



HAL
open science

Bad cycles and chaos in iterative Approval Voting

Benoît Kloeckner

► **To cite this version:**

| Benoît Kloeckner. Bad cycles and chaos in iterative Approval Voting. 2020. hal-02434670v2

HAL Id: hal-02434670

<https://hal.science/hal-02434670v2>

Preprint submitted on 19 Nov 2020 (v2), last revised 9 Feb 2022 (v4)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Bad cycles and chaos in iterative Approval Voting

Benoît R. Kloeckner ^{*†‡}

November 19, 2020

We consider synchronized iterative voting in the Approval Voting system. We give examples with a Condorcet winner where voters apply simple, sincere, consistent strategies but where cycles appear that can prevent the election of the Condorcet winner, or that can even lead to the election of a “consensual loser”, rejected in all circumstances by a majority of voters.

We conduct numerical experiments to determine how rare such cycles are. It turns out that when voters apply Laslier’s Leader Rule they are quite uncommon, and we prove that they cannot happen when voters’ preferences are modeled by a one-dimensional culture. However a slight variation of the Leader Rule accounting for possible draws in voter’s preferences witnesses much more bad cycle, *especially* in a one-dimensional culture.

Then we introduce a continuous-space model in which we show that these cycles are stable under perturbation. Last, we consider models of voters behavior featuring a competition between strategic behavior and reluctance to vote for candidates that are ranked low in their preferences. We show that in some cases, this leads to chaotic behavior, with fractal attractors and positive entropy.

1 Introduction

1.1 Iterative Voting

In any voting system, the choice of ballot by any voter is guided by her preferences between the candidates and the expected effect of each ballot she can cast. To estimate

*Univ Paris Est Creteil, CNRS, LAMA, F-94010 Creteil, France

†Univ Gustave Eiffel, LAMA, F-77447 Marne-la-Vallée, France

‡ORCID 0000-0002-4966-7864

the effect of any ballot, one needs to know the ballots cast by other voters, or at least to have some information on them. Assuming perfect information is neither realistic nor theoretically useful: in most cases, changing one’s ballot will not change the outcome of the election (hence, as is well-known, the notion of Nash equilibrium is ineffectively broad in the context of elections [Las09]). One thus often considers situations where information is imperfect, so that the ballot to be chosen by any given voter might change the outcome, with a tiny but non-zero probability. Another point of view is to consider the possibility for voters of aligned interests to form coalitions and decide together which ballots to cast. A third point of view, taking from both of the first two, is to consider that some (or all) information about voters’ intents is common knowledge, and that voters choose their ballot under the assumption that other voters that have the same preferences as theirs will make the same reasoning as they will – thus overcoming the insufficient weight of one ballot to change the outcome.

In any case, these considerations introduce a *dynamical* component to voting: after all voters adjusted their intentions, the information under which this adjustment was made is outdated; the new intents result in new information that could be shared, and itself result in new strategic adjustments of voters’ intents, etc. Natural questions are thus whether there is some dynamical equilibrium, i.e. intents of all voters that are preserved when all shared publicly; whether such equilibrium is unique; what properties equilibriums have, in particular what candidate can be elected at equilibrium; and whether any other set of intent converges to an equilibrium after successive adjustments. The field of *iterative voting* is the study of these question, under various modeling assumptions. One shall distinguish between *synchronized* and *asynchronous* iterative voting. In synchronized iterative voting, all voters are given the information at time t , and simultaneously adjust their intents to build the new situation at time $t + 1$.¹ In asynchronous iterative voting, some voters adjust their intents given the available information, and information is updated before some other voters make their adjustments, etc.

Setting and motivation. 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.

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). These new issues with the Approval Voting system complement in particular the examples provided in [SDL06]. We also introduce a continuous-space model in which finer voter behavior can be encoded; the most striking outcome is that simple behaviors restricted to sincere ballots can in some situation result in chaotic dynamics.

Note that Approval Voting is especially interesting in this context because several

¹Note that what is important is not that the voters’ intents change exactly at the same time, but that all voters have the opportunity to adjust their preferences before information is updated.

sincere ballot coexist for each order of preference, so that there exist strategies that are both sincere, simple, and non-trivial (we discuss this in more detail below); but it is likely that bad cycles and chaotic dynamics exist in other voting systems.

Considering synchronized iterative voting is interesting from a theoretical perspective, when trying to define a “best” ballot for each voter: to be efficient in her choice, a voter should not only vote strategically, but also anticipate that other voters will also do so, who will themselves try to anticipate all voters’ strategies, etc. A model for this situation is to replace instantaneous anticipations by iterations: each voter is given a fixed strategy, which dictates for each possible expected outcome which ballot to cast, and we inductively apply these strategies and update the expected outcome. Reaching a dynamical equilibrium thus means that inductively anticipating each other’s strategy result in a stable state. Conversely, the existence of a cycle means that in some circumstances there is no meaningful way for voters to inductively anticipate other’s strategies. Note that even when limiting the number of steps of counter-strategies, the game-theoretic analysis of voting rules can be quite deep, see [GHR19].

Synchronized iterative voting is also a model for political elections, which are usually preceded by polls informing the voters of others’ intents; we will thus name our dynamical model *Polling Dynamics* (but note that of course this model is imperfect, and that understanding the effect of actual polls on strategic voters is not the sole intent of our work). This point of view has been notably taken in [CPP04], and [RE12] where the emphasis is on the effect of the amount of information given by the polls (ranks, or scores, etc.)

Note also that while asynchronous voting is broader, it makes room for very many parameters and is thus more flexible; in consequence, producing examples in synchronous iterative voting is made in principle more difficult by its constraints.

We shall consider *myopic strategies*, where the behavior of voters only depend on the last poll (however, we will allow at some point voter’s choices to also depend upon the previous ballot they cast, in order to model change aversion). It would be interesting to explore strategies taking into account all past polls, but this is beyond the scope of the present article; let us simply conjecture that if voters compute an expected outcome based on a (possibly weighted) average of previous polls, convergence should be much more likely.

