Voting with Rank Dependent Scoring Rules

Judy Goldsmith, Jérôme Lang, Nicholas Mattei, and Patrice Perny

Abstract

Positional scoring rules in voting compute the score of an alternative by summing the scores for the alternative induced by every vote. This summation principle ensures that all votes contribute equally to the score of an alternative. We relax this assumption and, instead, aggregate scores by taking into account the rank of a score in the ordered list of scores obtained from the votes. This defines a new family of voting rules, *rank-dependent scoring rules (RDSRs)*, based on ordered weighted average (OWA) operators, which include all scoring rules, and many others, most of which of new. We study some properties of these rules, and show, empirically, that certain RDSRs are less manipulable than Borda voting, across a variety of statistical cultures and on real world data from skating competitions.

1 Introduction

Voting rules aim at aggregating the ordinal preferences of a set of individuals in order to produce a commonly chosen alternative. Many voting rules are defined in the following way: given a voting profile P, a collection of votes, where a vote is a linear ranking over alternatives, each vote contributes to the score of an alternative. The global score of the alternative is then computed by summing up all these contributed ("local") scores, and finally, the alternative(s) with the highest score win(s). The most common subclass of these scoring rules is that of *positional scoring rules*: the local score of x with respect to vote v depends only on the rank of x in v, and the global score of x is the sum, over all votes, of its local scores. Among prominent scoring rules we find the Borda rule as well as plurality, antiplurality and k-approval. However, there are occasionally undesirable features of scoring rules.

Example 1 *Four travelers have been asked to try six restaurants and to rank them for* `TripAdvisor.com`. *The resulting profile is $P = \langle acbdef, bcadef, dcaebf, ebadfc \rangle$, where $\langle acbdef \rangle$ means that the voter's preferred alternative is a, followed by c etc. The organizers of the competition feel that the highest and lowest ranks given to each alternative should count less than median scores. Therefore, they feel that c should win, followed by b, followed by a, then d, then by e, and finally by f. Neither Borda (which would elect a), nor k-approval for any k, gives this exact ranking.*

However, if we first compute the four local Borda scores of the six alternatives disregarding the two most extreme scores for each, then we get the desired ranking. More generally, we can weigh the scores according to their ranks in the ordered list of scores; for instance, the two extreme scores may have a weight $1/6$ each while the middle scores would have a weight $1/3$ each. This rank dependent weighting can be done in a natural way by coupling positional scoring rules together with ordered weighted average operators (OWAs) [34], to create *Rank Dependent Scoring Rules* (RDSRs). Each RDSR is characterized by the combination of a vector of both scores $\mathbf{s}, s_i \in \mathbf{s}, \geq 0$ and weights $\mathbf{w}, w_i \in \mathbf{w} \geq 0$.

RDSRs constitute an important class of aggregation procedures that are used quite commonly in everyday life. Artistic sports in the Summer and Winter Olympics, such as diving and skating, are judged by first removing the high and low scores and averaging the remaining scores achieved. Before recent changes, the London Interbank Offer Rate (LIBOR) inter-day bank rate, responsible for setting interest rates for most financial markets in the world, was computed (and manipulated) by soliciting 18 estimations of price, removing the high and low 4, and averaging the remaining

10 [1]. Biased aggregators such as RDSRs, a new area of study, are commonly used in internet recommendation settings such as Yelp! and TripAdvisor, and may affect rating behavior [16].

Order weighted averages have been studied in the context of cardinal utilities. In this paper we use OWAs to aggregate scores obtained by alternatives according to their ranks in the votes. This requires us to export these rank dependent functions from cardinal settings to ordinal settings. This allows us to apply rank dependent functions in setting where eliciting cardinal utilities is not feasible or expressible. Casting these functions as voting rules allows us to study these aggregation procedures with the same tools and techniques we use to study voting rules in social choice. Rank dependent functions have received much attention in multi-criteria decision making (see, *e.g.*, Yager et al. [35]) and decision under uncertainty (see, *e.g.*, Diecidue and Wakker [10]). OWAs have scarcely been used in social choice, with the recent exception of work by Skowron et al. [29], where they have been used in the context of proportional representation for multi-winner elections and group recommendation. Skowron et al. use OWAs to allow the flexibility to give a different importance to an agent's most preferred candidate/item among those that are selected, to her second most preferred candidate/item, etc.

Because $\mathbf{w}$ can give less weight to more extreme ranks given to an alternative,[1] we call these vectors *extreme-averse*. We expect that rules obtained for such extreme-averse vectors will typically be less frequently manipulable by small coalitions of voters than the corresponding rules obtained for a uniform $\mathbf{w}$. We confirm this expectation for a variety of statistical distributions and real world data from skating competitions in Section 6.

In the next section we formalize the notion of combining positional scoring rules with OWAs to create rank dependent scoring rules (RDSRs). We then provide background for and study axiomatic properties of this new class of voting rules. Next, we focus on a particular subclass of RDSRs, called the "Borda family", obtained by fixing the scoring vector $\mathbf{s}$ to Borda, and allowing the OWA vector to vary. Then we give experimental results that show that under several different distributions over profiles, some typical members of the Borda family are less frequently manipulable by a single voter than the Borda rule.

2 Formal Definitions

An election is a pair $E = (C, P)$ where C is the set of candidates or alternatives $\{c_1, \dots, c_m\}$, $|C| = m$, and P is a profile consisting of a set of voters indexed by their preference orders, $\{\succ_1, \dots, \succ_n\}$, $|P| = n$. Each voter is represented by a complete strict order (a vote) over the set of alternatives.

Many voting systems are *positional scoring rules* [30, 36], where there is a score vector $\mathbf{s} = \langle \alpha_1, \alpha_2, \dots, \alpha_m \rangle$, with $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_m$ and $\alpha_1 > \alpha_m$, that assigns points, for each vote, to the alternative placed at position α_i in that vote. The winner(s) are the alternative(s) maximizing the sum of points awarded over all the voters. Arguably the most famous positional scoring rules are Borda and plurality, with $\mathbf{s}_{\text{BORDA}} = \langle m-1, m-2, \dots, m-m \rangle$ and $\mathbf{s}_{\text{PLURALITY}} = \langle 1, 0, \dots, 0 \rangle$.

To combine positional scoring rules with OWAs [34], we introduce a weight vector $\mathbf{w} = \langle w_1, w_2, \dots w_n \rangle$ that is *normalized*, $(\sum_{i=1}^{n} w_i) = 1$. Through this construction we are able to maintain the property of anonymity in our voting rules while at the same time moderating the results based on the given weight vector over the ranks of alternatives. Elkind et al. [12] introduced a family of monotonic, non-homogenous rules, M-scoring rules, which are special cases of OWA operators where $M = m/2 + 1$ entries equal to weight 1 and all other entries are weight 0.

We formally define RDSRs in Definition 2 as irresolute social choice functions, that output a possibly empty set of (co)winners; as usual, irresolute rules can be made resolute by being combined with a tie-breaking priority mechanism.[2]

[1] Though, we may also give *more* weight to more extreme ranks given to an alternative, which is arguably much less desirable.

