Can Approximation Circumvent Gibbard-Satterthwaite?

Ariel D. Procaccia
Harvard SEAS
arielpro@seas.harvard.edu

Abstract

The Gibbard-Satterthwaite Theorem asserts that any reasonable voting rule cannot be strategyproof. A large body of research in AI deals with circumventing this theorem via computational considerations; the goal is to design voting rules that are computationally hard, in the worst-case, to manipulate. However, recent work indicates that the prominent voting rules are usually easy to manipulate.

In this paper, we suggest a new CS-oriented approach to circumventing Gibbard-Satterthwaite, using randomization and approximation. Specifically, we wish to design strategyproof randomized voting rules that are close, in a standard approximation sense, to prominent score-based (deterministic) voting rules. We give tight lower and upper bounds on the approximation ratio achievable via strategyproof randomized rules with respect to positional scoring rules, Copeland, and Maximin.

Introduction

In the last decade or so Computational Social Choice has taken center stage in AI research. Perhaps the most prominent issue in this field is the problem of manipulation in elections. An agent is said to *manipulate* the election if it misreports its true preferences over the alternatives in order to achieve a better outcome. A voting rule—a function that designates the winning alternative given the agents' preferences—is deemed *strategyproof (SP)* if it is immune to manipulation.

In the seventies Gibbard (1973) and Satterthwaite (1975) independently proved the following striking impossibility theorem: if there are at least three alternatives, no voting rule can satisfy the following three properties simultaneously: (i) any alternative can be elected; (ii) the rule is not dictatorial, that is, there is no single agent that determines the outcome of the election; (iii) the rule is strategyproof. In other words, any "reasonable" voting rule cannot be strategyproof.

A CS-oriented approach to circumventing the Gibbard-Satterthwaite Theorem was suggested by Bartholdi et al. (1989). They argued that although manipulation may be possible in theory, it may prove to be a computationally intractable problem under some voting rules. Indeed, it was established that the well-known Single Transferable Vote (STV) rule is NP-hard to manipulate (Bartholdi & Orlin 1991). Subsequent work has established that a variety of voting rules are computationally hard to manipulate under different assumptions on the manipulation setting (see, e.g., (Conitzer & Sandholm 2003; Conitzer, Sandholm, & Lang 2007; Hemaspaandra & Hemaspaandra 2007; Faliszewski *et al.* 2008; Meir *et al.* 2008)). Nevertheless, recent papers argue that worst-case complexity cannot realistically prevent manipulation, and in fact manipulation may be usually feasible in practice (see, e.g., (Conitzer & Sandholm 2006; Procaccia & Rosenschein 2007; Friedgut, Kalai, & Nisan 2008; Xia & Conitzer 2008; Dobzinski & Procaccia 2008; Zuckerman, Procaccia, & Rosenschein 2009)).

In the discussion of their influential paper, Conitzer and Sandholm (2006, pp. 632–633) proposed (in passing) a different approach to circumventing Gibbard-Satterthwaite: using randomization. Indeed, Gibbard (1977) has characterized the family of SP randomized voting rules, where an SP randomized rule has the property that an agent cannot increase its expected utility from lying, for any utility function. The characterization is quite restrictive but nevertheless allows some freedom in designing SP randomized voting rules. Conitzer and Sandholm use Gibbard's characterization to design an SP randomized rule that is "close" to the prominent Copeland rule in an informal sense.

Our approach and results. In this paper we take a step forward in the direction indicated by Conitzer and Sandholm. We suggest that although prominent voting rules are not SP, it is possible to design randomized SP voting rules that are close to the prominent ones in a formal sense. Specifically, we wish to *approximate* prominent *score-based* voting rules. By score-based we mean that the voting rule has some clear notion of score, and that it elects an alternative with maximum score. We say that a randomized voting rule has an *approximation ratio* of γ with respect to a score-based (deterministic) voting rule if the expected score of an alternative elected by the randomized rule is within a γ-fraction of the maximum score. The logic is that if a randomized rule provides a good approximation to a prominent rule then it would often elect an alternative that is optimal or nearly-optimal according to the target rule, and hence the randomized rule would make a desirable preference aggregation method in its own right. In essence, we study a new CS-oriented approach to preventing manipulation by employing randomization and approximation, two concepts that are prevalent in algorithmics.

We first consider the family of positional scoring rules. We show that any positional scoring rule can be approximated by an SP randomized voting rule to a factor of $\Omega(1/\sqrt{m})$, where m is the number of alternatives. This result holds in particular for the Plurality rule, where each agent votes for a single alternative. We demonstrate that no SP randomized rule can do asymptotically better. We further show that Borda—where each agent awards $m-k$ points to the alternative it ranks kth—can be approximated by an SP randomized rule to a factor of $1/2+\Omega(1/m)$, and that this bound is essentially tight.

We next investigate Copeland's rule, where the score of an alternative is the number of other alternatives it beats in pairwise elections. We show that Copeland can be approximated by an SP randomized rule within $1/2+\Omega(1/m)$, and provide a matching upper bound. Finally, we show that the Maximin rule cannot be approximated to any nontrivial factor by SP randomized rules.

We would like to stress that we do not use approximation to circumvent computational complexity, as in work on algorithmically approximating hard-to-compute voting rules, e.g., Dodgson's rule (Caragiannis *et al.* 2009). Rather, approximation is employed to achieve strategyproofness. Indeed, by the Gibbard-Satterthwaite Theorem the prominent voting rules are not SP, but by randomly deviating from the original rules we can achieve strategyproofness. The question is how close we can be while maintaining strategyproofness. In this sense this paper is related to the recent work on approximate mechanism design without money (Procaccia & Tennenholtz 2009).