1.2 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 – or in other words, if all candidates present on the ballot are preferred to all candidates not present on it [BF07]. Note that by design of this voting system, there are no reasons not to vote sincerely in this sense (at least when there is a known tie-breaking rule, otherwise see [End13]).

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 outcome rather than the fact that voters embrace this necessity (“*Manipuler c’est bien, la manipulabilité c’est mal*”: “Manipulation is good, manipulability is bad”). Indeed strategies that either need a lot of information on other’s plans or that are too contrived cannot be applied uniformly, creating asymmetries among voters. Even in the absence of manipulation or strategic voting, 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* [Gib73].

Gibbard defines a game form to be *straightforward* if for every player v and every order of preference on the outcomes, v has a strategy b that is dominant for these preferences, i.e. for all strategies of other player, no other strategy of v would yield a better outcome than b . We know that as soon as there are at least 3 candidates, no non-dictatorial voting system can be straightforward for all voters individually [Gib73, Sat75]. Since single voters often have little weight to an election, we can extend this definition: let us say that a situation is straightforward for a given voter or set of voters when they have a choice of ballots which, whatever the ballots cast by other voters are, yields an outcome at least as good (from all of their points of view simultaneously) than any other possible choice. Here “situation” refers to a given method of vote for a given set of candidates, a given set of voters, a given coalition with their possibilities of coordination, and possibly to additional restriction on the allowed strategy. For example, one can impose that ballots cast be sincere; in most voting systems, this leaves each voter with only one allowed strategy and thus trivially implies straightforwardness. Let us give another example: in an election using Plurality or Approval Voting, if a candidate α is the top-preferred of a majority of voters able to cooperate fully, then the situation is straightforward for this set of voters: they can all vote for α and get her elected. In a straightforward situation, a voter or coalition of voters can safely choose how to vote by considering only their preferences, in a way that they will not regret whatever the others choose to vote; in particular they can disregard any poll and any other information to make their choice. Since straightforwardness cannot be hoped in general, one can try to determine whether certain systems often offer weak forms of straightforwardness (see again [Dur15]).

In approval Voting, as soon as there are more than two candidates each voter has several sincere ballots, corresponding to the various point in her order of preference where she 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 straightforward Approval Voting can be in some circumstances.

1.3 Is Approval Voting a Condorcet system “in practice”?

While Approval Voting is known not to be a Condorcet system, several arguments have been raised that could seem to indicate it might be close to 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; this was improved in [SS04] where it is shown (for Approval Voting as well as other voting rules such as Plurality or Instant Runoff Voting) that only Condorcet Winners are elected at strong Nash equilibriums. Strong Nash equilibrium are said by Brams and Sanver to be “globally stable”, but one point not considered in either works above 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.² Let us give more detail on these results, explained in the framework we shall use here.

To be applied, the Leader Rule (LR) 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, this 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 *polling dynamics* (PD). 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 PD 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.

1.4 Description of the main results

The goal of this article is to construct examples showing that in Approval Voting, the PD can exhibit a “bad cycle” even in the presence of a Condorcet winner (the main

²A similar argument has been raised by Warren D. Smith and is used by advocates of Approval Voting and Range Voting, see <https://www.rangevoting.org/AppCW.html>. The argument of Smith is much less rigorous than Laslier’s, since the equilibrium assumption is kept implicit.

point shall not be the mere existence of cycles, a rather unsurprising 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 Section 3.1, we assume all voters apply the Leader Rule, thus additionally implying a rather strong form of rationality under Laslier’s uncertainty model. We show an example in which PD exhibits a cycle where the Condorcet winner cannot be elected. Moreover the *basin of attraction* of the cycle, i.e. the set of initial states leading to the cycle under PD, 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 3.2), we assume a slight modification of the Leader Rule: some voters have several candidates they decide never to approve; the strategy is kept simple, sincere, and consistent. In this example, there is a Condorcet winner and a *consensual loser*, ranked last by a majority of the electorate (actually almost two-third of it). 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 PD and where one of the state elects the consensual loser.

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 PD.

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). They mitigate the impression possibly given by previous works that Approval Voting is a Condorcet-in-practice method:

- 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 expectations 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 PD 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 always disapproved by a large majority of voters.

It could be expected that such bad cycles are in practice rare. In section 4 we explore numerically this question under various cultures modeling the construction of random voters' preferences. The main take-away is that bad cycles are rare when voters apply the LR, but can be more common, appearing in more than 15% of electorates, in its modification where voters give themselves a limit to the candidates they may approve of. Even more striking is that the culture where bad cycles are rarest for the LR is the one where they are most common for its modification.

We also consider more general behavior of voters by introducing in Section 6 *Continuous Polling Dynamics* (CPD), a continuous-space variation of the PD into which the latter can be embedded. We first show that the above discrete bad cycle are stable under perturbation of the discrete model; the considered perturbations model situations 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), and
- 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.

Last, in Section 7 we consider another kind of voter's behavior, where voters are reluctant to approve of a candidate who is ranked low in their preference order. The arbitration between a strategy such as the LR and this reluctance can, surprisingly, make the CPD a *chaotic dynamical system*, making the prediction of the next winners from the observation of a sequence of previous expected winners practically impossible. We introduce some tools from dynamical systems to quantify the amount of chaos in the CPD.

1.5 Brief overview of the literature

Let us end this section by describing some more previous works; with no pretense to exhaustivity, the aim is to describe some directions previously explored in iterative voting.

In the case of Plurality Voting, discussed in slightly more details below, 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. Restricted strategies,

where voters may be constrained are assumed not to use best or better responses, have been studied in [OMP⁺15] and [GLR⁺13].

Non-myopic strategies have been considered for example in [BHY14] and [AGP17]. Note that there voters change their ballots a random or fixed number of times, and the focus is on the quality of the last result: the dynamical aspect of iterative voting is not considered. Both articles show that voters learning from past information produce relatively good outcomes.

Note that in most of these results (the exceptions concerning Plurality), voters are assumed to adjust their ballot one by one (a specific case of asynchronous iterative voting). Other interesting results can be found in [OLP⁺15, ROL⁺15].