[2] As the composite scores also allow us to completely *rank* alternatives, RDSRs can also be defined as social welfare

Definition 2 *Given a scoring vector* $\mathbf{s} = \langle s_1, \ldots, s_m \rangle$ *and an OWA vector* $\mathbf{w} = \langle w_1, \ldots, w_n \rangle$*, where m is the number of alternatives and n the number of voters, we can define a voting rule* $F_{\mathbf{s},\mathbf{w}}(P)$[3] *associated with* $\mathbf{s}$ *and* $\mathbf{w}$*.*

Let $P = \langle \succ_1, \ldots, \succ_n \rangle$ *be a profile. For each voter* $\succ_i$ *and alternative* c_j*, let* $rank(c_j, \succ_i)$ *be the rank of* c_j *in vote* $\succ_i$*. Let* $\mathbf{r}(c_j, P) = \langle rank(c_j, \succ_1), \ldots, rank(c_j, \succ_n) \rangle$ *be the vector of ranks received by alternative* c_j *and* $\mathbf{r}^{\uparrow}(c_j, P)$ *be the sorting of* $\mathbf{r}(c_j, P)$ *in non-decreasing order such that the elements* $\mathbf{r}^{\uparrow}{}_1 \leq \mathbf{r}^{\uparrow}{}_2 \leq \ldots \leq \mathbf{r}^{\uparrow}{}_n$*.*

For alternative c_j *we create a vector of the scores associated with the ranks in all the votes to create the rank score vector,* $\mathbf{S}(c_j, P) = \langle s_{rank(c_j,\succ_1)}, \ldots, s_{rank(c_j,\succ_n)} \rangle$*. In order to apply the OWA operators we need to sort* $\mathbf{S}(c_j, P)$ *in* non-decreasing *order. Thus let* $\mathbf{S}^{\uparrow}(c_j, P)$ *be a reordering of* $\mathbf{S}(c_j, P)$ *where the elements* $\mathbf{S}^{\uparrow}{}_1 \leq \mathbf{S}^{\uparrow}{}_2 \leq \ldots \leq \mathbf{S}^{\uparrow}{}_n$*.*

We can now define the score for each alternative c_j *as:*

$$T_{\mathbf{s},\mathbf{w}}(c_j, P) = \mathbf{w} \cdot \mathbf{S}^{\uparrow}(c_j, P) = \sum_{i=1}^{n} w_i \times \mathbf{S}^{\uparrow}{}_i(c_j, P)$$

and $F_{\mathbf{s},\mathbf{w}}$ *selects the alternative(s) maximizing* $T_{\mathbf{s},\mathbf{w}}(x, P)$*.*

Thus, w_1 is associated with the *worst* score that an alternative receives, w_2 to the second worst score, etc. We use Pareto dominance to compare two score vectors: a vector $\mathbf{a} = \langle a_1, \ldots, a_n \rangle$ dominates another vector $\mathbf{b} = \langle b_1, \ldots, b_n \rangle$ if, for all i, $a_i \geq b_i$ and $a_i > b_i$ for some i.

Example 3 *As in Example 1, let* $m = 6$*,* $n = 4$ *and* $P = \langle acbdef, bcadef, dcaebf, ebadfc \rangle$*. Now, let* $\mathbf{s} = \mathbf{s}_{\text{BORDA}} = \langle 5, 4, 3, 2, 1, 0 \rangle$*, and* $\mathbf{w} = \langle 0, 1/4, 1/4, 1/2 \rangle$*.*

$\mathbf{w} =$	⟨ 0	1/4	1/4	1/2 ⟩	$T_{\mathbf{s},\mathbf{w}}(x)$
$\mathbf{S}^{\uparrow}(a) =$	⟨ 3	3	3	5 ⟩	4.0
$\mathbf{S}^{\uparrow}(b) =$	⟨ 1	3	4	5 ⟩	4.25
$\mathbf{S}^{\uparrow}(c) =$	⟨ 0	4	4	4 ⟩	4.0
$\mathbf{S}^{\uparrow}(d) =$	⟨ 2	2	2	5 ⟩	3.5
$\mathbf{S}^{\uparrow}(e) =$	⟨ 1	1	2	5 ⟩	3.25
$\mathbf{S}^{\uparrow}(f) =$	⟨ 0	0	0	1 ⟩	0.5

Therefore, the (unique) winner is b. If instead we choose $\mathbf{w}' = \mathbf{w}_{\text{OLYMPIC}} = \langle 0, 1/2, 1/2, 0 \rangle$*, then the scores are respectively* 3.0, 3.5, 4.0, 2.0, 1.5, 0.0 *and the winner is c (followed by b, a, d, e and f).*

So far, we have used weight vectors where we drop some extreme rankings. RDSRs are much more general than this. There are several interesting cases that occur based on settings to $\mathbf{w}$. We define two families of OWA vectors and then discuss a few specific cases of induced voting rules.

***k*-Uniform Interval ($\mathbf{w}_{k\text{-INTERVAL}}$):** Given parameter k, we drop k scores at the beginning and ending of the OWA operator: $\mathbf{w} = \langle 0_1, \ldots, 0_k, 1/n{-}2k, \ldots, 1/n{-}2k, 0_1, \ldots, 0_k \rangle$. This is a proper generalization of $\mathbf{w}_{\text{OLYMPIC}}$ and allows us to capture other rules that are used in practice such as the LIBOR interest rate setting aggregation rule [1]. As specific cases of k-uniform intervals we have: the *uniform vector* $\mathbf{w}_{\text{AVERAGE}} = \langle 1/n, \ldots, 1/n \rangle$, obtained for $k = 0$; the Olympic Average $\mathbf{w}_{\text{OLYMPIC}}\langle 0, 1/n{-}2, \ldots, 1/n{-}2, 0 \rangle$; and the *median* ($\mathbf{w}_{\text{MEDIAN}}$) $\mathbf{w}_{n+1/2} = 1$ when n is odd and $w_{(n/2)+1} = 1$ when n is even, with $w_i = 0$ for all other i. Using $\mathbf{w}_{\text{AVERAGE}}$ leads to recovering classical positional scoring rules.

***k*-Median ($\mathbf{w}_{k-\text{MEDIAN}}$):** Given $k \in \{1, \ldots, n\}$, let

$$\mathbf{w}_{k-\text{MEDIAN}} = \langle 0_1, 0_2, \ldots, 0_{k-1}, 1_k, 0_{k+1}, \ldots, 0_n \rangle.$$

functions, that produce a set of weak orders on the set of alternatives.

[3] We will often omit P when it is clear from context.

When $k = n$, then under the condition $s_1 > s_2$, we get the *nomination rule* where the co-winners are the alternatives that are ranked first (and thus have highest score) by at least one voter. More generally, if $s_1 = \ldots = s_i > s_{i+1}$, then the co-winners are the alternatives ranked among the top i alternatives by some voter. Note that $F_{\mathbf{s},\mathbf{w}_{\text{NOMINATION}}} = F_{\mathbf{t},\mathbf{w}_{\text{NOMINATION}}}$ for any two scoring vectors $\mathbf{s}$, $\mathbf{t}$ such that $s_1 > s_2$ and $t_1 > t_2$.

When $k = 1$, we recover a rule sometimes called "maximin" [7], that we prefer to call "maximin-score" (so that it is not confused with the Simpson-Kramer rule, which is also often called "maximin"), where all co-winners maximize the least score they receive, or equivalently, minimize the largest rank they receive. Note that this is independent of the setting of $\mathbf{s}$, that is, for any two strictly decreasing scoring vectors $\mathbf{s}$, $\mathbf{t}$, we have $F_{\mathbf{s},\mathbf{w}_{\text{MAXIMIN}}} = F_{\mathbf{t},\mathbf{w}_{\text{MAXIMIN}}}$.

When n is odd, for $k = \frac{n+1}{2}$ we obtain again the median rule, that for which the co-winners maximize their median rank; again, this is "almost" independent of the setting of $\mathbf{s}$ (and fully independent of the setting of $\mathbf{s}$ under the restriction that all scores of $\mathbf{s}$ are distinct).

When $k = \lfloor \frac{n}{2} \rfloor$ and $\mathbf{s} = \mathbf{s}_{\text{BORDA}}$, and more generally for any strictly decreasing scoring vector, we recover the Bucklin winner(s). Under the Bucklin rule, the alternative receiving a strict majority ($> \lfloor \frac{n}{2} \rfloor$) of votes in the first round is declared the winner. If no alternative has a strict majority, then the alternative(s) receiving $> \lfloor \frac{n}{2} \rfloor$ votes including the first place and second place votes is declared the winner(s). This procedure, adding the votes in the next position, continues until some set of alternatives is declared the winner. With $k = \lfloor \frac{n}{2} \rfloor$ and $\mathbf{s} = \mathbf{s}_{\text{BORDA}}$ note that any alternative receiving a majority of first place votes will be the only alternative with m points in the $k = \lfloor \frac{n}{2} \rfloor$ position of his sorted score vector ($\mathbf{S}^{\uparrow}$). If no alternative has this score, then the set of alternative(s) which have $m-1$ points in the $k = \lfloor \frac{n}{2} \rfloor$th position of $\mathbf{S}^{\uparrow}$ will be those alternatives who have a strict majority of votes placing them in first or second place, and on for the eventual Bucklin winner, in whatever "round" it may appear. Thus, With $k = \lfloor \frac{n}{2} \rfloor$ and $\mathbf{s} = \mathbf{s}_{\text{BORDA}}$ we recover the Bucklin voting rule.

