X

Computing the Margin of Victory for Various Voting Rules

Lirong Xia, SEAS, Harvard University

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 of 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

1. INTRODUCTION

Voting is a popular method used to aggregate agents' preferences, which not only has been widely used in political elections, but also, more recently, has found applications in multi-agent systems [Ephrati and Rosenschein 1991], building better recommendation systems [Ghosh et al. 1999] and web-search engines [Dwork et al. 2001], etc. An important measurement in voting for the robustness of outcome is the *margin of victory*, which 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 the voters 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 (see `http://en.wikipedia.org/wiki/United_Kingdom_Alternative_Vote_referendum,_2011`). To extend risk-limiting audit methods to other voting systems, naturally, the first step is to compute the margin of victory. However, some recent work suggests that at least for STV, computing the margin of victory is much harder than 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 some common 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 is fixed-parameter tractable w.r.t. the margin of victory.[1] 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.

Table I. Computing the margin of victory.

Voting rule	MOV	
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	FPT NPC $\Theta(\log m)$-apprx.	(Theorem 8) (Theorem 4) (Algorithm 4)
Maximin	FPT NPC 2-apprx.	(Theorem 9) (Theorem 3) (Algorithm 5)
STV (a.k.a. IRV, AV)	NPC for MOV_1 No 2-apprx.	(Theorem1) (Corollary 1)
Ranked pairs	NPC for MOV_1 No 2-apprx.	(Theorem2) (Corollary 1)

Finally, we ask typically how large the margin of victory is, when the votes are drawn i.i.d. according to some distribution. 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], 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 15 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 the number of alternatives,

[1] This implies that when the margin of victory is bounded above by a constant, then there is a polynomial-time algorithm to compute it.

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

We also show that all voting rules studied in this paper are continuous generalized scoring rules.

1.2. Related work and discussion

When the election has 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 have completely different motivations. For BRIBERY, hardness of computation is good news, because it means that 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 15).

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.

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 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 approve 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

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

of positional scoring rules are *Borda*, for which the scoring vector is $(m-1, m-2, \dots, 0)$; *plurality*, for which the scoring vector is $(1, 0, \dots, 0)$; and *veto*, for which the scoring vector is $(1, \dots, 1, 0)$. 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 $S_M(P, c) = \min\{D_P(c, c') : c' \in \mathcal{C}, c' \neq c\}$.
— **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 $D_P(c, c')$. 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.
— **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.
— **Single transferable vote (STV), a.k.a. instant runoff voting (IRV) or alternative vote (AV)**: The election has 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). The last-remaining alternative is the winner.

For ranked pairs, plurality with runoff, and STV, 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.

A voting rule r is *WMG-based*, if the winners only depends on the WMG of the input profile. That is, for any pair of profiles P_1, P_2, if WMG(P_1) = WMG(P_2), then $r(P_1) = r(P_2)$. Borda, Copeland, maximin, and ranked pairs are WMG-based.

We now recall the definition of *generalized scoring rules (GSRs)* [Xia and Conitzer 2008]. 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 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 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 1 *Let $K \in \mathbb{N}$, $f : L(\mathcal{C}) \to \mathbb{R}^K$ and $g : \mathcal{O}_K \to (2^{\mathcal{C}} \setminus \emptyset)$. f and g determine a* generalized scoring rule (GSR) *$GS(f, g)$ as follows. For any profile $P = (V_1, \dots, V_n) \in L(\mathcal{C})^n$, let $f(P) = \sum_{i=1}^n f(V_i)$, and let $GS(f, g)(P) = g(Ord(f(P)))$. 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 $(f(V))_{k_1} \neq (f(V))_{k_2}$. That is, the k_1th component is not always the same as the k_2th component (otherwise we can remove one of them without changing the rule). 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.

2.1. Margin of victory, Bribery, and Manipulation

Definition 2 *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 MoV(P) to denote MoV(P, r) when causing no confusion. We now define the computational problems.

Definition 3 *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 4 *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, destructive BRIBERY is a special case of MOV$_k$, where the given profile P has a unique winner. MOV is also closely related to the unweighted coalitional optimization problem [Zuckerman et al. 2009].

Definition 5 *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 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.*

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 position 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 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 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 observation 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 $n-1$.

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, if there is a polynomial-time 2-approximation algorithm for MOV, then we can use it to solve MOV_1 in polynomial-time. Therefore, we have the following inapproximability result.[5]

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 $c \rightarrow d$ is increased by 2 and the weight on $d \rightarrow c$ 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.

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

[5]The author thanks Emily Shen for this very nice observation.

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 we 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 be easily extended to general MOV problems. The corresponding ways that 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).

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 $\lceil m/2 \rceil + 1$.

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:

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

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

— **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 alternatives where e is ranked in the top position and where $d \succ c$; let B_e denote the set of alternatives 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 increases 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: MoVPluO

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. Again, for simplicity of presentation, we assume that the input profile has a unique winner.

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

$\text{MoV}(P^{NM})$. For every $c \neq d$, we compute a *relative margin* $\text{RM}(d, c)$ 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 $\text{MoV}(P^{NM})$ is at least $\text{RM}(d, c^*)$, 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 $2(\lceil \log m \rceil + 1) \cdot \text{RM}(d, c^*)$ votes in $\text{RM}(d, c^*)$ iterations, and then change all of them to $[c^* \succ \text{Others} \succ d]$. In each of the $\text{RM}(d, c^*)$ 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 beats c^* in pairwise elections by 2.

Algorithm 4: AppMoVCopeland

Input: A profile P of n votes.
Output: The margin of victory for Copeland$_\alpha$.
1 Let $\{d\} = r(P)$, $P^* = P$, $\mathcal{I} = \emptyset$, $Q = \emptyset$. $W = [c^* \succ \text{Others} \succ d]$
2 Compute $\text{RM}(d, c)$ for every $c \neq d$, and let $c^* = \arg\min_c \text{RM}(d, c)$.
3 **for** *RM*(d, c^*) *iterations* **do**
4   Let $\mathcal{C}_d = \{c : D_{P^* \cup Q}(d, c) > 0\}$.
5   **for** $\lceil \log m \rceil + 1$ *rounds* **do**
6   Find a vote $V_j \in P^*$ where $d \succ c$ holds for at least half of alternatives c in $\mathcal{C}_d$.
7         $\mathcal{C}_d \leftarrow \mathcal{C}_d \setminus \{c : d \succ_{V_j} c\}$, $P^* \leftarrow P^* \setminus \{V_j\}$, $Q \leftarrow Q \cup \{W\}$, $\mathcal{I} \leftarrow \mathcal{I} \cup \{j\}$.
8   **end**
9   Let $\mathcal{C}_* = \{c : D_{P^* \cup Q}(c, c^*) > 0\}$.
10   **for** $\lceil \log m \rceil + 1$ *rounds* **do**
11   Find a vote $V_j \in P^*$ where $c \succ c^*$ holds for at least half of alternatives c in $\mathcal{C}_*$.
12         $\mathcal{C}_d \leftarrow \mathcal{C}_d \setminus \{c : c \succ_{V_j} c^*\}$, $P^* \leftarrow P^* \setminus \{V_j\}$, $Q \leftarrow Q \cup \{W\}$, $\mathcal{I} \leftarrow \mathcal{I} \cup \{j\}$.
13   **end**
14 **end**
15 **return** $|\mathcal{I}|$.