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). In the case when strategies are restricted, [OMP⁺15] gives sufficient conditions for convergence in many voting systems.

The question of the quality of equilibriums have been considered for Plurality, Veto and Borda in [BCMP13], where the voters can adjust their votes one at a time, and starting from the state where each voter casts her sincere ballot (which is unique in these voting systems). Concerning Plurality, they find that Nash equilibriums that can be attained in this way are all very good; but in their model, the individual updates are made greedily and any candidate that would start with two or more votes less than the starting winner will never receive any new vote. One could expect real voters to vote not for immediate improvement of the winner, but in order to give a better position to a contender that might receive more support from others.

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); this corresponds to the “2-pragmatist” strategy mentioned e.g. in [RE12], and in [GLR⁺13] which studies restricted strategic voting in asynchronous iterated voting. If we assume that all voters apply a 2-pragmatist strategy, 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 equilibriums: every candidate that is not a Condorcet loser could be elected, depending on the results of the first poll. In a sense, this strategy reduces the PD 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.³ At least a Condorcet loser 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. Instant Runoff Voting) 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.

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 Voting Systems. This article benefited from relevant comments provided by Adrien Fauré, Jean-François Laslier, Reshef Meir and anonymous referees who I warmly thank.

2 Formalism and notations

Let us set up some notations for our discrete model and formalize the Polling Dynamics (PD).

2.1 Candidates, electorate, preferences and strategies.

We consider a finite set of *candidates*, named by lower-case letters from the beginning of the alphabet, $\mathcal{C} = \{a, b, c, \dots\}$. An *outcome*, the set of which is denoted by \mathcal{O} , is a result of an election (or poll, or anticipation of result by a voter) and can be of different nature depending on the setting: it can for example give the number of votes received by each candidate, or only their rankings (supposed to be a linear order: a tie-breaking rule is assumed, e.g. alphabetic order), or only the identity of the winner and the runner-up. A linear order on the set of candidates is written as a word in the candidates names (so that for three candidates, the possible linear orders are $abc, acb, bac, bca, cab, cba$), the first-ranked on the left. We assume that at least the winner and runner-up can be deduced from the outcome; if the outcome is o , we denote them by $w(o)$ and $r(o)$, respectively. We denote $\alpha \succ_o \beta$ to say that in the outcome o , α is ranked before β (so that for example $w(o) \succ_o r(o)$). Using this notation implicitly assumes that the information given by the outcome is sufficient to deduce the order between α and β .

The *preferences* of a voter is a linear order on \mathcal{C} , but with the anti-symmetric property relaxed to allow 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 is denoted in the same way than a linear order, with parentheses to group the tied candidates: $a(bc)d$ means a is the

³In some sense, multiple-round voting systems could be thought of as a way to counter the reluctance of some voters to vote strategically.

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 \mathcal{P} .

We assume Approval Voting is used, so that a *ballot* is an arbitrary subset of \mathcal{C} ; the set of ballots is denoted by \mathcal{B} . For a voter with preferences $\pi \in \mathcal{P}$, a ballot B is said to be *sincere* whenever for all $\alpha \in B$ and all $\beta \notin B$, $\alpha >_{\pi} \beta$.⁴ For example, if there are 4 candidates and $\pi = a(bc)d$, the possible sincere ballots are \emptyset (the empty ballot, aka blank vote, which we will mostly ignore in this article), $\{a\}$, $\{a, b, c\}$, $\{a, b, c, d\}$ (which is in practice equivalent to a blank vote, and which we will therefore also ignore). A sincere ballot is simply determined by a limit of acceptance, all names above this limit being written in the ballot. For each $\pi \in \mathcal{P}$, there are thus at most $|\mathcal{C}| - 1$ sincere ballots that are meaningful (in the sense that they are neither blank nor do they contain all names), and this bound is attained exactly when there are no ties in π .

A *consistent strategy* is a mapping $f : \mathcal{O} \rightarrow \mathcal{B}$; we shall often call them simply *strategies*. We say that a voter applies strategy f if whenever she expects the outcome $o \in \mathcal{O}$, she casts the ballot $f(o)$ (in either an election or in a poll). An *inconsistent* strategy would be given by a sequence $(f_n)_{n \geq 1}$ of such maps, with f_n applied in the n -th iteration. A strategy f is said to be *sincere* if for all $o \in \mathcal{O}$, $f(o)$ is a sincere ballot. Here, we will say that f is a *simple* strategy whenever $f(o)$ only depends on the ranking between the candidates.

A particular example of consistent, simple and sincere strategy, associated with preferences π without ties, is the *Leader Rule* f_{π}^{LR} introduced by Laslier [Las09], defined as follows: for all $o \in \mathcal{O}$ and all $\alpha \in \mathcal{C}$, $\alpha \in f_{\pi}^{\text{LR}}(o)$ if and only if either $\alpha >_{\pi} w(o)$ or $\alpha = w(o) >_{\pi} r(o)$. In other words, applying the Leader Rule consists in approving all candidates preferred to the expected winner, and to vote for the expected winner if and only if she is preferred to the runner-up.

We assume a finite set of *voter types*, named by upper-case letters from the end of the alphabet $\mathcal{T} = \{Z, Y, X, \dots\}$, with all voters of a given type applying the same strategy (this condition will be somewhat relaxed later in this article). 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 $(n_B)_{B \in \mathcal{B}}$ giving the number of each possible ballot cast. The set of all possible elections is denoted by \mathcal{E} . A map $F : \mathcal{O} \rightarrow \mathcal{E}$ is defined by applying the strategies of each voter type given an expected outcome o and compounding the resulting ballots:

$$F(o) = \left(\sum_{\Omega \text{ s.t. } f_{\Omega}(o)=B} n_{\Omega} \right)_{B \in \mathcal{B}}.$$

⁴Sincerity is often defined with respect to a linear order, i.e. without ties. When ties are allowed, there are two possible ways to extend the definition, ours or only asking $\alpha \geq_{\pi} \beta$ instead of $\alpha >_{\pi} \beta$. Choosing strict comparison makes the sincerity a stronger condition, and thus our results are made stronger.

