Computing the Margin of Victory for Various Voting Rules

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

The margin of victory of an election, defined as the smallest number k such that k voters can change the winner by voting differently, is an important measurement for robustness of the election outcome. It also plays an important role in implementing efficient post-election audits, which has been widely used in the United States to detect errors or fraud caused by malfunctions of electronic voting machines.

In this paper, we investigate the computational complexity and (in)approximability of computing the margin of victory for various 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. We also prove a dichotomy theorem, which states that for all *continuous generalized scoring rules*, including all voting rules studied in this paper, either with high probability the margin of victory is $\Theta(\sqrt{n})$, or with high probability the margin of victory is $\Theta(n)$, where n is the number of voters. Most of our results are quite positive, suggesting that the margin of victory can be efficiently computed. This sheds some light on designing efficient post-election audits for voting rules beyond the plurality rule.

Categories and Subject Descriptors: J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*; I.2.11 [**Distributed Artificial Intelligence**]: Multiagent Systems

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Computational social choice, margin of victory, generalized scoring rules

1. INTRODUCTION

The *margin of victory* is an important measurement in voting for robustness of the election outcome. Given a voting rule and a collection of votes, the margin of victory is the smallest number k such that changing k votes can change the winners. In that sense, an election with a large margin of victory is usually thought to be more robust than an election with a small margin of victory.

In addition to being interesting in its own right, margin of victory also plays an important role in conducting efficient *post-election audits*, which are nowadays a standard method used in the United States to detect incorrect outcome of electronic voting caused by software or hardware problems of voting machines [Norden et al. 2007]. When voters use voting machines to cast votes, their votes might not be correctly counted electronically, due to various problems including software bugs, programming mistakes, other machine output errors, or even some clip-on devices that manipulate the memory of voting machines [Wolchok et al. 2010]. In fact, according to [Norden et al. 2007], at least thirty states in the US have reported such problems by 2007. Post-election audits require *voter-verifiable paper records*, that is, when a voter casts her vote on a voting machine, a paper record of the vote is stored for possible future auditing. After all voters have cast their votes, some randomly selected electronic votes are compared to their paper copies manually to decide whether the election outcome is trustworthy. If there are too many mismatches indicating that with high probability the election outcome might be wrong, then a full recount (which is extremely costly) would occur. Since manually checking part of paper records may still be costly, *risk-limiting audit methods* have been proposed to audit as few votes as possible, while limiting the risk to a low level [Stark 2008a,b, 2009b,a, 2010]. In risk-limiting audit methods, the margin of victory is a critical parameter of risk guarantees. The larger the margin of victory is, the less votes are needed to be manually checked.

So far, most post-election auditing techniques were designed specifically for elections that use the *plurality rule*, where each voter votes for a single alternative, and the alternative with most votes (the number is called plurality score) wins. For the plurality rule, the margin of victory is very easy to compute, which is the difference between the highest and the second-highest plurality scores divided by 2. Therefore, risk-limiting audit methods can be applied with low computational costs for plurality. While the plurality rule is the most widely-used voting rule in political elections at the moment, it is important to extend the study to other popular voting rules. For example, *alternative vote (AV)* (a.k.a. *instant runoff voting (IRV)* or *single transferable vote (STV)*) is currently being used in political elections in Australia, India, and Ireland. Moreover, in 2011, the United Kingdom held a nation-wide referendum (the second nationwide referendum in the history) to change their voting system from plurality to STV.[1] To extend risk-limiting audit methods to other voting systems, naturally, a first step is to compute the margin of victory. For example, [Sarwate et al. 2011] proposed a risk-limiting audit method for STV based on the computation of the margin of victory. However, some recent work suggests that at least for STV, computing the margin of victory is much harder than for plurality [Magrino et al. 2011; Cary 2011], and it was conjectured that computing the margin of victory for STV is NP-complete [Magrino et al. 2011].

1.1. Our Contributions

In this paper, we investigate the computational complexity and (in)approximability of computing the margin of victory (MOV) for various 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. Our results are summarized in Table I. It can be seen from the table that deciding whether the margin of victory is 1 is NP-complete for ranked pairs and STV, which proves the conjecture proposed in [Magrino et al. 2011]. Deciding whether the margin of victory is larger than some given number is NP-complete for Copeland and maximin, but the problem can be solved in polynomial-time (in theory) if the margin of victory is bounded from above by a constant. For other rules, we design polynomial-time algorithms. For the four voting rules where exact computation is NP-hard (Copeland, maximin, STV, and ranked pairs), we show that for Copeland (respectively, maximin), there is a polynomial-time $\Theta(\log m)$ (respectively, 2)-approximation algorithm (where m is the number of alternatives), and there is no polynomial-time 2-approximation algorithm for STV or ranked pairs, unless P=NP. We also reveal a connection between the margin of victory and the destructive *unweighted coalitional optimization* (UCO) problem [Zuckerman et al. 2009], which allows us to convert an approximation algorithm for MOV to an approximation algorithm for destructive UCO, and vice versa. Finally, we ask typically how large the margin of victory is, when the votes are drawn i.i.d. Our main result is the following dichotomy theorem, which states that for a large class of voting rules called *continuous generalized scoring rules* [Xia and Conitzer 2008] (we will give formal definitions in Section 6.1), either with high probability the margin of victory is $\Theta(\sqrt{n})$, where n is the number of voters, or with high probability the margin of victory is $\Theta(n)$.

Theorem 14 Let r be a continuous generalized scoring rule and let π be a distribution over all linear orders, such that for each linear order V, $\pi(V) > 0$. Suppose we fix π and the number of alternatives, generate n votes i.i.d. according to π, and let P_n denote the profile. Then, one (and exactly one) of the following two observations holds.

(1) For any ϵ, there exists $\beta > 0$ such that $\lim_{n\to\infty} \Pr(\frac{1}{\beta}\sqrt{n} \leq \mathrm{MoV}(P_n) \leq \beta\sqrt{n}) \geq 1-\epsilon$.

(2) There exists $\beta > 0$ such that $\lim_{n\to\infty} \Pr(\mathrm{MoV}(P_n) \geq \beta n) = 1$.

It was shown in [Xia and Conitzer 2008] that all voting rules studied in this paper are generalized scoring rules. In this paper, we show that they are continuous generalized scoring rules. Due to the space constraint, some proofs are omitted. They can be found on the author's homepage.

[1] http://en.wikipedia.org/wiki/United_Kingdom_Alternative_Vote_referendum,_2011

Table I. Complexity of computing the margin of victory. In all NP-hardness proofs, m (the number of alternatives) and n (the number of voters) are not restricted.

Voting rule	Margin of Victory	
Approval voting	P	(Theorem 10)
Positional scoring rules (incl. Borda and plurality)	P	(Theorem 5)
Plurality with runoff	P	(Theorem 7)
Bucklin	P	(Theorem 6)
Copeland	$O(n^{\text{MOV}})$	(Theorem 8)
	NPC	(Theorem 4)
	$\Theta(\log m)$-apprx.	(Algorithm 4)
Maximin	$O(n^{\text{MOV}})$	(Theorem 9)
	NPC	(Theorem 3)
	2-apprx.	(Algorithm 5)
STV (a.k.a. IRV, AV)	NPC for MOV_1	(Theorem1)
	No 2-apprx.	(Corollary 1)
Ranked pairs	NPC for MOV_1	(Theorem2)
	No 2-apprx.	(Corollary 1)

1.2. Related Work and Discussion

When the voting rule always selects a single winner, deciding whether the margin of victory is larger than a given threshold k is identical to the destructive BRIBERY problem [Faliszewski et al. 2009], where we are given a profile, a disfavored alternative d, and a number k, and we are asked whether there is a way to change no more than k votes to make d not the unique winner. Despite the similarity in definitions, MOV and BRIBERY are motivated in completely different ways. For BRIBERY, hardness of computation is good news, because computational complexity can serve as a barrier against bribery.[2] On the other hand, we would like to compute MOV as fast as possible to extend risk-limiting audit methods beyond plurality. Therefore, negative results for "using computational complexity to protect elections" are positive results for computing the margin of victory. For example, it is commonly believed that strategic behavior (including manipulation and bribery) is typically easy to compute, which can be interpreted as evidence that computing the margin of victory is typically easy. Indeed, this is witnessed by our dichotomy theorem (Theorem 14). In addition, in BRIBERY, the briber usually only has partial information about voters' votes or preference, while in MOV, we have complete information about the electronic votes.

Moreover, from a technical point of view, to the best of our knowledge, little was known about the computational complexity of destructive BRIBERY. [3] Therefore, most of our results for MOV are also new results for destructive BRIBERY. Moreover, we also show that for STV and ranked pairs, constructive BRIBERY is NP-complete even for $k = 1$; and for Copeland and maximin, constructive BRIBERY is NP-complete.

The dichotomy theorem proved in this paper is related to two lines of research on different (yet closely related) problems. The first is the study of "frequency of manipulability", defined as the probability for a randomly generated profile to be manipulable by a group of manipulators, where the non-manipulators' votes are generated i.i.d. according to some distribution [Peleg 1979; Baharad and Neeman 2002; Slinko 2002, 2004; Procaccia and Rosenschein 2007; Xia and Conitzer 2008]. Specifically, in [Xia and Conitzer 2008], we showed that for generalized scoring rules, when the non-manipulators votes are drawn i.i.d., 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 for some distributions the probability

[2]This belongs to a popular research direction in Computational Social Choice, i.e., using computational complexity to protect elections. See [Faliszewski et al. 2010; Faliszewski and Procaccia 2010; Rothe and Schend 2012] for recent surveys.

[3]Constructive BRIBERY has been shown to be in P for plurality and veto [Faliszewski et al. 2009], which implies that destructive BRIBERY for plurality and veto is also in P.

that the manipulators can succeed goes to 1. This tackles the UCO problem, but it is not clear how this result can be generalized to MOV.