The median rank rule is reminiscent of the *majority judgment rule* proposed by Balinski and Laraki [2]. However, there is a crucial difference: majority judgment is defined for a cardinal profile where each voter gives a score to each alternative. In RDSRs we map from ordinal profiles, as it is common in voting — this is important, especially when it comes to position our voting rules with respect to others.

$\mathscr{M}$-scoring rules: Taking

$$\mathbf{w}_{\mathscr{M}} = \langle 1_1, 1_2, \ldots, 1_{\lfloor n/2+1 \rfloor}, 0_{\lfloor n/2+2 \rfloor}, \ldots, 0_n \rangle$$

we obtain the family of $\mathscr{M}$-scoring rules defined in [12].

3 Properties of RDSRs

There are many properties surveyed in the social choice literature. A rule is said to have or obey a property if the property holds for all possible profiles. We focus on properties important to us, and refer the reader to texts in the literature for a more comprehensive survey, *e.g.*, [23].

Some basic fairness criteria that most sensible voting rules obey are: *anonymity*, insensitivity to permutations of the set of voters in a profile P; *neutrality*, insensitivity to permutations of the set of alternatives C; and *universal domain*, every alternative in C can be a winner.

Condorcet consistency states that, when one alternative is majority pair-wise preferred to all other alternatives, that alternative is the unique winner. *Monotonicity* states that, given a profile P and winning alternative x, if we modify any set of votes in P to produce P' where the only change is promoting x, then x is still the winner of the election run on the profile P'.

Other properties concern the behavior of voting rules when splitting, combining, and expanding the given profiles. *Reinforcement* states, given two disjoint profiles P_1 and P_2, if $F(P_1) \cap F(P_2) \neq \emptyset$ then $F(P_1 \cup P_2) = F(P_1) \cap F(P_2)$. *Homogeneity* states, given a profile P, multiplying all voters in the profile any number of times should not change the result.

Reinforcement and homogeneity concern *variable electorates* and are not immediately applicable to RDSRs, which are defined for a fixed value of n. However, they apply to families of rules $\{w^{(n)}, n \geq 1\}$ of vectors (one for each possible number of votes), exactly like properties that concern variable sets of alternatives need scoring rules (typically defined for a fixed m) to be defined as families of rules for a varying m.

All RDSRs satisfy anonymity and neutrality. We show that monotonicity (satisfied by all scoring rules) extends to rank-dependent scoring rules.

Proposition 4 *For every* $\mathbf{w}$ *and* $\mathbf{s}$*,* $F_{\mathbf{s},\mathbf{w}}$ *is monotonic.*

Proof. Let P be a profile and $x \in F_{\mathbf{s},\mathbf{w}}(P)$. Let P' be obtained by raising x from rank i to rank $i-1$ in one of the votes, leaving everything else unchanged. Let j be the number of votes in P who rank x in the first $i-1$ positions. Then $\mathbf{S}^{\uparrow}(x,P) = \langle \mathbf{S}^{\uparrow}{}_1, \ldots, \mathbf{S}^{\uparrow}{}_{n-j}, \ldots, \mathbf{S}^{\uparrow}{}_n \rangle$, with $\mathbf{S}^{\uparrow}{}_{n-j} = s_i$, and $\mathbf{S}^{\uparrow}(x,P') = \langle \mathbf{S}^{\uparrow}{}'_1, \ldots, \mathbf{S}^{\uparrow}{}'_{n-j}, \ldots, \mathbf{S}^{\uparrow}{}'_n \rangle$ with $\mathbf{S}^{\uparrow}{}'_k = \mathbf{S}^{\uparrow}{}_k$ for all $k \neq n-j$ and $\mathbf{S}^{\uparrow}{}'_{n-j} = s_{i-1}$. Because $s_{i-1} \geq s_i$, $\mathbf{S}^{\uparrow}(x,P')$ weakly Pareto-dominates $\mathbf{S}^{\uparrow}(x,P')$, therefore $T_{\mathbf{s},\mathbf{w}}(x,P') \geq T_{\mathbf{s},\mathbf{w}}(x,P)$. Similarly, $T_{\mathbf{s},\mathbf{w}}(x',P') \leq T_{\mathbf{s},\mathbf{w}}(x',P)$ for any $x' \neq x$; therefore, the score of x remains maximal when moving from P to P', and $x \in F_{\mathbf{s},\mathbf{w}}(P)$. ❑

The following example shows that RDSRs do not necessary fulfill reinforcement nor homogeneity, even for natural collection of scoring vectors and OWA vectors; a similar result was shown for $\mathscr{M}$-scoring rules in [12].

Example 5 *Set* $\mathbf{w}_{\text{OLYMPIC}}$ *and* $\mathbf{s}_{\text{BORDA}}$*. Let* $C = \{a,b,c,d,e\}$ *and* $P = \langle abcde, bcade, deacb \rangle$*. This gives us* $\mathbf{S}^{\uparrow}(a,P) = \langle 2,2,4 \rangle$*,* $\mathbf{S}^{\uparrow}(b,P) = \langle 0,3,4 \rangle$*,* $\mathbf{S}^{\uparrow}(c,P) = \langle 1,2,3 \rangle$*,* $\mathbf{S}^{\uparrow}(d,P) = \langle 1,1,4 \rangle$*, and* $\mathbf{S}^{\uparrow}(e,P) = \langle 0,0,3 \rangle$*, thus,* $T_{\mathbf{s},\mathbf{w}}(a,P) = 2$*,* $T_{\mathbf{s},\mathbf{w}}(b,P) = 3$*,* $T_{\mathbf{s},\mathbf{w}}(c,P) = 2$*,* $T_{\mathbf{s},\mathbf{w}}(d,P) = 1$*,* $T_{\mathbf{s},\mathbf{w}}(e,P) = 0$*, and the winner is* b*.*

Now, let $3 \times P$ *be the 9-voter profile obtained by replacing each vote in* P *by three identical votes. We now have* $\mathbf{S}^{\uparrow}(a,P) = \langle 2,2,2,2,2,2,4,4,4 \rangle$*,* $\mathbf{S}^{\uparrow}(b,P) = \langle 0,0,0,3,3,3,4,4,4 \rangle$*, etc. Thus,* $T_{\mathbf{s},\mathbf{w}}(a,P) = 18/7$*,* $T_{\mathbf{s},\mathbf{w}}(b,P) = 17/7$*,* $T_{\mathbf{s},\mathbf{w}}(c,P) = 16/7$*,* $T_{\mathbf{s},\mathbf{w}}(d,P) = 11/7$*,* $T_{\mathbf{s},\mathbf{w}}(e,P) = 6/7$*, and the winner is* a*.*

Example 5 shows that some natural RDSRs are not homogeneous, and, *a fortiori*, do not satisfy reinforcement. This implies that the class of RDSR contains elements that are not generalized scoring rules [33].

