Journal of Artificial Intelligence Research 48 (2013) 923-951 Submitted 7/13; published 12/13

A Smooth Transition from Powerlessness to Absolute Power

Elchanan Mossel MOSSEL@STAT.BERKELEY.EDU
University of California, Berkeley
Berkeley, CA 94720 USA
Weizmann Institute of Science
Rehovot, Israel

Ariel D. Procaccia ARIELPRO@CS.CMU.EDU
Carnegie Mellon University
Pittsburgh, PA 15213 USA

Miklós Z. Rácz RACZ@STAT.BERKELEY.EDU
University of California, Berkeley
Berkeley, CA 94720 USA

Abstract

We study the phase transition of the coalitional manipulation problem for generalized scoring rules. Previously it has been shown that, under some conditions on the distribution of votes, if the number of manipulators is $o\left(\sqrt{n}\right)$, where n is the number of voters, then the probability that a random profile is manipulable by the coalition goes to zero as the number of voters goes to infinity, whereas if the number of manipulators is $\omega\left(\sqrt{n}\right)$, then the probability that a random profile is manipulable goes to one. Here we consider the critical window, where a coalition has size $c\sqrt{n}$, and we show that as c goes from zero to infinity, the limiting probability that a random profile is manipulable goes from zero to one in a smooth fashion, i.e., there is a smooth phase transition between the two regimes. This result analytically validates recent empirical results, and suggests that deciding the coalitional manipulation problem may be of limited computational hardness in practice.

1. Introduction

Finding "good" voting systems which satisfy some natural requirements is one of the main goals in social choice theory. This problem is increasingly relevant in the area of artificial intelligence and in computer science more broadly, where virtual elections are now an established tool for preference aggregation (see, e.g., Caragiannis & Procaccia, 2011).

A naturally desirable property of a voting system is *strategyproofness* (a.k.a. nonmanipulability): no voter should benefit from voting strategically, i.e., voting not according to her true preferences. However, Gibbard (1973) and Satterthwaite (1975) showed that no reasonable voting system can be strategyproof. Before stating their result, let us specify the problem more formally.

We consider n voters electing a single winner among m candidates. The voters specify their opinion by ranking the candidates, and the winner is determined according to some predefined *social choice function* (SCF) $f : S_m^n \to [m]$ of all the voters' rankings, where S_m denotes the set of all possible total orderings of the m candidates. We call a collection of rankings by the voters a *ranking profile*. We say that a SCF is *manipulable* if there exists a ranking profile where a voter can achieve a more desirable outcome of the election according to her true preferences by voting in a way that does not reflect her true preferences.

The Gibbard-Satterthwaite theorem states that any SCF which is not a dictatorship (i.e., not a function of a single voter), and which allows at least three candidates to be elected, is manipulable. This has contributed to the realization that it is unlikely to expect truthfulness in voting. Consequently, there have been many branches of research devoted to understanding the extent of the manipulability of voting systems, and to finding ways of circumventing the negative results.

One approach, introduced by Bartholdi, Tovey, and Trick (1989), suggests computational complexity as a barrier against manipulation: a SCF may not be manipulable in practice if it is hard for a voter to compute a manipulative vote. A significant body of work focuses on the worst-case complexity of manipulation (see the survey by Faliszewski and Procaccia, 2010). Here we are interested specifically in the *coalitional manipulation* problem, where a group of voters can change their votes in unison, with the goal of making a given candidate win. Various variations of this problem are known to be $\mathcal{NP}$-hard under many of the common SCFs (Conitzer, Sandholm, & Lang, 2007; Xia, Zuckerman, Procaccia, Conitzer, & Rosenschein, 2009; Betzler, Niedermeier, & Woeginger, 2011).

Crucially, this line of work focuses on worst-case complexity. While worst-case hardness of manipulation is a desirable property for a SCF to have, it does not tell us much about *typical* instances of the problem—is it *usually* easy or hard to manipulate? A recent line of research on average-case manipulability has been questioning the validity of such worst-case complexity results. The goal of this alternative line of work is to show that there are no "reasonable" voting rules that are computationally hard to manipulate *on average*. Specifically, the goal is to rule out the following informal statement: there are "good" voting rules that are hard to manipulate on average under any "sufficiently rich" distribution over votes.

Taking this point of view, showing easiness of manipulation under a restricted class of distributions—such as i.i.d. votes or even uniform votes (the impartial culture assumption)—is interesting, even if these do not necessarily capture all possible real-world elections. Specifically, if we show that manipulation is easy under such distributions, then any average-case hardness result would necessarily have to make some unnatural technical assumptions to avoid these distributions. Studying such restricted distributions over votes is indeed exactly what some recent papers have done.

For the coalitional manipulation problem, Procaccia and Rosenschein (2007a) first suggested that it is trivial to determine whether manipulation is possible for most coalitional manipulation instances, from a typical-case computational point of view; one can make a highly informed guess purely based on the number of manipulators. Specifically, they studied a setting where there is a distribution over votes (which satisfies some conditions), and concentrated on a family of SCFs known as positional scoring rules. They showed that if the size of the coalition is $o\left(\sqrt{n}\right)$, then with probability converging to 1 as $n \to \infty$, the coalition is powerless, i.e., it cannot change the outcome of the election. In contrast, if the size of the coalition is $\omega\left(\sqrt{n}\right)$ (and $o\left(n\right)$), then with probability converging to 1 as $n \to \infty$, the coalition is all-powerful, i.e., it can elect any candidate. Later Xia and Conitzer (2008b) proved an analogous result for so-called generalized scoring rules, a family that contains almost all common voting rules. See also related work by Peleg (1979), Slinko (2004), Pritchard and Slinko (2006), and Pritchard and Wilson (2009). We discuss additional related work in Section 1.2.

Our primary interest in this paper is to understand the critical window that these papers leave open, when the size of the coalition is $\Theta(\sqrt{n})$. Specifically, we are interested in the *phase transition* in the probability of coalitional manipulation, when the size of the coalition is $c\sqrt{n}$ and c varies from zero to infinity, i.e., the *transition from powerlessness to absolute power*.

In the past few decades there has been much research on the connection between phase transitions and computationally hard problems (see, e.g., Fu & Anderson, 1986; Cheeseman, Kanefsky, & Taylor, 1991; Achlioptas, Naor, & Peres, 2005). In particular, it is often the case that the computationally hardest problems can be found at critical values of a sharp phase transition (see, e.g., Gomes & Walsh, 2006, for an overview). On the other hand, smooth phase transitions are often found in connection with computationally easy (polynomial) problems, such as 2-coloring (Achlioptas, 1999) and 1-in-2 SAT (Walsh, 2002). Thus understanding the phase transition in this critical window may shed light on where the computationally hardest problems lie.

Recently, Walsh (2011) *empirically* analyzed two well-known voting rules—veto and single transferable vote (STV)—and found that there is a smooth phase transition between the two regimes. Specifically, Walsh studied coalitional manipulation with unweighted votes for STV and weighted votes for veto, and sampled from a number of distributions in his experiments, including i.i.d. distributions, correlated distributions, and votes sampled from real-world elections. Our main result complements and improves upon Walsh's analysis in two ways; while Walsh's results show how the phase transition looks like concretely for veto and STV, we *analytically* show that the phase transition is indeed smooth for *any generalized scoring rule* (including veto and STV) when the votes are i.i.d. This suggests that deciding the coalitional manipulation problem may not be computationally hard in practice.

1.1 Our Results

We now present our results, but first let us formally specify the setup of the problem. We denote a ranking profile by $\sigma = (\sigma_1, \dots, \sigma_n) \in S_m^n$, and for a candidate a, define $W_a = \{\sigma \in S_m^n \mid f(\sigma) = a\}$, the set of ranking profiles where the outcome of f is a. Our setup and assumptions are the following.

Assumption 1. *We assume that the number of candidates, m, is constant.*

Assumption 2. *We assume that the SCF f is* anonymous*, i.e., it treats each voter equally.*

Assumption 3. *We assume that the votes of voters are i.i.d., according to some distribution p on S_m. Furthermore, we assume that there exists $\delta > 0$ such that for every $\pi \in S_m$, $p(\pi) \geq \delta$ (necessarily $\delta \leq 1/m!$).*

If we were to assume only these, then our setup would include uninteresting cases, such as when f is a constant—i.e., no matter what the votes are, a specific candidate wins. Another less interesting case is when the probability of a given candidate winning vanishes as $n \to \infty$—we can then essentially forget about this candidate for large n (in the sense that a coalition of size $\Omega(n)$ would be necessary to make this candidate win). To exclude these and focus on the interesting cases, we make an additional assumption which concerns both the SCF and the distribution of votes.

Assumption 4. *We assume that there exists $\varepsilon > 0$ such that for every n and for every candidate $a \in [m]$, the probability of a being elected is at least $\varepsilon > 0$, i.e., $\mathbb{P}(W_a) \geq \varepsilon$ (necessarily $\varepsilon \leq 1/m$).*

All four assumptions are satisfied when the distribution is uniform (i.e., under the impartial culture assumption) and the SCF is close to being neutral (i.e., neutral up to some tie-breaking rules); in particular, they hold for all commonly used SCFs. The assumptions are somewhat more general than this, although the i.i.d. assumption remains a restrictive one. However, as discussed earlier, even showing easiness of manipulation under such a restricted class of distributions is interesting.

As mentioned before, we are interested in the case when the coalition has size $c\sqrt{n}$ for some constant c. Define the probabilities

$$\begin{aligned}
\underline{q}_n(c) &:= \mathbb{P}\left(\text{ some coalition of size } c\sqrt{n} \text{ can elect any candidate }\right),\\
\overline{q}_n(c) &:= \mathbb{P}\left(\text{ some coalition of size } c\sqrt{n} \text{ can change the outcome of the election }\right),\\
\underline{r}_n(c) &:= \mathbb{P}\left(\text{ a specific coalition of size } c\sqrt{n} \text{ can elect any candidate }\right),\\
\overline{r}_n(c) &:= \mathbb{P}\left(\text{ a specific coalition of size } c\sqrt{n} \text{ can change the outcome of the election }\right),
\end{aligned}$$

and let

$$\underline{q}(c) := \lim_{n\to\infty} \underline{q}_n(c), \quad \overline{q}(c) := \lim_{n\to\infty} \overline{q}_n(c), \quad \underline{r}(c) := \lim_{n\to\infty} \underline{r}_n(c), \quad \overline{r}(c) := \lim_{n\to\infty} \overline{r}_n(c),$$

provided these limits exist. Clearly $\underline{q}_n(c) \leq \overline{q}_n(c)$, $\underline{r}_n(c) \leq \overline{r}_n(c)$, $\underline{r}_n(c) \leq \underline{q}_n(c)$, and $\overline{r}_n(c) \leq \overline{q}_n(c)$.

Before we describe our results, which deal with these quantities, we first explain how these relate to the various variants of the coalitional manipulation problem. In the coalitional manipulation problem the coalition is fixed, and thus the relevant quantities are $\underline{r}_n(c)$ and $\overline{r}_n(c)$. Closely related is the problem of determining the *margin of victory*, which is the minimum number of voters who need to change their votes to change the outcome of the election. Also related is the problem of *bribery*, the minimum number of voters who need to change their votes to make a given candidate win. The main difference between these problems is that in coalitional manipulation the coalition is fixed, whereas in the latter two problems the coalition is not fixed. Hence the relevant quantities for studying the latter two are $\underline{q}_n(c)$ and $\overline{q}_n(c)$. Our tools also allow us to deal with other related quantities (such as *microbribery*, Faliszewski, Hemaspaandra, Hemaspaandra, & Rothe, 2009), but we focus our attention on the four quantities described above.

Our first result analyzes the case when the size of the coalition is $c\sqrt{n}$ for large c. We show that if c is large enough, then with probability close to 1, a specific coalition of size $c\sqrt{n}$ can elect any candidate. This holds for any SCF that satisfies the above (mild) restrictions.

Theorem 1.1. *Assume that Assumptions 1, 2, 3, and 4 hold. For any $\eta > 0$ there exists a constant $c = c(\eta, \delta, \varepsilon, m)$ such that $\underline{r}_n(c) \geq 1 - \eta$ for every n. In particular, we can choose*

$$c = (4/\delta)\log(2m!/\eta)\left[\sqrt{\log(2m/\eta)} + \sqrt{\log(2/\varepsilon)}\right].$$

It follows that

$$\lim_{c\to\infty} \liminf_{n} \underline{r}_n (c) = 1.$$

This result extends previous theorems of Procaccia and Rosenschein (2007a), and Xia and Conitzer (2008b), from scoring rules and generalized scoring rules, respectively, to anonymous SCFs.

Our second result deals with the case when the size of the coalition is $c\sqrt{n}$ for small c, and the transition as c goes from 0 to ∞. Here we assume additionally that f is a generalized scoring rule (to be defined in Section 3.1.1); this is needed because there exist (pathological) anonymous SCFs for which the result below does not hold (see the beginning of Section 3 for an example).

