Equilibria of Plurality Voting: Lazy and Truth-Biased Voters

Edith Elkind[1], Evangelos Markakis[2], Svetlana Obraztsova[3],
and Piotr Skowron[4(✉)]

[1] University of Oxford, Oxford, UK
[2] Athens University of Economics and Business, Athens, Greece
[3] Tel Aviv University, Tel Aviv, Israel
[4] University of Warsaw, Warsaw, Poland
p.skowron@mimuw.edu.pl

Abstract. We present a systematic study of Plurality elections with strategic voters who, in addition to having preferences over election winners, also have secondary preferences, governing their behavior when their vote cannot affect the election outcome. Specifically, we study two models that have been recently considered in the literature: *lazy* voters, who prefer to abstain when they are not pivotal, and *truth-biased* voters, who prefer to vote truthfully when they are not pivotal. For both lazy and truth-biased voters, we are interested in their behavior under different tie-breaking rules (lexicographic rule, random voter rule, random candidate rule). Two of these six combinations of secondary preferences and tie-breaking rules have been studied in prior work; for the remaining four, we characterize pure Nash equilibria (PNE) of the resulting strategic games and study the complexity of related computational problems. We then use these results to analyze the impact of different secondary preferences and tie-breaking rules on the election outcomes. Our results extend to settings where some of the voters are non-strategic.

1 Introduction

Plurality voting is a popular tool for collective decision-making in many domains, including both human societies and multiagent systems. Under this voting rule, each voter is supposed to vote for her most favorite candidate (or abstain); the winner is then the candidate that receives the highest number of votes. If several candidates have the highest score, the winner is chosen among them using a *tie-breaking rule*; popular tie-breaking rules include the *lexicographic rule*, which imposes a fixed priority order over the candidates; the *random candidate rule*, which picks one of the tied candidates uniformly at random; and the *random voter rule*, which picks the winner among the tied candidates according to the preferences of a randomly chosen voter.

In practice, voters are often *strategic*, i.e., they may vote non-truthfully if they can benefit from doing so. In that case, an election can be viewed as a game, where the voters are the players, and each player's space of actions includes voting for any candidate or abstaining. For deterministic rules (such as Plurality

with lexicographic tie-breaking), the behavior of strategic voters is determined by their preference ordering, i.e., a ranking of the candidates, whereas for randomized rules a common approach is to specify utility functions for the voters; i.e., the voters are assumed to maximize their *expected utility* under the lottery induced by tie-breaking. The outcome of the election can then be identified with a pure Nash equilibrium (PNE) of the resulting game.

However, under Plurality and with 3 or more voters, this approach fails to provide a useful prediction of voting behavior: for each candidate c there is a PNE where c is the unique winner, irrespective of the voters' preferences. Indeed, if there are at least 3 voters, the situation where all of them vote for c is a PNE, as no voter can change the election outcome. Such equilibria may disappear if we use a more refined model of voters' preferences that captures additional aspects of their decision-making. For instance, in practice, if a voter feels that her vote is unlikely to have any effect on the outcome, she may decide to abstain from the election. Also, voters may be averse to lying about their preferences, in which case they can be expected to vote for their top candidate unless there is a clear strategic reason to vote for someone else. By taking into account these aspects of voters' preferences, we can obtain a more faithful model of their behavior.

The problem of characterizing and computing the equilibria of Plurality voting, both for "lazy" voters (i.e., ones who prefer to abstain when they are not pivotal) and for "truth-biased" voters (ones who prefer to vote truthfully when they are not pivotal), has recently received a considerable amount of attention. However, it is difficult to compare the existing results, since they rely on different tie-breaking rules. In particular, Desmedt and Elkind [6], who study lazy voters, use the random candidate tie-breaking rule, and Obraztsova et al. [17] consider truth-biased voters and the lexicographic tie-breaking rule. Thus, it is not clear whether the differences between the results in these papers can be attributed to the voters' secondary preferences, or to the tie-breaking rule.

The primary goal of our paper is to tease out the effects of different features of these models, by systematically considering all the combinations of secondary preferences and tie-breaking rules. We consider two types of secondary preferences (lazy voters and truth-biased voters) and three tie-breaking rules (the lexicographic rule, the random voter rule, and the random candidate rule); while two of these combinations have been studied earlier by [6,17], to the best of our knowledge, the remaining four possibilities have not been considered before. For each of the new scenarios, we characterize the set of PNE for the resulting game; in doing so, we also fill in a gap in the characterization of [6] for lazy voters and random candidate tie-breaking. We then consider the problems of deciding whether a given game admits a PNE and whether a given candidate can be a co-winner/unique winner in some PNE of a given game. For all settings, we determine the computational complexity of each of these problems, classifying them as either polynomial-time solvable or NP-complete. Our characterization results enable us to analyze the impact of various features of our model on the election outcomes, and thereby evaluate the plausibility of our assumptions about voters' secondary preferences. Finally, we briefly discuss the implications of our results in the setting where some of the voters may be *principled*, i.e., always vote truthfully.

