Bad cycles in iterative Approval Voting

Benoît Kloeckner

Bad cycles in iterative Approval Voting

Benoît R. Kloeckner *

January 10, 2020

1 Introduction

This article is about synchronized *iterative voting* in the context of Approval Voting, the voting system in which a ballot can contain the names of any subset of candidates, and the candidate whose name is present in the most ballots is elected.

Assuming that, before the election, successive polls occur to which voters react strategically, we shall exhibit examples showing the possibility of cycles with strong negative properties (in particular, non election of an existing Condorcet winner, or possible election of a candidate strongly rejected by a majority of the electorate). We thus uncover new flaws in the Approval Voting system, which complement in particular the examples provided in [SDL06].

Let us mention some previous works in iterative voting. In the case of Plurality Voting, discussed in slightly more details in Section 6, very general convergence results have been obtained by Meir and co-authors [MLR14, Mei15, MPRJ17]. Many other voting rules –not Approval– have been considered by Lev and co-authors [LR16, KSLR17]. Their theoretical results are negative (no guaranteed convergence) but empirical tests seem to indicate that cycles are rare. Note that in most of these results (the exceptions concerning Plurality), voters are assumed to adjust their ballot one by one. Our setting will be different, *Successive Polling Dynamics* modeling a situation where all voters are informed of the expected result at the same time. Other interesting results can be found in [OLP$^+$, ROL$^+$15].

*Université Paris-Est, Laboratoire d'Analyse et de Matématiques Appliquées (UMR 8050), UPEM, UPEC, CNRS, F-94010, Créteil, France

1.1 On sincerity, strategic voting and straightforwardness

In Approval Voting, one can say that a ballot is *sincere* whenever any candidate preferred to another whose name present on the ballot, must also be on the ballot.

Durand notes in [Dur15] that in general the meaning of "sincerity" is open to interpretation, and that this word has often been used to argue against strategic voting. He makes a compelling point that strategic voting is to be expected, and even advised to voters, and that what causes a democratic problem is the *necessity* of resorting to strategies to get the best ouctome rather than the fact that voters embrace this necessity ("Manipuler c'est bien, la manipulabilité c'est mal"). Insufficient information, too contrived (and as a consequence non-uniformly applied) strategies create asymmetries among voters; and ballots deemed sincere can be cast but afterward regretted in view of the outcome of the election, thus lowering trust and confidence in the democratic system. A most important property of voters' behavior is thus *straightforwardness*, meaning their chosen ballot to depend only on one's preferences, not on the expected outcome of the election.

In approval Voting, as soon as there are more than two candidates there are several sincere ballots, corresponding to the various point in her order of preference where the voter can draw the line between acceptance and rejection. In particular, strategic voting (i.e. choosing one's ballot depending on the ballots expected to be cast by the other voters) can occur even when restricting to sincere votes. In other words, sincerity in the above sense does not imply straightforwardness. Our examples will in particular show how very far from straightforwardness Approval Voting can be in some circumstances.

1.2 Is Approval Voting a Condorcet system "in practice"?

While not a Condorcet system, several arguments have been raised that could seem to indicate it might be close Condorcet "in practice". Brams and Sanver [BS03] showed that when a Condorcet winner exists, her election is a strong Nash equilibrium, i.e no coalition of voters can organize a strategical vote so as to improve the outcome for each and every one of the members of the coalition (note that other outcomes can also be strong Nash equilibrium). Strong Nash equilibrium are said by Brams and Sanver to be "globally stable", but one absent point in their work is whether they are "attractive" equilibrium (in a sense to be made precise below), which has a strong bearing to the question whether they should be expected to be reached in practice.

Laslier [Las09] proved that under a large-electorate model with uncertainty in the recording of votes and perfect common information, the best course of action for voters results in a particular strategy, the "Leader Rule". Additionally, he proved that if there is a Condorcet winner and all voters apply the Leader Rule, then there is at least one equilibrium, and any equilibrium elects the Condorcet winner. Let us give more detail on these results, explained in the framework we shall use here.

To be applied, the Leader Rule needs voters to have a conception about which candidate is likely to win the election (the expected winner), and which candidate is likely to turn second (the expected runner-up). The LR then consists, given the preferences of the voter, in voting for all candidates preferred to the expected winner, to no candidate the expected winner is preferred to, and to vote for the expected winner if and only if she is preferred to the expected runner-up; in particular, strategic voting selects one particular case of the various sincere ways to vote. Laslier shows that under a certain small uncertainty on the recording of votes, this strategy maximizes the odds of improving the outcome of the election.

Assume that a perfectly accurate poll has been conducted (in particular voters answer the poll with the ballot they actually intend to cast); given the collection of initial ballot the voters intend to cast, assuming that after the poll is made public each voter applies the LR to adjust her ballot, we get a new set of intended ballots to be cast. This creates a dynamical system which we call the *successive polling dynamics* (SPD). An equilibrium (or "dynamical equilibrium", distinct from a Nash equilibrium) is then a fixed point of this dynamics, i.e. a state where if all voters would adjust their ballots according the announced results, the adjusted ballots would produce the same result). Laslier second important result can be phrased as follows: whenever there is a Condorcet winner, under SPD there exist at least one dynamical equilibrium and any dynamical equilibrium elects the Condorcet winner. Again, the question of the "attractivity" of the equilibrium is not addressed.

1.3 Description of the main results

The goal of this note is to construct examples showing that in Approval Voting, the SPD can exhibit a problematic cycle even in the presence of a Condorcet winner (the main point shall not be the mere existence of cycles, a rather unsurprizing phenomenon, but rather that such cycles can result in the election of a suboptimal candidate). As noted by Laslier, previous examples of cycles (notably in [BF07]) needed some voters to change their strategy at some iteration of the process; in our example the assumed strategic behaviors is consistent, i.e. constant in time; they are also sincere, and simple.

In one example (Section 3), we assume all voters apply the Leader Rule, thus additionally implying a rather strong form of rationality under Laslier's uncertainty model, and we show that SPD exhibits a cycle where the Condorcet winner cannot be elected. Moreover the *basin of attraction* of the cycle, i.e. the set of initial state leading to the cycle under SPD, is quite larger than the basin of attraction of the Condorcet-winner-electing equilibrium: in practice, it seems quite likely to get caught in a such cycle instead of converging to the stable equilibrium.