Meanwhile, the use of Approval Voting yields a fixed mapping $AV : \mathcal{E} \rightarrow \mathcal{O}$, which depending on the choice of outcomes gives the number of votes received by each candidate, or their ranking (ties broken by alphabetic order), or the winner and runner-up.

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

$$\begin{aligned} \text{PD} : \mathcal{O} &\rightarrow \mathcal{O} \\ o &\mapsto \text{AV}(F(o)). \end{aligned}$$

This model in particular assumes *consistent strategies* (strategies do not vary in time). Since \mathcal{O} is the space where PD operates, outcomes are also called *states*. PD models the situation where before the election a certain number k of perfect polls are conducted (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

$$o_0 = \text{AV} \left(\left(\sum_{\substack{\Omega \text{ s.t. } b_0(\Omega)=B} n_{\Omega} \right)_{B \in \mathcal{B}} \right)$$

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

We are thus interested in the dynamical properties of PD, i.e. of the behavior of its iterates PD^k and especially in its orbits (i.e. the families $(\text{PD}^k(o))_{k \in \mathbb{Z}_+}$ where $o \in \mathcal{O}$). A *periodic orbit* (also named *cycle*) is a family of distinct states o_1, \dots, o_p such that $\text{PD}(o_i) = o_{i+1}$ for all $i \in \{1, \dots, p-1\}$ and $\text{PD}(o_p) = o_1$. 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 equilibrium*) is a 1-cycle, i.e. a state o such that $\text{PD}(o) = o$. A cycle is said to be *trivial* if all its states have the same winner; otherwise it is said to be *non-trivial* (non-triviality implies that the cycle has period at least 2).

When \mathcal{O} is finite (e.g. when outcomes are the ranking of candidates), we can represent PD by the oriented graph with \mathcal{O} as set of vertices, and with exactly one outgoing edge for each $o \in \mathcal{O}$, with endpoint $\text{PD}(o)$. Cycles are then cycles of the oriented graph, and fixed points are states with a loop (to increase readability, we will omit loops in our representation of graphs of Polling Dynamics; fixed points are thus represented without an outgoing edge). In this case, the *basin of attraction* of a cycle o_1, \dots, o_p is the set of all states o' such that there exist $k \in \mathbb{Z}_+$ and $i \in \{1, \dots, p\}$ such that $\text{PD}^k(o') = o_i$.

Remark 2.1. This model is restrictive as it assumes all voters adjust their ballots simultaneously; see Section 6 for a generalization. While it feels unrealistic that all voters are informed of a poll and adjust their strategies before the next poll, for political elections it stays more realistic than voters adjusting their votes one by one, and the information being updated at each of these tiny steps.

2.3 Remarkable Candidates and the Polling Dynamics.

Note that the preferences play no role in PD; their importance is that they enable to define a sincere strategy (and not all strategies have a preference π for which they are sincere, so the sincerity 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 strictly more voters that strictly prefer α to β than voters that strictly prefer β to α (i.e. in the *majority graph*, assuming indifferent voters abstain, there is an arrow from α to β). 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. When preferences have no ties, this is the usual definition of a Strong Condorcet Winner; when there are ties, a stronger definition could be possible: to dominate, one could ask for a majority of all voters, including abstainers. When ties are allowed, we shall mention whenever our results still hold with this stronger definition.

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 linear order τ on the candidates is a *Condorcet order* whenever each candidate dominates all candidates ranked below them in τ . Last, we will use a stronger notion than Condorcet loser: a candidate is said to be a *consensual loser* whenever there is a strict majority of the electorate that ranks her last (possibly tied with others) in their preferences. A consensual loser rarely exist, and when she does she is a Condorcet loser.

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 Discrete examples of bad cycles

3.1 Non-convergence of the Leader Rule

The main goal of this Section is to prove the following.

Theorem A. *Using Approval Voting, there exists 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,
- PD 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 $\mathcal{C} = \{a, b, c, d\}$ 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 in the majority graph $b > c > d > b$ 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 cycle b, c, d in the majority graph. The precise numbers of classes U to Z are chosen, for the sake of fanciness, to exclude any perfect tie (different sums of distinct powers of 2 never agree).

To define the Leader Rule, the outcomes need only give the winner and runner-up, and they are denoted by ordered pairs wr or w, r . Consider the outcome $o_1 = ba$. 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 $o_2 := \text{PD}(o_1) = da$ – i.e. a stays second, while the previously unthreatening d comes in first position. 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 $o_3 := \text{PD}(o_2) = ca$. 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 $\text{PD}(o_3) = o_1$.

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

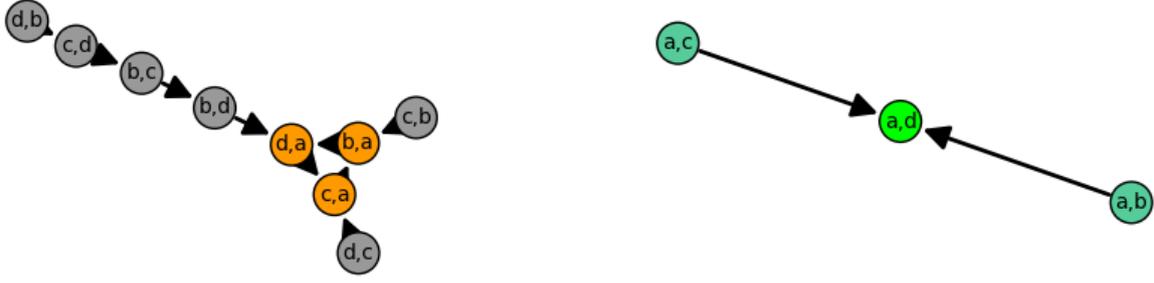


Figure 1: The PD of the first example. States where the Condorcet winner is elected are shown in green, light green for the dynamical equilibrium. The periodic states not electing the Condorcet winner are shown in orange.

