

Generalized Decision Scoring Rules: Statistical, Computational, and Axiomatic Properties

Lirong Xia, Rensselaer Polytechnic Institute

We pursue a *design by social choice, evaluation by statistics and computer science* paradigm to build a principled framework for discovering new social choice mechanisms with desirable statistical, computational, and social choice axiomatic properties.

Our new framework is called *generalized decision scoring rules (GDSRs)*, which naturally extend *generalized scoring rules* [Xia and Conitzer 2008] to arbitrary preference space and decision space, including sets of alternatives with fixed or unfixed size, rankings, and sets of rankings. We show that GDSRs cover a wide range of existing mechanisms including MLEs, Chamberlin and Courant rule, and resolute, irresolute, and preference function versions of many commonly studied voting rules. We provide a characterization of statistical consistency for any GDSR w.r.t. any statistical model and asymptotically tight bounds on the convergence rate. We investigate the complexity of winner determination and a wide range of strategic behavior called *vote operations* for all GDSRs, and prove a general phase transition theorem on the minimum number of vote operations for the strategic entity to succeed. We also characterize GDSRs by two social choice normative properties: *anonymity* and *finite local consistency*.

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

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Social choice; Generalized decision scoring rules; Statistical consistency; Vote operations; Finite local consistency

1. INTRODUCTION

Social choice theory studies how to aggregate agents' opinions or preferences to make a (joint) decision. Traditional social choice theory concerns how to reach a consensus that is evaluated by agents' *subjective* satisfaction of the decision. Ideally, we would like to respect agents' opinions and preferences and make a decision to satisfy all agents, which is often impossible due to their conflicting preferences. A typical example is political elections.

In many multi-agent scenarios, the goal of social choice is to aggregate agents' preferences to reveal the ground truth or make an *objectively* optimal decision for the decision maker, who may not be the group of agents. For example, online retailers (decision makers) aggregate reviewers' ratings of an item to provide an estimate to the true popularity of the item, to make decisions such as whether or not to continue selling this item. In such settings, instead of making a joint decision (e.g. the aggregated score) to

This work is supported by the National Science Foundation under grant IIS-1453542. The author thanks anonymous reviewers of EC-15, COMSOC-12, 14, AAMAS-14, and AIJ for helpful comments on (parts of) this paper.

Authors addresses: L. Xia, Computer Science Department, Rensselaer Polytechnic Institute; email: xial@cs.rpi.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC'15, June 15–19, 2015, Portland, OR, USA.

ACM 978-1-4503-3410-5/15/06 ...\$15.00.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

<http://dx.doi.org/10.1145/2764468.2764518>

satisfy the agents (e.g. the reviewers), we want to make a decision with high *objective* quality evaluated w.r.t. the ground truth. Still, we often need to respect agents’ preferences and use a good mechanism in the social choice sense, for scenarios with strong societal context, e.g. a group of friends vote to decide a restaurant for dinner. Even for scenarios with less societal context, e.g. meta-search engines [Dwork et al. 2001], recommender systems [Ghosh et al. 1999], crowdsourcing [Mao et al. 2013], semantic webs [Porello and Endriss 2013], and peer grading for MOOC [Raman and Joachims 2014], some social choice axiomatic properties are still desired.

These new applications of social choice in multi-agent scenarios naturally lead to the following challenging question: “*How can we design new social choice mechanisms with desirable statistical, computational, and social choice axiomatic properties?*”

1.1. Our Contributions

Recently, Azari Soufiani et al. [2014] follow a “design by statistic, evaluation by social choice and computer science” paradigm and propose a statistical decision-theoretic framework to discover new mechanisms as decision rules, then evaluate their computational and axiomatic properties. In this paper, we answer the above question by pursuing a *design by social choice, evaluation by statistics and computer science* paradigm. That is, we propose a novel framework where all mechanisms automatically satisfy some desired social choice axiomatic properties, and examine the statistical and computational aspects of them. More precisely, we propose *generalized decision scoring rules (GDSRs)* that naturally extend *generalized scoring rules* [Xia and Conitzer 2008] to arbitrary preference space and decision space, including sets of alternatives with fixed or unfixed size, rankings, and sets of rankings. We show that GDSRs cover a wide range of existing social choice mechanisms including MLEs, Chamberlin and Courant rule, and resolute, irresolute, and ranking versions of all positional scoring rules, STV, ranked pairs, etc. We obtain the following general results.

Statistical properties: We focus on *statistical consistency*, which is a desirable statistical property that requires the mechanism to reveal the ground truth with probability 1 as the data size goes to infinity. We obtain the following characterization of consistent GDSRs w.r.t. any given statistical model.

Theorem 1 (informal). *A GDSR is statistically consistent w.r.t. a statistical model if and only if for any ground truth d , the GDSR is robust against smaller perturbations around the profile corresponding to the probability distribution associated with d .*

We also provide asymptotically tight bounds on the convergence rate for all GDSRs.

Computational properties: We prove that winner determination and computing *vote operations* for any GDSR are fixed-parameter tractable w.r.t. the number of alternatives. Vote operations are a new and general class of agents’ strategic behavior covering coalitional manipulation, bribery, control by adding/deleting votes, etc. We also prove the following theorem on the phase transition of the number of necessary vote operations for the strategic entity to achieve her goal.

Theorem 7 (informal). *For any integer GDSR and any vote operations, the number of operations needed is one of the following four cases with probability that goes to 1 as the number of i.i.d. votes goes to infinity: (1) 0, (2) $\Theta(\sqrt{n})$, (3) $\Theta(n)$, and (4) ∞ .*

Axiomatic properties: We prove the following axiomatic characterization for GDSRs.

Theorem 8 (informal). *A social choice mechanism is a GDSR if and only if it satisfies anonymity and finite local consistency.*

Main conceptual contributions. The framework of GDSRs and the corresponding new paradigm towards the design and analysis of new social choice mechanisms are the main conceptual contributions. In addition, the notion of vote operations provides a novel framework for unifying various types of strategic behavior as well as obtaining new results on worst-case complexity and average-case complexity (via the phase transition theorem), as we will show in Corollary 2.

Main technical contributions. We feel that there are two main technical contributions, especially compared to existing results for generalized scoring rules. (1) The proof for the characterization of statistical consistency (Theorem 1) involves a novel application of a multivariate central limit theorem. (2) The proof for the phase transition on minimum number of vote operations (Theorem 7) involves a novel application of sensitivity analysis for ILPs.

1.2. Related Work and Discussions

The study of statistical properties of social choice mechanisms can be dated back to the *Condorcet Jury Theorem* in the 18th century [Condorcet 1785], which states that when there are two decisions, the majority rule is statistically consistent w.r.t. a simple statistical model. Most recent work on statistical approaches towards social choice focused on computation and characterization of the *maximum likelihood estimators (MLEs)* of various statistical models [Conitzer and Sandholm 2005; Braverman and Mossel 2008; Conitzer et al. 2009; Elkind et al. 2010; Xia and Conitzer 2011; Procaccia et al. 2012; Lu and Boutilier 2011; Azari Soufiani et al. 2012, 2013; Caragiannis et al. 2013]. While most previous research are case-by-case analysis of existing social choice mechanisms, GDSRs provides a novel and general framework to obtain *new* social choice mechanisms, and also to obtain general results on statistical properties, especially statistical consistency (Theorem 1).

The study of computational aspects of social choice mechanisms, especially winner determination and agents' strategic behavior, were initiated by the seminal work of Bartholdi et al. [1989a,b, 1992]. There is a large literature on these problems in the computational social choice community, for example [Hemaspaandra et al. 1997; Conitzer et al. 2007; Procaccia and Rosenschein 2007; Friedgut et al. 2008; Xia and Conitzer 2008; Faliszewski et al. 2009; Xia et al. 2009; Betzler et al. 2011; Obratsova and Elkind 2011; Mossel and Racz 2012; Chierichetti and Kleinberg 2012; Davies et al. 2014], among many others. See recent surveys by Faliszewski and Procaccia [2010], Faliszewski et al. [2010], and Rothe and Schend [2013]. While most previous results are case-by-case, our general theorems in Section 5 work for all GDSRs and all vote operations. The phase transition theorem (Theorem 7) suggests that the minimum number of vote operations is likely to be easy to compute in practice. Thus the theorem is negative if we want to use high complexity to protect elections, but it is positive in some other cases where fast computation is favored, for example to compute the margin of victory for post-election audits [Cary 2011; Magrino et al. 2011; Xia 2012]. This significantly extends the phase transition phenomenon for manipulation [Xia and Conitzer 2008; Mossel et al. 2013] to other types of strategic behavior and other mechanisms (GDSRs).

The study of axiomatic characterizations of social choice mechanisms has been popularized since Arrow's impossibility theorem [Arrow 1950]. In addition to voting rules, axiomatic characterizations have been extended to other application domains, for example recommender systems [Pennock et al. 2000] and ranking systems [Altman and Tennenholtz 2010]. The axiomatic characterization of GDSRs (Theorem 8) is similar to the axiomatic characterization of generalized scoring rules [Xia and Conitzer 2009], and can be used to show that some mechanisms are not GDSRs (e.g. Example 7). To

the best of our knowledge, our results on statistical consistency and vote operations proved are not previously known even for generalized scoring rules.

After all, why are GDSRs useful? We feel that there are at least two reasons. First, because GDSRs and vote operations are very general, the results obtained in this paper can be applied to existing social choice mechanisms to understand the important questions on statistical, computational, and axiomatic properties. Specifically, the applications of statistical consistency (Theorem 1) in Section 4 are new, and most applications of the phase transition theorem (Theorem 7) are new. Second, more importantly, the generality and structure of GDSRs allow us to design new social choice mechanisms along the new paradigm of design by social choice, evaluation by statistics and computation: we can easily obtain new mechanisms that satisfy anonymity and finite local consistency, and then use the general theorems in this paper to evaluate and compare them w.r.t. statistical and computational properties. The design may also be automated using machine learning, as briefly discussed by Xia [2013] for generalized scoring rules.

