



HAL
open science

Stategic Candidacy with Keen Candidates

Jérôme Lang, Vangelis Markakis, Nicolas Maudet, Svetlana Obraztsova, Maria Polukarov, Zinovi Rabinovich

► **To cite this version:**

Jérôme Lang, Vangelis Markakis, Nicolas Maudet, Svetlana Obraztsova, Maria Polukarov, et al.. Strategic Candidacy with Keen Candidates. 2019. hal-02294985

HAL Id: hal-02294985

<https://hal.science/hal-02294985>

Preprint submitted on 23 Sep 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Strategic Candidacy with Keen Candidates

Jérôme Lang
LAMSADE, Univ. Paris-Dauphine,
France
lang@lamsade.dauphine.fr

Svetlana Obraztsova
Nanyang Technological University
svetlana.obraztsova@gmail.com

Vangelis Markakis
Athens University of Economics and
Business
markakis@gmail.com

Maria Polukarov
King's College London
maria.polukarov@kcl.ac.uk

Nicolas Maudet
Sorbonne Université, CNRS, LIP6,
F-75005 Paris, France
nicolas.maudet@lip6.fr

Zinovi Rabinovich
Nanyang Technological University
zinovi.rabinovich@gmail.com

ABSTRACT

In strategic candidacy games, both voters and candidates have preferences over the set of candidates, and candidates make strategic decisions about whether to run an electoral campaign or withdraw from the election, in order to manipulate the outcome according to their preferences. In this work, we extend the standard model of strategic candidacy games to scenarios where candidates may find it harmful for their reputation to withdraw from the election and would only do so if their withdrawal changes the election outcome for the better; otherwise, they would be keen to run the campaign. We study the existence and the quality of Nash equilibria in the resulting class of games, both analytically and empirically, and compare them with the Nash equilibria of the standard model. Our results demonstrate that while in the worst case there may be none or multiple, bad quality equilibria, on average, these games have a unique, optimal equilibrium state.

KEYWORDS

Computational Social Choice; Algorithmic Game-Theory; Voting; Strategic Candidacy

1 INTRODUCTION

In spite of its now distant relevance to the current affairs of the US, the presidential campaign of the year 2000 continues to stir much debate and discussion. It is a common opinion that the presence of the “Green Party” candidate, Ralph Nader, had contributed to Al Gore’s loss to George W. Bush. Nader is thought to have “siphoned away” a pivotal amount of votes from Gore. Taking into account the fact that Gore’s Democratic party was politically closer to Nader’s own stance, the latter’s decision to campaign appears against this background to be counter-intuitive, counter-productive and, even, irrational. Perhaps more surprising is the observation, particularly by political science research, that Nader’s behaviour is not unique. For instance, Bol et al [2015] note that the number of candidates in political elections under Plurality or Plurality with runoff schemes is typically higher than what the equilibrium analysis of the standard game-theoretic model would predict.

However, this latter, formal, study also suggest that such behaviour may not be as irrational as first impressions suggest. Rather,

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 2019, Montreal, Canada

© 2019 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.
<https://doi.org/doi>

it is the matter of the standard models not addressing additional benefits candidates may reap from an election campaign. Benefits that do not depend on the election outcome. Both independent candidates and parties are, quoting from Bol et al [2015], “motivated by other goals than winning, such as the activation of a local party section, the possibility of raising public awareness about certain issues, (...) or increasing their visibility.” The last point is particularly important from a long-term perspective, as “because they participate in other elections, parties need constant visibility and are likely to endorse a candidate for election (...), even if he has no chance of winning”.

In this work, we take these “other goals” into account and define a game-theoretic model for strategic candidacy, termed the *keen candidacy game*. Our model prescribes that in addition to the utility that a candidate draws from an election outcome, she gets an additional reward if she runs the campaign. As we will show, this indeed rationalises keen, zealous candidate participation. Furthermore, we show that even smallest amounts of rewards for zeal significantly impact candidate participation. What’s more, even if a party system is considered instead of a set of independent candidates, rewarded zeal can stabilise the elections by making candidate participation an equilibrium strategy. Going beyond rationalising the behaviour of politicians, however, by presenting a novel modification of candidacy games we complement the research on one of the hotter topics in the AI community.

Indeed, strategic candidacy in voting scenarios has attracted much attention of recent. The problem arises when potential candidates have their own preferences over possible outcomes of the election and are able to strategically decide whether to run in the election or withdraw/abstain. The latter choice was adopted by the Socialist Party during one of the recent regional elections in France. Unable to win, the Socialist Party preferred to withdraw and ensure the success of a centrist candidate against the right-wing National Front.

Such scenarios have been formalised as strategic games [5] and their equilibrium properties have been studied in a number of works [3, 6–8, 10, 13, 14, 16]. A common point of all these works is that the game-theoretic model assumes that the payoff of a player (i.e., a potential candidate) depends only on the outcome of the election: a candidate prefers a state s of the game to another state t if and only if she prefers the candidate winning in s to the candidate winning in t . Unfortunately, as our motivating scenario of Nader-Bush-Gore and, more generally, political science research suggest, this outcome centred reward assumption appears inaccurate and insufficient for real world domains. Indeed, candidates

derive additional value from their participation, independently of the elections' outcome. This additional value from running an electoral campaign can be negative (the campaign incurs some cost) or positive (participating in the election gives the candidate an opportunity to advertise his party or political platform and thus raises his profile/reputation). The former (negative) bias for participation has been studied in [11] considering a strategic candidacy model with so-called "lazy" candidates. In this work, we complete the study by considering the opposite (positive) bias.