Related Work. Equilibria of Plurality voting have been investigated by a number of researchers, starting with [10]. However, most of the earlier works either consider solution concepts other than pure Nash equilibria, such as iterative elimination of dominated strategies [7,13], or assume that voters have incomplete information about each others' preferences [14]. Both types of secondary preferences (lazy voters and truth-biased voters) appear in the social choice literature, see, respectively, [2,3,19] and [8,11]. In computational social choice, truth-biased voters have been considered by Meir et al. [12] in the context of dynamics of Plurality voting; subsequently, Plurality elections with truth-biased voters have been investigated empirically by Thompson et al. [20] and theoretically by Obraztsova et al. [17]. To the best of our knowledge, the only paper to study computational aspects of Plurality voting with lazy voters is that of Desmedt and Elkind [6].

Our approach to tie-breaking is well-grounded in existing work. Lexicographic tie-breaking is standard in the computational social choice literature. The random candidate rule has been discussed by [6], and, more recently, by [15,16]. The random voter rule is used to break ties under the Schulze method [18]; the complexity of manipulation under this tie-breaking rule has been studied by [1].

2   Preliminaries

For any positive integer t, we denote the set $\{1, \ldots, t\}$ by $[t]$. We consider elections with a set of *voters* $N = [n]$ and a set of *alternatives*, or *candidates*, $C = \{c_1, \ldots c_m\}$. Each voter is associated with a *preference order*, i.e., a strict linear order over C; we denote the preference order of voter i by $\succ_i$. The list $(\succ_1, \ldots, \succ_n)$ is called a *preference profile*. For each $i \in N$, we set a_i to be the top choice of voter i, and let $\mathbf{a} = (a_1, \ldots, a_n)$. Given two disjoint sets of candidates X, Y and a preference order $\succ$, we write $X \succ Y$ if in $\succ$ all candidates from X are ranked above all candidates from Y.

We also assume that each voter $i \in N$ is endowed with a *utility function* $u_i : C \to \mathbb{N}$; $u_i(c_j)$ is the utility derived by voter i if c_j is the unique election winner. We require that $u_i(c) \neq u_i(c')$ for all $i \in N$ and all $c, c' \in C$ such that $c \neq c'$. The vector $\mathbf{u} = (u_1, \ldots, u_n)$ is called the *utility profile*. Voters' preference orders and utility functions are assumed to be consistent, i.e., for each $i \in N$ and every pair of candidates $c, c' \in C$ we have $c \succ_i c'$ if and only if $u_i(c) > u_i(c')$; when this is the case, we will also say that $\succ_i$ is *induced* by u_i. Sometimes, instead of specifying preference orders explicitly, we will specify the utility functions only, and assume that voters' preference orders are induced by their utility functions; on other occasions, it will be convenient to reason in terms of preference orders.

A *lottery* over C is a vector $\mathbf{p} = (p_1, \ldots, p_m)$ with $p_j \geq 0$ for all $j \in [m]$ and $\sum_{j \in [m]} p_j = 1$. The value p_j is the probability assigned to candidate c_j. The *expected utility* of a voter $i \in N$ from a lottery $\mathbf{p}$ is given by $\sum_{j \in [m]} u_i(c_j) p_j$.

In this work, we consider Plurality elections, where each voter $i \in N$ submits a *vote*, or *ballot*, $b_i \in C \cup \{\varnothing\}$; if $b_i = \varnothing$, voter i is said to *abstain*. The list

of all votes $\mathbf{b} = (b_1, \ldots, b_n)$ is also called a *ballot vector*. We say that a ballot vector is *trivial* if $b_i = \varnothing$ for all $i \in N$. Given a ballot vector $\mathbf{b}$ and a ballot b', we write $(\mathbf{b}_{-i}, b')$ to denote the ballot vector obtained from $\mathbf{b}$ by replacing b_i with b'. The *score* of an alternative c_j in an election with ballot vector $\mathbf{b}$ is given by $\mathrm{sc}(c_j, \mathbf{b}) = |\{i \in N \mid b_i = c_j\}|$. Given a ballot vector $\mathbf{b}$, we set $M(\mathbf{b}) = \max_{c \in C} \mathrm{sc}(c, \mathbf{b})$ and let $W(\mathbf{b}) = \{c \in C \mid \mathrm{sc}(c, \mathbf{b}) = M(\mathbf{b})\}$, $H(\mathbf{b}) = \{c \in C \mid \mathrm{sc}(c, \mathbf{b}) = M(\mathbf{b}) - 1\}$, $H'(\mathbf{b}) = \{c \in C \mid \mathrm{sc}(c, \mathbf{b}) = M(\mathbf{b}) - 2\}$. These sets are useful in our analysis in the next sections. The set $W(\mathbf{b})$ is called the *winning set*. Note that if $\mathbf{b}$ is trivial then $W(\mathbf{b}) = C$. If $|W(\mathbf{b})| > 1$, the winner is selected from $W(\mathbf{b})$ according to one of the following tie-breaking rules.

(1) Under the *lexicographic rule* R^L, the winner is the candidate $c_j \in W(\mathbf{b})$ such that $j \leq k$ for all $c_k \in W(\mathbf{b})$.
(2) Under the *random candidate rule* R^C, the winner is chosen from $W(\mathbf{b})$ uniformly at random.
(3) Under the *random voter rule* R^V, we select a voter from N uniformly at random; if she has voted for a candidate in $W(\mathbf{b})$, we output this candidate, otherwise we ask this voter to report her most preferred candidate in $W(\mathbf{b})$, and output the answer. This additional elicitation step may appear difficult to implement in practice; fortunately, we can show that in equilibrium it is almost never necessary.

