Comparing Ways of Obtaining Candidate Orderings from Approval Ballots

Théo Delemazure[1], Chris Dong[2],
Dominik Peters[1], and Magdaléna Tydrichová[3]

[1]CNRS, LAMSADE, Université Paris Dauphine - PSL
[2]Technical University of Munich
[3]CentraleSupélec, Paris Saclay University

Abstract

To understand and summarize approval preferences and other binary evaluation data, it is useful to order the items on an *axis* which explains the data. In a political election using approval voting, this could be an ideological left-right axis such that each voter approves adjacent candidates, an analogue of single-peakedness. In a perfect axis, every approval set would be an interval, which is usually not possible, and so we need to choose an axis that gets closest to this ideal. The literature has developed algorithms for optimizing several objective functions (e.g., minimize the number of added approvals needed to get a perfect axis), but provides little help with choosing among different objectives. In this paper, we take a social choice approach and compare 5 different axis selection rules axiomatically, by studying the properties they satisfy. We establish some impossibility theorems, and characterize (within the class of scoring rules) the rule that chooses the axes that maximize the number of votes that form intervals, using the axioms of ballot monotonicity and resistance to cloning. Finally, we study the behavior of the rules on data from French election surveys, on the votes of justices of the US Supreme Court, and on synthetic data.

1. Introduction

In this paper, we work on analyzing and understanding binary evaluation data. One source of such data is approval voting, and for simplicity we will generally use voting terminology, so that each evaluator is a *voter* who *approves* the candidates that have been assigned an evaluation of 1. Our aim is to obtain an ordering of the candidates (an *axis*) which we interpret to *perfectly depict* the data if every voter approves an interval of the axis (an approval version of single-peaked preferences). Usually, such axes do not exist, so we study rules that, given an approval profile, find the axes that best approximate the interval structure and that thereby provide a good (ordinal) one-dimensional embedding of the profile. Such rules have many applications for understanding and visualizing data, as well as direct use-cases where the axis itself plays a key role:

- *Ordering political candidates and parties.* In politics, if voters are asked to approve candidates, an axis corresponds to an ideological ordering of the candidates from left-wing to right-wing. In France, the major pollsters use many different axes, which they apparently construct ad hoc; our rules will find an axis in a principled way.

- *Ordering members of parliament.* Once elected, we can order politicians by interpreting each bill as a "voter" who approves those members who supported it.

- *Archaeological seriation.* A well-established approach in archaeology for ordering artefacts by their age is to let features that were temporarily "in fashion" (e.g., drawing styles) approve artefacts (Petrie, 1899; Baxter, 2003). In the true ordering by age, each feature is likely to induce an interval.

- *Scheduling*. A conference organizer could design a schedule by polling attendees about which talks they wish to see and arrange them so attendees can join for consecutive talks. A different way of using our rules to schedule the conference is for key terms to approve the papers that mention them, leading to a thematically coherent ordering of the talks.

Algorithmically, our task is well studied. To check whether a perfect axis exists, one needs to check whether the 0/1 approval matrix has the *consecutive ones property* (C1P), which can be done in linear time (Booth and Lueker, 1976). However, in all the applications discussed above, the 0/1 matrices are likely to only approximately satisfy C1P. The problem of finding an axis that makes as many votes as possible into an interval is NP-complete and already appears in the book of Garey and Johnson (1979, SR14) together with several similar problems about recognizing almost-C1P matrices like minimizing the number of approvals to add to satisfy C1P. However, this complexity theoretic work does not tell us which of these objective functions "work best".

We provide a framework for answering this question, using the axiomatic method of social choice. We interpret different objective functions as *rules* that take an approval profile as input and decide on an axis. We will compare these rules by identifying properties that they satisfy or fail. Given a context where some properties seem particularly desirable, this will help with selecting a good objective function.

The five rules that are the protagonists of our paper are:

- *Voter Deletion.* Minimize the number of votes that are not intervals of the axis.
- *Minimum Flips.* Minimize the number of approvals that need to be added or removed from ballots to make all votes intervals of the axis.
- *Ballot Completion.* Minimize the number of approvals that need to be added to ballots to make all votes intervals of the axis.
- *Minimum Swaps.* Minimize the average number of swaps within the axis that are needed to turn votes into intervals of the axis.

- *Forbidden Triples.* Minimize the total size of holes in a vote, weighted by how many approved candidates they separate.

On a high level, we find that Voter Deletion and Ballot Completion satisfy a desirable monotonicity property, while the last two rules use more information contained in the profile. We do not identify any positive features of Minimum Flips.

Besides introducing the rules and the axioms, we also prove an impossibility result saying that no scoring rule (which are rules that optimize a voter-additive objective function) can simultaneously satisfy two versions of the "clones" principle that identical candidates should be treated in reasonable ways. We also establish a characterization result that the Voter Deletion rule is the unique scoring rule that satisfies one of these versions (clone resistance) as well as ballot monotonicity.

We conclude the paper by applying our rules to different datasets, including French election surveys (ordering candidates left to right), votes of the justices of the US Supreme Court (ordering justices from conservative to progressive), and synthetic datasets. The simulations show how our rules differ, which perform best, and how they compare to rules that are based on rankings rather than approvals.

2. Related Work

The work of Escoffier et al. (2021), extended in the thesis of Tydrichová (2023, Sec. 4.4), is closest to ours, as it compares different methods for finding axes that make a profile of rankings of the candidates *nearly single-peaked*. Single-peaked ranking preferences (Black, 1948) are frequently studied in social choice because they can avoid impossibility theorems and computational hardness (Elkind et al., 2022). Escoffier et al. (2021) focus on computational complexity, but also consider axiomatic properties satisfied by different objective functions. However, they do not provide axiomatic characterization or impossibility results, and our experiments suggest that the approval approach may lead to better axes than the ranking approach. Nearly single-peaked preferences are well-studied algorithmically, both their recognition (Bredereck et al., 2016; Erdélyi et al., 2017; Elkind and Lackner, 2014) and their impact on the winner determination problem of computationally hard voting rules (Misra et al., 2017; Chen et al., 2023).

For approval ballots, structured preferences are studied by Elkind and Lackner (2015), who say that a profile satisfies *Candidate Interval* (CI) if there is a perfect axis for it [see also Terzopoulou et al., 2021]. Dietrich and List (2010) discuss a similar concept in judgement aggregation. The study of the algorithmic problem of recognizing profiles that are *nearly* C1P goes back to Booth (1975) and has received thorough attention since (e.g., Hajiaghayi and Ganjali, 2002; Tan and Zhang, 2007; Chauve et al., 2009; Dom et al., 2010; Narayanaswamy and Subashini, 2015). Our study uses axioms and experiments instead of computational complexity, and focusses on selecting a good axis rather than measuring nearly single-peakedness.

3. Preliminaries

Let C be a set of m *candidates*, and V a set of n *voters*. An *approval ballot* is a non-empty subset of candidates $A \subseteq C$. An *approval profile* P is a collection of n approval ballots $P = (A_i)_{i \in V}$. We denote by $\mathcal{P}$ the set of all approval profiles.

An *axis* $\lhd$ is a linear order of the candidates so that $a \lhd b$ means that candidate a is on the left of b on the axis. We write $a \unlhd b$ if $a \lhd b$ or $a = b$. For brevity, we will sometimes omit the $\lhd$ and write *abc* for the axis $a \lhd b \lhd c$. Let $\mathcal{A}$ be the set of all axes over C. The direction of an axis is irrelevant, so we will informally treat the axes *abc* and *cba* as being the same axis.

An approval ballot A_i is an *interval* of an axis $\lhd$ if for any two candidates $a, b \in A_i$ and every c such that $a \lhd c \lhd b$, we have $c \in A_i$. If instead $c \notin A_i$, we say that c is an *interfering candidate*. A profile P is *linear* if there exists an axis $\lhd$ such that all approval ballots in P are intervals of $\lhd$. We also say that this axis $\lhd$ is *consistent* with the profile P. We write $\text{con}(P) \subseteq \mathcal{A}$ for the set of all axes consistent with P.

For an approval ballot A and an axis $\lhd = c_1 c_2 \dots c_m$ with candidates relabeled by their axis position, we denote by $x_{A,\lhd} = (x^1_{A,\lhd}, \dots, x^m_{A,\lhd})$ the *approval vector* where $x^i_{A,\lhd} = 1$ if $c_i \in A$ and 0 otherwise. For instance, for the axis $\lhd = abcd$ and ballot $A = \{b, c\}$, we get the vector $(0, 1, 1, 0)$, while $A' = \{a, d\}$ gives the vector $(1, 0, 0, 1)$ (which has two interfering candidates). The *approval matrix* of a profile $P = (A_i)_i$ has $x_{A_i,\lhd}$ as its ith row. Thus, its (i, j)-entry is equal to 1 iff $c_j \in A_i$. Note that a profile is linear if and only if its approval matrix (for any axis $\lhd$) satisfies the *consecutive one property* (or C1P, see the survey by Dom (2009)), i.e., its columns can be reordered such that in each row, the "1"s form an interval.

An *axis rule* f is a function that takes as input an approval profile P and returns a non-empty set of axes $f(P) \subseteq \mathcal{A}$, such that for each $\lhd$ in $f(P)$ its reverse axis $\overleftarrow{\lhd}$ is also in $f(P)$, encoding the idea that the direction of the axis does not matter.

In this paper, we will focus on the family of *scoring rules*. Let $s : 2^C \times \mathcal{A} \to \mathbb{R}_{\geqslant 0}$ be a *cost function*, indicating the cost $s(A_i, \lhd)$ that a ballot $A_i \in P$ incurs when the axis $\lhd$ is chosen. By summing up these costs, we get the cost $s(P, \lhd) = \sum_{A_i \in P} s(A_i, \lhd)$ of an axis $\lhd$ for the profile P. An axis rule f is a *scoring rule* if there is a cost function s_f such that $f(P) = \arg\min_{\lhd \in \mathcal{A}} s_f(P, \lhd)$ for all profiles P.

A focus on the class of scoring rules can be justified as an analogue to scoring rules in voting theory, in that every scoring rule satisfies the *reinforcement* axiom (Young, 1975) which roughly says that if f chooses the same axis in two profiles P_1 and P_2, then it also chooses that axis in the combined profile $P_1 + P_2$. However, providing an axiomatic characterization of this class appears to be difficult since the neutrality axiom turns out to be quite weak in our setting. Another motivation for scoring rules is their natural interpretation as *maximum likelihood estimators* when there is a ground truth axis, as observed by Conitzer et al. (2009) in the voting setting. To see the connection, let $\lhd$ be the ground truth axis, and suppose voters obtain their approval ballots A_i i.i.d. from a probability distribution $\mathbb{P}(A_i \mid \lhd)$ (where intuitively ballots are more likely the closer they are to forming an interval of $\lhd$). Then, the likelihood of a profile P is $\mathbb{P}(P \mid \lhd) = \prod_i \mathbb{P}(A_i \mid \lhd)$. To find the axis inducing maximum likelihood, we solve $\text{MLE}(P) := \arg\max_{\lhd} \mathbb{P}(P \mid \lhd) = \arg\min_{\lhd} - \sum_i \log(\mathbb{P}(A_i \mid \lhd))$, which is a scoring rule with costs $s_f(A_i, \lhd) = -\log(\mathbb{P}(A_i \mid \lhd))$.

4. Axis Rules

In this section, we introduce five scoring rules. Many are inspired by objective functions proposed for near single-peakedness (Faliszewski et al., 2014; Escoffier et al., 2021).

The first and simplest rule is called *Voter Deletion (VD)*:

Voter Deletion This rule returns the axes that minimize the number of ballots to delete from the profile P in order to become consistent with it. This rule is a scoring rule based on the cost function s_{VD} such that $s_{\text{VD}}(A, \lhd) = 0$ if A is an interval of $\lhd$, and 1 otherwise.

The idea behind this rule is that perhaps some "maverick" voters are "irrational", and should hence be disregarded. The aim is to delete as few maverick voters as possible.

However, VD does not measure the *degree of incompatibility* of a given vote with an axis. For example, VD does not distinguish ballots that miss just one candidate to be an interval, and an approval ballot in which only the two extreme candidates of the axis are approved. For this reason, more gradual rules might do better.

The first rule in this direction is *Minimum Flips* (MF) which changes ballots by adding and removing candidates from them.

Minimum Flips This rule returns the axes that minimize the total number of candidates that need to be added/removed from approval ballots in order to make the profile linear. It is the scoring rule based on:

$$s_{\mathrm{MF}}(A, \lhd) = \min_{x,y \in A \,:\, x \unlhd y} |\{z \in A : z \lhd x \text{ or } y \lhd z\}| + |\{z \notin A : x \lhd z \lhd y\}|.$$

MF finds for each vote A_i the interval ballot closest to A_i in Hamming distance, with that distance being the cost of $\lhd$. Equivalently, the rule finds the linear profile of minimum total Hamming distance to the input profile, and returns its axes.

In many applications, adding approvals seems better motivated than removing them. For example, a voter i might not approve a candidate c because i does not know who c is; fixing this error corresponds to adding a candidate. The *Ballot Completion* (BC) rule implements this thought.

Ballot Completion This rule returns the axes $\lhd$ that minimize the number of candidates to add to approval ballots to make the profile consistent with it. It is the scoring rule based on:

$$s_{\mathrm{BC}}(A, \lhd) = |\{b \notin A : a \lhd b \lhd c \text{ for some } a, c \in A\}|.$$

Thus, given a ballot A and an axis $\lhd$, this rule counts all interfering candidates with respect to A and $\lhd$. To see the difference between MF and BC, observe that $s_{\mathrm{BC}}(\{a, d\}, abcd) = 2$ as we need to add b and c to obtain an interval, while $s_{\mathrm{MF}}(\{a, d\}, abcd) = 1$ as we can just remove a.

In the approval context, BC is the only rule we know of that has already been used in the literature to find an underlying political axis of voters, on the data of experiments conducted during the 2012 and 2017 French presidential elections (Lebon et al., 2017; Baujard and Lebon, 2022). The axes found by BC were close to the orderings discussed in the media.

The *Minimum Swaps* (MS) rule modifies the *axis* rather than the ballots. Given an approval ballot A, the MS rule asks how many candidate swaps we need to perform in an axis $\lhd$ until A becomes an interval of it: the cost $s_{\mathrm{MS}}(A, \lhd)$ is the minimum Kendall-tau distance between $\lhd$ and an axis $\lhd'$ (the number of swaps of adjacent candidates needed to go from $\lhd$ to $\lhd'$) such that A is an interval of $\lhd'$. For instance, $s_{\mathrm{MS}}(\{a, d\}, abcd) = 2$ because we need to have a next to d on any axis consistent with $\{a, d\}$, and we need at least two swaps to obtain this.

Minimum Swaps This scoring rule uses the cost function

$$s_{\mathrm{MS}}(A, \lhd) = \textstyle\sum_{x \notin A} \min(|\{y \in A \colon y \lhd x\}|, |\{y \in A \colon x \lhd y\}|).$$

To see why this formula implements our swapping description, note that to modify the axis $\lhd$ such that A becomes an interval of it, we need to "push outside" all $x \notin A$ such that there exist $y, z \in A$ with $y \lhd x \lhd z$. We can either push x to the left side or to the right side, and thus we must swap x with at least all candidates $y \in A$ to its right or to its left, which gives $s_{\mathrm{MS}}(A, \lhd) \geqslant \sum_{x \notin A} \min(|\{y \in A : y \lhd x\}|, |\{y \in A : x \lhd y\}|)$.

We now prove by induction that this bound is reached and $s_{\mathrm{MS}}(A, \lhd) \leqslant \sum_{x \notin A} \min(|\{y \in A : y \lhd x\}|, |\{y \in A : x \lhd y\}|)$. If there are 0 or 1 interfering candidates, the formula is trivially satisfied. Let us assume there are $k \geqslant 2$ interfering candidates. Consider x_l the "left-most" interfering candidate, i.e. there are no interfering candidate x such that $x \lhd x_l$. Similarly, define x_r the "right-most" interfering candidate. Then, push the one for which we need the minimal number of swaps to get outside of the interval. For instance, if $|\{y \in A, y \lhd x_l\}| < |\{y \in A, x_r \lhd y\}|$, we successively swap x_l with all candidates $y \in A$ such that $y \lhd x_l$ until x_l is pushed outside of the interval. It is easy to see that we required $\min(|\{y \in A, y \lhd x_l\}|, |\{y \in A, x_l \lhd y\}|)$ swaps, and only the position of x_l changed, since it was part of all the swaps. Thus, each interfering candidate (except x_l) have the same number of approved candidates to its right and to its left after than before x_l was moved. If $|\{y \in A, y \lhd x_l\}| > |\{y \in A, x_r \lhd y\}|$, we use the same reasoning, with x_r instead of x_l. The induction concludes the proof of the formula.

Our last rule is *Forbidden Triples* (FT), inspired by a proposal for rankings by Escoffier et al. (2021). It counts the number of violations of the interval condition as defined in Section 3.

Forbidden Triples This scoring rule uses the cost function

$$s_{\mathrm{FT}}(A, \lhd) = |\{(x, y, z) : x, z \in A, y \notin A, x \lhd y \lhd z\}|.$$

Note that there is one forbidden triple for each interfering candidate, and each pair of candidates respectively on its left and its right, so $s_{\mathrm{FT}}(A, \lhd) = \sum_{x \notin A} |\{y \in A : y \lhd x\}| \times |\{y \in A : x \lhd y\}|$. For instance, we have $s_{\mathrm{FT}}(\{a, b, d, e\}, abcde) = 2 \times 2 = 4$ while $s_{\mathrm{FT}}(\{a, b, c, e\}, abcde) = 3 \times 1 = 3$. Intuitively, this rule looks at the holes in a vote, with larger holes separating many approved candidates counting more.

Our five rules can be related via a chain of inequalities.

Proposition 1 *For any A and $\lhd$, we have $s_{\mathrm{VD}}(A, \lhd) \leqslant s_{\mathrm{MF}}(A, \lhd) \leqslant s_{\mathrm{BC}}(A, \lhd) \leqslant s_{\mathrm{MS}}(A, \lhd) \leqslant s_{\mathrm{FT}}(A, \lhd)$.*

Proof. To see $s_{\mathrm{VD}}(A, \lhd) \leqslant s_{\mathrm{MF}}(A, \lhd)$, note that if A is not an interval of $\lhd$ then $s_{\mathrm{VD}}(A, \lhd) = 1$ and at least one candidate must be flipped to make A an interval of $\lhd$, so $s_{\mathrm{MF}}(A, \lhd) \geqslant 1$. If A is an interval then $s_{\mathrm{MF}}(A, \lhd) = s_{\mathrm{VD}}(A, \lhd) = 0$.