In our *keen candidacy game (KCG)* model the utility that a candidate draws from an election outcome is augmented by an additional reward if she runs the campaign. This reward may be relatively small compared to the utility gap between two different outcomes, but need not always be so, as the Nader-Bush-Gore events demonstrate. We therefore consider several variants of our model, depending on whether the value of the bias for participation is *small* (more precisely, smaller than all differences between a candidate's values for two different outcomes), *large* (more precisely, larger than all differences between a candidate's values for two different outcomes), or *medium* (larger than some value differences but smaller than some others). We will also focus on a specific subclass of the medium bias model, where candidates are partitioned into parties, and are willing to run if and only if it does not prevent their party from winning.

We study the properties of pure strategy Nash equilibria in this class of games, and find significant differences with the standard (unbiased) model of strategic candidacy. Specifically, while in the standard model the existence of a Condorcet winner guarantees the existence of equilibria, in the keen candidacy games this is no longer true. Similarly, while for the Copeland voting rule equilibria always exist even in the absence of a Condorcet winner, keen candidacy games under Copeland may possess no equilibrium states.

Given this, we focus then on the question of how the system can be stabilised. Intuitively, by increasing the participation reward we should arrive at the state where everyone is keen to participate, so that "all in" is an equilibrium. Indeed, this intuition works for a large bias. In other cases, we provide bounds on the number of equilibria (for the multi-party case, in particular, these bounds are implied by a characterisation of its Nash equilibria), and on the quality of equilibria.

Finally, we perform simulations with random data, which demonstrate most encouraging results. Specifically, while in the worst case there may be instances with none, many and/or bad quality equilibria, on average (and in fact, for a vast majority of instances) we have a unique and optimal ("all in") equilibrium already for small values of the participation bias. Moreover, while larger rewards lead to predictions that are not very interesting (since it is typically the case that all candidates run), the setting with smaller rewards offers a better prediction power: first, it almost always leads to a small number of equilibria; second, it seems to agree with what is observed in real elections—the number of candidates at equilibrium is typically larger than in the standard model, but does not always converge into the "all in" state.

2 MODEL

There is a set of *potential candidates* $C = \{c_1, c_2, \dots, c_m\}$ and a set of *voters* $V = \{v_1, v_2, \dots, v_n\}$ such that $C \cap V = \emptyset$. Each voter $v \in V$ has a *preference order* $\succ_v \in \mathcal{L}(C)$ where $\mathcal{L}(C)$ is the set of all linear orders over C . The list of orders $P^V = (\succ_v)_{v \in V}$ is called the *voters' preference profile*.

An election proceeds as follows. First, a subset of candidates $A \subseteq C$ announce that they participate in the election; we refer to A as the *actual candidates*. Then, each voter $v \in V$ reports his preferences over the active candidates, which are obtained by restricting \succ_v to A . It is assumed that all voters report their preferences sincerely. Finally, a *voting rule* r takes the set A and the voters' preferences as input, and outputs a candidate $w \in A$; this candidate is called the *election winner*. Ties are broken according to a predetermined order of the candidates, denoted as \triangleleft . Common voting rules are:

Positional scoring rules. Such rules are associated with a scoring vector (s_1, \dots, s_m) where $s_1 \geq s_2 \geq \dots \geq s_m$ and $s_1 > s_m$. If a voter ranks a candidate at the j -th position, the candidate gets a score of s_j from this vote, and his total score is the sum of scores over all the votes. The candidate with the highest score (with ties broken according to \triangleleft) wins the election. The most popular representative of this family of rules is *Plurality* with the scoring vector $(1, 0, 0, \dots, 0)$.

Condorcet-consistent rules. A candidate x *beats* a candidate y , if a majority of voters rank x above y in P^V . A candidate x is a *Condorcet winner* if x beats y for all $y \in C \setminus \{x\}$. A voting rule r is *Condorcet-consistent* if $r(P^V) = \{x\}$ whenever there is a (unique) Condorcet winner x for P^V . The most popular representative of this family is *Copeland* that elects the candidate x maximising $|\{y \in C \mid x \text{ beats } y\}|$.

While voters are assumed to sincerely report their preferences, candidates make strategic decisions about whether to run in the election, based on their own preferences over the set C . Formally, each candidate $c \in C$ has two available actions: 1 (run) and 0 (abstain), and is endowed with a preference order \succ_c over C ; the list $P^C = (\succ_c)_{c \in C}$ is referred to as the *candidates' preference profile*. A candidate $c \in C$ has *self-supporting preferences* if $c \succ_c x$ for all $x \in C \setminus \{c\}$. Indeed, in settings with keen candidates it is natural to assume that they have such preferences; importantly, our negative results also hold for this special case.

A keen candidate would only withdraw his candidacy if that would change the election outcome for the better; however, if his withdrawal has no effect on the election outcome, he prefers to participate. In this paper, we will be particularly interested in the effect of the *rate* of this bias on the candidates' behaviour. To this end, it is convenient to equip the candidates with cardinal utilities, which are consistent with their preference orders. That is, each candidate c has a *utility function* $u_c : C \rightarrow \mathbb{R}$ such that $u_c(x) > u_c(y)$ if and only if $x \succ_c y$. Additionally, there is a *bias for participation*, $\epsilon \geq 0$.

Different value ranges of ϵ , visualised in Figure 1, define the following cases of our model:

Zero bias, $\epsilon = 0$: in this case the game is isomorphic to the standard model of strategic candidacy where a player prefers one strategy profile over the other solely based on his preference order over the candidates.




Figure 1: Possible values of ϵ

Small bias, $\epsilon < \min_{c, x, y \in C, x \neq y} |u_c(x) - u_c(y)|$: in this case the participation bias never overrides a candidate's preference between two outcomes of the election.

Medium bias, defined as $\min_{c, x, y \in C, x \neq y} |u_c(x) - u_c(y)| < \epsilon < \max_{c, x, y \in C, x \neq y} |u_c(x) - u_c(y)|$: the participation bias has an intermediate value that sometimes overrides a candidate's preference between two outcomes, but not always.

