Average-Case Tractability of Manipulation in Voting via the Fraction of Manipulators

Ariel D. Procaccia Jeffrey S. Rosenschein
School of Engineering and Computer Science
The Hebrew University of Jerusalem, Israel
{arielpro,jeff}@cs.huji.ac.il

ABSTRACT

Recent results have established that a variety of voting rules are computationally hard to manipulate in the worst-case; this arguably provides some guarantee of resistance to manipulation when the voters have bounded computational power. Nevertheless, it has become apparent that a truly dependable obstacle to manipulation can only be provided by voting rules that are *average-case* hard to manipulate.

In this paper, we analytically demonstrate that, with respect to a wide range of distributions over votes, the coalitional manipulation problem can be decided with overwhelming probability of success by simply considering the ratio between the number of truthful and untruthful voters. Our results can be employed to significantly focus the search for that elusive average-case-hard-to-manipulate voting rule, but at the same time these results also strengthen the case against the existence of such a rule.

Categories and Subject Descriptors

F.2 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity;
I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent Systems*;
J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms

Algorithms, Theory, Economics

Keywords

Computational complexity, Voting

1. INTRODUCTION

Voting is often used as a method of aggregating the preferences of heterogeneous, self-interested agents. Unfortunately, for a socially desirable outcome to emerge from an election, voters should reveal their true preferences—but precluding manipulation is impossible in general. The celebrated Gibbard-Satterthwaite Theorem [2] states that, with any reasonable voting rule (a function that determines the outcome of the election, given the voters' preferences), there are elections where some of the voters can benefit by voting untruthfully.

Computational complexity theory seemingly provides a way to circumvent the Gibbard-Satterthwaite Theorem. It has been suggested that, although in principle a voter may lie in order to improve its position, determining if it is possible in practice, given a specific setting, may be a computationally hard problem. Recent results imply that indeed manipulation is often $\mathcal{NP}$-hard. Nevertheless, it can be argued that *worst-case* hardness, although demonstrating some measure of resistance to manipulation, does not preclude it. Therefore, an *average-case* analysis is required.

Naturally, it is not possible to find a voting rule that is usually hard to manipulate with respect to any distribution over the instances. However, on the face of it, it is reasonable to hope for a voting rule that has this property at least under certain interesting distributions. Sadly, two recent papers presented evidence to the contrary [3, 1]. In this paper, we pursue the abovementioned line of research by establishing more results about the average-case tractability of manipulations.

2. PRELIMINARIES

An election consists of a set $V = \{v^1, v^2, \ldots\}$ of voters, and a set $C = \{c_1, c_2, \ldots, c_m\}$ of candidates; voters' indices usually appear in superscript, while candidates' indices usually appear in subscript. Each voter's preferences can be represented as a linear order;[1] let $\mathcal{L}$ be the set of linear orders on C. A *voting rule* is a function $F : \mathcal{L}^V \to C$, that maps the preferences of the voters to the winning candidate.

The voting rules we shall discuss in this paper are scoring rules. A scoring rule is defined by a vector $\alpha = \langle \alpha_1, \ldots, \alpha_m \rangle$, where the α_l are real numbers such that $\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_m \geq 0$. Each voter awards α_1 points to the candidate it ranks first, α_2 points to the candidate it ranks second, and in general α_l points to the candidate it ranks l'th. The candidate with the most points (summed over all the voters) wins the election. Some prominent scoring rules are:

- *Plurality*: $\vec{\alpha} = \langle 1, 0, \ldots, 0 \rangle$.
- *Borda*: $\vec{\alpha} = \langle m-1, m-2, \ldots, 0 \rangle$.

[1] A linear order is a binary relation that satisfies antisymmetry, transitivity, and totality.

AAMAS'07 May 14–18 2007, Honolulu, Hawai'i, USA.

- *Veto*: $\vec{\alpha} = \langle 1, \ldots, 1, 0 \rangle$.

