

Approval-Based Elections and Distortion of Voting Rules

Grzegorz Pierczynski, Piotr Skowron

University of Warsaw, Warsaw, Poland

g.pierczynski@gmail.com, p.skowron@mimuw.edu.pl

Abstract

We consider elections where both voters and candidates can be associated with points in a metric space and voters prefer candidates that are closer to those that are farther away. It is often assumed that the optimal candidate is the one that minimizes the total distance to the voters. Yet, the voting rules often do not have access to the metric space M and only see preference rankings induced by M . Consequently, they often are incapable of selecting the optimal candidate. The distortion of a voting rule measures the worst-case loss of the quality being the result of having access only to preference rankings. We extend the idea of distortion to approval-based preferences. First, we compute the distortion of Approval Voting. Second, we introduce the concept of acceptability-based distortion—the main idea behind is that the optimal candidate is the one that is acceptable to most voters. We determine acceptability-distortion for a number of rules, including Plurality, Borda, k -Approval, Veto, Copeland, Ranked Pairs, the Schulze’s method, and STV.

1 Introduction

We consider the classic election model: we are given a set of *candidates*, a set of *voters*—the voters have preferences over the candidates—and the goal is to select the *winner*, i.e., the candidate that is (in some sense) most preferred by the voters. The two most common ways in which the voters express their preferences is (i) by ranking the candidates from the most to the least preferred one, or (ii) by providing *approval sets*, i.e., subsets of candidates that they find acceptable. The collection of rankings (resp. approval sets), one for each voter, is called a ranking-based (resp. approval-based) profile. There exist a plethora of rules that define how to select the winner based on a given preference profile, and comparing these election rules is one of the fundamental questions of the social choice theory [Arrow *et al.*, 2002].

One such approach to comparing rules, proposed by Proccaccia and Rosenschein [2006], is based on the concept of *distortion*. Hereinafter, we explore its metric variant [Anshelevich *et al.*, 2018]: the main idea is to assume that the

voters and the candidates are represented by points in a metric space M called the *issue space*. The optimal candidate is the one that minimizes the sum of the distances to all the voters. However, the election rules do not have access to the metric space M itself but they only see the ranking-based profile induced by M : in this profile the voters rank the candidates by their distance to themselves, preferring the ones that are closer to those that are farther. Since the rules do not have full information about the metric space they cannot always find optimal candidates. The distortion quantifies the worst-case loss of the utility being effect of having only access to rankings. Formally, the distortion of a voting rule is the maximum, over all metric spaces, of the following ratio: the sum of the distances between the elected candidate and the voters divided by the sum of the distances between the optimal candidate and the voters.

The concept of distortion is interesting, yet—in its original form—it only allows to compare ranking-based rules. In this paper we extend the distortion-based approach so that it captures approval preferences. In the first part of the paper we analyze the distortion of Approval Voting (AV), i.e., the rule that for each approval-based profile A returns the candidate that belongs to the most approval sets from A . To formally define the distortion of AV one first needs to specify, for each metric space M , what is the approval-based profile induced by M . Here, we assume that each voter is the center of a certain ball and approves all the candidates within it. We can see that each metric space induces a (possibly large) number of approval-based profiles—we obtain different profiles for different lengths of radiuses of the balls. This is different from ranking-based profiles, where (up to tie-breaking) each metric space induced exactly one profile. Thus, the distortion of AV might depend on how many candidates the voters decide to approve. Indeed, it is easy to observe that if each voter approves all the candidates, then the rule can pick any of them, which results in an arbitrarily bad distortion. On the other hand, by an easy argument we will show that for each metric space M there exists an approval-based profile A consistent with M , such that AV for A selects the optimal candidate. In other words: AV can do arbitrarily well or arbitrarily bad, depending on how many candidates the voters approve.

Our first main contribution is that we fully characterize how the distortion of AV depends on the length of radiuses of approval balls. Specifically, we show that the distortion of




Figure 1: The relation between the fraction of voters approving the optimal candidate and the distortion of AV, for the case when the approval radiuses of the voters have all equal lengths.

AV is equal to 3, when the lengths of approval radiuses of the voters are all equal and such that the optimal candidate is approved by between $1/4$ and $1/2$ of the population of the voters (and this is the optimal distortion for the case of radiuses of equal length). The exact relation between the number of voters approving the optimal candidate and the distortion of AV is depicted in Figure 1.

In the second part of the paper we explore the following related idea: assume that the goal of the election rule is not to select the candidate minimizing the total distance to the voters, but rather to pick the one that is acceptable for most of them. E.g., AV perfectly implements this idea. A natural question is how good are ranking-based rules with respect to this criterion. To answer this question we introduce a new concept of *acceptability-based distortion* (in short, *ab-distortion*). We assume that each metric space, apart from the points corresponding to the voters and candidates, contains acceptability balls—one for each voter (as before, each voter is the center of the corresponding ball). The optimal candidate is the one that belongs to the most acceptability balls, and the ab-distortion distortion measures the normalized difference between the numbers of balls to which the elected and the optimal candidates belong. The ab-distortion is a real number between 0 and 1, where 0 corresponds to selecting the optimal candidate and 1 is the worst possible value¹.

