The Thirty-Second AAAI Conference
on Artificial Intelligence (AAAI-18)

Manipulative Elicitation — A New Attack on Elections with Incomplete Preferences

Palash Dey
palash@tifr.res.in
Tata Institute of Fundamental Research, Mumbai

Abstract

Lu and Boutilier (2011) proposed a novel approach based on "minimax regret" to use classical score based voting rules in the setting where preferences can be any partial (instead of complete) orders over the set of alternatives. We show here that such an approach is vulnerable to a new kind of manipulation which was not present in the classical (where preferences are complete orders) world of voting. We call this attack "manipulative elicitation." More specifically, it may be possible to (partially) elicit the preferences of the agents in a way that makes some distinguished alternative win the election who may not be a winner if we elicit every preference completely. More alarmingly, we show that the related computational task is polynomial time solvable for a large class of voting rules which includes all scoring rules, maximin, Copeland$^\alpha$ for every $\alpha \in [0, 1]$, simplified Bucklin voting rules, etc. We then show that introducing a parameter per pair of alternatives which specifies the minimum number of partial preferences where this pair of alternatives must be comparable makes the related computational task of manipulative elicitation NP-complete for all common voting rules including a class of scoring rules which includes the plurality, k-approval, k-veto, veto, and Borda voting rules, maximin, Copeland$^\alpha$ for every $\alpha \in [0, 1]$, and simplified Bucklin voting rules. Hence, in this work, we discover a fundamental vulnerability in using minimax regret based approach in partial preferential setting and propose a novel way to tackle it.

1 Introduction

Aggregating preferences of a set of agents over a set of alternatives is a fundamental problem in voting theory which has been used in many applications in AI for making various decisions. Prominent examples of such applications include collaborative filtering (Pennock, Horvitz, and Giles 2000), similarity search (Fagin, Kumar, and Sivakumar 2003), winner determination in sports competitions (Betzler, Bredereck, and Niedermeier 2014), etc. (Moulin et al. 2016). In a typical scenario of voting, we have a set of alternatives, a tuple of "preferences", called a profile, over the set of alternatives, and a voting rule which chooses a set of alternatives as winners based on the profile. Classically, preferences are often modeled as complete orders over the set of alternatives. However, in typical applications of voting in AI, collaborative filtering for example, the number of alternatives is huge and we have only partial orders over the set of alternatives as preferences.

There have been many attempts to extend the use of voting theory in settings with incomplete preferences. The approach of Konczak and Lang (2005) was to study the possible and necessary winner problems. In these problems, the input is a profile of partial preferences and we want to compute the set of alternatives who wins (under some fixed voting rule) in at least one completion of the profile for the possible winner problem; for the necessary winner problem, we want to compute the set of alternatives who wins in every completion of the profile. There have been substantial research effort in the last decade to better understand these two problems (Lang et al. 2007; Pini et al. 2007; Walsh 2007; Xia and Conitzer 2011; Betzler, Hemmann, and Niedermeier 2009; Chevaleyre et al. 2010; Betzler, Bredereck, and Niedermeier 2010; Baumeister, Roos, and Rothe 2011; Lang et al. 2012; Faliszewski et al. 2014; Dey, Misra, and Narahari 2016b; 2016a; 2017; 2015; Dey and Misra 2017). One of the main criticisms of this approach is that the definition of a necessary winner is so strong that none of the alternatives may satisfy it whereas the definition of a possible winner is so relaxed that a large number of alternatives may satisfy it. Moreover, the computational problem of finding the set of possible winners is NP-hard for most of the common voting rules (finding the set of necessary winners is also co-NP-hard for some voting rules, ranked pairs for example) (Xia and Conitzer 2011).

Lu and Boutilier (2011) took a completely different approach to handle incomplete preferences and proposed a worst case regret based approach for score based voting rules. These voting rules assign some score to every alternative based on the profile and select the alternatives with the maximum (or minimum) score as winners. Many popular voting rules, for example, scoring rules, maximin, Copeland, etc. are score based voting rules. For score based voting rules, intuitively speaking, the worst case regret, called maximum regret in (Lu and Boutilier 2011), of declaring an alternative w as a winner is the maximum possible difference between the score of w and the score of a winning alternative in any completion of the input partial profile; the winners of a partial profile are the set of alternatives with the minimum maximum (called minimax) regret. A completion of a partial profile is another profile where every preference is complete and it respects the orderings of the corresponding preference in the partial profile. The minimax regret based approach is not only theoretically robust as argued in (Lu and Boutilier 2011) but also practically appealing since computing winners is polynomial time solvable for all commonly used voting rules.

1.1 Motivation

Although the minimax regret based approach enjoys many exciting features, it introduces a new (which was not present in the classical setting with complete preferences) kind of attack on the election which we call "manipulative elicitation." That is, it may be possible to partially elicit the preferences in such a way that makes some favorable alternative win the election. For example, let us consider a plurality election $\mathcal{E}$ where an alternative, say w, is the top alternative of one preference and another alternative, say x, is the top alternative of every other preference. In a plurality election, the winners are the set of alternatives who appear as the top alternative in the largest number of preferences. Hence, x is the unique winner in $\mathcal{E}$. Let us now consider a partial profile where, in every partial preference, only w and every other alternative who is preferred less than w in the corresponding preference in $\mathcal{E}$ are comparable. Let us call the resulting partial profile $\mathcal{E}'$. If n is the number of preferences, then the minimax regret plurality score of w in $\mathcal{E}'$ is $(n-1)$ whereas the minimax regret plurality score of every other alternative is n which makes w the unique winner of $\mathcal{E}'$. We call this phenomenon manipulative elicitation. The problem of manipulative elicitation is even more alarming in AI since, in many applications (collaborative filtering for example), the parts of the preferences that will be elicited can often be influenced and controlled in such settings.

1.2 Our Contribution

Our main contribution in this paper is the discovery of the manipulative elicitation attack in regret based partial preferential setting. We also show that the corresponding computational problem for manipulative elicitation is polynomial time solvable for every *monotone* voting rule which includes all commonly used score based voting rules [Theorem 1 and Corollary 1]. Intuitively speaking, we call a score based voting rule monotone if improving the position of some alternative in any (complete) preference can only improve its score; we defer its formal definition till Section 2. To counter the negative result of Theorem 1, we introduce a parameter per pair of alternatives which specifies the minimum number of partial preferences where these two alternatives should be comparable. We establish success of our approach by showing that the new constraints make the corresponding computational task of manipulative elicitation NP-complete for a large class of scoring rules [Theorem 2] which includes the plurality [Theorem 3], veto [Theorem 4], k-approval for any k, and Borda voting rules [Corollary 2], maximin [Theorem 5], Copeland$^\alpha$ for every $\alpha \in [0,1]$ [Theorem 6], and simplified Bucklin [Theorem 7] voting rules. We remark that there could be various ways to enforce lower bounds on the number of partial preferences where a particular pair of alternatives is comparable. For example, this can be a feature in the applications which would allow users to generate these bounds from some distribution which would in turn overrule the possibility of such manipulation (due to our hardness results).