Hereinafter, we conceptually partition the set V of voters into two sets: $V = V_1 \uplus V_2$, where $|V_1| = N$ and $|V_2| = n$. $V_1 = \{v^1, \ldots, v^N\}$ is the set of nonmanipulators, while $V_2 = \{v^{N+1}, \ldots, v^{N+n}\}$ is the set of manipulators, who are colluding in an attempt to make a certain candidate p win the election.

Definition 1. In the Coalitional-Weighted-Manipulation (CWM) problem, we are given the set of voters $V = V_1 \uplus V_2$, the set of candidates C, the weights of all voters, and a preferred candidate $p \in C$. In addition, we are given the votes of the voters in V_1, and assume that the manipulators are aware of these votes. We ask whether it is possible for the manipulators in V_2 to cast their votes in a way that makes the preferred candidate p win the election.

3. RESULTS

We wish to study the relationship between the number of nonmanipulators versus manipulators (or, if you will, the fraction of manipulators relative to overall voters) and the chances that an instance of CWM is a "yes" or a "no" instance. Essentially, we suggest that in many cases the manipulators cannot affect the outcome of the election at all.

Definition 2. An instance of CWM is a *closed* instance if, no matter how the manipulators in V_2 cast their votes, the same candidate gets elected. An instance that is not a closed instance is called an *open* instance.

3.1 Fraction of Manipulators is Small

In this subsection we demonstrate that when the fraction of manipulators is small, that is $n = o(\sqrt{N})$, then usually instances of CWM are closed. This result holds for scoring rules, and requires only weak assumptions on the distribution of votes.

Lemma 1. *Consider an instance of the coalitional manipulation problem in a scoring rule with parameters $\vec{\alpha}$. Let S_k^i be the score given to candidate c_k by the voter v^i. If there exists a candidate c_k such that for all $c_l \neq c_k$, $\sum_{i=1}^N S_k^i - \sum_{i=1}^N S_l^i > \alpha_1 n$, then the instance is closed.*

Let D^i be a distribution over voter v^i's votes, $1 \leq i \leq N$; denote the joint distribution over votes by $D^N = \prod_{i=1}^N D^i$. D^i induces a random variable S_k^i, which determines the points voter v^i awards candidate c_k.

Theorem 1. *Let P be a scoring rule with parameters $\vec{\alpha}$, and assume that the number of manipulators and nonmanipulators satisfies:*

- $n = o(\sqrt{N})$.

Let D^i be voter i's distribution over the possible votes with $m = O(1)$ candidates, and denote $D^N = \prod_{i=1}^N D^i$. Let S_k^i, for each $v^i \in V_1$ and $c_k \in C$, be random variables, induced by the D^i, which determine the score of candidate c_k from voter v^i. Assume that the distributions over votes satisfy:

- ***(d1)*** *There exists a constant $d > 0$ such that for all $v^i \in V_1$ and $c_k, c_l \in C$, $d <$ Var$[S_k^i - S_l^i]$.*
- ***(d2)*** *The D^i are independently distributed.*

Then the probability that an instance is closed converges to 1 as the number of voters grows.

Proof of Theorem 1. By Lemma 1 we have:

$$\begin{aligned}
&\Pr_{D^N}[\text{instance is closed}] \\
&\geq \Pr_{D^N}\left[\exists c_k \in C, \forall c_l \neq c_k,\ \sum_{i=1}^N S_k^i - \sum_{i=1}^N S_l^i > \alpha_1 n\right] \\
&\geq \Pr_{D^N}\left[\forall c_k, c_l \neq c_k, |\sum_{i=1}^N S_k^i - \sum_{i=1}^N S_l^i| > \alpha_1 n\right] \\
&= 1 - \Pr_{D^N}\left[\exists c_k, c_l \in C\ s.t.\ 0 \leq \sum_{i=1}^N S_k^i - \sum_{i=1}^N S_l^i \leq \alpha_1 n\right].
\end{aligned}$$

Now, by the union bound, we have that

$$\begin{aligned}
&\Pr_{D^N}\left[\exists c_k, c_l \in C\ s.t.\ 0 \leq \sum_{i=1}^N S_k^i - \sum_{i=1}^N S_l^i \leq \alpha_1 n\right] \\
&\quad \leq \sum_{c_k, c_l \in C} \Pr_{D^N}\left[0 \leq \sum_{i=1}^N S_k^i - \sum_{i=1}^N S_l^i \leq \alpha_1 n\right].
\end{aligned} \tag{1}$$