Remark 3.1. It is worthwhile to observe (although but a bit tedious to check) that in this example, *all* states can be obtained from a certain choice of a sincere ballots for each voter type. For example, dc is obtained if all voters cast the ballot with only their preferred candidates. That all states are covered is hardly surprising: there are only 12 states, but $3^7 = 2187$ combinations of sincere ballots.

Remark 3.2. 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 candidates will all be sent by PD 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.

Remark 3.3. As pointed out by one of the anonymous referees, Theorem A does not hold for 3 candidates. Assume there is a Condorcet Winner a and two other candidates b, c . Assuming preferences have no ties (an implicit assumption for the LR), there are 6 possible voter types; let $T_{\alpha\beta}$ denotes the voters ranking α first and β second, and let t also denotes the proportion of such voters. We assume there is no equality in any partition into two groups of three voter types, which is a generic condition, to avoid dealing with ties. Since a is a Condorcet winner,

$$T_{ab} + T_{ac} + T_{ca} > \frac{1}{2} > T_{bc} + T_{ba} + T_{cb}$$

$$T_{ab} + T_{ac} + T_{ba} > \frac{1}{2} > T_{bc} + T_{ca} + T_{cb}$$

One of b and c would win a duel against the other, and without loss of generality we assume b does, i.e.

$$T_{ba} + T_{bc} + T_{ab} > \frac{1}{2} > T_{ca} + T_{cb} + T_{ac}.$$

We assume that all voters follow the Leader Rule and prove that the Polling Dynamic converges to one of the outcomes abc or acb .

First, whatever the expected outcome is, a receives either $T_{ab}+T_{ac}+T_{ca}$ or $T_{ab}+T_{ac}+T_{ba}$ votes, depending on whether b is above c or not, in any case more than half.

When a is expected winner, b and c receives $T_{bc} + T_{ba} + T_{cb} < \frac{1}{2}$ and $T_{bc} + T_{ca} + T_{cb} < \frac{1}{2}$ votes respectively, so that a stays ahead and either abc or acb is an equilibrium, attracting the other one of these two states.

Whenever a is second, none of the voters for a vote for the expected winner, and it follows that in the next round, the contender will be behind a . The third candidate will receive the votes of voters preferring her to the expected winner; in the case of c , it means either $T_{ca} + T_{cb} + T_{ac}$ or $T_{ca} + T_{cb} + T_{bc}$ votes, according to whether b or a is expected to win, and both are less than $\frac{1}{2}$. It remains to consider the outcome cab , where b receives the votes $T_{ba} + T_{bc} + T_{ab}$. In this case, we can have a sequence of outcomes cab, bac but then a wins in the next outcome. Whenever a gets at least second, she will therefore become first and stay first ever after.

Now, in every case where a is third, she will receive more than half the votes and c will receive less than half, so that a will get at least the second position in the next round, after which she will take the lead. This ends the proof of convergence to equilibrium.

3.2 Second example: the possible election of a consensual loser

This second example, at the small cost of introducing a modification of the LR accounting for ties, improves on the previous one on two accounts: it necessitates only 3 candidate, and it exhibits a cycle where a consensual loser could get elected.

Theorem B. *Using Approval Voting, there exists an electorate on 3 candidates, with ties in preferences allowed, such that:*

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

Proof. We consider the following types 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 $\alpha \in \mathcal{C}$, $\alpha \in f_{\pi}^{\text{LR}}(o)$ if and only if $\alpha >_{\pi} w(o)$ or $\alpha = w(o) >_{\pi} r(o)$; we call this the *modified* Leader Rule to emphasize the presence of ties since as shown by this example they prevent Laslier's Theorem [Las09] to hold. 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 loser, loosing to b by another landslide 104 to 201: about two-third of the electorate would *never* vote for c , making it a consensual loser (by quite a margin).

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

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

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:

$$Z : \{a\} \quad Y : \{a\} \quad X : \{b\} \quad W : \{c\} \\ a : 103 \quad b : 100 \quad c : 104.$$

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 PD 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 consensual loser c depending on whether the number of polls conducted before the election is odd or even, while the other 2 are equilibriums, one electing the Condorcet winner a the other the Condorcet runner-up b . \square

Remark 3.4. If we accept to depart further from the LR, we could avoid ties in preferences 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). This shows that Theorem B also holds with the stronger definition of Condorcet winner (see Section 2.3).

Remark 3.5. 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 numbers 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 .

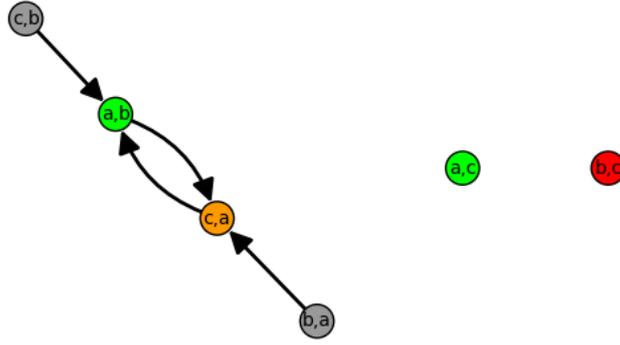


Figure 2: The PD map of the second example. Grey states cannot occur after a certain number of iterations (here, 1 iteration). Light green states are accessible after an arbitrary large number of iterations and elect the Condorcet winner; orange (periodic) and red (fixed point) states are accessible after an arbitrarily large number of iterations but do not elect the Condorcet winner. The orange state is arguably the most problematic, as it elects the consensual loser c .

Remark 3.6. 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 .

4 Experimental exploration of bad cycles

The examples of the previous section are interesting from a theoretical point of view, showing that the existence of equilibriums and the quality of candidates they elect, do not by themselves suffice to ensure a good outcome. But from a practical point of view, it could be that the bad cycles we exhibited are so rare that they do not matter too much.

In this section, we present experimental estimates of the frequency of bad cycle in electorate where a Condorcet winner exists. We continue to follow the setup of Section 2, using either the Leader Rule or the modified Leader Rule used in Section 3.2. This

leaves us to model the electorate, and for this we compare several “cultures”.

