

Selecting the Most Conflicting Pair of Candidates

Théo Delemazure¹, Łukasz Janeczko², Andrzej Kaczmarczyk² and Stanisław Szufa^{1,2}

¹CNRS, LAMSADE, Université Paris Dauphine - PSL

²AGH University

theo.delemazure@dauphine.eu, ljaneczko@agh.edu.pl, andrzej.kaczmarczyk@agh.edu.pl,
s.szufa@gmail.com

Abstract

We study committee elections from a perspective of finding the most conflicting candidates, that is, candidates that imply the largest amount of conflict, as per voter preferences. By proposing basic axioms to capture this objective, we show that none of the prominent multiwinner voting rules meet them. Consequently, we design committee voting rules compliant with our desiderata, introducing *conflictual voting rules*. A subsequent deepened analysis sheds more light on how they operate. Our investigation identifies various aspects of conflict, for which we come up with relevant axioms and quantitative measures, which may be of independent interest. We support our theoretical study with experiments on both real-life and synthetic data.

1 Introduction

Where a collective decision over a set of options based on a number of opinions has to be reached, conflict is inevitable. Reflected by differences in the opinions, it usually comes from different perspectives of the opinions (e.g., in the case of human opinions, a beginner investor would likely have completely different opinion on various asset classes than their professional counterpart). However, conflict might also be option-based and stem from diverse, sometimes even contradicting, inherent qualities of the options (e.g., potentially high-return assets typically have high risk levels).

We are interested in how to identify these conflicting options, based on the preferences. Since the options might represent multiple entities (e.g., sports players, societal issues, or marketing strategies), answering this question has numerous natural applications that include selecting competitors to organize engaging sport events (e.g., boxing matches), controversial topics to organize interesting political debates, disputable issues for socially-relevant deliberations, or conflicting ideas for boosting the creativity with passionate discussions.

A somewhat different application of identifying conflicting options is learning new insights about the options or the opinions' perspectives. Imagine a space agency that validates procedures (options) for landing on the Moon using a collection of complex simulations. Each simulation assesses the quality

of each procedure and ranks the procedures (expressing an opinion) from the ones that are most likely to the one that are the least likely to succeed. The existence of two significantly conflicting procedures can then offer additional insights. It might suggest that there is some (possibly unknown) feature that impact only some of the simulations and is ignored by the rest. Alternatively, the two procedures may differ in some operational detail that is a crucial success factor for some of the simulated scenarios. In both cases, a careful inspection of the procedures would help to recognize the source of the conflict and thus contribute to advancing the explainability of the simulations or the knowledge about the procedures.

To provide a big picture of our approach, we use a particularly illustrative application, which is finding polarizing issues. Having a collection of *ordinal preferences* of voters expressing their view on the importance of the issues, we aim at identifying *two* issues that are the *most conflicting* ones. This goal poses a significant conceptual challenge as conflict, which in our scenario can be associated with polarization, is actually of dual nature. Indeed, we want to find two issues that are (1) supported by two different, large, and, ideally, equal-sized groups of voters and that are (2) perceived ideologically as far from each other as possible in these two groups. As it turns out, balancing between these two conditions forms an important part of our modeling of conflict. A natural motivation behind our goal in this scenario is to reduce the polarization. This motivation combines the two aforementioned applications: of selecting two disputable issues and of fostering learning about the selected issues and the voters. Indeed, learning about the groups of voters can help in finding communication channels to improve dialogue between the groups. The conflicting issues, on the other hand, are the first ones to be addressed by specific policies. Additionally, the selected issues might indicate particular compromises that must be obeyed not to divide the society more.

We view our task of finding conflicting candidates naturally falling into the framework of *multiwinner voting rules*. Such rules, given *preferences* of voters over *candidates* and a desired committee size, select a winning *committee* — a subset of the candidates — of the desired size that aims at meeting certain objectives. Adapting this terminology to our scenario, we also have a number of candidates and a collection

	MaxSum	MaxNash	MaxPolar	MaxSwap
Reverse Stability	✓	✓	✓	✓
Conflict Consistency	✓	✓	✓	✓
Conflict Monotonicity	✗	✗	✗	✗
Antagonization Consistency	✓	✓	✓	✓
Matching Domination	✓	✓	✓	✗
Balance Preference	✗	✓	✗	✓

Table 1: Axiomatic properties of conflictual rules.

of voters¹ expressing preferences over the candidates. Our objective is to select the most conflictual candidates.

In their survey, Faliszewski *et al.* [2017] discuss three so far studied objectives of multiwinner voting rules: (1) individual excellence, focusing on selecting individually the best candidates, (2) diversity, which aims at representing as many voters as possible, and (3) proportional representation, which selects a committee that proportionally represents the voters. There seems to be a clear consensus that the two former goals are achieved by, respectively, the k -Borda [Debord, 1992] and the Chamberlin-Courant [Chamberlin and Courant, 1983; Elkind *et al.*, 2017b] rules. Understanding how various rules guarantee achieving proportionality is an active area of research [Monroe, 1995; Elkind *et al.*, 2017a; Elkind *et al.*, 2017b; Skowron *et al.*, 2017; Faliszewski *et al.*, 2019; Peters *et al.*, 2021].

Focusing on diverse candidates might appear as being able to fulfil our goal of seeking conflictual candidates. Intuitively, two most diverse candidates should be exactly those that divide the voters the most. However, as we show in Section 3.2, a fairly small example shows that this is not the case (in fact, the example shows that none of the standard three objectives meets our needs). We also verify this considering scenarios beyond the “worst-case” ones in our simulations.