Among the ranking-based rules that we consider in this paper, the best (and the optimal) ab-distortion is attained by Ranked Pairs and the Schulze’s method. It is an open question, whether Ranked Pairs is the only natural rule with this property. It is worth mentioning, that its ab-distortion is closely related to the size of the Smith set, so in case it is small (in particular, when the Condorcet winner exists) this rule has even better ab-distortion. We have found an interesting result for Copeland. Although in case of classic (distance-based) distortion most Condorcet rules are equally good, this

¹The reader might wonder why we define the ab-distortion as a difference rather than as a ratio (as it is done for the classic definition of the distortion). Indeed, we first used the ratios in our definition, but then it was very easy to construct instances where any rule had the distortion of $+\infty$. Further, we found that these results do not really speak of the nature of the rules but rather are artifacts of the used definition. Consequently, we found that the considering the difference gives more meaningful results.

Rule	Distance-based distortion ($[1; +\infty)$)	Ab-distortion ($[0; 1]$)
Each rule	≥ 3	$\geq \frac{\ell-1}{\ell}$ for $\ell > 1$ $\geq \frac{1}{2}$ otherwise
Plurality	$2m - 1$	$\frac{m-1}{m}$
Borda	$2m - 1$	$\frac{m-1}{m}$
k-approval	∞	1
Veto	∞	1
Copeland	5	1
Ranked Pairs	5	$\frac{\ell-1}{\ell}$ for $\ell > 1$ $\frac{1}{2}$ otherwise
Schulze’s Rule	5	$\frac{\ell-1}{\ell}$ for $\ell > 1$ $\frac{1}{2}$ otherwise
STV	$O(\ln m)$	$\frac{2^{m-1}-1}{2^{m-1}}$

Table 1: The comparison of the distortion for various ranking-based rules. The results in the left column (for the distance-based distortion) are known in the literature. The results for ab-distortion are new to this paper; here, m denotes the number of the candidates and ℓ is the size of the Smith set.

is no longer the case when acceptability is the criterion we primarily care about. The ab-distortion of Copeland is equal to 1, which is the worst possible value. This rule is optimal only if the Condorcet winner exists (e.g. when the metric space is one-dimensional). The distortion of scoring rules (Plurality, Borda, Veto, k-approval) is significantly worse than that for Ranked Pairs. Another surprising result is the distortion of STV—while this rule is known to achieve a very good distance-based distortion, its ab-distortion is even worse than for Plurality (denoting the number of candidates as m , STV and Plurality achieve the ab-distortion of $\frac{2^{m-1}-1}{2^{m-1}}$ and $\frac{m-1}{m}$, respectively). In case of all these rules the worst-case instances were obtained in one-dimensional Euclidean metric spaces. Our results are summarized in Table 1.

The motivation for studying the two types of distortion (the distance-based and the approval-based distortion) is the following. Very often it is not clear what should be the goal of an election. It has been argued that there is a tension between the desire to select a candidate judged “as highly as possible” and supported “by as many people as possible”. Ideally, one would like to have a rule that works well with respect to the both criteria. Our goal is to examine whether such an ideal rule exists, and if not, to assess the tension between the two criteria, focusing on the well-known rules.

Due to space constraints we redelegate all the proofs to the full version of the paper [Pierczynski and Skowron, 2019].

2 Preliminaries

For each set S by 2^S and $\Pi(S)$ we denote, respectively, the powerset of S and the set of all linear orders over S . By S^c we denote the complement set of S , and by S^* —the set of all vectors with the elements from S . For each two sets S_1, S_2

and a function $f: S_1 \rightarrow 2^{S_2}$ by $R_f: S_2 \rightarrow 2^{S_1}$ we denote the function defined as follows:

$$\forall y \in S_2 \quad R_f(y) = \{x \in S_1: y \in f(x)\}$$

For convenience we assume that $[-\infty; +\infty]$ denotes the affinely extended real number system (the set of real numbers \mathbb{R} with additional symbols $+\infty$, and $-\infty$). We take the following convention for arithmetical operations:

$$\forall a \in \mathbb{R} \quad \frac{a}{\pm\infty} = 0 \quad \forall a \in (0; +\infty] \quad \frac{\pm a}{0} = \pm\infty.$$

Expressions $\frac{0}{0}$, $\frac{\pm\infty}{\pm\infty}$, $0 \cdot \pm\infty$ and $\pm\infty - \pm\infty$ are undefined.