In a second example (Section 4), we assume a slight modification of the Leader Rule: some voters have several candidates they decide never to approve; the strategy is thus kept simple, sincere, and consistent. In this example, there is a Condorcet winner and a *worst candidate*, which not only is a Condorcet loser, but in fact has almost two-third of the electorate that would never approve of her. We get two equilibriums, one electing the Condorcet winner the other electing another candidate, and a cycle of order 2 which attracts two-third of the possible states in the SPD and where one of the state elects the worst candidate.

These examples, while specific, are quite stable: the polls creating transition from one expected order of candidates to another are not too close calls, so that their parameters can easily be changed in a range without changing the SPD. In addition, they are stable under perturbation of the model itself: in Section 5 we show the persistence of our examples in a continuous-space model where:

- at each iteration only a given fraction of voters adjust their ballots strategically (e.g. the other not being aware of the poll, or being reluctant to change their minds),
- a fixed strategy is assumed for each voter type only when the candidates scores are sufficiently far apart, without making any assumption on the voters behavior in case of almost equality between any two candidates.

These examples show that Approval Voting exhibits strong issues with respect to polling effects on election outcome (but we do not claim that this flaw does not appear in many other voting systems), in particular mitigating the Condorcet-in-practice stance that could be assumed in view of previous works:

- a slight modification of the Leader Rule can create equilibriums not electing a Condorcet winner, even when one exists,
- Condorcet-electing equilibriums can fail to attract most possible states, i.e. most initial expections on the outcome of the election could lead to cycle not ensuring the election of the Condorcet winner, even when applying the Leader Rule,
- polling can have an extreme impact on the election outcome: rigging *any one* poll can prevent the election of the Condorcet winner even if *all* subsequent polls are perfectly conducted and reported; and in fact, even if all polls are perfect but in the first one voters respond according to a pre-established expectation of the outcome of the election, this expectation can determine the outcome of the election even after arbitrary many polls: polling induces neither synchronization nor loss of memory,
- even with perfect unrigged polls, the sheer *number* of polls (e.g. its parity) can decide the outcome of the election,
- the SPD can get a majority of first poll results to lead after iteration to a cycle, some states of which elect a candidate fully disapproved by a large majority of voters,
- combining the last two items, the parity of the number of polls conducted can lead to elect a candidate fully disapproved by a large majority of voters.

Acknowledgments. I am indebted to Adrien Fauré `@AdrienGazouille` for a long debate on twitter (in French) that lead me to seek and design the examples presented here, and to François Durand for introducing me to the Social choice theory and for many discussions on the Voting Systems. This article benefited from relevant comments provided by Adrien Fauré, Jean-François Laslier and Reshef Meir who I warmly thank.

2 Formalism and notations

Let us set up some notations and formalize the SPD.

2.1 Candidates, electorate, preferences and strategies.

We consider a finite set of *candidates*, named by lower-case letter from the beginning of the alphabet, $\mathscr{C} = \{a, b, c, \dots\}$. An *outcome* is a total order on the set of candidates, written as a word in the candidates names (so that for three candidates, the possible outcomes are $abc, acb, bac, bca, cab, cba$) and represent the order in which the candidate are ranked by an election (or poll); we thus assume there is some tie-breaking rule (e.g. alphabetic order). The set of outcomes is denoted by $\mathscr{O}$. Given an outcome o, the first-ranked candidate is called the elected candidate or the winner ("expected winner" in the case of a poll) and is denoted by $w(o)$; the second-ranked candidate is called the runner-up (expected runner-up) and is denoted by $r(o)$. We denote $\alpha \succ_o \beta$ to say that α is ranked before β in o (so that for example $w(o) \succ_o r(o)$).

The *preferences* of a voter is an order on $\mathscr{C}$, possibly with ties (this will play a role in our second example); we denote $\alpha >_\pi \beta$ to say that α is strictly preferred to β in the preferences π, and $\alpha \geq_\pi \beta$ to say that α is preferred to β or tied with her. A set of preferences can be denoted in the same way than an outcome, with parentheses to group the tied candidates, e.g. $a(bc)d$ means a is the favorite candidate, d is the least preferred, and b and c are tied, both ranked between a and d. The set of possible preferences is denoted by $\mathscr{P}$.

Assuming Approval Voting is used, a *ballot* is an arbitrary subset of $\mathscr{C}$; the set of ballots is denoted by $\mathscr{B}$. A *strategy* is a mapping $f : \mathscr{O} \to \mathscr{B}$; we say that a voter applies strategy f if whenever she expects the outcome $o \in \mathscr{O}$, she casts the ballot $f(o)$ (in either an election or in a poll). For a voter with preferences $\pi \in \mathscr{P}$, a strategy f is said to be *sincere* whenever for all α, β with $\alpha \geq_\pi \beta$ and for all $o \in \mathscr{O}$, $\beta \in f(o) \implies \alpha \in f(o)$. We could have asked instead $\alpha >_\pi \beta$ in this definition, but the above choice makes the sincerity a stronger condition, and thus our results are made stronger.

A particular example of sincere strategy associated with preferences π without ties is the *Leader Rule* f^{LR}_π introduced by Laslier [Las09], defined as follows: for all $o \in \mathscr{O}$ and all $\alpha \in \mathscr{C}$, $\alpha \in f^{\mathrm{LR}}_\pi(o)$ if and only if either $\alpha >_\pi w(o)$ or $\alpha = w(o) >_\pi r(o)$.

We assume a finite set of *voter types*, named by upper-case letters from the end of the alphabet $\mathscr{T} = \{Z, Y, X, \dots\}$. The *electorate* is the data for each voter type Ω of preferences π_Ω, a strategy f_Ω and a number of voters n_Ω.

2.2 Elections and Successive Polling Dynamics.

An *election* (also modelling polls) is a tuple of non-negative integers $(n_\beta)_{\beta \in \mathscr{B}}$ giving the number of each possible ballot cast. The set of all possible elections is denoted by $\mathscr{E}$. The assumption of Approval Voting yields a fixed mapping $\mathrm{AV} : \mathscr{E} \to \mathscr{O}$, which ranks the candidate in decreasing order of the number of ballots their name appear of (ties broken by alphabetic order). The *winner* of an election E is the first ranked in $\mathrm{AV}(E)$, while the *runner-off* is the second ranked.