We have $s_{\mathrm{MF}}(A, \lhd) \leqslant s_{\mathrm{BC}}(A, \lhd)$ because in MF we can add and remove approvals, but in BC we can only add approvals. Thus, if a solution is optimal for BC, it is also a solution for MF with the same cost (but not necessarily optimal).

Finally, observe that for any interfering candidate x on A, $\min(|\{y \in A : y \lhd x\}|, |\{y \in A : x \lhd y\}|) \geqslant 1$. Moreover, as these are all natural numbers, $\min(|\{y \in A : y \lhd x\}|, |\{y \in A : x \lhd y\}|) \leqslant |\{y \in A : y \lhd x\}| \times |\{y \in A : x \lhd y\}|$. Thus, $s_{\mathrm{BC}}(A, \lhd) \leqslant s_{\mathrm{MS}}(A, \lhd) \leqslant s_{\mathrm{FT}}(A, \lhd)$ by the definitions of these rules. □

We say that two axis rules f_1 and f_2 are *equivalent* if for all profiles P we have $f_1(P) = f_2(P)$. Note that if $n \leqslant 2$ or $m \leqslant 2$, every profile is linear. Moreover, if there are $m = 3$ candidates, all the rules defined in this section are equivalent (as there is only one non-interval approval vector, so the only possible costs are 0 and 1). If there are $m = 4$ candidates, VD and MF are equivalent and BC and MS are equivalent. This is because the respective cost functions coincide for $m \leqslant 4$, which does not remain true for $m \geqslant 5$. Indeed, for $m \geqslant 5$, the rules are pairwise non-equivalent. Example 1 shows a profile with $m = 4$ for which VD, BC and FT all select different axes. We give another profile in Appendix A.1 with $m = 7$ for which no two rules select the same axes.

Example 1 *Consider the profile $P = (4 \times \{b, c, d\}, 4 \times \{a, b\}, 3 \times \{a, d\}, 1 \times \{a, c\}, 1 \times \{b, c\})$. On this profile, all rules agree that $a \lhd b \lhd c$, but they disagree on the position of d. Indeed, $\lhd_1 = abc\underline{d}$ is optimal for VD and MF, $\lhd_2 = \underline{d}abc$ for BC and MS, and $\lhd_3 = a\underline{d}bc$ and $\lhd_4 = ab\underline{d}c$ for FT. Table 1 shows the profile aligned according to the four possible axes. One can easily see that among these axes (1) the axis $\lhd_1$ on the left minimizes the VD cost with only 4 non-interval ballots, (2) the axis $\lhd_2$ in the middle minimize the BC cost with 5 red circles and (3) the axes $\lhd_3$ and $\lhd_4$ on the right minimizes the FT cost with 6 forbidden triplets.*

	a	b	c	$\underline{d}$	$\underline{d}$	a	b	c	a	$\underline{d}$	b	c	a	b	$\underline{d}$	c
4×		×	×	×	×	•	×	×		×	×	×		×	×	×
4×	×	×				×	×		×	•	×		×	×		
3×	×	•	•	×	×	×			×	×			×	•	×	
1×	×	•	×			×	•	×	×	•	•	×	×	•	•	×
1×		×	×				×	×			×	×		×	•	×

Table 1: Profile of Example 1 on 4 different axes. Red circles indicate interfering candidates.

As we already mentioned, problems about recognizing matrices that are almost C1P have long been known to be NP-complete. Hardness of VD and BC is explicitly known (see Booth (1975)), and the reductions only uses approval sets of size 2. The results for other rules directly follows that they are equivalent to either VD or BC when $\max_i |A_i| = 2$ (See Appendix A.2 for a detailed proof.)

Theorem 1 *The VD, MF, BC, MS and FT rules are NP-complete to compute, even if* $\max_i |A_i| = 2$.

A lot of other axis rules could be defined. However, in this paper, we focus on the five rules introduced above, and leave the study of potential other rules to further research. In particular, we think that greedy variants of the rules we introduced are of interest to circumvent computational hardness.

5. Axiomatic Analysis

In this section, we conduct an axiomatic analysis of the rules we introduced. Table 2 summarizes the results of this section.

We start with some basic axioms that all our rules satisfy. The first two are classic symmetry axioms: a rule f is *anonymous* if whenever two profiles P and P' are such that every ballot appears exactly as often in P as in P', then $f(P) = f(P')$. It is *neutral* if for every profile P, renaming the candidates in P leads to the same renaming in $f(P)$. The third basic property fundamentally captures the aim of an axis rule: if there are perfect axes, then the rule should return those.

Consistency with linearity A rule f is *consistent with linearity* if $f(P) = \mathrm{con}(P)$ for all linear profiles P.

These axioms allow us to assume that the cost function has certain structure, in particular that it attains its minimum value for consistent axes, that it is invariant under reversing the axis, and that it is symmetric.

Lemma 1 *Let f be a scoring rule. Then, f is neutral and consistent with linearity if and only if it is induced by a cost function s_f such that*

1. $s_f(A, \lhd) \geqslant 0$*, and* $s_f(A, \lhd) = 0$ *if and only if A is an interval of* $\lhd$*,*
2. $s_f(A, \lhd) = s_f(A, \overleftarrow{\lhd})$*, and*
3. *there exists a function* $g : \{0,1\}^m \to \mathbb{R}_{\geqslant 0}$ *such that* $s_f(A, \lhd) = g(x_{A,\lhd}) = g(x_{A,\overleftarrow{\lhd}})$ *(so s_f depends only on the induced approval vector $x_{A,\lhd}$).*

We provide the formal proof in Appendix B.1.

	VD	MF	BC	MS	FT
Stability	✓	✗	✗	✗	✗
Ballot monotonicity	✓	✗	✓	✗	✗
Clearance	✗	✗	✓	✓	✓
Veto-centrism	✗	✗	✗	✓	✓
Clone-proximity	✗	✗	✗	✗	✓
Clone-resistance	✓	✗	✗	✗	✗

Table 2: Properties satisfied by the axis rules.

5.1. Stability and Monotonicity

Some rules are more sensitive to changes in information than others. Intuitively, Voter Deletion rarely reacts to changes in the profile, as it only checks whether the ballots are intervals of the axis or not. Thus, a single voter will have little effect on the axes selected. Indeed, for VD, adding a new ballot to the profile cannot completely change the set of optimal solutions. For other rules, this is not the case.

Stability A rule f satisfies *stability* if for every profile P and approval ballot A, we have $f(P) \cap f(P \cup \{A\}) \neq \emptyset$.

A similar axiom is used by Ceron and Gonzalez (2021) to characterize Approval Voting. Whether stability is a desirable property depends on the context: while it implies that the rule is robust, it also means that the rule might disregard too much information.

Proposition 2 *Stability is satisfied by VD, but not by MF, BC, MS and FT.*

Proof. Let us prove that VD satisfies this axiom. Let P be a profile and A an approval ballot. Assume by contradiction that $\text{VD}(P) \cap \text{VD}(P \cup \{A\}) = \emptyset$. Let $\lhd \in \text{VD}(P)$ and $\lhd' \in \text{VD}(P \cup \{A\})$. This means $s_{\text{VD}}(P, \lhd) \leqslant s_{\text{VD}}(P, \lhd') - 1$. Moreover, by definition of VD, $0 \leqslant s_{\text{VD}}(A, \cdot) \leqslant 1$. Put together, we have:

$$\begin{aligned} s_{\text{VD}}(P \cup \{A\}, \lhd) &= s_{\text{VD}}(P, \lhd) + s_{\text{VD}}(A, \lhd) \\ &\leqslant (s_{\text{VD}}(P, \lhd') - 1) + 1 \\ &\leqslant s_{\text{VD}}(P, \lhd') + s_{\text{VD}}(A, \lhd') \\ &= s_{\text{VD}}(P \cup \{A\}, \lhd') \end{aligned}$$

Therefore, $s_{\text{VD}}(P \cup \{A\}, \lhd) \leqslant s_{\text{VD}}(P \cup \{A\}, \lhd')$, which contradicts $\lhd \notin f(P \cup \{A\})$.

For $f \in \{\text{MF}, \text{BC}, \text{MS}, \text{FT}\}$, let us consider the profile $P = (\{a, b, e\}, \{a, b, c, e\}, \{b, c, d, e, f\})$. By consistency with linearity, $f(P) = \{aebcfd, aebcdf, aebcdf, abecfd\}$ (up to the reversed axes). Now, consider the ballot $A = \{a, b, d, f\}$. For any $\lhd \in f(P)$, we have $s_{\text{MF}}(P \cup \{A\}, \lhd) = s_{\text{BC}}(P \cup \{A\}, \lhd) = 2$, $s_{\text{MS}}(P \cup A, \lhd) = 4$ and $s_{\text{FT}}(P \cup A, \lhd) = 8$.

However, let us consider the axis $\lhd' = ceabfd$. The only ballot not interval of $\lhd'$ is $\{b, c, d, e, f\}$, which yields the cost of 1 for MF and BC, 2 for MS and 6 for FT. Therefore, $f(P) \cap f(P \cup \{A\}) = \emptyset$ for $f \in \{\text{MF}, \text{BC}, \text{MS}, \text{FT}\}$. Thus, these rules do not satisfy stability. □

Monotonicity axioms say that if the input changes so as to more strongly support the current output, then the output should stay the same. For our setting, we define monotonicity to say that if some voters *complete* their ballots by approving all interfering candidates with respect to the current axis $\lhd$, then $\lhd$ should continue being selected.

Ballot monotonicity A rule f satisfies *ballot monotonicity* if for every profile P, ballot $A \in P$ and axis $\lhd \in f(P)$ such that A is not an interval of $\lhd$, we still have $\lhd \in f(P')$ for the profile P' obtained from P by replacing A by the interval $A' = \{x \in C : \exists y, z \in A \text{ s.t. } y \trianglelefteq x \trianglelefteq z\}$.

VD and BC satisfy this axiom, but the other rules do not.

Proposition 3 *Ballot monotonicity is satisfied by VD and BC, but not by MF, MS and FT.*

Proof. For VD, observe that by changing the ballot A to A' we decrease the cost of $\lhd$ by 1, and the cost of all other axes decreases by at most 1, so $\lhd$ is still among the selected axes. For BC, suppose we need to add k candidates to the ballot A as part of making the profile linear. Then the change to A' reduces the cost of $\lhd$ by k, and the cost of any other axis decreases by at most k (as we added only k candidates to A), so $\lhd$ is still selected.

Let us now prove that the other rules do not satisfy this axiom. We assume the opposite for contradiction. Consider a set of 6 candidates $C = \{a, b, c, d, e, f\}$ and a profile P containing each possible ballot of 4 candidates. As all rules satisfy neutrality, all axes are optimal for MF, MS and FT. More formally, there is $x \in \mathbb{R}$ such that for all $f \in \{\text{MF}, \text{MS}, \text{FT}\}$ and all $\lhd$, we have $s_f(\lhd, P) = x$. Consider now the axis $\lhd_1 = abcdef$, and the ballot $A = \{a, b, e, f\} \in P$. By ballot monotonicity, $\lhd_1$ is still optimal for P' which is identical to P in which A is replaced by $\{a, b, c, d, e, f\}$. Denote $\lhd_2 = abfcde$. Because $\lhd_1$ is optimal for P', we have

$$s_f(\lhd_1, P') - s_f(\lhd_2, P') \leqslant 0$$

For any axis $\lhd$ and $f \in \{\text{MF}, \text{MS}, \text{FT}\}$, we have $s_f(\lhd, P')$ rewrites as $s_f(\lhd, P) + s_f\lhd, \{a, b, c, d, e, f\} - s_f\lhd, A$. By simply noting that $s_f\lhd, \{a, b, c, d, e, f\} = 0$, we have

$$\begin{aligned} s_f(\lhd_1, P) - s_f(\lhd_1, A) - (s_f(\lhd_2, P) - s_f(\lhd_2, A)) &\leqslant 0 \\ x - s_f(\lhd_1, A) - (x - s_f(\lhd_2, A)) &\leqslant 0 \\ s_f(\lhd_2, A) - s_f(\lhd_1, A) &\leqslant 0 \end{aligned}$$

For MF, MS and FT, the cost of A on $\lhd_1$ is respectively 2, 4 and 8, while the cost on $\lhd_2$ is respectively 1, 2 and 6. Therefore, the last inequality is not satisfied, and we have proven by contradiction that these rules do not satisfy ballot monotonicity. [Théo: Went with Chris' idea. Proofread needed.] [Magdaléna: Proofread done, OK for me.] □

5.2. Centrists and Outliers

On a high level, good axes should place less popular candidates towards the extremes, where they are less likely to destroy intervals. Conversely, popular candidates are safer to place in the center. We will define two axioms that identify profiles where this expectation is strongest, and that require candidates to be correctly placed in center or extreme positions.

Our first axiom considers the placement of very unpopular candidates. The axiom is easiest to satisfy by placing them at the extremes, but it does not require doing so in all cases.

Clearance A rule f satisfies *clearance* if for every profile P and any candidate x that is never approved in P, for each $\lhd \in f(P)$, there is no $A \in P$ with $y, z \in A$ with $y \lhd x \lhd z$.

Thus, under clearance, never-approved candidates cannot be interfering.

Proposition 4 *The clearance property is satisfied by BC, MS and FT, but not by VD and MF.*

Proof. For BC, MS, and FT, assume for contradiction that there exist a profile P with a never approved candidate x, and $\lhd \in f(P)$ such that y and z are approved together in a ballot $A \in P$ while $y \lhd x \lhd z$. Consider the axis $\lhd'$ identical to $\lhd$ but in which x was moved to the left extreme. As x is interfering on A, we have $s_{\mathrm{BC}}(A, \lhd') = s_{\mathrm{BC}}(A, \lhd) - 1$, $s_{\mathrm{MS}}(A, \lhd') = s_{\mathrm{MS}}(A, \lhd) - \min(|\{y : y \lhd x\}|, |\{y : x \lhd y\}|)$ and $s_{\mathrm{FT}}(A, \lhd') = s_{\mathrm{FT}}(A, \lhd) - |\{y : y \lhd x\}| \cdot |\{y : x \lhd y\}|$. In any case, we have $s_f(A, \lhd') < s_f(A, \lhd)$. Moreover, as x is never approved, for all other ballots $A_i \in P$ we have $s_f(A, \lhd') \leqslant s_f(A, \lhd)$ for $f \in \{\mathrm{BC}, \mathrm{MS}, \mathrm{FT}\}$. Thus, we have found an axis $\lhd'$ with strictly lower cost than $\lhd$, a contradiction.

Consider the profile $P = (\{a, b\}, \{a, c\}, \{a, d\})$ on the set of candidates $C = \{a, b, c, d, e\}$, and let $f \in \{\mathrm{VD}, \mathrm{MF}\}$. For each $\lhd \in f(P)$, $s_f(P, \lhd) \geqslant 1$. Indeed, at most two of the candidates b, c, d can be placed next to a on $\lhd$, so at least one of the ballots of P is not an interval of $\lhd$. Let us now consider the following axis $\lhd = \mathit{baced}$. We have $s_f(P, \lhd) = 1$, so $\lhd \in f(P)$ and hence f does not satisfy the clearance property. Indeed, e is never approved, but it interferes with the ballot $\{a, d\}$ on $\lhd$. □

While VD and MF always choose *some* axis that satisfies the clearance condition, they can additionally choose axes which violate this condition, and hence they fail the axiom.

For another way of formalizing the intuition that unpopular candidates should be placed at the extremes, we consider *veto profiles* in which every ballot has size $m - 1$, so each voter approves all but one of the candidates. For a veto profile, the only voters who will approve an interval are those who veto a candidate at one extreme of the axis. Since veto profiles do not have any interesting structure, the best candidates to put at the left and right end of the axis are two candidates with the lowest approval score (i.e. the *most* vetoed candidates). All of our rules indeed choose only such outcomes.

We can extend this intuition to say that the *least* vetoed candidate should be placed in the center, so that as few ballots as possible have holes in the center.

Veto-centrism A rule f satisfies *veto-centrism* if for every veto profile P, the median candidate (or one of the two median candidates if the number of candidates is even) of any axis $\lhd \in f(P)$ has the highest approval score.

Among the rules studied in this paper, only MS and FT satisfy veto-centrism.

Proposition 5 *FT and MS satisfy veto-centrism, but VD, MF and BC fail it.*

Proof. Let us denote the candidates by $c_1, c_2, \dots, c_m$, and let A_{-i} be a ballot approving all candidates but c_i. For each $i \in \{1, \dots, m\}$, we denote by n_i the number of occurrences of A_{-i} in P. Since P is a veto profile, any ballot of P corresponds to A_{-i} for some $i \in \{1, \dots, m\}$, and thus $n = n_1 + n_2 + \ldots + n_m$.

We make the two following straightforward observations: (1) Given an axis $\lhd$ with left- and rightmost candidates c_l and c_r, only A_{-l} and A_{-r} are intervals of $\lhd$, (2) VD, BC and MF are equivalent on veto profiles. Indeed, their cost functions equal 1 for any A_{-i} that is not an interval of a given axis.

Assume first that $f \in \{\mathrm{VD}, \mathrm{MF}, \mathrm{BC}\}$. Then, each axis $\lhd$ with left- and rightmost candidates c_l and c_r has a cost of $n - n_l - n_r$. It follows that an axis is optimal if and only if its two outermost candidates correspond to the two least approved candidates. In particular, the optimality of a solution is independent of the position of the most approved candidate (as soon as it is not an extremity). Hence, if $m > 4$, for all veto profile P there exists an optimal axis such that the most approved candidate is not the median candidate. Therefore, VD, BC and MF do not satisfy veto-centrism.

Let us now prove that veto-centrism holds for FT and MS. For simplicity, we assume that m is odd, and so that there is only one median candidate on the axis, at position $\frac{m+1}{2}$ (however, note that the reasoning below also works for m even, with only a slight straightforward modification).

