Independence of Irrelevant Alternatives under the Lens of Pairwise Distortion

Théo Delemazure[1], **Jérôme Lang**[1], **Grzegorz Pierczyński**[2]

[1] CNRS, Paris Dauphine University, PSL
[2] University of Warsaw

Abstract

We give a quantitative analysis of the *independence of irrelevant alternatives* (IIA) axiom. IIA says that the society's preference between x and y should depend only on individual preferences between x and y: we show that, in several contexts, if the individuals express their preferences about additional (or "irrelevant") alternatives, this information helps to estimate better which of x and y has higher social welfare. Our contribution is threefold: (1) we provide a new tool to measure the impact of IIA on social welfare (*pairwise distortion*), based on the well-established notion of voting distortion, (2) we study the average impact of IIA in both general and metric settings, with experiments on synthetic data, and its impact with real datasets; and (3) we study the worst-case impact of IIA in the 1D-Euclidean metric space.

Introduction

Independence of irrelevant alternatives (IIA) states that a society's preference between two alternatives x and y depends only on how its individual members compare x and y (Arrow 1950), as opposed to taking into account the positions they give to other alternatives. IIA is the key axiom of Arrow's impossibility theorem: in ordinal settings, with at least three alternatives, IIA is incompatible with the unrestricted domain assumption, Pareto-efficiency, and non-dictatorship—three properties that are hard to give up.

Arrow's theorem has had, until today, a tremendous importance in social choice. It is often seen as negative, since there are many compelling arguments in favor of IIA (see for instance (Maskin 2020)). Still, IIA has been criticized for not taking into account preference intensities, and this may result in a loss of social welfare. To understand this argument, consider the following example:

$$51\% \text{ of voters: } x \succ y \succ z_1 \succ \ldots \succ z_{100}$$
$$49\% \text{ of voters: } y \succ z_1 \succ \ldots \succ z_{100} \succ x$$

What should be the collective preference between x and y here? If we exclude "irrelevant" alternatives $z_1, \ldots, z_{100}$, then the rational choice seems to use majority and conclude $x \succ y$.[1] On the other hand, we may have a strong intuition that it is not the right choice, because the additional alternatives provide us some implicit information: the preference $x \succ y$ of the first group of voters is most probably much *weaker* than the preference $y \succ x$ of the second group.

However this argument lacked until now a *quantitative* analysis. We argue that a recent research trend, *ordinal-cardinal voting distortion*, that aims at measuring the loss of social welfare (sum of individual utilities) caused by the use of ordinal instead of cardinal information, provides a suitable framework for quantifying the impact of IIA on social welfare. In our example above, if we had access to information about strength of preference, expressed as cardinal utilities, we would probably find that y has a higher social welfare than x. By denying access to information about strength of preference, IIA makes us choose x against y, which leads to what we call a *pairwise distortion* relative to x and y.

This being said, there are various ways of exploiting the additional information given by the irrelevant alternatives, and the question of which ones actually help reducing the loss of social welfare is not trivial. On the latter example, using plurality scores to choose between x and y still leads to choosing x; and using Copeland scores too. But using Borda scores leads this time to choosing y.

For choosing between two alternatives given a profile, we define *pairwise voting rules*. They take as input two alternatives (say, x and y) and a preference profile on a set of alternatives containing x and y, and output either x or y. Now the most natural pairwise voting rule satisfying IIA is the *pairwise majority rule*, that outputs the result of majority voting between x and y for each x and y (note that this is the only pairwise rule satisfying IIA and extending the majority rule for elections with two alternatives). Of course, applying pairwise majority to all pairs of alternatives sometimes returns a nontransitive relation over alternatives.

On the other hand, any social welfare function mapping a preference profile to a collective ranking of alternatives induces a (transitive) pairwise voting rule, where the choice between x and y is made by projecting the collective ranking on $\{x, y\}$. Thus, pairwise voting rules are a common frame-

[1] As a consequence of May's theorem (May 1952), IIA, together with a set of very weak conditions (neutrality, anonymity, and positive responsiveness), implies that the collective preference between x and y must be determined by majority voting between x and y.

work capturing both pairwise majority, which satisfies IIA, and (nondictatorial) social welfare functions, that do not.

Next, we define *pairwise distortion* of a pairwise voting rule f relative to x, y and a profile, as the loss of social welfare caused by choosing the alternative among $\{x, y\}$ determined by f. The details of the definition may vary—we focus on the two following settings: (1) in the first one, we consider *average pairwise distortion*, where the utilities or costs of the alternatives follow a given distribution, or are drawn from real datasets, (2) in the second one, we consider *worst-case pairwise distortion* in metric domains (voters and alternatives are located in a metric space and voters' preferences decrease with distance).

Now, the main question is the following: which pairwise rules give a lower pairwise distortion than pairwise majority, and how can we compare these rules according to their pairwise distortion? Our main findings are:

- When considering *average* distortion of synthetic data, as well as *empirical* distortion of real datasets, the overall picture is that (i) pairwise rules based on Borda and Copeland perform much better than pairwise majority, and their pairwise distortion decreases with the number of alternatives m; and (ii) the pairwise rule based on plurality scoring performs worse than pairwise majority, and its pairwise distortion increases with m.
- When considering *worst-case*, *metric* distortion, pairwise majority has distortion 3; when additional alternatives are placed by the election designer *cooperatively*, in the 1-dimensional (1D) Euclidean space, the plurality and Copeland pairwise rules perform just like pairwise majority, while the Borda pairwise rule and one of its variants do much better.

In the full version of this paper, we consider more rules and experiments. It also contains the proof of all our results.

Outline of the paper After discussing related work we give the necessary background and notations. Then, building on the classical literature on distortion, we define pairwise distortion of pairwise voting rules. We start by exploring the average pairwise distortion with experiments on real and synthetic data, then we look into bounds of worst-case pairwise distortion for the 1D-Euclidean metric space. We conclude by discussing further issues.

Related Work

Distortion

Distortion has been introduced by Procaccia and Rosenschein (2006) as a means to evaluate whether it is reasonable to make a collective decision after eliciting only ordinal preferences. Assuming that cardinal preferences are represented by utilities, the social welfare of an alternative is the sum of the utilities it provides to the agents. The distortion of a voting rule f for a given profile is then defined as the ratio between the maximum social welfare of an alternative, and the social welfare of the alternative selected by f; and the distortion of f is the maximum, over all profiles, of the distortion of f for that profile. *Metric* distortion (Anshelevich et al. 2018) aims at minimizing social cost instead of maximizing social welfare: voters and alternatives belong to a metric space, and the cost of an alternative to a voter is the distance between them. Voting distortion (metric or non-metric) has been the topic of a significant number of papers, too many for us to cite them all (and most of them are only moderately related to our concerns). See (Anshelevich et al. 2021) for an extensive survey of the literature until 2021.

Average-case analyses of distortion are far less common than worst-case analyses. For single-winner voting, Boutilier et al. (2015) show that the Borda rule is optimal for the uniform distribution, and Gonczarowski et al. (2023) show that a suitable positional scoring rule (*binomial voting*) performs well for all distributions. Caragiannis et al. (2017) consider average distortion for multi-winner rules, Filos-Ratsikas, Micha, and Voudouris (2019) for district-based elections and Benadè, Procaccia, and Qiao (2019) for social welfare functions.

Independence of Irrelevant Alternatives

The primary reason why Arrow imposed IIA was to prevent the implicit use of interpersonal comparisons (Arrow 1950). However, it also prevents the use of information about *intensities of preferences between two alternatives* revealed by the positions of thesealternatives with respect to *other* ("irrelevant") alternatives. This has been previously discussed(Coakley 2016; Pearce 2021; Maskin 2020; Sen 1970; Osborne 1976; Hillinger 2005; Lehtinen 2011; Maskin 2023), and examples such as the one presented in our introduction highlight its practical negative implications.

Because Arrow's theorem ruled out the existence of a social welfare function under "reasonable" conditions, it has had a negative impact on welfare economics (Fleurbaey and Mongin 2005; Pearce 2021; Igersheim 2019). As other properties stated in Arrow's theorem can hardly be given up, IIA is the most debatable of the conditions of Arrow's theorem, and is actually given up *de facto* when defining voting rules. Still, IIA is considered attractive for several reasons, such as avoiding vote splitting (Maskin 2020).

Pairwise Voting Rules

Let V be a set of n voters and A a set of m alternatives. A ranking $\succ$ of A is a linear order (irreflexive, antisymmetric, transitive and connected relation) of A. $\mathcal{L}(A)$ denotes the set of all rankings over A. A preference profile is a collection of rankings $P = (\succ_1, \ldots, \succ_n)$. For a ranking $\succ_i$, we denote by σ_i the corresponding rank function: for each alternative $x \in A$, $\sigma_i(x) = |\{y \in A \mid y \succ_i x\}| + 1$ the rank of x in $\succ_i$.

A *pairwise (voting) rule* is a function f that, given a preference profile P over A and two alternatives $x, y \in A$, outputs $f(P \mid x, y) \in \{x, y\}$. Equivalently, f associates with every preference profile P a tournament (an irreflexive, antisymmetric and connected relation, but not necessarily transitive).[2] A pairwise rule f satisfies IIA if $f(P \mid x, y) = f(P' \mid x, y)$ for all $P = (\succ_1, \ldots, \succ_n)$ and $P' = (\succ'_1, \ldots, \succ'_n)$ such that for all voters $i \in V$, $x \succ_i y$ if and only if $x \succ'_i y$.

[2] Yet another interpretation: $f(P|., .)$ is the restriction of a deterministic choice function (Sen 1971) on pairs of alternatives.

Among pairwise rules that satisfy IIA, the canonical one is the *pairwise majority* rule: $f_{maj}(P \mid x, y) = x$ (resp. y) if a majority of voters prefer x to y (resp. y to x). In case of a tie, we use a tie-breaking priority relation over alternatives. *We will use such a tie-breaking mechanism more generally for all pairwise rules.* Note that except for the treatment of ties, the graph induced from pairwise majority by $x \to y$ if $f_{maj}(P \mid x, y) = x$ is the majority graph associated with P.

Another prominent family of pairwise rules consists of those that output transitive tournaments, that is, if $f(P \mid x, y) = x$ and $f(P \mid y, z) = y$ then $f(P \mid x, z) = x$. In this case, f corresponds to a social welfare function g mapping every profile P to a ranking $g(P) \in \mathcal{L}(A)$ defined by $x \succ_{g(P)} y$ if and only if $f(P \mid x, y) = x$. Conversely, any social welfare function g induces a pairwise rule $g_{\mathbf{PW}}$.

Among pairwise rules of this class, we will mostly focus on those that are based on a score function Sc that maps every profile P and alternative x to a score $Sc(x, P)$. The pairwise rule f_{Sc} is then defined by $f_{Sc}(P|x, y) = \text{argmax}(Sc(x, P), Sc(y, P))$.

We will make use of the following pairwise rules, all based on some scoring function Sc. We include three positional scoring rules: Plurality because of its simplicity and wide usage, Borda because of its central role in voting and its optimal average distortion in some settings (Caragiannis and Procaccia 2011), Half-approval as it was shown to have a good average (classical) distortion for a large class of distributions (Gonczarowski et al. 2023); and one prominent Condorcet rule (Copeland), which is known to have good metric distortion guarantees.

Positional scoring rules Let $\vec{s} = (s_1, s_2, \ldots, s_m)$ be a non-increasing vector. It is *normalized* if $s_1 = 1$ and $s_m = 0$. Each $x \in A$ gets s_j points from each voter $i \in V$ who ranks x at position j. The score of $x \in A$ for profile P is:

$$Sc(x, P) = \sum_{j=1}^{m} s_j \, |\{i \in V | \sigma_i(x) = j\}|$$

We consider the following pairwise rules: $\mathbf{Borda_{PW}}$ ($\vec{s} = (m-1, m-2, \ldots, 1, 0)$), k-$\mathbf{approval_{PW}}$ ($\vec{s} = (1, \ldots, 1, 0, \ldots, 0)$ with 1 in the first k positions), with as special cases $\mathbf{Plurality_{PW}}$ ($k = 1$) and $\mathbf{Half\text{-}approval_{PW}}$ ($k = \lceil m/2 \rceil$).

$\mathbf{Copeland_{PW}}$ For $x, y \in A$, we say that x *dominates* y if a majority of voters prefer x to y. The Copeland score $Sc(x)$ is the number of alternatives y dominated by x.[3]

Pairwise Distortion

We now formally define *pairwise distortion*. Similarly as for standard distortion, we consider *unconstrained* distortion in which we assume voters gain unconstrained cardinal utilities from alternatives, and *metric* distortion, in which voters and alternatives are embedded in a metric space, and the cost of an alternative for a voter is the distance between them.

[3]The version of Copeland we use here is Copeland^0, where pairwise ties don't give any points.

Unconstrained Pairwise Distortion

In the unconstrained distortion setting, every voter $i \in V$ receives a utility $U_i(x) \in \mathbb{R}_{\geq 0}$ from alternative $x \in A$. A *utility profile* U is a collection $U = (U_i)_{i \in V}$.[4] We say that a preference profile P and a utility profile U are consistent with each other if for all $x, y \in A$ and all voters i, if $U_i(x) > U_i(y)$, then $x \succ_i y$ in P and we denote it $P \approx U$. The *social welfare* of an alternative $x \in A$ is $SW(x) = \sum_{i \in V} U_i(x)$.

The *pairwise distortion* of a pairwise rule f on a utility profile U for two alternatives $x, y \in A$ is the worst-case ratio over all $P \approx U$ between the social welfare of the optimal alternative and that of $f(P \mid x, y)$:

$$\text{dist}(f, U \mid x, y) = \max_{P : P \approx U} \frac{\max(SW_U(x), SW_U(y))}{SW_U(f(P \mid x, y))}$$

Metric Pairwise Distortion

In the *metric* distortion setting, we assume that both voters and alternatives are points in some pseudometric space $(V \cup A, d)$ with $d : (V \cup A)^2 \to \mathbb{R}_{\geq 0}$ a distance function, where $d(i, x)$ represents the cost of alternative x for voter i. The *social cost* of an alternative $x \in A$ for the pseudometric d is $SC_d(x) = \sum_{i \in V} d(i, x)$. As in the general setting, we can naturally induce a preference profile P based on d. We denote $P \approx d$ if for all $x, y \in A$ and all voters i, if $d(x, i) < d(y, i)$ then $x \succ_i y$ in P.

In the metric setting, the pairwise distortion of a pairwise rule f on a metric d for two alternatives $x, y \in A$ is the worst-case ratio over all $P \approx d$ between the social cost of $f(P \mid x, y)$ and that of the optimal alternative:

$$\text{dist}(f, d \mid x, y) = \max_{P : P \approx d} \frac{SC_d(f(P \mid x, y))}{\min(SC_d(x), SC_d(y))}$$

Average and Empirical Pairwise Distortion

We first focus on the average-case scenario, and define the *average pairwise distortion* given a probability distribution $\mathcal{D}$ over utility profiles U (for the unconstrained setting) or over pseudometrics d (for the metric setting). When the distribution is sampled based on a real dataset, we refer to it as *empirical distortion*.

Given a utility profile or a pseudometric, we obtain a pairwise distortion for each pair of alternatives, which we have then to aggregate; for this we consider two possibilities: taking the *maximum* or the *average* over all pairs.

The average pairwise distortion for $\delta \in \{\mathbf{avg}, \mathbf{max}\}$ and $\Delta \in \{U, d\}$ (respectively the unconstrained and the metric settings) is defined as:

$$\text{avg-dist}(f, \delta, \mathcal{D}) = \underset{\Delta \sim \mathcal{D}}{\mathbf{avg}} \; \underset{(x,y) \in A}{\delta} \text{dist}(f, \Delta \mid x, y)$$

Note that we use here two ways of averaging, which should not be confused: averaging over profiles sampled with $\mathcal{D}$ (for defining *average* pairwise distortion), and averaging over pairs of alternatives ($\delta = \mathbf{avg}$).

[4]In the literature, it is often assumed that utilities are normalized (for all $i \in V$, $\sum_{x \in A} U_i(x) = 1$), as otherwise the worst-case distortion is infinitely large. However, for the analysis of average distortion, such an assumption is not required, as we are already restricting ourselves to a probability distribution over utilities.

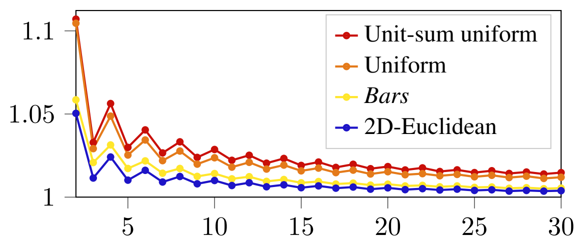

Figure 1: Average pairwise distortion (over 100,000 profiles) for $m = 2$ and several distributions. The x-axis corresponds to the number of voters n.

Distributions

Our study of average pairwise distortion relies on experiments. We first have to choose which probability distributions to use for generating profiles, with the aim of observing the behaviour of the rules when the parameters n and m vary. We choose the following distributions, two standard synthetic ones and one sampled from a real-world dataset.

Uniform In the unconstrained setting, we sample profiles according to the uniform distribution of utilities over $[0, 1]$.

2D-Euclidean uniform In the metric setting, we sample positions of voters and alternatives uniformly at random in the 2-dimensions (2D) Euclidean space.

Bars In the unconstrained setting, we investigate the empirical distortion of the *Bars* dataset (Lesser et al. 2017) from the *Preflib* database (Mattei and Walsh 2013). It contains ratings $r \in \{1, 2, 3, 4, 5\}$ of bars. We interpret these ratings as utilities, adding some small random noise to the ratings in order to remove ties. Then, for a given pair (n, m), profiles are sampled by selecting randomly n voters and m alternatives from the dataset (which originally contains 95 voters and 16 alternatives).

Moreover, for the analysis of the case with $m = 2$ alternatives, we also consider the **unit-sum uniform** distribution in the unconstrained setting: for each voter $i \in V$, the utility of x is selected uniformly at random in $[0, 1]$, and the utility of y is $U_i(y) = 1 - U_i(x)$.

Two Alternatives

In this first experiment, we focus on the case of two alternatives and investigate how the average distortion varies when we increase the number of voters n. This case is of particular interest, because it corresponds to the average distortion of the pairwise majority rule (for which the presence of additional alternatives has no influence). Note that when $m = 2$, the average pairwise distortion for $\delta = \mathbf{avg}$ and $\delta = \mathbf{max}$ are the same as there is only one pair of alternatives.

