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

Recent years have seen a surge of interest in the computational aspects of social choice theory. This attention is motivated by applications of social choice paradigms to fields such as electronic commerce and multiagent systems. Voting, in particular, is often used as a method of aggregating the preferences of heterogeneous, self-interested agents. For instance, voting has been employed to help agents reach an agreement regarding joint plans, schedules [10], and recommended movies [9].

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 [13] 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. This result holds if there are three or more candidates, if the voting rule is non-dictatorial (i.e., sensitive to more than the wishes of a single voter), and if the voting rule could conceivably enable any one of the candidates to win, under some alignment of voter preferences.

Computational complexity theory seemingly provides a way to circumvent the Gibbard-Satterthwaite Theorem. It has been suggested by Bartholdi et al. [3] 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. Indeed, Bartholdi and Orlin [2] demonstrated that the important Single Transferable Vote (STV) voting rule is $\mathcal{NP}$-hard to manipulate. More recent results imply that it is possible to make simple voting rules, which could otherwise be manipulated efficiently, hard to manipulate by adding a preround [6]. Other research shows that many voting rules are $\mathcal{NP}$-hard to manipulate by a coalition of weighted voters [5, 4]; these results hold even when the number of candidates is constant, though the number of manipulative voters must be unbounded.

The abovementioned computational results all imply *worst-case* hardness. Although it can be argued that these results demonstrate some measure of resistance to manipulation, by no means do they preclude it. An $\mathcal{NP}$-hard problem has an infinite number of "hard instances" [11], but it may still be the case that most instances are "easy". In our setting, this means that a strategic voter might usually be able to determine whether or not to reveal its true preferences. Therefore, an *average-case* analysis is required.

Naturally, it is not possible for 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. Procaccia and Rosenschein [12] defined the notion of *junta distribution*, and argued that if an algorithm can usually decide the manipulation problem with respect to a junta distribution, the same algorithm would usually succeed with respect to other natural distributions. The authors proceeded to show that there is such an algorithm when the voting rule is a *scoring* rule (see Section 2 for a formal definition). However, the relation between junta distributions and other distributions is not as yet formally established.

Conitzer and Sandholm [7] proposed a different approach: they demonstrated that if an instance of the manipulation problem has certain properties, it is easy to decide. Further, it was *empirically* shown that, in a variety of voting rules, a large fraction of instances possess these properties. The main drawback of this result is that the experiments only consider instances that are distributed according to a specific family of distributions, and in all experiments the number of manipulators and candidates is extremely small compared to the number of voters. In addition, the paper provides few *theoretical* guarantees.

In the current paper we notice that, for a variety of interesting distributions, it is either almost always possible to find a successful manipulation, or almost never possible. The correct alternative depends only on easily testable properties of the distribution, and on the fraction of manipulators. If n is the number of manipulators and N is the number of nonmanipulative voters, we demonstrate that, when $n = o(\sqrt{N})$, manipulation is almost never possible under almost any distribution where the voters vote independently. When $n = \omega(\sqrt{N})$, we characterize the distributions where manipulation is almost always possible, and the ones where it is almost never possible. We rigorously prove these results in the context of an important family of voting rules, but argue that they are also valid with respect to others. Ultimately, our results yield a generic algorithm that usually decides the manipulation problem under many natural distributions. This can be viewed as additional evidence against the existence of a voting rule that is hard to manipulate, but can also be used to focus the search for such a rule, should it exist.

Some recent research in economics has independently recognized that when the fraction of manipulators is small, manipulation is rarely possible [14, 1]. However, these papers consider only variations on the uniform distribution over possible elections; this is plausible from the economist's point of view, but in computer science we can preclude average-case hardness only by showing that the problem is average-case tractable under *any* distribution. Additionally, unlike the abovementioned work, we present our results and their implications from a computational point of view. In particular, the formulation of the manipulation problem that we consider is the one generally accepted in computer science, and not the one investigated in that research.

The paper proceeds as follows. In Section 2 we give some background concerning voting rules and the coalitional manipulation problem. In Section 3 we present our results. In Section 4 we discuss the meaning and significance of our results, and we conclude in Section 5.

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 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.

2.1 Scoring Rules

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

Remark 1. Plurality is in fact the voting rule that governs most real-life elections.

