Approximate Clones in Ordinal Preferences

Théo Delemazure[a]

[a]ILLC, University of Amsterdam, The Netherlands

Abstract.

We study different notions of *approximate clones* in elections with ordinal preferences. Informally, two candidates are approximate clones in a preference profile if they are close to being *perfect* clones, where perfect clones are candidates that are ranked adjacently by all voters. We discuss two measures to quantify this proximity, and we show under which conditions the voting rules that are known to be independent of clones (such as IRV, Ranked Pairs and Schulze) are also independent of approximate clones. In particular, we show that for more than four candidates, none of these rules are independent of approximate clones in the general case. We also study the maximal proximity to being clones that can be expected in preference profiles sampled from classical probabilistic models. Finally, we investigate the presence of approximate clones in real-world datasets, including political data from French presidential election experiments and local Scottish elections.

1 Introduction

A major issue in many electoral systems is what is called the *spoiler effect*: one of the candidates that is losing could have won if another candidate (the spoiler) had not run in the election. A typical example is the single-winner voting rule *plurality* in which every voter votes for one candidate and the candidate with the highest score wins: if several candidates from the same side of the political spectrum are running, voters from this side are *split* between these candidates, which might harm their chances to win. Sometimes, the spoiler effect can be mutual and two candidates could have won individually if the other one had withdrawn. This is what happens for some rules (such as plurality) when two candidates are *clones*, meaning that every voter ranks the two candidates next to each other in their preference order (for instance, if two candidates are from the same party).

To capture this effect, Tideman [31] introduced the axiom of *independence of clones* for elections based on ordinal preferences (i.e., rankings). This axiom says that when two such clones exist, removing either of them should not change the winner of the election (unless the removed clone is the winner, in which case the other clone should become the new winner). This axiom was extensively studied and discussed in the social choice literature since then, particularly for single-winner and multi-winner voting, and has been extended to other preference formats, such as approval preferences. With ordinal preferences, it has been shown that independence of clones is satisfied by a few single-winner voting rules (in particular Instant Runoff Voting, Ranked Pairs and Schulze's method) and multi-winner voting rules (in particular the Single Transferable Vote). Some of these rules (IRV and STV) are used in several countries for large-scale political elections (e.g., in Ireland, Australia, and some US states).

However, *perfect* clones are unlikely to occur in any large-scale political election. Nevertheless, this axiom informs us on the expected behavior of a voting rule in presence of *approximate* clones: if a rule is *independent of clones*, then we can assume that it is less sensitive to the spoiler effect. An interesting question is then: in practice, can we expect to find *approximate* clones? How should the proximity to being clones be defined for a pair of candidates? Moreover, do the rules that satisfy independence of clones also satisfy independence of approximate clones, if these candidates are close enough to being perfect clones?

1.1 Our contributions

We first recall the model of ordinal preferences, the definition of clones and of the independence of clones axiom, and some single-winner voting rules satisfying this axiom in Section 2.

We then introduce the two notions of approximate clones that are used in this paper in Section 3/ These notions aim at quantifying how close a pair of candidates is to being clones. The first one, which we call α-clones, is based on the size of the maximum subset of voters that see the two candidates as clones and rank them adjacently (in this case, α is the proportion of such voters, and is equal to 1 for perfect clones). The second one, that we call β-swap clones, is based on the number of swaps of adjacent candidates in voters' rankings that are required for the two candidates to be clones (in this case, β is the average number of swaps required by voters, and is equal to 0 for perfect clones). In Section 4, we show under which conditions and for which values of α and β the voting rules that are known to be independent of perfect clones are also independent of approximate clones. Notably, in the general case and for more than four candidates, we show that none of these rules satisfy independence of approximate clones for any value of $\alpha < 1$ or $\beta > 0$.

In Section 5, we compute the prominence of approximate clones in profiles sampled from classical probabilistic distributions, with a focus on the impartial culture (IC) and Mallows' model. Then, in Section 6 we investigate datasets of actual preferences and measure the proximity of the pairs of candidates to being clones in these datasets. This analysis is based on datasets from experiments on alternative voting methods conducted during French political elections, datasets of actual votes in local Scottish elections, and rankings of juries in sports competitions. We conclude and discuss potential future work in Section 7.

1.2 Related Work

Perhaps the closest work is that of Janeczko et al. [19], who study among other problems the computational complexity of identifying approximate clones in preference profiles. They also compute the

proximity to being clones of pairs of candidates in synthetic preference profiles (from the map of elections) and in real data, using a notion equivalent to our α-clones measure.

On the weakening of axioms through quantitative measures, Bardal et al. [2] conduct a similar study on proportionality axioms for committee elections. Delemazure et al. [13] introduce measures of how much conflict pairs of candidates induce, which can be seen as an opposite goal of what we are aiming in this work (in particular, the *discrepancy* measure they introduce can be linked to our notion of β-swap clones).

More generally, the independence of clones axiom has been studied in various contexts and has inspired a rich body of research in social choice theory. The notion, originally introduced in the context of single-winner voting with ordinal preferences [31] has later been generalized to weak orders [26, 18] and to approval preferences [9]. Some variants have also been proposed and discussed for the multi-winner setting [34, 3].

2 Preliminaries

Let $V = \{1, \dots, n\}$ be the set of voters and $C = \{c_1, \dots, c_m\}$ be the set of candidates. An ordinal preference profile P is a collection $P = (\succ_1, \dots, \succ_n)$ where $\succ_i$ is a ranking (a strict total order) on C representing the preferences of voter i over the candidates. Given a ranking $\succ_i$, we denote by $\sigma_i(x)$ the position of candidate x in $\succ_i$ (e.g., $\sigma_i(x) = 1$ if x is ranked first). For a preference profile P and a candidate $x \in C$, we denote by P_{-x} the preference profile obtained by removing x from every ranking in P (i.e., x is removed from the election).

A single-winner *voting rule* f is a function that takes as input an ordinal preference profile P and returns a non-empty set of (tied) winners $f(P) \subseteq C$. Note that we consider here *irresolute* voting rules, i.e., rules that can return several winners in case of a tie.

In an ordinal profile P, we say that two candidates are *clones* if they are next to each other in every voter's ranking. More formally, $x, y \in C$ are clones if for every voter i, for every candidate $z \in C \setminus \{x, y\}$, we have $x \succ_i z$ if and only if $y \succ_i z$ (or alternatively, $|\sigma_i(x) - \sigma_i(y)| = 1$ for every voter).