Thus, the outcome of an election is a lottery over C; however, for R^L this lottery is degenerate, i.e., it always assigns the entire probability mass to a single candidate. For each $X \in \{L, C, V\}$ and each ballot vector $\mathbf{b}$, let $\mathbf{p}^X(\mathbf{b})$ denote the lottery that corresponds to applying R^X to the set $W(\mathbf{b})$. From the definition of R^C, it follows that for every $c_j \in C$ it holds that if $p_j^C(\mathbf{b}) \neq 0$ then $p_j^C(\mathbf{b}) \geq \frac{1}{m}$. Similarly, for R^V, it follows that if $p_j^V(\mathbf{b}) \neq 0$ then $p_j^V(\mathbf{b}) \geq \frac{1}{n}$.

In what follows, we focus on two types of secondary preferences, namely, *lazy* voters, who prefer to abstain when their vote has no effect on the election outcome, and *truth-biased* voters, who never abstain, but prefer to vote truthfully when their vote has no effect on the election outcome. Formally, pick $\varepsilon < \min\{\frac{1}{m^2}, \frac{1}{n^2}\}$, and consider a utility profile $\mathbf{u}$ and a tie-breaking rule $R^X \in \{R^C, R^V, R^L\}$. Then

– if voter i is *lazy*, her utility in an election with ballot vector $\mathbf{b}$ under tie-breaking rule R^X is given by

$$U_i(\mathbf{b}) = \begin{cases} \sum_{j \in [m]} p_j^X(\mathbf{b}) u_i(c_j), & \text{if } b_i \in C, \\ \sum_{j \in [m]} p_j^X(\mathbf{b}) u_i(c_j) + \varepsilon, & \text{if } b_i = \varnothing. \end{cases}$$

– if voter i is *truth-biased*, her utility in an election with ballot vector $\mathbf{b}$ under tie-breaking rule R^X is given by

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

We consider settings where all voters are of the same type, i.e., either all voters are lazy or all voters are truth-biased; we refer to these settings as *lazy* or *truth-biased*, respectively, and denote the former by $\mathcal{L}$ and the latter by $\mathcal{T}$.

We investigate all possible combinations of settings $(\mathcal{L}, \mathcal{T})$ and tie-breaking rules (R^L, R^C, R^V). A combination of a setting $\mathcal{S} \in \{\mathcal{L}, \mathcal{T}\}$, a tie-breaking rule $R \in \{R^L, R^C, R^V\}$ and a utility profile $\mathbf{u}$ induces a strategic game, which we will denote by $(\mathcal{S}, R, \mathbf{u})$: in this game, the players are the voters, the action space of each player is $C \cup \{\varnothing\}$, and the players' utilities $U_1, \ldots, U_n$ for a vector of actions $\mathbf{b}$ are computed based on the setting and the tie-breaking rule as described above. We say that a ballot vector $\mathbf{b}$ is a *pure Nash equilibrium (PNE)* of the game $(\mathcal{S}, R, \mathbf{u})$ if $U_i(\mathbf{b}) \geq U_i(\mathbf{b}_{-i}, b')$ for every voter $i \in N$ and every $b' \in C \cup \{\varnothing\}$.

For each setting $\mathcal{S} \in \{\mathcal{L}, \mathcal{T}\}$ and each tie-breaking rule $R \in \{R^L, R^C, R^V\}$, we define three algorithmic problems, which we call $(\mathcal{S}, R)$-ExistNE, $(\mathcal{S}, R)$-TieNE, and $(\mathcal{S}, R)$-SingleNE. In each of these problems, we are given a candidate set C, $|C| = m$, a voter set N, $|N| = n$, and a utility vector $\mathbf{u} = (u_1, \ldots, u_n)$, where each u_i is represented by m numbers $u_i(c_1), \ldots, u_i(c_m)$; these numbers are positive integers given in binary. In $(\mathcal{S}, R)$-TieNE and $(\mathcal{S}, R)$-SingleNE we are also given the name of a target candidate $c_p \in C$. In $(\mathcal{S}, R)$-ExistNE we ask if $(\mathcal{S}, R, \mathbf{u})$ has a PNE. In $(\mathcal{S}, R)$-TieNE we ask if $(\mathcal{S}, R, \mathbf{u})$ has a PNE $\mathbf{b}$ with $|W(\mathbf{b})| > 1$ and $c_p \in W(\mathbf{b})$. In $(\mathcal{S}, R)$-SingleNE we ask if $(\mathcal{S}, R, \mathbf{u})$ has a PNE $\mathbf{b}$ with $W(\mathbf{b}) = \{c_p\}$. Each of these problems is obviously in NP, as we can simply guess an appropriate ballot vector $\mathbf{b}$ and check that it is a PNE.

In what follows, we omit some proofs due to space constraints; the omitted proofs can be found in the full version of the paper [9].

3   Lazy Voters

In this section, we study PNE in Plurality games with lazy voters. The case where the tie-breaking rule is R^C has been analyzed in detail by Desmedt and Elkind [6], albeit for a slightly different model; we complement their results by considering R^L and R^V.