It is also possible to consider weighted voting: voters are associated with weights, and a voter with weight k is taken into account as k voters casting identical votes. For example, if $m = 3$ and the voting rule is Borda, a voter with weight 2 awards 4 points to its first choice, 2 points to its second choice, and 0 points to its last choice.

2.2 The Coalitional Manipulation Problem

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. The standard assumption is that the manipulators know exactly how the nonmanipulators cast their votes.[2] In addition, we assume that m, the number of candidates, is constant: $m = O(1)$.

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

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

[2] In game-theoretic settings, it is usually assumed that strategic players know how the other players play.

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.

It is known that CWM is $\mathcal{NP}$-hard when the voting rule is Borda or Veto and the number of candidates satisfies $m \geq 3$, as well as when the voting rule is Copeland or Maximin and the number of candidates satisfies $m \geq 4$ [5, 4]. We shall argue below that CWM is average-case tractable in these rules (when the number of candidates is constant).

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 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.

Naturally, knowing whether an instance is closed goes a long way towards deciding CWM. For example, a closed instance is a "yes" instance if and only if the distinguished candidate in Definition 2 is the preferred candidate p.

Any result in average-case complexity clearly depends on the distribution over the possible instances of the manipulation problem. In fact, since we are mostly interested in whether an instance is open or closed, the only parameters of the problem that are not given are the votes of the nonmanipulators, and the weights of the voters. However, as in [7], we prove sufficient conditions for openness/closedness which depend only on the total weight of the manipulators — the individual weights of the manipulators are of no importance. Therefore, weights are a nonissue, and it is sufficient to consider distributions over the possible votes of the nonmanipulators in V_1.

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.

Given an instance of the manipulation problem in a scoring rule, consider the scores of candidates based only on the votes of the nonmanipulators. If there is a candidate whose score is higher than others' by more than $\alpha_1 n$, then the instance is surely closed: even if all manipulators ranked this candidate last and another candidate first, the difference in scores would decrease by at most $\alpha_1 n$, which is not enough to close the gap. Further, notice that the total score, based on the nonmanipulators' votes, of candidate c_k is given by $\sum_{i=1}^{N} S_k^i$. We have established the following sufficient condition for closedness:

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 $c_l \neq c_k$, $\sum_{i=1}^{N} S_k^i - \sum_{i=1}^{N} S_k^i > \alpha_1 n$, 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.

Example 1. Let the voting rule be Borda, and $m = 3$ — each voter awards 2 points to its first choice, 1 point to its second choice, and 0 points to its last. If D^i is the uniform distribution, then for $k = 1, 2, 3$, S_k^i is 2 with probability $\frac{1}{3}$, 1 with probability $\frac{1}{3}$, and 0 with probability $\frac{1}{3}$.

We are now ready to present our result.

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.

The proof relies heavily on the central limit theorem. For our purposes, this theorem implies that the probability that a sum of random variables obtains values in a very small segment is very small, as long as the variance of the random variables is nonzero.

Theorem 2 (Central Limit Theorem). *[8] Let*

$$X^1, X^2, \ldots, X^N, \ldots$$

be a sequence of independent discrete random variables. For each i, denote the mean and variance of X^i by μ^i and σ^i, *respectively, and assume that* $\sum_{i=1}^N \sigma^i \xrightarrow{N\to\infty} \infty$*, and that* $|X^i| \le A$ *for some constant* A *and all* i*. Then for* $a < b$*:*

$$\Pr\left[a < \frac{\sum_{i=1}^N X^i - \sum_{i=1}^N \mu^i}{\sqrt{\sum_{i=1}^N \sigma^i}} < b\right] \xrightarrow{N\to\infty} \frac{1}{\sqrt{2\pi}}\int_a^b e^{-\frac{x^2}{2}}\,dx.$$

Proof of Theorem 1. By Lemma 1 we have:

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

Now, by the union bound, we have that

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

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 $\sum_{i=1}^N \sigma^i \xrightarrow{N\to\infty} \infty$. In addition, for all $v^i \in V$, $|X^i| \le \alpha_1$. Therefore, we may apply Theorem 2 to the variables X^i.

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

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

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

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

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

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.

As before, we characterize instances of the manipulation problem in scoring rules. Notice that in the current setting, manipulators may often have enough power to sway the outcome of the election. Therefore, we require a sufficient condition for the openness of a manipulation instance.

