Strategyproof Approximations of Distance Rationalizable Voting Rules*

Travis C. Service
Department of Electrical Engineering and Computer Science
Vanderbilt University
travis.c.service@vanderbilt.edu

Julie A. Adams
Department of Electrical Engineering and Computer Science
Vanderbilt University
juile.a.adams@vanderbilt.edu

ABSTRACT

This paper considers randomized strategyproof approximations to distance rationalizable voting rules. It is shown that the *Random Dictator* voting rule (return the top choice of a random voter) nontrivially approximates a large class of distances with respect to unanimity. Any randomized voting rule that deviates too greatly from the Random Dictator voting rule is shown to obtain a trivial approximation (i.e., equivalent to ignoring the voters' votes and selecting an alternative uniformly at random).

The outlook for consensus classes, other than unanimity is bleaker. This paper shows that for a large number of distance rationalizations, with respect to the majority and Condorcet consensus classes that no strategyproof randomized rule can asymptotically outperform uniform random selection of an alternative. This paper also shows that veto cannot be approximated nontrivially when approximations are measured with respect to minimizing the number of vetoes an alternative receives.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent systems

General Terms

Theory

Keywords

voting, distance rationalization, strategyproof, approximation

1. INTRODUCTION

The Gibbard-Satterthwaite theorem [11, 16] states that any natural voting procedure can be manipulated. A growing body of work in computation social choice has investigated methods for circumventing the Gibbard-Satterthwaite theorem through hardness of computation [2, 3, 8]. For example, the manipulation problem for many common voting rules has been shown to be NP-hard. However, these are worst case results and say nothing about the manipulation problem on average. Conitzer and Sandholm [3] show that many voting rules can be manipulated in polynomial time for a large fraction of elections.

Recently, Procaccia [15] considered approximating score-based voting rules by strategyproof randomized voting rules. However, Procaccia's approach is limited to rules that have a natural measure of score. This paper studies the approximation of common voting rules with respect to the distance rationalization framework. Approaching the approximation of voting rules from the viewpoint of distance rationalization permits the approximation of voting rules that do not necessarily have a natural measure of score.

In many elections there is an alternative that is a clear winner. For example, if every voter prefers alternative w to every other alternative, then w is the clear winner. Similarly, if w is the Condorcet winner (i.e., w is preferred to every other alternative by a majority of the voters), then w is the clear winner of the election. Both of the previous examples are different notions of consensus in an election. In the second example, a consensus is said to exist in an election whenever there is a Condorcet winner. The Condorcet consensus class is the subset of elections in which their exists a Condorcet winner. The distance rationality framework casts voting in the context of selecting an alternative that is closest to being a consensus winner. The notion of closeness in the distance rationalization framework is formalized by employing distances over elections. The distances employed provide a natural means by which to measure the approximation ratio obtained by strategyproof randomized voting rules.

This paper shows that under the distance rationalizability framework, consensus with respect to Unanimity and arbitrary votewise distances with l_p norms can be approximated to a nontrivial factor[1]. Thus, all positional scoring rules and their (pseudo-)distance rationalizations, given by Elkind et al. [4] can be nontrivially approximated. Similarly, $2-\frac{2}{n}$ and $O(m)$ approximations are obtained for the standard distance rationalizations of plurality and Borda voting, respectively. These approximation ratios are significantly better than random selection of an alternative, which results in a $\Omega(n)$ and $\Omega(nm)$ approximation for plurality and Borda, respectively.

Surprisingly, the Random Dictator rule (select the first

*This research has been supported by an ONR DEPSCOR Award # N000140911161.

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), June, 4–8, 2012, Valencia, Spain.

[1]An approximation ratio is said to be trivial if it is achieved by selecting an alternative uniformly at random.

choice alternative of a random voter) is shown to nontrivially approximate all votewise distance rationalizations with respect to Unanimity and l_p norm. It is shown that deviating too much from the Random Dictator rule results in a trivial ratio under the l_1 norm.

Lower bounds are provided for a number of distance rationalizations. For example, a bound of $\Omega(m)$ is proven for approximating the standard distance rationalization of Borda.

Approximation ratios obtained for a given voting rule are highly dependent on the distance rationalization considered. For example, under unanimity and the discrete distance, plurality can be approximated to a factor of $2-\frac{2}{n}$. However, under the same distance and the majority consensus class, plurality cannot be approximated better than $\Omega(n)$.

The outcome for consensus classes, other than unanimity is bleaker. This paper shows for a number of other distance rationalizable voting rules that essentially one cannot do better than uniform random selection of an alternative.

The remainder of this paper is as follows. Section 2 presents preliminary definitions and related work. Section 3 presents a number of upper and lower bounds on the approximations obtainable by randomized strategyproof voting rules for a number of distance rationalizations. Section 4 concludes with some final remarks.

2. PRELIMINARIES

This section begins by presenting the basic notation and definitions employed throughout this paper and concludes by discussing related work.

2.1 Elections and Strategy-Proofness

An election $E = (A, V)$ consists of a set of alternatives $A = \{a_1, \cdots, a_m\}$ and a tuple of voters $V = \{v_1, \cdots, v_n\}$. Each voter v_i has a strict total preference order $\succ_i$ over the set of alternatives A. Let $v \in V$ and $a \in A$. Define $v(a)$ to be the rank of alternative a in v's preference order. A tuple of preference orders $(\succ_1, \cdots, \succ_n)$ is referred to as a preference profile. A voting rule f is a function that maps each preference profile to a winning alternative.

Informally, a voting rule is manipulable if there exists a preference profile under which some voters can benefit (possibly in expectation) by misrepresenting their true preferences. A voting rule is strategyproof if it is not manipulable. That is, f is manipulable if there exists a preference profile $P = (\succ_1, \cdots, \succ_n)$, a voter v_i and a preference order $\succ_i'$ such that $f(P') \succ_i f(P)$ where $P' = (\succ_1, \cdots, \succ_i', \cdots, \succ_n)$. If f is a randomized voting rule, then f is manipulable if a voter can increase its expected utility by misrepresenting her preferences under some preference profile.

The following two classes of deterministic voting rules are used throughout this paper.

Definition 1. *A deterministic voting rule is said to be unilateral if it is a function of exactly one voters' vote.*

Definition 2. *A deterministic voting rule is said to be duple if it always elects one of two fixed alternatives.*

The following result provides necessary conditions for a randomized voting rule to be strategyproof.