2.1 Our Metric Model

An *election instance* is a tuple (N, C, d, λ) , where $N = \{1, 2, \dots, n\}$ is the set of *voters*, $C = \{c_1, c_2, \dots, c_m\}$, is the set of *candidates*, $d: (N \cup C)^2 \rightarrow \mathbb{R}$ is a *distance function* (d allows us to view the candidates and the voters as points in a pseudo-metric space), and $\lambda: N \rightarrow 2^C$ is an *acceptability function*, mapping each voter $i \in N$ to a subset of candidates that i finds acceptable. We assume that λ is *nonempty*², i.e., for each $i \in N$, $\lambda(i) \neq \emptyset$, and that is *local consistent*—for each $i \in N$, $c_a, c_b \in C$, if $c_a \in \lambda(i)$ and $d(i, c_b) \leq d(i, c_a)$, then $c_b \in \lambda(i)$. Often we will also require that λ satisfies a stronger condition, called *global consistency*—for each $i, j \in N$, $c_a, c_b \in C$, if $c_a \in \lambda(i)$ and $d(j, c_b) \leq d(i, c_a)$, then $c_b \in \lambda(j)$. Intuitively, local-consistency means that for each voter $i \in N$ we can associate $\lambda(i)$ with a ball with the center at the point of this voter. A voter $i \in N$ considers a candidate c_j to be acceptable for him, $c_j \in \lambda(i)$, if and only if c_j lies within the ball. Such a ball will be further called the *acceptability ball* and its radius—the *acceptability radius*. Then, global consistency can be interpreted as an assumption that all the acceptability radiuses have equal lengths.

We will sometimes slightly abuse the notation: by saying that an instance satisfies local (global) consistency we will mean that the acceptability function in the instance satisfies the respective property.

By \mathbb{I} , we denote the set of all election instances. Since issue spaces are often argued to be Euclidean spaces with small numbers of dimensions, we additionally introduce the following notation: for each $k \in \mathbb{N}$ let \mathbb{E}^k denote the set of all the instances where the elements of N and C are associated with points from \mathbb{R}^k , and d is the Euclidean distance.

2.2 Preference Representation

In most cases, it is difficult for the voters to explicitly position themselves in the issue space, and often even the space itself is unknown. Therefore, we will consider voting rules that take as inputs preference profiles induced by election instances, instead of instances themselves. We consider two classic approaches to represent preferences.

²We add this assumption to avoid pathological situations where some voters like no candidates. This way we avoid having somehow degenerated hard cases for the distortion.

Ranking-based profiles. A *ranking-based profile* induced by an election instance $I = (N, C, d, \lambda)$ is the function $\succ_I: N \rightarrow \Pi(C)$, mapping each voter to a linear order over C such that for all $i \in N$ and all $c_a, c_b \in C$ if $d(i, c_a) < d(i, c_b)$ then $c_a \succ_i c_b$.³ For each voter $i \in N$, the relation $\succ_I(i)$ (for convenience also denoted as \succ_i , whenever the instance is clear from the context) is called the *preference order* of i . If for some $c_x, c_y \in C$ it holds that $c_x \succ_i c_y$, we say that i prefers c_x over c_y .

Approval-based profiles. An *approval-based profile* of an election instance $I = (N, C, d, \lambda)$ is a locally consistent acceptability function $A: N \rightarrow 2^C$. We say that a candidate c_x is *approved* by a voter $i \in N$ if $c_x \in A(i)$. We will say that the approval-based profile is *truthful* if for all $i \in N$ it holds that $A(i) = \lambda(i)$.

Additional notation. Let us introduce some additional useful notation. Let $P: C^* \rightarrow N$ be a function mapping vectors of distinct candidates to sets of voters as follows:

$$P((c_{i_1}, c_{i_2}, \dots, c_{i_k})) = \{v \in N : c_{i_1} \succ_v c_{i_2} \succ_v \dots \succ_v c_{i_k}\}$$

For convenience, we will write $P(c_{i_1}, c_{i_2}, \dots, c_{i_k})$ instead of $P((c_{i_1}, c_{i_2}, \dots, c_{i_k}))$.⁴ Note that for all $c_a, c_b \in C$ we have $P(c_a, c_b) \cap P(c_b, c_a) = \emptyset$ and $P(c_a, c_b) \cup P(c_b, c_a) = N$.

We say that c_a *dominates* c_b if $|P(c_a, c_b)| > \frac{n}{2}$ and that c_a *weakly dominates* c_b if $|P(c_a, c_b)| \geq \frac{n}{2}$. We say that a candidate c_x *Pareto-dominates* a candidate c_y if there holds that $|P(c_x, c_y)| = n$. A candidate c_y is *Pareto-dominated* if there exists a candidate c_x who Pareto-dominates c_y .

2.3 Definitions of Voting Rules

An *election rule* (also referred to as a *voting rule*) is a function mapping each preference profile to a set of tied *winners*. We distinguish *ranking-based* rules—taking ranking-based profiles as arguments, and *approval-based* rules—defined analogously. Among approval-based rules, we focus on *Approval Voting* (AV)—the rule that selects those candidates that are approved by most voters. In the remaining part of this subsection we recall definitions of the ranking-based rules that we study in this paper.