2. PRELIMINARIES

Let $\mathcal{C} = \{c_1, \dots, c_m\}$ denote a set of m alternatives and let \mathcal{S} denote the set of all possible preferences that the agents can report. In this paper, we assume that $\mathcal{S} = \mathcal{L}(\mathcal{C})$ for illustration. All results apply to choices of \mathcal{S} with finite elements. Each agent casts a *vote* in \mathcal{S} to represent her preferences. The vector of all agents' votes P is called a *profile*. Let $\mathcal{S}^* = \mathcal{S} \cup \mathcal{S}^2 \cup \dots$ denote the set of all profiles. In the literature, \mathcal{S} is often the set of all linear orders over \mathcal{C} , denoted by $\mathcal{L}(\mathcal{C})$, but can also be other sets such as subsets of alternatives as in approval voting.

Let \mathcal{D} denote the set of (*joint*) decisions. A *mechanism* (or *voting rule*) r is a mapping that assigns to each profile a single decision in \mathcal{D} . Common choices of \mathcal{D} are: (1) \mathcal{C} , where mechanisms are often called *resolute voting rules*; (2) $(2^{\mathcal{C}} \setminus \{\emptyset\})$, where mechanisms are often called *irresolute voting rules*; and (3) $\mathcal{L}(\mathcal{C})$, where mechanisms are often called *preference functions* (a.k.a. *social welfare functions*).

Many commonly-studied voting rules have resolute, irresolute, and preference function versions. For example, an irresolute *positional scoring rule* is characterized by a scoring vector $\vec{s} = (s_1, \dots, s_m)$ with $s_1 \geq s_2 \geq \dots \geq s_m$. For any alternative c and any linear order $V \in \mathcal{L}(\mathcal{C})$, we let $\bar{s}(V, c) = s_j$, where j is the rank of c in V . Given a profile P , the irresolute version of positional scoring rule chooses all alternatives c that maximize $\sum_{V \in P} \bar{s}(V, c)$, where P is viewed as a multi-set of votes. The resolute version chooses a single alternative by further applying a tie-breaking mechanism, and the preference function version ranks the alternatives w.r.t. their scores, and uses a tie-breaking mechanism when necessary.

As another example, the *single transferable vote (STV)* rule is naturally defined as a preference function that outputs a ranking in the following $m-1$ steps: in each step, the alternative ranked in the top positions least often is eliminated from the profile;¹ the outcome ranking is the inverse of the elimination order. The resolute version of STV simply outputs the top-ranked alternative in the winning ranking, and an irresolute version contains all alternatives that can be made the winner for some tie-breaking mechanisms (c.f. the *parallel-universes tiebreaking* [Conitzer et al. 2009]).

In the *Chamberlin and Courant rule* [Chamberlin and Courant 1983], we are given a satisfaction function $\vec{s} = (s_1, \dots, s_m)$ (c.f. the scoring function for positional scoring rules) and a number $k \in \mathbb{N}$. We want to choose k alternatives such that the total satisfaction is maximized, where the satisfaction of an agent w.r.t. a set of k alternatives

¹In case there is a tie, we apply a tie-breaking mechanism.

is her maximum satisfaction of any single alternative in the set evaluated by \vec{s} . The *Monroe rule* [Monroe 1995] further requires that each alternative can only be used to satisfy no more than $\lceil \frac{n}{k} \rceil$ agents.

In this paper, given \mathcal{S} , \mathcal{D} , and the number of agents n , a *statistical model* (model for short) $\mathcal{M} = (\mathcal{D}, \mathcal{S}^n, \vec{\pi})$ has three parts: a *parameter space* \mathcal{D} , which is the same as the decision space, a *sample space* \mathcal{S}^n , and a set of distributions over \mathcal{S} , denoted by $\vec{\pi} = \{\pi_d : d \in \mathcal{D}\}$. Intuitively, π_d is the distribution over i.i.d. votes when the ground truth is d .

Definition 1 (Mallows' model with fixed dispersion [Mallows 1957]) Given \mathcal{C} and $0 \leq \varphi \leq 1$, the model is $(\mathcal{L}(\mathcal{C}), \mathcal{L}(\mathcal{C})^n, \vec{\pi})$ such that for any profile P , we have $\pi_W(P) = \prod_{V \in P} \left(\frac{1}{Z_M} \varphi^{\text{Kendall}(V, W)} \right)$, where $\text{Kendall}(V, W)$ is Kendall-tau distance between V and W , that is, the number of different pairwise comparisons in V and W . Z_M is the normalization factor such that $Z_M = \sum_{U \in \mathcal{L}(\mathcal{C})} \varphi^{\text{Kendall}(U, W)}$.

Throughout the paper we let P_n denote an i.i.d.-generated profile of n votes from a distribution π_d that is clear from the context. We further require that $\pi_d(V) > 0$ for all $V \in \mathcal{S}$ and all $d \in \mathcal{D}$. Given $\mathcal{M} = (\mathcal{D}, \mathcal{S}^n, \vec{\pi})$, a mechanism r is *statistically consistent*² w.r.t. \mathcal{M} if for all $d \in \mathcal{D}$, $\lim_{n \rightarrow \infty} \Pr(r(P_n) = d) \rightarrow 1$, where the n votes in P_n are generated i.i.d. from π_d . That is, a statistically consistent mechanism correctly reveals the ground truth with probability 1 as n goes to infinity.

3. GENERALIZED DECISION SCORING RULES

For any $K \in \mathbb{N}$, let $\mathcal{B}_K = \{b_1, \dots, b_K\}$ represents the K dimensions of \mathbb{R}^K . A *total preorder* (preorder for short) is a reflexive, transitive, and total relation. Let $\text{Pre}(\mathcal{B}_K)$ denote the set of all preorders over \mathcal{B}_K . For any $\vec{p} \in \mathbb{R}^K$, we let $\text{Order}(\vec{p})$ denote the preorder \succeq over \mathcal{B}_K where $b_{k_1} \succeq b_{k_2}$ if and only if $[\vec{p}]_{k_1} \geq [\vec{p}]_{k_2}$. That is, the k_1 -th component of \vec{p} is at least as large as the k_2 -th component of \vec{p} . For any preorder \succeq , if $b \succeq b'$ and $b' \succeq b$, then we write $b =_{\succeq} b'$. Each preorder \succeq naturally induces a (partial) strict order \triangleright , where $b \triangleright b'$ if and only if $b \succeq b'$ and $b' \not\succeq b$.

Definition 2 (Generalized decision scoring rules) Given an decision space \mathcal{D} , $K \in \mathbb{N}$, $f : \mathcal{S} \rightarrow \mathbb{R}^K$ and $g : \text{Pre}(\mathcal{B}_K) \rightarrow \mathcal{D}$, we define a generalized decision scoring rule (GDSR), denoted by $\text{GDSR}_{(f, g)}$, to be a mapping so that for any profile P , $\text{GDSR}_{(f, g)}(P) = g(\text{Order}(f(P)))$, where $f(P) = \sum_{V \in P} f(V)$.³

Moreover, if $f(V) \in \mathbb{Z}^K$ for all $V \in \mathcal{S}$, then $\text{GDSR}_{(f, g)}$ is called an integer GDSR.

In words, a GDSR first uses f to transform the input profile P to a vector $f(P) = \sum_{V \in P} f(V)$ in \mathbb{R}^K , then use g to select the winner based on the preorder over the components in $f(P)$. For any $V \in \mathcal{S}$, $f(V)$ is called a *generalized scoring vector*, $f(P)$ is called the *total generalized score vector*. To simplify notation, we let $\text{Order}_f(P) = \text{Order}(f(P))$. We note that $\text{Order}_f(P)$ is a preorder over \mathcal{B}_K , which means that it may contain ties. For any distribution π over \mathcal{S} , we define $f(\pi) = \sum_{V \in \mathcal{S}} \pi(V) f(V)$ and $\text{Order}_f(\pi) = \text{Order}(f(\pi))$. In this paper, we assume that no components in the generalized scoring vectors are *redundant*. That is, for each pair $k_1 \neq k_2$, there always exists $V \in \mathcal{S}$ such that $[f(V)]_{k_1} \neq [f(V)]_{k_2}$. This is without loss of generality because if there is a redundant component, we can easily remove it without changing the GDSR.

Example 1 (Borda) Borda is the positional scoring rule with scoring vector $\vec{s} = (m - 1, m - 2, \dots, 0)$. $K = m$. f_B and g_B are defined as follows.

²This should not be confused with the *consistency* axiom in social choice.

³Equivalently, GDSRs can be defined geometrically similar to *hyperplane rules* [Mossel et al. 2013].

f_B : For any $V \in \mathcal{L}(\mathcal{C})$ and any $i \leq m$, $[f_B(V)]_i$ is the score of c_i in V . That is, $\bar{s}(V, c_i)$. Hence, for any profile P and $i \leq m$, $[f_B(P)]_i$ is the total Borda score of c_i in P .

g_B : For the resolute version ($\mathcal{D} = \mathcal{C}$), g_B selects the alternative that corresponds to the largest component in $f_B(P)$ (and uses a tie-breaking mechanism when necessary); for the irresolute version ($\mathcal{D} = (2^{\mathcal{C}} \setminus \emptyset)$), g_B selects all alternatives that correspond to the largest components in $f_B(P)$; for the preference function version ($\mathcal{D} = \mathcal{L}(\mathcal{C})$), g_B selects $\text{Order}(f_B(P))$ if it is a linear order, otherwise uses a tie-breaking mechanism to obtain the winning linear order. \square

Example 2 (STV) For STV, the generalized scoring vectors have exponentially many components. For every proper subset A of \mathcal{C} and every alternative c not in A , there is a component in the vector that contains the number of times that c is ranked first if all alternatives in A are removed. We define $\text{GDSR}_{(f,g)}$ as follows.