To the best of our knowledge, no rule suitable for finding such conflicting candidates has been introduced so far. Nonetheless, this line of research closely aligns with the concept of polarization² in elections extensively studied from axiomatic, computational, and experimental perspectives [Alcalde-Unzu and Vorsatz, 2013; Hashemi and Endriss, 2014; Alcalde-Unzu and Vorsatz, 2016; Can *et al.*, 2017; Colley *et al.*, 2023; Faliszewski *et al.*, 2023]. Yet, the focus has so far been on measuring a value of polarization of all voters or of a single candidate, which is a different problem than selecting a pair (or a subset) of conflicting candidates.

Our Contributions. Our primary contribution is conceptual, offering a novel perspective on the voting theory inspired by finding conflicting candidates. To this end, we establish

¹For convenience and increased readability, we decided to keep term “voters,” even though, as demonstrated in the introduction, our model and its applications are not limited to voting scenarios.

²As well as with agreement and diversity of votes, as argued by Faliszewski *et al.* [2023].

new foundational axioms and show that they differ from the traditional ones. Defining more demanding axioms, we arrive at an impossibility result that sets our expectations (Section 3). This result also guides us in developing new multiwinner voting rules, termed *conflictual voting rules*, in Section 4. They are specifically designed to identify the most conflicting pair(s) of candidates as required by the introduced axioms (Table 1 highlights our axiomatic analysis). To better understand the intricacies of the introduced rules, in Section 5, we provide multiple metrics capturing different aspects of conflict, which can be of independent interest.

Finally, in Section 6, we present the experimental evaluation, using both synthetic and real-life elections, empirically showing the behavior of the conflictual rules.³ Thus, we not only validate our theoretical insights but also provide insights into their real-world applications and implications.

Due to space constraint, missing proofs, details, and additional experiment results are deferred to the full version of this work [Delemazure *et al.*, 2024].

2 Preliminaries

For a set X , let $\#X$ denote the cardinality of that set. Let $E = (C, V)$ be an election with a set $C = \{c_1, \dots, c_m\}$ of candidates and a set $V = \{v_1, \dots, v_n\}$ of voters. A profile $P = (\succ_1, \dots, \succ_n)$ is a collection of rankings (total linear orders) over C such that each voter v_i is associated with \succ_i ; we denote the set of all possible profiles by \mathcal{P} . By $\overleftarrow{P} = (\overleftarrow{\succ}_1, \dots, \overleftarrow{\succ}_n)$, we denote profile P with all ballots being reversed. In particular, for all $v_i \in V$ and all $c, c' \in C$, $c \succ_i c'$ if and only if $c' \overleftarrow{\succ}_i c$. A committee voting rule \mathcal{R} is a function that takes as input a profile $P \in \mathcal{P}$ and a committee size $k \geq 2$ and outputs a non-empty set $\mathcal{R}(P, k)$ of committees, such that for each $W \in \mathcal{R}(P, k)$ we have $\#W = k$ and $W \subseteq C$. In this document, we focus on the case $k = 2$, and write $\mathcal{R}(P)$ for simplicity.

For brevity, we denote by $v(a)$ the position of candidate a in vote v . More formally, we have $v_i(a) = \#\{x \mid x \succ_i a\} + 1$. Moreover, let $v(a, b)$ denote the distance between a and b in vote v , that is, $v(a, b) = v(b) - v(a)$; for conciseness instead of $v(a, b)$, we write $v(ab)$. For instance, in the vote $v = a \succ b \succ c \succ d \succ e$, we have $v(b) = 2$, $v(e) = 5$ and $v(be) = v(b, e) = 5 - 2 = 3$. Analogously $v(eb) = -3$.

Furthermore, for all pairs $\{a, b\}$ of candidates from C , we denote $V^{a \succ b} = \{v_i \in V \mid a \succ_i b\}$ the set of voters preferring a to b in profile P . Finally, we say that a pair of candidates $\{a, b\}$ is conflictual if their ordering is not unanimous, i.e., $\#V^{a \succ b} > 0$ and $\#V^{b \succ a} > 0$. In other words, a pair of candidates is conflictual if neither of them Pareto-dominates the other one.

3 Properties of Conflictual Rules

Before defining our rules, we proceed with fundamental properties that we require from them. As we will show, already these somewhat weak, basic axioms prove that the objectives

³The experiments’ source code is freely available at <https://github.com/Project-PRAGMA/conflictual-rules--IJCAI-24>.

of conflictual rules are far from those of the rules commonly studied in the computational social choice literature.

3.1 Fundamental Axioms

We start by defining two fundamental axioms expected from conflictual rules. The first one assures that unless we have an identity election, only a conflicting pair could win.

Definition 1 (Conflict Consistency). *Rule \mathcal{R} is conflict consistent if it does not output non-conflicting pair if there exists a conflicting one.*

The second axiom comes from the observation that in conflicting voting rule, we want “love” and “hate” to have symmetrical purposes. Thus, the most polarizing pair in a profile should be as much polarizing if we reverse all voters’ preferences.

Definition 2 (Reverse Stability). *Rule \mathcal{R} is reverse stable if for each profile P it holds that $\mathcal{R}(P) = \mathcal{R}(\overleftarrow{P})$.*

3.2 Relation to Standard Axioms

The postulates of our axioms stand in stark contrast to all established objectives of multiwinner voting rules, that is, to individual excellence, proportionality, and diversity. In particular, conflictual rules are incompatible with the unanimity axiom. This is a very weak axiom guaranteeing effectiveness, saying that if a candidate is ranked first by every voter, it should be selected in the committee. This is the main difference with the classical paradigm of social choice: conflictual rules specifically avoid consensual candidates.

