Statistical Robustness of Voting Rules

Ronald L. Rivest* **and Emily Shen**†

Abstract

We introduce a notion of "statistical robustness" for voting rules. We say that a voting rule is *statistically robust* if the winner for a profile of ballots is most likely to be the winner of any random sample of the profile, for any positive sample size. We show that some voting rules, such as plurality, veto, and random ballot, are statistically robust, while others, such as approval, score voting, Borda, single transferable vote (STV), Copeland, and Maximin are not statistically robust. Furthermore, we show that any positional scoring rule whose scoring vector contains at least three different values (i.e., any positional scoring rule other than t-approval for some t) is not statistically robust.

Keywords: social choice, voting rule, sampling, statistical robustness

1 Introduction

It is well known that polling a sample of voters before an election may yield useful information about the likely outcome of the election, if the sample is large enough and the voters respond honestly.

It is less well known that the effectiveness of a sample in predicting an election outcome also depends on the voting rule (social choice function) used.

We say a voting rule is "statistically robust" if for any profile the winner of any random sample of that profile is most likely to be the same as the (most likely) winner for the complete profile. While the sample result may be "noisy" due to sample variations, if the voting rule is statistically robust the most common winner(s) for a sample will be the same as the winner(s) of the complete profile.

To coin some amusing terminology, we might say that a statistically robust voting rule is "weather resistant"—you expect to get the same election outcome if the election day weather is sunny (when all voters show up at the polls) as you get on a rainy day (when only some fraction of the voters show up). (We assume here that the chance of a voter showing up on a rainy day is independent of his preferences.)

We show that plurality voting is statistically robust, while—perhaps surprisingly—approval voting, STV, and most other familiar voting rules are not statistically robust.

We consider the property of being statistically robust a desirable one for a voting rule, and thus consider lack of such statistical robustness a defect in voting rules. In general, we consider a voting rule to be somewhat defective if applying the voting rule to a sample of the ballots may give misleading guidance regarding the likely winner for the entire profile.

One reason why statistical robustness may be desirable is for "ballot-auditing" (Lindeman and Stark 2012), which attempts to confirm the result of the election by checking that the winner of a sample is the same as the overall winner.

Similarly, in an AI system that combines the recommendations of expert subsystems according to some aggregation rule, it may be of interest to know whether aggregating the recommendations of a sample of the experts is most likely to yield the same result as aggregating the recommendations of all experts. In some situations, some experts may have transient faults or be otherwise temporarily unavailable (in a manner independent of their recommendations) so that only a sample of recommendations is available for aggregation.

Since our definition is new, there is little or no directly related previous work. The closest work may be that of Walsh and Xia (Walsh and Xia 2011), who study various "lot-based" voting rules with respect to their computational resistance to strategic voting. In their terminology, a voting rule of the form "Lottery-Then-X" (a.k.a. "LotThenX") first takes a random sample of the ballots, and then applies voting rule X (where X may be plurality, Borda, etc.) to the sample. Their work is not concerned, as ours is, with the fidelity of the sample winner to the winner for the complete profile. Amar (Amar 1984) proposes actual use of the "random ballot" method. Procaccia et al. (Procaccia, Rosenschein, and Kaminka 2007) study a related but different notion of "robustness" that models the effect of voter errors.

The rest of this paper is organized as follows. Section 2 introduces notation and the voting rules we consider. We define the notion of statistical robustness for a voting rule in Section 3, determine whether several familiar voting rules

*Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139 rivest@mit.edu

†Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139 eshen@csail.mit.edu

are statistically robust in Section 4, and close with some discussion and open questions.

2 Preliminaries

Ballots, Profiles, Alternatives. Assume a profile $P = (B_1, B_2, \ldots, B_n)$ containing n ballots will be used to determine a single winner from a set $\mathcal{A} = \{A_1, A_2, \ldots, A_m\}$ of m alternatives. The form of a ballot depends on the voting rule used. We may view a profile as either a sequence or a multiset; it may contain repeated items (identical ballots).

Social choice functions. Assume that a voting rule (social choice function) f maps profiles to a single outcome (one of the alternatives): for any profile P, $f(P)$ produces the *winner* for the profile P.