Then, we say that a voting rule f satisfies *independence of clones* if adding or removing a clone does not change (significantly) the result of the election. More formally, we have the following.

Definition 1. *A voting rule f satisfies* independence of clones *if for every preference profile P in which two candidates $x, y \in C$ are clones, we have (1) for every $z \in C \setminus \{x, y\}$, $z \in f(P)$ if and only if $z \in f(P_{-x})$ and (2) $f(P) \cap \{x, y\} \neq \emptyset$ if and only if $y \in f(P_{-x})$.*

This axiom was first introduced by Tideman [31], who showed that the *Instant Runoff Voting (IRV)* rule [27] satisfies this axiom, and introduced another rule called *Ranked Pairs* that also satisfies it. Most of the other well-studied ordinal voting rules do not satisfy independence of clones. A notable exception is *Schulze*'s method [26]. We now define these rules more formally.

Definition 2. *The* Instant Runoff Voting *(IRV) rule proceeds by eliminating candidates sequentially. At each round, the candidate with the fewest first-place votes among non-eliminated candidates is eliminated. The last candidate remaining is the winner.*

The next two rules (Ranked Pairs and Schulze's method) are based on the pairwise majority graph. In other words, their definitions rely on the notion of margin of victory, which is defined for a pair of candidates $\{x, y\} \in C$ as $M_{x,y} = n_{x \succ y} - n_{y \succ x}$ where $n_{x \succ y} = |\{i \in V \mid x \succ_i y\}|$ is the number of voters that prefer x to y.

Definition 3. *The* Ranked Pairs *rule starts by sorting the pairs of candidates $(x, y) \in C^2$ by decreasing order of their margin of victory $M_{x,y}$. Then, the rule constructs a domination graph by adding the pairs in this order. For each pair (x, y), it adds a directed edge from x to y as long as this does not create a cycle (in this case, the pair is skipped). The resulting graph defines a ranking of the candidates. The winner is the candidate at the top of this ranking.*

The next voting rule is based on the notion of *path strength* in the pairwise majority graph. A path is a sequence of candidates $\pi = (c_1, \dots, c_k)$. The *strength* of the path is the minimum margin of victory along the path, i.e., $S_\pi = \min_{i=1}^{k-1} M_{c_i, c_{i+1}}$.

Definition 4. *The* Schulze *rule is defined as follows. For every pair of candidates $x, y \in C$, we define the* beatpath strength $S_{x,y}$ *to be the maximum strength over all paths from x to y. Schulze's method elects the candidate x that beats every other candidate y in beatpath strength, i.e., the candidate x such that for every $y \in C \setminus \{x\}$, we have $S_{x,y} \geq S_{y,x}$. Such a candidate is guaranteed to exist.*

Note that these two rules are *Condorcet methods*, i.e., they elect the Condorcet winner if it exists. As a reminder, a Condorcet winner $x \in C$ is a candidate that beats every other candidate in pairwise comparison, i.e., for every $y \in C \setminus \{x\}$, we have $M_{x,y} > 0$.

In case of ties at any step of these rules, we consider all possible resolutions, and the set of winners is the union of the winners from all possible tie-broken outcomes. Note that this way of breaking ties is known as *parallel-universe tie-breaking* [11, Section 7]. However, in this work we will generally use examples without ties to simplify the presentation.

3 Approximate Clones

We now introduce our notions of approximate clones, which quantify how close to being clones are two candidates in a given preference profile.

The first notion, that we call *α-clones*, is based on the maximum proportion of voters for which two candidates are clones.

Definition 5. *Given a preference profile P, we say that two candidates $x, y \in C$ are* α-clones *if there exists a subprofile P' of P containing at least $\alpha \cdot n$ voters and in which x and y are clones. More formally:*

$$\alpha \leq \frac{1}{n} |\{i \in V \mid |\sigma_i(x) - \sigma_i(y)| = 1\}|.$$

For a pair of candidates, we generally consider the maximum value α for which they are α-clones. This notion is equivalent to the one introduced by Janeczko et al. [19].

One issue with this notion is that it focuses only on the voters for which the two candidates are clones, and ignores the voters for which they are not. In particular, the fact that two candidates are close in some rankings (without necessarily being adjacent) should also be taken into account. For instance, consider a first profile P_1 in which 70% of voters rank x and y next to each other and the remaining 30% rank them respectively first and last in their ranking (so with a rank difference of $m-1$), and a second profile P_2 in which 60% of voters rank x and y next to each other, but the remaining 40% rank them respectively first and third in their ranking (so with a rank difference of only 2). One could argue that x and y are closer to be clones in

the second profile than in the first one, even though they are perfect clones for a larger share of voters in the first profile. This motivates the notion of β-*swap clones*.

Definition 6. *Given an ordinal preference profile P, we say that two candidates $x, y \in C$ are* β-swap clones *if we can perform at most $\beta \cdot n$ swaps of adjacent candidates in voters' rankings to obtain a profile in which x and y are clones. More formally:*

$$\beta \geq \frac{1}{n} \sum_{i \in V} |\sigma_i(x) - \sigma_i(y)| - 1.$$

Again, for a pair of candidates, we generally consider the minimum value β for which they are β-swap clones.

To get a better understanding of these notions of approximate clones, we first consider the characteristic preference profiles that have been studied by Faliszewski et al. [15]: the *identity profile*, the *antagonistic profile* and the *uniform profile*.

In the *identity profile*, all voters have the same ranking $\succ$, so any two candidates are either 1-clones or 0-clones, depending on whether or not they are next to each other in $\succ$. The same is true in the *antagonistic profile*, in which half of the voters have the same ranking $\succ$ and the other half have the opposite ranking $\overleftarrow{\succ}$. In both identity and antagonistic profiles, the value of β for a pair $\{x, y\}$ is $|\sigma(x) - \sigma(y)| - 1$ where σ is the rank function associated with $\succ$.

However, in the *uniform profile* in which every ranking appears exactly once, every pair of candidates is α-clones for the same value of $\alpha = 2/m$. Similarly, every pair of candidates is β-swap clones for $\beta = (m-2)/3$. This can be deduced from the fact that the average values of α and β over all pairs of candidates in any profile are respectively $2/m$ and $(m-2)/3$, and by neutrality of the uniform profile, every pair of candidates has the same value of α and β.