Large bias, $\epsilon > \max_{c, x, y \in C, x \neq y} |u_c(x) - u_c(y)|$: in this case the participation bias overrides everything else.

As a particularly interesting case of the medium bias model, we consider a typical scenario in the political scene, where the candidates are divided into parties and only the election of a candidate that belongs to another party makes a difference for them. Formally, we say that the set of candidates C is partitioned into parties P_1, \dots, P_k such that $P_i \cap P_j = \emptyset$ and $\cup_{i=1}^k P_i = C$. We denote $x \sim y$ if candidates x and y belong to the same party, and for each candidate $c \in C$ let $P(c)$ be the party that contains c . Now, in addition to the self-supporting preference assumption, we assume that each candidate c ranks the members of his party over the other parties' members: i.e., $x \succ_c y$ for all $x \in P(c)$ and $y \notin P(c)$. Also, we assume that each candidate is keen to run the electoral campaign except if it ruins the chances of his own party to win. For all $c, x, y \in C$ we thus have:

$$0 < \epsilon < u_c(x) - u_c(y) \quad \forall x \in P(c), y \notin P(c) \quad (1)$$

$$\epsilon > |u_c(x) - u_c(y)| \quad \forall x, y \in P(c) \vee x, y \notin P(c) \quad (2)$$

Now, the tuple $\langle C, V, P^V, r, \triangleleft, (u_c)_{c \in C}, \epsilon \rangle$ defines what we term the *keen candidacy game* (KCG), as follows. KCG is a strategic game, in which the set of players is C and each player's set of actions is $\{0, 1\}$. We denote the action (strategy) of a player $c \in C$ by s_c ; the vector $\mathbf{s} = (s_c)_{c \in C}$ is called a *strategy profile*. We identify a strategy profile \mathbf{s} with the set of actual candidates $A(\mathbf{s}) = \{c \in C \mid s_c = 1\}$. Consequently, each profile defines an outcome $w(\mathbf{s}) \in C$, which is simply the outcome of the voting rule r by voters in V over candidates in $A(\mathbf{s})$ when $A(\mathbf{s})$ is not empty (with ties broken according to \triangleleft), and when $A(\mathbf{s}) = \emptyset$, the outcome is c_{\triangleleft} —the highest ranked candidate according to \triangleleft . For a candidate c and a strategy profile \mathbf{s} , we denote $\mathbf{s} + c = A(\mathbf{s}) \cup \{c\}$ if $s_c = 0$ and $\mathbf{s} - c = A(\mathbf{s}) \setminus \{c\}$ if $s_c = 1$.

The utility of a player c from a profile \mathbf{s} is given by

$$U_c(\mathbf{s}) = u_c(w(\mathbf{s})) + b_c(\mathbf{s}) \quad (3)$$

where $b_c(\mathbf{s}) = \epsilon$ if $s_c = 1$ and $b_c(\mathbf{s}) = 0$ if $s_c = 0$.

We will be interested in (pure strategy) Nash equilibria of KCGs. Recall that a strategy profile \mathbf{s} is said to be an equilibrium if no player in the game can profitably deviate from that profile. Formally, given a game $\Gamma = \Gamma(C, V, P^V, r, \triangleleft, (u_c)_{c \in C}, \epsilon)$, we say that a strategy profile \mathbf{s} is a *pure strategy Nash equilibrium* (PSNE) of Γ if for every

candidate $c \in C$, $U_c(\mathbf{s}) \geq U_c(\mathbf{t})$ where \mathbf{t} is the strategy profile given by $t_x = s_x$ for $x \in C \setminus \{c\}$, and $t_c = 1 - s_c$.

3 WORST CASE ANALYSIS

We start with exploring the properties of KCGs analytically. Specifically, we observe several crucial differences with the standard model, discuss the number of possible equilibria and consider the special case with a multi-party structure.

3.1 Small bias model vs. unbiased model

For the unbiased model, it was shown that under Condorcet-consistent rules any subset of candidates that contains a Condorcet winner (if one exists) defines a PSNE profile [10]. In fact, as we show below, if the set of candidates has a Condorcet winner, then the corresponding candidacy game has a PSNE under *any* voting rule, as long as this rule respects majorities for the case of two candidates; moreover, there is always an equilibrium where the Condorcet winner is elected.

OBSERVATION 1. *Let Γ be a standard candidacy game (i.e., $\epsilon = 0$), with an arbitrary voting rule r respecting majority when there are two candidates, and assume there is a Condorcet winner c . Then, the singleton $\{c\}$ is a pure strategy Nash equilibrium of Γ .*

However, when the candidates are keen to participate in the election, this is no longer true, even for the small participation bias model. Quite the contrary, as we demonstrate in the following Examples 3.1 and 3.2, not only a Condorcet winner may lose the election in an equilibrium profile, there may be no equilibria at all, despite the existence of a Condorcet winner.

Example 3.1. Consider a KCG with 6 voters and 4 candidates $\{a, b, c, d\}$, where r is Plurality, P^V and P^C are as follows (with any consistent utilities u_c , $c \in C$) and ϵ is small:¹

$$P^V: \begin{array}{c|cccc} 2 & a & c & b & d \\ 2 & b & c & a & d \\ 1 & d & c & a & b \\ 1 & d & c & b & a \end{array} \quad P^C: \begin{array}{c|cccc} a & a & b & d & c \\ b & b & a & d & c \\ c & c & a & b & d \\ d & d & a & b & c \end{array}$$

Here, candidate c is a Condorcet winner. However, the profile $\{c\}$ is not a PSNE, as any other candidate wants to join the election, even if this would not change the outcome. In fact, the only equilibrium is $\{a, b, c, d\}$ and the winner is a . That is, even though the game has a PSNE, there is no equilibrium profile in which the Condorcet winner wins the election!