Preliminaries

Let $N=\{1,\ldots,n\}$ be a set of agents, and let A be a set of alternatives, $|A|=m$. The preferences of agent $i\in N$ are given by a linear order $\succ_i\in\mathcal{L}$ over the alternatives, where $\mathcal{L}=\mathcal{L}(A)$ is the set of all linear orders over A. A collection of the agents' preferences $\succ=\langle\succ_1,\ldots,\succ_n\rangle\in\mathcal{L}^n$ is called a *preference profile*. Given $\succ_i$, we denote by $\succ_i(x)\in\{1,\ldots,m\}$ the position of $x\in A$ in $\succ_i$; to avoid confusion note that position 1 is the top position and position m is the bottom position.

A *voting rule* is a function $f:\mathcal{L}^n\to A$ that designates a winning alternative given a preference profile. A *randomized voting rule* (also known as a *decision scheme* (Gibbard 1977)) is a function $f:\mathcal{L}^n\to\Delta(A)$ that returns a probability distribution over A. In other words, a randomized voting rule is allowed to use randomization when selecting an alternative.

We informally define manipulation of randomized voting rules, following the definition of Gibbard (1977); a formal definition is not required for the purposes of this paper. Associate with agent $i\in N$ a utility function $u_i:A\to\mathbb{R}_+$. We say that $\succ_i$ is consistent with u_i if $x\succ_i y$ implies $u_i(x)\geq u_i(y)$. Agent i *manipulates* a randomized voting rule f if there exist preferences for the other agents such that the expected utility (with respect to u_i) from reporting a ranking that is inconsistent with u_i is greater than the expected utility from reporting a ranking that is consistent with u_i. A randomized voting rule is *strategyproof (SP)* if it cannot be manipulated for any $i\in N$ and u_i. This definition is quite strong as it requires nonmanipulability with respect to every utility function.

We next present a necessary condition for SP randomized rules. We say that a voting rule f is *unilateral* if it only depends on the vote of a single agent, that is, if there exists $i\in N$ such that for every $\succ,\succ'\in\mathcal{L}^n$ where $\succ_i=\succ'_i$ it holds that $f(\succ)=f(\succ')$. We say that a voting rule f is *duple* if its range is of size at most two, i.e., there exist $x,y\in A$ such that $f(\succ)\in\{x,y\}$ for all $\succ\in\mathcal{L}^n$. A randomized voting rule f is a *probability mixture* over voting rules $f_1,\ldots,f_k$ if there exist $\alpha_1,\ldots,\alpha_k$, with $\alpha_j\geq 0$ for all $j=1,\ldots,k$ and $\sum_{j=1}^k\alpha_j=1$, such that given $\succ\in\mathcal{L}^n$ it holds that $f(\succ)=f_j(\succ)$ with probability α_j.

Theorem 1 (Gibbard (1977))**.** *An SP randomized voting rule is a probability mixture over rules each of which is either unilateral or duple.*

Our goal in this paper is to construct SP randomized rules that approximate prominent score-based voting rules. We say that $f:\mathcal{L}^n\to A$ is a *score-based voting rule with score function* $\mathrm{sc}:A\times\mathcal{L}^n\to\mathbb{N}$ if for every $\succ\in\mathcal{L}^n$, $f(\succ)$ is an alternative $x\in A$ that maximizes the score $\mathrm{sc}(x,\succ)$. A randomized rule f is said to yield an *approximation ratio of* γ with respect to a score-based voting rule f' with score function sc if the expected score of the elected alternative is optimal up to a factor of γ, that is, for all $\succ\in\mathcal{L}^n$,

$$\mathbb{E}[\mathrm{sc}(f(\succ),\succ)]\geq\gamma\cdot\max_{x\in A}\mathrm{sc}(x,\succ)\ ,$$

where the expectation is taken over the randomization of f.

Positional Scoring Rules

Positional scoring rules are a central family of score-based voting rules, which have received significant attention in AI in recent years (see, e.g., (Hemaspaandra & Hemaspaandra 2007; Procaccia & Rosenschein 2007) for papers that focus exclusively on positional scoring rules). A positional scoring rule is defined by a vector of nonnegative real numbers $\vec{\alpha}=\langle\alpha_1,\ldots,\alpha_m\rangle$. The score of an alternative given a preference profile $\succ\in\mathcal{L}^n$ is defined as

$$\mathrm{sc}(x,\succ)=\sum_{i\in N}\alpha_{\succ_i(x)}\ .$$

Informally, the alternative is awarded α_j points by every agent that ranks it in the jth place. There are three prominent scoring rules:

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

In order to gain some intuition, let us consider the randomized voting rule that simply selects an alternative uniformly at random. This rule is clearly SP. However, in general it cannot guarantee a nontrivial approximation ratio with respect to positional scoring rules. Indeed, consider the Plurality rule and a preference profile where all the agents rank an alternative $a \in A$ first. It holds that the optimum is n, whereas the expected Plurality score of a random alternative is n/m. Hence, choosing a random alternative provides an approximation ratio of at most $1/m$.

Let us now consider a slightly more sophisticated method; the following generic randomized voting rule depends on the parameters of the positional scoring rule.

Rule 1. Select an agent $i \in N$ uniformly at random. Elect the winner according to the following probability distribution: the probability of alternative $x \in A$ is $\alpha_{\succ_i(x)}/\sum_{j=1}^{m}\alpha_j$.