Theorem 1 ([12])**.** *If R is a strategyproof randomized voting rule, then R is a probability distribution over unilateral and duple rules.*

2.2 Distance Rationalization

A consensus in an election $E = (A, V)$ is a clear winner[2] $w \in A$. For example, if every voter in V ranks alternative w first, then w can be considered the consensus winner. Formally, a consensus class $\mathcal{K} = (\mathcal{E}, \mathcal{W})$ is a tuple, where $\mathcal{E}$ is a set of elections and $\mathcal{W} : \mathcal{E} \to A$ is a function that determines for each election $E \in \mathcal{E}$ the consensus winner. We consider three consensus classes in which there exists a clear winner [4, 5, 6, 13, 14].

1. Unanimity ($\mathcal{U}$): Consists of all elections in which every voter ranks the same alternative first. The consensus winner is the alternative preferred by all voters.
2. Majority ($\mathcal{M}$): Consists of all elections in which some alternative is ranked first by a strict majority of the voters. The consensus winner is the unique alternative that more than half of the voters rank first.
3. Condorcet ($\mathcal{C}$): Consists of all elections in which there exists a Condorcet winner (i.e., an alternative that defeats every other alternative in a pairwise election). The Condorcet winner is the consensus winner.

A consensus class can be extended to a voting rule over arbitrary elections by defining the winning alternative in an election E to be the alternative that is closest to being a consensus winner. Such an extension requires a notion of distance between elections.

Informally, a distance function $d : X \times X \to \mathbb{R} \cup \{\infty\}$ on a set X is a function mapping pairs of elements in X to a non-negative real value representing the distance between the elements. A pseudo-distance function is similar to a distance function but there may be distance 0 between two distinct members of X. We are interested in (pseudo-)distances over the set of elections.

Definition 3. *Let $\mathcal{K} = (\mathcal{E}, \mathcal{W})$ be a consensus class and d a (pseudo-)distance function on the set of elections. A voting rule, f, is $(\mathcal{K}, d)$-rationalizable if for every election E*

$$f(E) \in \underset{a \in A}{\operatorname{argmin}} \left\{ \min_{E' \in \mathcal{E} \,:\, a \in \mathcal{W}(E')} \{d(E, E')\} \right\}.$$

Refer to value $d(a) = \min_{E' \in \mathcal{E} \,:\, a \in \mathcal{W}(E')} \{d(E, E')\}$ as the distance score a. Distance rationalizable voting rules select an alternative with the least distance score. In general, there may be multiple alternatives that tie for the best distance score. Voting rules break ties in such situations.

Some of the results in this paper apply to arbitrary distances of a particular form. Let d be a distance over preference orders and let N be a norm on $\mathbb{R}^n$. Then the function

$$\hat{d}(E, E') := \begin{cases} \infty & \text{if } A \neq A' \text{ or } |V| \neq |V'| \\ N(d(v_1, v_1'), \cdots, d(v_n, v_n')) & \text{otherwise} \end{cases}$$

is a distance over preference profiles. Distances of the above form are referred to as votewise distances [5].

Many of the norms used in votewise distances in the literature (and in the results presented in this paper) are the l_p norms ($p \in \mathbb{N} \cup \{\infty\}$) defined for each $r_1, \cdots, r_n \in \mathbb{R}$ as

$$l_p(r_1, \cdots, r_n) := \begin{cases} (r_1^p + \cdots + r_n^p)^{\frac{1}{p}} & \text{if } p \in \mathbb{N} \\ \max\{r_1, \cdots, r_n\} & \text{if } p = \infty. \end{cases}$$

[2]Sometimes multiple consensus winners are allowed.

Unless otherwise stated, it is assumed that the norm N in the definition of $\hat{d}$ is the l_1 norm (i.e., the sum of the distances over individual votes). For any norm N and any distance over preference orders d, define $N \circ d$ to be the corresponding votewise distance.

Several common votewise distances are employed throughout this paper.

1. $d_{swap}(v, v')$ represents the number of adjacent pairs that must be swapped in the preference order v in order to obtain v'.

2. $d_{disc}(v, v')$ is 0 if $v = v'$ and 1 otherwise.

3. Fix m numbers $\alpha_1 \geq \cdots \geq \alpha_m$. Define the distance: $d_\alpha(v, v') = \sum_{a \in A} |\alpha_{v(a)} - \alpha_{v'(a)}|$.

We also consider the following non-votewise distance, d_{ins}. d_{ins} is defined as follows. Let $E = (A, V)$ and $E' = (A, V')$ be elections. For each $v_i \in V$ let $\succ_{v_i}$ be v_i's preference order. Likewise, for $v_i \in V'$, let $\succ'_{v_i}$ be v_i's preference order. Then $d_{ins}(E, E') = |V \setminus V'| + |V' \setminus V| + 2|\{v_i \in V \cap V' : \succ_{v_i} \neq \succ'_{v_i}\}|$. Elkind et al. [7] show that d_{ins} under $\mathcal{C}$ is essentially equivalent to considering the number of voters that need be added to an election to make a given candidate the Condorcet winner. More formally, Elkind et al. show that given an election $E = (A, V)$ and an alternative $a \in A$, there exists an election $E_1 = (A, V \cup V_1)$ such that a is the Condorcet winner and $|V_1| \leq k$, if and only if there exists an election $E_2 = (A, V_2)$ such that $d_{ins}(E, E_2) \leq k$.

2.3 Voting Rules

Several common voting rules considered in this paper are now defined. The employed voting rules assign alternatives a score based upon the voters' preferences and then select the alternative with the best score. We present the definitions of the voting rules for situations where no ties in score occur. When multiple alternatives tie for the best score, an arbitrary, but fixed, tie breaking scheme is employed to select a single alternative. All of the presented results are unaffected by the actual tie breaking scheme employed.

1. Positional Scoring Rules: Let $\alpha = (\alpha_1, \cdots, \alpha_m)$ be a vector of m integers with $\alpha_1 \geq \cdots \geq \alpha_m$. Under the voting rule R_α, alternative a is awarded α_k points for each voter that ranks a in position k. The winner under R_α is the alternative with the largest score. Plurality is defined by $\alpha = (1, 0, \cdots, 0)$, Borda by $\alpha = (m-1, m-2, \cdots, 1, 0)$ and veto by $\alpha = (1, 1, \cdots, 1, 0)$.

   Elkind et al. [4] show that every positional scoring rule R_α is pseudo-distance rationalizable with respect to $(\mathcal{U}, \hat{d}_\alpha)$. Further, plurality and Borda are rationalizable with respect to $\mathcal{U}$ with the distances $\hat{d}_{disc}$ and $\hat{d}_{swap}$, respectively.

