Strategic Voting and Strategic Candidacy

Markus Brill and Vincent Conitzer

Abstract

Models of strategic candidacy analyze the incentives of candidates to run in an election. Most work on this topic assumes that strategizing only takes place among candidates, whereas voters vote truthfully. In this paper, we extend the analysis to also include strategic behavior on the part of the voters. (We also study cases where only candidates or only voters are strategic.) We consider two settings in which strategic voting is well-defined and has a natural interpretation: majority-consistent voting with single-peaked preferences and voting by successive elimination. In the former setting, we analyze the type of strategic behavior required in order to guarantee desirable voting outcomes. In the latter setting, we determine the complexity of computing the set of potential outcomes if both candidates and voters act strategically.

1 Introduction

When analyzing voting rules, the set of candidates is usually assumed to be fixed. In a pathbreaking paper, Dutta, Jackson, and Le Breton [8] have initiated the study of *strategic candidacy* by accounting for candidates' incentives to run in an election. They assumed that candidates have preferences over other candidates and defined a voting rule to be *candidate stable* if no candidate ever has an incentive not to run. In this model, it is assumed that every candidate prefers himself to all other candidates. Therefore, the winner of an election never has an incentive not to run. Non-winning candidates, on the other hand, might be able to alter the winner by leaving the election. Dutta et al. [8] showed that, under mild conditions, no non-dictatorial rule is candidate stable.

This result naturally leads to the question of how voting outcomes are affected by candidates' incentives. It is straightforward to model strategic candidacy as a two-stage game. At the first stage, each candidate decides whether to run in the election or not. At the second stage, each voter casts a ballot containing a ranking of the running candidates. When analyzing this game, an important ingredient is the assumed voter behavior. That is, what assumptions are made about the votes in the second stage, conditional on the set of running candidates?

Most papers on strategic candidacy (see Section 2 for an overview) assume that voters vote truthfully, i.e., their reported ranking for any given subset of candidates corresponds to their true preferences, restricted to that subset. However, it is well known that this is an unrealistic assumption [15, 30]. It is therefore natural to account for strategic behavior on the part of the voters as well. Thus, in the models we consider, both candidates and voters act strategically.

The technical problem in accounting for strategic voting is that, generally speaking, too many voting equilibria exist [23, 6]. If we only consider Nash equilibria, then any profile of votes for which no single voter can change the outcome is an equilibrium. In some cases, a straightforward refinement rules out many of the equilibria [7, 32, 24]. For example, in a majority election between two candidates, it is natural to rule out the strange equilibria where some voters play the weakly dominated strategy of voting for their less-preferred candidate. But this reasoning does not generally extend to more than two candidates. In this paper, we focus on two settings that admit natural equilibrium refinements.

The first setting is that of single-peaked preferences [4]. It is well known that, if the number of voters is odd, this domain restriction guarantees the existence of a Condorcet winner (namely, the median) and admits a strategyproof and Condorcet-consistent voting rule (namely, the median rule) [21]. Dutta et al. [8] observed that any Condorcet-consistent rule is candidate-stable in any domain that guarantees the existence of a Condorcet winner.[1] We study the effect of strategic candidacy with single-peaked preferences when the voting rule is *not* Condorcet-consistent. Our motivation is that the voting rules that are most widely used in practice, *plurality*, *plurality with runoff*, and *single transferable vote* (STV), may fail to select the Condorcet winner, even for single-peaked preferences. We consider the class of *majority-consistent* voting rules, which are rules that, if there is a candidate that is ranked first by more than half the voters, will select that candidate. This class includes all Condorcet-consistent rules, but also other rules such as plurality, plurality with runoff, STV, and Bucklin. For this class, we show that under some assumptions on strategic behavior, the Condorcet winner does in fact end up being elected (though for other assumptions this does not hold).

The second setting is voting by successive elimination. This voting rule, which is often used in committees, proceeds by holding successive pairwise elections. In this setting, there is a particularly natural notion of strategic voting known as *sophisticated voting* [13, 22, 20]. The outcomes of sophisticated voting (the so-called *sophisticated outcomes*) have been characterized by Banks [1] for the case when all candidates run. Dutta et al. [9] extended the characterization result by Banks to the case of strategic candidacy. We study the computational complexity of sophisticated outcomes in the latter case and show that computing the set of sophisticated outcomes is NP-complete.

The paper is organized as follows. We review related literature in Section 2 and introduce necessary concepts in Section 3. Sections 4 and 5 contain the results for the two settings described above, and Section 6 concludes.

2 Related Work

Strategic candidacy was introduced by Dutta, Jackson, and Le Breton [8], who showed that every non-dictatorial voting rule might give candidates incentives not to run. Subsequently, Ehlers and Weymark [10] and Samejima [28] came up with alternative proofs and extensions of some of the results of Dutta et al. [8]. Furthermore, models of strategic candidacy have been extended to set-valued [11, 26] and probabilistic [27] voting rules.

In a companion paper, Dutta, Jackson, and Le Breton [9] focussed on the effects of strategic candidacy on the class of binary voting rules. They completely characterized the set of equilibrium outcomes for the successive elimination procedure, a prominent member of this class. We will make use of this characterization when proving our computational intractability result in Section 5.

Samejima [29] studied strategic candidacy for single-peaked preferences and characterized the class of candidate stable voting rules for this domain. He showed that, under some mild conditions, a voting rule is candidate stable for single-peaked preferences if and only if it is a *k-th leftmost peak rule* for some $k \leq |V|$. A k-th leftmost peak rule fixes a single-peaked axis, identifies each voter with his most preferred candidate (his "peak"), and selects the peak of the k-th leftmost voter according to the ordering given by the axis. (The median rule is the special case for $k = \frac{n+1}{2}$.)

Also related are two papers that precede Dutta et al. [8]. Osborne and Slivinski [25] and Besley and Coate [3] study plurality equilibria in a candidacy game where all voters are potential candidates and running is costly. In both papers, preferences of voters and

[1] Lang et al. [16] extended this result by showing that, in this setting, no coalition of candidates ever has an incentive to change their strategies as long as the Condorcet winner is running.

candidates are defined via a spatial model (which, in the one-dimensional case, yields single-peaked preferences). However, the focus of these two papers is different from ours: They are mainly interested in how the number and spatial position of candidates that run in equilibrium is affected by parameters such as entry costs, preferences, and candidates' utilities for winning. There is also a number of technical differences to our paper. For example, Osborne and Slivinski [25] consider a continuum of voters and assume that voters vote truthfully. And Besley and Coate [3] add a third stage to the two-stage candidacy game by letting the selected candidate choose a policy from a given policy space. None of the two papers considers strong equilibria.

Finally, there is a recent paper by Lang, Maudet, and Polukarov [16], which studies for which voting rules the candidacy game admits pure equilibria under the assumption that voters vote truthfully. They also consider strong equilibria and show that, for every domain that guarantees the existence of a Condorcet winner and for every Condorcet-consistent voting rule, a set of running candidates forms a strong equilibrium if and only if the Condorcet winner is contained in the set.