Positional scoring rules. For a given vector $\vec{s} = (\alpha_1, \alpha_2, \dots, \alpha_m)$, the scoring rule implemented by \vec{s} works as follows. A candidate c_a gets α_i points from each voter j who puts c_a in the i th position in \succ_j . The rule elects the candidates whose total number of points, collected from all the voters, is maximal. Some well-known scoring rules which we will study in the further part of this work are the following:

Plurality: $\vec{s} = (1, 0, \dots, 0)$,

Veto: $\vec{s} = (1, \dots, 1, 0)$,

Borda: $\vec{s} = (m - 1, m - 2, \dots, 1, 0)$,

³We allow for ties in the metric space—the voter ranks the equidistant candidates using a fixed tie-breaking rule. Allowing for ties gives somehow simpler hard-instances, yet they can be tweaked so that they do not depend on tie-breaking at all. Also, the positive results (the upper-bounds) work even for adversarial tie-breaking.

⁴It will always be clear from the context whether in the inscription $P(x)$, x should be interpreted as a vector or as a candidate.

k-approval: $\vec{s} = (\underbrace{1, \dots, 1}_k, 0, \dots, 0)$ (for $1 \leq k \leq m$).

Copeland. The *Copeland* rule elects candidates c_w who dominate at least as many candidates as any other candidate. More formally, a candidate c_w is a winner if and only if:

$$\forall c_x \in C \quad |\{c: |P(c_w, c)| > \frac{n}{2}\}| \geq |\{c: |P(c_x, c)| > \frac{n}{2}\}|$$

Ranked Pairs. *Ranked Pairs* works as follows: first we sort the pairs of candidates (c_i, c_j) in the descending order of the values $|P(c_i, c_j)|$. Then, we construct a graph G where the vertices correspond to the candidates. We start with the graph with no edges; then we iterate over the sorted list of pairs—for each pair (c_i, c_j) we add an edge from c_i to c_j unless there is already a path from c_j to c_i in G . If such a path exists, we simply skip this pair. Clearly, the so-constructed graph G is acyclic. The source nodes of G are the winners.

Later on we will also use the well-known fact that Ranked Pairs always elects a subset of the *Smith set*. Recall that the Smith set of an instance I is the subset S of the candidates such that (i) every candidate from S dominates each candidate outside of S , and (ii) no non-empty proper subset of S satisfies the previous condition.

The Schulze’s Rule. The *Schulze’s rule* works as follows: let the *beatpath* of length k from candidate c_a to c_b be a sequence of candidates $c_{x_1}, c_{x_2}, \dots, c_{x_{k-1}}$ such that c_a dominates c_{x_1} , $c_{x_{k-1}}$ dominates c_b and for each $i \in \{1, \dots, k-2\}$, c_{x_i} dominates $c_{x_{i+1}}$. Let the *strength* of the beatpath be the minimum of values $P(c_a, c_{x_1}), P(c_{x_1}, c_{x_2}), \dots, P(c_{x_{k-1}}, c_b)$. By $p[c_a, c_b]$ we denote the maximum of strengths of all beatpaths from c_a to c_b . Candidate c_w is the winner if and only if for each candidate c it holds that $p[c_w, c] \geq p[c, c_w]$.

STV. *Single Transferable Vote (STV)* works iteratively as follows: if there is only one candidate, elect this candidate. Otherwise, eliminate the candidate who has the least points according to the Plurality rule and repeat the algorithm.

Note that the aforementioned rules are irresolute by definition. Further, we did not specify the tie-breaking rule used when sorting edges in Ranked Pairs and when eliminating candidates in STV. We will make all these rules resolute by using the lexicographical tie-breaking rule, denoted by \succ_{lex} .

2.4 Measuring the Quality of Social Choice

In this section we formalize the concept of distortion that, on the intuitive level, we already introduced in Section 1.

Distance-based approach. A natural idea to relate the quality of a candidate c with the sum of the distances from this candidate to all the voters. The lower this sum is, the higher the quality. Following this intuition, the *distortion* of a voting rule φ in instance $I \in \mathbb{I}$, is defined as follows (below, c_o denotes the optimal candidate for I):

$$D_I(\varphi) = \max_{P \in \mathcal{P}_I} \frac{\sum_{i \in N} d(i, \varphi(P))}{\sum_{i \in N} d(i, c_o)},$$

where \mathcal{P}_I is the set of profiles induced by I (either ranking or approval, depending on the domain of φ). $D(\varphi) \in [1; +\infty]$.

This approach can be applied to any rule discussed so far. For ranking-based rules it has already been widely studied in the literature, hence in the further part we will focus on AV.