2 Preliminaries and Problem Formulation

For a positive integer k, we denote the set $\{1, 2, \ldots, k\}$ by $[k]$. Let $\mathcal{A} = \{a_i : i \in [m]\}$ be a set of m *alternatives*. We denote the set of all subsets of $\mathcal{A}$ of cardinality 2 by $\binom{\mathcal{A}}{2}$. A complete order over the set $\mathcal{A}$ of alternatives is called a (complete) *preference*. We say that an alternative $a \in \mathcal{A}$ is placed at the ℓ^{th} position (from left or from top) in a preference $\succ$ if $|\{b \in \mathcal{A} : b \succ a\}| = \ell - 1$. We denote the set of all possible preferences over $\mathcal{A}$ by $\mathcal{L}(\mathcal{A})$. A tuple $\succ = (\succ_i)_{i\in[n]} \in \mathcal{L}(\mathcal{A})^n$ of n preferences is called a *profile*. An *election* $\mathcal{E}$ is a tuple $(\succ, \mathcal{A})$ where $\succ$ is a profile over a set $\mathcal{A}$ of alternatives. If not mentioned otherwise, we denote the number of alternatives and the number of preferences by m and n respectively. A map $r_c : \uplus_{n,|\mathcal{A}|\in\mathbb{N}^+} \mathcal{L}(\mathcal{A})^n \longrightarrow 2^{\mathcal{A}} \setminus \{\emptyset\}$ is called a *voting rule*. Given an election $\mathcal{E}$, we can construct from $\mathcal{E}$ a directed weighted graph $\mathcal{G}_\mathcal{E}$ which is called the *weighted majority graph* of $\mathcal{E}$. The set of vertices in $\mathcal{G}_\mathcal{E}$ is the set of alternatives in $\mathcal{E}$. For any two alternatives x and y, the weight of the edge (x, y) is $\mathcal{D}_\mathcal{E}(x, y) = \mathcal{N}_\mathcal{E}(x, y) - \mathcal{N}_\mathcal{E}(y, x)$, where $\mathcal{N}_\mathcal{E}(a, b)$ is the number of preferences where the alternative a is preferred over the alternative b for $a, b \in \mathcal{A}, a \neq b$. Examples of some common voting rules are as follows.

$\triangleright$**Positional scoring rules:** An m-dimensional vector $\alpha = (\alpha_1, \alpha_2, \ldots, \alpha_m) \in \mathbb{N}^m$ with $\alpha_1 \geqslant \alpha_2 \geqslant \ldots \geqslant \alpha_m$ and $\alpha_1 > \alpha_m$ for every $m \in \mathbb{N}$ naturally defines a voting rule — an alternative gets score α_i from a preference if it is placed at the i^{th} position, and the score of an alternative is the sum of the scores it receives from all the preferences. The winners are the alternatives with the maximum score. Scoring rules remain unchanged if we multiply every α_i by any constant $\lambda > 0$ and/or add any constant μ. Hence, we can assume without loss of generality that for any score vector α, we have $gcd((\alpha_i)_{i\in[m]}) = 1$ and there exists a $j < m$ such that $\alpha_\ell = 0$ for all $\ell > j$. We call such an α a *normalized* score vector. If α_i is 1 for $i \in [k]$ and 0 otherwise, then, we get the k-*approval* voting rule. The k-*approval* voting rule is also called the $(m-k)$-*veto* voting rule. The 1-approval voting rule is called the *plurality* voting rule and the 1-veto voting rule is called the *veto* voting rule. If $\alpha_i = m - i$ for every $i \in [m]$, then we get the *Borda* voting rule.
$\triangleright$**Maximin:** The maximin score of an alternative x is $\min_{y\neq x} \mathcal{N}_\mathcal{E}(x, y)$. The winners are the alternatives with the maximum maximin score.
$\triangleright$**Copeland$^\alpha$:** Given $\alpha \in [0, 1]$, the Copeland$^\alpha$ score of an alternative x is $|\{y \neq x : \mathcal{D}_\mathcal{E}(x, y) > 0\}| + \alpha|\{y \neq x : \mathcal{D}_\mathcal{E}(x, y) = 0\}|$. The winners are the alternatives with the maximum Copeland$^\alpha$ score.
$\triangleright$**Simplified Bucklin:** An alternative x's simplified Bucklin score is the minimum number ℓ such that x is placed within the top ℓ positions in more than half of the preferences. The

winners are the alternatives with the lowest simplified Bucklin score.

We call a voting rule *"score based"* if the voting rule assigns some score to every alternative based on the profile and chooses either the set of alternatives with the maximum score or the set of alternatives with the minimum score as winners. All the above mentioned voting rules are score based. We say that a score based voting rule s is *monotone* if, for every positive integer n, every two profiles $(\succ_i)_{i\in[n]}$ and $(\succ'_i)_{i\in[n]}$ over any finite set $\mathcal{A}$ of alternatives, and every alternative $x \in \mathcal{A}$ such that $\{y \in \mathcal{A} : x \succ_i y\} \subseteq \{y \in \mathcal{A} : x \succ'_i y\}$ for every $i \in [n]$, we have $s(x, (\succ_i)_{i\in[n]}) \leqslant s(x, (\succ'_i)_{i\in[n]})$. We call a voting rule r *neutral* if, for every positive integer n, every profile $(\succ_i)_{i\in[n]}$ over any finite set $\mathcal{A} = \{x_i : i \in [m]\}$ of m alternatives, and every permutation σ of $[m]$, we have $\sigma(r((\succ_i)_{i\in[n]})) = r((\sigma(\succ_i))_{i\in[n]})$ where $\sigma(\succ_i) = x_{\sigma(1)} \succ x_{\sigma(2)} \succ \cdots \succ x_{\sigma(m)}$ if $\succ_i = x_1 \succ x_2 \succ \cdots \succ x_m$. We call a voting rule *worst efficient* if the worst possible score with n preferences over m alternatives can be computed in a polynomial (in m and n) amount of time. We observe that all the voting rules mentioned above are neutral, worst efficient, and monotone if, for the case of simplified Bucklin voting rule, we replace the simplified Bucklin score with negative of that and choose the alternative with the maximum score.