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 $[c^* \succ$

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 UCO (DUCO) 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.

Definition 6 *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 1 *Any WMG-based rule and Borda satisfy* 2*-canceling-out. Any positional scoring rule satisfies m-canceling-out. STV satisfies $(m!)$-canceling-out.*

Proof of Proposition 1: 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 $Q(V) = \text{Rev}(V)$. Similarly, all alternatives have the same Borda score in $\{V, \text{Rev}(V)\}$, which means that Borda satisfies 2-canceling-out. 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 2 *For any voting rule r that satisfies L-canceling-out and any profile P^{NM},*

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

Proof of Proposition 2: 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 $OPT_{\text{DUCO}}(P^{NM}, r) \leq L \cdot \text{MoV}(P^{NM}, r)$. ■

Next, we show upper bounds for MOV in terms of solutions to DUCO.

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

— $MoV(P^{NM}, STV) \leq OPT_{\text{DUCO}}(P^{NM}, STV)$;
— $MoV(P^{NM}, maximin) \leq OPT_{\text{DUCO}}(P^{NM}, maximin)$;
— $MoV(P^{NM}, Copeland_\alpha) = O((\log m) \cdot OPT_{\text{DUCO}}(P^{NM}, Copeland_\alpha))$;
— $MoV(P^{NM}, RP) = O((\log m) \cdot OPT_{\text{DUCO}}(P^{NM}, RP))$.

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, 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 $k \geq$ 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 one manipulator's vote cannot change the weight difference between any pair of edges by more than 2. We first show the following claim, whose proof is relegated to the appendix.

Claim 1 *If in the WMG of P^{NM} the weights on all edges $a \to b$ incompatible with $\succ'$ are reduced by at least $\min(2k, max(D_{P^{NM}}(a, b), 0))$, 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 $a \to b$ 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 $a \to b$ 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 $\lceil \log(m^2) \rceil + 1 = \lceil 2 \log m \rceil + 1$ iterations. In each iteration, 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 $\lceil 2 \log m \rceil + 1$ iterations, by changing the vote to $\succ'$, the weight on each edge that is not compatible with $\succ'$ either is reduced by $2k$, or becomes 0 or lower. By Claim 1, ranked pairs would 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 is in the appendix.

Proposition 3 *For any $\alpha \neq 1$ and any $\epsilon > 0$, there exists a profile P^{NM} such that MoV$(P^{NM}$, Copeland$_\alpha) \geq (\log m)^{1-\epsilon} \cdot OPT_{\text{DUCO}}(P^{NM}$, Copeland$_\alpha)$.*

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 1, Proposition 2, 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. We first show that for any generalized scoring rule, if we fix the number of alternatives 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.

Theorem 14 *For any generalized scoring rule and any ϵ, there exists $\beta > 0$ such that $\lim_{n\to\infty} Pr(MoV(P_n) \geq \beta\sqrt{n}) \geq 1 - \epsilon$.*

Proof of Theorem 14: Let $r = \text{GS}(f, g)$ denote the generalized scoring rule. W.l.o.g. r is non-redundant (the theorem can be easily extended to redundant cases). We prove that there exists $\beta > 0$ such that $\lim_{n\to\infty} \Pr(\text{MoV}(P_n) \leq \beta\sqrt{n}) < \epsilon$. In words, the probability that changing $\beta\sqrt{n}$ votes can change the outcome is smaller than ϵ as n goes to infinity. To show this, we prove that there exists $\beta > 0$ such that the probability of changing $\beta\sqrt{n}$ votes can change the order between *any* pair of components of the total generalized score vector is below ϵ, when n is large enough.

Let d_{max} denote the maximum difference between any pair of components in generalized score vectors. That is, $d_{max} = \max_{t,t',L\in L(\mathcal{C})}\{(f(L))_t - (f(L))_{t'}\}$. Let $X_{(k_1,k_2)}$ denote the random variable that takes $(f(L))_{k_1} - (f(L))_{k_2}$ with probability $\pi(L)$ for any $L \in L(\mathcal{C})$. Because r is not redundant, $X_{(k_1,k_2)}$ is not always 0. If $X_{(k_1,k_2)}$ is a constant, then changing any number of votes (including $\beta\sqrt{n}$) cannot change the order between o_{k_1} and o_{k_2}.

Suppose $X_{(k_1,k_2)}$ is not always a constant. Let μ and σ denote the mean and standard deviation of $X_{(k_1,k_2)}$, respectively. That is, $\mu = E[X_{(k_1,k_2)}]$, $\sigma = \sqrt{\text{Var}(X_{(k_1,k_2)})}$. Let $Y_{(k_1,k_2)} = (f(P_n))_{k_1} - (f(P_n))_{k_2} = \underbrace{X_{(k_1,k_2)} + \cdots + X_{(k_1,k_2)}}_{n}$. By the Central Limit Theorem, for any z we have the following calculation, where Φ is the cumulative density function of the normal distribution whose mean is zero and whose variance is 1.

$$\lim_{n\to\infty} \Pr\left(\frac{Y_{(k_1,k_2)} - n\mu}{\sigma\sqrt{n}} \leq z\right) = \Phi(z)$$

Therefore, for any $\gamma > 0$, $\lim_{n\to\infty} \Pr(|Y_{(k_1,k_2)} - n\mu| \leq \gamma\sigma\sqrt{n}) = \Phi(\gamma) - \Phi(-\gamma)$. We choose an arbitrary $\beta_{(k_1,k_2)} > 0$ such that $\Phi(\beta_{(k_1,k_2)}/\sigma) - \Phi(-\beta_{(k_1,k_2)}/\sigma) < \epsilon/(2m^2)$. Let $N_{(k_1,k_2)}$ be a natural number such that for any $n \geq N_{(k_1,k_2)}$, $\Pr(|Y_{(k_1,k_2)} - n\mu| \leq \beta_{(k_1,k_2)}\sqrt{n}) < \epsilon/(2m^2)$. It follows that $\Pr(|Y_{(k_1,k_2)}| \leq \beta_{(k_1,k_2)}\sqrt{n}) < \epsilon/(m^2)$.