3 Preliminaries

This section introduces the concepts and notations that are used in the remainder of the paper. For a finite set X, let $\mathcal{L}(X)$ denote the set of rankings of X, where a ranking is a binary relation on X that is complete, transitive, and antisymmetric. For a ranking $R \in \mathcal{L}(X)$, $top(R)$ denotes the top-ranked element according to R.

3.1 Players and Preferences

Let C be a finite set of candidates and V a finite set of voters. Throughout this paper, we assume that $|V|$ is odd. The set P of *players* is given by $P = C \cup V$. We assume that $C \cap V = \emptyset$.[2] Each player $p \in P$ has preferences over the set of candidates, given by a ranking $R_p \in \mathcal{L}(C)$. For all candidates $c \in C$, we assume that the top-ranked candidate in R_c is c itself.[3] A preference profile $R = (R_p)_{p \in P} \in \mathcal{L}(C)^P$ contains preferences for all players. For a player $p \in P$ and two candidates $a, b \in C$, we write $a \succeq_p b$ if $(a, b) \in R_p$ and $a \succ_p b$ if $a \succeq_p b$ and $a \neq b$.

For a preference profile R and a candidate c, let $V_R(c)$ denote the set of voters that have c as their top-ranked candidate, i.e., $V_R(c) = \{v \in V : top(R_v) = c\}$. Moreover, for a candidate $d \neq c$, let $V_R(c, d)$ denote the set of voters that prefer c to d, i.e., $V_R(c, d) = \{v \in V : c \; R_v \; d\}$. Candidate c is a *majority winner in* R if $|V_R(c)| > |V|/2$, and c is a *Condorcet winner in* R if $|V_R(c, d)| > |V|/2$ for all $d \in C \setminus \{c\}$. Note that both concepts ignore the preferences of candidates. Every preference profile can have at most one majority winner and at most one Condorcet winner. If candidate is a majority winner in R, then it is also a Condorcet winner in R.

Let $\lhd \in C \times C$ be a strict ordering of the candidates. A preference profile $R = (R_p)_{p \in P}$ is *single-peaked with respect to* $\lhd$ if the following condition holds for all $a, b \in C$ and $p \in P$: if $a \lhd b \lhd top(R_p)$ or $top(R_p) \lhd b \lhd a$, then $b \; R_p \; a$. For a preference profile R that is single-peaked with respect to $\lhd$, the *median* of R is defined as the unique candidate c for which both $\sum_{a \in C : a \lhd c} |V_R(a)| < |V|/2$ and $\sum_{a \in C : c \lhd a} |V_R(a)| < |V|/2$. It is well known that the median is a Condorcet winner in R.

[2] See Dutta et al. [8, 9] for results without this assumption.

[3] This assumption is known as *narcissism*. Without it, scenarios can arise where no candidate has an incentive to run (see [8], page 1017). We also assume that each candidate prefers himself to the outcome $\top$, which corresponds to the case where no candidate runs.

Let $c_1 \lhd c_2 \lhd \ldots \lhd c_m$ and let R be a preference profile that is single-peaked with respect to $\lhd$. The *peak distribution* of R with respect to $\lhd$ is the vector of length m whose j-th entry is the number $|V_R(c_j)|$ of voters that rank c_j highest.

3.2 Voting Rules

A *voting rule* f maps a non-empty subset $B \subseteq C$ of candidates and a profile of votes $r = (r_v)_{v \in V} \in \mathcal{L}(B)$ to a candidate $f(B, r) \in B$. A voting rule f is *majority-consistent* if $f(B, (R_v)_{v \in V}) = c$ whenever c is a majority winner in $R|_B$, and f is *Condorcet-consistent* if $f(B, (R_v)_{v \in V}) = c$ whenever c is a Condorcet winner in $R|_B$. Because majority winners are always Condorcet winners, (perhaps confusingly) Condorcet-consistency implies majority-consistency.

A *scoring rule* is a voting rule that is defined by a sequence $s = (s^n)_{n \geq 1}$, where for each $n \in \mathbb{N}$, $s^n = (s_1^n, \ldots, s_n^n) \in \mathbb{R}^n$ is a *score vector* of length n. For a preference profile R on k candidates, the score vector s^k is used to allocate points to candidates: each candidate receives a score of s_j^k for each time it is ranked in position j by a voter. (Again, preferences of candidates are ignored.) The scoring rule then selects the candidate with maximal total score. In the case of a tie, a fixed tie-breaking ordering $\tau \in \mathcal{L}(C)$ is used. Prominent examples of scoring rules are *plurality* ($s^n = (1, 0, \ldots, 0)$), *Borda's rule* ($s^n = (n-1, n-2, \ldots, 0)$), and *veto* ($s^n = (0, \ldots, 0, -1)$).

The plurality winner is a candidate maximizing $|V_R(\cdot)|$. Plurality is majority-consistent, but not Condorcet-consistent. Borda's rule and veto are not majority-consistent and (hence) not Condorcet-consistent.

3.3 Candidacy and Voting as a Two-Stage Game

We consider the following two-stage game. At the first stage, each candidate decides whether to run in the election or not. At the second stage, each voter casts a ballot containing a ranking of the running candidates. Throughout, we consider *complete-information* games: the preferences of the candidates and voters are common knowledge among the candidates and voters. Hence, we do not need to model games as (pre-)Bayesian and strategies do not have to condition on the player's type.

Let S_p be the set of strategies of player p. Then for each candidate $c \in C$, the set S_c is given by $\{0, 1\}$, with the convention that 1 corresponds to "running" and 0 corresponds to "not running." For each voter $v \in V$, the set S_v consists of all functions

$$s_v : 2^C \to \bigcup_{B \subseteq C} \mathcal{L}(B)$$

that map a subset $B \subseteq C$ of candidates to a ranking $s_v(B) \in \mathcal{L}(B)$. The interpretation is that $s_v(B)$ is the vote of voter v when the set of running candidates is B. In particular, each S_v contains a strategy that corresponds to *truthful voting* for voter v: this strategy maps every set B to the ranking $R_v|_B$. In general, however, a voter can rank two candidates differently depending on which other candidates run.

We are now ready to define the outcomes of the game. A *strategy profile* $s = (s_p)_{p \in P}$ contains a strategy for every player. Given a strategy profile s and a voting rule f, define $C(s) = \{c \in C : s_c = 1\}$ (the set of running candidates[4]) and $r(s) = (s_v(C(s)))_{v \in V} \in \mathcal{L}(C(s))^V$ (the votes cast for this set of running candidates). The outcome $o_f(s)$ of s under f is then given by $o_f(s) = f(C(s), r(s))$.

[4]If $C(s) = \emptyset$, define $o_f(s) = \top$. The assumption that every candidate prefers himself to $\top$ ensures that at least one candidate will run whenever candidates act strategically.