The second is the study of the "minimum manipulation coalition size" problem, which is defined similarly to MOV, except that all voters who change their votes must prefer the new winner to the old one [Nitzan 1982; Chamberlin 1985; Pritchard and Slinko 2006; Pritchard and Wilson 2007, 2009]. Specifically, in [Pritchard and Wilson 2009], 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.

2. PRELIMINARIES

Let $\mathcal{C}$ denote the set of *alternatives*, $|\mathcal{C}| = m$. We assume strict preference orders.[4] 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 set of non-empty winning alternatives. That is, $r : L(\mathcal{C})^* \to (2^{\mathcal{C}} \setminus \emptyset)$. Throughout the paper, we let n denote the number of votes and let m denote the number of alternatives.

For any profile P and any pair of alternatives (c, d), let $D_P(c, d)$ denote the number of times that $c \succ d$ in P minus the number of times that $d \succ c$ in P. The *weighted majority graph (WMG)* of P, denoted by $\text{WMG}(P)$, is a directed graph whose vertices are the alternatives, and there is an edge between each pair of vertices, where the weight on $c \to d$ is $D_P(c, d)$. We note that in the WMG of any profile, all weights on the edges have the same parity (and whether it is odd or even depends on the parity of the number of votes n), and $D_P(c, d) = -D_P(d, c)$.

In this paper, we study the following voting rules.

— **Approval**: Each voter approves a subset of alternatives. The alternative that is approved by most voters is the winner.
— **(Positional) scoring rules**: Given a *scoring vector* $\vec{s}_m = (\vec{s}_m(1), \ldots, \vec{s}_m(m))$ of m integers, for any vote $V \in L(\mathcal{C})$ and any $c \in \mathcal{C}$, let $\vec{s}_m(V, c) = \vec{s}_m(j)$, where j is the rank of c in V. For any profile $P = (V_1, \ldots, V_n)$, let $\vec{s}_m(P, c) = \sum_{j=1}^{n} \vec{s}_m(V_j, c)$. The rule will select $c \in \mathcal{C}$ so that $\vec{s}_m(P, c)$ is maximized. We assume score components in $\vec{s}_m$ are nonincreasing. Some examples of positional scoring rules are *Borda*, for which the scoring vector is $(m-1, m-2, \ldots, 0)$; *plurality*, for which the scoring vector is $(1, 0, \ldots, 0)$; and *veto*, for which the scoring vector is $(1, \ldots, 1, 0)$. When there are only two alternatives, Borda, plurality, and veto (as well as all other voting rules introduced below) become *majority*.
— **Copeland**$_\alpha$ ($0 \leq \alpha \leq 1$): For any two alternatives c and c', we can simulate a *pairwise election* between them, by seeing how many votes prefer c to c', and how many prefer c' to c; the winner of the pairwise election is the one preferred more often. Then, an alternative receives one point for each win in a pairwise election, α points for each tie, and zero point for each loss. This is the Copeland score of the alternative. A Copeland winner maximizes the Copeland score.
— **Maximin**: A maximin winner c maximizes the maximin score $S_M(P, c) = \min\{D_P(c, c') : c' \in \mathcal{C}, c' \neq c\}$.
— **(Simplified) Bucklin**: The Bucklin score of an alternative c, denoted by $S_B(P, c)$, is the smallest number t such that more than half of the votes rank c somewhere in the top t positions. A Bucklin winner minimizes the Bucklin score.
— **Plurality with runoff**: The rule has two steps. In the first step, all alternatives except the two that are ranked in the top positions the most often are eliminated; in the second round, the majority rule is used to select the winner.

[4] Approval voting is an exception, where a voter's preferences are represented by a subset of alternatives she approves.

— **Single transferable vote (STV), a.k.a. instant runoff voting (IRV) or alternative vote (AV)**: The election has m rounds. Fix a tie-breaking order among the alternatives. 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). When there are ties, we use the fixed tie-breaking mechanism to eliminate an alternative. The last-remaining alternative is the winner.
— **Ranked pairs**: This rule first creates an entire ranking of all the alternatives. In each step, we will consider a pair of alternatives c, c' that we have not previously considered; specifically, we choose the remaining pair with the highest $D_P(c, c')$. When there are ties among multiple pairs, we use a fix a tie-breaking mechanism to select one pair.[5] We then fix the order $c \succ c'$, unless this contradicts previous orders that we fixed (that is, it violates transitivity). We continue until we have considered all pairs of alternatives (hence we have a full ranking). The alternative at the top of the ranking wins.

For plurality with runoff, in cases where we do not mention the tie-breaking mechanism, we adopt the *parallel-universes tie-breaking* [Conitzer et al. 2009]. That is, an alternative is a winner, if it wins w.r.t. *some* tie-breaking mechanism. We study the computation of the margin of victory for STV and ranked pairs with fixed tie-breaking mechanism mainly because computing the winner(s) for these two rules under the parallel-universes tie-breaking mechanism is already NP-complete [Conitzer et al. 2009; Brill and Fischer 2012].

A voting rule r is *WMG-based*, if the winners only depend on the WMG of the input profile. That is, for any pair of profiles P_1, P_2, if $\text{WMG}(P_1) = \text{WMG}(P_2)$, then $r(P_1) = r(P_2)$. Borda, Copeland, maximin, and ranked pairs are WMG-based. To improve readability, we defer the definition of generalized scoring rules to Section 6.1. This will not hinder the presentation of the complexity results.

2.1. Margin of Victory, Bribery, and Manipulation

Definition 1 *Given a voting rule r and a profile P, the* margin of victory (MoV) *of P, denoted by MoV(P, r), is the smallest number k such that the set of winners can be changed by changing k votes in P, while keeping the other votes unchanged.*

In this paper, we sometimes use $\text{MoV}(P)$ to denote $\text{MoV}(P, r)$ when causing no confusion. We now define the computational problems.

Definition 2 *In the* MOV *problem, we are given a voting rule r and a profile P. We are asked to compute MoV(P).* MOV$_k$ *is the decision variant of* MOV*, where we are given a natural number k, and we are asked whether the margin of victory is at most k.*

MOV$_k$ is closely related to the BRIBERY problem [Faliszewski et al. 2009], defined as follows.

Definition 3 *In a constructive (respectively, destructive)* BRIBERY *problem, we are given a profile P composed of n votes, a quota $k < n$, and a (dis)favored alternative $d \in \mathcal{C}$. We are asked whether the briber can change no more than k votes such that d is the unique winner (respectively, d is not the unique winner).*

More precisely, for voting rules that always outputs a unique winner, destructive BRIBERY is equivalent to MOV$_k$. MOV is also closely related to the unweighted coalitional optimization problem [Zuckerman et al. 2009], defined as follows.

Definition 4 *In a constructive (respectively, destructive)* UNWEIGHTED COALITIONAL OPTIMIZATION (UCO) *problem, we are given a voting rule r, a profile P^{NM} of the non-manipulators, and a (dis)favored alternative $d \in \mathcal{C}$. We are asked to compute the smallest number of manipulators*

[5]This is is similar to the fixed tie-breaking mechanism for STV. The difference is that in STV, the order is among alternatives, while in ranked pairs, the order is among pairs of alternatives.

who can cast votes P^M such that $\{d\} = r(P^{NM} \cup P^M)$ (respectively, $\{d\} \neq r(P^{NM} \cup P^M)$). Constructive (respectively, destructive) UNWEIGHTED COALITIONAL MANIPULATION (UCM_k) *is the decision variant of* UCO*, where we are given k manipulators, and we are asked whether they can cast votes P^M such that $\{d\} = r(P^{NM} \cup P^M)$ (respectively, $\{d\} \neq r(P^{NM} \cup P^M)$).*

3. HARDNESS RESULTS

For STV and ranked pairs, constructive UCM_1 is NP-complete [Bartholdi and Orlin 1991; Xia et al. 2009]. The next two theorems prove that for both rules MOV_1 and constructive (destructive) BRIBERY are NP-complete, by showing reductions from constructive UCM_1.

Theorem 1 *It is* NP*-complete to compute* MOV_1 *and constructive (destructive)* BRIBERY *for STV. This holds for* any *fixed tie-breaking mechanism.*

Proof of Theorem 1: It is easy to check that MOV_1 for STV is in NP. We prove the NP-hardness by a reduction from a special constructive UCM_1 problem for STV, where c is ranked in the top position in at least one vote in P^{NM}. This problem has been shown to be NP-complete [Bartholdi and Orlin 1991]. For any constructive UCM_1 instance (STV, P^{NM}, c) where c is ranked in the top positions in at least one vote in P^{NM} ($|P^{NM}| = n-1$), we construct the following MOV_1 instance. Let $\mathcal{C}' = \{c, c_1, \ldots, c_{m-1}\}$ denote the set of alternatives in the constructive UCM_1 instance.

Alternatives: $\mathcal{C}' \cup \{d\}$, where d is an auxiliary alternative.

Profile: Let P denote a profile of $2n-1$ votes as follows. The first $n-1$ votes are obtained from P^{NM} by putting d right below c. The next n votes ranks d in the first position (other alternatives are ranked arbitrarily).

It is easy to check that STV$(P) = \{d\}$. Suppose the constructive UCM_1 instance has a solution, denoted by V. Then, let V' denote the linear order over $\mathcal{C}' \cup \{d\}$ obtained from V by ranking d in the bottom position. Let P' denote the profile where voter n changes her vote to V'. We note that d is ranked in the top position for $n-1$ time in P'. Therefore, d is never eliminated in the first $|\mathcal{C}'| - 1$ rounds. Moreover, for any $j \leq |\mathcal{C}'| - 1$, the alternative that is eliminated in the jth round for P' is exactly the same as the alternative that is eliminated in the jth round for $P^{NM} \cup \{V\}$. In the last round, c is ranked in the top for n times, which means that STV$(P') = \{c\} \neq \{d\}$. Hence, $MoV(P) = 1$.

