Complexity of Finding Equilibria of Plurality Voting Under Structured Preferences

Edith Elkind
University of Oxford
Oxford, United Kingdom
elkind@cs.ox.ac.uk

Evangelos Markakis
Athens University of Economics and Business
Athens, Greece
markakis@gmail.com

Svetlana Obraztsova
I-CORE, Hebrew University of Jerusalem
Jerusalem, Israel
svetlana.obraztsova@gmail.com

Piotr Skowron
University of Oxford
Oxford, United Kingdom
p.k.skowron@gmail.com

ABSTRACT

We study the complexity of finding pure Nash equilibria in voting games over well-known restricted preference domains, such as the domains of single-peaked and single-crossing preferences. We focus on the Plurality rule, and, following the recent work of Elkind et al. [15], consider three popular tie-breaking rules (lexicographic, random-candidate, and random-voter) and two types of voters' attitude: lazy voters, who prefer to abstain when their vote cannot affect the election outcome, and truth-biased voters, who prefer to vote truthfully in such cases. Elkind et al. [15] have shown that for most of these combinations of tie-breaking rules and voters' attitudes finding a Nash equilibrium is NP-hard; in contrast, we demonstrate that in almost all cases this problem is tractable for preferences that are single-peaked or single-crossing, under mild technical assumptions.

General Terms

Algorithms, Economics, Theory

Keywords

Nash equilibrium; Plurality; algorithms; single-peaked; single-crossing

1. INTRODUCTION

Since the famous Arrow's impossibility theorem [1] and its further consequences pointed out by Gibbard [21] and Satterthwaite [27], we know that virtually every voting rule creates incentives for the voters to act strategically, i.e., to misreport their preferences in order to enforce a more preferable outcome. These results are often viewed as disappointing, since negative consequences of strategic voting may be disastrous [19, 8]. Consequently, in order to understand various voting rules and to make a conscious decision as to which voting rule to use in a given scenario, one needs to study how these rules operate in a strategic environment.

Appears in: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016), J. Thangarajah, K. Tuyls, C. Jonker, S. Marsella (eds.), May 9–13, 2016, Singapore.*

As first pointed out by Farquharson [20], outcomes of elections with strategic voters can be better understood through analyzing the corresponding voting games, and in particular by describing their Nash equilibria. This line of research has received a considerable amount of attention in the social choice literature, see, e.g., [31, 3, 10]. Consequently, it is important to investigate the computational complexity of finding Nash equilibria in voting games: efficient algorithms for this problem would allow to predict an outcome of a specific election instance or to derive general conclusions by analyzing the structure of equilibria for real data describing voters' preferences [23].

In this paper we focus on equilibria of the Plurality rule, which is one of the most popular tools for making collective decisions. Under Plurality, each voter specifies her most preferred candidate or abstains from voting. If there exists a unique candidate who received the highest number of votes, then she is the winner; otherwise, a tie-breaking rule is used to select the winner from the set of top-scoring candidates. Popular tie-breaking rules include the *lexicographic rule*, where we fix an order over the set of candidates and select the candidate with the lowest rank in this order, the *random candidate rule*, where the winning candidate is selected uniformly at random from the set of top-scorers, and the *random voter rule*, where the ties are resolved according to the preferences of a single voter who is selected uniformly at random (the random voter rule is used to break ties under, e.g., the Schulze method [28]).

Unfortunately, Plurality voting admits highly counterintuitive equilibria: e.g., if there are at least 3 voters, even if all voters have identical rankings, the profile where they all vote for a candidate they rank last is a Nash equilibrium. Therefore, a common approach is to endow voters with secondary preferences over outcomes and study the equilibria of the resulting game. Recent work focused on two types of secondary preferences: *lazy voters*, who prefer to abstain if their vote cannot influence the outcome of elections, and the *truth-biased voters*, who prefer to vote truthfully in such a case; see, e.g., [3, 6, 29, 13, 22, 24, 32]. In particular, Elkind et al. [15], building on the earlier work of Desmedt and Elkind [11] and Obraztsova et al. [26], characterized the Nash equilibria of games that correspond to all six possible combinations of secondary preferences and tie-breaking rules, and showed that for five of them finding Nash equilibria is NP-hard (the only easy combination is lazy voters and lexicographic tie-breaking).

While these intractability results are discouraging, they rely on building complicated preference profiles. Therefore, one may hope that they can be circumvented by imposing additional structural constraints on voters' preferences. In particular, many problems in computational social choice have recently been shown to become easier when the voters' preferences are single-peaked [5] or single-crossing [25]; see, e.g., the work by Faliszewski et al. [18], Brandt et al. [7] and Faliszewski et al. [17] for results on manipulation, control and bribery, and by Betzler et al. [4] and Skowron et al. [30] on committee selection rules.

In this paper, we pursue this direction for the problem of finding Nash equilibria of Plurality voting. Specifically, we consider the five hard combinations of secondary preferences and tie-breaking rules studied by Elkind et al. [15], and, for each of them, investigate the complexity of finding a Nash equilibrium under single-peaked or single-crossing preferences. For almost all of these scenarios, we derive polynomial-time algorithms for our problem, under a mild technical assumption, which we refer to as *tie-consistency*. We view our results as a first step towards designing algorithms for finding Nash equilibria in realistic voting scenarios: while we do not expect typical real-life elections to be single-peaked or single-crossing, it is plausible that they are often not too far from having these properties, for an appropriate notion of distance, and one may be able to develop an efficient procedure for handling such 'almost structured' instances by building on our ideas.

2. PRELIMINARIES

Preferences, utilities, ballots. For each positive integer t, we set $[t] = \{1, \dots, t\}$. Let $N = [n]$ be a set of voters and let $C = \{c_1, \dots c_m\}$ be a set of *candidates*. Each voter $i \in N$ is endowed with a *utility function* $u_i : C \to \mathbb{N}$: for each $c_j \in C$ the quantity $u_i(c_j)$ is the intrinsic utility that i gains when c_j becomes the unique election winner. We assume that each voter assigns different utilities to different candidates. A *utility profile* is a list $\mathbf{u} = (u_1, \dots, u_n)$ of utility functions of all voters. For each $i \in N$, the utility function u_i induces a *preference order* $\succ_i$ over C: we set $c \succ_i c'$ if and only if $u_i(c) > u_i(c')$. A *preference profile* is a list of preference orders of all voters. For each voter $i \in N$, let a_i denote i's most preferred candidate, and let $\mathbf{a} = (a_1, \dots, a_n)$.

Under the Plurality rule each voter submits her *ballot* b_i, which is the name of a single candidate or, if i chooses to abstain, $b_i = \varnothing$. A *ballot profile* $\mathbf{b} = (b_1, \dots, b_n)$ is a list of ballots of all voters. We write $(\mathbf{b}_{-i}, b')$ to denote the profile obtained from $\mathbf{b}$ by replacing the ballot of the i-th voter, b_i, with b'. Let $\mathrm{sc}(c_j, \mathbf{b}) = |\{i \in N \mid b_i = c_j\}|$ denote the *score* of candidate c_j. Let $M(\mathbf{b}) = \max_{c \in C} \mathrm{sc}(c, \mathbf{b})$ be the maximum score of a candidate in the ballot vector $\mathbf{b}$. The *winning set* $W(\mathbf{b}) = \{c \in C \mid \mathrm{sc}(c, \mathbf{b}) = M(\mathbf{b})\}$ is the set of all candidates who received the highest score. The winner is selected from the winning set according to a tie-breaking rule. Additionally, we define the following sets that will be useful in our analysis: $H(\mathbf{b}) = \{c \in C \mid \mathrm{sc}(c, \mathbf{b}) = M(\mathbf{b}) - 1\}$, $H'(\mathbf{b}) = \{c \in C \mid \mathrm{sc}(c, \mathbf{b}) = M(\mathbf{b}) - 2\}$.