2.1 Incomplete Election and Minimax Regret Extension of Score Based Voting Rules

Although preferences are classically modeled as complete orders, in many scenarios, preferences can be any partial order over $\mathcal{A}$. We often denote a partial order $\mathcal{R}$ by the set $\{(a, b) : a, b \in \mathcal{A}, a\mathcal{R}b\}$. Given a profile $\mathcal{P}$ of partial preferences (which we call a partial profile), we denote the set of all completions of $\mathcal{P}$ to complete orders by $\mathcal{C}(\mathcal{P})$. Lu and Boutilier (2011) proposed a novel approach to extend the use of score based voting rules for settings with partial profiles based on a notion of regret. Let s be a score based voting rule so that the winner is an alternative with the maximum score. Positional scoring rules, maximin, Copeland$^\alpha$ for every $\alpha \in [0, 1]$, etc. are prominent examples of such score based voting rules. Let us denote the score that a score based voting rule s assigns to an alternative $a \in \mathcal{A}$ in a profile $\succ \in \mathcal{L}(\mathcal{A})^n$ by $s(a, \succ)$. We denote the minimax regret voting rule based on a voting rule s by $\mathfrak{s}$. For a profile $\succ$, let $s(\succ) = \text{argmax}_{a\in\mathcal{A}}\{s(a, \succ)\}$. Given a partial profile $\mathcal{P}$ and a score based rule s, $\mathfrak{s}(\mathcal{P})$ is defined as follows.

$$
\begin{aligned}
s - Regret(a, \succ) &= |s(s(\succ), \succ) - s(a, \succ)| \\
s - MR(a, \mathcal{P}) &= \max_{\succ \in \mathcal{C}(\mathcal{P})} s - Regret(a, \succ) \\
\mathfrak{s}(\mathcal{P}) &= \operatorname*{argmin}_{a\in\mathcal{A}} s - MR(a, \mathcal{P})
\end{aligned}
$$

For a partial profile $\mathcal{P}$ and a minimax regret (MR for short) voting rule $\mathfrak{s}$, we say that an alternative $a \in \mathcal{A}$ co-wins if $a \in \mathfrak{s}(\mathcal{P})$ and wins uniquely if $\mathfrak{s}(\mathcal{P}) = \{a\}$. For an alternative $a \in \mathcal{A}$, if $s - MR(a, \mathcal{P}) = s - Regret(a, \succ)$ for some $\succ \in \mathcal{C}(\mathcal{P})$, then we call an alternative in $s(\succ)$ a *competing alternative* of a in $\mathcal{P}$.

We now formally define manipulative elicitation and the basic computational problem of manipulative elicitation for a score based voting rule s.

Definition 1 ($\mathfrak{s}$-manipulative elicitation). *For a profile $\succ$ over a set $\mathcal{A}$ of alternatives, we say that a partial profile $\mathcal{P}$ is called a manipulative elicitation if $\succ \in \mathcal{C}(\mathcal{P})$ and $\mathfrak{s}(\mathcal{P}) = \{c\}$.*

Definition 2 ($\mathfrak{s}$-MANIPULATIVE ELICITATION). *Given a set $\mathcal{A}$ of alternatives, a profile $\succ \in \mathcal{L}(\mathcal{A})^n$ of n preferences, and an alternative $c \in \mathcal{A}$, compute if there exists a partial profile $\mathcal{P}$ such that $\succ \in \mathcal{C}(\mathcal{P})$ and $\mathfrak{s}(\mathcal{P}) = \{c\}$?*

We will see in Theorem 1 that the $\mathfrak{s}$-MANIPULATIVE ELICITATION problem is polynomial time solvable for every neutral, monotone, and worst efficient score based voting rule. This shows that all the commonly used voting rule considered here are vulnerable under manipulative elicitation. In the hope to counter this drawback, we extend the basic problem in Definition 2 to MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT in Definition 3. We will indeed see in Section 4 that the MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT problem is NP-complete for all the voting rules that we consider in this paper. For a partial profile $\succ = (\succ_i)_{i\in[n]}$ and $\{a, b\} \in \binom{\mathcal{A}}{2}$, we denote the number of partial preferences in $\succ$ where a and b are comparable by $\mathfrak{p}_{\{a,b\}}(\succ)$.

Definition 3 ($\mathfrak{s}$-MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT). *Given a set $\mathcal{A}$ of alternatives, a profile $\succ \in \mathcal{L}(\mathcal{A})^n$ of n voters, a function $\mathfrak{f} : \binom{\mathcal{A}}{2} \longrightarrow \mathbb{N}$ such that $0 \leqslant \mathfrak{f}(\{a, b\}) \leqslant n$ for every $\{a, b\} \in \binom{\mathcal{A}}{2}$, and an alternative $x \in \mathcal{A}$, compute if there exists a partial profile $\mathcal{P}$ such that $\succ \in \mathcal{C}(\mathcal{P})$, $\mathfrak{f}(\{a, b\}) \leqslant \mathfrak{p}_{\{a,b\}}(\succ)$ for every $\{a, b\} \in \binom{\mathcal{A}}{2}$ and $\mathfrak{s}(\mathcal{P}) = \{x\}$?*

We remark that both the computational problems in Definition 2 and 3 have been defined for the unique winner case; we could as well define these problems for the co-winner case also. It turns out that all our proofs (except Theorem 1) can be easily modified for the co-winner counterpart and our choice for defining these problems in the unique winner setting is only a matter of exposition.

3 Polynomial Time Algorithm for MANIPULATIVE ELICITATION

Our first result is Theorem 1 which shows that the MANIPULATIVE ELICITATION problem is polynomial time solvable for a large class of voting rules.

Theorem 1. *The* MANIPULATIVE ELICITATION *problem is polynomial time solvable for every monotone, neutral, and worst efficient score based voting rule s.*

Proof. Let $(\mathcal{A}, \succ = (\succ_i)_{i\in[n]}, c)$ be an arbitrary instance of $\mathfrak{s}$-MANIPULATIVE ELICITATION. Our algorithm is as follows. If c receives the worst possible score in $\succ$, then we output NO; otherwise we output YES. Our algorithm runs in polynomial time since s is worst efficient. To prove the correctness of our algorithm, we begin with Claim 1 below.

Claim 1. *If the score of c in $\succ$ is the worst possible score (say β_n) that any alternative in $\mathcal{A}$ can possibly receive in any profile with n preferences under s, then the $\mathfrak{s}$-*MANIPULATIVE ELICITATION *instance is a* NO *instance.*

Proof. Suppose not, then let us assume that $\mathcal{R} = (R_i)_{i\in[n]}$ be a partial profile such that $\succ \in \mathcal{C}(\mathcal{R})$ and $\mathfrak{s}(\mathcal{R}) = \{c\}$. Let $s - MR(c,\mathcal{R}) = s(s(\succ'),\succ') - s(c,\succ')$ for some $\succ' = (\succ_i')_{i\in[n]} \in \mathcal{C}(\mathcal{R})$. We now claim the following.

Claim 2. $s(c,\succ') = \beta_n$.