Definition 3 (Unanimity). *Rule \mathcal{R} is unanimous if a candidate that appears on top of every ranking is always in the winning committee.*

Proposition 1. *There is no rule \mathcal{R} that satisfies both unanimity and conflict-consistency.*

Proof. Consider the profile $\{a \succ b \succ c, a \succ c \succ b\}$. Unanimity implies a belongs to the committee, but the only conflicting pair in this profile is $\{b, c\}$. \square

The incompatibility stated by Proposition 1 has a strong implication. Namely, conflictual rules are different from the very prominent family of committee scoring rules. The observation follows from the fact that the latter rules must be unanimous [Skowron *et al.*, 2019].

3.3 Further Properties of Conflictual Rules

In this section, we consider more axioms, that sound desirable but are unfortunately not always possible to satisfy.

The first axiom we want to introduce is a kind of Pareto-domination axiom, since it is based on the domination of a pair of candidates by another one. However, the pair needs not to be dominated in every ranking, as in classic Pareto properties. Given profile P and two conflicting pairs of candidates $\{a, b\}$ and $\{x, y\}$ we say that $\{a, b\}$ is *matching-dominating* $\{x, y\}$ if there exists a bijective function $f : V \rightarrow V$ that (1) maps all voters from $V^{a \succ b}$ to $V^{x \succ y}$ and from $V^{b \succ a}$ to $V^{y \succ x}$, (2) for all voters v , $|v(ab)| \geq |(f(v))(xy)|$ and (3) there exists a voter v such that $|v(ab)| > |(f(v))(xy)|$.

Example 1. *Consider profile $\{v_1 : a \succ x \succ y \succ b, v_2 : a \succ y \succ x \succ b, v_3 : b \succ x \succ a \succ y, v_4 : b \succ y \succ a \succ x\}$ and pairs $\{a, b\}$ and $\{x, y\}$ of candidates, which gives sets $V^{a \succ b} = \{v_1, v_2\}$ and $V^{x \succ y} = \{v_1, v_3\}$. A matching $f = \{v_1 \rightarrow v_1, v_2 \rightarrow v_3, v_3 \rightarrow v_2, v_4 \rightarrow v_4\}$ meets condition (1). Recalling that, e.g., $(f(v_2))(xy) = v_3(xy) = 2$, it is easy to verify that $|v_i(ab)| \geq |(f(v_i))(xy)|$ for all $i \in \{1, 2, 3, 4\}$ (condition (2)). Finally, since $|v_1(ab)| = 3 > |(f(v_1))(xy)| = |v_1(xy)| = 1$, f meets condition (3). Hence, the pair $\{a, b\}$ dominates pair $\{x, y\}$.*

Definition 4 (Matching Domination). *Rule \mathcal{R} satisfies matching domination if matching-dominated pairs are never selected.*

The next axiom is inspired by the monotonicity notion from the literature of classical social choice. The idea is the following: If one pair is the most conflicting in a given profile, adding more conflict between the two candidates should not make another pair even more conflicting, and be selected instead of the original one.

Definition 5 (Conflict Monotonicity). *Rule \mathcal{R} is conflict monotonic if for each profile P and each selected pair $\{a, b\} \in \mathcal{R}(P)$, it holds that if we increase the distance between a and b by swapping one of them with its adjacent candidate in one of the votes, $\{a, b\}$ is still selected.*

Unfortunately, we can show that this property is incompatible with conflict consistency and matching domination.

Theorem 1. *No rule satisfies matching-domination, conflict consistency, and conflict monotonicity.*

Proof. Let f be a rule satisfying these 3 axioms, and consider the following profile: $\{v_1 : a \succ b \succ c \succ d, v_2 : b \succ a \succ d \succ c\}$. By conflict consistency, the only pairs that can be selected are $\{a, b\}$ and $\{c, d\}$. Assume that $\{a, b\}$ is selected. The proof if $\{c, d\}$ is selected is almost the same. Now, consider the profile $\{v_1 : a \succ b \succ c \succ d, v_2 : b \succ d \succ c \succ a\}$, in which we increased the conflict between a and b in the second vote by swapping a with its neighbors. By conflict monotonicity, $\{a, b\}$ should still be selected. Now, the values of $v(ab)$ are $(1, -3)$ and the values of $v(ad)$ are $(-3, 2)$. With the matching $f = \{v_1 \rightarrow v_2, v_2 \rightarrow v_1\}$, we obtain that $\{a, d\}$ is dominating $\{a, b\}$, thus $\{a, b\}$ cannot be selected by matching-domination. This is a contradiction. \square

This example highlights why conflict monotonicity is quite hard to achieve, but an extreme (and weaker) version of it can actually be satisfied by conflictual rules. The idea is that instead of increasing the conflict in only one ranking, we increase it in every ranking at once, and we increase it as much as we can in every ranking.

By *antagonization* of profile P with respect to pair $\{a, b\}$, denoted by P^{ab} , we refer to a profile P in which, in all votes from $V^{a \succ b}$ we shift a to the first position and b to the last position, and in all votes from $V^{b \succ a}$ we shift b to the first position and a to the last position. However, the relative order of all the other candidates remains the same.

Definition 6 (Antagonization Consistency). *Rule \mathcal{R} is antagonization consistent if, given a profile P and a selected pair $\{a, b\} \in \mathcal{R}(P)$, $\{a, b\}$ is also selected in P^{ab} .*

4 Conflictual Rules

After we made explicit what we expect from conflictual rules by the means of axioms, we can start searching for such rules.