Tie-breaking rules. A *lottery* over a candidate set X is a vector $\mathbf{p} = (p_j)_{j \in X}$ with $p_j \geq 0$ for each $j \in X$ and $\sum_{j \in X} p_j = 1$. A *tie-breaking rule* returns a (possibly degenerate) lottery over the winning set. The *expected utility* of a voter $i \in N$ in an election with ballot profile $\mathbf{b}$ and tie-breaking rule R is computed as $\sum_{j \in W(\mathbf{b})} u_i(c_j) p_j$, where $\mathbf{p}$ is the lottery returned by R.

We consider three ways of resolving ties. The *lexicographic rule* R^L outputs the candidate $w \in W(\mathbf{b})$ with the lowest index, i.e., if $W(\mathbf{b}) = \{c_{j_1}, \dots, c_{j_k}\}$, $j_1 < \dots < j_k$, it sets $p_{j_1} = 1$, $p_{j_\ell} = 0$ for $\ell = 2, \dots, k$. The *random candidate rule* R^C returns the lottery $\mathbf{p}$ such that $p_j = 1/|W(\mathbf{b})|$ for each $c_j \in W(\mathbf{b})$. Under the *random voter rule* R^V we ask a random voter to pick her most preferred candidate from $W(\mathbf{b})$. Formally, the random voter rule returns the lottery $\mathbf{p}$ with $p_j = |B(c_j, W(\mathbf{b}))|/n$ for each $c_j \in W(\mathbf{b})$, where $B(c_j, W(\mathbf{b}))$ is the number of voters whose favorite candidate in $W(\mathbf{b})$ is c_j.

Lazy and truth-biased voters. Following Elkind et al. [15] (see also references therein), we consider *lazy* voters, who prefer to abstain when their vote has no effect on the election outcome, and *truth-biased* voters, who prefer to vote truthfully in such a case. These attitudes are formally captured by the formulas defining agents' overall utilities. Let ε be a constant satisfying $0 < \varepsilon < \min\{\frac{1}{m^2}, \frac{1}{n^2}\}$. Let $\mathbf{p}$ be the lottery produced by the tie-breaking rule on the candidate set $W(\mathbf{b})$. We say that a voter $i \in N$ is *lazy* if her utility at $\mathbf{b}$ is given by

$$U_i(\mathbf{b}) = \begin{cases} \sum_{c_j \in W(\mathbf{b})} p_j u_i(c_j) & \text{if } b_i \in C, \\ \sum_{c_j \in W(\mathbf{b})} p_j u_i(c_j) + \varepsilon & \text{if } b_i = \varnothing; \end{cases}$$

we say that she is *truth-biased* if her utility at $\mathbf{b}$ is given by

$$U_i(\mathbf{b}) = \begin{cases} \sum_{c_j \in W(\mathbf{b})} p_j u_i(c_j) & \text{if } b_i \in C \setminus \{a_i\}, \\ \sum_{c_j \in W(\mathbf{b})} p_j u_i(c_j) + \varepsilon & \text{if } b_i = a_i, \\ -\infty & \text{if } b_i = \varnothing. \end{cases}$$

Computational problems. From the above discussion we see that the voting game is determined by three parameters: the voters' attitude (lazy or truth-biased), the tie-breaking rule, and the voters' utility profile. Formally, a voting game is a triple $(\mathcal{S}, R, \mathbf{u})$, where $\mathcal{S} \in \{\mathcal{L}, \mathcal{T}\}$, $R \in \{R^C, R^V, R^L\}$, and $\mathbf{u}$ is the utility profile. The action space of each voter is $C \cup \{\varnothing\}$. A ballot vector $\mathbf{b}$ is a *pure Nash equilibrium (PNE)* of the game $(\mathcal{S}, R, \mathbf{u})$ if for each voter $i \in N$ and each action $b' \in C \cup \{\varnothing\}$ it holds that $U_i(\mathbf{b}) \geq U_i(\mathbf{b}_{-i}, b')$ (where the definition of U_i depends on $\mathcal{S}$ and R, as specified above). We are now ready to define the computational problems we consider.

Definition 1 (ExistNE, TieNE, SingleNE). *Fix $\mathcal{S} \in \{\mathcal{L}, \mathcal{T}\}$ and $R \in \{R^C, R^V, R^L\}$. In $(\mathcal{S}, R)$-*ExistNE *we are given a utility profile $\mathbf{u}$ and ask whether the voting game $(\mathcal{S}, R, \mathbf{u})$ admits a PNE. In $(\mathcal{S}, R)$-*TieNE *and $(\mathcal{S}, R)$-*SingleNE *we are given a utility profile $\mathbf{u}$ and a distinguished candidate $c_j \in C$; in $(\mathcal{S}, R)$-*TieNE *we ask whether G admits a PNE $\mathbf{b}$ with $|W(\mathbf{b})| > 1$ and $c_j \in W(\mathbf{b})$, and in $(\mathcal{S}, R)$-*SingleNE *we ask whether G admits a PNE $\mathbf{b}$ with $W(\mathbf{b}) = \{c_j\}$.*

Note that polynomial-time algorithms for SingleNE and TieNE can be used to solve ExistNE in polynomial time.

Structured preference domains. We will consider elections that are single-dimensional in the following sense: voters or candidates can be ordered along a single axis so that voters' preferences are consistent with this axis. Formalizing this idea gives rise to single-peaked [5] and single-crossing [25] profiles.

Definition 2. *Let $\lhd$ be a strict linear order over a candidate set C. A preference order $\succ$ over C is* single-peaked *with respect to $\lhd$ if for every triple of distinct candidates $a, b, c \in C$ it holds that if $a \lhd b \lhd c$ then $b \succ c$ or $b \succ a$. A preference profile $(\succ_1, \dots, \succ_n)$ over C is* single-peaked *if there exists a strict linear order $\lhd$ over C such that for every $i = 1, \dots, n$ the preference order $\succ_i$ is single-peaked with respect to $\lhd$.*

Definition 3. *A preference profile $(\succ_1, \dots, \succ_n)$ over a candidate set C is called* single-crossing *with respect to a voter order $(i_1, \dots, i_n)$ if for every pair of candidates $a, b \in C$ such that $a \succ_{i_1} b$ there exists a $p \in \{1, \dots, n\}$ such that $a \succ_{i_k} b$ for all $k \leq p$ and $b \succ_{i_k} a$ for all $k > p$.*

There exist polynomial-time algorithms that detect whether a given preference profile is single-peaked [2, 12, 16] or single-crossing [12, 14, 9], and, if so, find an order of candidates/voters that witnesses this. Therefore, in what follows, when considering single-peaked profiles over a candidate set $\{c_1, \dots, c_m\}$, we assume that they are single-peaked with respect to the candidate order $c_1 \lhd \dots \lhd c_m$, and when considering single-crossing profiles, we assume that they are single-crossing with respect to the voter order $(1, \dots, n)$.