Theorem 1.2. *Assume that Assumptions 1, 2, 3, and 4 hold, and furthermore that f is a generalized scoring rule. Then:*

(1) The limits $\underline{q}(c)$, $\overline{q}(c)$, $\underline{r}(c)$ and $\overline{r}(c)$ exist.

(2) There exists a constant $K = K(f, \delta) < \infty$ such that $\overline{q}(c) \leq Kc$; in particular, $\lim_{c\to 0} \overline{q}(c) = 0$.

(3) For all $0 < c < \infty$, $0 < \underline{q}(c) \leq \overline{q}(c) < 1$ and $0 < \underline{r}(c) \leq \overline{r}(c) < 1$, and furthermore $\underline{q}(c)$, $\overline{q}(c)$, $\underline{r}(c)$ and $\overline{r}(c)$ are all continuously differentiable in c with bounded derivative.

In words, Part 2 means that if c is small enough then with probability close to 1 no coalition of size $c\sqrt{n}$ can change the outcome of the election, and the statements about $\overline{r}$ and $\underline{r}$ in Part 3 mean that the coalitional manipulation problem has a smooth phase transition: as the number of manipulators increases, the probabilities that a coalition has some power, and that it has absolute power, increase smoothly. Parts 1 and 2 of the theorem simply make a result of Xia and Conitzer (2008b) more precise, by extending the analysis to the $\Theta(\sqrt{n})$ regime. More importantly, in the proofs of these statements we introduce the machinery needed to establish Part 3, which is our main result.

Since the coalitional manipulation problem does *not* have a sharp phase transition, Theorem 1.2 can be interpreted as suggesting that realistic distributions over votes are likely to yield coalitional manipulation instances that are tractable in practice, even if the size of the coalition concentrates on the previously elusive $\Theta(\sqrt{n})$ regime; this is true for any generalized scoring rule, and in particular for almost all common social choice functions (an exception is Dodgson's rule). This interpretation has a negative flavor in further strengthening the conclusion that worst-case complexity is a poor barrier to manipulation.

However, the complexity glass is in fact only half empty. The probability that the margin of victory is at most $c\sqrt{n}$ is captured by the quantity $\overline{q}_n$, hence Part 3 of Theorem 1.2 also implies that the margin of victory problem has a smooth phase transition. As recently pointed out by Xia (2012a), efficiently solving the margin of victory problem could help in *post-election audits*—used to determine whether electronic elections have resulted in an incorrect outcome due to software or hardware bugs—and its tractability is in fact desirable.

The methods we use are flexible, and can be extended to various setups of interest that do not directly satisfy our assumptions above, for instance single-peaked preferences. Consider a one-dimensional political spectrum represented by the interval $[0, 1]$, and fix the location of the candidates. Assume voters are uniformly distributed on the interval, independently of each other. For technical reasons, this distribution does not satisfy our assumptions, since there will be rankings $\pi \in S_m$ such that $p(\pi) = 0$; however, our tools allow us to deal with this setting as well. For instance, if the locations of the m candidates are $\left\{\frac{1}{2m}, \frac{3}{2m}, \dots, \frac{2m-1}{2m}\right\}$, then our results hold (with appropriate quantitative modifications). Similarly, if the locations were something else, then there would exist a subset of candidates who have an asymptotically nonvanishing probability of winning, and the same results hold restricted to this subset of candidates.

Finally, we discuss the role of tie-breaking in our setup, since this is often an important issue when studying manipulation. However, since we consider manipulation by coalitions of size $c\sqrt{n}$, ties where there exist a constant number of voters such that if their votes are changed appropriately there is no longer a tie, are not relevant. Indeed, our tools allow us to extend the results of Theorem 1.2 to a class of SCFs slightly beyond generalized scoring rules, and, in particular, these allow for arbitrary tie-breaking rules (see Section 3.2.1 for details).

1.2 Additional Related Work

A recent line of research with an average-case algorithmic flavor also suggests that manipulation is indeed typically easy; see, e.g., the work of Kelly (1993), Conitzer and Sandholm (2006), Procaccia and Rosenschein (2007b), and Zuckerman et al. (2009) for results on certain restricted classes of SCFs. A different approach, initiated by Friedgut, Kalai, Keller and Nisan (2008, 2011), who studied the fraction of ranking profiles that are manipulable, also suggests that manipulation is easy on average; see further the work of Xia and Conitzer (2008a), Dobzinski and Procaccia (2008), Isaksson, Kindler and Mossel (2012), and Mossel and Rácz (2012). We refer to the survey by Faliszewski and Procaccia (2010) for a detailed history of the surrounding literature. See also related literature in economics, e.g., the work of Good and Mayer (1975), Chamberlain and Rothschild (1981), and Myatt (2007).

Recent work by Xia (2012a) is independent from, and closely related to, our work. As mentioned above, Xia's paper is concerned with computing the margin of victory in elections. He focuses on computational complexity questions and approximation algorithms, but one of his results is similar to Parts 1 and 2 of Theorem 1.2. However, our analysis is completely different; our approach facilitates the proof of Part 3 of the theorem, which is our main contribution. An even more recent (and also independent) manuscript by Xia (2012b) considers similar questions for generalized scoring rules and captures additional types of strategic behavior (such as control), but again, crucially, this work does not attempt to understand the phase transition (nor does it subsume our Theorem 1.1).

2. Large Coalitions

Without further ado, we prove Theorem 1.1. The main idea is to observe that for i.i.d. distributions, the Hamming distance of a random ranking profile from a fixed subset of ranking profiles concentrates around its mean. The theorem follows from standard concentration inequalities.

Proof of Theorem 1.1. For $\sigma, \sigma' \in S_m^n$, define

$$d\left(\sigma, \sigma'\right) = \frac{1}{n} \sum_{i=1}^{n} \mathbf{1}\left[\sigma_i \neq \sigma_i'\right],$$

i.e., $d(\sigma, \sigma')$ is $1/n$ times the Hamming distance of σ and σ'. If U is a subset of ranking profiles and σ is a specific ranking profile then define $d_U(\sigma) = \min_{\sigma' \in U} d(\sigma, \sigma')$. The function d_U is Lipschitz with constant $1/n$, and therefore by McDiarmid's inequality we have the following concentration inequality:

$$\mathbb{P}\left(\left|d_U(\sigma) - \mathbb{E} d_U\right| \geq c\right) \leq 2 \exp\left(-2c^2 n\right) \tag{1}$$

for any $c > 0$ and $U \subseteq S_m^n$. Suppose $U \subseteq S_m^n$ has measure at least ε, i.e., U is such that $\mathbb{P}(\sigma \in U) \geq \varepsilon$, and take γ such that $2\exp\left(-2\gamma^2 n\right) < \varepsilon$, e.g., let $\gamma = \sqrt{\log(2/\varepsilon)}/\sqrt{n}$. Then (1) implies that there exists $\sigma \in U$ such that $\left|d_U(\sigma) - \mathbb{E} d_U\right| \leq \gamma$, but since $d_U(\sigma) = 0$, this means that $\mathbb{E} d_U \leq \gamma$. So for such a set U, we have

$$\mathbb{P}\left(d_U(\sigma) > \gamma + c\right) \leq \exp\left(-2c^2 n\right)$$

for any $c > 0$. Choosing $c = B/\sqrt{n}$ and defining $B' = B + \sqrt{\log(2/\varepsilon)}$ we get that

$$\mathbb{P}\left(d_U(\sigma) > B'/\sqrt{n}\right) \leq \exp\left(-2B^2\right). \tag{2}$$

In the language of the usual Hamming distance, this means that the probability that the ranking profile needs to be changed in at least $B'\sqrt{n}$ coordinates to be in U is at most $\exp\left(-2B^2\right)$, which can be made arbitrarily small by choosing B large enough.

By our assumption, $\mathbb{P}(\sigma \in W_a) \geq \varepsilon$ for every a, and therefore by (2) and a union bound we get

$$\mathbb{P}\left(\exists a : d_{W_a}(\sigma) > B'/\sqrt{n}\right) \leq m \exp\left(-2B^2\right).$$

By choosing $B = \sqrt{\log(2m/\eta)}$, this probability is at most $\eta/2$.

Consider a specific coalition of size $DB'\sqrt{n}$, where $D = D(\delta, m)$ will be chosen later. Using Chernoff's bound and a union bound, with probability close to one, for every possible ranking π the coalition has at least $B'\sqrt{n}$ voters with the ranking π:

$$\begin{aligned}
&\mathbb{P}\left(\exists \pi \in S_m : \text{ coalition of size } DB'\sqrt{n} \text{ has less than } B'\sqrt{n} \text{ voters with ranking } \pi\right) \\
&\qquad \leq m!\mathbb{P}\left(\operatorname{Bin}\left(DB'\sqrt{n}, \delta\right) < B'\sqrt{n}\right) \leq m! \exp\left(-\left(1 - 1/D\delta\right)^2 DB'\sqrt{n}\delta/2\right) \\
&\qquad \leq m! \exp\left(-\left(1 - 1/D\delta\right)^2 D\delta/2\right),
\end{aligned}$$

where $\operatorname{Bin}(DB'\sqrt{n}, \delta)$ denotes a binomial random variable with parameters $DB'\sqrt{n}$ and δ, and where we used our assumption that for every voter the probability for every ranking is at least $\delta > 0$. Choosing $D = (4/\delta)\log(2m!/\eta)$, this probability is at most $\eta/2$.

By the anonymity of f, the outcome only depends on the number of voters voting according to each ranking. Consequently, if σ is such that it is at a distance of at most $B'/\sqrt{n}$ away from each W_a, and where for each ranking π there are at least $B'\sqrt{n}$ voters in the coalition with ranking π, then the coalition is able to achieve any outcome. Using the above and a union bound this happens with probability at least $1 - \eta$. $\square$

3. Small Coalitions and the Phase Transition

This section is almost entirely devoted to the proof of Theorem 1.2, but it also includes some helpful definitions, examples, and intuitions.

Consider the following example of a SCF. For $a \in [m]$ let $n_a(\sigma)$ denote the number of voters who ranked candidate a on top in the ranking profile σ. Define the SCF f by $f(\sigma) = \sum_{a=1}^{m} a n_a(\sigma) \mod m$. This SCF is clearly anonymous (since it only depends on the number of voters voting according to specific rankings), and moreover it is easy to see that *any* single voter can *always* elect any candidate.

This example shows that, in general, we cannot have a matching lower bound for the size of the manipulating coalition on the order of $\sqrt{n}$. However, this is an artificial example (one would not consider such a voting system in real life), and we expect that a matching lower bound holds for most reasonable SCFs.

Xia and Conitzer (2008b) introduced a large class of SCFs called generalized scoring rules, which include most commonly occurring SCFs. In the following we introduce this class of SCFs, provide an alternative way of looking at them (as so-called "hyperplane rules"), and show that for this class of SCFs if the coalition has size $c\sqrt{n}$ for small enough c, then the probability of being able to change the outcome of the election can be arbitrarily close to zero. At the end of the section we then prove the smooth transition as stated in Part 3 of Theorem 1.2.

3.1 Generalized Scoring Rules, Hyperplane Rules, and their Equivalence

In this section we introduce generalized scoring rules and hyperplane rules and show their equivalence.

3.1.1 Generalized Scoring Rules

We now define generalized scoring rules.

Definition 1. *For any $y, z \in \mathbb{R}^k$, we say that y and z are* equivalent*, denoted by $y \sim z$, if for every $i, j \in [k]$, $y_i \geq y_j$ if and only if $z_i \geq z_j$.*

Definition 2. *A function $g : \mathbb{R}^k \to [m]$ is* compatible *if for any $y \sim z$, $g(y) = g(z)$.*

That is, for any function g that is compatible, $g(y)$ is completely determined by the total preorder of $\{y_1, \dots, y_k\}$ (a total preorder is an ordering in which ties are allowed).

Definition 3 (Generalized scoring rules)**.** *Let $k \in \mathbb{N}$, $f : S_m \to \mathbb{R}^k$ (called a* generalized scoring function*), and $g : \mathbb{R}^k \to [m]$ where g is compatible (g is called a* decision function*). The functions f and g determine the* generalized scoring rule $\mathrm{GS}(f, g)$ *as follows: for $\sigma \in S_m^n$, let*

$$\mathrm{GS}(f, g)(\sigma) := g\left(\sum_{i=1}^{n} f(\sigma_i)\right).$$

From the definition it is clear that every generalized scoring rule (GSR) is anonymous.