Proof. The idea of the proof is that if $s(c,\succ') > \beta_n$, then we can construct another profile which can be used to calculate worse regret for c than $\succ'$ and this will contradict the choice of $\succ'$. Formally, let us define another profile $\bar{\succ} = (\bar{\succ}_i)_{i\in[n]}$ where $\bar{\succ}_i$ is obtained from $\succ_i'$ by "moving" c immediately to the right of the alternatives that are on the left of c in either $\succ_i'$ or $\succ_i$ for $i \in [n]$; that is, for every $i \in [n]$, $\bar{\succ}_i$ is defined as follows.

$$\begin{aligned}\bar{\succ}_i = \{(a,b) : a,b \in \mathcal{A}\setminus\{c\}, a \succ_i' b\}\\ \cup \{(c,a) : a \in \mathcal{A}, c \succ_i' a, c \succ_i a\}\\ \cup \{(a,c) : a \in \mathcal{A}, a \succ_i' c \text{ or } a \succ_i c\}\end{aligned}$$

The profile $\bar{\succ} \in \mathcal{C}(\mathcal{R})$ since $\succ \in \mathcal{C}(\mathcal{R})$ and $\succ' \in \mathcal{C}(\mathcal{R})$. Due to monotonicity of s, the score of c in $\bar{\succ}$ is at most the score of c in $\succ'$ and the score of every other alternative in $\bar{\succ}$ is at least their score in $\succ'$. However, $\succ'$ has been used to calculate the MR score of c under s. Hence, we have the following:

$$s(s(\succ'),\succ') = s(s(\bar{\succ}),\succ'), \quad s(c,\succ') = s(c,\bar{\succ})$$

We now have the following:

$$\beta_n \leqslant s(c,\succ') = s(c,\bar{\succ}) \leqslant s(c,\succ) = \beta_n$$

The first inequality follows from the definition of β_n and the second inequality follows from monotonicity of s. □

Let $y \in s(\succ')$ and $s - MR(y,\mathcal{R}) = s(s(\succ''),\succ'') - s(y,\succ'')$ for some $\succ'' = (\succ_i'')_{i\in[n]} \in \mathcal{C}(\mathcal{R})$. We now have the following claim.

Claim 3. $s(s(\succ''),\succ'') \leqslant s(s(\succ'),\succ')$

We now combine Claim 2 and 3 as follows to prove the main claim.

$$\begin{aligned}s - MR(y,\mathcal{R}) &= s(s(\succ''),\succ'') - s(y,\succ'')\\ &\leqslant s(s(\succ'),\succ') - s(y,\succ'')\\ &\leqslant s(s(\succ'),\succ') - \beta_n\\ &= s(s(\succ'),\succ') - s(c,\succ')\\ &= s - MR(c,\mathcal{R})\end{aligned}$$

The second line follows from Claim 3, the third line follows from the definition of β_n, and the fourth line follows from Claim 2. Hence we have $s - MR(y,\mathcal{R}) \leqslant s - MR(c,\mathcal{R})$ which contradicts our assumption that $\mathfrak{s}(\mathcal{R}) = \{c\}$. □

We now show that if c does not receive the worst possible score with n preferences over $\mathcal{A}$ under s from the profile $\succ$, then the instance is a YES instance. To see this, let us consider the partial profile $\mathcal{P} = (P_i)_{i\in[n]}$ as $P_i = \{c \succ y : c \succ_i y\}$ for every $i \in [n]$. Let $\mathcal{R}$ be any profile in $\mathcal{C}(\mathcal{P})$. Since, the alternative c does not receive the worst possible score with n preferences over $\mathcal{A}$ under s from the profile $\succ$, $s(c,\mathcal{R}) < \beta_n$. Hence, if α is the best possible score with n preferences over $\mathcal{A}$ under s, we have $s - MR(c,\mathcal{P}) < \alpha - \beta_n$. On the other hand, for any alternative $y \in \mathcal{A}\setminus\{c\}$, let us consider the profile $\mathcal{Q}_y = (Q_i)_{i\in[n]}$ where $Q_i = c \succ \cdots \succ y$ for every $i \in [n]$. Now due to monotonicity of s, we have $s(c,\mathcal{Q}_y) = \alpha$ and $s(y,\mathcal{Q}_y) = \beta_n$. Hence, we have $s - MR(y,\mathcal{P}) = \alpha - \beta_n$ for every $y \in \mathcal{A}\setminus\{c\}$ and thus $\mathfrak{s}(\mathcal{P}) = \{c\}$. □

We remark that the proof of Theorem 1 for the co-winner case is trivial: every instance is a YES instance since a partial profile where every preference is empty makes every alternative win due to neutrality. Since scoring rules, maximin, Copeland$^\alpha$ for every $\alpha \in [0,1]$, and simplified Bucklin voting rules are monotone, neutral, and worst efficient, Theorem 1 immediately implies the following corollary.

Corollary 1. *The* MANIPULATIVE ELICITATION *problem is polynomial time solvable for scoring rules, maximin, Copeland$^\alpha$ for every $\alpha \in [0,1]$, and simplified Bucklin voting rules.*

4 Hardness Results for MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT

In this section we show that the MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT problem is NP-complete for maximin, Copeland$^\alpha$ for every $\alpha \in [0,1]$, simplified Bucklin, and a large class of scoring rules which includes the k-approval voting rule for every k and the Borda voting rule. In the interest of space, we omit some of our proofs. They are available here (Dey 2017).

Let us define a restricted version of the classical set cover problem which we call SET COVER FREQUENCY TWO. We will see in Lemma 1 that this problem is NP-complete by reducing from the vertex cover problem which is well known to be NP-complete (Garey and Johnson 1979). Most of our NP-hardness reductions are from this problem.

Definition 4 (SET COVER FREQUENCY TWO)**.** *Given a universe $\mathcal{U}$ of cardinality q, a family $\mathcal{S} = \{S_i : i \in [t]\}$ of t subsets of $\mathcal{U}$ such that for every $a \in \mathcal{U}$, we have $|\{i \in [t] : a \in S_i\}| = 2$, and a positive integer ℓ, compute if there exists a subset $\mathcal{G} \subseteq \mathcal{S}$ containing at most ℓ sets such that $\cup_{A\in\mathcal{G}} A = \mathcal{U}$. We denote an arbitrary instance of* SET COVER FREQUENCY TWO *by $(\mathcal{U},\mathcal{S},\ell)$.*

Lemma 1. SET COVER FREQUENCY TWO *is* NP-*complete.*

We begin with showing that the MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT problem is NP-complete for a large class of scoring rules which included the k-approval voting rule for every $3 \leqslant k \leqslant \gamma m$ for any constant $0 < \gamma < 1$ and the Borda voting rule. While describing a (complete) preference, if we do not mention the order of any two alternatives, they can be ordered arbitrarily. On the other hand, if we are describing a partial preference and we do not mention the order of any two alternatives, then they should be assumed to be incomparable.

Theorem 2. *Let r be a normalized scoring rule such that there exists a function $g : \mathbb{N} \longrightarrow \mathbb{N}$ such that for every $m \in \mathbb{N}$, we have $3m \leqslant g(m) \leqslant poly(m)$ and if $\alpha = (\alpha_i)_{i\in[g(m)]}$, then there exists a positive integer $\mathfrak{p}$ such that $3 \leqslant \mathfrak{p} \leqslant g(m) - m + 3, \alpha_{\mathfrak{p}} > \alpha_{\mathfrak{p}+1}$ and $\alpha_{\mathfrak{p}-1} = poly(m)$. Then the* MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT *problem is* NP*-complete for the scoring rule r.*

Proof. The MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT problem clearly belongs to NP. To prove NP-hardness, we reduce from SET COVER FREQUENCY TWO to MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT for the scoring rule r. Let $(\mathcal{U} = \{u_1, \ldots, u_q\}, \mathcal{S} = \{S_i : i \in [t]\}, \ell)$ be an arbitrary instance of SET COVER FREQUENCY TWO. Let us consider the following instance $(\mathcal{A}, \mathcal{P}, c, \mathfrak{f})$ of MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT where $\mathcal{A}$ is defined as follows.

$$\mathcal{A} = \{a_i : i \in [q]\} \cup \{c, d\} \cup W,$$
$$\text{where } W = \{w_1, \ldots, w_{g(q)-q-2}\}$$

The profile $\mathcal{P}$ consists of the following preferences. For an integer $0 \leqslant k \leqslant g(q)-q-2$, we denote the set $\{w_i : i \in [k]\}$ by W_k. Let $\kappa = \max\{i \in [g(q)] : \alpha_i \neq 0\}$; we observe that $\alpha_\kappa = 1$ since r is normalized. For $X \subseteq \mathcal{U}$, let us denote the set $\{a_j : u_j \in X\}$ of alternatives by X to simplify notation.