Let $\beta = \min_{k_1,k_2}\{\beta_{(k_1,k_2)}/(4d_{max})\}$ and $N = \max\{N_{(k_1,k_2)}\}$, where the minimum is taken across over all pairs of (k_1, k_2) such that $X_{(k_1,k_2)}$ is not a constant. For any $n \geq N$, we have the following calculation.

$$
\begin{aligned}
& \Pr(\text{changing } \beta\sqrt{n} \text{ votes can change the winner}) \\
\leq & \sum_{k_1,k_2} \Pr(\text{changing } \beta\sqrt{n} \text{ votes can change the order between } o_{k_1} \text{ and } o_{k_2}) \\
\leq & \sum_{k_1,k_2} \Pr(|Y_{(k_1,k_2)}| \leq 4d_{max}\beta\sqrt{n}) \qquad (\text{where } X_{(k_1,k_2)} \text{ is non-constant}) \\
\leq & \sum_{k_1,k_2} \Pr(|Y_{(k_1,k_2)}| \leq \beta_{(k_1,k_2)}\sqrt{n}) \qquad (\text{because } 4d_{max}\beta \leq \beta_{(k_1,k_2)}) \\
\leq & \frac{m(m-1)}{2} \cdot \frac{\epsilon}{m^2} < \epsilon
\end{aligned}
$$

This completes the proof. ■

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) $o' \not\triangleright_1 o$ and $o' \triangleright_2 o$. That is, $(\trianglerighteq_1 \oplus \trianglerighteq_2)$ is the preorder where $\triangleright_2$ is used to break ties in $\trianglerighteq_1$.

Definition 7 *Given a generalized scoring rule $r = GS(f, g)$, let $H(f) = \{Ord(f(P_1) - f(P_2)) : P_1, P_2 \in L(\mathcal{C})^*, |P_1| = |P_2|\}$. For any preorder $\trianglerighteq$, we let Nbr($\trianglerighteq$) denote the set of all* linear *orders that are $(\trianglerighteq \oplus \trianglerighteq')$ for some $\trianglerighteq' \in H(f)$. A generalized scoring rule is* continuous, *if 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($f(P)$)).*

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 now show the dichotomy theorem for continuous generalized scoring rule.

Theorem 15 *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 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(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) $\geq \beta n$) $= 1$.

Observation (1) states that with high probability (that can be arbitrarily close to 1), the margin of victory is $O(\sqrt{n})$. If Observation (1) holds, then together with Theorem 14, with high probability the margin of victory is $\Theta(\sqrt{n})$. If Observation (2) holds, then the margin of victory is $\Theta(n)$ with high probability.

Proof of Theorem 15: The idea of the proof is the following. Let $P_\pi = \sum_{V\in L(\mathcal{C})} \pi(V) \cdot V$. By the Central Limit Theorem, $f(P_n)/n$ is approximately $f(P_\pi)$, whose induced preorder over $\mathcal{O}_K$ is denoted by $\trianglerighteq^\pi$. 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 $o \triangleright^\pi o'$. Therefore, 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 $\omega(\sqrt{n})$ 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 2 *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 2: 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 $\trianglerighteq' = \text{Ord}(f(P^1) - f(P^2))$ 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 Theorem 14, 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 $4\beta n d_{max} = n d_{min}/2$. 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 from Claim 2. ■

Finally, we show that all voting rules studied in this paper are continuous generalized scoring rules. We first present a lemma, whose proof (as well as the proof of Theorem 16) are relegated to the appendix.

Lemma 1 *Let $r = GS(f, g)$ be non-redundant. For any profile P and any $\trianglerighteq_* \in H(f)$, there exists a preorder $\trianglerighteq'_* \in H(f)$ that refines $\trianglerighteq_*$ and $(Ord(f(P)) \oplus \trianglerighteq'_*)$ is a linear order.*

Lemma 1 implies that for any non-redundant generalized scoring rule GS(f, g) and any profile P, Nbr(Ord$(f(P))$) is non-empty. It will be used to prove that all voting rules studied in this paper are continuous generalized scoring rules in the following way. Given a profile P with induced preorder $\trianglerighteq$, we first find $\trianglerighteq_* \in H(f)$ that breaks some "critical" ties to ensure the properties we want. Note that $(\trianglerighteq \oplus \trianglerighteq_*)$ might not be a linear order. Then, we apply Lemma 1 on $\trianglerighteq_*$ to obtain a preorder $\trianglerighteq'_*$, where the "critical" ties are broken in the same way as in $\trianglerighteq_*$, and $(\trianglerighteq \oplus \trianglerighteq'_*)$ is a linear order.

Theorem 16 *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 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 $\Theta(n)$. Most of our results are quite positive, suggesting that margin of victory can be efficiently computed.

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 be extend to other types of strategic behavior, e.g., control by adding and control by deleting voters [Bartholdi et al. 1992]. Also, how to extend risk-limiting audit methods beyond plurality is an important topic for future research.

8. ACKNOWLEDGMENTS

Lirong Xia 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, comments, and feedback. The author also wants to thank Piotr Faliszewski for helpful suggestions and comments.

REFERENCES

BARTHOLDI, III, J. AND ORLIN, J. 1991. Single transferable vote resists strategic voting. *Social Choice and Welfare 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.

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

CONITZER, V., ROGNLIE, M., AND XIA, L. 2009. Preference functions that score rankings and maximum likelihood estimation. In *Proc. of IJCAI-09*. 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.

DWORK, C., KUMAR, R., NAOR, M., AND SIVAKUMAR, D. 2001. Rank aggregation methods for the web. In *Proc. of WWW-01*. 613–622.

EPHRATI, E. AND ROSENSCHEIN, J. S. 1991. The Clarke tax as a consensus mechanism among automated agents. In *Proc. of AAAI-91*. 173–178.

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.

GHOSH, S., MUNDHE, M., HERNANDEZ, K., AND SEN, S. 1999. Voting for movies: the anatomy of a recommender system. In *Proc. of the third annual conference on Autonomous Agents*. 434–435.

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

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

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).

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 2009)*.

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. of 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. of CCS-10*. 1–14.

XIA, L. AND CONITZER, V. 2008. Generalized scoring rules and the frequency of coalitional manipulability. In *Proc. of EC-08*. 109–118.

XIA, L. AND CONITZER, V. 2009. Finite local consistency characterizes generalized scoring rules. In *Proc. of IJCAI-09*. 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. of IJCAI-09*. 348–353.

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

Online Appendix to: Computing the Margin of Victory for Various Voting Rules

Lirong Xia, SEAS, Harvard University

A. PROOFS

Claim 1 If in the WMG of P^{NM} the weights on all edges $a \to b$ that are not compatible with $\succ'$ are reduced by at least $\min(2k, max(D_{P^{NM}}(a,b),0))$, then ranked pairs will output $\succ'$ in the resulting WMG.

Proof of Claim 1: Let G' denote the WMG obtained from WMG(P^{NM}) by reducing the weight on all edges $a \to b$ by at least $\min(2k, max(D_{P^{NM}}(a,b),0))$. For the sake of contradiction, suppose RP(G') does not output $\succ'$ for any tie-breaking mechanism. Then, we choose an output that contains an edge $a \to b$ not compatible with $\succ'$ (meaning that $b \succ' a$) with the highest weight. Because $b \succ' a$, there must exists a path from b to a in WMG(P^*) that precludes fixing $a \succ b$. Let $\mathbf{p}$ denote the path. Then, in WMG(P^*) the weight of each edge along $\mathbf{p}$ must be at least as high as the weight on $a \succ b$. We next show that in G' the weight of each edge along $\mathbf{p}$ must be at least as high as the weight on $a \to b$. Let $w_{G'}(\cdot)$ denote the weight of the edges in G'. For each edge $\vec{e}$ along $\mathbf{p}$ we discuss the following two cases.