We start by extending a result of [6] to all three tie-breaking rules considered here.

Proposition 1. *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 $(\mathcal{L}, R, \mathbf{u})$ then for every voter $i \in N$ either $b_i = \varnothing$ or $b_i \in W(\mathbf{b})$. If $|W(\mathbf{b})| = 1$, there is exactly one voter $i \in N$ with $b_i \neq \varnothing$.*

Proof. Suppose that $b_i \neq \varnothing$, $b_i \notin W(\mathbf{b})$ for some voter $i \in N$. Then if i changes her vote to $\varnothing$, the set $W(\mathbf{b})$ will not change, so i's utility would improve by ε, a contradiction with $\mathbf{b}$ being a PNE of $(\mathcal{L}, R, \mathbf{u})$. Similarly, suppose that $|W(\mathbf{b})| = 1$ and there are two voters $i, i' \in N$ with $b_i \neq \varnothing$, $b_{i'} \neq \varnothing$. It has to be the case that $b_i = b_{i'} = c_j$ for some $c_j \in C$, since otherwise $|W(\mathbf{b})| > 1$. But then if voter i changes her vote to $\varnothing$, c_j will remain the election winner, so i's utility would improve by ε, a contradiction.

Lexicographic Tie-breaking. The scenario where voters are lazy and ties are broken lexicographically turns out to be fairly easy to analyze.

Theorem 1. *For any utility profile $\mathbf{u}$ the game $G = (\mathcal{L}, R^L, \mathbf{u})$ has the following properties:*

(1) *If $\mathbf{b}$ is a PNE of G then $|W(\mathbf{b})| \in \{1, m\}$. Moreover, $|W(\mathbf{b})| = m$ if and only if $\mathbf{b}$ is the trivial ballot and all voters rank c_1 first.*
(2) *If $\mathbf{b}$ is a PNE of G then there exists at most one voter i with $b_i \neq \varnothing$.*
(3) *G admits a PNE if and only if all voters rank c_1 first (in which case c_1 is the unique PNE winner) or there exists a candidate c_j with $j > 1$ such that (i) $\mathrm{sc}(c_j, \mathbf{a}) > 0$ and (ii) for every $k < j$ it holds that all voters prefer c_j to c_k. If such a candidate exists, he is unique, and wins in all PNE of G.*

The following corollary is directly implied by Theorem 1.

Corollary 1. *$(\mathcal{L}, R^L)$-EXISTNE, $(\mathcal{L}, R^L)$-SINGLENE and $(\mathcal{L}, R^L)$-TIENE are in P.*

Remark 1. The reader may observe that, counterintuitively, while the lexicographic tie-breaking rule appears to favor c_1, it is impossible for c_1 to win the election unless he is ranked first by all voters. In contrast, c_2 wins the election as long as he is ranked first by at least one voter and no voter prefers c_1 to c_2. In general, the lexicographic tie-breaking rule favors lower-numbered candidates with the exception of c_1. As for c_1, his presence mostly has a destabilizing effect: if some, but not all voters rank c_1 first, no PNE exists. This phenomenon is an artifact of our treatment of the trivial ballot vector: it disappears if we assume (as [6] does) that when $\mathbf{b} = (\varnothing, \ldots, \varnothing)$ the election is declared invalid and the utility of each voter is $-\infty$: under this assumption c_1 is the unique possible equilibrium winner whenever he is ranked first by at least one voter.

Randomized Tie-breaking. We now consider R^C and R^V. Desmedt and Elkind [6] give a characterization of utility profiles that admit a PNE for lazy voters and R^C. However, there is a small difference between our model and theirs regarding the trivial ballot vector, as explained in Remark 1 above. Further, their results implicitly assume that the number of voters n exceeds the number of candidates m; if this is not the case, Theorem 2 in their paper is incorrect (see Remark 2).

Thus, we will now provide a full characterization of utility profiles $\mathbf{u}$ such that $(\mathcal{L}, R^C, \mathbf{u})$ admits a PNE, and describe the corresponding equilibrium ballot profiles. Our characterization result remains essentially unchanged if we replace R^C with R^V: for almost all utility profiles $\mathbf{u}$ and ballot vectors $\mathbf{b}$ it holds that $\mathbf{b}$ is a PNE of $(\mathcal{L}, R^C, \mathbf{u})$ if and only if it is a PNE of $(\mathcal{L}, R^V, \mathbf{u})$; the only exception is the case of full consensus (all voters rank the same candidate first).

Theorem 2. *Let $\mathbf{u} = (u_1, \ldots, u_n)$ be a utility profile over C, $|C| = m$, and let $R \in \{R^C, R^V\}$. 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 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 set of candidates $X = \{c_{\ell_1}, \ldots, c_{\ell_k}\}$ with $2 \leq k \leq \min(n/2, m)$
    and a partition of the voters into k groups $N_1, \ldots, N_k$ of size n/k each such
    that for each $j \in [k]$ and each $i \in N_j$ we have $c_{\ell_j} \succ_i c$ for all $c \in X \setminus \{c_{\ell_j}\}$,
    and, moreover, $\frac{1}{k}\sum_{c\in X} u_i(c) \geq \max_{c\in X\setminus\{c_{\ell_j}\}} u_i(c)$.