Proposition 1. *In any preference profile, the average value of α for which two candidates are α-clones over all pairs of candidates is equal to $2/m$, and the average value of β for which two candidates are β-swap clones over all pairs of candidates is equal to $(m-2)/3$.*

Proof. Let us first prove that the average value of α for which two candidates are α-clones is $\alpha = 2/m$. Let $k_{a,b}$ be the number of voters that rank a and b adjacently. For one voter, there exist exactly $m-1$ pairs of candidates that are next to each other in their ranking. The sum of the $k_{a,b}$ over all voters is equal to $n \cdot (m-1)$ and the average value of $k_{a,b}$ over all $m(m-1)/2$ pairs is $\frac{n(m-1)}{m(m-1)/2} = \frac{2n}{m}$. Therefore, the average value of $\alpha_{a,b} = k_{a,b}/n$ is $2/m$.

The proof for the average value of β follows the same methodology and can be found in the Appendix.[1] □

Proposition 1 directly implies that in every preference profile, there exists at least one pair of candidates that are α-clones for $\alpha \geq 2/m$, and one pair of candidates that are β-swap clones for $\beta \leq (m-2)/3$. This means for instance that for $m = 4$, there always exist candidates that are clones for at least half of the voters.

4 Independence of Approximate Clones

Motivated by the idea that perfect clones rarely occur in practice, we are now interested in assessing whether voting rules that are independent of perfect clones are also independent of approximate clones.

We say that a rule is independent of α-clones if we can safely remove any of the α-clones without changing the outcome of the election, similarly to the case of perfect clones.

[1] theo.delemazure.fr/more/approximate-clones-appendix.pdf

Definition 7. *We say that a voting rule f is independent of α-clones if for a profile P and two candidates $x, y \in C$ that are α-clones in P, we have (1) for all $z \in C \setminus \{x, y\}$, $z \in f(P)$ if and only if $z \in f(P_{-x})$, and (2) $f(P) \cap \{x, y\} \neq \emptyset$ if and only if $y \in f(P_{-x})$.*

Note that the roles of x and y are symmetric in this definition. When $\alpha = 1$, independence of α-clones corresponds to the classical independence of clones. When $\alpha = 0$, we obtain an axiom much stronger than the well-known (and already very strong) axiom of *independence of losers*, which is the single-winner voting version of the *independence of irrelevant alternatives* axiom [1] and which states that the outcome of the election should not change when we remove any candidate that is not a winner.[2]

Unfortunately, a simple example is enough to show that any reasonable rule fails this axiom for any $\alpha < 1$.

Theorem 2. *Any voting rule that coincides with the majority rule when $m = 2$ fails independence of α-clones for any $\alpha < 1$.*

Proof. Consider the following profile P:

$$k : a \succ b \succ c \qquad k : c \succ b \succ a \qquad 1 : a \succ c \succ b$$

In this profile, b anc c are perfect clones (and thus α-clones) and a and b are $2k/(2k+1)$-clones, which tends to 1 as k tends to infinity. Thus, for any $\alpha < 1$, there exists k such that $2k/(2k+1) > \alpha$.

Let us assume that $a \in f(P)$. By independence of α-clones applied to the pair $\{a, b\}$, $b \in f(P_{-a})$. However, c wins the majority vote against b in P_{-a}, a contradiction. If $b \in f(P)$, by independence of α-clones applied to the pair $\{b, c\}$, $c \in f(P_{-b})$, which again contradicts the majority vote. Finally, if $c \in f(P)$, we use the same argument by reversing the roles of b and c. □

In particular, IRV, Ranked Pairs and Schulze's method fail this property. It is not too surprising to obtain such a negative result: the voters for whom the two candidates are not clones might put one of the clone first and the other last (like in the example of the proof). In this case, if the first clone (the one ranked first) is the winner, it may be too demanding to ask that when we remove it, the other clone (who is ranked last by some voters) should replace it.

4.1 Weak independence of approximate clones

A weaker and perhaps easier to satisfy version of this axiom is to ask that at least one of the two α-clones can be removed without changing the outcome of the election. Here, only one of the clones is treated as the potential spoiler; and its presence or absence should not impact the result.

Definition 8. *We say that a voting rule f is* weakly *independent of α-clones if for a profile P and two candidates $x, y \in C$ that are α-clones in P, the following holds for* either *$P' = P_{-x}$ or $P' = P_{-y}$: (1) for all $z \in C \setminus \{x, y\}$, $z \in f(P)$ if and only if $z \in f(P')$ and (2) $f(P) \cap \{x, y\} \neq \emptyset$ if and only if $f(P') \cap \{x, y\} \neq \emptyset$.*

Clearly, independence of α-clones implies weak independence of α-clones. One could think of intermediate axioms, but we will see that even our weak notion leads to quite negative results, which reduce the need for a strengthening.

Note that weak independence of clones prevents a mutual spoiler effect (i.e., cases in which any of the two clones win when running

[2] In particular, for $\alpha = 0$, if we remove one of the winning candidate, then all other candidates should become winners.

alone, but both clones lose when they both run) but it does not prevent all kinds of spoiler effect. For instance, weak independence allows for a one-way spoiler effect. It could be that candidate x wins when running alone, but both x and y lose when running together. The axiom would still be satisfied as long as y also loses when running alone. Unfortunately, this phenomenon can happen with IRV, Ranked Pairs and Schulze's method for any $\alpha < 1$.

We can define (weak) independence of β-swap clones in a similar way than for α-clones. In the remainder of this section, we will assess whether rules that are known to be independent of perfect clones (and that we introduced in Section 2) are also weakly independent of α-clones or β-swap clones.

We can first show that when $m \geq 4$, they are not weakly independent of α-clones for any $\alpha < 1$.

Theorem 3. *For $m \geq 4$, IRV and Ranked Pairs are not weakly independent of α-clones for any $\alpha < 1$.*

Proof. Let $\alpha < 1$ and $k \in \mathbb{N}$. We start with IRV, and consider the following profile P:

$$\begin{array}{ccc} k : a \succ b \succ c \succ d & k : b \succ a \succ c \succ d & k : d \succ a \succ b \succ c \\ k+1 : c \succ d \succ a \succ b & 2 : a \succ d \succ b \succ c & 2 : b \succ d \succ a \succ c \end{array}$$

In this profile, a and b are $(4k+1)/(4k+5)$-clones, which tends to 1 when k tends to infinity. Thus, there exists k such that $\alpha < (4k+1)/(4k+5)$. With IRV, d is eliminated first, then c, then b, and finally a wins. By independence of α-clones, a should win when we remove b or b when we remove a. However, if we remove a, then c is eliminated first, giving $k+1$ points to d, and b is eliminated next, thus d wins. Similarly, if we remove b, d wins. This contradicts weak independence of α-clones.