Tie-consistent utilities. If a ballot profile $\mathbf{b}$ results in a tie, a voter i may change her vote to some $c \in W(\mathbf{b}) \setminus \{b_i\}$ to make c the unique election winner. In what follows, we will consider voters who make all such decisions in the same way, i.e., always prefer the lottery over $W(\mathbf{b})$ induced by the tie-breaking rule or always prefer their top candidate in $W(\mathbf{b}) \setminus \{b_i\}$. Formally, we say that a voter i is *tie-loving* if for every set of candidates X, $|X| \geq 3$, she weakly prefers the uniform lottery over X to her second most preferred candidate in X being the unique winner; we say that i is *tie-hating* if for every set of candidates X, $|X| \geq 3$, she weakly prefers her second most preferred candidate in X being the unique winner to the uniform lottery over X. A voter is *tie-consistent* if she is tie-loving or tie-hating. For one of our proofs we need a stronger notion: we say that a voter i is *strongly tie-loving* if for every set of candidates X, $|X| \geq 3$, she weakly prefers any lottery over X that assigns the probability at least $1/m$ to her top candidate in X over her second most preferred candidate in X being the unique winner.

There are many utility functions that ensure tie-consistency. For example, if a voter's utility for a candidate declines quickly with the candidate's position in that voter's preference order, she is likely to be tie-loving. More precisely, we say that a utility function u over a candidate set C, $|C| = m$, is *exponentially decreasing* if for every pair of candidates $c, c' \in C$ it holds that $u(c) > u(c')$ implies $u(c) \geq m \cdot u(c')$. If a voter's utility function u is exponentially decreasing, then she is strongly tie-loving: indeed, if she ranks the candidates in X as $c \succ c' \succ \dots$ then under the lottery over X that assigns probability at least $\frac{1}{m}$ to c her utility is at least $\frac{1}{m}u(c)$ and if c' wins, her utility is $u(c') \leq \frac{1}{m}u(c)$. Conversely, if a voter's utility declines very slowly with the candidate's ranking, she is likely to be tie-hating. For instance, consider a voter whose utility function is given by $u(c_j) = 2^m - 2^j$ for $j \in [m]$. If her second most preferred candidate in a set X, $|X| \geq 3$, is c_k, then her utility from the uniform lottery over X is at most $2^m - \frac{1}{|X|}(2 + 2^k + (|X| - 2)2^{k+1}) < 2^m - 2^k = u(c_k)$. However, if a voter's utility function is linear (i.e., she assigns a utility of $m - i$ to her i-th most preferred candidate) she is neither tie-loving nor tie-hating: if $|C| = 4$ and she ranks the candidates as $a \succ b \succ c \succ d$, she prefers b to the uniform lottery over $\{a, b, d\}$, but she prefers the uniform lottery over $\{a, c, d\}$ to c.

For most of our results, we assume that all voters in the input profile are tie-consistent; note that we allow profiles where tie-loving and tie-hating voters coexist. We need this assumption for technical reasons; while it is clearly restrictive, as argued above, it is satisfied by some interesting utility functions. We believe that (strongly) tie-loving voters are particularly natural: such voters are optimistic about the coin tosses of the tie-breaking rule, and do not want to deny their preferred candidate a chance of winning. Importantly, all hardness results of Elkind et al. [15] hold for voters with exponentially decreasing utilities, so tie-consistency alone does not make our computational problems easy.

3. LAZY VOTERS

Elkind et al. [15] have shown that for lazy voters and lexicographic tie-breaking, all our computational problems are in P, and that $(\mathcal{L}, R)$-SingleNE is easy even for randomized tie-breaking rules. Therefore, we focus on ExistNE and TieNE, and on randomized tie-breaking rules.

We say that a candidate set $X = \{c_{\ell_1}, \dots, c_{\ell_k}\}$ with $2 \leq k \leq \min(n/2, m)$ is *balanced with respect to a utility profile* $\mathbf{u}$ if the voters can be split into k groups $N_1, \dots, N_k$ of size n/k each so that for all $j \in [k]$, $i \in N_j$ and $c \in X \setminus \{c_{\ell_j}\}$ we have $c_{\ell_j} \succ_i c$; we say that it is *strongly balanced* if, moreover, $\frac{1}{k}\sum_{c \in X} u_i(c) \geq \max_{c \in X \setminus \{c_{\ell_j}\}} u_i(c)$. Elkind et al. [15] characterize PNE of Plurality voting for lazy voters and randomized tie-breaking.

Theorem 1 (Elkind et al. [15], Theorem 2). *Let $\mathbf{u} = (u_1, \dots, u_n)$ be a utility profile over C, $|C| = m$, and let $R \in \{R^C, R^V\}$. Then the game $G = (\mathcal{L}, R, \mathbf{u})$ admits a PNE if and only if one of the following conditions holds:*

1. *all voters rank some candidate $c_j \in C$ first;*
2. *each candidate is ranked first by at most one voter, and $\forall \ell \in N$: $\frac{1}{n}\sum_{i \in N} u_\ell(a_i) \geq \max_{i \in N \setminus \{\ell\}} u_\ell(a_i)$.*
3. *there exists a strongly balanced set of candidates for $\mathbf{u}$.*

Elkind et al. [15] use this characterization to show that $(\mathcal{L}, R)$-ExistNE and $(\mathcal{L}, R)$-TieNE with $R \in \{R^C, R^V\}$ are NP-hard for general preferences, even if voters' utilities are exponentially decreasing. In contrast, we will now show that for tie-consistent voters these problems become easy if the voters' preferences are single-crossing or single-peaked.

Theorem 2. *$(\mathcal{L}, R)$-ExistNE and $(\mathcal{L}, R)$-TieNE with $R \in \{R^C, R^V\}$ admit polynomial-time algorithms if voters are known to be tie-consistent and the preference profile induced by $\mathbf{u}$ is single-crossing with respect to the voter order $(1, \dots, n)$.*

Proof. We will show that for single-crossing preferences the conditions of Theorem 1 can be verified in polynomial time. Conditions (1) and (2) are straightforward to check. Further, it is easy to check if any of the voters is tie-hating (given that a voter is tie-consistent, we can check if she is tie-hating by checking her utilities for at most two triples of candidates), and if that is the case, no set is strongly balanced. Finally, for tie-loving voters every balanced set is strongly balanced. Hence, in the rest of the proof we will show how to check if $\mathbf{u}$ admits a balanced set of a fixed size $k \in \{2, \dots, \min(n/2, m)\}$. We assume that k divides n, since otherwise the answer is obviously "no".

We first prove that the voter partition witnessing that X is balanced has to be such that the voters in each group form a block in $(1, \dots, n)$.

LEMMA 1. *If a voter partition $N_1, N_2, \dots, N_k$ witnesses that a subset of candidates X is balanced, then the groups of voters can be renumbered so that for each $j \in [k]$ we have $N_j = \{(j-1)k+1, \dots, jk\}$.*

PROOF. Assume for the sake on contradiction that there exist $c, c' \in X$ and $1 \le i < i' < i'' \le n$ such that voters i and i'' prefer c to all other candidates in X, but voter i' prefers c' to all other candidates in X. But then the pair (c, c') violates the condition in Definition 3, a contradiction. □

We will now use Lemma 1 to develop a graph-theoretic algorithm for our problem. Our solution characterizes all size-k winning sets for PNE ballot vectors as paths in a certain graph.