Figure 1 shows the pairwise distortion for the different distributions. Interestingly, the average pairwise distortion is very close to 1. It reach its highest value for $n = 2$ and the unit-sum uniform distribution, with an average (pairwise) distortion of $5/2 - ln(4) \approx 1.11$. We also observe that pairwise distortion gets asymptotically smaller (with some parity effect due to tie-breaking) when the number of voters n increases. This is not surprising: because of Hoeffding's inequality, the social welfare (or the social cost) of the two alternatives get closer to each other when n increases, thus reducing distortion. These observations suggest that the average distortion of the pairwise majority rule is usually very far from its worst-case distortion (3 in the metric setting and the normalized unconstrained setting)

Increasing the Number of Alternatives

In this section, we investigate how the average pairwise distortion varies with the number of alternatives m. For all experiments, we use profiles of 30 voters and up to 15 alternatives. We compare distortion for pairwise majority (which satisfies IIA) and four transitive pairwise rules: $\text{Plurality}_{\mathbf{PW}}$, $\text{Half-approval}_{\mathbf{PW}}$, $\text{Borda}_{\mathbf{PW}}$, and $\text{Copeland}_{\mathbf{PW}}$.

The first row of Figure 2 shows the average *mean* pairwise distortion ($\delta = \mathbf{avg}$). For the pairwise majority rule, it remains constant with the number of alternatives m, as it satisfies IIA. The transitive pairwise rules considered here behave differently: the average pairwise distortion of $\text{Plurality}_{\mathbf{PW}}$ increases with m, as each voter gives information about only one of the m alternatives (its preferred one). On the opposite, $\text{Borda}_{\mathbf{PW}}$ and $\text{Copeland}_{\mathbf{PW}}$ both seem to really take advantage of the extra information brought by the additional alternatives, as their average pairwise distortion decreases with m (and particularly quickly in the unconstrained setting);[5]; $\text{Half-approval}_{\mathbf{PW}}$ lays in the middle, and its distortion is not far from that of pairwise majority. Note that the average pairwise distortion of pairwise majority varies from one distribution to another. In particular, it is higher for distributions in the unconstrained setting than in the metric setting. This suggests that in the unconstrained setting, using other pairwise rules already have more potential to do better than pairwise majority.

Our conclusion is that using information about additional alternatives can help a lot, provided that the way to use it is carefully chosen, and that $\text{Borda}_{\mathbf{PW}}$ and $\text{Copeland}_{\mathbf{PW}}$ seem both to be good choices.

The second row of Figure 2 shows the variation with m of the average *max* pairwise distortion ($\delta = \mathbf{max}$). It increases with m for all distributions and all pairwise rules, even pairwise majority. It is an intuitively expected behavior: we consider all pairwise distortion values obtained for each pair of alternatives, *and keep only the largest value*. The more alternatives there are, the more pairs to be considered, and the more likely a *bad* pair is sampled. However, the relative order between the curves is the same as for $\delta = \mathbf{avg}$. In particular, $\text{Borda}_{\mathbf{PW}}$ and $\text{Copeland}_{\mathbf{PW}}$ always show a better average max pairwise distortion than pairwise majority, and the gap is wider than for $\delta = \mathbf{avg}$: that is, the average loss of social welfare caused by IIA is even larger if we only look at the worst pair of alternatives in the profile. These results align with our theoretical findings of the next section.

[5]There is an exception however: with the 1D-Euclidean uniform distribution, $\text{Borda}_{\mathbf{PW}}$ has higher average pairwise distortion than pairwise majority. This can be explained by the specific structure of 1D preferences. (Still, average distortion decreases with the number of alternatives.)

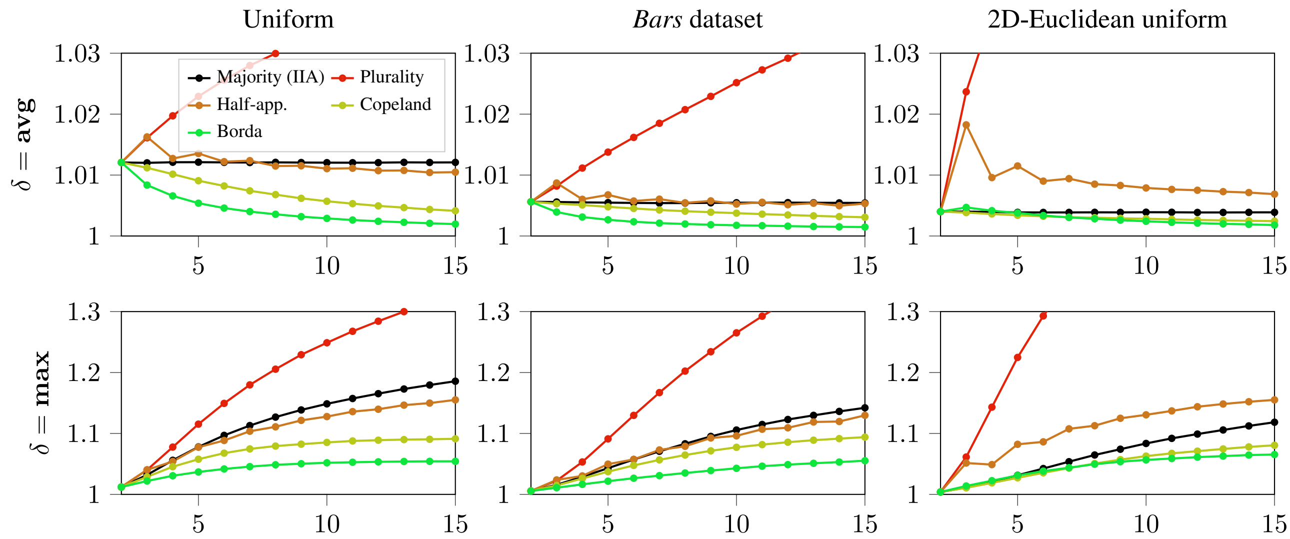

Figure 2: Average *mean* pairwise distortion ($\delta = \mathbf{avg}$, first row) and average *max* pairwise distortion ($\delta = \mathbf{max}$, second row) over 10,000 random profiles. The x-axis corresponds to the number of alternatives m.

Worst-Case Metric Pairwise Distortion

In this section we only consider the metric setting, since the research on standard distortion suggests that it leads to more positive results than the unconstrained setting (see, *e.g.*, (Anshelevich et al. 2021)). We consider *worst-case* pairwise distortion, by assuming that voters are placed in the metric space so as to maximize the pairwise distortion of a specific pair of alternatives (x, y), given the positions of all the alternatives.

A key question is how to choose the positions of the other alternatives when determining the worst-case pairwise distortion of a pair (x, y). Our aim is to compute tight lower and upper bounds of the worst-case distortion for each pairwise rule. This problem can be seen as a game: a first agent selects the positions of the alternatives, and a second agent responds in an adversarial manner by choosing the positions of the voters that maximize pairwise distortion. A *cooperative* (resp. *adversarial*) first agent that places the alternatives so as to minimize (resp. maximize) the worst-case pairwise distortion gives us a lower (resp. upper) bound.[6]

The intuition behind the lower bound with a cooperative agent is that the designer of the game can choose the positions of alternatives so as to maximize the gain of information obtained from voters' rankings, and thus ease preference elicitation. For instance, in a facility location context, if the designer wants to know which of x and y is a better collective choice (perhaps because they are the only possible choices), they can ask the voters to rank x, y, as well as other carefully chosen additional *fake* alternatives, used as reference points. The adversarial case may have less intuitive appeal but is in line with standard worst-case assumptions made when defining distortion in various settings.

It is known that the classical worst-case (pairwise) distortion for $m = 2$ alternatives $\{x, y\}$ is 3, which can be seen on this well-known instance: half of the voters prefer x and the other half y. Assume the tie is broken in favor of x. The y-voters are all located at the same position as y; the x-voters are half-way between x and y. The (pairwise) distortion is $\frac{n/2 \cdot 1 + n/2 \cdot 1/2}{n/2 \cdot 0 + n/2 \cdot 1/2} = 3$. In this instance, adding alternatives looks promising: if we add an alternative exactly between x and y, we will be able to see that the voters who prefer x to y only have a slight preference, while the others have a strong preference for y. Now, how much can we improve pairwise distortion from 3 (obtained with $m = 2$) when we increase the number of alternatives with a cooperative agent? And how bad can it get against an adversarial agent?

Clearly, for any pairwise rule satisfying IIA, pairwise distortion remains constant, independently from the number and positions of other alternatives. In particular, the worst-case pairwise distortion of pairwise majority is always 3.

Without loss of generality, fix two alternatives $x, y \in A$. Denoting $d|_A$ the restriction of a pseudometric d to the set of alternatives A, for $\gamma \in \{\inf, \sup\}$ (respectively the cooperative and adversarial cases) we define

$$\mathrm{dist}(f, \gamma, m) = \underset{d|_A}{\gamma} \sup_{d} \mathrm{dist}(f, d \mid x, y)$$

which we call γ-pairwise distortion for m alternatives.

From now on, *we focus on the case of a 1D-Euclidean space*: voters and alternatives are associated with positions on a line. We make this choice as it is a natural setting to start this study, and it is known to have a lot of practical interpretations, such as facility location; we leave the study of general metric spaces for further research.[7] We denote

[6] The terminology 'cooperative/adversarial' is consistent with our game-theoretic abstract interpretation, but we do not mean that the additional alternatives are chosen strategically. Instead we could say 'optimistic/pessimistic'; this decision-theoretic terminology is not perfect either since it suggests a behaviour towards risk.

[7] Still, Theorems 5 and 6 can be generalized to any metric space.

	IIA (majority$_{\mathbf{PW}}$)	Borda$_{\mathbf{PW}}$	OddBorda$_{\mathbf{PW}}$	k-Approval$_{\mathbf{PW}}$	Plurality$_{\mathbf{PW}}$
inf-pairwise distortion	3	$\frac{m+1}{m-1}$	$\frac{2m-1}{2m-3}$	2	3
sup-pairwise distortion	3	$2m-1$	$4m-5$	∞	∞

Table 1: γ-pairwise distortions of several pairwise voting rules in the 1d-Euclidean metric space.

$p(e) \in \mathbb{R}$ the position of $e \in V \cup A$ on the line. We will assume without loss of generality that the positions of the alternatives x and y with respect to which we study pairwise distortion are $p(x) = 0$ and $p(y) = 1$.

With only two alternatives ($A = \{x, y\}$), all the rules that we consider boil down to the majority rule, whose pairwise distortion is 3. From the well-known fact that the pairwise majority rule outputs a transitive order of alternatives in a 1D-Euclidean space (because of single-peakedness), in this setting Copeland$_{PW}$ is equivalent to pairwise majority. Therefore, its inf-pairwise and sup-pairwise distortions are 3 for any $m \geq 2$ in the 1D-Euclidean space.

In the remaining of this section, we compute the inf and sup-pairwise distortion of various pairwise rules. Table 1 summarize our results.

Lower Bound

We first focus on the cooperative case: alternatives are positioned in order to minimize worst-case pairwise distortion (however, recall that voters will still be positioned in an adversarial way). The question is here whether the lower bound of 3 obtained when we impose IIA can decrease if we use some other pairwise rules.

Without much surprise, for Plurality$_{\mathbf{PW}}$ we cannot obtain a better pairwise distortion than 3, even if we can choose the position of the alternatives.

Theorem 1. *The* inf*-pairwise distortion of Plurality*$_{\mathbf{PW}}$ *in the 1D-Euclidean metric space is* 3 *for any* $m \geq 2$.

For k-approval with $k \notin \{1, m-1\}$, we get a better inf-pairwise distortion, as we can reduce it from 3 to 2.

Theorem 2. *For any* $m \geq 4$ *and* $2 \leq k \leq m-2$*, the* inf*-pairwise distortion of* k*-approval*$_{\mathbf{PW}}$ *in the 1D-Euclidean metric space is* 2.

The result for Borda$_{\mathbf{PW}}$ is even more positive: the inf-pairwise distortion quickly tends to 1 as we add alternatives.

Theorem 3. *The* inf*-pairwise distortion of the Borda*$_{\mathbf{PW}}$ *pairwise rule in the 1D-Euclidean metric space is* $^{m+1}/_{m-1}$ *for any* $m \geq 2$.

To prove this, we first show that the inf-pairwise distortion of most positional scoring rule for given positions of alternatives *between* x and y can be computed easily. In the following proposition, we assume that all alternatives $z_j \in A$ are between x and y. We denote $A = \{z_1, \ldots, z_m\}$ such that $p(z_1) \leq p(z_2) \leq \cdots \leq p(z_{m-1}) \leq p(z_m)$. We naturally have $x = z_1, y = z_m$. For simplicity, the positions of the alternatives are noted $p_j = p(z_j) = d(x, z_j)$ with $p_1 = 0$ and $p_m = 1$.

Lemma 1. *For the 1D-Euclidean metric space, the* inf*-pairwise distortion of a positional scoring rule associated with the normalized scoring vector* $s = (s_1, \ldots, s_m)$ *with* $s_1 = 1$ *and* $s_{m-1} > s_m = 0$*, for fixed positions of the alternatives* $p_1, \ldots, p_m$ *between* x *and* y*, is equal to*

$$\max\left(\max_{1 \leq i,j < m} \frac{K_{i,j}}{1 - K_{i,j}}, \max_{1 \leq i,j < m} \frac{K'_{i,j}}{1 - K'_{i,j}}\right)$$

where:

$$K_{i,j} = \frac{p_{i+1} \cdot s_j + (1 + p_{m+1-j}) \cdot s_i}{2(s_i + s_j)}$$

$$K'_{i,j} = \frac{(1 - p_{m+1-(i+1)}) \cdot s_j + (2 - p_j) \cdot s_i}{2(s_i + s_j)}$$

Note that when the positions of the alternatives are symmetrical (for all j, $p_j = 1 - p_{m+1-j}$), we have $K_{i,j} = K'_{i,j}$.

We now give a proof sketch of Theorem 3.

Proof sketch of Theorem 3. The normalized Borda scoring vector is defined by $s_j = \frac{m-j}{m-1}$ for all j. We first show that there exist positions of the alternatives that achieve distortion $\frac{m+1}{m-1}$. We place alternatives at equal distance from each other, *i.e.*, such that $p_j = \frac{j-1}{m-1}$. Note that these positions are symmetrical (for all j, $p_j = 1 - p_{m+1-j}$), thus, $K'_{i,j} = K_{i,j}$. Using values of the Borda vector and these positions of the alternatives with Lemma 1, we can compute the value of $K_{i,j}$ for all i, j. Consider now the function h defined by $h(i,j) = K_{i,j}$. By straightforward calculations we can show that the partial derivative of h with respect to i is always non-negative, while the one with respect to j is always non-positive. It implies that h reaches its maximum for $i = m-1$ and $j = 1$, which corresponds to the case $K_{m-1,1}$, giving a distortion of $\frac{m+1}{m-1}$.

Let us now prove that we cannot obtain a better worst-case pairwise distortion. For this we consider the following profile: a fraction $\frac{m-1}{m}$ of voters are at distance $\frac{1}{2}$ of x and y and prefer x to y ($x \succ_i y$), and the remaining $\frac{1}{m}$ fraction of voters are at the same position as y, and obviously prefer y to x. Now, let $m_{(1,2]}$ be the number of alternatives $z_j \in A$ with $1 < p(z_j) \leq 2$ (they are closer to y than x is), $m_{>2}$ the number of alternatives with $p(z_j) > 2$ (x is closer to y than they are), and $m_{<0}$ the number of alternatives with $p(z_j) < 0$. The normalized score of x (i.e., its score divided by the number of voters) is $Sc(x) = \frac{1}{m} \cdot \frac{m_{>2}+m_{<0}}{m-1} + \frac{m-1}{m} \cdot \frac{m_{(1,2]}+m_{>2}+m_{<0}+1}{m-1}$ and the one of y is $Sc(y) = \frac{1}{m} \cdot 1 + \frac{m-1}{m} \cdot \frac{m_{(1,2]}+m_{>2}+m_{<0}}{m-1}$. Thus $Sc(x) - Sc(y) = \frac{m_{>2}+m_{<0}}{m(m-1)} \geq 0$, meaning w.l.o.g. that Borda$_{\mathbf{PW}}$ selects x. However, y is the better alternative in this profile, which gives a distortion $\frac{^{m-1}/_m \cdot {}^1/_2 + {}^1/_m \cdot 1}{^{m-1}/_m \cdot {}^1/_2} = \frac{m+1}{m-1}$. This shows that with this profile, any positions of the alternatives leads to a worst-case distortion $\geq \frac{m+1}{m-1}$. $\square$

One may wonder whether $\text{Borda}_{\mathbf{PW}}$ is the optimal rule for $\inf$-pairwise distortion. It appears that a better bound is obtained by a slightly different, yet very similar scoring rule, defined by the scoring vector $\vec{s} = (\ldots, 9, 7, 5, 3, 1, 0)$. We call it *Odd* $\textit{Borda}_{\mathbf{PW}}$, because (except for the last position) it contains consecutive odd numbers. We prove that in the 1D-Euclidean metric space, with m alternatives, the $\text{OddBorda}_{\mathbf{PW}}$ rule has a $\inf$-pairwise distortion of ${}^{2m-1}/_{2m-3}$. We also show that if we add the additional assumption that all alternatives can be placed only between x and y, this bound is tight among all positional scoring rules and achievable only by $\text{OddBorda}_{\mathbf{PW}}$.

Theorem 4. *The* $\inf$*-pairwise distortion of* $\textit{OddBorda}_{\mathbf{PW}}$ *in the 1D-Euclidean metric space is* ${}^{2m-1}/_{2m-3}$ *for any* $m \geq 2$*. This is the lowest distortion among all positional scoring rules if we assume that all alternatives are between* x *and* y*.*

We conjecture that this bound remains tight without the assumption that all alternatives are between x and y. Another question is whether this bound applies to rules besides scoring rules. This can already be partly answered: for given positions of the alternatives between x and y, we can compute a lower bound on the worst-case pairwise distortion, independently of the rule used. We found that for $m < 6$, over $1,000,000$ different positions of the alternatives, the lower bound is always $\geq \frac{2m-1}{2m-3}$, suggesting that $\text{OddBorda}_{\mathbf{PW}}$ is optimal when all alternatives are between x and y. However, when $m \geq 6$, we found that $\text{OddBorda}_{\mathbf{PW}}$ is not the optimal pairwise rule for the 1D-Euclidean metric space.

Upper Bound

We now focus on the pessimistic case, in which both alternatives and voters are placed in an adversarial way.