- $K = \sum_{i=0}^{m-1} \binom{m}{i}(m-i)$; the elements of \mathcal{B}_K are indexed by (A, j) , where A is a proper subset of \mathcal{C} and $j \leq m, c_j \notin A$.
- $[f(V)]_{(A,j)} = 1$, if after removing A from V , c_j is at the top of the modified V ; otherwise, $[f(V)]_{(A,j)} = 0$.
- g mimics the process of STV to select a winner (for the resolute version), a set of winners (for the irresolute version, using the parallel-universes tiebreaking [Conitzer et al. 2009]), or a ranking (for the preference function version). \square

Example 3 (MLE) MLE with a fixed-order tie-breaking⁴ of any model \mathcal{M} is a GDSR. Let $K = |\mathcal{D}|$, $\mathcal{D} = \{d_1, d_2, \dots, d_K\}$, and for any $V \in \mathcal{S}$ and any $i \leq K$, $[f(V)]_i = \log \pi_{d_i}(V)$. g outputs the decision that corresponds to the largest component and uses the same tie-breaking mechanism as in the MLE. Not all MLEs are GDSRs. For example, the MLE that uses the first agent's vote to break ties is not a GDSR as it violates anonymity (see Theorem 8). \square

Example 4 (The Chamberlin and Courant rule) The Chamberlin and Courant rule is a GDSR where each component of $f(V)$ is indexed by a set of k alternatives, and its value is the satisfaction of V via the given satisfaction function. \square

MLEs and the Chamberlin and Courant rule are not generalized scoring rules because the decisions are not single winners. We will show that the Monroe rule is not a GDSR in Example 7 after presenting the axiomatic characterization (Theorem 8).

In addition, we can prove by construction that many other commonly studied social choice mechanisms are integer GDSRs. The constructions are similar to those for generalized scoring rules [Xia and Conitzer 2008].

Proposition 1 *Generalized scoring rules [Xia and Conitzer 2008] are GDSRs with $\mathcal{D} = \mathcal{C}$. The Chamberlin and Courant rule (for integer satisfaction function) is a GDSR. The resolute version, irresolute version, and preference function version of positional scoring rules (for integer score vectors), Bucklin, Copeland, maximin, ranked pairs, STV are integer GDSRs.⁵*

4. STATISTICAL CONSISTENCY

We first introduce some definitions to present the results.

Definition 3 (Extension of a preorder) *We say that $\succeq' \in \text{Pre}(\mathcal{B}_K)$ is an extension of $\succeq \in \text{Pre}(\mathcal{B}_K)$, if for all $b, b' \in \mathcal{B}_K$, we have $(b \succ b') \Rightarrow (b \succ' b')$. For any $\succeq, \succeq' \in \text{Pre}(\mathcal{B}_K)$, we let $\succeq \oplus \succeq'$ denote the preorder in $\text{Pre}(\mathcal{B}_K)$ obtained from \succeq by using \succeq' to break ties.*

⁴A fixed-order tie-breaking break ties among alternatives w.r.t. a fixed linear order over all alternatives.

⁵Definitions of these rules can be found in [Nurmi 1987].

That is, b_i is strictly preferred to b_j in $(\succeq \oplus \succeq')$ if and only if (1) $b_i \succ b_j$, or (2) $b_i =_{\succeq} b_j$ and $b_i \succ' b_j$.

For example, $[c_1 \succ' c_2 \succ' c_3]$ is an extension of $[c_1 =_{\succeq} c_2 \succ c_3]$, but is not an extension of $[c_1 =_{\succeq} c_3 \succ c_2]$.

Definition 4 (Possible linear orders) Given a generalized scoring function f , we define the set of possible linear orders, denoted by $PL(f)$, to be the linear orders over \mathcal{B}_K that are the orders of the total score vector of some profile. Formally, $PL(f) = \{\text{Order}_f(P) : P \in \mathcal{S}^*\} \cap \mathcal{L}(\mathcal{B}_K)$.

Definition 5 (Neighborhood) For any $\succeq \in \text{Pre}(\mathcal{B}_K)$, we define the neighborhood of \succeq w.r.t. f , denoted by $\text{Nbr}_f(\succeq)$, to be all linear orders over \mathcal{B}_K that can be obtained from \succeq by using a linear order in $PL(f)$ to break ties. That is, $\text{Nbr}_f(\succeq) = \{\succeq \oplus \triangleright^* : \triangleright^* \in PL(f)\}$. Given f , the neighborhood of a distribution π , denoted by $\text{Nbr}_f(\pi)$, is the neighborhood of $f(\pi)$. That is, $\text{Nbr}_f(\pi) = \text{Nbr}_f(\text{Order}_f(\pi))$.

For example, let $m = 3$, let f_B be the function for Borda as in Example 1, and let π be the distribution where $\pi(c_1 \succ c_2 \succ c_3) = \pi(c_1 \succ c_3 \succ c_2) = 0.5$. Then $\text{Order}_{f_B}(\pi) = [c_1 \succ c_2 =_{\succeq} c_3]$ and $\text{Nbr}_{f_B}(\pi) = \{[c_1 \succ c_2 \succ c_3], [c_1 \succ c_3 \succ c_2]\}$. We note that the definition of neighborhood does not involve the g function.

Theorem 1 Given $\mathcal{M} = (\mathcal{D}, \mathcal{S}^n, \bar{\pi})$, f , and g , $\text{GDSR}_{(f,g)}$ is consistent w.r.t. \mathcal{M} if and only if for all $d \in \mathcal{D}$ and all $\triangleright \in \text{Nbr}_f(\pi_d)$, we have $g(\triangleright) = d$.

Proof: The theorem is mainly based on the following lemma, which characterizes the asymptotic behavior of $\text{Order}_f(P_n)$. It states that for any distribution π over \mathcal{S} , if we generate votes in P_n i.i.d. from π , then $\text{Order}_f(P_n)$ asymptotically almost surely (a.a.s.) falls in the neighborhood of π w.r.t. f . We recall that it is assumed that the generalized scoring vectors have no redundant components, which means that no pair of components in the total generalized score vector are always equal for all profiles.

Lemma 1 Given a generalized scoring function f , for any distribution π that is positive everywhere on \mathcal{S} , we have:

- (1) for any $\triangleright \in \text{Nbr}_f(\pi)$, there exists a constant $\delta_{\triangleright} > 0$ so that for sufficiently large n , $\Pr(\text{Order}_f(P_n) = \triangleright) > \delta_{\triangleright}$.
- (2) for any $\succeq \notin \text{Nbr}_f(\pi)$, $\lim_{n \rightarrow \infty} \Pr(\text{Order}_f(P_n) = \succeq) = 0$.

Proof: We first illustrate the idea behind the proof in a very special case where $\text{Order}_f(\pi)$ is a linear order. In this case $\text{Nbr}_f(\pi) = \{\text{Order}_f(\pi)\}$. Then, the Central Limit Theorem tells us that for each linear order V , the frequency of V in P_n goes to $\pi(V)$ as $n \rightarrow \infty$, and the noise is $O(n^{-1})$. That is, with probability that goes to 1, votes in P_n are distributed as $\pi + O(n^{-1})\pi_{\text{noise}}$. Then, when n is sufficiently large, the $O(n^{-1})\pi_{\text{noise}}$ part cannot affect $\text{Order}_f(\pi)$. Hence, as n goes to ∞ , $\text{Order}_f(P_n) = \text{Order}_f(\pi) \in \text{Nbr}_f(\pi)$.

The proof for the general case is more involved, because if $\text{Order}_f(\pi)$ is not a linear order, then the noise part $O(n^{-1})\pi_{\text{noise}}$ acts as a tie-breaker and thus cannot be overlooked even for large n . Our main mathematical tool for estimating π_{noise} is the Multivariate Lindeberg-Lévy Central Limit Theorem [Greene 2011, Theorem D.18A], which states that for i.i.d. generated vector-valued random variables X_i , if the covariance matrix Σ for the components of X_i is nonsingular, then $(\sum_{i=1}^n X_i - nE(X_i))/\sqrt{n}$ converges in probability to a multivariate normal distribution $\mathcal{N}(0, \Sigma)$. However, this theorem cannot be directly applied to analyze the asymptotic frequencies of the linear orders because the resulting covariance matrix is singular. This is because for any

given n , the number of occurrences of all $m!$ linear orders must sum up to n , which means that they are linearly correlated.