We allow f to be randomized, in order for "ties" to be handled reasonably. Our development could alternatively have allowed f to output the *set* of tied winners; we prefer allowing randomization, so that f always outputs a single alternative. In our analysis, however, we do consider the set $\mathbf{ML}(f(P))$ of most likely winners for a given profile.

Thus, we say that A is a "most likely winner" of P if no other alternative is more likely to be $f(P)$. There may be several most likely winners of a profile P. For most profiles and most voting rules, however, we expect f to act deterministically, so there is a single most likely winner.

Often the social choice function f will be *neutral*—symmetric with respect to the alternatives—so that changing the names of the alternatives won't change the outcome distribution of f on any profile. While there is nothing in our development that requires that f be neutral, we shall restrict attention in this paper to neutral social choice functions. Thus, for example, a tie-breaking rule used by f in this paper will not depend on the names of the alternatives; it will pick one of the tied alternatives uniformly at random.

We do assume that social-choice function f is *anonymous*—symmetric with respect to voters: reordering the ballots of a profile leaves the outcome unchanged.

We will consider the following voting rules. (For more details on voting rules, see (Brams and Fishburn 2002), for example.)

Many of the voting rules are preferential voting rules; that is, each B_i gives a linear order $A_{i1} \succ A_{i2} \succ \ldots \succ A_{im}$. (In the rest of the paper, we will omit the $\succ$ symbols and just write $A_{i1}A_{i2}\ldots A_{im}$, for example.)

- A *positional scoring rule* is defined by a vector $\vec{\alpha} = \langle \alpha_1, \ldots, \alpha_m \rangle$; we assume $\alpha_i \geq \alpha_j$ for $i \leq j$.

  Alternative i gets α_j points for every ballot that ranks alternative i in the jth position. The winner is the alternative that receives the most points.

  Some examples of positional scoring rules are:

  Plurality: $\vec{\alpha} = \langle 1, 0, \ldots, 0 \rangle$
  Veto: $\vec{\alpha} = \langle 1, \ldots, 1, 0 \rangle$
  Borda: $\vec{\alpha} = \langle m-1, m-2, \ldots, 0 \rangle$

- *Single-transferable vote (STV)* (also known as *instant-runoff voting (IRV)*): The election proceeds in m rounds. In each round, the alternative with the fewest votes is eliminated. Each ballot is counted as a vote for its highest-ranked alternative that has not yet been eliminated. The winner of the election is the last alternative remaining.
- *Plurality with runoff*: The winner is the winner of the pairwise election between the two alternatives that receive the most first-choice votes.
- *Copeland*: The winner is an alternative that maximizes the number of alternatives it beats in pairwise elections.
- *Maximin*: The winner is an alternative whose lowest score in any pairwise election against another alternative is the greatest among all the alternatives.

Other (non-preferential) voting rules we consider are:

- *Score voting* (also known as *range voting*): Each allowable ballot type is associated with a vector that specifies a score for each alternative. The winner is the alternative that maximizes its total score.
- *Approval* (Brams and Fishburn 1978; Laslier and Sanver 2010): Each ballot gives a score of 1 or 0 to each alternative. The winner is an alternative whose total score is maximized.
- *Random ballot* (Gibbard 1977) (also known as *random dictator*): A single ballot is selected uniformly at random from the profile, and the alternative named on the selected ballot is the winner of the election.

3 Sampling and Statistical Robustness

Sampling. The profile P is the universe from which the sample will be drawn.

We define a *sampling process* to be a randomized function G that takes as input a profile P of size n and an integer parameter k ($1 \leq k \leq n$) and produces as output a sample S of P of expected size k, where S is a subset (or sub-multiset) of P.

We consider three kinds of sampling:

- *Sampling without replacement.* Here $G_{WOR}(P, k)$ produces a set S of size exactly k chosen uniformly without replacement from P.
- *Sampling with replacement.* Here $G_{WR}(P, k)$ produces a multiset S of size exactly k chosen uniformly with replacement from P.
- *Binomial sampling.* Here $G_{BIN}(P, k)$ produces a sample S of expected size k by including each ballot in P in the sample S independently with probability $p = k/n$.

