Finite Local Consistency Characterizes Generalized Scoring Rules

Lirong Xia
Department of Computer Science
Duke University
Durham, NC 27708, USA
lxia@cs.duke.edu

Vincent Conitzer
Department of Computer Science
Duke University
Durham, NC 27708, USA
conitzer@cs.duke.edu

Abstract

An important problem in computational social choice concerns whether it is possible to prevent manipulation of voting rules by making it computationally intractable. To answer this, a key question is how *frequently* voting rules are manipulable. We [Xia and Conitzer, 2008] recently defined the class of *generalized scoring rules (GSRs)* and characterized the frequency of manipulability for such rules. We showed, by examples, that most common rules seem to fall into this class. However, no natural axiomatic characterization of the class was given, leaving the possibility that there are natural rules to which these results do not apply. In this paper, we characterize the class of GSRs based on two natural properties: it is equal to the class of rules that are *anonymous* and *finitely locally consistent*.

Generalized scoring rules also have other uses in computational social choice. For these uses, the order of the GSR (the dimension of its score vector) is important. Our characterization result implies that the order of a GSR is related to the minimum number of locally consistent components of the rule. We proceed to bound the minimum number of locally consistent components for some common rules.

1 Introduction

Computational social choice is a rapidly growing research area within artificial intelligence and multiagent systems. Often, agents have different preferences over a set of alternatives, and a common decision must be made. One way to resolve this problem is to ask each agent to report her preferences (rank the alternatives), and then apply a *voting rule* to the vector of rankings to select the winning alternative.

Sometimes, an agent can make herself better off by reporting false preferences. In this case, we say that the voting rule is *manipulable*. If a voting rule is not manipulable, then it is said to be *strategy-proof*. Unfortunately, no voting rule that satisfies certain extremely natural properties can be strategy-proof, due to the celebrated Gibbard-Satterthwaite Theorem [Gibbard, 1973; Satterthwaite, 1975]. In recent years, researchers in the computational social choice community have tried to circumvent this impossibility result by investigating whether there are desirable voting rules that are computationally hard to manipulate. The idea is that even if a successful manipulation exists, it will not occur unless the agent can discover the manipulation. It has been shown that, in various circumstances—that is, for various rules, numbers of manipulators, manipulation goals, *etc.*—finding a successful manipulation is NP-hard [Bartholdi *et al.*, 1989; Bartholdi and Orlin, 1991; Conitzer and Sandholm, 2003; Elkind and Lipmaa, 2005; Conitzer *et al.*, 2007; Hemaspaandra and Hemaspaandra, 2007; Faliszewski *et al.*, 2008; Xia *et al.*, 2009].

Unfortunately, all of these results are *worst-case* results. That is, there are *some* cases where manipulation is hard (if P≠NP), but manipulation might be easy most of the time. Some recent research suggests that this is indeed the case [Procaccia and Rosenschein, 2007b; Conitzer and Sandholm, 2006; Procaccia and Rosenschein, 2007a; Zuckerman *et al.*, 2009; Xia and Conitzer, 2008]. Specifically, we [Xia and Conitzer, 2008] introduced a new class of voting rules called *generalized scoring rules (GSRs)*, and studied the probability that the manipulators are able to influence the outcome of the election, given a distribution over the nonmanipulators' votes. Under some minor conditions on the rule and distribution, we showed the following results (generalizing earlier work for positional scoring rules [Procaccia and Rosenschein, 2007a]). When the number of manipulators is relatively small ($o(n^p)$, where n is the number of voters, and $p < \frac{1}{2}$), this probability goes to zero. On the other hand, when the number of manipulators is relatively large ($O(n^p)$ and $o(n)$, where $\frac{1}{2} < p$), the probability that the manipulators can make *any* alternative win goes to one (under some conditions on the distribution and the rule). This only leaves a knife-edge case where there are $\Theta(\sqrt{n})$ manipulators; see [Walsh, 2009] for an empirical study of this case.

Moreover, to argue the relevance of our results, we showed that a variety of common rules fall into the category of GSRs, and remarked that we did not know of any commonly studied rule that is not a GSR. (However, we did not give any formal result about the generality of this class of rules.) The apparent wide applicability of GSRs makes this class potentially interesting from the perspective of other problems in computational social choice. Indeed, some such uses are quite obvious. GSRs map every vote to a vector of scores (which are not necessarily associated with alternatives), and the outcome of the rule is based strictly on the sum of these vectors. As a result, the votes of a subset of the electorate can be summarized completely by the sum of their score vectors. (The problem of summarizing the votes of a subelectorate was recently introduced and studied [Chevaleyre *et al.*, 2009].) In fact, the definition of GSRs is even more restrictive: the final outcome only depends on direct comparisons among the components of the summed score vector. For example, the outcome may depend on a comparison between the first component and the third component of the summed vector; then, it does not matter (for this comparison) whether these components are 42 and 50, respectively, or 101 and 967, because in both cases component 1 is smaller. Because of this, the GSR framework is also useful for preference elicitation, specifically, for determining whether enough information has been elicited from the voters to declare the winner. In particular, if it becomes clear that the remaining (not yet elicited) information about the voters' preferences can no longer change any of the comparisons in scores, then we can terminate elicitation.

One important open problem posed in [Xia and Conitzer, 2008] is to give a natural *axiomatic* characterization of the class of GSRs, based on some criteria that we would like voting rules to have. Such axiomatizations are important because they give us deeper insight into the rules, and can often be used to prove important results about these rules. For GSRs, having an axiomatic characterization is especially important in order to know how the frequency-of-manipulability results in [Xia and Conitzer, 2008], which are negative results for the agenda of making manipulation computationally hard, might be circumvented. Axiomatic characterization of voting rules is a common topic in the social choice literature. For two alternatives, the majority rule has been characterized in [May, 1952]. Young characterized positional scoring rules by consistency, neutrality, and anonymity [Young, 1975]. When ties are not possible (the voting rule always produces a single winner), we say that an anonymous voting rule r is *consistent* if the following holds: for any profiles (multisets of votes) P_1 and P_2 with $r(P_1) = r(P_2) = c$, we must have $r(P_1 \cup P_2) = c$. Young's paper [Young, 1975] also allows for ties, so that r is a set-valued function; the more general definition of consistency is that if $r(P_1) \cap r(P_2) \neq \emptyset$, then $r(P_1 \cup P_2) = r(P_1) \cap r(P_2)$. (We note that the union of two multisets adds the multiplicities.) However, in this paper we will only consider rules that always produce a single winner.