Let $\mathcal{S} = \{l_1, \dots, l_{m!}\}$ denote the set of all $m!$ linear orders. To avoid the singularity, our analysis will focus on $l_1, \dots, l_{m!-1}$. For any $j \leq m! - 1$, let \vec{v}_j denote the vector in $\{0, 1\}^{m!-1}$ where the j -th component is 1 and all other components are zeros. We then define i.i.d. multivariate random variables X_1, \dots, X_n , where each X_i takes \vec{v}_j with probability $\pi(l_j)$, and takes $\vec{0}$ with probability $\pi(l_{m!})$. It is not hard to verify that the mean of X_1 is $E(X_1) = (\pi(l_1), \dots, \pi(l_{m!-1}))$ and the covariance matrix is the following.

$$\Sigma_\pi = \begin{bmatrix} \pi(l_1) - \pi(l_1)^2 & -\pi(l_1)\pi(l_2) & \cdots & -\pi(l_1)\pi(l_{m!-1}) \\ -\pi(l_2)\pi(l_1) & \pi(l_2) - \pi(l_2)^2 & \cdots & -\pi(l_2)\pi(l_{m!-1}) \\ \vdots & \vdots & \ddots & \vdots \\ -\pi(l_{m!-1})\pi(l_1) & -\pi(l_{m!-1})\pi(l_2) & \cdots & \pi(l_{m!-1}) - \pi(l_{m!-1})^2 \end{bmatrix}$$

Since each diagonal element is strictly larger than the sum of the absolute values of other elements in the same row, Σ_π is non-singular according to the Levy-Desplanques Theorem [Horn and Johnson 1985]. Let $Y_n = X_1 + \dots + X_n$. Each Y_n naturally corresponds to a profile P_n of n votes, where for all $j \leq m! - 1$, $[Y_n]_j$ is the number of occurrences of l_j , and $n - \sum_{j=1}^{m!-1} [Y_n]_j$ is the number of occurrences of $l_{m!}$. By the multivariate Central Limit Theorem, $Y_{noise} = \frac{Y_n - nE(X_1)}{\sqrt{n}}$ converges in distribution to the multivariate normal distribution $\mathcal{N}(0, \Sigma_\pi)$.

Part (1) of the lemma. For any $\triangleright \in \text{Nbr}_f(\pi)$, there exists a profile P such that (1) $\text{Order}_f(P) \in \mathcal{S}$, and (2) $\text{Order}_f(\pi) \oplus \text{Order}_f(P) = \triangleright$. We define $\vec{p} \in \mathbb{R}^{m!}$ such that for all $j \leq m!$, $[\vec{p}]_j = P(l_j)/|P| - 1/m!$, where $P(l_j)$ is the number of occurrences of l_j in P . That is, $\sum_j [\vec{p}]_j = 0$ and for all j , $|[\vec{p}]_j| \leq 1$. Since $\text{Order}(\vec{p}) = \text{Order}_f(P)$ and is a strict order, there exist positive numbers $\delta_1, \dots, \delta_{m!-1}$ such that for any vector $\vec{q} \in \mathbb{R}^{m!}$ with (1) $\sum_j [\vec{q}]_j = 0$ and (2) for all $j \leq m! - 1$, $|[\vec{p}]_j - [\vec{q}]_j| < \delta_j$, we have $\text{Order}_f(\vec{q}) = \text{Order}_f(\vec{p})$.

Let $S = \prod_{j=1}^{m!-1} [\pi(l_j) - \frac{1}{m!} - \delta_j, \pi(l_j) - \frac{1}{m!} + \delta_j]$ denote a hypercube in $\mathbb{R}^{m!-1}$. When \vec{x} is generated from $\mathcal{N}(0, \Sigma_\pi)$, the probability that $\vec{x} \in S$ is strictly positive because $\mathcal{N}(0, \Sigma_\pi)$ has full support. It is not hard to prove that for any Y_n , if $Y_{noise} \in S$, then for the corresponding profile P_n we have $\text{Order}_f(P_n) = \text{Order}_f(\pi) \oplus \text{Order}_f(\vec{p}) = \triangleright$. Hence the probability for $\text{Order}_f(P_n) = \triangleright$ is at least $\Pr(Y_{noise} \in S)$, which converges to $\Pr(\vec{x} \in S)$ when \vec{x} is generated from $\mathcal{N}(0, \Sigma_\pi)$. This proves part (1).

Part (2) of the lemma. For any $\triangleright \notin \text{Nbr}_f(\pi)$, we prove the lemma in the following three cases.

Case 1: \triangleright does not extend $\text{Order}_f(\pi)$. Following a similar argument with the case where $\text{Order}_f(\pi)$ is a linear order, if b_i is strictly preferred to b_j in $\text{Order}_f(\pi)$, then with probability that goes to 1, b_i is strictly preferred to b_j in $\text{Order}_f(P_n)$. So the probability for $\text{Order}_f(P_n) = \triangleright$ goes to 0.

Case 2: \triangleright is not a linear order. We recall that for any pair of $k_1, k_2 \leq K$ with $k_1 \neq k_2$, there exists a linear order l such that $[f(l)]_{k_1} \neq [f(l)]_{k_2}$. Therefore, following the Berry-Esseen theorem, the probability of a tie between the k_1 -th component and k_2 -th component of $f(P_n)$ for i.i.d. generated P_n is $O(n^{-0.5})$, which goes to 0 as $n \rightarrow \infty$.

Case 3: \triangleright is a linear order and extends $\text{Order}_f(\pi)$, but there is no profile $\triangleright \in \text{PL}(f)$ such that $\triangleright = \text{Order}_f(\pi) \oplus \triangleright$. It follows that $\triangleright \notin \text{PL}(f)$, otherwise $\triangleright = \text{Order}_f(\pi) \oplus \triangleright$, which is a contradiction. Hence for any profile P_n , $f(P_n) \neq \triangleright$, which means that $\Pr(f(P_n) = \triangleright) = 0$. ■

The “if” direction follows after Lemma 1. To prove the “only if” direction, if there exists o and $\triangleright \in \text{Nbr}_f(\pi_d)$ with $g(\triangleright) \neq o$, then by Lemma 1, as $n \rightarrow \infty$, the probability for the order over the components of the total generalized score vector to be \triangleright is non-

negligible. Hence, with non-negligible probability $\text{GDSR}_{(f,g)}$ will not output o when P_n is generated i.i.d. from π_d , which means that $\text{GDSR}_{(f,g)}$ is not consistent. ■

Theorem 1 can be applied to prove that STV is statistically consistent w.r.t. the Mallows model and the model by Conitzer and Sandholm [2005].

Proposition 2 *STV (preference function) is consistent w.r.t. the Mallows model. STV (resolute rule) is consistent w.r.t. all models for positional scoring rules proposed by Conitzer and Sandholm [2005].*

Proof: We will use the following straightforward corollary of Theorem 1.

Corollary 1 *Given a model $\mathcal{M} = (\mathcal{D}, \mathcal{S}^n, \vec{\pi})$, if for all $d \in \mathcal{D}$ and all extensions $\triangleright \in \mathcal{L}(\mathcal{B}_K)$ of $\text{Order}_f(\pi_d)$ we have $g(\triangleright) = d$, then $\text{GDSR}_{(f,g)}$ is consistent w.r.t. \mathcal{M} .*

The GDSR formulation of STV (as a resolute rule and as a preference function) is the same as in Example 2. W.l.o.g. suppose the ground truth parameter $d = [c_1 \succ \dots \succ c_m]$. By Corollary 1, to show that STV (preference function) is consistent w.r.t. \mathcal{M}_φ , it suffices to show that for any $2 \leq k \leq m$, after removing $C_k = \{c_{k+1}, \dots, c_m\}$, c_k has strictly the lowest expected plurality score, where the expectation is taken over a randomly generated ranking from \mathcal{M}_φ given d . To this end, for any $i < k$, we consider the following one-one mapping. For any rankings V where c_i is ranked at the top after all alternatives in C_k are removed, we switch the positions of c_i and c_{i+1} . This will give us another ranking V' where c_{i+1} is ranked in the top position if alternatives in C_k are removed. It is easy to check that $\text{Kendall}(V, d) = \text{Kendall}(V', d) - 1$, which means that the expected plurality score of c_i is higher than the expected plurality score of c_{i+1} after C_k is removed. This shows that if C_k is removed, then the expected score of c_k is strictly smaller than all other remaining alternatives. Hence for any \triangleright that is an extension of $\text{Order}_f(\pi_d)$, $g(\triangleright) = d$. By Corollary 1, STV is consistent w.r.t. \mathcal{M}_φ .

To prove that STV (resolute rule) is consistent w.r.t. $\mathcal{M}_{\vec{s}}$, w.l.o.g. suppose $d = c_1$, it suffices to show that for any $C \subseteq \mathcal{C}$, after removing C , c_1 has the strictly largest expected plurality score. This can be proved by a similar argument to the proof for \mathcal{M}_φ : for any other alternative $c \neq c_1$, for any linear order V where c_1 is ranked in the top after removing C , we can obtain another linear order V' by switching the positions of c_1 and c . Since the position of c_1 in V is strictly higher than the position of c_1 in V' , we have $\pi_d(V) \geq \pi_d(V')$, and the inequality is strict for some V . The proposition follows after a similar argument as for \mathcal{M}_φ . ■

The next theorem characterizes all GDSRs that are consistent w.r.t. *some* models.

Theorem 2 *A GDSR is consistent w.r.t. some model if and only if for all $d \in \mathcal{D}$, $g^{-1}(d) \cap \text{PL}(f) \neq \emptyset$.*

We note that the condition in Theorem 2 is stronger than requiring that the GDSR is an onto mapping.

Proof: The “if” direction: For any profile P' with $\text{Order}_f(P') \in g^{-1}(d) \cap \text{PL}(f)$, since $\text{Order}_f(P')$ is a linear order, there exists $t \in \mathbb{N}$ so that $\text{Order}_f(tP' \cup S) = \text{Order}_f(P') \triangleright$, where $tP' \cup S$ is the profile composed of t copies of P' plus each linear order in S . We note that $P_d = tP' \cup S$ is a profile that contains all types of linear orders. Then, we define a distribution π_d such that for any linear order V , $\pi_d(V) = \frac{P_d(V)}{|P_d|}$, where $P_d(V)$ is the number of occurrences of V in P_d . Consistency follows after Theorem 1, because the neighborhood of P_d only contains $f(P_d)$.

The “only if” direction. Suppose there exists $d \in \mathcal{D}$ such that $g^{-1}(d) \cap \text{PL}(f) = \emptyset$. We prove the following lemma.