2. Maximin: The Maximin score of an alternative a in an election E is

$$sc(a) = \min_{b \in A \setminus \{a\}} \{|\{v_i \in V : a \succ_i b\}|\}.$$

   That is, the Maximin winner of an election E is the alternative that performs the best in their worst pairwise election against any other alternative. Maximin is $(\mathcal{C}, d_{ins})$-rationalizable [4].

3. Dodgson: The Dodgson score of an alternative a in an election $E = (V, A)$ is the minimum number of swaps of adjacent alternatives in the preferences of voters to make a the Condorcet winner. The winner is the alternative with the smallest Dodgson score. Dogdson is $(\mathcal{C}, \hat{d}_{swap})$-rationalizable.

2.4 Related Work

The manipulation problem has received much attention [1, 10, 15]. For a recent survey, see Faliszewski and Procaccia [9].

Recent work has investigated the approximation of common voting rules by randomized voting rules that are strategyproof (or almost strategyproof). Procaccia [15] recently quantified the level of approximation that can be obtained for a number of score based voting rules.

Birell and Pass [1] relax the requirement that the randomized approximation be strategyproof. Rather, Birell and Pass consider randomized voting rules such that no voter can improve their expected utility by more than ϵ by voting untruthfully. On a positive note, Birell and Pass show that for ϵ sufficiently large ($\omega(\frac{1}{n})$), every deterministic voting rule can be approximated by an ϵ-strategyproof randomized rule. However, for $\epsilon = o(\frac{1}{n})$, every ϵ-strategyproof voting rule is a distribution over unilateral and duple rules (just as strategyproof randomized voting rules are).

One drawback of Procaccia's approach is that it is limited to score based voting rules. Birell and Pass [1] measure the quality of an approximation by the minimum number of votes that must be changed in order to elect the approximate winner. Hence, Birell and Pass' results apply to arbitrary voting rules. This paper investigates the use of the distance rationalization framework to measure the quality of approximations obtained by randomized strategyproof voting rules. Like Birell and Pass' approach, our approach allows for defining approximations to non-score based rules.

3. RESULTS

All of the results in this paper are stated for particular distance rationalizations. For example, rather than saying plurality can be approximated well, we state the distance rationalization by which we are approximating. It will be shown that the same voting rule can be approximated to different degrees, depending upon the rationalization used. Hence, it is meaningless to say, in this framework, that, for example, plurality can be approximated well, since it depends on which rationalization of plurality is employed.

Let $(\mathcal{K}, d)$ be a distance rationalization. Let E be an election and let w be the winning alternative under $(\mathcal{K}, d)$. A natural measure of the approximation ratio of a strategyproof randomized voting rule R with respect to $(\mathcal{K}, d)$ is $\mathbb{E}\left(\frac{d(R(E))}{d(w)}\right)$, with the understanding that $\frac{0}{0} = 1$ and $\frac{n}{0} = +\infty$ for any $n > 0$. It is straightforward to observe that for the Majority and Condorcet consensus classes that no strategyproof voting rule is guaranteed to return a consensus winner when one exists. Hence, the approximation ratio obtained by any strategyproof voting rule to distance rationalizations with respect to Majority or Condorcet is $+\infty$. However, simply stating an approximation ratio of $+\infty$ for these consensus classes is less than desirable as it may be the case that, for example, when there is a majority winner, the given strategyproof voting rule always selects an alternative that is close to being a consensus winner. Therefore, our lower bound results consider the approximation ratio obtained by randomized strategyproof voting rules on elections where there is no consensus winner.

The presented upper bounds, with respect to $\mathcal{U}$ all employ a strategyproof randomized voting rule that is consistent with $\mathcal{U}$. A unanimous winner, if it exists, is always selected. Hence, if E is a member of the consensus class, then the result is a perfect approximation.

3.1 Upper Bounds

Surprisingly the following strategyproof randomized voting rule obtains a nontrivial approximation ratio with respect to $(\mathcal{U}, \hat{d})$, when d is any votewise distance and $\hat{d} = l_p \circ d$.

Random Dictator. *Uniformly at random, select a voter v. Return v's first choice.*

Theorem 2 shows that Random Dictator obtains a good approximation to many distance rationalizations.

Theorem 2. *Let d be a distance over preference orders, $p \in \mathbb{N} \cup \{\infty\}$, and let $\hat{d} = l_p \circ d$ be the corresponding distance over elections. For preferences over m alternatives and $x \in A$, let $d_{Max}(m, x)$ be the maximum distance between any preference order $\succ$, that does not rank x first, to the closest preference order to $\succ$ that does rank x first. Let $d_{Max}(m) = \max_{x \in A}\{d_{Max}(m, x)\}$. Let $d_{Min}(m, x)$ be the minimum distance between any preference order $\succ$, that does not rank x first, to any preference order that does rank x first. Let $d_{Min}(m) = \min_{x \in A}\{d_{Min}(m, x)\}$.*

Random Dictator approximates $(\mathcal{U}, \hat{d})$ to within a factor of

$$\left(1 - \frac{1}{n}\right) \cdot \left(\frac{d_{Max}(m)}{d_{Min}(m)} + 1\right).$$

Proof. Let $E = (A, V)$ be an election with n voters and m alternatives. Let $w \in A$ be a winning alternative in $(\mathcal{U}, \hat{d})$ and let x be the number of voters that rank w first. If w is the unanimous winner, then Random Dictator selects w and obtains a perfect approximation. So assume that $x \leq n - 1$.

For $a \in A$, let p_a be the probability that a is selected by Random Dictator. That is, p_a is the ratio of the number of voters that rank a first to the total number of voters n. Hence, $p_w = \frac{x}{n}$.