On the other hand, suppose $MoV(P) = 1$. Then, there exists a voter n' and a vote $V'_{n'}$ such that $STV(P_{-n'}, V'_{n'}) \neq \{d\}$. Let $P' = (P_{-n'}, V'_{n'})$. We note that in STV for P', d must be eliminated in the last round, because d is ranked in the top position for at least $n-1$ times. Moreover, we recall that c is ranked in the first position in at least one vote in P^{NM}, and d is ranked right below c in the corresponding vote in P'. Therefore, d beats all alternatives in $\mathcal{C}' \setminus \{c\}$ in their pairwise elections, which means that in the last round the only remaining alternatives must be c and d. This only happens when $d \succ c$ in $V_{n'}$ and $c \succ d$ in $V'_{n'}$. Therefore, w.l.o.g. we can let $n' = n$. Let V be a linear order obtained from $V'_{n'}$ by removing d. It follows that V is a solution to the constructive UCM_1 instance. This shows that it is NP-complete to compute MOV_1.

For constructive (respectively, destructive) BRIBERY we ask whether we can bribe one voter to make c win (respectively, to make d not the unique winner). ■

Theorem 2 *It is* NP*-complete to compute* MOV_1 *and constructive (destructive)* BRIBERY *for ranked pairs with some fixed tie-breaking mechanism.*

Proof of Theorem 2: It is easy to check that MOV_1 for ranked pairs is in NP. We prove the NP-hardness by a reduction from a special constructive UCM_1 problem for ranked pairs, where n is odd (we note that $|P^{NM}| = n-1$), and no weight in the majority graph is larger than $n-5$. Let $\mathcal{C}' = \{c, c_1, \ldots, c_{m-1}\}$ denote the set of alternatives in the constructive UCM_1 instance. Let $I = [c_1 \succ c_2 \succ \cdots \succ c_{m-1}]$ and $R = [c_{m-1} \succ c_{m-2} \succ \cdots \succ c_1]$. If the weight on some edge is larger than $n-5$, then we tweak the instance by adding two pairs of $\{[c \succ I], [R \succ c]\}$. This special constructive UCM_1 has been shown to be NP-complete [Xia et al. 2009]. For any such a constructive UCM_1 instance (RP, P^{NM}, c), we construct the following instance of MOV_1.

Alternatives: $\mathcal{C}' \cup \{d, e\}$, where d and e are auxiliary alternatives.

Profile: Let P denote a profile of $3n - 2$ votes as follows.

• The first $n - 1$ votes are obtained from P^{NM} by putting $d \succ e$ right below c.

• The remaining votes are defined in the following table.

Number of votes	Preferences
n	$d \succ e \succ I \succ c$
$(n-1)/2$	$d \succ c \succ e \succ R$
$(n-1)/2$	$e \succ c \succ d \succ R$

Let P' denote the profile obtained from P by removing one vote of $d \succ e \succ I \succ c$. We make the following observations on the weighted majority graph of P'.

• The sub-graph for alternatives in $\mathcal{C}'$ is the same as the weighted majority graph of the constructive UCM_1 instance.

• There is an edge from d to e with weight $2(n-1)$.

• There are no edges between d and c, and e and c.

• The weights on the edges from d or e to $\mathcal{C}' \setminus \{c\}$ is at least $n - 1$.

Therefore, in the final ranking, it is fixed that $d \succ e \succ (\mathcal{C}' \setminus \{c\})$. Suppose ties among edges are broken in the order where $e \to c$ is fixed before $c \to d$. It is not hard to see that the winner under ranked pairs for P is d. If the constructive UCM_1 instance has a solution, denoted by V, then, the margin of victory is 1. This is because if one voter whose vote was $[d \succ e \succ I \succ c]$ switches to $[c \succ e \succ d \succ V]$, then the winner becomes c.

On the other hand, suppose the margin of victory is 1, where voter j can cast a different vote V' to change the winner. We next show that voter j's vote must be $[d \succ e \succ I \succ c]$. We note that in P, $[d \succ e \succ c]$ in only one type of votes, which are $[d \succ e \succ I \succ c]$. Therefore, this is the only type of votes from which voter j can make c beats both d and e in their pairwise elections. In the tie-breaking mechanism we described above, if d or e beats c in their pairwise elections, then d will be the winner under ranked pairs. This proves that j's vote was $[d \succ e \succ I \succ c]$. Then, c is the winner under ranked pairs if and only if (1) both d and e are ranked below c in V', and (2) the vote obtained from V' by removing d and e is a solution to the constructive UCM_1 instance. Therefore, MOV_1 for ranked pairs is NP-complete to compute.

For constructive (respectively, destructive) BRIBERY we ask whether we can bribe one voter to make c win (respectively, to make d not the unique winner). ■

For STV and ranked pairs, a polynomial-time 2-approximation algorithm for MOV can be used to solve MOV_1 in polynomial-time. Therefore, we have the following inapproximability result.

Corollary 1 *There is no polynomial-time* 2*-approximation algorithm for* MOV *for STV or ranked pairs, unless* P= NP.

For maximin and Copeland, we will use *McGarvey's trick* [McGarvey 1953], which constructs a profile whose WMG is the same as some targeted weighted directed graph. This will be helpful because when we present the proof, we only need to specify the WMG instead of the whole profile, and then by using McGarvey's trick, a profile can be constructed in polynomial time. The trick works as follows. For any pair of alternatives (c, d), if we add the following pair of votes

$$[c_3 \succ \cdots \succ c_{\lceil m/2 \rceil + 1} \succ c \succ d \succ c_{\lceil m/2 \rceil + 2} \succ \cdots \succ c_m]$$
$$[c_m \succ \cdots \succ c_{\lceil m/2 \rceil + 2} \succ c \succ d \succ c_{\lceil m/2 \rceil + 1} \succ \cdots \succ c_3]$$

to a profile P, then in the WMG the weight on $c \to d$ is increased by 2 and the weight on $d \to c$ is decreased by 2, and the weights on other edges do not change. Moreover, c, d, and $c_{\lceil m/2 \rceil + 1}$ are ranked within top $\lceil m/2 \rceil + 2$ positions in the two votes. This observation (that we can ensure that for any given alternative e, when we apply McGarvey's trick, e is always ranked among top $\lceil m/2 \rceil + 2$ positions) will be useful in the proofs.

Theorem 3 *When k is a part of input, it is* NP*-complete to compute* MOV$_k$ *and constructive (destructive)* BRIBERY *for maximin.*

Proof of Theorem 3: We prove the theorem by a reduction from EXACT COVER BY 3-SETS (X3C) [Garey and Johnson 1979]. In an X3C instance, we are given two sets $\mathcal{A} = \{a_1, \ldots, a_q\}$ (where q is a multiple of 3) and $\mathcal{E} = \{E_1, \cdots, E_t\}$, where for each $E \in \mathcal{E}$, $E \subseteq \mathcal{A}$ and $|E| = 3$. We are asked whether there exist $q/3$ elements $\mathcal{E}' = \{E_{j_1}, \ldots, E_{j_{q/3}}\}$ in $\mathcal{E}$ such that each element in $\mathcal{A}$ appears in one and exactly one element in $\mathcal{E}'$. Given an X3C instance where w.l.o.g. $q > 16$, we construct the following MOV$_k$ instance for maximin.

Alternatives: $\mathcal{A} \cup \{c, d\}$. Let $k = q/3$.

Profile: The profile is composed of two parts P_1 and P_2, where P_1 encodes the X3C instance and P_2 is used to implement the McGarvey's trick. P_1 is composed of the following t votes: for each $j \leq t$, there is a vote $V_j = [d \succ (\mathcal{A} \setminus E_j) \succ c \succ E_j]$. P_2 is the profile such that in the WMG of $P_1 \cup P_2$, we have the following edges.

- $d \to c$ with weight $2q/3 - 1$.
- For every $i \leq q$, $d \to a_i$ with weight $2q/3 + 3$ and $a_i \to c$ with weight $2q/3 - 1$.

Moreover, when applying McGarvey's trick to obtain P_2, we always ensure that c is ranked within top $\lceil m/2 \rceil + 2 = \lceil q/2 \rceil + 4$ positions. Let $P = P_1 \cup P_2$.

Since d is the Condorcet winner in P, d is the maximin winner. Suppose the X3C instance has a solution, w.l.o.g. $\{E_1, \ldots, E_{q/3}\}$. Then, we change $V_1, \ldots, V_{q/3}$ to $[c \succ \text{Others}]$, and in the resulting profile, the maximin score of d is -1 (via c) and the maximin score of c is -1 (via any alternative in $\mathcal{A}$), which means that d is not the unique winner. Hence, the MoV is at most $q/3$.

If the MoV is at most $q/3$, then there is a way to change the outcome by changing $q/3$ votes. By changing $q/3$ votes, the weights on each edge cannot be changed by more than $2q/3$. Therefore, the maximin score of alternatives in $\mathcal{A}$ is no more than -3 (via d), and the maximin score of d is at least -1 (only possible via c), which means that only c can end up in a tie with d. For the maximin score of c to be -1, the weights on the edges from $\mathcal{A}$ to c must be reduced to -1 or less. Because when applying McGarvey's trick, c is always ranked within top $\lceil q/2 \rceil + 4 < q - 3$ positions, if any of the $q/3$ votes are in P_2, then there must exist an alternative $a \in \mathcal{A}$ such that the weight of $a \to c$ is at least 3. It follows that that $q/3$ votes correspond to an exact cover of $\mathcal{A}$. This proves the theorem.

For constructive (respectively, destructive) BRIBERY we ask whether we can bribe $q/3$ voters to make c win (respectively, to make d not the unique winner). ■

Theorem 4 *When k is a part of input, it is* NP*-complete to compute* MOV$_k$ *and constructive (destructive)* BRIBERY *for Copeland$_\alpha$ (for all $0 \leq \alpha \leq 1$).*

Proof of Theorem 4: The proof is similar to the proof of Theorem 3. The difference is that in the MOV$_k$ instance the set of alternatives is $\mathcal{A} \cup \{c, d, e\}$, and in the WMG of $P_1 \cup P_2$, we have the following edges.