Lemma 2 *For any distribution π and any generalized scoring function f , $\text{Nbr}_f(\pi) \subseteq \text{PL}(f)$.*

Proof: For any $\triangleright \in \text{Nbr}_f(\pi)$, let P_{\triangleright} be a profile such that $\triangleright = \text{Order}_f(\pi) \oplus \text{Order}_f(P_{\triangleright})$. The existence of P_{\triangleright} is guaranteed by the definition of $\text{Nbr}_f(\pi)$ (Definition 5). For any n , we let $Q_n = Q_n^1 \cup Q_n^2$ be a profile composed of the following two parts.

- (1) The first part Q_n^1 contains the following votes: for any $V \in \mathcal{S}$, there are $\lfloor \pi(V) \cdot n \rfloor$ copies of V .
- (2) The second part Q_n^2 contains $\lfloor \sqrt{n} \rfloor$ copies of P_{\triangleright} .

By the Central Limit Theorem, for any i, j such that b_i is strictly preferred to b_j in $\text{Order}_f(\pi)$, as n goes to infinity $[f(Q_n^1)]_i - [f(Q_n^1)]_j = \Theta(n)$ a.s.; for any i, j such that b_i is tied with b_j in $\text{Order}_f(\pi)$, as n goes to infinity $|[f(Q_n^1)]_i - [f(Q_n^1)]_j| = O(1)$ a.s. Therefore, Q_n^2 effectively acts as a tie-breaker for $\text{Order}_f(\pi)$ in the same way as $\text{Order}_f(P_{\triangleright})$. This shows that there exists n such that $\text{Order}_f(Q_n) = \triangleright$ and proves the lemma. ■

By this lemma, because $g^{-1}(d) \cap \text{PL}(f) = \emptyset$, for any distribution π and any $\triangleright \in \text{Nbr}_f(\pi)$, $g(\triangleright) \neq d$. By Theorem 1, $\text{GDSR}_{(f,g)}$ is not consistent w.r.t. any model. ■

4.1. Convergence Rate

Given $\mathcal{M} = (\mathcal{D}, \mathcal{S}^n, \vec{\pi})$ and a consistent $\text{GDSR}_{(f,g)}$, we next give an upper bound on the convergence rate for $\text{GDSR}_{(f,g)}$ to reveal the ground truth with high probability under i.i.d. generated votes. Let s_{max} denote the maximum absolute value of the components in all generalized scoring vectors. That is, $s_{max} = \max_{V,j} |[f(V)]_j|$. Let s_{min} denote the minimum non-zero absolute value of the components in all generalized scoring vectors. Let d_{min} denote the smallest non-zero difference between the components in all $f(\pi_d)$. That is, $d_{min} = \min_{i,j \leq K, d} |[f(\pi_d)]_i - [f(\pi_d)]_j| : [f(\pi_d)]_i \neq [f(\pi_d)]_j$. Let p_{min} denote the minimum non-zero probability of any linear order under any parameter, that is, $p_{min} = \min_{V,d} \pi_d(V)$.

Theorem 3 Suppose $\text{GDSR}_{(f,g)}$ is a consistent estimator for $\mathcal{M} = (\mathcal{D}, \mathcal{S}^n, \vec{\pi})$. For any $d \in \mathcal{D}$ and $n \in \mathbb{N}$, we have:

$$\Pr(\text{GDSR}_{(f,g)}(P_n) \neq d) < K \cdot \exp\left(-n \cdot \frac{d_{min}}{8s_{max}^2}\right) + \frac{(K(K-1)s_{max})^3}{(2p_{min})^{1.5}(s_{min})^3\sqrt{n}} = O(n^{-0.5})$$

Proof: Let $\text{Strict}(\pi_d)$ denote the set of strict pairwise comparisons in $\text{Order}_f(\pi_d)$, that is, $(b_i, b_j) \in \text{Strict}(\pi_d)$ if and only if b_i is strictly preferred to b_j in $\text{Order}_f(\pi_d)$. For any profile P , if $\text{Order}_f(P)$ is a linear order that extends $\text{Order}_f(\pi_d)$, then $\text{Order}_f(P) \in \text{Nbr}_f(\pi_d)$. By Theorem 1, $\text{GDSR}_{(f,g)}(P) = d$. Hence, if $\text{GDSR}_{(f,g)}(P) \neq d$, then there are only two possibilities: (1) for some $(b_i, b_j) \in \text{Strict}(\pi_d)$, $[f(P)]_j \geq [f(P)]_i$, or (2) there exist $i \neq j$ with $[f(P)]_i = [f(P)]_j$.

For case (1), for any pair of $(b_i, b_j) \in \text{Strict}(\pi_d)$, we let X_1, \dots, X_n denote i.i.d. variables that represent $[f(l)]_i - [f(l)]_j$ for randomly generated l from π_d . Let $Y_n = (X_1 + \dots + X_n)/n$. We have $E(X_i) \geq d_{min}$, $\text{Var}(X_i) < 2s_{max}$, and each X_i takes a value in $[-2s_{max}, 2s_{max}]$. By Hoeffding's inequality [Hoeffding 1963], we have: $\Pr(Y_n \leq 0) = \Pr(Y_n - E(X_1) \leq -E(X_1)) \leq \exp\left(-\frac{2n^2 E(X_1)}{n(4s_{max})^2}\right) \leq \exp\left(-n \cdot \frac{d_{min}}{8s_{max}^2}\right)$.

For case (2), for any pair of i, j with $[f(\pi_d)]_i = [f(\pi_d)]_j$, we define X_i and Y_n similarly as in case (1). The third moment of X_1 is no more than s_{max}^3 and $\text{Var}(X_i) \geq p_{min}s_{min}^2$. By Berry-Esseen theorem, the probability for $Y_n = 0$ is no more than $\frac{(s_{max})^3}{(p_{min})^{1.5}(s_{min})^3\sqrt{n}}$.

Combining the above calculations, for (1) we only need to consider adjacent pairs in $\text{Strict}(\pi_d)$ and for (2) we need to consider all pairs of tied components. Hence the probability that either (1) or (2) holds is at most $K \cdot \exp\left(-n \cdot \frac{d_{min}}{8s_{max}^2}\right) + \frac{(K(K-1)s_{max})^3}{(2p_{min})^{1.5}(s_{min})^3\sqrt{n}}$, which proves the theorem. ■

The next theorem states that the $O(n^{-0.5})$ bound in Theorem 3 is asymptotically tight for some model and GDSR.

Theorem 4 *There exists a model \mathcal{M} where $\mathcal{D} = \mathcal{C}$ and a GDSR r such that (1) r is consistent w.r.t. \mathcal{M} , and (2) there exists $d \in \mathcal{D}$ such that for all even numbers n , $\Pr(r(P_n) \neq d) = \Omega(n^{-0.5})$, where votes in P_n are generated i.i.d. from π_d .*

Proof: Let there be three alternatives $\{c_1, c_2, c_3\}$, $\mathcal{D} = \mathcal{C}$, and $\mathcal{S} = \mathcal{L}(\mathcal{C})$. We define the probability distributions in model \mathcal{M} as follows:

$$\begin{aligned}\pi_{c_1}(c_1 \succ c_2 \succ c_3) &= 0.5, & \pi_{c_1}(c_1 \succ c_3 \succ c_2) &= 0.5 \\ \pi_{c_2}(c_2 \succ c_1 \succ c_3) &= 0.5, & \pi_{c_2}(c_2 \succ c_3 \succ c_1) &= 0.5 \\ \pi_{c_3}(c_3 \succ c_1 \succ c_2) &= 0.5, & \pi_{c_3}(c_3 \succ c_2 \succ c_1) &= 0.5\end{aligned}$$

Let r be the Borda rule with fixed order tie-breaking $c_1 \succ c_2 \succ c_3$, except in one case: if c_1 's total score is strictly the largest, and the total scores of c_2 and c_3 are exactly the same, then the winner is c_2 (instead of c_1 for Borda). It is not hard to verify that r is a GDSR, using the same f_B in Example 1, and a slightly different g'_B that selects c_2 when the preorder is $c_1 \triangleright c_2 =_{\triangleright} c_3$, otherwise g'_B is the same as g_B . By Theorem 1, r is consistent w.r.t. \mathcal{M} .

For any profile P and alternative c , let $s_B(P, c)$ denote the Borda score of c in P . For any even n , when the ground truth is c_1 , the probability for $s_B(P_n, c_2) = s_B(P_n, c_3)$ is $\binom{n}{n/2}/2^n$. By Stirling's formula, we have

$$\frac{\binom{n}{n/2}}{2^n} = \frac{n!}{(\frac{n}{2}!)^2 2^n} \approx \frac{\sqrt{2\pi n} (\frac{n}{e})^n}{(\sqrt{\pi n} (\frac{n}{2e})^{n/2})^2 2^n} = \frac{\sqrt{2}}{\sqrt{\pi n}} = \Omega(n^{-0.5})$$

Similar to the proof of Theorem 3, it is not hard to show that the probability for the total score of c_1 to be the highest is $1 - \exp(-\Omega(n)) = 1 - o(n^{-0.5})$. So the probability for $s_B(P_n, c_1) > s_B(P_n, c_2) = s_B(P_n, c_3)$ is $\Omega(n^{-0.5})$. In all such cases $r(P_n) = c_2 \neq c_1$, which proves the theorem. \blacksquare

For specific distributions and GDSRs we can improve the convergence rate as in the following proposition. The proof follows after a straightforward application of the Hoeffding's inequality and is thus omitted.