Regarding FT, given an axis $\lhd$, let us denote by k_i the position of c_i on $\lhd$. Then each copy of A_{-i} in P creates $t_{k_i} = (k_i - 1) \cdot (m - k_i)$ forbidden triples and $s_{\mathrm{FT}}(P, \lhd) = \sum_{i \leqslant m} n_i t_{k_i}$. t_{k_i} is maximal when $k_i = \frac{m+1}{2}$ (i.e. when c_i is the median candidate of $\lhd$) Assume wlog that the most approved candidate is c_1. Let us suppose for contradiction that $\lhd$ is an optimal axis such that c_2 is its median candidate with strictly less approvals, i.e., $n_2 > n_1$. Let $\lhd'$ be an axis obtained from $\lhd$ by swapping the positions of c_1 and c_2. We claim that $s_{\mathrm{FT}}(P, \lhd') < s_{\mathrm{FT}}(P, \lhd)$. For each $i \neq 1, 2$, the number of triples caused by ballots of type A_{-i} is the same for both axes,as this number only depends on the position of c_i on the axis. Hence, $s_{\mathrm{FT}}(P, \lhd')$ and $s_{\mathrm{FT}}(P, \lhd)$ only differ in triples caused by ballots A_{-1} and A_{-2}. We have

$$s_{\mathrm{FT}}(P, \lhd') - s_{\mathrm{FT}}(P, \lhd) = (n_1 t_{\frac{m+1}{2}} + n_2 t_{k_1}) - (n_1 t_{k_1} + n_2 t_{\frac{m+1}{2}}) = (n_1 - n_2)(t_{\frac{m+1}{2}} - t_{k_1}) < 0,$$

as $n_1 < n_2$ and $t_{\frac{m+1}{2}} > t_{k_1}$. This show the claim and contradicts the optimality of $\lhd$.

Finally, let us prove that Veto-centrism holds for MS. Similarly to the case of FT, we can note that each copy of A_{-i} in P generates $t_{k_i} = \min\{k_i - 1, m - k_i\}$ swaps. Indeed, c_i is the unique non-approved candidate in A_{-i}, so it needs to be swapped with all the candidates on its left, or its right. This value is maximal if c_i is the median candidate of the axis, i.e., if $k_i = \frac{m+1}{2}$. It is now easy to see that the proof for FT also holds for MS. □

In fact, both rules always place candidates so that the approval scores are single-peaked.

Clearance and veto-centrism suggest that MS and FT use the information in a profile well by correctly placing popular and unpopular candidates. This tendency to put low-approval candidates towards the ends is also confirmed by our experiments in Section 6. While this generally seems sound, in the political context it can lead to wrong answers: there can be ideologically centrist candidates who don't get many votes due to not being well-known. We leave for future work whether there are rules that can correctly place candidates in these contexts.

5.3. Clones and Resistance to Cloning

We now focus on the behaviour of rules in the presence of essentially identical candidates. We say that $a, b \in C$ are *clones* if for each voter $i \in V$, $a \in A_i$ if and only if $b \in A_i$. While perfect clones are rare, two candidates may have very similar sets of supporters, and studying clones gives insights for how rules handle similar candidates.

Intuitively, one would expect clones to be next to each other on any optimal axis. This is captured by the following axiom:

> **Clone-proximity** A rule f satisfies *clone-proximity* if for every profile P in which $a, a' \in C$ are clones, for every axis $\lhd \in f(P)$ and any candidate x such that $a \lhd x \lhd a'$ or $a' \lhd x \lhd a$, we have $x \in A$ whenever $a, a' \in A$ for every $A \in P$.

Note that in the definition, x and a are not necessarily clones: x can be approved even if a is not approved.

Surprisingly, only FT satisfies clone-proximity. All of our rules choose at least one axis where the clones are next to each other, but the rules other than FT may choose extra axes with a violation. For instance, in $P = (2 \times \{b, c\}, 2 \times \{c, d\}, 1 \times \{a, a', b, d\})$ the axis $a \lhd b \lhd c \lhd d \lhd a'$ in which the clones a and a' are at opposite extremes is an optimal axis for VD, MF and BC.

Proposition 6 *Clone-proximity is satisfied by FT, but not by VD, MF, BC and MS.*

Proof. We first prove that FT satisfy the axiom. Let $P = (A_1, \dots, A_n)$ be a profile where a and a' are clones. Let $\lhd$ be an axis. We denote by $T_\lhd$ the set of all forbidden triples (i, l, c, r) such that $l \lhd c \lhd r$ and $l, r \in A_i$ but $c \notin A_i$. Then, $s_{\mathrm{FT}}(P, \lhd) = |T_\lhd|$.

First note that we cannot have a forbidden triple (i, l, c, r) with one of the clone as c (in the center) and the other on one side (l or r), as a and a' are always approved together. Thus, the only triples involving both a and a' are those for which both sides l and r are one of the clones (e.g. $l = a$ and $r = a'$). For an axis $\lhd$, let us denote $S_{\lhd}^{(a,a')}$ the number of such triples. Moreover, let us denote $S_{\lhd}^{a}$ the number of triples involving a and not a' and $S_{\lhd}^{a'}$ the number of triples involving a' and not a. Finally, let us denote $S_{\lhd}^{0}$ the number of triples involving neither a nor a'. For any axis $\lhd$, we have $s_{\mathrm{FT}}(P, \lhd) = S_{\lhd}^{(a,a')} + S_{\lhd}^{a} + S_{\lhd}^{a'} + S_{\lhd}^{0}$.

Assume now by contradiction that the clones are not next to each other on the optimal axis $\lhd$, i.e. there exists $x \in C$ such that $a \lhd x \lhd a'$ or $a' \lhd x \lhd a$ and a ballot $A_i \in P$ such that $a, a' \in A_i$ and $x \notin A_i$. Thus, we have $S_{\lhd}^{(a,a')} \geqslant 1$.

Assume without loss of generality that $S_{\lhd}^{a} \leqslant S_{\lhd}^{a'}$. Let us consider the axis $\lhd'$ obtained by moving a' next to a on $\lhd$, i.e. there is no $x \in C$ such that $a \lhd x \lhd a'$ or $a' \lhd x \lhd a$. Thus, we have $S_{\lhd'}^{(a,a')} = 0$, and $S_{\lhd'}^{a'} = S_{\lhd'}^{a} = S_{\lhd}^{a}$, as all triples that does not involve a' will not be affected by the move, and the triples involving a' will be the same as those involving a now that they are next to each other. For the same reason, $S_{\lhd'}^{0} = S_{\lhd}^{0}$. Thus, we have the following.

$$
\begin{aligned}
s_{\mathrm{FT}}(P, \lhd) &= S_{\lhd}^{(a,a')} + S_{\lhd}^{a} + S_{\lhd}^{a'} + S_{\lhd}^{0} \\
&\geqslant 0 + S_{\lhd}^{a} + S_{\lhd}^{a} + S_{\lhd}^{0} \\
&\geqslant 0 + S_{\lhd'}^{a} + S_{\lhd'}^{a'} + S_{\lhd'}^{0} \\
&= s_{\mathrm{FT}}(P, \lhd')
\end{aligned}
$$

Axis $\lhd'$ has a lower FT cost than $\lhd$, a contradiction. This proves that the FT rule satisfies this property.

We now show that other rules do not satisfy this property. Consider the following profile:

$$2 : \{a_1, a_2\}, \quad 2 : \{a_2, a_3\}, \quad 1 : \{x, x', a_1, a_3\}$$

Because of the cycle a_1, a_2, a_3, this profile is not linear so all axes have cost at least 1. Now observe that the axis $x \lhd a_1 \lhd a_2 \lhd a_3 \lhd x'$ has cost 1 for VD, BC and MF. On this axis, a_2 is between the clones x and x' but is never approved with them. Thus, these three rules fail this axiom.

For MS, consider the following profile:

$$1 : \{a, a', b, b'\}, \quad 1 : \{b, b', x, x'\}, \quad 1 : \{x, x', a, a'\}$$

By neutrality of Minimal Swaps, the cost of all axes in which clones are next to each other is the same as the score of $a \lhd a' \lhd x \lhd x' \lhd b \lhd b'$, which is 4. However, another axis has cost 4 for Minimum Swaps: $x \lhd a \lhd a' \lhd x' \lhd b \lhd b'$. On this axis, a is between the clones x and x' but a is not approved in the ballot $\{b, b', x, x'\}$, containing x and x'. Thus, Minimum Flips also fails this property. □

Inspired by axioms from voting theory (Tideman, 1987), we could require that removing or adding a clone to the profile would not change the result. More precisely, if we remove a clone from a profile, the restriction of any optimal axis should remain optimal, and adding a clone to a profile should not modify the relative order of the other candidates on any optimal axis. To formally define this, we need some notation. For a profile P defined on a set C of candidates, we denote $P_{C'}$ the restriction of P to a subset of candidates $C' \subseteq C$. We also denote P_{-c} the reduction of the profile to $C \setminus \{c\}$ where $c \in C$ is a given candidate. Similarly, we define $\lhd_{C'}$ and $\lhd_{-c}$. We can now state the axiom:

Resistance to cloning A rule f is *resistant to cloning* if for every profile P in which $a, a' \in C$ are clones, (1) for all axes $\lhd \in f(P)$, we have $\lhd_{-a'} \in f(P_{-a'})$ and (2) for all axes $\lhd^* \in f(P_{-a'})$, there is an axis $\lhd \in f(P)$ with $\lhd_{-a'} = \lhd^*$.

Among the rules studied in this paper, only VD is resistant to cloning.

Proposition 7 *VD satisfies resistance to cloning, but not MF, BC, MS and FT.*

Proof. We start by proving that VD satisfies resistance to cloning. Let $\lhd \in f(P)$; we will prove that $\lhd_{-a'} \in f(P_{-a'})$. It is easy to see that $s_{\text{VD}}(\lhd_{-a'}, P_{-a'}) \leqslant s_{\text{VD}}(\lhd, P)$ as all interval ballots of P on $\lhd$ will remains interval ballots of $P_{-a'}$ on $\lhd_{-a'}$. Now, assume that $\lhd_{-a'} \notin f(P_{-a'})$ and instead some axis $\lhd' \in f(P_{-a'})$ is optimal with cost $s_{\text{VD}}(\lhd', P_{-a'}) < s_{\text{VD}}(\lhd_{-a'}, P_{-a'})$. Consider the axis $\lhd'_{+a'}$ which is equivalent to $\lhd'$ with a' put next to a. Clearly, an approval ballot of P is an interval of $\lhd'_{+a'}$ if and only if its restriction in $P_{-a'}$ is an interval of $\lhd'$. Thus, $s_{\text{VD}}(\lhd'_{+a'}, P) = s_{\text{VD}}(\lhd', P_{-a'})$. Combining all of this, we have:

$$s_{\text{VD}}(\lhd'_{+a'}, P) = s_{\text{VD}}(\lhd', P_{-a'}) < s_{\text{VD}}(\lhd_{-a'}, P_{-a'}) \leqslant s_{\text{VD}}(\lhd, P)$$

which contradicts the optimality of $\lhd$ for P. This proves that $\lhd_{-a'} \in f(P_{-a'})$.

The opposite direction uses the same reasoning. Let $\lhd \in f(P_{-a'})$ and $\lhd_{+a'}$ the equivalent axis in which we put a next to a'. Again, $s_{\text{VD}}(\lhd_{+a'}, P) = s_{\text{VD}}(\lhd, P_{-a'})$. Now assume by contradiction that there is $\lhd' \in f(P)$ with a lower cost than $\lhd_{+a'}$: $s(\lhd', P) < s(\lhd_{+a'}, P)$. As explained above, we have $s_{\text{VD}}(\lhd'_{-a'}, P_{-a'}) \leqslant s_{\text{VD}}(\lhd', P)$. Combining these three inequalities gives $s_{\text{VD}}(\lhd'_{-a'}, P_{-a'}) < s_{\text{VD}}(\lhd, P_{-a'})$, which contradicts the optimality of $\lhd$. This proves that $\lhd_{+a'}$ is optimal for VD in P.

To prove that BC does not satisfy resistance to cloning, let us consider the profile $P = (3 \times \{b, a, a'\}, 4 \times \{c, a, a'\}, 2 \times \{b, c\})$. It is easy to check that the unique optimal axis (up to reversal, and permutation of a and a') is $b \lhd c \lhd a \lhd a'$ with $s_{\text{BC}}(P, \lhd) = 3$. Indeed, if a and a' are not next to each other, at least two types of ballot will not be interval of the axis, which will yield a cost of at least 5, and the axes on which b and c are the extremities a cost of at least 4.

However, if we remove the candidate a', the cost of $\lhd_{-a'} = bca$ is 3. It is hence no more optimal, as the axis $\lhd^* = bac$ achieves a lower cost of 2. We use a very similar idea to prove that MF does not satisfy resistance to cloning. We consider the profile $P = (1 \times \{b, d\}, 2 \times \{b, a, a'\}, 2 \times \{c, a, a'\}, 1 \times \{a, e\}, 3 \times \{b, c, d, e\})$. We can check that the axis $\lhd = dbceaa'$ is optimal for MF with $s_{\text{MF}}(P, \lhd) = 4$. If a and a' are not next to the other on the axis, at least two of the ballot types $\{b, a, a'\}$, $\{c, a, a'\}$ and $\{b, c, d, e\}$ are not intervals, which yields a cost of at least 4.

Any axis of form (up to reversal) $\{d, b\} \lhd \{a, a'\} \lhd \{c, e\}$ has a cost greater or equal than 6 because of ballots $\{b, c, d, e\}$. Any axis with one candidate on the left of $\{a, a'\}$ and three candidate on the right of $\{a, a'\}$ (up to reversal) has a cost of at least 4: the ballots $\{b, c, d, e\}$ generates at least 3 flips, and at least one of the ballots $\{b, d\}, \{c, e\}$ is not an interval either.

However, $\lhd_{-a'} = dbcea$ is not optimal for $P_{-a'}$: $s_{\text{MF}}(P_{-a'}, \lhd_{-a'}) = 4$ (ballots $\{b, a\}$ and $\{c, a\}$ are not intervals). The axis $\lhd^* = dbace$ has a lower cost of 3 (as each copy of ballot $\{d, b, c, e\}$ generates one flip and all other ballots are intervals of the axis).

To prove that FT and MS do not satisfy resistance to cloning, let us consider the profile $P = (3 \times \{a, b\}, 3 \times \{b, c\}, 1 \times \{a, c, d\})$, and let $f \in \{\text{FT}, \text{MS}\}$. We have $f(P) = \{\lhd^1, \lhd^2\}$ with $\lhd^1 = abcd$ and $\lhd^2 = dabc$ (up to the reversed axes). Indeed, $s_{\text{FT}}(P, \lhd^i) = 2$ and $s_{\text{MS}}(P, \lhd^i) = 1$ for $i \in \{1, 2\}$. These are the only axes on which both $\{a, b\}$ and $\{b, c\}$ are intervals – in other words, the score of any other axis will be at least 3. Let us now consider a profile $P' = (3 \times \{a', a, b\}, 3 \times \{b, c\}, 1 \times \{a', a, c, d\})$. We note that it is the profile P to which we have added a candidate a', clone of a. Under resistance to cloning, there should be an axis $\lhd \in f(P')$ such that $\lhd_{-a} = \lhd^1$. Among all possible axes generalizing $\lhd^1$, the best one for MF and FT (up to the permutation of a and a') is $\lhd = aa'bcd$, with a cost of 4 for FT and 2 for MS. However, we can find an axis $\lhd^* = daa'bc$ with cost of 3 for FT and 1 for MS. Hence, there is no $\lhd \in f(P')$ such that $\lhd_{-a} = \lhd^1$. Thus, FT and MS do not satisfy resistance to cloning. □

These two clone axioms are quite strong: each excludes all but one of our rules. Indeed, we now show that if a scoring rule satisfies neutrality and consistency with linearity, then clone-proximity and resistance to cloning are actually incompatible.[1]

Theorem 2 *No neutral scoring rule satisfies resistance to cloning, clone proximity, and consistency with linearity.*

Proof. Let f be a scoring rule satisfying all four axioms, and s_f its cost function. As proven in Lemma 1, by neutrality there is a function $g_f : \{0,1\}^m \to \mathbb{R}_{\geqslant 0}$ such that $s_f(A, \lhd) = g_f(x_{A,\lhd}) = g_f(x_{A,\overleftarrow{\lhd}})$, where $x_{A,\lhd}$ is the approval vector of A and $x_{A,\overleftarrow{\lhd}}$ is the reversed vector.

Let $y = g_f((1,0,1,0)) = g_f((0,1,0,1))$ and $y' = g_f((1,0,1,1)) = g_f((1,1,0,1))$ be the cost of the respective vectors. By Lemma 1 (using consistency with linearity), $y > 0$. Let $q \in \mathbb{N}$ with $q > y'/y$ and consider the profile $P = (q \times \{b,c\}, q \times \{c,d\}, 1 \times \{a,b,d\})$. For $\lhd \in \{\lhd^1, \lhd^2\}$ with $\lhd^1 = abcd$ and $\lhd^2 = bcda$, we have $s_f(\{a,b,d\}, \lhd) = y'$. All other axes break one of the pairs $\{b,c\}$, $\{c,d\}$, thus ensuring a cost of at least $q \cdot y > y'$. Therefore, $\lhd_1, \lhd_2 \in f(P)$.

Consider now the profile P' in which we add a clone b' of b: $P' = (q \times \{b,b',c\}, q \times \{c,d\}, 1 \times \{a,b,b',d\})$. By clone-proximity, b and b' are next to each other on any $\lhd \in f(P')$. By resistance to cloning, there exists $\lhd^3$ (resp. $\lhd^4$) in $f(P')$ extending $\lhd^1$ (resp. $\lhd^2$). Combining this with neutrality, $f(P')$ contains $\lhd^3 = abb'cd$ and $\lhd^4 = bb'cda$, which thus must have the same cost. Since the ballots $\{b,b',c\}$ and $\{c,d\}$ are intervals of both of these axes and the rule is consistent with linearity, they contribute a cost of 0 and thus the cost difference of the two axes only depends on the remaining ballot. This implies $s_f(\{a,b,b',d\}, \lhd^3) = s_f(\{a,b,b',d\}, \lhd^4)$, i.e. $g_f((1,1,1,0,1)) = g_f((1,1,0,1,1))$.