Further, when condition (1) holds for some $c_j \in C$ and $R = R^C$, then for each
$i \in N$ the game G has a PNE where i votes for c_j and all other voters abstain,
whereas if $R = R^V$, the game G has a PNE where all voters abstain; if condition
(2) holds, then G has a PNE where each voter votes for her top candidate; and
if condition (3) 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.

Remark 2. Desmedt and Elkind [6] claim (Theorems 1 and 2) that for R^C and
lazy voters, a PNE exists if and only if the utility profile satisfies either condition
(1) or (3) with the constraint $k \leq n/2$ removed. To see why this is incorrect,
consider a 2-voter election over $C = \{x, y, z\}$, where the voters' utility functions
are consistent with preference orders $x \succ y \succ z$ and $x \succ z \succ y$, respectively.
According to [6], the vector (y, z) is a PNE. This is obviously not true: each of
the voters would prefer to change her vote to x. Note, however, that the two
characterizations differ only when $m \geq n$, and in practice the number of voters
usually exceeds the number of candidates.

   Desmedt and Elkind [6] show that checking condition (3) of Theorem 2 is
NP-hard; in their proof $n > m$, and the proof does not depend on how the
trivial ballot is handled. Further, their proof shows that checking whether a
given candidate belongs to some such set X is also NP-hard. On the other
hand, Theorem 2 shows that PNE with singleton winning sets only arise if some
candidate is unanimously ranked first, and this condition is easy to check. We
summarize these observations as follows.

Corollary 2. *For $R \in \{R^C, R^V\}$, the problems $(\mathcal{L}, R)$-EXISTNE and $(\mathcal{L}, R)$-
TIENE are NP-complete, whereas $(\mathcal{L}, R)$-SINGLENE is in P.*

4   Truth-Biased Voters

For truth-biased voters, our exposition follows the same pattern as for lazy vot-
ers: we present some general observations, followed by a quick summary of the
results for lexicographic tie-breaking, and continue by analyzing randomized tie-
breaking. The following result is similar in spirit to Proposition 1.

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 $i \in N$ we have $b_i = a_i$
or $b_i \in W(\mathbf{b})$.*

Lexicographic Tie-breaking. Obraztsova et al. [17] characterize the PNE of the game $(\mathcal{T}, R^L, \mathbf{u})$. Their characterization is quite complex, and we will not reproduce it here. However, for the purposes of comparison with the lazy voters model, we will use the following description of *truthful* equilibria.

Proposition 3 (Obraztsova et al. [17], Theorem 1). *Consider a utility profile $\mathbf{u}$, let $\mathbf{a}$ be the respective truthful ballot vector, and let $j = \min\{r \mid c_r \in W(\mathbf{a})\}$. Then $\mathbf{a}$ is a PNE of $(\mathcal{T}, R^L, \mathbf{u})$ if and only if neither of the following conditions holds:*

(1) $|W(\mathbf{a})| > 1$, and there exists a candidate $c_k \in W(\mathbf{a})$ and a voter i such that $a_i \neq c_k$ and $c_k \succ_i c_j$.
(2) $H(\mathbf{a}) \neq \emptyset$, and there exists a candidate $c_k \in H(\mathbf{a})$ and a voter i such that $a_i \neq c_k$, $c_k \succ_i c_j$, and $k < j$.

We will also utilize a crucial property of non-truthful PNE. For this, we first need the following definition.

Definition 1. Consider a ballot vector $\mathbf{b}$, where candidate c_j is the winner under R^L. A candidate $c_k \neq c_j$ is called a *threshold candidate with respect to* $\mathbf{b}$ if either (1) $k < j$ and $\mathrm{sc}(c_k, \mathbf{b}) = \mathrm{sc}(c_j, \mathbf{b}) - 1$ or (2) $k > j$ and $\mathrm{sc}(c_k, \mathbf{b}) = \mathrm{sc}(c_j, \mathbf{b})$. We denote the set of threshold candidates with respect to $\mathbf{b}$ by $T(\mathbf{b})$.

That is, a threshold candidate is someone who could win the election if he had one additional vote. A feature of all non-truthful PNE is that there must exist at least one threshold candidate. The intuition for this is that, since voters who are not pivotal prefer to vote truthfully, in any PNE that arises under strategic voting, the winner receives just enough votes so as to beat the required threshold (as set by the threshold candidate) and not more. Formally, we have the following lemma.

Lemma 1 (Obraztsova et al. [17], Lemma 2). *Consider a utility profile $\mathbf{u}$, let $\mathbf{a}$ be the respective truthful ballot vector, and let $\mathbf{b} \neq \mathbf{a}$ be a non-truthful PNE of $(\mathcal{T}, R^L, \mathbf{u})$. Then $T(\mathbf{b}) \neq \emptyset$. Further, $\mathrm{sc}(c_k, \mathbf{b}) = \mathrm{sc}(c_k, \mathbf{a})$ for every $c_k \in T(\mathbf{b})$, i.e., all voters whose top choice is c_k vote for c_k.*

The existence of a threshold candidate is an important observation about the structure of non-truthful PNE, and we will use it repeatedly in the sequel. Note that the winner in $\mathbf{a}$ does not have to be a threshold candidate in a non-truthful PNE $\mathbf{b}$.