In fact, in contrast with Observation 1, a single candidate can never define an equilibrium of a keen candidacy game.

OBSERVATION 2. *In a KCG Γ , an equilibrium profile must contain at least two active candidates.*

Moreover, in contrast with Observation 1 for standard games, a KCG may not have a PSNE, even in the presence of a Condorcet winner, and even for a small participation bias.

Example 3.2. Consider a KCG with 4 voters and 4 candidates $\{a, b, c, d\}$, where r is Plurality and P^V , P^C are as follows (with any

¹In our examples, when the tie-breaking ordering \triangleleft is not specified it is assumed to be lexicographic/alphabetical. The first column in P^V indicates the number of voters casting the different ballots.

consistent utilities u_c , $c \in C$, and a small ϵ):

$$P^V: \begin{array}{c|cccc} 1 & a & d & c & b \\ 1 & b & a & d & c \\ 1 & c & b & a & d \\ 1 & d & a & b & c \end{array} \quad P^C: \begin{array}{c|cccc} a & a & d & c & b \\ b & b & a & d & c \\ c & c & b & a & d \\ d & d & a & b & c \end{array}$$

In this game, candidate a is a Condorcet winner. However, no strategy profile is stable under unilateral player deviations. Indeed, by Observation 2, singleton sets are not PSNE. Now, candidate a would like to join the election at any of the profiles $\{b, c\}$, $\{b, d\}$, $\{c, d\}$, $\{b, c, d\}$; candidate b would join $\{a, c\}$, $\{a, d\}$, $\{a, c, d\}$; c would join $\{a, b\}$ but leave from $\{a, b, c, d\}$; and d would join $\{a, b, c\}$ but leave from $\{a, b, d\}$. Hence, there is no Nash equilibrium in this game!

Even more interestingly, there is an analogous example for the Copeland rule, for which in the standard model a Nash equilibrium is guaranteed to exist even in the absence of a Condorcet winner [10].

Example 3.3. Consider a KCG with 4 candidates $\{a, b, c, d\}$, where r is the Copeland rule, and P^V is such that in pairwise elections candidate a beats candidate b ; candidate b beats candidate c ; c beats a and d ; and d beats a and b . The candidates' preference profile P^C is as follows (as before, assume any consistent utilities u_c , $c \in C$, and a small ϵ):

$$P^C: \begin{array}{c|cccc} a & a & d & b & c \\ b & b & c & a & d \\ c & c & d & b & a \\ d & d & a & b & c \end{array}$$

The score of candidates a and b is 1, and of c and d is 2; hence, candidate c wins the election by the lexicographic tie-breaking. Here, all strategy profiles are prone to deviations, and this is including the set $\{a, c, d\}$ containing the Copeland winner c and all the candidates beaten by him in pairwise elections, which was used to prove the PSNE existence in the standard model. Indeed, candidate b wants to join $\{a, c, d\}$ and obtain the reward for participation, even though his presence does not change the election's winner.

Moreover, even when a PSNE does exist, the Copeland winner of the full profile may not be an equilibrium winner.

Example 3.4. Consider a KCG with 4 candidates $\{a, b, c, d\}$, where r is Copeland, and P^V is such that in pairwise elections candidate a beats b and c ; b beats c and d ; and d beats a and c . Hence, the score of candidates a, b and d is 2, the score of candidate c is 0, and candidate a wins the election. As for the candidates' preferences, we only need to specify that candidate b prefers d over a , to show that the full set $\{a, b, c, d\}$ is not a PSNE (as b wants to leave), and the only equilibrium is $\{a, c, d\}$ where d (but not a) wins the election.

3.2 Medium bias: multi-party elections

We now move towards the multi-party model with a medium bias, as described in Section 2.

We first observe that in an equilibrium profile (if one exists), each party must be actively represented.

OBSERVATION 3. *If \mathbf{s} is a PSNE of a KCG Γ satisfying (1) and (2), then for any party P , $A(\mathbf{s}) \cap P \neq \emptyset$.*

Furthermore, as we show below, a profile \mathbf{s} is a PSNE of a KCG with a multi-party structure, if and only if all losing parties are represented in \mathbf{s} in full and every inactive candidate makes his party lose the election should he decide to run.

THEOREM 3.5. *Let Γ be a KCG satisfying (1) and (2). Then, \mathbf{s} is a PSNE of Γ iff the following conditions are met:*

- (i) $\forall P$ such that $w(\mathbf{s}) \notin P$, $P \subseteq A(\mathbf{s})$;
- (ii) $\forall P$ such that $w(\mathbf{s}) \notin P$ and $\forall c \in P$, $w(\mathbf{s} - c) \notin P$;
- (iii) $\forall c \notin A(\mathbf{s})$, $w(\mathbf{s} + c) \approx c$.

PROOF. First, let \mathbf{s} be a PSNE of Γ . Then, condition (i) is implied by the same arguments as Observation 3.

Assume on the contrary that condition (ii) is not satisfied: i.e., there are P and $c \in P$ such that $w(\mathbf{s}) \notin P$ and $w(\mathbf{s} - c) \in P$. But then, by (1), $U_c(\mathbf{s} - c) = u_c(w(\mathbf{s} + c)) > u_c(w(\mathbf{s})) + \epsilon = U_c(\mathbf{s})$; that is, c has an incentive to withdraw from the election, a contradiction.

Assume on the contrary that condition (iii) is not satisfied: i.e., there is $c \notin A(\mathbf{s})$ such that $w(\mathbf{s} + c) \sim c$. By condition (i), we have that $w(\mathbf{s}) \sim c$. But then, by (2), $U_c(\mathbf{s} + c) = u_c(w(\mathbf{s} + c)) + \epsilon > u_c(w(\mathbf{s})) = U_c(\mathbf{s})$; that is, c has an incentive to join the election, a contradiction.