Proposition 6 *For every* $m \geq 3$ *and* $n \geq 5$*, no rule* $F_{\mathbf{s},\mathbf{w}}$ *is Condorcet-consistent.*

Proof. Assume $n \geq 5$ and $n \neq 8$. Let $X = \{x_1, \ldots, x_m\}$. Let $k = \lfloor \frac{n}{3} \rfloor$ and $q = n - 3k$ (note that $q \leq 2$); let P be the following profile: we have k votes $x_1 \succ x_2 \succ \ldots \succ x_m$, k votes $x_m \succ x_1 \succ \ldots \succ x_{m-1}$ and $n - 2k = k + q$ votes $x_2 \succ \ldots \succ x_{m-2} \succ x_1 \succ x_m$. Because $n \geq 5$ and $n \neq 8$, we have $2k > \frac{n}{2}$ and, *a fortiori*, $2k + q > \frac{n}{2}$, therefore x_1 is a Condorcet winner. Now, the nondecreasingly reordered score vector for x_1 is $\langle n - 2k \times s_{m-1}, k \times s_2, k \times s_1 \rangle$ and that of x_2 is $\langle k \times s_3, k \times s_3, n - 2k \times s_1 \rangle$, therefore the scores of x_1 and x_2 are

$$\begin{aligned} T_{\mathbf{s},\mathbf{w}}(x_1) &= \textstyle\sum_{i=1}^{k+q} w_i s_{m-1} + \sum_{i=k+q+1}^{2k+q} w_i s_2 + \sum_{i=2k+q+1}^{n} w_i s_1 \\ T_{\mathbf{s},\mathbf{w}}(x_2) &= \textstyle\sum_{i=1}^{k} w_i s_3 + \sum_{i=k+1}^{2k} w_i s_2 + \sum_{i=2k+1}^{n} w_i s_1. \end{aligned}$$

$$\begin{aligned} \text{and} \quad & T_{\mathbf{s},\mathbf{w}}(x_1) - T_{\mathbf{s},\mathbf{w}}(x_2) \\ = \quad & \textstyle\sum_{i=1}^{k} w_i (s_{m-1} - s_3) + \sum_{i=k+1}^{k+q} w_i (s_{m-1} - s_2) \\ & \textstyle + \sum_{i=k+q+1}^{2k} w_i (s_2 - s_2) + \sum_{i=2k+1}^{2k+q} w_i (s_2 - s_1) \\ & \textstyle + \sum_{i=2k+q+1}^{n} w_i (s_1 - s_1). \end{aligned}$$

None of the five terms can be strictly positive, therefore $T_{\mathbf{s},\mathbf{w}}(x_1) - T_{\mathbf{s},\mathbf{w}}(x_2) \leq 0$, which entails $F_{\mathbf{s},\mathbf{w}}(P) \neq \{x_1\}$, which shows that whatever the value of $\mathbf{w}$ and $\mathbf{s}$, $F_{\mathbf{s},\mathbf{w}}(P)$ is not Condorcet-consistent. The proof for $n = 8$ is similar, but taking 2 votes $x_1 \succ x_2 \succ \ldots \succ x_m$, 3 votes $x_m \succ x_1 \succ \ldots \succ x_{m-1}$, and 3 votes $x_2 \succ \ldots \succ x_{m-2} \succ x_1 \succ x_m$. ❑

This result generalizes the known result from Fishburn [15] and Moulin [23] that no scoring rule is Condorcet-consistent.

4 The Borda Family

This section focuses on the subclass of RDSRs obtained by fixing the scoring vector to match the Borda scoring vector $\mathbf{s}_{\text{BORDA}} = \langle m-1, m-2, \ldots, m-m\rangle$. Maximizing an OWA applied to scores is equivalent to minimizing an OWA applied to ranks, hence this family (all RDSRs realizable using a Borda scoring rule) is particularly meaningful (besides the importance of the Borda rule in voting). A first question is, are there any positional scoring rules, apart from Borda, which belong to the Borda family? The answer is, somewhat surprisingly, positive, when n and m are both fixed.

Proposition 7 *Let n and m be fixed, and define:*

$$\boldsymbol{w}_{\text{LEXIMIN}} = \langle {}^{m^{n-1}}\!/_{W}, {}^{m^{n-2}}\!/_{W}, \ldots, {}^{m}\!/_{W}, {}^{1}\!/_{W}\rangle$$

and

$$\boldsymbol{w}_{\text{LEXIMAX}} = \langle {}^{1}\!/_{W}, {}^{m}\!/_{W}, \ldots, {}^{m^{n-2}}\!/_{W}, {}^{m^{n-1}}\!/_{W}\rangle,$$

where $W = 1 + m + \ldots + m^{n-1}$. Then $F_{\boldsymbol{s}_{\text{BORDA}},\boldsymbol{w}_{\text{LEXIMIN}}}$ and $F_{\boldsymbol{s}_{\text{BORDA}},\boldsymbol{w}_{\text{LEXIMAX}}}$ are classical scoring rules, associated with the scoring vectors:

$$\boldsymbol{s}_{\text{LEXPL}} = \langle n^{m-1}, n^{m-2}, \ldots, n^2, n, 1\rangle$$

and

$$\boldsymbol{s}_{\text{LEXAPL}} = \langle n^{m-1}, n^{m-1} - n, \ldots, n^{m-1} - n^{m-2}, 0\rangle.$$

Proof. Let us start with $F_{\mathbf{s}_{\text{BORDA}},\mathbf{w}_{\text{LEXIMIN}}}$. For any profile P and integer k, let $A_k(x,P)$ be the number of votes in P in which x is ranked in position k, and $B_k(x,P) = \sum_{j\leq k} A_k(x,P)$ be the number of votes in P in which x is ranked in position at most k. Recall: $\mathbf{r}_i(x)$ is the ith best rank given to x, and $m - \mathbf{r}_i(x)$ the ith best Borda score given to x. Note that we have $\mathbf{r}_i(x) = k$ if and only if (1) $B_{k-1}(x,P) < i$ and (2) $B_k(x,P) \geq i$.

We claim that (1) for any x, y, we have $T_{\mathbf{s}_{\text{BORDA}},\mathbf{w}_{\text{LEXIMIN}}}(x) > T_{\mathbf{s}_{\text{BORDA}},\mathbf{w}_{\text{LEXIMIN}}}(y)$ if and only if there is a $k \leq m-1$ such that (a) for all $i < k$, $A_i(x,P) = A_i(y,P)$ and (b) $A_k(x,P) > A_k(y,P)$.

Assume (a) and (b) hold for some k. Let $i^* = B_k(y,P) + 1$. Then we have $\mathbf{r}_{i^*}(x) = k$ and $\mathbf{r}_{i^*}(y) \geq k+1$, and for all $i \leq i^*$, $\mathbf{r}_i(x) = \mathbf{r}_i(y)$. Now,

$$\begin{array}{ll} & T_{\mathbf{s}_{\text{BORDA}},\mathbf{w}_{\text{LEXIMIN}}}(x) - T_{\mathbf{s}_{\text{BORDA}},\mathbf{w}_{\text{LEXIMIN}}}(y) \\ = & \frac{1}{W}\sum_{i=1}^{n} m^{n-i}(m - \mathbf{r}_i(x)) - (m - \mathbf{r}_i(y)) \\ = & \frac{1}{W}\sum_{i=1}^{n} m^{n-i}(\mathbf{r}_i(y) - \mathbf{r}_i(x)) \\ = & \frac{1}{W}\left(m^{n-i^*}(\mathbf{r}_{i^*}(y) - \mathbf{r}_{i^*}(x)) + \sum_{i>i^*} m^{n-i}(\mathbf{r}_i(y) - \mathbf{r}_i(x))\right) \\ \geq & \frac{1}{W}\left(m^{n-i^*} - \sum_{i>i^*} m^{n-i}(m-1)\right) \\ > & 0. \end{array}$$

Conversely, if (a) and (b) do not hold then for all k, we have $B_k(x,P) \leq B_k(y,P)$, therefore, for all i, $\mathbf{r}_i(x) \geq \mathbf{r}_i(y)$, which implies $T_{\mathbf{s}_{\text{BORDA}},\mathbf{w}_{\text{LEXIMIN}}}(x) \leq T_{\mathbf{s}_{\text{BORDA}},\mathbf{w}_{\text{LEXIMIN}}}(y)$.

Now, we claim that (2) the total score according to the scoring rule associated with $\mathbf{s}_{\text{LEXPL}}$, $T_{\mathbf{s}_{\text{LEXPL}}}(x) > T_{\mathbf{s}_{\text{LEXPL}}}(y)$ if and only if there is a $k \leq m-1$ such that (a) for all $i < k$, $A_i(x,P) = A_i(y,P)$ and (b) $A_{k,P}(x) > A_{k,P}(y)$.