3.1.2 Hyperplane Rules

Preliminaries and notation. In the following, for a SCF let us write $f \equiv f_n$, i.e., let us explicitly note that f is a function on n voters; also let us write $\sigma \equiv \sigma^n$. Since the SCF f_n is anonymous, the outcome only depends on the numbers of voters who vote according to particular rankings. Let D_n denote the set of points in the probability simplex $\Delta^{m!}$ for which all coordinates are integer multiples of $1/n$. Let us denote a typical element of the probability simplex $\Delta^{m!}$ by $x = \{x_\pi\}_{\pi \in S_m}$. For a ranking profile σ^n, let us denote the corresponding element of the probability simplex by $x(\sigma^n)$, i.e., for all $\pi \in S_m$,

$$x(\sigma^n)_\pi = \frac{1}{n}\sum_{i=1}^{n} \mathbf{1}[\sigma_i = \pi].$$

By our assumptions the outcome of f_n only depends on $x(\sigma^n)$, so by abuse of notation we may write that $f_n : \Delta^{m!}|_{D_n} \to [m]$ with $f_n(\sigma^n) = f_n(x(\sigma^n))$.

We are now ready to define hyperplane rules.

Definition 4 (Hyperplane rules)**.** *Fix a finite set of affine hyperplanes of the simplex $\Delta^{m!}$: $H_1, \ldots, H_\ell$. Each affine hyperplane partitions the simplex into three parts: the affine hyperplane itself and two open halfspaces on either side of the affine hyperplane. Thus the affine hyperplanes $H_1, \ldots, H_\ell$ partition the simplex into finitely many (at most 3^ℓ) regions. Let $F : \Delta^{m!} \to [m]$ be a function which is constant on each such region. Then the sequence of SCFs $\{f_n\}_{n\geq 1}$, $f_n : S_m^n \to [m]$, defined by*

$$f_n(\sigma^n) = F(x(\sigma^n))$$

is called a hyperplane rule induced by the affine hyperplanes $H_1, \ldots, H_\ell$ and the function F.

A function $F : \Delta^{m!} \to [m]$ naturally partitions the simplex $\Delta^{m!}$ into m parts based on the outcome of F. (For hyperplane rules this partition is coarser than the partition of $\Delta^{m!}$ induced by the affine hyperplanes $H_1, \ldots, H_\ell$.) We abuse notation and denote these parts by $\{W_a\}_{a\in[m]}$. The following definition will be useful for us.

Definition 5 (Interior and boundaries of a partition induced by F)**.** *We say that $x \in \Delta^{m!}$ is an interior point of the partition $\{W_a\}_{a\in[m]}$ induced by F if there exists $\alpha > 0$ such that for all $y \in \Delta^{m!}$ for which $|x - y| \leq \alpha$, we have $F(x) = F(y)$. Otherwise, we say that $x \in \Delta^{m!}$ is on the boundary of the partition, which we denote by B.*

For a hyperplane rule the boundary B is contained in the union of the corresponding affine hyperplanes. Conversely, suppose $F : \Delta^{m!} \to [m]$ is an arbitrary function and the sequence of (anonymous) SCFs $\{f_n\}_{n\geq 1}$, $f_n : S_m^n \to [m]$ is defined by $f_n(\sigma^n) = F(x(\sigma^n))$. If the boundary B of F is contained in the union of finitely many affine hyperplanes of $\Delta^{m!}$, then F is not necessarily a hyperplane rule, but there exists a hyperplane rule $\hat{F}$ such that F and $\hat{F}$ agree everywhere except perhaps on the union of the finitely many affine hyperplanes.

3.1.3 Equivalence

Xia and Conitzer (2009) gave a characterization of generalized scoring rules: a SCF is a generalized scoring rule if and only if it is anonymous and finitely locally consistent (see Xia & Conitzer, 2009, Definition 5). This characterization is related to saying that generalized scoring rules are the same as hyperplane rules, yet we believe that spelling this out explicitly is important, because the geometric viewpoint of hyperplane rules is somewhat different, and in this probabilistic context it is also more flexible.

Lemma 3.1. *The class of generalized scoring rules coincides with the class of hyperplane rules.*

Proof. First let us show that every hyperplane rule is a generalized scoring rule. Let us consider the hyperplane rule induced by affine hyperplanes $H_1, \ldots, H_\ell$ of the simplex $\Delta^{m!}$, and the function $F : \Delta^{m!} \to [m]$. The affine hyperplanes of $\Delta^{m!}$ can be thought of as hyperplanes of $\mathbb{R}^{m!}$ that go through the origin—abusing notation we also denote these by $H_1, \ldots, H_\ell$. Let $u_1, \ldots, u_\ell$ denote unit normal vectors of these hyperplanes.

We need to define functions f and g such that for every ranking profile $\sigma^n \in S_m^n$, $\mathrm{GS}(f,g)(\sigma^n) = F(x(\sigma^n))$. We will have $f : S_m \to \mathbb{R}^{\ell+1}$ and $g : \mathbb{R}^{\ell+1} \to [m]$. Coordinates $1, \ldots, \ell$ of f correspond to hyperplanes $H_1, \ldots, H_\ell$, while the last coordinate of f will always be 0 (this is a technical necessity to make sure that the function g is compatible). Let us look at the coordinate corresponding to hyperplane H_j with normal vector u_j. For $\pi \in S_m$ define

$$(f(\pi))_j \equiv (f(\pi))_{H_j} \equiv (f(\pi))_{u_j} := (u_j)_\pi ,$$

where the coordinates of $\mathbb{R}^{m!}$ are indexed by elements of S_m. Then

$$(f(\sigma^n))_j := \sum_{i=1}^{n} (f(\sigma_i))_j = \sum_{i=1}^{n} (u_j)_{\sigma_i} = n\,(u_j \cdot x(\sigma^n)) .$$

The sign of $(f(\sigma^n))_j$ thus tells us which side of the hyperplane H_j the point $x(\sigma^n)$ lies on. We define $g(y)$ for all $y \in \mathbb{R}^{\ell+1}$ such that $y_{\ell+1} = 0$; then the requirement that g be compatible defines g for all $y \in \mathbb{R}^{\ell+1}$. For $x \in \mathbb{R}$, define $\mathrm{sgn}(x)$ to be 1 if $x > 0$, -1 if $x < 0$, and 0 if $x = 0$.

To define $g(y_1, \ldots, y_\ell, 0)$, look at the vector $(\mathrm{sgn}(y_1), \ldots, \mathrm{sgn}(y_\ell))$. This vector determines a region in $\Delta^{m!}$ in the following way: if $\mathrm{sgn}(y_j) = 1$, then the region lies in the same open halfspace as u_j, if $\mathrm{sgn}(y_j) = -1$ then the region lies in the open halfspace which does not contain u_j, and finally if $y_j = 0$, then the region lies in the hyperplane H_j. Now we define $g(y_1, \ldots, y_\ell, 0)$ to be the value of F on the region of $\Delta^{m!}$ defined by $(\mathrm{sgn}(y_1), \ldots, \mathrm{sgn}(y_\ell))$. The value of $g(y_1, \ldots, y_\ell, 0)$ is well-defined since F is constant in each such region. Moreover, if we take $y \sim z$ with $y_{\ell+1} = z_{\ell+1} = 0$, then necessarily $(\mathrm{sgn}(y_1), \ldots, \mathrm{sgn}(y_\ell)) = (\mathrm{sgn}(z_1), \ldots, \ sgn(z_\ell))$, and thus $g(y) = g(z)$: so g is compatible (this is where we used the extra coordinate).

Now let us show that every generalized scoring rule is a hyperplane rule. Suppose a generalized scoring rule is given by functions $f : S_m \to \mathbb{R}^k$ and $g : \mathbb{R}^k \to [m]$. For a ranking profile $\sigma^n \in S_m^n$, define $f(\sigma^n) := \sum_{i=1}^n f(\sigma_i) = n \sum_{\pi \in S_m} f(\pi)\,(x(\sigma^n))_\pi$; in this way we can view f as a function mapping $\mathbb{N}_{\geq 0}^{m!} \setminus \{0\}$ to $\mathbb{R}^k$ (and hence can also view $\mathrm{GS}(f,g)$ as a function mapping $\mathbb{N}_{\geq 0}^{m!} \setminus \{0\}$ to $[m]$). Since this mapping is homogeneous, we may extend the domain of f (and hence that of $\mathrm{GS}(f,g)$) to $\mathbb{Q}_{\geq 0}^{m!} \setminus \{0\}$ in the natural way.

For a total preorder $\mathcal{O}$, let $R_{\mathcal{O}} = \left\{x \in \mathbb{Q}_{\geq 0}^{m!} \setminus \{0\} : f(x) \sim \mathcal{O}\right\}$. By definition, if $x, y \in R_{\mathcal{O}}$ then $g(f(x)) = g(f(y))$, i.e., $\mathrm{GS}(f,g)$ is constant in each region $R_{\mathcal{O}}$. Each region $R_{\mathcal{O}}$ is a $\mathbb{Q}$-convex cone, i.e. if $x, y \in R_{\mathcal{O}}$ and $\lambda \in \mathbb{Q} \cap [0,1]$, then $\lambda x + (1-\lambda) y \in R_{\mathcal{O}}$, and furthermore if $\mu \in \mathbb{Q}_{>0}$, then $\mu x \in R_{\mathcal{O}}$ (both of these properties follow immediately from Definition 1). Thus we can write $\mathbb{Q}_{\geq 0}^{m!} \setminus \{0\}$ as the disjoint union of the $\mathbb{Q}$-convex cones $\{R_{\mathcal{O}}\}_{\mathcal{O}}$. The only way to do this is by taking finitely many hyperplanes of $\mathbb{R}^{m!}$ and cutting $\mathbb{Q}_{\geq 0}^{m!} \setminus \{0\}$ using these hyperplanes; a precise statement of this can be found in Appendix A. This essentially follows from a result by Kemperman (1986, Thm. 2)—to keep the paper self-contained we reproduce in Appendix A his results and proof, and show how the statement above follows from his results. Since our function is homogeneous, we need only look at the values of $\mathrm{GS}(f,g)$ on the simplex $\Delta^{m!}$. By the above, the simplex is divided into regions $\left\{R_{\mathcal{O}} \cap \Delta^{m!}\right\}_{\mathcal{O}}$ via affine hyperplanes of $\Delta^{m!}$, and the function $\mathrm{GS}(f,g)$ is constant on $R_{\mathcal{O}} \cap \Delta^{m!}$ for each total preorder $\mathcal{O}$, so $\mathrm{GS}(f,g)$ is indeed a hyperplane rule. $\square$

3.1.4 Examples

Most commonly used SCFs are generalized scoring rules / hyperplane rules, including all positional scoring rules, instant-runoff voting, Coombs' method, contingent vote, the Kemény-Young method, Bucklin voting, Nanson's method, Baldwin's method, Copeland's method, maximin, and ranked pairs. Some of these examples were already shown by Xia and Conitzer (2008b, 2009), but nevertheless in Appendix B we detail explanations of many of these examples. The main reason for this is that the perspective of a hyperplane rule arguably makes these explanations simpler and clearer. A rule that does not fit into this framework is Dodgson's rule, which is not homogeneous (see, e.g., Brandt, 2009), and therefore it is not a hyperplane rule.

3.2 Small Coalitions for Generalized Scoring Rules

We now show that for generalized scoring rules, a coalition of size $c\sqrt{n}$ for small enough c can only change the outcome of the election with small probability. By the equivalence above, we can work in the framework of hyperplane rules.

We consider two metrics on $\Delta^{m!}$: the L^1 metric, denoted by d_1 or $\|\cdot\|_1$, and the L^2 metric, denoted by d_2 or $\|\cdot\|_2$. The L^1 metric is important in this setting, since changing the votes of voters corresponds to moving in the L^1 metric on $\Delta^{m!}$; this connection is formalized in the following lemma.