- ▷ $\forall i \in [t] : W_{\mathfrak{p}-3} \succ S_i \succ d \succ c \succ (\mathcal{U} \setminus S_i) \succ (W \setminus W_{\mathfrak{p}-3})$
- ▷ $t - 2$ copies of $W_{\mathfrak{p}-1} \succ a_i \succ d \succ (\mathcal{U} \setminus \{a_i\}) \succ (W \setminus W_{\mathfrak{p}-1}) \succ c$
- ▷ $t(\alpha_{\mathfrak{p}} - \alpha_{\mathfrak{p}+2})$ copies of $W_2 \succ \cdots \succ c \succ \cdots d$ where the alternative c is placed at κ position from left.
- ▷ If $\alpha_{g(q)-1} = 1$, then we add $(q+1)t\alpha_{\mathfrak{p}-1}$ copies of $\cdots \succ c \succ d$
- ▷ Otherwise:
  - – If $\kappa \leqslant g(q) - q - 1$, then we add $(q + 1)t\alpha_{\mathfrak{p}-1}$ copies of $W_2 \succ \cdots \succ c \succ d \succ \mathcal{U} \succ \cdots$ where the alternative d is placed at $\kappa + 1$ position from left and we add, for every $i \in [q]$, $(q+1)t\alpha_{\mathfrak{p}-1}$ copies of $W_2 \succ \cdots \succ a_i \succ d \succ c \succ (\mathcal{U} \setminus \{a_i\}) \succ \cdots$ where the alternative d is placed at $\kappa + 1$ position from the left.
  - – Otherwise we add $(q + 1)t\alpha_{\mathfrak{p}-1}$ copies of $W_2 \succ \cdots \succ \mathcal{U} \succ c \succ d \succ \cdots$ where the alternative d is placed at $\kappa + 1$ position from the left.

For ease of reference, we call the above four groups as $\mathcal{G}_1, \mathcal{G}_2, \mathcal{G}_3$, and $\mathcal{G}_4$ respectively. Let n be the number of preferences in $\mathcal{P}$. We observe that $n = poly(m)$ since $\alpha_{\mathfrak{p}-1} = poly(m)$. The function $\mathfrak{f}$ is defined as follows: $\mathfrak{f}(\{d, x\}) = n$ for every $x \in \mathcal{A} \setminus (\{c, d\}), \mathfrak{f}(\{d, c\}) = n - \ell$; the value of $\mathfrak{f}$ be 0 for all other pairs of alternatives. This finishes the description of our reduced instance. We now claim that the two instances are equivalent.

In one direction, let us assume that the SET COVER FREQUENCY TWO instance is a YES instance; without loss of generality, let us assume (by renaming) that $S_1, \ldots, S_\ell$ forms a set cover of $\mathcal{U}$. Let us consider the following partial profile $\mathcal{Q}$ with $\mathcal{P} \in \mathcal{L}(\mathcal{Q})$.

- ▷ Preferences in $\mathcal{G}_1$: $\forall i \in [\ell] : ((W_{\mathfrak{p}-3} \cup S_i) \succ d \succ ((\mathcal{U} \setminus S_i) \cup (W \setminus W_{\mathfrak{p}-3})) \bigcup c \succ ((\mathcal{U} \setminus S_i)) \cup ((W \setminus W_{\mathfrak{p}-3})))$
- ▷ Preferences in $\mathcal{G}_1$: $\forall i$ with $\ell + 1 \leqslant i \leqslant t : (W_{\mathfrak{p}-3} \cup S_i) \succ d \succ c \succ ((\mathcal{U} \setminus S_i) \cup (W \setminus W_{\mathfrak{p}-3})$
- ▷ Preferences in $\mathcal{G}_2$: $t - 2$ copies of $(W_{\mathfrak{p}-1} \cup a_i) \succ d \succ ((\mathcal{U} \setminus \{a_i\}) \cup (W \setminus W_{\mathfrak{p}-1}) \cup \{c\})$
- ▷ Preferences in $\mathcal{G}_3$: $t(\alpha_{\mathfrak{p}} - \alpha_{\mathfrak{p}+2})$ copies of $c \succ X$ where the alternative $X = \{b \in \mathcal{A} : c \succ b \text{ in } \mathcal{G}_3\}$.
- ▷ Preferences in $\mathcal{G}_4$: for every preference in $\mathcal{G}_4$, we add $c \succ Y$ where the alternative $Y = \{b \in \mathcal{A} : c \succ b \text{ in the corresponding preference in } \mathcal{G}_4\}$.

Let Δ be the score that the alternative w_1 receives in $\mathcal{P}$. We observe that the minimum scores that the alternatives c and $a_i, i \in [q]$ receive in profile $\mathcal{R}$ with $\mathcal{R} \in \mathcal{L}(\mathcal{Q})$ are all the same; let it be λ. We summarize the MR score (based on r) of every alternative from $\mathcal{Q}$ in Table 1. Hence the alternative c wins uniquely in $\mathcal{Q}$.

Alternative	MR-r score from $\mathcal{Q}$	Competing alternative
c	$\Delta - t\alpha_\ell - \lambda$	w_1 (or w_2)
$a_i, \forall i \in [q]$	$\Delta - t\alpha_\ell + \alpha_{\ell+1} - \lambda$	w_1 (or w_2)
$w_1(w_2)$	Δ	$w_2(w_1)$
$w \in W \setminus W_2$	Δ	w_1
d	$> D - t\alpha_\ell$	w_1 (or w_2)

Table 1: Summary of MR scores (based on r) of all the alternatives from the partial profile $\mathcal{Q}$ in the proof of Theorem 2.