Assume (a) and (b). We have $T_{\mathbf{s}_{\text{LEXPL}}}(x) = \sum_{i=1}^{m} A_i(x,P) \cdot n^{m-i}$. Note that, for any i, $|A_i(x,P) - A_i(y,P)| \leq n$.

$$\begin{array}{ll} \text{Then} & T_{\mathbf{s}_{\text{LEXPL}}}(x,P) - T_{\mathbf{s}_{\text{LEXPL}}}(y) \\ = & \sum_{i=1}^{m} A_i(x,P) \cdot n^{m-i} - \sum_{i=1}^{m} A_i(y,P) \cdot n^{m-i} \\ = & (A_k(x,P) - A_k(y,P))n^{m-k} + \sum_{i=k+1}^{m} (A_i(x,P) - A_i(y,P))n^{m-i} \\ \geq & n^{m-k} + (n) \cdot n^{m-k+1} \\ > & 0. \end{array}$$

Conversely, if (a) and (b) do not hold then for all $k \leq m-1$, we have $A_k(x,P) \leq A_k(y,P)$; this means that either there is a $k \leq m-1$ such that (a) for all $i \leq k$, $A_i(x,P) = A_i(y,P)$ and (b) $A_k(y,P) > A_k(x,P)$, in which case $T_{\mathbf{s}_{\text{LEXPL}}}(y) - T_{\mathbf{s}_{\text{LEXPL}}}(x,P) \geq 0$, or that for all k, $A_k(x,P) = A_k(y,P)$, in which case $T_{\mathbf{s}_{\text{LEXPL}}}(y) - T_{\mathbf{s}_{\text{LEXPL}}}(x) \geq 0$ as well.

(1) and (2) together imply that $F_{\mathbf{s}_{\text{BORDA}},\mathbf{w}_{\text{LEXIMIN}}}$ is the scoring rule associated with scoring vector $\mathbf{s}_{\text{LEXPL}}$. The proof that $F_{\mathbf{s}_{\text{BORDA}},\mathbf{w}_{\text{LEXIMAX}}}$ is the scoring rule associated with scoring vector $\mathbf{s}_{\text{LEXAPL}}$ is similar. ❑

Example 8 *Let $m = 4$, $n = 6$, and the profile P be composed of two votes $x \succ t \succ z \succ y$, 2 votes $y \succ t \succ x \succ z$, one vote $z \succ y \succ x \succ t$ and one vote $t \succ z \succ x \succ y$. The vectors of ranks, reordered non-decreasingly, are $\boldsymbol{r}^{\uparrow}(x) = \langle 1,1,3,3,3,3 \rangle$; $\boldsymbol{r}^{\uparrow}(y) = \langle 1,1,2,4,4,4 \rangle$; $\boldsymbol{r}^{\uparrow}(z) = \langle 1,2,2,2,4,4 \rangle$; $\boldsymbol{r}^{\uparrow}(t) = \langle 1,2,2,3,3,4 \rangle$. We have $A_1(y,P) = A_1(x,P)$ and $A_2(y,P) > A_2(x,P)$, therefore $T_{\boldsymbol{s}_{\text{LEXPL}}}(y) > T_{\boldsymbol{s}_{\text{LEXPL}}}(x)$; and we have $A_1(y,P) > A_1(z,P)$ and $A_1(y,P) > A_1(t,P)$, therefore $T_{\boldsymbol{s}_{\text{LEXPL}}}(y) > T_{\boldsymbol{s}_{\text{LEXPL}}}(z)$ and $T_{\boldsymbol{s}_{\text{LEXPL}}}(y) > T_{\boldsymbol{s}_{\text{LEXPL}}}(t)$: the winner for $\boldsymbol{s}_{\text{LEXPL}}$ is y. We can also check that the winner for $\boldsymbol{s}_{\text{LEXAPL}}$ is x.*

Note that if n is not fixed, then $F_{\mathbf{s}_{\text{BORDA}},\mathbf{w}_{\text{LEXIMIN}}}$ and $F_{\mathbf{s}_{\text{BORDA}},\mathbf{w}_{\text{LEXIMAX}}}$ are not scoring rules in the usual sense, because all weights but one would have to be infinitesimals.

Therefore, when n and m are fixed, at least three rules are in the intersection of the family of scoring rules and the Borda family (Borda, lexicopraphic plurality and lexicographic antiplurality), whereas when only m is fixed, only Borda is known to be both in the family of scoring rules and in the Borda family. We conjecture that the intersection (on both cases, n fixed and n not fixed) do not contain any other rules than these, but did not come up with a proof.

5 RDSRs and Fairness

The use of the OWA operator in RDSRs allows an election designer to favor a fair distribution of satisfaction among voters, whenever this is desirable. The score $T_{\mathbf{s},\mathbf{w}}(c_j,P)$ can act as an inequality measure (see, *e.g.*, [32]) taking into account the distribution of scores $s_{rank(c_j,\succ_k)}, k = 1,\ldots,n$ whenever weights satisfy $w_1 > w_2 > \ldots > w_n > 0$.

The intuition behind choosing strictly decreasing weights can be given as follows: one puts more weight on the least satisfied voter (smallest score), then on the second least satisfied voter and so on. This is a natural extension of the *min* and *leximin* operators. These operators allow for more compensation between scores assigned to alternatives by the voters. With a proper choice of weights, there remains some possibility for the election designer to compensate the dissatisfaction of one agent by the satisfaction of some others, while still preserving a somewhat egalitarian notion of fairness by favoring alternatives that have a well-balanced scoring profile. Specifically, we want to favor alternatives whose vectors of scores do not contain too many extreme scores.

This can be stated more formally using transfers that reduce societal inequality, also known as Pigou-Dalton tranfers [24], by the following proposition.

Proposition 9 *Let $P = (\succ_1,\ldots,\succ_n)$ be a preference profile and c an alternative such that $rank(c,\succ_k) < rank(c,\succ_i)$ for some pair of voters (i,k). Then for any alternative c' such that vector $r(c',P)$*

and $r(c,P)$ satisfies:

$$\begin{array}{rcl} s_{rank(c',\succ_k)} &=& s_{rank(c,\succ_k)} - \varepsilon \\ s_{rank(c',\succ_i)} &=& s_{rank(c,\succ_i)} + \varepsilon \\ s_{rank(c,\succ_j)} &=& s_{rank(c',\succ_j)},\ \forall j \in N \setminus \{i,k\} \end{array}$$

for some $\varepsilon \in (0, s_k - s_i)$, then $T_{\mathbf{s},\mathbf{w}}(c',P) > T_{\mathbf{s},\mathbf{w}}(c,P)$ whenever $\mathbf{w}$ is strictly decreasing.

Proof. Let L and L' be the two vectors of $\mathbb{R}^n$ defined by $L_j = \sum_{k=1}^{j} \mathbf{S}^{\uparrow}{}_k(c,P)$ and $L'_j = \sum_{k=1}^{j} \mathbf{S}^{\uparrow}{}_k(c',P)$ for all $j \in N$. Since we pass from $\mathbf{S}^{\uparrow}(c,P)$ to $\mathbf{S}^{\uparrow}(c',P)$ using a Pigou-Dalton transfer of size ε from component $s_{rank(c,\succ_k)}$ to component $s_{rank(c,\succ_i)}$ then we know that L' Pareto-dominates L [19, 28].

Now, let $\mathbf{w}'$ be the vector derived from $\mathbf{w}$ by setting: $\mathbf{w}'_n = \mathbf{w}_n$ and $\mathbf{w}'_j = \mathbf{w}_j - \mathbf{w}_{j+1}$ for $j = \{1,\ldots,n-1\}$, we observe that $T_{\mathbf{s},\mathbf{w}}(c,P) = \mathbf{w}' \cdot L$ and $T_{\mathbf{s},\mathbf{w}}(c',P) = \mathbf{w}' \cdot L'$. Then, due to the strictly decreasing property on $\mathbf{w}$, we know that $\mathbf{w}'_j > 0$ for all $j \in N$. Hence $\mathbf{w}'_j \cdot L'_j \geq \mathbf{w}'_j \cdot L_j$ for all $j \in N$, one of these inequalities being strict by Pareto dominance. Hence $\mathbf{w}' \cdot L' > \mathbf{w}' \cdot L$ and therefore $T_{\mathbf{s},\mathbf{w}}(c',P) > T_{\mathbf{s},\mathbf{w}}(c,P)$. ❑

Hence, when using strictly decreasing weights, an alternative c maximizing an OWA score $T_{\mathbf{s},\mathbf{w}}(c,P)$ over the set of alternatives has a scoring vector $\mathbf{S}^{\uparrow}(c,P)$ that cannot be improved in terms of Pigou-Dalton transfer by another vector $\mathbf{S}^{\uparrow}(c',P)$. This is a way of rewarding fairness in score aggregation as illustrated in the following Example.