Lemma 3.2. *Let $\sigma^n, \tau^n \in S_m^n$. Then $d_1(x(\sigma^n), x(\tau^n)) \leq \frac{2}{n} d_H(\sigma^n, \tau^n)$, where d_H denotes Hamming distance, i.e., $d_H(\sigma^n, \tau^n) = \sum_{i=1}^n \mathbf{1}[\sigma_i \neq \tau_i]$. Furthermore, if $y \in D_n$, then there exists $\hat{\tau}^n \in S_m^n$ such that $x(\hat{\tau}^n) = y$ and $d_1(x(\sigma^n), y) = \frac{2}{n} d_H(\sigma^n, \hat{\tau}^n)$.*

Proof. Let $\pi^0 = \sigma^n$, and for $i = 1, \ldots, n$, define the ranking profile π^i as

$$\pi^i = (\tau_1, \ldots, \tau_i, \sigma_{i+1}, \ldots, \sigma_n).$$

By definition, $\pi^n = \tau^n$. The desired inequality then follows from the triangle inequality:

$$d_1\left(x\left(\sigma^n\right), x\left(\tau^n\right)\right) = d_1\left(x\left(\pi^0\right), x\left(\pi^n\right)\right) \leq \sum_{i=1}^{n} d_1\left(x\left(\pi^{i-1}\right), x\left(\pi^i\right)\right)$$
$$= \sum_{i=1}^{n} \frac{2}{n} \mathbf{1}\left[\sigma_i \neq \tau_i\right] = \frac{2}{n} d_H\left(\sigma^n, \tau^n\right).$$

For the second part of the lemma, construct $\hat{\tau}^n$ as follows. For each $\pi \in S_m$, let $I_\pi := \{i \in [n] : \sigma_i = \pi\}$. If $x\left(\sigma^n\right)_\pi \leq y_\pi$, then for every $i \in I_\pi$, let $\hat{\tau}_i = \pi$. If $x\left(\sigma^n\right)_\pi > y_\pi$, then choose a subset of indices $I'_\pi \subset I_\pi$ of size $|I'_\pi| = n y_\pi$, and for every $i \in I'_\pi$, let $\hat{\tau}_i = \pi$. Finally, define the rest of the coordinates of $\hat{\tau}^n$ so that $x\left(\hat{\tau}^n\right) = y$. The construction guarantees that then $d_1\left(x\left(\sigma^n\right), y\right) = \frac{2}{n} d_H\left(\sigma^n, \hat{\tau}^n\right)$. □

It is therefore natural to define distances from the boundary B using the L^1 metric:

Definition 6 (Blowup of boundary)**.** *For $\alpha > 0$, we define the blowup of the boundary B by α to be*

$$B^{+\alpha} = \left\{ y \in \Delta^{m!} : \exists x \in B \text{ such that } \left\|x - y\right\|_1 \leq \alpha \right\}.$$

In order for some coalition to be able to change the outcome of the election at a given ranking profile, the point on the simplex corresponding to this ranking profile needs to be sufficiently close to the boundary B; this is formulated in the following lemma.

Lemma 3.3. *Suppose we have n voters, a coalition of size k, and the ranking profile is $\sigma^n \in S_m^n$, which corresponds to the point $x\left(\sigma^n\right) \in \Delta^{m!}$ on the probability simplex. A necessary condition for the coalition to be able to change the outcome of the election from this position is that $x\left(\sigma^n\right) \in B^{+2k/n}$. Conversely, if $x\left(\sigma^n\right) \in B^{+(2k-m!)/n}$, then there exists a coalition of size k that can change the outcome of the election.*

Proof. For any ranking profile τ^n that the coalition can reach, we have $d_H\left(\sigma^n, \tau^n\right) \leq k$, and so by Lemma 3.2 we have $d_1\left(x\left(\sigma^n\right), x\left(\tau^n\right)\right) \leq \frac{2k}{n}$. If $x\left(\sigma^n\right) \notin B^{+2k/n}$, then for every ranking profile τ^n which the coalition can reach, $x\left(\sigma^n\right)$ and $x\left(\tau^n\right)$ are in the same region determined by the hyperplanes, and so $F\left(x\left(\tau^n\right)\right) = F\left(x\left(\sigma^n\right)\right)$, i.e., the coalition cannot change the outcome of the election.

Now suppose that $x\left(\sigma^n\right) \in B^{+(2k-m!)/n}$. Then by definition there exists $y \in B$ such that $d_1\left(x\left(\sigma^n\right), y\right) \leq \frac{2k-m!}{n}$. Since $y \in B$, there exists $\hat{y} \in D_n$ such that $d_1\left(y, \hat{y}\right) \leq \frac{m!}{n}$ and $F\left(\hat{y}\right) \neq F\left(x\left(\sigma^n\right)\right)$. By the triangle inequality, $d_1\left(x\left(\sigma^n\right), \hat{y}\right) \leq \frac{2k}{n}$, and then by the second part of Lemma 3.2 there exists $\hat{\tau}^n \in S_m^n$ such that $x\left(\hat{\tau}^n\right) = \hat{y}$ and $d_H\left(\sigma^n, \hat{\tau}^n\right) \leq k$. The coalition consisting of voters with indices in $I := \{i \in [n] : \sigma_i \neq \hat{\tau}_i\}$ can thus change the outcome of the election. □

Corollary 3.4. *If we have n voters, the probability that some coalition of size k can change the outcome of the election is bounded from below by $\mathbb{P}\left(x\left(\sigma^n\right) \in B^{+(2k-m!)/n}\right)$ and from above by $\mathbb{P}\left(x\left(\sigma^n\right) \in B^{+2k/n}\right)$, where σ^n is drawn according to the probability distribution satisfying the conditions of the setup.*

Gaussian limit. Due to the i.i.d.-ness of the votes, the multinomial random variable $x\left(\sigma^{n}\right)$ concentrates around its expectation, and the rescaled random variable

$$\tilde{x}\left(\sigma^{n}\right):=\sqrt{n}\left(x\left(\sigma^{n}\right)-\mathbb{E}\left(x\left(\sigma^{n}\right)\right)\right)$$

converges to a normal distribution, with zero mean and specific covariance structure. For our analysis it is better to use this Gaussian picture, and thus we will reformulate the preliminaries above in this limiting setting. First, let us determine the limiting distribution.

Lemma 3.5. *We have* $\tilde{x}\left(\sigma^{n}\right) \Rightarrow_{n} N(0, \Sigma)$*, where the covariance structure is given by* $\Sigma=\operatorname{diag}(p)-p p^{T}$*, where recall that* p *is the distribution of a vote.*

Proof. It is clear that $\mathbb{E}\left(\tilde{x}\left(\sigma^{n}\right)\right)=0$. Computing the covariance structure, we first have that $\mathbb{E}\left(x_{\pi}^{2}\right)=\frac{1}{n^{2}} \sum_{i, j=1}^{n} \mathbb{P}\left(\sigma_{i}=\pi, \sigma_{j}=\pi\right)=\left(1-\frac{1}{n}\right) p(\pi)^{2}+\frac{1}{n} p(\pi)$, from which we have $\operatorname{Var}\left(x_{\pi}\right)=\frac{1}{n}\left(p(\pi)-p(\pi)^{2}\right)$ and thus $\operatorname{Var}\left(\tilde{x}_{\pi}\right)=p(\pi)-p(\pi)^{2}$. Then similarly for $\pi \neq \pi^{\prime}$ we have

$$\mathbb{E}\left(x_{\pi} x_{\pi^{\prime}}\right)=\frac{1}{n^{2}} \sum_{i, j=1}^{n} \mathbb{P}\left(\sigma_{i}=\pi, \sigma_{j}=\pi^{\prime}\right)=\frac{1}{n^{2}} \sum_{i \neq j} p(\pi) p\left(\pi^{\prime}\right)=\left(1-\frac{1}{n}\right) p(\pi) p\left(\pi^{\prime}\right),$$

from which we have that $\operatorname{Cov}\left(x_{\pi}, x_{\pi^{\prime}}\right)=-\frac{1}{n} p(\pi) p\left(\pi^{\prime}\right)$ and thus

$$\operatorname{Cov}\left(\tilde{x}_{\pi}, \tilde{x}_{\pi^{\prime}}\right)=-p(\pi) p\left(\pi^{\prime}\right). \qquad \square$$

Note: because of the concentration of $x\left(\sigma^{n}\right)$ around its mean, and our assumption that for every n and for every candidate $a \in[m]$, $\mathbb{P}\left(f\left(\sigma^{n}\right)=a\right) \geq \varepsilon$, it is necessary that for every $\alpha>0$ and for every candidate $a \in[m]$ there exists $y \in \Delta^{m!}$ such that $\left\|y-\mathbb{E}\left(x\left(\sigma_{1}\right)\right)\right\|_{1} \leq \alpha$ and $F(y)=a$.

Denote by μ the distribution of $N(0, \Sigma)$ and let $\tilde{X}$ denote a random variable distributed according to μ. Note that μ is a degenerate multivariate normal distribution, as the support of μ concentrates on the hyperplane H_{0} where the coordinates sum to zero. (This is because $\sum_{\pi \in S_{m}} \tilde{x}\left(\sigma^{n}\right)_{\pi}=0$.)

The underlying function $F: \Delta^{m!} \rightarrow[m]$ corresponds to a function $\tilde{F}: \mathbb{R}^{m!}|_{H_{0}} \rightarrow[m]$ in the Gaussian limit, and this function $\tilde{F}$ partitions $\mathbb{R}^{m!}|_{H_{0}}$ into m parts based on the outcome of $\tilde{F}$. We denote these parts by $\left\{\tilde{W}_{a}\right\}_{a \in[m]}$. We will need the following definitions and properties of boundaries, analogous to those above.

Definition 7 (Interior and boundaries of a partition). *We say that* $\tilde{x} \in \mathbb{R}^{m!}|_{H_{0}}$ *is an interior point of the partition* $\left\{\tilde{W}_{a}\right\}_{a \in[m]}$ *induced by* $\tilde{F}$ *if there exists* $\alpha>0$ *such that for all* $\tilde{y} \in \mathbb{R}^{m!}|_{H_{0}}$ *for which* $\|\tilde{x}-\tilde{y}\|_{1} \leq \alpha$*, we have* $\tilde{F}(\tilde{x})=\tilde{F}(\tilde{y})$*. Otherwise, we say that* $\tilde{x} \in \mathbb{R}^{m!}|_{H_{0}}$ *is on the boundary of the partition, which we denote by* $\tilde{B}$*.*

Lemma 3.6. *If the boundary* B *comes from a hyperplane rule, i.e.,* B *is contained in the union of* ℓ *affine hyperplanes in* $\Delta^{m!}$*, then* $\tilde{B}$ *is contained in the union of* $\tilde{\ell}$ *hyperplanes of* $\mathbb{R}^{m!}|_{H_{0}}$*, where* $\tilde{\ell} \leq \ell$*.*

Proof. Two things can happen to an affine hyperplane H of $\Delta^{m!}$ when we take the Gaussian limit: (1) if $\mathbb{E}(x(\pi)) \in H$, then translation by $\mathbb{E}(x(\pi))$ takes H into a hyperplane $\tilde{H}$ of $\mathbb{R}^{m!}|_{H_0}$, and since $\tilde{H}$ goes through the origin, scaling (in particular by $\sqrt{n}$) does not move this hyperplane; (2) if $\mathbb{E}(x(\pi)) \notin H$, then translation by $\mathbb{E}(x(\pi))$ takes H into an affine hyperplane $\tilde{H}$ of $\mathbb{R}^{m!}|_{H_0}$ that does not go through the origin, and then scaling by $\sqrt{n}$ moves $\tilde{H}$ to an affine hyperplane of $\mathbb{R}^{m!}|_{H_0}$ whose L^2 distance from the origin is proportional to $\sqrt{n}$, so in the $n \to \infty$ limit this affine hyperplane "vanishes". $\square$

Definition 8 (Blowup of boundary)**.** *For $\alpha > 0$, we define the blowup of the boundary $\tilde{B}$ by α to be*

$$\tilde{B}^{+\alpha} = \left\{ \tilde{y} \in \mathbb{R}^{m!}|_{H_0} : \exists \tilde{x} \in \tilde{B} \text{ such that } \|\tilde{x} - \tilde{y}\|_1 \leq \alpha \right\}.$$

Let us focus specifically on a coalition of size $c\sqrt{n}$ for some (small) constant c. Corollary 3.4 implies the following.

Corollary 3.7. *For hyperplane rules the limit of the probability that in an election with n voters some coalition of size $c\sqrt{n}$ can change the outcome of the election is $\mu\left(\tilde{X} \in \tilde{B}^{+2c}\right)$.*

The following claim, together with Corollary 3.7, tells us that for hyperplane rules a coalition of size $c\sqrt{n}$ can change the outcome of the election with only small probability, given that c is sufficiently small, proving Part 2 of Theorem 1.2.