Note that this rule is a probability mixture over unilateral rules. Indeed, for every $i \in N$, $j \in \{1,\ldots,m\}$, let f_{ik} be the unilateral rule that, given $\succ \in \mathcal{L}^n$, selects the alternative ranked in the kth position of $\succ_i$. Then Rule 1 uses each f_{ik} with probability $\alpha_k/(n \cdot \sum_{j=1}^{m}\alpha_j)$.

Next we observe that Rule 1 selects an alternative $x \in A$ with probability

$$\sum_{i\in N}\frac{1}{n}\cdot\frac{\alpha_{\succ_i(x)}}{\sum_{j=1}^{m}\alpha_j}=\frac{\mathrm{sc}(x,\succ)}{\sum_{y\in A}\mathrm{sc}(y,\succ)}\quad . \tag{1}$$

Put another way, the rule selects each alternative with probability proportional to its score.

Furthermore, notice that Rule 1 is SP. This is rather obvious, as if agent i reported $\succ_i$ and is selected by Rule 1 then the expected utility with respect to u_i is

$$\sum_{x\in A}\frac{\alpha_{\succ_i(x)}}{\sum_{j=1}^{m}\alpha_j}\cdot u_i(x)\quad ,$$

and this expression is clearly maximized when $\succ_i$ is consistent with u_i. If, on the other hand, agent i is not selected by the rule, then its vote is irrelevant to the outcome.

Finally, we provide a simple analysis that establishes the approximation ratio yielded by Rule 1 with respect to positional scoring rules.

Theorem 2. *Let f be a positional scoring rule with parameters $\vec{\alpha}$. Then the approximation ratio of Rule 1 with respect to f is $\Omega(1/\sqrt{m})$.*

Proof. Let $\succ \in \mathcal{L}^n$ be a preference profile. Assume without loss of generality that $a \in A$ maximizes the score according to f, and let $\mathrm{OPT} = \mathrm{sc}(a,\succ)$. We also denote

$$\mathrm{SUM}=\sum_{x\in A}\mathrm{sc}(x,\succ)=n\cdot\sum_{j=1}^{m}\alpha_j\quad .$$

By Equation (1), alternative x is selected by Rule 1 with probability $\mathrm{sc}(x,\succ)/\mathrm{SUM}$. Therefore, the expected score of the winner according to the rule is

$$\begin{aligned}&\frac{\mathrm{OPT}}{\mathrm{SUM}}\cdot\mathrm{OPT}+\sum_{x\in A\setminus\{a\}}\frac{\mathrm{sc}(x,\succ)}{\mathrm{SUM}}\cdot\mathrm{sc}(x,\succ)\\ &\geq\frac{\mathrm{OPT}^2}{\mathrm{SUM}}+\frac{1}{\mathrm{SUM}}\sum_{x\in A\setminus\{a\}}\left(\frac{\mathrm{SUM}-\mathrm{OPT}}{m-1}\right)^2\\ &=\frac{1}{\mathrm{SUM}}\cdot\left[\mathrm{OPT}^2+\frac{(\mathrm{SUM}-\mathrm{OPT})^2}{m-1}\right]\quad ,\end{aligned}$$

where the first transition follows from the fact that $\sum_{x\in A\setminus\{a\}}\mathrm{sc}(x,\succ)=\mathrm{SUM}-\mathrm{OPT}$. The approximation ratio achieved by the algorithm is then at least

$$\frac{1}{\mathrm{SUM}}\cdot\frac{\mathrm{OPT}^2+\frac{(\mathrm{SUM}-\mathrm{OPT})^2}{m-1}}{\mathrm{OPT}}\quad . \tag{2}$$

By differentiating with respect to OPT we conclude that the expression in Equation 2 is minimized when $\mathrm{OPT} = \mathrm{SUM}/\sqrt{m}$. Substituting this expression back into Equation (2) gives the announced bound. ☐

It is natural to ask whether the analysis in the proof of Theorem 2 is tight. The answer depends on the parameters of the scoring rule, but the answer is "yes" with respect to some positional scoring rules. In particular, the theorem below establishes a powerful and general lower bound with respect to Plurality: no SP randomized voting rule can approximate Plurality to a factor asymptotically better than the one given in the statement of Theorem 2, namely $\omega(1/\sqrt{m})$. The theorem's proof exploits Gibbard's Theorem (Theorem 1), and is given in the appendix.

Theorem 3. *No SP randomized voting rule can approximate Plurality to a factor of $\omega(1/\sqrt{m})$.*

Next we observe that the upper bound given in Theorem 3 does not hold with respect to some prominent positional scoring rules, e.g., Borda. Returning to the proof of Theorem 2, note that for Borda it holds that $\mathrm{OPT} \leq n(m-1)$, whereas $\mathrm{SUM} = nm(m-1)/2$. Therefore, $\mathrm{OPT}/\mathrm{SUM} \leq 2/m$, and the minimum of the expression in Equation (2) is achieved when $\mathrm{OPT} = n(m-1)$. Substituting these values back into the expression, we obtain that Rule 1 yields an approximation ratio of $1/2+\Omega(1/m)$.[1] We have obtained the following corollary .

Corollary 4. *Rule 1 gives a $(1/2+\Omega(1/m))$-approximation with respect to Borda.*

It turns out that for the case of Borda even randomly choosing an alternative (which is, as noted above, SP) in fact gives a similar, but slightly worse, approximation ratio. Indeed, the expected Borda score of a random alternative in a profile $\succ \in \mathcal{L}^n$ is

$$\sum_{x\in A}\frac{1}{m}\cdot\mathrm{sc}(x,\succ)=\frac{1}{m}\cdot\frac{nm(m-1)}{2}=\frac{n(m-1)}{2}\quad .$$

Since the optimum may be at most $n(m-1)$, this gives an approximation ratio of 1/2.