In this paper, we introduce a new criterion for voting rules: finite local consistency. A voting rule satisfies *finite local consistency* (FLC) if the set of all profiles can be partitioned into finitely many parts, such that the voting rule is consistent within each part. The minimum number of parts for a rule is the *degree of consistency* for the rule. For example, a consistent rule has degree of consistency 1. We then characterize generalized scoring rules by anonymity and finite local consistency, and we show that the order of a GSR (that is, the dimension of the score vector) is related to the degree of consistency of the rule. It follows that Dodgson's rule is not a GSR, because it does not satisfy homogeneity [Brandt, 2009], and FLC is a stronger property than homogeneity. Finally, we give lower and upper bounds on the degree of consistency for some common voting rules.

2 Preliminaries

Basics of voting. Let $\mathcal{C} = \{c_1, \ldots, c_m\}$ be the set of *alternatives* (or *candidates*). A linear order on $\mathcal{C}$ is a transitive, antisymmetric, and total relation on $\mathcal{C}$. The set of all linear orders on $\mathcal{C}$ is denoted by $L(\mathcal{C})$. An n-voter profile P on $\mathcal{C}$ consists of n linear orders on $\mathcal{C}$. That is, $P = (V_1, \ldots, V_n)$, where for every $i \leq n$, $V_i \in L(\mathcal{C})$. The set of all profiles on $\mathcal{C}$ is denoted by $P(\mathcal{C})$. In the remainder of the paper, m denotes the number of alternatives and n denotes the number of voters (agents). A *(voting) rule* r is a function from the set of all profiles on $\mathcal{C}$ to $\mathcal{C}$, that is, $r : P(\mathcal{C}) \to \mathcal{C}$. A rule r is *anonymous* if the output of the rule is insensitive to the names of the voters; in this case, we can think of a profile as a multiset rather than a vector. An anonymous rule r is *homogenous* if for any profile P and any $k \in \mathbb{N}$, $k > 0$, $r(P) = r(kP)$, where kP is the profile composed of k copies of P.

Common rules. Common rules are anonymous and homogenous, including the following example rules:

Copeland: For any two alternatives c_i and c_j, we can simulate a *pairwise election* between them, by seeing how many votes prefer c_i to c_j, and how many prefer c_j to c_i. Then, an alternative receives one point for each win in a pairwise election. Typically, an alternative also receives half a point for each pairwise tie. The winner is the alternative who has the highest score.

Bucklin: An alternative c's Bucklin score is the smallest number k such that more than half of the votes rank c among the top k alternatives. The winner is the alternative that has the smallest Bucklin score. Ties are broken by the number of votes that rank an alternative among the top k.

Other common voting rules include *positional scoring rules*, *STV*, *maximin*, and *ranked pairs*. Due to space constraint, we omit their definitions here. All the above rules need some tiebreaking mechanism. In this paper, we assume nothing about this mechanism, except that in each voting rule, ties are broken in a consistent way.

Generalized scoring rules [Xia and Conitzer, 2008]. Let $K = \{1, \ldots, k\}$.

Definition 1 *For any $\vec{a}, \vec{b} \in \mathbb{R}^k$, we say that $\vec{a}$ and $\vec{b}$ are* equivalent with respect to K,[1] *denoted by $\vec{a} \sim_K \vec{b}$, if for any $i, j \in K$, $a_i \geq a_j \Leftrightarrow b_i \geq b_j$ (where a_i denotes the ith component of the vector $\vec{a}$,* etc.*).*

Definition 2 *A function $g : \mathbb{R}^k \to \mathcal{C}$ is* compatible with K *if for any $\vec{a}, \vec{b} \in \mathbb{R}^k$, $\vec{a} \sim_K \vec{b} \Rightarrow g(\vec{a}) = g(\vec{b})$.*

That is, for any mapping g that is compatible with K, $g(\vec{a})$ is completely determined by comparisons within K. Generalized scoring rules are defined as follows.

Definition 3 *Let $k \in \mathbb{N}$, $f : L(\mathcal{C}) \to \mathbb{R}^k$ (a* generalized scoring function*), and $g : \mathbb{R}^k \to \mathcal{C}$ where g is compatible with*

[1] In the original definition of generalized scoring rules, $\vec{a}$ and $\vec{b}$ are allowed to be equivalent w.r.t. any partition of K. In this paper, we only consider the special case where the partition consists of a single element K.

K *(a* decision function*). f and g determine the* (unweighted) generalized scoring rule $GS(f,g)$ *as follows. For any profile of votes* $V_1,\ldots,V_n \in L(\mathcal{C})$, $GS(f,g)(V_1,\ldots,V_n) = g(\sum_{i=1}^n f(V_i))$. *We say that* $GS(f,g)$ *is* of order k.

That is, every vote results in a vector of scores according to f, and g decides the winner based on comparisons between the total scores. For example, Copeland can be modeled as a generalized scoring rule in the following way. The total generalized score vector will consist of the scores in the pairwise elections. Let
– $k_{Copeland} = m(m-1)$; the components are indexed by pairs (i,j) such that $i,j \leq m$, $i \neq j$.
– $(f_{Copeland}(V))_{(i,j)} = \begin{cases} 1 & \text{if } c_i \succ_V c_j \\ 0 & \text{otherwise} \end{cases}$
– $g_{Copeland}$ selects the winner based on $f_{Copeland}(P)$ as follows. For each pair $i \neq j$, if $(f_{Copeland}(P))_{(i,j)} > (f_{Copeland}(P))_{(j,i)}$, then add 1 point to i's Copeland score; if $(f_{Copeland}(P))_{(j,i)} > (f_{Copeland}(P))_{(i,j)}$, then add 1 point to j's Copeland score; if tied, then add 0.5 to both i's and j's Copeland scores. The winner is the alternative that gets the highest Copeland score. The other voting rules we studied can also be modeled as GSRs:

Proposition 1 ([Xia and Conitzer, 2008]) *All positional scoring rules, Copeland, STV, maximin, and ranked pairs are generalized scoring rules.*

As for Bucklin, the results in this paper imply that it is also a generalized scoring rule.