For Ranked Pairs, consider the following profile P:

$$\begin{array}{ccc} k+7 : d \succ a \succ b \succ c & k : c \succ a \succ b \succ d & 10 : b \succ c \succ a \succ d \\ 2 : a \succ b \succ d \succ c & 3 : c \succ a \succ d \succ b & 4 : d \succ c \succ a \succ b \end{array}$$

In this profile, a and b are $(2k+7)/(2k+26)$-clones, which tends to 1 when k tends to infinity. Therefore, there exists k such that $\alpha < (2k+7)/(2k+26)$. The margins of victory are as follows:

	a	b	c	d
a	0	$2k+6$	-8	4
b	$-(2k+6)$	0	12	-2
c	8	-12	0	0
d	-4	2	0	0

The pairs are added in the following order: first $a \succ b$, then $b \succ c$. Then, $c \succ a$ is skipped because it would create a cycle. Then $a \succ d$ and $d \succ b$ are added. However, the final ranking is $a \succ d \succ b \succ c$ and the winner is a. Now, if we remove a, the pairs $b \succ c$ and $d \succ b$ are added, and the winner is d. If we remove b, the pairs $a \succ d$ and $c \succ a$ are added and the winner is c. Thus, a clone is winning when both clones are in the election, but not when we remove one, which contradicts weak independence of α-clones.

Using the same profile, we can also show that Schulze's method is not weakly independent of α-clones for $\alpha < 1$.

Indeed, a beats every other candidate in terms of beatpath strength, and wins the election. However, if we remove b, then c beats a and d in terms of beatpath strength, and wins, and if we remove a, d beats b and c in terms of beatpath strength, and wins. Again, this contradicts weak independence of α-clones. □

The same proof can be used to show that these rules are not independent of β-swap clones for any $\beta > 0$. An open question is thus: is there a (sensible) ordinal voting rule that is independent of α-clones for some $\alpha < 1$ when $m \geq 4$? Clearly, the trivial rule that returns all candidates is independent of α-clones for any $\alpha \geq 0$, but it is not an interesting rule.

The situation is different when there are $m = 3$ candidates. In this case, IRV is weakly independent of α-clones for $\alpha \geq 2/3$, and Ranked Pairs and Schulze for any $\alpha \geq 0$.

Theorem 4. *For $m = 3$, IRV is weakly independent of α-clones for $\alpha \geq 2/3$, and not for $\alpha < 2/3$. Ranked Pairs and Schulze are weakly independent of α-clones for any $\alpha \geq 0$.*

Proof. Let us first show that IRV is not weakly independent of α-clones for $\alpha < 2/3$. Consider the following profile with 3 candidates a, b, c:

$$\begin{array}{ccc} 2k : c \succ a \succ b & k+1 : a \succ c \succ b & k+1 : b \succ c \succ a \\ k : a \succ b \succ c & k : b \succ a \succ c & \end{array}$$

In this profile, a and b are $4k/(6k+2)$-clones, which tends to $2/3$ when k tends to infinity. Thus, there exists k such that $4k/(6k+2) > \alpha$ for any $\alpha < 2/3$. With IRV, c is eliminated first, then a wins. However, if we remove either a or b, c wins, which contradicts weak independence of α-clones.

Now, let us prove that IRV is independent of α-clones for $\alpha \geq 2/3$. If in the profile with both clones, one of the clone is eliminated first, then removing it from the election will not change the outcome. Thus, the only elections for which removing either one of the clones can change the result is when the first eliminated candidate is not a clone. In this case, removing either of the clones should make this third alternative win, like in the example above. Let us assume without loss of generality that a and b are α-clones, and that c is the first eliminated candidate. There are only 6 possible rankings of the three candidates. Because c is eliminated first, its plurality score is necessarily lower than $n/3$. Therefore, there are at most $n/3$ rankings of the form $c \succ a \succ b$ or $c \succ b \succ a$. Let us denote $n_1 \leq n/3$ the number of such rankings. Moreover, since a and b are $2/3$-clones, there are at most $n/3$ rankings of the form $a \succ c \succ b$ or $b \succ c \succ a$. Let us denote n_2 the number of such rankings. Combining these two facts, we obtain that there are *at least* $n - n_1 - n_2 \geq n/3$ rankings of the form $a \succ b \succ c$ or $b \succ a \succ c$. Moreover, since there are n_2 rankings of the form $a \succ c \succ b$ or $b \succ c \succ a$, one of the two rankings appears more than $n_2/2$ times by pigeonhole principle. Let us assume without loss of generality that $a \succ b \succ c$ appears more than $n_2/2$ times. Then, if we remove b, the number of rankings of the form $a \succ c$ is greater than $n_2/2 + (n - n_1 - n_2) = (n - n_1 - n_2/2) \geq n/2$, which is enough to make a win. This proves that IRV is independent of α-clones for $\alpha \geq 2/3$.

Let us now consider Ranked Pairs. Clearly, if there are no cycles involving the three candidates, then the outcome will not change when we remove a clone. Let us assume that there is a cycle between the three candidates a, b and c. Without loss of generality, assume that a and b are the α-clones and the cycle is $a \succ b \succ c \succ a$. If c wins in this profile, it is still winning when we remove b, and if either a or b wins, then b is winning when we remove a. In both cases, the axiom is satisfied. Since Ranked Pairs and Schulze coincide when $m = 3$, the above proof also works for Schulze. □

The above proof can also be used to show that IRV is independent of β-swap clones for $\beta \leq 1/3$, but not for $\beta > 1/3$, and that Ranked Pairs and Schulze are independent of β-swap clones for any $\beta \geq 0$.

4.2 Subclasses of profiles

Let us now consider subclasses of profiles for which rules are weakly independent of α-clones. Clearly, any Condorcet method is weakly independent of α-clones for $\alpha \geq 0$ when there exists a Condorcet winner, since at most one of the clone can be the winner, and removing a loser does not change the Condorcet winner. This include single-peaked profiles, which necessarily have a Condorcet winner [4].