Acceptability-based approach. Now we present an alternative way to measure the quality of candidates, based on the acceptability function. Intuitively, the more voters a candidate c is acceptable for, the higher his quality. Besides, we would like the maximal possible quality not to depend on the number of voters. Therefore, we define the *acceptability-based distortion (ab-distortion, in short)* of a voting rule φ in instance $I \in \mathbb{I}$ as the following expression:

$$D_I(\varphi) = \max_{P \in \mathcal{P}_I} \frac{R_\lambda(c_o) - R_\lambda(\varphi(P))}{n},$$

where \mathcal{P}_I is the set of profiles induced by I (either ranking-based or approval-based, depending on the domain of φ). Clearly, the ab-distortion is always a value from $[0; 1]$. By definition, in truthful profiles Approval Voting always elects an optimal candidate in terms of ab-distortion. Thus, we will consider our acceptability-based measure only for ranking-based rules.

Let E be a mathematical expression that can depend on some characteristics of an instance (e.g., on the number of candidates, or the size of the Smith set). We say that the (acceptability-based) distortion of a voting rule φ is E , if for each instance I , $D_I(\varphi) \leq E$ and for each E and each $\epsilon > 0$ there is a instance I with $D_I(\varphi) > E - \epsilon$.

3 Distortion of Approval Voting

In this section we analyze the distance-based distortion of Approval Voting (AV)—hereinafter we denote AV by φ_{AV} .

We start by showing that in the most general case, if we do not make any additional assumptions about the acceptability function, the distortion of AV can be arbitrarily bad.

Proposition 3.1. *There exists an instance $I \in \mathbb{I}^1$ such that $D_I(\varphi_{AV}) = +\infty$.*

This result is rather pessimistic. However, one could ask a somehow related question—does there for each instance I always exist an approval profile consistent with I that would result in a good distortion? In contrast to Proposition 3.1, here the answer is much more positive.

Proposition 3.2. *For each instance $I \in \mathbb{I}$, there is an approval based profile A of I such that φ_{AV} is the optimal candidate (minimizing the total distance to voters).*

Propositions 3.1 and 3.2 show that for each metric space M there always exists two approval profiles A_1, A_2 locally consistent with M such that for A_1 AV selects the worst possible candidate, and for A_2 it selects the optimal one—since A_1 and A_2 are both consistent with M , they only differ in the sizes of approval balls. This formally shows that the performance of AV strongly depends on how many candidates the voters decide to approve. Below, we provide our main result of this section—assuming that all the acceptability balls have radiuses of the same length, we show the exact relation between this length of approval radiuses and the distance-based distortion of AV. In particular, we show that the best approval

radius is such that the optimal candidate is approved by between $1/4$ and $1/2$ fraction of all the voters.

Definition 3.3. An approval-based profile A induced by an instance I is p -efficient for $p \in [0; 1]$ if $R_A(c_o) = pn$.

In words, a profile is p -efficient if the number of voters who approve the optimal candidate is the p fraction of n .

Theorem 3.4. For each globally consistent p -efficient instance I , we have the following results:

$$D_I(\varphi_{AV}) \leq \begin{cases} +\infty & \text{for } p \in \{0, 1\} \\ \frac{1-p}{p} & \text{for } p \in (0; \frac{1}{4}] \\ 3 & \text{for } p \in [\frac{1}{4}, \frac{1}{2}] \\ \frac{2-p}{1-p} & \text{for } p \in [\frac{1}{2}, 1). \end{cases}$$

The above function is depicted in Figure 1.

All these bounds are attained for instances in \mathbb{E}^1 . The detailed construction is provided in the full version of the paper [Pierczynski and Skowron, 2019].

Finally, for completeness, we give an analogue of Proposition 3.2, but for globally-consistent instances.

Proposition 3.5. For each instance $I \in \mathbb{I}$, there exists an approval profile A globally consistent with I , such that

$$\frac{\sum_{i \in N} d(i, \varphi(A))}{\sum_{i \in N} d(i, c_o)} \leq \frac{11}{3}.$$

4 AB-Distortion of Ranking Rules

Recall that the ab-distortion of a voting rule is a value from $[0; 1]$, proportional to the difference between the number of voters accepting the optimal candidate and the number of voters accepting the winner. By definition, this value equals 0 for AV (provided the approval profile is truthful). In this section we analyze the ab-distortion of ranking-based rules.

We start by proving the lower bound on the ab-distortion of any ranking-based voting rule.

Theorem 4.1. For each $\ell \in \mathbb{N}$ and each ranking-based rule φ , there exists a globally consistent instance I such that:

1. the size of the Smith set in the ranking-based profile induced by I equals ℓ ,
2. $D_I(\varphi) = \begin{cases} \frac{\ell-1}{\ell} & \text{for } \ell \geq 2 \\ \frac{1}{2} & \text{for } \ell = 1. \end{cases}$

In the subsequent part of this section we will assess the distortion of specific voting rules, specifically looking for one that meets the lower-bound from Theorem 4.1.

4.1 Condorcet Rules

We start by looking at Condorcet consistent rules. Note that the lower bound found in Theorem 4.1 is promising, as it depends on the size of the Smith set. In particular, if $\ell = 1$, this bound equals $\frac{1}{2}$. Our first goal is to determine, whether Condorcet rules meet this bound.