3 Finite local consistency

In this section, we introduce *finite local consistency*.

Definition 4 *Let S be a set of profiles. r is* locally consistent *on S if for any $P_1, P_2 \in S$ with $r(P_1) = r(P_2)$, we have $P_1 \cup P_2 \in S$ and $r(P_1 \cup P_2) = r(P_1) = r(P_2)$.*

Definition 5 *For any natural number t, a voting rule r is t-*consistent *if there exists a partition $\{S_1,\ldots,S_t\}$ of all profiles such that for all $i \leq t$, r is locally consistent within S_i. A voting rule r is* finitely locally consistent *if it is t-consistent for some natural number t.*

We note that if a voting rule satisfies FLC, then it must also satisfy homogeneity, because FLC requires that for any two profiles in the same component with the same winner, the union of those two profiles is also in the component (and has the same winner)—so any integer multiple of a profile must be in the same component (and have the same winner).[2]

We emphasize that in this definition, a rule is defined for a fixed number m of alternatives, but for profiles of arbitrarily many voters. Later, we will show that some common rules are finitely locally consistent for *every* $m \in \mathbb{N}$; however, in those cases, t depends on m, which is allowed, as long as t is finite. We note that this finiteness condition is important: for *any* homogeneous voting rule, there exists a partition that has infinitely many elements such that the voting rule is locally consistent on each element of the partition. Here, each element has the form $\{kP : k \in \mathbb{N}\}$ for some P that is not composed of multiple copies of any other profile.

The *degree of consistency* of a voting rule r (for a particular m) is the smallest number of elements in a locally consistent partition of profiles. That is, the degree of consistency of r is t if r is t-consistent, and for any $t' < t$, r is not t'-consistent. (We note that the partition corresponding to this lowest t is not necessarily unique.) The degree of consistency can be seen as an approximation to consistency: the lower the degree of consistency of a voting rule, the more "consistent" it is, and 1-consistency is equivalent to the standard definition of consistency. We will be interested in the exact degree of consistency (rather than just whether it is finite or not), because, as we will show, this degree is related to the order of a GSR equivalent to the rule, which in turn is important for the summarization and elicitation problems that we mentioned in the introduction.

4 Finite local consistency characterizes generalized scoring rules

We now present our main result. Let $\mathcal{P}(k)$ be the number of *total preorders* over k elements, that is, the total number of ways to rank k elements, allowing for ties.

Theorem 1 *r is a generalized scoring rule if and only if r is anonymous and finitely locally consistent. Moreover, for any t-consistent voting rule r, there exists a GSR of order $(\frac{t(t-1)m(m-1)}{4})m!+1$ that is equivalent to r; conversely, for any GSR $GS(f,g)$ of order k, there exists a $\mathcal{P}(k)$-consistent voting rule r that is equivalent to $GS(f,g)$.*[3]

Proof of Theorem 1: We prove the "if" part by a geometrical representation of a voting rule that is anonymous and homogenous, similarly to [Young, 1975]. Let $L(\mathcal{C}) = \{l_1,\ldots,l_{m!}\}$ be the set of all linear orders over $\mathcal{C}$. Let r be a voting rule that satisfies anonymity and FLC. It follows that r is anonymous and homogenous, so that profiles can be represented as multisets of votes. Hence, there is a one-to-one correspondence between the set of all profiles and the set of all points in $\mathbb{N}^{m!}$: any profile $P = \sum_{x=1}^{m!} w_x l_x$, $w_x \in \mathbb{N}$ is associated with the point $\vec{p} = (w_1,\ldots,w_{m!})$, that is, $\vec{p} \in \mathbb{N}^{m!}$, and for any $j \leq m!$, the jth component of $\vec{p}$ is exactly the number of voters whose preferences are l_j in P. Therefore, r can also be seen as a mapping from $\mathbb{N}^{m!}$ to $\mathcal{C}$, defined as follows: for any $\vec{p} \in \mathbb{N}^{m!}$, $r(\vec{p}) = r(P)$, where P is the profile that $\vec{p}$ corresponds to. In the remainder of the proof, we will not distinguish between the point $\vec{p}$ and the profile P. Also, because r is homogenous, the domain of r can be extended to ${\mathbb{Q}_{\geq 0}}^{m!}$ (vectors of nonnegative rationales) in the following way. For any $\vec{p} \in {\mathbb{Q}_{\geq 0}}^{m!}$, let $h \in \mathbb{N}$ be such that $h\vec{p} \in \mathbb{N}^{m!}$; then, let $r(\vec{p}) = r(h\vec{p})$. (This is well defined because by homogeneity, the choice of h does not matter.)

Because r is t-consistent, there exists a partition $(S_1,\ldots,S_t)$ of $\mathbb{N}^{m!}$ such that r is locally consistent within each S_i. We note that $\vec{p} \in S_i$ implies $h\vec{p} \in S_i$ for each

[2] We thank Felix Brandt and Markus Brill for pointing out the connection between FLC and homogeneity to us.

[3] The $\mathcal{P}(k)$ bound can be improved if more information about the structure of the GSR is taken into account. For the sake of simplicity, we omit discussion of this in this paper.