Assume that $N_j = \{(j-1)k+1, \dots, jk\}$ for $j \in [k]$. We construct a directed graph $\mathcal{G}^k = (\mathcal{V}^k, \mathcal{E}^k)$ as follows. Let

$$\mathcal{V}^k = \{s, t\} \cup \{(j, \ell) \mid j \in [k], c_\ell \in C\}.$$

The set $\mathcal{E}^k$ contains edges $(s, (1, \ell))$ and $((k, \ell), t)$ for each $c_\ell \in C$. Further, for each $j \in [k-1]$ and every $c_\ell, c_r \in C$ there is an edge from (j, ℓ) to $(j+1, r)$ if and only if the last voter in N_j prefers c_ℓ to c_r, but the first voter in N_{j+1} prefers c_r to c_ℓ. This completes the description of $\mathcal{G}^k$.

By construction, every s–t path in $\mathcal{G}^k$ is of the form

$$P = (s, (1, \ell_1), \dots, (k, \ell_k), t).$$

We claim that for any such path $r \neq j$ implies $\ell_r \neq \ell_j$, i.e., P corresponds to a subset of candidates $X(P) = \{c_{\ell_1}, \dots, c_{\ell_k}\}$ of size k. Moreover, for every set $X \subseteq C$ of size k there is a path P in $\mathcal{G}^k$ such that $X = X(P)$ if and only if there exists a one-to-one mapping π between $\{N_1, \dots, N_k\}$ and X such that for each $j \in [k]$ all voters in N_j prefer $\pi(N_j)$ to all other candidates in X.

Indeed, consider a path $P = (s, (1, \ell_1), \dots, (k, \ell_k), t)$ and a mapping π given by $\pi(N_j) = c_{\ell_j}$ for all $j \in [k]$. We will argue that for each $j \in [k]$ all voters from the set N_j prefer c_{ℓ_j} to c_{ℓ_r} for all $r \in [k] \setminus \{j\}$; this also shows that $\ell_r \neq \ell_j$ whenever $r \neq j$. Indeed, pick an arbitrary $j \in [k]$ and a voter $i \in N_j$. Consider a candidate c_{ℓ_r} with $r < j$. Since $((r, \ell_r), (r+1, \ell_{r+1}))$ is an edge of P, the last voter in N_r prefers c_{ℓ_r} to $c_{\ell_{r+1}}$, but the first voter in N_{r+1} prefers $c_{\ell_{r+1}}$ to c_{ℓ_r}. Since our election is single-crossing, it follows that voter i prefers $c_{\ell_{r+1}}$ to c_{ℓ_r}. As this holds for any $r < j$ and i's preference order is transitive, it follows that i prefers c_{ℓ_j} to c_{ℓ_r} for all $r < j$. The case $r > j$ is analyzed similarly.

Conversely, consider a set $X = \{c_{\ell_1}, \dots, c_{\ell_k}\}$ and a one-to-one mapping π such that for each $j \in [k]$ all voters from the set N_j prefer $\pi(N_j)$ to all other candidates in X. By renumbering the candidates in X, we can assume that $\pi(N_j) = c_{\ell_j}$ for all $j \in [k]$. We will argue that $(s, (1, \ell_1), \dots, (k, \ell_k), t)$ is a path in $\mathcal{G}^k$. Indeed, by construction $\mathcal{G}^k$ contains edges $(s, (1, \ell_1))$ and $((k, \ell_k), t)$. Now, consider some $r \in [k-1]$. The voters in N_r prefer c_{ℓ_r} to all other candidates in X. Thus, in particular, the last voter in N_r prefers c_{ℓ_r} to $c_{\ell_{r+1}}$. On the other hand, the voters in N_{r+1} prefer $c_{\ell_{r+1}}$ to all other candidates in X. Thus, in particular, the first voter in N_{r+1} prefers $c_{\ell_{r+1}}$ to c_{ℓ_r}. This means that $\mathcal{G}$ contains the edge $((r, \ell_r), (r+1, \ell_{r+1}))$. As this holds for every $r \in [k-1]$, our claim is proved.

To find a balanced set of candidates of size k, we construct the graph $\mathcal{G}^k$ as described above and check if it has a path from s to t. We output "yes" if this is the case for at least one value of $k \in \{2, \dots, \min(n/2, m)\}$. For each k, this check can be performed in polynomial time. This completes the proof for $(\mathcal{L}, R)$-EXISTNE.

For $(\mathcal{L}, R)$-TIENE, we modify our algorithm as follows. It is straightforward to check if c_p can be in $W(\mathbf{b})$ for some equilibrium ballot vector $\mathbf{b}$ corresponding to condition (2). To see whether c_p is contained in a strongly balanced set, we first verify that all voters are tie-loving. Then, for each $k = 2, \dots, \min(n/2, m)$ we construct the graph $\mathcal{G}^k$ as described above and, for each $i \in [k]$, check whether $\mathcal{G}^k$ contains an s–t path passing through (i, p). We output "yes" if the answer is positive for some values of k and i. Clearly, this procedure can be implemented in polynomial time. □

THEOREM 3. *$(\mathcal{L}, R)$-EXISTNE and $(\mathcal{L}, R)$-TIENE with $R \in \{R^C, R^V\}$ admit polynomial-time algorithms if voters are known to be tie-consistent and the preference profile induced by $\mathbf{u}$ is single-peaked with respect to the candidate order $c_1 \lhd \dots \lhd c_m$.*

PROOF. Just as in the proof of Theorem 2, it suffices to show how to find a balanced set of candidates of size k, for a value of k that divides n. Fix $k \ge 2$. For each $s \in [k] \setminus \{1\}$, $i \in [m]$, $j \in [m] \setminus [i]$, set $A[s, i, j] = 1$ if there is a set of candidates $Y = \{c_{\ell_1}, c_{\ell_2}, \dots, c_{\ell_s}\}$ with $\ell_1 < \ell_2 < \dots < \ell_s$ such that $c_{\ell_{s-1}} = c_i$, $c_{\ell_s} = c_j$, and each candidate in $Y \setminus \{c_j\}$ is preferred to all other candidates in Y by exactly n/k voters; otherwise, set $A[s, i, j] = 0$. Clearly, there exists a balanced set X of size k if and only if $A[k, i, j] = 1$ for some $i, j \in [m]$; indeed, if Y is a set witnessing that $A[k, i, j] = 1$, we can set $X = Y$. We will now show how to compute the quantities $A[s, i, j]$ inductively.

Clearly, we have $A[2, i, j] = 1$ if and only if exactly n/k voters prefer c_i to c_j. For the inductive step, we use the following lemma.

LEMMA 2. *For any $s = 2, \dots, k$, any $i = 1, \dots, m$ and any $j = i+1, \dots, m$ we have $A[s+1, i, j] = 1$ if and only if there exists an $r < i$ such that $A[s, r, i] = 1$ and exactly n/k voters prefer c_i to both c_r and c_j.*

PROOF. If $A[s+1, i, j] = 1$, let $Y = \{c_{\ell_1}, c_{\ell_2}, \dots, c_{\ell_{s+1}}\}$ be a set that witnesses this, where $c_{\ell_s} = c_i$, $c_{\ell_{s+1}} = c_j$. Let $r = \ell_{s-1}$. By construction exactly n/k voters prefer c_i to all candidates in Y, and, in particular, to c_r and c_j. Further, the set $Y \setminus \{c_j\}$ witnesses that $A[s, r, i] = 1$.

