

How Many Vote Operations Are Needed to Manipulate A Voting System?

Lirong Xia
School of Engineering and Applied Sciences
Harvard University
Cambridge, MA 02138, USA
lxia@seas.harvard.edu

Abstract

In this paper, we propose a framework to study a general class of strategic behavior in voting, which we call *vote operations*. Our main theorem is the following: if we fix the number alternatives, generate n votes i.i.d. according to a distribution π , and let n go to infinity. Then, for any $\epsilon > 0$, with probability at least $1 - \epsilon$, the minimum number of operations that are necessary for the strategic to achieve her goal falls into one of the following four categories: (1) 0, (2) $\Theta(\sqrt{n})$, (3) $\Theta(n)$, and (4) ∞ . This theorem holds for any set of vote operations, any individual vote distribution π , and any generalized scoring rule, which includes (but not limited to) most commonly studied voting rules, e.g., approval voting, all positional scoring rules (including Borda, plurality, and veto), plurality with runoff, Bucklin, Copeland, maximin, STV, and ranked pairs.

We also show that many well-studied types of strategic behavior fall under our framework, including constructive and destructive manipulation, bribery, and control by adding/deleting votes, and margin of victory and minimum manipulation coalition size. Therefore, our main theorem naturally applies to these problems.

1 Introduction

Voting is a popular method used to aggregate voters’ preferences to make a joint decision. One of the most desired properties for voting rules is *strategy-proofness*, that is, no voter has incentive to misreport her preferences to obtain a better outcome of the election. Unfortunately, strategy-proofness is not compatible with some other natural desired properties, due to the celebrated Gibbard-Satterthwaite theorem [14, 22], which states that when there are at least three alternatives, no strategy-proof voting rule satisfy the following two natural properties: non-imposition (every alternative can win) and non-dictatorship (no voter is a dictator whose top ranked alternative is always the winner).

Even though manipulation is inevitable, researchers have set out to investigate whether computational complexity can serve as a barrier against various types of strategic behavior, including manipulation. The idea is, if we can prove that it is computationally too costly for a strategic individual to find a beneficial operation, she may give up doing so. Initiated by Bartholdi, Tovey, and Trick [2], a fair amount of work has been done in the computational social choice community to characterize the computational complexity of various types strategic behavior, including the following. See [12, 10, 21] for recent surveys.

- *Manipulation*: a voter or coalition of voters cast false vote(s) to make the winner more preferable.
- *Bribery*: a strategic individual changes some votes by bribing the voters to make the winner preferable to her [9]. The bribery problem is closely related to the problem of computing the *margin of victory* [5, 16, 31].
- *Control*: a strategic individual adds or deletes votes to make the winner more preferable to her [3]. Control by adding votes is equivalent to *false-name manipulation* [6].

Most previous results in “using computational complexity as a barrier against strategic behavior” are worst-case analyses of computational complexity. Recently, increasing number of results show that manipulation, as a particular type of strategic behavior, is typically not hard to compute. One direction, mainly pursued in the theoretical computer science community, is to obtain a quantitative version of the Gibbard-Satterthwaite theorem, showing that for any given voting rule that is “far” enough from any dictatorships, an instance of manipulation can be found easily with high probability. This line of research was initiated by Friedgut, Kalai, and Nisan [13], where they proved the theorem for 3 alternatives and neutral voting rules. The theorem was extended to arbitrary number of alternatives by Isaksson, Kindler, and Mossel [15], and finally, the neutrality constraint was removed by Mossel and Racz [17]. Other extensions include Dobzinski and Procaccia [8] and Xia and Conitzer [33].

Another line of research is to characterize the “frequency of manipulability”, defined as the probability for a randomly generated preference-profile to be manipulable by a group of manipulators, where the non-manipulators’ votes are generated i.i.d. according to some distribution (for example, the uniform distribution over all possible types of preferences). Peleg [18], Baharad and Neeman [1], and Slinko [23, 24] studied the asymptotic value of the frequency of manipulability for positional scoring rules when the non-manipulators’ votes are drawn i.i.d. uniformly at random. Procaccia and Rosenschein [20] showed that for positional scoring rules, when the non-manipulators votes are drawn i.i.d. according to any distribution that satisfies some natural conditions, if the number of manipulators is $o(\sqrt{n})$, where n is the number of non-manipulators, then the probability that the manipulators can succeed goes to 0 as n goes to infinity; if the number of manipulator is $\omega(n)$, then the probability that the manipulators can succeed goes to 1. This dichotomy theorem was generalized to a class of voting rules called *generalized scoring rules (GSRs)* by Xia and Conitzer [32]. GSRs are believed to be a general class of voting rules, because many commonly studied voting rules are GSRs, including (but not limited to) approval voting, all positional scoring rules (which include Borda, plurality, and veto), plurality with runoff, Bucklin, Copeland, maximin, STV, and ranked pairs.¹ Another evidence is that GSRs admit a natural axiomatic characterization [34].

¹We will formally define generalized scoring rules in Section 2. The definition of these commonly studied voting rules can be

While most of the aforementioned results are about manipulation, in this paper, we focus the optimization variants of various types of strategic behavior, including manipulation, bribery, and control. Despite being natural, to the best of our knowledge, such optimization variants have been investigated for only three types of strategic behavior. The first is the *unweighted coalitional optimization (UCO)* problem, where we are asked to compute the minimum number of manipulators who can make a given alternative win [37]. Approximation algorithms have been proposed for UCO for specific voting systems, including positional scoring rules and maximin [37, 35, 36]. The second is the *margin of victory* problem, where we are asked to compute the smallest number of voters who can change their votes to change the winner [16, 5, 31]. The third is the *minimum manipulation coalition size* problem, which is similar to the margin of victory, except that all voters who change their votes must prefer the new winner to the old winner [19].

1.1 Our Contributions

In this paper, we introduce a unified framework to study a class of strategic behavior for generalized scoring rules, which we call *vote operations*. In our framework, a strategic individual seeks to change the winner by applying some operations, which are modeled as vectors in a multidimensional space. We study three goals of the strategic individual: (1) making a favored alternative win, called *constructive vote operation (CVO)*, (2) making a disfavored alternative lose, called *destructive vote operation (DVO)*, and (3) change the winner of the election, called *change-winner vote operation (CWVO)*. The framework will be formally defined in Section 3. This is our main conceptual contribution.

Our main technical contribution is the following asymptotical characterization of the number of operations that are necessary for the strategic individual to achieve her goal.

Theorem 1 (informally put) Fix the number of alternatives and the set of vote operations. For any generalized scoring rules and any distribution π over votes, we generate n votes i.i.d. according to π and let n go to infinity. Then, for any VO $\in \{CVO, DVO, CWVO\}$ and any $\epsilon > 0$, with probability at least $1 - \epsilon$, the minimum number of operations that are necessary for the strategic individual to achieve VO falls into one of the following four categories: (1) 0, (2) $\Theta(\sqrt{n})$, (3) $\Theta(n)$, and (4) ∞ .

More informally, Theorem 1 states that in large elections, to achieve a specific goal (one of the three goals described above), with probability that can be infinitely close to 1 the strategic individual need to either do nothing (the goal is already achieved), apply $\Theta(\sqrt{n})$ vote operations, apply $\Theta(n)$ vote operations, or the goal cannot be achieved no matter how many vote operations are applied. This characterization holds for any generalized scoring rule, any set of vote operations, and any distribution π for individual votes.