Claim 3.8. *Suppose our SCF is a hyperplane rule, and in particular let $\left\{\tilde{H}_i\right\}_{i=1}^{M}$ be a collection of hyperplanes in $\mathbb{R}^{m!}|_{H_0}$ such that $\tilde{B} \subseteq \bigcup_{i=1}^{M} \tilde{H}_i$. Then*

$$\mu\left(\tilde{X} \in \tilde{B}^{+c}\right) \leq \sqrt{\frac{2}{\pi}} \frac{Mc}{\sqrt{\delta}}.$$

Proof. By our condition and a union bound we have

$$\mu\left(\tilde{X} \in \tilde{B}^{+c}\right) \leq \sum_{i=1}^{M} \mu\left(\tilde{X} \in \tilde{H}_i^{+c}\right).$$

For a hyperplane $\tilde{H}$ in $\mathbb{R}^{m!}|_{H_0}$, denote (one of) the corresponding unit normal vector(s) (in the hyperplane H_0) by u. Then

$$\tilde{H} = \left\{ \tilde{x} \in \mathbb{R}^{m!}|_{H_0} : u \cdot \tilde{x} = 0 \right\}$$

and since L^1 distance is always greater than L^2 distance, we have

$$\tilde{H}^{+c} \subseteq \left\{ \tilde{x} \in \mathbb{R}^{m!}|_{H_0} : \exists \tilde{y} \in \tilde{H} \text{ such that } \|\tilde{x} - \tilde{y}\|_2 \leq c \right\} = \left\{ \tilde{x} \in \mathbb{R}^{m!}|_{H_0} : |u \cdot \tilde{x}| \leq c \right\}.$$

Since $\tilde{X}$ is a multidimensional Gaussian r.v., $u \cdot \tilde{X}$ is a one-dimensional Gaussian r.v. (which is centered). Therefore

$$\mu\left(\tilde{X} \in \tilde{H}^{+c}\right) \leq \mu\left(u \cdot \tilde{X} \in [-c, c]\right) \leq \frac{2c}{\sqrt{2\pi \operatorname{Var}\left(u \cdot \tilde{X}\right)}}.$$

We have that

$$\mathrm{Var}\left(u \cdot \tilde{X}\right) = \mathbb{E}\left(u \cdot \tilde{X}\right)^2 = \mathbb{E}\left(u^T \tilde{X} \tilde{X}^T u\right) = u^T \Sigma u,$$

and so all that remains to show is that

$$\min_{u : \|u\|=1, u \perp \mathbf{1}} u^T \Sigma u \geq \delta,$$

where $\mathbf{1}$ is the $m!$-dimensional vector having 1 in every coordinate.

Let $\lambda_1(\Sigma) \geq \lambda_2(\Sigma) \geq \cdots \geq \lambda_{m!}(\Sigma)$ denote the eigenvalues of Σ. Since Σ is positive semidefinite, all eigenvalues are nonnegative. We know that 0 is an eigenvalue of Σ (the corresponding eigenvector is $\mathbf{1}$), so $\lambda_{m!}(\Sigma) = 0$. By the variational characterization of eigenvalues we have

$$\min_{u : \|u\|=1, u \perp \mathbf{1}} u^T \Sigma u = \lambda_{m!-1}(\Sigma),$$

and so we need to show that $\lambda_{m!-1}(\Sigma) \geq \delta$. To do this we use Weyl's inequalities.

Lemma 3.9 (Weyl's inequalities)**.** *For an $n \times n$ matrix M let $\lambda_1(M) \geq \lambda_2(M) \geq \cdots \geq \lambda_n(M)$ denote its eigenvalues. If A and C are $n \times n$ symmetric matrices then*

$$\begin{aligned} \lambda_j(A + C) &\leq \lambda_i(A) + \lambda_{j-i+1}(C) \qquad && \text{if } i \leq j, \\ \lambda_j(A + C) &\geq \lambda_i(A) + \lambda_{j-i+n}(C) \qquad && \text{if } i \geq j. \end{aligned}$$

We use Weyl's inequality for $A = \mathrm{diag}(p)$ and $C = -pp^T$. The eigenvalues of A are $\{p(\pi)\}_{\pi \in S_m}$, all of which are no less than δ. Since C has rank 1, all its eigenvalues but one are zero, and the single nonzero eigenvalue is $\lambda_{m!}(C) = -p^T p$. Since $\Sigma = \mathrm{diag}(p) - pp^T = A + C$, Weyl's inequality tells us that

$$\lambda_{m!-1}(\Sigma) \geq \lambda_{m!}(\mathrm{diag}(p)) + \lambda_{m!-1}\left(-pp^T\right) \geq \delta + 0 = \delta. \qquad \square$$

This implies that we have a lower bound of $\Omega(\sqrt{n})$ for the size of the coalition needed in order to change the outcome of the election for hyperplane rules. As mentioned before, most commonly occurring SCFs are in this class of rules: see Appendix B for many examples.

3.2.1 "Almost" Hyperplane Rules

Furthermore, the Gaussian limiting setting above is not sensitive to small changes to the voting rule for finite n. Consequently, for SCFs that are "almost" hyperplane rules (in a sense we make precise below), the same conclusion holds: a coalition of size $\Omega(\sqrt{n})$ is needed in order to be able to change the outcome of the election with non-negligible probability. In particular, the same result holds for SCFs with arbitrary tie-breaking rules for ranking profiles which lie on one of the hyperplanes (e.g., the tie-breaking rule can depend on the number of voters n).

Definition 9 ("Almost" hyperplane rules)**.** *Fix a finite set of affine hyperplanes of the simplex $\Delta^{m!}$: $H_1, \ldots, H_\ell$. These partition the simplex into finitely many regions. Let $F : \Delta^{m!} \to [m]$ be a function which is constant on each such region, and let B denote the* *induced boundary. Then the sequence of SCFs* $\{f_n\}_{n\geq 1}$, $f_n : S_m^n \to [m]$, *is called an "almost" hyperplane rule if for every* σ^n *such that* $x(\sigma^n) \notin B^{+o\left(1/\sqrt{n}\right)}$, *we have*

$$f_n(\sigma^n) = F(x(\sigma^n)).$$

This SCF is called an "almost" hyperplane rule induced by the affine hyperplanes $H_1, \dots, H_\ell$ *and the function* F.

Lemma 3.10. *Suppose the sequence of SCFs* $\{f_n\}_{n\geq 1}$, $f_n : S_m^n \to [m]$, *is an "almost" hyperplane rule defined by* ℓ *hyperplanes. Then in the Gaussian limiting setting the boundary* $\tilde{B}$ *is contained in the union of* $\tilde{\ell}$ *hyperplanes of* $\mathbb{R}^{m!}|_{H_0}$, *where* $\tilde{\ell} \leq \ell$.

Proof. For finite n, the induced boundary of f_n in the simplex $\Delta^{m!}$ is contained in $B^{+o\left(1/\sqrt{n}\right)}$, by definition. Since in the Gaussian limit we scale by $\sqrt{n}$, the blowup by $o\left(1/\sqrt{n}\right)$ of the boundary B disappears in the limit, and hence we are back to the situation of Lemma 3.6. Consequently, the affine hyperplanes corresponding to our "almost" hyperplane rule either "disappear to infinity" or become hyperplanes of $\mathbb{R}^{m!}|_{H_0}$. □

Corollary 3.11. *Corollary 3.7 and Claim 3.8 hold for "almost" hyperplane rules as well.*

3.3 Smoothness of the Phase Transition

In this final subsection our goal is to show Parts 1 and 3 of Theorem 1.2. The existence of the limits in Part 1 follows immediately from the Gaussian limit described above; we do not detail this, but rather give formulas for these limiting probabilities. These then imply the properties described in Part 3 of the theorem.

In the following let the hyperplane rule be given by affine hyperplanes $H_1, \dots, H_\ell$ of $\Delta^{m!}$ and the function $F : \Delta^{m!} \to [m]$; in the limiting setting denote by $\tilde{H}_1, \dots, \tilde{H}_{\tilde{\ell}}$ the corresponding hyperplanes of $\mathbb{R}^{m!}|_{H_0}$ and denote by $\tilde{F} : \mathbb{R}^{m!}|_{H_0} \to [m]$ the corresponding function.

3.3.1 The Quantities $\overline{q}$ and $\underline{q}$

For $\tilde{x} \in \mathbb{R}^{m!}|_{H_0}$ define

$$\alpha(\tilde{x}) := \inf_{\tilde{y}:\tilde{F}(\tilde{y})\neq\tilde{F}(\tilde{x})} d_1(\tilde{x}, \tilde{y}), \qquad \beta(\tilde{x}) := \max_{a\in[m]} \inf_{\tilde{y}:\tilde{F}(\tilde{y})=a} d_1(\tilde{x}, \tilde{y}).$$

From the previous subsection it is then immediate that we can write

$$\overline{q}(c) = \mu\left(\tilde{X} : \alpha\left(\tilde{X}\right) \leq 2c\right),$$
$$\underline{q}(c) = \mu\left(\tilde{X} : \beta\left(\tilde{X}\right) \leq 2c\right).$$

It is important to note that the boundary $\tilde{B}$ is contained in the union of finitely many hyperplanes, $\tilde{H}_1, \dots, \tilde{H}_{\tilde{\ell}}$, and thus the regions where $\tilde{F}$ is constant are convex cones which are the intersection of finitely many halfspaces. Consequently $\alpha(\tilde{x})$ is either $d_1(\tilde{x}, 0)$, where 0 denotes the origin of $\mathbb{R}^{m!}$, or it is $d_1\left(\tilde{x}, \tilde{H}_j\right)$ for some $1 \leq j \leq \tilde{\ell}$, where $d_1\left(\tilde{x}, \tilde{H}_j\right) =$

$\inf_{\tilde{y}\in\tilde{H}_j} d_1(\tilde{x},\tilde{y})$. If we scale $\tilde{x}$ by some positive constant λ, then the distance from the origin and from every hyperplane scales as well (i.e., $d_1(\lambda\tilde{x},0) = \lambda d_1(\tilde{x},0)$ and $d_1\left(\lambda\tilde{x},\tilde{H}_j\right) = \lambda d_1\left(\tilde{x},\tilde{H}_j\right)$), and thus for every $\lambda > 0$, we have $\alpha(\lambda\tilde{x}) = \lambda\alpha(\tilde{x})$. Consequently, if we write $\tilde{x} = \|\tilde{x}\|_2\,\tilde{s}$, where $\tilde{s} \in S^{m!-1}$, and $S^{m!-1}$ denotes the $(m!-1)$-sphere (not to be confused with S_m^n, the set of ranking profiles on n voters and m candidates), then we have $\alpha(\tilde{x}) = \|\tilde{x}\|_2\,\alpha(\tilde{s})$.

The same scaling property holds for β as well, and hence we have

$$\overline{q}(c) = \mu\left(\tilde{X} : \left\|\tilde{X}\right\|_2 \alpha\left(\tilde{S}\right) \leq 2c\right), \tag{3}$$

$$\underline{q}(c) = \mu\left(\tilde{X} : \left\|\tilde{X}\right\|_2 \beta\left(\tilde{S}\right) \leq 2c\right). \tag{4}$$

Recall that our condition that for every $a \in [m]$, $\mathbb{P}(f(\sigma) = a) \geq \varepsilon$, implies that for every $\eta > 0$ and for every $a \in [m]$ there exists $\tilde{x} \in \mathbb{R}^{m!}|_{H_0}$ such that $\|\tilde{x}\|_2 \leq \eta$ and $\tilde{F}(\tilde{x}) = a$. Consequently for every $\tilde{x} \in \mathbb{R}^{m!}|_{H_0}$ we must have $\alpha(\tilde{x}) \leq d_1(\tilde{x},0)$ and $\beta(\tilde{x}) \leq d_1(\tilde{x},0)$. In particular, for $\tilde{s} \in S^{m!-1}$ we have $d_1(\tilde{s},0) \leq \sqrt{m!}\,d_2(\tilde{s},0) = \sqrt{m!}$ and so $\alpha(\tilde{s}), \beta(\tilde{s}) \leq \sqrt{m!}$. This immediately implies that for every $c > 0$ we have

$$\underline{q}(c) \geq \mu\left(\tilde{X} : \left\|\tilde{X}\right\|_2 \leq \frac{2c}{\sqrt{m!}}\right) > 0.$$

To show that $\overline{q}(c) < 1$, note that since the boundary is contained in the union of finitely many hyperplanes, there exists $\tilde{s}^* \in S^{m!-1}$ such that $\alpha(\tilde{s}^*) > 0$. By continuity of α, there exists a neighborhood $U \subseteq S^{m!-1}$ of $\tilde{s}^*$ such that for every $\tilde{s} \in U$, $\alpha(\tilde{s}) \geq \alpha(\tilde{s}^*)/2$. For any $\tilde{x}$ such that $\tilde{x}/\|\tilde{x}\|_2 \in U$ and $\|\tilde{x}\|_2 > \frac{4c}{\alpha(\tilde{s}^*)}$, we have

$$\alpha(\tilde{x}) = \|\tilde{x}\|_2\,\alpha\left(\tilde{x}/\|\tilde{x}\|_2\right) > \frac{4c}{\alpha(\tilde{s}^*)}\frac{\alpha(\tilde{s}^*)}{2} = 2c.$$

So consequently

$$\overline{q}(c) \leq 1 - \mu\left(\tilde{X} : \tilde{X}/\left\|\tilde{X}\right\|_2 \in U, \left\|\tilde{X}\right\|_2 > \frac{4c}{\alpha(\tilde{s}^*)}\right) < 1.$$

Finally, the fact that $\underline{q}(c)$ and $\overline{q}(c)$ are continuously differentiable follows from the formulas (3) and (4), since $\underline{q}(c)$ and $\overline{q}(c)$ are both written as the Gaussian volume of a subset of $\mathbb{R}^{m!}|_{H_0}$, and in both cases this subset grows continuously as c increases. The derivative of both $\underline{q}(c)$ and $\overline{q}(c)$ is bounded at zero (by Corollary 3.7 and Claim 3.8), while as $c \to \infty$ the derivative approaches zero, and since the derivative is continuous, it must be bounded by a constant for the whole half-line.