Thus, the output of the voting rule on a sample might be denoted as $f(S)$, or $f(G(P, k))$, depending on the situation.

Statistically Robust Voting Rules. We now give our main definitions.

Definition 1 *If X is a discrete random variable (or more generally, some function defined on a finite sample space), we let $\mathbf{ML}(X)$ denote the set of values that X takes with maximum probability. That is,*

$$\mathbf{ML}(X) = \{x \mid \Pr(X = x) \textit{ is maximum}\}$$

denotes the set of "most likely" possibilities for the value of X.

For example, $\mathbf{ML}(f(P))$ contains the "most likely winner(s)" for a (possibly randomized) voting rule f and profile P; typically this will contain just a single alternative. Similarly, $\mathbf{ML}(f(G(P,k)))$ contains the most likely winner(s) of a sample of size k. Note that $\mathbf{ML}(f(P))$ involves randomization only within f (if any), whereas $\mathbf{ML}(f(G(P,k)))$ also involves the randomization of sampling by G.

Definition 2 *We say that a social choice function f is* statistically robust *for sampling rule G if for any profile P of size n and for any sample size $k \in \{1, 2, ..., n\}$,*

$$\mathbf{ML}(f(G(P,k))) = \mathbf{ML}(f(P)) .$$

That is, an alternative is a most likely winner for a sample of size k if and only if it is a most likely winner for the entire profile P.

Note that this definition works smoothly with ties: if the original profile P was tied (i.e., there is more than one most likely winner of P), then the definition requires that all most likely winners of P have maximum probability of being a winner in a sample (and that no other alternatives will have such maximum probability).

Having a statistically robust voting rule is something like having an "unbiased estimator" in classical statistics. However, we are not interested in estimating some linear combination of the individual elements (as with classical statistics), but rather in knowing which alternative is most likely (i.e., which is the winner), a computation that may be a highly nonlinear function of the ballots.

A simple plurality example. Suppose we have a plurality election with 10 votes: 6 for A, 3 for B, and 1 for C. We try all three sampling methods, all possible values of k, and see how often each alternative is a winner in 1000 trials; Figure 1 reports the results, illustrating the statistical robustness of plurality voting, a fact we prove in Section 4.3.

For brevity in this paper, we will generally assume that the three kinds of sampling will yield equivalent results; we don't expect differences in the results depending on which sampling process is used. (In a longer paper we would not make this assumption.) Thus, to show a method is not statistically robust, it suffices here to show that it is not statistically robust for one of the three sampling methods. However, we do take care to show that plurality is robust under all three sampling methods.

In fact, statistical robustness under sampling without replacement implies statistical robustness under binomial sampling, as shown in the following theorem.

Theorem 1 *If a voting rule f is statistically robust under sampling without replacement, then it is statistically robust under binomial sampling.*

Proof: When binomial sampling returns an empty sample, then, with a neutral tie-breaking rule, no alternative gains any advantage. For non-empty samples, by the assumption of statistical robustness under sampling without replacement, for any $k > 0$, $f(P)$ is the most likely winner of a uniform random sample of size k. Therefore, $f(P)$ is the most likely winner of a sample produced by binomial sampling with any positive probability p. ■

4 Statistical Robustness of Various Voting Rules

In this section, we analyze whether various voting rules are statistically robust.

4.1 Random Ballot

Theorem 2 *The random ballot method is statistically robust, with sampling methods G_{WR}, G_{WOR}, and G_{BIN}.*

Proof: Each ballot is equally likely to be chosen as the one to name the winner. ■

4.2 Score Voting

Theorem 3 *Score voting is not statistically robust.*

Proof: By means of a counterexample. Consider the following profile:

(1) $A_1 : 100, \quad A_2 : 0$
(99) $A_1 : 0, \quad A_2 : 1$

There is one vote that gives scores of 100 for A_1 and 0 for A_2, and 99 votes that gives scores of 0 for A_1 and 1 for A_2.