3.4 Equilibrium Concepts

Let $s = (s_p)_{p \in P}$ be a strategy profile. For a subset $\tilde{P} \subseteq P$ and a profile of strategies $s'_{\tilde{P}} = (s'_p)_{p \in \tilde{P}}$ for players in $\tilde{P}$, let $(s'_{\tilde{P}}, s_{-\tilde{P}})$ denote the strategy profile where each player $p \in \tilde{P}$ plays strategy s'_p and all remaining players play the same strategy as in s. Fix a voting rule f and a preference profile R. For a strategy profile s and a subset $\tilde{P} \subseteq P$ of players, say that s is *(R, f)-deviation-proof w.r.t.* $\tilde{P}$ if for all $s'_{\tilde{P}}$, there exists $p \in \tilde{P}$ such that

$$o_f(s) \succeq_p o_f(s'_{\tilde{P}}, s_{-\tilde{P}}).$$

For a strategy profile $s = (s_p)_{p \in P}$, we sometimes write $s = (s_C, s_V)$, where $s_C = (s_c)_{c \in C}$ is the profile of candidate strategies and $s_V = (s_v)_{v \in V}$ is the profile of voter strategies. We can now define equilibrium behavior for both candidates and voters.

Definition 1. *Let R be a preference profile and let f be a voting rule. A strategy profile $s = (s_C, s_V)$ is*

- *a* C-equilibrium for R under f *if s is (R, f)-deviation-proof w.r.t. $\{c\}$ for all $c \in C$;*
- *a* strong C-equilibrium for R under f *if s is (R, f)-deviation-proof w.r.t. C' for all $C' \subseteq C$;*
- *a* V-equilibrium for R under f *if for every $s'_C \in \{0, 1\}^C$, (s'_C, s_V) is (R, f)-deviation-proof w.r.t. $\{v\}$ for all $v \in V$;*
- *a* strong V-equilibrium for R under f *if for every $s'_C \in \{0, 1\}^C$, (s'_C, s_V) is (R, f)-deviation-proof w.r.t. V' for all $V' \subseteq V$.*

We omit the reference to R and f if the preference profile or the voting rule is known from the context. In a C-equilibrium, no candidate can achieve a more preferred outcome by unilaterally changing his strategy. In a strong C-equilibrium, no coalition of candidates can change the outcome in such a way that every player in the coalition prefers the new outcome to the original one. Thus, (strong) C-equilibria correspond to (strong) Nash equilibria when strategies of voters are assumed to be fixed. For voters, the equilibrium notions are more demanding: In order to be considered a (strong) V-equilibrium, the strategies of voters are required to form a (strong) Nash equilibrium for every subset $B \subseteq C$ of running candidates.

It is instructive to relate these definitions to established game-theoretic solution concepts for extensive-form games, such as subgame-perfect equilibrium and subgame-perfect strong equilibrium. A strategy profile s is a *subgame-perfect equilibrium* of a game G if for any subgame $G' \subseteq G$, the restriction of s to G' is a Nash equilibrium of G', and it is a *subgame-perfect strong equilibrium* if for any subgame $G' \subseteq G$, the restriction of s to G' is a strong Nash equilibrium of G'. In the candidacy game, every subgame (other than the game itself) corresponds to a voting game that takes place after the candidates have decided whether or not to run. Thus, a proper subgame can be identified with the set of candidates that run in this subgame.

For candidates, playing a subgame-perfect equilibrium is not a stronger requirement than playing a Nash equilibrium, because the only subgame in which they play is the entire game itself. For voters, on the other hand, playing a subgame-perfect equilibrium entails playing a Nash equilibrium for every possible set of running candidates. Therefore, we have the following.

Fact 1. *A strategy profile is a subgame-perfect equilibrium of the candidacy game if and only if it is both a C-equilibrium and a V-equilibrium.*

For subgame-perfect *strong* equilibria, one implication is straightforward.

Fact 2. *Every subgame-perfect strong equilibrium of the candidacy game is both a strong C-equilibrium and a strong V-equilibrium.*

However, the following example shows that the other direction does not hold in general, because even if coalitions of either one type of players cannot successfully deviate, it is possible that a mixed coalition including players of both types can.

Example 1. *Consider a preference profile with candidates a, b, c and a single voter with preferences $a \succ b \succ c$. The preferences of candidate b are given by $b \succ_b c \succ_b a$. The voting rule f selects the candidate ranked first by the voter whenever all three candidates run; if, however, at most two candidates run, the lexicographically last one is chosen, ignoring the voter's vote. Let s be the strategy profile in which a and c run and the voter votes truthfully. The outcome of s under f is $o_f(s) = c$. We claim that s is (1) a strong C-equilibrium and (2) a strong V-equilibrium, but (3) not a subgame-perfect strong equilibrium (in fact not even a strong equilibrium).*

For (1), observe that c has no incentive to participate in any deviation. The same holds for a, because the outcome will still be c if a deviates (whether b runs or not). And when all three candidates run, the outcome is a, making candidate b—the only deviator—worse off. For (2), s is a strong V-equilibrium because the voter makes his favorite candidate win in the only case where his vote has any influence. For (3), consider the following deviation. Candidate b deviates to running and the voter deviates to ranking b first whenever b runs. The outcome will change to b, and both deviators (candidate b and the voter) prefer b to c.

Splitting up the equilibrium definitions into separate requirements for C and V allows us to capture scenarios in which only players of one type (candidates or voters) act according to the corresponding equilibrium notion. In Section 4 we will analyze which combinations of equilibrium notions yield desirable outcomes. We will present both positive results, stating that a desirable outcome will be selected whenever a strategy profile meets a certain combination of equilibrium conditions, and negative results, stating that undesirable outcomes may be selected even if certain equilibrium conditions hold.

In sufficiently general settings, the existence of solutions is not guaranteed for any of the equilibrium concepts in Definition 1.[5] However, for all the positive results in Section 4, we also show that every preference profile admits a strategy profile that meets the corresponding equilibrium conditions.

4 Majority-Consistent Voting Rules and Single-Peaked Preferences

In this section, we assume that preference profiles are single-peaked and that the order $\lhd$ witnessing single-peakedness is given. (If the order is not part of the input, it can be computed in polynomial time [2, 12].) Note that our definition of single-peakedness in Section 3.1 also requires the preferences of *candidates* to be single-peaked with respect to $\lhd$. Given that the preferences of voters are single-peaked with respect to $\lhd$, this does not appear to be an unreasonable assumption.

We are interested in the following question: which requirements on the strategies of players are sufficient for the Condorcet winner (which is guaranteed to exist) to be the outcome? For Condorcet-consistent rules, the answer to this question is relatively straightforward [16].

[5]Subgame-perfect equilibria (i.e., strategy profiles that are simultaneously a C-equilibrium and a V-equilibrium) are guaranteed to exist if one allows for *mixed strategies* and extends the preferences of players to the set of all probability distributions over $C \cup \{\top\}$ in an appropriate way.

Therefore, we are mainly interested in voting rules that are majority-consistent, but not Condorcet-consistent. The simplest and most important such rule is plurality.

Since plurality is not a k-th leftmost peak rule, the result by Samejima [29] (see Section 2) implies that there exist profiles where some candidates have an incentive not to run (assuming truthful voting). Indeed, it is easy to construct such a profile.[6]

Example 2. *Consider a single-peaked preference profile with candidates $a \lhd b \lhd c$ and peak distribution $(3, 2, 4)$. Under truthful voting, the plurality winner is c. However, if candidate a does not run, the three voters in $V_R(a)$ rank candidate b first, making b the plurality winner. By single-peakedness, candidate a prefers b to c.*

This example also shows that plurality can fail to select the Condorcet winner when all candidates run and all voters vote truthfully. The next example shows that requiring both candidates and voters to play subgame-perfect equilibrium strategies is still not sufficient for the Condorcet winner to be chosen.