The proof of Theorem 1 is based on the Central Limit Theorem and sensitivity analyses of integer linear programmings (ILPs). It works as follows. We will formulate each of the strategic individual's three goals as a set of ILPs in Section 4. By applying Central Limit Theorem, we show that with probability that goes to 1 the random generated preference-profile satisfies a desired property. Then, for each such preference-profile we apply the sensitivity analyses in [7] to show that with high probability the number of operations that are necessary is either 0, $\Theta(\sqrt{n})$, $\Theta(n)$, or ∞ .

While Theorem 1 looks quite abstract, we show later in the paper that many well-studied types of strategic behavior fall under our vote operation framework, including constructive (destructive) unweighted coalitional optimization, bribery, control by adding/deleting voters, and margin of victory and minimum manipulation coalition size.² Therefore, we naturally obtain corollaries of Theorem 1 for these types of strategic behavior. Of course our theorem applies to other types of strategic behavior, for example the mixture of any types mentioned above, which is known as *multimode control attacks* [11].

found in, e.g., [32].

²We defer the definition of these types of strategic behavior to Section 6.

1.2 Related Work and Discussion

To the best of our knowledge, we are the first to do the following in the voting setting: (1) study manipulation, bribery, and control under a unified framework, (2) model the strategic individual’s goals as ILPs, and (3) conduct sensitivity analyses for them. Our main theorem applies to any generalized scoring rule for destructive manipulation, constructive and destructive bribery and control by adding/deleting votes, where no similar results were obtained even for specific voting rules. Three previous papers obtained similar results for specific types of strategic behavior. The applications of our main theorem to these types of strategic behavior are slightly weaker, but we stress that our main theorem is significantly more general.

First, the dichotomy theorem in [32] implies that, (informally) when the votes are drawn i.i.d. from *some* distribution, with probability that goes to 1 the solution to constructive and destructive UCO is either 0 or approximately \sqrt{n} for *some* favored alternatives. However, this result only works for the UCO problem and some distributions over the votes.

Second, it was proved in [31] that for any non-redundant generalized scoring rules that satisfy a continuity condition, when the votes are drawn i.i.d. and let the number of voters n go to infinity, either with probability that can be arbitrarily close to 1 the margin of victory is $\Theta(\sqrt{n})$, or with probability that can be arbitrarily close to 1 the margin of victory is $\Theta(n)$. It is easy to show that for non-redundant voting rules, the margin of victory is never 0 or ∞ . Though it was shown in [31] that many commonly studied voting rules are GSRs that satisfy such continuity condition, in general it is not clear how restrictive the continuity condition is. More importantly, the result only works for the margin of victory problem.

Third, in [19], the authors investigated the distribution over the minimum manipulation coalition size for positional scoring rules when the votes are drawn i.i.d. from the uniform distribution. However, it is not clear how their techniques can be extended beyond the uniform distributions and positional scoring rules, which are a very special case of generalized scoring rules. Moreover, the paper only focused on the minimum manipulation coalition size problem.

Our results has both negative and positive implications. On the negative side, our results provide yet another evidence that computational complexity is not a strong barrier against strategic behavior, because the strategic individual now has some information about the number of operations that are needed, without spending any computational cost or even without looking at the input instance. Although the estimation of our theorem may not be very precise (because we do not know which of the four cases a given instance belongs to), such estimation may be explored to designing effective algorithms that facilitate strategic behavior. On the positive side, this easiness of computation is not always a bad thing: sometimes we want to do such computation in order to test how robust a given preference-profile is. For example, computing the margin of victory is an important component in designing novel *risk-limiting audit methods* [16, 5, 31, 25, 26, 28, 27, 29].

While being quite general, our results have mainly three limitations. First, they are asymptotical results, where we fix the number of alternatives and let the number of voters go to infinity. We do not know the convergence rate, or equivalently, how many voters are needed for the observation to hold. In fact, this is a standard setting in previous work, especially in the studies of “frequency of manipulability”. We feel that our results work well in settings where there are small number of alternatives and large number of voters, e.g., political elections. Second, our results show that with high probability one of the four cases holds (0 , $\Theta(\sqrt{n})$, $\Theta(n)$, ∞), but we do not know which case holds more often. We will briefly discuss this issue in Appendix B. It is possible to refine our study for specific voting rules and specific types of strategic behavior that fall under our framework, which we leave as future work.

2 Preliminaries

Let \mathcal{C} denote the set of *alternatives*, $|\mathcal{C}| = m$. We assume strict preference orders. That is, a vote is a linear order over \mathcal{C} . The set of all linear orders over \mathcal{C} is denoted by $L(\mathcal{C})$. A *preference-profile* P is a collection of n votes for some $n \in \mathbb{N}$, that is, $P \in L(\mathcal{C})^n$. Let $L(\mathcal{C})^* = \bigcup_{n=1}^{\infty} L(\mathcal{C})^n$. A *voting rule* r is a mapping that assigns to each preference-profile a single winner. That is, $r : L(\mathcal{C})^* \rightarrow \mathcal{C}$. Throughout the paper, we let n denote the number of votes and m denote the number of alternatives.

We now recall the definition of *generalized scoring rules (GSRs)* [32]. For any $K \in \mathbb{N}$, let $\mathcal{O}_K = \{o_1, \dots, o_K\}$. A *total preorder* (*preorder* for short) is a reflexive, transitive, and total relation. Let $\text{Pre}(\mathcal{O}_K)$ denote the set of all preorders over \mathcal{O}_K . For any $\vec{p} \in \mathbb{R}^K$, we let $\text{Ord}(\vec{p})$ denote the preorder \succeq over \mathcal{O}_K where $o_{k_1} \succeq o_{k_2}$ if and only if $p_{k_1} \geq p_{k_2}$. That is, the k_1 -th component of \vec{p} is as large as the k_2 -th component of \vec{p} . For any preorder \succeq , if $o \succeq o'$ and $o' \succeq o$, then we write $o =_{\succeq} o'$. Each preorder \succeq naturally induces a (partial) strict order \triangleright , where $o \triangleright o'$ if and only if $o \succeq o'$ and $o' \not\succeq o$.

Definition 1 Let $K \in \mathbb{N}$, $f : L(\mathcal{C}) \rightarrow \mathbb{Z}^K$ and $g : \text{Pre}(\mathcal{O}_K) \rightarrow \mathcal{C}$. f and g determine a generalized scoring rule (GSR) $GS(f, g)$ as follows. For any preference-profile $P = (V_1, \dots, V_n) \in L(\mathcal{C})^n$, let $f(P) = \sum_{i=1}^n f(V_i)$, and let $GS(f, g)(P) = g(\text{Ord}(f(P)))$. We say that $GS(f, g)$ is of order K .

For any $V \in L(\mathcal{C})$, $f(V)$ is called a *generalized scoring vector*, $f(P)$ is called a *total generalized scoring vector*, and $\text{Ord}(f(P))$ is called the *induced preorder* of P . Most commonly studied voting rules are generalized scoring rule, including (but not limited to) approval voting, Bucklin, Copeland, maximin, plurality with runoff, ranked pairs, and multi-stage voting rules that use GSRs in each stage to eliminate alternatives (including Nanson's and Baldwin's rule).

As an example, we recall the proof from [32] that the *single transferable vote (STV)* rule (a.k.a. *instant-runoff voting* or *alternative vote* for single-winner elections) is a generalized scoring rule.