— If $w_{G'}(a \to b) = 0$, then $w_{G'}(\vec{e}) \geq 0$. Otherwise suppose $w_{G'}(\vec{e}) < 0$, we have the following calculation.
$$\begin{aligned} D_{P^*}(\vec{e}) &\leq D_{P^{NM}}(\vec{e}) + k \leq w_{G'}(\vec{e}) - 2k + k \\ &< -k = w_{G'}(a \to b) - k \leq D_{P^{NM}}(a,b) - k \leq D_{P^*}(a,b), \end{aligned}$$
which is a contradiction.

— If $w_{G'}(a \to b) > 0$, then $w_{G'}(\vec{e}) \geq w_{G'}(a \to b)$. Otherwise suppose $w_{G'}(\vec{e}) < w_{G'}(a \to b)$, we have the following calculation.
$$\begin{aligned} D_{P^*}(\vec{e}) &\leq D_{P^{NM}}(\vec{e}) + k \leq w_{G'}(\vec{e}) + k \\ &< w_{G'}(a \to b) + 2k - k = D_{P^{NM}}(a,b) - k \leq D_{P^*}(a,b), \end{aligned}$$
which is a contradiction.

Therefore, $\mathbf{p}$ is a path that is as strong as $a \to b$ in G'. Because, $b \succ a$ is not fixed for any tie-breaking mechanism in G', at least one edge $b' \to a'$ along $\mathbf{p}$ is precluded before $a \succ b$ is considered. That is, there exists a path from a' to b' that is strictly stronger than $b' \to a'$. This means that there exists an edge that is not compatible with $\succ'$ and has a strictly higher weight than the weight on $b' \to a'$, which is as high as the weight on $a \to b$. However, this contradicts the maximality of the weight on $a \to b$. ■

Proposition 3 For any $\alpha \neq 1$ and any $\epsilon > 0$, there exists a profile P^{NM} such that MoV(P^{NM}, Copeland$_\alpha$) $\geq (\log m)^{1-\epsilon} \cdot OPT_{\text{DUCO}}(P^{NM}, \text{Copeland}_\alpha)$.

Proof of Proposition 3: Given an even number t, we construct an election with $m = \binom{t}{t/2} + 2$ alternatives. The set of alternatives is $\{d, c\} \cup \mathcal{E}$, where alternatives in $\mathcal{E}$ are indexed by subsets of size $t/2$ of $\{1, \ldots, t\}$. That is, $\mathcal{E} = \{E_S : S \subseteq \{1, \ldots, t\}, |S| = t/2\}$. For convenience, alternatives in $\mathcal{E}$ are given an (arbitrary) lexicographic order, which will be useful in constructing the votes.

We now define the profile $P^{NM} = P_1 \cup P_2$ as follows. For each $j \leq n$, let A_j denote the set of alternatives E_S with $j \in S$. It can be verified that $|A_j| = \binom{t-1}{t/2-1} = \binom{t}{t/2}/2$. Let

$$V_j = [d \succ A_j \succ e \succ c \succ \text{Rev}(\mathcal{E} \setminus A_j)]$$

DOI 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

$$W_j = [(\mathcal{E} \setminus A_j) \succ e \succ c \succ d \succ \text{Rev}(A_j)]$$

Let $P_1 = \{V_j, W_j : j \leq t\}$. P_2 is used to ensure that if no more than $t/2$ votes are changed, only c may have a higher Copeland score than d. To this end, let M_1 (respectively, M_2) denote a cyclic permutation over A_1 (respectively, $(\mathcal{E} \setminus A_2)$). Let

$$V = [A_1 \succ c \succ d \succ (\mathcal{E} \setminus A_1) \succ e]$$

$$W = [(\mathcal{E} \setminus A_1) \succ d \succ e \succ c \succ A_1]$$

and let $P_2 = \{M_1^i(M_2^i(\{V, W\})) : \forall i \leq |A_1|\}$.

We note that in the WMG of P, there is an edge from d to each alternative in $\mathcal{E}$ with weight t, an edge from d to e with weight $\binom{t}{t/2}$, an edge from each alternative in $(\mathcal{E} \setminus A_1)$ to e with weight $\binom{t}{t/2}$, and an edge from e to c with weight $2t$. We make the following observation on the Copeland scores of alternatives in P^{NM}.

— $S_C(P^{NM}, d) = \binom{t}{t/2} + 1 + \alpha$, because d ties with c and beats all other alternatives.
— $S_C(P^{NM}, c) = (\binom{t}{t/2} + 1)\alpha$, because c loses to e and ties with all other alternatives.
— $S_C(P^{NM}, e) \leq \binom{t}{t/2}/2 + 1$, because e loses to d and all alternatives in $(\mathcal{E} \setminus A_1)$.
— The Copeland score of any alternative $E \in \mathcal{E}$ is no more than $\frac{3}{4}\binom{t}{t/2} + 2$ due to P_2.

We note that the optimal solution to DUCO is 1, because if the manipulator votes for $[c \succ \text{Others}]$, then the Copeland score of c is the same as that of d, which means that d is not the unique winner. When t is large enough, we prove that MoV is at least $t/4$, that is, changing any $t/4$ votes cannot make d lose. When t is large enough, e and alternatives in $\mathcal{E}$ cannot beat d because each of them have at least two incoming edges with high weight ($> t/2$). For c, suppose for the sake of contradiction there exists a way to change $t/4$ votes to make the Copeland score of c at high as the Copeland score of d. Let $\{A_{j_1}, \ldots, A_{j_{t/4}}\}$ denote the subsets of $\mathcal{E}$ that are ranked higher than c in these votes. It follows that $A_{j_1} \cup \ldots \cup A_{j_{t/4}}$ must covers at least $|\mathcal{E}| - 2$ elements in $\mathcal{E}$, otherwise c would have tied with too many alternatives in $\mathcal{E}$ for d not to be the unique winner. However, for any $S \subseteq \{1, \ldots, t\}$ with $|S| = t/2$ and $S \cap \{j_1, \ldots, j_{t/4}\} = \emptyset$, S is not in any of $\{A_{j_1}, \ldots, A_{j_{t/4}}\}$, and the number of such sets S is larger than 2, which is a contradiction. Hence, MoV is at least $t/4$. When t is large enough we have $t/4 > (t \log t)^{1-\epsilon} > (\log m)^{1-\epsilon}$. ■

Lemma 1 Let $r = \text{GS}(f, g)$ be a non-redundant generalized scoring rule. For any profile P and any $\trianglerighteq \in H(f)$, there exists a preorder $\trianglerighteq_* \in H(f)$ that refines $\trianglerighteq$ and $(\text{Ord}(f(P)) \oplus \trianglerighteq_*)$ is a linear order.