These data together determine the *Successive Polling dynamics* (SPD), which is the mapping taking an argument an outcome assumed by all voters, and returning the outcome from Approval Voting when each voters casts a ballot following her strategy:

$$\begin{aligned}\mathrm{SPD} : \mathscr{O} &\to \mathscr{O}\\ o &\mapsto \mathrm{AV}\Big(\sum_{\Omega \in \mathscr{T}} n_\Omega \cdot \big(\mathbf{1}(\alpha \in f_\Omega(o)) \big)_{\alpha \in \mathscr{C}} \Big).\end{aligned}$$

where $\mathbf{1}(A)$ is 1 whenever the assertion A is true, and 0 whenever it is false.

This model in particular assumes *consistent strategies* (strategies do not vary in time), and moreover that strategies only depend on the expected outcome of the election; outcomes will therefore also be called *states*. In a generalized framework, states would be the possible arguments of the mapping determining a strategy. SPD models the situation where before the election a certain number k of perfect polls are conducted before the election (perfect meaning that they are answered sincerely, i.e. each voter tells the ballot he or she would cast given the current expected outcome, they are made publicly available, and all voters are polled). Then, if we assume initial intended ballots $b_0(\Omega)$ for each voter type Ω, the first poll result is

$$o_0 = \mathrm{AV}\Big(\sum_{\Omega \in \mathscr{T}} n_\Omega \cdot \big(\mathbf{1}(\alpha \in b_0(\Omega)) \big)_{\alpha \in \mathscr{C}} \Big)$$

and the final outcome of the election is $\mathrm{SPD}^k(o_0)$.

We are thus interested in the dynamical properties of SPD, i.e. of the behavior of its iterates SPD^k and especially in its orbits (i.e. the families $(\mathrm{SPD}^k(o))_{k \in \mathbb{Z}_+}$ where $o \in \mathscr{O}$). A *periodic orbit* (also named *cycle*) is a family of distinct states $o_1, \dots, o_p$ such that $\mathrm{SPD}(o_i) = o_{i+1}$ for all $i \in \{1, \dots, p-1\}$ and $\mathrm{SPD}(o_p) = o_1$, i.e. it. The number p is called the *period* of the periodic orbit, which is also called a *cycle*, or p-cycle to precise the period. A *fixed point* (also named *dynamical equilibirum*) is a 1-cycle, i.e. a state o such that $\mathrm{SPD}(o) = o$. A cycle is said to be *trivial* if all its states have the same winner; otherwise it said to be *non-trivial* (non-triviality implies that the cycle has period at least 2). Given a cycle $o_1, \dots, o_p$, its *basin of attraction* is the set of all states o' such that there exist $k \in \mathbb{Z}_+$ and $i \in \{1, \dots, p\}$ such that $\mathrm{SPD}^k(o') = o_i$. We can represent SPD by the oriented graph with $\mathscr{O}$ as set of vertices, and with exactly one outgoing edge for each $o \in \mathscr{O}$, with endpoint $\mathrm{SPD}(o)$. Cycles are then cycle of the oriented graph, and fixed point are states with a loop.

Remark 2.1. This model is restrictive as it assumes all voters adjust their ballots simultaneously (but see Section 5 that broadens its relevance). If one considers instead arbitrary groups of voters adjusting their ballots in some way (beneficial to them given the current state of affairs, i.e. better-replies), then several definition of acyclicity have been defined (see e.g. [MPRJ17]); in particular *strong acyclicity* means that whatever the order of adjustments is and whatever better-replies are chosen by the voters, the system converges to an equilibrium while *weak acyclicity* means that there exist an order and better-replies leading to an equilibrium. With this synchronized setting, our examples will exclude strong acyclicity but not weak acyclicity.

2.3 Remarkable Candidates and the SPD.

Note that the preferences play no role in SPD; their importance is that they enable to define a sincere strategy (and not all strategies have a preferences π for which they are sincere, so the assumption is a restriction even if the preferences are unrestricted), and they enable to define particular candidates with particular significance with respect to a given electorate.

When a cycle occurs, an important point is whether the elected candidates in the various states of the cycle are the same, and whether they have a particular quality. Given an electorate, a candidate α is said to *dominate* a candidate β (sometimes written $\alpha > \beta$, but beware that this is not a transitive relation) whenever there are more voters that strictly prefer α to β than voters that prefer β to α or are indifferent between the two. Given we allowed for ties in preferences, there are several possible definitions and we choose the stronger one; in particular, it cannot happen that at same time α dominates β and β dominates α.

A candidate α is then said to be a *Condorcet winner* whenever she dominates every other candidate; a Condorcet winner may or may not exist, but if it exist it is unique. A candidate β is said to be a *Condorcet loser* whenever she is dominated by every other candidate; again, a Condorcet loser may or may not exist and is unique if she exist. Similarly, one says that a total order $o \in \mathscr{O}$ is a *Condorcet order* whenever each candidate dominates all candidates ranked below them in o. Last, we will use a stronger notion than Condorcet loser: a candidate is said to be a *worst candidate* whenever there is a strict majority of the electorate that ranks her last (possibly tied with others) in their preferences.

Given a set of candidates and an electorate, there are many questions of interest: are there equilibrium or trivial cycles, and which candidate do they elect? are there non-trivial cycles, and which candidates are elected in the various states of the cycles? when there are several cycles, how big are their respective basin of attraction? which cycle has a particular state, e.g. the state where each voter votes only for her favorite candidate, in its basin of attraction, and which candidates can be elected in this cycle? how do the answers to all these questions depend on the electorate?

3 First example

In this Section, we prove the following.