In order to define the rules, we should first define what could be a *conflict*, in particular between two voters. Intuitively, there is a conflict induced by a pair of candidates $\{a, b\}$ between two voters if they disagree on the ordering between a and b . Moreover, the more distant a and b are in the rankings of the voters, the greater the conflict. Starting with this, we define the sum-conflict (conf^+) and the Nash-conflict (conf^\times) as follows:

Definition 7 (Pairwise conflict). *For $\circ \in \{+, \times\}$, let the conflict induced by a pair of candidates a and b between two votes v and v' be:*

$$\text{conf}_{v,v'}^\circ(a, b) = \begin{cases} 0 & \text{if } v(ab) \cdot v'(ab) > 0 \\ |v(ab)| \circ |v'(ba)| & \text{otherwise} \end{cases}$$

For instance, the Nash (resp. sum) pairwise conflict induced by a and b between votes $a \succ b \succ c$ and $b \succ c \succ a$ is $\text{conf}_{v,v'}^\times(a, b) = 1 \times 2 = 2$ (resp. $\text{conf}_{v,v'}^+(a, b) = 1 + 2 = 3$). By extension, for two candidates a and b , the conflict is defined as the sum of the pairwise conflict over all possible pairs of voters: $\text{conf}^\circ(a, b) = \sum_{v,v' \in V} \text{conf}_{v,v'}^\circ(a, b)$.

Then, we can define the rules that select the pairs of candidates that maximize this value. Therefore, we define the *MaxSumConflict* rule based on conf^+ , $\text{MaxSum}(P) = \text{argmax}_{a,b \in C} \text{conf}^+(a, b)$, which is equivalent to

$$\text{argmax}_{a,b \in C} \#V^{b \succ a} \sum_{v \in V^{a \succ b}} v(ab) + \#V^{a \succ b} \sum_{v \in V^{b \succ a}} v(ba).$$

Similarly, we define the *MaxNashConflict* rule, $\text{MaxNash}(P) = \text{argmax}_{a,b \in C} \text{conf}^\times(a, b)$, which is equivalent to

$$\text{argmax}_{a,b \in C} \sum_{v \in V^{a \succ b}} v(ab) \cdot \sum_{v \in V^{b \succ a}} v(ba).$$

Remark 1. *Looking at the second rule, it is tempting to define an alternative rule in which we maximize the sum $\sum_{v \in V^{a \succ b}} v(ab) + \sum_{v \in V^{b \succ a}} v(ba)$ instead of the product. However, such a rule does not satisfy conflict-consistency, and could elect a non-conflicting pair.⁴*

The following example demonstrates the defined rules.

Example 2. *Let profile P over 6 candidates be:*

$$\begin{aligned} a \succ x \succ c \succ d \succ y \succ b, \\ c \succ y \succ b \succ a \succ x \succ d. \end{aligned}$$

*The only conflicting pairs of candidates are $\{a, b\}$ and $\{x, y\}$. We have that $\text{conf}^+(a, b) = 6 = \text{conf}^+(x, y)$, $\text{conf}^\times(a, b) = 5$, and $\text{conf}^\times(x, y) = 3 \times 3 = 9$. Hence, while in the *MaxSum* rule both pairs tie, *MaxNash* clearly selects $\{x, y\}$.*

⁴Selecting only from conflicting pairs makes this rule conflict-consistent, yet somehow non-monotonic. Indeed, consider a non-conflicting pair A that scores higher than a conflicting pair B . After adding a voter equally increasing the pairs' scores and making A conflicting, the rule selects A , which is intuitively "less conflicting." So, we do not consider rules naturally failing conflict-consistency.

Proposition 2. *MaxSum and MaxNash satisfy reverse stability, conflict consistency, antagonization consistency, and matching-domination.*

Another approach is to select pairs of candidates $\{a, b\}$ that maximize the minimum number $\text{nonconf}(a, b)$ of swaps of adjacent candidates to make $\{a, b\}$ a non-conflicting pair in the profile P . We call the respective rule *MaxSwap* and remark that formally *MaxSwap* selects pairs $\{a, b\}$ such that:

$$\text{argmax}_{a,b \in C} \min \left(\sum_{v \in V^{a \succ b}} v(ab), \sum_{v \in V^{b \succ a}} v(ba) \right).$$

For demonstration recall Example 2, where *MaxSwap* would select $\{x, y\}$ as we need 3 swaps to make this pair non-conflicting, compared to one swap required for $\{a, b\}$.

Proposition 3. *MaxSwap rule satisfies reverse-stability, conflict consistency, and antagonization consistency. MaxSwap fails matching-domination.*

5 Interpretation

So far, we have introduced several conflictual rules and provided their axiomatic analysis. We now look further, beyond the somehow worst-case study that axioms usually offer. We want to understand what are the practical differences between our rules. To this end, we introduce various notions that help us interpret the rules' behavior.

5.1 Partitioning and Discrepancy

Intuitively, a pair of candidates is perfectly polarizing if there are two groups of equal sizes that have conflicting preferences on this pair, and all voters have very strong opinions of the candidates. We adapt these two features to describe our committees in the notions of social *partitioning ratio* and candidates' *discrepancy*. On the one hand, candidates a and b provide maximal partitioning ratio if exactly half of the voters prefer a to b . Formally, we define partitioning ratio $\alpha(a, b) \in [0, 1]$ for candidates a and b as

$$\alpha(a, b) = \frac{2}{n} \min(\#V^{a \succ b}, \#V^{b \succ a}).$$

On the other hand, candidates a and b have high discrepancy if every voter strongly prefers one candidate to the other. In particular, if each voter ranks either a or b first, and the other candidate last, then the pair $\{a, b\}$ has maximal discrepancy. Formally, we define the discrepancy $\beta(a, b) \in [0, 1]$ as

$$\beta(a, b) = \frac{1}{n \cdot (m - 1)} \sum_{v \in V} |v(ab)|.$$

Given a profile P , by $\alpha_{\max}(P)$ and $\beta_{\max}(P)$ we denote the $\max_{a,b} \alpha(a, b)$ and $\max_{a,b} \beta(a, b)$, respectively. First note that $\beta_{\max} \in (1/3, 1]$ since there is always a pair of candidates $\{a, b\}$ such that $\beta(a, b) > 1/3$. We defer complete calculations and the α and β values of characteristic profiles studied by Faliszewski *et al.* [2023] (i.e., identity, uniformity, and antagonism) to the full version [Delemazure *et al.*, 2024].