Proposition 3 *Suppose $\text{GDSR}_{(f,g)}$ is consistent w.r.t. $\mathcal{M} = (\mathcal{D}, \mathcal{S}^n, \vec{\pi})$. For any $d \in \mathcal{D}$ and $n \in \mathbb{N}$, if for all extensions \triangleright of $\text{Order}_f(\pi_d)$, $g(\triangleright) = d$, then:*

$$\Pr(\text{GDSR}_{(f,g)}(P_n) \neq d) < K \cdot \exp\left(-n \cdot \frac{d_{\min}}{8s_{\max}^2}\right)$$

Example 5 *The bound in Proposition 3 applies to STV (preference function) w.r.t. \mathcal{M}_φ for all φ and STV (resolute rule) w.r.t. $\mathcal{M}_{\vec{s}}$ for all \vec{s} , following the proof of Proposition 2.*

5. COMPUTATIONAL PROPERTIES

We start with a simple observation showing that winner determination is *fixed-parameter tractable* w.r.t. the number alternatives.⁶ A problem is fixed-parameter tractable w.r.t. parameter p if there exists an algorithm that solves the problem and runs in time $h(p)|I|^{O(1)}$ for some function h of p , where $|I|$ is the input size. One natural interpretation is that when $h(p)$ is small, the problem can be solved in polynomial time.

Theorem 5 *Computing the decision for any integer GDSR is fixed-parameter tractable w.r.t. the number of alternatives.*

⁶We assume that f and g in $\text{GDSR}_{(f,g)}$ can be computed in polynomial time in m .

Proof sketch: The theorem is proved by observing that (1) computing $f(P)$ takes $h(m) \cdot n$ steps for some function h , (2) computing the preorder over \mathcal{B}_K takes no more than $K^2 h'(m) \cdot n$ steps, and (3) computing g takes time that only depends on K . ■

5.1. Vote Operations and an ILP Formulation

In many types of strategic behavior investigated in the (computational) social choice literature, the strategic entity (e.g. a group of manipulators, a briber, or a chairman) affects the outcome of the mechanism by changing the votes in the profile. For GDSRs, any such action can be uniquely represented by changes in the total generalized scoring vector. We first recall the definitions of two well-studied types of agents' strategic behavior, then formally define *vote operations* for integer GDSRs.

Definition 6 ([Zuckerman et al. 2009]) *In a constructive (respectively, destructive) UNWEIGHTED COALITIONAL OPTIMIZATION (UCO) problem, we are given a mechanism r , a non-manipulators' profile P^{NM} , and a (dis)avored decision $d \in \mathcal{D}$. 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, $d \neq r(P^{NM} \cup P^M)$).*

Definition 7 ([Faliszewski et al. 2009]) *In a constructive (respectively, destructive) OPT-BRIBERY problem, we are given a profile P and a (dis)avored decision $d \in \mathcal{D}$. We are asked to compute the smallest number k such that the strategic individual can change no more than k votes such that d is the winner (respectively, d is not the winner).*

Definition 8 *Given an integer GDSR $_{(f,g)}$, a set of vote operations is denoted by $\Delta = [\vec{\delta}_1 \cdots \vec{\delta}_T]$, where for each $i \leq T$, $\vec{\delta}_i \in \mathbb{Z}^K$ is the column vector that represents the changes made to the total generalized score vector by applying the i -th vote operation. For each $l \leq K$, let Δ_l denote the l -th row of Δ .*

Example 6 Actions in UCO are vote operations where Δ is $\{f(V) : V \in \mathcal{S}\}$ (the order of the generalized score vectors in Δ does not matter). That is, the group of manipulators is the strategic individual, and each vote cast by a manipulator is a vote operation.

Actions in OPT-BRIBERY are vote operations where Δ is $\{f(W) - f(V) : V, W \in \mathcal{S}\}$. That is, each action of "changing a vote from V to W " is a vote operation. □

Next, we present an integer linear program ILP_{\triangleright} to compute the minimum number of vote operations for the strategic entity to change the preorder of the components of the total generalized score vector to \triangleright .

Definition 9 *Given GDSR $_{(f,g)}$, a profile P , the vote operations Δ , and a preorder \triangleright over \mathcal{B}_K , we define ILP_{\triangleright} as:*

$$\begin{aligned} & \min \quad \|\vec{v}\|_1 \\ \text{s.t.} \quad & \forall o_i =_{\triangleright} o_j : (\Delta_i - \Delta_j) \cdot \vec{v} = [f(P)]_j - [f(P)]_i \\ & \forall o_i \triangleright o_j : (\Delta_i - \Delta_j) \cdot \vec{v} \geq [f(P)]_j - [f(P)]_i + 1 \\ & \forall i : v_i \geq 0 \text{ and are integers} \end{aligned} \quad (ILP_{\triangleright})$$

In ILP_{\triangleright} , \vec{v} is a column vector, where for each $i \leq T$, v_i represents the number of the i -th operation (corresponding to $\vec{\delta}_i$) taken by the strategic entity. $\|\vec{v}\|_1 = \sum_{i=1}^T v_i$ is the total number of operations. $\Delta \cdot \vec{v}$ is the change in the total generalized scoring vector introduced by the strategic entity, where for any $l \leq K$, $\Delta_l \cdot \vec{v}$ is the change in the l -th component of the total generalized score vector.

Next, we define the strategic entity's three goals and the corresponding computational problems studied in this paper.

Definition 10 In the CONSTRUCTIVE VOTE OPERATION (CVO) problem, we are given $GDSR_{(f,g)}$, a profile P , a favored decision d , and a set of vote operations $\Delta = [\vec{\delta}_1 \cdots \vec{\delta}_T]$, and we are asked to compute the smallest number k , denoted by $CVO(P, d)$, such that there exists a vector $\vec{v} \in \mathbb{N}_{\geq 0}^T$ with $\|\vec{v}\|_1 = k$ and $g(\text{Order}(f(P) + \Delta \cdot \vec{v})) = d$. If such \vec{v} does not exist, then we denote $CVO(P, d) = \infty$.

The DESTRUCTIVE VOTE OPERATION (DVO) problem is defined similarly, where d is the disfavored decision, and we are asked to compute the smallest number k , denoted by $DVO(P, d)$, such that there exists a vector $\vec{v} \in \mathbb{N}_{\geq 0}^T$ with $\|\vec{v}\|_1 = k$ and $g(\text{Order}(f(P) + \Delta \cdot \vec{v})) \neq d$.

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

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

The relationship between ILP_{\triangleright} and CVO, DVO, CWVO is revealed in the following Lemma, whose proof is straightforward.

Lemma 3 Given $GDSR_{(f,g)}$, a decision d , and a profile P , we have:

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

Moreover, the solution to each of the three problems is the optimal solution to the corresponding ILP. For example, if $CVO(P, d) < \infty$, then

$$CVO(P, d) = \min\{\|\vec{v}\|_1 : \vec{v} \text{ is the solution to some } ILP_{\triangleright} \text{ where } g(\triangleright) = d\}$$

5.2. Complexity and Asymptotic Properties of Vote Operations

Theorem 6 For any integer $GDSR$ and any vote operation, CVO, DVO, and CWVO are fixed-parameter tractable w.r.t. the number of alternatives.

Proof sketch: We show the idea behind the CVO case. By Lemma 3, d can be made to win by using no more than k vote operations if and only if ILP_{\triangleright} has a feasible solution for some \triangleright with $g(\triangleright) = d$. Since K is a function that only depends on m and there are no more than 3^{K^2} total preorders over \mathcal{B}_K to make a given alternative win, the total number of ILP_{\triangleright} is no more than 3^{K^2} . For each ILP_{\triangleright} , the number of variables is T and the size of the ILP is $h'(K)n$ for some function h' . By Lenstra's theorem [Lenstra 1983], checking whether the ILP has a feasible solution takes time in $h(K)n$, which can be combined with binary search to compute the feasible solution with the smallest $\|\vec{v}\|_1$. Since K can be computed from m , the problem is fixed-parameter tractable w.r.t. m . ■

We now present the main theorem of this section, which establishes the asymptotic property of the minimum number of vote operations for the strategic entity to succeed.

Theorem 7 Let $GDSR_{(f,g)}$ be an integer $GDSR$, let π be a distribution over \mathcal{S} , and let Δ be a set of vote operations. Suppose we fix the number of alternatives, generate n votes i.i.d. according to π , and let P_n denote the profile. Then, for any alternative c , any $VO \in \{CVO, DVO, CWVO\}$, and any $\epsilon > 0$, there exists $\beta^* > 1$ such that as $n \rightarrow \infty$, the probability for the following four events sum up to more than $1 - \epsilon$: (1) $VO(P_n, d) = 0$, (2) $\frac{1}{\beta^*} \sqrt{n} < VO(P_n, d) < \beta^* \sqrt{n}$, (3) $\frac{1}{\beta^*} n < VO(P_n, d) < \beta^* n$, and (4) $VO(P_n, d) = \infty$.

⁷When $VO = CWVO$, we let $VO(P_n, d)$ denote $CWVO(P_n)$.

Proof: We recall that $f(\pi) = \sum_{V \in \mathcal{S}} \pi(V) \cdot f(V)$. Let $\triangleright_\pi = \text{Order}(f(\pi))$. The theorem is proved in two steps. Step 1: we show that as n goes to infinity, with probability that goes to one the difference between any pair of components in $f(P_n)$ is either $\Theta(\sqrt{n})$ or $\Theta(n)$. Step 2: we apply sensitivity analysis to ILPs that are similar to $\text{ILP}_{\triangleright}$ (Definition 9) to prove the theorem. The idea behind Step 2 is, for any profile P_n , if the difference between a pair of components in $f(P_n)$ is $\Theta(\sqrt{n})$, then we consider this pair of components to be “almost tied”; if the difference is $\Theta(n)$, then we consider them to be “far away”. Take CVO for example, we can easily identify the cases where $\text{CVO}(P_n, d)$ is either 0 (when $\text{GDSR}_{(f,g)} = d$) or ∞ (by Lemma 3). Then, we will first try to break these “almost tied” pairs by using ILPs, and show that if there exists an integer solution \vec{v} , then the objective value $\|\vec{v}\|_1$ is $\Theta(\sqrt{n})$. Otherwise, we have to change the orders between some “far away” pairs by using $\text{ILP}_{\triangleright}$ ’s, and show that if there exists an integer solution to some $\text{ILP}_{\triangleright}$ with $g(\triangleright) = d$, then the objective value is $\Theta(n)$.