- $d \to e$, $e \to c$, $c \to d$, $a_1 \to e$ with weight $2q/3 + 1$.
- For every $i \leq q$, there is an edge $d \to a_i$ with weight $2q/3 + 3$ and an edge $a_i \to c$ with weight $2q/3 - 1$.
- For every $i \leq q$, $d \to a_i$ with weight $2q/3 + 3$ and $a_i \to c$ with weight $2q/3 - 3$; and for a_i there exits $\lfloor q/2 \rfloor - 1$ incoming edges from other elements in $\mathcal{A}$ with weight $2q/3 + 3$.

It follows that d is not the unique winner when the $q/3$ votes correspond to an exact cover of $\mathcal{A}$. For constructive (respectively, destructive) BRIBERY we ask whether we can bribe $q/3$ voters to make c win (respectively, to make d not the unique winner). ■

4. POLYNOMIAL-TIME ALGORITHMS

For simplicity, we present polynomial-time algorithms for destructive BRIBERY, which is equivalent to MOV with unique winner. The algorithms can easily extend to general MOV problems. The corresponding ways to change the winners can also be computed easily.

Theorem 5 *Let r be a positional scoring rule. Algorithm 1 runs in polynomial time and computes* MOV *for r.*

The idea behind the algorithm is to check (for each k less than n and each "adversarial" alternative c) whether k voters can change their votes to reduce the score difference between the current winner d and c to 0. We note that c might not be a winner. If c's score is as high as d's, then d is not the unique winner (and the winner can be an alternative different from c and d).

Algorithm 1: MoVScoring

Input: A position scoring rule r and a profile P of n votes.
Output: The margin of victory for r.

Let $\{d\} = r(P)$.
for *any number* $k = 1 \to n$ **do**
    for *any alternative* $c \neq d$ **do**
        Rank the votes in P by the score of d minus the score of c in non-increasing order.
        Choose the top k votes and change them to $[c \succ \text{Others} \succ d]$.
        if *in the resulting profile d is not the winner* **then**
            Output that the margin of victory is k and terminate the algorithm.
        end
    end
end

Theorem 6 *Algorithm 2 runs in polynomial time and computes* MOV *for Bucklin.*

The idea behind the algorithm is to check (for each k less than n, each "adversarial" alternative c, and each targeted position $l \leq \lceil m/2 \rceil + 1$) whether k voters can change their votes to make the Bucklin score of c at most l while making the Bucklin score of the current winner at least l. We only need to check l up until $\lceil m/2 \rceil + 1$ because the Bucklin score of the Bucklin winner is at most $\lceil m/2 \rceil + 1$.

Algorithm 2: MoVBucklin

Input: A profile P of n votes.
Output: The margin of victory for Bucklin.

Let $\{d\} = r(P)$.
for *any* $k = 1 \to n$ *and any* $l = 1 \to \lceil m/2 \rceil + 1$ **do**
    for *any alternative* $c \neq d$ **do**
        For each vote in P, compute whether changing it to $[c \succ \text{Others} \succ d]$ increases the number of times c is ranked within top l positions and/or decreases the number of times d is ranked within top $l - 1$ positions.
        Compute whether k voters can make c ranked in top l positions in more than half of the votes, while making d ranked in top $l - 1$ positions in less than half of the votes.
        if *there exist such votes* **then**
            Output that the margin of victory is k and terminate the algorithm.
        end
    end
end

Theorem 7 *Algorithm 3 runs in polynomial time and computes* MOV *for plurality with runoff.*

Given a profile, for every alternative e, we let T_e denote the set of votes where e is ranked in the top positions. Let d denote the current winner. For each k from 1 to n, the algorithm does the following two checks in sequence:

— **Check 1:** We first check whether there is a way to convert k votes to make d not in the runoff. It suffices to focus on converting k votes in T_d. If there is a way to do so, then the margin of victory is at most k, and we can skip Check 2 below.

— **Check 2:** Otherwise d must be in the runoff. Then, the algorithm checks that for each adversarial c and the "threshold" of plurality scores $l \leq Plu(c) + k$, whether k voters can change their votes such that in the new profile, the following three conditions are satisfied.
  (1) The plurality scores of c is at least l.
  (2) The plurality score of any other alternative is no more than l.
  (3) c beats d in their pairwise election.

For Check 2, we partition the votes as follows. For any alternative e, let A_e denote the set of votes where e is ranked in the top position and where $d \succ c$; let B_e denote the set of votes where e is ranked in the top position and where $c \succ d$. In the algorithm, we try to first convert votes in A_e to meet Condition (2) (the plurality score of e in the new profile is no more than l). Given a profile, for any alternative e ($e \neq d$) and any l, we define $t_e^l = \max\{Plu(e) - l, 0\}$. That is, t_e^l is the least number of times that votes in T_e must be changed to meet Condition (2). For every k, every adversarial c, and every l such that $l \leq Plu(c) + k$, the algorithm always makes sure that Condition (2) is satisfied, and tries to increase the weight on the edge $c \to d$ in the WMG as much as possible. This is achieved by first trying to change votes in A_e to rank c in the top position, because each of such changes increases the weight on $c \to d$ by 2. If all votes in A_e are used up, then we have to change votes in B_e to reduce the plurality score of e, but changing votes in B_e does not increase the weight on $c \to d$. Let $k' = k - \sum_{e \in \mathcal{C} \setminus \{c,d\}} t_e^l$. If $k' \geq 0$, then we can change k' votes in T_d to rank c in the top position. Because Check 1 failed, c and d enter the runoff. Finally, let Q denote the resulting profile. We check whether c beats d in their pairwise election in the following way.

$$D_Q(c,d) = D_P(c,d) + 2 \cdot (k' + \sum_{e \in \mathcal{C} \setminus \{c,d\}} \min\{t_e^l, |A_e|\}) \tag{1}$$

If $D_Q(c,d) \geq 0$, then the margin of victory is at most k. If both checks fail for some specific k, then the margin of victory is more than k, and the algorithm continues to check $k+1$. Formally, we have the following algorithm.

Algorithm 3: MoVPluRunOff

Input: A profile P of n votes.
Output: The margin of victory for plurality with runoff.

Let $\{d\} = r(P)$.
Compute the partition A_e and B_e for each alternative e.
for *any* $k = 1 \to n$ **do**
  Check 1: Let c_1, c_2 denote the alternatives who have the highest plurality scores in $\mathcal{C} \setminus \{d\}$. If $\max\{Plu(d) - k - Plu(c_1), 0\} + \max\{Plu(d) - k - Plu(c_2), 0\} \leq k$, then output that the margin of victory is k and terminate the algorithm.
  Check 2: for *each* $c \neq d$ *and each* $l \leq Plu(c) + k$ **do**
    Compute $D_Q(c,d)$ as defined in Equation (1).
    if $D_Q(c,d) \geq 0$ **then**
      Output that the margin of victory is k and terminate the algorithm.
    end
  end
end

Theorem 8 *For any fixed number k,* MOV_k *for Copeland is in* P.

Theorem 9 *For any fixed number k,* MOV$_k$ *for maximin is in* P.

Both problems are solved by the following algorithm: for each adversarial alternative c, we enumerate all subsets of k voters and check if we can change the winner by ranking c in the top position and the current winner d in the bottom position.

Theorem 10 *There exists a polynomial-time algorithm that computes* MOV *for approval.*

The algorithm is very similar to Algorithm 1. We check for each adversarial c whether changing k votes to $\{c\}$ can make the current winner no longer a unique winner. We note that constructive BRIBERY is NP-complete [Faliszewski et al. 2009] for approval.

5. APPROXIMATION ALGORITHMS

In this section, we first present approximation algorithms for Copeland and maximin. Then, we study the relationship between MOV and destructive UCO. For simplicity of presentation, we assume that the input profile has a unique winner.

Algorithm 4: AppMoVCopeland

Input: A profile P of n votes.
Output: The margin of victory for Copeland$_\alpha$.

Let $\{d\} = r(P)$, $P^* = P, \mathcal{I} = \emptyset, Q = \emptyset$. $W = [c^* \succ \text{Others} \succ d]$
Compute RM(d, c) for every $c \neq d$, and let $c^* = \arg\min_c \text{RM}(d, c)$.
for *RM*(d, c^*) *iterations* **do**
    Let $\mathcal{C}_d = \{c : D_{P^* \cup Q}(d, c) > 0\}$.
    for $\lceil \log m \rceil + 1$ *rounds* **do**
        Find a vote $V_j \in P^*$ where $d \succ c$ holds for at least half of alternatives c in $\mathcal{C}_d$.
        $\mathcal{C}_d \leftarrow \mathcal{C}_d \setminus \{c : d \succ_{V_j} c\}, P^* \leftarrow P^* \setminus \{V_j\}, Q \leftarrow Q \cup \{W\}, \mathcal{I} \leftarrow \mathcal{I} \cup \{j\}$.
    end
    Let $\mathcal{C}_* = \{c : D_{P^* \cup Q}(c, c^*) > 0\}$.
    for $\lceil \log m \rceil + 1$ *rounds* **do**
        Find a vote $V_j \in P^*$ where $c \succ c^*$ holds for at least half of alternatives c in $\mathcal{C}_*$.
        $\mathcal{C}_d \leftarrow \mathcal{C}_d \setminus \{c : c \succ_{V_j} c^*\}, P^* \leftarrow P^* \setminus \{V_j\}, Q \leftarrow Q \cup \{W\}, \mathcal{I} \leftarrow \mathcal{I} \cup \{j\}$.
    end
end
return $|\mathcal{I}|$.