Then A_1 wins the complete profile.

Under binomial sampling with probability p, A_1 wins with probability about p — that is, with about the probability A_1's vote is included in the sample. (The probability is not exactly p because the binomial sampling may produce an empty sample, in which case A_1 and A_2 will be equally likely to be selected as the winner.)

For $p < 1/2$, A_2 wins more than half the time; thus score voting is not robust under binomial sampling. ■

4.3 Plurality

Throughout this section, we let n_i denote the number of votes alternative A_i receives, with $\sum_i n_i = n$.

Theorem 4 *Plurality voting is statistically robust, with sampling without replacement.*

Proof: Assume $n_1 > n_2 \geq \ldots \geq n_m$, so A_1 is the unique winner of the complete profile. (The proof below can easily be adapted to show that plurality is statistically robust when the complete profile has a tie for the winner.)

Let $K = (k_1, k_2, ..., k_m)$ denote the number of votes for the various alternatives within the sample of size k.

Let $\binom{a}{b}$ denote the binomial coefficient "a choose b." Thus, there are $\binom{n}{k}$ ways to choose a sample of size k from the profile of size n.

The probability of a given configuration K is equal to $\Pr(K) = (\prod_i \binom{n_i}{k_i})/\binom{n}{k}$.

Let $\gamma(i)$ denote the probability that A_i wins the election, and let $\gamma(i, k_{max})$ denote the probability that A_i receives k_{max} votes and wins the election.

G_{WOR}			
k	A	B	C
1	594	303	103
2	625	258	117
3	727	217	56
4	794	206	0
5	838	162	0
6	868	132	0
7	920	80	0
8	1000	0	0
9	1000	0	0
10	1000	0	0

G_{WR}			
k	A	B	C
1	597	299	104
2	569	325	106
3	676	260	64
4	718	256	26
5	749	219	32
6	764	212	24
7	804	181	15
8	818	171	11
9	842	146	12
10	847	150	3

G_{BIN}			
k	A	B	C
1	507	315	178
2	619	277	104
3	698	235	67
4	763	212	25
5	822	161	17
6	879	117	4
7	930	70	0
8	973	27	0
9	993	7	0
10	1000	0	0

Figure 1: Plurality voting with three sampling schemes on a profile with ten votes: 6 for A, 3 for B, and 1 for C. 1000 trials were run for each sample size, with sample sizes running from $k = 1$ to $k = 10$. The entry indicates how many trials each alternative won, with ties broken by uniform random selection. (The perhaps surprisingly large value for C of 178 for G_{BIN} results from the likelihood of an empty sample when $k = 1$; such ties are broken randomly.) Note that $ML(G(P, k)) = A$ for all three sampling methods G and all sample sizes k.

Then $\gamma(i) = \sum_{k_{max}} \gamma(i, k_{max})$, and $\gamma(i, k_{max}) = \sum_{K \in \mathcal{K}} \Pr(K)/\text{Tied}(K)$, where $\mathcal{K}$ is the set of configurations K such that $k_i = k_{max}$ and $k_j \leq k_{max}$ for all $j \neq i$, and Tied(K) is the number of alternatives tied for the maximum score in K. (Note that this equation depends on the tie-breaking rule being neutral.)

For any k_{max}, consider now a particular configuration K used in computing $\gamma(1, k_{max})$: $K = (k_1, k_2, ..., k_m)$, where $k_1 = k_{max}$ and $k_i \leq k_{max}$ for $i > 1$.

Now consider the corresponding configuration K' used in computing $\gamma(2, k_{max})$, where k_1 and k_2 are switched: $K' = (k_2, k_1, k_3, ..., k_m)$.

Each configuration K' used in computing $\gamma(2, k_{max})$ has such a corresponding configuration K used in computing $\gamma(1, k_{max})$.

Then, by Lemma 1 below, $\Pr(K) > \Pr(K')$. Thus, $\gamma(1, k_{max}) > \gamma(2, k_{max})$.

Since $\gamma(1, k_{max}) > \gamma(2, k_{max})$ for any k_{max}, we have that $\gamma(1) > \gamma(2)$; that is, A_1 is more likely to be the winner of the sample than A_2.