For $\text{Plurality}_{\mathbf{PW}}$, and more generally k-$\text{approval}_{\mathbf{PW}}$ for any k, this upper bound is infinitely large when $m > 2$. We saw earlier that these rules (especially $\text{Plurality}_{\mathbf{PW}}$) do not use enough information to take advantage of the additional alternatives, we now know that they might actually lose all important information if we add alternatives.

Theorem 5. *The* $\sup$*-pairwise distortion of* $\textit{Plurality}_{\mathbf{PW}}$ *and more generally* k*-*$\textit{approval}_{\mathbf{PW}}$ *for all* $k \leq n$*, in the 1D-Euclidean metric space, is* $+\infty$ *for any* $m \geq 3$*.*

The results are slightly better for $\text{Borda}_{\mathbf{PW}}$ and $\text{OddBorda}_{\mathbf{PW}}$.

Theorem 6. *The* $\sup$*-pairwise distortion of* $\textit{Borda}_{\mathbf{PW}}$ *(resp.* $\textit{OddBorda}_{\mathbf{PW}}$*) in the 1D-Euclidean metric space is equal to* $2m-1$ *(resp.* $4m-5$*) for all* $m \geq 2$*.*

Proof. We show the bound for $\text{Borda}_{\mathbf{PW}}$. We divide voters into two groups: those who prefer x to y, and those who prefer y to x. We assume without loss of generality that $\text{Borda}_{\mathbf{PW}}$ selects y but x has a lower social cost. Observe that every voter who prefers x to y gives at least one more point to x than to y, and every voter who prefers y to x gives at most $m-1$ more points to y than to x. Therefore, if we denote $\alpha \in [0,1]$ the proportion of voters who prefer x to y, we have $Sc(x) \geq n\alpha$ and $Sc(y) \leq n(1-\alpha)(m-1)$. To have $Sc(y) \geq Sc(x)$ (as y is preferred), we need $\alpha \leq (1-\alpha)(m-1)$, which implies $\alpha \leq \frac{m-1}{m}$. Now, observe that voters preferring x to y maximize distortion by being at the same position as x if they could, and voters preferring y to x can maximize distortion by being exactly between x and y (${}^1/_2$ of both). Therefore, distortion cannot be higher than

$$\frac{\alpha + (1-\alpha)\frac{1}{2}}{(1-\alpha)\frac{1}{2}} = \frac{1+\alpha}{1-\alpha} \leq \frac{1+\frac{m-1}{m}}{1-\frac{m-1}{m}} = 2m-1$$

Moreover, this bound is reached. Let ${}^1/_4 > \varepsilon > 0$ and consider the profile in which all alternatives $z_j \neq x, y$ are at position $p(z_j) = \varepsilon$. Set $m-1$ voters at position 2ε, ranking y last and x second last, and one voter at position ${}^1/_2 + \varepsilon$, ranking y first and x last. In this profile $Sc(x) = Sc(y) = m-1$ and assume w.l.o.g. that ties are broken in favor of y. The distortion for this profile is $\frac{(m-1)(1-2\varepsilon)+{}^1/_2-\varepsilon}{{}^1/_2+\varepsilon}$ and it tends to $2m-1$ when ε tends to 0. □

This implies that for these two rules, if distortion quickly gets close to 1 in the best case, it also quickly becomes very large in the worst case. However, by comparing the proofs of the $\inf$-pairwise and $\sup$-pairwise distortions, we notice that the positions of alternatives in the best-case scenario (equidistant alternatives) seems more natural than the positions of alternatives in the worst-case scenario (almost all alternatives at the same position). In particular, the equidistant positions of alternatives are their expected positions if they are uniformly distributed on the line. This uniformity of the positions partly explains why average distortion decreases with the number of alternatives, but also with the dimension. This intuition that pairwise distortion is in average closer to the best case and decreases with the number of alternatives is also supported by the experiments reported in the previous section.

Conclusion

We have introduced pairwise distortion as a tool for the quantitative analysis of the impact on social welfare of the Independence of Irrelevant Alternatives (IIA) axiom. Our conclusions are mixed:

- using information about additional alternatives may help reducing average distortion, but it crucially depends on the choice of the pairwise voting rule used. We found out that — among the rules we studied — the Copeland and Borda pairwise rules are particularly good at decreasing average distortion, but the Plurality pairwise rule has the opposite effect and leads to a larger distortion than sticking to IIA and using pairwise majority.
- when it comes to worst-case distortion, a crucial parameter is the origin of additional alternatives. If they are chosen by the election designer, then the Borda pairwise rule is quite good, and its variant OddBorda (a rule that may be interesting on its own) is even better. However, if they are chosen adversarially, then better stick to IIA.

Among further issues, it is worth looking at average pairwise distortion under distributions in which votes are correlated (such as Mallows or mixtures thereof), and proving or disproving our conjecture about the optimality of OddBorda in the 1D-Euclidean space.

Acknowledgements

Grzegorz Pierczyński was supported by Poland's National Science Center grant no. 2022/45/N/ST6/00271. Théo Delemazure and Jérôme Lang were supported by the French government under management of Agence Nationale de la Recherche as part of the "Investissements d'avenir" program, reference ANR-19-P3IA-0001 (PRAIRIE 3IA Institute). We thank Dominik Peters and the reviewers for helpful comments.

References

Anshelevich, E.; Bhardwaj, O.; Elkind, E.; Postl, J.; and Skowron, P. 2018. Approximating optimal social choice under metric preferences. *Artificial Intelligence*, 264: 27–51.

Anshelevich, E.; Filos-Ratsikas, A.; Shah, N.; and Voudouris, A. A. 2021. Distortion in Social Choice Problems: The First 15 Years and Beyond. In *IJCAI*, 4294–4301. ijcai.org.

Arrow, K. J. 1950. A difficulty in the concept of social welfare. *Journal of political economy*, 58(4): 328–346.

Benadè, G.; Procaccia, A. D.; and Qiao, M. 2019. Low-Distortion Social Welfare Functions. In *AAAI*, 1788–1795. AAAI Press.

Boutilier, C.; Caragiannis, I.; Haber, S.; Lu, T.; Procaccia, A. D.; and Sheffet, O. 2015. Optimal social choice functions: A utilitarian view. *Artificial Intelligence*, 227: 190–213.

Caragiannis, I.; Nath, S.; Procaccia, A. D.; and Shah, N. 2017. Subset Selection via Implicit Utilitarian Voting. *J. Artif. Int. Res.*, 58(1): 123–152.

Caragiannis, I.; and Procaccia, A. D. 2011. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence*, 175(9-10): 1655–1671.

Coakley, M. 2016. Interpersonal comparisons of the good: Epistemic not impossible. *Utilitas*, 28(3): 288–313.

Filos-Ratsikas, A.; Micha, E.; and Voudouris, A. A. 2019. The distortion of distributed voting.

Fleurbaey, M.; and Mongin, P. 2005. The news of the death of welfare economics is greatly exaggerated. *Soc. Choice Welf.*, 25(2-3): 381–418.

Gonczarowski, Y. A.; Kehne, G.; Procaccia, A. D.; Schiffer, B.; and Zhang, S. 2023. The Distortion of Binomial Voting Defies Expectation. *CoRR*, abs/2306.15657.

Hillinger, C. 2005. The Case for Utilitarian Voting. *Homo Oeconomicus*, 23: 295–321.

Igersheim, H. 2019. The Death of Welfare Economics: History of a Controversy. *History of Political Economy*, 51(5): 827–865.

Kizilkaya, F. E.; and Kempe, D. 2023. Plurality Veto: A Simple Voting Rule Achieving Optimal Metric Distortion. arXiv:2206.07098.

Lehtinen, A. 2011. A welfarist critique of social choice theory. *Journal of Theoretical Politics*, 23(3): 359–381.

Lesser, O.; Dery, L.; Kalech, M.; and Elovici, Y. 2017. Group Decision Support for Leisure Activities Using Voting and Social Networks. *Group Decision and Negotiation*, 26.

Maskin, E. 2020. A modified version of Arrow's IIA condition. *Soc. Choice Welf.*, 54(2-3): 203–209.

Maskin, E. 2023. Arrow's Theorem, May's Axiom, Borda's Rule.

Mattei, N.; and Walsh, T. 2013. PrefLib: A Library of Preference Data HTTP://PREFLIB.ORG. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT 2013)*, Lecture Notes in Artificial Intelligence. Springer.

May, K. O. 1952. A set of independent necessary and sufficient conditions for simple majority decisions. *Econometrica*, 20(4): 680–684.

Osborne, D. K. 1976. Irrelevant Alternatives and Social Welfare. *Econometrica*, 44(5): 1001–1015.

Pearce, D. 2021. Individual and social welfare: a bayesian perspective. In *Frisch Lecture delivered to the World Congress of the Econometric Society*.

Procaccia, A. D.; and Rosenschein, J. S. 2006. The distortion of cardinal preferences in voting. In *International Workshop on Cooperative Information Agents*, 317–331. Springer.

Sen, A. 1970. *Collective Choice and Social Welfare*. Harvard University Press.

Sen, A. K. 1971. Choice Functions and Revealed Preference. *The Review of Economic Studies*, 38(3): 307–317.

Examples of Computing Pairwise Distortion

Example 1 (pairwise distortion given a utility profile)**.** *Let the set of alternatives be $A = \{a, b, c\}$. There are three voters whose utilities are depicted below. (Note that utilities are normalized.)*

x	a	b	c	*ranking*
$U_1(x)$	13	6	1	$a \succ b \succ c$
$U_2(x)$	10	1	11	$c \succ a \succ b$
$U_3(x)$	5	8	7	$b \succ c \succ a$
$SW(x)$	28	15	19	

Let us consider all pairs of alternatives. The pairwise majority rule selects a against b, b against c and c against a. Therefore:

- $SW(a) = 28$*;* $SW(b) = 15$*;* $SW(c) = 19$*.*
- $f_{maj}(P \mid a, b) = a$*; dist*$(f, U \mid a, b) = \frac{\max(28,15)}{28} = 1$
- $f_{maj}(P \mid a, c) = c$*; dist*$(f, U \mid a, c) = \frac{\max(28,19)}{19} = \frac{28}{19}$
- $f_{maj}(P \mid b, c) = b$*; dist*$(f, U \mid b, c) = \frac{\max(15,19)}{15} = \frac{19}{15}$

Example 2 (pairwise distortion given a distance profile)**.** *Let the set of alternatives be $A = \{a, b, c\}$, located in the line as follows: $p(a) = 0$; $p(b) = 0.5$; $p(c) : 1$. 45 voters are located at position 0, 30 at position 0.6 and 25 at position 0.8. We have:*

x	a	b	c	*ranking*
$45\times$	0	0.5	1	$a \succ b \succ c$
$30\times$	0.6	0.1	0.4	$b \succ c \succ a$
$25\times$	0.8	0.3	0.2	$c \succ b \succ a$
$SC(x)$	38	33	62	

Therefore:

- $f_{maj}(P \mid a, b) = b$*; dist*$(f, d \mid a, b) = \frac{33}{\min(38,33)} = 1$
- $f_{maj}(P \mid a, c) = c$*; dist*$(f, d \mid a, c) = \frac{62}{\min(38,62)} = \frac{62}{38}$
- $f_{maj}(P \mid b, c) = b$*; dist*$(f, d \mid b, c) = \frac{33}{\min(33,62)} = 1$

Average distortion for $m = 2$

In our experiments with $m = 2$ alternatives, we observed that the highest average distortion was obtained for $n = 2$ voters and the unit-sum uniform distribution of utilities in the unconstrained setting. In this section, we compute the exact value of this distortion.

Proposition 1. *The average (pairwise) distortion of the unit-sum distribution of utilities in the unconstrained setting for $m = 2$ alternatives and $n = 2$ voters is equal to $5/2 - ln(4)$.*

Proof. Let us denote the candidates by x and y. As in the main document, we assume ties are broken by some order on the alternatives. Without loss of generality, we assume ties are broken in favor of x.

We denote the two voters i_1 and i_2 with $U_1(x) = t_1$ and $U_2(x) = t_2$. Thus, $U_1(y) = 1 - t_1$ and $U_2(y) = 1 - t_2$ It is clear that if $t_1 < 1/2$ and $t_2 < 1/2$ or $t_1 > 1/2$ and $t_2 > 1/2$ the two voters vote for the same alternative, and the distortion is 1. Moreover, if the two voters disagree, there is a tie. We assumed that x wins in case of tie. Thus, there are two cases:

- If $t_1 + t_2 \geq 1$ then x is the optimal alternative: the distortion is 1. This corresponds to $t_1 \in [0, 1/2]$ and $t_2 \in [1 - t_1, 1]$ or $t_2 \in [0, 1/2]$ and $t_1 \in [1 - t_2, 1]$.
- If $t_1 + t_2 < 1$ then y is the optimal alternative: the distortion is $\frac{1-t_1+1-t_2}{t_1+t_2}$. This corresponds to $t_1 \in [0, 1/2]$ and $t_2 \in [1/2, 1 - t_1]$ or $t_2 \in [0, 1/2]$ and $t_1 \in [1/2, 1 - t_2]$

We can now compute the average distortion

$$
\begin{aligned}
d^* = &\left(\int_{t_1=0}^{1/2}\int_{t_2=0}^{1/2} 1 dt_1 dt_2\right) + \left(\int_{t_1=1/2}^{1}\int_{t_2=1/2}^{1} 1 dt_1 dt_2\right)\\
&+ \left(\int_{t_1=0}^{1/2}\int_{t_2=1-t_1}^{1} 1 dt_1 dt_2\right) + \left(\int_{t_2=0}^{1/2}\int_{t_1=1-t_2}^{1} 1 dt_1 dt_2\right)\\
&+ \left(\int_{t_1=0}^{1/2}\int_{t_2=1/2}^{1-t_1} \frac{(1-t_2)+(1-t_1)}{t_2+t_1} dt_1 dt_2\right)\\
&+ \left(\int_{t_2=0}^{1/2}\int_{t_1=1/2}^{1-t_2} \frac{(1-t_2)+(1-t_1)}{t_2+t_1} dt_1 dt_2\right)
\end{aligned}
$$

Note that we can group the integrals by two:

$$
\begin{aligned}
d^* =&2\left(\int_{t_1=0}^{1/2}\int_{t_2=0}^{1/2} 1 dt_1 dt_2\right) + 2\left(\int_{t_1=0}^{1/2}\int_{t_2=1-t_1}^{1} 1 dt_1 dt_2\right)\\
&+ 2\left(\int_{t_1=0}^{1/2}\int_{t_2=1/2}^{1-t_1} \frac{(1-t_2)+(1-t_1)}{t_2+t_1} dt_1 dt_2\right)
\end{aligned}
$$

The first integral is equal to:

$$
2\left(\int_{t_1=0}^{1/2}\int_{t_2=0}^{1/2} 1 dt_1 dt_2\right) = 2\left(\frac{1}{4}\right) = \frac{1}{2}
$$

The second integral is equal to:

$$
2\left(\int_{t_1=0}^{1/2}\int_{t_2=1-t_1}^{1} 1 dt_1 dt_2\right) = 2\left(\int_{t_1=0}^{1/2} t_1 dt_1\right) = \frac{1}{4}
$$

We now compute the third integral, which is the only case in which we have a distortion greater than 1. We change variable and define $u = 1 - t_1 - t_2$.

$$
\begin{aligned}
&2\left(\int_{t_1=0}^{1/2}\int_{t_2=1/2}^{1-t_1} \frac{(1-t_2)+(1-t_1)}{t_2+t_1} dt_1 dt_2\right)\\
=&2\left(\int_{t_1=0}^{1/2}\int_{u=1/2-t_1}^{0} -\frac{(1-(1-t_1-u))+(1-t_1)}{(1-t_1-u)+t_1} dt_1 du\right)\\
=&2\left(\int_{t_1=0}^{1/2}\int_{u=0}^{1/2-t_1} \frac{1+u}{1-u} dt_1 du\right)
\end{aligned}
$$

$$=2\left(\int_{t_1=0}^{1/2}\left[-u-2ln(|1-u|)\right]_{u=0}^{1/2-t_1}dt_1\right)$$
$$=2\left(\int_{t_1=0}^{1/2}(-(1/2-t_1)-2ln(1/2+t_1))dt_1\right)$$
$$=2\left[-1/2t_1+1/2t_1^2-2((1/2+t_1)ln(1/2+t_1)-t_1)\right]_0^{1/2}$$
$$=2(-1/4+1/8-2(-1/2)-(-ln(1/2)))$$
$$=2(7/8+ln(2))=7/4-ln(4)$$

If we sum all the integrals, we obtain:

$$1/2+1/4+7/4-ln(4)=5/2-ln(4)$$

□

Additional Pairwise Rules

In this section, we will present the results for worst-case distortion for three additional pairwise rules based on well-known social welfare functions. The experiments for average distortion including these rules are presented in the next section.

Veto$_{\mathbf{PW}}$ The k-approval$_{\mathbf{PW}}$ rule with $k = m-1$ ($\vec{s} = (1,1,\ldots,1,0)$).

Single Transferable Vote (STV$_{\mathbf{PW}}$) In each step, the alternative with the smallest plurality score (the number of voters who rank it first) is eliminated. Then STV$_{\mathbf{PW}}(P \mid x,y)$ returns the alternative among $\{x,y\}$ that was eliminated later.

PluralityVeto$_{\mathbf{PW}}$ A new rule, called PluralityVeto, was introduced in (Kizilkaya and Kempe 2023); is is the simplest known rule whose worst-case metric distortion is 3. It works as follows: every alternative z starts with a score $Sc(z)$ equal to its plurality score. Then, voters are picked sequentially according to a fixed sequence, and every voter decrease the score of her least preferred alternative z that still has a positive score $Sc(z) > 0$ at this point. When an alternative reach $Sc(z) = 0$, it is eliminated. Then PluralityVeto$_{\mathbf{PW}}(P \mid x,y)$ returns the alternative among $\{x,y\}$ that was eliminated later.

The results of inf-pairwise distortion and sup-pairwise distortion for all these rules are actually the same as for Plurality (Theorems 7 and 8 are analogous to Theorems 1 and 5). We first prove a small lemma that will be useful in our proofs.