Proof of Lemma 1: The proof is similar to the proof of Theorem 3 in [Xia and Conitzer 2008]. The idea is, we start with a pair of profiles that give us $\trianglerighteq$, then add votes sequentially to break ties in $\trianglerighteq$. Let d_{max} denote the maximum difference in pairs of components of any generalized score vector. That is, $d_{max} = \max_{L \in L(\mathcal{C}), k_1, k_2}\{|(f(L))_{k_1} - (f(L))_{k_2}|\}$

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 for every $L \in L(\mathcal{C})$, $(f(L))_{k_1} - (f(L))_{k_2}$ is unique. Otherwise $\{k_1, k_2\} \in \mathcal{K}_2$. That is, $\mathcal{K}_1$ is composed of pairs that are tied in every preorder in $H(f)$. Let P_1^1, P_2^1 be two profiles with the same size such that $\trianglerighteq = \text{Ord}(f(P_1^1) - f(P_2^1))$. We next find a refinement of $\trianglerighteq$ in $H(f)$ that does not contain a tie between any pairs in $\mathcal{K}_2$.

Let d_{min}^1 denote the minimum positive difference in pairs of components of $f(P_1^1) - f(P_2^1)$. That is, $d_{min}^1 = \min_{k_1, k_2}\{|(f(P_1^1) - f(P_2^1))_{k_1} - (f(P_1^1) - f(P_2^1))_{k_2}| > 0\}$. Suppose there exist $\{k_1, k_2\} \in \mathcal{K}_2$ such that the k_1th component of $f(P_1^1) - f(P_2^1)$ is the same as the k_2th component of $f(P_1^1) - f(P_2^1)$. Because $\{k_1, k_2\} \in \mathcal{K}_2$, there exist two votes V_1 and V_2 such that $(f(V_1) - f(V_2))_{k_1} \neq (f(V_1) - f(V_2))_{k_2}$. We let $P_1^2 = \lceil \frac{4d_{max}}{d_{min}^1} + 1 \rceil P_1^1 \cup \{V_1\}$ and $P_2^2 = \lceil \frac{4d_{max}}{d_{min}^1} +$

$1]P_2^1 \cup \{V_2\}$. It follows that $\text{Ord}(f(P_1^2) - f(P_2^2))$ strictly refines $\text{Ord}(f(P_1^1) - f(P_2^1))$, because the k_1th component is different from the k_2th component in $f(P_1^2) - f(P_2^2)$. We continue to find $P_1^3, P_2^3, P_1^4, P_2^4 \ldots$ until all pairs of components in $\mathcal{K}_2$ have different values. Let $\trianglerighteq_*$ denote the induced preorder of the final profiles.

Because r is non-redundant, if $\{k_1, k_2\} \in \mathcal{K}_1$, then in any profile P we have $(f(P))_{k_1} \neq (f(P))_{k_2}$. Therefore, $(\text{Ord}(f(P)) \oplus \trianglerighteq_*)$ is a linear order. ■

Theorem 16 All positional scoring rules, maximin, Coepalnd$_\alpha$, ranked pairs, plurality with runoff, Bucklin, STV, and approval are continuous generalized scoring rules.

Proof of Theorem 16: For each of these voting rules, we first define the f and g functions, then we show that they are continuous. Most of the f and g functions are the same as defined in [Xia and Conitzer 2008].

• **Positional scoring rules:** Suppose the scoring vector is $\vec{s}_m = (\vec{s}_m(1), \ldots, \vec{s}_m(m))$. The total generalized score vector will simply consist of the total scores of the individual alternatives. Let $K_{\vec{s}_m} = m$, $f_{\vec{s}_m}(V) = (\vec{s}_m(V, c_1), \ldots, \vec{s}_m(V, c_m))$, and $g_{\vec{v}}(\trianglerighteq)$ selects the top elements in $\trianglerighteq$.

For any profile P with order $\trianglerighteq$, we will show that either no refinement of $\trianglerighteq$ has different winner(s), or we can use two preorders in $H(f)$ to break ties in $\trianglerighteq$ and obtain different sets of winners. Suppose $g_{\vec{s}_m}(\trianglerighteq)$ contains a unique winner d. No matter how ties are broken in $\trianglerighteq$, d is still the unique winner, because the d-component is already strictly larger than other components. This means that the winner for any refinement of $\trianglerighteq$ are the same.

Suppose $g_{\vec{s}_m}(\trianglerighteq)$ contains at least two alternatives $\{c, d\}$. We next show that any of these alternatives can be made the unique winner by using some preorders in $H(f_{\vec{s}_m})$ to break ties. For any alternative d, let M_d denote a cyclic permutation among $\mathcal{C} \setminus \{d\}$. Let $P_1 = \{[d \succ M_d^i(\text{Others})] : i \leq m-1\}$, $P_2 = \{[M_d^i(\text{Others}) \succ d] : i \leq m-1\}$. Let $\trianglerighteq_d = \text{Ord}(f(P_1) - f(P_2))$. Let $\trianglerighteq'_d$ denote the preorder in $H(f_{\vec{s}_m})$ obtained by applying Lemma 1 on $\trianglerighteq_d$. It follows that d is strictly preferred to all other alternatives in $\trianglerighteq'_d$, which means that $g_{\vec{s}_m}(\trianglerighteq \oplus \trianglerighteq'_d) = \{d\}$. Similarly, there exists $\trianglerighteq'_c \in H(f_{\vec{s}_m})$ such that $(\trianglerighteq \oplus \trianglerighteq'_c)$ is a linear order and $g_{\vec{s}_m}(\trianglerighteq \oplus \trianglerighteq'_c) = \{c\}$, which means that we can always find two preorders in $\text{Nbr}(\trianglerighteq)$ that have different winners. Therefore, $\text{GS}(f_{\vec{s}_m}, g_{\vec{s}_m})$ is continuous.

• **Copeland$_\alpha$:** For Copeland, the total generalized score vector will consist of the scores in the pairwise elections. Let

— $K_C = m(m-1)$; the components are indexed by pairs (i, j) such that $i, j \leq m$, $i \neq j$.

— $(f_C(V))_{(i,j)} = \begin{cases} 1 & \text{if } c_i \succ_V c_j \\ 0 & \text{otherwise} \end{cases}$.

— g_C selects the winner based on $\trianglerighteq$ as follows. For each pair $i \neq j$, if $o_{(i,j)} \triangleright o_{(j,i)}$, then add 1 point to i's Copeland score; if $o_{(j,i)} \triangleright o_{(i,j)}$, then add 1 point to j's Copeland score; if tied, then add α to both i's and j's Copeland scores. The winners are the alternatives with the highest Copeland score.

Let P be a profile whose induced preorder is $\trianglerighteq$. We will show that if there exists a refinement $\trianglerighteq'$ of $\trianglerighteq$ such that $g_C(\trianglerighteq') \neq g_C(\trianglerighteq)$, then it is impossible that $g_C(\trianglerighteq \oplus \trianglerighteq_*)$ is the same for all $\trianglerighteq_* \in H(f_C)$ where $(\trianglerighteq \oplus \trianglerighteq_*)$ is a linear order. For Copeland, refinements of $\trianglerighteq$ effectively break ties in pairwise elections in $\trianglerighteq$. For any pair of alternative c_1, c_2, we let $\trianglerighteq_{(c_1,c_2)}$ denote the preorder obtained by applying Lemma 1 to $\text{Ord}(f_C(c_1 \succ \text{Others} \succ c_2) - f_C(c_2 \succ \text{Others} \succ c_1))$. It follows that for every $c' \neq c_1$, $o_{(c_1,c')} \triangleright_{(c_1,c_2)} o_{(c',c_1)}$, and for every $c' \neq c_2$, $o_{(c',c_2)} \triangleright_{(c_1,c_2)} o_{(c_2,c')}$. That is, if we use $\trianglerighteq_{(c_1,c_2)}$ to break ties in pairwise elections, c_1 wins all tied pairwise elections and c_2 loses all tied pairwise elections. Suppose $\trianglerighteq'$ is a refinement of $\trianglerighteq$ and $g_C(\trianglerighteq') \neq g_C(\trianglerighteq)$. We first show that there exits $\trianglerighteq'_* \in H(f_C)$ such that $(\trianglerighteq \oplus \trianglerighteq'_*)$ is a linear order and $g_C(\trianglerighteq) \neq g_C(\trianglerighteq \oplus \trianglerighteq'_*)$, by discussing the following two cases.