By a similar argument, for every $i > 1$, $\gamma(1, k_{max}) \geq \gamma(i, k_{max})$ for any k_{max}, so $\gamma(1) > \gamma(i)$. Therefore, A_1 is the most likely to win the sample. ■

Lemma 1 *If $n_1 > n_2$, $k_1 > k_2$, $n_1 \geq k_1$, and $n_2 \geq k_2$, then $\binom{n_1}{k_1}\binom{n_2}{k_2} > \binom{n_1}{k_2}\binom{n_2}{k_1}$.*

Proof: We wish to show that

$$\binom{n_1}{k_1}\binom{n_2}{k_2} > \binom{n_1}{k_2}\binom{n_2}{k_1}. \tag{1}$$

If $n_2 < k_1$, then $\binom{n_1}{k_2}\binom{n_2}{k_1} = 0$, so (1) is trivially true.

If $n_2 \geq k_1$, then we can rewrite (1) as $\binom{n_1}{k_1}/\binom{n_2}{k_1} > \binom{n_1}{k_2}/\binom{n_2}{k_2}$. So it suffices to show that for $n_1 > n_2$, $\binom{n_1}{k}/\binom{n_2}{k}$ is increasing with k, which is easily verified. ■

Theorem 5 *Plurality voting is statistically robust, under binomial sampling.*

Proof: Follows from Theorems 4 and 1.

Theorem 6 *Plurality is statistically robust, under sampling with replacement.*

Proof: The proof follows the same structure as for sampling without replacement. Again, assume $n_1 > n_2 \geq \ldots \geq n_m$.

For each configuration K used in computing $\gamma(1, k_{max})$ and the corresponding configuration K' used in computing $\gamma(2, k_{max})$, we show that $\Pr(K) > \Pr(K')$.

Under sampling with replacement, the probability of a configuration $K = (k_1, \ldots, k_m)$ is equal to

$$\Pr(K) = \binom{k}{k_1, \ldots, k_m} \prod_{i=1}^{m} \left(\frac{n_i}{n}\right)^{k_i}.$$

For any configuration K used in computing $\gamma(1, k_{max})$, consider the corresponding configuration K', obtained by swapping k_1 and k_2, used in computing in $\gamma(2, k_{max})$: $K' = (k_2, k_1, k_3, \ldots, k_m)$. Then

$$\Pr(K') = \binom{k}{k_1, \ldots, k_m} \left(\frac{n_2}{n}\right)^{k_1} \left(\frac{n_1}{n}\right)^{k_2} \prod_{i=3}^{m} \left(\frac{n_i}{n}\right)^{k_i}.$$

So $\Pr(K) = (n_1/n_2)^{(k_1-k_2)} \Pr(K')$. If $n_1 > n_2$ and $k_1 > k_2$, then $\Pr(K) > \Pr(K')$. If $n_1 > n_2$ and $k_1 = k_2$, then $\Pr(K) = \Pr(K')$. Thus, $\gamma(1, k_{max}) > \gamma(2, k_{max})$ for every k_{max}, and therefore, A_1 is more likely than A_2 to win a sample without replacement.

By a similar argument, for every $i > 1$, $\gamma(1, k_{max}) \geq \gamma(i, k_{max})$ for any k_{max}, so $\gamma(1) > \gamma(i)$. Therefore, A_1 is most likely to win the sample.

■

4.4 Veto

Theorem 7 *Veto is statistically robust.*

Proof: Each ballot can be thought of as a vote for the least-preferred alternative; the winner is the alternative who receives the *fewest* votes.

For plurality, we showed that the alternative who receives the *most* votes in the complete profile is the most likely to receive the most votes for a random sample. By symmetry, the same arguments can be used to show that the alternative who receives the *fewest* votes in the complete profile is the most likely to receive the fewest votes in a random sample. Thus, the winner of a veto election is the most likely to win in a random sample. ■

4.5 Approval Voting

We were surprised to discover the following.