Lemma 2. *For $n_1, n_2, m_1, m_2 \in \mathbb{N}_{>0}$ we have:*

$$\frac{n_1+n_2}{m_1+m_2} \leq \max(\frac{n_1}{m_1}, \frac{n_2}{m_2})$$

Proof. Let us assume that $n_1/m_1 \geq n_2/m_2$. There exists $\alpha \in \mathbb{R}_{>0}$ such that $m_2 = \alpha m_1$. Because $n_1/m_1 \geq n_2/m_2$, we know that $n_2 \leq \alpha n_1$. Thus,

$$\frac{n_1+n_2}{m_1+m_2} \leq \frac{(1+\alpha)n_1}{1+\alpha)m_1} \leq \frac{n_1}{m_1}$$

□

Theorem 7. *The* inf*-pairwise distortion of Veto$_{\mathbf{PW}}$, STV$_{\mathbf{PW}}$ and PluralityVeto$_{\mathbf{PW}}$ in the 1D-Euclidean metric space is 3 for any $m \geq 2$.*

Proof. Let us assume without loss of generality that alternative x is at position $p(x) = 0$ and alternative y at position $p(y) = 1$. Therefore, $d(x,y) = 1$. We denote the other alternatives by $z_1, \ldots, z_{m-2} \in A$, with positions $p_j = p(z_j)$ and we assume they are placed on the 1D-Euclidean space such that $p(z_1) \leq \cdots \leq p(z_{m-2})$.

We first prove that 3 is a lower bound of the inf-pairwise distortion for these three rules. In the proof for all the rules we consider, we use repeatedly two metric profiles, that is, two sets of n voters together with their positions. For both profiles, the cost of y will tends to three times the cost of x when n tends to $+\infty$. Therefore, if the non-optimal alternative y is selected, then the distortion on the profile will be 3. We will show that for all the pairwise rules we consider there is always at least one of these two profiles in which y is preferred over x (for Veto$_{\mathbf{PW}}$, we use a third profile).

Because of the symmetry between x and y, we assume without loss of generality that for every pairwise rule considered, ties are broken in favor of y.

Profile $P_1 \approx d_1$ This profile is well-known in the literature on distortion. N voters are at the same position as x, i.e., $p(i) = 0$; and $N+1$ voters are equidistant from x and y, that is, $p(i) = 1/2$. Moreover, we consider the profile $P_1 \approx d_1$ in which they prefer y to x. In this profile, if y is preferred over x, the distortion is

$$\frac{1 \cdot N + 1/2 \cdot (N+1)}{1/2 \cdot (N+1)} = \frac{3N+1}{N+1}$$

which tends to 3 as N tends to infinity. This profile is depicted in Figure 3.

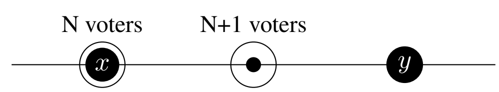

Figure 3: The illustration of profile P_1.

Profile $P_2 \approx d_2$ In this profile, among the $m-2$ additional alternatives, there is an alternative $z \in A$ such that $p(z) \leq 1/2$. We place N voters at position $p(i) = p(z)/2$ (they prefer x to y) and one voter i at the same position as y (that is, $p(i) = 1$). If y is preferred over x, then the distortion is

$$\frac{N(1-p(z)/2)}{N(p(z)/2)+1} \geq \frac{1-1/4}{1/4+1/N} = \frac{3}{1+4/N}$$

which tends to 3 as N tends to infinity. This profile is depicted in Figure 3.

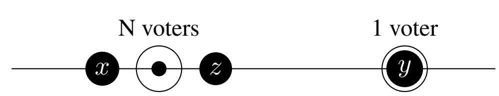

Figure 4: The illustration of profile P_2.

We now prove that 3 is a lower bound for the considered rules.

Veto$_{\mathbf{PW}}$ Let f be Veto$_{\mathbf{PW}}$. Let us first assume that all alternatives are between x and y on the line: $p(x) \le p(z_1) \le \cdots \le p(z_{m-2}) \le p(y)$. For each voter, the last ranked alternative is either x or y, so alternatives other than x and y will never be vetoed by anyone. Therefore in P_1, x get a veto score of $Sc(x) = N$ and y of $Sc(y) = N + 1$ and we have $f(P_1 \mid x, y) = y$.

Let us now assume there exists $z \in A$ such that $p(z) > 1 = p(y)$ on the line. In profile P_1, z is ranked last by all voters, so x and y both get a point from all voters, and obtain a Veto score of $Sc(x) = Sc(y) = 2N+1$, and by tie-breaking $f(P_1 \mid x, y) = y$. In both cases, the selected alternative is y, giving distortion 3.

Let us now assume that for all $z \in A$, we have $p(z) \le p(y) = 1$, but there exists at least one $z \in A$ such that $p(z) < p(x) = 0$. We take the z which is the most on the left on the line (i.e., $z = \mathrm{argmin} p$). Let us consider another profile $P_3 \approx d_3$ in which we have N voters at the same position than y ($p(i) = 1$), who give one point to both x and y as they prefer both of them to z, and one voter at the same position than z ($p(i) = p(z)$), who thus vetoes y. d_3 is depicted in Figure 5.

Figure 5: The illustration of profile P_3.

In this profile, y gets a Veto score of $Sc(y) = N$ and x of $Sc(x) = N + 1$, therefore $f(P_3 \mid x, y) = x$. The distortion is $\frac{N+p(z)/2}{1+p(z)/2}$, which tends to infinity with N. This concludes the proof for Veto$_{\mathbf{PW}}$: we cannot do better than 3 for the worst case distortion.

STV$_{\mathbf{PW}}$ Let f be STV$_{\mathbf{PW}}$. If there are no $z \in A$ such that $p(x) = 0 \le p(z) \le 1 = p(y)$ on the line, then in P_1 every voter puts either x or y first, so all $z \ne x, y$ get eliminated in the first $m - 2$ rounds, and the last round is the majority vote. Thus, y wins and $f(P_1 \mid x, y) = y$.

Let us now assume there is a $z \in A$ such that $p(x) = 0 \le p(z) \le 1 = p(y)$ on the line, and without loss of generality, that $p(z) < 1/2$. In profile $P_2 \approx d_2$, we can assume without loss of generality that all voters at position $p(i) = p(z)/2$ prefer z over x (as we take the worst profile P_2 that is consistent with d_2). Thus, z has a plurality score $Sc(z) = N$, y of $Sc(y) = 1$ and all other alternatives (including x) of $Sc(x) = 0$. The score of x can only increase if z is eliminated, which cannot happen before x is eliminated. Therefore, x is eliminated before y and $f(P_2 \mid x, y) = y$. In both cases, the selected alternative is y, so the worst-case distortion is at least 3.

PluralityVeto$_{\mathbf{PW}}$ Let f be PluralityVeto$_{\mathbf{PW}}$. We use the same proof as for STV. In the first case (there is no z such that $p(x) = 0 \le p(z) \le 1 = p(y)$) with P_1, y is ranked on top in more than half of the votes, therefore it is above x in the output ranking. In the second case with P_2 ($0 \le p(z) \le 1/2$), x is eliminated before y as its plurality score is $Sc(x) = 0$ (so it is eliminated from the beginning), and the plurality score of y is $Sc(y) = 1$. In both cases, $f(P \mid x, y) = y$, so the worst-case distortion is at least 3.

We now prove that 3 is an upper bound on the inf-pairwise distortion, *i.e.*, that the positions of the alternatives can be chosen so that the worst-case pairwise distortion is 3.

For Veto$_{\mathbf{PW}}$, we place all alternatives $z_j \ne x, y$ exactly at the midpoint between x and y: $p(z_j) = 1/2$. In this configuration, veto is equivalent to the majority rule between y and x, as the only thing that matters for each voter is her last ranked alternative, which is always either x or y. Therefore, the worst-case pairwise distortion is 3, as for the majority rule.

For PluralityVeto$_{\mathbf{PW}}$ and STV$_{\mathbf{PW}}$, let us place all the alternatives $z_j \ne x, y$ at position $p(z_j) = -3$ on the line: we have $d(x, z_j) = 3$ and $d(y, z_j) = 4$.

Assume first that all voters in the profile have x and y as their top two alternatives. In that case, all three rules boil down to the majority rule and the distortion is at most 3.

Assume now that at least one voter has one of her top-2 alternatives which is neither x nor y. Because of the positions of the alternatives, this voter is placed at a position $p(i) \le -1$ on the line. We show that the pairwise distortion cannot be more than 3 in that case. Suppose for the sake of contradiction, that the profile contains such nonempty set of voters $V_{out} \subseteq V$ (i.e., there positions are all before -1: $p(i) \le -1$). The situation is depicted in Figure 6.

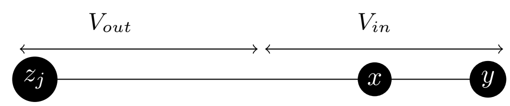

Figure 6: The illustration of the upper bound proof for STV$_{\mathbf{PW}}$ and PluralityVeto$_{\mathbf{PW}}$.

It means that these voters prefer x to y and the distance between them and x is at least 1.

Suppose now that $f(P \mid x, y) = x$, while y has lower social cost than x. In such a case, for all the rules we consider, we would still have $f(P \mid x, y) = x$ if all the voters from V_{out} are moved to the same position as x (i.e., $p(i) = 0$), because it will increase the plurality score of x without increasing the score of y. The distance between each voter $i \in V_{out}$ and x and y will be reduced by the same amount $\sum_{i \in V_{out}} d(i, x)$ and the distortion in such a case will increase, due to the following inequality, which is a direct consequence of Lemma 2.

$$\frac{n_1 + n_2}{n_1 + n_3} \le \frac{n_2}{n_3}, \text{ for all } n_1, n_2, n_3 > 0, n_2 \ge n_3. \qquad (1)$$

with $n_1 = \sum_{i \in V_{out}} p(i)$.

Thus, we can increase distortion by moving all voters from V_{out} to the same position as x, having x and y as their top two alternatives, and we go back to the previous situation with distortion at most 3.

Let us assume now that x has lower social cost than y and $f(P \mid x, y) = y$. All the other alternatives are on the left of x on the line, so they are closer to x than to y. Therefore, under all the rules we consider here, y is preferred to x only if the plurality score of y is larger than the plurality score of x. Indeed, y being the only alternative on the right of x on the line, the only way for y to increase its plurality score for both $\text{STV}_{\mathbf{PW}}$ and $\text{PluralityVeto}_{\mathbf{PW}}$ is that x gets eliminated. Thus, if y is preferred to x, it needs to have a higher plurality score than x from the beginning.

In other words, we want that among the voters from $V_{in} = V \setminus V_{out}$, the majority prefers y to x (because all the other have an alternative $z_j \neq x, y$ as their preferred one). If $i \in V_{out}$ then $p(i) \leq -1$ and $d(i,x) \geq 1$. Let us denote $Y = \sum_{i \in V_{in}} d(i,y), X = \sum_{i \in V_{in}} d(i,x), Z = \sum_{i \in V_{out}} d(i,x) \geq |V_{out}|$. Note that we have $\sum_{i \in V_{out}} d(i,y) = \sum_{i \in V_{out}} (d(i,x) + d(x,y)) = Z + |V_{out}|$. The distortion is equal to:

$$\frac{Y + Z + |V_{out}|}{X + Z}$$

We know that $Y/X \leq 3$ as all voters $i \in V_{in}$ either vote for x or y, which means that we are in the case of a majority vote with only two alternatives. Moreover, $Z+|V_{out}|/Z \leq 2$ as $Z \geq |V_{out}|$.

From Lemma 2, we obtain:

$$\frac{Y + Z + |V_{out}|}{X + Z} \leq \max(\frac{Y}{X}, \frac{Z + |V_{out}|}{Z}) \leq 3$$

Therefore, we cannot get a worst-case pairwise distortion larger than 3 by having voters with position $p(i) \leq -1$ on the line. This proves that with these positions of the alternatives, the worst-case pairwise distortion of these rules is at most 3, which completes the proof. □

Theorem 8. *The* sup*-pairwise distortion of* $\textit{Veto}_{\mathbf{PW}}$*,* $\textit{STV}_{\mathbf{PW}}$ *and* $\textit{PluralityVeto}_{\mathbf{PW}}$ *in the 1D-Euclidean metric space is* $+\infty$ *for any* $m \geq 3$*.*

Proof. The result for $\text{Veto}_{\mathbf{PW}}$ follows directly from Theorem 5. For $\text{STV}_{\mathbf{PW}}$ and $\text{PluralityVeto}_{\mathbf{PW}}$ assume, as before, that x (resp. y) is at position $p(x) = 0$ (resp. $p(y) = 1$). Let $1/4 > \varepsilon > 0$. Set all alternatives $z_j \neq x, y$ at position $p(z_j) = \varepsilon$ and all voters at position $p(i) = 2\varepsilon$. All voters rank y last and x second last, so x and y both have plurality score 0. We assume without loss of generality that because of tie-breaking, $f(P \mid x, y) = y$, thus the pairwise distortion of x and y in this profile is $\frac{1-2\varepsilon}{2\varepsilon}$, which tends to infinity when ε tends to 0. □

Additional Experiments on Average Distortion

In this section, we extend the experiments presented in the main text by considering more distributions for *average pairwise distortion* and more datasets for *empirical distortion*. Besides, we include here the new pairwise rules introduced in the previous section. Finally, in addition to the two ways of aggregating pairwise distortions (with $\delta = \mathbf{max}$ and $\delta = \mathbf{avg}$) presented in the main text, we also consider the third one, denoted by $\delta = \mathbf{win}$. It corresponds to the classical idea of distortion, where we only compare the winning alternative of the voting rule (i.e., the top ranked alternative) to the one with the best social welfare (or social cost). Naturally, this metrics concerns only transitive pairwise rules f (associated to a social welfare function g) because we need to be able to determine the top ranked alternative in $g(P)$ for each profile P (it is denoted by $\hat{g}(P)$), while there is no clear winner under the pairwise majority rule.

More formally, the average distortion for $\delta = \mathbf{win}$ in the unconstrained setting is defined as follows:

$$\begin{aligned}&\text{avg-dist}(f, \delta, \mathcal{D})\\ &= \underset{U \sim \mathcal{D}}{\text{avg}} \max_{y \in A} \max_{P \approx U} \frac{\max(SW_U(\hat{g_f}(P)), SW_U(y)}{SW_U(f(P \mid \hat{g_f}(P), y))}\end{aligned}$$

The definition for the metric setting is analogous.

In the next subsections, we describe our distributions and comment the results for each setting.

Unconstrained Average Distortion

In the unconstrained average distortion setting, we consider the following distributions:

- **Uniform distribution** of utilities in $[0, 1]$.
- **Normalized uniform distribution** of utilities: for all voters i, we draw utilities uniformly at random in $[0, 1]$, and then we normalize utilities: $\sum_{x \in A} U_i(x) = 1$.
- **Impartial Gaussian distribution** of utilities with standard deviation $\sigma = 0.2$ and center of the distribution equal to 0.5 for each alternative x (we also make sure that each utility is ≥ 0).
- **Ordered Gaussian distribution**: this distribution is very similar to the previous one, but instead of having a utility 0.5 as the center of the distribution for all alternatives, the center of each alternative is selected uniformly at random in $[0, 1]$, so each alternative has a different "true" value. In that sense, it is close to a Mallows model, with a central ranking (the one induced by the center of the distribution of each alternative).

Figure 7 shows the results for these distributions. The first observation is that there is no major difference between the uniform and the normalized uniform distributions. The impartial Gaussian is also very close to the uniform distribution, but its distortion is closer to 1 for all rules. This makes sense: with a Gaussian distribution, the gap between the social welfare of two alternative is more likely to be smaller than for the uniform distribution. However, the ordered Gaussian distribution is quite different from the others. Interestingly, with $\delta = \mathbf{avg}$ and $\delta = \mathbf{win}$, we obtain very low distortions for the "good" rules, but it is not the case for $\delta = \mathbf{max}$. Moreover, for this distribution, Plurality, Veto, STV and PluralityVeto all perform even worse than for the other distributions.

Generally, we can divide the rules into 3 groups:

- *Majoritarian rules* (Plurality, Veto, STV, PluralityVeto) for which the average pairwise distortion increases (and even faster when $\delta = \mathbf{max}$). Note that PluralityVeto is the best of this family, and Veto is the worst, except for the ordered Gaussian distribution, in which Veto is the best of this family.

- *Half-approval*, which generally has a distortion very close to the pairwise majority rule, except with the ordered Gaussian distribution, for which the distortion of Half-Approval is clearly higher than the one of Majority (but still lower than the one of the rules of the first group).
- *Copeland and Borda*, for which the average pairwise distortion decreases for $\delta = \mathbf{avg}$ or increases more slowly than the one of pairwise majority for $\delta = \mathbf{max}$. Among the two, Borda is better than Copeland.

Note that we observe the same relative order of the rules for $\delta = \mathbf{win}$, which correspond more to the idea of the *classical* voting distortion than the two other metrics. However, we cannot compare the social welfare functions to pairwise majority here, because we need the rule to output a clear winner (which might not be the case of pairwise majority). It is very interesting to note that even in this setting, the average distortions of Borda and Copeland decrease when the number of alternatives increases, which suggest that even if we focus on the standard distortion setting (i.e., where we only care about the winner), we observe that Borda takes advantage of the extra information that the additional alternatives bring.

Empirical distortion

For the empirical distortion, we used the *Bars* and *Restaurants* datasets (Lesser et al. 2017) from the *Preflib* database (Mattei and Walsh 2013). They contain ratings $r \in \{1,2,3,4,5\}$ of bars and restaurants. We interpret these ratings as utilities, adding some random noise to the ratings in order to remove ties. Then, for a given pair (m, n), profiles are sampled by selecting randomly n voters and m alternatives from the dataset (the *Bars* dataset originally contains 95 voters and 16 alternatives, and the *Restaurant* dataset 93 voters and 23 alternatives).

Figure 8 shows the results for these two datasets. We will not comment much on them, as they are very close to the results for the distributions of the unconstrained setting. The results are not exactly the same: for instance, the distortion of the pairwise majority rule is far smaller for empirical distortion than with the uniform distribution (note that the scale of the y-axis is different). Because of this, Borda and Copeland cannot improve that much, which explains why the gap between these rules and the pairwise majority rule is less impressive for empirical distortion (but still important). Note that the distortion of the pairwise majority rule is smaller here because voters in these datasets tend to have similar opinions on the restaurants and bars (their opinions are correlated), so it is easier to select the "correct" alternative for each pair.

Metric average distortion

In this subsection, we consider all d-Euclidean spaces with all combinations of the following parameters:

- **Dimension:** $d \in \{1,2,3,10\}$ (we add 10 to have an idea of what happen with a large dimension)
- **Distribution:** We used the uniform distribution in $[0,1]^d$ and the Gaussian distribution centered in $(0.5,\dots,0.5)$ with standard deviation $\sigma = 0.5$ in all dimensions.

Figures 9 and 10 show the results for these distributions. We first notice that the uniform and Gaussian distributions seem to give very similar results, so the only important parameter is the number of dimensions d. Another observation is that in comparison to the unconstrained setting and the empirical distortion, the average pairwise distortion of pairwise majority is here quite low (around 1.002, while it is larger than 1.01 in the unconstrained setting). And the higher the dimension, the lower the average pairwise distortion. For this reason, it is hard to reduce even more the average distortion for the other pairwise rules.