Example 1 *STV selects the winner in m rounds. In each round, the alternative that gets the lowest plurality score (the number of times that the alternative is ranked in the top position) drops out, and is removed from all of the votes (so that votes for this alternative transfer to another alternative in the next round). Ties are broken alphabetically. The last-remaining alternative is the winner.*

To see that STV is a GSR, we will use generalized score vectors with many components. For every proper subset S of alternatives, for every alternative c outside of S , there is a component in the vector that contains the number of times that c is ranked first if all of the alternatives in S are removed. Let

- $K_{STV} = \sum_{i=0}^{m-1} \binom{m}{i} (m-i)$; the components are indexed by (S, j) , where S is a proper subset of \mathcal{C} and $j \leq m, c_j \notin S$.
- $(f_{STV}(V))_{(S,j)} = 1$, if after removing S from V , c_j is at the top; otherwise, let $(f_{STV}(V))_{(S,j)} = 0$.
- g_{STV} selects the winners based on \succeq as follows. In the first round, let j_1 to be the index such that $o_{(\emptyset, j_1)}$ is ranked the lowest in \succeq among all $o_{(\emptyset, j)}$ (if there are multiple such j 's, then we break ties alphabetically to select the least-preferred one). Let $S_1 = \{c_{j_1}\}$. Then, for any $2 \leq i \leq m-1$, define S_i recursively as follows: $S_i = S_{i-1} \cup \{j_i\}$, where j_i is the index such that $o_{(S_{i-1}, j_i)}$ is ranked the lowest in \succeq among all $o_{(S_{i-1}, j)}$; finally, the winner is the unique alternative in $(\mathcal{C} \setminus S_{m-1})$.

Another evidence on the generality of GSRs is that GSRs admit a natural axiomatic characterization [34]. That is, GSRs are the class of voting rules that satisfy *anonymity*, *homogeneity*, and *finite local consistency*. Anonymity says that the winner does not depend on the name of the voters, homogeneity says that if we duplicate the preference-profile multiple times, the winner does not change, and finite local consistency is an approximation to the well-studied *consistency* axiom. Not all voting rules are GSRs, for example, *Dodgson's rule* is not a GSR because it does not satisfy homogeneity [4], and the following skewed majority rule is also not a GSR because it also violates homogeneity.

Example 2 For any $0 < \gamma < 1$, the γ -majority rule is defined for two alternatives $\{a, b\}$ as follows: b is the winner if and only if the number of voters who prefer b is more than the number voters who prefer a by at least n^γ .

Admittedly, these γ -majority rules are quite artificial. Later in this paper we will see that the observation made for GSRs in our main theorem (Theorem 1) does not hold for γ -majority rules for any $\gamma \neq 1/2$.

3 The Unified Framework

All types of strategic behavior mentioned in the introduction have the following characteristics in common. The strategic individual (who can be a group of manipulators, a briber, or a controller, etc.) changes the winner by changing the votes in the preference-profile. Therefore, for generalized scoring rules, any such an operation is uniquely represented by changes in the total generalized scoring vector. This is in contrast to some other types of strategic behavior where the strategic individual changes the set of alternatives or the voting rule [3, 30]. In this section, we first define the set of operations the strategic individual can apply, then define her goals. Given a generalized scoring rule of order K , we model the strategic behavior, called *vote operations*, as a set of vectors, each of which has K dimensions, representing the changes made to the total generalized scoring vector if the strategic individual apply this operation.

Definition 2 Given a GSR $GS(f, g)$ of order K , let $\Delta = [\vec{\delta}_1 \cdots \vec{\delta}_T]$ denote the vote operations, where for each $i \leq T$, $\vec{\delta}_i \in \mathbb{Z}^K$ represents the changes made to the generalized scoring vector by applying the i -th vote operation. For each $l \leq K$, let Δ_l denote the l -th row of Δ .

We will show examples of these vote operations for some well-studied types of strategic behavior in Section 6. Given the set of available operations Δ , the strategic individual's behavior is characterized by a vector $\vec{v} \in \mathbb{N}_{\geq 0}^T$, where for each $i \leq T$, v_i represents the number of i -th operation (corresponding to $\vec{\delta}_i$) that she applies. Let $(\vec{v})'$ denote the transpose of \vec{v} and let $\|\vec{v}\|_1 = \sum_{i=1}^T v_i$ denote the total number of operations in \vec{v} , which is the L1-norm of \vec{v} . It follows that $\Delta \cdot (\vec{v})'$ is the change in the total generalized scoring vector introduced by the strategic individual, where for any $l \leq K$, $\Delta_l \cdot (\vec{v})'$ is the change in the l -th component.

Next, we give definitions of the strategic individual's three goals and the corresponding computational problems studied in this paper.

Definition 3 In the CONSTRUCTIVE VOTE OPERATION (CVO) problem, we are given a generalized scoring rule $GS(f, g)$, a preference-profile P , a favored alternative c , and a set of vote operations $\Delta = [\vec{\delta}_1 \cdots \vec{\delta}_T]$, and we are asked to compute the smallest number k , denoted by $\text{CVO}(P, c)$, such that there exists a vector $\vec{v} \in \mathbb{N}_{\geq 0}^T$ with $\|\vec{v}\|_1 = k$ and $g(\text{Ord}(f(P) + \Delta \cdot (\vec{v})')) = c$. If such \vec{v} does not exist, then we denote $\text{CVO}(P, c) = \infty$.

The DESTRUCTIVE VOTE OPERATION (DVO) problem is defined similarly, where c is the disfavored alternative, and we are asked to compute the smallest number k , denoted by $\text{DVO}(P, c)$, such that there exists a vector $\vec{v} \in \mathbb{N}_{\geq 0}^T$ with $\|\vec{v}\|_1 = k$ and $g(\text{Ord}(f(P) + \Delta \cdot (\vec{v})')) \neq c$.

In the CHANGE-WINNER VOTE OPERATION (CWVO) problem, we are not given c and we are asked to compute $\text{DVO}(P, GS(f, g)(P))$, denoted by $\text{CWVO}(P)$.

In CVO, the strategic individual seeks to make c win; in DVO, the strategic individual seeks to make c lose; and in CWVO, the strategic individual seeks to change the current winner.

For a given instance (P, r) , CWVO is a special case of DVO, where $c = GS(f, g)(P)$. We distinguish these two problems because in this paper, the input preference-profiles are generated randomly, so the winners of these preference-profiles might be different. Therefore, when the preference-profiles are randomly generated, the distribution for the solution to DVO does not immediately give us a distribution for the solution to CWVO.

4 The ILP Formulation

Let us first put aside the strategic individual's goal for the moment (i.e., making a favored alternative win, making a disfavored alternative lose, or changing the winner) and focus on the following question: given a preference-profile P and a preorder \succeq over the K components of the generalized scoring vector, that is, $\succeq \in \text{Pre}(\mathcal{O}_K)$, how many vote manipulations are need to change the order of the total generalized scoring vector to \succeq ? Formally, given a $\text{GS}(f, g)$, a preference-profile P and $\succeq \in \text{Pre}(\mathcal{O}_K)$, we are interested in $\min\{\|\vec{v}\|_1 : \vec{v} \in \mathbb{N}_{\geq 0}^K, \text{Ord}(f(P) + \Delta \cdot (\vec{v})') = \succeq\}$.

This can be computed by the following integer linear programming ILP_{\succeq} , where v_i represents the i th component in \vec{v} , which must be a nonnegative integer. We recall that Δ_l denotes the l -th row vector of Δ .

$$\begin{aligned} \min \quad & \|\vec{v}\|_1 \\ \text{s.t.} \quad & \forall o_i =_{\succeq} o_j : (\Delta_i - \Delta_j) \cdot (\vec{v})' = [f(P)]_j - [f(P)]_i \\ & \forall o_i \succ o_j : (\Delta_i - \Delta_j) \cdot (\vec{v})' \geq [f(P)]_j - [f(P)]_i + 1 \\ & \forall i : v_i \geq 0 \end{aligned} \quad (\text{LP}_{\succeq})$$

Now, we take the strategic individual's goal into account. We immediately have the following lemma as a warmup, whose proofs are straightforward and are thus omitted.