Obraztsova et al. show that, given a candidate $c_p \in C$ and a score s, it is computationally hard to decide whether the game $(\mathcal{T}, R^L, \mathbf{u})$ has a PNE $\mathbf{b}$ where c_p wins with a score of s. This problem may appear to be "harder" than $(\mathcal{T}, R^L)$-TieNE or $(\mathcal{T}, R^L)$-SingleNE, as one needs to ensure that c_p obtains a specific score; on the other hand, it does not distinguish between c_p being the unique top-scorer or being tied with other candidates and winning due to tie-breaking. We now complement this hardness result by showing that all three problems we consider are NP-hard for $\mathcal{T}$ and R^L.

Theorem 3. $(\mathcal{T}, R^L)$-SINGLENE, $(\mathcal{T}, R^L)$-EXISTNE, and $(\mathcal{T}, R^L)$-TIENE are NP-*complete.*

The proof is by reduction from MAXIMUM k-SUBSET INTERSECTION (MSI); see [9] for a formal definition of this problem. Surprisingly, the complexity of MSI was very recently posed as an open problem by Clifford and Popa [5]; subsequently, MSI was shown to be hard under Cook reductions in [21]. In our proof we first establish NP-hardness of MSI under Karp reductions, which may be of independent interest, and then show NP-hardness of our problems by constructing reductions from MSI.

Randomized Tie-breaking.   It turns out that for truth-biased voters, the tie-breaking rules R^C and R^V induce identical behavior by the voters; unlike for lazy voters, this holds even if all voters rank the same candidate first.

For clarity, we present our characterization result for randomized tie-breaking in three parts. We start by considering PNE with winning sets of size at least 2; the analysis for this case turns out to be very similar to that for lazy voters.

Theorem 4. *Let* $\mathbf{u} = (u_1, \ldots, u_n)$ *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,* $\frac{1}{n} \sum_{i \in N} u_\ell(a_i) \geq \max_{i \in N \setminus \{\ell\}} u_\ell(a_i)$ *for each* $\ell \in N$.
(2) *there exists a set of candidates* $X = \{c_{\ell_1}, \ldots, c_{\ell_k}\}$ *with* $2 \leq k \leq \min(n/2, m)$ *and a partitioning of the voters into* k *groups* $N_1, \ldots, N_k$ *of size* n/k *each such that for each* $j \in [k]$ *and each* $i \in N_j$ *we have* $c_{\ell_j} \succ_i c$ *for all* $c \in X \setminus \{c_{\ell_j}\}$, *and, moreover,* $\frac{1}{k} \sum_{c \in X} u_i(c) \geq \max_{c \in X \setminus \{c_{\ell_j}\}} u_i(c)$.

Further, if condition (1) holds, then G *has a PNE where each voter votes for her top candidate, and if condition (2) holds for some* X, *then* G *has a PNE where each voter votes for her favorite candidate in* X. *The game* G *has no other PNE.*

The case where the winning set is a singleton is surprisingly complicated. We will first characterize utility profiles that admit a truthful PNE with this property.

Theorem 5. *Let* $\mathbf{u} = (u_1, \ldots, u_n)$ *be a utility profile over* C, *let* $R \in \{R^C, R^V\}$, *and suppose that* $W(\mathbf{a}) = \{c_j\}$ *for some* $c_j \in C$. *Then* $\mathbf{a}$ *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* $c_j \succ_i c_k$.

Finally, we consider elections that have non-truthful equilibria with singleton winning sets.

Theorem 6. *Let* $\mathbf{u} = (u_1, \ldots, u_n)$ *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* $c_j \in C$ *and* $b_r \neq a_r$ *for some* $r \in N$. *Then* $\mathbf{b}$ *is a PNE of the game* $G = (\mathcal{T}, R, \mathbf{u})$ *if and only if all of the following conditions hold:*

(1) $b_i \in \{a_i, c_j\}$ for all $i \in N$;

(2) $H(\mathbf{b}) \neq \emptyset$;

(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 every candidate $c_\ell \in H'(\mathbf{b})$ and each voter $i \in N$ with $b_i = c_j$, i prefers c_j to the lottery where a candidate is chosen from $H(\mathbf{b}) \cup \{c_j, c_\ell\}$ according to R.

We now consider the complexity of EXISTNE, TIENE, and SINGLENE for truth-biased voters and randomized tie-breaking. The reader may observe that the characterization of PNE with ties in Theorem 4 is essentially identical to the one in Theorem 2. As a consequence, we immediately obtain that $(\mathcal{T}, R^C)$-TIENE and $(\mathcal{T}, R^V)$-TIENE are NP-hard. For EXISTNE and SINGLENE, a simple modification of the proof of Theorem 3 shows that these problems remain hard under randomized tie-breaking. These observations are summarized in the following corollary.

Corollary 3. *For $R \in \{R^C, R^V\}$, $(\mathcal{T}, R)$-SINGLENE, $(\mathcal{T}, R)$-TIENE, and $(\mathcal{T}, R)$-EXISTNE are NP-complete.*

5   Comparison

We are finally in a position to compare the different models considered in this paper.