Theorem 8 *Approval voting is* not *statistically robust.*

Proof: Proof by counterexample. Consider the following profile:

(r) $\{A_1\}$
(r) $\{A_2, A_3\}$

There are r ballots that approve of A_1 only and r ballots that approve of A_2 and A_3. Each alternative receives r votes, and each wins the election with probability $1/3$.

However, in a sample of size 1, A_1 wins with probability $1/2$, while A_2 and A_3 each win with probability $1/4$.

Similarly, in a sample without replacement of size $n-1$, A_1 wins with probability $1/2$ (when one of the ballots for $\{A_2, A_3\}$ is the ballot excluded from the sample), while A_2 and A_3 each win with probability $1/4$. ■

Note that the example above shows not only that approval is not (fully) statistically robust, but also that there does not exist a threshold τ such that for any sample of size at least τ-fraction of n, approval is statistically robust.

4.6 Borda

Theorem 9 *Borda voting is* not *statistically robust.*

Proof: Proof by counterexample. Consider the following profile:

(n_1) $A_1\ A_2\ A_3$
(n_2) $A_2\ A_3\ A_1$
(n_3) $A_3\ A_1\ A_2$

Suppose $n_1 > n_2$ and $n_1 > n_3$. Then in a sample of size 1, each A_i wins with probability n_i/n, and A_1 is the most likely winner.

In the complete profile, A_1 gets a Borda score of $2n_1+n_3$, A_2 gets $2n_2+n_1$, and A_3 gets $2n_3+n_2$. If $2n_2-n_3 > n_1$ (e.g., $n_1 = 100, n_2 = 70, n_3 = 30$), then A_2 beats A_1 in the complete profile.

Thus, Borda is not statistically robust with sampling with or without replacement. ■

Borda is a special case of positional scoring rules. Section 4.10 shows more generally that any positional scoring rule whose scoring vector contains at least 3 distinct values is not robust.

4.7 Single Transferable Vote (STV)

Theorem 10 *STV is* not *statistically robust.*

Proof:

We give two proofs that STV is not statistically robust.

Proof 1: For a sample of size 1, the most likely winner will be the alternative with the most first-choice votes in the complete profile. However, it is well known that STV does not always elect the alternative with the most first-choice votes.

Proof 2: We give a sketch of a counterexample with a larger sample size k.

We construct a profile for which the winner is very unlikely to be the winner in any smaller sample.

Choose m (the number of alternatives) and r (a "replication factor") both as large integers.

The profile will consist of $n = mr$ ballots:

$(r+1)$ $A_1\ A_m \ldots$
(r) $A_2\ A_m\ A_1 \ldots$
(r) $A_3\ A_m\ A_1 \ldots$
...
(r) $A_{m-1}\ A_m\ A_1 \ldots$
$(r-1)$ $A_m\ A_1 \ldots$

where the specified alternatives appear at the front of the ballots, and "..." indicates that the order of the other lower-ranked alternatives is irrelevant.

In this profile, A_m is eliminated first, then $A_2, \ldots, A_{m-1}$ in some order, until A_1 wins.

Suppose now that binomial sampling is performed, with each ballot retained with probability p. Let n_i be the number of ballots retained that list A_i first. Each n_i is a binomial random variable with mean $(r+1)p$ (for A_1), rp (for $A_2 \ldots A_{m-1}$), and $(r-1)p$ (for A_m).

Claim 1: The probability that $n_m = 0$ is effectively zero, for any fixed p, as $r \to \infty$.

Claim 2: The probability that there exists an i, $1 \le i < m$, such that $n_i < n_m$ goes to 1 as $m, r \to \infty$.

Note that as r gets large, then n_i and n_m are very nearly identically distributed, so the probability that $n_i < n_m$ goes to $1/2$. The probability that *some* n_i will be smaller than n_m goes to 1.