Theorem A. *Using Approval Voting, there exist an electorate on a set of* 4 *candidates such that:*

- *there is a Condorcet winner,*

- *each voter has preferences without ties and follows as strategy the Leader Rule,*
- SPD *has a cycle, whose basin of attraction contains a majority (actually two-third) of the states, and none of whose states elects the Condorcet winner.*

Proof. We set $\mathscr{C} = \{a, b, c, d\}$ and consider an electorate with 7 types of voters, with the following preferences and numbers (strategies are given by the Leader Rule):

$T : abcd$	100	$U : bacd$	1000	$X : bcad$	1004
		$V : cadb$	1001	$Y : cdab$	1008
		$W : dabc$	1002	$Z : dbac$	1016

The voters in classes U, V, W like a but each prefers one of b, c, d better, while the voters in classes X, Y, Z do not like a too much but distaste one of b, c, d even more, creating a cycle of collective preferences $b > c > d > b$ (with a close to tie with each of b, c, d). Meanwhile, voters in T prefers a to any other player, and their moderate number suffice to make a a Condorcet winner, while maintaining the cyclic preference between b, c, d. The precise numbers of classes U to Z are chosen to exclude any perfect tie (different sums of distinct powers of 2 never agree); this is only for the sake of fanciness, to avoid needing tiebreakers.

Consider the outcome $o_1 = bacd$. Under the Leader Rule, it leads to the following ballots and results:

$T : \{a\}$	$U : \{b\}$	$X : \{b\}$	$a : 3111$	$b : 3020$
	$V : \{c, a, d\}$	$Y : \{c, d, a\}$		$c : 2009$
	$W : \{d, a\}$	$Z : \{d, b\}$		$d : 4027$

so that $o_2 := \mathrm{SPD}(o_1) = dabc$ – i.e. a stays second, while the seemingly unthreatening d gets first position in a cyclic rotation of b, c, d. By symmetry, we can already tell that a cycle will occur, but for the sake of completeness here are the computations. The strategic adjustments triggered by o_2 are as follows:

$T : \{a, b, c\}$	$U : \{b, a, c\}$	$X : \{b, c, a\}$	$a : 3105$	$b : 2104$
	$V : \{c, a\}$	$Y : \{c, d\}$		$c : 4113$
	$W : \{d\}$	$Z : \{d\}$		$d : 3026$

so that $o_3 := \mathrm{SPD}(o_2) = cadb$. The corresponding strategic adjustments are then:

$T : \{ab\}$	$U : \{b, a\}$	$X : \{b, c\}$	$a : 3118$	$b : 4122$
	$V : \{c\}$	$Y : \{c\}$		$c : 3013$
	$W : \{d, a, b\}$	$Z : \{d, b, a\}$		$d : 2018$

so that $\mathrm{SPD}(o_3) = o_1$.

Similar computations gives the graph of SPD, represented in Figure 1. □

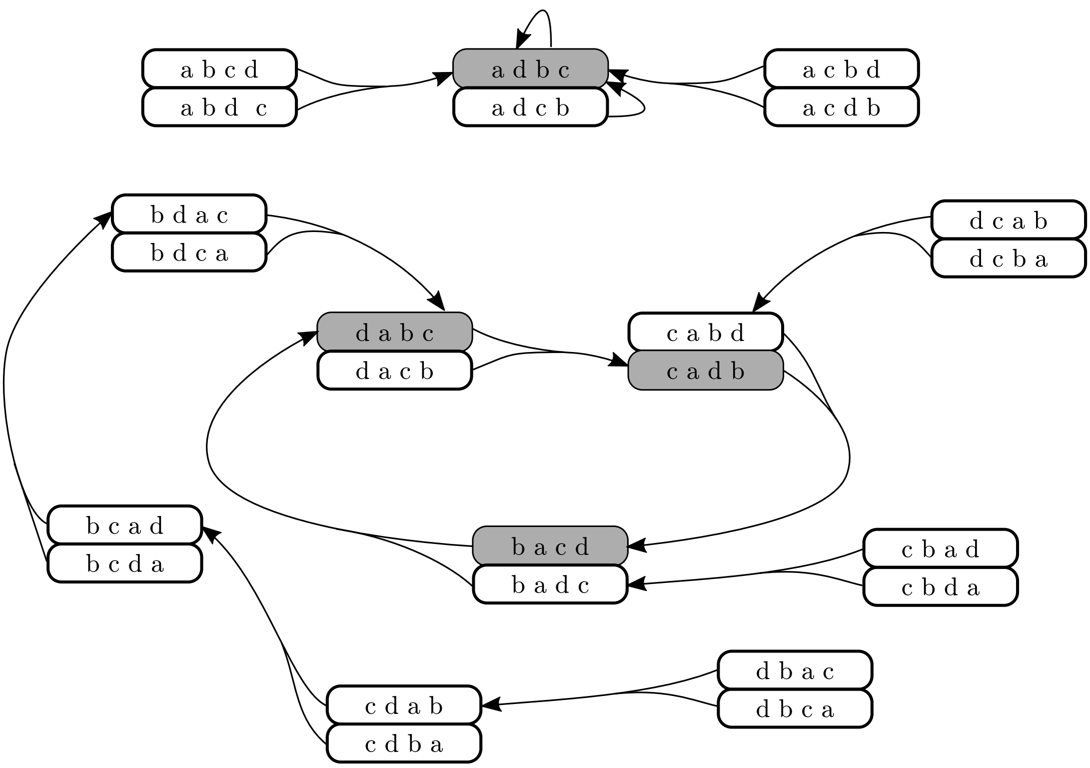

Figure 1: The SPD of the first example. When two states agree on the winner and runner-up, they must have the same image state and are thus represented together. States that can be observed after arbitrary long iterations of the map are represented in gray (one fixed point and a 3-cycle).

Remark 3.1. In the cases of $n > 4$ candidates, we can take the above example and add $n-4$ dummy candidates that appear at the end of all voters preferences. The only property that may not be preserved in this operation is the size of the basin of attraction of the 3-cycle: for example the states starting with one of the dummy variables will all be sent by SPD to a state starting with a, since voters would vote for all of a, b, c, d. However this is easily fixed by adding a voter type, for example in the case of a fifth candidate e one could take $S : becda, 50$. Indeed, this voters will break the tie between a, b, c, d whenever e is the expected winner, in favor of b, thus leading to the basin of attraction of the 3-cycle.

We do not know whether Theorem A holds for 3 candidates; a tedious case by case analysis shows that 3 candidates and 3 voters type cannot lead to both a cycle and a Condorcet winner.

4 Second example

This second example, at the small cost of considering a slightly broader set of strategies, improves on the previous one on two accounts: it necessitates only 3 candidate, and it exhibits a cycle where a worst candidate could get elected.