5.1. Copeland

We first present a polynomial-time $2(\lceil \log m \rceil + 1)$-approximation algorithm for MOV for Copeland$_\alpha$. The idea is the following. Let d denote the current winner whose Copeland score is denoted by $s_C(P^{NM}, d)$. For any profile, any alternative c, and any number t, we let

$$s'_t(P^{NM}, c) = |\{c' : c' \neq c, D_P(c', c) < 2t\}| + \alpha \cdot |\{c' : c' \neq c, D_P(c', c) = 2t\}|.$$

That is, $s'_t(P^{NM}, c)$ is the Copeland score of c if for every $c' \neq c$, we increase the weight on $c \to c'$ in the WMG by $2t$, where t can be negative. We will use this value as a lower bound on MoV(P^{NM}), as it captures the effect of ranking c first instead of last in t votes. For every $c \neq d$, we compute a *relative margin* RM(d, c) between d and c, defined as the minimum t such that $s'_{-t}(P^{NM}, d) \leq s'_t(P^{NM}, c)$. Let c^* be the alternative that has the smallest relative margin from d. It follows that MoV(P^{NM}) is at least RM(d, c^*), because by changing t votes, the Copeland score of d cannot be less than $s'_{-t}(P^{NM}, d)$ and the Copeland score of c cannot be more than $s'_t(P^{NM}, c)$. Moreover, in the WGM the weight on any edge cannot be changed by more than $2t$. Algorithm 4 finds $2(\lceil \log m \rceil + 1) \cdot$ RM(d, c^*) votes in RM(d, c^*) iterations, and then change all of them to $[c^* \succ \text{Others} \succ d]$. In each of the RM$(d, c^*)$ iterations, we first find $\lceil \log m \rceil + 1$ votes where for each alternative $c \neq d$ that loses to d in pairwise election, $d \succ c$ in at least one of these votes. Therefore, if we change these $(\lceil \log m \rceil + 1)$ votes to $[c^* \succ \text{Others} \succ d]$, then for each alternative c that loses to d in pairwise election, the weight on $d \to c$ is reduced by at least 2. In each of the $\text{RM}(d, c^*)$ iterations, we find $(\lceil \log m \rceil + 1)$ votes to increase the weights on the outgoing edges from c^* to alternatives that beat c^* in pairwise elections by 2.

Theorem 11 *Algorithm 4 runs in polynomial time and computes a* $2(\lceil \log m \rceil + 1)$*-approximation for* MOV *for Copeland*$_\alpha$*.*

Proof of Theorem 11: It suffices to show that in each of the $\text{RM}(d, c^*)$ iterations (Step 3), there always exists a vote that "covers" at least half of the alternatives in $\mathcal{C}_d$ (Step 6) and a vote that "covers" at least half of the alternatives in $\mathcal{C}_*$ (Step 11). We recall that for each alternative $c \in \mathcal{C}_d$, $D_{P^* \cup Q}(d, c) > 0$ and $c \succ d$ in all votes in Q. Therefore, $d \succ c$ in at least half of the votes in P^*, which means that there exists a vote $V_j \in P$ where $d \succ c$ for at least half of alternatives $c \in \mathcal{C}_d$. Similarly, Step 11 always successfully finds a vote. When the algorithm returns, $|\mathcal{I}| \leq 2(\lceil \log m \rceil + 1) \cdot \text{RM}(d, c^*) \leq 2(\lceil \log m \rceil + 1) \cdot \text{MoV}(P)$, which proves the theorem. ■

We feel that $\Theta(\log m)$ is a good approximation ratio in practice, because in most political elections, the number of alternatives is not very large.

5.2. Maximin

The idea behind the algorithm is the following. Changing one vote cannot change the maximin score of any alternative by more than 2. Therefore, given a profile P, let $\{d\} = \text{Maximin}(P)$ and let c^* denote the alternative that has the largest maximin score among all alternatives different from d. It follows that the MoV is at least $(S_M(P, d) - S_M(P, c^*))/4$. Let d' denote an arbitrary alternative such that $D_P(d, d') = S_M(P, d)$. Algorithm 5 finds $(S_M(P, d) - S_M(P, c^*))/2$ votes where $d \succ d'$, and then change all of them to $[c^* \succ \text{Others} \succ d]$. In the new profile, the maximin score of d is no more than $S_M(P, c^*)$ and the maximin score of c is at least $S_M(P, c^*)$, which means that d is not the unique winner.

Algorithm 5: AppMoVMaximin

Input: A profile P of n votes.
Output: The margin of victory for maximin.
Let $\{d\} = r(P)$, $c^* = \arg\max_{c \neq d} S_M(P, c)$, $d' = \arg\min_{c \neq d} D_P(d, c)$, $W = [c^* \succ \text{Others} \succ d]$.
Let $\mathcal{I}$ denote the indices of $(S_M(P, d) - S_M(P, c^*))/2$ votes in P where $d \succ d'$.
return $|\mathcal{I}|$.

Theorem 12 *Algorithm 5 runs in polynomial time and computes a* 2*-approximation for* MOV *for maximin.*

5.3. Connection between MOV and Destructive UCO

In this section, we reveal a connection between the optimal solution to MOV and the optimal solution to destructive unweighted coalition optimization problem (DUCO, see Definition 4) for the four voting rules studied in this paper where MOV is hard to compute (i.e., STV, ranked pairs, maximin, and Copeland). These bounds will tell us how to convert an approximation algorithm for DUCO to an approximation algorithm for MOV w.r.t. the same rule, and vice versa.

We first show upper bounds for MOV in terms of solutions to DUCO.

Theorem 13 *For any profile* P^{NM}*,*

- *MoV*$(P^{NM}, \textit{STV}) \leq OPT_{\text{DUCO}}(P^{NM}, \textit{STV})$*;*
- *MoV*$(P^{NM}, \textit{maximin}) \leq OPT_{\text{DUCO}}(P^{NM}, \textit{maximin})$*;*
- *MoV*$(P^{NM}, \textit{Copeland}_\alpha) = O((\log m) \cdot OPT_{\text{DUCO}}(P^{NM}, \textit{Copeland}_\alpha))$*;*
- *MoV*$(P^{NM}, \textit{RP}) = O((\log m) \cdot OPT_{\text{DUCO}}(P^{NM}, \textit{RP}))$.

Proof of Theorem 13: Let $\{d\} = r(P^{NM})$, $k = OPT_{\text{DUCO}}(P^{NM}, r)$ and P^M denote the manipulators' votes in the DUCO problem. Let $P^* = P^{NM} \cup P^M$.

STV: W.l.o.g. the plurality score of d is more than k (otherwise we can easily make the plurality score of d to be zero, which means that d is not the unique winner.) Let c be an arbitrary alternative that is still in the election when d is eliminated in round T when we apply STV on P^*. We choose arbitrary k voters in P^{NM} who rank d in their top positions, and change their votes to P^M. Let P' denote the profile obtained in this way. We next show that d is eliminated no later than round T when we apply STV to P'. Suppose for the sake of contradiction, d is eliminated later than round T. Then, when we apply STV to P', the only changes in the first $T-1$ rounds are that the plurality score of c increases and the plurality score of d decreases. Since c remains in round T for P, in each of the first $T-1$ rounds for P', the alternative that drops out is the same as the alternative that drops out in the same round for P. Therefore, d must drop out in the Tth round for P', which is a contradiction. Therefore, the MoV is no more than k.

maximin: The proof is similar to Algorithm 5. Let c^* be an arbitrary alternative that minimizes $D_{P^*}(d, c^*)$. We choose arbitrary k votes in P^{NM} where $d \succ c^*$ and change them to P^M.

Copeland: The proof is similar to Algorithm 4. Let c^* denote the alternative that has the largest relative margin from d, that is, $\text{RM}(d, c^*)$ is maximized. Because a single manipulator's vote cannot reduce the relative margin between any pair of alternatives by more than one, we have $k \geq \text{RM}(d, c^*)$. We recall that Algorithm 4 outputs a MoV no more than $2(\lceil \log m \rceil + 1) \cdot \text{RM}(d, c^*) \leq 2(\lceil \log m \rceil + 1)k$. This proves the upper bound.

Ranked pairs: Let $\succ'$ denote the order output by ranked pairs for P^*. Let c^* denote the alternative ranked on top of $\succ'$. We note that changing one vote will change the weight on any edge by no more than 2. Therefore, changing one manipulator's vote cannot change the weight difference between any pair of edges by more than 4. We first show the following claim, whose proof is omitted and can be found on the author's homepage.

Claim 1 *If in the WMG of P^{NM} the weights on all edges $a \rightarrow b$ incompatible with $\succ'$ are reduced by at least $\min(4k, max(D_{P^{NM}}(a, b), 0))$, then ranked pairs will output $\succ'$ in the resulting WMG.*

We are now ready to prove the upper bound for ranked pairs. The idea is similar to Algorithm 4 for Copeland. We find $\Theta(k \log m)$ votes in P^{NM} and change them to $\succ'$. These votes are found in a way such that each edge $a \rightarrow b$ that is not compatible with $\succ'$ is in at least k such votes, unless its weight has already become zero. We note that when the weight on $a \rightarrow b$ is larger than zero, $a \succ b$ in at least half of the votes in P^{NM}. Therefore, we can apply the following greedy algorithm to find the $\Theta(k \log m)$ votes. The algorithm has $2k$ iterations, each of which contains $\lceil \log(m^2) \rceil + 1 = \lceil 2 \log m \rceil + 1$ rounds. In each round, we first compute the remaining edges (that have positive weights and are not compatible with $\succ'$), and then we find a vote that covers at least half of them. After $\lceil 2 \log m \rceil + 1$ rounds, by changing the votes to $\succ'$, the weight on each edge that is not compatible with $\succ'$ either is reduced by 2, or becomes 0 or lower. Therefore, after $2k$ iterations, we can apply Claim 1, concluding that ranked pairs will output $\succ'$, where d is not ranked in the top. This proves the upper bound. ■

Obviously the upper bounds for STV and maximin shown in Theorem 13 are tight. The next proposition shows that the $\Theta(\log m)$ upper bound for Copeland is almost asymptotically tight, whose proof can be found on the author's homepage.