Now, consider the profile P'' which is a copy of P but with a clone a' of a: $P'' = (q \times \{b,c\}, q \times \{c,d\}, 1 \times \{a,a',b,d\})$. Using the same arguments as in the case of P' yields two optimal axes $\lhd^5 = aa'bcd$ and $\lhd^6 = bcdaa'$. However, let us now compare $\lhd^5$ to $\lhd^7 = abcda'$. The ballots $\{b,c\}$ and $\{c,d\}$ are intervals of both axes, and the cost of $\{a,a',b,d\}$ is the same on both, as we already showed that $g_f((1,1,1,0,1)) = g_f((1,1,0,1,1))$. Thus, $\lhd^7$ is also an optimal axis, which is in contradiction with clone-proximity, since a and a' are not next to each other. □

We can show that resistance to cloning and ballot monotonicity in fact characterize VD among scoring rules. This not only distinguishes VD from the other introduced rules, but shows its normative appeal among the entire class of scoring rules. The full proof is in Appendix B.2, where we also show that the axioms are logically independent, assuming neutrality.

Theorem 3 *Let $m \geqslant 6$, and let f be a neutral scoring rule. Then f satisfies consistency with linearity, ballot monotonicity, and resistance to cloning if and only if it is VD.*

Proof sketch. Let f be a scoring rule satisfying neutrality, consistency with linearity, resistance to cloning and ballot monotonicity. As shown in Appendix B.1, f is induced by a symmetric cost function s with $s(A, \lhd) = 0$ iff A forms an interval in $\lhd$. Further, s only depends on the approval vector $x_{A,\lhd}$, i.e. there exists a function $g : \{0,1\}^m \to \mathbb{R}_{\geqslant 0}$ such that $s(A, \lhd) = g(x_{A,\lhd})$ for all ballots A and axis $\lhd$.

The steps of the proof are as follows:

[1] There are rules that are not scoring rules which satisfy both clone resistance and proximity. For example, consider the rule that takes a profile, identifies all maximal clone sets, and replaces each by a single representative candidate. Then apply a rule to the collapsed profile and de-replace the representatives.

1. Using ballot monotonicity, we show that there is a function h such that for all A and $\lhd$ such that A is not an interval of $\lhd$, $s(A, \lhd) = h(m, k_{\text{app}}, k_{\text{int}})$, where m is the number of candidates, $k_{\text{app}} = |A|$ is the number of approved candidates and k_{int} is the number of interfering candidates.
2. Using resistance to cloning, we show that for A not interval of $\lhd$, $s(A, \lhd)$ only depends on the sum $k_{\text{app}} + k_{\text{int}}$, i.e. there is h such that $s(A, \lhd) = h(m, k_{\text{app}} + k_{\text{int}})$.
3. We show that for A not interval of $\lhd$, $s(A, \lhd)$ can only take two values: $s(A, \lhd) = h^*_m$ if $k_{\text{app}} + k_{\text{int}} = m$ and $s(A, \lhd) = h_m$ otherwise.
4. Finally, we show that $h^*_m = h_m$ and thus the rule is VD. □

If we drop ballot monotonicity, we can use similar ideas to show that resistance to cloning characterizes (under mild conditions) the class of *topological rules*. These are scoring rules with a function h such that $s_f(A, \lhd) = h(k)$, where k is the number of contiguous holes that the axis $\lhd$ creates in A (Appendix B.3), such as the "genus rule" with $h(k) = k$ that counts the total number of contiguous holes. For instance, on $\lhd = abcde$, the cost induced by $\{a, e\}$ is 1, but the one induced by $\{a, c, e\}$ is 2.

Resistance to cloning can be strengthened to *heredity*, a kind of independence of irrelevant alternatives axiom. It states that if we remove *any* candidate (not just a clone), the rule should return the original axes with that candidate omitted.

> **Heredity** A rule f satisfies *heredity* if for every profile P and every subset of candidates $C' \subseteq C$, we have that for each axis $\lhd \in f(P)$, there exists $\lhd^* \in f(P_{C'})$ such that $\lhd_{C'} = \lhd^*$.

However, an easy impossibility theorem shows that no reasonable axis rule can satisfy this axiom.

Proposition 8 *No axis rule satisfies heredity and consistency with linearity.*

Proof. Let f be an axis rule satisfying heredity and let $P = (\{a, b\}, \{a, c\}, \{a, d\})$. Let $\lhd \in f(P)$. In $\lhd$, there must be at least two candidates on the same side of a (as there are two sides and three candidates b, c and d), wlog b and c. By heredity, if we remove d, in $f(P_{-d})$ there must be an axis where a is in an extreme position. However by consistency with linearity, $f(P_{-d}) = \{bac, cab\}$, a contradiction. □

6. Experiments

In this section, we investigate the rules from Section 4 using an experimental analysis. While the rules are hard to compute, for m up to about 12, we can find the best axes in reasonable time. We describe two strategies in Appendix C.1: brute force (using pruning and heuristics) and ILP solvers.

Our main aims are: (1) to compare axis rules for approval profiles to two known rules for nearly single-peakedness for ranking profiles, *Voter Deletion* (VD-rank, Erdélyi et al. (2017)) and *Forbidden Triples* (FT-rank, Escoffier et al. (2021)), and (2) to compare axis rules to each other. For this, we use both synthetic and real datasets.

6.1. Synthetic Data

To better understand how different rules behave, we tested them on several synthetic data models (see Appendix C.2 for detailed descriptions and results) which sample a linear profile on a ground truth axis and add random noise to it. We then measured the distance of a rule's output to the

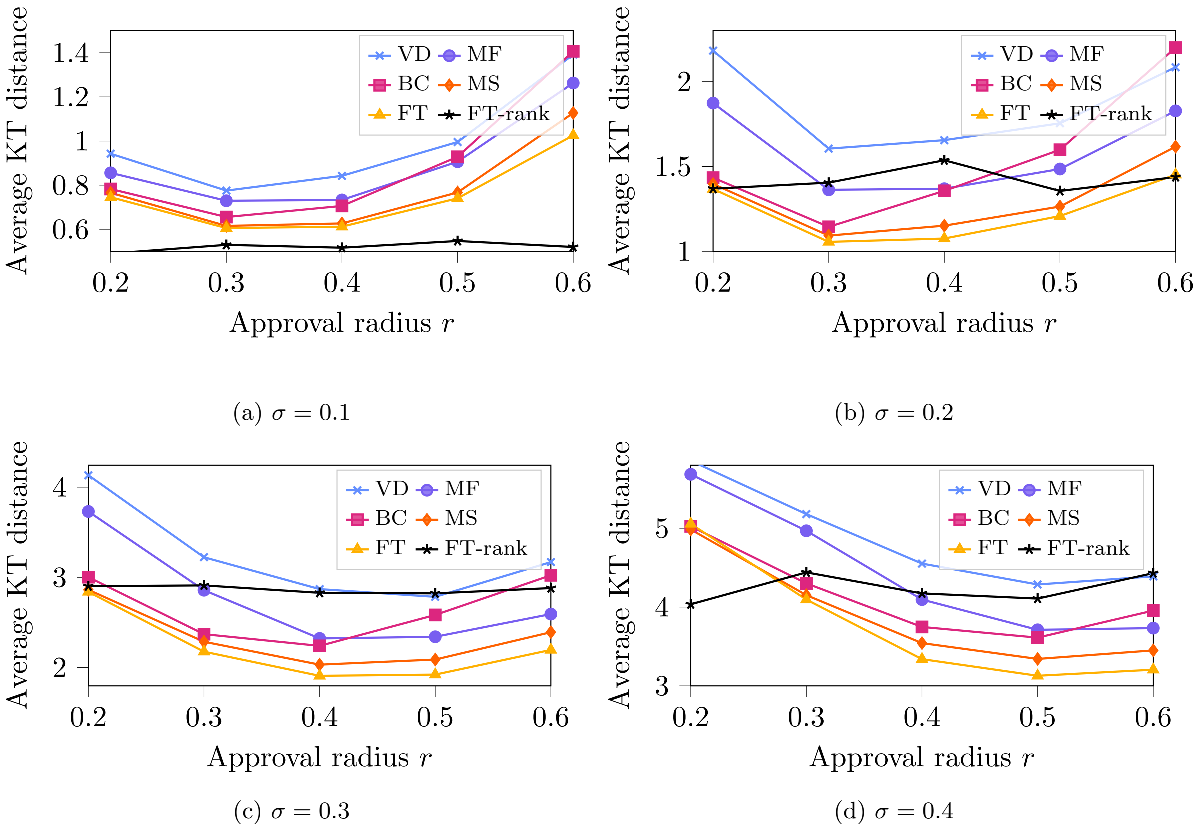

(a) $\sigma = 0.1$

(b) $\sigma = 0.2$

(c) $\sigma = 0.3$

(d) $\sigma = 0.4$

Figure 1: Evolution of the average KT distance between the axes returned by the rules and the actual axes for $r \in [0.2, 0.6]$, averaged over 1,000 random samples.

ground truth. Some of our rules are in fact the MLEs of these noise models, so as predicted they perform well in those cases. However some rules adapted better than others to different noise models. We observed that for all models, our rules tend to push the least approved candidates towards the extremes.

To compare approval-based and ranking-based rules, we introduce the *noisy observation model*, inspired by random utility models such as the Thurstone-Mosteller model. Each candidate and voter $x \in C \cup V$ is associated with a position $p(x) \in \mathbb{R}$ on the line. Each voter v estimates the position of each candidate c under independent normal noise: $p_v(c) = p(c) + \mathcal{N}(0, \sigma)$ with σ a parameter of the model. Voters approve (resp. rank) candidates based on their estimations. More precisely, the approval set of voter v contains all candidates such that $|p(v) - p_v(c)| \leqslant r$, where the approval radius r is a parameter of the model. The ranking of v is given by decreasing distances between $p(v)$ and $p_v(c)$.

The positions $p(c)$ of the candidates describe a ground truth axis $\triangleleft = c_1 c_2 \dots c_m$ such that $p(c_1) \leqslant p(c_2) \leqslant \dots \leqslant p(c_m)$. Figure 1 shows the Kendall-tau (KT) swap distance between the axes output by different rule results and the ground truth for $\sigma \in \{0.1, 0.2, 0.3, 0.4\}$ and $r \in \{0.2, 0.3, 0.4, 0.5, 0.6\}$. We conducted experiments with $m = 7$ candidates, $n = 100$ voters and 1000 random profiles for each set of parameters. We find that VD-rank is always far from the true axes (at distance 7–8, too much to fit in the chart), and that for most values of σ and r, approval rules actually perform better than FT-rank, returning axes with a lower average KT distance to the ground truth. This is surprising, as intuitively rankings provide more information than approvals. We note however that FT-rank is better than approval methods when r is very small or very large, so many approval sets are of size 0 or 1 (or m), and thus provide no information on candidates proximity. FT-rank is also slightly better when σ is small, but in this

Rule						◁						Min KT	Avg KT
VD	R	LO	NPA	LFI	PS	EM	LR	DLF	FN	UPR	SP	5	7.71
MF	LO	NPA	LFI	PS	EM	LR	DLF	FN	UPR	R	SP	**1**	4.43
BC	LO	NPA	LFI	PS	EM	LR	DLF	FN	R	UPR	SP	2	4.0
MS	LO	NPA	LFI	PS	EM	LR	DLF	FN	R	UPR	SP	2	4.0
FT	LO	NPA	PS	LFI	EM	R	LR	DLF	FN	UPR	SP	**1**	**3.71**
VD-rank	FN	DLF	R	LO	NPA	LFI	PS	EM	SP	UPR	LR	22	24.0
FT-rank	LO	NPA	R	LFI	PS	EM	LR	DLF	FN	UPR	SP	3	5.71

Table 3: Optimal axis of each rule for the 2017 French presidential election

case all approval rules also have very good performance, with their average KT distances all below 1. We also observe that for all values of parameters, the axes returned by the rules using more information (e.g., FT) are closer to the ground truth axes than those returned by the rules using less information (e.g., VD).

6.2. The French Presidential Election

We now present the results of our rules on two political datasets: the 2017 and 2022 edition of the online experiment *Voter Autrement* conducted during the French presidential elections (Bouveret et al., 2018). In parallel to the actual elections, the participants were invited to express their opinions on candidates using various voting methods, including approval and ranking-based ones. This allows us to compare our axis rules for both settings. After data cleaning, for the 2017 [2022] dataset, we obtained approval preferences of 20 076 voters [1379 voters] and preference rankings of 5 796 voters [412 voters] over 11 candidates [12 candidates]. Details on how the data was gathered and the experiments conducted can be found in Appendix C.3, together with our detailed results. There, we also explain how we reweighted votes to counteract response bias and to match the distribution of official election results.

Regarding approval rules, we note that they all returned very similar axes. They mostly differ on the position of less popular candidates (often placed at one of the extremes), and the relative order of candidates within their ideological subgroup (e.g., left-wing candidates). We computed the KT distance between the axes returned by our rules and the ones used by the main 7 polling institutes. All rules return an axis that has a KT distance of less than 5 to at least one poll institute axis (while the worst possible KT distance are 27 and 33 for $m = 11$ and 12). For instance, the ordering obtained with FT is very similar to the one of the *Ipsos institute*:

FT: LO, NPA, PS, LFI, EM, R, LR, DLF, FN, UPR, SP

Ipsos: LO, NPA, LFI, PS, EM, R, LR, DLF, FN, SP, UPR

The KT distance between them is 2. Note that most of the small parties (LO, NPA, R, UPR, SP), displayed using small font, are placed at one of the extremes. In Table 3, we present the axes returned by all the tested rules, as well as the minimum and average Kandall-tau distance to the poll institute axes.

Regarding ranking-based methods, the quality of the axes returned by FT-rank seems comparable to the axes returned by approval rules. Again, the VD-rank axes were much less convincing. This corroborates other observations in the literature. For instance, Sui et al. (2013) ran experiments on 2002 Irish General Election data and found that the VD-optimal axis only fit 0.4%–2.9% of voters. Escoffier et al. (2021) ran experiments on a similar French presidential election dataset and also observed that the optimal axis found using Voter Deletion was very different from the orderings discussed in French media. In our experiments, the optimal VD-rank

Rule	Avg KT	Correct Median
VD	4.94	53.8 %
MF	4.22	58.5 %
BC	3.68	56.9 %
MS	3.55	64.6 %
FT	**3.43**	**66.2** %

Table 4: Average Kendall-tau distance to MQ axis, and % of time the axis has the same median candidate than the MQ axis, averaged over 65 terms.

axes only cover less then 4% of voters. For comparison, the approval version of VD returned axes covering more than 60% of voters.

Finally, we observe that all rules violate the heredity property on our dataset. Removing even the least approved candidate could change the returned axis. However, these changes are marginal, like a less popular candidate being pushed towards an extreme or two left-wing candidates being inverted.

6.3. Supreme Court of the United States

Finally, we used our rules to obtain an ideological ordering of the 9 justices of the Supreme Court of the United States. The dataset is based on the opinions authored and joined by the justices. Each opinion, concurrence, or dissent becomes a ballot "approving" the justices that joined in it. The intuition is that justices joining the same opinion share an ideology so should be placed close together. This data is derived from the SCDB database, see Appendix C.4 for details and results.

The problem of ordering the justices has been extensively studied; the standard method used by political analysts is the *Martin-Quinn* (MQ) method, which uses a dynamic item response theory model (Martin and Quinn, 2002). A limitation of this model is that it can only use the vote data (whether a justice agreed with the majority or not), while our model can use more fine-grained data from which opinions were joined, which includes the *reasons* for the vote. We compare the axes returned by our rules for 65 terms between 1946 and 2021, removing the years having more than 9 justices involved (e.g. if one is replaced mid-term).

Table 4 shows the average KT distance of the axes returned by our rules to the Martin-Quinn axis. We see that these distances are in average quite low. Moreover, we observe that the FT rule comes closest, while the VD rule is relatively far away. We also checked how often the axes computed by our rules agreed with the Martin-Quinn axis on which justice is placed in the median position. This is of particular interest since the median justice tends to be pivotal. Again we see that the FT rule agrees most commonly with the Martin-Quinn axis, choosing the same median justice in 66% of terms. For future work, we see potential in adapting our rules to obtain methods perhaps more interesting than the Martin-Quinn method (as they will satisfy axiomatic properties).

7. Future Work

There are many promising directions for future work, such as considering methods that output other types of structures, like circular axes (in which the first and last candidates on the axis are next to each other) or embeddings into multiple dimensions, or introducing metric distances between candidates on the axis. An axiomatic approach could provide novel insights for all these problems. Moreover, the methods we present not only return a set of optimal axes, but also their "cost", which provides an indicator of how close a profiles is to be linear. One could try to analyze these methods as rules measuring the degrees of linearity of approval profiles.

Technically, several open questions remain. It would be interesting to obtain an axiomatic characterization of the class of scoring rules using the reinforcement axiom, though this is made challenging by the neutrality axiom being quite weak in our setting. It would also be useful to design polynomial-time computable rules that produce good outputs, to be able to deal with many candidates. Greedy versions of our rules are a natural starting point, but maybe better techniques exist.

Acknowledgments

This work was funded in part 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). The authors would like to thank Isabelle Lebon for her help in the early stages of this project, and Jérôme Lang for his helpful comments.