Note that if an alternative $a \in A \setminus \{w\}$ is selected by Random Dictator, then a must be ranked first by at least one voter. If $p = \infty$, then the maximum possible distance score of any candidate is $d_{Max}(m)$ and the minimum possible distance score of w is $d_{Min}(m)$. Thus, the expected distance score is

$$\begin{aligned}
\mathbb{E}\left(\frac{\hat{d}(R(E))}{\hat{d}(w)}\right) &\leq \frac{x}{n} + \frac{n-x}{n} \cdot \frac{d_{Max}(m)}{d_{Min}(m)} \\
&\leq \left(1 - \frac{1}{n}\right)\left(\frac{d_{Max}(m)}{d_{Min}(m)} + 1\right).
\end{aligned}$$

Similarly, if $p \in \mathbb{N}$, then the maximum distance score of a possible winner under Random Dictator is $(n-1)^{\frac{1}{p}} d_{Max}(m)$. Likewise, the minimum possible distance score of w is $(n - x)^{\frac{1}{p}} d_{Min}(m)$. Thus, the approximation ratio obtained by Random Dictator is

$$\begin{aligned}
\mathbb{E}\left(\frac{\hat{d}(R(\succ))}{\hat{d}(w)}\right) &= \frac{1}{\hat{d}(w)}\left[\sum_{a \in A} p_a \hat{d}(a)\right] \\
&= \frac{1}{\hat{d}(w)}\left[\frac{x}{n}\hat{d}(w) + \sum_{a \in A \setminus \{w\}} p_a \hat{d}(a)\right] \\
&\leq \frac{x}{n} + \frac{1}{\hat{d}(w)} \sum_{a \in A \setminus \{w\}} p_a (n-1)^{\frac{1}{p}} d_{Max}(m) \\
&\leq \frac{x}{n} + \frac{n-x}{n} \cdot \frac{(n-1)^{\frac{1}{p}} d_{Max}(m)}{(n-x)^{\frac{1}{p}} d_{Min}(m)} \\
&= \frac{x}{n} + \frac{(n-x)^{1-\frac{1}{p}}}{n} \cdot \frac{(n-1)^{\frac{1}{p}} d_{Max}(m)}{d_{Min}(m)} \\
&\leq \frac{n-1}{n} + \frac{(n-1)^{1-\frac{1}{p}}}{n} \cdot \frac{(n-1)^{\frac{1}{p}} d_{Max}(m)}{d_{Min}(m)} \\
&\leq \left(1 - \frac{1}{n}\right)\left(\frac{d_{Max}(m)}{d_{Min}(m)} + 1\right).
\end{aligned}$$

□

Recall that every positional scoring rule is pseudo-distance rationalizable with respect to unanimity under the votewise distance $\hat{d}_\alpha = l_1 \circ d_\alpha$ [4]. Note that if $\alpha_i \neq \alpha_j$ whenever $i \neq j$, then $d_{Min}(m) = 2(\alpha_1 - \alpha_2)$ and $d_{Max}(m) = 2(\alpha_1 - \alpha_m)$. Corollary 1 shows that all positional scoring rules with $\alpha_i \neq \alpha_j$ are approximated by Random Dictator to within a nontrivial factor.

Corollary 1. *If R_α is a positional scoring rule, such that $\alpha_i \neq \alpha_j$ whenever $i \neq j$, then Random Dictator approximates $(\mathcal{U}, \hat{d}_\alpha)$ to within a factor of*

$$\left(1 - \frac{1}{n}\right) \cdot \left(\frac{\alpha_1 - \alpha_m}{\alpha_1 - \alpha_2} + 1\right).$$

If $\alpha_1 = \alpha_2 \neq \alpha_3$, then Random Dictator does not approximate $(\mathcal{U}, \hat{d}_\alpha)$ well. Consider the election in which each voter ranks alternative w second and no other alternative is ranked first more than once. The distance score of w is 0, but the distance score of every other alternative is at least $(n-1)(\alpha_1 - \alpha_3)$. However, Random Dictator never selects w in such elections.

Theorem 2 allows one to obtain nontrivial approximation ratios with respect to the standard distance rationalizations of common voting rules. For example, Plurality is known to be $(\mathcal{U}, \hat{d}_{disc})$-rationalizable. Since, under $\hat{d}_{disc}$, $d_{Min}(m) = d_{Max}(m) = 1$, Random Dictator approximates $(\mathcal{U}, \hat{d}_{disc})$ to within a factor of $2 - \frac{2}{n}$. Hence, Random Dictator is significantly better than uniform random selection of a candidate, which obtains an approximation ratio of $\Omega(n)$ in the worst case (e.g., when some $w \in A$ is ranked first by all but one voter).

Similarly, Borda is $(\mathcal{U}, \hat{d}_{swap})$-rationalizable. Notice that in this case $d_{Max}(m) = \Theta(m)$ and $d_{Min}(m) = 1$. Hence, Random Dictator obtains a $O(m)$ approximation. Note that uniformly at random selecting an alternative obtains an approximation ratio of $\Theta(nm)$ to $(\mathcal{U}, \hat{d}_{swap})$ (e.g., when one alternative is ranked first by all but one voter and every other alternative obtains the same distance score of $\Theta(nm)$).

It may be expected that the voting rules presented by Procaccia [15] will outperform Random Dictator, since an alternative is selected with probability proportional to its score. However, electing an alternative with probability proportional to its Borda score obtains a $\Theta(nm)$ approximation to $(\mathcal{U}, \hat{d}_{swap})$. Consider the election in which each voter has the same preference order. Let A' be the set of alternatives ranked in the lower $\frac{m}{2}$ positions by each voter. The sum of the scores of alternatives in A' is $\Omega(nm^2)$. Thus, the probability that some alternative in A' is selected is $\Theta(1)$. Therefore, selecting an alternative with probability proportional to its Borda score results in a $\Omega(nm)$ approximation. In the case of Plurality, the strategyproof randomized voting rule given by Procaccia [15] is the Random Dictator rule.

3.2 Lower Bounds

All of the presented lower bounds employ Yao's Minimax principle [17]. Consider the following two player game. The space of the first player's strategies consist of all duple and unilateral voting rules and the second player's strategies consist of all preference profiles. Given a choice of a pure strategy for each player, the outcome is defined to be the approximation ratio obtained by the unilateral or duple rule selected by the first player on the preference profile selected by the second player. Since the first player's pure strategies consist of all unilateral and duple rules, the first player's mixed strategies contain all strategyproof randomized voting rules.

Let P be any probability distribution over preference profiles. In this setting, Yao's Minimax principle states that the approximation ratio obtained by any strategyproof randomized rule is at most the approximation ratio obtained by the best deterministic duple or unilateral rule over P. Hence, the performance of any strategyproof randomized voting rule can be lower bounded by constructing a probability distribution over preference profiles on which no unilateral or duple rule performs well in expectation.

When employing Yao's principle for proving lower bounds, there are two cases to be considered: unilateral rules and duple rules. For our purposes, it will suffice to treat a duple rule as a set of two alternatives. Hence, for a duple rule D, we may treat D as a set of two alternatives as we are indifferent to how D selects a winner.