Example 10 *Let $m = 4$, $n = 3$, $P = \langle acbd, cbad, dbac \rangle$, $\mathbf{s} = \mathbf{s}_{\text{BORDA}}$, and $\mathbf{w} = \langle 1/2, 1/3, 1/6 \rangle$.*

$\mathbf{w} =$	$\langle 1/2$	$1/3$	$1/6 \rangle$	$T_{\mathbf{s},\mathbf{w}}(x)$
$\mathbf{S}^{\uparrow}(a) =$	$\langle$ 1	1	3 $\rangle$	8/6
$\mathbf{S}^{\uparrow}(b) =$	$\langle$ 1	2	2 $\rangle$	9/6
$\mathbf{S}^{\uparrow}(c) =$	$\langle$ 0	2	3 $\rangle$	7/6
$\mathbf{S}^{\uparrow}(d) =$	$\langle$ 0	0	3 $\rangle$	3/6

Here b is the winner whereas a,b,c would remain indifferent under the Borda rule, while the maximin-score rule (cf. Footnote 4) would also be indifferent between a and b. Note that the Leximin refinement of the maximin-score rule would yield the same ranking $b \succ a \succ c \succ d$ as $T_{s,w}$, but this is not always the case. Consider the scoring vectors $\mathbf{S}^{\uparrow}(x) = \langle 0,3,3 \rangle$ and $\mathbf{S}^{\uparrow}(y) = \langle 1,1,1 \rangle$. We get $T_{\mathbf{s},\mathbf{w}}(x) = 3/2$ whereas $T_{\mathbf{s},\mathbf{w}}(y) = 1$. In such drastic cases where fairness is strongly conflicting with overall efficiency (measured by the sum of scores), RDSRs allow the election designer the possibility of sacrificing a minority of opinions so as to preserve a high average score, thus departing from the Leximin refinement of the maximin-score rule.

6 Manipulation: Empirical Experiments

We conjecture that RDSRs that drop the extreme ranks may be, on average, less manipulable than standard scoring rules. Since all voting rules are manipulable we can only hope that by dropping some of the extreme ranks we have defined a class of voting rules that is manipulable less often in expectation. Since RDSRs are used in practice in situations with small numbers of voters, such as Olympic artistic scoring and interest rates, we investigate settings that contain one manipulator and just a handful of voters.

Commonly, manipulation in computational social choice is defined so its computational complexity can be studied cleanly as a decision problem [4]: "given a profile, and an alternative x, can the voter change her vote so that x is the winner?" Here we use the definition of manipulation from classical social choice [3]: "given a profile, can the manipulator change her vote so that the

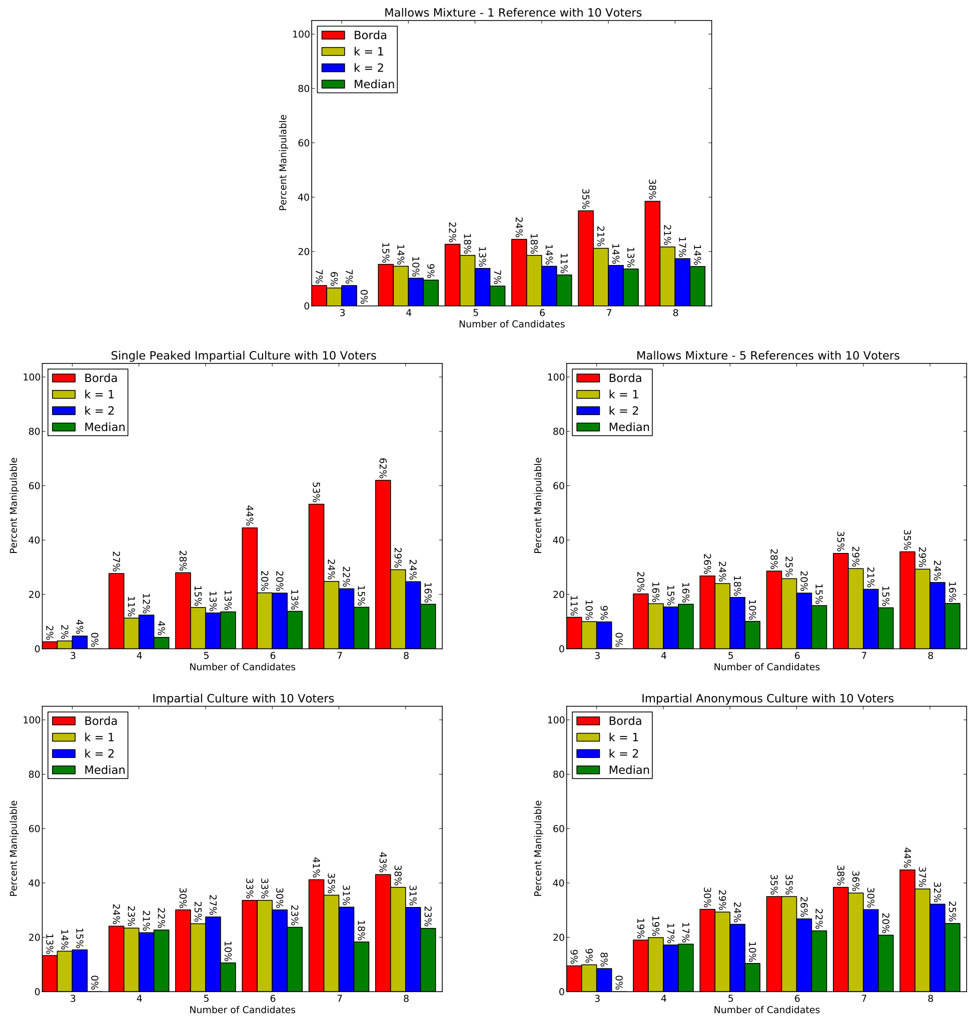

Figure 1: Graphs showing the frequency of manipulation for OWAs using the $\mathbf{w}_{k\text{-INTERVAL}}$ weight vector versus normal Borda scoring for instances with 10 voters. Generally, as we increase k (the number of dropped ranks) towards the median we have less opportunity for manipulation. This relationship becomes particularly strong as we increase the correlation among the votes.

outcome is better than with her original, sincere, vote?" This requires lifting preferences over alternatives to sets of alternatives. We use the definition from Duggan and Schwartz [11] known as the optimistic manipulator assumption (also known as the non-unique winner model in computational social choice). Formally, a manipulation by voter i exists if if there is a vote $\succ_i'$ and alternative x such that $x \in F_{\mathbf{s},\mathbf{w}}(\{P \setminus \succ_i\} \cup \succ_i')$ and $rank(x, \succ_i) > rank(x, \succ_i)$ for all $j \in F_{\mathbf{s},\mathbf{w}}(P)$.

Worst-case results about the hardness of manipulation abound in social choice [4, 8, 14] but these results may not reflect the cost in practice to compute manipulations [21, 26, 31]. The complexity of manipulation for the Borda rule, which we study here, is polynomial for a single manipulator but NP-complete for two or more manipulators [4, 9] Many such analyses assume that all preferences are equally likely, but that is not supported by studies in behavioral social choice [25, 27] or studies on real data [20, 27]. In order to understand how the manipulability RDSRs changes with respect to the underlying distribution of votes we use five generative statistical cultures to create profiles for our testing as well as testing on real world data.