Proposition 1 *For any $\alpha \neq 1$ and any $\epsilon > 0$, there exists a profile P^{NM} such that MoV$(P^{NM}, \text{Copeland}_\alpha) \geq (\log m)^{1-\epsilon} \cdot OPT_{\text{DUCO}}(P^{NM}, \text{Copeland}_\alpha)$.*

We next give a general lower bound for voting rules that satisfy some canceling-out condition, defined as follows.

Definition 5 *For $L \in \mathbb{N}$, a voting rule r satisfies L-*canceling-out* if for any vote V, there exists a profile $Q(V)$ composed of no more than $L-1$ votes, such that for any profile P, $r(P) = r(P \cup \{V\} \cup Q(V))$. That is, the votes in $Q(V)$ cancels out V.*

Proposition 2 *Any WMG-based rule (including Borda) satisfy 2-canceling-out. Any positional scoring rule satisfies m-canceling-out. STV satisfies $(m!)$-canceling-out.*

Proof of Proposition 2: For any V, WMG$(V, \text{Rev}(V))$ does not have any edge with positive weight. Therefore, any WMG-based rule satisfies 2-canceling-out, where $Q(V) = \text{Rev}(V)$. Let M denote a cyclic permutation among alternatives. That is, $M : c_1 \to c_2 \to \cdots \to c_m \to c_1$. Then, for any positional scoring rule r, any vote V, and any profile P, we have $r(P) = r(P \cup \{V, M(V), \ldots, M^{m-1}(V)\})$, which means that it satisfies m-canceling-out. Finally, it is not hard to verify that for any profile P, STV$(P \cup L(\mathcal{C})) = $ STV(P), which means that STV satisfies $(m!)$-canceling-out. ∎

Proposition 3 *For any voting rule r that satisfies L-canceling-out and any profile P^{NM},*

$$OPT_{\text{DUCO}}(P^{NM}, r)/L \leq \textit{MoV}(P^{NM}, r)$$

Proof of Proposition 3: Let $\{d\} = r(P^{NM})$, $k = \text{MoV}(P^{NM}, r)$. k votes $P \subseteq P^{NM}$ are changed to P' such that d is not the unique winner. For DUCO, let the P^M denote the following Lk votes for the manipulators. There are k votes for P', and for each $V \in P$, there are $L - 1$ votes $Q(V)$. It follows that $r((P^{NM} \setminus P) \cup P') = r(P^{NM} \cup P^M)$, which means that $OPT_{\text{DUCO}}(P^{NM}, r) \leq L \cdot \text{MoV}(P^{NM}, r)$. ∎

We can use these bounds to convert approximation algorithms for DUCO to approximation algorithms for MOV, and vice versa. The following corollary follows from Proposition 2, Proposition 3, and Theorem 13.

Corollary 2 *For STV, a polynomial-time β-approximation algorithm for* DUCO *(respectively,* MOV*) can be used to compute an $(m!\beta)$-approximation for* MOV *(respectively,* DUCO*) in polynomial time. For ranked pairs, a polynomial-time β-approximation algorithm for* DUCO *(respectively,* MOV*) can be used to compute an $\Theta(\beta \log m)$-approximation for* MOV *(respectively,* DUCO*) in polynomial time.*

As we have mentioned, in political elections usually m is not large. Therefore, the $(m!\beta)$-approximation ratio may not be as unacceptable as it seems to be.

For Copeland and maximin, it is known that DUCO is in P [Conitzer et al. 2007], which means that there exist polynomial-time 1-approximation algorithms. However, a similar corollary does not yield a better bound than Algorithm 4 and Algorithm 5.

6. TYPICALLY HOW LARGE IS THE MARGIN OF VICTORY?

Let us start with a simple example for the majority rule for two alternatives $\{a, b\}$.

Example 1 *Suppose there are n voters, whose votes are drawn i.i.d. from a distribution π over all possible votes (i.e., voting for a with probability $\pi(a)$ or voting for b with probability $\pi(b)$). Let Y_a (respectively, Y_b) denote random variable that represents the total number of voters for a (respectively, for b). The margin of victory is thus a random variable $(Y_a - Y_b)/2$. Let X denote the random variable that takes 1 with probability $\pi(a)$ and takes -1 with probability $\pi(b)$. It follows that $(Y_a - Y_b)/2 = (\underbrace{X + \cdots + X}_{n})/2$. By the Central Limit Theorem, $(Y_a - Y_b)/2$ converges to a normal distribution with mean $n \cdot E(X)/2$ and variance $n \cdot Var(X)/2$.*

We are interested in usually how large is $(X_a - X_b)/2$. Not surprisingly, the answer depends on the distribution π. If $\pi(a) = \pi(b) = 1/2$, then the mean of $(X_a - X_b)/2$ is zero, and the probability that it is a few standard deviations away from the mean is small. For example, the probability that its absolute value is larger than $4\sqrt{n \cdot Var(X)/2}$ is less than 0.01, which means that with 99% probability the margin of victory is no more than $4\sqrt{n \cdot Var(X)/2}$. On the other hand, if $\pi(a) \neq \pi(b)$, w.l.o.g. $\pi(a) > \pi(b)$, then the mean of $(X_a - X_b)/2$ is $n(\pi(a) - \pi(b))/2$, which means that with high probability the margin of victory is very close to $n(\pi(a) - \pi(b))/2$.

We see in the above example that for the majority rule, depending on the distribution π over possible votes, the margin of victory is either $\Theta(\sqrt{n})$ or $\Theta(n)$, when we fix the number of alternatives and let the number of voters go to infinity. In this section, we prove this dichotomy theorem for a large class of generalized scoring rules, defined in the next subsection.

6.1. Generalized Scoring Rules

First, we recall the definition of *generalized scoring rules (GSRs)* [Xia and Conitzer 2008]. For any $K \in \mathbb{N}$, let $\mathcal{O}_K = \{o_1, \ldots, o_K\}$. A *total preorder* (preorder for short) is a reflexive, transitive, and total relation. Let Pre($\mathcal{O}_K$) denote the set of all preorders over $\mathcal{O}_K$. For any $\vec{p} \in \mathbb{R}^K$, we let Ord($\vec{p}$) denote the preorder $\trianglerighteq$ over $\mathcal{O}_K$ where $o_{k_1} \trianglerighteq o_{k_2}$ if and only if $p_{k_1} \geq p_{k_2}$. That is, the k_1th component of $\vec{p}$ is at least as large as the k_2th component of $\vec{p}$. For any preorder $\trianglerighteq$, if $o \trianglerighteq o'$ and $o' \trianglerighteq o$, then we write $o =_{\trianglerighteq} o'$. Each preorder $\trianglerighteq$ naturally induces a (partial) strict order $\triangleright$, where $o \triangleright o'$ if and only if $o \trianglerighteq o'$ and $o' \ntrianglerighteq o$. A preorder $\trianglerighteq'$ is a *refinement* of another preorder $\trianglerighteq$, if $\triangleright'$ extends $\triangleright$. That is, $\triangleright \subseteq \triangleright'$. We note that $\trianglerighteq$ is a refinement of itself. When $\triangleright \subsetneq \triangleright'$, we say that $\trianglerighteq'$ is a *strict refinement* of $\trianglerighteq$.

Definition 6 (Generalized scoring rules [Xia and Conitzer 2008]) *Let $K \in \mathbb{N}$, $f : L(\mathcal{C}) \to \mathbb{R}^K$ and $g : Pre(\mathcal{O}_K) \to (2^{\mathcal{C}} \setminus \emptyset)$. f and g determine a* generalized scoring rule (GSR) *$GS(f, g)$ as follows. For any profile $P = (V_1, \ldots, V_n) \in L(\mathcal{C})^n$, let $f(P) = \sum_{i=1}^n f(V_i)$, and let $GS(f, g)(P) = g(Ord(f(P)))$. We say that $GS(f, g)$ is* of order K.

For any $V \in L(\mathcal{C})$, $f(V)$ is called a *generalized score vector*, $f(P)$ is called a *total generalized score vector*, and Ord($f(P)$) is called the *induced preorder* of P. A generalized scoring rule is *non-redundant*, if for any pair $k_1, k_2 \leq K$ with $k_1 \neq k_2$, there exists a vote $V \in L(\mathcal{C})$ such that $(f(V))_{k_1} \neq (f(V))_{k_2}$. All voting rules studied in this paper are generalized scoring rules, which admit a natural axiomatic characterization [Xia and Conitzer 2009]. In this paper, we assume w.l.o.g. that all generalized scoring rules are non-redundant. Our results can be naturally extended to redundant generalized scoring rules.

For any pair of preorders $\trianglerighteq_1$, $\trianglerighteq_2$, we let $(\trianglerighteq_1 \oplus \trianglerighteq_2)$ denote the preorder $\trianglerighteq$ where $o \triangleright o'$ if and only if (1) $o \triangleright_1 o'$ or (2) $o' \ntriangleright_1 o$ and $o' \triangleright_2 o$. That is, $(\trianglerighteq_1 \oplus \trianglerighteq_2)$ is the preorder where $\triangleright_2$ is used to break ties in $\trianglerighteq_1$. Given a generalized scoring rule $r = \text{GS}(f, g)$, let $H(f) = \{\text{Ord}(f(P_1) - f(P_2)) : P_1, P_2 \in L(\mathcal{C})^*, |P_1| = |P_2|\}$. For any preorder $\trianglerighteq$, we let Nbr($\trianglerighteq$) denote the set of all *linear* orders that are $(\trianglerighteq \oplus \trianglerighteq')$ for some $\trianglerighteq' \in H(f)$. We now define *continuous* generalized scoring rules.