8. Bibliography

Antoinette Baujard and Isabelle Lebon. 2022. Retelling the story of the 2017 French presidential election: The contribution of approval voting. *Homo Oeconomicus* (2022), 1–22. https://doi.org/10.1007/s41412-022-00134-7 [→ p. 5]

Michael J. Baxter. 2003. *Statistics in archaeology.* Arnold. [→ p. 2]

Duncan Black. 1948. On the rationale of group decision-making. *The Journal of Political Economy* 56, 1 (1948), 23–34. https://doi.org/10.1086/256633 [→ p. 3]

Kellogg Speed Booth. 1975. *PQ-tree algorithms.* Ph. D. Dissertation. University of California, Berkeley, and Lawrence Livermore Laboratory. https://dominik-peters.de/archive/booth1975.pdf [→ p. 3, 7, 23]

Kellogg S. Booth and George S. Lueker. 1976. Testing for the consecutive ones property, interval graphs, and graph planarity using PQ-tree algorithms. *J. Comput. System Sci.* 13, 3 (1976), 335–379. https://doi.org/10.1016/S0022-0000(76)80045-1 [→ p. 2]

Sylvain Bouveret, Renaud Blanch, Antoinette Baujard, François Durand, Herrade Igersheim, Jérôme Lang, Annick Laruelle, Jean-François Laslier, Isabelle Lebon, and Vincent Merlin. 2018. Voter Autrement 2017 - Online Experiment. Dataset and companion article on Zenodo. https://doi.org/10.5281/zenodo.1199545 [→ p. 17]

Robert Bredereck, Jiehua Chen, and Gerhard J. Woeginger. 2016. Are there any nicely structured preference profiles nearby? *Mathematical Social Sciences* 79 (2016), 61–73. https://doi.org/10.1016/j.mathsocsci.2015.11.002 [→ p. 3]

Federica Ceron and Stéphane Gonzalez. 2021. Approval voting without ballot restrictions. *Theoretical Economics* 16, 3 (2021), 759–775. https://doi.org/10.3982/te4087 [→ p. 8]

Cedric Chauve, Ján Maňuch, and Murray Patterson. 2009. On the gapped consecutive-ones property. *Electronic Notes in Discrete Mathematics* 34 (2009), 121–125. https://doi.org/10.1016/j.endm.2009.07.020 [→ p. 3]

Jiehua Chen, Christian Hatschka, and Sofia Simola. 2023. Efficient algorithms for Monroe and CC rules in multi-winner elections with (nearly) structured preferences. In *Proceedings of the 26th European Conference on Artificial Intelligence (ECAI)*. 397–404. https://doi.org/10.3233/FAIA230296 [→ p. 3]

Vincent Conitzer, Matthew Rognlie, and Lirong Xia. 2009. Preference functions that score rankings and maximum likelihood estimation. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*. 109–115. https://www.cs.cmu.edu/~conitzer/preferenceIJCAI09.pdf [→ p. 4]

Franz Dietrich and Christian List. 2010. Majority voting on restricted domains. *Journal of Economic Theory* 145, 2 (2010), 512–543. https://doi.org/10.1016/j.jet.2010.01.003 [→ p. 3]

Michael Dom. 2009. Algorithmic aspects of the consecutive-ones property. *Bulletin of the European Association for Theoretical Computer Science* 98 (2009), 27–59. http://www.mdom.de/fsujena/publications/Dom09_s.pdf [→ p. 4, 23]

Michael Dom, Jiong Guo, and Rolf Niedermeier. 2010. Approximation and fixed-parameter algorithms for consecutive ones submatrix problems. *J. Comput. System Sci.* 76, 3-4 (2010), 204–221. https://doi.org/10.1016/j.jcss.2009.07.001 [→ p. 3]

Edith Elkind and Martin Lackner. 2014. On detecting nearly structured preference profiles. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*. 661–667. https://doi.org/10.1609/aaai.v28i1.8823 [→ p. 3]

Edith Elkind and Martin Lackner. 2015. Structure in dichotomous preferences. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*. 2019–2025. https://arxiv.org/pdf/1505.00341.pdf [→ p. 3]

Edith Elkind, Martin Lackner, and Dominik Peters. 2022. Preference restrictions in computational social choice: A survey. arXiv:2205.09092 [cs.GT] https://arxiv.org/abs/2205.09092 [→ p. 3]

Gábor Erdélyi, Martin Lackner, and Andreas Pfandler. 2017. Computational aspects of nearly single-peaked electorates. *Journal of Artificial Intelligence Research (JAIR)* 58 (2017), 297–337. https://doi.org/10.1613/jair.5210 [→ p. 3, 15]

Bruno Escoffier, Olivier Spanjaard, and Magdaléna Tydrichová. 2021. Measuring nearly single-peakedness of an electorate: Some new insights. In *Proceeding of the 7th International Conference on Algorithmic Decision Theory (ADT)*. 19–34. https://doi.org/10.1007/978-3-030-87756-9_2 [→ p. 3, 4, 6, 15, 17]

Piotr Faliszewski, Edith Hemaspaandra, and Lane A. Hemaspaandra. 2014. The complexity of manipulative attacks in nearly single-peaked electorates. *Artificial Intelligence* 207 (2014), 69–99. https://doi.org/10.1016/j.artint.2013.11.004 [→ p. 4]

Michael R. Garey and David S. Johnson. 1979. *Computers and Intractability: A Guide to the Theory of NP-completeness*. W. H. Freeman and Company. [→ p. 2]

Michael R. Garey, David S. Johnson, and Larry Stockmeyer. 1976. Some simplified NP-complete graph problems. *Theoretical Computer Science* 1, 3 (1976), 237–267. https://doi.org/10.1016/0304-3975(76)90059-1 [→ p. 23]

Mohammad Taghi Hajiaghayi and Yashar Ganjali. 2002. A note on the consecutive ones submatrix problem. *Inform. Process. Lett.* 83, 3 (2002), 163–166. https://doi.org/10.1016/s0020-0190(01)00325-8 [→ p. 3]

Richard M. Karp. 1972. Reducibility among Combinatorial Problems. In *Complexity of Computer Computations*. 85–103. https://doi.org/10.1007/978-1-4684-2001-2_9 [→ p. 23]

Isabelle Lebon, Antoinette Baujard, Frédéric Gavrel, Herrade Igersheim, and Laslier Jean-François. 2017. What approval voting reveals about the preferences of French voters. *Revue économique* 68 (2017), 1063–1076. https://doi.org/10.3917/reco.pr2.0084 [→ p. 5]

Andrew D. Martin and Kevin M. Quinn. 2002. Dynamic ideal point estimation via Markov chain Monte Carlo for the US Supreme Court, 1953–1999. *Political analysis* 10, 2 (2002), 134–153. https://www.jstor.org/stable/25791672 [→ p. 18]

Neeldhara Misra, Chinmay Sonar, and P. R. Vaidyanathan. 2017. On the complexity of Chamberlin–Courant on almost structured profiles. In *Proceedings of the 5th International Conference on Algorithmic Decision Theory (ADT)*. 124–138. https://doi.org/10.1007/978-3-319-67504-6_9 [→ p. 3]

N. S. Narayanaswamy and R. Subashini. 2015. Obtaining matrices with the consecutive ones property by row deletions. *Algorithmica* 71 (2015), 758–773. https://doi.org/10.1007/s00453-014-9925-1 [→ p. 3]

W. M. Flinders Petrie. 1899. Sequences in prehistoric remains. *Journal of the Anthropological Institute of Great Britain and Ireland* (1899), 295–301. https://doi.org/10.2307/2843012 [→ p. 2]

Xin Sui, Alex Francois-Nienaber, and Craig Boutilier. 2013. Multi-dimensional single-peaked consistency and its approximations. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI)*. 375–382. https://www.cs.toronto.edu/~cebly/Papers/SuiEtAl_singlePeaked_ijcai13.pdf [→ p. 17]

Jinsong Tan and Louxin Zhang. 2007. The consecutive ones submatrix problem for sparse matrices. *Algorithmica* 48 (2007), 287–299. https://doi.org/10.1007/s00453-007-0118-z [→ p. 3]

Zoi Terzopoulou, Alexander Karpov, and Svetlana Obraztsova. 2021. Restricted domains of dichotomous preferences with possibly incomplete information. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*. 5726–5733. https://doi.org/10.1609/aaai.v35i6.16718 [→ p. 3]

T. Nicolaus Tideman. 1987. Independence of clones as a criterion for voting rules. *Social Choice and Welfare* 4 (1987), 185–206. https://doi.org/10.1007/bf00433944 [→ p. 12]

Magdaléna Tydrichová. 2023. *Structural and algorithmic aspects of preference domain restrictions in collective decision making: contributions to the study of single-peaked and Euclidean preferences*. Ph. D. Dissertation. Sorbonne Université. https://theses.hal.science/tel-04143294v1/document [→ p. 3]

H. Peyton Young. 1975. Social choice scoring functions. *SIAM J. Appl. Math.* 28, 4 (1975), 824–838. https://doi.org/10.1137/0128067 [→ p. 4]

	VD	MF	BC	MS	FT
aefgbcd	**36**	38	124	126	132
efgabcd	37	**37**	99	119	163
gfabcde	42	42	**88**	108	244
agfbcde	39	39	99	**99**	195
eagfbcd	40	40	122	122	**128**

Table 5: Five axes on the profile P defined in Appendix A.1, and their cost for the different scoring rules. For each axis, we give its score for all axis rules. The optimal values for each rule are given in bold.

A. Appendix of Section 4

A.1. Non-equivalence of Axis Rules

In Section 4 (Example 1), we discussed an example with $m = 4$ for which 3 rules returned different axes. In this section, we provide another example where every pair of rules select different axes. To this purpose, let us consider the profile $P = (18 \times \{a, b\}, N \times \{b, c\}, N \times \{c, d\}, 15 \times \{d, e\}, 4 \times \{e, f\}, 1 \times \{a, g\}, 20 \times \{b, c, f, g\}, 15 \times \{a, e, f, g\}, 2 \times \{a, g, d\})$ on $C = \{a, b, c, d, e, f, g\}$, where N is an integer big enough to ensure bcd is an interval of any optimal axis for any rule (note that such N can always be found).

We give in Table 5 five axes and their respective cost for each scoring rule introduced in Section 4. Note that each of these axes minimizes the cost for a distinct one of the five rules, and it is possible to verify this, either computationally, or by a case-distinction proof (surprisingly, assuming bcd an interval of axis, and using several observations on rules behaviour, the search space can be reduced significantly so there is not so many cases to consider). Instead of giving all details of this proof, we try to put forward some behavioural tendencies of rules in order to better understand their differences.

First, we note that VD and MF often yield the same score. This is because many ballots of P only approve two candidates, in which case the VD score and MF score are equal. More generally, given a ballot A and and axis $\triangleleft$, $s_{\text{VD}}(A, \triangleleft) = s_{\text{MF}}(A, \triangleleft)$ if and only if (1) A is an interval of $\triangleleft$, (2) $\triangleleft$ creates a unique contiguous hole of size one in A, or (3) $\triangleleft$ creates a unique contiguous hole in A and there is only one approved candidate on the left (or on the right) of this hole. To distinguish VD and MF, we have added a ballot $\{a, g, d\}$ to P which creates two contiguous holes on $\underline{a}efgbc\underline{d}$. This ensures that this axis is only optimal for VD.

Let us now focus on differences between MF and BC. Roughly speaking, MF seems somehow more sensitive to the number of holes, while BC seems more sensitive to the sum of the holes sizes. For instance, a ballot associated to the approval vector $(1, 0, 0, 0, 0, 0, 1)$ achieves higher BC-cost than a ballot associated to the approval vector $(1, 0, 0, 0, 0, 1, 0)$, while MF gives them both the same score. On the other hand, BC gives the same score to $(0, 1, 1, 0, 0, 0, 1)$ and $(0, 1, 0, 1, 0, 0, 1)$ while MF gives a better score to the first one. This observation is what allows to distinguish MF and BC, by finding suitable weights of $\{d, e\}$, $\{e, f\}$, and $\{a, d, g\}$ in P.

We then note that BC and MS seems to give similar scores quite often. Actually, given a ballot A and and axis $\triangleleft$, $s_{\text{BC}}(A, \triangleleft) = s_{\text{MS}}(A, \triangleleft)$ if and only if (1) A is an interval of $\triangleleft$, (2) $\triangleleft$ creates a unique continuous hole in A and there is only one approved candidate on the left (or on the right) of this hole, or (3) $\triangleleft$ creates two continuous holes in A and there is a unique approved candidate on the left of the left-most hole and on the right of the right-most hole. As a consequence, the BC and MS cost function return the same values for any approval ballot of size $|A| \leqslant 3$. Thus, only $\{b, c, f, g\}$ and $\{a, e, f, g\}$ are able to distinguish BC and MS.

Finally, FT seems to give more importance to bigger ballots. Indeed, as for each interfering candidate we multiply the number of approved candidates on its left by the number of approved candidates on its right, the FT score becomes more important with the increasing number of approved candidates. The same as in the case of MS, the ballots $\{b, c, f, g\}$ and $\{a, e, f, g\}$ (with suitable weights) were added to P to put into evidence this observation, and to help differentiate FT from other rules. Actually, we note that these ballots are intervals of $eagfbcd$.

A.2. Complexity

In this section we provide the proof that the rules defined in Section 4 are all hard to compute. Most results have already been proven in the literature. In particular, framed as *near-C1P matrices problem*, the VD, MF and BC rules are known to be NP-complete, as stated in (Dom, 2009). Here, we give the reduction for VD and BC and use a simple argument to generalize to all other rules.

We first give the proof for *Voter Deletion*. For this, we recall that a profile is linear if and only if its approval matrix satisfies the C1P. Thus, computing VD is equivalent to the *consecutive ones submatrix* problem, which was already shown to be NP-complete (Booth, 1975). For convenience, we provide a proof here.

Theorem 4 *The VD problem is NP-complete even if each voter approves at most two candidates (i.e.,* $\max_i |A_i| = 2$*).*

Proof. We use a polynomial time reduction from the Hamiltonian path problem, known to be NP-complete Karp (1972). Let $G = (X, E)$ be an undirected graph with $|X| = n$ and $|E| = m$. An Hamiltonian path is a path that visits each vertex exactly once. The Hamiltonian path problem consists in deciding whether such a path exists. The Voter Deletion problem consists in deciding, given as input a profile P and $k \in \mathbb{N}$, whether an axis of score at most k exists. We now show that we can reduce the VD problem from the Hamiltonian path problem.

We create an election with $C = X$ as the set of candidates. Then, we define the profile P as follows: for each edge $(u, v) \in E$, there exist a voter v_e approving $\{u, v\}$. Thus, all voters are distinct. Since the size of all approval ballots is 2, any axis can satisfy at most $n - 1$ pairwise distinct voters, and if $n - 1$ pairwise distinct voters are satisfied by the axis $c_1 \lhd \ldots \lhd c_m$, then $(c_1, \ldots, c_m)$ is a Hamiltonian path. Conversely, if $(c_1, \ldots, c_m)$ is a Hamiltonian path of G, the axis $c_1 \lhd \ldots \lhd c_m$ satisfies $n - 1$ voters with approval ballots of the form $\{c_i, c_{i+1}\}$. Thus there exists a Hamiltonian path if and only if there exist an axis with a Voter Deletion score of $n - 1$ in the election P. This proves that VD is NP-hard. The completeness comes from the fact that it takes polynomial time to compute the VD cost. □

Similarly, one can show that the ballot completion rule is equivalent to the *consecutive ones matrix augmentation* problem, which is also NP-complete. For convinience, we provide a proof here.

Theorem 5 *The BC problem is NP-complete even if each voter approves at most two candidates (i.e.,* $\max_i |A_i| = 2$*).*

Proof. We use a polynomial time reduction from the Optimal Arrangement problem, known to be NP-complete Garey et al. (1976). Let $G = (X, E)$ be an undirected graph with $|X| = n$ and $|E| = m$. The Optimal Arrangement problem decides, given an integer k, whether there is a one-to-one function $f : X \to [1, n]$ such that $\sum_{(u,v) \in E} |f(u) - f(v)| \leqslant k$. The Ballot Completion problem decides, given an integer k, whether there exists an axis of score $\leqslant k$. We create an election with $C = X$ the set of candidates. The set of voters is defined as follows: for each edge $(u, v) \in E$, we introduce a voter v_e with ballot $\{u, v\}$. For an axis $\lhd$, let $f_\lhd(c)$ correspond to the position of the candidate c on the axis (e.g. 1 for the left-most candidate). Then given a ballot $A = \{u, v\}$ and an axis $\lhd$, the Ballot Completion score equals $s_{\text{BC}}(A, \lhd) = |f_\lhd(u) - f_\lhd(v)| - 1$. Thus, the ballot completion score of an axis $\lhd$ with this profile is equal to $\sum_{\{u,v\} \in E} |f_\lhd(u) - f_\lhd(v)| - 1 = \left(\sum_{\{u,v\} \in E} |f_\lhd(u) - f_\lhd(v)|\right) - |E|$. Therefore, there exists an arrangement of cost $\leqslant k$ if and only if there exists an axis of BC cost $\leqslant k - |E|$. Thus, BC is NP-hard. The completeness comes from the fact that the BC cost is computable in polynomial time. $\square$

Now, observe that when $\max_i |A_i| = 2$, VD and MF cost functions are identical. If the ballot A is not an interval of the axis, then it always cost one flip to make it an interval (by flipping one of the two approved candidates). Moreover, observe that MS, FT and BC cost functions are equivalent when $\max_i |A_i| = 2$. It is clear from the formulas, as for all interfering candidates $x \notin A$, $|\{y \in A, y \lhd x\}| = |\{y \in A, x \lhd y\}|) = 1$. This gives the result for all the rules.

B. Omitted Proofs of Section 5

B.1. Neutrality and Consistency with Linearity

Lemma 1 *Let f be a scoring rule. Then, f is neutral and consistent with linearity if and only if it is induced by a cost function s_f such that*

1. *$s_f(A, \lhd) \geqslant 0$, and $s_f(A, \lhd) = 0$ if and only if A is an interval of $\lhd$,*
2. *$s_f(A, \lhd) = s_f(A, \overleftarrow{\lhd})$, and*
3. *there exists a function $g : \{0,1\}^m \to \mathbb{R}_{\geqslant 0}$ such that $s_f(A, \lhd) = g(x_{A,\lhd}) = g(x_{A,\overleftarrow{\lhd}})$ (so s_f depends only on the induced approval vector $x_{A,\lhd}$).*

Proof. Let f be a scoring rule induced by the cost function s_f. Let's show (1) first. Assume for contradiction that there is an axis $\lhd$ and a ballot A such that $s_f(A, \lhd)$ attains its minimal value but A is not an interval of $\lhd$. Then, on the linear profile $P = \{A\}$, we have $\lhd \in f(P)$, which is a contradiction with consistency with linearity. Similarly, if A is an interval of $\lhd$ but $s_f(A, \lhd)$ is not minimal, then on the linear profile $P = \{A\}$, we do not have $\lhd \in f(P)$ while $\lhd$ is consistent with P, a contradiction.

Moreover, note that the cost function s'_f such that for all ballots A and axis $\lhd$, $s'_f(A, \lhd) = s_f(A, \lhd) - \min_{\lhd'} s_f(A, \lhd')$ yields a cost function still inducing f. Indeed, for all profiles P, $s'_f(P, \lhd) = s_f(P, \lhd) - \sum_{A_i \in P} \min_{\lhd'} s_f(A_i, \lhd')$ so the optimal axes are the same for both functions. Thus, we assume wlog that $s_f(A, \lhd) = 0$ if and only if A is an interval of $\lhd$.