Theorem 4.2. Let I be an instance with a Condorcet winner and φ be a Condorcet consistent rule. Then, $D_I(\varphi) \leq \frac{1}{2}$. This bound is achievable for a globally consistent $I \in \mathbb{E}^2$.

Proof. Let c_w be the winner and c_o be the optimal candidate. Since we assumed that the Condorcet candidate exists in I , we have that c_w weakly dominates c_o . Then, we have:

$$\frac{n}{2} \leq |P(c_w, c_o)| = n - |P(c_o, c_w)|$$

Thus, $|P(c_o, c_w)| \leq \frac{n}{2}$. Further, we have that:

$$\begin{aligned} |R_\lambda(c_o)| - |R_\lambda(c_w)| &= |R_\lambda(c_o) \setminus R_\lambda(c_w)| + |R_\lambda(c_o) \cap R_\lambda(c_w)| \\ &\quad - |R_\lambda(c_w) \setminus R_\lambda(c_o)| - |R_\lambda(c_w) \cap R_\lambda(c_o)| \\ &= |R_\lambda(c_o) \setminus R_\lambda(c_w)| - |R_\lambda(c_w) \setminus R_\lambda(c_o)| \\ &\leq |R_\lambda(c_o) \setminus R_\lambda(c_w)| \leq P(c_o, c_w) \leq \frac{n}{2}. \end{aligned}$$

This completes the first part of the proof.

For the hard instance, we have four candidates c_x, c_y, c_z , and c_c , placed in points $(3, 3)$, $(0, 0)$, $(6, 0)$, and $(3, 2)$, respectively. There are $\frac{n}{2} - 1$ voters in $(3, 4)$ with preferences $c_x \succ c_c \succ c_y, c_z$, approving only c_x , $\frac{n}{4} + 1$ voters in $(1, 0)$ with rankings $c_y \succ c_c \succ c_x, c_z$, approving only c_y , and $\frac{n}{4}$ voters in $(5, 0)$ with preferences $c_z \succ c_c \succ c_x, c_y$, approving only c_z . Candidate c_c is the Condorcet winner and the optimal candidate is c_x . The distortion of each rule electing c_c is $\frac{1}{2} - \frac{1}{n}$, which is arbitrarily close to $\frac{1}{2}$. \square

From the above theorem, we get that for $\ell = 1$ each Condorcet election method matches the lower bound from Theorem 4.1. Now we will prove that there exists election rules, namely Ranked Pairs and the Schulze's rule, which match this bound for each ℓ .

Theorem 4.3. The ab-distortion of Ranked Pairs and of the Schulze's rule is equal to:

- $\frac{\ell-1}{\ell}$ for $\ell \geq 2$,
- $\frac{1}{2}$ for $\ell = 1$

where ℓ is the size of the Smith set of considered instance.

As we can see, there is no rule with a better ab-distortion than Ranked Pairs. Yet, it is not a feature of all the Condorcet methods. As we will see, even for the well-known Copeland rule, the possible pessimistic distortion is much worse.

Theorem 4.4. The ab-distortion of Copeland is 1. This bound is achieved for globally consistent instances in \mathbb{E}^2 .

4.2 Scoring Rules

Let us now move to positional scoring rules. Here, we obtain significantly worse results than for Ranked Pairs and the Schulze's rule. A general tight upper bound for the ab-distortion of any scoring rule remains an open problem. Below we provide bounds that are tight for certain specific scoring rules.

Theorem 4.5. For a scoring rule φ defined by vector $\vec{s} = (s_1, \dots, s_m)$ the ab-distortion of φ satisfies:

1. $D_I(\varphi) = 1$, if $s_1 = \dots = s_m$,
2. $D_I(\varphi) \leq \frac{\max_{i,j} |s_i - s_j|}{\max_{i,j} |s_i - s_j| + \min_{i,j} |s_i - s_j|}$, otherwise.

The bound obtained in Theorem 4.5 is not tight in general. For example, for Plurality we have a tighter estimation.

Theorem 4.6. *The ab-distortion of Plurality is $\frac{m-1}{m}$. This bound is achieved for globally consistent instances in \mathbb{E}^1 .*

Yet, for a number of scoring rules the bound from Theorem 4.5 is tight. Below, we give some sufficient conditions.

Proposition 4.7. *The bound from Theorem 4.5 is tight for each scoring rule satisfying the following conditions:*

1. $s_1 \geq \dots \geq s_m$,
2. $\forall_{1 \leq i \leq m-1} s_1 - s_2 \leq s_i - s_{i+1}$

even for globally consistent instances in \mathbb{E}^1 .

Theorem 4.5 combined with Theorem 4.7 imply the ab-distortion for a number of scoring rules.

Corollary 4.8. *There exists a globally consistent instance $I \in \mathbb{E}^1$, for which:*

1. the ab-distortion of k -approval is $\frac{1}{1+k} = 1$,
2. the ab-distortion of Veto is $\frac{1}{1+0} = 1$,
3. the ab-distortion of Borda is $\frac{m-1}{m-1+1} = \frac{m-1}{m}$.