Formally, given $n \in \mathbb{N}$ and $\beta > 1$, let \mathcal{P}_β denote the set of all n -vote profiles P that satisfy the following two conditions: for any pair $i, j \leq K$,

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

The following lemma is a straightforward application of the Central Limit Theorem.

Lemma 4 *For any $\epsilon > 0$, there exists β such that $\lim_{n \rightarrow \infty} \Pr(P_n \in \mathcal{P}_\beta) > 1 - \epsilon$.*

For any given ϵ , in the rest of the proof we fix β to be a constant guaranteed by Lemma 4. The next lemma will be frequently used in the proof.

Lemma 5 *Fix an integer matrix \mathbf{A} . There exists a constant $\beta_{\mathbf{A}}$ that only depends on \mathbf{A} , such that if the following LP*

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

has an integer solution, then the solution is no more than $\beta_{\mathbf{A}} \cdot \|\vec{b}\|_\infty$.

Proof: We apply the result by Cook et al. [1986] on the sensitivity analysis of ILPs. Let \mathbf{A} be a $m^* \times n^*$ integer matrix, which includes the constraints $\vec{x} \geq \vec{0}$. Suppose LP (1) has a (non-negative) integer solution. Then, it follows from Theorem 5 (ii) in [Cook et al. 1986] that LP (1) has a (non-negative) integer solution \vec{z} such that $\|\vec{z} - \vec{0}\|_\infty \leq n^* \cdot M(\mathbf{A}) \cdot (\|\vec{b} - \vec{0}\|_\infty + 2)$, where $M(\mathbf{A})$ is the maximum of the absolute values of the determinants of the square sub-matrices of \mathbf{A} . Since \mathbf{A} is fixed, the right hand side becomes a constant, that is, $\|\vec{z}\|_\infty = O(\|\vec{b}\|_\infty)$. Therefore, there exists $\beta_{\mathbf{A}}$ such that the optimal value for integer solutions of LP (1) is no more than $\vec{1} \cdot (\vec{z})' \leq n^* \|\vec{z}\|_\infty \leq \beta_{\mathbf{A}} \cdot \|\vec{b}\|_\infty$. ■

We cannot directly apply Lemma 5 to $\text{ILP}_{\triangleright}$ because sometime $\|\vec{b}\|$ is $\Theta(n)$. Let $\triangleright' \ominus \triangleright$ denote the strict orders that are in \triangleright' but not in \triangleright . That is, $(o_i, o_j) \in (\triangleright' \ominus \triangleright)$ if and only if $o_i \triangleright' o_j$ and $o_i =_{\triangleright} o_j$. We define the following ILP that is similar to $\text{ILP}_{\triangleright}$ to check whether there is a way to break “almost tied” pairs of components so that d is the winner. For any preorder \triangleright and any extension of \triangleright , denoted by \triangleright' , we define $\text{ILP}_{\triangleright' \ominus \triangleright}$ as follows.

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

We note that some constraints in $\text{ILP}_{\triangleright' \ominus \triangleright}$ depend on \triangleright (not only depend on the pairwise comparisons in $(\triangleright' \ominus \triangleright)$). This will not cause confusion because we will always

indicate \succeq in the subscript. It is easy to see that $ILP_{\succeq' \ominus \succeq}$ has a solution \vec{v} if and only if the strategic entity can make the order between any pairs of o_i, o_j with $o_i =_{\succeq} o_j$ to be the one in \succeq' by applying l -th operation for v_l times, and the total number of vote operations is $\|\vec{v}\|_1$.

Below we firsts prove the theorem for CVO, then show how to extend the proof to DVO and CWVO. The following two claims identify the profiles in \mathcal{P}_β for which CVO is $\Theta(\sqrt{n})$ and $\Theta(n)$, respectively, whose proofs are relegated to the appendix.

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

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

Lastly, for any $P \in \mathcal{P}_\beta$ such that $\text{GDSR}_{(f,g)}(P) \neq d$, the only case not covered by Claim 1 and Claim 2 is that there is no \succeq with $\text{GDSR}_{(f,g)}(\succeq) = d$ such that ILP_{\succeq} has an integer solution. It follows from Lemma 3 that in this case $\text{CVO}(P, d) = \infty$. We note that β' in Claim 1 and Claim 2 does not depend on n . Let β^* be an arbitrary number that is larger than the two β' s. This proves the theorem for CVO.

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

Remarks on the non-triviality of the proof. Lemma 4 is quite straightforward and naturally corresponds to a random walk in multidimensional space. However, we did not find an obvious connection between random walk theory and the observation made in Theorem 7. The main difficulty in proving Theorem 7 is handling ties among components of the total generalized score vector for GDSRs. Exact ties only happen with negligible probability for randomly generated profile P_n , but it is not clear when some components are close (but not exactly the same), how often the strategic entity can make some components equivalent to achieve her goal. To tackle this difficulty, we convert the vote manipulation problem to multiple ILPs and apply Lemma 5.

It is not hard to see that control by adding/deleting votes [Bartholdi et al. 1992], margin of victory [Cary 2011; Magrino et al. 2011; Xia 2012], and minimum manipulation coalition size [Pritchard and Wilson 2009] are vote operations. Therefore, we have the following corollary of Theorem 7.

Corollary 2 *For any integer GDSR, any distribution π over S , suppose n agents' votes are i.i.d. generated from π . The number of actions needed for constructive (or destructive) unweighted coalitional manipulation (or bribery, control by adding/deleting votes), margin of victory, or minimum manipulation coalition size is one of the following with probability 1 as $n \rightarrow \infty$. (1) 0, (2) $\Theta(\sqrt{n})$, (3) $\Theta(n)$, and (4) ∞ .*

For some cases the proof needs a minor modification. For example, for bribery it is only possible to choose operation $f(W) - f(V)$ when some agents cast V in the profile. This is handled by restricting the set of vote operations to $\{f(W) - f(V) : \pi(V) > 0\}$, and adding more constraints to the ILPs. We omit the definitions of some types of strategic behavior in Corollary 2 and the formal proof due to the space constraint.

6. AN AXIOMATIC CHARACTERIZATION

A mechanism is *anonymous*, if it is insensitive to permutations over agents' votes.

Definition 11 A set \mathcal{P} of profiles is consistent, if for any $P_1, P_2 \in \mathcal{P}$ with $r(P_1) = r(P_2)$, we have $P_1 \cup P_2 \in \mathcal{P}$. A mechanism r is locally consistent on a consistent set \mathcal{P} if for any $P_1, P_2 \in \mathcal{P}$ with $r(P_1) = r(P_2)$, we have $r(P_1 \cup P_2) = r(P_1) = r(P_2)$.

Definition 12 For any natural number t , a mechanism r is t -consistent if there exists a partition $\{\mathcal{P}_1, \dots, \mathcal{P}_t\}$ of all profiles such that for all $i \leq t$, r is locally consistent within \mathcal{P}_i . A voting rule r is finitely locally consistent if it is t -consistent for some finite t .

Finite local consistency naturally generalizes the *consistency* axiom in social choice, which is the case for $t = 1$, and is a weak constraint because as we have shown, many commonly studied mechanisms satisfy it. We feel that this is desirable because it does not put too much constraint on the new mechanisms that can be discovered in the GDSR framework. Meanwhile, GDSRs have a nice mathematical structure that facilitates further exploration. For example, the g function can be equivalently defined as a decision tree for which there are many learning algorithms.

Theorem 8 Given an decision space, a mechanism is a GDSR if and only if it satisfies anonymity and finite local consistency.

The proof is similar to the proof of the axiomatic characterization for generalized scoring rules [Xia and Conitzer 2009] and can be found on Arxiv. Similar to generalized scoring rules, finite local consistency implies *homogeneity*, which requires that for any profile P and any natural number l , we have $r(P) = r(lP)$. Therefore, any social choice mechanisms that does not satisfy homogeneity is not a GDSR. Because Dodgson's rule does not satisfy homogeneity [Fishburn 1977; Brandt 2009], which means that they are neither generalized scoring rule nor GDSR. We next show that the Monroe rule does not satisfy homogeneity, which means that it is not a GDSR.

Example 7 Let $m = 6$ and $k = 2$, $V = [a_1 \succ a_2 \succ \text{others} \succ a_3]$, and $V^* = [a_3 \succ a_2 \succ \text{others} \succ a_1]$. Let $P = (V, V, V^*)$. The Monroe winner for P is $\{a_1, a_3\}$, where a_1 is assigned to both voters whose preferences are V and a_3 is assigned to the voter whose preferences are V^* . However, the Monroe winner for $2P$ is $\{a_1, a_2\}$, where a_1 is assigned to three voters whose preferences are V , and a_2 is assigned to two voters whose preferences are V^* and one voter whose preferences are V . This means that the Monroe rule does not satisfy homogeneity and is not a GDSR. \square

7. FUTURE WORK

A natural (and ongoing) direction for future research is to design and deploy new application-specific social choice mechanisms under the GDSR framework, and further understand their statistical, computational, and axiomatic properties, including for example statistical minimaxity, (in)approximability, and other natural axiomatic properties such as monotonicity. How to incorporate GDSR framework and machine learning algorithms is also an important and promising direction.