We now show (2). If A is an interval of $\lhd$, it is also an interval of $\overleftarrow{\lhd}$, so we clearly have $s_f(A, \lhd) = s_f(A, \overleftarrow{\lhd})$. Assume now that A is not an interval of $\lhd$. Thus, $y = s_f(A, \lhd) > 0$ and $y' = s_f(A, \overleftarrow{\lhd}) > 0$. Assume for contradiction that $y \neq y'$, and wlog that $y < y'$. Let us denotes the candidates $c_1, \dots, c_m$ such that $\lhd = c_1 c_2 \dots c_m$. Moreover, let z be the minimal value of $s_f(A, \lhd)$ for a ballot A that is not an interval of $\lhd$. We know that $z > 0$. Take $q \in \mathbb{N}$ such that $q > y/z$ and consider the profile P which contains A and for each $i \in [1, m-1]$, q ballots $\{c_i, c_{i+1}\}$. Clearly, any axis $\lhd' \notin \{\lhd, \overleftarrow{\lhd}\}$ is breaking at least one pair, inducing a cost greater that $q \cdot z > y$. The cost of $\lhd$ is y and the cost of $\overleftarrow{\lhd}$ is $y' > y$. Thus, $f(P) = \{\lhd\}$ which contradicts the definition of axis rules. Therefore, $y = y'$ and $s_f(A, \lhd) = s_f(A, \overleftarrow{\lhd})$.

Finally, we show (3), i.e. that f only depends on the approval vectors $x_{A,\lhd}$ of the ballots in the profile. For this, we show that f is induced by a cost function s^*_f such that $s^*_f(A, \lhd)$ only depends on $x_{A,\lhd}$. Let Π be the set of all candidate permutations, and define $s^*_f(A, \lhd) = \sum_{\pi \in \Pi} s_f(\pi(A), \pi(\lhd))$ for all A and $\lhd \in \mathcal{A}$, with $\pi(A) = \{\pi(a) : a \in A\}$ and $\pi(\lhd) = \pi(c_1) \dots \pi(c_m)$ for $\lhd = c_1 \dots c_m$. s^*_f clearly satisfies conditions (1) and (2) of the Lemma. We now show that it also satisfies the condition (3).

Take any $A, \lhd$ and $A', \lhd'$ with the same approval vector, i.e., $x_{A,\lhd} = x_{A',\lhd'}$. Then, there exists a permutation $\tau \in \Pi$ with $\tau(A) = A'$ and $\tau(\lhd) = \lhd'$. Thus, we obtain that $s^*_f(A', \lhd') = s^*_f(\tau(A), \tau(\lhd)) = \sum_{\pi \in \Pi} s(\pi(\tau(A)), \pi(\tau(\lhd))) = \sum_{\pi' \in \Pi} s(\pi'(A), \pi'(\lhd)) = s^*_f(A, \lhd)$.

To show that f is still induced by this rule, let $\lhd \in f(P)$ be an optimal axis for profile P. Then, by neutrality, $\pi(\lhd) \in f(\pi(P))$ for all $\pi \in \Pi$. This implies that $s_f(\pi(\lhd), \pi(P)) \geqslant s_f(\pi(\lhd'), \pi(P))$ for all axes $\lhd' \in \mathcal{A}$. Since this inequality carries over to the sum over all $\pi \in \Pi$, this implies $s^*_f(\lhd, P) \geqslant s^*_f(\lhd', P)$ for all $\lhd'$. For the other direction, let $\lhd' \notin f(P)$ and fix some $\lhd \in f(P)$. With the same argument, we obtain $s^*_f(\lhd, P) > s^*_f(\lhd', P)$, which shows that, for all profiles, an axis $\lhd$ has minimal cost w.r.t. s^*_f iff it is chosen by f. Thus, we can assume wlog that the cost function s_f inducing f only depends on the approval vectors.

For the other direction, assume that $s_f(A, \lhd) = 0$ if and only if A is an interval of $\lhd$. Let us show that f satisfy consistency with linearity. Let P be a linear profile. If $\lhd$ is consistent with P, for all $A_i \in P$, $s_f(A_i, \lhd) = 0$, so $s_f(P, \lhd) = 0$. If $\lhd$ is not consistent with P, there is some $A_i \in P$ such that $s_f(A_i, \lhd) > 0$, so $s_f(P, \lhd) > 0$. Thus, $f(P) = \mathit{con}(P)$. To show neutrality, we simply need to use (3). If we rename the candidates on A and $\lhd$ to obtain A' and $\lhd'$ we clearly have the same approval vector $x_{A,\lhd} = x_{A',\lhd'}$, so $s_f(A, \lhd) = s_f(A', \lhd')$. This implies neutrality, and concludes the proof. □

B.2. Characterization of Voter Deletion

Theorem 3 *Let $m \geqslant 6$, and let f be a neutral scoring rule. Then f satisfies consistency with linearity, ballot monotonicity, and resistance to cloning if and only if it is VD.*

Proof. We already showed that VD satisfies all the axioms. For the other direction, let f be a scoring rule satisfying neutrality, consistency with linearity, resistance to cloning and ballot monotonicity. As shown in Appendix B.1, f is induced by a symmetric cost function s with $s(A, \lhd) = 0$ iff A forms an interval in $\lhd$. Further, s only depends on the approval vector $x_{A,\lhd}$, i.e. there exists a function $g : \{0,1\}^m \to \mathbb{R}_{\geqslant 0}$ such that $s(A, \lhd) = g(x_{A,\lhd})$ for all ballots A and axis $\lhd$.

The steps of the proof are as follows:

1. Using ballot monotonicity, we show that there is a function h such that for all A and $\lhd$ such that A is not an interval of $\lhd$, $s(A, \lhd) = h(m, k_{\text{app}}, k_{\text{int}})$, where m is the number of candidates, $k_{\text{app}} = |A|$ is the number of approved candidates and k_{int} is the number of interfering candidates.

2. Using resistance to cloning, we show that for A not interval of $\lhd$, $s(A, \lhd)$ only depends on the sum $k_{\text{app}} + k_{\text{int}}$, i.e. there is h such that $s(A, \lhd) = h(m, k_{\text{app}} + k_{\text{int}})$.

3. We show that for A not interval of $\lhd$, $s(A, \lhd)$ can only take two values: $s(A, \lhd) = h^*_m$ if $k_{\text{app}} + k_{\text{int}} = m$ and $s(A, \lhd) = h_m$ otherwise.

4. Finally, we show that $h^*_m = h_m$ and thus the rule is VD.

In all of the following, A and $\lhd$ are chosen such that A is not an interval of $\lhd$, and thus we already know that $s(A, \lhd) > 0$. We start with **Step 1** by applying ballot monotonicity to a very symmetric profile.

Lemma 2 *There is a function h such that $s(A, \lhd) = h(m, k_{app}, k_{int})$ where m is the number of candidates, $k_{app} = |A|$ and k_{int} is the number of interfering candidates.*

Proof. Assume for contradiction that there are A, A', $\lhd$ and $\lhd'$ such that $s(A, \lhd) \neq s(A', \lhd')$ but $|A| = |A'| = k_{\text{app}}$ and the number of interfering candidates k_{int} is the same in both cases. Denote $x = x_{A,\lhd}$ and $x' = x_{A',\lhd'}$ the respective approval vectors. Obviously, $x \neq x'$.

Consider the set of candidates $\mathcal{C} = \{c_1, \dots, c_{k_{\text{app}}}\} \cup \{d_1, \dots, d_{k_{\text{int}}}\} \cup \{b_1, \dots, b_{m-(k_{\text{app}}+k_{\text{int}})}\} = C \cup D \cup B$. Let P^* be the profile in which every ballot A of size $|A| = k_{\text{app}}$ is approved by one voter. By neutrality, all axes are chosen.

Now, let P be the profile where the voter with ballot C changes it to $C \cup D$. By ballot monotonicity, for any $\lhd$ such that the set of interfering candidates of C on $\lhd$ is D, then $\lhd$ must remain chosen. In other words, these are axes such that $\{d \notin C : \exists c_i, c_j \in C, c_i \lhd d \lhd c_j\} = D$.

Let $\lhd^1$ and $\lhd^2$ be such axes, but such that the approval vectors of C on these axes are $x_{C,\lhd^1} = x$ and $x_{C,\lhd^2} = x'$. This is possible since there are k_{app} approved candidates in C and k_{int} interfering candidates with respect to both axes. Then we have that $s(P^*, \lhd^1) = s(P^*, \lhd^2)$ and $s(P, \lhd^1) = s(P, \lhd^2)$. However, for any axis $\lhd$ in that case, we have $s(P^*, \lhd) = s(P, \lhd) + s(C, \lhd)$. Thus, $s(C, \lhd^1) = s(C, \lhd^2)$ which means $g(x) = g(x')$. This contradicts $s(A, \lhd) \neq s(A', \lhd')$, and concludes the proof of the lemma. □

In **Step 2**, we use resistance to cloning to show that the cost actually only depends on $k_{\text{app}} + k_{\text{int}}$

Lemma 3 *There is h such that $s(A, \lhd) = h(m, k_{app} + k_{int})$.*

Proof. For $m = 3$, all scoring rules are equivalent. Let $m \geqslant 4$. Let $3 \leqslant k_{\text{app}} + k_{\text{int}} \leqslant m$. We know from step 1 that there exists h such that $s(A, \lhd) = h(m, k_{\text{app}}, k_{\text{int}})$ for all A and $\lhd$. In this proof, we show that for all k_{app} and k_{int}, we have $h(m, k_{\text{app}}, k_{\text{int}}) = h(m, k_{\text{app}} + k_{\text{int}} - 1, 1)$, implying that the cost function only depends on $k_{\text{app}} + k_{\text{int}}$.

This is clearly true for $k_{\text{app}} + k_{\text{int}} = 3$ as in this case the only possibility is $k_{\text{app}} = 2$ and $k_{\text{int}} = 1$, otherwise the ballot is an interval. For $k_{\text{app}} + k_{\text{int}} > 3$, let the set of candidates be $\mathcal{C} = \{c_1, \dots, c_{k_{\text{app}}}\} \cup \{d_1, \dots, d_{k_{\text{int}}}\} \cup \{b_1, \dots, b_{m-(k_{\text{app}}+k_{\text{int}})}\} = C \cup D \cup B$. Note that if $k_{\text{app}} + k_{\text{int}} = m$, $B = \emptyset$.

We assume that $k_{\text{app}} + k_{\text{int}} < m$. The proof if it is equal to m is the same – if not simpler. Let z be the minimal cost $s(A, \lhd)$ over all approval ballots A and axes $\lhd$ such that A is not an interval of $\lhd$ (for less than m candidates), and $y = \max(h(4, 2, 1), h(m, k_{\text{app}}, k_{\text{int}}), h(m, k_{\text{app}} + k_{\text{int}} - 1, 1))$. Let $q \in \mathbb{N}$ such that $q > y/z$ and consider the following profile P on $\{c_1, c_2, d_1, b_1\}$.

$$\begin{aligned}
&q \times \{c_1, d_1\} \\
&q \times \{c_1, d_1, c_2\} \\
&q \times \{d_1\} \\
&q \times \{c_2\} \\
&q \times \{b_1\} \\
&1 \times \{c_1, c_2\} \\
&1 \times \{d_1, c_2\}
\end{aligned}$$

Clearly, any axis such that $\{c_1, d_1\}$ or $\{c_1, d_1, c_2\}$ do not form an interval has cost greater than $q \cdot z > y$. The other axes (up to reversal) are:

$$\begin{aligned}
&\underline{c_1} d_1 \underline{c_2} b_1 \\
&b_1 \underline{c_1} d_1 \underline{c_2} \\
&\underline{d_1} c_1 \underline{c_2} b_1 \\
&b_1 \underline{d_1} c_1 \underline{c_2}
\end{aligned}$$

They each break one of $\{c_1, c_2\}$ or $\{d_1, c_2\}$ with cost $h(4, 2, 1) \leqslant y$, thus they all are optimal.

Now, clone c_2 into $\{c_2, \dots, c_{k_{\text{app}}}\}$, d_1 into $\{d_1, \dots, d_{k_{\text{int}}}\}$ and b_1 into $\{b_1, \dots, b_{m-(k_{\text{app}}+k_{\text{int}})}\}$. Clearly, all clones of each category need to be next to each other on the axis, otherwise the q

ballots containing these clones (obtained from $\{d_1\}$, $\{c_2\}$ or $\{b_1\}$) would induce a cost greater than $q \cdot z > y$. Combining this with resistance to cloning gives that the axes should be of one of the following forms (up to change of the positions of the clones):

$$\underline{c_1} d_1 \dots d_{k_{\text{int}}} \underline{c_2 \dots c_{k_{\text{app}}}} b_1 \dots b_{m-(k_{\text{app}}+k_{\text{int}})}$$
$$b_1 \dots b_{m-(k_{\text{app}}+k_{\text{int}})} \underline{c_1} d_1 \dots d_{k_{\text{int}}} \underline{c_2 \dots c_{k_{\text{app}}}}$$
$$\underline{d_1 \dots d_{k_{\text{int}}}} c_1 \underline{c_2 \dots c_{k_{\text{app}}}} b_1 \dots b_{m-(k_{\text{app}}+k_{\text{int}})}$$
$$b_1 \dots b_{m-(k_{\text{app}}+k_{\text{int}})} \underline{d_1 \dots d_{k_{\text{int}}}} c_1 \underline{c_2 \dots c_{k_{\text{app}}}}$$

Indeed, for all these axes, there is only one ballot that is not an interval. The first two axes break the one obtained from $\{c_1, c_2\}$ (by adding clones) with k_{app} approved candidates and k_{int} interfering ones, while the last two break the ballot obtained from $\{d_1, c_2\}$ with $k_{\text{app}} + k_{\text{int}} - 1$ approved candidates and 1 interfering one (c_1). Thus, the first two have cost $h(m, k_{\text{app}}, k_{\text{int}}) < y$ and the last two have cost $h(m, k_{\text{app}} + k_{\text{int}} - 1, 1) < y$. Thus, they are respectively the axes with lowest cost that can be reduced to the axes obtained with 4 candidates $\{c_1, c_2, d_1, b_1\}$. By resistance to cloning, this means that each of these axes should be among the optimal ones. This directly implies that $h(m, k_{\text{app}}, k_{\text{int}}) = h(m, k_{\text{app}} + k_{\text{int}} - 1, 1)$. The proof if $m = k_{\text{app}} + k_{\text{int}}$ is exactly the same, but without candidates b_i.

This prove that given $k \leqslant m$, for all k_{app} and k_{int} such that $k_{\text{app}} + k_{\text{int}} = k$, we have $h(m, k_{\text{app}}, k_{\text{int}}) = h(m, k-1, 1)$. And thus, we can say that the function actually only depends on k, and that there is h such that $s(A, \lhd) = h(m, k_{\text{app}} + k_{\text{int}})$ for all A and $\lhd$. □

We now proceed to **Step 3**, in which we show that for given m, the cost only depends on whether both extremes of the axis are approved.

Lemma 4 *There are constants such that $s(A, \lhd) \in \{h^*_m, h_m\}$, where the score is equal to h^*_m if A contains both extremes of $\lhd$ and is equal to h_m otherwise.*

Proof. We know that there is h such that $s(A, \lhd) = h(m, k_{\text{app}} + k_{\text{int}})$. In the rest of the proof, we denote $k = k_{\text{app}} + k_{\text{int}}$ We can simply set $h^*_m = h(m, m)$, as A contains both extremes of $\lhd$ if and only if $k = m$. For $4 \leqslant k < m$, we will show that $h(m, k) = h(m, k-1)$. Since we know that $k \geqslant 3$ (otherwise the ballot is an interval), this implies that there is h_m such that for all $k \in [3, m-1]$, $h(m, k) = h_m$. Let $k \in [4, m-2]$ (the proof for $k = m-1$ is identical, but without the candidates b_i).

As in the proof of the previous lemma, let z be the minimum cost for a non-interval ballot of an axis (for less than m candidates) and $y = \max(h(m, 3), h(m, k), h(m, k-1))$. Let $q \in \mathbb{N}$ such that $q > y/z$ and consider the following profile P over the set of candidates $C = \{a_1, b_1, c_1, c_2, d_1\}$:

$$q \times \{d_1, c_2\}$$
$$q \times \{c_2, a_1\}$$
$$q \times \{c_1, d_1, c_2, a_1\}$$
$$q \times \{d_1\}$$
$$q \times \{c_2\}$$
$$q \times \{b_1\}$$
$$1 \times \{c_1, c_2\}$$

On this profile, any axis breaking one of the ballots of the first three categories induce a cost greater than $q \cdot z > y$. The only axes that do not break these ballots are the following:

$$\underline{c_1} d_1 \underline{c_2} a_1 b_1$$
$$b_1 \underline{c_1} d_1 \underline{c_2} a_1$$
$$d_1 \underline{c_2} a_1 \underline{c_1} b_1$$
$$b_1 d_1 \underline{c_2} a_1 \underline{c_1}$$

Note that the only ballot that is not an interval of these axes is $\{c_1, c_2\}$ with a cost of $h(5, 3) \leqslant y$. Thus, all these axes are optimal for f.

Now, clone c_2 into $\{c_2, \dots, c_{k-2}\}$, d_1 into $\{d_1, d_2\}$ and b_1 into $\{b_1, \dots, b_{m-k-1}\}$. Clearly, all clones of each category need to be next to each other on the axis, otherwise the q ballots containing these clones (obtained from $\{d_1\}$, $\{c_2\}$ or $\{b_1\}$) would induce a cost greater than $q \cdot z > y$. Combining this with resistance to cloning gives that the axes should be of one of the following forms (up to change of the positions of the clones):

$$\underline{c_1} d_1 d_2 \underline{c_2 \dots c_{k-2}} a_1 b_1 \dots b_{m-k-1}$$
$$b_1 \dots b_{m-k-1} \underline{c_1} d_1 d_2 \underline{c_2 \dots c_{k-2}} a_1$$
$$d_1 d_2 \underline{c_2 \dots c_{k-2}} a_1 \underline{c_1} b_1 \dots b_{m-k-1}$$
$$b_1 \dots b_{m-k-1} d_1 d_2 \underline{c_2 \dots c_{k-2}} a_1 \underline{c_1}$$

Indeed, for all these axes, there is only one ballot that is not an interval: $\{c_1, c_2, \dots, c_{k-2}\}$ with cost $h(m, k)$ for the first two axes and $h(m, k-1)$ for the last two axes. In both cases, this cost is $< y$, so they are respectively the axes with lowest cost that can be reduced to the axes obtained with 5 candidates $\{a_1, b_1, c_1, c_2, d_1\}$. By resistance to cloning, this means that each of these axes should be among the optimal ones. This implies that $h(m, k) = h(m, k-1)$. The proof if $k = m-1$ is exactly the same, but without candidates b_i.