4.3 Iterative Rules

All scoring rules that we considered have poor ab-distortion, and in particular are considerably worse than Condorcet rules (especially for instances with Condorcet winners).

Interestingly, STV in terms of acceptability, behaves worse even than Plurality. This is somehow surprising since for distance-based distortion, STV is better than any positional scoring rules, and only slightly worse than Condorcet rules.

Theorem 4.9. *The ab-distortion of STV is $\frac{2^{m-1}-1}{2^{m-1}}$.*

The above bound is tight even in one-dimensional Euclidean spaces. It is also tight if we restrict ourselves to global consistent instances. There, the hard instances that we found use $(m-2)$ -dimensional Euclidean space.

Proposition 4.10. *The bound from Theorem 4.9 is tight for locally consistent instances from \mathbb{E}^1 and globally consistent instances from \mathbb{E}^{m-2} .*

5 Related Work

The spatial model of preferences is quite popular in the social choice and political science literature. For example seminal works studying spacial models we refer the reader to [Davis and Hinich, 1966; Plott, 1967; Enelow and Hinich, 1984; Enelow and Hinich, 1990; McKelvey and Ordeshook, 1990; Merrill and Grofman, 1999; Schofield, 2007].

The concept of distortion was first introduced by Procaccia and Rosenschein [2006]. In their work they did not assume the existence of a metric space, but rather used a generic cardinal utility model (where the voters can have arbitrarily utilities for candidates). This model was later studied by Caragiannis and Procaccia [2011] and Boutilier *et al.* [2015]. Recently, Benadé *et al.* [2019] introduced the concept of distortion for social welfare functions, i.e., functions mapping voters preferences to rankings over candidates, and Benadé *et al.* [2017] adapted and used the concept of distortion in the context of participatory budgeting to evaluate different methods of preference elicitation.

The studies of the distortion in the metric model were initiated by Anshelevich *et al.* [2018], and then continued by Anshelevich and Postl [2017], Feldman *et al.* [2016], Goel *et al.* [2017], and Gross *et al.* [2017].

The analysis of the distortion forms a part of a broader trend in social choice stemming from the utilitarian perspective. For classic works in welfare economics that discuss the utilitarian approach we refer the reader to the article of Ng [1997] and the book of Roemer [1998]. This approach has also recently received a lot of attention from the computer science community. Apart from the papers that directly study the concept of distortion that we discussed before, examples include the works of Filos-Ratsikas and Miltersen [2013], Branzei *et al.* [2013], and Chakrabarty and Swamy [2014].

6 Conclusion

In this paper we have extended the concept of distortion of voting rules to approval-based preferences. This extension allows to compare rules that take different types of input: approval sets and rankings over the candidates. To the best of our knowledge, only very few formal methods are known that allow for such a comparison. We are aware of only one work that formally relates these two models: Laslier and Sanver [2010] proved that in the strong Nash equilibrium Approval Voting selects the Condorcet winner, if such exists. There are also some works that model elections in which voters cast both rankings and (consistent) approval ballots [Camps *et al.*, 2014].

Our contribution is twofold. First, we have determined the distortion of Approval Voting, and explained how this distortion depends on voters' approval sets. Specifically, we have shown that the socially best outcome is obtained when the voters approve not too many and not too few candidates. If the lengths of voters' acceptability radiuses are all equal, the best distortion is obtained when the approval sets are such that between $\frac{1}{4}$ and $\frac{1}{2}$ of the voters approve the optimal candidate.

Second, we have defined a new concept of acceptability-based distortion (ab-distortion). Here, we assume that the voters have certain acceptability thresholds; the ab-distortion of a given rule φ measures how many voters (in the worst-case) would be satisfied from the outcomes of φ . We have determined the ab-distortion for a number of election rules (our results are summarized in Table 1), and reached the following conclusions. The analysis of the classic and the acceptability-based distortions both suggest that Condorcet rules perform better than scoring and iterative ones. Further, our acceptability-based approach suggests that Ranked Pairs and the Schulze's rule are particularly good rules, in particular significantly outperforming the Copeland's rule. Thus, our study recommends Ranked Pairs or the Schulze's method as a rule that robustly performs well for both criteria (total distance, and acceptability). The question whether Ranked Pairs is the only natural ranking-based rule performing well for both criteria is open. Approval Voting is also a very good rule that can be considered an appealing alternative to Ranked Pairs, provided the sizes of the approval sets of the voters are appropriate.

Acknowledgements

The authors were supported by the Foundation for Polish Science within the Homing programme (Project title: "Normative Comparison of Multiwinner Election Rules").