In the other direction, let us assume that the MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT instance $(\mathcal{A}, \mathcal{P}, c, \mathfrak{f})$ is a YES instance. Let $\mathcal{Q}$ be a partial profile such that $\mathcal{P} \in \mathcal{C}(\mathcal{Q})$ and the alternative c wins uniquely in $\mathcal{Q}$ under the MR scoring rule based on r. We observe that if a preference profile $\mathcal{R}_c$ with $\mathcal{R}_c \in \mathcal{L}(\mathcal{Q})$ is used to calculate the MR score of the alternative based on r, then the MR score of c based on r is at least $\Delta - t\alpha_\ell - \lambda$ using the alternative w_1 as a competing alternative where λ and Δ are as defined above. Let $J \subseteq [t]$ be the set of $i \in [t]$ such that the corresponding partial preferences in the group $\mathcal{G}_1$ in $\mathcal{Q}$ leave the alternatives c and d incomparable. Since $\mathfrak{f}(\{d, c\}) = n - \ell$, we have $|J| \leqslant \ell$. We claim that $\{S_j : j \in J\}$ forms a set cover of $\mathcal{U}$. Suppose not, then let $u_k \in \mathcal{U} \setminus (\cup_{j\in J} S_j)$. Then we observe that the MR score of a_k based on r is at least $\Delta - t\alpha_\ell - \lambda$ using the alternative w_1 as a competing alternative.However, this contradicts our assumption that c is the unique MR-r winner of $\mathcal{Q}$. Hence $\{S_j : j \in J\}$ forms a set cover of $\mathcal{U}$ and thus the SET COVER FREQUENCY TWO is a YES instance. This concludes the proof of the theorem. □

Theorem 2 immediately gives us the following corollary.

Corollary 2. *The* MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT *problem is* NP*-complete for the Borda and* k*-approval voting rules for every* $3 \leqslant k \leqslant \gamma m$ *for any constant* $0 < \gamma < 1$*.*

A drawback of Theorem 2 is that it does not cover the plurality, 2-approval, and the k-veto voting rules for $k = o(m)$. We will show that the MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT problem is NP-complete for the k-veto voting rule for any $1 \leqslant k \leqslant \gamma m$ for any constant $0 < \gamma < 1$ in Theorem 4. We now show in Theorem 3 that the MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT problem is NP-complete for the plurality and 2-approval voting rules by reducing it from the X3C problem which is defined as follows and known to be NP-complete (Garey and Johnson 1979).

Definition 5 (X3C). *Given a universe* $\mathcal{U}$ *of cardinality* q *such that* q *is divisible by* 3*, a family* $\mathcal{S} = \{S_i : i \in [t]\}$ *of* t *subsets of* $\mathcal{U}$ *each of cardinality* 3*, compute if there exists a subset* $\mathcal{G} \subseteq \mathcal{F}$ *of* $q/3$ *sets such that* $\cup_{A \in \mathcal{G}} A = \mathcal{U}$*. We denote an arbitrary instance of X3C by* $(\mathcal{U}, \mathcal{S})$*.*

Theorem 3. *The* MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT *problem is* NP*-complete for the plurality and the* 2*-approval voting rules.*

Proof. Let us first consider the plurality voting rule. The MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT problem for the plurality voting rule clearly belongs to NP. To prove NP-hardness, we reduce from X3C to MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT for the plurality voting rule. Let $(\mathcal{U} = \{u_1, \ldots, u_q\}, \mathcal{S} = \{S_i : i \in [t]\})$ be an arbitrary instance of X3C. For every $i \in [q]$ let us define $f_i = |\{j \in [t] : u_i \in S_j\}|$. Let us assume, without loss of generality, that $f_i < t - q/2$ (if not, then we add $3t$ new elements in $\mathcal{U}$ and t sets in $\mathcal{S}$ each of size 3 and collectively covering these new $3t$ elements). Let us assume, without loss of generality, that q is divisible by 6; if not then we add 3 new elements in $\mathcal{U}$ and a set consisting of these three new elements in $\mathcal{S}$. Let us consider the following instance $(\mathcal{A}, \mathcal{P}, c, \mathfrak{f})$ of MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT where $\mathcal{A}$ is defined as follows.

$$\mathcal{A} = \{a_i : i \in [q]\} \cup \{c, d, w\}$$

The profile $\mathcal{P}$ consists of the following preferences. For $X \subseteq \mathcal{U}$, let us also denote, for the sake of simplicity of notation, the set $\{a_j : u_j \in X\}$ of alternatives by X.

- ▷ $\forall i \in [t] : d \succ S_i \succ c \succ (\mathcal{U} \setminus S_i) \succ w$
- ▷ $q/6 + 1$ copies of $c \succ (\mathcal{A} \setminus \{c, w\})$
- ▷ 1 copy of $d \succ w \succ c \succ (\mathcal{A} \setminus \{d, w\})$

For ease of reference, we call the above three groups as $\mathcal{G}_1, \mathcal{G}_2$, and $\mathcal{G}_3$ respectively. Let n be the number of preferences in $\mathcal{P}$. That is, $n = t + q/6 + 2$. The function $\mathfrak{f}$ is defined as follows: $\mathfrak{f}(\{c, a_i\}) = n - f_i + 1$ for every $i \in [q]$, $\mathfrak{f}(\{w, d\}) = n - 1$, $\mathfrak{f}(\{w, x\}) = n$ for every $x \in \mathcal{A} \setminus \{w, d\}$; the value of $\mathfrak{f}$ be 0 for all other pairs of alternatives. This finishes the description of our reduced instance. We now claim that the two instances are equivalent.

In one direction, let us assume that the X3C instance is a YES instance; without loss of generality, let us assume (by renaming) that $S_1, \ldots, S_{q/3}$ forms a set cover of $\mathcal{U}$. Let us consider the following partial profile $\mathcal{Q}$ with $\mathcal{P} \in \mathcal{L}(\mathcal{Q})$.

- ▷ Preferences in $\mathcal{G}_1$: $\forall i \in [q/3] : (S_i \succ c \succ (\mathcal{U} \setminus S_i) \succ w) \bigcup (d \succ w)$
- ▷ Preferences in $\mathcal{G}_1$: $\forall i$ with $q/3 + 1 \leqslant i \leqslant t : (c \succ (\mathcal{U} \setminus S_i) \succ w) \bigcup (d \succ w)$
- ▷ Preferences in $\mathcal{G}_2$: $q/6 + 1$ copies of $c \succ (\mathcal{A} \setminus \{c, w\})$
- ▷ Preferences in $\mathcal{G}_3$: 1 copy of $w \succ c \succ (\mathcal{A} \setminus \{d, w\})$

We summarize the MR-plurality score of every alternative from $\mathcal{Q}$ in Table 2. Hence the alternative c wins uniquely in $\mathcal{Q}$.

Alternative	MR-plurality score from $\mathcal{Q}$	Competing alternative
c	$t - q/6$	d
$a_i, \forall i \in [q]$	t	d
d	$t - q/6 + 1$	c

Table 2: Summary of MR-plurality scores of all the alternatives from the partial profile $\mathcal{Q}$ in the proof of Theorem 3.

