2010] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. Using complexity to protect elections. *Communications of the ACM*, 53(11):74–82, 2010.
- [Farquharson, 1969] Robin Farquharson. *Theory of Voting*. Yale University Press, 1969.
- [Feddersen *et al.*, 1990] Timothy J. Feddersen, Itai Sened, and Stephen G. Wright. Rational voting and candidate entry under plurality rule. *American Journal of Political Science*, 34(4):1005–1016, 1990.
- [John G. Kemény and Snell, 1976] James Laurie Snell John G. Kemény John G. Kemény, James Laurie Snell John G. Kemény and J. Laurie Snell. *Finite Markov Chains*. Springer, 1976.
- [Lang *et al.*, 2013] Jerome Lang, Nicolas Maudet, and Maria Polukarov. New results on equilibria in strategic candidacy. In *SAGT*, pages 13–25, 2013.
- [Meir *et al.*, 2010] Reshef Meir, Maria Polukarov, Jeffrey S. Rosenschein, and Nicholas R. Jennings. Convergence to equilibria of plurality voting. In *AAAI*, pages 823–828, 2010.
- [Messner and Polborn, 2007] M. Messner and M. K. Polborn. Strong and coalition-proof political equilibria under plurality and runoff rule. *International Journal of Game Theory*, 35(2):287–314, 2007.
- [Nash, 1951] John Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951.
- [Rabinovich *et al.*, 2014] Zinovi Rabinovich, Svetlana Obratsova, Omer Lev, Evangelos Markakis, and Jeffrey S. Rosenschein. Analysis of equilibria in iterative voting schemes. In *COMSOC*, 2014.