Lemma 2. *Consider an instance of the coalitional manipulation problem in a scoring rule with parameters* $\vec{\alpha}$*, and assume* $n \ge m$*. Let* S_k^i *be the score given to candidate* c_k *by the voter* v^i*. Let* $C' \subseteq C$ *such that for any two candidates* $c_k, c_l \in C'$ *it holds that* $\sum_{i=1}^N S_k^i - \sum_{i=1}^N S_l^i < \frac{\alpha_1-\alpha_m}{2m}\cdot n$*, and for any* $c_k \in C'$ *and* $c_l \notin C'$*,* $\sum_{i=1}^N S_k^i - \sum_{i=1}^N S_l^i \ge 0$*. Then the manipulators can make any candidate in* C' *win.*

Proof. Assume w.l.o.g. that the manipulators wish to make candidate $c_m \in C'$ win; the manipulators vote as follows. The i'th manipulator, $i = 1, \ldots, n$, ranks c_m first, $c_{i \bmod (m-1)}$ last, and the other candidates in some arbitrary order. Each candidate other than c_m is ranked last by at least $\left\lfloor\frac{n}{m-1}\right\rfloor$ manipulators, and the rest of the manipulators award it at most α_1 points. Therefore, the difference in the points awarded by the manipulators to c_m and any other candidate is at least $\left\lfloor\frac{n}{m-1}\right\rfloor\cdot(\alpha_1-\alpha_m) \ge \frac{\alpha_1-\alpha_m}{2m}\cdot n$, where the inequality holds whenever $n \ge m$. □

Our theorems regarding the current setting are weaker than the ones in Section 3.1, in the sense that the voters' votes are (independent and) identically distributed. The following theorem differentiates two cases: if there are at least two candidates whose expected score is at least as large as any other candidate's, then the instance is open; otherwise, the instance is closed. Intuitively, whenever the first case holds, one of the candidates with a large expected score will surely win, but the manipulators are powerful enough to decide between them. However, if there is a candidate whose expected score is greater than any other's, even a large fraction of manipulators cannot prevent this candidate from winning.

THEOREM 3. *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.*

The proof of the theorem relies on Chernoff's bounds. Informally, these bounds assert that the probability that the average of i.i.d. random variables will be far from their common expectation is very small.

LEMMA 3 (CHERNOFF'S BOUNDS). *Let $X^1, \ldots, X^N$ be i.i.d. random variables such that $a \leq X^i \leq b$ and $\mathbb{E}[X^i] = \mu$. Then for any $\epsilon > 0$, it holds that:*

- $\Pr[\frac{1}{N}\sum_{i=1}^{N} X^i \geq \mu + \epsilon] \leq e^{-2N\frac{\epsilon^2}{(b-a)^2}}$
- $\Pr[\frac{1}{N}\sum_{i=1}^{N} X^i \leq \mu - \epsilon] \leq e^{-2N\frac{\epsilon^2}{(b-a)^2}}$

PROOF OF THEOREM 3.

1. Assume $|C'| \geq 2$. Using Lemma 2 with $d' = \frac{\alpha_1 - \alpha_m}{2m}$, and Lemma 1, we obtain:

$$\begin{aligned}
&\Pr_{D^N}[\text{Instance is open}] \\
&\geq \Pr_{D^N}[c \text{ can be made to win iff } c \in C'] \\
&\geq \Pr_{D^N}\left[\left(\forall c_{l_1}, c_{l_2} \in C', \ \sum_{i=1}^{N} S_{l_1}^i - \sum_{i-1}^{N} S_{l_2}^i < d'n\right)\right. \\
&\left.\wedge \left(\forall c_k \in C', c_l \in C \setminus C', \ \sum_{i=1}^{N} S_k^i - \sum_{i=1}^{N} S_l^i > \alpha_1 n\right)\right] \\
&= 1 - \Pr_{D^N}\left[\left(\exists c_{l_1}, c_{l_2} \in C' \ s.t. \ \sum_{i=1}^{N} S_{l_1}^i - \sum_{i-1}^{N} S_{l_2}^i \geq d'n\right)\right. \\
&\left.\vee \left(\exists c_k \in C', c_l \in C \setminus C' \ s.t. \ \sum_{i=1}^{N} S_k^i - \sum_{i=1}^{N} S_l^i \leq \alpha_1 n\right)\right]
\end{aligned} \tag{2}$$