Conversely, assume that all three conditions hold and show that \mathbf{s} is a PSNE. Let $w(\mathbf{s}) = w$ and let $c \in P(w)$, that is, $c \sim w$. If $c = w$, then $U_c(\mathbf{s}) = u_c(c) + \epsilon$ is maximal and c has no incentive to withdraw. If $c \neq w$ and $c \in A(\mathbf{s})$, then $U_c(\mathbf{s}) = u_c(w) + \epsilon > u_c(w(\mathbf{s} - c)) = U_c(\mathbf{s} - c)$, where the inequality holds by (1) if $w(\mathbf{s} - c) \approx w$ or by (2) if $w(\mathbf{s} - c) \sim w$; again, candidate c has no incentive to withdraw. If $c \notin A(\mathbf{s})$, then c has no incentive to join the election by condition (iii) and equation (1). Now, let $c \notin P(w)$. In this case, $c \in A(\mathbf{s})$ by condition (i), and by condition (ii) and equation (1), he has no incentive to withdraw from the election. \square

However, yet there are instances of multi-party KCGs where the three conditions stated in Theorem 3.5 are not satisfied and hence, the existence of PSNE is not guaranteed.

Example 3.6. Consider a multi-party KCG with two parties $P_1 = \{a, b\}$ and $P_2 = \{c, d\}$. The voting rule is (partially) described by the following choice function:

$$\begin{array}{c|ccc} & a & b & ab \\ \hline c & a & c & c \\ d & d & b & d \\ \hline cd & a & b & d \end{array}$$

That is, if the set of active candidate is given by $\{a, c\}$ then the winner of the election is candidate a ; if the active candidates are $\{a, b, c, d\}$ then d is elected, and so on.

We now check that for any strategy profile \mathbf{s} , one of the conditions (i), (ii) or (iii) of Theorem 3.5 is violated: (1) $\{a, c\}$, $\{b, c\}$, $\{a, d\}$ and $\{b, d\}$ violate condition (i) (for instance, in $\{a, c\}$, candidate d has an incentive to join); (2) $\{a, c, d\}$, $\{b, c, d\}$, $\{a, b, d\}$ and $\{a, b, c, d\}$ violate condition (ii): in the former two profiles, candidate c wants to withdraw, in the latter two, a wants to withdraw; (3) $\{a, b, c\}$ violates condition (iii) (d wants to join). Hence, there is no PSNE in this game.

Note that the characterization in Theorem 3.5 does not provide us with an efficient algorithm to decide whether a PSNE exists in a given KCG, as each party P has exponentially many subsets. However, if there is a constant bound on the size of the parties, this can be decided in poly-time.

COROLLARY 3.7. *In a multi-party KCG Γ , all PSNE can be found in time $O\left(\sum_{i=1}^k 2^{|P_i|}\right)$. Therefore, finding all PSNE (and a fortiori*

deciding whether there exists one) is fixed-parameter tractable in $\max_i |P_i|$.

3.3 Large bias for participation

Finally, we observe that if the value of the bias is large, a PSNE is guaranteed to exist; moreover, it is unique.

THEOREM 3.8. *Let Γ be a KCG with $\epsilon > \max_{c,x,y \in C} |u_c(x) - u_c(y)|$. Then, the full strategy profile $(1, \dots, 1)$ where all the candidates are running, is the unique PSNE of Γ .*

The proof of Theorem 3.8 is straightforward, since the bias for participation prevails the preferences over the candidates.

However, while the existence and uniqueness of PSNE is a very desirable property, Theorem 3.8 assumes an unrealistically large bias for participation. A natural question is therefore, by how much this value can be reduced, yet not eliminating or multiplying equilibrium profiles. As it turns out though, in the worst case this is as good as one can get.

3.4 Number of equilibria

To investigate further the question of the number of equilibria, it is useful to remark that the game can be encoded by the hypercube of dimension m . In the unbiased case, it is easy to see that there can be up to 2^{m-1} equilibria (no more, due to some structural properties of the hypercube, with the bound being reached when a candidate wins whenever he runs).

Even a small participation bias decreases this bound.

THEOREM 3.9. *Let Γ be a KCG with $m \geq 1$ candidates and $\epsilon > 0$. Then, the number of PSNE is at most $\frac{(m-2)2^{m-1}+1}{m-1}$.*

PROOF. The deviation graph constitutes an *oriented* hypercube where (1) each node has exactly m neighbours in the hypercube, (2) each equilibrium has only ingoing edges, (3) any node (except the one representing the empty state) must have *at least one* ingoing edge (it cannot be the case that all the candidates want to leave, since at least the winner is willing to stay). Now, using a counting argument, assume δ equilibria: this induces $m\delta + 2^m - \delta - 1$ edges. As there are overall $m2^{m-1}$ edges in the hypercube, the result follows. \square

For a medium bias, we focus again on the multi-party elections. According to Theorem 3.5, a Nash equilibrium in this case is of the form $S_t \cup (\cup_{t' \neq t} P_{t'})$, where the winner is of party P_t and conditions (ii) and (iii) are met. Let $|P_t| = m_t$, and δ_t the number of equilibria where the winning party is P_t . When $S_t \subseteq P_t$, we denote $(S_t)^+ = S_t \cup (\cup_{t' \neq t} P_{t'})$.