Conversely, suppose that $A[s, r, i] = 1$ for some r such that exactly n/k voters prefer c_i to both c_r and c_j. Let Y' be a set witnessing that $A[s, r, i] = 1$ and set $Y = Y' \cup \{c_j\}$. We need to show that each candidate in Y' is preferred to all candidates in Y by exactly n/k voters. Consider first candidate c_i. There are exactly n/k voters who prefer c_i to c_r and c_j. Since our election is single-peaked and $r < i$, each of these voters prefers c_i to any candidate c_p with $p < r$, and hence to all candidates in Y. Now, consider a candidate $c_p \in Y' \setminus \{c_i\}$. There are exactly n/k voters who prefer c_p to any candidate in Y', and, in particular, to c_i. Since our election is single-peaked and $i < j$, each of these voters has to prefer c_i to c_j, so by transitivity each of these voters prefers c_p to any candidate in Y. This completes the proof. $\square$

Using Lemma 2, we can efficiently compute $A[\ell, i, j]$ for all $s \in [k] \setminus \{1\}$, $i \in [m]$, $j \in [m] \setminus [i]$. We output "yes" if $A[k, i, j] = 1$ for some $i, j \in [m]$. This completes the proof for $(\mathcal{L}, R)$-ExistNE.

We now consider $(\mathcal{L}, R)$-TieNE. As argued in the proof of Theorem 2, it suffices to design an efficient algorithm that, given a value of $k \in \{2, \dots, \min(n/2, m)\}$, checks whether $c_p \in X$ for some set balanced X of size k. To solve this problem, we modify our dynamic program as follows. For each $q \in [p]$, we set $A^q[s, i, j] = 1$ if there exists a set $Y = \{c_{\ell_1}, c_{\ell_2}, \dots, c_{\ell_s}\}$ with $\ell_1 < \ell_2 < \dots < \ell_s$ such that $c_{\ell_{s-1}} = c_i$, $c_{\ell_s} = c_j$, each candidate in $Y \setminus \{c_j\}$ is preferred to all other candidates in Y by exactly n/k voters, and if $q \le s$ then $\ell_q = p$; otherwise, set $A^q[s, i, j] = 0$. The quantities $A^q[s, i, j]$ can be computed similarly to $A[s, i, j]$. Specifically, for $s < q$ we have $A^q[s, i, j] = A[s, i, j]$, for $s = q$ we have $A^q[s, i, j] = 1$ if $A[s, i, j] = 1$ and $j = p$ and $A^q[s, i, j] = 0$ otherwise, and for $s > q$ we have $A^q[s, i, j] = 1$ if $A^q[s - 1, r, i] = 1$ for some $r < i$ and exactly n/k voters prefer c_i to both c_r and c_j, and $A^q[s, i, j] = 0$ otherwise. Now it is easy to see that $c_p \in X$ for some balanced set X of size k if and only if $A^q[k, i, j] = 1$ for some $q \le p$ and some $i, j \in [m]$. $\square$

4. TRUTH-BIASED VOTERS

We will now consider truth-biased voters. We first present our results for randomized tie-breaking. Again, we build on the work of Elkind et al. [15], who characterize PNE of Plurality voting for this case. Their characterization are presented below (Theorem 4 deals with ties, and Theorem 5 describes equilibria with a unique winner).

Theorem 4 (Elkind et al. [15], Theorem 4). *Let $\mathbf{u} = (u_1, \dots, u_n)$ be a utility profile over C, $|C| = m$, and let $R \in \{R^C, R^V\}$. The game $G = (\mathcal{T}, R, \mathbf{u})$ admits a PNE with a winning set of size at least 2 if and only if one of the following conditions holds:*

1. *each candidate is ranked first by at most one voter, and, moreover, $\frac{1}{n}\sum_{i \in N} u_\ell(a_i) \ge \max_{i \in N \setminus \{\ell\}} u_\ell(a_i)$ for each $\ell \in N$;*
2. *there exists a strongly balanced set of candidates for $\mathbf{u}$.*

Further, if condition (1) holds, then G has the truthful ballot vector $\mathbf{a}$ as a PNE, and if condition (2) holds for some set X, then G has a PNE where each voter votes for her favorite candidate in X. The game G has no other PNE.

Theorem 5 (Elkind et al. [15], Theorems 5 & 6). *Let $\mathbf{u} = (u_1, \dots, u_n)$ be a utility profile over C, let $R \in \{R^C, R^V\}$, and consider a ballot vector $\mathbf{b}$ with $W(\mathbf{b}) = \{c_j\}$ for some $c_j \in C$. If $\mathbf{b} = \mathbf{a}$, then $\mathbf{b}$ is a PNE of the game $G = (\mathcal{T}, R, \mathbf{u})$ if and only if for every $i \in N$ and every $c_k \in H(\mathbf{a}) \setminus \{a_i\}$ it holds that $c_j \succ_i c_k$. On the other hand, if $\mathbf{b} \ne \mathbf{a}$, then $\mathbf{b}$ is a PNE of $G = (\mathcal{T}, R, \mathbf{u})$ if and only if all of the following conditions hold:*

1. *$b_i \in \{a_i, c_j\}$ for all $i \in N$;*
2. *$H(\mathbf{b}) \ne \varnothing$;*
3. *$c_j \succ_i c_k$ for all $i \in N$ and all $c_k \in H(\mathbf{b}) \setminus \{b_i\}$;*
4. *for each candidate $c_\ell \in H'(\mathbf{b})$ each voter $i \in N$ with $b_i = c_j$ prefers c_j to the lottery where a candidate is chosen from $H(\mathbf{b}) \cup \{c_j, c_\ell\}$ according to R.*

Finally, Proposition 1 establishes a useful property of PNE for truth-biased voters.

Proposition 1 (Elkind et al. [15], Proposition 2). *For every $R \in \{R^L, R^C, R^V\}$ and every utility profile $\mathbf{u}$, if a ballot vector $\mathbf{b}$ is a PNE of $(\mathcal{T}, R, \mathbf{u})$ then for every voter $i \in N$ either $b_i = a_i$ or $b_i \in W(\mathbf{b})$.*

Elkind et al. [15] show that ExistNE, TieNE and SingleNE are NP-complete for truth-biased voters and randomized tie-breaking, even if voters' utilities are exponentially decreasing. We will now show that these hardness results disappear when preferences are single-peaked or single-crossing, under an appropriate assumption on the voters' attitude to ties. Interestingly, for TieNE it suffices to assume that voters are tie-consistent, whereas for SingleNE (which is not even concerned with ties!) we need the stronger assumption that all voters are (strongly) tie-loving.

Theorem 6. $(\mathcal{T}, R)$-TieNE *with $R \in \{R^C, R^V\}$ (respectively, $(\mathcal{T}, R^C)$-SingleNE and $(\mathcal{T}, R^V)$-SingleNE) admits a polynomial-time algorithm if all voters are tie-consistent (respectively, tie-loving and strongly tie-loving) and the preference profile induced by $\mathbf{u}$ is single-peaked or single-crossing.* $(\mathcal{T}, R)$-ExistNE *with $R \in \{R^C, R^V\}$ is in* P *whenever* $(\mathcal{T}, R)$-SingleNE *is in* P.