Fix two candidates $c_k, c_l \in C$, and denote $X^i = S_k^i - S_l^i$. Let $\mu^i = \mathrm{E}[X^i]$, $\sigma^i = \mathrm{Var}[X^i]$. Notice that $\sum_{i=1}^N S_k^i - \sum_{i=1}^N S_l^i = \sum_{i=1}^N X^i$. In addition, observe that by assumption (d1) $d < \sigma^i$, and thus $\sum_{i=1}^N \sigma^i \overset{N\to\infty}{\longrightarrow} \infty$. In addition, for all $v^i \in V$, $|X^i| \leq \alpha_1$. Therefore, we may apply the Central Limit Theorem to the variables X^i.

$$\begin{aligned}
&\Pr_{D^N}\left[0 \leq \sum_{i=1}^N X^i \leq \alpha_1 n\right] \\
&= \Pr_{D^N}\left[\frac{-\sum_{i=1}^N \mu^i}{\sqrt{\sum_{i-1}^N \sigma^i}} \leq \frac{\sum_{i=1}^N X^i - \sum_{i=1}^N \mu^i}{\sqrt{\sum_{i-1}^N \sigma^i}} \right.\\
&\qquad\qquad \left. \leq \frac{\alpha_1 n - \sum_{i=1}^N \mu^i}{\sqrt{\sum_{i-1}^N \sigma^i}}\right] \\
&\overset{N\to\infty}{\longrightarrow} \frac{1}{\sqrt{2\pi}} \int_{\frac{-\sum_{i=1}^N \mu^i}{\sqrt{\sum_{i-1}^N \sigma^i}}}^{\frac{\alpha_1 n - \sum_{i=1}^N \mu^i}{\sqrt{\sum_{i-1}^N \sigma^i}}} e^{-\frac{x^2}{2}} dx \\
&\leq \int_{\frac{-\sum_{i=1}^N \mu^i}{\sqrt{\sum_{i-1}^N \sigma^i}}}^{\frac{\alpha_1 n - \sum_{i=1}^N \mu^i}{\sqrt{\sum_{i-1}^N \sigma^i}}} 1\, dx \\
&= \frac{\alpha_1 n}{\sqrt{\sum_{i-1}^N \sigma^i}} \leq \frac{\alpha_1 n}{\sqrt{dN}} = O\left(\frac{n}{\sqrt{N}}\right).
\end{aligned}$$

Plugging this result into Equation (1), we have that

$$\begin{aligned}
&\Pr_{D^N}\left[\exists c_k, c_l \in C\ s.t.\ 0 \leq \sum_{i=1}^N S_k^i - \sum_{i=1}^N S_l^i \leq \alpha_1 n\right] \\
&\leq m(m-1) \cdot O\left(\frac{n}{\sqrt{N}}\right) \\
&= O\left(\frac{n}{\sqrt{N}}\right),
\end{aligned}$$

where the second transition follows from the fact that m is constant. Rolling back, we have that

$$\Pr_{D^N}[\text{instance is open}] \geq 1 - O\left(\frac{n}{\sqrt{N}}\right).$$

Under the assumption that $n = o(\sqrt{N})$, this expression converges to 1 as the number of voters grows. □

3.2 Fraction of Manipulators is Large

In this subsection, we tackle a setting where the number of manipulators is large, i.e., $n = \omega(\sqrt{N})$, but not excessively so, i.e., $n = o(N)$. The mathematical techniques we use here differ from the ones applied in Section 3.1.

THEOREM 2. *Let P be a scoring rule with parameters $\vec{\alpha}$, and assume that the number of manipulators and nonmanipulators satisfies:*

- $n = \omega(\sqrt{N})$ *and* $n = o(N)$.