3.3.2 The Quantities $\overline{r}$ and $\underline{r}$

In the previous setup when the coalition of size $c\sqrt{n}$ was not specified, the ranking profile could be changed arbitrarily within a Hamming ball of radius $c\sqrt{n}$. On the probability simplex $\Delta^{m!}$ this corresponded to an L^1 ball of radius $2c/\sqrt{n}$, and in the rescaled limiting setting it corresponded to an L^1 ball in $\mathbb{R}^{m!}|_{H_0}$ of radius $2c$. When the coalition of size $c\sqrt{n}$ is specified, things are slightly different. In particular, when we look at the probability distribution on the probability simplex $\Delta^{m!}$ induced by the distribution on ranking profiles (or, in the limiting setting, the Gaussian distribution on $\mathbb{R}^{m!}|_{H_0}$), then we have lost track of the votes of any specific coalition. Nonetheless, the Gaussian limiting setting still provides formulas for the limiting probabilities $\underline{r}(c)$ and $\overline{r}(c)$.

We can first draw a random ranking profile for the other $n - c\sqrt{n}$ voters not in the coalition, $\sigma^{n-c\sqrt{n}}$, and then the voters in the coalition can set their votes arbitrarily. The question is, how can the coalition affect the outcome of the vote? In particular, (a) can they change the outcome of the election, and (b) can they elect any candidate?

The ranking profile $\sigma^{n-c\sqrt{n}}$ corresponds to a point $x\left(\sigma^{n-c\sqrt{n}}\right)$ on the probability simplex $\Delta^{m!}$, and by setting their votes the coalition can move this point on the probability simplex in some neighborhood of $x\left(\sigma^{n-c\sqrt{n}}\right)$. We omit the calculation for finite n and only present the result in the limiting setting.

Suppose the limiting ranking profile of the voters other than the coalition corresponds to the point $\tilde{x} \in \mathbb{R}^{m!}|_{H_0}$. Then the set of points the coalition can reach is the following:

$$R_c(\tilde{x}) := \left\{\tilde{y} \in \mathbb{R}^{m!}|_{H_0} : \forall \pi \in S_m : \tilde{y}_\pi - \tilde{x}_\pi + cp(\pi) \geq 0\right\}.$$

We can then define

$$\begin{aligned}
\varphi(\tilde{x}) &:= \inf\left\{\gamma : \exists \tilde{y} \in R_\gamma(\tilde{x}) \text{ such that } \tilde{F}(\tilde{y}) \neq \tilde{F}(\tilde{x})\right\},\\
\psi(\tilde{x}) &:= \inf\left\{\gamma : \forall a \in [m]\, \exists \tilde{y} \in R_\gamma(\tilde{x}) \text{ such that } \tilde{F}(\tilde{y}) = a\right\},
\end{aligned}$$

and it follows immediately that we can then write

$$\begin{aligned}
\overline{r}(c) &= \mu\left(\tilde{X} : \varphi\left(\tilde{X}\right) \leq c\right),\\
\underline{r}(c) &= \mu\left(\tilde{X} : \psi\left(\tilde{X}\right) \leq c\right).
\end{aligned}$$

In the same way as in Section 3.3.1 one can argue that φ and ψ scale: if $\lambda > 0$ then $\varphi(\lambda\tilde{x}) = \lambda\varphi(\tilde{x})$ and $\psi(\lambda\tilde{x}) = \lambda\psi(\tilde{x})$. Hence we have

$$\overline{r}(c) = \mu\left(\tilde{X} : \left\|\tilde{X}\right\|_2 \varphi\left(\tilde{S}\right) \leq c\right), \tag{5}$$

$$\underline{r}(c) = \mu\left(\tilde{X} : \left\|\tilde{X}\right\|_2 \psi\left(\tilde{S}\right) \leq c\right). \tag{6}$$

For every $0 < c < \infty$ we have $\overline{r}(c) \leq \overline{q}(c) < 1$ (using Section 3.3.1). Let us now show that also $\underline{r}(c) > 0$. We claim that for all $\tilde{s} \in S^{m!-1}|_{H_0}$, $\psi(\tilde{s}) \leq \frac{2}{\delta}$. This follows from the fact that if $\tilde{s} \in S^{m!-1}|_{H_0}$ then $S^{m!-1}|_{H_0} \subseteq R_{\frac{2}{\delta}}(\tilde{s})$, which is true because if $\tilde{y} \in S^{m!-1}|_{H_0}$ then for all $\pi \in S_m$, $\tilde{y}_\pi - \tilde{s}_\pi + \frac{2}{\delta}p(\pi) \geq -1 - 1 + \frac{2}{\delta}\delta = 0$. Thus we have

$$\underline{r}(c) \geq \mu\left(\tilde{X} : \left\|\tilde{X}\right\|_2 \leq \frac{c\delta}{2}\right) > 0$$

as claimed.

Finally, the fact that $\underline{r}(c)$ and $\overline{r}(c)$ are continuously differentiable follows from the formulas (5) and (6) using an argument given above: $\underline{r}(c)$ and $\overline{r}(c)$ are written as the Gaussian volume of subsets of $\mathbb{R}^{m!}|_{H_0}$, and these subsets grow continuously as c increases. The derivative of both $\underline{r}(c)$ and $\overline{r}(c)$ is bounded at zero (by Corollary 3.7 and Claim 3.8), while as $c \to \infty$ the derivative approaches zero, and since the derivative is continuous, it must be bounded by a constant for the whole half-line. $\square$

Acknowledgments

We thank anonymous referees for helpful comments. E.M. is supported by NSF (DMS 1106999) and by DOD ONR grant N000141110140, and M.Z.R. is supported by a UC Berkeley Graduate Fellowship and by NSF (DMS 0548249).

Appendix A. Decomposing $\mathbb{R}^d$ as the Disjoint Union of Finitely Many Convex Cones: Only Via Hyperplanes

In order for the paper to be self-contained, we reproduce here the main definitions and results of Kemperman (1986) that make precise the claim used in the proof of Lemma 3.1 that the only way to decompose $\mathbb{Q}^d_{\geq 0} \setminus \{0\}$ into the disjoint union of finitely many $\mathbb{Q}$-convex cones is via hyperplanes. Kemperman's paper deals with convex sets in general, but here we summarize the results about convex *cones* that are relevant to us. Kemperman's results pertain to finite dimensional linear spaces and we will state them in this form; in the end we show how results for $\mathbb{R}^d_{\geq 0}$ follow immediately from these, and as a consequence we also obtain the claim used in the proof of Lemma 3.1.

Let us start with the main definitions. In the following, all linear spaces are over the reals and are finite dimensional. Let X be a linear space. A convex cone is a subset $K \subseteq X$ such that $x, y \in K$ and $\lambda > 0$ imply $x + y \in K$ and $\lambda x \in K$. (We do not require that $0 \in K$.) For a set $A \subseteq X$, denote its affine hull by $\operatorname{aff}(A)$, its convex hull by $\operatorname{cvx}(A)$, and its closure by $\operatorname{cl}(A)$. Note that if $K \subseteq X$ is a convex cone, then $\operatorname{aff}(K)$ is a linear subspace of X.

We define two special types of convex cones: basic convex cones and elementary convex cones.

Definition 10 (Basic convex cone)**.** *Let K be a convex cone in a finite dimensional linear space X. We say that K is a* basic convex cone *(in X) if K is a member $K = K_0$ of some partition*

$$X = K_0 \dot\cup K_1 \dot\cup \ldots \dot\cup K_r$$

of X into finitely many disjoint convex cones $\{K_i\}_{i=0}^r$.

Note that any linear subspace Y of X is a basic convex cone, from which it immediately follows that K is a basic convex cone in X if and only if it is a basic convex cone in $\operatorname{aff}(K)$.

In order to define elementary convex cones, we need a few more definitions.

Definition 11 (Open polyhedral convex cone)**.** *Let K be a convex cone in a finite dimensional linear space X. We say that K is an* open polyhedral convex cone *relative to X if K*

can be expressed as the intersection of finitely many open halfspaces $H_1, \ldots, H_\ell$ of X, each of which has the origin on its boundary. The whole linear space X is an open polyhedral convex cone with $\ell = 0$.

Definition 12 (Relatively open polyhedral convex cone)**.** *Let K be a convex cone in a finite dimensional linear space X. Then K is a* relatively open polyhedral convex cone *if either $K = \emptyset$ or K is an open polyhedral convex cone relative to* $\mathrm{aff}\,(K)$*.*

Definition 13 (Elementary convex cone)**.** *Let K be a convex cone in a finite dimensional linear space X. We say that K is an* elementary convex cone *if K can be represented as a disjoint union of finitely many relatively open polyhedral convex cones.*

The main result of Kemperman concerning convex cones is the following (Kemperman, 1986, Thm. 2).

Theorem A.1. *Let K be a convex cone in $\mathbb{R}^d$. Then K is a basic convex cone if and only if it is an elementary convex cone.*

In Lemma 3.1 we only use the "only if" direction, and we thus leave the proof of the "if" direction as an exercise for the reader.

Proof of "only if" direction. Let X be a finite dimensional linear space and let K be a basic convex cone in X of dimension $d = \dim(K) = \dim(Y)$, where $Y = \mathrm{aff}\,(K)$. We prove by induction on d the following:

(i) The relative interior of K, denoted by K^0, is a relatively open polyhedral convex cone.

(ii) If $K^0 \neq Y$, then denote by $F_1, \ldots, F_\ell$ the $(d-1)$-dimensional hyperplanes in Y corresponding to the finitely many faces of the polyhedron $\mathrm{cl}\,(K) = \mathrm{cl}\left(K^0\right)$. Then the convex cones $F_i \cap K$, $i = 1, \ldots, \ell$, are elementary convex cones of dimension at most $d-1$ (but they need not be disjoint).

(iii) The convex cone K is also an elementary convex cone.

If $K = \emptyset$, then properties (i) - (iii) hold. If $d = 0$, then necessarily $K = \{0\}$, since K is a convex cone, and again K satisfies properties (i) - (iii) above.

So we may assume that $d \geq 1$ and that each basic convex cone of dimension at most $d-1$ satisfies properties (i) - (iii) above. Since K is a basic convex cone, there exists a partition

$$Y = K_0 \dot{\cup} K_1 \dot{\cup} \ldots \dot{\cup} K_r \tag{7}$$

of Y into finitely many disjoint convex cones $\{K_j\}_{j=0}^r$, with $K_0 = K$. We may assume that $r \geq 0$ is minimal, and hence the K_j are non-empty. Note that K^0 is also non-empty since $\dim(K) = \dim(Y)$.

If $r = 0$ then $K = K_0 = Y$ and the properties (i) - (iii) above are immediately satisfied, so we may assume that $r \geq 1$. For $j = 1, \ldots, r$, let H_j be a hyperplane in Y which separates the convex cone $K = K_0$ with non-empty interior K^0 from the non-empty convex cone K_j. (Such hyperplanes exist by the hyperplane separation theorem, and, moreover, each such hyperplane goes through the origin, because each K_j contains at least one point from every open ball around the origin, since each K_j is a cone.) Let H_j^0 be the associated open half space in Y which contains the interior K^0 of K. Let

$$L^0 = H_1^0 \cap \cdots \cap H_r^0.$$

Then L^0 is a polyhedral convex cone, which is open relative to Y, and contains the interior K^0 of K.

We claim that $L^0 = K^0$. It is enough to show that $L^0 \subseteq K$, because then $L^0 \subseteq K^0$ follows from the definition of K^0. Suppose on the contrary that there exists $x \in L^0$ such that $x \notin K$. Then from the partition (7) there must exist an index $1 \leq j \leq r$ with $x \in K_j$. This implies that $x \notin H_j^0$ and thus $x \notin L^0$, which is a contradiction. This proves (i).

Now let us show (ii). By (7), we can write the linear space F_i as the disjoint union of the convex cones $F_i \cap K_j$, $j = 0, \ldots, r$, and thus $F_i \cap K$ is a basic convex cone and hence, by induction, an elementary convex cone.

Finally, let us show that K is an elementary convex cone. Since K^0 is a polyhedral convex cone which is open relative to Y, it only remains to show that $K \setminus K^0$ can be written as a finite disjoint union of relatively open polyhedral convex cones. By (ii), we can write $K \setminus K^0$ as the finite union of elementary convex cones:

$$K \setminus K^0 = \cup_{i=1}^{\ell} (F_i \cap K),$$

so what remains is to show that we can write this as a finite *disjoint* union of relatively open polyhedral convex cones. We may assume w.l.o.g. that $F_i \cap K \neq \emptyset$ for all i and that $(F_i \cap K) \nsubseteq (F_j \cap K)$ for all $i \neq j$ (otherwise we can leave out $F_i \cap K$ from the union).