Proof. Theorem 4 is very similar (though not identical) to Theorem 1, so the algorithms described in the proofs of Theorem 2 (for single-crossing preferences) and Theorem 3 (for single-peaked preferences) can be used to solve $(\mathcal{T}, R)$-TieNE for $R \in \{R^C, R^V\}$ in polynomial time when voters are tie-consistent. Thus, from now on we focus on $(\mathcal{T}, R)$-SingleNE.

We first consider single-peaked preferences. Fix $R \in \{R^C, R^V\}$. We will describe an efficient algorithm that, given a utility profile $\mathbf{u}$, a candidate $c_j \in C$ and a score s, decides whether the game $G = (\mathcal{T}, R, \mathbf{u})$ has a PNE $\mathbf{b}$ with $W(\mathbf{b}) = \{c_j\}$ such that $\mathrm{sc}(c_j, \mathbf{b}) = s$ and $a_i \ne b_i$ for some $i \in N$. To solve $(\mathcal{T}, R)$-SingleNE, we check whether the truthful profile $\mathbf{a}$ is a PNE (which is easy), and then run our algorithm for c_p and for all $s \in [n]$.

We will first describe a few key properties of PNE for truth-biased voters. Suppose that $G = (\mathcal{T}, R, \mathbf{u})$ has a PNE $\mathbf{b}$ with $W(\mathbf{b}) = \{c_j\}$ such that $\mathrm{sc}(c_j, \mathbf{b}) = s$ and $a_i \ne b_i$ for some $i \in N$. Then by Theorem 5 $H(\mathbf{b}) = \{c \mid \mathrm{sc}(c, \mathbf{b}) = s - 1\} \ne \varnothing$. Further, $\mathrm{sc}(c_j, \mathbf{a}) < s$ by Proposition 1.

Let $X = \{c \mid \mathrm{sc}(c, \mathbf{a}) \ge s\}$, $Y = \{c \mid \mathrm{sc}(c, \mathbf{a}) = s - 1\} \setminus \{c_j\}$, $Z = \{c \mid \mathrm{sc}(c, \mathbf{a}) = s - 2\} \setminus \{c_j\}$.

Lemma 3. *$\mathrm{sc}(c, \mathbf{b}) \le s - 3$ for all $c \in X$, $\mathrm{sc}(c, \mathbf{b}) \le s - 3$ or $\mathrm{sc}(c, \mathbf{b}) = s - 1$ for all $c \in Y$, $\mathrm{sc}(c, \mathbf{b}) \le s - 2$ for all $c \in Z$.*

PROOF. In $\mathbf{b}$ candidate c_j is the unique winner with s points, and, by Proposition 1, all voters that vote non-truthfully vote for him. This implies that $\mathrm{sc}(c, \mathbf{b}) \leq s - 2$ for every $c \in Z$. $\mathrm{sc}(c, \mathbf{b}) \leq s - 1$ for every $c \in Y$, and $\mathrm{sc}(c, \mathbf{b}) \leq s - 1$ for every $c \in X$. Now, if $c \in X$ and $\mathrm{sc}(c, \mathbf{b}) = s - 1$, then there exists a voter i with $a_i = c$, $b_i = c_j$. But then $c \in H(\mathbf{b})$ and $c \succ_i c_j$, a contradiction with condition (3) of Theorem 5. Similarly, if $c \in X$ and $\mathrm{sc}(c, \mathbf{b}) = s - 2$, then $c \in H'(\mathbf{b})$ and there exists a voter i with $a_i = c$, $b_i = c_j$. If this voter changes her vote back to c, the outcome will be chosen from $H(\mathbf{b}) \cup \{c, c_j\}$ according to R. If $R = R^C$, then i has to choose between the uniform lottery over $H(\mathbf{b}) \cup \{c, c_j\}$ and c_j. Since i is tie-loving, she would choose the lottery, a contradiction with $\mathbf{b}$ being a PNE. Now, suppose that $R = R^V$. In the new ballot vector candidate c would be ranked first by $s - 1$ voters, and all other candidates would be ranked first by at most $s - 1$ voters, so $n \leq (s - 1)m$, and hence under R^V candidate c is selected with probability at least $1/m$. Thus, since i is strongly tie-loving, she would choose the lottery, a contradiction with $\mathbf{b}$ being a PNE. By the same argument, for every $c \in Y$ we have either $\mathrm{sc}(c, \mathbf{b}) = s - 1$ or $\mathrm{sc}(c, \mathbf{b}) \leq s - 3$. □

The proof of Lemma 3 uses the assumption that voters are (strongly) tie-loving. This lemma implies that $H(\mathbf{b}) \subseteq Y$, $H'(\mathbf{b}) \subseteq Z$. Now, let $C_1 = \{c_1, \dots, c_{j-1}\}$, $C_2 = \{c_{j+1}, \dots, c_m\}$, $Y_1 = Y \cap C_1$, $Y_2 = Y \cap C_2$, $Z_1 = Z \cap C_1$, $Z_2 = Z \cap C_2$, $H_1 = H(\mathbf{b}) \cap C_1$, $H_2 = H(\mathbf{b}) \cap C_2$, $H'_1 = H'(\mathbf{b}) \cap C_1$, $H'_2 = H'(\mathbf{b}) \cap C_2$.

For two candidate sets $A, B \subseteq C$ with $A \subseteq B$, we say that A forms a *prefix* of B if $\ell < r$, $c_\ell, c_r \in B$ and $c_r \in A$ implies $c_\ell \in A$; we say that A forms a *suffix* of B if $\ell < r$, $c_\ell, c_r \in B$ and $c_\ell \in A$ implies $c_r \in A$. The proof of the following lemma again uses the assumption that voters are (strongly) tie-loving.

LEMMA 4. *H_1 forms a prefix of Y_1, H_2 forms a suffix of Y_2, H'_1 forms a prefix of Z_1, and H'_2 forms a suffix of Z_2.*

PROOF. We provide the proof for H_1 and Y_1; other cases are similar (for H'_1 and Z_1 and for H'_2 and Z_2 we need the assumption that voters are (strongly) tie-loving). If $c_r \in H_1$, but $c_\ell \notin H_1$, there exists a voter i with $a_i = c_\ell$, $b_i = c_j$. Since i's preferences are single-peaked with respect to $\lhd$, she ranks c_ℓ first, and $\ell < r < j$, it follows that $c_r \succ_i c_j$; since $c_r \in H(\mathbf{b})$, voter i can profitably deviate to voting c_r, a contradiction. □

We are now ready to describe our algorithm. Given a profile $\mathbf{u}$, a candidate c_j, and a score s, we check that $\mathrm{sc}(c_j, \mathbf{a}) < s$, and reject if this is not the case. We then determine the sets Y_1, Y_2, Z_1, and Z_2. Then, for each $y_1 = 0, \dots, |Y_1|$, $y_2 = 0, \dots, |Y_2|$, $z_1 = 0, \dots, |Z_1|$, $z_2 = 0, \dots, |Z_2|$, we perform the following procedure.