Now, it holds that:

$$\begin{aligned}
&\Pr_{D^N}\left[\exists c_{l_1}, c_{l_2} \in C' s.t. \ \sum_{i=1}^{N} S_{l_1}^i - \sum_{i=1}^{N} S_{l_2}^i \geq d'n\right] \\
&\leq \sum_{c_{l_1}, c_{l_2} \in C'} \Pr_{D^N}\left[\sum_{i=1}^{N} S_{l_1}^i - \sum_{i=1}^{N} S_{l_2}^i \geq d'n\right] \\
&= \sum_{c_{l_1}, c_{l_2} \in C'} \Pr_{D^N}\left[\sum_{i=1}^{N}(S_{l_1}^i - S_{l_2}^i)\right. \\
&\qquad\qquad \left.\geq \mathbb{E}\left[\sum_{i=1}^{N}(S_{l_1}^i - S_{l_2}^i)\right] + d'n\right] \\
&\leq |C'| \cdot (|C'| - 1) \cdot e^{\frac{-2N\left(\frac{d'n}{N}\right)^2}{(2\alpha_1)^2}} \\
&\leq m(m-1)e^{-d_1\frac{n^2}{N}} \\
&= O(e^{-d_1\frac{n^2}{N}})
\end{aligned} \tag{3}$$

for some constant $d_1 > 0$. The first transition follows from the union bound, the second from the fact that all candidates in C' have maximal expected score and the linearity of expectation, and the third from Chernoff's bounds (where we use the fact that the difference between the scores given to two candidates by a voter is in the range $[-\alpha_1, \alpha_1]$, and that the S_k^i are i.i.d. for a fixed k if the D^i are i.i.d.).

Further, we have that:

$$\begin{aligned}
&\Pr_{D^N}\left[\exists c_k \in C', c_l \in C \setminus C' \ s.t. \ \sum_{i=1}^{N} S_k^i - \sum_{i=1}^{N} S_l^i \leq \alpha_1 n\right] \\
&\leq \sum_{c_k \in C', c_l \in C \setminus C'} \Pr_{D^N}\left[\sum_{i=1}^{N} S_k^i - \sum_{i=1}^{N} S_l^i \leq \alpha_1 n\right] \\
&= \sum_{c_k \in C', c_l \in C \setminus C'} \Pr_{D^N}\left[\sum_{i=1}^{N}(S_k^i - S_l^i)\right. \\
&\left.\leq \mathbb{E}\left[\sum_{i=1}^{N}(S_k^i - S_l^i)\right] - \left(\mathbb{E}\left[\sum_{i=1}^{N}(S_k^i - S_l^i)\right] - \alpha_1 n\right)\right] \\
&\leq |C'| \cdot (m - |C'|) \cdot e^{\frac{-2N\left(\frac{d''N - \alpha_1 n}{N}\right)^2}{(2\alpha_1)^2}} = O(e^{-d_2 N}).
\end{aligned} \tag{4}$$

The first transition follows from the union bound. The third transition is entailed by Chernoff's bounds, where d'' is a constant such that $\mathbb{E}[S_k^i - S_l^i] \geq d''$ for all $c_k \in C', c_l \in C \setminus C'$.[3] The last transition follows from the assumption that $n = o(N)$; $d_2 > 0$ is a constant.

Combining Equations (2), (3), and (4), and applying

[3]It is safe to state that such a constant d'' exists, as we assumed that $\mathbb{E}[S_k^i - S_l^i] > 0$, and D is (implicitly) a distribution that is dependent only on the number of candidates m, and not on the number of voters $N + n$.

the union bound, yields:

$$\begin{aligned}
&\Pr_{D^N}[\text{The instance is open}] \\
&\quad \geq 1 - \left(O(e^{-d_1 \frac{n^2}{N}}) + O(e^{-d_2 N}) \right) \\
&\quad = 1 - O(e^{-d \frac{n^2}{N}}),
\end{aligned}$$

for $d = d_1$, under the assumption that $n = o(N)$. When $n = \omega(\sqrt{N})$, this expression converges to 1 as the number of voters grows.

2. Assume $C' = \{c_k\}$. By Lemma 2, we have:

$$\begin{aligned}
&\Pr_{D^N}[\text{instance is closed}] \\
&\quad \geq \Pr_{D^N}\left[\forall c_l \neq c_k,\ \sum_i S_k^i - \sum_i S_l^i > \alpha_1 n\right] \\
&\quad = 1 - \Pr_{D^N}\left[\exists c_l \neq c_k\ s.t. \sum_i S_k^i - \sum_i S_l^i \leq \alpha_1 n\right].
\end{aligned} \tag{5}$$

Similarly to Equation (4), it holds that

$$\Pr_{D^N}\left[\exists c_l \neq c_k\ s.t. \sum_i S_k^i - \sum_i S_l^i \leq \alpha_1 n\right] \leq O(e^{-dN}).$$

Plugging this into Equation (5) gives the desired result.