Example 3. *Consider a single-peaked preference profile with candidates $a \lhd b \lhd c \lhd d \lhd e$ and peak distribution $(11, 3, 3, 3, 3)$. The Condorcet winner is b. Let s be the strategy profile in which $s_x = 1$ for all $x \in \{a, b, c, d, e\}$ and s_v is "truthful voting" for all voters v. Then $o_{plurality}(s) = a$ and no candidate other than a can change that outcome by unilaterally deviating. Therefore, s is a C-equilibrium. To see that s is also a V-equilibrium, we need to check that "truthful voting" is deviation-proof for every subset of running candidates. Deviation-proofness clearly holds whenever at most two candidates run. If at least three candidates run, single-peakedness implies that the leftmost among the running candidates has a plurality score of at least 11, whereas each other running candidate has a score of at most 9. Thus, no voter can change the outcome by unilaterally deviating.*

We go on to show that the Condorcet winner *will* be chosen if we require stronger equilibrium notions. We first analyze strong V-equilibria. Our result does not require single-peaked preferences; it holds for a strictly larger class of preference profiles.[7]

Theorem 1. *Let R be a preference profile with Condorcet winner c^* and let f be a majority-consistent voting rule.*

(i) If $R|_B$ has a Condorcet winner for every nonempty subset $B \subseteq C$, then there exists a subgame-perfect strong equilibrium (and hence a strategy profile that is both a strong C-equilibrium and a strong V-equilibrium) for R under f in which all candidates run.

(ii) If s is a strong V-equilibrium for R under f and $s_{c^*} = 1$, then $o_f(s) = c^*$.

Proof. For (i), denote by $c_B \in B$ the Condorcet winner in $R|_B$. Let s be a strategy profile where all candidates run and all voters rank c_B first whenever the set of running candidates is given by B. Hence, $o_f(s) = c^*$. We claim that s is a subgame-perfect strong equilibrium for R under f. In order to prove this claim, we need to show that for every subgame, there is no beneficial deviation for any coalition.

First, consider the subgame that is given by the entire game itself. Suppose, for the sake of contradiction, that there is a coalition $\tilde{P} = C' \cup V'$ of candidates and voters that can change the outcome to some $a \neq c^*$ and that all players in $\tilde{P}$ prefer a to c^*. Let $s' = (s'_{\tilde{P}}, s_{-\tilde{P}})$ denote the strategy profile that results from this deviation. Observe that $c^* \notin \tilde{P}$, because $c^* \succ_{c^*} a$. Therefore, c^* is still running under s' and all non-deviating

[6]We often simplify examples by specifying the peak distribution only. This piece of information is clearly sufficient to identify both the Condorcet winner and, in the absence of ties, the plurality winner.

[7]In particular, note that Theorem 1 does not make any assumptions on the preferences of candidates (other than narcissism).

voters $V \setminus V'$ still rank c^* highest under s'. That means that the number $|V'|$ of deviating voters has to be greater than $|V|/2$, as otherwise majority-consistency of f would yield $o_f(s') = c^*$. But then V' is a majority of voters, each preferring a over c^*. This contradicts the assumption that c^* is a Condorcet winner.

Second, consider a subgame that arises after the candidates have chosen whether or not to run. Let $B \subseteq C$ be the set of candidates that run in this subgame. If $B = \emptyset$, the outcome is $\top$ and no coalition of voters can change the outcome. If $B \neq \emptyset$, all voters rank c_B first by the definition of s. By an argument analogous to the one above, the existence of a successfully deviating coalition of voters would violate the assumption that c_B is the Condorcet winner in $R|_B$. Therefore, s is a subgame-perfect strong equilibrium.

For *(ii)*, let s be a strong V-equilibrium for R under f with $s_{c^*} = 1$. Assume for the sake of contradiction that $o_f(s) = a \neq c^*$. We will show that s is not a strong V-equilibrium, by means of the following deviation. Let $\tilde{P} = V_R(c^*, a)$ be the set of voters that prefer c^* over a and let $s' = (s'_{\tilde{P}}, s_{-\tilde{P}})$ be the strategy profile in which all voters in $\tilde{P}$ rank c^* first whenever c^* runs. Since $s_{c^*} = 1$ and $|\tilde{P}| = |V_R(c^*, a)| > |V|/2$, majority-consistency of f implies $o_f(s') = c^*$. Moreover, $c^* \succ_p a$ for all $p \in \tilde{P}$ by the definition of $\tilde{P}$. Therefore, s is not (R, f)-deviation-proof w.r.t. $\tilde{P}$, contradicting the assumption that s is a strong V-equilibrium. □

We remark that part *(ii)* of Theorem 1 can be generalized by observing that it is sufficient for f to satisfy the following condition, which is considerably weaker than majority-consistency:

> Whenever a set $V' \subseteq V$ of voters forms a majority (i.e., $|V'| > |V|/2$), then for every candidate $a \in C$ that is running and every profile of votes for voters in $V \setminus V'$, the voters in V' can vote in such a way that candidate a is chosen.

It can be shown that all unanimous C2 functions [14] satisfy this property.

The following corollary summarizes the consequences of Theorem 1 for single-peaked preference profiles.

Corollary 1. *Let R be a single-peaked preference profile with Condorcet winner c^* and let f be a majority-consistent voting rule.*

(i) There exists a subgame-perfect strong equilibrium (and hence a strategy profile that is both a strong V-equilibrium and a strong C-equilibrium) for R under f.

(ii) If s is a strong V-equilibrium and a C-equilibrium (strong or not) for R under f, then $o_f(s) = c^*$.

Proof. For every subset $B \subseteq C$ of candidates, $R|_B$ is still single-peaked and thus has a Condorcet winner. Therefore, *(i)* immediately follows from the first part of Theorem 1. For *(ii)*, assume that s is a strong V-equilibrium. Then, the second part of Theorem 1 implies that the Condorcet winner c^* will be the outcome whenever he runs. It follows that, holding the strategies of the voters fixed, the strategy "running" ($s_{c^*} = 1$) strictly dominates "not running" ($s_{c^*} = 0$) for c^*. Since in a C-equilibrium (strong or not), no candidate plays a strategy that is strictly dominated when the voters' strategies are held fixed, c^* is running (and winning) in any C-equilibrium. □

Thus, the Condorcet winner will be chosen if voter strategies form a strong V-equilibrium and candidate strategies satisfy a minimal degree of rationality. Messner and Polborn [19] show a similar result for the plurality rule when all candidates are assumed to run.

We provide two examples that show that the statements of Corollary 1 do not hold for voting rules that are not majority-consistent.