We claim that then for every i,

$$\operatorname{rel\,int}(F_i \cap K) \subseteq (F_i \cap K) \setminus \bigcup_{j \neq i} (F_j \cap F_i \cap K), \tag{8}$$

from which it immediately follows that $\operatorname{rel\,int}(F_i \cap K) \cap \operatorname{rel\,int}(F_j \cap K) = \emptyset$ for $i \neq j$. To show (8), let the two open halfspaces on either side of the hyperplane F_j be denoted by F_j^+ and F_j^-. W.l.o.g. assume that $K \cap F_j^- = \emptyset$. Since $(F_i \cap K) \nsubseteq (F_j \cap K)$, we must have $(F_i \cap K) \cap F_j^+ \neq \emptyset$. Let $x \in (F_i \cap K) \cap F_j^+$ and let $y \in F_j \cap F_i \cap K$. Since $F_i \cap K$ is convex, the interval from x to y is contained in $F_i \cap K$, but because $(F_i \cap K) \cap F_j^- = \emptyset$, no points on this line past the point y can be in $F_i \cap K$; hence $y \notin \operatorname{rel\,int}(F_i \cap K)$.

Since $F_i \cap K$ is a basic convex cone, $\operatorname{rel\,int}(F_i \cap K)$ is a relatively open polyhedral convex cone by induction. If $F_i \cap K = \operatorname{aff}(F_i \cap K)$ then $\operatorname{rel\,int}(F_i \cap K) = F_i \cap K$. If not, then denote by $F_{i,1}, \ldots, F_{i,\ell_i}$ the hyperplanes in $\operatorname{aff}(F_i \cap K)$ corresponding to the finitely many faces of the polyhedron $\operatorname{cl}(F_i \cap K)$. By induction, the convex cones $F_{i,j} \cap F_i \cap K$, $j = 1, \ldots, \ell_i$, are elementary convex cones, and we can write

$$K \setminus K^0 = \left(\dot{\cup}_{i=1}^{\ell} \operatorname{rel\,int}(F_i \cap K) \right) \dot{\bigcup} \left(\cup_{i=1}^{\ell} \cup_{j=1}^{\ell_i} (F_{i,j} \cap F_i \cap K) \right).$$

What remains to be shown is that $\cup_{i=1}^{\ell} \cup_{j=1}^{\ell_i} (F_{i,j} \cap F_i \cap K)$ can be written as a finite disjoint union of relatively open polyhedral convex cones; this follows by iterating the previous argument. $\square$

Let us now show that $\mathbb{R}^d_{\geq 0}$ is a basic convex cone in $\mathbb{R}^d$. For $i = 1, \ldots, d$, define the closed halfspace $H_i^{\geq 0} = \left\{x \in \mathbb{R}^d : x_i \geq 0\right\}$ and its complement $H_i^{<0} = \left\{x \in \mathbb{R}^d : x_i < 0\right\}$, and from these define the convex cones

$$K_i = H_1^{\geq 0} \cap \cdots \cap H_{i-1}^{\geq 0} \cap H_i^{<0}, \qquad i = 1, \ldots, d.$$

Then we can write $\mathbb{R}^d$ as the disjoint union of the convex cones $\mathbb{R}^d_{\geq 0}$ and $K_1, \ldots, K_d$, showing that indeed $\mathbb{R}^d_{\geq 0}$ is a basic convex cone. This implies that if we can write $\mathbb{R}^d_{\geq 0}$ as the disjoint union of the convex cones $C_1, \ldots, C_r$, then each C_i is a basic convex cone, and hence, by Theorem A.1, an elementary convex cone.

Now let us turn to the claim in the proof of Lemma 3.1. In Lemma 3.1, we write $\mathbb{Q}^{m!}_{\geq 0} \setminus \{0\}$ as the disjoint union of finitely many $\mathbb{Q}$-convex cones: $\mathbb{Q}^{m!}_{\geq 0} \setminus \{0\} = C_0 \dot\cup C_1 \dot\cup \ldots \dot\cup C_r$. For $i = 0, \ldots, r$, let $\tilde{C}_i = \operatorname{cvx}(C_i)$. It is known (see, e.g., Young, 1975) that $C_i = \mathbb{Q}^{m!} \cap \tilde{C}_i$. The $\tilde{C}_i$ are therefore disjoint convex cones which satisfy

$$\tilde{C}_0 \dot\cup \tilde{C}_1 \dot\cup \ldots \dot\cup \tilde{C}_r \subseteq \mathbb{R}^{m!}_{\geq 0} \tag{9}$$

and

$$\operatorname{cl}\left(\tilde{C}_0\right) \cup \operatorname{cl}\left(\tilde{C}_1\right) \cup \cdots \cup \operatorname{cl}\left(\tilde{C}_r\right) = \mathbb{R}^{m!}_{\geq 0}. \tag{10}$$

Our goal is to show that each $\tilde{C}_i$ is an elementary convex cone. Conditions (10) and (9) are very similar to the definition of a basic convex cone; in this spirit let us introduce the following definition.

Definition 14 (Basic convex cone up to closure)**.** *Let K_0 be a convex cone in a finite dimensional linear space X. We say that K_0 is a* basic convex cone up to closure *(in X) if there exist disjoint convex cones $K_1, \ldots, K_r$ such that*

$$K_0 \dot\cup K_1 \dot\cup \ldots \dot\cup K_r \subseteq X$$

and

$$\operatorname{cl}(K_0) \cup \operatorname{cl}(K_1) \cup \cdots \cup \operatorname{cl}(K_r) = X.$$

Since $\mathbb{R}^d_{\geq 0}$ is a basic convex cone, the $\tilde{C}_i$ above are basic convex cones up to closure.

In fact, every basic convex cone up to closure is an elementary convex cone; the proof is exactly the same as the one shown above for the "only if" direction of Theorem A.1, one just needs to replace "basic convex cone" with "basic convex cone up to closure" everywhere in the proof, and make the appropriate changes. Moreover, the other direction of Theorem A.1 implies that actually every basic convex cone up to closure is a basic convex cone.

Hence the $\tilde{C}_i$ are elementary convex cones, which is what we need in Lemma 3.1.

Appendix B. Most Voting Rules are Hyperplane Rules: Examples

In the following we show that all positional scoring rules, instant-runoff voting, Coombs' method, contingent vote, the Kemény-Young method, Bucklin voting, Nanson's method, Baldwin's method, and Copeland's method are all hyperplane rules.

- **Positional scoring rules.** Let $w \in \mathbb{R}^m$ be a weight vector. Given a ranking profile vector σ, the (normalized) *score* of candidate $a \in [m]$ is $s_a = \frac{1}{n} \sum_{i=1}^{n} w\left(\sigma_i^{-1}(a)\right)$. The *positional scoring rule* associated to the weight vector w elects the candidate who has the highest score. (In case of a tie, there is some tie-breaking rule, but we do not care about this here.) We denote such a SCF on n voters by f_n^w. Examples include plurality (with weight vector $w = (1, 0, 0, \ldots, 0)$), Borda count (with weight vector $w = (m-1, m-2, \ldots, 0)$) and veto (with weight vector $w = (1, 1, \ldots, 1, 0)$).

  To a sequence of SCFs $\{f_n^w\}_{n \geq 1}$ we can associate a function $F^w : \Delta^{m!} \to [m]$ in the following way. For a candidate $a \in [m]$ and $x \in \Delta^{m!}$, define the (normalized) score $s_a(x) = \sum_{\pi \in S_m} x_\pi w\left(\pi^{-1}(a)\right)$, and let

  $$F^w(x) := \arg\max_{a \in [m]} s_a(x),$$

  if this arg max is unique, and if it is not unique, then there is some tie-breaking rule. This construction guarantees that $f_n^w = F^w|_{D_n}$. For candidates $a \neq b$, define

  $$H_{a,b} := \left\{ x \in \Delta^{m!} : s_a(x) = s_b(x) \right\},$$

  which is an affine hyperplane of the probability simplex $\Delta^{m!}$. Clearly the boundary B^w is contained in the union of $\binom{m}{2}$ such affine hyperplanes:

  $$B^w \subseteq \bigcup_{a \neq b \in [m]} H_{a,b}.$$

- **Instant-runoff voting.** If a candidate receives absolute majority of first preference votes, then that candidate wins. If no candidate receives an absolute majority, then the candidate with fewest top votes is eliminated. In the next round the votes are counted again, with each ballot counted as one vote for the advancing candidate who is ranked highest on that ballot. This is repeated until the winning candidate receives a majority of the vote against the remaining candidates.

  The boundary corresponds to two kinds of situations: either (1) there is a tie at the top at the end, when only two candidates remain; or (2) there is a tie for eliminating a candidate at the end of one of the rounds. Technically situation (1) is also contained in situation (2), since at the very end one can view choosing a winner as eliminating the second placed candidate. One can see that if candidates a and b are tied for elimination after candidates $C \subseteq [m] \setminus \{a, b\}$ (where $C = \emptyset$ is allowed) have been eliminated, then necessarily

  $$\sum_{C' \subseteq C} \sum_{\substack{\{\pi(1), \ldots, \pi(|C'|)\} = C', \\ \pi(|C'|+1) = a}} x_\pi = \sum_{C' \subseteq C} \sum_{\substack{\{\pi(1), \ldots, \pi(|C'|)\} = C', \\ \pi(|C'|+1) = b}} x_\pi. \tag{11}$$

  Consequently, denoting by $s_{a,C}(x)$ the quantity on the left hand side of (11), the boundary B is contained in the union of at most $m^2 2^m$ affine hyperplanes:

  $$B \subseteq \bigcup_{a \neq b} \bigcup_{C \subseteq [m] \setminus \{a,b\}} \left\{ x \in \Delta^{m!} : s_{a,C}(x) = s_{b,C}(x) \right\}.$$

- **Coombs' method.** This is similar to IRV, but the elimination rule is different. If a candidate receives absolute majority of first preference votes, then that candidate wins. If no candidate receives an absolute majority, then the candidate who is ranked last by the most voters is eliminated. In the next round the votes are counted again, with each ballot counted as one vote for the advancing candidate who is ranked highest on that ballot. This is repeated until the winning candidate receives a majority of the vote against the remaining candidates.

  The boundary corresponds to two kinds of situations: either (1) there is a tie at the top at the end, when only two candidates remain; or (2) there is a tie for eliminating a candidate at the end of one of the rounds. Technically situation (1) is also contained in situation (2), since at the very end one can view choosing a winner as eliminating the second placed candidate. One can see that if candidates a and b are tied for elimination after candidates $C \subseteq [m] \setminus \{a,b\}$ (where $C = \emptyset$ is allowed) have been eliminated, then necessarily

$$\sum_{C' \subseteq C} \sum_{\substack{\{\pi(m),\dots,\pi(m-|C'|+1)\}=C', \\ \pi(m-|C'|)=a}} x_\pi = \sum_{C' \subseteq C} \sum_{\substack{\{\pi(m),\dots,\pi(m-|C'|+1)\}=C', \\ \pi(m-|C'|)=b}} x_\pi. \tag{12}$$

  Consequently, denoting by $s_{a,C}(x)$ the quantity on the left hand side of (12), the boundary B is contained in the union of at most $m^2 2^m$ affine hyperplanes:

$$B \subseteq \bigcup_{a \neq b} \bigcup_{C \subseteq [m] \setminus \{a,b\}} \left\{ x \in \Delta^{m!} : s_{a,C}(x) = s_{b,C}(x) \right\}.$$

- **Contingent vote.** This is also similar to IRV, except here all but two candidates get eliminated after the first round. If a candidate receives absolute majority of first preference votes, then he/she wins. If no candidate receives an absolute majority, then all but the top two leading candidates are eliminated and there is a second count, where the votes of those who supported an eliminated candidate are redistributed among the two remaining candidates. The candidate who then achieves absolute majority wins.