This proves that for a given m, there exist some value h_m such that for all non interval ballot A on $\lhd$ with $k_{\text{app}} + k_{\text{int}} < m$, we have $s(A, \lhd) = h_m$. □

Finally, in **Step 4** we prove that for all $m \geqslant 4$, $h_m = h^*_m$. It consists in three substeps: (i) for all $m \geqslant 4$, $h^*_m \leqslant h_m$, (ii) for all $m \geqslant 6$, $h^*_m \geqslant h_m$ and (iii) for all $m \geqslant 4$, if $h^*_{m+1} = h_{m+1}$, then $h^*_m = h_m$. As for $m = 3$ all rules are equivalent, this is enough to characterize VD for all m.

Lemma 5 *$h^*_m \leqslant h_m$ for all $m \geqslant 4$.*

Proof. First, we show $h^*_m \leqslant h_m$ for all $m \geqslant 4$. Let $m \geqslant 4$, and assume for contradiction that $h^*_m > h_m$. Again, $z > 0$ indicates the minimal cost of a non-interval ballot on any axis for at most m candidates. Let $q \in \mathbb{N}$ such that $q > \max(h_m, h_{m+1})/z$, and consider the profile P on m candidates $\mathcal{C} = \{c_1, c_2\} \cup \{b_1\} \cup \{d_1, \dots, d_{m-3}\}$ with $D = \{d_1, \dots, d_{m-3}\}$:

$$q \times D \cup \{c_2\}$$
$$q \times D \cup \{c_1, b_1\}$$
$$1 \times \{c_1, c_2\}$$

Note that all axes that breaks one of the first two ballots have cost at least $q \cdot z > h_m$. The other axes are of the following form (up to change of positions among candidates of D):

$$\underline{c_1} b_1 d_1 \dots d_{m-3} \underline{c_2}$$
$$b_1 \underline{c_1} d_1 \dots d_{m-3} \underline{c_2}$$

The only ballot that is not an interval of these axis is $\{c_1, c_2\}$ with a cost of h^*_m on the first axis and h_m on the second one. Since $h_m < h^*_m$, this means only the second axis is optimal.

Now, let's clone b_1 into $\{b_1, b_2\}$. Again, all axes that do not comply with the ballots of the first two categories have cost higher than $q \cdot z > h_{m+1}$. The only other axes that generalizes $b_1 c_1 d_1 \dots d_{m-3} c_2$ and do not break the ballots of the first two categories are:

$$\underline{b_2} b_1 c_1 d_1 \dots d_{m-3} c_2$$
$$b_1 \underline{b_2} c_1 d_1 \dots d_{m-3} c_2$$
$$b_1 c_1 \underline{b_2} d_1 \dots d_{m-3} c_2$$

The cost of $\{c_1, c_2\}$ on any of these axes is h_{m+1}. By resistance to cloning, at least one of them should be among the optimal axes. Since they all have the same cost, they all are optimal for this profile.

Now, let us remove the clone b_1 of b_2. By resistance to cloning, the following two axes should be among the optimal axes:

$$\underline{b_2} c_1 d_1 \dots d_{m-3} c_2$$
$$c_1 \underline{b_2} d_1 \dots d_{m-3} c_2$$

Since the only ballot that is not an interval of these axes is $\{c_1, c_2\}$, there respective costs are h_m and h^*_m. However, we assumed that $h^*_m > h_m$, so the second axis cannot be optimal, a contradiction. □

Lemma 6 *$h^*_m \geqslant h_m$ for all $m \geqslant 6$.*

Proof. Let $m \geqslant 5$ and assume by contradiction that $h^*_{m+1} < h_{m+1}$. Consider the set of candidates $\mathcal{C} = \{c_1, c_2\} \cup \{d_1, d_2\} \cup \{b_1, \dots b_{m-4}\}$ with $B = \{b_1, \dots b_{m-4}\} \neq \emptyset$. Again, let z be the lowest cost of any non-interval approval ballot on any axis for less than m candidates. Let $y = \max(h^*_m, h^*_{m+1})$ and $q \in \mathbb{N}$ such that $q > y/z$, and consider the following profile:

$$q \times \{d_1, c_2\}$$
$$q \times \{d_2, c_2\}$$
$$q \times \{b_1, \dots, b_{m-4}\}$$
$$1 \times \{c_1, c_2\}$$

Again, any axis breaking one of the ballots of the first three categories induces a cost of at least $q \cdot z > y$. These ballots are intervals of an axis $\lhd$ if $\lhd$ contains the interval $d_1 \lhd c_2 \lhd d_2$ and the set B forms an interval. It can have B before or after $d_1 \lhd c_2 \lhd d_2$ on the axis, and c_1 between the two intervals or on one extremity of the axis. In any axis of this kind, the only ballot that is not an interval is $\{c_1, c_2\}$ with cost $h_m \leqslant y$, since c_2 is not an extremity of the axis. Thus, all these axes are selected by the rule. In particular, the axis $\lhd^* = d_1 \underline{c_2} d_2 \underline{c_1} b_1 \dots b_{m-4}$ is selected.

Let us now clone c_1 into $\{c_1, c_3\}$, we obtain $m + 1 \geqslant 6$ candidates. Any axis generalizing $\lhd^*$ breaks at least the ballot $\{c_1, c_2, c_3\}$, and at least one extreme of the axis are not part of the ballot (since c_1 and c_2 are not on the extremes). Thus, the cost of any axis generalizing $\lhd^*$ is at least h_{m+1}. However, consider the axis $\lhd' = \underline{c_1} d_1 \underline{c_2} d_2 b_1 \dots b_{m-4} \underline{c_3}$. The only ballot that is not an interval of this axis is $\{c_1, c_2, c_3\}$, and both extremes of the axis are approved, so the cost is $h^*_{m+1} < h_{m+1}$. This imply that no axis generalizing $\lhd^*$ can be selected as they do not have lowest cost. This contradicts resistance to cloning. Therefore, $h^*_{m+1} \geqslant h_{m+1}$. □

We now know that $h_m = h^*_m$ for all $m \geqslant 6$ by combining the last two lemmas. Finally, we show that for $m > 4$, $h_m = h^*_m$ implies $h_{m-1} = h^*_{m-1}$. For this, take $m > 4$ and consider the profile P

defined in Lemma 5. We remind that the two axes that do not break the ballots appearing q times are of the following form (up to change of positions among candidates):

$$\underline{c_1} b_1 d_1 \dots d_{m-3} \underline{c_2}$$
$$b_1 \underline{c_1} d_1 \dots d_{m-3} \underline{c_2}$$

The cost of these axes are respectively h^*_m and h_m so they both are optimal since $h_m = h^*_m$. If we remove the clone d_{m-3} of $\{d_1, \dots, d_{m-4}\}$, by resistance to cloning the following two axes should be selected by the rule:

$$\underline{c_1} b_1 d_1 \dots d_{m-4} \underline{c_2}$$
$$b_1 \underline{c_1} d_1 \dots d_{m-4} \underline{c_2}$$

The cost of these two axes are respectively h^*_{m-1} and h_{m-1}. Thus, $h^*_{m-1} = h_{m-1}$.

This implies that for all $m \geqslant 4$, the cost is 0 if the ballot is an interval of the axis, and h_m otherwise. Without loss of generality, we take $h_m = 1$. For $m = 3$, the only approval vector that is induced by non interval ballots is $(1, 0, 1)$, and we can assume without loss of generality that its cost is 1. Thus, f is equal to VD. □

For sharpness of the characterization among neutral scoring rules, the trivial rule TRIV returning all axes satisfies every axiom but consistency with linearity, the genus rule returning the number of continuous holes as score[2] only fails ballot monotonicity, and the BC rule only violates resistance to cloning. Note that if we allow an infinitely large ground set of candidates, then resistance to cloning and consistency with linearity imply neutrality for scoring rules.

B.3. Supplementary Result: Resistance to cloning implies almost topological.

In this section, we investigate the class of rules satisfying resistance to cloning. A scoring rule f belongs to the class of *topological rules*, if there is a monotone function h such that $s_f(A, \lhd) = h(k)$ for all $A, \lhd$, where k is the number of continuous holes that A creates in $\lhd$.

The following axiom is the (very mild) counterpart to clearance: While the latter demands that only axes can be chosen in which unapproved candidates are not interfering, the following axiom demands that many such axes must be included in the choice set. In contrast to clearance, inclusion clearance is satisfied by all five introduced rules in this paper.

Inclusion Clearance We say that a rule f satisfies *inclusion clearance* if the following holds: let X be the set of candidates that are never approved in P. Then there is $\lhd \in f(P)$ such that there is no $A \in P$ with $y, z \in A$, $x \in X$ and $y \lhd x \lhd z$. Further, all other $\lhd'$ that have X on the extremes and coincide with $\lhd$ on $C \setminus X$ are chosen too.

Theorem 6 *Let f be a faithful and neutral scoring rule. If f satisfies cloning consistency and inclusion-clearance, then there are h^*, h, such that $s(A, \lhd) = h^*(n)$ if A approves both extremes of $\lhd$, and $s(A, \lhd) = h(n)$ else, where n is the number of continuous holes A creates in $\lhd$. Further, $h^*(n) \leqslant h(n) \leqslant h^*(n + 1)$ for all n.*

Since many arguments remain similar to the ones in Appendix B.2, we only provide the outline of the proof.

Proof. The steps are as follows:

1. There is h such that $s(A, \lhd) = h(n, x_1, \dots, x_{n+1}, y_1, \dots y_n, i, m)$ (or $s(A, \lhd) = 0$).

[2]This is the scoring rule with $s_G(A, \lhd) = |\{(x, y) \in A : \exists z, x \lhd z \lhd y \text{ and } \forall z \text{ s.t. } x \lhd z \lhd y, z \notin A\}|$

- where n is the number of holes m is the number of candidates present in the axis $\lhd$, x_i is the cardinality of the $i-th$ approved interval, and y_i is the cardinality of the $i-th$ hole.

2. There is h such that $s(A, \lhd) = h(n, x, y_1, \dots y_n, m)$
   - where $x = |A| = x_1 + \cdots + x_{n+1}$
3. There is h such that $s(A, \lhd) = h(n, x + y, m)$
   - where $y = y_1 + \dots y_n$ is the number of interfering candidates.
4. There is h such that $s(A, \lhd) = h(n, i, m)$
   - where $i = 1$ if A contains both extremes of $\lhd$ and $i = 0$ else.
5. There is h such that $s(A, \lhd) = h(n, i)$.
6. For all n, we have $h(n, 1) \leqslant h(n, 0) \leqslant h(n + 1, 1)$.

Step 1 follows from inclusion-clearance and neutrality.

Step 2 follows from two lemmas that work similarly to the characterization of VD. The first shows that $h(n, x_1, \dots, x_m, y_1, \dots y_n, m) = h(n, x_1 \pm 1, \dots, x_m \mp 1, y_1, \dots y_n, m)$, while the second shows that we can invert the first r approved intervals, i.e., $h(n, x_1, \dots, x_m, y_1, \dots y_n, m) = h(n, x_r, \dots, x_1, \dots, x_m, y_{r-1}, \dots, y_1, y_r, \dots y_n, m)$.

Step 3 follows from a lemma which works similarly to the characterization of VD. There, we can flip x_1 and y_1 and still obtain the same score. Thus, further combined with the previous two lemmas, we obtain $h(n, x, y_1, \dots y_n, m) = h(n, x + y - n, n, m)$.

Step 4 works again exactly as in the characterization of VD. We take one hole of size 2 and one of size 1.

Step 5 This uses a new construction. First, norm $h(1, 0, m) = 1$ for all m. Then show (by induction) that $h(n + 1, 0, m) - h(1, 0, m) = h(n + 1, 0, m + 1) - h(1, 0, m + 1)$ for all n, m.

For this, Consider two axes with a single swap difference $12 \dots m$, $21 \dots m$ and take the ballot $\{1, 3, 5, \dots\}$. Then, assume that the differences are not equal, use this to create a profile where one of these two axes is chosen but then the other and thus cloning consistency is violated.

Step 6 The right inequality follows from weak clearance and the left is obtained exactly as in the VD characterization. □

We can further restrict the class of scoring rules to the class of *local* scoring rules. For these, the score does not depend on non-interfering candidates. Formally, a neutral scoring rule is local if $g(x_{A,\lhd}) = g(x')$ for all $A, \lhd$, where x' is the subvector of $x_{A,\lhd}$ where the non-interfering 0's are cut off. Note that as long as we use the same rule for all feasible sets, all five introduced rules satisfy locality.

Clearly, among local scoring rules, Theorem 6 turns into a characterization. It is open whether locality is required.

C. Details of the Experiments

C.1. Implementation

In this section, we explain the methods we used for implementing the rules. We focus here on explaining our approach to reduce the runtime. First, we present how we improved the brute-force method to be usable in all our experiments. Then, we explain the implementation of the Integer Linear Programming (ILP) encodings which we used for two rules: Voter Deletion and Ballot Completion.

C.1.1. Brute-force Method

The brute-force method is straightforward: compute the cost of all the axes for the given profile, and return the ones with minimal cost. However, this approach takes time exponential with m, and hence is not usable in practice even for relatively small values of m. Thus, we used pruning methods and heuristics.

We start by pre-processing the approval profile. We assume that each voter has a weight w_i (potentially all weights are equal to one). Then, we aggregate the weights of all the voters whose approval ballots are identical. For instance, if two voters i and j have the same ballot $A_i = A_j$, we replace them by a unique voter having the ballot $A_i = A_j$ and the weight $w_i + w_j$. Moreover, we remove all ballots that are intervals of any axis, and hence do not help to identify the axes with minimal cost: namely, these are empty ballots, singletons and full ballots ($A = \mathcal{C}$).

We keep a variable containing the lowest cost found *so far*, as well as a variable containing all axes with this cost. Every time we compute the cost of an axis for a given profile by adding up the costs of the ballots, the sum might surpass this value before we read the whole profile. We can then move to the next axis. To save as much running time as possible, we order the ballots by decreasing weights so that we start by the ballots of highest weight.

Another similar method that reduces the running time is the following: for any axis on m candidates, we can compute a lower bound of its cost by removing two candidates from the profile and computing the cost of the reduced axis on the $m - 2$ remaining candidates. In practice, we group axes into sets of $(m - 1)(m - 2)/2$ axes (one for each position of the missing pair of candidates) and if we observe that the cost of their common reduced axis is higher than the current lowest cost, this means that no axis of this set will be optimal and that we can completely skip all of them.

Finally, we can initialize the current lowest cost axis with an axis we suppose to be good. For instance, for political datasets, we can use the axes adopted by the media.

Combining all these strategies, we have never needed more than one hour to find optimal axes for profiles on up to 12 candidates. For less than 7 candidates, the result was always returned in less than one second.

C.1.2. ILP Encoding

We also implemented an ILP formulation of Voter Deletion rule and Ballot Completion rule. In this section we briefly describe them.

First, we conduct the same pre-processing on the approval profile by merging the weights of identical ballots. Then, we create a binary variable $x_{a,b}$ such that $x_{a,b} = 1$ if and only if $a \lhd b$ on the axis. Then, the formulation is different for VD and BC:

- **VD:** For each voter i, we add a binary variable y_i such that $y_i = 1$ if and only if the ballot A_i is an interval of the axis (using an appropriate inequality on the variables $x_{a,b}$). Finally, the cost is equal to the sum of the $w_i \cdot (1 - y_i)$ where w_i is the weight of ballot i.

- **BC:** For each candidate $a \in C$, we introduce an integer variable $p_a \in [0, m-1]$ that encodes the position of the candidate on the axis ($p_a = \sum_b x_{b,a}$). Then, for each ballot A_i, we define two variables M_i and m_i respectively for the right-most and left-most position of candidates approved in A_i. The BC cost of the ballot A_i is then given by $M_i - m_i - |A_i| + 1$. Finally, we sum this cost over all ballots (multiplied by the weights w_i) to obtain the overall cost.

C.2. Synthetic Data

In this section, we present more experiments and results on synthetic data models. In particular, we compare our rules on a variety of models. For each model, and a given *ground truth axis* (drawn uniformly at random), the approval ballots are sampled i.i.d. according to this axis.

As we already mentioned in Section 3, any scoring rule is actually the Maximum Likelihood Estimator of some probability distribution model over the approval ballots, so the performance of the rules are likely to reflect simply how similar they are to the MLE of the models used. However, these experiments can give an idea of how well the rules can generalize to different models. In this section, we study four models, each inspired by one of our rules (but the rules are not necessarily exactly the MLE of the models).

- **Maverick Voters:** In this model, we sample a *maverick voter* (i.e. an approval ballot uniformly at random) with probability $p \in [0, 1]$. We sample an *interval voter* (i.e. an approval ballot that is an interval of $\lhd$ uniformly at random) with probability $1 - p$.

- **Random Flips:** In this model, we first sample for each voter an approval ballot that is an interval of the axis $\lhd$ uniformly at random (among all interval ballots). Then for each candidate, we switch its status (from approved to non approved, or conversely) with probability $p \in [0, 1]$.

- **Random Omissions:** In this model, we first sample for each voter an approval ballot that is an interval of the axis $\lhd$ uniformly at random (among all interval ballots). Then for each *approved* candidate, we switch its status (from approved to non approved) with probability $p \in [0, 1]$.

- **Random Swaps:** In this model, for each voter, we sample an axis $\lhd'$ using the Mallows model with center $\lhd$ and parameter $\phi \in [0, 1]$. As a reminder, the probability of $\lhd'$ in this model is $\frac{\phi^{KT(\lhd,\lhd')}}{C}$ where KT is the Kendall-tau distance and C a constant. Once $\lhd'$ is sampled, we sample uniformly at random an approval ballot which is an interval of $\lhd'$.

For all these models, we take $p < 1/2$. Note that except for FT, we can link each model to an axis rule. However, a deeper analysis shows that only the Maverick Voters model actually corresponds to the model of which VD is the MLE.