When selecting conflicting pairs, rules must do a trade-off between these two notions. To build our intuition, we look at

discrepancy and partitioning ratio in a specific case. With the new insights we obtain, we develop a family of voting rules based on α and β . Let us fix some candidates a and b and consider the case in which the value $|v(ab)|$ is the same for all voters $v \in V$; hence, we have $\beta(a, b) = |v(ab)|/(m-1)$ (for ease of presentation, we omit a and b and use α and β in this paragraph). Then, using some constant C_1 independent of a and b , we can express the sum-conflict between a and b as follows:

$$\begin{aligned} \text{conf}^+(a, b) &= \#V^{b>a} \sum_{v \in V^{a>b}} v(ab) + \#V^{a>b} \sum_{v \in V^{b>a}} v(ba) \\ &= 2 \#V^{b>a} \#V^{a>b} \beta \\ &= 2 \left(n \frac{\alpha}{2}\right) \left(n \left(1 - \frac{\alpha}{2}\right)\right) \beta = C_1 \alpha (2 - \alpha) \beta. \end{aligned}$$

Analogously, $\text{conf}^\times(a, b) = C_2 \alpha (2 - \alpha) \beta^2$ and $\text{nonconf}(a, b) = C_3 \alpha \beta$. The expressions clearly illustrate the tension between discrepancy and partitioning ratio: while some rules give more weight to pairs that divide the society more equally, other prefer those that have higher discrepancy.

The observed trade-off forms in fact a flexible framework for a family of rules covering the whole spectrum of possible behavior. Let \mathcal{R} be some rule endowed with a scoring function $S(\cdot, \cdot) : [0, 1]^2 \rightarrow \mathbb{R}_{\geq 0}$ that selects pairs of candidates $\{a, b\}$ which maximize value $S(\alpha(a, b), \beta(a, b))$. Each such rule \mathcal{R} naturally satisfies reverse-stability. Moreover, the rule satisfies conflict consistency if and only if $S(\alpha, \beta) = 0$ when $\alpha = 0$ and $S(\alpha, \beta) > 0$ otherwise. If S is strictly increasing with both α and β when $\alpha > 0$, then the rule based on S also satisfies antagonism-consistency and matching-domination. In particular, these conditions are met by the family of scoring functions $S(\alpha, \beta) = \alpha \beta^p$ for $p > 0$, that we named p -MaxPolarization rules (p -MaxPolar in short). Note that, the higher the value of p , the more weight is put on discrepancy in comparison to partitioning ratio. The natural dependence on discrepancy and partitioning ratio motivates us to include MaxPolar rules into our further analysis.

We visualize how the outcome of our rules depends on discrepancy and partitioning ratio using the following example. Consider a preference profile over four candidates in which a pair $\{a_1, b_1\}$ has maximal discrepancy and a pair $\{a_2, b_2\}$ has maximal partitioning ratio; hence, we have $\beta(a_1, b_1) = 1$ and $\alpha(a_2, b_2) = 1$. Understanding the interplay between discrepancy and partition ratio boils down to the question: Which pair is preferred in our scenario for different values of the non-fixed values $\beta(a_2, b_2)$ and $\alpha(a_1, b_1)$? For clarity, we once again consider our simplified case in which $\beta(a_2, b_2) = |v(a_2, b_2)|/m-1$ for all voters v . Figure 1 depicts how the preferred pair depends on the non-fixed values for various rules. For example, for 2-MaxPolar, the pair $\{a_1, b_1\}$ is preferred if $\alpha(a_2, b_2) \beta^2(a_2, b_2) < \alpha(a_1, b_1) \beta^2(a_1, b_1)$, which boils down to $\alpha(a_2, b_2) < \alpha(a_1, b_1)$ for our simplified scenario. Hence, the figure contains the plot of $\alpha(a_1, b_1) = \beta^2(a_2, b_2)$, where pair $\{a_1, b_1\}$ wins for each point over the curve. Similarly, for MaxNash, where $S(\alpha, \beta) = \alpha(2 - \alpha)\beta^2$ in our simplified scenario, we see the plot of $\alpha(2 - \alpha) = \beta^2$ (omitting the arguments for α and β for readability).

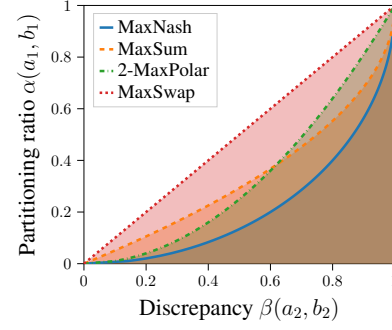


Figure 1: Area where $\{a_2, b_2\}$ is preferred to $\{a_1, b_1\}$ for different rules. We assume $\alpha(a_2, b_2) = 1$ and $\beta(a_1, b_1) = 1$.

5.2 Polarization Balance

Partitioning ratio and discrepancy enable us to interpret the rules we introduced in Section 4 in the special case described above. By this, we learn whether these rules give more importance to voters having strong opinions or to be evenly split. Yet, there is another phenomenon distinguishing these rules from the MaxPolar rules, and that is the discrepancy balance.