Case 1: there exists $d \in g_C(\trianglerighteq)$ such that $d \notin g_C(\trianglerighteq')$. Let c be an arbitrary alternative in $g_C(\trianglerighteq')$. Let $\trianglerighteq'_* = \trianglerighteq_{(c,d)}$. It follows that the Copeland score of c in $(\trianglerighteq \oplus \trianglerighteq'_*)$ is at least the Copeland score of c in $\trianglerighteq'$, and the Copeland score of d in $(\trianglerighteq \oplus \trianglerighteq'_*)$ is at most the Copeland score of d in $\trianglerighteq'$. Because $d \notin g_C(\trianglerighteq')$, we have $d \notin g_C(\trianglerighteq \oplus \trianglerighteq'_*)$, which means that $g(\trianglerighteq \oplus \trianglerighteq'_*) \neq g(\trianglerighteq)$.
Case 2: there exists $d \in g_C(\trianglerighteq')$ such that $d \notin g_C(\trianglerighteq)$. Similarly to Case 1, let c denote an arbitrary alternative in $g_C(\trianglerighteq)$. Because $d \in g_C(\trianglerighteq')$, the Copeland score of c is no more than the Copeland score of d in $\trianglerighteq'$. Let $\trianglerighteq'_* = \trianglerighteq_{(d,c)}$. It follows that the Copeland score of d (respectively, c) in $(\trianglerighteq \oplus \trianglerighteq'_*)$ is at least (respectively, no more than) the Copeland score of d (respectively, c) in $\trianglerighteq'$. Therefore, the Copeland score of c is no more than the Copeland score of d in $(\trianglerighteq \oplus \trianglerighteq'_*)$. If the Copeland score of c is smaller than the Copeland score of d in $(\trianglerighteq \oplus \trianglerighteq'_*)$, then $c \notin g_C(\trianglerighteq \oplus \trianglerighteq'_*)$. If the Copeland score of c equals to the Copeland score of d in $(\trianglerighteq \oplus \trianglerighteq'_*)$, then either both of them are in $g_C(\trianglerighteq \oplus \trianglerighteq'_*)$, or none of them is in $g_C(\trianglerighteq \oplus \trianglerighteq'_*)$. In both cases we have $g_C(\trianglerighteq \oplus \trianglerighteq'_*) \neq g_C(\trianglerighteq)$.

Finally, we show that it is impossible that $g_C(\trianglerighteq \oplus \trianglerighteq_*)$ is the same for all $\trianglerighteq_* \in H(f_C)$ where $(\trianglerighteq \oplus \trianglerighteq_*)$ is a linear order. Suppose for the sake of contradiction, for all $\trianglerighteq_* \in H(f_C)$, $g_C(\trianglerighteq \oplus \trianglerighteq_*)$ is the same. We have shown above that $g_C(\trianglerighteq \oplus \trianglerighteq_*) \neq g_C(\trianglerighteq)$. We discuss the following two cases.

Case 1: there exists $d \in g_C(\trianglerighteq)$ such that $d \notin g_C(\trianglerighteq \oplus \trianglerighteq_*)$. Let c be an arbitrary alternative in $g_C(\trianglerighteq \oplus \trianglerighteq_*)$. Then, the Copeland score of d (respectively, c) in $(\trianglerighteq \oplus \trianglerighteq_{(d,c)})$ is as high as (respectively, no more than) the Copeland score of d (respectively, c) in $\trianglerighteq$. Therefore, the Copeland score of d is as high as the Copeland score of c in $(\trianglerighteq \oplus \trianglerighteq_{(d,c)})$, which means that either $d \in g_C(\trianglerighteq \oplus \trianglerighteq_{(d,c)})$, or both $\{c, d\}$ are not in $g_C(\trianglerighteq \oplus \trianglerighteq_{(d,c)})$. Both cases contradicts the assumption that $g_C(\trianglerighteq \oplus \trianglerighteq_{(d,c)}) = g_C(\trianglerighteq \oplus \trianglerighteq_*)$.
Case 2: there exists $d \in g_C(\trianglerighteq \oplus \trianglerighteq_*)$ such that $d \notin g_C(\trianglerighteq)$. Let c be an arbitrary alternative in $g_C(\trianglerighteq)$. We note that the Copeland score of c is strictly higher than the Copeland score of d in $(\trianglerighteq \oplus \trianglerighteq_{(c,d)})$, which means that d is not a winner for $(\trianglerighteq \oplus \trianglerighteq_{(c,d)})$. This contradicts the assumption that $g_C(\trianglerighteq \oplus \trianglerighteq_{(c,d)}) = g_C(\trianglerighteq \oplus \trianglerighteq_*)$.

Therefore, GS(f_C, g_C) is continuous.
• **STV:** For STV, we will use a total generalized score vector 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 $\trianglerighteq$ as follows. Fix a tie-breaking mechanism over the alternatives. In the first round, let j_1 to be the index such that $o_{(\emptyset,j_1)}$ is ranked the lowest in $\trianglerighteq$ among all $o_{(\emptyset,j)}$ (if there are multiple such j's, then we use the tie-breaking mechanism 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 $\trianglerighteq$ among all $o_{(S_{i-1},j)}$; finally, the winner is the unique alternative in $(\mathcal{C} \setminus S_{m-1})$. g_{STV} selects all alternatives that are winners w.r.t. some tie-breaking mechanism.