□

The next corollary establishes a useful connection between the proof of Theorem 3 and deciding the coalitional manipulation problem.

Corollary 4. *Under the conditions of Theorem 3, 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 3 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 4 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 $c_k, c_l \in C$, $d < \mathrm{Var}[S_k^i - S_l^i]$. This is certainly a condition that seems very

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 $n = o(\sqrt{N})$ **then** ▷ Theorem 1
  choose arbitrary manipulators' vote; c is the winner
  if $p = c$ **then**
    return true
  else
    return false
  end if
else if $n = \omega(\sqrt{N})$ and $n = o(N)$ **then** ▷ Theorem 3
  if p has maximal expected score **then**
    return true
  else
    return false
  end if
else ▷ $n = \Theta(\sqrt{N})$ or $n = \Omega(N)$
  return ?
end if

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, we argue below that interesting distributions satisfy all three assumptions.

As an example, we shall consider the family of distributions that Conitzer and Sandholm used to obtain empirical evidence regarding the nonexistence of voting rules that are hard to manipulate [7]; these distributions are due to the Marquis de Condorcet himself. The starting point is that there is a "correct" ranking of candidates t, and voters disagree with this ranking over pairs of candidates with probability q. More formally, the probability of a voter casting a vote r is proportional to

$$q^{a(r,t)}(1-q)^{m(m-1)/2 - a(r,t)}, \tag{6}$$

where $a(r,t)$ is the number of pairs of candidates on whose relative ranking r and t agree. The parameter q can take values in $[1/2, 1]$: if $q = 1$ then the voters always agree with the correct ranking, and if $q = 1/2$ then voters vote randomly. Of course, the expression in Equation (6) above has to be normalized in order to obtain a probability distribution.

Proposition 5. *Let $m = O(1)$. Condorcet's distribution with any $0.5 \leq q < 1$ satisfies (d1), (d2), and (d3).*

Proof. By definition, the distribution satisfies (d2) and (d3), so it is sufficient to prove that (d1) is satisfied. Let $v^i \in V$, $c_k, c_l \in C$, and let $d' > 0$ be a constant such that dividing the expression in Equation (6) by d' yields a probability distribution. Let t be the "correct" ranking of candidates, and observe the restriction of t to all candidates other than c_k, c_l — $t \downarrow_{C \setminus \{c_k, c_l\}}$. Now, let r_1 be the expansion of this ranking such that c_k is ranked first and c_l last, and let r_2 be the expansion such that c_l is ranked first and c_k is ranked last. Under r_1 it holds that $S_k^i - S_l^i = \alpha_1 - \alpha_m$, while under r_2 it holds that $S_k^i - S_l^i = \alpha_m - \alpha_1$. Therefore, with respect to at least one of r_1 and r_2, it is true that $|(S_k^i - S_l^i) - \mathbb{E}[S_k^i - S_l^i]| \geq \alpha_1 - \alpha_m$ — w.l.o.g. with respect to r_1. By the construction of r_1, this ranking can differ from t only on pairs of candidates which include c_k or c_l, i.e., $a(r_1, t) \leq 2(m-1)$. Therefore, $\Pr[v^i\text{'s ballot is } r_1] \geq \frac{p^{2(m-1)}(1-p)^{(\frac{m}{2}-2)(m-1)}}{d'}$; denote this (constant) expression by d. To conclude, we have obtained:

$$\begin{aligned} \mathrm{Var}[S_k^i - S_l^i] &= \mathbb{E}[((S_k^i - S_l^i) - \mathbb{E}[S_k^i - S_l^i])^2] \\ &\geq \Pr[v^i\text{'s ballot is } r_1] \cdot (\alpha_1 - \alpha_m)^2 \\ &\geq d(\alpha_1 - \alpha_m)^2. \end{aligned}$$

□

Remark 2. Incidentally, the so-called *junta distribution*, which Procaccia and Rosenschein present in [12], satisfies (d2) and (d3) but not (d1): the votes of the nonmanipulators are distributed in an interval which is proportional to the number of manipulators, and thus the variance can be very small in terms of the number of nonmanipulators. This of course makes the distribution harder to manipulate — otherwise Theorem 1 could have been used to decide instances distributed with respect to the distribution.

