

TCER Working Paper Series

A FRAMEWORK FOR THE QUANTITATIVE EVALUATION OF VOTING RULES

Michael Munie

Yoav Shoham

February 2011

Working Paper E-27

<http://tcer.or.jp/wp/pdf/e27.pdf>



TOKYO CENTER FOR ECONOMIC RESEARCH
1-7-10-703 Iidabashi, Chiyoda-ku, Tokyo 102-0072, Japan

©2011 by Michael Munie and Yoav Shoham.

All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including ©notice, is given to the source.

Abstract

In order to characterize the set of desirable social choice functions, researchers have proposed axioms that all social choice functions should satisfy. However, it has been shown that achieving these axioms is impossible. Instead of viewing this impossibility result as a limitation, it can be viewed as an opportunity. In this paper, we develop a means of comparing various social choice functions with regard to a desired axiom by quantifying how often the axiom is violated. To this end, we offer a new framework for measuring the quality of social choice functions that builds from and provides a unifying framework for previous research. This framework takes the form of what we call a “violation graph.” Graph properties have natural interpretations as metrics for comparing social choice functions. Using the violation graph we present new metrics, such as the minimal domain restriction, for assessing social choice functions and provide exact and probabilistic results for voting rules including plurality, Borda, and Copeland. Motivated by the empirical results, we also prove asymptotic results for scoring voting rules. First, these results suggest that voting rules based on pairwise comparison (ex: Copeland) are better than scoring rules (ex: Borda count). Second, these results also suggest that although we can never fulfill our desired set of axioms, the frequency of violation is so small that with even a modest number of voters we can expect to never violate our axioms.

Michael Munie
Stanford University
Computer Science Department
*
munie@stanford.edu

Yoav Shoham
Stanford University
Computer Science Department
*
shoham@stanford.edu

A Framework for the Quantitative Evaluation of Voting Rules

Michael Munie
Computer Science Department
Stanford University, CA
munie@stanford.edu

Pingzhong Tang
Computer Science Department
Hong Kong University
of Science and Technology
Clear Water Bay, Hong Kong
kenshin@cse.ust.hk

Yoav Shoham
Computer Science Department
Stanford University, CA
shoham@stanford.edu

May 28, 2009

Abstract

In order to characterize the set of desirable social choice functions, researchers have proposed axioms that all social choice functions should satisfy. However, it has been shown that achieving these axioms is impossible. Instead of viewing this impossibility result as a limitation, it can be viewed as an opportunity. In this paper, we develop a means of comparing various social choice functions with regard to a desired axiom by quantifying how often the axiom is violated.

To this end, we offer a new framework for measuring the quality of social choice functions that builds from and provides a unifying framework for previous research. This framework takes the form of what we call a “violation graph.” Graph properties have natural interpretations as metrics for comparing social choice functions. Using the violation graph we present new metrics, such as the minimal domain restriction, for assessing social choice functions and provide exact and probabilistic results for voting rules including plurality, Borda, and Copeland. Motivated by the empirical results, we also prove asymptotic results for scoring voting rules.

First, these results suggest that voting rules based on pairwise comparison (ex: Copeland) are better than scoring rules (ex: Borda count). Second, these results also suggest that although we can never fulfill our desired set of axioms, the frequency of violation is so small that with even a modest number of voters we can expect to never violate our axioms.

1 Introduction

Social choice theory is concerned with the extent to which individual preferences can be aggregated into a social decision in a satisfactory manner. The idea of preference aggregation can be implemented by designing a *social choice function*, which takes a preference profile (a collection of the orderings of each individual over alternatives that express the preference of that individual) as input and returns a socially optimal alternative. The satisfaction of our goal to aggregate individual preferences is then described by the fulfillment of a set of socially desirable conditions on the social choice function.

Unfortunately, social choice theory is sometimes called a “science of the impossible” because of the domination of various impossibility theorems. For example, the Muller-Satterthwaite theorem [20] states that every social choice function that is weakly Pareto efficient and monotonic has a dictator, whose preferred alternative is always chosen by the social choice function. One can test by a small computer program [17] that out of 3^{36} possible social choice functions with 2 voters and 3 candidates, there are only 17 monotonic functions, all of which have extremely unsatisfactory properties such as dictatorship or constant output.

One perspective is that monotonicity is too strong and should be dropped as a desirable conditions. However, the strength of this condition can be viewed as a positive or a negative, depending on if the framework being used is qualitative or quantitative. Qualitatively, we are unable to differentiate between the quality of social choice functions because they all violate our criteria. In our quantitative work, the strength is a positive because it better allows us to discriminate between social choice functions which we use to ask to what extent do different social choice functions violate this condition.