For a given model and a given rule, we sample a profile according to the model and we compute the Kendall-tau distance between the axis returned by the rule and the ground truth axis. In case of a tie, we take the average KT over all returned axes. For all our experiments, we set $m = 7$ candidates (for bigger m the computation takes too long), $n = 100$ voters, and we average over 1000 random profiles.

We used the following parameters for the models: $p = 0.2$ for Maverick Voters, $p = 0.3$ for Random Flips, $p = 0.45$ for Random Omissions, and $\phi = 0.5$ for Random Swaps. Figure 2a presents the KT distances. We normalized the distances such that the maximum is 1 for each model and looked at the average over all models. The results are displayed in Figure 2b. The main conclusion seems to be that no rule really generalizes to all models, but VD is particularly bad at generalizing beyond the Maverick Voters model.

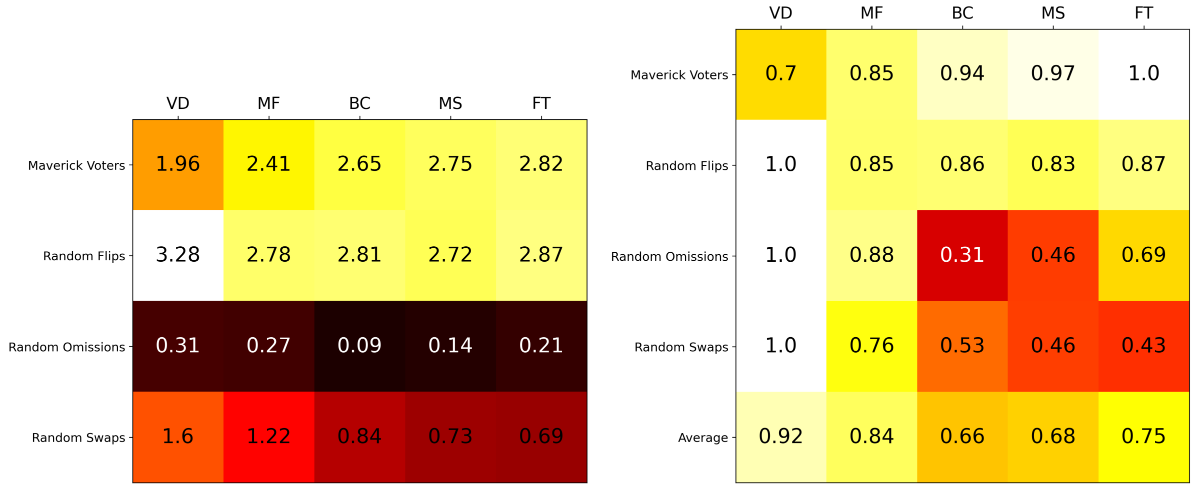

(a) Not normalized. (b) Normalized.

Figure 2: Average Kandall-tau distance to the ground truth axes for different rules and models, averaged over 1000 profiles. The last row of the figure (b) "*Average*" shows the average over all models.

C.3. The French Presidential Election

In this section, we present the results for the French presidential elections datasets. These datasets were gathered in parallel to the actual presidential elections of 2017 and 2022 and part of the *"Voter Autrement"* project.[3] During one month, anyone could answer an online survey, which was promoted on social networks and mailing lists. Participants were asked what would have been their vote for various alternative voting methods, such as approval voting, score voting, Borda, instant runoff voting and the majority judgement rule.

In our experiments, we only need approvals and rankings. For the approval preferences, no preprocessing is required, as we can simply use the approval votes of participants. For the ranking preferences, it is more complicated. Indeed, participants were allowed to rank only a subset of the candidates, for instance their four most favorite ones. However, in our experiments we need full rankings, so we removed all incomplete rankings from the datasets.

We added weights to the voters so that the sample is more representative. Indeed, the set of participants is heavily biased towards the left. Luckily, we know for each participant for which candidate they actually voted at the election (if they agreed to answer this question). Thus, we can adapt the weights of the participants based on this information, so that the distribution of opinions reflects the actual election result. For instance, participants who voted for the main candidate from the left are over-represented, so they get a weight smaller than 1, while participants who voted for the far-right candidates are under-represented, so they get a larger weight. Note that, obviously, this does not completely eliminate the bias.

As a benchmark, we used axes developed by the main polling institutes operating in France. They use these axes (1) when asking the participants which candidate they support and (2) when they present the results. We collected these axes from documents published online by the institutes, see Tables 6 and 8. Note that the axes differ by polling institute. The main differences are (i) the positions of the "small" candidates, as these are hard to place since they often have no obvious classification as left-wing nor right-wing, and (ii) the positions of candidates inside an ideological subgroup (e.g. the far-left candidates or the far-right candidates).

Tables 3 and 7 show the axes returned by each of our rules (including ranking rules), their minimal KT distance to the axes of polling institutes (i.e., the distance to the closest of those

[3]See https://www.gate.cnrs.fr/vote/.

Institute	◁										
BVA	LO	NPA	LFI	PS	EM	LR	DLF	FN	UPR	SP	R
Opinionway	LO	NPA	LFI	PS	EM	LR	DLF	FN	UPR	SP	R
IFOP	LO	NPA	LFI	PS	EM	R	LR	DLF	FN	UPR	SP
IPSOS	LO	NPA	LFI	PS	EM	R	LR	DLF	FN	SP	UPR
Harris Interactive	LO	NPA	LFI	PS	EM	R	LR	DLF	SP	UPR	FN
Odoxa	LO	NPA	LFI	PS	EM	R	LR	DLF	UPR	SP	FN
Elabe	NPA	LO	LFI	PS	EM	R	LR	UPR	DLF	FN	SP

Table 6: Axes used by poll institutes for the 2017 French presidential election

Rule	◁												Min KT	Avg KT
VD	PCF	LO	NPA	LFI	EELV	PS	EM	LR	DLF	REC	RN	R	4	5.62
MF	LO	NPA	LFI	PCF	PS	EELV	EM	LR	R	RN	REC	DLF	4	5.38
BC	LO	NPA	PCF	LFI	EELV	PS	EM	LR	R	RN	REC	DLF	3	5.12
MS	LO	NPA	PCF	LFI	PS	EELV	EM	LR	R	RN	REC	DLF	3	4.88
FT	LO	NPA	LFI	PCF	PS	EELV	EM	LR	R	RN	REC	DLF	4	5.38
VD-rank	DLF	R	PCF	LO	NPA	LFI	EELV	PS	EM	LR	RN	REC	18	20.62
FT-rank	LO	NPA	PCF	LFI	PS	EELV	EM	LR	R	RN	DLF	REC	**2**	**3.88**

Table 7: Optimal axis of each rule for the 2022 French presidential election

Institute	◁											
BVA	LO	NPA	LFI	PCF	PS	EELV	EM	LR	DLF	REC	RN	R
Opinionway	LO	NPA	PCF	LFI	PS	EELV	EM	LR	R	DLF	REC	RN
IFOP	LO	NPA	PCF	LFI	PS	EELV	EM	LR	DLF	RN	REC	R
IPSOS	NPA	LO	LFI	PCF	EELV	PS	EM	LR	R	RN	DLF	REC
Harris Interactive	LO	NPA	PCF	LFI	PS	EELV	EM	LR	DLF	RN	REC	R
Cluster17	LO	NPA	PCF	LFI	EELV	PS	EM	R	LR	DLF	RN	REC
Odoxa	LO	NPA	PCF	LFI	EELV	PS	EM	R	LR	DLF	RN	REC
Elabe	NPA	LO	PCF	LFI	PS	EELV	EM	LR	DLF	RN	REC	R

Table 8: Axes used by poll institutes for the 2022 French presidential election

axes), and their average KT distance to polling institutes. Note that the axes show the parties of the candidates (not the candidate names), and for the colors we followed the choices made by editors of *Wikipedia*.[4]

The axes returned by the different rules are very similar, and they are also close to the axes used by the institutes (except for the VD-rank rule). The differences mainly concern the positions of the less popular candidates (e.g. R) and the positions of the candidates inside each ideological subgroup (e.g. between PS and EELV for the 2022 election).

C.4. Supreme Court of the United States

We derived this dataset from the Supreme Court Database (http://scdb.wustl.edu/), which contains data for Supreme Court decisions starting in 1946. The Court consists of 9 justices who *vote* on each case about which of the two parties to the case wins. The Court then publishes a *majority opinion* explaining the Court's reasoning. Justices can also submit *concurring opinions* and *dissenting opinions*, and *join* any of the opinions submitted by others. Concurring opinions explain additional or alternative reasons, written by justices who voted with the majority. Dissenting opinions explain why a justice did not vote with the majority.

The Martin-Quinn method for deriving an axis of justices uses only the binary vote data (i.e. whether a justice voted for or against the winning party), and its underlying model assumes that a decision divides the axis of justices in the middle, with all justices to one side of the cutoff voting the same way. One issue with this approach is that justices may vote for the same party but have different reasons for it. It could be for example that the most progressive and most conservative justices vote the same way, while the centrist justices vote the other way, for example due to procedural reasons. This is not well-captured by the model. In addition, the model does not use some relevant information. For example, if two justices very frequently join each other in their concurring or dissenting opinions, this suggests that these justices should be placed near each other on the axis.

In our experiments, we discarded all terms with more than 9 justices (e.g. if one is replaced mid-term), giving us 65 terms and thus 65 profiles of approval ballots. We compared our rules to the axes obtained by the established Martin-Quinn method, by computing the KT distance between the axes.

Figures 3 to 8 show the evolution of the positions of the justices on the axes for the last 20 terms, according to the axes produced by the Martin-Quinn method and by our rules. It is very clear that the Martin-Quinn method is smoother over time, which is by the rule's design, since it takes the justice positions of the last term as a prior for their positions in the next term. Our rules are less stable.

[4] https://fr.wikipedia.org/wiki/Modèle:Infobox_Parti_politique_français/couleurs

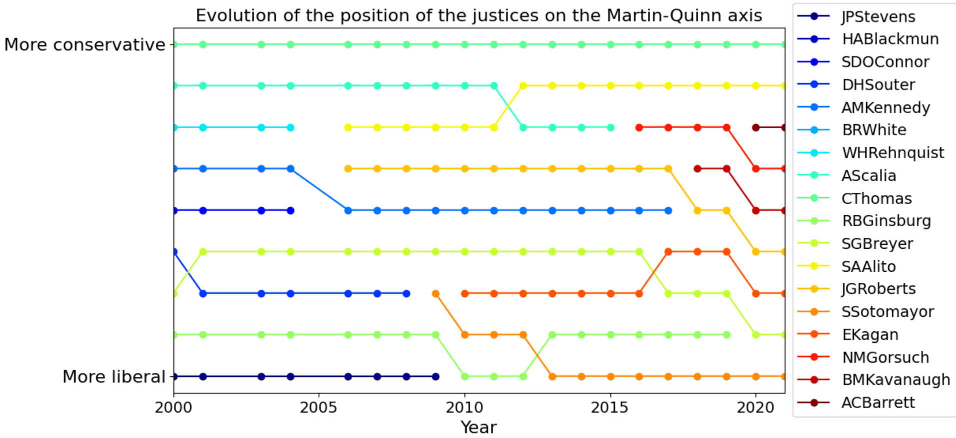

Figure 3: Positions of the justices for terms between 2000 and 2021 for the MQ method.

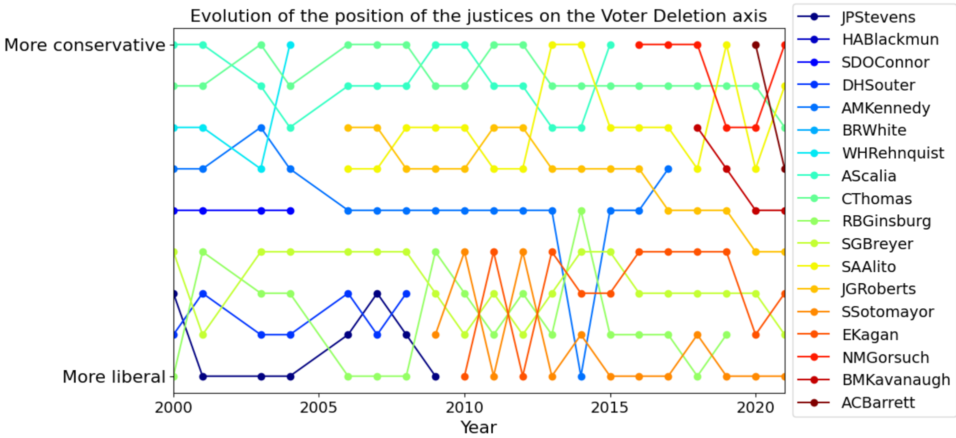

Figure 4: Positions of the justices for terms between 2000 and 2021 for the VD rule.

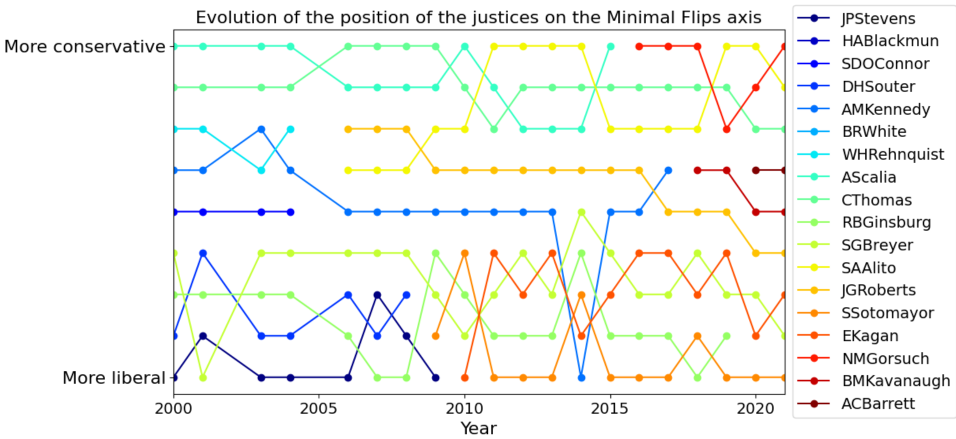

Figure 5: Positions of the justices for terms between 2000 and 2021 for the MF rule.

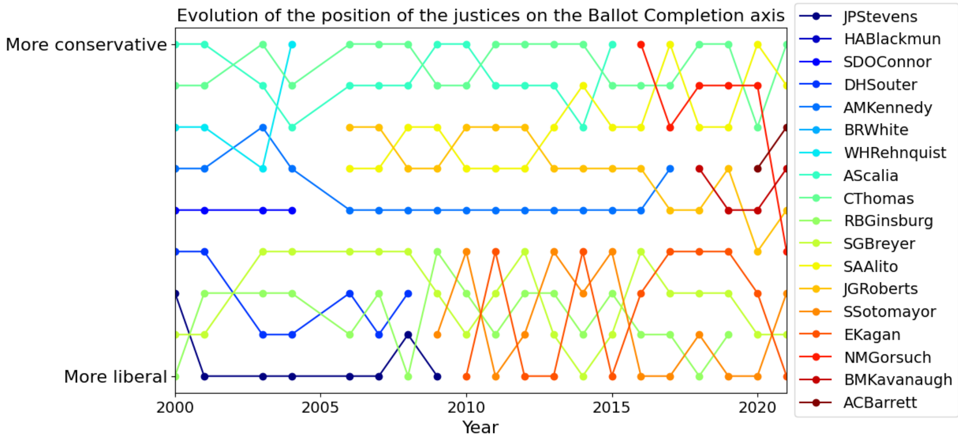

Figure 6: Positions of the justices for terms between 2000 and 2021 for the BC rule.

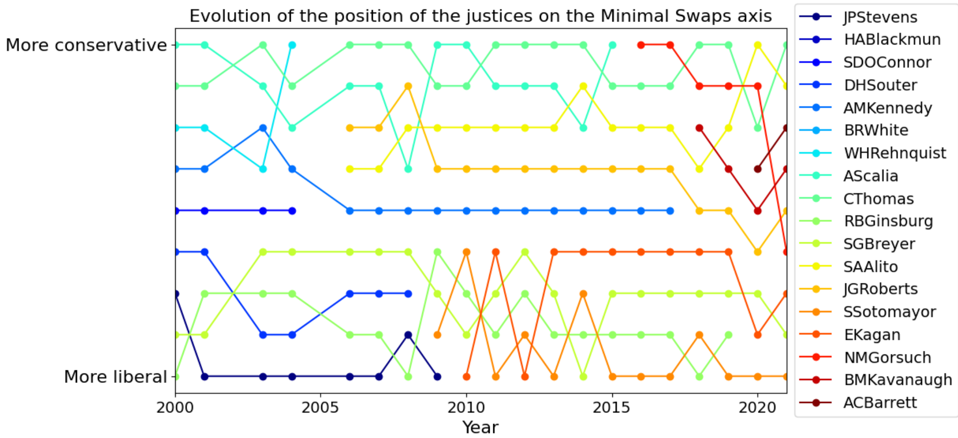

Figure 7: Positions of the justices for terms between 2000 and 2021 for the MS rule.

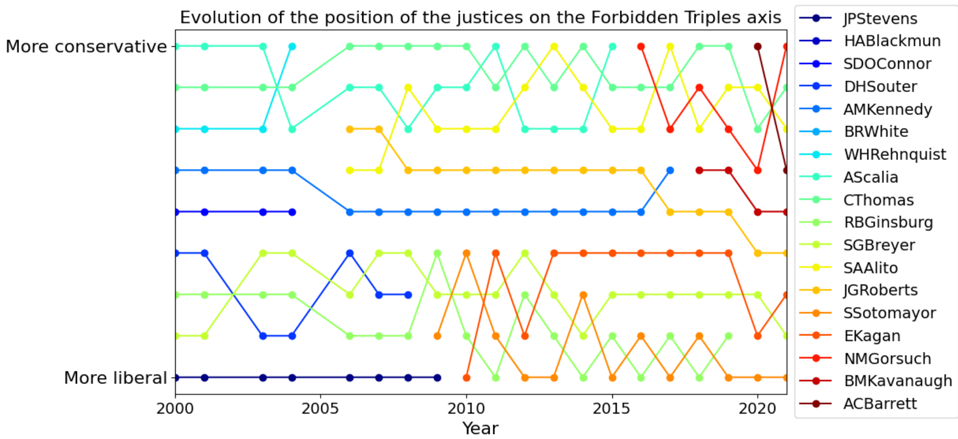

Figure 8: Positions of the justices for terms between 2000 and 2021 for the FT rule.