LEMMA 3.10. *Let Γ be a multi-party KCG. Then, $\delta_t \leq 2^{m_t-1}$.*

PROOF. Let $S_t \subseteq P_t$ and $c \in P_t \setminus S_t$. Now, (1) if $w((S_t)^+) \notin P_t$, then c joining $(S_t)^+$ is a profitable deviation, (2) if $w((S_t)^+) \in P_t$ and $w((S_t)^+ \cup \{c\}) \in P_t$, then c joining $(S_t)^+$ is a profitable deviation, and (3) if $w((S_t)^+) \in P_t$ and $w((S_t)^+ \cup \{c\}) \notin P_t$, then c leaving $(S_t)^+ \cup \{c\}$ is a profitable deviation. Therefore, there is either a deviation from $(S_t)^+$ to $(S_t)^+ \cup \{c\}$ or vice versa.

Consider the graph whose set of edges is 2^{P_t} and that contains a vertex from S to S' if there is a profitable deviation from $(S)^+$ to $(S')^+$. Due to the above observation, the number of edges in the

graph is $m_t \cdot 2^{m_t-1}$. Now, if S is a Nash equilibrium, then there are no outgoing edges from S and, because of (1), there are m_t ingoing edges from S . Thus, we have $m_t \delta_t \leq m_t 2^{m_t-1}$, that is, $\delta_t \leq 2^{m_t-1}$. \square

COROLLARY 1. *A multi-party KCG has $\leq \sum_{t=1}^k 2^{m_t-1}$ PSNE.*

We now show that for Plurality (but more generally for all scoring rules identified in [15], for which every choice function can be implemented by some voting profile) this upper bound is reached asymptotically for $m \geq 2$.

PROPOSITION 3.11. *Let Γ be a multi-party KCG with $m \geq 2$. For the Plurality rule, there are profiles for which the number of PSNE is at least $(\sum_{t=1}^k 2^{m_t-1}) - 2k$.*

PROOF. We first specify a profile for which the bound is reached asymptotically. We define the following constraints:

A(t, S_t) : $w((S_t)^+) \in P_t, \forall t = 1 \dots k, S_t \subseteq P_t, m_t - |S_t|$ even.

B(t, S_t) : $w((S_t)^+) \notin P_t, \forall t = 1 \dots k, S_t \subseteq P_t, m_t - |S_t|$ odd.

C($t_1, t_2, S_{t_1}, x_{t_2}$) : $w((S_{t_1})^+ \setminus \{x_{t_2}\}) \in P_{t_1}, \forall t_1, t_2 = 1 \dots k, t_2 \neq t_1, S_{t_1} \subseteq P_{t_1}, x_{t_2} \in P_{t_2}$.

It is easy (even if tedious) to verify that these constraints are globally consistent (we omit this part of the proof).

Now, recall that with Plurality (and most scoring rules, with the noticeable exception of Borda) we can implement every choice function by some profile [15]. Thus, we can build a profile satisfying all (A-B-C) constraints. From Theorem 2, for all $t = 1 \dots k, S_t \subseteq P_t, m_t - |S_t|$ even, $(S_t)^+$ is a Nash equilibrium: condition (i) holds because $w((S_t)^+) \in P_t$ due to (A), condition (ii) holds because if any candidate from P_t joins, then the winner is no longer in P_t due to the (B), and (iii) holds because of the (C) constraint. The number of PSNE is therefore $\sum_{t=1}^k \sum_{i=1}^{\lfloor \frac{m_t-1}{2} \rfloor} \binom{m_t}{m_t-2i}$.

Now, for m_t odd, $\sum_{i=1}^{\lfloor \frac{m_t-1}{2} \rfloor} \binom{m_t}{m_t-2i} = 2^{m_t-1} - 1$, and for m_t even, $\sum_{i=1}^{\lfloor \frac{m_t-1}{2} \rfloor} \binom{m_t}{m_t-2i} = 2^{m_t-1} - 2$. Hence, there are at least $\sum_{t=1}^k (2^{m_t-1} - 2) = (\sum_{t=1}^k 2^{m_t-1}) - 2k$ PSNE. \square

3.5 Quality of equilibria

The quality of equilibria is usually measured by the *price of anarchy* (PoA) [9], that compares the (worst possible) value of some social objective function at equilibrium and optimum states. In the context of strategic candidacy, inspired by [2], it is natural to compare the *initial scores* (i.e., when *all* the candidates are present) of the winners at the (worst) PSNE and the OPT ("all in") profiles respectively. However, because this absolute value of difference of score may highly depend on the voting rule used, we normalize it by dividing it by the maximum score that can be attained with the rule, $ms(\Gamma)$ —that is, n for plurality and $m - 1$ for Copeland.² Formally, let $S_c(s)$ denote the score obtained by a candidate $c \in A(s)$. The (additive) price of anarchy of a KCG Γ is then defined by:

$$PoA(\Gamma) = \max_{PSNE \text{ of } \Gamma} \{S_w(OPT)(OPT) - S_w(PSNE)(OPT) / ms(\Gamma)\} \quad (4)$$

THEOREM 3.12. *Let Γ be a KCG with n voters. Then, $PoA(\Gamma) = 1/2$ for Plurality; and $PoA(\Gamma) = 1$ for Copeland, as n and m grow. This holds for any ϵ smaller than a large bias.*

²Note that using a multiplicative version of PoA has its own caveat, as scores can be null.

PROOF. We prove for odd n under Plurality, the even case is similar. The upper bound is straightforward. If $PoA(\Gamma) > \frac{1}{2}$ then the winner of the OPT state gets the majority of votes under Plurality, and if he does not run in the equilibrium profile, he will join the election and win—a contradiction.