We offer a new framework for quantitatively evaluating and comparing social choice functions. There are three key dimensions on which our evaluation depends: the type of violation we seek to quantify, the specific metric by which we measure the violation, and the population distribution over which the violation exists. We choose the first dimension to be monotonicity here and will motivate this choice in the next section. As we will discuss in the paper, our three dimensions have simple and natural graph interpretations. By introducing a model called the “violation graph,” we provide a means of easily varying the dimensions of the problem while applying the same algorithms to produce the computational results.

We then demonstrate our framework using different voting rules as target social choice functions and compare them via the aforementioned dimensions using violation graphs. For small and medium sized domains, we provide both exact and sampled results. Interestingly, these results show that (among other things) voting rules based on pairwise comparison, such as Copeland and Maxmin, perform better than score-based voting rules, such as plurality and Borda. Motivated by these empirical results, we also prove certain asymptotic results that help us better understand the satisfaction of monotonicity for scoring rules when the number of voters tends to infinity.

The structure of our paper is as follows. In Section 2, we discuss the foundational material needed to proceed. In Sections 3 and 4 we review the related work in social choice theory and formally present the violation graph approach. Then in Section 5, we provide empirical results and algorithms. In Section 6 we prove asymptotic results

that support these findings. Finally, we conclude and discuss future work in Section 7.

2 Background

To lay a foundation, we will first briefly cover the necessary material on social choice.

2.1 Social Choice Functions: Maskin Monotonicity versus Strategy Proofness

Formally, let $N = \{1, 2, \dots, n\}$ denote a set of voters, and O denote a finite set of alternatives. Let L be the set of strict total orders over O and denote voter i 's preference as \succ_i . A social choice function is a function $C : L^n \mapsto O$.

Not all social choice functions are equally desirable. There are several principles that social choice theorists have argued that social choice functions should ideally adhere to. As we mentioned, Maskin's monotonicity (sometimes called "strong monotonicity") is such a desirable property. It states that when a social choice function chooses a candidate based on some preference profile of all voters, this candidate should remain the selection under preference changes where the winner's position does not fall in each voter's profile. More formally,

Definition 1 (*Maskin Monotonicity*) *A social choice function C is Maskin monotonic if for any $o \in O$ and any preference profile $\succ \in L^n$ with $C(\succ) = o$, then for any other preference profile \succ' with the property that $\forall i \in N, \forall o' \in O, o \succ'_i o'$ if $o \succ_i o'$, it must be that $C(\succ') = o$.*

There are two critical points to note about our use of Maskin's Monotonicity. First, as we mentioned in the introduction, all non-trivial social choice functions violate Maskin monotonicity at some profile. Although it may seem that the requirement of Maskin monotonicity is too strict, it is precisely the strictness that gives us great power of discrimination. If we instead used a weaker requirement like standard monotonicity [19], many social choice functions would not violate this property at all, giving us limited power to speak of the differences between voting rules. Secondly, Maskin's monotonicity also plays an important role in implementation theory [18] in that it serves as a necessary condition of Nash implementability.

A closely related condition is strategy-proofness.

Definition 2 *A social choice function C is manipulable at profile \succ by individual i via \succ'_i if $C(\succ_{-i}, \succ'_i) \succ_i C(\succ)$, where (\succ_{-i}, \succ'_i) is the profile resulting from replacing \succ_i with \succ'_i in \succ .*

A social choice function is *strategy-proof* if it is not manipulable by any individual at any profile.

It is well known [20] that strategy-proofness and Maskin monotonicity are equivalent in unrestricted domains and that strategy-proofness is stronger in restricted domain. Strategy-proofness is a reasonable property on the basis of which to compare

social choice functions, and in the next section we discuss the previous work on measuring its violation. However, Maskin Monotonicity is no less good of a candidate, and it is important to realize that it is a different diagnostic tool because the relationship of the number of profiles violating these different properties is unknown. Maskin monotonicity is especially crucial to study given its importance in social choice theory and implementation theory as we have mentioned.

Before we continue, we will describe the voting rules we will later analyze.

2.2 Voting Rules

The specific voting rules we consider in this paper are Plurality, Borda, Copeland, Maxmin, and Plurality with Runoff.