We can extend this result using the notion of *Smith set*, which is the set of candidates of minimum size such that every candidate inside the set beats every candidate outside the set in pairwise comparison. More formally, it is the set S of minimum size such that for all $x \in S$ and $y \in C \setminus S$, we have $M_{x,y} > 0$. If the Smith set contains exactly one candidate, then this candidate is the Condorcet winner. A voting rule is said to be *independent of Smith-dominated alternatives* if removing a candidate that is not in the Smith set does not change the outcome of the election. Additionally, a voting rule is said to satisfy the *Smith criterion* if the winner is necessarily from the Smith set. Both axioms are satisfied by Ranked Pairs and Schulze's method, as well as Copeland [10] and Kemeny-Young [22]. We can show the following result.

Theorem 5. *Any voting rule that is independent of Smith-dominated alternatives and satisfies the Smith criterion is weakly independent of α-clones for $\alpha \geq 0$ when the size of the Smith set is at most 3.*

The proof of this result follows the same idea than the proof of Theorem 4 and is provided in the Appendix.

This result is particularly relevant, since it is known that in practice, when the number of voters is large, most preference profiles admit a Condorcet winner, and profiles that do not admit a Condorcet winner generally have a Smith set of size at most 3 [17, 32]. However, this result holds for *weak* independence of α-clones, but not anymore for independence of α-clones (in particular, the profile used in the proof of Theorem 2 contains a Condorcet winner).

5 Probabilistic Models of Preferences

Let us now investigate what values of α can be expected in profiles sampled from classical probabilistic models of preferences, in particular the maximum value over all pairs in the profile. We focus here on the *impartial culture* and the *Mallows model*, which are among the most used models in social choice [7].

5.1 Impartial Culture

In the *impartial culture* model, every ranking is equally likely to be sampled for every voter. More formally, for every ranking $\succ$ and any voter $i \in V$, we have $\mathbb{P}(\succ_i = \succ) = 1/m!$.

In this model, the expected value of α for which two candidates are α-clones is the same for every pair of candidates by neutrality. From Proposition 1, we know that this expected value of α is $2/m$ for any pair of candidates. Thus, as the number of voters n tends to infinity, the expected maximum value of α over all pairs of candidates also tends to $2/m$.

Similarly, the expected value of β for which two candidates are β-swap clones is $(m-2)/3$ for every pair of candidates, and the expected maximum value of β over all pairs of candidates tends to $(m-2)/3$ as the number of voters n tends to infinity.

5.2 Mallows Model

The Mallows model [23] is a very popular probabilistic model for sampling synthetic preference profiles. It takes two parameters: a *central ranking* $\succ^*$ and a *dispersion parameter* $\phi \in [0, 1]$ which controls the concentration of the probability distribution around the central ranking. More precisely, the probability of sampling a ranking $\succ$ is given by $\mathbb{P}(\succ | \succ^*, \phi) = \frac{1}{Z_m} \phi^{d(\succ, \succ^*)}$ where $d(\succ, \succ^*) = |\{a, b \in C \mid a \succ b \text{ and } b \succ^* a\}|$ is the *swap distance* (or number of inversions) between the ranking $\succ$ and the central ranking $\succ^*$, and Z_m is a normalization constant that ensures that the probabilities sum to 1. It is known that $Z_m = \sum_{\succ} \phi^{d(\succ, \succ^*)} = \prod_{i=1}^{m-1}(1 + \phi + \ldots + \phi^i)$ [6]. Note that for $\phi = 0$, the model samples the central ranking with probability 1, and for $\phi = 1$, every ranking is equally likely to be sampled, and we fall back to the impartial culture model.

Let us now compute the expected maximum value of α for which two candidates are α-clones in the Mallows model. Let us assume that the central ranking is $a_1 \succ^* a_2 \succ^* \ldots \succ^* a_m$, and let us compute the probability of sampling a ranking $\succ$ in which the candidates a_1 and a_2 are next to each other. Empirical evidence shows that this pair is the one which maximizes the expected value of α.

Proposition 6. *The probability of sampling a ranking $\succ$ in which the candidates a_1 and a_2 are next to each other in the Mallows model with parameter $\phi < 1$ is equal to $\frac{(1-\phi^{2m-2})(1-\phi)}{(1-\phi^{m-1})(1-\phi^m)}$.*

Proof. We want to compute the probability to sample a ranking $\succ$ in which a_1 and a_2 are adjacent. There are two possibilities for the sampled ranking $\succ$: either $a_1 \succ a_2$ or $a_2 \succ a_1$. Let us start by considering the first case.

In this case, for a ranking $\succ$ in which k candidates are ranked above a_1 and a_2, we have exactly $2k$ inversions involving either a_1 or a_2. Now, the relative order of all other candidates does not matter, so we can consider all possible permutations of the remaining $m-2$ candidates $a_3, \ldots, a_m$. Thus, the probability to sample a ranking $\succ$ in which a_1 and a_2 are next to each other with $a_1 \succ a_2$ and k candidates ranked above them is equal to $\frac{1}{Z_m}\phi^{2k} \cdot Z_{m-2}$, where $Z_{m-2} = \sum_{\succ} \phi^{d(\succ, \succ^*)}$ is the normalization constant for rankings over $m-2$ candidates. If we sum this over all possible values of k, the probability to sample a ranking in which a_1 and a_2 are next to each other with $a_1 \succ a_2$ is equal to $\frac{1}{Z_m} \sum_{k=0}^{m-2} \phi^{2k} \cdot Z_{m-2}$.

In the second case, for a ranking $\succ$ in which k candidates are ranked above a_1 and a_2, we have exactly $2k$ inversions involving either a_1 or a_2 and another candidate, and 1 inversion involving a_1 and a_2. Thus, the probability to sample a ranking $\succ$ in which a_1 and a_2 are next to each other with $a_2 \succ a_1$ and k candidates ranked above them is equal to $\frac{1}{Z_m}\phi^{2k+1} \cdot Z_{m-2}$, and the probability to sample a ranking in which a_1 and a_2 are next to each other with $a_2 \succ a_1$ is equal to $\frac{1}{Z_m} \sum_{k=0}^{m-2} \phi^{2k+1} \cdot Z_{m-2}$.