To show the lower bound, we use the following example with $\frac{n+5}{2}$ candidates $\{w_1, w_2, a_1, \dots, a_{\frac{n+1}{2}}\}$:

p^V					p^C					
1	...	1	...	1	w_1	w_2	a_1	...	$a_{\frac{n+1}{2}}$	
w_1	...	w_1	a_1	...	$a_{\frac{n+1}{2}}$	w_1	w_2	a_1	...	$a_{\frac{n+1}{2}}$
a_1	...	$a_{\frac{n-1}{2}}$	w_2	...	w_2	w_2	...	w_2
...	w_1	...	w_1
...
w_2	...	w_2	w_2	...	w_1	...	w_1

In this game, w_1 is the OPT winner with the score of $\frac{n-1}{2}$. Yet, he is the least favourite candidate for $a_1, \dots, a_{\frac{n+1}{2}}$, and in the PSNE $\{w_1, w_2\}$, w_2 (with initial score of 0) wins.

As for Copeland, remark that w_1 beats $m - 1$ candidates, wins in all-in, and is only beaten by w_2 , with score 1. But (w_1, w_2) is an equilibrium.

Finally, to see that our construction remains valid for any ϵ smaller than a large bias, observe that w_1 is the very last choice of candidates $a_1 \dots a_{\frac{n+1}{2}}$, the ones who could enter in equilibrium (w_1, w_2) . \square

So essentially, in terms of *scores*, strategic candidacy can yield bad outcomes, irrespective of the value of ϵ : an equilibrium where the winner would have a very low compared to the ones of the winner in the initial situation. Another natural measure would also be to compare the *rankings* obtained by candidates in these situations. The conclusions would be similar: by inspecting our construction, it can be seen that the winner in equilibrium (w_1, w_2) may be the one with the worst score, the last one in terms of ranking.

4 AVERAGE CASE ANALYSIS

Given that the existence and uniqueness of KCG equilibria can only be guaranteed by a large participation bias, it is crucial to understand how often instances with none or multiple PSNE can be expected in practice. To this end, we performed extensive simulations with randomly generated data.

4.1 Setup

Voting rules: We used Plurality and Copeland.

Number of candidates and voters: The number of candidates m ranged from 3 to 8. The number of voters n is set to 101, although we obtain the same results for larger n (note also, that increasing the number of voters does not significantly affect the computation time, as the crucial parameter for finding the number of equilibria is m but not n).

Generation of ordinal preferences: Both voters' and candidates' preferences were generated using an *impartial culture* (i.e. drawn uniformly from all possible orderings), with the constraint of being self-supporting for the candidates.

Cardinal utility functions for the candidates: We first generated random orderings as the preferences of both voters and candidates. Then, to determine cardinal utility functions for the candidates, we used Borda utilities: a candidate c has a utility of $m - 1$ points if his top choice is elected (i.e., himself), $m - 2$ points if his second best choice is elected, and so on, and has a zero utility if his last choice wins. In addition, he has a reward of ϵ if he participates in the election.

Values for ϵ : Since we used the integer Borda utilities, for any $\epsilon \in (0, 1)$, our model will have exactly the same set of equilibria, since all these values affect the final utility of each player in the same way. Similarly, all the values of ϵ in $(1, 2)$ induce the same game. Hence, we only needed to pick a single value from each such interval, so we used the set of values $\{0, 0.5, 1.5, 2.5, \dots, m - 1 + 0.5\}$.

4.2 Results

We report our findings for the Plurality and the Copeland rules. In both cases, and for ease of presentation and readability, we present figures for instances with $m = 5$ candidates. The same conclusions can be drawn for all other instances we explored, with different number of candidates.

Plurality. Figure 2 shows for each value of ϵ the distribution in terms of the number of equilibria. Recall that with $m = 5$ and with Borda utilities, as noted in the previous subsection, it only makes sense to consider values of ϵ below 4, thus we have used $\epsilon \in \{0, 0.5, 1.5, 2.5, 3.5\}$. The second axis shows the number of equilibria in each instance. Recall that we can have at most $16 = 2^{m-1}$ PSNE. Finally, the third axis shows, for each ϵ and for each possible number of equilibria, the percentage of instances that had these many equilibria, under this ϵ (averaged over 10,000 instances for each value of ϵ).




Figure 2: Plurality, 5 candidates, impartial culture

When there is no bias for participation, i.e., $\epsilon = 0$, we see that a quite likely situation (which occurs about 25% of the time) is to have the maximum possible number of equilibria (in this case, 16), with a bell-shaped distribution for other values. It is also particularly unlikely to have an instance without an equilibrium (although theoretically possible).

As soon as we introduce the smallest possible bias ϵ , we see an interesting phase transition. The distribution now concentrates massively on a single equilibrium (typically, it concentrates on the equilibrium where all the candidates participate). At the same time, the proportion of instances without equilibria at all becomes significant (about 10%). Clearly, as the bias increases, the figure shows that it becomes more and more likely to have a single equilibrium. E.g, when $\epsilon = 1.5$, this holds for more than 92% of the instances).

Copeland. Figure 3 shows an analogous picture for the equilibria under the Copeland rule. We also observe a phase transition occurring again as soon as ϵ becomes positive. However when $\epsilon = 0$, the likelihood of having the maximum possible number of equilibria is much higher—75% compared to 25% under Plurality (remember that the existence of equilibria is guaranteed for Copeland when $\epsilon = 0$). On the other hand, a small bias leads in most cases to a single equilibrium.




Figure 3: Copeland, 5 candidates, impartial culture

Finally, we comment on the additive price of anarchy observed in our experiments. Recall that this is the largest difference in the score (either Plurality or Copeland score) between the winner if all candidates are in and a winner at an equilibrium state (the—very few—instances without PSNE were ignored). We see that the PoA is consistently very low on average, tending to 0 as ϵ grows:

Bias ϵ	0	0.5	1.5	2.5	3.5
Plurality	0.022	0.011	0.005	0.001	0
Copeland	0.007	0.007	0.004	0.002	0

Thus, despite the higher theoretical bounds, in the vast majority of instances our model significantly refines the set of possible PSNE, converging to states with desirable properties, such as (typically) full participation in the election, that leads to socially preferable outcomes.