Let $\hat{d} = l_1 \circ d$ be a votewise distance. Let $\succ$ be a preference profile over m alternatives and let $a \in A$. Let $d_a(\succ)$ be the minimum distance between $\succ$ and any other preference profile $\succ'$ that ranks a first. We say that d and $(\mathcal{U}, \hat{d})$ are *rank based* if $d_a(\succ)$ depends only on the rank of a in $\succ$. Thus, $d_a(\succ)$ is independent of how the alternatives in $A \setminus \{a\}$ are ranked and is also independent of the alternative a. That is, if $\succ_a$ is any preference order that ranks a in position k and $\succ_b$ is any preference order that ranks b in position k, then $d_a(\succ_a) = d_b(\succ_b)$. Hence, there exists a function $r_d : \mathbb{N} \to \mathbb{R}^{\geq 0}$ such that $d_a(\succ) = r_d(k)$, where k is the rank of a in $\succ$. Note that all the votewise distances defined in this paper are rank based. Define $d_{Ave}(m) = \frac{1}{m}\sum_{k=1}^{m} r_d(k)$.

The Random Dictator rule is not very desirable as it ignores the preferences of all but one randomly selected voter. However, as the next result shows, it is not possible to deviate too greatly from the Random Dictator rule and still obtain a nontrivial approximation to $(\mathcal{U}, \hat{d})$, where $\hat{d}$ is any rank based votewise distance.

Theorem 3. *Let $(\mathcal{U}, \hat{d})$ be rank based and let R be a strategyproof randomized voting rule, such that with probability p, R selects a duple rule and with probability q, R selects a unilateral rule that does not select the voter's first choice alternative. If $p = \Omega(1)$ or $q = \Omega(1)$, then R obtains an approximation ratio of $\Omega(nd_{Ave}(m))$.*

Proof. The proof of Theorem 3 employs Yao's Minimax principle, to that end we define a randomized procedure for generating preference profiles. The procedure constructs a preference profile as follows:

1. Let $n = n' + 1$, $m \geq 4$, and let $m - 1$ divide n'. Select $w \in A$ uniformly at random.
2. Choose permutations π over $A \setminus \{w\}$ and σ over V uniformly at random.
3. Define the preferences of the voters as follows. For $i \in \{1, \cdots, n'\}$, $v_{\sigma(i)}$ ranks alternative $\pi(k)$ in position $\pi(k+i)+1$ (where the addition $k+i$ is modulo $m-1$) and ranks w first. Voter $v_{\sigma(n)}$ ranks w second and ranks the other alternatives arbitrarily.

Let E be a random election drawn from the above distribution. By construction, there is some $w \in A$ that is ranked first by all but one voter. Hence, the distance score of w is $r_d(2)$, as w is ranked second by one voter and the norm is l_1. For fixed distance function d, $r_d(2) = \Theta(1)$. Also, since $m - 1$ divides n', every other alternative is ranked in position k by at least $\frac{n'}{m-1}$ voters, for each $k = 2, \cdots, m$, and is ranked first by at most one voter. Hence, every alternative other than w achieves a distance score of at least $r_d(1) + \frac{n'}{m-1} \cdot \sum_{k=2}^{m} r_d(k) = \Theta(n \cdot d_{Ave}(m))$, since $r_d(1) = 0$.

Consider first the case where $p = \Omega(1)$. That is, when R selects a duple rule D with at least constant probability. Since w is selected from A at random, the probability that $w \in D$ is at most $\frac{2}{m}$. Hence, with probability at least $\frac{m-2}{m} \geq \frac{1}{3}$, an alternative with distance score $\Theta(n \cdot d_{Ave})$ is returned. Therefore the approximation ratio of R is at least $\Theta(pn \cdot d_{Ave}) = \Omega(nd_{Ave}(m))$.

Now consider the case where $q = \Omega(1)$. Since w is ranked first by all but one alternative and the voter that does not rank w first is selected uniformly at random, with probability $q \cdot \frac{n-1}{n}$, R selects an alternative other than w with a distance score of $\Theta(n \cdot d_{Ave})$. Thus, again, the approximation ratio of R is at least $\Theta(q \cdot \frac{n-1}{n} \cdot nd_{Ave}) = \Omega(nd_{Ave})$. ◻

Roughly speaking, Theorem 3 shows that for a strategyproof randomized voting rule R to obtain a nontrivial approximation ratio to $(\mathcal{U}, \hat{d})$, R must, with probability tending toward 1 for increasing numbers of voters and alternatives, select an alternative that is ranked first by some voter.

Theorem 4. *No strategyproof randomized voting rule approximates $(\mathcal{U}, \hat{d}_{disc})$ (plurality) to a ratio less than $2 - \frac{2}{n}$.*

Proof. Define the following procedure for constructing random preference profiles:

1. Select $w \in A$ uniformly at random. For ease of exposition, assume the members of $A \setminus \{w\}$ are the first $m - 1$ integers: $0, \cdots, m - 2$.
2. Choose permutations π over $A \setminus \{w\}$ and σ over V uniformly at random.
3. For $i \in \{1, \cdots, n - 1\}$, $v_{\sigma(i)}$ ranks w first and ranks alternative k in position $\pi(k+i)+1$ (where the addition is modulo $m-1$). $v_{\sigma(n)}$ ranks alternative k in position $\pi(k)$ and ranks w last.

In any elections drawn from the given distribution, the distance score of w is 1 (as all but one voter ranks w first). The distance score of any other alternative is at least $n-1$.

First, consider a duple rule D. Since w is selected uniformly at random, the probability that $w \in D$ is at most $\frac{2}{m}$. Thus, the probability that an alternative other than w is selected is at least $\frac{m-2}{m} \geq \frac{1}{3}$. Therefore, with probability at least $\frac{1}{3}$, D selects an alternative other than w with distance score at least $n-1$, resulting in an approximation ratio of at least $\frac{n-1}{3}$.

Now consider a unilateral rule, U. Since U is unilateral, there exists a single voter v_i that determines the winner of the election under U. The probability that v_i ranks w first is $\frac{n-1}{n}$. Since w is ranked first or is ranked last by every voter, any unilateral rule maximizes the probability of selecting w when it always returns either the first or last ranked alternative of v_i. Hence, U maximizes the probability of selecting w when it always returns v_i's first choice alternative. Therefore, any unilateral rule U selects w with probability at most $\frac{n-1}{n}$ and with probability at least $\frac{1}{n}$, U selects an alternative other than w. Hence, the approximation ratio of U is at least $\frac{n-1}{n} + \frac{1}{n} \cdot \frac{n-1}{1} = 2 - \frac{2}{n}$. Thus, no strategyproof randomized voting rule approximates $(\mathcal{U}, \hat{d}_{disc})$ to a factor less than $2 - \frac{2}{n}$. □

The remainder of the presented lower bound proofs use the following procedure (or a slight variation) for creating a distribution, $\mathcal{P}$, on preference profiles.