For one dimension ($d = 1$), we see a spike in average pairwise distortion for Borda (and $\delta = \mathbf{avg}$) at $m = 3$, then distortion decreases down to its original level. To understand why this is the case, observe that if the third alternative z is far closer to x than y, it will "steals" a lot of voters from x. And the smaller the number of dimensions, the higher the probability that the third alternative is far closer to one alternative than the other (because for more dimensions, z might be close to x in one dimension, but at the same time far from it in another). This is why Borda performs so bad for all δ when $d = 1$.

We observe the same effect, but heavily reduced for $d = 2$. In particular, Borda is now the best rule for $\delta = \mathbf{avg}$ and $\delta = \mathbf{max}$. For $d \geq 3$, the effect almost disappeared.

Finally, we also notice that the relative order of the "bad" rules is changed in the metric setting: the worst rules are now Plurality and STV, followed by Veto (which is the worst one in the unconstrained setting) and then by PluralityVeto.

Omitted Proofs

In this section, we provide the complete proofs of the missing or drafted proofs from the section on worst-case metric pairwise distortion in the main text.

Lower Bound

Theorem 1. *The* $\inf$*-pairwise distortion of Plurality*$_{\mathbf{PW}}$ *in the 1D-Euclidean metric space is* 3 *for any* $m \geq 2$*.*

Proof. The proof is analogous to the proofs of STV and PluralityVeto from Theorem 7. ☐

Theorem 2. *For any* $m \geq 4$ *and* $2 \leq k \leq m-2$*, the* $\inf$*-pairwise distortion of* k*-approval*$_{\mathbf{PW}}$ *in the 1D-Euclidean metric space is* 2*.*

Proof. Let f be k-approval$_{\mathbf{PW}}$. We have $2 \leq k \leq m-2$ In this proof, we assume without loss of generality that $f(P \mid x, y) = x$ and that y has lower social cost. We also assume for now that all other alternatives $z_j \neq x, y$ are between x and y on the line, and we denote $A = \{z_1, \dots, z_m\}$ such that $x = z_1$, $y = z_m$ and $0 = p(x) = p(z_1) \leq p(z_2) \leq \dots \leq p(z_{m-1}) \leq p(z_m) = p(y) = 1$. For simplicity, we write $p_j = p(z_j)$.

Because there are only two possible scores in the scoring vector (0 and 1) and at least one of x and y has score of 0 (because one of them is ranked last), there are only two profiles (or any mixture of them) that could maximize distortion (with x preferred to y and y with lower social cost):

- In the first profile $P_1 \approx d_1$, all voters give a score 0 to both x and y. Distortion is maximized if all voters are at position $p(i) = d(i,x) = 1 - \frac{1-p_{(m+1)-(k+1)}}{2}$, as they rank x last and y at position $(k+1)$, so both will get a k-approval score of 0. This corresponds to the following profile (green alternatives are the approved one by the n voters):
  The distortion is then
$$D_1 = \frac{1 - \frac{1-p_{(m+1)-(k+1)}}{2}}{\frac{1-p_{(m+1)-(k+1)}}{2}} = \frac{1+p_{m-k}}{1-p_{m-k}}$$
  If we switch the roles of x and y (x and y being still at positions 0 and 1 on the line but now x has a lower social cost and y is preferred), the distortion is then
$$D'_1 = \frac{1 - \frac{p_{k+1}}{2}}{\frac{p_{k+1}}{2}} = \frac{2-p_{k+1}}{p_{k+1}}$$
- In the second profile $P_2 \approx d_2$, half of the voters give a point to x, and half of the voters give a point to y. The former are placed at the farthest position from x in which they still give it a point, which is $p(i) = \frac{p_{k+1}}{2}$ (in which case x is ranked at position k), and the latter are placed at the same position as y, such that $p(i) = 1$ and $d(i,y) = 0$. This corresponds to the following profile:
  The pairwise distortion is then
$$D_2 = \frac{\frac{1}{2}\cdot 1 + \frac{1}{2}\cdot\frac{p_{k+1}}{2}}{\frac{1}{2}\cdot(1-\frac{p_{k+1}}{2})} = \frac{2+p_{k+1}}{2-p_{k+1}}$$
  Again, if we switch the roles of x and y, the pairwise distortion is
$$D'_2 = \frac{\frac{1}{2}\cdot 1 + \frac{1}{2}\cdot\frac{(1-p_{m+1-(k+1)})}{2}}{\frac{1}{2}\cdot(1-\frac{1-p_{m+1-(k+1)}}{2})} = \frac{3-p_{m-k}}{1+p_{m-k}}$$

We do not need to consider mixtures of the two types of profiles, as it cannot increase distortion, by Lemma 2.

We want the positions of alternatives to minimize $\max(D_1, D'_1, D_2, D'_2)$. We observe that only the values of p_{m-k} and p_{k+1} actually matter. Moreover, we can minimize $\max(D_1, D'_2)$ and $\max(D'_1, D_2)$ independently, as they do not have common parameters. When p_{m-k} increases, D_1 increases and D'_2 decreases. Thus, to minimize $\max(D_1, D'_2)$, we need:
$$\frac{1+p_{m-k}}{1-p_{m-k}} = \frac{3-p_{m-k}}{1+p_{m-k}}$$
and simple calculations give $p_{m-k} = 1/3$.

Similarly, to minimize $\max(D'_1, D_2)$, we need:
$$\frac{2-p_{k+1}}{p_{k+1}} = \frac{2+p_{k+1}}{2-p_{k+1}}$$
and simple calculations give $p_{k+1} = 2/3$. Thus, we can minimize $\max(D_1, D'_1, D_2, D'_2)$ if $p_{m-k} = 1/3$ and $p_{k+1} = 2/3$. This is possible only if $m - k < k + 1$, i.e, if $k > m-1/2$. With these values for p_{m-k} and p_{k+1}, we have $D_1 = D'_1 = D_2 = D'_2 = 2$, which is the worst-case pairwise distortion. We cannot reduce this bound by putting some other alternatives outside the segment $[0,1]$, as one of D_1, D'_1, D_2 and D'_2 will always be at least 2.

Now, let us focus on the case where $k \leq m-1/2$. As $p_{k+1} \leq p_{m-k}$, we cannot directly use the proof above. However, we can artificially 'exclude' some alternatives by placing them far away from x and y. Indeed, let us assume that $k \leq m-1/2$, and place $m-k-2$ alternatives far away from both x and y (for instance at position $p(z) = d(x,z) = 10$ on the line). This way, there are $m' = k+2$ alternatives between x and y (including them), and all the excluded alternatives will get a score of 0 from all voters positioned between x and y. In that sense, we fall back to the case described above, but for m' alternatives. We then place the alternatives such that $p_{k+1} = 2/3$ and $p_{m'-k} = p_2 = 1/3$, which is possible as $m' - k < k + 1$. If some voters give a point to one of the excluded alternatives, then they either give no point to both x and y, or they give one point to y. In both cases, the pairwise distortion would be reduced if these voters moved towards x and y and take one of the positions described in the first part of the proof. Thus, all voters have to be between x and y to reduce distortion, and this bring us back to the case explained in the first part of the proof. Therefore, we also have positions of the alternatives such that the worst-case pairwise distortion is 2 when $k \leq m-1/2$. Again, this bound cannot be reduced by having other positions of the alternatives as $\max(D_1, D'_1, D_2, D'_2) \geq 2$ in any case. □

Lemma 1. *For the 1D-Euclidean metric space, the* inf-*pairwise distortion of a positional scoring rule associated with the normalized scoring vector $s = (s_1, \ldots, s_m)$ with $s_1 = 1$ and $s_{m-1} > s_m = 0$, for fixed positions of the alternatives $p_1, \ldots, p_m$ between x and y, is equal to*
$$\max\left(\max_{1\leq i,j<m}\frac{K_{i,j}}{1-K_{i,j}}, \max_{1\leq i,j<m}\frac{K'_{i,j}}{1-K'_{i,j}}\right)$$
where:
$$K_{i,j} = \frac{p_{i+1}\cdot s_j + (1+p_{m+1-j})\cdot s_i}{2(s_i+s_j)}$$
$$K'_{i,j} = \frac{(1-p_{m+1-(i+1)})\cdot s_j + (2-p_j)\cdot s_i}{2(s_i+s_j)}$$

Proof. We assume for now that all other alternatives $z_j \neq x, y$ are between x and y on the line, and we denote $A = \{z_1, \ldots, z_m\}$ such that $x = z_1$ and $y = z_m$ and $0 = p(x) = p(z_1) \leq p(z_2) \leq \cdots \leq p(z_{m-1}) \leq p(z_m) = p(y) = 1$. For simplicity, we write $p_j = p(z_j)$. Let f be a positional scoring rule with normalized vector $\vec{s} = (s_1, \ldots, s_m)$.

We also assume that $f(P \mid x, y) = x$, and that y has lower social cost. We denote $\alpha^x_j \in [0,1]$ the proportion of voters giving rank j to x. α^y_j is defined similarly for y. Since all alternatives are between x and y, every voter places either x or y in last position. We have three conditions on the vectors $\alpha^x = (\alpha^x_1, \ldots, \alpha^x_{m-1})$ and $\alpha^y = (\alpha^y_1, \ldots, \alpha^y_{m-1})$, the first one is that the score of x is larger than the score of y (since

$f(P \mid x, y) = x$):

$$Sc(x) = \sum_{j=1}^{m-1} \alpha_j^x s_j \geq \sum_{j=1}^{m-1} \alpha_j^y s_j = Sc(y) \qquad (2)$$

Since we want to maximize distortion, we can consider that the inequality in this equation is an equality, because lowering the score of x can only increase its social cost. The second condition is that the total number of voters is equal to n, with proportions summing to 1. Since every alternative is between x and y, every voter either has x or y at the last position. Thus, each voter is counter in exactly one α_j^x or one α_j^y, and we have:

$$\sum_{j=1}^{m-1} \alpha_j^x + \sum_{j=1}^{m-1} \alpha_j^y = 1 \qquad (3)$$

Finally, the last condition is that every proportion should be non-negative.

$$\forall j, \alpha_j^x \geq 0 \text{ and } \alpha_j^y \geq 0 \qquad (4)$$

Note that for all voters included in the group α_j^x (*i.e.*, giving rank j to x and m to y), there exists a unique position that will maximize distortion. This position is the closest possible to y (with lower social cost), and the farthest possible from x. This position is $p(i) = d(i, x) = \frac{p_{j+1}}{2}$. If the position is closer to x, the distortion decreases, and if the position is further from x, then the voters do not rank x at position j anymore. Analogously, voters in α_j^y will be at position $1 - \frac{1-p_{m+1-j}}{2}$ to maximize distortion. For instance, voters who place x first should be at position $\frac{p_2}{2}$ and voters who place y first at position 1 (i.e., the same position as y).

Therefore, the distortion is equal to

$$\begin{aligned}
&\frac{\sum_{j=1}^{m-1} \frac{p_{j+1}}{2}\alpha_j^x + \sum_{j=1}^{m-1}\left(1 - \frac{1-p_{m+1-j}}{2}\right)\alpha_j^y}{\sum_{j=1}^{m-1}\left(1 - \frac{p_{j+1}}{2}\right)\alpha_j^x + \sum_{j=1}^{m-1} \frac{1-p_{m+1-j}}{2}\alpha_j^y} \\
=&\frac{\sum_{j=1}^{m-1} \frac{p_{j+1}}{2}\alpha_j^x + \sum_{j=1}^{m-1}\left(1 - \frac{1-p_{m+1-j}}{2}\right)\alpha_j^y}{1 - \left(\sum_{j=1}^{m-1} \frac{p_{j+1}}{2}\alpha_j^x + \sum_{j=1}^{m-1}\left(1 - \frac{1-p_{m+1-j}}{2}\right)\alpha_j^y\right)} \\
=&\frac{K}{1-K}
\end{aligned}$$

with

$$K = \sum_{j=1}^{m-1} \frac{p_{j+1}}{2}\alpha_j^x + \sum_{j=1}^{m-1}\left(1 - \frac{1-p_{m+1-j}}{2}\right)\alpha_j^y$$

Therefore, maximizing distortion is equivalent to maximizing $K \in [0, 1]$. Note that we also need to maximize K' which is the equivalent to K when we switch the roles of x and y.

We will show that K is maximal when there exists exactly one rank $j \leq m-1$ such that $\alpha_j^x > 0$, and exactly one rank $j \leq m-1$ such that $\alpha_j^y > 0$, and all other proportions are equal to 0. Obviously, at least one of the α_j^x is positive, otherwise x would not be preferred (since $s_{m-1} > 0$). Similarly, at least one of the α_j^y is positive, otherwise y would not have a lower social cost than x.

We assume by contradiction that there exist $j, j', k \in [1, m-1]$ such that $\alpha_j^x > 0$, $\alpha_{j'}^x > 0$ and $\alpha_k^y > 0$. We will prove that we can change these proportions (and only these proportions) such that the pairwise distortion does not decrease, and one of these proportions is equal to 0. We assume that everything else is fixed, and we have in this profile $\alpha_j^x > 0$, $\alpha_{j'}^x > 0$ and $\alpha_k^y > 0$. We assume that $s_j \neq s_{j'}$, otherwise we can put all voters of α_j^x and $\alpha_{j'}^x$ at the same position (the one further from x) without decreasing distortion, and it would make either α_j^x or $\alpha_{j'}^x$ equal to 0. We will now let these three proportions vary. We use the three variables β_j^x, $\beta_{j'}^x$ and β_k^y to denote the varying proportion respectively corresponding to voters ranking x in position j, x in position j', and y in position k, and we will show that the distortion is not lower when one of these three variables is equal to 0 than when $\beta_j^x = \alpha_j^x$, $\beta_{j'}^x = \alpha_{j'}^x$ and $\beta_k^y = \alpha_k^y$.

We have the following conditions, because of conditions (3) and (2) (the sums of proportions and the scores remain equal):

$$\beta_j^x + \beta_{j'}^x + \beta_k^y = \alpha_j^x + \alpha_{j'}^x + \alpha_k^y$$

$$s_j\beta_j^x + s_{j'}\beta_{j'}^x + s_k\beta_k^y = s_j\alpha_j^x + s_{j'}\alpha_{j'}^x + s_j\alpha_k^y$$

By combining these two conditions, we obtain that all three proportions are linearly dependent. Indeed, if we multiply the first equation by $s_{j'}$ and combine it with the second equation, we get

$$\beta_j^x = \beta_k^y \frac{s_{j'} - s_k}{s_j - s_{j'}} + \frac{\alpha_j^x(s_j - s_{j'}) + \alpha_y^k(s_k - s_{j'})}{s_j - s_{j'}}$$

which is properly defined since $s_j \neq s_{j'}$. Now, $\beta_{j'}^x = \alpha_j^x + \alpha_{j'}^x + \alpha_k^y - \beta_j^x - \beta_k^y$, so these three values are linearly dependent. We denote $a_1, a_2, b_1, b_2 \in \mathbb{R}$ such that $\beta_{j'}^x = a_1\beta_j^x + b_1$ and $\beta_k^y = a_2\beta_j^x + b_2$. Thus, K is a linear function of β_j^x, which means that it reaches its maximum when β_j^x reaches its minimal or maximal value. In both cases, this implies that the maximum pairwise distortion is reached when at least one of β_j^x, $\beta_{j'}^x$ and β_k^y is equal to 0.

We can show the same way that when $j, k, k' \in [1, m-1]$ such that $\alpha_j^x > 0$, $\alpha_k^y > 0$ and $\alpha_{k'}^y > 0$, we can change one of these proportions to be 0 without decreasing the pairwise distortion.

As a consequence, there cannot exist more than one $i \leq m-1$ and one $j \leq m-1$ such that $\alpha_i^x > 0$ and $\alpha_j^y > 0$. Thus, if we denote i and j these two indices with positive proportions, we have $\alpha_j^y = 1 - \alpha_i^x$ and we can write $\alpha_i^x s_i = \alpha_j^y s_j = (1 - \alpha_i^x)s_j$. This directly gives us $\alpha_i^x = \frac{s_j}{s_j + s_i}$ and $\alpha_j^y = \frac{s_i}{s_i + s_j}$. By considering all possible pairs $i, j \in [m-1]$ of indices, we obtain the following formula for K:

$$
\begin{aligned}
K &= \max_{i,j} \frac{d_{i+1}}{2}\frac{s_j}{s_i+s_j} + \left(1 - \frac{1-d_{m+1-j}}{2}\right)\frac{s_i}{s_i+s_j} \\
&= \max_{i,j} \frac{d_{i+1}s_j + (1+d_{m+1-j})s_i}{2(s_i+s_j)} \\
&= \max_{i,j} K_{i,j}
\end{aligned}
$$

We define the formula of K' in a similar way by switching the roles of x and y (i.e., $z'_j = z_{m+1-j}$). This gives us the formula stated in the proposition and completes the proof. □

Theorem 3. *The* inf*-pairwise distortion of the Borda*$_{\mathbf{PW}}$ *pairwise rule in the 1D-Euclidean metric space is* $^{m+1}/_{m-1}$ *for any* $m \geq 2$*.*

Proof. Note that the normalized Borda scoring vector is defined by $s_j = \frac{m-j}{m-1}$ for all j. We first show that there exists positions of the alternatives $p_1, \ldots, p_m$ that achieves a distortion $^{m+1}/_{m-1}$. Let us place alternatives at equal distance from each other, such that $p_j = p(z_j) = \frac{j-1}{m-1}$. Note that these positions are symmetrical (for all $j, p_j = 1 - p_{m+1-j}$), thus, $K'_{i,j} = K_{i,j}$ for all i, j. Using values of the Borda vector and these positions of the alternatives with Lemma 1, we obtain

$$
\begin{aligned}
K'_{i,j} = K_{i,j} &= \frac{\frac{i}{m-1}\frac{m-j}{m-1} + (1 + \frac{m+1-(j+1)}{m-1})\frac{m-i}{m-1}}{2(\frac{m-i}{m-1} + \frac{m-j}{m-1})} \\
&= \frac{m(2m-1) - i(m-1) - jm}{2(m-1)(2m-i-j)}
\end{aligned}
$$

Define the function h by $h(i,j) = K_{i,j}$. We show that the partial derivative of h with respect to i is always non-negative, while the one with respect to j is always non-positive. We have that:

$$
\begin{aligned}
\frac{\partial h(i,j)}{\partial i} =& \beta(-(2(m-1)(2m-i-j)) \\
&+ 2\,(m(2m-1) - i(m-1) - jm))
\end{aligned}
$$

with $\beta = \frac{(m-1)}{(2(m-1)(2m-i-j))^2}$, which is always positive, so it does not have any impact on the sign of $\frac{\partial h(i,j)}{\partial i}$.

$$
\begin{aligned}
\frac{\partial h(i,j)}{\partial i} =& \beta(-2(m-1)(2m) + 2i(m-1) + 2j(m-1) \\
&+ 2m(2m-1) - 2i(m-1) - 2jm) \\
=& \beta(-4(m-1)m - 2j + 2m(2m-1)) \\
=& \beta\,(2m-2j) \\
\geq& 0 \text{ because } j \leq m
\end{aligned}
$$

Similarly:

$$
\begin{aligned}
\frac{\partial h(i,j)}{\partial j} =& \beta(-2m(2m-i-j) \\
&+ 2\,(m(2m-1) - i(m-1) - jm)) \\
=& \frac{\partial h(i,j)}{\partial i} - \beta 2(2m-i-j) \\
=& \beta((2m-2j) - 2(2m-i-j)) \\
=& \beta(2i-2m) \\
\leq& 0 \text{ because } i \leq m
\end{aligned}
$$

It implies that h reaches its maximum for $i = m-1$ and $j = 1$, which corresponds to the case $K_{m-1,1}$, giving a distortion of $^{m+1}/_{m-1}$.