Interestingly, it is possible to establish a matching lower bound, that is, no SP randomized voting rule can give an approximation ratio that is bounded away from 1/2 with respect to Borda. In other words, it is essentially impossible to do better than randomly choose an alternative. The rather straightforward technique used in the proof of Theorem 3 falls short here. Instead, we employ Yao's Minimax principle (Yao 1977).

[1] In general the approximation ratio is small for any positional scoring rule where SUM is far greater than OPT.

Theorem 5. *No SP randomized voting rule can approximate Borda to a factor of* $1/2 + \omega(1/\sqrt{m})$.

Proof. We consider the zero-sum game where the strategies of the row player are unilateral and duple rules, and the strategies of the column player are preference profiles. The payoff in the cell that corresponds to a profile $\succ \in L^n$ and the rule $f : \mathcal{L}^n \to A$ is the ratio between the Borda scores $\mathrm{sc}(f(\succ), \succ)/(\max_{x\in A} \mathrm{sc}(x, \succ))$. The row player wishes to maximize the payoff, whereas the column player wishes to minimize the payoff. Note that, by Gibbard's Theorem, the set of mixed strategies of the row player contains all the SP randomized voting rules. By the Minimax theorem, an upper bound on the approximation ratio of the best randomized SP voting rule is given by the expected approximation ratio achieved by the best (deterministic) unilateral or duple rule under some specific distribution over preference profiles. In the following we therefore construct a "bad" distribution over profiles.

Let $n = m-1$, and assume for ease of exposition that $\sqrt{m}$ is an integer. We consider the distribution over preference profiles induced by the following procedure.

1. Select an alternative $x^* \in A$ uniformly at random.
2. For each $i \in N$, select a position $k_i \in \{1, \ldots, \sqrt{m}\}$ independently and uniformly at random. Agent i ranks x^* in position k_i.
3. Choose a permutation π_1 over $\{1, \ldots, m-1\}$ uniformly at random. For every $i = 1, \ldots, n-1$, let $\pi_{i+1}(j) = \pi_i(j+1)$, where the addition is cyclic, i.e., $(m-1)+1 = 1$.
4. Denote $A \setminus \{x^*\} = \{x_1, \ldots, x_{m-1}\}$. Intuitively, the agents rank the alternatives in $A \setminus \{x^*\}$ cyclically in the positions $1, \ldots, k_i - 1, k_i + 1, \ldots, m$. Formally, for $j = 1, \ldots, k_i - 1$, agent i ranks alternative $x_{\pi_i(j)}$ in position j. For $j = k_i + 1, \ldots, m$, i ranks alternative $x_{\pi_i(j-1)}$ in position j. See Table 1 for an example.

1	2	3
c	b	d
b	a	b
a	d	c
d	c	a

Table 1: An example of a profile generated by the procedure in the proof of Theorem 5. We have that $x^* = b$, which is placed in positions $k_1 = 2, k_2 = 1, k_3 = 2$ in the rankings of the different agents. The other three alternatives are ranked cyclically in the remaining positions.

Let us compute the Borda scores of the different alternatives in a preference profile $\succ \in \mathcal{L}^n$ generated according to the above procedure. The alternative x^* is always ranked in the first $\sqrt{m}$ positions, hence

$$\mathrm{sc}(x^*, \succ) \geq (m-1) \cdot (m - \sqrt{m}) \ . \tag{3}$$

On the other hand, clearly $\mathrm{sc}(x^*, \succ) \leq (m-1)^2$. Furthermore, an upper bound on the score of each alternative $x_j \in A \setminus \{x^*\}$ is certainly given by its score when it is ranked once in each of the first $m-1$ positions, that is,

$$\mathrm{sc}(x_j, \succ) \leq \sum_{k=1}^{m-1} (m-k) = \frac{m(m-1)}{2} \ . \tag{4}$$

We first consider the case where $f : \mathcal{L}^n \to A$ is a unilateral rule with respect to agent i. We claim that f chooses the special alternative x^* with small probability. Intuitively this is true since, when only $\succ_i$ is considered, there is nothing that distinguishes x^* from the other alternatives in positions $1, \ldots, \sqrt{m}$.

Formally, let $\succ \in \mathcal{L}^n$ be a preference profile generated according to our procedure. As before, let $x^* \in A$ be the special alternative with high Borda score with respect to $\succ$. Then we claim that

$$\Pr[f(\succ) = x^* \mid \succ_i] \leq \frac{1}{\sqrt{m}} \ , \tag{5}$$

where the probability is taken over profiles generated by the procedure, given that the ranking of agent i is $\succ_i$. Indeed, observe that since π_1 is a random permutation of $\{1, \ldots, m-1\}$, π_i is also a random permutation of $\{1, \ldots, m-1\}$. In other words, our procedure for generating profiles is equivalent to the following procedure: let $\succ_i$ be a random permutation of the alternatives; choose x^* at random among the first $\sqrt{m}$ alternatives of $\succ_i$, then complete the profile as before, that is, by choosing a position for x^* in every vote and shifting the rest of the alternatives cyclically. Viewed this way, it is easy to see that the procedure has the following property: for every two alternatives $y, y' \in A$ in the first $\sqrt{m}$ positions of $\succ_i$ we have that

$$\Pr[y = x^* \mid \succ_i] = \Pr[y' = x^* \mid \succ_i] \ . \tag{6}$$

Assume that $f(\succ)$ is an alternative in the first $\sqrt{m}$ positions of $\succ_i$. By Equation (6), for every alternative $y \in A$ in the first $\sqrt{m}$ positions of $\succ_i$ it holds that

$$\Pr[y = x^* \mid \succ_i] = \frac{1}{\sqrt{m}} \ ,$$

and in particular this is true for $f(\succ)$. In addition, if $f(\succ)$ is not an alternative in the first $\sqrt{m}$ positions of $\succ_i$ then it cannot be x^*, that is,

$$\Pr[f(\succ) = x^* \mid \succ_i] = 0 \ .$$

Equation (5) directly follows.