Theorem B. *Using Approval Voting, there exist an electorate on* 3 *candidates such that:*

- *there are a Condorcet winner and a worst candidate,*
- *each voter follows a consistent, sincere strategy,*
- SPD *has a* 2*-cycle, whose basin of attraction contains a majority (actually two-third) of the states, and one of whose state elects the worst candidate,*
- *there is an equilibrium not electing the Condorcet winner.*

Proof. We consider the following type of voters:

$$Z : abc \quad 101 \qquad Y : a(bc) \quad 2 \qquad X : bac \quad 100 \qquad W : c(ab) \quad 104$$

with as strategy the Leader Rule as it was introduced above (for all $\alpha \in \mathscr{C}$, $\alpha \in f_\pi^{\mathrm{LR}}(o)$ if and only if $\alpha >_\pi w(o)$ or $\alpha = w(o) >_\pi r(o)$); note that with the ties some new situations appear: voters of type W will not choose between a and b, thus always casting the ballot $\{c\}$, no matter which outcome is expected. Similarly, voters of type Y always cast the ballot a (this last type is only introduced here for tie-breaking).

Note that a is a Condorcet winner, beating b with a score of 103 to 100 (voters of type C abstaining) and c with a score of 203 to 104. Moreover c is a Condorcet looser, loosing to b by another landslide 104 to 201: about two-third of the electorate would *never* vote for c, making a worst candidate (by quite a margin).

Assume as starting expected outcome the result obtained if each voters votes for every candidates she does not rank last:

$$Z : \{a, b\} \qquad Y : \{a\} \qquad X : \{b, a\} \qquad W : \{c\}$$
$$a : 203 \qquad b : 201 \qquad c : 104$$

leading to a being expected winner and b expected runner-up (corresponding to the Condorcet order). This leads voters of type Z and X to adjust their votes: their favorite candidate is either threatened by their second-favorite (for Z) or have a shot at winning the election from a current runner-up position (for X). Consistently with their strategies they choose to vote only for their favorite candidate:

$$Z : a \qquad Y : a \qquad X : b \qquad W : c$$
$$a : 103 \qquad b : 100 \qquad c : 104.$$

The second poll thus results in a close-call win of c with a as runner-up. This result induces voters of type Z and X to resume approving both a and b, in order not to let c be elected (again, this is a consistent application of the modified Leader Rule). This results in the the same ballots being cast as in the first poll, so we get a 2-cycle, one of the states electing the worse candidate.

Drawing the full graph of SPD in this case (Figure 2), we see that of 6 states, 4 lead to the cycle that can elect either the Condorcet winner a or the worst candidate c depending on whether the number of polls conducted before the election is odd or even, while the other 2 are stable, one electing the Condorcet winner a the other the Condorcet runner-up b. □

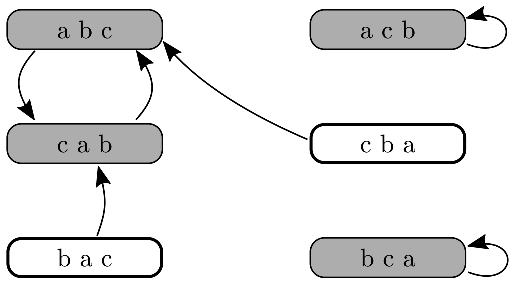

Figure 2: The SPD map of the second example. States that can be observed after arbitrary long iterations of the map are represented in gray (two fixed point and a 2-cycle).

Remark 4.1. We could avoid ties and preserve the features of the examples by splitting W into two types of voters of equal size, with respective preferences *cab* and *cba* and strategy to always vote $\{c\}$ (with the interpretation that these voters prefer c to the other two by far, but still have a slight preference between a and b).

Remark 4.2. It could be argued that in a real-life situation, some of the electors would fail to adjust their votes and the close-call situation where c receive 104 votes and a only 103 might not happen. However, we made these choice of number to make c the worst possible candidate we could; we can preserve the essence of the example by taking $|W|$ anywhere between 103 and 200, trading stability of the 2-cycle against unpopularity of c.

Remark 4.3. Another argument that could be raised against this example is that it needs that a large proportion of voters having c as favorite candidate would never vote for any other candidate. While this is indeed a crucial feature of the voters preferences in this example, there are two counter-arguments. First, this situation seems not all that unlikely: far-right candidates with a strong anti-establishment discourse can have many supporters who would consider all other candidates (or at least those with a chance of being elected) as part of the very same "establishment" and thus would only approve c. Second, this can be a textbook case of manipulation by a coalition: if the minority of all voters who prefer c (with preferences cb or ca say) gather in a coalition and decide to vote only for c, they get a good chance to have c elected against the will of a two-third majority! Actually, these counter-arguments feed on each other: an anti-establishment discourse can serve the purpose of forming a coalition-in-practice of voters who will not express their preferences between a and b in order to favor c.

5 Stability of cycles in a continuous model

The above discrete-time, discrete space model is quite crude and thus can fail to convince that cycle could appear in practice. Discreteness of time is relevant, since polls occur at precise times, triggering some ballot adjustments. Discreteness of space is more of a weakness of the model; more precisely, the assumptions that all voters of a given type adjust their ballots at each poll, and that the strategies depend only on the expected outcome without taking into account almost ties are unrealistic. In this section, we propose a discrete-time, continuous-space setting and show that the cycles in the above examples are stable: they persist even when we assume only a (large enough) fraction of voters to adjust their ballot at each poll, and in the continuous model they are attractive: all states near enough to the cycle are attracted to it. Attractiveness is obviously preserved if we perturb the dynamics far from the cycle, so that the cycles persist even if we only assume the strategies above to be applied away from almost ties.

5.1 A continuous-state setting

We expand the setting of Section 2 in the following way. To each voter type Ω is associated a set $\mathscr{B}_\Omega \subset \mathscr{B}$ of *admissible ballots*, representing the ballots that could be cast by voters of this type. It could be the set of sincere ballots according to their preferences π_Ω, thus assuming mere sincerity; or if we are embedding a discrete example as we will, it could be the image set $f_\Omega(\mathscr{O})$ of the strategy f_Ω of this type.