Lemma 1 *Given a GSR $\text{GS}(f, g)$, an alternative c , and a preference-profile P ,*

- $\text{CVO}(P, c) < \infty$ if and only if there exists \succeq such that $g(\succeq) = c$ and LP_{\succeq} has an integer solution;
- $\text{DVO}(P, c) < \infty$ if and only if there exists \succeq such that $g(\succeq) \neq c$ and LP_{\succeq} has an integer solution;
- $\text{CWVO}(P) < \infty$ if and only if there exists \succeq such that $g(\succeq) \neq \text{GS}(f, g)(P)$ and LP_{\succeq} has an integer solution. (We do not need the input c for this problem.)

Moreover, the solution to each of the three problems is the minimum objective value in all LPs corresponding to the problem. For example, if $\text{CVO}(P, c) < \infty$, then

$$\text{CVO}(P, c) = \min_{\|\vec{v}\|_1} \{ \vec{v} \text{ is the solution to some } \text{LP}_{\succeq} \text{ where } g(\succeq) = c \}$$

5 The Main Theorem

In this section we prove the main theorem, which states that fix m , for any generalized scoring rules and any set of vote operations Δ , if n votes are generated i.i.d., then for CVO (respectively, DVO, CWVO), with probability that can be infinitely close to 1, the solution is either 0, $\Theta(\sqrt{n})$, $\Theta(n)$, or ∞ .

Theorem 1 *Let $\text{GS}(f, g)$ be a generalized scoring rule, let π be a distribution over all linear orders, and let Δ be a set of vote operations. Suppose we fix the number of alternatives, generate n votes i.i.d. according to π , and let P_n denote the preference-profile. Then, for any alternative c , $\text{VO} \in \{\text{CVO}, \text{DVO}, \text{CWVO}\}^3$, and any $\epsilon > 0$, there exists $\beta^* > 1$ such that as n goes to infinity, the total probability for the following four events sum up to more than $1 - \epsilon$: (1) $\text{VO}(P_n, c) = 0$, (2) $\frac{1}{\beta^*} \sqrt{n} < \text{VO}(P_n, c) < \beta^* \sqrt{n}$, (3) $\frac{1}{\beta^*} n < \text{VO}(P_n, c) < \beta^* n$, and (4) $\text{VO}(P_n, c) = \infty$.*

Proof of Theorem 1: Let $P_\pi = \sum_{V \in \mathcal{L}(\mathcal{C})} \pi(V) \cdot V$, $f(P_\pi) = \sum_{V \in \mathcal{L}(\mathcal{C})} \pi(V) \cdot f(V)$, and $\succeq_\pi = \text{Ord}(f(P_\pi))$. We first prove the theorem for CVO, and then show how to adjust the proof for DVO and CWVO. The theorem is proved in the following two steps. Step 1: we show that as n goes to infinity, with probability that goes to one we have the following: in a randomly generated P_n , the difference between any pair of components in $f(P_n)$ is either $\Theta(\sqrt{n})$ or $\Theta(n)$. Step 2: we apply sensitivity analyses to ILPs that are similar to the

³When $\text{VO} = \text{CWVO}$, we let $\text{VO}(P_n, c)$ denote $\text{CWVO}(P_n)$.

ILP given in Section 4 to prove that for any such preference-profile and any $\text{VO} \in \{\text{CVO}, \text{DVO}, \text{CWVO}\}$, $\text{VO}(P_n, c)$ is either 0, $\Theta(\sqrt{n})$, $\Theta(n)$, or ∞ . The idea behind Step 2 is, for any preference-profile P_n , if the difference between a pair of components in $f(P_n)$ is $\Theta(\sqrt{n})$, then we consider this pair of components (not alternatives) are “almost tied”; if the difference is $\Theta(n)$, then we consider them to be “far away”. Take CVO as an example, we can easily identify the cases where $\text{CVO}(P_n, c)$ is either 0 (when $\text{GS}(f, g) = c$) or ∞ (by Lemma 1). Then, we will first try to break these “almost tied” pairs by using LPs that are similar to LP_{\succeq} introduced in Section 4, and show that if there exists an integer solution \vec{v} , then the objective value $\|\vec{v}\|_1$ is $\Theta(\sqrt{n})$. Otherwise, we have to change the orders between some “far away” pairs by using LP_{\succeq} ’s, and show that if there exists an integer solution to some LP_{\succeq} with $g(\succeq) = c$, then the objective value is $\Theta(n)$.

Formally, given $n \in \mathbb{N}$ and $\beta > 1$, let \mathcal{P}_β denote the set of all n -vote preference-profile P that satisfies the following two conditions (we recall that $f(P_\pi) = \sum_{V \in L(c)} \pi(V) \cdot f(V)$): for any pair $i, j \leq K$,

1. if $[f(P_\pi)]_i = [f(P_\pi)]_j$ then $\frac{1}{\beta}\sqrt{n} < |[f(P)]_i - [f(P)]_j| < \beta\sqrt{n}$;
2. if $[f(P_\pi)]_i \neq [f(P_\pi)]_j$ then $\frac{1}{\beta}n < |[f(P)]_i - [f(P)]_j| < \beta n$.

The following lemma was proved in [31], which follows from the Central Limit Theorem.

Lemma 2 *For any ϵ , there exists β such that $\lim_{n \rightarrow \infty} \Pr(P_n \in \mathcal{P}_\beta) > 1 - \epsilon$.*

For any given ϵ , in the rest of the proof we fix β to be a constant guaranteed by Lemma 2. The next lemma will be frequently used in the rest of the proof, whose proof is deferred to the appendix.

Lemma 3 *Fix an integer matrix \mathbf{A} . There exists a constant $\beta_{\mathbf{A}}$ that only depends on \mathbf{A} , such that if the following LP has an integer solution, then the solution is no more than $\beta_{\mathbf{A}} \cdot \|\vec{b}\|_\infty$.*

$$\min \|\vec{x}\|_1, \text{ s.t. } \mathbf{A} \cdot \vec{x} \geq \vec{b}$$

To prove that with high probability $\text{CVO}(P_n, c)$ is either 0, $\Theta(\sqrt{n})$, $\Theta(n)$, or ∞ , we introduce the following notation. A preorder \succeq' is a *refinement* of another preorder \succeq , if \succeq' extends \succeq . That is, $\succeq \subseteq \succeq'$. We note that \succeq is a refinement of itself. Let $\succeq' \ominus \succeq$ denote the strict orders that are in \succeq' but not in \succeq . That is, $(o_i, o_j) \in (\succeq' \ominus \succeq)$ if and only if $o_i \succ' o_j$ and $o_i =_{\succeq} o_j$. We define the following LP that is similar to LP_{\succeq} defined in Section 4, which will be used to check whether there is a way to break “almost tied” pairs of components to make c win. For any preorder \succeq and any of its refinement \succeq' , we define $\text{LP}_{\succeq' \ominus \succeq}$ as follows.

$$\begin{aligned} & \min \|\vec{v}\|_1 \\ \text{s.t.} \quad & \forall o_i =_{\succeq} o_j : (\Delta_i - \Delta_j) \cdot (\vec{v})' = [f(P)]_j - [f(P)]_i \\ & \forall (o_i, o_j) \in (\succeq' \ominus \succeq) : (\Delta_i - \Delta_j) \cdot (\vec{v})' \geq [f(P)]_j - [f(P)]_i + 1 \\ & \forall i : v_i \geq 0 \end{aligned} \quad (\text{LP}_{\succeq' \ominus \succeq})$$

$\text{LP}_{\succeq' \ominus \succeq}$ is defined with a little abuse of notation because some of its constraints depend on \succeq (not only the pairwise comparisons in $(\succeq' \ominus \succeq)$). This will not cause confusion because we will always indicate \succeq in the subscript. We note that there is a constraint in $\text{LP}_{\succeq' \ominus \succeq}$ for each pair of components o_i, o_j with $o_i =_{\succeq} o_j$. Therefore, $\text{LP}_{\succeq' \ominus \succeq}$ is used to find a solution that breaks ties in \succeq . It follows that $\text{LP}_{\succeq' \ominus \succeq}$ has an integer solution \vec{v} if and only if the strategic individual can make the order between any pairs of o_i, o_j with $o_i =_{\succeq} o_j$ to be those in \succeq' by applying the i -th operation v_i times, and the total number of vote operations is $\|\vec{v}\|_1$.