Example 4. *Let R be a single-peaked preference profile with candidates $a \lhd b \lhd c$ and peak distribution $(5, 0, 4)$. If f is Borda's rule, there does not exist a strong V-equilibrium (and hence no subgame-perfect strong equilibrium). To see this, consider the case where all candidates run. Observe that in any strong V-equilibrium, the outcome would have to be a. (Suppose the outcome is not a. Then, the five voters in $V_R(a)$ can jointly deviate and change the outcome to a. They can do this by having one voter voting $a \succ b \succ c$, and the remaining four voters voting exactly the opposite rankings of the voters in $V_R(c)$.) However, there is no strong V-equilibrium that yields outcome a. This is because the voters in $V_R(c)$ prefer both other alternatives to a, and—no matter how the voters in $V_R(a)$ vote—the voters in $V_R(c)$ can jointly deviate and achieve an outcome other than a. (One of b and c will obtain a score of at least 3 from the voters in $V_R(a)$. Without loss of generality, suppose it is b. Then the voters in $V_R(c)$ can all vote $b \succ c \succ a$, making b the winner.)*

Example 5. *Let R be a single-peaked preference profile with candidates $a \lhd b \lhd c$ and five voters: three voters have preferences $a \succ b \succ c$ and two voters have preferences $b \succ c \succ a$. The Condorcet winner is a. Let f be the voting rule veto*[8] *and let s be the strategy profile where all candidates run and all voters vote truthfully. Then, $o_f(s) = b$. Moreover, s is a strong C-equilibrium and a strong V-equilibrium. The former holds because any deviation involving a does not change the outcome (provided b still runs), and c can only change the outcome to the less preferred alternative a. For the latter, the only interesting case is when all three candidates run. In this case, the two voters in $V_R(b)$ have no incentive to deviate from truthful voting (their favorite candidate is winning) and there is no way for the three voters in $V_R(a)$ to jointly deviate and achieve outcome a. (They can change the outcome to c by voting $a \succ c \succ b$, but they prefer b to c.) It can furthermore be shown that, when all candidates run,* every *strong V-equilibrium yields outcome b.*

We now move to the case where *candidates* play a strong equilibrium. If voters vote truthfully, the outcome will be the Condorcet winner.

Theorem 2. *Let R be a single-peaked preference profile with Condorcet winner c^* and let f be a majority-consistent voting rule.*

(i) There exists a strong C-equilibrium for R under f where all voters vote truthfully.

(ii) If s is a strong C-equilibrium for R under f where all voters vote truthfully, then $o_f(s) = c^*$.

Proof. For (*i*), let s be the strategy profile in which only c^* runs and all voters vote truthfully. We show that this is a strong C-equilibrium for R under f. Suppose, for the sake of contradiction, that $\tilde{C} \subseteq C$ is a coalition of candidates that can, by changing its strategies, make alternative $a \neq c^*$ win, and moreover that all candidates in $\tilde{C}$ prefer a to c^*. Define $C^- = \{c \in C : c \lhd c^*\}$ and $C^+ = \{c \in C : c^* \lhd c\}$, and without loss of generality suppose that $a \in C^-$. Because candidates' preferences are single-peaked and they rank themselves first, it follows that $\tilde{C} \subseteq C^-$. But this implies that still, no candidate in C^+ runs. Hence, all voters with $top(R_v) \in C^+ \cup \{c^*\}$ still rank c^* first (since they vote truthfully), and because f is majority-consistent, it follows that c^* wins. This gives us the desired contradiction.

For (*ii*), let s be a strong C-equilibrium for R under f where all voters vote truthfully. Consider the set $C(s)$ of candidates that are running under s. Define $C_s^- = \{c \in C(s) : c \lhd c^*\}$ and $C_s^+ = \{c \in C(s) : c^* \lhd c\}$. Assume for the sake of contradiction that $o_f(s) = a \neq c^*$. Without loss of generality, suppose that $a \in C_s^-$. Consider the set $\tilde{C}$ of candidates

[8] Veto does not only violate majority-consistency, but also the weaker property defined after Theorem 1.

given by $\tilde{C} = C_s^+ \cup \{c^*\}$. Define $s'_{\tilde{C}} = (s'_c)_{c\in\tilde{C}}$ by

$$s'_c = \begin{cases} 1 & \text{if } c = c^* \\ 0 & \text{if } c \in C_s^+ \end{cases}$$

and observe that $o_f(s'_{\tilde{C}}, s_{-\tilde{C}}) = c^*$. The reason for the latter is that (1) the set of voters v with $top(R_v) = c^*$ or $c^* \lhd top(R_v)$ forms a majority, (2) all of these voters satisfy $top(R_v|_{C(s'_{\tilde{C}}, s_{-\tilde{C}})}) = c^*$, and (3) all voters vote truthfully by assumption. Moreover, single-peakedness implies that all candidates in $\tilde{C}$ prefer c^* to a. Therefore, s is not (R, f)-deviation-proof w.r.t. $\tilde{C}$, contradicting the assumption that s is a strong C-equilibrium. □

Similar to the case of Theorem 1, we now provide examples that show that Theorem 2 cannot be generalized in certain ways. The first example shows that Theorem 2 does not hold for Borda's rule (which is not majority-consistent).

Example 6. *Consider a single-peaked preference profile with candidates $a \lhd b \lhd c$ and five voters: three voters have preferences $a \succ b \succ c$ and two voters have preferences $b \succ c \succ a$. The Condorcet winner is a. Let s be the strategy profile where $s_a = s_b = s_c = 1$ and s_v is "truthful voting" for all voters v. It is easily verified that s is a strong C-equilibrium and $o_{Borda}(s) = b$. In fact, it can be checked that the Condorcet winner is not chosen in* any *strong C-equilibrium with truthful voting. (The only other strong C-equilibrium under truthful voting has candidates b and c running and also yields outcome b.)*

The second example shows that Theorem 2 does not hold if the preferences of candidates are not single-peaked with respect to the given order.

Example 7. *Consider the following preference profile with candidates a, b, c and 14 voters.*

4	4	6
a	b	c
b	a	b
c	c	a

The preferences of the candidates are such that a prefers c over b and b prefers c over a. Whereas the preferences of the voters are single-peaked with respect to the ordering $a \lhd b \lhd c$, this is not true for the preferences of the candidates. (Therefore, this profile is not single-peaked according to the definition in Section 3.1.) The Condorcet winner is b and the Condorcet loser is c. Let s be the strategy profile where all candidates run and all voters vote truthfully. It is easily verified that s is a strong C-equilibrium and $o_{plurality}(s) = c$. In fact, "everybody running" is the only strong C-equilibrium under truthful voting.

Since Theorem 1 already covers the case where *both* voters and candidates play a strong (subgame-perfect) equilibrium, only one case is left to consider: candidates playing a strong C-equilibrium, and voters merely playing a V-equilibrium. The following example shows that these requirements are *not* sufficient for the Condorcet winner to be chosen.