Thus, in any sample $G_{BIN}(P, k)$, we expect to see some A_i other than A_m eliminated first. Since all of A_i's votes then go to A_m, A_m will with high probability never be eliminated, and will be the winner. ■

4.8 Plurality with Runoff

Theorem 11 *Plurality with runoff is* not *statistically robust.*

Proof: Proof by counterexample. Consider the following profile:

(n_1) $A_1\ A_2\ A_3$
(n_2) $A_2\ A_3\ A_1$
(n_3) $A_3\ A_2\ A_1$

with $n_1 > n_2, n_1 > n_3$, and $n_2+n_3 > n_1$. Then A_1 is most likely to win a sample of size 1, but A_2 wins the complete election. ■

4.9 Copeland and Maximin

Theorem 12 *Copeland and Maximin are* not *statistically robust.*

Proof: Proof by counterexample:

(n_1) $A_1\ A_2\ A_3$
(n_2) $A_2\ A_1\ A_3$
(n_3) $A_3\ A_1\ A_2$

Suppose $n_1 + n_3 > n_2$ and $n_1 + n_2 > n_3$, and $n_2 > n_1$ and $n_2 > n_3$. (For example, let $n_1 = 30, n_2 = 40, n_3 =$

20.) Then A_1 is the Condorcet winner (and therefore the Copeland and maximin winner), but A_2 is the most likely winner of a sample of size 1. ■

4.10 Positional Scoring Rules

Theorem 13 *Let $\vec{\alpha} = \langle \alpha_1, \ldots, \alpha_m \rangle$ be any positional scoring rule with integer α_i's such that $\alpha_1 > \alpha_i > \alpha_m$ for some $1 \leq i \leq m$. Then the positional scoring rule defined by $\vec{\alpha}$ is* not *robust.*

Proof: We will show that a counterexample exists for any $\vec{\alpha}$ for which $\alpha_1 > \alpha_i > \alpha_m$ for some i.

We construct a profile as follows. Start with r copies of each of the $m!$ possible ballots, for some large r. Clearly, for this profile, all m alternatives are tied, and each alternative wins the election with equal probability, $1/m$.

We will show that this profile can be "tweaked" so that in the resulting profile, all m alternatives are again tied and each win with equal probability. However, the number of first-choice votes will no longer be equal for all alternatives, so for a sample of size 1, the alternatives will not all win with equal probability.

The "tweak" is performed as follows. Take a single ballot type b (i.e., a permutation of the alternatives) that has A_1 in position 1 on the ballot, A_2 in position i, and A_3 in position m. Consider the 6 ballot types obtained by permuting A_1, A_2, A_3 within b (while keeping the other alternatives' positions fixed). We will change the number of ballots of each of these 6 types by $\delta_1, \ldots, \delta_6$ (the δ_i's may be positive, negative, or zero).

That is, starting from a ballot type $A_1[\ldots]A_2[\ldots]A_3$, we will change the counts of the following 6 ballot types:

$A_1[\ldots]A_2[\ldots]A_3$ by δ_1
$A_1[\ldots]A_3[\ldots]A_2$ by δ_2
$A_2[\ldots]A_1[\ldots]A_3$ by δ_3
$A_2[\ldots]A_3[\ldots]A_1$ by δ_4
$A_3[\ldots]A_1[\ldots]A_2$ by δ_5
$A_3[\ldots]A_2[\ldots]A_1$ by δ_6

where the "[...]" parts are the same for all 6 ballot types.

In order to keep the scores of $A_4, \ldots, A_m$ unchanged, we require $\delta_1 + \ldots + \delta_6 = 0$.

Next, in order to keep the scores of $A_1, \ldots, A_3$ unchanged, we write one equation for each of the three alternatives:

$$(\delta_1 + \delta_2)\alpha_1 + (\delta_3 + \delta_5)\alpha_i + (\delta_4 + \delta_6)\alpha_m = 0,$$
$$(\delta_3 + \delta_4)\alpha_1 + (\delta_1 + \delta_6)\alpha_i + (\delta_2 + \delta_5)\alpha_m = 0,$$
$$(\delta_5 + \delta_6)\alpha_1 + (\delta_2 + \delta_4)\alpha_i + (\delta_1 + \delta_3)\alpha_m = 0.$$

Finally, to ensure that the number of first-choice votes changes (so that the probability of winning a sample of size 1 changes) for at least one of A_1, A_2, A_3, we add an additional equation, $\delta_1 + \delta_2 = 1$, for example.