Tie-breaking Rules. We have demonstrated that in equilibrium the two randomized tie-breaking rules (R^C and R^V) induce very similar behavior, and identical election outcomes, both for lazy and for truth-biased voters. This is quite remarkable, since under truthful voting these tie-breaking rules can result in very different lotteries. In contrast, there is a substantial difference between the randomized rules and the lexicographic rule. For instance, with lazy voters, EXISTNE is NP-hard for R^C and R^V, but polynomial-time solvable for R^L. Further, R^L is, by definition, not neutral, and Theorem 1 demonstrates that candidates with smaller indices have a substantial advantage. For truth-biased voters the impact of tie-breaking rules is less clear: while we have NP-hardness results for all three rules, it appears that, in contrast with lazy voters, PNE induced by randomized tie-breaking are "simpler" than those induced by R^L.

Lazy vs. Truth-Biased Voters. Under lexicographic tie-breaking, the sets of equilibria induced by the two types of secondary preferences are incomparable: there exists a utility profile $\mathbf{u}$ such that the sets of candidates who can win in PNE of $(\mathcal{L}, R^L, \mathbf{u})$ and $(\mathcal{T}, R^L, \mathbf{u})$ are disjoint.

Example 1. Let $C = \{c_1, c_2, c_3\}$, and consider a 4-voter election with one vote of the form $c_2 \succ c_3 \succ c_1$, and three votes of the form $c_3 \succ c_2 \succ c_1$. The only PNE of $(\mathcal{L}, R^L, \mathbf{u})$ is $(c_2, \emptyset, \emptyset, \emptyset)$, where c_2 wins, whereas the only PNE of $(\mathcal{T}, R^L, \mathbf{u})$ is (c_2, c_3, c_3, c_3), where c_3 wins.

For randomized tie-breaking, the situation is more interesting. For concreteness, let us focus on R^C. Note first that the utility profiles for which there exist PNE with winning sets of size 2 or more are the same for both voter types. Further, if $(\mathcal{L}, R^C, \mathbf{u})$ has a PNE $\mathbf{b}$, with $|W(\mathbf{b})| = 1$ (which happens only if there is a unanimous winner), then $\mathbf{b}$ is also a PNE of $(\mathcal{T}, R^C, \mathbf{u})$. However, $(\mathcal{T}, R^C, \mathbf{u})$ may have additional PNE, including some non-truthful ones. In particular, for truth-biased voters, the presence of a strong candidate is sufficient for stability: Proposition 3 implies that if there exists a $c \in C$ such that $\mathrm{sc}(c, \mathbf{a}) \geq \mathrm{sc}(c', \mathbf{a}) + 2$ for all $c' \in C \setminus \{c\}$, then for any $R \in \{R^L, R^C, R^V\}$ the truthful ballot vector $\mathbf{a}$ is a PNE of $(\mathcal{T}, R, \mathbf{u})$ with $W(\mathbf{a}) = \{c\}$.

Existence of PNE. For truth-biased voters, one can argue that, when the number of voters is large relative to the number of candidates, under reasonable probabilistic models of elections, the existence of a strong candidate (as defined in the previous paragraph) is exceedingly likely. Thus, elections with truth-biased voters typically admit stable outcomes; this is corroborated by the experimental results of [20]. In contrast, for lazy voters stability is more difficult to achieve, unless there is a candidate that is unanimously ranked first: under randomized tie-breaking rules, there needs to be a very precise balance among candidates that end up being in $W(\mathbf{b})$, and under R^L the eventual winner has to Pareto-dominate all candidates that lexicographically precede him.

Quality of PNE. In all of our models, a candidate ranked last by all voters cannot be elected, in contrast to the basic game-theoretic model for Plurality voting. However, not all non-desirable outcomes are eliminated: under R^V and R^C both lazy voters and truth-biased voters can still elect a Pareto-dominated candidate with non-zero probability in PNE. This has been shown for lazy voters and R^C (Example 1 in [6]), and the same example works for truth-biased voters and R^V. A similar construction shows that a Pareto-dominated candidate may win under R^L when voters are truth-biased. In contrast, lazy voters cannot elect a Pareto-dominated candidate under R^L: Theorem 1 shows that the winner has to be ranked first by some voter.

We can also measure the quality of PNE by analyzing the Price of Anarchy (PoA) in both models. The study of PoA in the context of voting has been

Table 1. Complexity results: P stands for "polynomial-time solvable", NPc stands for "NP-complete".

	SINGLENE	TIENE	EXISTNE
$(\mathcal{L}, R^L)$	P (Corollary 1)	P (Corollary 1)	P (Corollary 1)
$(\mathcal{L}, R^C)$	P (Corollary 2)	NPc (Corollary 2)	NPc (Corollary 2)
$(\mathcal{L}, R^V)$	P (Corollary 2)	NPc (Corollary 2)	NPc (Corollary 2)
$(\mathcal{T}, R^L)$	NPc (Theorem 3)	NPc (Theorem 3)	NPc (Theorem 3)
$(\mathcal{T}, R^C)$	NPc (Corollary 3)	NPc (Corollary 3)	NPc (Corollary 3)
$(\mathcal{T}, R^V)$	NPc (Corollary 3)	NPc (Corollary 3)	NPc (Corollary 3)

recently initiated by Branzei et al. [4]. The additive version of PoA, which was considered in [4], is defined as the worst-case difference between the score of the winner under truthful voting and the truthful score of a PNE winner. It turns out that PoA can be quite high, both for lazy and for truth-biased voters: in the full version of the paper we show that PoA $= \Omega(n)$ in all of our models. Even though these results are not encouraging, PoA is only a worst-case analysis and we expect a better performance on average. Indeed, for the truth-biased model, this is supported by the experimental evaluation in [20].