6.1 Synthetic Data

We study manipulation under several assumptions about the distribution of preferences over the m alternatives. The **Impartial Culture (IC)** assumes the probability of observing any of the $m!$ preference orders is equally likely for each voter; namely $p = \frac{1}{m!}$. This culture is a kind of worst case assumption, we do not know anything about the feelings of the underlying voters so we assume there is no bias in the generation process. The **Impartial Anonymous Culture (IAC)** is a strict generalization of IC which assumes the probability of observing any probability distribution over the $m!$ possible orders is equally likely. That is, each vector of length $m!$ that has sum 1, describing the distribution over the $m!$ possible votes, is equally likely to occur.

The **Mallows Mixture Models (MM)** makes the underlying assumption that there is a true ranking and that individuals deviate from the ground truth with decreasing probability as the ranking moves away from the reference. Given reference rankings $\sigma_1, \dots \sigma_n$, probabilities $\phi_1, \dots, \phi_n$, and mixture model (discrete probability distribution) $\pi_1, \dots, \pi_n$, we generate rankings that have a Kendall Tau distance $\tau = (\sigma, \sigma')$ from the reference ranking that is proportional to ϕ_i^τ. We select among the n models by selecting one randomly from the given probability distribution [17, 18]. We use two flavors of Mallows Models in our experiments: a pure Mallows model with one reference order and a Mallows Mixture with five.

Finally we examine a distribution which creates a correlation between the shape of the individual preference profiles. The **Single Peaked Impartial Culture (SPIC)** assumes that each single peaked preference profile compatible with m alternatives is equally likely. Single-peakedness is is an important domain restriction introduced by Black [5] and is widely studied in the computational social choice community for its strategic [13] and empirical properties [20]. Intuitively, single-peakedness is the idea that all voters have a point along a shared axis where they would be happiest, and rank alternatives farther from this point worse.

In Figures 1 we compare the manipulability of the Borda scoring vector with OWAs using variants of the $\mathbf{w}_{k\text{-INTERVAL}}$ weight vectors. For each of the statistical cultures mentioned, we generate 1000 random instances and test, via brute force search, whether a single agent that is randomly drawn from the set of voters can successfully manipulate the instance. Any election where the outcome is the same as the manipulator's honest preference was discarded and a new instance generated. Thus, in all 1000 elections, the results are never the same as the manipulators true preference.

We see that, as we induce more correlation between the voters, we decrease the opportunities for manipulation. Thus, in the limit for a fixed dispersion parameter, a Mallows Mixture with a large number of reference orders tend more towards the IC model (all orders are increasingly, equally likely), while a Mallows model with a single reference will have a more tightly correlated set of votes, tending towards profiles that exhibit the Single-Peaked domain restriction.

Even with the decreased opportunities for manipulation in these correlated models, RDSRs do better when we drop a small percentage of the extreme ranks. This is probably because, in these small settings, one extreme voter can move a particular alternative up or down based on an extreme rank. If a particular alternative is receiving 1's and 2's on average and we give him a 9, then this score is very out of line with the feelings of the group. However, using $\mathbf{w}_{k\text{-INTERVAL}}$ vectors we can downplay these extreme scores and move more towards the median view of all the voters. Similar results were shown by Cervone et al. [6] in their work on voting rules that use the mediancenter to aggregate preferences.

We ran the same experiment for settings with 20 and 30 voters. As with most results on manipulation, as the pool of voters grows larger, the opportunities for manipulation decrease. In the uncorrelated models there is still a (relatively) large chance for manipulation; when we go to the correlated models we eliminate these opportunities. This may be why variants of $\mathbf{w}_{k\text{-INTERVAL}}$ are used for artistic sports in the Olympics and other places where there is some general consensus

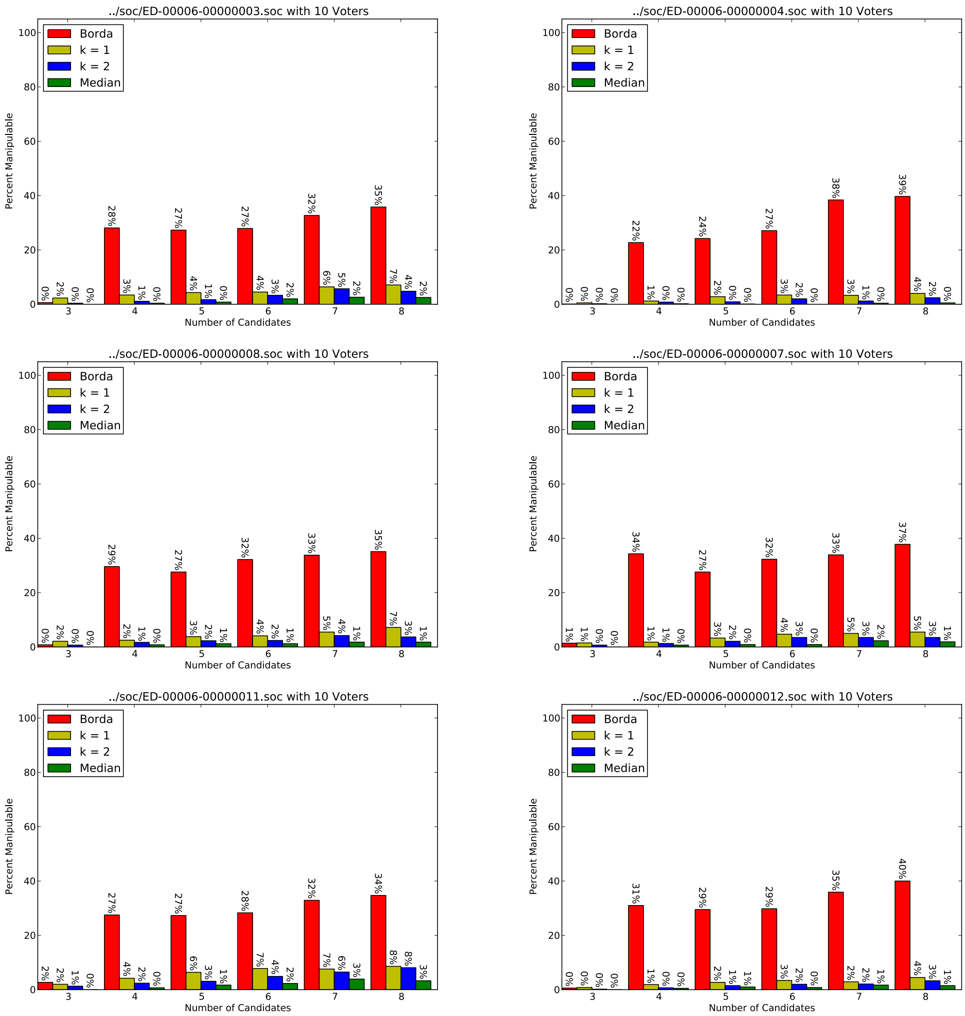

Figure 2: Graphs showing the frequency of manipulation for OWAs using the $\mathbf{w}_{k\text{-INTERVAL}}$ weight vector versus normal Borda scoring for instances with 10 voters redrawn from various datasets in ED-00006 [22]. The results on the real-world data are even more striking than the synthetic data. Only dropping the top and bottom ranks (as is often the case in Olympic artistic sports) virtually eliminates the opportunities for manipulation by a single voter

about technical ability with small perturbations in the final orderings of the individual voters. In these settings, as we can see from our experiments, mixing scoring rules with OWA vectors can help to eliminate incentives for individuals to misreport their preferences.

6.2 Olympic Skating Data

In order to test if our results about manipulation hold in the real world we decided to test against real world data. As the basis for this experiment we use the *Skate* dataset (ED-00006) from PrefLib [22]. This dataset contains figure skating rankings from various competitions during the 1998 season including the World Juniors, World Championships, and the Olympics. These data sets have 10–25 alternatives (skaters) and 8–10 judges (voters).

In order to make the results computable and comparable with our results on the synthetic data we perform resampling on these data sets. Using techniques from bootstrap sampling in behavorial social choice [25, 27] we redraw, with replacement, 10 votes from the distribution of votes given by the empirical data. We then project these votes down to the required number of alternatives by randomly removing a set of alternatives from all the votes. We repeat this with independent draws for 1000 samples per setting to k, m, and n. This process gives us information about the replicability of observing an event — the more times we observe an event happening due to small perturbations in the data, the more likely it may occur in practice.