Example 8. *Consider a single-peaked preference profile with candidates $a \lhd b \lhd c$ and peak distribution $(1, 1, 1)$. The Condorcet winner is b. Let s be a strategy profile with $s_c = 1$ and voter strategies s_v that satisfy*

$$top(s_v(B)) = \begin{cases} c & \text{if } c \in B \\ top(R_v|_B) & \text{otherwise} \end{cases}$$

	strong V-equilibrium	V-equilibrium	truthful voting ($s_v(B) = R_v|_B$)
strong C-equilibrium	yes (Corollary 1)	no (Example 8)	yes (Theorem 2)
C-equilibrium	yes (Corollary 1)	no (Example 3)	no (Example 3)
naive candidacy ($s_c = 1$)	yes (Theorem 1)	no (Example 3)	no (Examples 2 & 3)

Table 1: Overview of results. A table entry is "yes" if every strategy profile that satisfies the corresponding (row and column) conditions yields the Condorcet winner under every majority-consistent voting rule. Moreover, for every "yes" entry, a strategy profile satisfying the conditions is guaranteed to exist.

for each $B \subseteq C$. That is, all three voters rank c first whenever c runs, and vote truthfully otherwise. Obviously, $o_{plurality}(s) = c$. We claim that s is both a V-equilibrium and a strong C-equilibrium. For the former, we distinguish two cases: If c runs, then all voters rank c first and no voter can change the outcome by unilaterally deviating. If c does not run, then at most two candidates run and no voter can benefit by voting for his less preferred candidate. For the latter, no coalition of candidates can change the outcome in such a way that all members of the coalition prefer the new outcome to c. (Such a coalition would need to include candidate c, who has no incentive to deviate.)

The phenomenon illustrated in this example is perhaps somewhat surprising: Assuming that candidates play a strong C-equilibrium, both truthful voting and strong V-equilibrium voting yields the desirable outcome; however, V-equilibrium voting—a notion of sophistication that might appear to be "in between" the other two notions—does not. Table 1 summarizes the results of this section.

5 Computing the Candidate Stable Set

In this section, we study a voting rule known as *voting by successive elimination (VSE)*. In particular, we will be interested in the computational complexity of computing outcomes under VSE if both candidates and voters act strategically. We do not require single-peaked preferences, but in order to avoid majority ties, we still assume that the number of voters is odd. VSE takes as input an ordering $\sigma \in \mathcal{L}(C)$ of the candidates. The rule proceeds by holding successive pairwise elections. In a pairwise election, there are two candidates a and b and every voter $v \in V$ votes for exactly one of the two candidates. Candidate a wins the pairwise election if the number of voters voting for a is strictly greater than $|V|/2$.

For a given subset $B \subseteq C$ of candidates with $|B| \geq 2$, VSE works as follows. Label the candidates such that $\sigma|_B = (c_1, c_2, \dots, c_{|B|})$. In the first round, there is a pairwise election between c_1 and c_2. The winner of this election proceeds to the second round, where he faces c_3. The winner of this election then faces c_4, and so on. VSE selects the winner of round $|B| - 1$.

Truthful voting for a voter v with preferences R_v corresponds to the strategy that, in every pairwise election between two candidates a and b, the voter votes for $top(R_v|_{\{a,b\}})$. It is well known that, under VSE, voters can benefit from voting strategically. Moreover, there is a particularly natural notion of strategic voting called *sophisticated voting* [13, 22, 20]. Sophisticated voting assumes that voters' preferences are common knowledge and applies a backward induction argument: In the last round of VSE, there is no incentive to vote strategically and thus the majority winner of the remaining two candidates will be chosen. Anticipating that, in the second-to-last round, voters are able to compare which outcome would eventually result from either one of the current candidates winning this round, and vote accordingly; etc. In the absence of majority ties, sophisticated voting yields a unique winning candidate, the *sophisticated outcome*. The sophisticated outcome corresponds to the outcome that results when voters iteratively eliminate weakly dominated strategies.

In order to determine both the truthful outcome and the sophisticated outcome, it is sufficient to know the truthful outcome of pairwise elections between all pairs of the candidates. This information is captured by the majority relation. For a preference profile R, the *majority relation* $R_M \subset C \times C$ is defined by

$$a \; R_M \; b \quad \text{if and only if} \quad V_R(a,b) > \frac{|V|}{2}.$$

The majority relation of a preference profile R (with an odd number of voters) can be conveniently represented as a *tournament*, i.e., a directed graph $T = (C, \succ)$ with $a \succ b$ if and only if $a \; R_M \; b$.

Shepsle and Weingast [31] defined an algorithm that, given a majority relation R_M, an ordering σ, and a subset $B \subseteq C$ of the candidates, computes the sophisticated outcome when the set of running candidates is given by B. Moreover, Banks [1] characterized the set of candidates that, for given R_M and $B \subseteq C$, are the sophisticated outcome for *some* ordering σ. This set is known as the *Banks set* $BA(B, R_M)$. In the notation[9] developed in this paper, $BA(B, R_M)$ corresponds to $\bigcup_\sigma o_{VSE(\sigma)}(s)$, where $s_c = 1$ if $c \in B$ and s_v is "sophisticated voting" for all voters $v \in V$.

Dutta et al. [9] analyzed how the set of sophisticated outcomes changes when strategic candidacy is accounted for. Consider a strategy profile $s = (s_C, s_V)$, where $s_C = (s_c)_{c \in C}$ and $s_V = (s_v)_{v \in V}$ and say that s is an *entry equilibrium* if it is a C-equilibrium and s_v is "sophisticated voting" for all voters $v \in V$. The *candidate stable set (CS)* is defined as the set of all candidates that are the sophisticated outcome for some collection of candidate preferences and for some ordering σ, when the set of running candidates is given by $C(s)$ for some entry equilibrium s.

More formally, for a preference profile $R = (R_p)_{p \in P}$, define $R_C = (R_c)_{c \in C}$ and $R_V = (R_v)_{v \in V}$. For an ordering σ, let $E(R, \sigma) = E(R_C, R_V, \sigma)$ denote the set of entry equilibria of R when the order is σ. Then, the candidate stable set of R_V is given by

$$CS(R_V) = \bigcup_{\sigma} \bigcup_{R_C} \bigcup_{s \in E(R_C, R_V, \sigma)} o_{VSE(\sigma)}(s).$$

Thus, the candidate stable set is the analog of the Banks set when strategic candidacy is taken into account. Since $CS(R_V)$ only depends on the majority relation R_M of R, we usually write $CS(R_M)$.[10]

Dutta et al. [9] have provided an elegant characterization of the candidate stable set in terms of the majority relation R_M. In order to present this characterization, we need some

[9]Strategies, outcomes, and equilibrium notions for VSE can be defined similarly to the definitions in Section 3. We omit the details since they are not important for our result. For formal definitions of the concepts considered in this section, we refer to Dutta et al. [9].

[10]Recall that the majority relation is independent of the preferences of candidates.

notation. Let $H(a, R_M)$ be the set of all subsets $B \subseteq C$ such that $R_M|_{B\times B}$ is transitive and $a \in B$ is the Condorcet winner in $R|_B$. Furthermore, say that a *covers* b if $a\ R_M\ b$ and for all $c \in C \setminus \{a, b\}$, $b\ R_M\ c$ implies $a\ R_M\ c$.