4.1 Description of the experimental setup

We fix the number of candidates to 6, $\mathcal{C} = \{a, b, c, d, e, f\}$ and the number of voter types to 20 (to be realistic, the number of voters type should be quite larger than the number of candidates so that many preferences are represented, but we kept the number limited for computational power reasons). The influence of these numbers is tested at the end of the section.

Each voter type is assigned a number of voters (pseudo-)uniformly drawn in $[0, 1]$ (since we use floats, this is equivalent to using very large integers; using smaller, more realistic integers would only increase the use of the tie-breaking rule, which would still be rare anyway). Each voter type is assigned a preference, which is modeled by a linear order on $\overline{\mathcal{C}} = \mathcal{C} \cup \{\text{limit}\}$; if we use the Leader Rule, `limit` is simply ignored, but if we use its modified version, the candidates ranked beyond `limit` are considered tied and all ranked last (and thus never put on this voters’ ballots). For example, the order $a > b > \text{limit} > c > d > e > f$ corresponds to the preference $abcdef$ when the LR is used, and to $ab(cdef)$ when its modified version is used. This choice enables us to completely decouple the assumed strategies from the culture, i.e. the construction of the preferences.

We first consider the Impartial Culture: each voter type draws uniformly a linear order on $\overline{\mathcal{C}}$. This is the most commonly studied culture and often serves as a benchmark.

We also consider more structured political cultures where various independent axes are assumed to represent various criteria on which the position of each candidate is evaluated by each voter. The dimension d , representing the number of those axes, parametrizes this family of cultures. The “political landscape” is then represented by the hypercube $[0, 1]^d$, equipped with the ℓ^1 distance

$$d((x_1, \dots, x_d), (y_1, \dots, y_d)) = \sum_{i=1}^d |x_i - y_i|.$$

This choice of distance corresponds to summing disagreements along each axis, seems at least as good as any other, and is slightly less demanding than the Euclidean distance for computations. Assuming a *d-dimensional culture* means we draw uniformly a position in $[0, 1]^d$ for each candidate and for each voter type. The preferences of a voter type is obtained by sorting the candidates by increasing distance to the voter type; `limit` is inserted by drawing a random “distance” to the voter type, with the same law than the distance between two uniform points in $[0, 1]^d$ (it would make no sense to assign a position for `limit` common to all voters, and various voter types should be allowed to be more or less inclusive in their preferences). This choice ensures that `limit` is on average in the middle of the preferences, independently of d and similarly to the Impartial Culture.

4.2 Experimental results

We ran extensive experiments for each combination of a set of strategies (Leader Rule or its modification) and each political culture among the Impartial Culture and the d -dimensional cultures with $d \in \{1, 2, 3, 400\}$. In each case, we generated 100 000 independent electorates and recorded whether a Condorcet winner exists, and if yes whether the Polling Dynamic has a “bad cycle” (or bad equilibrium), some of whose states do not elect the Condorcet winner. Each run took a couple of hours on a single core of a modern CPU. The results are provided in Table 1.

Strategy \ Culture	Impartial	$d = 1$	$d = 2$	$d = 3$	$d = 400$
Leader Rule	(70%) 1.3%	(100%) 0.0%	(90%) 0.2%	(87%) 0.2%	(85%) 0.3%
modified Leader Rule	(75%) 6.3%	(92%) 15%	(89%) 7.8%	(88%) 6.0%	(87%) 3.0%

Table 1: Experimental results with 6 candidates and 20 voter types. In small, the proportion of electorates where a Condorcet winner exists; in normal size, the proportion of electorates with a bad cycle or equilibrium, among electorates having a Condorcet winner (both rounded). The 95% confidence Wilson score interval gives a deviation of less than 0.23 percentage point for the 15% value and less than 0.2 percentage point for the other values; the deviation is less than 0.04 percentage point for the value 0.3% and below.

We see that the Leader Rule is very effective in electing the Condorcet winner in practice, with only the Impartial Culture witnessing slightly larger odds of not electing her. Structured landscape very rarely produce bad cycles, with the linear ($d = 1$) landscape having produced none in 100 000 attempts (this is no accident, see Section 5).

On the contrary, when voters apply the relaxed rule and, thus, refuse to vote to candidate below their personal threshold no matter what, bad cycles are more common. They never dominate, but they culminate *precisely* in the case of a linear landscape, with 15% of electorates leading to a bad cycle in the Polling Dynamic. The situation that is most favorable to the Leader Rule, is also the worst one for its modification! This means that when some voters expect a candidate they do not like to be a Condorcet winner, their best interest can be not to apply the Leader Rule. At first sight, this may seem difficult to reconcile with the optimality of the Leader Rule proven by Laslier [Las09]. The point is that Laslier assumes a voter thinking she will be the only one to change her mind, and trying to optimize the expected outcome. Our experimental result show that, as soon as a voter anticipates that other voters may also change their votes in view of a poll, the Leader Rule may not be their optimal course of action anymore.

This experiment also gave the opportunity to find various kind of problematic Polling Dynamics, presented in Figures 3, 4, 5, 6.

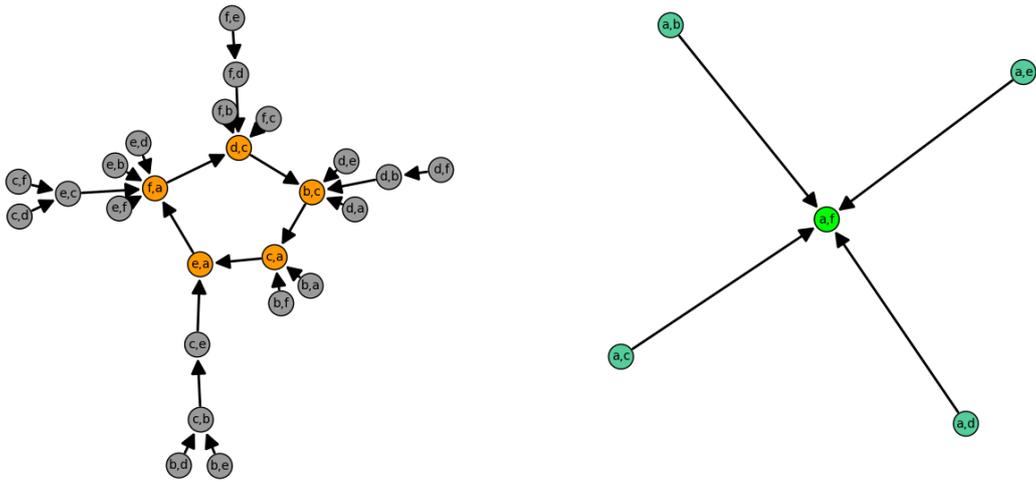


Figure 3: A bad cycle of period 5 under the Leader Rule, produced with the $d = 400$ landscape.