$h \in \mathbb{N}$, because each S_i must be closed under the union of vectors that produce the same result, and we can take the union of h vectors $\vec{p}$. Now, for any $i \leq t$, we define $S_i^{\mathbb{Q}} = \{q\vec{p} : q \in \mathbb{Q}_{\geq 0}, \vec{p} \in S_i\}$. It follows that ${\mathbb{Q}_{\geq 0}}^{m!} = \bigcup_{i=1}^t S_i^{\mathbb{Q}}$, and for any $i_1 \neq i_2$, $S_{i_1}^{\mathbb{Q}} \cap S_{i_2}^{\mathbb{Q}} = \{0\}$. For any $i \leq t$, any $j \leq m$, we define $S_i^j = S_i^{\mathbb{Q}} \cap r^{-1}(c_j)$. That is, S_i^j is the set of points (equivalently, profiles) in $S_i^{\mathbb{Q}}$ whose winner is c_j. It follows that for any $\vec{p}_1, \vec{p}_2 \in S_i^j \cap \mathbb{N}^{m!}$, we have $\vec{p}_1 + \vec{p}_2 \in S_i^j$; for any $\vec{p} \in S_i^j$, any $q \in \mathbb{Q}_{\geq 0}$, we must have $q\vec{p} \in S_i^j$. For any $S \subseteq {\mathbb{R}_{\geq 0}}^{m!}$, we say that S is $\mathbb{Q}$-*convex* if for any $\lambda \in \mathbb{Q} \cap [0, 1]$, any $\vec{p}_1, \vec{p}_2 \in S$, we have $\lambda\vec{p}_1 + (1-\lambda)\vec{p}_2 \in S$. We say a $\mathbb{Q}$-convex set S is a $\mathbb{Q}$-*convex cone*, if for any $q \in \mathbb{Q}_{\geq 0}$, any $\vec{p} \in S$, we have $q\vec{p} \in S$. We make the following easy claim (proof omitted).

Claim 1 *For any $i \leq t$, any $j \leq m$, S_i^j is a $\mathbb{Q}$-convex cone.*

For any $S \subseteq {\mathbb{R}_{\geq 0}}^{m!}$, we let $conv(S)$ be the *convex hull* of S in ${\mathbb{R}_{\geq 0}}^{m!}$. That is, $conv(S) = \{\sum_{i=1}^h \alpha_i \vec{p}_i : h = 1, 2, \ldots, \sum_{i=1}^h \alpha_i = 1, (\forall i \leq h)\ \alpha_i > 0, \alpha_i \in \mathbb{R}, \vec{p}_i \in S\}$.

Lemma 1 (proved in [Young, 1975]) *$S \subseteq \mathbb{Q}^{m!}$ is $\mathbb{Q}$-convex if and only if $S = conv(S) \cap \mathbb{Q}^{m!}$.*

Let $d \in \mathbb{N}$, $S_1, S_2 \subseteq \mathbb{R}^d$, and for any $x \in \mathbb{R}$, let $\delta(x) = 1$ if $x > 0$, $\delta(x) = -1$ if $x < 0$, and $\delta(0) = 0$. We say that S_1 and S_2 are *separated* by a finite set of vectors $I = \{\vec{p}_1, \ldots, \vec{p}_o\}$, in which $\vec{p}_i \in \mathbb{R}^l$ for all $i \leq o$, if there exist two sets $O_1, O_2 \subseteq \{-1, 0, 1\}^I$ such that $O_1 \cap O_2 = \{0\}$, and for any $\vec{p} \in S_1$ ($\vec{p} \neq 0$), we have $\delta(\vec{p}, I) = (\delta(\vec{p} \cdot \vec{p}_1), \ldots, \delta(\vec{p} \cdot \vec{p}_o)) \in O_1$; for any $\vec{p} \in S_2$ ($\vec{p} \neq 0$), we have $(\delta(\vec{p} \cdot \vec{p}_1), \ldots, \delta(\vec{p} \cdot \vec{p}_o)) \in O_2$. In this case we also say that I separates S_1 from S_2 *via* O_1, O_2.

$S \subseteq \mathbb{R}^d$ is called an *affine space* if for any $\vec{p}_1, \vec{p}_2 \in S$, any $q_1, q_2 \in \mathbb{R}$, we have $q_1\vec{p}_1 + q_2\vec{p}_2 \in S$. For any $S' \subseteq \mathbb{R}^d$, we let $\mathrm{aff}(S')$ denote the *affine extension* of S' as follows: $\mathrm{aff}(S') = \{\sum_{i=1}^h \alpha_i \vec{p}_i : h = 1, 2, \ldots, (\forall i \leq h)\ \alpha_i \in \mathbb{R}, \vec{p}_i \in S'\}$. That is, $\mathrm{aff}(S')$ is the smallest affine space in $\mathbb{R}^d$ that contains S'. We let $relint(conv(S))$ denote the *relative interior* of $conv(S)$, defined as follows. $relint(conv(S))$ is the set of all vectors $\vec{p} \in \mathbb{R}^d$ such that there exists $\epsilon > 0$ such that $B(\vec{p}, \epsilon) \cap \mathrm{aff}(S) \subseteq conv(S)$, where $B(\vec{p}, \epsilon)$ is the ball centered on $\vec{p}$ with radius ϵ.

Lemma 2 *Let $S \subseteq \mathbb{R}^{m!}$ be an affine space, and let $S_1, S_2 \subseteq S \cap {\mathbb{Q}_{\geq 0}}^{m!}$ be two $\mathbb{Q}$-convex cones such that $S_1 \neq S_2$, $S_1 \cap S_2 = \{0\}$. There exists a finite set of vectors $I \subseteq \mathbb{R}^{m!}$ that separates S_1 from S_2, and $|I| \leq dim(S)$.*

Proof. We prove the claim by induction on $dim(S)$. When $dim(S) = 1$, it must be the case that one of S_1 and S_2 is $\{0\}$, and the other has an element $\vec{p}' \neq 0$. Without loss of generality, we let $S_1 = \{0\}$, $S_2 \neq \{0\}$. In this case, we let $I = \{\vec{p}'\}$, $O_1 = \{0\}$, and $O_2 = \{0, 1\}$.

Suppose Lemma 2 holds for $dim(S) \leq d$. Without loss of generality, we assume $dim(\mathrm{aff}(S_1)) \geq dim(\mathrm{aff}(S_2))$. When $dim(S) = d + 1$, there are two cases.