Let P be a profile whose induced preorder is $\trianglerighteq$. For STV, refinements of $\trianglerighteq$ effectively break ties in each round. Therefore, suppose for some refinement $\trianglerighteq'$ of $\trianglerighteq$ we have $g_{STV}(\trianglerighteq') \neq g_{STV}(\trianglerighteq)$, then $g_{STV}(\trianglerighteq') \subset g_{STV}(\trianglerighteq)$, which means that $|g_{STV}(\trianglerighteq)| \geq 2$. We next show that for each $d \in g_{STV}(\trianglerighteq)$, there exits $\trianglerighteq'_* \in H(f_{STV})$ such that $(\trianglerighteq \oplus \trianglerighteq'_*)$ is a linear order and whose unique winner is d. Suppose for some tie-breaking mechanism, the order of elimination is, w.l.o.g. $c_1, \ldots c_{m-1}, d$. Let $\trianglerighteq'_*$ denote the preorder obtained by Lemma 1 on Ord$(f_{STV}([d \succ \text{Others}]) - f_{STV}([c_1 \succ c_2 \succ$

$\cdots \succ c_{m-1} \succ d]))$. If we apply STV on $(\trianglerighteq \oplus \trianglerighteq'_*)$, then in the first round c_1 should be eliminated, because for any alternative $c \neq c_1$, $o_{(\emptyset,c)} \triangleright'_* o_{(\emptyset,c_1)}$; in the second round c_2 should be eliminated; etc. It follows that $g_{STV}(\trianglerighteq \oplus \trianglerighteq'_*) = \{d\}$. Therefore, there exist two preorders $\trianglerighteq_1, \trianglerighteq_2 \in H(f_{STV})$ such that (1) $(\trianglerighteq \oplus \trianglerighteq_1)$ and $(\trianglerighteq \oplus \trianglerighteq_2)$ are linear orders, and (2) $g_{STV}(\trianglerighteq \oplus \trianglerighteq_1) \neq g_{STV}(\trianglerighteq \oplus \trianglerighteq_2)$. Therefore, GS$(f_{STV}, g_{STV})$ is continuous.

• **Maximin:** For maximin, we use the same total generalized score vector as for Copeland, that is, the vector of all scores in pairwise elections. Let

— $K_M = m(m-1)$; the components are indexed by pairs (i,j) such that $i,j \leq m, i \neq j$.

— $(f_M(V))_{(i,j)} = \begin{cases} 1 & \text{if } c_i \succ_V c_j \\ 0 & \text{otherwise} \end{cases}$

— $g_M(\trianglerighteq)$ chooses an alternative c_i such that for any $i' \leq m$, $i' \neq i$, there exists $j' < m$, $j' \neq i'$ such that for any $j \leq m$, $j \neq i$, we have $o_{(i,j)} \trianglerighteq o_{(i',j')}$.

The proof of continuity for GS(f_M, g_M) is similar to the proof of continuity for positional scoring rules. Let P be a profile whose induced preorder is $\trianglerighteq$. If $|g_M(\trianglerighteq)| = 1$, then in any refinement of $\trianglerighteq$ the winner is unique. Suppose $|g_M(\trianglerighteq)| \geq 2$. Let $\trianglerighteq_d$ denote the preorder obtained by Lemma 1 on Ord$(f_M([d \succ \text{Others}]) - f_M([\text{Others} \succ d]))$. It follows that $(\trianglerighteq \oplus \trianglerighteq_d)$ is a linear order and $g_M(\trianglerighteq \oplus \trianglerighteq_d) = \{d\}$. This proves that GS$(f_M, g_M)$ is continuous.

• **Ranked pairs:** We use the same total generalized score vector as for Copeland and maximin, that is, the vector of all scores in pairwise elections. Let

— $K_{RP} = m(m-1)$; the components are indexed by pairs (i,j) such that $i,j \leq m, i \neq j$.

— $(f_{RP}(V))_{(i,j)} = \begin{cases} 1 & \text{if } c_i \succ_V c_j \\ 0 & \text{otherwise} \end{cases}$

— g_{RP} selects the winner based on $\trianglerighteq$ as follows. We fix a tie-breaking mechanism over pairs of alternatives. In each step, we consider a pair of alternatives c_i, c_j that we have not previously considered; specifically, we choose the remaining pair that is ranked highest in $\trianglerighteq$ (if there are more than one of such pairs, we use the tie-breaking mechanism to select the most preferred one). We then fix the order $c_i \succ c_j$, unless this contradicts previous orders that we fixed (that is, it violates transitivity). We continue until we have considered all pairs of alternatives. The alternative at the top of the ranking wins. g_{RP} selects all alternatives that are winners w.r.t. some tie-breaking mechanism.

The proof of continuity for GS(f_{RP}, g_{RP}) is similar to the proof of continuity for STV. Let P be a profile whose induced preorder is $\trianglerighteq$. Since effectively refinements of $\trianglerighteq$ are only used to break ties among edges, for any refinement $\trianglerighteq'$ of $\trianglerighteq$, if $g_{RP}(\trianglerighteq') \neq g_{RP}(\trianglerighteq)$, then $|g_{RP}(\trianglerighteq)| \geq 2$. For any alternative d in $g_{RP}(\trianglerighteq)$, let L_d denote the output order, where d is ranked in the top position. Let $\trianglerighteq_d$ denote the preorder obtained by Lemma 1 on Ord$(f_{RP}([L_d]) - f_{RP}([\text{Rev}(L_d)]))$. We next prove that d is the unique winner in $g_{RP}(\trianglerighteq \oplus \trianglerighteq_d)$. The proof is similar to the proof of Claim 1. Suppose for the sake of contradiction d is not the unique winner in $g_{RP}(\trianglerighteq \oplus \trianglerighteq_d)$, then there exists a tie-breaking mechanism such that the output order L' is different from L_d. Let $a \to b$ be an edge with heaviest weight that is compatible with L' but not L_d. There must exists a path with same or higher weight from b to a in the WMG of $\trianglerighteq$. It follows that there must exists a path with strictly higher weight from b to a in the WMG of $(\trianglerighteq \oplus \trianglerighteq_d)$. Therefore, this path cannot be fixed before $a \succ b$ is considered, but that means that there exists an edge $a' \to b'$ that is compatible with L' but not compatible with L_d, and whose weight is strictly higher than the weight of $a \to b$. This contradicts the maximality of $a \to b$. Hence, $(\trianglerighteq \oplus \trianglerighteq_d)$ is a linear order and $g_{RP}(\trianglerighteq \oplus \trianglerighteq_d) = \{d\}$. This proves that GS$(f_{RP}, g_{RP})$ is continuous.

• **Bucklin:** For Bucklin, the generalized score vector consists of number of times that each alternative is ranked among top j positions for all $j \leq m$, and a threshold whose value is always $n/2$. Let

— $K_B = m^2 + 1$. The first type of components are index by (i, j) where $i, j \leq m$. The second type has one component indexed by 0.

— $(f_{\text{Pluo}}(V))_{(i,j)} = \begin{cases} 1 & \text{if } c_i \text{ is ranked within top } j \text{ in } V \\ 0 & \text{otherwise} \end{cases}$, $(f_{\text{Pluo}}(V))_0 = 1/2$.

— The Bucklin score of an alternative c_i is the smallest j such that $o_{(i,j)} \rhd o_0$. $g_{\text{Pluo}}(\unrhd)$ chooses alternatives with the smallest Bucklin score.

Let P be a profile whose induced preorder is $\unrhd$. We discuss the following three cases. For Case 1 and Case 2, we show that there exists $\unrhd_1, \unrhd_2 \in H(f_B)$ such that $(\unrhd \oplus \unrhd_1)$ and $(\unrhd \oplus \unrhd_2)$ are linear orders and $g_B(\unrhd \oplus \unrhd_1) \neq g_B(\unrhd \oplus \unrhd_2)$. For Case 3, we show that for any refinement $\unrhd'$ of $\unrhd$, $g_B(\unrhd') = g_B(\unrhd)$.