Let us now prove that we cannot obtain a better worst-case pairwise distortion. For this we consider the profile of the case $K_{m-1,1}$. In this profile, a fraction $^{m-1}/_{m}$ of voters are at distance $p(i) = {}^1/_2$ of x and y and ranking x before y (*i.e.*, $x \succ_i y$), and the remaining $^1/_m$ fraction of voters are at the same position as y ($p(i) = 1$), ranking it first. Now, let $m_{(1,2]}$ be the number of alternatives with position in $(1,2]$ (they are closer to y than x is), $m_{>2}$ the number of alternatives with position > 2 (x is closer to y than they are), and $m_{<0}$ the number of alternatives with position < 0. The situation is as follows (the number below the empty circle represent the proportion of voters at one given position):

The normalized score of x (*i.e.*, the score divided by the number of voters) is:

$$
Sc(x) = \frac{1}{m}\cdot\frac{p_{>2} + p_{<0}}{m-1} + \frac{m-1}{m}\cdot\frac{m_{(1,2]} + m_{>2} + m_{<0} + 1}{m-1}
$$

and the one of y is:

$$
Sc(y) = \frac{1}{m}\cdot 1 + \frac{m-1}{m}\cdot\frac{m_{(1,2]} + m_{>2} + m_{<0}}{m-1}
$$

Thus:

$$
Sc(x) - Sc(y) = \frac{m_{>2} + m_{<0}}{m(m-1)} \geq 0
$$

This means without loss of generality that x if the preferred alternative between x and y for the Borda rule. However, y is the better alternative in this profile, which gives a distortion of:

$$
\frac{\frac{m-1}{m}\cdot\frac{1}{2} + \frac{1}{m}\cdot 1}{\frac{m-1}{m}\cdot\frac{1}{2}} = \frac{m+1}{m-1}
$$

This prove that with this profile of voters, any positions of the alternatives lead to a worst-case distortion greater than $^{m+1}/_{m-1}$. □

Theorem 4. *The* inf*-pairwise distortion of OddBorda*$_{\mathbf{PW}}$ *in the 1D-Euclidean metric space is* $^{2m-1}/_{2m-3}$ *for any* $m \geq 2$*. This is the lowest distortion among all positional scoring rules if we assume that all alternatives are between* x *and* y*.*

Proof. We first show the second part of the theorem. Thus, we assume that every alternative is between x and y, so we can use Lemma 1. Analogously to the proof of Theorem 3, we denote $A = \{z_1, \dots, z_m\}$ with $x = z_1$, $y = z_m$ and $p(z_1) \leq \dots \leq p(z_m)$. For simplicity, we denote $p_j = p(z_j)$ the position of z_j.

We want to show that $\max_{i,j} \frac{K_{i,j}}{1-K_{i,j}} \leq \frac{2m-1}{2m-3}$ so we need that for all $i, j \in [1, m-1]$, $K_{i,j} \leq \frac{2m-1}{4(m-1)}$.

We first show by induction that for all j, the distortion can be lower or equal to $\frac{2m-1}{2m-3}$ if and only if $p_j = \frac{j-1}{m-1}$. The induction hypothesis is the following: for all $j \in [1, m]$, we have $p_j \leq \frac{j-1}{m-1}$ and $p_{(m+1)-j} \geq \frac{(m+1-j)-1}{m-1}$. Combining the two will give us that for all $j \in [1, m]$, $p_j = \frac{j-1}{m-1}$.

It is true for $j = 1$, as $z_1 = x$ and $z_m = y$ so $p_1 = 0$ and $p_m = 1$. Assume it is true for $j \geq 1$, we will show that it is true for $j+1$. We need $K_{j,j} \leq \frac{2m-1}{4(m-1)}$:

$$\frac{p_{j+1}s_j + (1 + p_{m+1-j})s_j}{4s_j} \leq \frac{2m-1}{4(m-1)}$$
$$p_{j+1} + p_{m+1-j} + 1 \leq \frac{2m-1}{m-1}$$
$$p_{j+1} \leq \frac{m}{m-1} - p_{m+1-j}$$

But we know by our induction hypothesis that $p_{m+1-j} \geq \frac{m-j}{m-1}$. Thus

$$p_{j+1} \leq \frac{m}{m-1} - \frac{m-j}{m-1} \leq \frac{j}{m-1}$$

By switching the roles of x and y, we get $1 - p_{m+1-(j+1)} = d(y, z_{m+1-(j+1)}) \leq \frac{j}{m-1}$, and thus

$$p_{m+1-(j+1)} \geq \frac{m-(j+1)}{m-1}$$

This proves the induction hypothesis, and this means that for all j, we have a tight bound for p_j: the distortion $\frac{2m-1}{2m-3}$ can only be reached if for all $j \in [1, m]$, $d_j = \frac{j-1}{m-1}$. We can rewrite $K_{i,j}$:

$$K_{i,j} = \frac{\frac{i}{m-1}s_j + (1 + \frac{m-j}{m-1})s_i}{2(s_i + s_j)} = \frac{is_j + (2m-1-j)s_i}{2(m-1)(s_i + s_j)}$$

Note that these positions of the alternatives are symmetrical, thus $K'_{i,j} = K_{i,j}$.

Because $K_{1,j} \leq \frac{2m-1}{4(m-1)}$ for all $j \in [1, m-1]$, we have $s_j \geq \frac{2(m-1-j)+1}{2m-3}$ for all j, and because $K_{j,1} \leq \frac{2m-1}{4(m-1)}$ for all $j \in [1, m-1]$, we have $s_j \leq \frac{2(m-1-j)+1}{2m-3}$ for all j.

Thus, the only scoring rule that could have a inf-pairwise distortion of $\frac{2m-1}{2m-3}$ (if we restrict ourselves to positions of alternatives between x and y) is the positional scoring rule with $s_j = \frac{2(m-1-j)+1}{2m-3}$ for all $j \in [1, m-1]$, and $s_m = 0$, which is the OddBorda scoring vector. This means that if a scoring rule reaches this worst-case pairwise distortion, it must be OddBorda.

Next, we prove that this distortion is actually reached by OddBorda$_{\mathbf{PW}}$. Using the OddBorda vector and equidistant positions of the alternatives (i.e., $p_j = \frac{j-1}{m-1}$), we obtain for all i, j, $K_{i,j} = \frac{2m-1}{4(m-1)}$. By symmetry of the positions of the alternatives, $K_{i,j} = K'_{i,j}$. Finally, Lemma 1 implies that the inf-pairwise distortion of OddBorda is at most $\frac{2m-1}{2m-3}$.

Let us now prove that this is also a lower bound, and we cannot achieve a better worst-case distortion with positions of alternatives not between x and y. The proof works the same way as for Borda$_{\mathbf{PW}}$. For this we consider the profile of the case $K_{m-1,1}$. In this profile, a fraction $\frac{2m-3}{2m-2}$ of voters are at distance $p(i) = \frac{1}{2}$ of x and y and ranking x before y (i.e. $x \succ_i y$), and the remaining $\frac{1}{2m-2}$ fraction of voters are at the same position than y ($p(i) = 1$), ranking it first. Now, let $m_{(1,2]}$ be the number of alternatives with position $1 < p(i) \leq 2$ (they are closer to y than x is), $m_{>2}$ the number of alternatives with position $p(i) > 2$ (x is closer to y than they are), and $m_{<0}$ the number of alternatives with position $p(i) < 0$. This corresponds to the following situation:

We can give a lower bound of the normalized score of x (*i.e.*, its score divided by the number of voters), by only considering the score given by the alternatives at position $p(i) = 1/2$:

$$Sc(x) \geq \frac{2m-3}{2m-2} \cdot \frac{2 \cdot (m_{(1,2]} + m_{>2} + m_{<0} + 1) + 1}{2m-3}$$

We can also give an upper bound of the normalized score of y (if $m_{(1,2]} + m_{>2} + m_{<0} = 0$, then it is strictly lower, otherwise it is equal):

$$Sc(y) \leq \frac{1}{2m-2} \cdot 1 + \frac{2m-3}{2m-2} \cdot \frac{2 \cdot (m_{(1,2]} + m_{>2} + m_{<0}) + 1}{2m-3}$$

Thus $Sc(x) - Sc(y) \geq \frac{1}{2m-2} > 0$, meaning that x if the preferred alternative between x and y for OddBorda$_{\mathbf{PW}}$. However, y is the better alternative in this profile, which gives a distortion of:

$$\frac{\frac{2m-3}{2m-2} \cdot \frac{1}{2} + \frac{1}{2m-2} \cdot 1}{\frac{2m-3}{2m-2} \cdot \frac{1}{2}} = \frac{2m-1}{2m-3}$$

This prove that with this profile of voters, any positions of the alternatives lead to a worst-case distortion greater than ${}^{2m-1}/_{2m-3}$ for OddBorda$_{\mathbf{PW}}$. ☐

Upper Bound

Theorem 5. *The* sup*-pairwise distortion of Plurality*$_{\mathbf{PW}}$ *and more generally k-approval*$_{\mathbf{PW}}$ *for all $k \leq n$, in the 1D-Euclidean metric space, is $+\infty$ for any $m \geq 3$.*

Proof. The proof for Plurality$_{\mathbf{PW}}$ is identical to the proofs of STV$_{\mathbf{PW}}$ and PluralityVeto$_{\mathbf{PW}}$ from Theorem 8. Now assume x (resp. y) is at position $p(x) = 0$ (resp. $p(y) = 1$). The example for k-approval for $k \geq 2$ is the following: set all alternatives $z_j \neq x, y$ at position $p(z_j) = 3$ and all voters at position $p(i) = 0$ (same position as x). All voters rank

x first and y second, thus both will have a k-approval score of $Sc(x) = Sc(y) = n$. Because of tie-breaking, we assume without loss of generality that k-approval$_{\mathbf{PW}}(P \mid x, y) = y$, giving us the distortion of $+\infty$. □

Theorem 6. *The* sup-*pairwise distortion of Borda$_{\mathbf{PW}}$ (resp. OddBorda$_{\mathbf{PW}}$) in the 1D-Euclidean metric space is equal to $2m-1$ (resp. $4m-5$) for all $m \geq 2$.*

Proof. We already proved the result for Borda$_{\mathbf{PW}}$. We now prove it for OddBorda$_{\mathbf{PW}}$ (the proof is very similar to the one for Borda$_{\mathbf{PW}}$).

Recall the OddBorda scoring vector: $(2m-3, 2m-5, \ldots, 3, 1, 0)$. We divide voters into two groups: those who prefer x to y, and those who prefer y to x. Assume that y is preferred to x in the output ranking but x has a lower social cost. Observe that every voter who prefers x to y gives at least one more point to x than to y, and every voter who prefers y to x gives at most $2m-3$ more points to y than to x. Therefore, if we denote $\alpha \in [0, 1]$ the proportion of voters who prefer x to y, we have $Sc(x) \geq n\alpha$ and $Sc(y) \leq n(1-\alpha)(2m-3)$. To have $Sc(y) \geq Sc(x)$, we need $\alpha \leq (1-\alpha)(2m-3)$, which implies $\alpha \leq \frac{2m-3}{2m-2}$. Now, observe that the voters preferring x to y would maximize distortion by being at position $p(i) = 0$ if they could, and voters preferring y to x can maximize distortion by being at position $p(i) = 1/2$ (equal distance from x and y). Therefore, the worst-case distortion cannot be higher than

$$\frac{\alpha + (1-\alpha)\frac{1}{2}}{(1-x)\frac{1}{2}} = \frac{1+\alpha}{1-\alpha} \leq \frac{1 + \frac{2m-3}{2m-2}}{1 - \frac{2m-3}{2m-2}} = 4m-5$$

Moreover, this bound is reached. Let $1/4 > \varepsilon > 0$ and consider the profile in which all alternatives $z_j \neq x, y$ are at position $p(z_j) = \varepsilon$. Set $2m-3$ voters at position $p(i) = 2\varepsilon$, ranking y last and x second last, and one voter at position $p(i) = 1/2 + \varepsilon$, ranking y first and x last. Both x and y have OddBorda score $Sc(x) = Sc(y) = 2m-3$. Assume without loss of generality that ties are broken in favor of y. The distortion for this profile is:

$$\frac{(2m-3)(1-2\varepsilon) + \frac{1}{2} - \varepsilon}{\frac{1}{2} + \varepsilon}$$

When ε tends to 0, this ratio tends to $4m-5$. This concludes the proof. □

Optimality of Odd Borda

In this section, we extend the discussion about the optimility of OddBorda$_{\mathbf{PW}}$ among all pairwise rules (not only positional scoring rules) in the 1D-Euclidean metric space. We assume in all this section that all alternatives $z_j \in A$ are placed between x and y on the line (i.e., $0 \leq p(z_j) \leq 1$).

We will first show that given positions of the alternatives, we can compute a lower bound on the worst-case distortion by considering a set of profiles of voters. This bound will hold for all pairwise rules, and rely on a simple idea: two pseudometrics d and d' such that $d \approx P$ and $d' \approx P$ will lead to the same preference between x and y, i.e. $f(P \mid x, y) = f(P' \mid x, y)$. In particular, we might have a distortion strictly greater than 1 in either d or d' because x can have a lower social cost than y in d and a higher social cost in d'.

For this, assume that we denote the alternatives $A = \{z_1, \ldots, z_m\}$ with $x = z_1$, $y = z_m$ and $p(z_1) \leq \cdots \leq p(z_m)$. We also denote the positions of the alternatives $p_j = p(z_j) = d(x, z_j)$. There is a finite set of positions p on the line at which the preference of the voters change (*i.e.*, the preference ranking associated to $p-\varepsilon$ and $p+\varepsilon$ are different for any $\varepsilon > 0$). These positions are all the points which are exactly half-way between two different alternatives. For instance, if these two alternatives are z_j and z_k with $p_j \leq p_k$, then the middle point is $p^{1/2}_{j,k} = \frac{p_j+p_k}{2}$ and for all voters i with $p(i) < p^{1/2}_{j,k}$, we have $z_j \succ_i z_k$ and for all voters i with $p(i) > p^{1/2}_{j,k}$, we have $z_k \succ_i z_j$. Except if this point is also a middle point for another pair of alternatives, the preferences between all other pairs remain the same before and after this point. Note that there are $\frac{m(m-1)}{2}$ pairs of alternatives, so at most $\frac{m(m-1)}{2}$ such middle points (at most, because some of these points might coincide, for instance if the alternatives are symmetrical). Let's denote $q_1, \ldots, q_t$ these points such that $q_1 < q_2 < \cdots < q_{t-1} < q_t$. To these points we add $q_0 = 0 = p(x)$ and $q_{t+1} = 1 = p(y)$. Now, observe that for all $0 \leq j \leq t$, voters have same preferences along all the interval (q_j, q_{j+1}), with position q_j if they want to maximize social cost of y and position q_{j+1} if they want to maximize social cost of x.

We now split the middle points into two groups: $Q_{<0.5} = \{q_j \mid q_j < 0.5\}$ and $Q_{\geq 0.5} = \{q_j \mid q_j \geq 0.5\}$ (respectively points that are closer to x and closer y). Let's now take $q_j \in Q_{\leq 0.5}$ and $q_k \in Q_{\geq 0.5}$ we consider two metrics:

- In d, a proportion $\alpha \in [0, 1]$ of voters are at position q_j and a proportion $1-\alpha$ of voters at position q_k.
- In d', a proportion $\alpha \in [0, 1]$ of voters are at position q_{j+1} and a proportion $1-\alpha$ of voters at position q_{k+1}.

As we explained, there exists a preference profile P such that $P \approx d$ and $P \approx d'$ (because there are no middle points in the two intervals (q_j, q_{j+1}) and (q_k, q_{k+1})). Let f be a pairwise rule.

- If $f(P \mid x, y) = x$, then the worst-case distortion is when all voters are either at q_{j+1} or q_{k+1}. Thus, we have a distortion of

  $$\frac{q_{j+1} \cdot \alpha + q_{k+1} \cdot (1-\alpha)}{(1-q_{j+1}) \cdot \alpha + (1-q_{k+1}) \cdot (1-\alpha)} = \frac{K_1(\alpha)}{1-K_1(\alpha)}$$

  with $K_1(\alpha) = q_{j+1} \cdot \alpha + q_{k+1} \cdot (1-\alpha)$. Then, maximizing the distortion is equivalent to maximizing $K_1(\alpha)$. Note that because $q_{k+1} > q_{j+1}$, $K_1(\alpha)$ is decreasing when α increases.

- If $f(P \mid x, y) = y$, then the worst-case distortion is when all voters are either at q_j or q_k. Thus, we have a distortion of

  $$\frac{(1-q_j) \cdot \alpha + (1-q_k) \cdot (1-\alpha)}{q_j \cdot \alpha + q_k \cdot (1-\alpha)} = \frac{K_2(\alpha)}{1-K_2(\alpha)}$$

with $K_2(\alpha) = (1-q_j)\cdot\alpha + (1-q_k)\cdot(1-\alpha)$. Then, maximizing distortion is equivalent to maximizing $K_2(\alpha)$. Note that because $q_k > q_j$, $K_2(\alpha)$ is increasing when α increases.