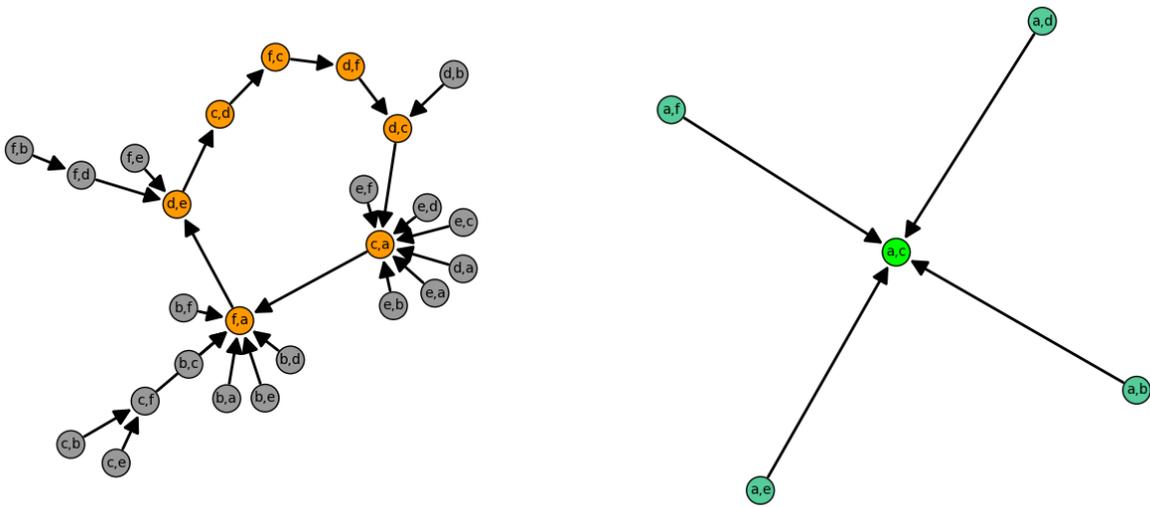


Figure 4: A very exceptional bad cycle of period 7 under the Leader Rule, produced with the Impartial Culture.

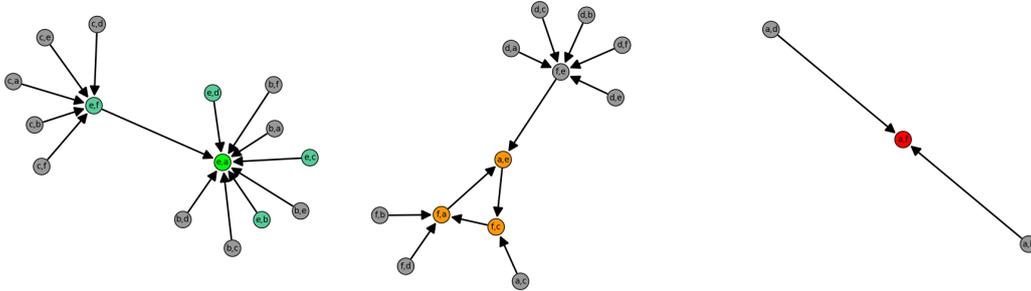


Figure 5: Both a bad cycle and an equilibrium not electing the Condorcet winner under the modified Leader Rule, produced with the Impartial Culture.

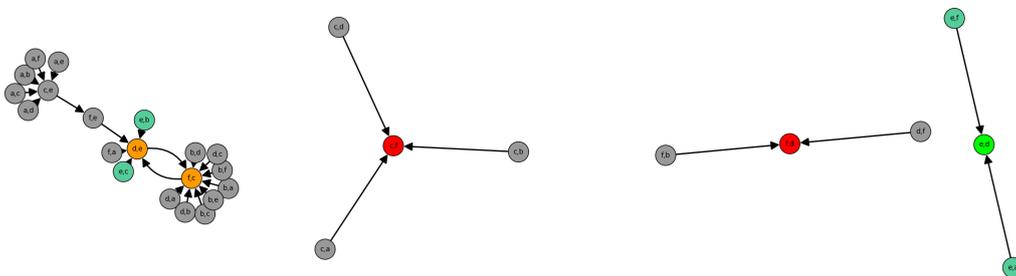


Figure 6: A bad cycle and two equilibria not electing the Condorcet winner under the modified Leader Rule, produced with the Impartial Culture. Note two remarkable features of this example: the equilibrium that does elect the Condorcet winner has a very small basin of attraction, only 3 states out of 30; and there are states where the Condorcet winner is elected, but that lie in the basin of attraction of the bad cycle (this never happens with the Leader Rule).

4.3 Effect of the numbers of candidates and voter types

We report in Table 2 more modest numerical experiment to assert the influence of the number of candidates and the number of voters. Here we only draw 6 250 electorates in each condition, yielding typically 4 times larger confidence intervals than samples of size 100, 000, leaving the expected deviations under the percentage point.