We can now conclude that for every $\succ_i \in \mathcal{L}$,

$$\begin{aligned} \mathbb{E}\left[\frac{\mathrm{sc}(f(\succ), \succ)}{\mathrm{sc}(x^*, \succ)} \middle| \succ_i \right] &\leq \frac{\frac{1}{\sqrt{m}}(m-1)^2 + \frac{\sqrt{m}-1}{\sqrt{m}} \cdot \frac{m(m-1)}{2}}{(m-1)(m-\sqrt{m})} \\ &\leq \frac{1}{2} + \frac{\sqrt{m}}{m - \sqrt{m}} \\ &= \frac{1}{2} + \mathcal{O}\left(\frac{1}{\sqrt{m}}\right) \ , \end{aligned}$$

where the first inequality is obtained by combining Equations (3), (4), and (5). Since this result holds for every $\succ_i \in \mathcal{L}$ we have that

$$\mathbb{E}\left[\frac{\mathrm{sc}(f(\succ), \succ)}{\mathrm{sc}(x^*, \succ)} \right] \leq \frac{1}{2} + \mathcal{O}\left(\frac{1}{\sqrt{m}}\right) \ .$$

We next consider the easier case where $f : \mathcal{L}^n \to A$ is a duple rule. Since x^* is selected at random, the probability that x^* is in the range of f is at most $2/m$, hence the expected ratio is even worse. By Yao's Minimax principle, we conclude that $1/2 + \mathcal{O}(1/\sqrt{m})$ is an upper bound on the approximation ratio that can be achieved via an SP randomized voting rule $\square$

We finally note that with respect to Veto, Rule 1 gives a nearly perfect approximation ratio of $1 - \mathcal{O}(1/m)$; this can be seen by substituting OPT $= n$ and SUM $= n(m-1)$ into Equation (2). However, this may be a strange consequence of Veto's representation as a positional scoring rule. In fact, one is intuitively interested in minimizing the number of vetoes, that is, the number of agents that dislike a given alternative above all others. However, this target is difficult to approximate as the optimum may be zero. We elaborate further in the discussion.

Copeland and Lull

For a preference profile $\succ \in \mathcal{L}^n$ and a pair of alternatives $x, y \in A$ let

$$P(x,y) = |\{i \in N : x \succ_i y\}| \ .$$

Given $\alpha \in [0,1]$, define the Copeland$_\alpha$ score of $x \in A$ as

$$\begin{aligned} \mathrm{sc}(x,\succ) = {} & |\{y \in A \setminus \{x\} :\ P(x,y) > n/2\}| \\ & + \alpha \cdot |\{y \in A \setminus \{x\} :\ P(x,y) = n/2\}| \ . \end{aligned}$$

Put another way, an alternative receives a point for each other alternative that it beats in a *pairwise election*, and receives an α-fraction of a point for every tie. The importance of the tie-breaking issue was discussed in a number of recent papers (see, e.g., (Faliszewski *et al.* 2008; 2009)). Nowadays Copeland$_{1/2}$ is attributed to A. H. Copeland, whereas Copeland$_1$ is credited to the 13th Century Majorcan philosopher Ramon Llull.

We consider the following randomized voting rule.

Rule 2. Choose a pair of alternatives uniformly at random. If one is preferred to the other by a majority of agents then it is the winner. Otherwise, flip a fair coin.

This rule is a probability mixture over duple rules. It is furthermore straightforward that Rule 2 is SP: for a fixed pair of alternatives that is selected by the rule, the manipulator is better off if the one with higher utility is elected, and hence can only benefit by reporting its preferences truthfully with respect to the pair. Since the argument holds for every pair, we conclude that the manipulator's best option is to report truthfully. This was also observed in (Conitzer & Sandholm 2006, Theorem 3).

Similarly to Rule 1 in the context of positional scoring rules, Rule 2 elects an alternative with probability proportional to its Copeland$_{1/2}$ score. Indeed, the probability of electing alternative $x \in A$ is

$$\begin{aligned} & \sum_{y \in A:\ P(x,y) > n/2} \frac{1}{\binom{m}{2}} + \sum_{y \in A:\ P(x,y) = n/2} \frac{1}{\binom{m}{2}} \cdot \frac{1}{2} \\ & = \frac{\mathrm{sc}(x,\succ)}{\binom{m}{2}} = \frac{\mathrm{sc}(x,\succ)}{\sum_{y \in A} \mathrm{sc}(y,\succ)} \ . \end{aligned}$$

We have the following general result.

Theorem 6. *Rule 2 gives a $(1/2 + \Omega(1/m))$-approximation with respect to Copeland$_\alpha$, for every $\alpha \geq 1/2$.*

The proof of this theorem and the two subsequent proofs are relegated to the appendix. Once again, it is possible to show that selecting an alternative uniformly at random provides a slightly worse approximation ratio of exactly 1/2 with respect to Copeland$_{1/2}$. We elaborate in the discussion section.

To complete the picture, we next show that it it impossible to to achieve an approximation ratio bounded away from 1/2 using an SP randomized voting rule.