- **Positional scoring rules:** Let s_j be the point value of an outcome ranked in position j and refer to $s = (s_1, s_2, \dots, s_m)$ as a scoring vector. Let $d(\succ_i, o) = s_{p(\succ_i, o)}$, where $p(\succ_i, o)$ is the position of an outcome o in an individual preference \succ_i , and for preference profile $\succ = (\succ_1, \dots, \succ_n)$, let $d(\succ, o) = \sum_{i \in N} d(\succ_i, o)$. The winner is the outcome o^* that maximizes $d(\succ, o)$.
 - **Plurality** is a positional scoring rule with scoring vector $(1, 0, \dots, 0)$.
 - **Borda** is a positional scoring rule with scoring vector $(m-1, m-2, \dots, 0)$.
- **Plurality with Runoff** is a variation of plurality where the plurality scoring vector is used to determine the two outcomes with the highest scores. Then, after all other outcomes but the top two have been removed from all the preference orderings, the plurality scoring rule is used again to determine the winner.
- **Copeland** is a pairwise comparison based voting rule where for each pair of outcomes the plurality rule is used to determine the winner of this pairwise comparison after removing all outcomes but the two in question. Each outcome gets one point for every pairwise comparison it wins and loses one point for each comparison it loses. The winner is the outcome with highest score.
- **Maxmin** is a voting rule based on pairwise comparison. Let $N(o, o')$ be the number of voters that prefer o to o' and the score of outcome o , equals $\min_{o' \in O \setminus o} N(o, o')$. The winner is then the outcome with the highest score.

We assume alternatives can be ordered lexicographically and that ties are broken lexicographically for all voting rules.

3 Related Work

Social choice functions are qualitatively characterized¹ by two strong and well known results, the Gibbard-Satterthwaite [12, 27] and the Muller-Satterthwaite [20] theorems. The Gibbard-Satterthwaite theorem states that for every social choice function that

¹Social welfare functions were similarly characterized by Arrow's [1] seminal paper.

is non-dictatorial and whose range has at least three alternatives, there exists at least one profile that can be manipulated by a single voter. In a similar result, the Muller-Satterthwaite theorem states that for every social choice function that is weakly Pareto efficient and non-dictatorial there is some profile where a monotonic change with respect to the winner would cause that alternative to lose.

While there has been research by Fedrizzi [9] and others on comparing voting systems by determining which of a set of desirable requirements for social choice functions are satisfied, the previous work most relevant to this paper has focused on quantifying the proportion of profiles that are manipulable. Although we focus on monotonicity instead of manipulability, the quantitative nature of this work is related. The first conjecture that the proportion of manipulable profiles was small was made by Pattanaik in [24]. An asymptotic bound on unstable profiles under the plurality rule was proved by Peleg in [25], and this work was extended and sharpened by Frisrup and Kleiding [11] and also Slinko [28] where the proportion of manipulable profiles was shown to be $\Theta(\frac{1}{\sqrt{n}})$, with n being the number of voters and the constant factor depending on the number of candidates m . In a similar line of work, Nurmi [23] focuses on quantifying how often voting rules select the same winners and how often the set of winners overlap.

Other research approached the problem using computer simulations, with early work by Nitzan [22] studying situations with up to 90 voters on plurality and Borda count by estimating the proportion of manipulable profiles. By looking at three metrics in addition to the proportion of manipulable profiles, Smith [29] simulated common social choice functions looking for strategy-proofness violations. A slightly different vein of work by Kelly [14, 15] examined the distribution of manipulable profile proportions for random social choice functions that were onto and non-dictatorial and, using computer simulation, found that the average proportion of manipulable profiles was almost one.

Much of the related work has assumed that each profile is equally likely (referred to as the Impartial Culture (IC) assumption) and thus interprets this proportion as a probability, as we do in this paper. However, other work has assumed that the identities of voters are anonymous as well as impartial (IAC) giving us a different distribution over profiles. Working in this distribution, Favardin [8, 7] examines both individual and coalitional manipulation (while also allowing for counter-threats) and finds that Borda count is more vulnerable to strategic manipulation than Copeland's rule. Also considering coalitional manipulation, Lepelley [16] generalizes the special cases of IAC and IC to the space of Polya-Eggenberger probability models and runs simulations to calculate violation probabilities. Slinko [26] evaluates some voting rules including approval voting, Borda, and plurality, by looking at the asymptotic average threshold coalition size and provides analytic results.

Because the domain of all profiles violates strategy-proofness, there is a rich body of work by Kelly and others [4, 2, 21] among others, looking at various restricted domains where strategy-proofness holds. Simple majority voting has been shown by Maskin [5] to be strategy-proof in the restricted domain where the Condorcet winner exists, but the Condorcet winner does not always exist, and the probability of its existence even tends towards 0 as the number of alternatives increases, as shown

in Fishburn [10]. Recent work by Bochet [3] gives the minimal domain restriction for a non-dictatorial, Pareto optimal, and strategy-proof (or Maskin monotonic) social choice function to exist. However, the social choice functions that satisfy these requirements in this minimal domain have a very strong dictatorial flavor. A different approach is to take a known domain where strategy-proofness holds and calculate the minimal monotonic extension as in Thomson [30] and Erden [6].

We focus on completing the picture in light of previous research. While most previous work has concentrated on strategy-proofness, we focus on Maskin’s monotonicity and quantifying how much it is violated. In addition to using metrics used in previous work, we also introduce the metric of domain restriction for measuring the quality of a social choice function.