References

- [Anshelevich and Postl, 2017] E. Anshelevich and J. Postl. Randomized social choice functions under metric preferences. *Journal of Artificial Intelligence Research*, 58:797–827, 2017.
- [Anshelevich *et al.*, 2018] E. Anshelevich, O. Bhardwaj, E. Elkind, J. Postl, and P. Skowron. Approximating optimal social choice under metric preferences. *Artificial Intelligence*, 264:27–51, 2018.
- [Arrow *et al.*, 2002] K. Arrow, A. Sen, and K. Suzumura, editors. *Handbook of Social Choice and Welfare, Volume 1*. Elsevier, 2002.
- [Benadé *et al.*, 2017] G. Benadé, S. Nath, A. Procaccia, and N. Shah. Preference elicitation for participatory budgeting. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, pages 376–382, 2017.
- [Benadé *et al.*, 2019] G. Benadé, A. Procaccia, and M. Qiao. Low-distortion social welfare functions. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence*, 2019. To appear.
- [Boutilier *et al.*, 2015] C. Boutilier, I. Caragiannis, S. Haber, T. Lu, A. D. Procaccia, and O. Sheffet. Optimal social choice functions: A utilitarian view. *Artificial Intelligence*, 227:190–213, 2015.
- [Branzei *et al.*, 2013] S. Branzei, I. Caragiannis, J. Morgenstern, and A. Procaccia. How bad is selfish voting? In *Proceedings of the 27th AAAI Conference on Artificial Intelligence*, pages 138–144, 2013.
- [Camps *et al.*, 2014] R. Camps, X. Mora, and L. Saumell. Social choice rules driven by propositional logic. *Annals of Mathematics and Artificial Intelligence*, 70(3):279–312, 2014.
- [Caragiannis and Procaccia, 2011] I. Caragiannis and A. D. Procaccia. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence*, 175(9–10):1655–1671, 2011.
- [Chakrabarty and Swamy, 2014] D. Chakrabarty and C. Swamy. Welfare maximization and truthfulness in mechanism design with ordinal preferences. In *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science*, pages 105–120, 2014.
- [Davis and Hinich, 1966] O. A. Davis and M. J. Hinich. A mathematical model of preference formation in a democratic society. In J. L. Bernd, editor, *Mathematical Applications in Political Science II*, pages 175–208. Southern Methodist University Press, 1966.
- [Enelow and Hinich, 1984] J. M. Enelow and M. J. Hinich. *The spatial theory of voting: An introduction*. CUP Archive, 1984.
- [Enelow and Hinich, 1990] J. M. Enelow and M. J. Hinich. *Advances in the spatial theory of voting*. Cambridge University Press, 1990.
- [Feldman *et al.*, 2016] M. Feldman, A. Fiat, and I. Golomb. On voting and facility location. In *Proceedings of the 17th ACM Conference on Economics and Computation*, pages 269–286, 2016.
- [Filos-Ratsikas and Miltersen, 2013] A. Filos-Ratsikas and P. B. Miltersen. Truthful approximations to range voting. In *Proceedings of the 10th International Conference on Web and Internet Economics*, pages 175–188, 2013.
- [Goel *et al.*, 2017] A. Goel, A. K. Krishnaswamy, and K. Munagala. Metric distortion of social choice rules: Lower bounds and fairness properties. In *Proceedings of the 18th ACM Conference on Economics and Computation*, pages 287–304, 2017.
- [Gross *et al.*, 2017] S. Gross, E. Anshelevich, and L. Xia. Vote until two of you agree: Mechanisms with small distortion and sample complexity. In *Proceedings of the 31st Conference on Artificial Intelligence*, pages 544–550, 2017.
- [Laslier and Sanver, 2010] J.-F. Laslier and M. Sanver. *The Basic Approval Voting Game*, pages 153–163. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010.
- [McKelvey and Ordeshook, 1990] R. D. McKelvey and P. C. Ordeshook. A decade of experimental research on spatial models of elections and committees. *Advances in the spatial theory of voting*, pages 99–144, 1990.
- [Merrill and Grofman, 1999] S. Merrill and B. Grofman. *A unified theory of voting: Directional and proximity spatial models*. Cambridge University Press, 1999.
- [Ng, 1997] Y.-K. Ng. A case for happiness, cardinalism, and interpersonal comparability. *The Economic Journal*, 107(445):1848–1858, 1997.
- [Pierczynski and Skowron, 2019] G. Pierczynski and P. Skowron. Approval-based elections and distortion of voting rules. Technical Report arXiv:1901.06709 [cs.GT], arXiv.org, 2019.
- [Plott, 1967] C. R. Plott. A notion of equilibrium and its possibility under majority rule. *The American Economic Review*, 57(4):787–806, 1967.
- [Procaccia and Rosenschein, 2006] A. D. Procaccia and J. S. Rosenschein. The distortion of cardinal preferences in voting. In *Proceedings of the 10th International Workshop on Cooperative Information Agents (CIA-2006)*, pages 317–331, 2006.
- [Roemer, 1998] J. E. Roemer. *Theories of distributive justice*. Harvard University Press, 1998.
- [Schofield, 2007] N. Schofield. *The spatial model of politics*. Routledge, 2007.