  Here the boundary B corresponds to two kinds of situations: either (1) there are two distinct top candidates, and when the votes of the voters who voted for other candidates are redistributed, then the two top candidates are in a dead heat; or (2) there are two or more candidates who receive an equal number of votes in the first round. Both of these situations can be described as subsets of affine hyperplanes, and so B is contained in the union of at most $m(m-1)$ affine hyperplanes:

$$B \subseteq \bigcup_{a \neq b} \left\{ x \in \Delta^{m!} : \sum_{\pi : \pi(1) = a} x_\pi + \sum_{\pi : \pi(1) \notin \{a,b\}, a \overset{\pi}{>} b} x_\pi = \sum_{\pi : \pi(1) = b} x_\pi + \sum_{\pi : \pi(1) \notin \{a,b\}, b \overset{\pi}{>} a} x_\pi \right\}$$
$$\cup \bigcup_{a \neq b} \left\{ x \in \Delta^{m!} : \sum_{\pi : \pi(1) = a} x_\pi = \sum_{\pi : \pi(1) = b} x_\pi \right\}.$$

- **Kemény-Young method.** Denote by K the Kendall tau distance, which is a metric on permutations which counts the number of pairwise disagreements between the two permutations, i.e.,

$$K(\tau_1, \tau_2) = \sum_{\{a,b\}} \mathbf{1}[a \text{ and } b \text{ are in the opposite order in } \tau_1 \text{ and } \tau_2],$$

where the sum is over all unordered pairs of distinct candidates. Given a ranking profile σ^n, the Kemény-Young method selects the ranking which minimizes the sum of Kendall tau distances from the votes:

$$\tau = \arg\min \sum_{i=1}^{n} K(\sigma_i, \tau),$$

and then the winner of the election is declared to be $\tau(1)$. For us it will be convenient to write τ as

$$\tau = \arg\min \sum_{\pi} x_\pi(\sigma^n) K(\pi, \tau).$$

Here if we are on the boundary B then there must exist two rankings τ_1 and τ_2 such that $\tau_1(1) \neq \tau_2(1)$ and $\sum_\pi x_\pi K(\pi, \tau_1) = \sum_\pi x_\pi K(\pi, \tau_2)$. Thus B is contained in the union of at most $(m!)^2$ affine hyperplanes:

$$B \subseteq \bigcup_{\tau_1 \neq \tau_2} \left\{ x \in \Delta^{m!} : \sum_\pi x_\pi K(\pi, \tau_1) = \sum_\pi x_\pi K(\pi, \tau_2) \right\}.$$

- **Bucklin voting.** First every candidate gets a point from all the voters who ranked them at the top. If there is a candidate who has a majority (i.e., more than $n/2$ points), then that candidate wins. If not, then every candidate gets a point from all the voters who ranked them second. If there is a candidate who has more than $n/2$ points after this, then the candidate with the most points wins (there might be multiple candidates with more than $n/2$ points after a given round). This process is iterated until there is a candidate with more than $n/2$ points.

Here a point on the boundary B corresponds to a situation where some pair of candidates have the same number of points after some number of rounds. Therefore B is contained in the union of at most $m^2(m-1)/2$ affine hyperplanes:

$$B \subseteq \bigcup_{a \neq b} \bigcup_{k=1}^{m} \left\{ x \in \Delta^{m!} : \sum_{i=1}^{k} \sum_{\pi:\pi(i)=a} x_\pi = \sum_{i=1}^{k} \sum_{\pi:\pi(i)=b} x_\pi \right\}.$$

- **Nanson's method.** This is Borda count combined with a variation of the instant-runoff voting procedure. First, the Borda scores of all candidates are computed, and then those candidates with Borda score no greater than the average Borda score are eliminated. Then the Borda scores of each remaining candidate are recomputed, as if the eliminated candidates were not on the ballot. This is repeated until there is a final candidate left.

The boundary corresponds to situations when a candidate's Borda score exactly equals the average score after some candidates have been eliminated. For $C \subseteq [m]$, denote by $s_{a,C}(x)$ the score of candidate a after exactly the candidates in C have been eliminated ($s_{a,C}(x)$ is a linear function of $\{x_\pi\}_{\pi \in S_m}$), and denote by $\bar{s}_C(x)$ the average score of remaining candidates after exactly the candidates in C have been eliminated. The boundary B is contained in the union of at most $m2^m$ affine hyperplanes:

$$B \subseteq \bigcup_{a \in [m]} \bigcup_{C \subseteq [m] \setminus \{a\}} \left\{ x \in \Delta^{m!} : s_{a,C}(x) = \bar{s}_C(x) \right\}.$$

- **Baldwin's method.** This is essentially Borda count combined with the instant-runoff voting procedure. First, the Borda scores of all candidates are computed, and then the candidate with the lowest score is eliminated. Then the Borda scores of each remaining candidate are recomputed, as if the eliminated candidate were not on the ballot. This is repeated until there is a final candidate left.

  The boundary corresponds to ties for eliminating a candidate at the end of one of the rounds. Borrow the notation $s_{a,C}(x)$ from the previous example. The boundary B is thus contained in the union of at most $m^2 2^m$ affine hyperplanes:

$$B \subseteq \bigcup_{a \neq b} \bigcup_{C \subseteq [m] \setminus \{a,b\}} \left\{ x \in \Delta^{m!} : s_{a,C}(x) = s_{b,C}(x) \right\}.$$

- **Copeland's method.** This is a pairwise aggregation method: every candidate gets 1 point for each other candidate it beats in a pairwise majority election, and 1/2 a point for each candidate it ties with in a pairwise majority election. The winner is the candidate who receives the most points. This method corresponds to cutting the simplex $\Delta^{m!}$ up into finitely many regions via $\binom{m}{2}$ affine hyperplanes, and in each region the winner is the candidate with the most points.

  While in the previous examples tie-breaking rules were not an issue, here it does become important. We do not care about tie-breaking rules when we are on an affine hyperplane where two candidates tie each other in a pairwise majority election. However, there are open regions in the intersection of halfspaces defined by the affine hyperplanes where candidates are tied at the top with having the same scores. In this case, in order for Copeland to be a hyperplane rule, we need to break ties in favor of the same candidate for the whole region. (This is also how ties are broken for Copeland's method in Xia & Conitzer, 2008b.)

  Using this tie-breaking rule Copeland's method is indeed a hyperplane rule, since the boundary is contained in the union of at most $\binom{m}{2}$ affine hyperplanes:

$$B \subseteq \bigcup_{a \neq b} \left\{ x \in \Delta^{m!} : \sum_{\pi : a \overset{\pi}{>} b} x_\pi = \sum_{\pi : b \overset{\pi}{>} a} x_\pi \right\}.$$

References

Achlioptas, D. (1999). *Threshold phenomena in random graph colouring and satisfiability.* Ph.D. thesis, Department of Computer Science, University of Toronto.

Achlioptas, D., Naor, A., & Peres, Y. (2005). Rigorous location of phase transitions in hard optimization problems. *Nature*, *435*(7043), 759–764.

Bartholdi III, J., Tovey, C., & Trick, M. (1989). The Computational Difficulty of Manipulating an Election. *Social Choice and Welfare*, *6*(3), 227–241.

Betzler, N., Niedermeier, R., & Woeginger, G. J. (2011). Unweighted coalitional manipulation under the Borda rule is NP-hard. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 55–60.

Brandt, F. (2009). Some remarks on Dodgson's voting rule. *Mathematical Logic Quarterly*, *55*(4), 460–463.

Caragiannis, I., & Procaccia, A. D. (2011). Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence*, *175*(9–10), 1655–1671.

Chamberlain, G., & Rothschild, M. (1981). A note on the probability of casting a decisive vote. *Journal of Economic Theory*, *25*(1), 152–162.

Cheeseman, P., Kanefsky, B., & Taylor, W. (1991). Where the really hard problems are. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 331–337.

Conitzer, V., & Sandholm, T. (2006). Nonexistence of Voting Rules That Are Usually Hard to Manipulate. In *Proceedings of the 21st National Conference on Artificial Intelligence*, Vol. 21, pp. 627–634.

Conitzer, V., Sandholm, T., & Lang, J. (2007). When are elections with few candidates hard to manipulate?. *Journal of the ACM*, *54*(3), 1–33.

Dobzinski, S., & Procaccia, A. (2008). Frequent Manipulability of Elections: The Case of Two Voters. In *Proceedings of the 4th International Workshop on Internet and Network Economics*, pp. 653–664. Springer.

Faliszewski, P., Hemaspaandra, E., Hemaspaandra, L., & Rothe, J. (2009). Llull and Copeland Voting Computationally Resist Bribery and Constructive Control. *Journal of Artificial Intelligence Research*, *35*, 275–341.

Faliszewski, P., & Procaccia, A. (2010). AI's War on Manipulation: Are We Winning?. *AI Magazine*, *31*(4), 53–64.

Friedgut, E., Kalai, G., Keller, N., & Nisan, N. (2011). A Quantitative Version of the Gibbard-Satterthwaite Theorem for Three Alternatives. *SIAM J. Comput.*, *40*(3), 934–952.

Friedgut, E., Kalai, G., & Nisan, N. (2008). Elections can be manipulated often. In *Proceedings of the 49th Annual Symposium on Foundations of Computer Science*, pp. 243–249. IEEE.

Fu, Y., & Anderson, P. (1986). Application of statistical mechanics to NP-complete problems in combinatorial optimisation. *Journal of Physics A: Mathematical and General*, *19*, 1605–1620.

Gibbard, A. (1973). Manipulation of Voting Schemes: A General Result. *Econometrica: Journal of the Econometric Society*, 587–601.

Gomes, C., & Walsh, T. (2006). Randomness and Structure. In Rossi, F., van Beek, P., & Walsh, T. (Eds.), *Handbook of Constraint Programming*, Foundations of Artificial Intelligence, pp. 639–664. Elsevier.

Good, I., & Mayer, L. (1975). Estimating the efficacy of a vote. *Behavioral Science*, *20*(1), 25–33.

Isaksson, M., Kindler, G., & Mossel, E. (2012). The Geometry of Manipulation: A Quantitative Proof of the Gibbard-Satterthwaite Theorem. *Combinatorica*, *32*(2), 221–250.

Kelly, J. (1993). Almost all social choice rules are highly manipulable, but a few aren't. *Social Choice and Welfare*, *10*(2), 161–175.

Kemperman, J. (1986). Decomposing $\mathbb{R}^d$ into finitely many semigroups. In *Indagationes Mathematicae (Proceedings)*, Vol. 89, pp. 71–78. Elsevier.

Mossel, E., & Rácz, M. (2012). A quantitative Gibbard-Satterthwaite theorem without neutrality. In *Proceedings of the 44th ACM Symposium on Theory of Computing (STOC)*, pp. 1041–1060. ACM. Full version to appear in *Combinatorica*, available as arXiv preprint at arXiv:1110.5888.

Myatt, D. (2007). On the theory of strategic voting. *The Review of Economic Studies*, *74*(1), 255–281.

Peleg, B. (1979). A note on manipulability of large voting schemes. *Theory and Decision*, *11*(4), 401–412.

Pritchard, G., & Slinko, A. (2006). On the average minimum size of a manipulating coalition. *Social Choice and Welfare*, *27*(2), 263–277.

Pritchard, G., & Wilson, M. (2009). Asymptotics of the minimum manipulating coalition size for positional voting rules under impartial culture behaviour. *Mathematical Social Sciences*, *58*(1), 35–57.

Procaccia, A., & Rosenschein, J. (2007a). Average-case tractability of manipulation in voting via the fraction of manipulators. In *Proceedings of the 6th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pp. 718–720.

Procaccia, A., & Rosenschein, J. (2007b). Junta Distributions and the Average-case Complexity of Manipulating Elections. *Journal of Artificial Intelligence Research*, *28*, 157–181.

Satterthwaite, M. (1975). Strategy-proofness and Arrow's Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions. *Journal of Economic Theory*, *10*(2), 187–217.

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

Walsh, T. (2002). The Interface between P and NP: COL, XOR, NAE, 1-in-k, and Horn SAT. In *Proceedings of the 17th National Conference on AI (AAAI 2002)*, pp. 695–700.

Walsh, T. (2011). Where Are the Hard Manipulation Problems?. *Journal of Artifical Intelligence Research*, *42*, 1–29.

Xia, L. (2012a). Computing the margin of victory for various voting rules. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*, pp. 982–999. ACM.

Xia, L. (2012b). How Many Vote Operations Are Needed to Manipulate A Voting System?. *Arxiv preprint arXiv:1204.1231*.

Xia, L., & Conitzer, V. (2008a). A Sufficient Condition for Voting Rules to be Frequently Manipulable. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC)*, pp. 99–108. ACM.

Xia, L., & Conitzer, V. (2008b). Generalized Scoring Rules and the Frequency of Coalitional Manipulability. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC)*, pp. 109–118. ACM.

Xia, L., & Conitzer, V. (2009). Finite Local Consistency Characterizes Generalized Scoring Rules. In *Proceedings of the 9th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 336–341.

Xia, L., Zuckerman, M., Procaccia, A. D., Conitzer, V., & Rosenschein, J. S. (2009). Complexity of unweighted coalitional manipulation under some common voting rules. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 348–353.

Young, H. (1975). Social choice scoring functions. *SIAM Journal on Applied Mathematics*, 824–838.

Zuckerman, M., Procaccia, A. D., & Rosenschein, J. S. (2009). Algorithms for the coalitional manipulation problem. *Artificial Intelligence*, *173*(2), 392–412.