Definition 7 (Continuous generalized scoring rules) *A voting rule r is a* continuous generalized scoring rule*, if there exist f, g such that (1) $r = GS(f, g)$, and (2) for every profile P the following condition holds. If for all $\trianglerighteq \in Nbr(Ord(f(P)))$, $g(\trianglerighteq)$ is the same, then for any refinement $\trianglerighteq'$ of Ord($f(P)$), $g(\trianglerighteq')$ is also the same (as Ord($f(P)$)).*

We note that continuity is a property for a GSR w.r.t. a specific (f, g) pair. This is because multiple (f, g) pairs may correspond to the same voting rule (meaning that for each profile, these GS(f, g)'s select the same winner), but it is possible that GS(f, g) is not continuous for some of these (f, g) pairs. In the above definition, Nbr($\trianglerighteq$) represents the set of profiles "around" P. Nbr($\trianglerighteq$) consists of all profiles whose induced preorders over $\mathcal{O}_K$ are linear orders, and are obtained from Ord($f(P)$) by using preorders in $H(f)$ to break ties. We require that all preorders in Nbr(Ord($f(P)$)) are linear orders for a technical reason, which will be clear in the proof of the dichotomy theorem. Later we will show that Nbr(Ord($f(P)$)) is always non-empty, and all voting rules studied in this paper are continuous generalized scoring rules. We do not know any commonly studied non-continuous GSR.

6.2. The Dichotomy Theorem

Theorem 14 *Let $r = GS(f, g)$ be a continuous generalized scoring rule and let π be a distribution over $L(\mathcal{C})$ such that for every $V \in L(\mathcal{C})$, $\pi(V) > 0$. Suppose we fix π and the number of alternatives, generate n votes i.i.d. according to π, and let P_n denote the profile. Then, one (and exactly one) of the following two observations holds.*

(*1*) *For any* ϵ*, there exists* $\beta > 0$ *such that* $\lim_{n\to\infty} Pr(\frac{1}{\beta}\sqrt{n} \leq MoV(P_n) \leq \beta\sqrt{n}) \geq 1-\epsilon$.
(*2*) *There exists* $\beta > 0$ *such that* $\lim_{n\to\infty} Pr(MoV(P_n) \geq \beta n) = 1$.

In Theorem 14, Observation (1) states that with high probability (that can be arbitrarily close to 1), the margin of victory is $\Theta(\sqrt{n})$. Following the intuition presented in Example 1, this happens when all alternatives are "tied". However, this intuition only works at a high level—the notion of being tied largely depends on the representation of f and g, which is not clear for voting rules beyond positional scoring rules. If Observation (2) holds, then the margin of victory is $\Theta(n)$ with high probability, because the margin of victory is always no more than n.

Proof of Theorem 14: We first show that for any generalized scoring rule, if we fix the number of alternatives and the distribution π, and let the number of votes go to infinity, then the probability that the margin of victory is $\omega(\sqrt{n})$ is arbitrarily close to 1.

Lemma 1 *For any generalized scoring rule and any* ϵ*, there exists* $\beta > 1$ *such that* $\lim_{n\to\infty} Pr(MoV(P_n) \geq \frac{1}{\beta}\sqrt{n}) \geq 1-\epsilon$.

Lemma 1 follows after the following claim, which is proved by the Central Limit Theorem. Let $P_\pi = \sum_{V\in L(\mathcal{C})} \pi(V)\cdot V$. The induced preorder of $f(P_\pi)$ over $\mathcal{O}_K$ is denoted by $\trianglerighteq^\pi$.

Claim 2 *For any generalized scoring rule, fix* m *and* π*, let* P_n *denote the profile of* n *votes generated i.i.d. from* π*, and let* n *go to infinity. For any* ϵ*, there exists* $\beta > 1$ *and* $N \in \mathbb{N}$ *such that when* $n > N$*, with probability at least* $1-\epsilon$ *all following observations hold.*

- *For any* i, j *with* $o_i =_{\trianglerighteq^\pi} o_j$, $\frac{1}{\beta}\sqrt{n} \leq |(f(P_n))_i - (f(P_n))_j| \leq \beta\sqrt{n}$.
- *For any* i, j *with* $o_i \triangleright^\pi o_j$, $\frac{1}{\beta}\sqrt{n} \leq |(f(P_n))_i - (f(P_n))_j - n[(f(P_\pi))_j - (f(P_\pi))_i]| \leq \beta\sqrt{n}$.

The idea for the rest of the proof is the following. By Claim 2, with high probability $f(P_n)/n$ is approximately $f(P_\pi)$. Therefore, with high probability if we only change $o(n)$ votes, then we cannot change the order between any pair of $o, o' \in \mathcal{O}_K$ where $o \triangleright^\pi o'$. It follows that changing no more than $o(n)$ votes only plays the role of a tie-breaker to obtain a refinement of $\trianglerighteq^\pi$. If we are allowed to change $\Theta(\sqrt{n})$ votes, then all linear orders in Nbr($\trianglerighteq^\pi$) can be obtained. Then, the continuity of r guarantees that to decide whether changing $\Theta(\sqrt{n})$ votes can change the winners, it suffices to check whether the set of winners is the same for all linear orders in Nbr($\trianglerighteq^\pi$). Formally, we prove the following claim.

Claim 3 *If there exist* $\trianglerighteq_1, \trianglerighteq_2 \in Nbr(\trianglerighteq^\pi)$ *such that* $g(\trianglerighteq_1) \neq g(\trianglerighteq_2)$*, then Observation (1) holds; otherwise Observation (2) holds.*

Proof of Claim 3: Suppose there exist $\trianglerighteq_1, \trianglerighteq_2 \in$ Nbr($\trianglerighteq^\pi$) such that $g(\trianglerighteq_1) \neq g(\trianglerighteq_2)$. We first show that for any $\epsilon > 0$, there exists $\beta > 0$ such that the probability that changing no more than $\beta\sqrt{n}$ votes can change the order to $\trianglerighteq_1$ is larger than $1-\epsilon/2$. Let T_{min} denote the minimum number of votes that are sufficient to induce any order in $H(f)$. That is,

$$T_{min} = \arg\min_T\{\forall \trianglerighteq \in H(f), \exists P^1, P^2 \text{ s.t. } |P^1| = |P^2| \leq T, \trianglerighteq = \text{Ord}(f(P^1) - f(P^2))\}$$

Because $|H(f)|$ is finite, T_{min} is also finite. Let $\trianglerighteq_1 = (\trianglerighteq^\pi \oplus \trianglerighteq')$, where $\trianglerighteq' \in H(f)$ such that $\trianglerighteq' =$ Ord$(f(P^1) - f(P^2))$ with $|P^1| = |P^2| \leq T_{min}$. We partition $\{\{k_1, k_2\} : k_1 \neq k_2\}$ into two sets $\mathcal{K}_1 \cup \mathcal{K}_2$. $\{k_1, k_2\} \in \mathcal{K}_1$ if and only if $(f(P_\pi))_{k_1} = (f(P_\pi))_{k_2}$, that is, $o_{k_1} =_{\trianglerighteq^\pi} o_{k_2}$. Otherwise $\{k_1, k_2\} \in \mathcal{K}_2$. For all $\{k_1, k_2\} \in \mathcal{K}_1$, $E[(f(P_n))_{k_1} - (f(P_n))_{k_2}] = 0$. For all $\{k_1, k_2\} \in \mathcal{K}_1$, $|E[(f(P_n))_{k_1} - (f(P_n))_{k_2}]| = \Theta(n)$.

Because $\trianglerighteq_1 = (\trianglerighteq^\pi \oplus \trianglerighteq')$ is a linear order, for any pair $\{k_1, k_2\} \in \mathcal{K}_1$, either $o_{k_1} \triangleright' o_{k_2}$ or $o_{k_2} \triangleright' o_{k_1}$. Similar to the proof of Lemma 1, by the Central Limit Theorem, for each pair $\{k_1, k_2\} \in \mathcal{K}_1$, there exist $N_{(k_1,k_2)}$ and $\beta_{(k_1,k_2)}$ such that for any $n \geq N_{(k_1,k_2)}$, the probability for $|(f(P_n))_{k_1} - (f(P_n))_{k_2}|$ to be larger than $\beta_{(k_1,k_2)}\sqrt{n}|(f(P^1) - f(P^2))_{k_1} - (f(P^1) - f(P^2))_{k_2}|$ is no more than $\epsilon/(4m^2)$. Then, we let $\beta_1 = \max\{\beta_{(k_1,k_2)}\}$, and we choose N_1 that satisfies the following conditions.

(1) N_1 is larger than all $N_{(k_1,k_2)}$.

(2) For any $n \geq N_1$ and any $\{k_1, k_2\} \in \mathcal{K}_2$, the probability for $\beta\sqrt{n}|(f(P^1) - f(P^2))_{k_1} - (f(P^1) - f(P^2))_{k_2}|$ to be larger than $|(f(P_n))_{k_1} - (f(P_n))_{k_2}|$ is no more than $\epsilon/(4m^2)$. This follows from the Central Limit Theorem and $|E[(f(P_n))_{k_1} - (f(P_n))_{k_2}]| = \Theta(n)$.

(3) For any $n \geq N_1$ and any $L \in L(\mathcal{C})$, the probability for the number of L-votes in P_n to be smaller than $\beta_1 T_{min}\sqrt{n}$ is smaller than $\epsilon/(4m!)$. This follows from the assumption that π takes positive probability on every $L \in L(\mathcal{X})$, and again, by the Central Limit Theorem, with high probability the number of L-votes in P_n is $\Theta(n)$, which is asymptotically larger than $\Theta(\sqrt{n})$.

Condition (1) ensures that with high probability, changing no more than $\beta_1 T_{min}\sqrt{n}$ votes (changing $\beta_1\sqrt{n}$ copies of P^2 to $\beta_1\sqrt{n}$ copies of P^1) can break ties in $\trianglerighteq^\pi$ as in $\trianglerighteq'$. Condition (2) ensures that with high probability changing no more than $\beta_1 T_{min}\sqrt{n}$ votes does not change any strict order in $\trianglerighteq^\pi$. Condition (3) ensures that with high probability P_n contains $\beta_1\sqrt{n}$ copies of P^2 for us to change. Therefore, when $n \geq N_1$, with probability larger than $1 - \epsilon/2$ we can change no more than $\beta_1 T_{min}\sqrt{n}$ votes to change the order to $\trianglerighteq_1$.