If we sum the two probabilities, we obtain the total probability of sampling a ranking $\succ$ in which a_1 and a_2 are adjacent:

$$\begin{aligned}
&\frac{1}{Z_m} \sum_{k=0}^{m-2} \phi^{2k} Z_{m-2} + \sum_{k=0}^{m-2} \phi^{2k+1} Z_{m-2} = \frac{Z_{m-2}}{Z_m} \sum_{k=0}^{2m-3} \phi^k \\
=&\frac{1}{(1+\phi+\cdots+\phi^{m-2})(1+\phi+\cdots+\phi^{m-1})} \cdot \frac{1-\phi^{2m-2}}{1-\phi} \\
=&\frac{(1-\phi)^2}{(1-\phi^{m-1})(1-\phi^m)} \cdot \frac{1-\phi^{2m-2}}{1-\phi} \\
=&\frac{(1-\phi^{2m-2})(1-\phi)}{(1-\phi^{m-1})(1-\phi^m)},
\end{aligned}$$

which concludes the proof. □

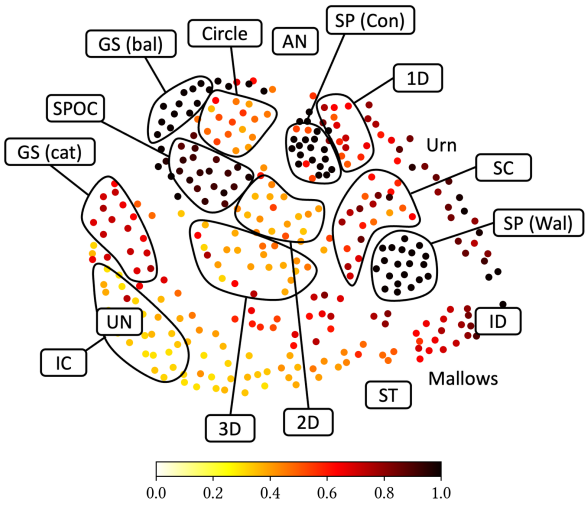


Figure 1: Maximum α value for pairs of candidates in profiles with $m = 10$ candidates and $n = 50$ voters, visualized on the *map of elections* [29].

Note that as m becomes large, the formula tends to $\alpha_{\max}(\phi) = 1 - \phi$. We provide a plot of this formula for several values of m in the Appendix. Our simulations show that this is indeed the maximum value of α for which two candidates are α-clones in the Mallows model. Thus, the expected value of α is highest for the pairs of candidates at the extremity of the ranking (the first two or the last two candidates). This means that if real data are close to preferences sampled from the Mallows model, candidates that are more likely to be approximate clones are either pairs of very popular candidates or pairs of very unpopular candidates.

5.3 The Map of Elections

Finally, we computed the maximum value of α for which two candidates are α-clones in profiles sampled from various models using the framework of the *map of elections* [29, 5]. On a map of elections, each point corresponds to a preference profile (in this case with $m = 10$ candidates and $n = 50$ voters), and the positions of the points are based on some distance between the profiles and some embedding algorithm. Intuitively, points that are close to each other on the map represent preference profiles with similar properties. In Figure 1, we used the swap distance and the embedding provided by the Kamada-Kawai algorithm [20] to obtain the positions of the profiles, and we computed for each profile the maximum value of α for which two candidates are α-clones. Note that this particular experiment has also been conducted by Janeczko et al. [19, Figure 1].

We observe that for the impartial culture (IC), we obtain low values of α, while we obtain high values for profiles close to the identity profile (ID), confirming our previous results. It appears that the Urn model (another popular model in social choice), gives elections with high maximum values of α as well. This is also the case of highly structured preferences (in particular single-peaked and single-crossing preferences, but also 1D-Euclidean preferences). Interestingly, we do not obtain so high values of α for Euclidean preferences in higher dimensions (e.g., 2D, 3D or Circle).

The map of elections for β-swap clones is very similar and can be found in the Appendix.

6 Approximate Clones in Practice

We now study actual datasets of ordinal preferences to see whether approximate clones exist in practice, and if so, for which values of α and β. In the following, we study three datasets of preference profiles. The first one contains data collected during voting experiments that took place during French presidential elections. The second one contains actual votes from local Scottish elections conducted with the Single Transferable Vote (STV) method, which is the equivalent of IRV for committee voting. The third one contains rankings of juries in figure skating competitions.

6.1 French presidential elections

We first consider preference profiles collected during voting experiments that took place during French presidential elections between 2007 and 2022. In these experiments, participants were asked (generally on the day of the election or the weeks before) to vote using alternative voting methods, such as IRV or Borda, thus providing us with ordinal preferences. Note that other voting methods were sometimes tested in these experiments, such as approval voting or range voting. We distinguish two kinds of experiments: *in situ* experiments, which are conducted directly at the polling station on the day of the election, and *online* experiments, which are conducted online and can reach a larger audience.

In this paper, we will use data collected in five such experiments: two *in situ* experiments conducted in 2007 in the city of Faches-Thumesnil [16] and in 2017 in the city of *Fort-de-France* [21].[3], and three *Online* experiments conducted in 2012 ("*Vote Pluriel*", [33]), in 2017 ("*Voter Autrement*", [8]) and in 2022 ("*Un Autre Vote*", [12]). Note that some of these experiments allowed voters to cast *truncated* rankings (i.e., rankings in which some candidates are not ranked), but for our experiments we will only consider the rankings in which all candidates are ranked. These five preference profiles contain between 10 and 12 candidates and between 370 and 5 755 voters (with full rankings). We detail these statistics in the Appendix.

Since the samples of participants of these experiments are not representative of the whole population (especially for the online experiments), we assign *weights* to the voters. To do so, we use the *voting intentions* of the participants, that were collected in the experiments, except in the 2007 one. Then, we compute the weights of the voters such that the weighted proportion of voters supporting a given candidate matches the candidate's actual vote share in the election. For example, if 20% of the participants voted for candidate $a \in C$ in our sample but this candidate actually received 10% of the votes in the first round of the presidential election, then we assign a weight of 0.5 to all voters who voted for a (or intended to vote for them). This does not remove all biases, but it makes the results more robust.

We observe that in these preference profiles, the values of α for which pairs of candidates are α-clones remain relatively low, with a maximum of 0.66 and only 8 pairs obtaining a value of α greater than 0.4. These pairs are listed in Table 1. The pair that seems to consistently have high values of α is the pair of candidates consisting of *Arthaud* and *Poutou*, who both ran in 2012, 2017 and 2022 elections, and have a value of $\alpha > 0.5$ in every dataset.[4] This is not too surprising as these two candidates are both from marginal far-left trotskyist

[3] Note that we did not use the 2017 dataset collected in Grenoble that is often used in social choice (e.g., by Janeczko et al. [19]) because in this experiment the participants did not provide rankings but assigned ratings between 0 and 1 to candidates (by placing a dot on a small line on their ballot), and rankings were derived from these ratings.