To see this, recall profile P from Example 2. Here, pairs $\{a, b\}$ and $\{x, y\}$ have the same partitioning ratio $\alpha = 1$ and discrepancy value $\beta = 3/5$. However, the discrepancy between a and b is not balanced between the two voters: a 's supporter is very extreme in its preferences, while b 's supporter is almost indifferent between the two. In comparison, the discrepancy between x and y is perfectly balanced. MaxPolar rules are agnostic to this phenomenon, while rules from Section 4 take it into account and would select pair $\{x, y\}$.

This leads us to introducing the third metric, the *discrepancy balance* γ . By $\mu(a, b) = \frac{\sum_{v \in V^{a>b}} v(ab)}{\#V^{a>b}}$ we denote the average discrepancy between a and b among supporters of a . For the discrepancy to be balanced, we want $\mu(a, b)$ and $\mu(b, a)$ to be as similar as possible. Thus, we define it as follows:

$$\gamma(a, b) = \min(\mu(a, b)/\mu(b, a), \mu(b, a)/\mu(a, b))$$

Measures α , β and γ are not independent from each other. For instance, $\beta(a, b) = 1$ and $\alpha(a, b) > 0$ implies $\gamma(a, b) = 1$ (but $\gamma(a, b) = 1$ does not imply anything for $\beta(a, b)$). Intuitively, the value of γ has more influence on rules like MaxSwap and MaxNash than on MaxSum, as the former are more egalitarian than the latter. This intuition is supported by the experiments presented in Section 6.

Note that α measures if the electorate is divided evenly between two candidates and γ measures if the discrepancy is balanced among supporters of each group. By combining these two ideas, we obtain the last measure, called *group discrepancy imbalance* ϕ , defined as:

$$\phi(a, b) = \frac{|\sum_{v \in V} v(ab)|}{\sum_{v \in V} |v(ab)|}.$$

It is 0 when the total discrepancies of the groups are equal and 1 when they are totally imbalanced (in that case $\alpha = 0$). As expected, contrary to γ , metric ϕ is sensitive to the size

of the groups. Particularly, if all rank differences are equal ($\beta = |v(ab)|/m-1$ for all voters v), then $\gamma = 1$, but $\phi = 1 - \alpha$.

With group discrepancy imbalance, we can rewrite two of our rules in the general case: MaxNash and MaxSwap. For MaxNash, we have that the score of the pair is proportional to $\beta^2(1 - \phi^2)$ and for MaxSwap to $\beta(1 - \phi)$, so they both care about this imbalance between the two groups. The other rules cannot be defined using ϕ in general, but in the particular case in which all rank differences are the same, we can simply replace α by $1 - \phi$, which gives $\beta(1 - \phi^2)$ for MaxSum.

Because of this, we introduce a new property that enable us to distinguish some rules.

Definition 8 (Balance Preference). *A rule satisfies balance preference if given two pairs $\{a, b\}$ and $\{x, y\}$, if there exists a perfect matching from voters to voters $\phi : V \rightarrow V$ such that $|v(ab)| = |\phi(v)(xy)|$ (i.e., their vectors of absolute rank differences are the same), and if $|\sum_{v \in V} v(ab)| < |\sum_{v \in V} v(xy)|$ (i.e., discrepancy between supporters of a and of b is more balanced), then $\{x, y\}$ cannot be selected.*

To better understand this axiom, consider the profile $P = \{2 \times x \succ a \succ b \succ y, 1 \times a \succ y \succ x \succ b, 1 \times b \succ y \succ x \succ a\}$. In P , voters are evenly divided between x and y , but x supporters are extreme while y supporters are quite indifferent. On the other hand, a is preferred to b by 3 voters, but they each have at least one extreme supporter. In this case, $\{x, y\}$ is more balanced in terms of group size, but $\{a, b\}$ is more balanced in terms of average polarization. More precisely, $\phi(x, y) = \frac{6-2}{6+2} = \frac{1}{2}$, and $\phi(a, b) = \frac{5-3}{5+3} = \frac{1}{4}$. Moreover, the pairs satisfy the condition of the balance preference axiom, so if a rule satisfies it, $\{x, y\}$ should not be selected. This is why MaxSwap and MaxNash select $\{a, b\}$. However, $\beta(a, b) = \beta(x, y)$ and $\alpha(a, b) < \alpha(x, y)$, so all MaxPolar rules select $\{x, y\}$, and we can also show that MaxSum selects $\{x, y\}$. More generally, we can state the following proposition.

Proposition 4. *MaxNash and MaxSwap satisfy balance preference and MaxSum and MaxPolar rules fail it.*

6 Experiments

The goals of our experiments are to compare conflictual rules to traditional ones, and to compare the introduced rules with each other. The full version of our paper [Delemazure *et al.*, 2024] contains a detailed description of the used data as well as further experimental results for the omitted synthetic data, rules, and polarization metrics we introduced.

6.1 Comparison with Standard Committee Rules

First, we compare MaxNash and two well-known committee rules Chamberlin-Courant and Borda, which respectively aim at achieving diversity and individual excellence. We chose MaxNash as a representative of all conflictual rules because all of them gave qualitatively same results.

Under *Borda* rule, each voter assigns $m - 1$ points to her favorite candidate, $m - 2$ points to the second one, $m - 3$ to the third one, and so on. Finally, we select the candidates with the highest scores. Under the *Chamberlin-Courant* rule (CC) and $k = 2$, each voter assigns $m - \min(v(x), v(y))$ points to the pair of candidates $\{x, y\}$, and the pair with the




Figure 2: Distribution of the positions of the winning candidates for different rules and distributions of positions. Each pair of colored points correspond to the winners of a single election.

highest score is selected. Thus, only the preferred candidate in the pair is taken into account.