1. Let n be even, $m \geq 4$, and let $m-1$ divide $\frac{n}{2}$. Select $w \in A$ uniformly at random. For ease of exposition, assume the members of $A \setminus \{w\}$ are the first $m-1$ integers: $0, \cdots, m-2$.

2. Choose permutations π over $A \setminus \{w\}$ and σ over V, uniformly at random.

3. For $i \in \{1, \cdots, \frac{n}{2}\}$, let $v_{\sigma(i)}$ rank w first and rank alternative k in position $\pi(k+i)+1$ (where the addition $k+i$ is modulo $m-1$). For $i \in \{\frac{n}{2}+1, \cdots, n\}$, let $v_{\sigma(i)}$ rank w second and rank alternative k in position $\pi(k+i)+2$ if $\pi(k+i) \neq 1$ and in position 1 otherwise.

Based on the above procedure, it is observed that the selected alternative w is "close" (under most natural definitions of distance) to being a consensus winner, since half of the voters rank it first and the other half second. Since every other alternative is ranked cyclically by the voters and $m-1$ divides evenly into $\frac{n}{2}$, each alternative other than w is ranked in position k by the same number of voters. Let $\mathcal{P}'$ be the distribution over preference profiles identical to $\mathcal{P}$, except that for $i \in \{\frac{n}{2}+1, \cdots, n\}$, $v_{\sigma(i)}$ ranks w last. Lemma 1 will be employed in many of the lower bound proofs.

Lemma 1. *Let $\mathcal{K}$ be a consensus class and let d be a distance over elections. Let E be an election drawn from $\mathcal{P}$ or $\mathcal{P}'$. If for each alternative $a \in A \setminus \{w\}$, $d(a) \geq d_{min}$, then the approximation ratio achieved by any strategyproof randomized voting rule to $(\mathcal{K}, d)$ is at least $\frac{1}{2} \cdot \frac{d_{min}}{d(w)}$.*

Proof. Let R be any strategyproof randomized voting rule. Let p be the probability that R selects a duple rule and $1-p$ the probability that R selects a unilateral rule.

Consider first a duple rule, D. Since w is selected uniformly at random, the probability that $w \in D$ is at most $\frac{2}{m}$. Hence, the expected distance approximation of the alternative selected is at least

$$\frac{2}{m} + \frac{m-2}{m}\frac{d_{min}}{d(w)} \geq \frac{1}{2}\frac{d_{min}}{d(w)},$$

since $m \geq 4$.

Now consider a unilateral rule, U. Since w is ranked first by exactly half of the voters and second by the other half (last by the other half in the case of $\mathcal{P}'$) and those voters that rank w first are randomly distributed amongst all voters, the probability that U selects w is at most $\frac{1}{2}$. Hence, the probability that U selects an alternative other than w is at least $\frac{1}{2}$. The expected distance approximation is at least:

$$\frac{1}{2} \cdot \frac{d(w)}{d(w)} + \frac{1}{2} \cdot \frac{d_{min}}{d(w)} > \frac{1}{2} \cdot \frac{d_{min}}{d(w)}.$$

Therefore, the approximation ratio obtained by R is at least $\frac{1}{2} \cdot \frac{d_{min}}{d(w)}$. □

Lemma 1 allows one to lower bound the approximation ratio achievable for a number of distance rationalizations by strategyproof randomized voting rules. In particular, for scoring rules with $\alpha_1 \neq \alpha_2$, Theorem 5 obtains a lower bound close to the upper bound obtained by Corollary 1.

Theorem 5. *If R_α is a positional scoring rule with $\alpha_1 \neq \alpha_2$, then no strategyproof randomized voting rule approximates $(\mathcal{U}, \hat{d}_\alpha)$ to within a factor less than*

$$\frac{\alpha_1}{\alpha_1 - \alpha_2} - \frac{S - \alpha_2}{(m-1)(\alpha_1 - \alpha_2)},$$

where $S = \sum_{k=1}^{m} \alpha_k$.

Proof. Let E be an election drawn from $\mathcal{P}$. Note that the distance score of w is $n(\alpha_1 - \alpha_2)$. Let $k \in \mathbb{N}$ such that $k(m-1) = \frac{n}{2}$. Such a k exists, as $m-1$ divides $\frac{n}{2}$. Every alternative $a \in A \setminus \{w\}$ is ranked first k times, second k times, and in position $i \geq 3$, $2k$ times. If voter v_i ranks a in position r, then v_i contributes $2(\alpha_1 - \alpha_r)$ to the distance score of a, since the distance to the closest preference order to v_i in which a is ranked first is $2(\alpha_1 - \alpha_r)$. Thus, for every $a \in A \setminus \{w\}$

$$\begin{aligned} d(a) &= k\left[2(\alpha_1 - \alpha_2)\right] + 2k\left[\sum_{i=3}^{m} 2(\alpha_1 - \alpha_i)\right] \\ &\geq 2k\left[2(m-2)\alpha_1 - 2\sum_{i=2}^{m}\alpha_i\right] \\ &= 2k\left[2(m-2)\alpha_1 - 2(S - \alpha_1)\right] \\ &= 4k(m-1)\alpha_1 - 4kS \\ &= 2n\alpha_1 - 2\left(\frac{n}{m-1}\right)S \end{aligned}$$

By Lemma 1, the approximation ratio of any strategyproof randomized voting rule is lower bounded by

$$\frac{1}{2}\frac{2n\alpha_1 - 2\left(\frac{n}{m-1}\right)S}{n(\alpha_1 - \alpha_2)} = \frac{\alpha_1}{\alpha_1 - \alpha_2} - \frac{S - \alpha_2}{(m-1)(\alpha_1 - \alpha_2)}.$$

□

Theorem 6. *The approximation ratio obtained by any strategyproof randomized voting rule to $(\mathcal{U}, \hat{d}_{swap})$ (Borda) is $\Omega(m)$.*