Case 1: There exists an alternative $c_i \notin g_B(\unrhd)$ such that $o_{(c_i,j)} =_{\unrhd} o_0$ for some j no more than the minimum Bucklin score. Let M denote a cyclic permutation among $\mathcal{C} \setminus \{c_i\}$. Let $\unrhd_1$ (respectively, $\unrhd_2$) denote the preorder obtained by Lemma 1 on Ord $\left(f_B(\{[c_i \succ M^j(\text{Others})] : j \leq m-1\}) - f_B(\{[M^j(\text{Others}) \succ c_i] j \leq m-1\})\right)$ (respectively, Ord$(f_B([\text{Others} \succ c_i]) - f_B([c_i \succ \text{Others}])))$. It follows that $c_i \in g_B(\unrhd \oplus \unrhd_1)$ and $c_i \notin g_B(\unrhd \oplus \unrhd_2)$, which means that $g_B(\unrhd \oplus \unrhd_1) \neq g_B(\unrhd \oplus \unrhd_2)$.

Case 2: $|g_B(\unrhd)| \geq 2$ and there exists an alternative $c_i \in g_B(\unrhd)$ such that $o_{(i,j)} =_{\unrhd} o_0$ for some j strictly smaller than the minimum Bucklin score in $\unrhd$. For each $d \in g_B(\unrhd)$ with $d \neq c_i$, we define $V_d = [c_i \succ \text{Others} \succ d]$ and $W_d = [d \succ \text{Others} \succ c_i]$. Let $P_1 = \{V_d : d \in g_B(\unrhd), d \neq c_i\}$ and $P_2 = \{W_d : d \in g_B(\unrhd), d \neq c_i\}$. Let $\unrhd_1$ (respectively, $\unrhd_2$) denote the preorder obtained by Lemma 1 on Ord$(f_B(P_1) - f_B(P_2))$ (respectively, Ord$(f_B(P_2) - f_B(P_1))$). It follows that $\{c_i\} = g_B(\unrhd \oplus \unrhd_1)$ and $\{c_i\} \neq g_B(\unrhd \oplus \unrhd_2)$, which means that $g_B(\unrhd \oplus \unrhd_1) \neq g_B(\unrhd \oplus \unrhd_2)$.

Case 3: Let $\unrhd'$ be an arbitrary refinement of $\unrhd$. Since Case 1 does not hold, for $\unrhd'$, the Bucklin score of any alternative not in $g_B(\unrhd)$ cannot be the same or lower than the minimum Bucklin score in $\unrhd$, which means that $g_B(\unrhd') \subseteq g_B(\unrhd)$. Since Case 2 also does not hold, there are following two possibilities. Possibility 1: $|g_B(\unrhd)| \geq 2$ and there does not exists $c_i \in g_B(\unrhd)$ such that $o_{i,j}$ is tied with o_0 for some j strictly smaller than the minimum Bucklin score in $\unrhd$. Therefore, the minimum Bucklin score in $\unrhd'$ is the same as the minimum Bucklin score in $\unrhd$, which means that $g_B(\unrhd') = g_B(\unrhd)$. Possibility 2: $|g_B(\unrhd)| = 1$, which means that $g_B(\unrhd') = g_B(\unrhd)$ because $g_B(\unrhd') \subseteq g_B(\unrhd)$.

This proves that GS(f_B, g_B) is continuous.

• **Plurality with runoff:** For plurality with runoff, the generalized score vector contains two types of components. The first type is indexed by $i \leq m$, which represents the number of times that each alternative is ranked in the top position. The second type is indexed by pairs (i, j) where $i, j \leq m$ and $i \neq j$, which encode the results of pairwise elections as for Copeland and maximin.

— $K_{\text{Pluo}} = m + m(m-1)$.

— $(f_{\text{Pluo}}(V))_i = \begin{cases} 1 & \text{if } c_i \text{ is at the top of } V \\ 0 & \text{otherwise} \end{cases}$, $(f_{\text{Pluo}}(V))_{(i,j)} = \begin{cases} 1 & \text{if } c_i \succ_V c_j \\ 0 & \text{otherwise} \end{cases}$

— $g_{\text{Pluo}}(\unrhd)$ chooses an alternative c_i if there exists $j \neq i$ such that the o_i and the o_j are ranked highest among the first type of components, and $o_{(i,j)} \unrhd o_{(j,i)}$.

The proof of continuity for GS$(f_{\text{Pluo}}, g_{\text{Pluo}})$ is similar to the proof of continuity for STV and ranked pairs. Let P be a profile whose induced preorder is $\unrhd$. Since effectively refinements of $\unrhd$ are only used to break ties, for any refinement $\unrhd'$ of $\unrhd$, if $g_{\text{Pluo}}(\unrhd') \neq g_{\text{Pluo}}(\unrhd)$, then $|g_{\text{Pluo}}(\unrhd)| \geq 2$. For any alternative d in $g_{\text{Pluo}}(\unrhd)$ that beats c in the runoff, let $\unrhd_d$ denote the preorder obtained by Lemma 1 on Ord$(f_{\text{Pluo}}([d \succ c \succ \text{Others}], [c \succ d \succ \text{Others}]) - f_{\text{Pluo}}([\text{Others} \succ c \succ d], [\text{Others} \succ c \succ d]))$. It follows that $\{c, d\}$ are the only pair of alternatives entering the runoff, where d beats c. Therefore, each alternative in $g_{\text{Pluo}}(\unrhd)$ can be made the unique winner by using some preorders in $H(f_{\text{Pluo}})$ to break ties. This proves that GS$(f_{\text{Pluo}}, g_{\text{Pluo}})$ is continuous.

• **Approval:** For approval, we need to slightly change the definition of generalized scoring rules. All results proved in this section about generalized scoring rules also applies to this definition. The definition is similar to the definition for positional scoring rules. Let $K_{App} = m$. Each component of the generalized score vector corresponds to an alternative. Let f_{App} be a mapping from all sets of subsets of $\mathcal{C}$ (that is, all possible sets of "approved" alternatives) to $\mathbb{R}^m$. For any set $S \subseteq \mathcal{C}$, the ith component of $f_{\text{App}}(S)$ is 1 if and only if $c_i \in S$, and g_{App} selects the top elements in $\trianglerighteq$.

Similar to the proof for positional scoring rules, we show that given a profile P whose induced preorder is $\trianglerighteq$, if $g_{\text{App}}(\trianglerighteq)$ contains only one winner, then any refinement cannot change it. If $g_{\text{App}}(\trianglerighteq)$ contains two or more winners, then for each winner d we let $\trianglerighteq'_*$ denote the preorder obtained by Lemma 1 on $\text{Ord}(f_{\text{App}}(\{d\}) - f_{\text{App}}(\mathcal{C} \setminus \{d\}))$. Then, $g_{\text{App}}(\trianglerighteq \oplus \trianglerighteq'_*) = \{d\}$. Therefore, each alternative in $g_{\text{App}}(\trianglerighteq)$ can be made the unique winner by using some preorders in $H(f_{\text{App}})$ to break ties. This proves that $\text{GS}(f_{\text{App}}, g_{\text{App}})$ is continuous.

■