The following two claims identify the preference-profiles in \mathcal{P}_β for which CVO is $\Theta(\sqrt{n})$ and $\Theta(n)$, respectively, whose proofs are deferred to the appendix.

Claim 1 *There exists $N \in \mathbb{N}$ and $\beta' > 1$ such that for any $n \geq N$, any $P \in \mathcal{P}_\beta$, if (1) c is not the winner for P , and (2) there exists a refinement \succeq^* of $\succeq_\pi = \text{Ord}(f(P_\pi))$ such that $g(\succeq^*) = c$ and $\text{LP}_{\succeq^* \ominus \succeq_\pi}$ has an integer solution, then $\frac{1}{\beta'}\sqrt{n} < \text{CVO}(P, c) < \beta'\sqrt{n}$.*

Claim 2 *There exists $\beta' > 1$ such that for any $P \in \mathcal{P}_\beta$, if (1) c is not the winner for P , (2) there does not exist a refinement \succeq^* of $\succeq_\pi = \text{Ord}(f(P_\pi))$ such that $LP_{\succeq^* \ominus \succeq_\pi}$ has an integer solution, and (3) there exists \succeq such that $g(\succeq) = c$ and LP_{\succeq} has an integer solution, then $\frac{1}{\beta'}n < \text{CVO}(P, c) < \beta'n$.*

Lastly, for any $P \in \mathcal{P}_\beta$ such that $\text{GS}(f, g)(P) \neq c$, the only case not covered by Claim 1 and Claim 2 is that there does not exist \succeq with $\text{GS}(f, g)(\succeq) = c$ such that LP_{\succeq} has an integer solution. It follows from Lemma 1 that in this case $\text{CVO}(P, c) = \infty$. We note that β' in Claim 1 and Claim 2 does not depend on n . Let β^* be an arbitrary number that is larger than the two β' 's. This proves the theorem for CVO.

For DVO, we only need to change $g(\succeq^*) = c$ to $g(\succeq^*) \neq c$ in Claim 1, and change $g(\succeq) = c$ to $g(\succeq) \neq c$ in Claim 2. For CWVO, $\text{CWVO}(P)$ is never 0 and we only need to change $g(\succeq^*) = c$ to $g(\succeq^*) \neq \text{GS}(f, g)(P)$ in Claim 1, and change $g(\succeq) = c$ to $g(\succeq) \neq \text{GS}(f, g)(P)$ in Claim 2. ■

Remark. The intuition in Lemma 2 is quite straightforward and naturally corresponds to random walk in multidimensional space. However, we did not find an obvious connection between random walk theory and observation made in Theorem 1 for general voting rules. We believe that it is unlikely that an obvious connection exists. One evidence is that the observation made in Theorem 1 does not hold for some voting rules. For example, consider the γ -majority rule defined in Example 2. It is not hard to see that as n goes to infinity, with probability that goes to 1 we have $\text{CVO}(P_n, b) = \text{DVO}(P_n, a) = \text{CWVO}(P_n) = n^\gamma/2$, which is not any of the four cases described in Theorem 1 if $\gamma \neq 1/2$. (This means that for any $\gamma \neq 1/2$, γ -majority is not a generalized scoring rule, which we already know because they do not satisfy homogeneity.)

The main difficulty in proving Theorem 1 is, for generalized scoring rules we have to handle the cases where some components of the total generalized scoring vector are equivalent. This only happens with negligible probability for the randomly generated preference-profile P_n , but it is not clear how often the strategic individual can make some components equivalent in order to achieve her goal. This is the main reason for us to convert the vote manipulation problem to multiple ILPs and conduct sensitivity analyses.

6 Applications of The Main Theorem

In this section we show how to apply Theorem 1 to some well-studied types of strategic behavior, including constructive and destructive unweighted coalitional optimization, bribery and control, and margin of victory and minimum manipulation coalition size. In the sequel, we will use each subsection to define these problems and describe how they fit in our vote operation framework, and how Theorem 1 applies. In the end of the section we present a unified corollary for all these types of strategic behavior.

6.1 Unweighted Coalitional Optimization

Definition 4 *In a constructive (respectively, destructive) UNWEIGHTED COALITIONAL OPTIMIZATION (UCO) problem, we are given a voting rule r , a preference-profile P^{NM} of the non-manipulators, and a (dis)favoured alternative $c \in \mathcal{C}$. We are asked to compute the smallest number of manipulators who can cast votes P^M such that $c = r(P^{NM} \cup P^M)$ (respectively, $c \neq r(P^{NM} \cup P^M)$).*

To see how UCO fits in the vote operation model, we view the group of manipulators as the strategic individual, and each vote cast by a manipulator is a vote operation. Therefore, the set of operations is exactly the set of all generalized scoring vectors $\{f(V) : V \in L(\mathcal{C})\}$. To apply Theorem 1, for constructive UCO we let $\text{VO} = \text{CVO}$ and for destructive UCO we let $\text{VO} = \text{DVO}$.

6.2 Bribery

In this paper we are interested in the optimization variant of the bribery problem [9].

Definition 5 *In a constructive (respectively, destructive) OPT-BRIBERY problem, we are given a preference-profile P and a (dis)avored alternative $c \in \mathcal{C}$. We are asked to compute the smallest number k such that the strategic individual can change no more than k votes such that c is the winner (respectively, c is not the winner).*

To see how OPT-BRIBERY falls under the vote operation framework, we view each action of “changing a vote” as a vote operation. Since the strategic individual can only change existing votes in the preference-profile, we define the set of operations to be the difference between the generalized scoring vectors of all votes and the generalized scoring vectors of votes in the support of π , that is, $\{f(V) - f(W) : V, W \in L(\mathcal{C}) \text{ s.t. } \pi(W) > 0\}$. Then, similarly the constructive variant corresponds to CVO and the destructive variant corresponds to DVO. In both cases Theorem 1 cannot be directly applied, because in the ILPs we did not limit the total number of each type of vote operations that can be used by the strategic individual. Nevertheless, we can still prove a similar proposition by taking a closer look at the relationship between CVO (DVO) and OPT-BRIBERY as follows: For any preference-profile, the solution to CVO (respectively, DVO) is a lower bound on the solution to constructive (respectively, destructive) OPT-BRIBERY, because in CVO and DVO there are no constraints on the number of each type of vote operations. We have the following four cases.