Similarly, there exist β_2 and N_2 such that when $n \geq N_2$, with probability larger than $1 - \epsilon/2$ we can change no more than $\beta_2 T_{min}\sqrt{n}$ votes to change the order to $\trianglerighteq_2$. We let $\beta = \max\{\beta_1 T_{min}, \beta_2 T_{min}\}$ and $N = \max\{N_1, N_2\}$. It follows that for any $n \geq N$, with probability larger than $1-\epsilon$ we can change no more than $\beta\sqrt{n}$ votes to change the order to $\trianglerighteq_1$ and meanwhile, we can change no more than $\beta\sqrt{n}$ votes to change the order to $\trianglerighteq_2$. Because $g(\trianglerighteq_1) \neq g(\trianglerighteq_2)$, one of them must be different from $r(P_n)$, which means that with probability that is at least $1 - \epsilon$, the margin of victory is no more than $\beta\sqrt{n}$.

We now show that if for all $\trianglerighteq_1, \trianglerighteq_2 \in$ Nbr($\trianglerighteq^\pi$), $g(\trianglerighteq_1) = g(\trianglerighteq_2)$, then Observation (2) holds. Let d_{min} denote the minimum positive difference between the k_1th component of $f(P_\pi)$ and k_2th component of $f(P_\pi)$ for all pairs (k_1, k_2). Because the mean of $f(P_n)$ is $nf(P_\pi)$, by the Central Limit Theorem, the probability for the following condition to hold goes to 1 as n goes to infinity: The difference between any pair of components of $f(P_n)$ is either 0 or larger than $nd_{min}/2$. Let d_{max} denote the largest component of $f(L)$ for all $L \in L(\mathcal{X})$. Let $\beta = \frac{d_{min}}{8d_{max}}$. It follows that by changing any βn votes, the difference between any pair of components in the total generalized score vector cannot be changed by more than $4\beta n d_{max} = nd_{min}/2$. Therefore, with high probability (that goes to 1), changing any βn votes are only tie-breakers for $\trianglerighteq^\pi$. Because r is continuous and for all $\trianglerighteq_1, \trianglerighteq_2 \in$ Nbr($\trianglerighteq^\pi$), $g(\trianglerighteq_1) = g(\trianglerighteq_2)$, no matter how ties in $\trianglerighteq^\pi$ are broken, the winner is always $g(\trianglerighteq_1)$. This completes the proof. ■

The theorem follows directly after Claim 3. ■

Finally, all voting rules studied in this paper are continuous generalized scoring rules, which means that Theorem 14 applies to these rules.

Theorem 15 *All positional scoring rules, maximin, Coepalnd$_\alpha$, ranked pairs, plurality with runoff, Bucklin, STV, and approval are continuous generalized scoring rules.*

7. SUMMARY AND FUTURE WORK

In this paper, we investigate the computational complexity and (in)approximability of computing the margin of victory for various commonly studied voting rules. We also prove a dichotomy theorem on the margin of victory when the votes are drawn i.i.d. Most of our results are quite positive, suggesting that margin of victory can be efficiently computed or approximated.

For future work, we plan to continue working on designing practical approximation or randomization algorithms to compute the margin of victory for STV and ranked pairs. It seems that our dichotomy theorem can extend to other types of strategic behavior, e.g., control by adding and control by deleting voters [Bartholdi et al. 1992]. More importantly, how to extend risk-limiting audit methods beyond plurality is an important topic for future research.

8. ACKNOWLEDGMENTS

The author is supported by NSF under Grant #1136996 to the Computing Research Association for the CIFellows Project. The author is deeply indebted to Ronald Rivest and Emily Shen for very inspiring and detailed discussions and comments. The author also thanks Piotr Faliszewski and Mark Wilson for helpful suggestions and comments.

REFERENCES

BAHARAD, E. AND NEEMAN, Z. 2002. The asymptotic strategyproofness of scoring and condorcet consistent rules. *Review of Economic Design 4*, 331–340.

BARTHOLDI, III, J. AND ORLIN, J. 1991. Single transferable vote resists strategic voting. *SCW 8,* 4, 341–354.

BARTHOLDI, III, J., TOVEY, C., AND TRICK, M. 1992. How hard is it to control an election? *Math. Comput. Modelling 16,* 8-9, 27–40. Formal theories of politics, II.

BRILL, M. AND FISCHER, F. 2012. The Price of Neutrality for the Ranked Pairs Method. In *Proc. AAAI*.

CARY, D. 2011. Estimating the margin of victory for instant-runoff voting. In *Proc. 2011 EVT/WOTE Conference*.

CHAMBERLIN, J. R. 1985. An investigation into the relative manipulability of four voting systems. *Behavioral Science 30,* 4, 195—203.

CONITZER, V., ROGNLIE, M., AND XIA, L. 2009. Preference functions that score rankings and maximum likelihood estimation. In *Proc. IJCAI).* 109–115.

CONITZER, V., SANDHOLM, T., AND LANG, J. 2007. When are elections with few candidates hard to manipulate? *Journal of the ACM 54,* 3, 1–33.

FALISZEWSKI, P., HEMASPAANDRA, E., AND HEMASPAANDRA, L. A. 2009. How hard is bribery in elections? *Journal of Artificial Intelligence Research 35*, 485–532.

FALISZEWSKI, P., HEMASPAANDRA, E., AND HEMASPAANDRA, L. A. 2010. Using complexity to protect elections. *Commun. ACM 53*, 74–82.

FALISZEWSKI, P. AND PROCACCIA, A. D. 2010. AI's war on manipulation: Are we winning? *AI Magazine 31,* 4, 53–64.

GAREY, M. AND JOHNSON, D. 1979. *Computers and Intractability*. W. H. Freeman and Company.

MAGRINO, T. R., RIVEST, R. L., SHEN, E., AND WAGNER, D. 2011. Computing the Margin of Victory in IRV Elections. In *Proc. 2011 EVT/WOTE Conference*.

MCGARVEY, D. C. 1953. A theorem on the construction of voting paradoxes. *Econometrica 21,* 4, 608–610.

NITZAN, S. 1982. The vulnerability of point-voting schemes to preference variation and strategic manipulation. *Public Choice 47*, 349–370.

NORDEN, L., BURSTEIN, A., HALL, J. L., AND CHEN, M. 2007. Post-election audits: Restoring trust in elections. Brennan Center for Justice at The New York University School of Law and The Samuelson Law, Technology and Public Policy Clinic at the University of California, Berkeley School of Law (Boalt Hall).

PELEG, B. 1979. A note on manipulability of large voting schemes. *Theory and Decision 11*, 401–412.

PRITCHARD, G. AND SLINKO, A. 2006. On the Average Minimum Size of a Manipulating Coalition. *SCW 27,* 2, 263–277.

PRITCHARD, G. AND WILSON, M. 2007. Exact results on manipulability of positional voting rules. *SCW 29,* 3, 487–513.

PRITCHARD, G. AND WILSON, M. C. 2009. Asymptotics of the minimum manipulating coalition size for positional voting rules under impartial culture behaviour. *Mathematical Social Sciences 1*, 35–57.

PROCACCIA, A. D. AND ROSENSCHEIN, J. S. 2007. Average-case tractability of manipulation in voting via the fraction of manipulators. In *Proc. AAMAS*. 718–720.

ROTHE, J. AND SCHEND, L. 2012. Typical-Case Challenges to Complexity Shields That Are Supposed to Protect Elections Against Manipulation and Control: A Survey. In *Proc. ISAIM*.

SARWATE, A., CHECKOWAY, S., AND SHACHAM, H. 2011. Risk-limiting Audits for Nonplurality Elections. Tech. Rep. CS2010-0967, UC San Diego. June.

SLINKO, A. 2002. On asymptotic strategy-proofness of classical social choice rules. *Theory and Decision 52*, 389–398.

SLINKO, A. 2004. How large should a coalition be to manipulate an election? *Mathematical Social Sciences 47,* 3, 289–293.

STARK, P. B. 2008a. Conservative statistical post-election audits. *The Annals of Applied Statistics 2,* 2, 550–581.

STARK, P. B. 2008b. A sharper discrepancy measure for post-election audits. *The Annals of Applied Statistics 2,* 3, 982–985.

STARK, P. B. 2009a. Efficient post-election audits of multiple contests: 2009 California tests. In *4th Annual Conference on Empirical Legal Studies (CELS)*.

STARK, P. B. 2009b. Risk-limiting post-election audits: P-values from common probability inequalities. *IEEE Transactions on Information Forensics and Security 4*, 1005–1014.

STARK, P. B. 2010. Super-simple simultaneous single-ballot risk-limiting audits. In *Proc. 2010 EVT/WOTE Conference*.

WOLCHOK, S., WUSTROW, E., HALDERMAN, J. A., PRASAD, H. K., KANKIPATI, A., SAKHAMURI, S. K., YAGATI, V., AND GONGGRIJP, R. 2010. Security analysis of india's electronic voting machines. In *Proc. CCS*. 1–14.

XIA, L. AND CONITZER, V. 2008. Generalized scoring rules and the frequency of coalitional manipulability. In *Proc. EC*. 109–118.

XIA, L. AND CONITZER, V. 2009. Finite local consistency characterizes generalized scoring rules. In *Proc. IJCAI*. 336–341.

XIA, L., ZUCKERMAN, M., PROCACCIA, A. D., CONITZER, V., AND ROSENSCHEIN, J. 2009. Complexity of unweighted coalitional manipulation under some common voting rules. In *Proc. IJCAI*. 348–353.

ZUCKERMAN, M., PROCACCIA, A. D., AND ROSENSCHEIN, J. S. 2009. Algorithms for the coalitional manipulation problem. *Artificial Intelligence 173,* 2, 392–412.