Theorem 7. *Let $\alpha \in [0,1]$. No SP randomized voting rule can approximate Copeland$_\alpha$ to a factor of $1/2 + \omega(1/m)$.*

Maximin

The Maximin rule is also known as Simpson's rule. The score of an alternative under Maximin is defined as

$$\mathrm{sc}(x,\succ) = \min_{y \in A \setminus \{x\}} P(x,y) \ ,$$

where as before we define

$$P(x,y) = |\{i \in N : x \succ_i y\}| \ .$$

Put another way, the score of x is the outcome of its worst pairwise election.

It is easy to see that the natural SP randomized voting rules, such as the ones employed above, fail in providing a nontrivial approximation ratio with respect to Maximin. As it turns out, this is the case with respect to any SP randomized voting rule.

Theorem 8. *No SP randomized voting rule can approximate Maximin to a factor of $\omega(1/m)$.*

The intuition behind this result is that in some profiles the Maximin score depends on a very specific pairwise competition. This makes the score hard to determine using a random duple rule (in stark contrast to Copeland). Moreover, unilateral rules are in general unhelpful when it comes to approximating voting rules that depend on pairwise elections. Hence, probability mixtures over unilateral and duple rules cannot approximate Maximin.

Discussion

Table 2 summarizes our results; the table gives the lower and upper bounds achievable by SP randomized rules with respect to the voting rules in the left column. All our bounds are tight (up to lower order terms).

Despite the tightness of the results, one might wonder whether the lower bounds (that is, the possibility results) are meaningful. In particular, in the context of Borda and Copeland the "proportional rules", Rules 1 and 2, give only slightly better bounds than choosing an alternative at random. We wish to point out that the differences in the approximation ratio might be meaningful when the number of alternatives is small. More importantly, we would like to design randomized voting rules that are "reasonable" in terms

Rule	Lower bound	Upper bound
Plurality	$\Omega(1/\sqrt{m})$	$\mathcal{O}(1/\sqrt{m})$
Borda	$1/2 + \Omega(1/m)$	$1/2 + \mathcal{O}(1/\sqrt{m})$
Copeland$_\alpha$	$1/2 + \Omega(1/m)$*	$1/2 + \mathcal{O}(1/m)$
Maximin	0	$\mathcal{O}(1/m)$

Table 2: A summary of the results. *Only $\alpha \in [1/2, 1]$.

of their social choice properties. In this sense, Rules 1 and 2 seem much better than the uniformly random election. For example, a Condorcet loser (an alternative that loses to every other alternative in a pairwise election) has no chance of being elected under Rule 2.

A minimal axiom that one might ask randomized voting rules to satisfy is *positive response*, informally defined as follows: for every agent there is some profile where the agent can increase an alternative's probability of being elected by pushing it upwards in its vote. This property is satisfied by Rules 1 and 2, and is not satisfied when a uniformly random alternative is elected.

More generally, future research should specify the desiderata that randomized voting rules must satisfy; it will then be possible to ask whether there exist randomized voting rules that are strategyproof, satisfy the additional properties, and yield a good approximation ratio with respect to prominent voting rules.

We have obtained results with respect to natural, prominent score-based voting rules. However, there are some rules, which are formally score-based[2], where it is unclear which notion of score is the most natural one. One example is Veto; we have seen that, as a scoring rule, Veto is easy to approximate, but it remains open whether it is possible to approximate Veto when the winner is the alternative that minimizes the number of vetoes. Bucklin is potentially another interesting example. Given a preference profile and an alternative x, let j be the minimum position such that a majority of agents rank x in position j or above; define the Bucklin score to be $m - j$. According to this definition choosing a random alternative provides an approximation ratio of 1/4, but using a surprising SP randomized voting rule one can obtain an approximation ratio of at least 13/48, that is, a ratio that is bounded away from what can be obtained trivially.

We finally return to our original question, "Can approximation circumvent Gibbard-Satterthwaite?" Unfortunately it is difficult to argue that the results in this paper provide a positive answer. Nevertheless, our results do not preclude good SP approximations with respect to notions of score that were not considered above. In conclusion, we believe that our approach offers new insights with respect to the manipulation problem, and may direct future work towards the design of practical, applicable SP randomized rules.

Acknowledgments

Vincent Conitzer has made important contributions to this paper, especially during a visit to the Hebrew University of Jerusalem. Conitzer's visit was supported by the United States-Israel Binational Science Foundation grant 2006-216.

[2]In fact every voting rule is score-based with respect to some *ad hoc* score function.

References

Bartholdi, J., and Orlin, J. 1991. Single Transferable Vote resists strategic voting. *Social Choice and Welfare* 8:341–354.

Bartholdi, J.; Tovey, C. A.; and Trick, M. A. 1989. The computational difficulty of manipulating an election. *Social Choice and Welfare* 6:227–241.

Caragiannis, I.; Covey, J. A.; Feldman, M.; Homan, C. M.; Kaklamanis, C.; Karanikolas, N.; Procaccia, A. D.; and Rosenschein, J. S. 2009. On the approximability of Dodgson and Young elections. In *Proc. of 20th SODA*, 1058–1067.

Conitzer, V., and Sandholm, T. 2003. Universal voting protocol tweaks to make manipulation hard. In *Proc. of 18th IJCAI*, 781–788.

Conitzer, V., and Sandholm, T. 2006. Nonexistence of voting rules that are usually hard to manipulate. In *Proc. of 21st AAAI*, 627–634.

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

Dobzinski, S., and Procaccia, A. D. 2008. Frequent manipulability of elections: The case of two voters. In *Proc. of 4th WINE*, 653–664.

Faliszewski, P.; Hemaspaandra, E.; ; and Schnoor, H. 2008. Copeland voting: Ties matter. In *Proc. of 7th AAMAS*, 983–990.