Proof. Notice that in an election E drawn from $\mathcal{P}$, $d(w) = \frac{n}{2}$ since w is ranked first by $\frac{n}{2}$ voters and second by $\frac{n}{2}$ voters. Let $k(m-1) = \frac{n}{2}$. For any other alternative $a \in A$, a is ranked in position 1 and 2, k times and in position $r = 3, \cdots, m$, $2k$ times. Each voter that ranks a in position r contributes $r-1$ to $d(a)$, as a must be swapped with at least $r-1$ alternatives in order for a to be ranked first. Hence

$$\begin{aligned} d(a) &= k \cdot 1 + 2k \sum_{r=3}^{m} (r-1) \\ &\geq \frac{1}{2} \cdot \frac{n}{m-1} \sum_{r=2}^{m} (r-1) \\ &= \frac{1}{2} \cdot \frac{n}{m-1} \frac{(m-1)m}{2} = \Omega(nm). \end{aligned}$$

By Lemma 1 every strategyproof randomized voting rule achieves an approximation ratio of $\Omega(m)$ on $(\mathcal{U}, \hat{d}_{swap})$. □

Lemma 2. *Let E be an election drawn from $\mathcal{P}'$ and let $a \in A \setminus \{w\}$. There exists a set $A_a \subseteq A \setminus \{w\}$ with $|A_a| = \Theta(m)$ such that at least $\frac{3n}{4}$ voters prefer all members of A_a to a.*

Proof. Let E be an election drawn from $\mathcal{P}'$. Let π and σ be the random permutations over $A \setminus \{w\}$ and V used to construct E. Let $k \in \mathbb{N}$ such that $k(m-1) = \frac{n}{2}$. Since $m-1$ divides $\frac{n}{2}$, such a k exists. Among the voters, $v_{\sigma(1)}, \cdots, v_{\sigma(\frac{n}{2})}$, a is ranked last k times. Let $v_{\sigma(i_1)}, \cdots, v_{\sigma(i_k)}$ be the k voters that rank a last among the voters $v_{\sigma(1)}, \cdots, v_{\sigma(\frac{n}{2})}$. By construction, all the voters $v_{\sigma(i_j)}$ have the same preference order. Let A_a be the set of $\frac{m}{10} - 1$ alternatives that immediately precede a in the voters $v_{\sigma(i_j)}$'s preference order. That is, A_a consists of those alternatives that are ranked among the bottom $\frac{m}{10}$ positions, other than a. For each i_j, by construction, voter $v_{\sigma(i_j - l)}$ ranks a in position $m-l$, for $l < m-1$. Thus, all alternatives in A_a are ranked above a by the voters $v_{\sigma(i_j)-l}$, for each $0 \leq l \leq m - \frac{m}{10} = \frac{9m}{10}$. Therefore among the voters $v_{\sigma(1)}, \cdots, v_{\sigma(\frac{n}{2})}$, there are $k \cdot \frac{9m}{10} \geq \frac{9(m-1)k}{10} = \frac{9n}{20}$ voters that rank all alternatives of A_a above a.

A similar argument applies to the voters among $v_{\sigma(\frac{n}{2}+1)}, \cdots, v_{\sigma(n)}$ that rank a second to last. Hence, among the voters $v_{\sigma(\frac{n}{2}+1)}, \cdots, v_{\sigma(n)}$, all members of A_a are ranked above a by at least $\frac{9n}{20}$ voters. Hence, the number of voters that prefer all members of A_a to a is greater than $\frac{3n}{4}$. □

Theorem 7. *The approximation ratio obtained by any strategyproof randomized voting rule to $(\mathcal{C}, \hat{d}_{disc})$ is $\Omega(n)$.*

Proof. Let E be an election drawn from $\mathcal{P}'$. By Lemma 2, for any alternative $a \in A \setminus \{w\}$ to become the Condorcet winner, $\Omega(n)$ voters must change their vote. However, for w to become the Condorcet winner, only one voter must change their vote. By Lemma 1, the approximation ratio of any strategyproof randomized voting rule is $\Omega(n)$. □

Theorem 7 roughly shows that no strategyproof randomized voting rule can outperform uniform random selection of an alternative, since under $\hat{d}_{disc}$ the maximum distance score of any alternative is at most n.

Theorem 8. *The approximation ratio obtained by any strategyproof randomized voting rule to $(\mathcal{C}, \hat{d}_{swap})$ (Dodgson) is $\Omega(n)$.*

Proof. Let E be an election drawn from $\mathcal{P}'$. To make a the Condorcet winner, a must be swapped with at least $|A_a| = \Theta(m)$ alternatives in the preferences of $\Omega(n)$ voters. To make w the Condorcet winner, it suffices that w be swapped with all the alternatives in the preference order of one voter that ranks w last. Thus, $d(w) = O(m)$. By Lemma 1, every strategyproof randomized voting rule obtains an approximation ratio of $\Omega(n)$. □

Theorem 9. *The approximation obtained by any strategyproof randomized voting rule to $(\mathcal{C}, d_{ins})$ (Maximin) is $\Omega(n)$.*

Proof. Recall that the d_{ins} score of an alternative a is equal to the minimum number of voters that must be added to make a the Condorcet winner. Let E be an election drawn from $\mathcal{P}'$. The d_{ins} score of w is 1, since the addition of a single voter that ranks w first will make w the Condorcet winner. However, the d_{ins} score of every other alternative is $\Omega(n)$, since to make any alternative $a \in A \setminus \{w\}$ the Condorcet winner $\Omega(n)$ voters must be added that rank a the members of A_a. Hence, by Lemma 1, the approximation ratio obtained by any strategyproof randomized voting rule is $\Omega(n)$. □

Notice that any alternative a can be made the Condorcet winner by the addition of at most $n+1$ voters that rank a first. Hence, $(\mathcal{C}, d_{ins})$ cannot be nontrivially approximated by any strategyproof randomized voting rule.

Theorem 10. *The approximation ratio obtained by any strategyproof randomized voting rule to $(\mathcal{M}, \hat{d}_{disc})$ is $\Omega(n)$.*

Proof. Let E be an election drawn from $\mathcal{P}$. To make w the majority winner, a single voter that ranks w second must change his vote. Every other alternative is ranked first by $\frac{n}{2(m-1)}$ voters. Hence, to make any other alternative the majority winner, at least $\frac{n}{2} + 1 - \frac{n}{2(m-1)} = \Omega(n)$ votes must be changed. By Lemma 1, every strategyproof randomized voting rule obtains an approximation ratio of $\Omega(n)$. □

Theorem 11. *Plurality is $(\mathcal{M}, \hat{d}_{disc})$-rationalizable.*

Proof. Clearly, a majority winner is also the plurality winner. Assume there is no majority winner. For alternative $a \in A$, the distance from E to the closest election in which a is a majority winner is $\lfloor \frac{n}{2} \rfloor + 1 - sc(a)$ (where $sc(a)$ is the plurality score of alternative a). Hence, the alternative with the highest plurality score also has the lowest distance. □

The distance rationalizations $(\mathcal{U}, \hat{d}_{disc})$ and $(\mathcal{M}, \hat{d}_{disc})$ have a large difference in approximation ratios achievable by strategyproof randomized voting rules even though they implement the same voting rule.