Case 1: $dim(\mathrm{aff}(S_1)) = dim(\mathrm{aff}(S_2)) = d + 1$. In this case $S = \mathrm{aff}(S_1) = \mathrm{aff}(S_2)$. First we prove that $relint(conv(S_1)) \cap relint(conv(S_2)) = \emptyset$. If not, suppose $\vec{p} \in relint(conv(S_1)) \cap relint(conv(S_2))$. Let $\vec{p} = \sum_{j=1}^h \alpha_j \vec{p}_j$, where $\sum_{j=1}^h \alpha_j = 1$, for all $j \leq h$, $\vec{p}_j \in S_1$ and $\alpha_j \geq 0$, and $B(\vec{p}, \epsilon) \cap S \subseteq conv(S_1)$, $B(\vec{p}, \epsilon) \cap S \subseteq conv(S_2)$. There exist $\beta_j \in \mathbb{Q}_{\geq 0}$ ($j \leq h$) such that $\vec{p}^* = \sum_{j=1}^h \beta_j \vec{p}_j \neq 0$, and the distance between $\vec{p}^*$ and $\vec{p}$ is less than ϵ (by setting the β_j sufficiently close to the α_j). We note that S_1 is $\mathbb{Q}$-convex, which means that $\vec{p}^* \in S_1$. It follows that $\vec{p}^* \in conv(S_2)$, because $\vec{p}^* \in B(\vec{p}, \epsilon) \cap S$. From Lemma 1 we have that $S_2 = conv(S_2) \cap {\mathbb{Q}_{\geq 0}}^{m!}$. Therefore, $\vec{p}^* \in conv(S_2) \cap {\mathbb{Q}_{\geq 0}}^{m!} = S_2$. This contradicts the assumption that $S_1 \cap S_2 = \{0\}$.

Because $relint(conv(S_1)) \cap relint(conv(S_2)) = \emptyset$, we apply the separating hyperplane theorem: there exists a hyperplane $H_{\vec{p}^*}$ characterized by $\vec{p}^* \in \mathbb{R}^{m!}$, such that for any $\vec{p}_1 \in S_1$, $\vec{p}_1 \cdot \vec{p}^* \leq 0$; for any $\vec{p}_2 \in S_2$, $\vec{p}_2 \cdot \vec{p}^* \geq 0$; and at least one of S_1 and S_2 is not contained in $H_{\vec{p}^*}$. We let $S' = S \cap H_{\vec{p}^*}$, and $S_1' = S_1 \cap S'$, $S_2' = S_2 \cap S'$. $H_{\vec{p}^*}$ does not contain S, so it follows that $dim(S') < dim(S) = d+1$. Applying Lemma 2 on S', S_1', S_2' (using the induction assumption), there exists a set of vectors I' that separates S_1' from S_2' via O_1', O_2', $|I'| \leq d$. Let $I = \{\vec{p}^*\} \cup I'$ and $O_1 = \{\vec{a} \in \{-1, 0, 1\}^I : \vec{a}|_{\{\vec{p}^*\}} = -1 \vee (\vec{a}|_{\{\vec{p}^*\}} = 0 \wedge \vec{a}|_{I'} \in O_1')\}$ (here, for $J \subseteq I$, let $\vec{a}|_J$ be the components of $\vec{a}$ corresponding to the vectors in J). This works because for any $\vec{p} \in S_1$, either $\vec{p}$ is in the open halfspace $\{\vec{p}' : \vec{p}' \cdot \vec{p}^* < 0\}$, or $\vec{p}$ is in $S_1 \cap H_{\vec{p}^*}$. Similarly, let $O_2 = \{\vec{a} \in \{-1, 0, 1\}^I : \vec{a}|_{\{p^*\}} = 1 \vee (\vec{a}|_{\{p^*\}} = 0 \wedge \vec{a}|_{I'} \in O_2')\}$. It follows that I separates S_1 from S_2 via O_1, O_2, and $|I| = |I'| + 1 \leq d + 1$.

Case 2: $dim(\mathrm{aff}(S_2)) < d + 1$. If $\mathrm{aff}(S_1) = \mathrm{aff}(S_2)$, then let $S' = \mathrm{aff}(S_1)$, $|S'| < d + 1$. Applying Lemma 2 on S', S_1, S_2 (by the induction assumption), we can conclude that there exists $I' \subseteq {\mathbb{Q}_{\geq 0}}^{m!}$ that separates S_1 from S_2, and $|I'| \leq d < d + 1$. If $\mathrm{aff}(S_1) \neq \mathrm{aff}(S_2)$, then there exists a hyperplane $H_{\vec{p}^*}$ (orthogonal to $\vec{p}^*$) such that $0 \in H_{\vec{p}^*}$, $S_2 \subseteq H_{\vec{p}^*}$, and $S_1 \not\subseteq H_{\vec{p}^*}$ (because the intersection of all hyperplanes that contains S_2 is S_2). Let $S' = \mathrm{aff}(S_2)$, and $S_1' = S_1 \cap S'$. S' is an affine space whose dimension is $dim(\mathrm{aff}(S_2)) < d + 1$. For any $\vec{p}_1, \vec{p}_2 \in S_1'$, any $\lambda \in \mathbb{Q}_{\geq 0}$, we have that $\lambda\vec{p}_1 + (1-\lambda)\vec{p}_2 \in S_1$ (because S_1 is $\mathbb{Q}$-convex), and $\lambda\vec{p}_1 + (1-\lambda)\vec{p}_2 \in S'$ (because S' is an affine space); hence, $\lambda\vec{p}_1 + (1-\lambda)\vec{p}_2 \in S_1'$. Hence S_1' is a $\mathbb{Q}$-convex cone.