[4] Janeczko et al. [19] also find this pair as the closest to being clones for the *Grenoble* dataset.

Dataset	Candidate 1	Candidate 2	α	β
2007-*in situ*	Besancenot	Buffet	0.46	1.54
2017-*in situ*	Cheminade	Lassalle	0.53	1.14
2017-*in situ*	Arthaud	Poutou	0.51	1.34
2017-*in situ*	Macron	Hamon	0.40	1.73
2012-*online*	Arthaud	Poutou	0.66	0.56
2017-*online*	Arthaud	Poutou	0.57	0.87
2017-*online*	Mélenchon	Hamon	0.43	1.45
2022-*online*	Arthaud	Poutou	0.50	0.97
2022-*online*	Zemmour	Le Pen	0.49	1.32

Table 1: Pairs of candidates that are α-clones for $\alpha \geq 0.4$ in the French voting experiments datasets.

parties, and are thus likely to be considered close by voters. More generally, we observe that all pairs of candidates in Table 1 are close on the political spectrum (see for instance [14]). One notable exception is the pair *Cheminade* and *Lassalle* in the 2017-*in situ* dataset, which were both fringe candidates that were very often ranked last by voters, and thus were likely to be adjacent in the rankings.

Generally, we observe a strong correlation between the values of α and β for pairs of candidates in these datasets. We also observe that the plot of the (α, β) values of the pairs of candidates more closely resembles the plots for a 2 or 3-Mallows model (i.e., a Mallows model in which voters are randomly split between $k \in \{2, 3\}$ central rankings) than the ones for a 1-Mallows model. We discuss this further in the Appendix.

The fact that the values of α are relatively low in these datasets is not surprising, as the French presidential election is based on *plurality with runoff*, a two-round system well known to be very sensitive to the spoiler effect and to clones. Thus, similar candidates are incentivized to not run separately in these elections, which might partly explain why we do not observe so many approximate clones.

6.2 *Scottish elections*

We now focus on another dataset, which contains actual votes in local Scottish elections. This dataset was collected by McCune and Graham-Squire [25] and analyzed by Bardal et al. [2, Section 3]. These are multi-winner elections in which the Single Transferable Vote (STV) method is used, which is the equivalent of IRV for committee voting and is known to be independent of clones [30]. Because these are multi-winner elections, in many instances several candidates from the same party are running. There is a total of 1 070 preference profiles, with between 3 and 14 candidates and between 664 and 14 217 voters. Again, voters were not required to cast full rankings and could give truncated rankings. However, we only kept the votes with a full ranking, giving us between 43 and 2 905 voters per profile. The details of these profiles are summarized in the Appendix.

We computed the maximum value of α for which two candidates are α-clones in each profile, and averaged these values over all profiles with the same number of candidates m. The results are shown in Figure 2. There appears to be a large number of approximate clones in these profiles, with an average maximum α around 0.75 for all values of m, and even 95 pairs having an α value greater than 0.9.

However, these high values of α are mostly driven by pairs of candidates from the same party, which are much more likely to be ranked close to each other (e.g., due to party's instructions). For instance, 99.4% of the 664 pairs of candidates that are α-clones for $\alpha \geq 0.8$ are from the same party, and the average values of α for pairs of candidates from the same party are much higher than the global average (see Figure 7 in the Appendix). If we now consider only pairs of candidates from *different* parties, we observe that the average *maximum*

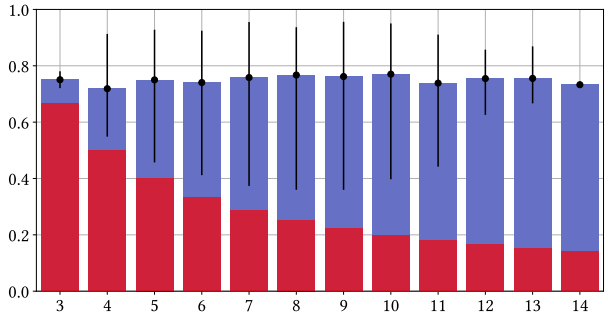


Figure 2: Average maximum value of α for which two candidates are α-clones over all instances in the Scottish elections dataset, depending on the number of candidates m (in blue). The red bars indicate the theoretical minimum $2/m$, and the black lines show the extreme values over all profiles for a given m.

value of α is around 0.55 when $5 \leq m \leq 13$ (see Figure 8 in the Appendix), which is more in line with what we observed for the French datasets. However, we still observe a few (60) pairs of candidates from different parties that have values of α greater than 0.7.

6.3 *Figure skating competitions*

Finally, we study a dataset of rankings of juries in figure skating competitions, collected by Smith [28] (based on juries numerical scores) and available on Preflib [24]. There are a total of 48 profiles, with between 7 and 9 voters (the juries) and between 14 and 30 candidates. Over all profiles, the average maximum value of α for which two candidates are α-clones is 0.91 (and the average minimum value of β is 0.135), and we observe perfect clones in 46% of profiles (22 out of 48). This is partly due to the low number of voters, and to the fact that these profiles are very close to identity [5] since the rankings are mostly based on participants' performances and skills, which can be seen as an objective measure in this kind of competitions.

7 Discussion and Further Remarks

We proposed and analyzed two measures, α-clones and β-swap clones, to identify approximate clones in ordinal preference profiles, and quantify their proximity. We saw that in practice, perfect clones rarely, if ever, exist in large-scale political elections, and even approximate clones are relatively rare, especially if they are from different parties. For votes of juries in competitions however, clones (perfect or approximate) are far more likely. Finally, our results showed that independence of clones results that exist for some voting rules such as IRV or Ranked Pairs generally do not hold when the candidates are not perfect clones. This finding challenges the practical relevance of the standard independence of clones axiom, suggesting that its guarantees do not robustly extend to the more realistic scenario of approximate clones. A key open question is whether any sensible voting rule can satisfy independence of approximate clones for non-trivial thresholds of α or β, or if such robustness is fundamentally unattainable.

Future work could extend these concepts in several directions. First, our notions of pairwise approximate clones could be generalized to sets of three or more candidates. Second, a similar analysis of approximate clones could be conducted for other preference formats, such as approval or cardinal preferences.