4 Violation Graph

The violation graph is a natural tool for comparing different social choice functions.

Definition 3 (*Violation Graph*) A violation graph is a tuple $M = (C, A, G)$. C is a social choice function. A , the violation type, is a function $V \times V \times C \mapsto \{True, False\}$ that encodes our desired violation property (such as Maskin’s monotonicity) and is true exactly when the two nodes and social choice function C form a violation. $G = (V, E)$ is a graph. V is the set of all possible preference profiles, and E is the set of edges² such that if $x, y \in V$, and $A(x, y, C)$ is true, then the edge $(x, y) \in E$.

We use three metrics to quantify how much a social choice function violates the property A . The first is the proportion of nodes with degree at least one. If we assume a uniform distribution over profiles, this is the probability that a random profile is involved in a violation of A . The second metric we use is the edge ratio, i.e. the ratio of the edges in the graph to the number of edges in a fully connected graph with the same number of nodes. This can be viewed as the probability that two successive elections would provide a violation of A , or in other words, would publicly demonstrate the flaw in our voting rule. The third metric we use, the minimal domain restriction required to make the violation graph disconnected, to our knowledge has not been used before. This value shows how many profiles are truly causing the violations but is a difficult metric to compute because it requires solving the NP-complete problem of independent set [13], which has no constant-factor approximation unless $P=NP$.

5 Empirical Results

To compute our desired metrics (edge ratio, node ratio, and domain restriction) on the violation graph, we apply two different computational strategies: exact computation and sampling. We purposely keep our metrics and model simple because of the computational challenges of generating the violation graph. With m alternatives and n

²We also note that the edges in our violation graph can be extended to hyper edges for axioms that require more than two profiles to provide an instance of a violation.

voters, the graph grows as $O(m!^{2n})$. For everything beyond small examples this is not feasible. However, to work around this problem, we sample the space and also prove asymptotic results in the next section.

We compute the edge ratio, node ratio, and domain restriction heuristics exactly when there are three candidates and seven or fewer voters. For larger numbers of voters, we generate 1,000,000 pairs of random profiles, and for each pair, we check for an edge between the two profiles. Checking for edges is a quick operation and we are only limited by the number of edges we wish to sample. Sampling nodes is more difficult because after first choosing a random node, we must check all other nodes to see if an edge exists. This is bounded by $O(m!^n)$. To make the process more efficient, we stop checking for neighbors of a node once we find the first one. As we will see in the next section, since the edges are sparse, we check most of the nodes. Because of this, the node ratio is computed with 1000 samples.

In order to compute the domain restriction, we tested four different heuristics for domain restriction (minimum degree, maximum degree, random edge, and random node) because the NP-Complete nature of this problem made it impossible to compute exactly. In our results section, we include the results from the best heuristic—the minimum degree. This heuristic iteratively removes all the neighbors of the node with the minimum degree. The results of this heuristic are guaranteed to be an upper bound for the optimal minimum domain restriction.

To give the reader a sense of what the violation graphs look like and their wildly varying structure, we include a sample of them for three candidates and two, three, and four voters, in Figures 1, 2 and 3. Note that we only depict the nodes with degree one or greater. We include these small cases since the graphs are suggestive of emerging structures. However, our experimental results that follow refer to the coarser measures of domain restriction and edge and node ratios.

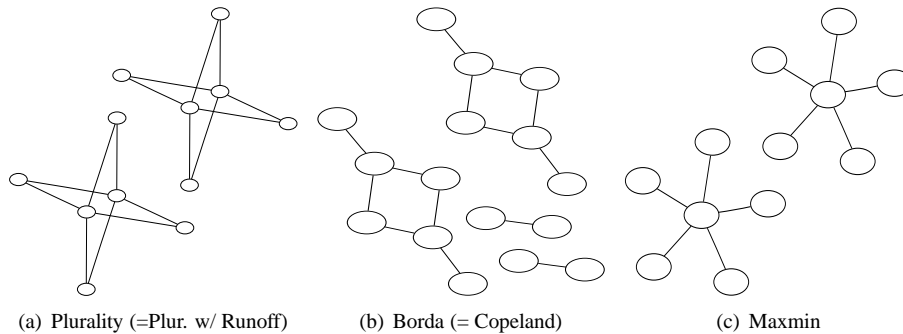


Figure 1: Violation Graphs for Two Voters and Three Candidates

5.1 Edge Ratio Results

In Figure 4, we graph the comparative behavior of the voting rules we tested. There are two key take-aways from the graphs. First, the evidence suggests that as the number




Figure 2: Violation Graphs for Three Voters and Three Candidates

of voters grows, the edge ratio goes to zero for all the voting rules we tested. As we will prove for scoring rules in Section 6, this ratio does in fact converge to zero as the number of voters grows to infinity. The second key point is that pairwise voting rules (Copeland and Maxmin) perform better than scoring rules (Plurality, Borda, and Plurality with Runoff) for all numbers of voters. Between the pairwise rules, Maxmin performs better than Copeland by a small amount. We do not show the graphs for other numbers of candidates due to space reasons, but these relationships also hold for candidates up to at least 20.