The 5 equations above in 6 variables will always be satisfiable with integer δ_i's (details omitted). We can choose the replication factor r to be large enough so that the numbers of each ballot type are non-negative. Thus, there always exists a counterexample to statistical robustness as long as $\alpha_1 > \alpha_i > \alpha_m$. ■

Note that this theorem implies the specific case of Borda given in Section 4.6.

5 Discussion and Open Questions

We have introduced and motivated a new property for voting rules: "statistical robustness," and provided an initial suite of results on the statistical robustness of several well-known voting rules.

The research reported here represents only the first steps towards a full understanding of the statistical robustness of voting rules, however, and many interesting open problems remain, some of which are given below.

It is perhaps surprising that plurality (and its complement, veto) and random ballot are the only interesting voting rules that appear to be statistically robust. Being statistically robust seems to be a somewhat fragile property, and a small amount of nonlinearity appears to destroy it.

For example, even plurality with weighted ballots (which one might have in an expert system with different experts having different weights) is not statistically robust: this is effectively the same as score voting.

Open Problem 1 *Do some voting rules become statistically robust for large enough sample sizes? For each interesting voting rule, and each kind of sampling, determine precisely for which values of k it is statistically robust. (Many of our proofs only look at the simple case $k = 1$.) For which voting rules is there a "threshold" $\tau(n) < n$ such that the voting rule is statistically robust for $k \geq \tau(n)$?*

We note that we can easily show that for approval and score voting, there does not exist such a threshold $\tau(n)$.

Open Problem 2 *Determine how the property of being (or not being) statistically robust relates to other well-studied voting rule properties.*

Conjecture 1 *Show that a voting rule cannot be statistically robust if the number of distinct meaningfully-different ballot types is greater than m, the number of alternatives.*

Conjecture 2 *Show that plurality and veto are the only statistically robust voting rules among those where each ballot "approves t" for some fixed t.*

Conjecture 3 *Show that a score voting rule cannot be robust if there are two profiles P and P' that have the same total score vectors, but which generate different distributions when sampled.*

Open Problem 3 *Determine how best to utilize the information contained in a sample of ballots to predict the overall election outcome, when the voting rule is not statistically robust. (There may be something better to do than merely applying the voting rule to the sample.)*

Open Problem 4 *The voting rules studied by Walsh and Xia (Walsh and Xia 2011) of the form "Lottery-Then-X" seem plausible alternatives for statistically robust voting rules, since their first step is to perform a lottery (take a sample of the profile). Determine which, if any, Lottery-Then-X voting rules are statistically robust.*

Acknowledgments

This work was supported in part by the Center for Science of Information (CSoI), an NSF Science and Technology Center, under grant agreement CCF-0939370.

References

Amar, A. R. 1984. Choosing representatives by lottery voting. *Yale Law Journal* 93 Yale L. J. 1283.

Brams, S., and Fishburn, P. 1978. Approval voting. *American Political Science Review* 72(3):831–847.

Brams, S. J., and Fishburn, P. C. 2002. Voting procedures. In Arrow, K. J.; Sen, A. K.; and Suzumura, K., eds., *Handbook of Social Choice and Welfare*. North-Holland. chapter 4.

Gibbard, A. 1977. Manipulation of schemes that mix voting with chance. *Econometrica* 45(3):665–681.

Laslier, J.-F., and Sanver, M. R., eds. 2010. *Handbook of Approval Voting*. Studies in Choice and Welfare. Springer.

Lindeman, M., and Stark, P. B. 2012. A gentle introduction to risk-limiting audits. *IEEE Security and Privacy* submitted. See `http://statistics.berkeley.edu/~stark/Preprints/gentle11.pdf`.

Procaccia, A. D.; Rosenschein, J. S.; and Kaminka, G. A. 2007. On the robustness of preference aggregation in noisy environments. In *Proceedings AAMAS'07*, 416–422.

Walsh, T., and Xia, L. 2011. Venetian elections and lot-based voting rules. In *IJCAI Workshop on Social Choice and Artificial Intelligence*.