In the other direction, let us assume that the MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT instance $(\mathcal{A}, \mathcal{P}, c, \mathfrak{f})$ is a YES instance. Let $\mathcal{Q}$ be a partial profile such that $\mathcal{P} \in \mathcal{C}(\mathcal{Q})$ and the alternative c wins uniquely in $\mathcal{Q}$ under the MR-plurality voting rule. Let $J \subseteq [t]$ be the set of $i \in [t]$ such that the corresponding partial preferences in the group $\mathcal{G}_1$ in $\mathcal{Q}$ leave the alternatives c and at least one alternative in S_i incomparable. A key observation is that since $\mathfrak{f}(\{c, a_i\}) = n - f_i + 1$ for $i \in [q]$, we have $\cup_{j \in J} S_j = \mathcal{U}$. Hence we have $|J| \geqslant q/3$. We now claim that $|J| \leqslant q/3$. Suppose not, then the MR-plurality score of d is at most $(t - q/3 - 1) + q/6 + 1 = t - q/6$ using c as competing alternative (we observe that, since $\mathfrak{f}(\{c, a_i\}) = n - f_i + 1$ for $i \in [q]$, using the alternative a_i as a competing alternative for any i will lead to MR-plurality score of d at most $2f_i < t - q/6$). Hence the MR-plurality score of d is at most $t - q/6$. However the MR-plurality score of c is at least $t + 1 - (q/6 + 1) = t - q/6$. This contradicts our assumption that c is the unique MR-plurality winner of $\mathcal{Q}$. Hence $\{S_j : j \in J\}$ forms a set cover of $\mathcal{U}$ and thus the X3C is a YES instance. This concludes the proof of the theorem.

For the 2-approval voting rule, we can introduce n dummy alternatives each of which appears at the first position in exactly one preference and in the rest $(n - 1)$ preferences, it appears in the bottom $(n - 1)$ positions. All other parameters of the reduction remain same. It is easy to see that a similar argument will prove the result for the 2-approval voting rule. □

We now show our hardness result for the k-veto voting rule by reducing from X3C.

Theorem 4. *The* MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT *problem is* NP*-complete for the* *k-veto voting rule for every* $1 \leqslant k \leqslant \gamma m$ *for any constant* $0 < \gamma < 1$.

We now show our hardness result for the maximin voting rule by reducing from SET COVER FREQUENCY TWO.

Theorem 5. *The* MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT *problem is* NP*-complete for the maximin voting rule.*

Proof. The MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT problem clearly belongs to NP. To prove NP-hardness, we reduce from SET COVER FREQUENCY TWO to MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT for the maximin voting rule. Let $(\mathcal{U} = \{u_1, \ldots, u_q\}, \mathcal{S} = \{S_i : i \in [t]\}, \ell)$ be an arbitrary instance of SET COVER FREQUENCY TWO. Let us consider the following instance $(\mathcal{A}, \mathcal{P}, c, \mathfrak{f})$ of MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT where $\mathcal{A}$ is defined as follows.

$$\mathcal{A} = \{a_i : i \in [q]\} \cup \{c, w_1, w_2, d\}$$

The profile $\mathcal{P}$ consists of the following preferences. For $X \subseteq \mathcal{U}$, let us also denote, for the sake of simplicity of notation, the set $\{a_j : u_j \in X\}$ of alternatives by X.

▷ $\forall i \in [t] : w_1 \succ w_2 \succ S_i \succ d \succ c \succ (\mathcal{U} \setminus S_i)$

▷ 2 copies of $c \succ \mathcal{U} \succ d \succ w_1 \succ w_2$

For ease of reference, we call the above two groups as $\mathcal{G}_1$ and $\mathcal{G}_2$ respectively. Let n be the number of preferences in $\mathcal{P}$. The function $\mathfrak{f}$ is defined as follows: $\mathfrak{f}(\{d, x\}) = n$ for every $x \in \mathcal{A} \setminus (\{c, d\}), \mathfrak{f}(\{d, c\}) = n - \ell$; the value of $\mathfrak{f}$ be 0 for all other pairs of alternatives. This finishes the description of our reduced instance. We now claim that the two instances are equivalent.

In one direction, let us assume that the SET COVER FREQUENCY TWO instance is a YES instance; without loss of generality, let us assume (by renaming) that $S_1, \ldots, S_\ell$ forms a set cover of $\mathcal{U}$. Let us consider the following partial profile $\mathcal{Q}$ with $\mathcal{P} \in \mathcal{L}(\mathcal{Q})$.

▷ Preferences in $\mathcal{G}_1$: $\forall i \in [\ell] : ((\{w_1, w_2\} \cup S_i) \succ d \succ \{\mathcal{U} \setminus S_i\}) \bigcup (c \succ \{\mathcal{U} \setminus S_i\})$

▷ Preferences in $\mathcal{G}_1$: $\forall i$ with $\ell + 1 \leqslant i \leqslant t : (\{w_1, w_2\} \cup S_i) \succ d \succ c \succ \{\mathcal{U} \setminus S_i\}$

▷ Preferences in $\mathcal{G}_2$: 2 copies of $c \succ \mathcal{U} \succ d \succ \{w_1, w_2\}$

We summarize the MR-maximin score of every alternative from $\mathcal{Q}$ in Table 3. Hence the alternative c wins uniquely in $\mathcal{Q}$.

In the other direction, let us assume that the MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT instance $(\mathcal{A}, \mathcal{P}, c, \mathfrak{f})$ is a YES instance. Let $\mathcal{Q}$ be a partial profile such that $\mathcal{P} \in \mathcal{C}(\mathcal{Q})$ and the alternative c wins uniquely in $\mathcal{Q}$ under the MR-maximin voting rule. We observe that for every $\mathcal{R} \in \mathcal{L}(\mathcal{Q})$ which can be used for calculating the MR-maximin score of the alternative c, we have $N_{\mathcal{R}}(c, d) \leqslant 2$. Also, there are only two preferences (the preferences in $\mathcal{G}_2$) where there exist some alternatives which are preferred over the alternative w_1. Hence the MR-maximin score of the alternative c in $\mathcal{Q}$ is at least $t - 2$. Let $J \subseteq [t]$ be the set of $i \in [t]$ such that the corresponding partial preferences in the group $\mathcal{G}_1$ in $\mathcal{Q}$ leave the alternatives c and d incomparable. Since $\mathfrak{f}(\{d, c\}) = n - \ell$, we have $|J| \leqslant \ell$. We claim that $\{S_j : j \in J\}$ forms a set cover of $\mathcal{U}$. Suppose not, then let $u_k \in \mathcal{U} \setminus (\cup_{j \in J} S_j)$. We observe that for every $\mathcal{R}' \in \mathcal{L}(\mathcal{Q})$, we have $N_{\mathcal{R}'}(a_k, c) = 2$. We also observe that $N_{\mathcal{R}'}(a_k, d) = N_{\mathcal{R}'}(a_k, w_i) = 2$ for every $i \in [2]$ since $\mathfrak{f}(\{d, u_i\}) = n$ for every $i \in [q]$ and $\mathfrak{f}(\{d, w_1\}) = \mathfrak{f}(\{d, w_2\}) = n$. Hence, the MR-maximin score of the alternative a_k is $t - 2$ where the alternative w_1 plays the role of a competing alternative. However, this contradicts our assumption that c is the unique MR-maximin winner of $\mathcal{Q}$. Hence $\{S_j : j \in J\}$ forms a set cover of $\mathcal{U}$ and thus the SET COVER FREQUENCY TWO is a YES instance. This concludes the proof of the theorem. □

Alternative	MR-maximin score from $\mathcal{Q}$	Comments
c	$t-2$	$N(w_1, c) - N(c, d)$
$a_i, \forall i \in [q]$	$t-1$	$N(w_1, c) - N(a_i, c)$
$w_1(w_2)$	t	$N(w_1, c) - N(w_2, w_1)$
d	t	$N(w_1, c) - N(d, w_1)$

Table 3: Summary of MR-maximin scores of all the alternatives from the partial profile $\mathcal{Q}$ in the proof of Theorem 5.