In addition, the data also shows another relationship. When fixing the number of voters as we increase the number of candidates, the edge ratio also goes to zero. In Figure 5, the edge ratio converges, but it does so at a slower rate than when we fix the number of candidates and increase the number of voters.

5.2 Node Ratio Results

In Figure 6, we graph the node ratio for Borda, Plurality, Plurality with run-off, Copeland, and Maxmin. In contrast to our edge results, as the number of voters increases, the node ratio increases. A more striking comparison is that in previous work on the node ratio for strategy-proofness it was shown empirically and theoretically that the node ratio goes to zero as the number of voters increases. From this we can see that even though the existence of at least one violation of strategy-proofness is equivalent to the exist-

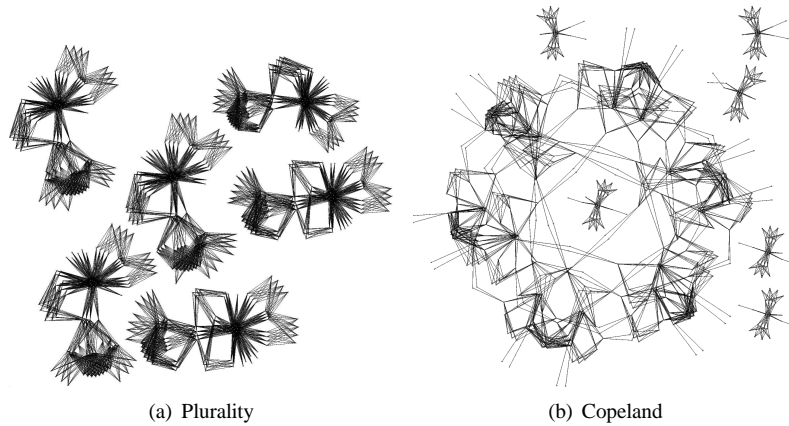


Figure 3: Violation Graphs for Four Voters and Three Candidates

tence of at least one violation of monotonicity on the unrestricted domain, the actual node ratio for the two axioms is different, and may converge to opposite values. We will see in the next section a closer look into the meaning of the node ratio by means of the domain restriction heuristic.


We also note that according to the node ratio, pairwise rules are better than scoring rules. This is not surprising since pairwise rules such as Copeland and Maxmin are well-known as extensions of *Condorcet procedures*, which choose the Condorcet winner when it exists. One can verify that among any pair of profiles where Condorcet winners exist, there are no monotonic violation.

5.3 Domain Restriction Results

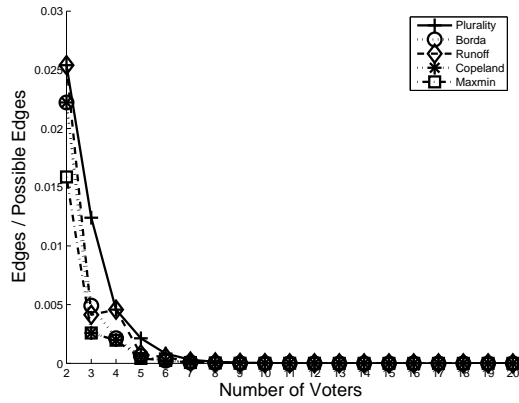
Domain restriction is the final metric we will examine. This allows us to refine our results on the node ratio. If a voting rule has a high node ratio but only very few nodes need to be removed to disconnect the graph (restrict the domain so that no violations occur), then the high node ratio is an artifact of a few central and well-connected nodes. If we can prohibit these profiles from occurring in practice, or we can rule them out due to low probability, a voting rule with a high node ratio but a low domain restriction might be considered a good voting rule.

Looking at the data, we find that pairwise rules are better than scoring rules, as can be seen in Figure 7. Although the node ratio of Copeland is lower than Maxmin, the structure of the graphs are such that Maxmin scores lower than Copeland using our domain restriction techniques. This can be seen in a small example by looking at the violation graphs of Maxmin in Figure 1(c), where by only deleting the two centered nodes, all the violation edges are eliminated.

Most importantly, even though the node ratio increases rapidly as the number of voters increases, most nodes do not need to be removed from the graph, and the proportion of nodes that needs to be removed tends toward zero for all the voting rules we tested.