Given a finite set $\mathscr{X}$, we consider the simplex of vertex set $\mathscr{X}$:

$$\Delta(\mathscr{X}) := \Big\{(p_x)_{x\in\mathscr{X}} \in [0,1]^{\mathscr{X}} \;\Big|\; \sum_{x\in\mathscr{X}} p_x = 1\Big\}.$$

An element of of $\Delta(\mathscr{X})$ can be interpreted as a probability vector on $\mathscr{X}$ or simply as a the proportions of a distribution of some quantity over the elements of $\mathscr{X}$. A state is said to be *extreme* when for each Ω, exactly one of the proportion is 1 and the others are 0; the set of extreme states is finite and can be identified with $\prod_{\Omega\in\mathscr{T}} \mathscr{B}_\Omega$.

A *state* is now an element of

$$\mathscr{S} := \prod_{\Omega\in\mathscr{T}} \Delta(\mathscr{B}_\Omega)$$

i.e. a state $s \in \mathscr{S}$ gives for each voter type and each admissible ballot the proportion of the voters of that type planing to cast this ballot. The use of a continuous model assumes voters are sufficiently numerous that we can consider each type arbitrarily finely divisible; the voters' counts n_Ω now only represent their respective proportion of all voters, rather than absolute numbers. There is a natural mapping which associates to each state the election that would result from the casting of the planned ballots:

$$\begin{array}{rccc} v: & \mathscr{S} & \to & \mathscr{E} \\ & s = \Big((p_b^\Omega)_{\beta\in\mathscr{B}_\Omega}\Big)_{\Omega\in\mathscr{T}} & \mapsto & \Big(\sum_\Omega n_\Omega \cdot p_\beta^\Omega\Big)_{\beta\in\mathscr{B}} \end{array}$$

where p_β^Ω is considered zero whenever $\beta \notin \mathscr{B}_\Omega$. The mapping $\mathrm{AV}\circ v : \mathscr{S} \to \mathscr{O}$ thus sends a state to the outcome that would result from it (recall that we assumed a tie-breaking rule, e.g. by alphabetical order).

A *general poll dynamics* (GPD) is a map $\Phi : \mathscr{S} \to \mathscr{S}$; such a generality is meant to allow more modeling possibilities, but the interesting GPDs are those that are grounded in a natural way in the other elements of the model.

Given an electorate, we choose as suggested above $\mathscr{B}_\Omega = f_\Omega(\mathscr{O})$ and consider the particular GPD which corresponds to all voters to apply the strategy of her type given the expected outcome, i.e.

$$\Phi_0(s) = \Big(\mathbf{1}(f_\Omega \circ \mathrm{AV} \circ v(s) = \beta)_{\beta\in\mathscr{B}_\Omega}\Big)_{\Omega\in\mathscr{T}}.$$

The mapping Φ_0 takes its values in the set of extreme states, and after the first iteration does not convey any more information than the mapping SPD given by the electorate.

We consider the supremum norm $|\cdot|$ on $\mathscr{S}$, i.e. given two states $s = (p_\beta^\Omega)$ and $\bar{s} = (\bar{p}_\beta^\Omega)$ we set

$$|s - \bar{s}| = \sup_{\Omega\in\mathscr{T},\beta\in\mathscr{B}_\Omega} \Big|p_\beta^\Omega - \bar{p}_\beta^\Omega\Big|$$

Given two GPDs Φ, Ψ and a set of states $A \in \mathscr{S}$, we consider

$$D_A(\Phi,\Psi) := \sup_{s\in A} \Big|\Phi(s) - \Psi(s)\Big|$$

and use D as a shortcut for $D_{\mathscr{S}}$, which is a metric on the set of all GPDs. More generally, $D_A(\Phi,\Psi)$ quantifies how close Φ and Ψ are on A; a smaller A makes upper bounds on D_A more lenient.

5.2 A continuous-state example with a cycle

Using a slight modification of our second example above, we get the following result where choices have been made to make all constants explicit (but they are not optimal).

Theorem C. *There exist an electorate on three candidates such that:*

- *there are a Condorcet winner and a worst candidate,*
- *each voter type is assigned a consistent, sincere strategy,*
- *for any GPD Φ where 85% of the voters of each type adjust their ballot according to their type's strategy whenever the expected election gives an outcome with at least 4 percentage points of margin between candidates (the reminding 15% keep their previous ballot and the GPD is arbitrary when margins are lower than 4%), Φ has an attractive 2-cycle one of whose states elects the worst candidate.*

Proof. We consider the following four-types electorate

$$Z : abc \quad 3 \qquad\qquad Y : a(bc) \quad 1 \qquad\qquad X : bac \quad 3 \qquad\qquad W : c(ab) \quad 5$$

with the modified Leader Rule as in Section 4 and the corresponding sets of admissible ballots. Again, a is a Condorcet winner and c a worst candidate (now refused by 7/12th of the electorate: we traded some badness of c for more stability). Observe that

- $\Delta(\mathscr{B}_Y) = \Delta(\big\{\{a\}\big\})$ and $\Delta(\mathscr{B}_W)$ are both singletons and can be ignored in the product,
- $\Delta(\mathscr{B}_X) = \Delta(\big\{\{b\}, \{a,b\}\big\})$ can be identified with $[0,1]$, denoting by x the proportion of voters of type X that cast the ballot $\{a,b\}$,
- $\Delta(\mathscr{B}_Z) = \Delta(\big\{\{a\}, \{a,b\}\big\})$ can be identified with $[0,1]$ by representing a probability vector by the proportion z of voters of type Z that cast the ballot $\{a,b\}$.

In this way, we can identify $\mathscr{S}$ and $[0,1]^2$ with coordinates (x,z).