If $y_1 + y_2 = 0$, we discard the 4-tuple (y_1, y_2, z_1, z_2). Otherwise, we let H_1 be the size-y_1 prefix of Y_1, let H_2 be the size-y_2 suffix of Y_2, let H'_1 be the size-z_1 prefix of Z_1, let H'_2 be the size-z_2 suffix of Z_2, and set $H = H_1 \cup H_2$, $H' = H'_1 \cup H'_2$. We then check that for each $i \in N$ we have $c_j \succ_i c_\ell$ for every $c_\ell \in H \setminus \{a_i\}$, and reject if this is not the case. Next, for each $c_r \in X \cup (Y \setminus H) \cup (Z \setminus H')$ we try to select $\mathrm{sc}(c_r, \mathbf{a}) - s + 3$ voters who rank c_r first and, for every candidate $c_\ell \in H'$, prefer c_j to the lottery where a candidate is chosen from $H \cup \{c_j, c_\ell\}$ according to R. We reject if for some c_r there are not enough voters whose preferences satisfy this requirement. Otherwise, let s' be the number of voters picked so far. If $s' > s - \mathrm{sc}(c_j, \mathbf{a})$, we reject, and if $s' < s - \mathrm{sc}(c_j, \mathbf{a})$, we try to pick additional $s - \mathrm{sc}(c_j, \mathbf{a}) - s'$ voters who rank some candidate in $C \setminus (H \cup H' \cup \{c_j\})$ first and, for every candidate $c_\ell \in H'$, prefer c_j to the lottery where a candidate is chosen from $H \cup \{c_j, c_\ell\}$ according to R. If we succeed, we terminate and report success, and otherwise we proceed to the next 4-tuple (y_1, y_2, z_1, z_2).

There are at most m^4 tuples to consider, and for each of them our procedure can be performed in polynomial time. Thus, the running time of our algorithm is polynomial.

It is not hard to see that if we report success, then there exists a non-truthful PNE where c_j wins with s points. Indeed, consider the 4-tuple (y_1, y_2, z_1, z_2) on which the algorithm terminates and the sets H and H' constructed for these values of y_1, y_2, z_1, and z_2, and let S be the set of voters picked in the respective iteration of our algorithm. Set $b_i = c_j$ for $i \in S$, $b_i = a_i$ for $i \notin S$. By construction we have $H(\mathbf{b}) = H \neq \varnothing$, $H'(\mathbf{b}) = H'$, $\mathrm{sc}(c_j, \mathbf{b}) = s$, and $b_i \neq a_i$ for some $i \in N$. Further, our procedure ensures that all conditions of Theorem 5 are satisfied.

Conversely, suppose that there exists a non-truthful PNE $\mathbf{b}$ where c_j wins with s points. Let $H_1 = H(\mathbf{b}) \cap C_1$, $H_2 = H(\mathbf{b}) \cap C_2$, $H'_1 = H'(\mathbf{b}) \cap C_1$, $H'_2 = H'(\mathbf{b}) \cap C_2$, and set $y_1 = |H_1|$, $y_2 = |H_2|$, $z_1 = |H'_1|$, $z_2 = |H'_2|$. Lemmas 3 and 4 imply that, when considering the 4-tuple (y_1, y_2, z_1, z_2), our procedure would report success.

For single-crossing preferences, we can prove an analogue of Lemma 4 for the ordering of the candidates in $A = \{a_i \mid i \in N\}$ induced by the order of the voters, which enables us to use essentially the same algorithm as above; we omit the details. □

It remains to consider truth-biased voters and lexicographic tie-breaking. For general preferences, this setting has been studied in detail by Obraztsova et al. [26], and subsequently Elkind et al. [15] have shown that the computational problems $(\mathcal{T}, R^L)$-EXISTNE, $(\mathcal{T}, R^L)$-TIENE and $(\mathcal{T}, R^L)$-SINGLENE are NP-complete. Our results for this model differ from the results in the rest of the paper: on the one hand, for single-peaked preferences we obtain a polynomial-time algorithm even *without* assuming that voters are tie-consistent, and on the other hand, we have not been able to obtain efficient algorithms for our problems for single-crossing preferences.

THEOREM 7. *$(\mathcal{T}, R^L)$-EXISTNE, $(\mathcal{T}, R^L)$-TIENE and $(\mathcal{T}, R^L)$-SINGLENE admit polynomial-time algorithms if the input profile is single-peaked.*

PROOF SKETCH. We only give a brief sketch of the algorithm for the TIENE problem. Suppose the target winner is candidate c_ℓ. Recall that we want to check if there exists a PNE $\mathbf{b}$ such that $c_\ell \in W(\mathbf{b})$ and $W(\mathbf{b}) > 1$. Starting from the truthful preference profile, the algorithm will try to construct the desirable PNE $\mathbf{b}$ if one exists.

First, we can assume that we seek a PNE where the winner has at least 4 points (we can check in polynomial time if there exists a PNE $\mathbf{b}$ where the score of c_ℓ is at most 3).

The algorithm will iterate through all possible values for the winning score of c_ℓ. Let s be the score we examine in a fixed iteration ($s \geq 4$). The algorithm performs the following steps in each iteration:

Step 1: Building the set of potential threshold candidates. We construct the set S, which consists of all potential threshold candidates with score s (the possible candidates who could tie with c_ℓ at equilibrium). This is essentially the s-eligible threshold set, as defined by Obraztsova et al. [26]. We then examine each candidate from S and check if any of them can be made to be an actual threshold candidate at the equilibrium we are looking for. We start by considering the lowest-indexed candidate. In each subsequent iteration we examine the candidate from S who has the lowest index among the ones we have not examined yet. Let t be the candidate we examine in a given iteration over S. We call t the *main threshold candidate*. Note that the tie-breaking rule should favor c_ℓ over t, since they both have the same score.

Step 2a: Reducing the support of candidates with high enough score. Given the score s and the main threshold candidate t, we will now start blocking other candidates from being winners at the PNE $\mathbf{b}$ that we are trying to construct and also from threatening the stability of $\mathbf{b}$. Given a candidate c, we say that we "block" c when we modify the votes of her supporters one by one so as to vote for c_ℓ , always starting from the votes where c_ℓ is ranked higher in the switched vote than in any other vote with the same top choice (ties are broken arbitrarily). This blocking procedure continues until c has $s-2$ points. We apply this procedure to each candidate who cannot become the winner with a unilateral deviation if he reaches $s-2$ points. We denote by T the set of all s-eligible threshold candidates who are either beaten by t in tie-breaking or have $s-1$ points. We apply the blocking procedure as follows:

- Block all candidates who have more than s points in the truthful ballot vector.
- Block all candidates with exactly s points that are not included in $T \cup \{t\}$.
- Block all candidates who have exactly $s-1$ points, beat c_ℓ in tie-breaking, and are not included in $T \cup \{t\}$.

If $\mathbf{b}'$ is the profile that has been obtained after blocking candidates as above, we refer to a vote as a *modified* vote in $\mathbf{b}'$ if we have already changed it from its initial value at this point in the algorithm.

Step 2b: remaining candidates with $s-1$ points. Given $\mathbf{b}'$, we denote by $D(\mathbf{b}')$ the set of all candidates c in the profile $\mathbf{b}'$ such that

- $\mathrm{sc}(c, \mathbf{a}) = s-1$;
- the tie-breaking rule favors c_ℓ over c and c over t;
- c appears above c_ℓ in at least one modified vote;
- c is not blocked in $\mathbf{b}'$.

We will now block the candidates from $D(\mathbf{b}')$, updating this set at every step until it becomes empty. If after this step the score of c_ℓ is larger than s, the algorithm reports that there is no PNE with t being the threshold candidate. If the score is exactly s, the algorithm returns "yes" and outputs the current ballot vector.