Proposition 1 (Dutta et al. [9])**.** *The candidate stable set is characterized as*

$$CS(R_M) = \{a \in C : \exists H \in H(a, R_M) \text{ s.t. } \forall b \notin H\ \exists c \in H \text{ s.t. } b \text{ does not cover } c\}.$$

We use this characterization to show that computing the candidate stable set is intractable. More precisely, we show that the following decision problem is NP-complete: Given a preference profile R and a candidate $c \in C$, is it the case that $c \in CS(R_M)$?

Theorem 3. *Computing the candidate stable set is NP-complete.*

Proof. Membership in NP is straightforward: for a fixed candidate, we can simply guess a set H and verify whether it satisfies the conditions in Proposition 1. For hardness, we give a reduction from 3SAT and adapt a construction that was used by Brandt et al. [5] to show that the Banks set is NP-hard to compute.

An instance[11] of *3SAT* is given by a Boolean formula $\varphi = (x_1^1 \vee x_1^2 \vee x_1^3) \wedge \cdots \wedge (x_m^1 \vee x_m^2 \vee x_m^3)$, where each $x \in \{x_i^1, x_i^2, x_i^3 : 1 \le i \le m\}$ is a literal. We assume the literals to be indexed and by X_i we denote the set $\{x_i^1, x_i^2, x_i^3\}$. Formula φ is *satisfiable* if there is a tuple $(x_1, \dots, x_m)$ in $\bigtimes_{1\le i\le m} X_i$ such that $v' = \bar{v}$ for no $v, v' \in \{x_1, \dots, x_m\}$.

Given a formula $\varphi = (x_1^1 \vee x_1^2 \vee x_1^3) \wedge \cdots \wedge (x_m^1 \vee x_m^2 \vee x_m^3)$, we define a tournament $T_\varphi = (C, \succ)$. (We will later invoke McGarvey's theorem [17], which guarantees the existence of a preference profile whose majority relation coincides with $\succ$.)

The set of nodes is given by $C = \{c^*\} \cup A \cup B \cup U_1 \cup \cdots \cup U_{2m-1}$, where $A = \{a_1, \dots, a_{2m-1}\}$, $B = \{b_1, \dots, b_{2m-1}\}$, and for all $j \le 2m - 1$,

$$U_j = \begin{cases} X_i & \text{if } j = 2i - 1, \\ \{y_j\} & \text{if } j = 2i. \end{cases}$$

The relation $\succ$ satisfies the following properties for all $u_i \in U_i$ and $u_j \in U_j$:

- $a_i \succ a_j$ if and only if $i < j$,
- $b_i \succ b_j$ if and only if $i < j$,
- $a_i \succ b_j$ if and only if $i = j$,
- $a_i \succ c^*$ and $b_i \succ c^*$ for all $i \le 2m - 1$,
- $c^* \succ u_i$ for all $i \le 2m - 1$,
- $u_i \succ a_j$ if and only if $i = j$,
- $b_i \succ u_j$ for all $i, j \le 2m - 1$,
- $u_i \succ u_j$ if $i < j$ and at least one of i and j is even.

For all $x \in X_i$ and $x' \in X_j$ with $i < j$, we furthermore have

- $x \succ x'$ if $x' \neq \bar{x}$, and
- $x' \succ x$ if $x' = \bar{x}$.

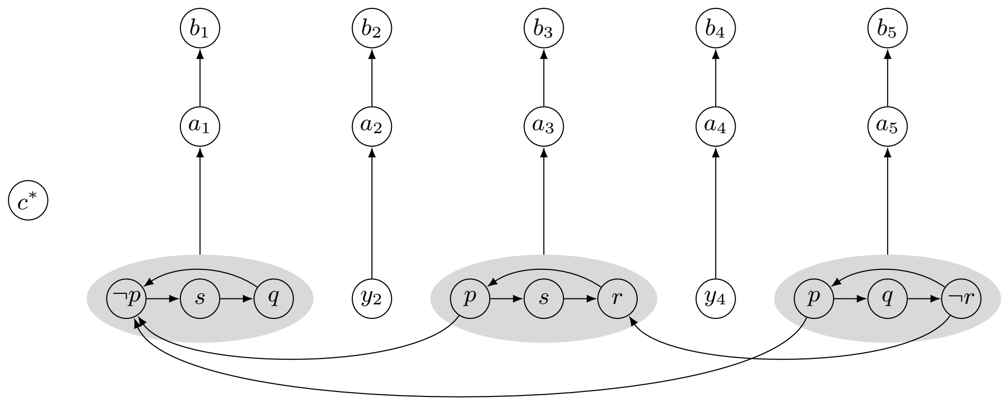

Figure 1: Tournament T_φ for the formula $\varphi = (\neg p \vee s \vee q) \wedge (p \vee s \vee r) \wedge (p \vee q \vee \neg r)$. A directed edge from node u to node v denotes $u \succ v$. All omitted edges point downwards or (in the case where two nodes are located at the same height) to the left.

Finally, there is a $\succ$-cycle $x_i^1 \succ x_i^2 \succ x_i^3 \succ x_i^1$ for all $i \leq m$. An example of a tournament T_φ for a specific formula φ is shown in Figure 1.

We now apply McGarvey's theorem [17] and let R^φ be a preference profile on candidate set C such that the majority relation R_M^φ coincides with $\succ$. Note that McGarvey's theorem is constructive and that the size of R^φ is polynomial in the size of φ.

We now show that the formula φ is satisfiable if and only if $c^* \in CS(R_M^\varphi)$. For the direction from left to right, assume that φ is satisfiable. Then there is a tuple $(x_1, \ldots, x_m)$ in $\bigtimes_{1 \leq i \leq m} X_i$ such that such that $x' = \bar{x}$ for no $x, x' \in \{x_1, \ldots, x_m\}$. Define

$$H = \{x_1, \ldots, x_m\} \cup \{y_2, y_4, \ldots, y_{2m-2}\} \cup \{c^*\}.$$

By the definition of $\succ$, c^* is the Condorcet winner in $R^\varphi|_H$ and H does not contain any cycles. Therefore, H is an element of $H(c^*, R_M^\varphi)$. Furthermore, no candidate in $C \setminus H$ covers all candidates in H. To see this, observe that such a candidate, call him d, would have to satisfy $d \succ h$ for all $h \in H$. This implies that $d \in \{b_1, \ldots, b_{2m-1}\}$. But no candidate in $\{b_1, \ldots, b_{2m-1}\}$ covers all candidates in H because, for each $i \leq 2m-1$, candidate b_i does not cover the unique candidate in the set $H \cap U_i$. Therefore, Proposition 1 yields that $c^* \in CS(R_M^\varphi)$.

For the direction from right to left, assume that $c^* \in CS(R_M^\varphi)$. By Proposition 1, there exists $H \in H(c^*, R_M^\varphi)$ such that no candidate outside H covers all candidates in H. We claim that H satisfies the following properties.

(*i*) $H \subseteq \{c^*\} \cup U_1 \cup \ldots \cup U_{2m-1}$,

(*ii*) $H \cap U_j \neq \emptyset$ for all $j \leq 2m-1$, and

(*iii*) $x' \neq \bar{x}$ for all $x, x' \in H$ (i.e., there does not exist a literal x such that H contains both a node corresponding to x and a node corresponding to $\bar{x}$).