By applying Lemma 2 on S', S_1', S_2 (using the induction assumption), there exists $I' \subset {\mathbb{Q}_{\geq 0}}^{m!}$ ($|I'| \leq d$) that separates S_1' from S_2 via O_1', O_2'. We let $I = I' \cup \{\vec{p}^*\}$; $O_1 = \{\vec{a} \in \{-1, 0, 1\}^I : \vec{a}|_{\{\vec{p}^*\}} \neq 0 \vee (\vec{a}|_{\{\vec{p}^*\}} = 0 \wedge \vec{a}|_{I'} \in O_1')\}$. This works because for any $\vec{p} \in S_1$, either $\vec{p} \cdot \vec{p}^* \neq 0$ (meaning that $\vec{p}$ is not in S'), or $\vec{p} \cdot \vec{p}^* = 0$, and $\delta(\vec{p}, I') \in O_1'$ (meaning that $\vec{p}$ is in $S_1 \cap S'$). Similarly we define $O_2 = \{\vec{a} \in \{-1, 0, 1\}^I : \vec{a}|_{\{\vec{p}^*\}} = 0 \wedge a|_{I'} \in O_2'\}$. It follows that I separates S_1 from S_2, and $|I| = |I'| + 1 \leq d + 1$. This completes the proof of Lemma 2. □

For any $i_1, i_2 \leq t$, $j_1, j_2 \leq m$, where either $i_1 \neq i_2$ or $j_1 \neq j_2$, $S_{i_1}^{j_1} \cap S_{i_2}^{j_2} = \{0\}$. (We recall that S_i^j is the set of points in $S_i^{\mathbb{Q}}$ whose winner is c_j.) From Lemma 2, there exists a finite set $I_{i_1 j_1, i_2 j_2}$ of vectors that separates $S_{i_1}^{j_1}$ from $S_{i_2}^{j_2}$ via $O^1_{i_1 j_1, i_2 j_2}, O^2_{i_1 j_1, i_2 j_2}$, where $|I_{i_1 j_1, i_2 j_2}| \leq m!$. Now we can define a corresponding generalized scoring rule, as follows.

• $k = |\bigcup_{(i_1,j_1)\neq(i_2,j_2)} I_{i_1j_1,i_2j_2}| + 1$, and the components are indexed by vectors in some $I_{i_1j_1,i_2j_2}$, and a 0 component (which is always 0). Because $|I_{i_1j_1,i_2j_2}| \leq m!$, we have $k \leq (\frac{t(t-1)m(m-1)}{4})m! + 1$.

• For any $(i_1,j_1) \neq (i_2,j_2)$, any $\vec{p} = (p_1,\ldots,p_{m!}) \in I_{i_1j_1,i_2j_2}$, any $b \leq m!$, the $\vec{p}$ component of the generalized score vector given vote (ranking) l_b is $f(l_b) = p_b$. We note that for any profile $\vec{p} = (w_1,\ldots,w_{m!})$, any $\vec{p}^* = (p^*_1,\ldots,p^*_{m!}) \in I_{i_1j_1,i_2j_2}$, the $\vec{p}^*$ component of $f(\vec{p})$ is $\sum_{x=1}^{m!} w_x p^*_x = \vec{p}\cdot\vec{p}^*$.

• For any $\vec{a} \in {\mathbb{Q}_{\geq 0}}^k$ with $\vec{a} \neq 0$, $g(\vec{a}) = c_j$ if and only if there exists $i \leq t$ such that for any $i' \leq t$, $j' \leq m$, there exists $o \in O^1_{ij,i'j'}$ such that for any $\vec{p}^* \in I_{ij,i'j'}$, the following three conditions hold: (1) $\vec{a}|_{\vec{p}^*}$ is strictly larger than 0 (the value of the 0 component), if and only if $o|_{\vec{p}^*} = 1$; (2) $\vec{a}|_{\vec{p}^*}$ is equal to 0, if and only if $o|_{\vec{p}^*} = 0$; and (3) $\vec{a}|_{\vec{p}^*}$ is strictly smaller than 0, if and only if $o|_{\vec{p}^*} = -1$. That is, $g(\vec{a}) = c_j$ if and only if there exists $i \leq t$ such that for any i', j', we always have $\vec{a} \notin S^{j'}_{i'}$ by using the set of separation vectors $I_{ij,i'j'}$. (That g is well defined will follow from the following argument.)

Next, we prove that $GS(f,g) = r$. For any profile $\vec{p} \in {\mathbb{Q}_{\geq 0}}^{m!}$, suppose $\vec{p} \in S^j_i$. For any $(i,j) \neq (i',j')$, since $\vec{p} \in S^j_i$, by using the separation vectors $I_{ij,i'j'}$ and $O^1_{ij,i'j'}, O^2_{ij,i'j'}$, $\vec{p}$ should be classified as "not in $S^{j'}_{i'}$". That is, there exists $o \in O^1_{ij,i'j'}$ such that for any $\vec{p}^* \in I_{ij,i'j'}$, $o|_{\vec{p}^*} = \delta(\vec{p}\cdot\vec{p}^*)$; and for any $o' \in O^2_{ij,i'j'}$, there exists $\vec{p}^* \in I_{ij,i'j'}$ such that $o'|_{\vec{p}^*} \neq \delta(\vec{p}\cdot\vec{p}^*)$. It follows that $GS(f,g)(\vec{p}) = c_j$.

The "only if" part is straightforward. For any total preorder $\mathcal{O}$ over $\{1,\ldots,k\}$, we let $S_{\mathcal{O}} = \{\vec{p} \in {\mathbb{Q}_{\geq 0}}^{m!} : f(\vec{p}) \sim \mathcal{O}\}$. $\{S_{\mathcal{O}}\}$ is a finitely locally consistent partition for the rule, of size $\mathcal{P}(k)$. □

We are not aware of any closed-form formula for $\mathcal{P}(k)$, though there exist recursive formulas. We now give a simple upper bound on $\mathcal{P}(k)$. Any total preorder V can be represented by a strict order $(c_{i_1} \succ c_{i_2} \succ \ldots \succ c_{i_m})$ and a string $\vec{s} = (s_1,\ldots,s_{m-1}) \in \{0,1\}^{m-1}$, as follows: if $s_l = 0$ then $c_{i_l} \succ_V c_{i_{l+1}}$, and if $s_l = 1$ then $c_{i_l} \approx_V c_{i_{l+1}}$. This implies $\mathcal{P}(k) \leq k!2^{k-1}$.