To get the intuition about differences between these rules and conflictual ones, we use the 2D-Euclidean framework [Elkind *et al.*, 2017a]. Here, voters and candidates are represented by points in the 2D-Euclidean space, and the preferences of voters are based on their distances to candidates: a voter v prefers a to b if v is closer to a than to b .

We sampled the positions on $[0, 1]^2$ of all voters and candidates using: (1) the uniform distribution and (2) the normal distribution centered in $(0.5, 0.5)$. Figure 2 presents the results for 10,000 sampled instances. For each instance, we marked on the plane the two selected candidates, highlighting the pairs for three random instances for reference.

Supporting the axiomatic analysis, the results show that the conflictual rules behave significantly different from the classical ones. Indeed, they select candidates that are far away from the center of the plane, while all other rules select those that are around the center. We also computed partitioning ratio and discrepancy for all pairs of candidates (see the long version). Their values confirm that, on average, the pairs selected by Borda and CC have lower values than those selected by the conflictual rules. As expected, this tendency is particularly strong for discrepancy — in our scenario, if two candidates are close to each other on the plane, then they are relatively close to each other in every ranking.

6.2 Comparisons of Conflictual Rules

We include next experimental results solely for the real-life datasets, deferring the discussion of the results for synthetic models to the paper’s full version [Delemazure *et al.*, 2024].

Our goal now is to compare the conflictual rules with each other, particularly focusing on MaxSwap, MaxNash, MaxSum, and 2-MaxPolar. We study 4 datasets: (i) preferences over 11 candidates gathered for experiments during French presidential elections in 2017 and 2022 [Bouveret *et al.*, 2018; Delemazure and Bouveret, 2022], which we expect to be conflictual, and that we discuss in more extent in Section 6.3, (ii) preferences over 10 sushi types [Kamishima, 2003] and (iii) juries’ ranking of contestant performances in figure skating competitions, from Preflib [Mattei and Walsh, 2013], which we expect to be less conflictual.

In this set of experiments, we sampled rankings from the




Figure 3: Metrics values of the selected pairs of candidates for real data. Each pair $\{a, b\}$ is a ball at point at $(\alpha(a, b), \beta(a, b))$, with radius $\gamma(a, b)$, and colored redder with increasing $\phi(a, b)$. Each plot represents 1000 profiles of 100 voters and 10 candidates.

dataset over a subset of candidates in the case of real data or from a probability model in the case of synthetic data. This way, we always generated the same number $n = 100$ of voters and $m = 10$ of candidates. For each generated profile, we simulated our rules. Then, we compared the values of the select pairs’ polarization metrics (partitioning ratio, discrepancy, discrepancy balance and group discrepancy imbalance). We repeated this experiment for 1000 random profiles. Figure 3 shows each selected pair $\{a, b\}$ as a dot with coordinate $(\alpha(a, b), \beta(a, b))$; its size is proportional to $\gamma(a, b)$ and the redder the dot the higher the value of $\phi(a, b)$. The first row depicts uniformly sampled pairs of candidates, providing an estimate of the actual distribution of the pairs in the datasets.

The first observation, based on the first row of the figure in which we plot points associated to random pairs of candidates, is that the three datasets are very different in terms of conflict. The most conflictual seems to be the political one, with most pairs having both a high partitioning ratio α and discrepancy β . The sushi dataset is already less conflictual: Some pairs have higher α , but their β is quite low, and conversely. This is even worse for the ice skating dataset, as rankings rely a lot on the quality of participants’ performance.

The experiments confirm our theoretical analysis of the for-

	2017		2022	
MaxSwap	Far-left	Far-right	Far-left	Far-right
MaxNash	Socialist	Far-right	Left	Far-right
MaxSum	Socialist	Far-right	Far-left	Far-right
2-MaxPolar	Far-left	Far-right	Far-left	Far-right
Borda	Left	Liberal	Left	Green
CC	Left	Conservative	Green	Far-right

Table 2: Selected pairs for different rules.

mula based on α and β , that we obtained assuming a simplified special case. In particular, we clearly see that MaxNash puts the most emphasis on β , while MaxSwap prefers an equal division between a and b supporters (i.e., maximizing α). Yet, group discrepancy imbalance also influences the results for MaxSwap. On the contrary, 2-MaxPolar ignores the discrepancy balance, which explains why it might return pairs of candidates with smaller and bluer dots.

6.3 Results on a Conflictual Election

We now consider the political data, featuring the most conflictual preferences, that was gathered by the *Voter Autrement initiative* during the 2017 and 2022 presidential elections.⁵

Table 2 summarizes the candidates selected by each rule and compares them to classic voting rules outcomes. We chose to put the political labelization of candidates instead of their names. While the non-conflictual rules return two *popular* candidates (e.g., the main candidates of each the left and the right side of the spectrum), conflictual rules tend to select at least one extreme candidate, if not two. Rules that put more emphasis on discrepancy like MaxNash would select more well-known candidates, sacrificing equal dividing of the society, as voters have strong preferences on them. On the contrary, MaxSwap might select fewer well-known candidates, but who divide the society more evenly. MaxSum and 2-MaxPolar lie in between these two extremes.

7 Conclusions and Future Work

We proposed and analyzed rules that aim at selecting the most conflicting candidates, and have shown that these rules fundamentally differ from the standard ones. Together with the proposed metrics, our rules allow us to better understand the structures causing conflict in the electorate.