1. If the solution to CVO (DVO) is 0, then the solution to constructive (destructive) OPT-BRIBERY is also 0.
2. If the solution to CVO (DVO) is $\Theta(\sqrt{n})$, as n become large enough, with probability that goes to 1 each type of votes in the support of π will appear $\Theta(n)$, which is $> \Theta(\sqrt{n})$, times in the randomly generated preference-profile, which means that there are enough votes of each type for the strategic individual to change.
3. If the solution to CVO (DVO) is $\Theta(n)$, then the solution to constructive (destructive) OPT-BRIBERY is either $\Theta(n)$ (when the strategic individual can change *all* votes to achieve her goal), or ∞ .
4. If the solution to CVO (DVO) is ∞ , then the solution to constructive (destructive) OPT-BRIBERY is also ∞ .

It follows that the observation made in Theorem 1 holds for OPT-BRIBERY.

6.3 Margin of Victory (MoV)

Definition 6 *Given a voting rule r and a preference-profile P , the margin of victory (MoV) of P is the smallest number k such that the winner can be changed by changing k votes in P . In the MOV problem, we are given r and P , and are asked to compute the margin of victory.*

For a given instance (P, r) , MOV is equivalent to destructive OPT-BRIBERY, where $c = r(P)$. However, when the input preference-profiles are generated randomly, the winners in these profiles might be different. Therefore, the corollary of Theorem 1 for OPT-BRIBERY does not directly imply a similar corollary for MOV. This relationship is similar to the relationship between DVO and CWVO.

Despite this difference, the formulation of MOV in the vote operation framework is very similar to that of OPT-BRIBERY: The set of all operations and the argument to apply Theorem 1 are the same. The only difference is that for MOV, we obtain the corollary from the CWVO part of Theorem 1, while the corollary for OPT-BRIBERY is obtained from the CVO and DVO parts of Theorem 1.

6.4 Minimum Manipulation Coalition Size (MMCS)

The MINIMUM MANIPULATION COALITION SIZE (MMCS) problem is similar to MOV, except that in MMCS the winner must be improved for all voters who change their votes [19].

Definition 7 In an MMCS problem, we are given a voting rule r and a preference-profile P . We are asked to compute the smallest number k such that a coalition of k voters can change their votes to change the winner, and all of them prefer the new winner to $r(P)$.

Unlike MOV, MMCS falls under the vote operation framework in the following dynamic way. For each preference-profile, suppose c is the current winner. For each adversarial $d \neq c$, we use $\{f(V) - f(W) : V, W \in L(\mathcal{C}) \text{ s.t. } d \succ_W c \text{ and } \pi(W) > 0\}$ as the set of operations. That is, we only allow voters who prefer d to c to participate in the manipulative coalition. We also replace each of LP_{\succeq} and $\text{LP}_{\succeq^* \ominus \ominus_{\succeq} \pi}$ by multiple LPs, each of which is indexed by a pair of alternatives (d, c) and the constraints are generated by using the corresponding set of operations. Then, the corollary for MMCS follows after a similar argument to that of CVO in Theorem 1.

6.5 Control by Adding/Deleting Votes (CAV/CDV)

Definition 8 In a constructive (respectively, destructive) OPTIMAL CONTROL BY ADDING VOTES (OPT-CAV) problem, we are given a preference-profile P , a (dis)favoured alternative $c \in \mathcal{C}$, and a set N' of additional votes. We are asked to compute the smallest number k such that the strategic individual can add k votes in N' such that c is the winner (respectively, c is not the winner).

For simplicity, we assume that $|N'| = n$ and the votes in N' are drawn i.i.d. from a distribution π' . To show how OPT-CAV falls under the vote operation model, we let the set of operations to be the generalized scoring vectors of all votes that are in the support of π' , that is, $\{f(V) : V \in L(\mathcal{C}) \text{ and } \pi'(V) > 0\}$. Then, the corollary follows from the CVO and DVO parts of Theorem 1 via a similar argument to the argument for OPT-BRIBERY.

Definition 9 In a constructive (respectively, destructive) OPTIMAL CONTROL BY DELETING VOTES (OPT-CDV) problem, we are given a preference-profile P and a (dis)favoured alternative $c \in \mathcal{C}$. We are asked to compute the smallest number k such that the strategic individual can delete k votes in P such that c is the winner (respectively, c is not the winner).

To show how OPT-CDV falls under the vote operation framework, we let the set of operations to be the negation of generalized scoring vectors of votes in the support of π' , that is, $\{-f(V) : V \in L(\mathcal{C}) \text{ and } \pi'(V) > 0\}$. Then, the corollary follows from the CVO and DVO parts of Theorem 1 via a similar argument to the argument for OPT-BRIBERY.

6.6 A Unified Corollary

The next corollary of Theorem 1 summarizes the results obtained for all types of strategic behavior studied in this section.

Corollary 1 For any generalized scoring rule (or equivalently, any voting rule that satisfies anonymity, homogeneity, and finite local consistency), any distribution π over votes, and any $X \in (\{\text{constructive, destructive}\} \times \{\text{UCO, OPT-BRIBERY, OPT-CAV, OPT-CDV}\}) \cup \{\text{MOV, MMCS}\}$, suppose the input preference-profiles are generated i.i.d. from π .⁴ Then, for any alternative c and any $\epsilon > 0$, there exists $\beta^* > 1$ such that the total probability for the solution to X to be one of the following four cases is more than $1 - \epsilon$ as n goes to infinity: (1) 0, (2) between $\frac{1}{\beta^*} \sqrt{n}$ and $\beta^* \sqrt{n}$, (3) between $\frac{1}{\beta^*} n$ and $\beta^* n$, and (4) ∞ .

⁴For CAV, the distribution over the new votes can be generated i.i.d. from a different distribution π' .

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] John Bartholdi, III, Craig Tovey, and Michael Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.
- [3] John Bartholdi, III, Craig Tovey, and Michael Trick. How hard is it to control an election? *Math. Comput. Modelling*, 16(8-9):27–40, 1992. Formal theories of politics, II.
- [4] Felix Brandt. Some remarks on Dodgson’s voting rule. *Mathematical Logic Quarterly*, 55:460–463, 2009.
- [5] David Cary. Estimating the margin of victory for instant-runoff voting. In *Proceedings of 2011 EVT/WOTE Conference*, 2011.
- [6] Vincent Conitzer and Makoto Yokoo. Using mechanism design to prevent false-name manipulations. *AI Magazine*, 31(4):65–77, 2010.
- [7] William J. Cook, Albertus M. H. Gerards, Alexander Schrijver, and Eva Tardos. Sensitivity theorems in integer linear programming. *Mathematical Programming*, 34(3):251–264, 1986.
- [8] Shahar Dobzinski and Ariel D. Procaccia. Frequent manipulability of elections: The case of two voters. In *Proceedings of the Fourth Workshop on Internet and Network Economics (WINE)*, pages 653–664, Shanghai, China, 2008.
- [9] Piotr Faliszewski, Edith Hemaspaandra, and Lane A. Hemaspaandra. How hard is bribery in elections? *Journal of Artificial Intelligence Research*, 35:485–532, 2009.
- [10] Piotr Faliszewski, Edith Hemaspaandra, and Lane A. Hemaspaandra. Using complexity to protect elections. *Commun. ACM*, 53:74–82, 2010.
- [11] Piotr Faliszewski, Edith Hemaspaandra, and Lane A. Hemaspaandra. Multimode control attacks on elections. *Journal of Artificial Intelligence Research*, 40:305–351, 2011.
- [12] Piotr Faliszewski and Ariel D. Procaccia. AI’s war on manipulation: Are we winning? *AI Magazine*, 31(4):53–64, 2010.
- [13] Ehud Friedgut, Gil Kalai, and Noam Nisan. Elections can be manipulated often. In *Proceedings of the Annual Symposium on Foundations of Computer Science (FOCS)*, pages 243–249, 2008.
- [14] Allan Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41:587–601, 1973.
- [15] Marcus Isaksson, Guy Kindler, and Elchanan Mossel. The Geometry of Manipulation: A Quantitative Proof of the Gibbard-Satterthwaite Theorem. In *Proceedings of the 51st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 319–328, Washington, DC, USA, 2010.
- [16] Thomas R. Magrino, Ronald L. Rivest, Emily Shen, and David Wagner. Computing the Margin of Victory in IRV Elections. In *Proceedings of 2011 EVT/WOTE Conference*, 2011.
- [17] Elchanan Mossel and Miklos Z. Racz. A quantitative Gibbard-Satterthwaite theorem without neutrality. In *Proceedings of the The 44th ACM Symposium on Theory of Computing (STOC)*, 2012.