(a) Up to 6 voters



(b) Up to 20 voters

Figure 4: Comparative Edge Ratio Results for Three Candidates (Two views of the same data)

6 Asymptotic Results

We can see from our experimental results that by fixing the number of candidates and increasing the number of voters, the ratio of violation edges decreases. One may start to wonder: does this ratio converge, and if so, what is the limit? We prove in this section that this ratio converges to 0 when $n \rightarrow \infty$.

Suppose the number of violation edges for some scoring rules is $VE(n, m)$ and the number of violation nodes is $VN(n, m)$, where the number of all possible edges and nodes is $E(n, m)$ and $N(n, m)$.

Our start point is a known result of Slinko[28] concerning the number of unstable³

³A profile is unstable if a unilateral deviation from one voter can lead to the change of outcome.

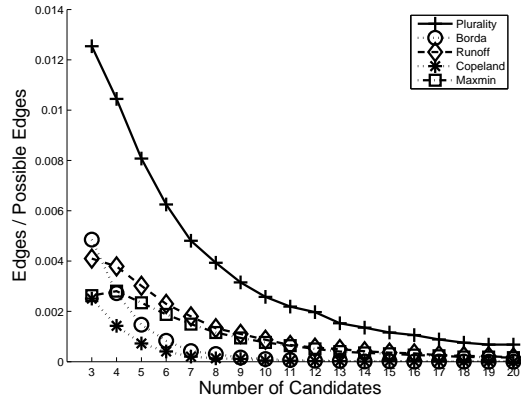


Figure 5: Comparative Edge Ratio Results for Three Voters

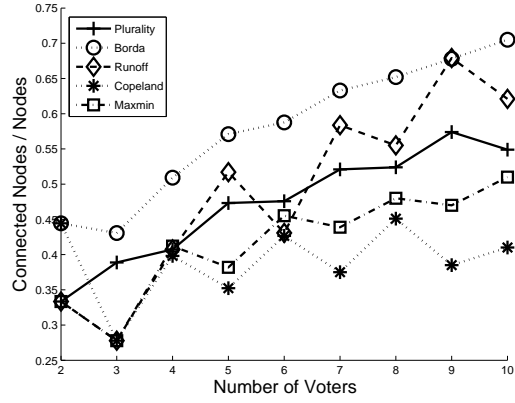


Figure 6: Comparative Node Ratio Results for Three Candidates

profiles $U(n, m)$. Let $L(n, m) = \frac{U(n, m)}{N(n, m)}$.

Theorem 1 (from Slinko [28])

$$\frac{d_m}{\sqrt{n}} \leq L(n, m) \leq \frac{D_m}{\sqrt{n}},$$

where d_m and D_m are constants that depend only on m .

Theorem 2 When $n \rightarrow \infty$, we have

- $\lim_{n \rightarrow \infty} \frac{VE(n, m)}{E(n, m)} = 0$,
- $\lim_{n \rightarrow \infty} \frac{VN(n, m)}{N(n, m)} = c$, where $0 \leq c \leq 1$.

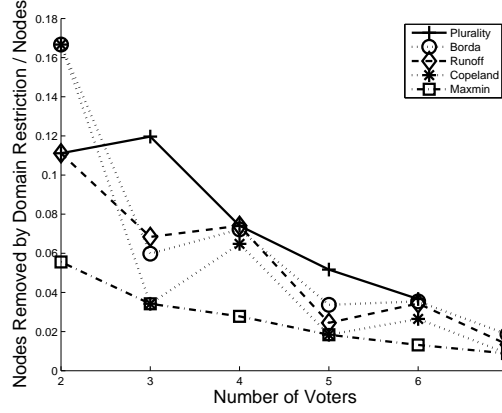


Figure 7: Proportion of Nodes Removed with Domain Restriction for Three Candidates

Let $C(n, m)$ be the number of all profile $>$ in $\{1, \dots, n\}$ such that when adding a vote $>_{n+1}$ from an additional voter $n + 1$, the outcome will be changed. Then the following lemma holds:

Lemma 1 $\frac{C(n, m)}{N(n, m)} = L(n + 1, m) \leq \frac{D_m}{\sqrt{n+1}}$.

Proof: Suppose $>$ is a profile in $\{1, \dots, n\}$ such that when adding a vote $>_{n+1}$ (note that there are $m!$ possible $>_{n+1}$) from an additional voter $n + 1$, the outcome will be changed. It is easy to see that $(>, >_{n+1})$ is an unstable profile in $\{1, \dots, n, n + 1\}$ and all the unstable profiles in $\{1, \dots, n, n + 1\}$ can be generated from such $>$. We have

$$\begin{aligned} C(n, m) &= \frac{U(n + 1, m)}{m!} = \frac{L(n + 1, m)N(n + 1, m)}{m!} \\ &= L(n + 1, m)N(n, m). \end{aligned}$$

Thus, according to Theorem 1, we have

$$\frac{C(n, m)}{N(n, m)} = L(n + 1, m) \leq \frac{D_m}{\sqrt{n+1}}.$$

■

Thus when $n \rightarrow \infty$, the probability that “adding an additional vote to a random profile changes the outcome” is 0.