For (*i*), observe that $c^* \succ h$ is necessary for $h \in H \setminus \{c^*\}$. For (*ii*), suppose that $H \cap U_j = \emptyset$ for some $j \leq 2m-1$. Then, b_j covers all candidates in H. For (*iii*), suppose $x' = \bar{x}$ for some $x, x' \in H$. Let $i, j \leq m$ be such that $x \in X_i$ and $x' \in X_j$. The assumption

[11] Following [5], we assume that for any two literals x and y in the same clause, neither $x = y$ nor $x = \bar{y}$.

in Footnote 11 implies $i \neq j$, and we can without loss of generality assume $i < j$. Then, there is a $\succ$-cycle $x \succ y_{2i} \succ x' \succ x$, violating transitivity of $R^\varphi|_{H \times H}$.

Define the tuple $(x_1, \ldots, x_m) \in \bigtimes_{1 \leq i \leq m} X_i$ by $x_i = H \cap X_i$ for all $i \leq m$. Property (*ii*) ensures that $(x_1, \ldots, x_m)$ is well defined and (*iii*) implies that $(x_1, \ldots, x_m)$ satisfies φ. $\square$

6 Conclusion

We have analyzed the combination of strategic candidacy and strategic voting in two settings that allow meaningful voting equilibria. In both settings, the set of equilibrium outcomes under strategic candidacy (given that voters are sufficiently sophisticated) has an elegant characterization: the Condorcet winner (in the single-peaked, majority-consistent rule setting with strong V-equilibria or with truthful voting and strong C-equilibria) and the candidate stable set (in the VSE setting with sophisticated voting). Whereas Condorcet winners are easy to compute, we have shown that the candidate stable set is computationally intractable.

It seems likely that the positive results in Section 4 extend to settings where preferences are single-peaked on a tree. It would also be interesting to check whether similar results can be obtained for related domain restrictions such as single-crossing or value-restricted preferences.

The positive results in Section 4 rely on finding the right level of equilibrium refinement (strong V-equilibrium, or strong C-equilibrium with truthful voting). If we move away from restricted domains, is there another type of equilibrium refinement [7, 32, 24] that allows us to arrive at meaningful equilibria by ruling out "unnatural" ones?

Equilibrium *dynamics* [18] is another topic for future research. For example, in the setting with single-peaked preferences and a majority-consistent rule, are there natural dynamics that are guaranteed to lead us to an equilibrium choosing the Condorcet winner?

On a higher level, one might wonder to what extent the phenomena exhibited in candidacy games can be related to other problems that involve altering the set of candidates, such as control problems, cloning, and nomination of alternatives.

Acknowledgments

We thank three anonymous reviewers for helpful comments. This work was supported by NSF and ARO under grants CCF-1101659, IIS-0953756, CCF-1337215, W911NF-12-1-0550, and W911NF-11-1-0332, and by a Feodor Lynen research fellowship of the Alexander von Humboldt Foundation.

References

[1] J. S. Banks. Sophisticated voting outcomes and agenda control. *Social Choice and Welfare*, 3:295–306, 1985.

[2] J. Bartholdi, III and M. Trick. Stable matching with preferences derived from a psychological model. *Operations Research Letters*, 5(4):165–169, 1986.

[3] T. Besley and S. Coate. An economic model of representative democracy. *The Quarterly Journal of Economics*, 112(1):85–114, 1997.

[4] D. Black. On the rationale of group decision-making. *Journal of Political Economy*, 56 (1):23–34, 1948.

[5] F. Brandt, F. Fischer, P. Harrenstein, and M. Mair. A computational analysis of the tournament equilibrium set. *Social Choice and Welfare*, 34(4):597–609, 2010.

[6] Yvo Desmedt and Edith Elkind. Equilibria of plurality voting with abstentions. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 347–356, 2010.

[7] B. Dutta and J.-F. Laslier. Costless honesty in voting. Presented at the 10th International Meeting of the Society for Social Choice and Welfare, 2010.

[8] B. Dutta, M. O. Jackson, and M. Le Breton. Strategic candidacy and voting procedures. *Econometrica*, 69(4):1013–1037, 2001.

[9] B. Dutta, M. O. Jackson, and M. Le Breton. Voting by successive elimination and strategic candidacy. *Journal of Economic Theory*, 103(1):190–218, 2002.

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

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

[12] B. Escoffier, J. Lang, and M. Öztürk. Single-peaked consistency and its complexity. In *Proceedings of the 18th European Conference on Artificial Intelligence (ECAI)*, pages 366–370. IOS Press, 2008.

[13] R. Farquharson. *Theory of Voting*. Yale University Press, 1969.

[14] P. C. Fishburn. Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, 33(3):469–489, 1977.

[15] A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41:587–602, 1973.

[16] J. Lang, N. Maudet, and M. Polukarov. New results on equilibria in strategic candidacy. In *Proceedings of the 6th International Symposium on Algorithmic Game Theory (SAGT)*, volume 8146 of *Lecture Notes in Computer Science (LNCS)*, pages 13–25. Springer-Verlag, 2013.

[17] D. C. McGarvey. A theorem on the construction of voting paradoxes. *Econometrica*, 21(4):608–610, 1953.

[18] R. Meir, M. Polukarov, J. S. Rosenschein, and N. R. Jennings. Convergence to equilibria in plurality voting. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*, pages 823–828, 2010.

[19] 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.

[20] N. R. Miller. *Committees, Agendas and Voting*. Harwood Academic, 1995.

[21] H. Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4):437–455, 1980.

[22] H. Moulin. *The Strategy of Social Choice*. North-Holland, 1983.

[23] R. B. Myerson and R. J. Weber. A theory of voting equilibria. *The American Political Science Review*, 87(1):102–114, 1993.

[24] S. Obraztsova, E. Markakis, and D. R. M. Thomson. Plurality voting with truth-biased agents. In *Proceedings of the 6th International Symposium on Algorithmic Game Theory (SAGT)*, pages 26–37, 2014.

[25] M. J. Osborne and A. Slivinski. A model of political competition with citizen-candidates. *The Quarterly Journal of Economics*, 111(1):65–96, 1996.

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

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

[28] Y. Samejima. Strategic candidacy, monotonicity, and strategy-proofness. *Economic Letters*, 88(2):190–195, 2005.

[29] Y. Samejima. Strategic candidacy and single-peakedness. *The Japanese Economic Review*, 58(4):423–332, 2007.

[30] M. A. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975.

[31] K. A. Shepsle and B. R. Weingast. Uncovered sets and sophisticated outcomes with implications for agenda institutions. *American Journal of Political Science*, 28(1):49–74, 1984.

[32] D. R. M. Thomson, O. Lev, K. Leyton-Brown, and J. Rosenschein. Empirical analysis of plurality election equilibria. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 391–398, 2013.

Markus Brill
Department of Computer Science
Duke University
Durham, NC, USA
Email: `brill@cs.duke.edu`

Vincent Conitzer
Department of Computer Science
Duke University
Durham, NC, USA
Email: `conitzer@cs.duke.edu`