Step 3: gaining additional support for c_ℓ. The challenging case is when the current score of c_ℓ is less than s.

	R^L	R^C	R^V
$\mathcal{L}$	any [15]	SP/SC (tc)	SP/SC (tc)
$\mathcal{T}$	SP	SP/SC (tl)	SP/SC (stl)

Table 1: Summary of our results: the cells of the table indicate sufficient conditions for the existence of polynomial-time algorithms for ExistNE. SP stands for 'single-peaked', 'SC' stands for 'single-crossing', tc stands for 'tie-consistent', tl stands for 'tie-loving', stl stands for 'strongly tie-loving'.

In this case, we need to further modify the votes in $\mathbf{b}'$ in favor of c_ℓ so as to make her a winner with score s. This is a more subtle task, for which we need to use the fact that preferences are single-peaked.

Towards this, we define a set of *safe* votes, which are the pool of votes that we will be allowed to change in favor of c_ℓ in the attempt to construct the desired equilibrium $\mathbf{b}$. Let B be the set of candidates that have been blocked so far. We say that a vote is *safe* if its top-ranked candidate is neither c_ℓ nor t and the intersection of the set of candidates above c_ℓ in the vote and $D(\mathbf{a}) \setminus B$ is empty; we denote the set of safe votes by $S(\mathbf{b}')$.

The rest of the algorithm proceeds by trying to find a large enough set of safe votes to change in favor of c_ℓ. This is done by identifying an appropriate pair of candidates in the axis of the single-peaked preferences, one to the left of c_ℓ and one to the right of c_ℓ, say, c_L and c_R, respectively. We then block all the candidates within the interval between c_L and c_R who have not already been blocked. Doing so provides us with a set of safe votes that we can transform in favor of c_ℓ, if this is a "yes"-instance. We omit further details of identifying this appropriate interval from this version. □

5. CONCLUSIONS AND FUTURE WORK

We have shown that for single-peaked/single-crossing preferences PNE of Plurality voting games can be found in polynomial time for almost all scenarios considered by Elkind et al. [15], under the assumption of tie-consistency and its variants; our results are summarized in Table 1. The most obvious open problems suggested by our work are relaxing the assumption of tie-consistency, and extending our results to preferences that are *almost* single-peaked/single-crossing, to capture a broader and more realistic range of social choice scenarios.

From the point of view of social choice theory, single-peaked and single-crossing preferences are considered to be 'good' as they ensure that the majority relation is transitive. Our work, together with the recent work on manipulation, control, bribery and committee selection rules (see references in Section 1) provides a new perspective on these preference domains, by demonstrating that they are also attractive for purely algorithmic reasons (which appear to be unrelated to transitivity of majority preferences).

Acknowledgements. Edith Elkind and Piotr Skowron were supported by ERC-StG 639945. Svetlana Obraztsova was supported by Israeli Science Foundation (ISF) via I-CORE postdoctoral scholarship. This work was partially supported by COST Action IC1205 on Computational Social Choice.

REFERENCES

[1] K. J. Arrow. A difficulty in the concept of social welfare. *Journal of Political Economy*, 58(4):328–346, 1950.

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

[3] M. Battaglini. Sequential voting with abstention. *Games and Economic Behavior*, 51:445–463, 2005.

[4] N. Betzler, A. Slinko, and J. Uhlmann. On the computation of fully proportional representation. *Journal of AI Research*, 47:475–519, 2013.

[5] D. Black. *The Theory of Committees and Elections*. Cambridge University Press, 1958.

[6] T. Borgers. Costly voting. *American Economic Review*, 94(1):57–66, 2004.

[7] F. Brandt, M. Brill, E. Hemaspaandra, and L. Hemaspaandra. Bypassing combinatorial protections: Polynomial-time algorithms for single-peaked electorates. *Journal of Artificial Intelligence Research*, 53:439–496, 2015.

[8] S. Brânzei, I. Caragiannis, J. Morgenstern, and A. D. Procaccia. How bad is selfish voting? In *AAAI'13*, pages 138–144, 2013.

[9] R. Bredereck, J. Chen, and G. Woeginger. A characterization of the single-crossing domain. *Social Choice and Welfare*, 41(4):989–998, 2013.

[10] E. Dekel and M. Piccione. Sequential voting procedures in symmetric binary elections. *Journal of Political Economy*, 108(1):34–55, 2000.

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

[12] J. Doignon and J. Falmagne. A polynomial time algorithm for unidimensional unfolding representations. *Journal of Algorithms*, 16(2):218–233, 1994.

[13] B. Dutta and A. Sen. Nash implementation with partially honest individuals. *Games and Economic Behavior*, 74(1):154–169, 2012.

[14] E. Elkind, P. Faliszewski, and A. Slinko. Clone structures in voters' preferences. In *ACM EC'12*, pages 496–513, 2012.

[15] E. Elkind, E. Markakis, S. Obraztsova, and P. Skowron. Equilibria of plurality voting: Lazy and truth-biased voters. In *SAGT'15*, 2015.

[16] B. Escoffier, J. Lang, and M. Öztürk. Single-peaked consistency and its complexity. In *ECAI'08*, pages 366–370, 2008.

[17] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. The complexity of manipulative attacks in nearly single-peaked electorates. *AI Journal*, 207:69–99, 2014.

[18] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. The shield that never was: Societies with single-peaked preferences are more open to manipulation and control. *Information and Computation*, 209(2):89–107, 2011.

[19] P. Faliszewski and A. Procaccia. AI's war on manipulation: Are we winning? *AI Magazine*, 31(4):52–64, 2010.

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

[21] A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41(4):587–601, 1973.

[22] M. Lombardi and N. Yoshihara. A full characterization of Nash implementation with strategy space reduction. *Economic Theory*, 54(1):131–151, 2013.

[23] N. Mattei and T. Walsh. PrefLib: A library of preference data. In *ADT'13*, pages 259–270, 2013.

[24] R. Meir, M. Polukarov, J. S. Rosenschein, and N. R. Jennings. Convergence to equilibria of plurality voting. In *AAAI'10*, pages 823–828, 2010.

[25] J. A. Mirrlees. An Exploration in the Theory of Optimum Income Taxation. *Review of Economic Studies*, 38(114):175–208, April 1971.

[26] S. Obraztsova, E. Markakis, and D. R. M. Thompson. Plurality voting with truth-biased agents. In *SAGT'13*, pages 26–37, 2013.

[27] M. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, 1975.

[28] M. Schulze. A new monotonic, clone-independent, reversal symmetric, and Condorcet-consistent single-winner election method. *Social Choice and Welfare*, 36(2):267–303, 2011.

[29] F. D. Sinopoli and G. Iannantuoni. On the generic strategic stability of Nash equilibria if voting is costly. *Economic Theory*, 25(2):477–486, 2005.

[30] P. Skowron, L. Yu, P. Faliszewski, and E. Elkind. The complexity of fully proportional representation for single-crossing electorates. *Theoretical Computer Science*, 569:43–57, 2015.

[31] B. Sloth. The theory of voting and equilibria in noncooperative games. *Games and Economic Behavior*, 5(1):152–169, 1993.

[32] D. R. M. Thompson, O. Lev, K. Leyton-Brown, and J. S. Rosenschein. Empirical analysis of plurality election equilibria. In *AAMAS'13*, pages 391–398, 2013.