A natural (and fairly non-trivial) follow up direction is to analyze extending these rules and axioms to selecting more than just two candidates. Since our rules provide a score for each pair of candidates, one polynomial-time computable possibility would be to simply select a committee maximizing the smallest score among all pairs of candidates. Nonetheless, approaches evaluating the whole committee at the same time, rather than pairs of candidates might turn out to be superior.

⁵The data comes from an online poll that asked voters to try alternative voting methods using rankings. We removed partial rankings (voters could sometimes rank only their top k alternatives) and re-weighted the voters (based on their actual vote in the election) getting a distribution closer to the actually observed one. Effectively, we obtained $n = 5755$ (resp. 412) voters and $m = 11$ (resp. 12) candidates for the 2017 (resp. 2022) dataset.

Acknowledgements

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 101002854), and from 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). The research presented in this paper is supported in part from the funds assigned by Polish Ministry of Science and Technology to AGH University.



References

- [Alcalde-Unzu and Vorsatz, 2013] J. Alcalde-Unzu and M. Vorsatz. Measuring the cohesiveness of preferences: An axiomatic analysis. *Social Choice and Welfare*, 41(4):965–988, 2013.
- [Alcalde-Unzu and Vorsatz, 2016] J. Alcalde-Unzu and M. Vorsatz. Do we agree? Measuring the cohesiveness of preferences. *Theory and Decision*, 80(2):313–339, 2016.
- [Bouveret *et al.*, 2018] S. Bouveret, R. Blanch, A. Baujard, F. Durand, H. Igersheim, J. Lang, A. Laruelle, J.-F. Laslier, I. Lebon, and V. Merlin. Voter autrement 2017 - online experiment. Dataset and companion article on Zenodo, July 2018.
- [Can *et al.*, 2017] B. Can, A. Ozkes, and T. Storcken. Generalized measures of polarization in preferences. Technical report, HAL, 2017.
- [Chamberlin and Courant, 1983] B. Chamberlin and P. Courant. Representative deliberations and representative decisions: Proportional representation and the Borda rule. *American Political Science Review*, 77(3):718–733, 1983.
- [Colley *et al.*, 2023] R. Colley, U. Grandi, C. Hidalgo, M. Macedo, and C. Navarrete. Measuring and controlling divisiveness in rank aggregation. In *Proceedings of IJCAI-2023*, pages 2616–2623, 2023.
- [Debord, 1992] B. Debord. An axiomatic characterization of Borda’s k -choice function. *Social Choice and Welfare*, 9(4):337–343, 1992.
- [Delemazure and Bouveret, 2022] T. Delemazure and S. Bouveret. Voter autrement 2022 - online experiment (“Un autre vote”). Dataset and companion article on Zenodo, April 2022.
- [Delemazure *et al.*, 2024] T. Delemazure, Ł. Janeczko, A. Kaczmarczyk, and S. Szufa. Selecting the most conflicting pair of candidates. Technical Report arXiv.2405.05870 [cs.GT], arXiv.org, 2024.
- [Elkind *et al.*, 2017a] E. Elkind, P. Faliszewski, J. Laslier, P. Skowron, A. Slinko, and N. Talmon. What do multiwinner voting rules do? An experiment over the two-dimensional euclidean domain. In *Proceedings of AAAI-2017*, pages 494–501, 2017.
- [Elkind *et al.*, 2017b] E. Elkind, P. Faliszewski, P. Skowron, and A. Slinko. Properties of multiwinner voting rules. *Social Choice and Welfare*, 48(3):599–632, 2017.
- [Faliszewski *et al.*, 2017] P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Multiwinner voting: A new challenge for social choice theory. In U. Endriss, editor, *Trends in Computational Social Choice*. AI Access Foundation, 2017.
- [Faliszewski *et al.*, 2019] P. Faliszewski, P. Skowron, S. Szufa, and N. Talmon. Proportional representation in elections: STV vs PAV. In *Proceedings of AAMAS-2011*, pages 1946–1948, 2019.
- [Faliszewski *et al.*, 2023] P. Faliszewski, A. Kaczmarczyk, K. Sornat, S. Szufa, and T. Wąs. Diversity, agreement, and polarization in elections. In *Proceedings of IJCAI-2023*, pages 2684–2692, 2023.
- [Hashemi and Endriss, 2014] V. Hashemi and U. Endriss. Measuring diversity of preferences in a group. In *Proceedings of ECAI-2014*, pages 423–428, 2014.
- [Kamishima, 2003] T. Kamishima. Nantonac collaborative filtering: Recommendation based on order responses. In *Proceedings of KDD-03*, pages 583–588, 2003.
- [Mattei and Walsh, 2013] N. Mattei and T. Walsh. Preflib: A library of preference data [HTTP://PREFLIB.ORG](http://preflib.org). In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT 2013)*, Lecture Notes in Artificial Intelligence. Springer, 2013.
- [Monroe, 1995] Burt L. Monroe. Fully proportional representation. *The American Political Science Review*, 89(4):925–940, 1995.
- [Peters *et al.*, 2021] D. Peters, G. Pierczyński, and P. Skowron. Proportional participatory budgeting with additive utilities. In *Proceedings of NeurIPS-2021*, pages 12726–12737, 2021.
- [Skowron *et al.*, 2017] P. Skowron, M. Lackner, M. Brill, D. Peters, and E. Elkind. Proportional rankings. In *Proceedings of IJCAI-2017*, pages 409–415, 2017.
- [Skowron *et al.*, 2019] P. Skowron, P. Faliszewski, and A. Slinko. Axiomatic characterization of committee scoring rules. *Journal of Economic Theory*, 180:244–273, 2019.