Let D^i be voter i's distribution over the possible votes with $m = O(1)$ candidates, and denote $D^N = \prod_{i=1}^N D^i$. Let S_k^i, for each $v^i \in V_1$ and $c_k \in C$, be random variables, induced by the D^i, which determine the score of candidate c_k from voter v^i. Assume that the distributions over votes satisfy:

- ***(d2)*** *The D^i are independently distributed*
- ***(d3)*** *The D^i are identically distributed.*

Let $C' = \{c_k \in C : \forall c_l \neq c_k,\ \mathbb{E}[S_k^1] \geq \mathbb{E}[S_l^1]\}$ be the subset of candidates with maximal expected score.

1. *If $|C'| \geq 2$, then the probability of drawing an open instance converges to 1 as the number of voters grows.*
2. *If $|C'| = 1$ then the probability of drawing a closed instance converges to 1 as the number of voters grows.*

COROLLARY 3. *Under the conditions of Theorem 2, if C' is the set of candidates with maximal expected score, then with probability that converges to 1 it holds that any candidate from C' can be made to win, and no other candidate can be made to win.*

4. DISCUSSION

Consider Algorithm 1, which instantly decides instances of the manipulation problem, drawn according to some distribution, on the basis of the ratio between the number of manipulators and nonmanipulators. Theorems 1 and 2 directly imply that for any distribution that satisfies assumptions (d1), (d2), and (d3), Algorithm 1 is almost never wrong when the number of voters is large. Indeed, when $n = o(\sqrt{N})$, Theorem 1 asserts that instances are almost always closed—and therefore p can be made to win iff p wins for any arbitrary vote of the manipulators. In case $n = \omega(\sqrt{N})$, Corollary 3 states that it is usually true that the manipulators can only make candidates with maximal expected score win the election.

But how restrictive are the assumptions (d1), (d2), and (d3)? Assumption (d1) requires that there exist a constant $d > 0$ such that for all $v^i \in V_1$ and $c_k, c_l \in C$, $d < \text{Var}[S_k^i - S_l^i]$. This is certainly a condition that seems very reasonable: the demand is that according to each voter's distribution, there are no two candidates that *always* have the same difference in scores. That is, we simply require a seemingly minimal element of randomness in the votes. Granted, requiring that the votes of the nonmanipulators be distributed i.i.d.—the union of assumptions (d2) and (d3)—is a much stricter assumption. Nevertheless, it is possible to provide interesting distributions that satisfy all three assumptions.

Algorithm 1 Deciding the coalitional manipulation problem in scoring rules via the fraction of manipulators. The input is a voting instance drawn according to a distribution over the votes of the nonmanipulators; p is the manipulators' preferred candidate.

```
if n = o(√N) then                                    ▷ Theorem 1
    choose arbitrary manipulators' vote; c is the winner
    if p = c then
        return true
    else
        return false
    end if
else if n = ω(√N) and n = o(N) then                  ▷ Theorem 2
    if p has maximal expected score then
        return true
    else
        return false
    end if
else                                                 ▷ n = Θ(√N) or n = Ω(N)
    return ?
end if
```

Although our results apply only to scoring rules, we believe that similar results hold for all other important voting rules. Preliminary results indicate that this is certainly true for the Copeland and Maximin rules—but the probabilistic properties behind the proofs can be found in other voting rules as well.

To conclude, there are two ways to interpret our results. A positive interpretation would be that an average-case-hard-to-manipulate voting rule and distribution exist, and the results may simply help focus the search for such a distribution. Interpreted negatively, these results strengthen the case against the existence of voting rules that are hard to manipulate. Indeed, they imply that the manipulation problem in many important voting rules can usually be trivially decided, with respect to a wide range of distributions.

5. ACKNOWLEDGMENT

This work was partially supported by Israel Science Foundation grant #898/05.

6. REFERENCES

[1] V. Conitzer and T. Sandholm. Nonexistence of voting rules that are usually hard to manipulate. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence*, 2006.

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

[3] A. D. Procaccia and J. S. Rosenschein. Junta distributions and the average-case complexity of manipulating elections. *Journal of Artificial Intelligence Research*, 27, 2006.