The three lines corresponding to ties (of equation $(z = x + \frac{1}{3})$ for a and b; $(x = \frac{1}{3})$ for a and c; $(z = \frac{2}{3})$ for b and c) are concurrent at the point where all three candidates are tied, and define six regions in $\mathscr{S}$ corresponding to the six possible outcomes (the lines themselves are attributed to one outcome according to the tie-breaking rule), as show in Figure 3. The region A_1 delimited by the lines of equations $(z < x + \frac{1}{6})$ and $(z > \frac{5}{6})$ results in the outcome abc with margins of $\frac{1}{24}$th of the electorate, i.e. slightly over 4%. Similarly, the region A_2 delimited by the lines of equations $(z < x + \frac{1}{6})$ and $(x < \frac{1}{6})$ result in the outcome cab with the same margins. Assuming Φ is a GPD with the property assumed in the third item, we have $\Phi(x,z) = (0.15x, 0.15z)$ whenever $(x,z) \in R_1$ and $\Phi(x,z) = (0.85{+}0.15x, 0.85{+}0.15z)$ whenever $(x,z) \in R_2$. In particular, $\Phi(R_1) \subset R_2$ and $\Phi(R_2) \subset R_1$. It follows that $\Phi^2(R_1) \subset R_1$, and since Φ^2 is a contraction on R_1 (of ratio 0.15^2) it must have a fixed point $(x_1, z_1) \in R_1$. Then the orbit of

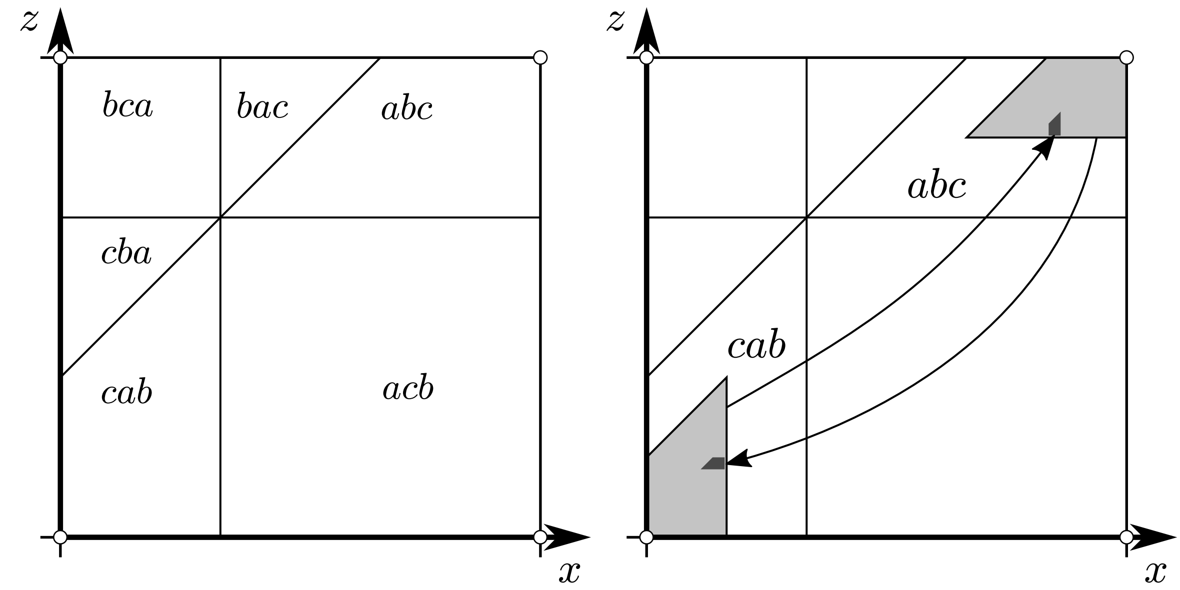

Figure 3: Left – a continuous-space example: the corners are the extreme states, corresponding to the four outcomes that are possible given the assigned strategies; if all voters adjust their ballot after a poll (GPD Φ_0), the whole *abc* and *bac* regions are sent to the bottom-left corner, the *cab* and *cba* regions to the upper-right corner, the *bca* region to the upper-right corner and the *acb* region to the lower-right corner. Right – a GPD where 85% of voters adjust their ballot when candidates are separated by 4% margins will send the light-grey regions into each other, thus ensuring a 2-cycle with one state near the upper-rigth corner, one near the lower-left corner.

(x_1, z_1) is a 2-cycle with one state inducing the outcome *abc* and the other inducing *cab*. Moreover any element of $R_1 \cup R_2$ is attracted to this cycle exponentially fast: e.g. if $(x, z) \in R_1$, for all $n \in \mathbb{N}$ we have $|\Phi^{2n}(x, z) - (x_1, z_1)| \leq 0.15^{2n}|(x, z) - (x_1, z_1)|$ and $|\Phi^{2n+1}(x, z) - \Phi(x_1, z_1)| \leq 0.15^{2n+1}|(x, z) - (x_1, z_1)|$. □

5.3 A general stability result

The above example was meant to be tangible and explicit, but the underlying phenomenon is quite general. The following result is easily proven with the same ideas.

Theorem D. *Consider an electorate on a given set $\mathscr{C}$ of candidates, such that the SPD has a cycle $o_1, \dots, o_k$. We denote by $B_i = (f_\Omega(o_{i-1}))_{\Omega \in \mathscr{T}}$ the vector of the ballots cast by the different types along this cycle.*

Consider first the GPD Φ_0, i.e. assume temporarily that all voters apply their type's strategy at each poll, and let ε_0 be the largest number such that in every election following from any B_i, the scores of candidates are separated by at least a fraction ε_0 of the electorate.

We assume $\varepsilon_0 > 0$ (i.e. the tie-breaking rule is not needed in the considered cycle), and for all $\varepsilon \in (0, \varepsilon_0)$ we denote by A_i^ε the largest region of $\mathscr{S}$ corresponding to B_i where all candidates are separated by at least a fraction ε of the electorate.

Then for all $\varepsilon \in (0, \varepsilon_0)$ there is a $\delta > 0$ with the following property: every GPD Φ that are continuous on each A_i^ε and such that $D_{A_i^\varepsilon}(\Phi, \Phi_0) < \delta$ has a k-cycle whose elements are in the A_i^ε, in particular with corresponding outcomes $o_1, \dots, o_k$.

This result means that any cycle of the discrete-space model that does not rely on the tie-breaking rule is stable: any GPD that is close enough to Φ_0 near the cycle exhibit a similar cycle. In particular, if we only assume at least a given fraction of each voter's type adjust their ballot at each polls, and if we only assume the prescribe strategies far away from ties, we still have a cycle.