As for the synthetic data, we discard (redraw) any instance that is strategy proof. In this experiment we ended up discarding and redrawing a large number of instances (occasionally requiring 10,000 samples to draw 1,000 instances for testing) because the votes are highly correlated, so there was often no manipulation possible. We tested with only those datasets from ED-00006 where there was a strict linear order over all the alternatives (20 total distributions).

The results of the experiment for 6 of the 20 datasets is shown in Figure 2. The results are even more striking than those for the synthetic data. There is almost no replicability of manipulable instances when we drop just the high and low ranks. When using regular Borda scoring, manipulable instances are drawn with high replicability, and thus may occur often in practice. The results for the other 14 distributions are similar to those seen in Figure 2. The highest replicability for manipulation for the Borda rule is $\approx 40\%$ with an average of $\approx 33\%$ while dropping just the two most extreme ranks reduces the replicability of manipulation to a maximum of $\approx 13\%$ with an average of $\approx 5\%$.

7 Conclusion

We have defined and analyzed a broad class of voting rules that take into account the weighted rank that an alternative receives in the ordered list of scores obtained from a profile of voters. This new family of voting rules, RDSRs, include many frequently used rules, including positional scoring rules and Olympic style scoring. We have shown that some of these rules, which drop some extreme ranks, appear to be less manipulable in practice than traditional scoring rules. We would like to have a complete axiomatic characterization of this class of rules so that we can correctly position it with respect to traditional scoring rules and other families of aggregation procedures.

Acknowledgments: Thanks to Edith Elkind and Vincent Conitzer for helpful discussions regarding this paper. This work is supported by the National Science Foundation, under grant IIS-1107011. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. NICTA is funded by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program.

References

[1] Anonymous. The rotten heart of finance. *The Economist*, July 2012. 7 July.

[2] M. Balinski and R. Laraki. A theory of measuring, electing, and ranking. *Proc. of the National Academy of Sciences*, 104(21):8720–8725, 2007.

[3] S. Barberà. Strategyproof social choice. In *Handbook of Social Choice and Welfare*, volume 2, chapter 25, pages 731–831. 2011.

[4] J. Bartholdi, C. Tovey, and M. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.

[5] D. Black. On the rationale of group decision-making. *The Journal of Political Economy*, 56(1), 1948.

[6] D. P. Cervone, R. Dai, D. Gnoutcheff, G. Lanterman, A. Mackenzie, A. Morse, N. Srivastava, and W. S. Zwicker. Voting with rubber bands, weights, and strings. *Mathematical Social Sciences*, 64(1):11–27, 2012.

[7] R. Congar and V. Merlin. A characterization of the maximin rule in the context of voting. *Theory and Decision*, 72(1):131–147, 2012.

[8] V. Conitzer, T. Sandholm, and J. Lang. When are elections with few candidates hard to manipulate? *Journal of the ACM (JACM)*, 54(3), 2007.

[9] J. Davies, G. Katsirelos, N. Narodytska, and T. Walsh. Complexity of and algorithms for Borda manipulation. In *Proc. of the Twenty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2011)*, pages 657–662, 2011.

[10] E. Diecidue and P. P. Wakker. On the intuition of rank-dependent utility. *Journal of Risk and Uncertainty*, 23(3):281–98, November 2001.

[11] J. Duggan and T. Schwartz. Strategic manipulability without resoluteness or shared beliefs: Gibbard-Satterthwaite generalized. *Social Choice and Welfare*, 17(1):85–93, 2000.

[12] E. Elkind, P. Faliszewski, and A. Slinko. Homogeneity and monotonicity of distance-rationalizable voting rules. In *Proc. 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, pages 821–828, 2011.

[13] P. Faliszewski, E. Hemaspaandra, L. A. Hemaspaandra, and J. Rothe. The shield that never was: Societies with single-peaked preferences are more open to manipulation and control. *Information and Computation*, 209(2):89–107, 2011.

[14] P. Faliszewski and A. D. Procaccia. AI's war on manipulation: Are we winning? *AI Magazine*, 31(4):53–64, 2010.

[15] P. Fishburn. Paradoxes of voting. *The American Political Science Review*, 68(2):537–546, 1974.

[16] F. Garcin, L. Xia, and B. Faltings. How aggregators influence human rater behavior? In *Proc. Workshop @ 14th ACM Conference on Electronic Commerce (EC-13)*, 2013.

[17] T. Lu and C. Boutilier. Learning Mallows models with pairwise preferences. In *Proc. of the 28th International Conference on Machine Learning (ICML-11)*, pages 145–152, 2011.

[18] C. Mallows. Non-null ranking models. *Biometrika*, 44(1):114–130, 1957.

[19] W. Marshall and I. Olkin. *Inequalities: Theory of Majorization and its Applications*. Academic Press, London, 1979.

[20] N. Mattei. Empirical evaluation of voting rules with strictly ordered preference data. In *Proc. of the 2nd International Conference on Algorithmic Decision Theory (ADT 2011)*, pages 165–177. Springer, 2011.

[21] N. Mattei, J. Forshee, and J. Goldsmith. An empirical study of voting rules and manipulation with large datasets. In *Proc. of the 4th International Workshop on Computational Social Choice (COMSOC 2012)*. 2012.

[22] N. Mattei and T. Walsh. Preflib: A library of preference data. In *Proc. of Third International Conference on Algorithmic Decision Theory (ADT 2013)*, pages 259–270, 2013.

[23] H. Moulin. *Axioms of Cooperative Decision Making*. Cambridge University Press, 1991.

[24] H. Moulin. *Fair Division and Collective Welfare*. MIT Press, 2003.

[25] A. Popova, M. Regenwetter, and N. Mattei. A behavioral perspective on social choice. *Annals of Mathematics and Artificial Intelligence*, 68(1–3):135–160, 2013.

[26] A. D. Procaccia and J. S. Rosenschein. Average-case tractability of manipulation in voting via the fraction of manipulators. In *Proc. of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-07)*, pages 718–720, 2007.

[27] M. Regenwetter, B. Grogman, A. A. J. Marley, and I. M. Testlin. *Behavioral Social Choice: Probabilistic Models, Statistical Inference, and Applications*. Cambridge Univ. Press, 2006.

[28] A. Shorrocks. Ranking income distributions. *Economica*, 50:3–17, 1983.

[29] P. Skowron, P. Faliszewski, and J. Lang. Finding a collective set of items: From proportional multirepresentation to group recommendation. In *Proc. of the 5th International Workshop on Computational Social Choice (COMSOC 2014)*, 2014.

[30] J. Smith. Aggregation of preferences with variable electorate. *Econometrica*, 41(6):1027–1041, 1973.

[31] T. Walsh. Where are the hard manipulation problems? *Journal of Artificial Intelligence Research*, 42:1–39, 2011.

[32] J. Weymark. Generalized Gini inequality indices. *Mathematical Social Sciences*, 1:409–430, 1981.

[33] L. Xia and V. Conitzer. Finite local consistency characterizes generalized scoring rules. *Proc. of the 21st International Joint Conference on Artificial Intelligence (IJCAI 2009)*, pages 336–341, 2009.

[34] R. Yager. On ordered weighted averaging aggregation operators in multicriteria decisionmaking. *IEEE Transactions on Systems, Man and Cybernetics*, 18(1):183–190, 1988.

[35] R. R. Yager, J. Kacprzyk, and G. Beliakov, editors. *Recent Developments in the Ordered Weighted Averaging Operators: Theory and Practice*, volume 265 of *Studies in Fuzziness and Soft Computing*. Springer, 2011.

[36] H. Young. Social choice scoring functions. *SIAM Journal on Applied Mathematics*, 28(4):824–838, 1975.

Judy Goldsmith
University of Kentucky
Lexington, KY, USA
Email: goldsmit@cs.uky.edu

Jérôme Lang
LAMSADE, Université Paris-Dauphine
Paris, France
Email: lang@lamsade.dauphine.fr

Nicholas Mattei
NICTA and UNSW
Sydney, Australia
Email: nicholas.mattei@nicta.com.au

Patrice Perny
LIP6-CNRS, UPMC
Paris, France
Email: patrice.perny@lip6.fr