- [18] Bezalel Peleg. A note on manipulability of large voting schemes. *Theory and Decision*, 11:401–412, 1979.
- [19] Geoffrey Pritchard and Mark C. Wilson. Asymptotics of the minimum manipulating coalition size for positional voting rules under impartial culture behaviour. *Mathematical Social Sciences*, 1:35–57, 2009.
- [20] Ariel D. Procaccia and Jeffrey S. Rosenschein. Average-case tractability of manipulation in voting via the fraction of manipulators. In *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 718–720, Honolulu, HI, USA, 2007.
- [21] Jörg Rothe and Lena Schend. Typical-Case Challenges to Complexity Shields That Are Supposed to Protect Elections Against Manipulation and Control: A Survey. In *International Symposium on Artificial Intelligence and Mathematics*, 2012.
- [22] Mark 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.
- [23] Arkadii Slinko. On asymptotic strategy-proofness of classical social choice rules. *Theory and Decision*, 52:389–398, 2002.
- [24] Arkadii Slinko. How large should a coalition be to manipulate an election? *Mathematical Social Sciences*, 47(3):289–293, 2004.
- [25] Philip B. Stark. Conservative statistical post-election audits. *The Annals of Applied Statistics*, 2(2):550–581, 2008.
- [26] Philip B. Stark. A sharper discrepancy measure for post-election audits. *The Annals of Applied Statistics*, 2(3):982–985, 2008.
- [27] Philip B. Stark. Efficient post-election audits of multiple contests: 2009 California tests. In *4th Annual Conference on Empirical Legal Studies (CELS)*, 2009.
- [28] Philip B. Stark. Risk-limiting post-election audits: P-values from common probability inequalities. *IEEE Transactions on Information Forensics and Security*, 4:1005–1014, 2009.
- [29] Philip B. Stark. Super-simple simultaneous single-ballot risk-limiting audits. In *Proceedings of 2010 EVT/WOTE Conference*, 2010.
- [30] T. Nicolaus Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4(3):185–206, 1987.
- [31] Lirong Xia. Computing the margin of victory for various voting rules. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, Valencia, Spain, 2012.
- [32] Lirong Xia and Vincent Conitzer. Generalized scoring rules and the frequency of coalitional manipulability. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 109–118, Chicago, IL, USA, 2008.
- [33] Lirong Xia and Vincent Conitzer. A sufficient condition for voting rules to be frequently manipulable. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 99–108, Chicago, IL, USA, 2008.

- [34] Lirong Xia and Vincent Conitzer. Finite local consistency characterizes generalized scoring rules. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI)*, pages 336–341, Pasadena, CA, USA, 2009.
- [35] Lirong Xia, Vincent Conitzer, and Ariel D. Procaccia. A scheduling approach to coalitional manipulation. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 275–284, Cambridge, MA, USA, 2010.
- [36] Michael Zuckerman, Omer Lev, and Jeffrey S. Rosenschein. An algorithm for the coalitional manipulation problem under maximin. In *Proceedings of the Tenth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, Taipei, Taiwan, 2011.
- [37] Michael Zuckerman, Ariel D. Procaccia, and Jeffrey S. Rosenschein. Algorithms for the coalitional manipulation problem. *Artificial Intelligence*, 173(2):392–412, 2009.

A Proofs

Lemma 3 Fix an integer matrix \mathbf{A} . There exists a constant $\beta_{\mathbf{A}}$ that only depends on \mathbf{A} , such that if the following LP has an integer solution, then the solution is no more than $\beta_{\mathbf{A}} \cdot \|\vec{b}\|_{\infty}$.

$$\min \|\vec{x}\|_1, \text{ s.t. } \mathbf{A} \cdot \vec{x} \geq \vec{b} \quad (1)$$

Proof of Lemma 3: Let \mathbf{A} be a $m^* \times n^*$ integer matrix, which includes the constraints $\vec{x} \geq \vec{0}$. Suppose LP (1) has a (non-negative) integer solution. We note that $\vec{0}$ is an optimal integer solution to $\min \vec{1} \cdot (\vec{x})'$ s.t. $\mathbf{A} \cdot \vec{x} \geq \vec{0}$. Then, it follows from Theorem 5 (ii) in [7] that LP (1) has a (non-negative) integer solution \vec{z} such that

$$\|\vec{z} - \vec{0}\|_{\infty} \leq n^* \cdot M(\mathbf{A}) \cdot (\|\vec{b} - \vec{0}\|_{\infty} + 2),$$

where $M(\mathbf{A})$ is the maximum of the absolute values of the determinants of the square sub-matrices of \mathbf{A} . Since \mathbf{A} is fixed, the right hand side becomes a constant, that is, $\|\vec{z}\|_{\infty} = O(\|\vec{b}\|_{\infty})$. Therefore, there exists $\beta_{\mathbf{A}}$ such that the optimal value in the ILP (1) is no more than $\vec{1} \cdot (\vec{z})' \leq n^* \|\vec{z}\|_{\infty} \leq \beta_{\mathbf{A}} \cdot \|\vec{b}\|_{\infty}$. ■

Claim 1 There exists $N \in \mathbb{N}$ and $\beta' > 1$ such that for any $n \geq N$, any $P \in \mathcal{P}_{\beta}$, if (1) c is not the winner for P , and (2) there exists a refinement \succeq^* of $\succeq_{\pi} = \text{Ord}(f(P_{\pi}))$ such that $g(\succeq^*) = c$ and $LP_{\succeq^* \ominus \succeq_{\pi}}$ has an integer solution, then $\frac{1}{\beta'} \sqrt{n} < \text{CVO}(P, c) < \beta' \sqrt{n}$.

Proof of Claim 1: Let $\succeq = \text{Ord}(f(P))$. Because $\text{GS}(f, g)(P) \neq c$, $g(\succeq) \neq c$. Therefore, the strategic individual has to change the order of some components in the generalized scoring vector to make c win. We note that $P \in \mathcal{P}_{\beta}$, which means that the difference between any pair of components of $f(P)$ is more than $\frac{1}{\beta} \sqrt{n}$. Let d_{max} denote the maximum difference between any pair of components in generalized score vectors. That is, $d_{max} = \max_{t, t', L \in L(C)} \{(f(L))_t - (f(L))_{t'}\}$. In order for c to win, the number of vote operations must be at least $\frac{1}{\beta} \sqrt{n} / d_{max}$. Therefore, $\text{CVO}(P, c) > \frac{1}{\beta d_{max}} \sqrt{n}$.

