Independence of Irrelevant Alternatives under the Lens of Pairwise Distortion

Théo Delemazure
CNRS, Paris Dauphine University – PSL

Jérôme Lang
CNRS, Paris Dauphine University – PSL

Grzegorz Pierczyński
University of Warsaw

Problem: preference aggregation

1^{st}	2^{nd}	3^{rd}	4^{th}
Bob	**Ann**	**Dan**	**Carol**
Dan	**Bob**	**Ann**	**Carol**
Carol	**Ann**	**Bob**	**Dan**
Bob	**Ann**	**Carol**	**Dan**
⬇			
Bob	**Ann**	**Dan**	**Carol**

Independence of Irrelevant Alternatives (IIA)

The collective choice between two candidates depends only on the preferences of voters over them.

If we blindly follow IIA, **Bob** should be preferred to **Ann** in the following example, even if **half of the voters hate him**:

51% **Bob** ≻ **Ann** ≻ ··· ≻ **Carol** ≻ **Dan**

49% **Ann** ≻ **Carol** ≻ ··· ≻ **Dan** ≻ **Bob**

Arrow's theorem: No reasonable rule satisfies IIA.

Our goal: **Measuring** the *(negative)* impact of **IIA** on social welfare.

We need a **framework** that can capture both *ranking rules* and *IIA*.

Pairwise voting rule: for *each pair of candidates*, decide which one is preferred.

Ranking rules are **transitive** pairwise rules and IIA is represented by the **majority rule**.

We need a **tool** to **quantify** the *loss* in social welfare induced by IIA.

We adapt the notion of **distortion** to pairwise voting rules, in both *utility model*, and *metric model* with costs.

Formally, it is *the ratio* between the utility/cost of the **selected** candidate and of the **best** candidate.

Pairwise distortion

Can be defined based on **utilities** or **costs** (for instance in the metric space).

If the total utilities are...

$$U(\textbf{Ann}) = 6000$$
$$U(\textbf{Bob}) = 9400$$

...but the rule selects **Ann** over **Bob**, the distortion is

$$\frac{9400}{6000} \approx 1.56$$

The lower the pairwise distortion, **the better.**

1. Average pairwise distortion

We conducted experiments on **synthetic** and **real** data, with both models. For instance, with **uniform distribution** of utilities, we obtain:

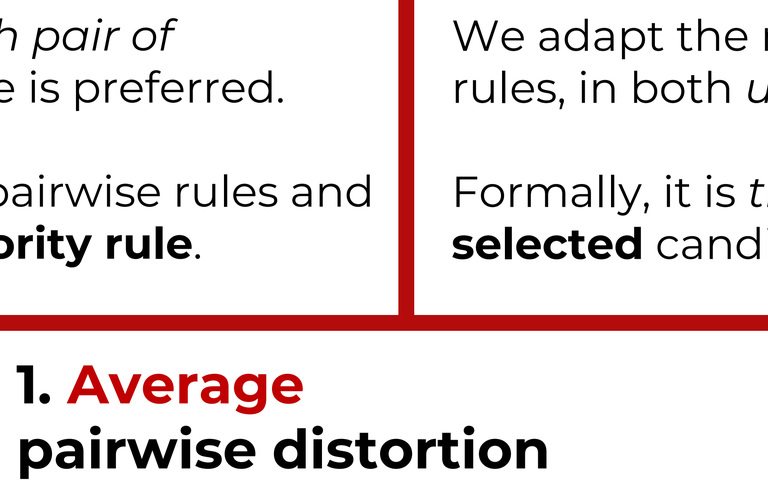

2. Worst-case pairwise distortion

In the **metric** model, if we assume worst-case positions of the voters, **how much can we decrease the distortion** by strategically placing the "*irrelevant*" alternatives?

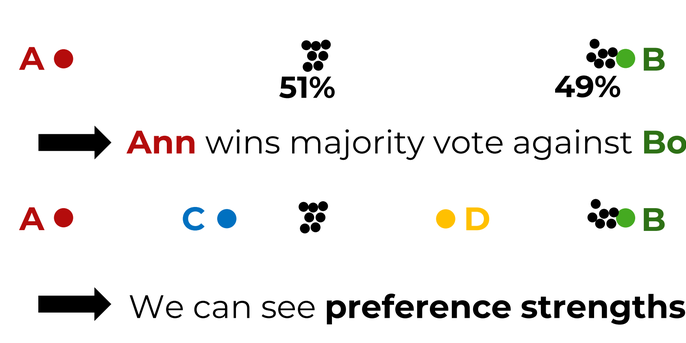

➡ **Ann** wins majority vote against **Bob**.

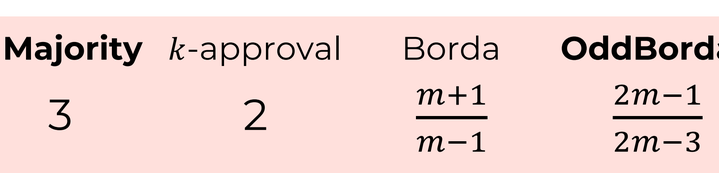

➡ We can see **preference strengths.**

Majority	k-approval	Borda	**OddBorda***
3	2	$\frac{m+1}{m-1}$	$\frac{2m-1}{2m-3}$

Scan to read the full paper!

The information provided by *irrelevant* alternatives can be used wisely by a rule, and in some contexts, **it may help to improve the social welfare of the society.**