The best possible rule will always take the alternative that minimizes worst-case distortion, so it will take the distortion for the minimum of $K_1(\alpha)$ and $K_2(\alpha)$. However, the voters will select the α that maximizes the worst-case. If we see it as a two players game, then the first player selects α the second player selects which of x or y is preferred (and thus if we look at $K_1(\alpha)$ or $K_2(\alpha)$). The first player wants to maximize the distortion and the second player wants to minimize the distortion. Thus, the first player wants to select α which maximizes the expression $\min(K_1(\alpha), K_2(\alpha))$. As $K_1(\alpha)$ decreases and $K_2(\alpha)$ increases when α increases, the expression $\min(K_1(\alpha), K_2(\alpha))$ reaches its highest value when $K_1(\alpha) = K_2(\alpha)$. We can now compute the corresponding value of α:

$$\begin{aligned} q_{j+1}\alpha + q_{k+1}(1-\alpha) &= (1-q_j)\alpha + (1-q_k)(1-\alpha) \\ \alpha(q_{j+1} - q_{k+1} + q_j - q_k) &= 1 - q_k - q_{k+1} \\ \alpha &= \frac{1 - q_k - q_{k+1}}{q_j + q_{j+1} - q_k - q_{k+1}} \end{aligned}$$

We now replace this value of α in the distortion of one of the two profiles (this gives the same result for both, as $K_1(\alpha) = K_2(\alpha)$):

$$\begin{aligned} &\frac{q_{j+1}\frac{1-q_k-q_{k+1}}{q_j+q_{j+1}-q_k-q_{k+1}} + q_{k+1}(1 - \frac{1-q_k-q_{k+1}}{q_j+q_{j+1}-q_k-q_{k+1}})}{1 - (q_{j+1}\frac{1-q_k-q_{k+1}}{q_j+q_{j+1}-q_k-q_{k+1}} + q_{k+1}(1 - \frac{1-q_k-q_{k+1}}{q_j+q_{j+1}-q_k-q_{k+1}}))} \\ =&\frac{q_{j+1}(1-q_k) + q_{k+1}(q_j - 1)}{q_j + q_{j+1} - q_k - q_{k+1} - (q_{j+1}(1-q_k) + q_{k+1}(q_j-1))} \\ =&\frac{q_{j+1}(1-q_k) + q_{k+1}(q_j-1)}{q_k(q_{j+1}-1) + q_j(1-q_{k+1})} \end{aligned}$$

We denote this value $\text{dist}_{j,k}$. Now the worst-case distortion for given position of the alternative is at least equal to:

$$\text{lb-dist}(p_1, \ldots, p_m) = \max_{\substack{q_j \in Q_{<0.5} \\ q_k \in Q_{\geq 0.5}}} \text{dist}_{j,k}$$

Note that we did not assume anything on the pairwise rule so this lower bound stands for every pairwise rule.

Simulations

Our approach was the following: we sampled positions of alternatives $p_1, \ldots, p_m$ uniformly at random between x and y and for each, we computed lb-dist$(p_1, \ldots, p_m)$. If, for a given m, we fail to find positions of alternatives such that lb-dist$(p_1, \ldots, p_m) \leq \frac{2m-1}{2m-3}$, this suggests that there are no such positions (especially for low values of m). This is of course not a formal proof, but a strong evidence. Moreover, we believe that the above computation can be used to obtain a formal proof (in particular, it is quite easy for $m = 3$ but seems hard to generalize).

To summarize, our experiments work as follow:

1. Sample positions of the additional alternatives uniformly at random between x and y (*i.e.*, $p_j \in [0, 1]$).
2. Compute the positions of middle points $q_1, \ldots, q_t$, and divide them into two groups $Q_{<0.5}$ and $Q_{\geq 0.5}$.
3. Compute
$$\text{lb-dist}(p_1, \ldots, p_m) = \max_{\substack{q_j \in Q_{<0.5} \\ q_k \in Q_{\geq 0.5}}} \text{dist}_{j,k}$$
4. Repeat.

Table 2 summarizes our results.

For each $m \leq 5$ (*i.e.*, up to 3 additional alternatives), we sampled 1,000,000 different positions of the additional alternatives. For none of them could we find a lower-bound of the worst-case distortion lower than $\frac{2m-1}{2m-3}$. And we found that this bound is reached only for equidistant positions of the alternatives. As explained before, this suggests that OddBorda$_{\mathbf{PW}}$ is optimal among all pairwise rules *if we assume that all alternatives are between x and y*. As it seems counter-intuitive that positions of alternatives outside $[0, 1]$ would lead to a better distortion, we conjecture that OddBorda$_{\mathbf{PW}}$ is optimal among all pairwise rules for the inf-pairwise distortion in the 1D-Euclidean metric space and $m \leq 5$

For $m \geq 6$, however, we found that for some positions of alternatives $p_1, \ldots, p_m$, the lower bound of the worst-case distortion is lower than the inf-pairwise distortion of OddBorda$_{\mathbf{PW}}$. This fact alone does not necessarily mean that OddBorda$_{\mathbf{PW}}$ is not optimal, as we do not try every profile of voters, but it suggests that there might be a better rule. We found this better rule, which is not a positional scoring rule, but in some sense is a variant of OddBorda$_{\mathbf{PW}}$, which is why we name it here 1D-OddBorda$_{\mathbf{PW}}$ (as it is tailored to be optimal specifically in a 1D-Euclidean metric space). We describe this rule in the next subsection.

The 1D-OddBorda Rule

In this section, we describe the 1D-OddBorda$_{\mathbf{PW}}$ rule and show its inf-pairwise distortion is equal to $\frac{4m-11}{4m-13}$. First, to understand this rule, let us look at the positions of alternatives leading to this inf-pairwise distortion. The first additional alternative z_2 (as $z_1 = x$) is positioned at $p_2 = \frac{1}{2m-6}$ (as $p_1 = p(x) = 0$). The alternatives $z_3, \ldots, z_{m-2}$ are at position $p_j = \frac{2j-4}{2m-6}$ and the last alternative z_{m-1} (as $z_m = y$ and $p_m = 1$) is at position $p_{m-1} = \frac{2m-7}{2m-6}$. For instance, for $m = 6$, this gives the vector $(0, 1/6, 2/6, 4/6, 5/6, 1)$, and for $m = 8$ this gives $(0, 1/10, 2/10, 4/10, 6/10, 8/10, 9/10, 1)$. With these positions, the midpoints between alternatives are the $q_j = \frac{j}{2(2m-6)} = \frac{j}{4m-12}$ for $1 \leq j \leq 4m-13$, to which we add $q_0 = 0$ the position of x and $q_{4m-12} = 1$ the position of y. Note that these q_j are equidistant. These are actually the same midpoints as for the equidistant positions of m' alternatives with $m' = 2m-5$. For instance, Figure 11 shows that for $m = 7$, there are far more equidistant *midpoints* with the positions of 1D-OddBorda$_{\mathbf{PW}}$ than with equidistant positions of alternatives.

For OddBorda$_{\mathbf{PW}}$ with m' alternatives, the inf-pairwise distortion use these midpoints $q_0, \ldots, q_{2m'-2}$. Note that

m	3	4	5	6	7	8	9
lower bound	≈ 1.667	≈ 1.400	≈ 1.286	≈ 1.185	≈ 1.143	≈ 1.112	≈ 1.091
OddBorda$_{\mathbf{PW}}$	≈ 1.667	1.4	≈ 1.286	≈ 1.222	≈ 1.182	≈ 1.154	≈ 1.133
1D-OddBorda$_{\mathbf{PW}}$				≈ 1.182	≈ 1.133	≈ 1.105	≈ 1.087

Table 2: Comparison of the lower worst-case distortion obtained in the experiments and the $\inf$-pairwise distortion of OddBorda$_{\mathbf{PW}}$ and 1D-OddBorda$_{\mathbf{PW}}$.

with m' alternatives, the position of x in the preference ranking changes at all midpoint q_j with $j \leq m' - 1$, thus x receives a different score in each interval (q_j, q_{j+1}) for $j \leq m' - 1$. The same thing happens for y for the q_j with $j > m' - 1$. What we want for 1D-OddBorda$_{\mathbf{PW}}$ to do is attribute the same scores to x and y as OddBorda$_{\mathbf{PW}}$ does on each interval (q_j, q_{j+1}). However, no positional scoring rule can make this work, as the score of x (resp. y) changes only when we cross a midpoint between x (resp. y) and another alternative, which is the case for all midpoints with equidistant positions, but not for the positions of alternatives described above. However, the overall preference ranking is changing when we cross each midpoint, and is different on each interval (q_j, q_{j+1}). For instance, we can go from $z_2 \succ z_3 \succ x \succ z_4 \succ \cdots \succ y$ to $z_3 \succ z_2 \succ x \succ z_4 \succ \cdots \succ y$: x and y remained at the same positions but the first two alternatives switched.

Therefore, if we associate each preference ranking to the correct score for x and y, we can simulate OddBorda$_{\mathbf{PW}}$ for $m' = 2m - 5$. This means that the score of x will be higher in the preference ranking $z_2 \succ z_3 \succ x \succ z_4 \succ \cdots \succ y$ than in the preference ranking $z_3 \succ z_2 \succ x \succ z_4 \succ \cdots \succ y$ where only the top-2 alternatives are switched, but not the position of x.

Thus, the $\inf$-pairwise distortion of 1D-OddBorda$_{\mathbf{PW}}$ for m alternatives is equal to the $\inf$-pairwise distortion of OddBorda$_{\mathbf{PW}}$ for $m' = 2m - 5$ alternatives, i.e. its $\inf$-pairwise distortion will be:

$$\frac{2(2m-5)-1}{2(2m-5)-3} = \frac{4m-11}{4m-13}$$

As one can see on Table 2, we did not find any better lower bound for $m \in [6, 9]$ over 1,000,000 different positions of the alternatives. Moreover, we conducted experiments with greater values of m (up to 20) and a lower number of tested positions and the lower bound was always greater than the $\inf$-distortion of 1D-OddBorda$_{\mathbf{PW}}$. We conjecture that this is the optimal pairwise rule in the 1D-Euclidean setting (with 1D-OddBorda$_{\mathbf{PW}}$ being equal to OddBorda$_{\mathbf{PW}}$ when $m \leq 5$), and leave the proof to further research. Note however that this rule is very specific to the 1D-Euclidean metric space.

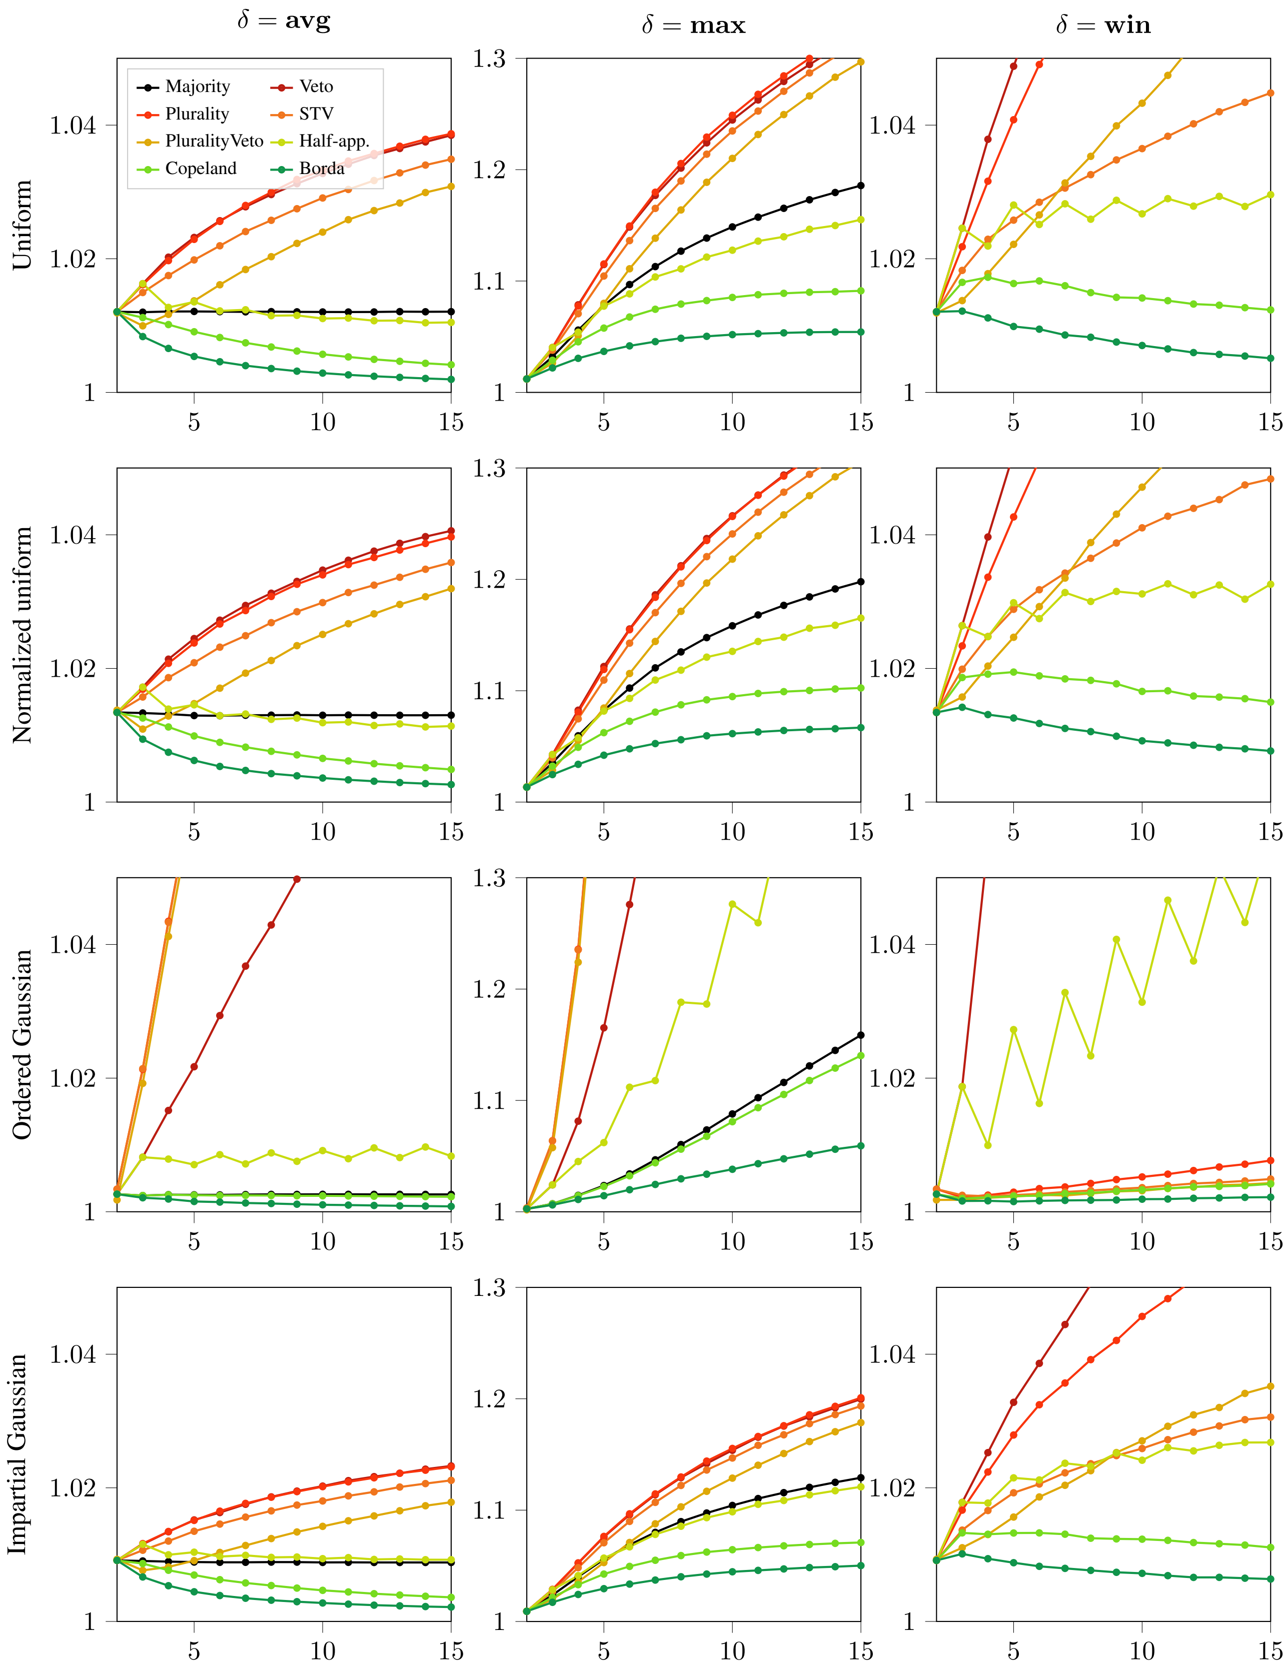

Figure 7: Results of the experiments for average pairwise distortion in the unconstrained setting over 10,000 random profiles. The x-axis corresponds to the number of voters m and the y-axis to the average pairwise distortion.