We next show the upper bound. Because $P \in \mathcal{P}_{\beta}$, for any pair o_i, o_j with $o_i =_{\succeq_{\pi}} o_j$, $|[f(P)]_i - [f(P)]_j| < \beta \sqrt{n}$. Therefore, the right hand side of each (in)equality in $LP_{\succeq^* \ominus \succeq_{\pi}}$ is no more than $\beta \sqrt{n}$. Applying Lemma 3 to $LP_{\succeq^* \ominus \succeq_{\pi}}$, we have that there exists a constant $\beta_{\succeq^*, \succeq_{\pi}}$ that only depends on \succeq^* and \succeq_{π} , and an integer solution \vec{v} with $\|\vec{v}\|_1 \leq \beta_{\succeq^*, \succeq_{\pi}} \sqrt{n}$ (the \mathbf{A} matrix in Lemma 3 is fixed because we fix the number of alternatives m , and the left hand side of each (in)equality in $LP_{\succeq^* \ominus \succeq_{\pi}}$ does not depend on n). let β' be the maximum of $d_{max} \beta$ and all $\beta_{\succeq^*, \succeq_{\pi}}$ (since we fix the number of alternatives, there are finite number of $\beta_{\succeq^*, \succeq_{\pi}}$'s). Since $\beta_{\succeq^*, \succeq_{\pi}} > \frac{1}{d_{max} \beta}$, $\beta' > 1$. It follows that $\|\vec{v}\|_1 < \beta' \sqrt{n}$. We next show that

for sufficient large n , if the strategic individual apply \vec{v} , then the order over components of the total scoring vector will become \succeq^* . That is, c can be made win.

The idea is, $LP_{\succeq^* \ominus \succeq_\pi}$ ensures that by applying \vec{v} , ties between the “almost tied” components are broken as in \succeq^* . Since $\|\vec{v}\|_1 = O(\sqrt{n})$, when n is large enough the order between any pair of “far away” components will not be affected. Formally, let $\vec{x} = f(P) + \Delta \cdot (\vec{v})'$. That is, \vec{x} is the total generalized scoring vector after the strategic individual applied \vec{v} . Because \vec{v} is a solution to $LP_{\succeq^* \ominus \succeq_\pi}$, for any pair o_i, o_j with $o_i \simeq_{\succeq_\pi} o_j$, the order between o_i and o_j in \succeq^* is the same as the order between o_i and o_j in $\text{Ord}(\vec{x})$. Since \succeq^* is an extension of \succeq_π , if $o_i \triangleright_\pi o_j$, then we must have $o_i \triangleright^* o_j$. Therefore, we only need to check that for any $o_i \triangleright_\pi o_j$, we have $[\vec{x}]_i > [\vec{x}]_j$. Because $P \in \mathcal{P}_\beta$, $[f(P)]_i - [f(P)]_j > \frac{1}{\beta}n$. We note that $\|\vec{v}\|_1 < \beta' \sqrt{n}$, which means that by applying \vec{v} , the strategic individual can only change the difference between any pair of components by no more than $d_{max} \beta' \sqrt{n}$. Let $N = (d_{max} \beta' \beta)^2 + 1$. When $n \geq N$, $d_{max} \beta' \sqrt{n} < \frac{1}{\beta}n$, which means that for any $o_i \triangleright_\pi o_j$, applying \vec{v} will not change the order between o_i and o_j in the total generalized scoring vector. This means that by applying \vec{v} , the strategic individual can make c win. Therefore, $\text{CVO}(P, c) \leq \|\vec{v}\|_1 < \beta' \sqrt{n}$. It follows that for any $n \geq N$, $\frac{1}{\beta'} \sqrt{n} < \text{CVO}(P, c) < \beta' \sqrt{n}$. ■

Claim 2 *There exists $\beta' > 1$ such that for any $P \in \mathcal{P}_\beta$, if (1) c is not the winner for P , (2) there does not exist a refinement \succeq^* of $\succeq_\pi = \text{Ord}(f(P_\pi))$ such that $LP_{\succeq^* \ominus \succeq_\pi}$ has an integer solution, and (3) there exists \succeq such that $g(\succeq) = c$ and LP_\succeq has an integer solution, then $\frac{1}{\beta'}n < \text{CVO}(P, c) < \beta'n$.*

Proof of Claim 2: Let $\succeq = \text{Ord}(f(P))$. Because the premises of Claim 1 do not hold, the strategic individual has to change the order of some pair of components that are “far away” (that is, the difference between them is $\Theta(n)$ before the strategic individual applies vote operations) to make c win. We note that one operation can only change the difference between a pair of components by at most d_{max} . Therefore, $\text{CVO}(P, c) \geq \frac{1}{\beta}n/d_{max}$.

On the other hand, it follows from Lemma 1 and condition (3) in the statement of the claim that $\text{CVO}(P, c) < \infty$. The only thing left to show is that there exists $\beta' > 1$ such that $\text{CVO}(P, c) < \beta'n$ for all $P \in \mathcal{P}_\beta$ and for all n . Because $P \in \mathcal{P}_\beta$, for any pair o_i, o_j , $|[f(P)]_i - [f(P)]_j| < \beta n$. Applying Lemma 3 to LP_\succeq and condition (3) in the statement of the claim, we have that for any \succeq , there exists β_\succeq that only depends on \succeq , and an optimal integer solution \vec{v} such that $\|\vec{v}\|_1 \leq \beta_\succeq n$. Let β' be the maximum of $d_{max} \beta$ and all β_\succeq (again, there are finite number of β_\succeq 's). It follows that $\frac{1}{\beta'}n < \text{CVO}(P, c) < \beta'n$. ■

B Discussion: How often the solution is 0 or ∞ ?

One important question is: how large is the probability that the solution to problems studied in this section is 0 or ∞ ? Not surprising the answer depends on both the voting rule and the type of vote operations. The probability can be large for some voting rules. For example, for any voting rule that always selects a given alternative d as the winner, the solution to CVO (respectively, DVO) is always 0 (respectively ∞) for $c = d$ and is always ∞ (respectively 0) for $c \neq d$. However, for common voting rules the alternatives are treated almost equally (except for cases with ties). Therefore, we may expect that for a preference-profile whose votes are generated i.i.d., each alternative has almost the same probability of being selected as the winner. This is indeed the case in all commonly used voting rules, including approval voting, all positional scoring rules (which include Borda, plurality, and veto), plurality with runoff, Bucklin, Copeland, maximin, STV, and ranked pairs. Therefore, for these voting rules, when the votes are drawn i.i.d. uniformly at random, the probability for CVO is approximately $\frac{1}{m}$ and the probability for DVO is approximately $\frac{m-1}{m}$. For CWVO, the answer is never 0 because changing 0 vote cannot change the winner.

We would also expect for common voting rules, for some types of strategic behavior studied in this

section, with low probability the solution is ∞ . For UCO, the strategic individual can introduce many (but finite) votes such that the non-manipulators' votes are negligible. For OPT-BRIBERY and MOV, the strategic individual can change all votes to achieve her goal. For MMCS, CAV, and CDV, it is not clear how large such probability is. The following table summarizes folklore results for common voting rules when votes are drawn i.i.d. uniformly at random.

Optimal solution is	0	∞
$\{\text{Cons.}\} \times \{\text{UCO, OPT-BRIBERY}\}$	$\approx \frac{1}{m}$	0
$\{\text{Des.}\} \times \{\text{UCO, OPT-BRIBERY}\}$	$\approx \frac{m-1}{m}$	0
MOV	0	0

Table 1: Probability for solutions to some problems to be 0 or ∞ for common voting rules.

More generally, it would be very interesting to study the probability for the solution to the problems studied in this section to fall in each of the four cases $(0, \Theta(\sqrt{n}), \Theta(n), \infty)$. We do not see much hope for a general results that apply to all types of strategic behavior and all generalized scoring rules. We leave this as a future direction.