5 DISCUSSION AND FUTURE WORK

As we have started with a real world scenario, it was important to us to arrive at a model family that supports and rationalises the Nader-Bush-Gore example. Akin to the way that [12] follows the hunch exhibited by [18], we have incorporated the insight of

Bol et al [2015] into the the model of candidacy games. We have, however, gone far beyond the simple rationalisation of zealous candidate behaviour.

In this paper we have presented a study of strategic candidacy games augmented with a positive participation bias, terming the model *keen candidacy games (KCGs)*. We break the range of possible bias values into four categories, depending on how it relates to the maximal and the minimal difference between election outcomes. We find that even within the category of small biases, the set of equilibria in KCGs is critically different from the standard non-biased model. We present results on the number and the quality of the equilibria in terms of price of anarchy. The latter explicitly addresses the dangers of an overzealous candidate participation. However, as our experiments suggest, in practice political zeal is more likely to be a source of stability in the elections process – ensuring the existence of a single “all in” equilibrium behaviour with a socially preferable outcome. In this sense, our model suggests that Nader’s participation was a sign of a stable political system.

Now, it is important to notice a certain parallel with voting bias games (e.g. [4, 17]). In these games voters’ reward is also augmented by a contextual bias, and negative bias is implemented in terms that lead to abstention. It is natural, therefore, to expect that the positive bias would also create a conceptual simile between the voter and candidate game biases. However, the expectation is amiss. Positively or “truth” biased voters, receive additional reward when they can not influence the elections outcome and revert to expressing their innate preferences over candidates. In other words, the bias concerns preference expression, but not the preference itself. In case of KCGs the bias is at the core of a candidate’s preference system. In a sense, KCG preferences are over a two dimensional grid, and each candidate has to order all pairs of the form (winner, participation bonus).

Now, although we have performed an extensive study of positively biased candidacy games, and our results have their merit, they call for further research in new directions. In particular, we would like to investigate parallel, simultaneous elections, such as those that occur during primary elections. Besides a more complex structure for a participation bias to interfere with, such linked elections open the possibility to discuss bias as an externality. That is, a situation where the presence of a candidate in internal elections of one party may play an invigoration or suppressing role on the chances of candidates in other party elections.

ACKNOWLEDGMENTS

J. L. and N. M. are partially supported by the ANR project 14-CE24-0007-01 - CoCoRiCo-CoDec. S.O. and Z.R. are in part supported by NTU SUG M4081985 and NTU SUG M4082008 grants.

REFERENCES

- [1] Bol, D., Blais, A., Laslier, J.-F., Macé, A., 2015. Electoral system and number of candidates: Candidate entry under plurality and majority runoff.
- [2] Branzei, S., Caragiannis, I., Morgenstern, J., Procaccia, A. D., 2013. How bad is selfish voting? In: AAAI.
- [3] Brill, M., Conitzer, V., 2015. Strategic voting and strategic candidacy. In: AAAI. p. forthcoming.
- [4] Desmedt, Y., Elkind, E., June 2010. Equilibria of plurality voting with abstentions. In: EC. Cambridge, Massachusetts, pp. 347–356.
- [5] Dutta, B., Le Breton, M., Jackson, M. O., 2001. Strategic candidacy and voting procedures. *Econometrica* 69, 1013–1037.
- [6] Dutta, B., Le Breton, M., Jackson, M. O., 2002. Voting by successive elimination and strategic candidacy in committees. *Journal of Economic Theory* 103, 190–218.
- [7] Ehlers, L., Weymark, J. A., 2003. Candidate stability and nonbinary social choice. *Economic Theory* 22 (2), 233–243.
- [8] Eraslan, H., McLennan, A., 2004. Strategic candidacy for multivalued voting procedures. *Journal of Economic Theory* 117 (1), 29–54.
- [9] Koutsoupias, E., Papadimitriou, C., 1999. Worst-case equilibria. In: Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science. pp. 404–413.
- [10] Lang, J., Maudet, N., Polukarov, M., 2013. New results on equilibria in strategic candidacy. In: SAGT. pp. 13–25.
- [11] Obraztsova, S., Elkind, E., Polukarov, M., Rabinovich, Z., 2015. Strategic candidacy games with lazy candidates. In: Proceedings of IJCAI-15. pp. 610–616.
- [12] Obraztsova, S., Polukarov, M., Rabinovich, Z., Elkind, E., 2017. Doodle poll games. In: Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2017). pp. 876–884.
- [13] Polukarov, M., Obraztsova, S., Rabinovich, Z., Kruglyi, A., Jennings, N., 2015. Convergence to equilibria in strategic candidacy. In: Proceedings of IJCAI-15. pp. 624–630.
- [14] Rodriguez-Alvarez, C., 2006. Candidate stability and voting correspondences. *Social Choice and Welfare* 27 (3), 545–570.
- [15] Saari, D. G., 1989. A dictionary for voting paradoxes. *Journal of Economic Theory* 48 (2), 443 – 475.
- [16] Sabato, I., Obraztsova, S., Rabinovich, Z., Rosenschein, J. S., 2017. Real candidacy games: a new model for strategic candidacy. In: Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2017). pp. 867–875.
- [17] Thompson, D. R. M., Lev, O., Leyton-Brown, K., Rosenschein, J. S., May 2013. Empirical aspects of plurality election equilibria. In: AAMAS. pp. 391–398.
- [18] Zou, J., Meir, R., Parkes, D., 2015. Strategic voting behavior in doodle polls. In: Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work and Social Computing. CSCW '15. ACM, New York, NY, USA, pp. 464–472. URL <http://doi.acm.org/10.1145/2675133.2675273>