For the Copeland$^\alpha$ voting rule, we have the following result for every $\alpha \in [0, 1]$.

Theorem 6. *The* MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT *problem is* NP*-complete for the Copeland*$^\alpha$ *voting rule for every* $\alpha \in [0, 1]$.

For the simplified Bucklin voting rule, we have the following result.

Theorem 7. *The* MANIPULATIVE ELICITATION WITH CANDIDATE PAIR LIMIT *problem is* NP*-complete for the simplified Bucklin voting rule.*

5 Conclusion and Future Work

In this work, we have discovered an important vulnerability, namely manipulative elicitation, in the use of minimax regret based extension of classical voting rules in the incomplete preferential setting. Moreover, we have shown that the related computational task is polynomial time solvable for many commonly used voting rules including all scoring rules, maximin, Copeland$^\alpha$ for every $\alpha \in [0, 1]$, simplified Bucklin voting rules, etc. Then we have shown that by introducing a parameter per pair of alternatives which specifies the minimum number of partial preferences where this pair of alternatives must be comparable makes the computational task of manipulative elicitation NP-complete for all the above mentioned voting rules.

A drawback of our approach is that the parameters can be non-uniform – their values do not need to be the same for every pair of alternatives. It would be interesting to study the computational complexity of the problem when the values of the parameters are all the same. In another direction, it would be interesting to conduct extensive experimentation to study usefulness of our approach in practice. This is specially important since computational intractability is known to provide only a weak barrier in other forms of election manipulation (Procaccia and Rosenschein 2007).

References

Baumeister, D.; Roos, M.; and Rothe, J. 2011. Computational complexity of two variants of the possible winner problem. In *Proc. 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 853–860.

Betzler, N.; Bredereck, R.; and Niedermeier, R. 2010. Partial kernelization for rank aggregation: theory and experiments. In *Proc. International Symposium on Parameterized and Exact Computation (IPEC)*. Springer. 26–37.

Betzler, N.; Bredereck, R.; and Niedermeier, R. 2014. Theoretical and empirical evaluation of data reduction for exact kemeny rank aggregation. *Auton. Agent Multi Agent Syst.* 28(5):721–748.

Betzler, N.; Hemmann, S.; and Niedermeier, R. 2009. A Multivariate Complexity Analysis of Determining Possible Winners given Incomplete Votes. In *Proc. International Joint Conference on Artificial Intelligence (IJCAI)*, volume 9, 53–58.

Chevaleyre, Y.; Lang, J.; Maudet, N.; and Monnot, J. 2010. Possible winners when new candidates are added: The case of scoring rules. In *Proc. International Conference on Artificial Intelligence (AAAI)*, 762–767.

Dey, P., and Misra, N. 2017. On the exact amount of missing information that makes finding possible winners hard. In *Proc. 42nd International Symposium on Mathematical Foundations of Computer Science*.

Dey, P.; Misra, N.; and Narahari, Y. 2015. Detecting possible manipulators in elections. In *Proc. 2015 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2015, Istanbul, Turkey, May 4-8, 2015*, 1441–1450.

Dey, P.; Misra, N.; and Narahari, Y. 2016a. Complexity of manipulation with partial information in voting. In *Proc. Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, 229–235.

Dey, P.; Misra, N.; and Narahari, Y. 2016b. Kernelization complexity of possible winner and coalitional manipulation problems in voting. *Theor. Comput. Sci.* 616:111–125.

Dey, P.; Misra, N.; and Narahari, Y. 2017. Frugal bribery in voting. *Theor. Comput. Sci.* 676:15–32.

Dey, P. 2017. Manipulative elicitation – a new attack on elections with incomplete preferences. *CoRR* abs/1711.03948.

Fagin, R.; Kumar, R.; and Sivakumar, D. 2003. Efficient similarity search and classification via rank aggregation. In *Proc. 2003 ACM SIGMOD International Conference on Management of Data*, SIGMOD '03, 301–312. New York, NY, USA: ACM.

Faliszewski, P.; Reisch, Y.; Rothe, J.; and Schend, L. 2014. Complexity of manipulation, bribery, and campaign management in bucklin and fallback voting. In *Proc. 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 1357–1358. International Foundation for Autonomous Agents and Multiagent Systems.

Garey, M. R., and Johnson, D. S. 1979. *Computers and Intractability*, volume 174. freeman New York.

Konczak, K., and Lang, J. 2005. Voting procedures with incomplete preferences. In *Proc. International Joint Conference on Artificial Intelligence-05 Multidisciplinary Workshop on Advances in Preference Handling*, volume 20.

Lang, J.; Pini, M. S.; Rossi, F.; Venable, K. B.; and Walsh, T. 2007. Winner determination in sequential majority voting. In *Proc. 20th International Joint Conference on Artificial Intelligence (IJCAI)*, volume 7, 1372–1377.

Lang, J.; Pini, M. S.; Rossi, F.; Salvagnin, D.; Venable, K. B.; and Walsh, T. 2012. Winner determination in voting trees with incomplete preferences and weighted votes. *Auton. Agent Multi Agent Syst.* 25(1):130–157.

Lu, T., and Boutilier, C. 2011. Robust approximation and incremental elicitation in voting protocols. In *Proc. 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, 287–293.

Moulin, H.; Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D. 2016. *Handbook of Computational Social Choice*. Cambridge University Press.

Pennock, D. M.; Horvitz, E.; and Giles, C. L. 2000. Social choice theory and recommender systems: Analysis of the axiomatic foundations of collaborative filtering. In *Proc. Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on on Innovative Applications of Artificial Intelligence, July 30 - August 3, 2000, Austin, Texas, USA.*, 729–734.

Pini, M. S.; Rossi, F.; Venable, K. B.; and Walsh, T. 2007. Incompleteness and incomparability in preference aggregation. In *Proc. 20nd International Joint Conference on Artificial Intelligence (IJCAI)*, volume 7, 1464–1469.

Procaccia, A. D., and Rosenschein, J. S. 2007. Junta distributions and the average-case complexity of manipulating elections. *J. Artif. Intell. Res.* 28:157–181.

Walsh, T. 2007. Uncertainty in preference elicitation and aggregation. In *Proc. International Conference on Artificial Intelligence (AAAI)*, volume 22, 3–8.

Xia, L., and Conitzer, V. 2011. Determining possible and necessary winners under common voting rules given partial orders. *J. Artif. Intell. Res.* 41(2):25–67.