5 The degree of consistency of common rules

In this section, we prove upper and lower bounds for the degree of consistency of some common voting rules (which Theorem 1 relates to the order of a GSR representation of that rule). Positional scoring rules have degree of consistency 1, but other rules do not. To prove a lower bound L on the degree of consistency of a voting rule, we construct a set of L profiles $P_1,\ldots,P_L$ that all have the same winner, but, for any $i \neq j$, the winner for $P_1 \cup P_2$ is another alternative (hence, none of these profiles can belong to the same element of a partition). To prove an upper bound U, for any alternative c_j, we partition the set of all profiles whose winner is c_j into U parts, denoted by $S^j_1,\ldots,S^j_U$, and prove that the rule is locally consistent within each part. Then, we let the partition of all profiles be $\{S_1,\ldots,S_U\}$, where $S_i = \bigcup_{j=1}^m S^j_i$. It is easy to see that the rule must be locally consistent within each S_i. Due to space constraint, we only include the proof for Bucklin (which is the shortest).

Theorem 2 *The degree of consistency of Bucklin is at least $\lfloor\frac{m}{2}\rfloor$, and at most $\lfloor\frac{m}{2}\rfloor + 1$.*

Proof. First we prove the lower bound. For any $k \leq \lfloor\frac{m}{2}\rfloor$, define $P_k = (V_1,\ldots,V_{8(m-2)})$ as follows.

• For $1 \leq j \leq 4(m-2)+1$, let c_1 be in the kth position of V_j, and let c_2 be in the $(k+1)$th position of V_j.

• For $4(m-2)+2 \leq j \leq 8(m-2)$, let c_1 be in the mth position of V_j, and let c_2 be in the kth position of V_j.

• Let $c_3,\ldots,c_m$ be ranked in a cyclic way in $V_1,\ldots,V_{8(m-2)}$—for example, $c_3 \succ c_4 \succ \ldots \succ c_m$ in V_1; $c_m \succ c_3 \succ c_4 \succ \ldots \succ c_{m-1}$ in V_2; and so on.

In each P_k, c_1 is ranked among the top k positions $4(m-2)+1 > 4(m-2)$ times; c_2 is ranked among the top k positions $4(m-2)-1 < 4(m-2)$ times. For any $3 \leq i \leq m$, c_i is ranked among the top k positions $8(k-1)$ times. Because $k \leq \frac{m}{2}$, $8(k-1) \leq 4(m-2)$. Hence, $Bucklin(P_k) = c_1$. Now, for any $k_1 < k_2 \leq \lfloor\frac{m}{2}\rfloor$, in $P_{k_1} \cup P_{k_2}$, c_1 is ranked among the top $k_2 - 1$ positions $4(m-2)+1 < 8(m-2)$ times; c_1 is ranked among the top k_2 positions $8(m-2)+2 > 8(m-2)$ times; and c_2 is ranked among the top k_2 positions $12(m-2)-1 > 8(m-2)+2$ times. Hence, in $P_{k_1} \cup P_{k_2}$, c_2 performs better than c_1, and therefore, $Bucklin(P_{k_1} \cup P_{k_2}) \neq c_1$. It follows that the degree of consistency of Bucklin is at least $\lfloor\frac{m}{2}\rfloor$.

Next, we prove the upper bound. Without loss of generality, we focus on the alternative c_m. Let P be a profile such that $Bucklin(P) = c_m$. For any k, let S^m_k be the set of profiles P such that $Bucklin(P) = c_m$, and the Bucklin score of c_m in P is k. For any $P, P' \in S^m_k$, any $i \leq m$, c_i is ranked among the top $k-1$ positions at most $\frac{|P|}{2}$ (resp., $\frac{|P'|}{2}$) times in P (resp., P'). Therefore, c_i is ranked among the top $k-1$ positions at most $\frac{|P|+|P'|}{2}$ times in $P \cup P'$. By similar reasoning, c_m is ranked among the top k positions strictly more than $\frac{|P|+|P'|}{2}$ times in $P \cup P'$. It is possible that some other c_i is also ranked among the top k positions strictly more than $\frac{|P|+|P'|}{2}$ times in $P \cup P'$; however, c_i must be ranked among the top k at most as many times as c_m in both P and P', and hence also in $P \cup P'$. If c_m is ranked among the top k more often than c_i in $P \cup P'$, then c_m wins; if they are ranked among the top k equally often, then c_m still wins, by consistent tiebreaking. It follows that Bucklin is locally consistent within each S^m_k. Moreover, we note that S^m_k is empty for $k > \lfloor\frac{m}{2}\rfloor + 1$. So, we only need the S^m_k with $k \leq \lfloor\frac{m}{2}\rfloor + 1$. By similar analysis for the other alternatives, it follows that the degree of consistency of Bucklin is at most $\lfloor\frac{m}{2}\rfloor + 1$. □

Theorem 3 *The degree of consistency of STV is $(m-1)!$.*

Theorem 4 *The degree of consistency of maximin is at least $(m-3)^{m-1}$, and at most $(m-1)^{m-1}$.*

Because $\lim_{m\to\infty} \frac{(m-3)^{m-1}}{(m-1)^{m-1}} = \frac{1}{e^2}$, the bounds on the degree of consistency of maximin are asymptotically tight up to a multiplicative constant.