REFERENCES

- ALTMAN, A. AND TENNENHOLTZ, M. 2010. An axiomatic approach to personalized ranking systems. *Journal of the ACM* 57, 4. Article 26.
- ARROW, K. J. 1950. A difficulty in the concept of social welfare. *Journal of Political Economy* 58, 4, 328–346.
- AZARI SOUFIANI, H., PARKES, D. C., AND XIA, L. 2012. Random utility theory for social choice. In *Proceedings of NIPS*. Lake Tahoe, NV, USA, 126–134.

- AZARI SOUFIANI, H., PARKES, D. C., AND XIA, L. 2013. Preference Elicitation For General Random Utility Models. In *Proceedings of UAI*. Bellevue, Washington, USA.
- AZARI SOUFIANI, H., PARKES, D. C., AND XIA, L. 2014. Statistical decision theory approaches to social choice. In *Proceedings of NIPS*. Montreal, Quebec, Canada.
- BARTHOLDI, III, J., TOVEY, C., AND TRICK, M. 1989a. The computational difficulty of manipulating an election. *Social Choice and Welfare* 6, 3, 227–241.
- BARTHOLDI, III, J., TOVEY, C., AND TRICK, M. 1989b. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare* 6, 157–165.
- 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.
- BETZLER, N., NIEDERMEIER, R., AND WOEGINGER, G. 2011. Unweighted coalitional manipulation under the Borda rule is NP-hard. In *Proceedings of IJCAI*. 55–60.
- BRANDT, F. 2009. Some remarks on Dodgson’s voting rule. *Mathematical Logic Quarterly* 55, 460–463.
- BRAVERMAN, M. AND MOSSEL, E. 2008. Noisy Sorting Without Resampling. In *Proceedings of SODA*. Philadelphia, PA, USA, 268–276.
- CARAGIANNIS, I., PROCACCIA, A., AND SHAH, N. 2013. When do noisy votes reveal the truth? In *Proceedings of EC*. Philadelphia, PA.
- CARY, D. 2011. Estimating the margin of victory for instant-runoff voting. In *Proceedings of 2011 EVT/WOTE Conference*.
- CHAMBERLIN, J. R. AND COURANT, P. N. 1983. Representative Deliberations and Representative Decisions: Proportional Representation and the Borda Rule. *The American Political Science Review* 77, 3, 718–733.
- CHIERICHETTI, F. AND KLEINBERG, J. 2012. Voting with Limited Information and Many Alternatives. In *Proceedings of SODA*. Kyoto, Japan, 1036–1055.
- CONDORCET, M. D. 1785. *Essai sur l’application de l’analyse à la probabilité des décisions rendues à la pluralité des voix*. Paris: L’Imprimerie Royale.
- CONITZER, V., ROGNLIE, M., AND XIA, L. 2009. Preference functions that score rankings and maximum likelihood estimation. In *Proceedings of IJCAI*. Pasadena, CA, USA, 109–115.
- CONITZER, V. AND SANDHOLM, T. 2005. Common voting rules as maximum likelihood estimators. In *Proceedings of UAI*. Edinburgh, UK, 145–152.
- 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.
- COOK, W. J., GERARDS, A. M. H., SCHRIJVER, A., AND TARDOS, E. 1986. Sensitivity theorems in integer linear programming. *Mathematical Programming* 34, 3, 251–264.
- DAVIES, J., KATSIRELOS, G., NARODYTSKA, N., WALSH, T., AND XIA, L. 2014. Complexity of and Algorithms for the Manipulation of Borda, Nanson and Baldwin’s Voting Rules. *Artificial Intelligence* 217, 20–42.
- DWORK, C., KUMAR, R., NAOR, M., AND SIVAKUMAR, D. 2001. Rank aggregation methods for the web. In *Proceedings of WWW*. 613–622.
- ELKIND, E., FALISZEWSKI, P., AND SLINKO, A. 2010. Good rationalizations of voting rules. In *Proceedings of AAI*. Atlanta, GA, USA, 774–779.
- 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. *Communications of the 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.
- FISHBURN, P. C. 1977. Condorcet social choice functions. *SIAM Journal on Applied Mathematics* 33, 3, 469–489.
- FRIEDGUT, E., KALAI, G., AND NISAN, N. 2008. Elections can be manipulated often. In *Proceedings of FOCS*. 243–249.
- GHOSH, S., MUNDHE, M., HERNANDEZ, K., AND SEN, S. 1999. Voting for movies: the anatomy of a recommender system. In *Proceedings of AGENTS*. 434–435.
- GREENE, W. H. 2011. *Econometric Analysis* 7th Ed. Prentice Hall.
- HEMASPAANDRA, E., HEMASPAANDRA, L. A., AND ROTHE, J. 1997. Exact analysis of dodgson elections: Lewis carroll’s 1876 voting system is complete for parallel access to np. *Journal of*

- the *ACM* 44, 806–825.
- HOEFFDING, W. 1963. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association* 58, 301, 13–30.
- HORN, R. A. AND JOHNSON, C. R. 1985. *Matrix Analysis*. Cambridge University Press.
- LENSTRA, H. W. 1983. Integer programming with a fixed number of variables. *Mathematics of Operations Research* 8, 538–548.
- LU, T. AND BOUTILIER, C. 2011. Learning mallows models with pairwise preferences. In *Proceedings of ICML*. Bellevue, WA, USA, 145–152.
- MAGRINO, T. R., RIVEST, R. L., SHEN, E., AND WAGNER, D. 2011. Computing the Margin of Victory in IRV Elections. In *Proceedings of 2011 EVT/WOTE Conference*.
- MALLOWS, C. L. 1957. Non-null ranking model. *Biometrika* 44, 1/2, 114–130.
- MAO, A., PROCACCIA, A. D., AND CHEN, Y. 2013. Better human computation through principled voting. In *Proceedings of AAAI*. Bellevue, WA, USA.
- MONROE, B. L. 1995. Fully Proportional Representation. *The American Political Science Review* 89, 4, 925–940.
- MOSSEL, E., PROCACCIA, A. D., AND RACZ, M. Z. 2013. A Smooth Transition From Powerlessness to Absolute Power. *Journal of Artificial Intelligence Research* 48, 1, 923–951.
- MOSSEL, E. AND RACZ, M. Z. 2012. A quantitative Gibbard-Satterthwaite theorem without neutrality. In *Proceedings of STOC*.
- NURMI, H. 1987. *Comparing voting systems*. Springer.
- OBRAZTSOVA, S. AND ELKIND, E. 2011. On the complexity of voting manipulation under randomized tie-breaking. In *Proceedings of IJCAI*. Barcelona, Catalonia, Spain, 319–324.
- PENNOCK, D. M., HORVITZ, E., AND GILES, C. L. 2000. Social choice theory and recommender systems: Analysis of the axiomatic foundations of collaborative filtering. In *Proceedings of AAAI*. Austin, TX, USA, 729–734.
- PORELLO, D. AND ENDRISS, U. 2013. Ontology Merging as Social Choice: Judgment Aggregation under the Open World Assumption. *Journal of Logic and Computation*.
- PRITCHARD, G. AND WILSON, M. C. 2009. Asymptotics of the minimum manipulating coalition size for positional voting rules under impartial culture behaviour. *Mathematical Social Sciences* 1, 35–57.
- PROCACCIA, A. D., REDDI, S. J., AND SHAH, N. 2012. A maximum likelihood approach for selecting sets of alternatives. In *Proceedings of UAI*.
- PROCACCIA, A. D. AND ROSENSCHEIN, J. S. 2007. Junta distributions and the average-case complexity of manipulating elections. *Journal of Artificial Intelligence Research* 28, 157–181.
- RAMAN, K. AND JOACHIMS, T. 2014. Methods for Ordinal Peer Grading. In *Proceedings of KDD*. New York, NY, USA, 1037–1046.
- ROTHER, J. AND SCHEND, L. 2013. Challenges to Complexity Shields That Are Supposed to Protect Elections Against Manipulation and Control: A Survey. *Annals of Mathematics and Artificial Intelligence* 68, 1–3, 161–193.
- XIA, L. 2012. Computing the margin of victory for various voting rules. In *Proceedings of EC*. Valencia, Spain, 982–999.
- XIA, L. 2013. Generalized Scoring Rules: A Framework That Reconciles Borda and Condorcet. *ACM SIGecom Exchanges* 12, 1, 42–48.
- XIA, L. AND CONITZER, V. 2008. Generalized scoring rules and the frequency of coalitional manipulability. In *Proceedings of EC*. Chicago, IL, USA, 109–118.
- XIA, L. AND CONITZER, V. 2009. Finite local consistency characterizes generalized scoring rules. In *Proceedings of IJCAI*. Pasadena, CA, USA, 336–341.
- XIA, L. AND CONITZER, V. 2011. A maximum likelihood approach towards aggregating partial orders. In *Proceedings of IJCAI*. Barcelona, Catalonia, Spain, 446–451.
- XIA, L., ZUCKERMAN, M., PROCACCIA, A. D., CONITZER, V., AND ROSENSCHEIN, J. 2009. Complexity of unweighted coalitional manipulation under some common voting rules. In *Proceedings of IJCAI*. Pasadena, CA, USA, 348–353.
- YOUNG, H. P. 1975. Social choice scoring functions. *SIAM Journal on Applied Mathematics* 28, 4, 824–838.
- ZUCKERMAN, M., PROCACCIA, A. D., AND ROSENSCHEIN, J. S. 2009. Algorithms for the coalitional manipulation problem. *Artificial Intelligence* 173, 2, 392–412.