Proof. By definition of ε_0, each A_i^ε is the intersection of $[0,1]^{\mathscr{T}}$ with a polyhedron containing in its interior the extreme point corresponding to B_i, and Φ_0 maps it entirely to the extreme point B_{i+1}. In particular, Φ_0^k sends the whole of A_1^ε to the extreme point B_1. If δ is small enough and $D_{A_i^\varepsilon}(\Phi, \Phi_0) < \delta$ for all i, then Φ^k is close enough to Φ_0^k to map A_1^ε into itself. By assumption, Φ^k is continuous on A_1^ε which is homeomorphic to a closed ball. By Brouwer's fixed point theorem, Φ^k has a fixed point in A_1^ε, leading to the desired cycle of Φ. □

Observe that without the continuity hypothesis, we would still get a cycle of outcomes $o_1, \dots, o_k$, but possibly not a cycle of states nor of elections. If Φ is contracting on each A_i^ε, the cycle of states is attracting.

6 Comparison with Plurality Voting

A reasonable question is to ask whether the flaws unveiled by the above examples are avoided in other voting systems. Considering Plurality Voting, the situation is complicated by the rigidity of the single-name ballot, which forces voters to choose a trade-off between preferences and probability to improve the outcome of the election. The works [MLR14] and [Mei15] have studied in depth models taking into account the scores of the candidates and a level of uncertainty to define the possible voters' strategies. They obtained several results proving under some assumptions convergence to equilibrium (the result closest to our present setting is Theorem 5 in [Mei15], where at each iteration an arbitrary subset of voters adjust their votes according to the current poll results, thus including the case studied here where all voters adjust their votes at each iteration). Presence of cycles in a SPD is thus not a universal feature (or rather bug) of voting systems.

Note that if we tried to design a very simple strategical model inspired by Leader Rule for Plurality voting, we could consider the case when every voter votes for either the expected winner or the expected runner-up, whoever comes first in her order of preference (the rationale is that when the electorate is large, a vote to any other candidate is orders of magnitude less likely to change the outcome of the election). If we assume that voteres apply this "Plurality Leader Rule", then of course the first poll is decisive in the polling dynamics and the outcome of the election, since the expected winner and expected runner-up are the only ones staying in competitions. Then there is always convergence to an equilibrium, and a very fast one at that (the polls are constant after the second iteration!) However there are very many equilibria: every candidate that is not a Condorcet looser could be elected, depending on the results of the first poll. In a sense, this strategy reduces the SPD to a two-rounds voting system where the first two candidates in the first round make it to the second one, which is decided by majority.[1] At least a Condorcet loser (in particular, a worst candidate) cannot be elected, since if she makes it to the second poll she is defeated in each poll after the first. Note that this reasoning shows that any system with an ultimate two-candidates round (e.g. IRV) will avoid electing a candidate ranked last by a majority of voters, including after successive polling; this is a very weak but positive feature that Approval Voting lacks.

7 Conclusion

We have considered the Successive Polling Dynamics under Approval Voting, assuming voters apply *simple, consistent, sincere* strategies. This dynamics can be thought of as a model of strategic voting, when voters try to anticipate the outcome of the election to decide their votes, while not being able to form a large-scale coalition.

In this context, we exhibited two examples showing that Approval Voting is far less of a Condorcet-in-practice voting system than could have been expected; specifically, we showed that:

i. assuming voters apply Laslier's Leader Rule, successive polls can lead from a majority of initial expected outcomes to a cycle failing to elect an existing Condorcet winner,

ii. letting voters have only slightly more general strategies, successive polls can lead to the election of a worst candidate.

We proved that these cycles are stable under natural perturbations: even assuming only a proportion of voters adjust their ballots at each iteration, and assuming different strategies in case of almost ties, similar cycles persist. This shows that under Approval Voting, not only convergence to equilibrium may not happen, but cycles can turn individually sound strategies into sub-optimal or even worst possible outcomes.

Last we made a brief comparison with Plurality Voting. Previous works (in particular [Las09]) showed that Approval Voting has quite better equilibrium properties than Plurality Voting (e.g. when there is a Condorcet winner, assuming voters apply the simple and rational Leader Rule, there exist an equilibrium and any equilibrium elects the Condorcet winner); what we showed is that Approval Voting, in counterpart, lacks general convergence to equilibrium, making its equilibrium qualities far less relevant in practice.

[1] In some sense, multiple-round voting systems could be thought of as a way to counter the reluctance of some voters to vote strategically.

References

[BF07] S. Brams & P. C. Fishburn – *Approval voting*, Springer Science & Business Media, 2007. 1.3

[BS03] S. J. Brams & M. R. Sanver – "Voter sovereignty and election outcomes", 2003. 1.2

[Dur15] F. Durand – "Vers des modes de scrutin moins manipulable", Thèse, Université Pierre et Marie Curie, 2015. 1.1

[KSLR17] A. Koolyk, T. Strangway, O. Lev & J. S. Rosenschein – "Convergence and quality of iterative voting under non-scoring rules", 2017. 1

[Las09] J.-F. Laslier – "The Leader Rule: a model of strategic approval voting in a large electorate", *Journal of Theoretical Politics* **21** (2009), no. 1, p. 113–136. 1.2, 2.1, 7

[LR16] O. Lev & J. S. Rosenschein – "Convergence of iterative scoring rules", *J. Artif. Intell. Res.* **57** (2016), p. 573–591. 1

[Mei15] R. Meir – "Plurality voting under uncertainty", in *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI'15, AAAI Press, 2015, p. 2103–2109. 1, 6

[MLR14] R. Meir, O. Lev & J. S. Rosenschein – "A local-dominance theory of voting equilibria", in *Proceedings of the Fifteenth ACM Conference on Economics and Computation* (New York, NY, USA), EC '14, ACM, 2014, p. 313–330. 1, 6

[MPRJ17] R. Meir, M. Polukarov, J. S. Rosenschein & N. R. Jennings – "Iterative voting and acyclic games", *Artificial Intelligence* **252** (2017), p. 100–122. 1, 2.1

[OLP+] S. Obraztsova, O. Lev, M. Polukarov, Z. Rabinovich & J. S. Rosenschein – "Farsighted voting dynamics". 1

[ROL+15] Z. Rabinovich, S. Obraztsova, O. Lev, E. Markakis & J. S. Rosenschein – "Analysis of equilibria in iterative voting schemes", in *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015. 1

[SDL06] F. D. Sinopoli, B. Dutta & J.-F. Laslier – "Approval voting: three examples", *Int. J. Game Theory* **35** (2006), p. 27–38. 1