6 Conclusions

We have characterized PNE of Plurality voting for several combinations of secondary preferences and tie-breaking rules. Our complexity results are summarized in Table 1.

Our results extend to the setting where some of the voters are *principled*, i.e., always vote truthfully (and never abstain). Due to space constraints, we are unable to fully describe these extensions (see [9]). Briefly, the presence of principled voters has the strongest effect on lazy voters and lexicographic tie-breaking, as illustrated by the following example, whereas for other settings the effect is less pronounced.

Example 2. Consider an election over a candidate set $C = \{c_1, \ldots, c_m\}$, $m > 1$, where there are two principled voters who both vote for c_m, and two lazy voters who both rank c_m last. Then the ballot vector where both lazy voters abstain is a PNE (with winner c_m). Moreover, for every $j \in [m-1]$ the ballot vector where both lazy voters vote for c_j is a PNE as well (with winner c_j).

In the absence of principled voters, PNE for lazy voters require very precise coordination among the voters and seem to be very different from what we observe in real life. In contrast, for truth-biased voters the presence of a strong candidate implies the existence of a truthful equilibrium, which requires little coordination among the players. It is therefore tempting to conclude that truth bias has a greater explanatory power than laziness. However, we demonstrated that the presence of principled voters changes this equation. Extending our analysis to a mixture of all three voter types is perhaps the most prominent open problem suggested by our work.

References

1. Aziz, H., Gaspers, S., Mattei, N., Narodytska, N., Walsh, T.: Ties matter: complexity of manipulation when tie-breaking with a random vote. In: AAAI 2013, pp. 74–80 (2013)
2. Battaglini, M.: Sequential voting with abstention. Games Econ. Behav. **51**, 445–463 (2005)
3. Borgers, T.: Costly voting. Am. Econ. Rev. **94**(1), 57–66 (2004)
4. Branzei, S., Caragiannis, I., Morgenstern, J., Procaccia, A.D.: How bad is selfish voting? In: AAAI 2013, pp. 138–144 (2013)
5. Clifford, R., Popa, A.: Maximum subset intersection. Inf. Process. Lett. **111**(7), 323–325 (2011)
6. Desmedt, Y., Elkind, E.: Equilibria of plurality voting with abstentions. In: ACM EC 2010, pp. 347–356 (2010)
7. Dhillon, A., Lockwood, B.: When are plurality rule voting games dominance-solvable? Games Econ. Behav. **46**, 55–75 (2004)
8. Dutta, B., Sen, A.: Nash implementation with partially honest individuals. Games Econ. Behav. **74**(1), 154–169 (2012)
9. Elkind, E., Markakis, E., Obraztsova, S., Skowron, P.: Equilibria of plurality voting: Lazy and truth-biased voters. arXivabs/1409.4132 (2014)
10. Farquharson, R.: Theory of Voting. Yale University Press, New Haven (1969)
11. Laslier, J.F., Weibull, J.W.: A strategy-proof condorcet jury theorem. Scand. J. Econ. (2012)
12. Meir, R., Polukarov, M., Rosenschein, J.S., Jennings, N.R.: Convergence to equilibria in plurality voting. In: AAAI 2010, pp. 823–828 (2010)
13. Moulin, H.: Dominance solvable voting schemes. Econometrica **47**, 1337–1351 (1979)
14. Myerson, R., Weber, R.: A theory of voting equilibria. Am. Polit. Sci. Rev. **87**(1), 102–114 (1993)
15. Obraztsova, S., Elkind, E.: On the complexity of voting manipulation under randomized tie-breaking. In: IJCAI 2011, pp. 319–324 (2011)
16. Obraztsova, S., Elkind, E., Hazon, N.: Ties matter: Complexity of voting manipulation revisited. In: AAMAS 2011, pp. 71–78 (2011)
17. Obraztsova, S., Markakis, E., Thompson, D.R.M.: Plurality voting with truth-biased agents. In: Vöcking, B. (ed.) SAGT 2013. LNCS, vol. 8146, pp. 26–37. Springer, Heidelberg (2013)
18. Schulze, M.: A new monotonic, clone-independent, reversal symmetric, and condorcet-consistent single-winner election method. Soc. Choice Welfare **36**(2), 267–303 (2011)
19. Sinopoli, F.D., Iannantuoni, G.: On the generic strategic stability of Nash equilibria if voting is costly. Econ. Theory **25**(2), 477–486 (2005)
20. Thompson, D.R.M., Lev, O., Leyton-Brown, K., Rosenschein, J.S.: Empirical analysis of plurality election equilibria. In: AAMAS 2013, pp. 391–398 (2013)
21. Xavier, E.: A note on a maximum k-subset intersection problem. Inf. Process. Lett. **112**(12), 471–472 (2012)