Faliszewski, P.; Hemaspaandra, E.; Hemaspaandra, L. A.; and Rothe, J. 2009. Llull and Copeland voting computationally resist bribery and constructive control. *Journal of Artificial Intelligence Research* 35:275–341.

Friedgut, E.; Kalai, G.; and Nisan, N. 2008. Elections can be manipulated often. In *Proc. of 49th FOCS*, 243–249.

Gibbard, A. 1973. Manipulation of voting schemes. *Econometrica* 41:587–602.

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

Hemaspaandra, E., and Hemaspaandra, L. A. 2007. Dichotomy for voting systems. *Journal of Computer and System Sciences* 73(1):73–83.

Meir, R.; Procaccia, A. D.; Rosenschein, J. S.; and Zohar, A. 2008. Complexity of strategic behavior in multi-winner elections. *Journal of Artificial Intelligence Research* 33:149–178.

Procaccia, A. D., and Rosenschein, J. S. 2007. Junta distributions and the average-case complexity of manipulating elections. *Journal of Artificial Intelligence Research* 28:157–181.

Procaccia, A. D., and Tennenholtz, M. 2009. Approximate mechanism design without money. In *Proc. of 10th EC*, 177–186.

Satterthwaite, M. 1975. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory* 10:187–217.

Xia, L., and Conitzer, V. 2008. Generalized Scoring Rules and the frequency of coalitional manipulability. In *Proc. of 9th EC*, 109–118.

Yao, A. C. 1977. Probabilistic computations: Towards a unified measure of complexity. In *Proc. of 17th FOCS*, 222–227.

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

Appendix: Missing Proofs

Proof of Theorem 3. Assume for ease of exposition that $m = k^2$ for some $k \in \mathbb{N}$, and let $n = m$. Let f be an SP randomized voting rule; then by Theorem 1 f is a probability mixture over rules that are either unilateral or duple. Let p be the probability that f selects a duple rule; the probability that f selects a unilateral rule is $1-p$.

Assume that f selects a duple D. We can simply treat D as a set of two alternatives, since we are indifferent to how D chooses between the two. Define

$$q_x = \Pr[x \in D \mid f \text{ is duple } D] \ .$$

It holds that $\sum_{x \in A} q_x = 2$. Hence, there is a set of k alternatives $A' \subseteq A$ such that

$$\sum_{x \in A'} q_x \leq \frac{2}{k} \ . \tag{7}$$

By the union bound, the probability that at least one member of A' is a member of D is upper-bounded by $\sum_{x \in A'} q_x$, hence by Equation (7)

$$\Pr[\exists x \in A' \text{ s.t. } x \in D \mid f \text{ is duple } D] \leq \frac{2}{k} \ . \tag{8}$$

Next, assume that f selects a unilateral rule. There is a subset $N' \subseteq N$ of k agents such that a unilateral rule over an agent in N' is selected with probability at most $1/k$ (since $n = m = k^2$). We construct a partial preference profile $\succ_{A \setminus N'} \in \mathcal{L}^{m-k}$ for the agents in $N \setminus N'$ that satisfies the following constraint: each alternative in $A \setminus A'$ is ranked first by exactly one agent (the other alternatives are ranked arbitrarily). This partial preference profile is sufficient to determine the probability that an alternative is elected, given that a unilateral rule over $N \setminus N'$ is selected by f. Since $|A'| \geq k$, we have that there exists $x^* \in A'$ such that

$$\Pr[f(\succ) = x^* \mid f \text{ unilateral over } N \setminus N'] \leq \frac{1}{k} \ .$$

We can now complete the preference profile $\succ$, by letting each of the agents in N' rank x^* first, and rank the other alternatives arbitrarily. Note that in our constructed preference profile the Plurality score of x^* is k, whereas the Plurality score of every other alternative is at most one.

To conclude the proof, we wish to upper-bound the probability that x^* is elected when f is applied to $\succ$. We first compute this probability given that f selects a unilateral rule. We have that

$$\begin{aligned}
&\Pr[f(\succ) = x^* | f \text{ unilateral}] \\
&= \Pr[f(\succ) = x^* | f \text{ unilateral over } N'] \\
&\qquad \cdot \Pr[f \text{ unilateral over } N' \mid f \text{ unilateral}] \\
&+ \Pr[f(\succ) = x^* | f \text{ unilateral over } N \setminus N'] \\
&\qquad \cdot \Pr[f \text{ unilateral over } N \setminus N' \mid f \text{ unilateral}] \\
&\leq 1 \cdot \frac{1}{k} + \frac{1}{k} \cdot 1 = \frac{2}{k} \ .
\end{aligned} \tag{9}$$

Hence, it holds that

$$\begin{aligned}
&\Pr[f(\succ) = x^*] \\
&= \Pr[f(\succ) = x^* \ | f \text{ duple}] \cdot p \\
&+ \Pr[f(\succ) = x^* \ | f \text{ unilateral}] \cdot (1-p) \\
&\leq \Pr[\exists x \in A' \text{ s.t. } x \in D \mid f \text{ is duple } D] \cdot p + \frac{2}{k} \cdot (1-p) \\
&\leq \frac{2}{k} \cdot p + \frac{2}{k} \cdot (1-p) = \frac{2}{k} \ .
\end{aligned}$$

It follows that the expected Plurality score of the winner with respect to $f(\succ)$ is at most

$$\frac{2}{k} \cdot k + \left(1 - \frac{2}{k}\right) \cdot 1 \leq 3 \ ,$$

We conclude that the approximation ratio of f is not greater than $3/k = 3/\sqrt{m}$. □

Proof sketch of Theorem 6. It is sufficient to prove the theorem for $\alpha = 1/2$, since for any larger value of α the scores of the alternatives are only larger, but the alternatives are nevertheless elected with probability proportional to their Copeland$_{1/2}$ score.