Procaccia [15] showed that veto can be approximated well with respect to maximizing the selected alternative's score. We show that veto cannot be approximated well with respect to minimizing the number of vetoes.

Let $\mathcal{V}$ be the consensus class consisting of elections in which some alternative is not vetoed. The consensus winner(s) are those alternatives that receive no vetoes.

Theorem 12. *Veto is $(\mathcal{V}, \hat{d}_{disc})$-rationalizable.*

Proof sketch. It suffices to observe that the distance score of an alternative a is simply the number of vetoes that it receives. If alternative a has fewer vetoes than alternative b, then the distance score of a is less than that of b. □

For $a \in A$, let $v(a)$, be the number of voters that veto a. Since the total number of vetoes by all voters is n, uniform randomly selecting an alternative results in an approximation ratio of $\frac{1}{m}\sum_{a \in A} v(a) = \frac{n}{m}$.

Theorem 13. *The approximation ratio obtained by any strategyproof randomized voting rule to* $(\mathcal{V}, \hat{d}_{disc})$ *is* $\Omega(\frac{n}{m})$.

Proof sketch. Let $\mathcal{P}''$ be a distribution over preference profiles constructed similarly to $\mathcal{P}$, except a single voter, selected uniformly at random, ranks w last. Let E be an election drawn from $\mathcal{P}''$. Since w is vetoed by a single voter, $d(w) = 1$. However, each $a \in A \setminus \{w\}$ is vetoed by $\Omega(\frac{n}{m})$ voters. The Theorem then follows by Lemma 1. □

4. CONCLUSIONS

This paper explores the idea of measuring the approximation ratio achieved by a strategyproof randomized voting rule with respect to a particular distance rationalization. Indeed, if a particular voting rule is employed in a given domain due to a domain specific distance rationalization, then the most natural measure of approximation is with respect to that rationalization.

This paper shows that the unanimity consensus class can be approximated well for a large class of distances by a single strategyproof randomized voting rule. It is shown that the Random Dictator voting rule (select the first choice alternative of a randomly selected voter) nontrivially approximates a large number of distance rationalizations with respect to unanimity. For a number of these distances, nearly tight lower bounds are presented. It is shown that deviating too greatly from the Random Dictator rule results in a trivial approximation ratio (i.e., the ratio obtained by ignoring the preference profile and selecting a random alternative).

The outlook for consensus classes, other than unanimity is bleaker. It is also shown that no strategyproof randomized voting rule nontrivially approximates many distance rationalizations with respect to the Majority and Condorcet consensus classes.

There exist a number of other distance rationalizations of common voting rules other than those considered in this paper. For example, the Copeland rule selects the alternative that maximizes the number of pairwise elections it wins and is rationalizable with respect to the Condorcet consensus class [13]. Future work will investigate the quality of strategyproof approximations obtainable to such rationalizations.

Another line of future work is to consider rationalization frameworks other than distance rationalization. For example, under the maximum likelihood estimation framework, one can measure the approximation ratio achieved by a randomized voting rule as a function of the likelihood of the alternative selected by that rule to being the true winner compared to the likelihood of the actual winner.

5. REFERENCES

[1] Eleanor Birrell and Rafael Pass. Approximately strategy-proof voting. In *Proceedings of the 22nd International Joint Conference on Artificial intelligence*, pages 67–72, 2011.

[2] Vincent Conitzer and Tuomas Sandholm. Universal voting protocol tweaks to make manipulation hard. In *Proceedings of the 18th International Joint Conference on Artificial intelligence*, pages 781–788, 2003.

[3] Vincent Conitzer and Tuomas Sandholm. Nonexistence of voting rules that are usually hard to manipulate. In *Proceedings of the 21st National Conference on Artificial intelligence*, pages 627–634. AAAI Press, 2006.

[4] Edith Elkind, Piotr Faliszewski, and Arkadii Slinko. On distance rationalizability of some voting rules. In *Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge*, pages 108–117. ACM, 2009.

[5] Edith Elkind, Piotr Faliszewski, and Arkadii Slinko. On the role of distances in defining voting rules. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, pages 375–382, 2010.

[6] Edith Elkind, Piotr Faliszewski, and Arkadii Slinko. Homogeneity and monotonicity of distance-rationalizable voting rules. In *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, pages 821–828, 2011.

[7] Edith Elkind, Piotr Faliszewski, and Arkadii Slinko. Rationalizations of condorcet-consistent rules via distances of hamming type. *Social Choice and Welfare*, pages 1–15, 2011.

[8] Edith Elkind and Helger Lipmaa. Hybrid voting protocols and hardness of manipulation. In *In Proceedings of the 16th International Symposium on Algorithms and Computation*, pages 206–215. Springer-Verlag, 2005.

[9] Piotr Faliszewski and Ariel D. Procaccia. Ai's war on manipulation: Are we winning. *AI Magazine*, 31(4):53–64, 2010.

[10] Ehud Friedgut, Gil Kalai, and Noam Nisan. Elections can be manipulated often. In *Proceedings of the 2008 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 243–249, 2008.

[11] Allan Gibbard. Manipulation of voting schemes. *Econometrica*, 41(4):587–601, 1973.

[12] Allan Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 25(3), 1977.

[13] Christian Klamler. The Copeland rule and condorcet's principle. *Economic Theory*, 25:745–749, 2005.

[14] Tommi Meskanen and Hannu Nurmi. Closeness counts in social choice. In *Power, Freedom, and Voting*. Springer-Verlag, 2008.

[15] Ariel Procaccia. Can approximation circumvent Gibbard-Satterthwaite? In *Proceedings of the 24th Conference on Artificial Intelligence*, pages 836–841, 2010.

[16] 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, 1975.

[17] Andrew Yao. Probabilistic computations: Toward a unified measure of complexity. In *Proceedings of the 18th Annual Symposium on Foundations of Computer Science*, pages 222–227, 1977.