Theorem 5 *The degree of consistency of Copeland is at least* $(2^{-m'}\binom{2m'}{m'})^m$*, where* $m' = \lfloor \frac{m-1}{2} \rfloor$*, and at most* $3^{\frac{m(m-1)}{2}}$*.*

Theorem 6 *The degree of consistency of ranked pairs is at least* $(m-2)!$*, and at most* $(m(m-1))!2^{m^2-m-1}$*.*

6 Conclusion

The frequency of manipulability of a rule is an important problem in computational social choice. Recently, this frequency has been characterized—except for a knife-edge case—for the class of generalized scoring rules [Xia and Conitzer, 2008]. While it was previously observed (based on examples) that this class seems very general, no formal argument for this generality has been given. In this paper we gave such an argument, by characterizing generalized scoring rules as the class of rules that satisfy finite local consistency (in addition to anonymity), a very natural weakening of the standard notion of consistency. This definition also leads to a quantitative degree of consistency of a rule, which turns out to be closely related to the order (dimension of the score vector) of a generalized scoring rule. Our axiomatic characterization also shows that Dodgson's rule is not a generalized scoring rule, because it does not satisfy homogeneity [Brandt, 2009]. We provided lower and upper bounds on the degree of consistency for some common voting rules, summarized in Table 1. These bounds imply corresponding bounds on the order of a generalized scoring rule representation of these rules, per Theorem 1.

Rule	Lower Bound	Upper Bound
Pos. scoring	1	1
Bucklin	$\lfloor \frac{m}{2} \rfloor$	$\lfloor \frac{m}{2} \rfloor + 1$
STV	(m-1)!	(m-1)!
Maximin	$(m-3)^{m-1}$	$(m-1)^{m-1}$
Copeland	$(2^{-\lfloor \frac{m-1}{2} \rfloor}\binom{2\lfloor \frac{m-1}{2} \rfloor}{\lfloor \frac{m-1}{2} \rfloor})^m$	$3^{\frac{m(m-1)}{2}}$
Ranked pairs	$(m-2)!$	$(m(m-1))!2^{m^2-m-1}$

Table 1: Lower and upper bounds on the degree of consistency of some common voting rules.

One direction for future research is to tighten the above bounds, and derive bounds for other rules. A related question is whether the bounds from Theorem 1 can be improved. Another direction is to study alternative uses of generalized scoring rules—we briefly mentioned the possibility of using them for elicitation and summarization purposes—as well as the properties that they possess.

Acknowledgments

We thank Felix Brandt, Markus Brill, and anonymous reviewers for helpful comments and suggestions. This work is supported by NSF under award number IIS-0812113. Lirong Xia is supported by a James B. Duke Fellowship, and Vincent Conitzer is supported by an Alfred P. Sloan Research Fellowship.

References

[Bartholdi and Orlin, 1991] John Bartholdi, III and James Orlin. Single transferable vote resists strategic voting. *Social Choice and Welfare*, 8(4):341–354, 1991.

[Bartholdi *et al.*, 1989] John Bartholdi, III, Craig Tovey, and Michael Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.

[Brandt, 2009] Felix Brandt. Some remarks on Dodgson's voting rule. *Mathematical Logic Quarterly*, 55(4), 2009.

[Chevaleyre *et al.*, 2009] Yann Chevaleyre, Jérôme Lang, Nicolas Maudet, and Guillaume Ravilly-Abadie. Compiling the votes of a subelectorate. In *Proc. IJCAI-09*, 2009.

[Conitzer and Sandholm, 2003] Vincent Conitzer and Tuomas Sandholm. Universal voting protocol tweaks to make manipulation hard. In *Proc. IJCAI-03*, pages 781–788, 2003.

[Conitzer and Sandholm, 2006] Vincent Conitzer and Tuomas Sandholm. Nonexistence of voting rules that are usually hard to manipulate. In *Proc. AAAI-06*, pages 627–634, 2006.

[Conitzer *et al.*, 2007] Vincent Conitzer, Tuomas Sandholm, and Jérôme Lang. When are elections with few candidates hard to manipulate? *Journal of the ACM*, 54(3):Article 14, 1–33, 2007.

[Elkind and Lipmaa, 2005] Edith Elkind and Helger Lipmaa. Hybrid voting protocols and hardness of manipulation. In *Proc. ISAAC-05*, 2005.

[Faliszewski *et al.*, 2008] Piotr Faliszewski, Edith Hemaspaandra, and Henning Schnoor. Copeland voting: ties matter. In *Proc. AAMAS-08*, pages 983–990, 2008.

[Gibbard, 1973] Allan Gibbard. Manipulation of voting schemes: a general result. *Econometrica*, 41:587–602, 1973.

[Hemaspaandra and Hemaspaandra, 2007] Edith Hemaspaandra and Lane A. Hemaspaandra. Dichotomy for voting systems. *Journal of Computer and System Sciences*, 73(1):73–83, 2007.

[May, 1952] K. May. A set of independent necessary and sufficient conditions for simple majority decision. *Econometrica*, 20:680–684, 1952.

[Procaccia and Rosenschein, 2007a] Ariel D. Procaccia and Jeffrey S. Rosenschein. Average-case tractability of manipulation in voting via the fraction of manipulators. In *Proc. AAMAS-07*, 2007.

[Procaccia and Rosenschein, 2007b] Ariel D. Procaccia and Jeffrey S. Rosenschein. Junta distributions and the average-case complexity of manipulating elections. *Journal of Artificial Intelligence Research (JAIR)*, 28:157–181, 2007.

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

[Walsh, 2009] Toby Walsh. Where are the really hard manipulation problems? The manipulation phase transition. In *Proc. IJCAI-09*, 2009.

[Xia and Conitzer, 2008] Lirong Xia and Vincent Conitzer. Generalized scoring rules and the frequency of coalitional manipulability. In *Proc. EC-08*, pages 109–118, 2008.

[Xia *et al.*, 2009] Lirong Xia, Michael Zuckerman, Ariel Procaccia, Vincent Conitzer, and Jeffrey Rosenschein. Complexity of unweighted coalitional manipulation under some common voting rules. In *Proc. of IJCAI-09*, 2009.

[Young, 1975] H. Peyton Young. Social choice scoring functions. *SIAM Journal of Applied Mathematics*, 28(4):824–838, 1975.

[Zuckerman *et al.*, 2009] Michael Zuckerman, Ariel D. Procaccia, and Jeffrey S. Rosenschein. Algorithms for the coalitional manipulation problem. *Artif. Intelli.*, 173(2):392–412, 2009.