Still, there remain gray areas, even when considering distributions that satisfy all three conditions: Theorems 1 and 3 do not apply to situations where $n = \Theta(\sqrt{N})$ or $n = \Omega(N)$. The second case, where $n = \Omega(N)$, does not seem very interesting: it is quite clear that when the number of manipulators is that large, the manipulators can usually do whatever they want, and thus an algorithm might always answer "true" in this case. However, the first case, where $n = \Theta(\sqrt{N})$, persists as a wide-open question. In fact, if an average-case-hard-to-manipulate distribution were to exist (especially in the context of scoring rules), our belief is that the distribution would be over instances that satisfy $n = \Theta(\sqrt{N})$.

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.

Theorem 1 deserves a short aside, as its conditions on the distribution are much weaker than those of Theorem 3, and an analogous version can be shown to be true for several other rules. This theorem, in the spirit of [14, 1], implies that manipulation may be a nonissue when the number of manipulators is small compared to the number of nonmanipulators: under almost any distribution on the votes (where the nonmanipulators vote independently), the manipulators can rarely even affect the outcome of the election.

5. CONCLUSIONS

In this paper, we have discussed how the ratio between the number of manipulators and nonmanipulators in an election using scoring rules can help decide the coalitional manipulation problem. We have shown that when the number of manipulators n satisfies $n = o(\sqrt{N})$, where N is the number of nonmanipulators, it holds that the manipulators can very rarely affect the outcome of the election. This result is true for almost any distribution over the votes (requiring merely independence and some minimal measure of randomness in the votes).

We have also discussed the case where $n = \omega(\sqrt{N})$; we have shown that in this case, it is almost always true that the manipulators can make any candidate with maximal expected score win the election, but no other candidate. However, this result requires that the votes be distributed i.i.d.

We have demonstrated how the abovementioned results can be applied to construct an algorithm that usually decides the coalitional manipulation problem correctly, when the distribution over the votes satisfies certain conditions. We have argued that the conditions are not very restrictive by demonstrating that an interesting family of distributions satisfies them. Finally, we have noted that our results strengthen the case against the existence of voting rules that are hard to manipulate in the average-case, but should one exist, our results may help zero-in on such a rule.

6. REFERENCES

[1] Eyad Baharad and Zvika Neeman. The asymptotic strategyproofness of scoring and Condorcet consistent rules. *Review of Economic Design*, 4:331–340, 2002.

[2] J. Bartholdi and J. Orlin. Single transferable vote resists strategic voting. *Social Choice and Welfare*, 8:341–354, 1991.

[3] J. Bartholdi, C. A. Tovey, and M. A. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6:227–241, 1989.

[4] V. Conitzer, J. Lang, and T. Sandholm. How many candidates are needed to make elections hard to manipulate? In *Proceedings of the International Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 201–214, 2003.

[5] V. Conitzer and T. Sandholm. Complexity of manipulating elections with few candidates. In *Proceedings of the National Conference on Artificial Intelligence*, pages 314–319, 2002.

[6] V. Conitzer and T. Sandholm. Universal voting protocol tweaks to make manipulation hard. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 781–788, 2003.

[7] 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.

[8] W. Feller. *Introduction to Probability Theory and its Applications*, volume 1. John Wiley, 3rd edition, 1968.

[9] S. Ghosh, M. Mundhe, K. Hernandez, and S. Sen. Voting for movies: the anatomy of a recommender system. In *Proceedings of the Third Annual Conference on Autonomous Agents*, pages 434–435, 1999.

[10] T. Haynes, S. Sen, N. Arora, and R. Nadella. An automated meeting scheduling system that utilizes user preferences. In *Proceedings of the First International Conference on Autonomous Agents*, pages 308–315, 1997.

[11] P. Orponen, K. Ko, U. Schoning, and O. Watanabe. Instance complexity. *Journal of the ACM*, 41(1):96–121, 1994.

[12] A. D. Procaccia and J. S. Rosenschein. Junta distributions and the average-case complexity of manipulating elections. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 497–504, 2006.

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

[14] Arkadii Slinko. How large should a coalition be to manipulate an election? *Mathematical Social Sciences*, 47(3):289–293, 2004.