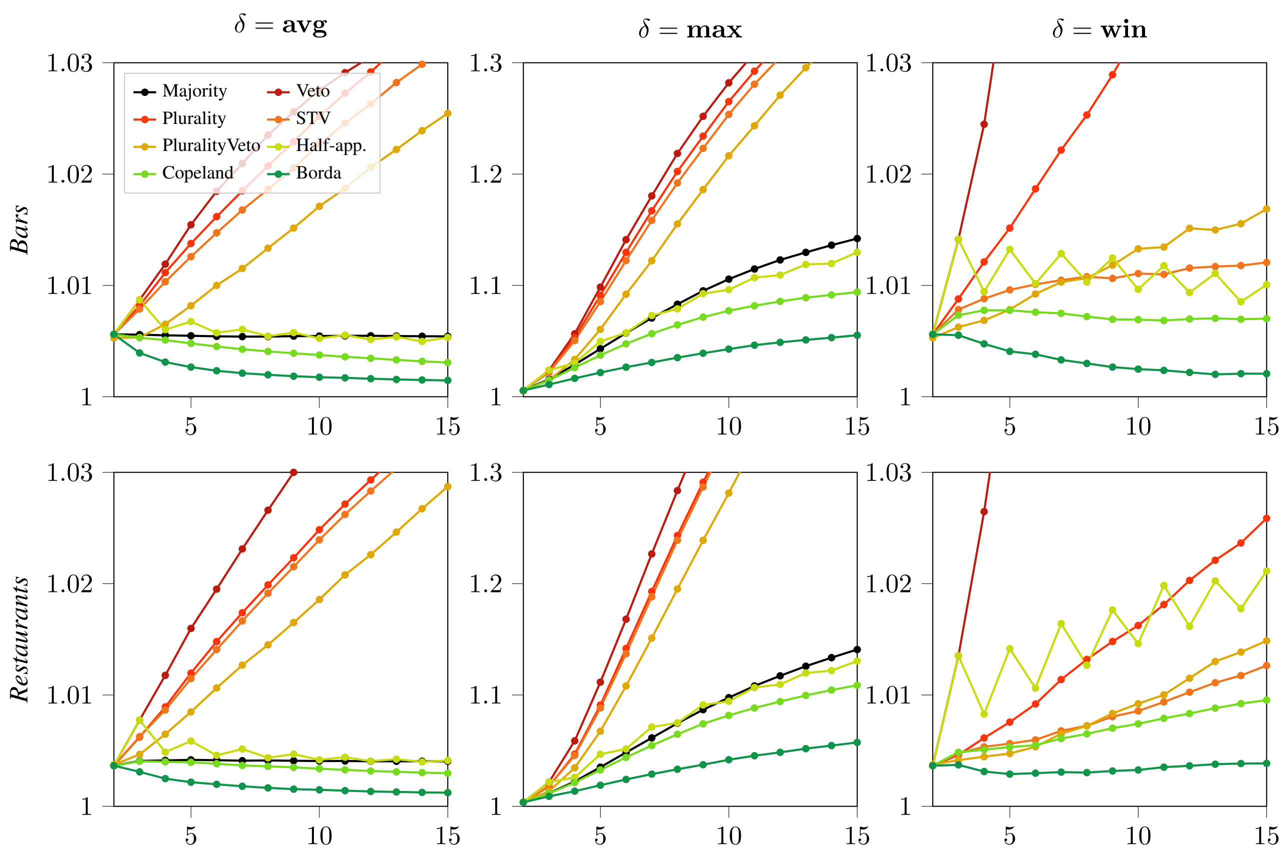

Figure 8: Results of the experiments for empirical distortion over 10,000 random profiles. The x-axis corresponds to the number of voters m and the y-axis to the average pairwise distortion.

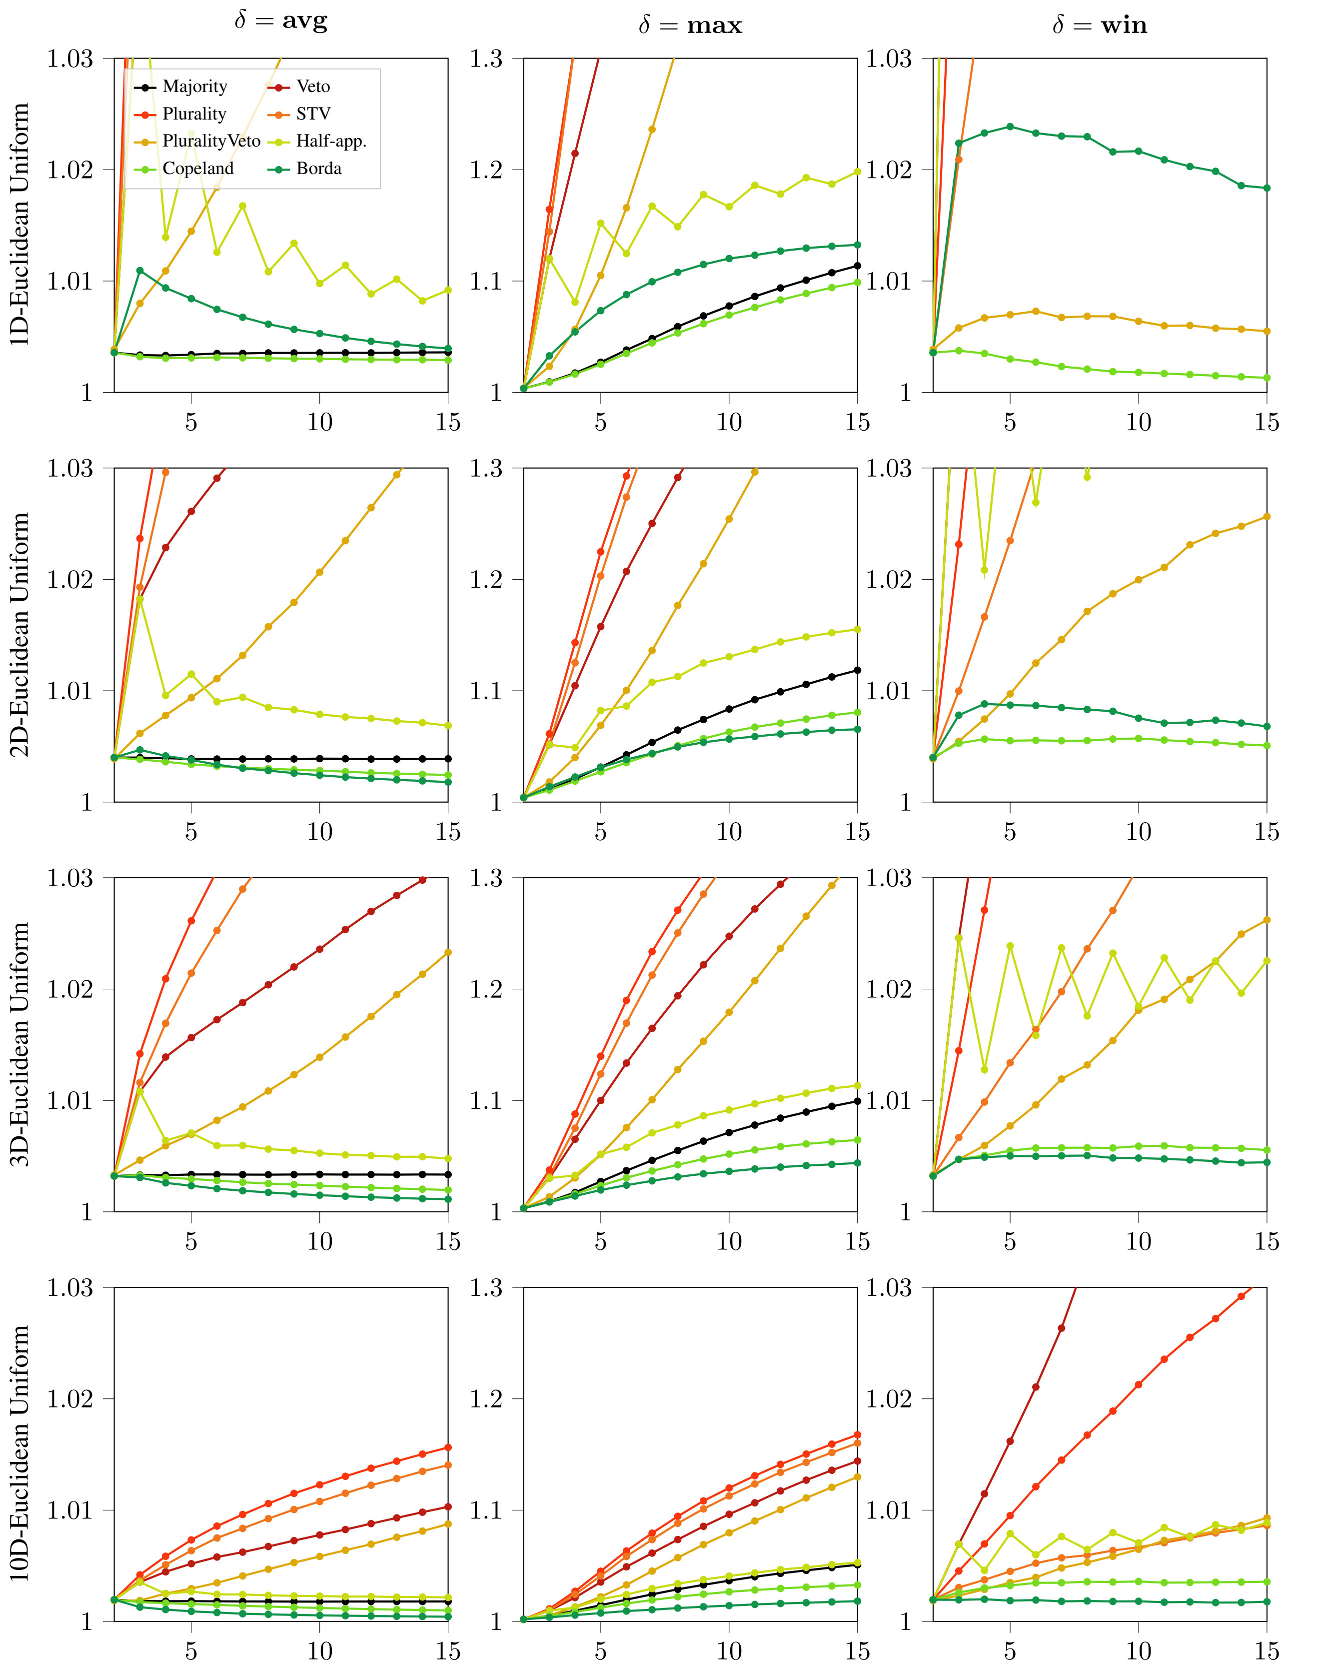

Figure 9: Results of the experiments for average pairwise distortion in the metric setting over 10,000 random profiles. The x-axis corresponds to the number of voters m and the y-axis to the average pairwise distortion.

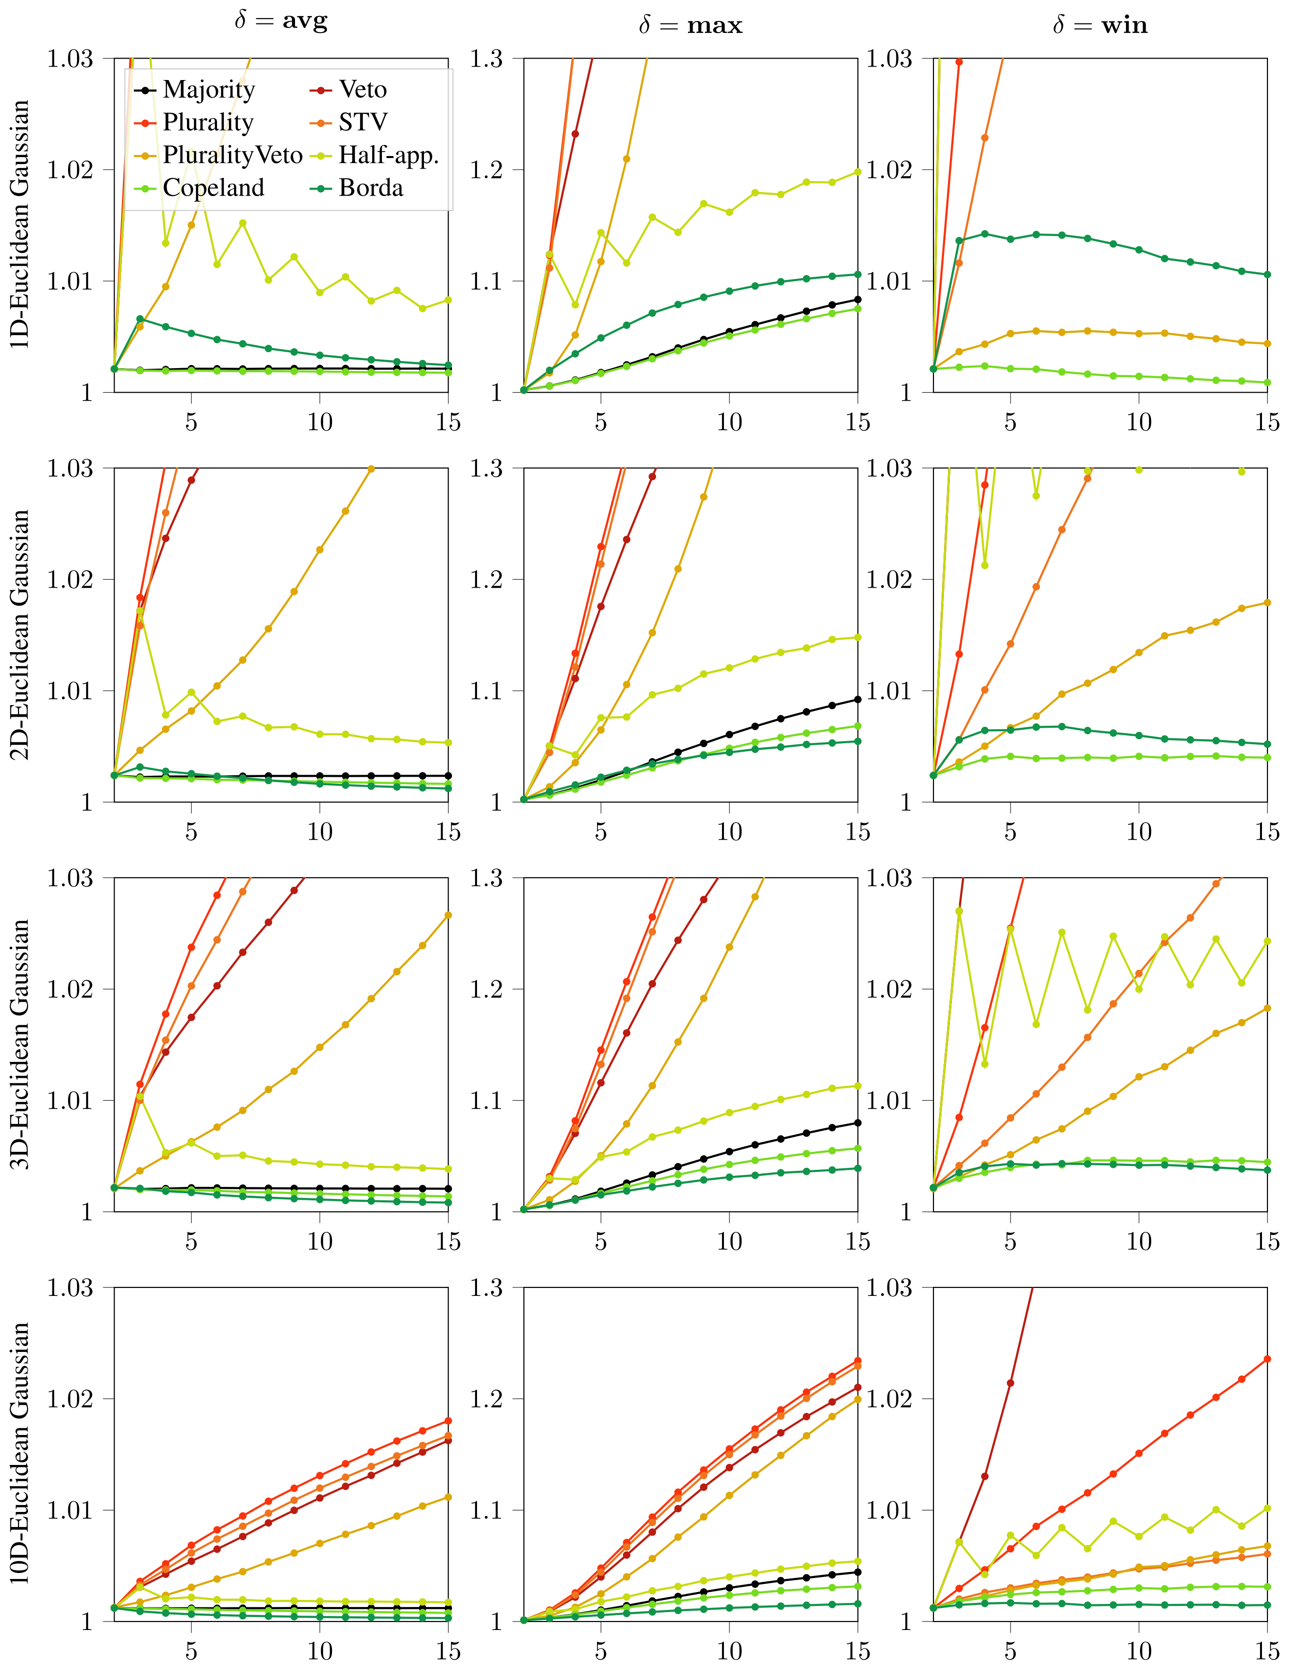

Figure 10: Results of the experiments for average pairwise distortion in the metric setting over 10,000 random profiles. The x-axis corresponds to the number of voters m and the y-axis to the average pairwise distortion.

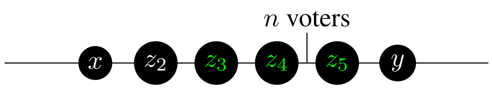

n voters
x
z_2
z_3
z_4
z_5
y

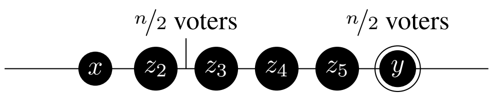

$n/2$ voters
$n/2$ voters
x
z_2
z_3
z_4
z_5
y

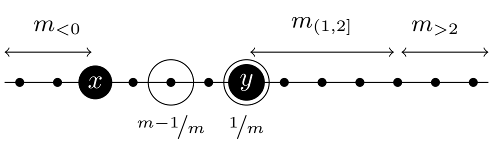

$m_{<0}$
$m_{(1,2]}$
$m_{>2}$
x
y
$m-1/m$
$1/m$

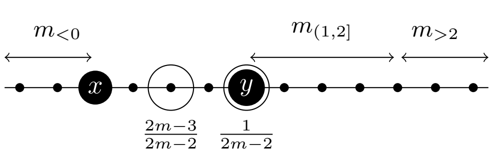

$m_{<0}$
$m_{(1,2]}$
$m_{>2}$
x
y
$\frac{2m-3}{2m-2}$
$\frac{1}{2m-2}$

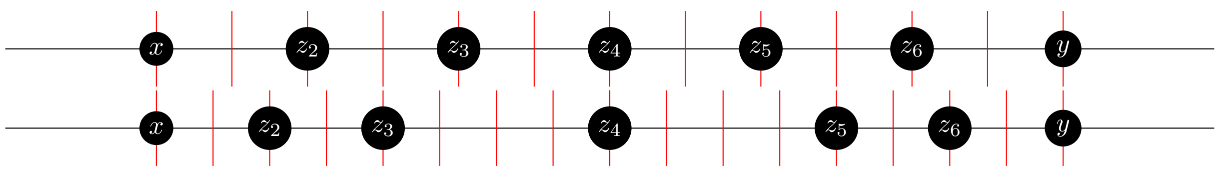

Figure 11: The positions of the midpoints q_j (in red) for two different positions of the alternatives: equidistant (top) and positions for 1D-OddBorda$_{\mathbf{PW}}$ (bottom).