Proof: (Theorem 2)

- We now prove the first part of Theorem 2, the edge ratio. Suppose otherwise, there are two possibilities: $\lim_{n \rightarrow \infty} \frac{VE(n, m)}{E(n, m)} \neq 0$ or $\frac{VE(n, m)}{E(n, m)}$ does not converge. Either way, we have $VE(n, m) = cE(n, m)$ for $c \neq 0$ for some sufficiently large n . One can prove this implies $VN(n, m) = c'N(n, m)$ for $c' \neq 0$.

Now consider a monotonic violation for n voters where for voting rule f we have $f(>') = a, f(>'') = b \neq a$ and $>''$ is an improvement of $>'$ w.r.t a . When adding an additional voter $n + 1$, by the comments following Lemma 1, we know that $f(>', >'_{n+1}) = a$ and $f(>'', >''_{n+1}) = b$ (there are a small proportion of profiles that change the outcome, but this reduces the amount of violations). Further, $(>'', >''_{n+1})$ is still monotonic from $(>', >'_{n+1})$ only if the rank of a in $>''_{n+1}$ is at least the same as in $>'_{n+1}$. The number of such pairs $(>'_{n+1}, >''_{n+1})$ is strictly less than $(m!)^2$, the factor by which the number of edges increases from n voters to $n + 1$ voters.

The other case that can generate monotonic violation with $n + 1$ voters is when a previous pair of profiles without violation now becomes one by adding voter $n + 1$. One can count that the number of such pairs is at most $(\frac{2Dm}{\sqrt{n}} + \frac{D^2m}{n})(1 - c)E(n, m)(m!)^2$. Clearly, its ratio to $E(n + 1, m)$ goes to 0 as $n \rightarrow \infty$.

To sum up, we have $\frac{VE(n+1,m)}{E(n+1,m)} = c'' \frac{VE(n,m)}{E(n,m)}$ for sufficiently large n and some $0 < c'' < 1$, which leads to $\lim_{n \rightarrow \infty} \frac{VE(n,m)}{E(n,m)} = 0$, a contradiction.

- Similarly, for the node ratio, consider a monotonic violation for n voters where for voting rule f we have $f(>') = a, f(>'') = b \neq a$ and $>''$ is an improvement of $>'$ w.r.t a . When adding an additional voter $n + 1$, we still have that $f(>', >'_{n+1}) = a$ and $f(>'', >''_{n+1}) = b$ (with a small proportion of deviations which tends to 0 divided by the number of all profiles). We now prove that, for such a violation node $>'$ for n voters, $(>', >'_{n+1})$ is a violation node for any $>'_{n+1}$ for $n + 1$ voters. In other words, we only need to prove that for any $>'_{n+1}$, we can find its monotonic improvement $>''_{n+1}$ w.r.t a . This can be easily achieved by fixing the ranking of a and permuting other candidates. Similarly, we can prove that for such a violation node $>''$ for n voters, $(>'', >''_{n+1})$ is still a violation node for any $>''_{n+1}$, since this amounts to say that we can find some vote that can be improved to $>''_{n+1}$ w.r.t to a . Thus, we conclude that for any violation node for n voters, we can generate $m!$ violation nodes corresponding to it for $n + 1$ voters.

The same as the edge ratio, the other case that can generate monotonic violation with $n + 1$ voters is when a previous pair of profiles without violation now becomes one by adding voter $n + 1$. However we can similarly exclude this since this proportion tends to 0.

To sum up, $\frac{VN(n+1,m)}{N(n+1,m)} = \frac{m!VN(n,m)}{m!N(n,m)} = \frac{VN(n,m)}{N(n,m)}$ for sufficiently large n , which leads to $\lim_{n \rightarrow \infty} \frac{VN(n,m)}{N(n,m)} = c$ for some constant c .

■

7 Conclusion and Future Work

The contribution of this paper is two-fold. First, we offer a new viewpoint on the comparative analysis of voting rules. Second, we provide empirical results on the

application of our new approach to Maskin’s monotonicity.

Our results show that pairwise voting rules, like Maxmin and Copeland, are well-behaved. On the other hand, the commonly used scoring rules of Plurality and Borda have a higher degree of monotonic violation by all the measures we tested. In addition, our results show that even though the majority of profiles are involved in a monotonic violation (the node ratio is increasing in voters), the violations can be eliminated by removing a small proportion of profiles. This domain restriction is not only very low, but it also tends to zero as the number of voters increase. Finally, our asymptotic results provide the first proof of the asymptotic behavior of the edge ratio, which goes to zero as the number of voters increases, and also show that the node ratio converges.