References

[1] K. J. Arrow. A difficulty in the concept of social welfare. *Journal of political economy*, 58(4):328–346, 1950. 3

[2] T. Bardal, M. Brill, D. McCune, and J. Peters. Proportional representation in practice: Quantifying proportionality in ordinal elections. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 13581–13588, 2025. 2, 7

[3] D. Baumeister, P. Faliszewski, J. Rothe, and P. Skowron. Multiwinner voting. In *Economics and Computation: An Introduction to Algorithmic Game Theory, Computational Social Choice, and Fair Division*, pages 403–465. Springer, 2024. 2

[4] D. Black. On the rationale of group decision-making. *Journal of political economy*, 56(1):23–34, 1948. 5

[5] N. Boehmer, R. Bredereck, P. Faliszewski, R. Niedermeier, and S. Szufa. Putting a compass on the map of elections. *arXiv preprint arXiv:2105.07815*, 2021. 6, 7

[6] N. Boehmer, P. Faliszewski, and S. Kraiczy. Properties of the mallows model depending on the number of alternatives: A warning for an experimentalist. In *International Conference on Machine Learning*, pages 2689–2711. PMLR, 2023. 5

[7] N. Boehmer, P. Faliszewski, Ł. Janeczko, A. Kaczmarczyk, G. Lisowski, G. Pierczyński, S. Rey, D. Stolicki, S. Szufa, and T. Wąs. Guide to numerical experiments on elections in computational social choice. *arXiv preprint arXiv:2402.11765*, 2024. 5

[8] S. Bouveret, R. Blanch, A. Baujard, F. Durand, H. Igersheim, J. Lang, A. Laruelle, J.-F. Laslier, I. Lebon, and V. Merlin. Voter autrement 2017 - Online experiment, 2018. URL https://doi.org/10.5281/zenodo.1199545. Zenodo, DOI: 10.5281/zenodo.1199545. 6

[9] F. Brandl and D. Peters. Approval voting under dichotomous preferences: A catalogue of characterizations. *Journal of Economic Theory*, 205:105532, 2022. 2

[10] J. S. Coleman. The possibility of a social welfare function. *The American Economic Review*, 56(5):1105–1122, 1966. 5

[11] V. Conitzer, M. Rognlie, and L. Xia. Preference functions that score rankings and maximum likelihood estimation. In *IJCAI*, volume 9, pages 109–115, 2009. 2

[12] T. Delemazure and S. Bouveret. Voter Autrement 2022 - The Online Experiment ("Un Autre Vote"), 2024. URL https://doi.org/10.5281/zenodo.10568799. Zenodo, DOI: 10.5281/zenodo.10568799. 6

[13] T. Delemazure, Ł. Janeczko, A. Kaczmarczyk, and S. Szufa. Selecting the most conflicting pair of candidates. *arXiv preprint arXiv:2405.05870*, 2024. 2

[14] T. Delemazure, C. Dong, D. Peters, and M. Tydrichová. Comparing ways of obtaining candidate orderings from approval ballots. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI)*, 2025. ISBN 978-1-956792-04-1. doi: 10.24963/ijcai.2024/305. URL https://doi.org/10.24963/ijcai.2024/305. 7

[15] P. Faliszewski, A. Kaczmarczyk, K. Sornat, S. Szufa, and T. Wąs. Diversity, agreement, and polarization in elections. *arXiv preprint arXiv:2305.09780*, 2023. 3

[16] E. Farvaque, H. Jayet, and L. Ragot. Quel mode de scrutin pour quel «vainqueur»? une expérience sur le vote préférentiel transférable. *Revue d'économie politique*, 119(2):221–246, 2009. 6

[17] A. Graham-Squire and D. McCune. An examination of ranked-choice voting in the united states, 2004–2022. *Representation*, 61(1):1–19, 2025. 5

[18] W. H. Holliday and E. Pacuit. Split cycle: a new condorcet-consistent voting method independent of clones and immune to spoilers. *Public Choice*, 197(1):1–62, 2023. 2

[19] Ł. Janeczko, J. Lang, G. Lisowski, and S. Szufa. Discovering consistent subelections. *arXiv preprint arXiv:2407.18767*, 2024. 1, 2, 6

[20] T. Kamada, S. Kawai, et al. An algorithm for drawing general undirected graphs. *Information processing letters*, 31(1):7–15, 1989. 6

[21] E. Kamwa, G. Joseph, and A. Roger. Une expérimentation de deux modes de scrutin en martinique lors du premier tour de l'élection présidentielle de 2017. *Revue d'économie politique*, 130(5):759–798, 2020. 6

[22] J. G. Kemeny. Mathematics without numbers. *Daedalus*, 88(4):577–591, 1959. 5

[23] C. L. Mallows. Non-null ranking models. i. *Biometrika*, 44(1/2):114–130, 1957. 5

[24] N. Mattei and T. Walsh. Preflib: A library of preference data. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT 2013)*, 2013. URL http://preflib.org. 7

[25] D. McCune and A. Graham-Squire. Monotonicity anomalies in scottish local government elections. *Social Choice and Welfare*, 63(1):69–101, 2024. 7

[26] M. Schulze. A new monotonic, clone-independent, reversal symmetric, and condorcet-consistent single-winner election method. *Social choice and Welfare*, 36(2):267–303, 2011. 2

[27] J. H. Smith. Aggregation of preferences with variable electorate. *Econometrica: Journal of the Econometric Society*, pages 1027–1041, 1973. 2

[28] W. D. Smith. Analysis of the skate data, 2000. https://rangevoting.org/Skating1998.html. 7

[29] S. Szufa, P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Drawing a map of elections in the space of statistical cultures. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems*, pages 1341–1349, 2020. 6

[30] N. Tideman. The single transferable vote. *Journal of Economic Perspectives*, 9(1):27–38, 1995. 7

[31] T. N. Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4(3):185–206, 1987. 1, 2

[32] T. N. Tideman and F. Plassmann. Modeling the outcomes of vote-casting in actual elections. In *Electoral systems: Paradoxes, assumptions, and procedures*, pages 217–251. Springer, 2011. 5

[33] K. Van der Straeten, J.-F. Laslier, and A. Blais. Vote au pluriel: how people vote when offered to vote under different rules. *PS: political science & politics*, 46(2):324–328, 2013. 6

[34] D. Woodall. Properties of preferential election rules. *Voting Matters*, 3 (Paper 4), 1994. 2