As in the proof of Theorem 2, we get that the approximation ratio is at least

$$\frac{1}{\binom{m}{2}} \cdot \frac{\text{OPT}^2 + \frac{\left(\binom{m}{2} - \text{OPT}\right)^2}{m-1}}{\text{OPT}} \ .$$

By differentiating with respect to OPT, we get that the function is monotone decreasing whenever

$$\text{OPT}^2 - \frac{m(m-1)^2}{4}$$

is nonpositive. The domain of the function is restricted to $(m-1)/2 \leq \text{OPT} \leq m-1$, therefore the function is monotone decreasing on the entire domain whenever $m \geq 4$. We conclude that the minimum is obtained when $\text{OPT} = m-1$. Substituting this value of OPT gives the announced approximation ratio.

Note that if $m < 4$ a simple calculation shows that Rule 2 yields an approximation ratio significantly better than 1/2. □

Proof sketch of Theorem 7. The proof is similar to the proof of Theorem 3. Assume that m is even, and let $n = m! + (m-1)$; since the number of agents is odd, the value of α is irrelevant. f selects a duple with probability p, and a unilateral rule with probability $1-p$. Given that a duple is selected, there is a set A', $|A'| = m/2$, such that an alternative in A' appears in a duple with probability at most $4/m$.

Given that a unilateral rule is selected, there is a set of agents N', $|N'| = m-1$, that is selected with probability at most

$$\frac{m-1}{m! + m - 1} \leq \frac{1}{m} \ .$$

For each permutation of A we have an agent in $N \setminus N'$ that ranks the alternatives according to the permutation. Therefore, based on the votes of $N \setminus N'$ all the alternatives are tied in their pairwise elections. There is an alternative $x^* \in A'$ that is elected with probability at most $2/m$ given that a unilateral rule over $N \setminus N'$ is selected.

We complete the profile by letting the agents in N' rank x^* first, and rank all the other alternatives cyclically. We get that the alternatives in $A \setminus \{x^*\}$ all have a score of $(m-2)/2$, whereas x^* has a score of $m-1$. The probability that x^* is elected is at most

$$p \cdot \frac{4}{m} + (1-p)\left(\left(1 - \frac{1}{m}\right) \cdot \frac{2}{m} + \frac{1}{m}\right) \leq \frac{4}{m} \ .$$

Hence the approximation ratio is at most

$$\frac{\frac{4}{m} \cdot (m-1) + \left(1 - \frac{4}{m}\right) \cdot \frac{m-2}{2}}{m-1} = \frac{1}{2} + \mathcal{O}\left(\frac{1}{m}\right) \ .$$

□

Proof sketch of Theorem 8. The proof is similar to the proof of Theorem 5, and again relies on Yao's Minimax principle. Let $n = m-1$, and assume for ease of exposition that m is even. We consider the distribution over preference profiles induced by the following procedure, which is identical to the one in the proof of Theorem 5 except that x^* is randomly placed in the first $m/2$ positions instead of the first $\sqrt{m}$.

1. Select an alternative $x^* \in A$ uniformly at random.
2. For each $i \in N$, select a position $k_i \in \{1, \ldots, m/2\}$ independently and uniformly at random. Agent i ranks x^* in position k_i.
3. Choose a permutation π_1 over $\{1,\ldots,\text{m-1}\}$ uniformly at random. For every $i = 1, \ldots, n-1$, let $\pi_{i+1}(j) = \pi_i(j+1)$, where the addition is cyclic.
4. Denote $A \setminus \{x^*\} = \{x_1, \ldots, x_{m-1}\}$. Intuitively, the agents rank the alternatives in $A \setminus \{x^*\}$ cyclically in the positions $1, \ldots, k_i - 1, k_i + 1, \ldots, m$. Formally, for $j = 1, \ldots, k_i - 1$, agent i ranks alternative $x_{\pi_i(j)}$ in position j. For $j = k_i + 1, \ldots, m$, i ranks alternative $x_{\pi_i(j-1)}$ in position j.

We compute the maximin scores of the different alternatives in a preference profile $\succ \in \mathcal{L}^n$ generated according to the above procedure. First, note that each alternative $x_j \in A \setminus \{x^*\}$ is ranked below position $m/2$ by exactly $m/2$ agents, hence $\text{sc}(x^*, \succ) \geq m/2$. On the other hand, for every $j \in \{1, \ldots, m-1\}$ there is $j' \in \{1, \ldots, m-1\}$ such that x_j is ranked below $x_{j'}$ by all agents save one, therefore $\text{sc}(x_j, \succ) = 1$.

Consider the case where $f : \mathcal{L}^n \to A$ is a unilateral rule with respect to agent i. By the same arguments as in the proof of Theorem 5 we have that

$$\Pr[f(\succ) = x^* \mid \succ_i] \leq 2/m \ ,$$

and therefore

$$\mathbb{E}\left[\left.\frac{\text{sc}(f(\succ), \succ)}{\text{sc}(x^*, \succ)}\right| \succ_i\right] \leq \frac{\frac{2}{m} \cdot m + \frac{m-2}{m} \cdot 1}{\frac{m}{2}} \leq \frac{6}{m} \ .$$

If $f : \mathcal{L}^n \to A$ is a duple rule, the probability that x^* is in the range of f is at most $2/m$, and we once again get that the expected approximation ratio is at most $6/m$. By Yao's Minimax principle, we conclude that $6/m$ is an upper bound on the approximation ratio that can be achieved via an SP randomized voting rule. □