One future direction is to prove a tight bound on the node ratio when the number of voters tends to infinity and also prove a bound on the edge ratio as the number of candidates tends to infinity. We are also investigating how voting rules behave with respect to the violation of the conjunction (disjunction) of two or more axioms. Finally, the most promising direction is going beyond the metrics of node and edge ratios and leveraging the richer structure (seen in Figures 1, 2, and 3) of the violation graph itself.

References

- [1] K. Arrow. A difficulty in the concept of social welfare. *Journal of Political Economy*, 58(4):328–346, 1950.
- [2] M. Barbie, C. Puppe, and A. Tasnadi. Non-manipulable domains for the Borda count. *Economic Theory*, 27:411–430, 2006.
- [3] O. Bochet and T. Storcken. Maximal domains for strategy-proof or Maskin monotonic choice rules. Research Memoranda 002, Maastricht, 2008.
- [4] D. E. Campbell and J. S. Kelly. Information and preference aggregation. *Social Choice and Welfare*, 17:3–24, 2000.
- [5] P. Dasgupta and E. Maskin. On the robustness of majority rule. *Journal of the European Economic Association*, 6(5):949–973, 09 2008.
- [6] O. Erdem and M. Sanver. Minimal monotonic extensions of scoring rules. *Social Choice and Welfare*, 25(1):31–42, October 2005.
- [7] P. Favardin and D. Lepelley. Some further results on the manipulability of social choice results. *Social Choice and Welfare*, 26:485–509, 2006.
- [8] P. Favardin, D. Lepelley, and J. Serais. Borda rule, Copeland method and strategic manipulation. *Review of Economic Design*, 7:213–228, 2002.
- [9] M. Fedrizzi, J. Kacprzyk, and H. Nurmi. How different are social choice functions : a rough sets approach. *Quality and Quantity*, 30(1):87–99, 1996.
- [10] P. Fishburn. *The Theory of Social Choice*. Princeton University Press, 1973.
- [11] P. Fristrup and H. Kleiding. A note on asymptotic strategy-proofness. *Economics Letters*, 31:307–312, 1989.

- [12] A. Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, July 1973.
- [13] D. Hochbaum. *Approximation Algorithms for NP-Hard Problems*. PWS Publishing Co., 1995.
- [14] J. Kelly. Conjectures and unsolved problems: 8. interjacency. *Social Choice and Welfare*, 6:331–335, 1989.
- [15] J. S. Kelly. Almost all social choice rules are highly manipulable, but a few aren't. *Social Choice and Welfare*, 10(2):161–175, April 1993.
- [16] D. Lepelley and F. Valognes. Voting rules, manipulability and social homogeneity. *Public Choice*, 116:165–184, 2003.
- [17] F. Lin and P. Tang. Computer-aided proofs of Arrow's and other impossibility theorems. In *AAAI*, pages 114–119, 2008.
- [18] E. Maskin and T. Sjoström. *Implementation theory*, volume 1 of *Handbook of Social Choice and Welfare*. Elsevier, March 2002.
- [19] H. Moulin. *Axioms of Cooperative Decision Making*. Cambridge University Press, 1988.
- [20] E. Muller and M. A. Satterthwaite. The equivalence of strong positive association and strategy-proofness. *Journal of Economic Theory*, 14(2):412–418, April 1977.
- [21] K. Nehring and C. Puppe. On the possibility of strategy-proof social choice: Non-dictatorship, anonymity and neutrality. *Economic Theory*, 135:269–305, 2007.
- [22] S. Nitzan. The vulnerability of point-voting schemes to preference variation and strategic manipulation. *Public Choice*, 47(2):349–370, January 1985.
- [23] H. Nurmi. Discrepancies in the outcomes resulting from different voting schemes. *Theory and Decision*, 25(2):193–208, 1998.
- [24] P. K. Pattanaik. Strategic voting without collusion under binary and democratic group decision rules. *The Review of Economic Studies*, 42, 1975.
- [25] B. Peleg. A note on manipulability of large voting schemes. *Theory and Decision*, 11:401–412, 1979.
- [26] G. Pritchard and A. Slinko. On the average minimum size of a manipulating coalition. *Social Choice and Welfare*, 27:263–277, 2006.
- [27] M. Satterthwaite. 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, April 1975.
- [28] A. Slinko. On asymptotic strategy-proofness of the plurality and run-off rules. *Social Choice and Welfare*, 19:313–324, 2002.

- [29] D. A. Smith. Manipulability measure of common social choice functions. *Social Choice and Welfare*, 16:639–661, 1999.
- [30] W. Thomson. Monotonic extensions on economic domains. *Review of Economic Design*, 4(1):13–33, 1999.