$ \mathcal{C} $ & $ \mathcal{T} $ Cul. & Strat.	3, 10	4, 10	6, 10	6, 20	6, 30	8, 20
Impartial, LR	(92%) 0	(84%) 0.3%	(71%) 1.0%	(70%) 1.3%	(69%) 1.3%	(60%) 2.0%
1-dim., LR	(100%) 0	(100%) 0	(100%) 0	(100%) 0	(100%) 0	(100%) 0
2-dim., LR	(97%) 0	(94%) 0.05%	(90%) 0.3%	(90%) 0.2%	(91%) 0.05%	(87%) 0.5%
3-dim., LR	(96%) 0	(92%) 0.1%	(85%) 0.4%	(87%) 0.2%	(89%) 0.1%	(84%) 0.4%
400-dim., LR	(95%) 0	(90%) 0.1%	(81%) 0.3%	(85%) 0.3%	(89%) 0.3%	(81%) 0.6%
Impartial, modified LR	(94%) 2, 1%	(88%) 3.5%	(77%) 5, 9%	(75%) 6.3%	(75%) 6.4%	(66%) 8.0%
1-dim., modified LR	(98%) 2.6%	(96%) 5.4%	(90%) 10%	(92%) 15%	(92%) 19%	(88%) 20%
2-dim., modified LR	(97%) 1.5%	(95%) 3.4%	(89%) 6.2%	(89%) 7.8%	(91%) 7.8%	(84%) 11%
3-dim., modified LR	(97%) 1.4%	(93%) 2.7%	(86%) 5.3%	(88%) 6%	(90%) 5.9%	(84%) 8.2%
400-dim., modified LR	(96%) 0.9%	(92%) 2.1%	(84%) 3.4%	(87%) 3.0%	(89%) 2.6%	(82%) 4.1%

Table 2: Complementary experimental results to check robustness. Samples have 6 250 pseudo-random electorates (except for 6 candidates and 20 voter types, where we reported the previous results with larger samples). In small, the proportion of electorates where a Condorcet winner exists; in normal size, the proportion of electorates with a bad cycle or equilibrium, among electorates having a Condorcet winner (both rounded). We observe that the number of candidates influences results, more candidates making bad cycles more likely. The number of voter types seems to have little influence in most conditions.

We see that the number of voter types has relatively small influence, at least in the range tested, with the exception of the 1-dimensional culture with the modified LR, where a larger voter type makes bad cycles and equilibriums more likely. The main take

away, valid in all considered conditions except the 1-dimensional culture with the LR, is that:

A larger number of candidates seems to make bad cycles and equilibriums more common among situations where a Condorcet Winner exists.

Computational power prevented us to explore larger number of candidates with samples large enough to draw conclusions.

5 Convergence of the Leader Rule in one-dimensional cultures

In view of the above experiments, it seems that the one-dimensional culture combines well with the Leader Rule. Black's Median Voter Theorem [Bla86] implies the existence of a Condorcet winner, and from [Las09] it follows that there exist equilibriums and that all equilibriums elect ω . Here we prove that in this particular culture, we moreover have convergence of the Polling Dynamics.

Definition 5.1. We say that an electorate *can be modeled by a one-dimensional culture* when there exist a *positional mapping*

$$x : \mathcal{C} \cup \mathcal{T} \rightarrow \mathbb{R}$$

with the following property: the preferences π of any voter type $\Omega \in \mathcal{T}$ is given by

$$\forall \alpha \neq \beta \in \mathcal{C}, \quad \alpha >_{\pi} \beta \quad \text{if and only if} \quad |x(\alpha) - x(\Omega)| < |x(\beta) - x(\Omega)|.$$

For simplicity, we assume further that $|x(\alpha) - x(\Omega)| \neq |x(\beta) - x(\Omega)|$ for all $\alpha \neq \beta$ and all Ω , so that all preferences are linear orders (without ties); and we assume that there is no partition of the voter types in two groups that have equal total number of voters, so that no duel between two candidates would end up in a tie (these are *generic* conditions: they are stable under small perturbation of the data and any data can be approximated arbitrarily close by data satisfying those conditions).

The goal of this section is to prove the following result.

Theorem 5.2. *If the electorate can be modeled by a one-dimensional culture and applies the Leader Rule, then the Polling Dynamic converges to an equilibrium (which elects the Condorcet winner, as prescribed by Laslier's Theorem).*

The end of the section is dedicated to the proof of Theorem 5.2. We use outcomes reduced to the winner and runner-up, we assume that the electorate can be modeled by a one-dimensional culture and we fix a positional mapping x . Denoting by $N = \sum_{\Omega \in \mathcal{T}} n_{\Omega}$ the total number of voters, a *median* is a value $m \in \mathbb{R}$ such that

$$\sum_{\Omega: x(\Omega) \leq m} n_{\Omega} \geq \frac{N}{2} \quad \text{and} \quad \sum_{\Omega: x(\Omega) \geq m} n_{\Omega} \geq \frac{N}{2}$$

i.e., at least half the voters have a position on the left of m , and at least half have position on the right of m . The assumption that there is no partition of voter types in two groups of equal size ensures that there is a unique median m , coinciding with the position of some voter type: $m = x(X_m)$, where the voter type X_m is called the *median type*. By assumption, there is a single candidate μ whose position is closest to $x(X_m)$.

We start with the following particular instance of Black's Median Voter Theorem, which we prove for the sake of completeness.

Lemma 5.3. *The candidate μ is a Condorcet winner.*

Proof. Let α be any other candidate. Voters preferring μ to α are those positioned on the half line L of endpoint $\frac{1}{2}(x(\alpha) + x(\mu))$ and containing $x(\mu)$. Since $x(\mu)$ is closer from $x(X_m)$ than $x(\alpha)$, L contains $x(X_m)$ and thus contains at least half the voters. Since there are no possible ties, μ dominates α . \square

Lemma 5.4. *Let $\alpha\beta$ be an outcome such that $\alpha \neq \mu$ and let $\alpha'\beta' = \text{PD}(\alpha\beta)$. Then either $|x(\alpha') - x(\mu)| < |x(\alpha) - x(\mu)|$, or $\alpha' = \alpha$ and $|x(\beta') - x(\mu)| < |x(\beta) - x(\mu)|$.*

Proof. The positions $x(\alpha)$ and $x(\beta)$ divide the real line in three component: the open bounded interval I between them, the open half-line A with extremity $x(\alpha)$ avoiding I , and the open half-line B with extremity $x(\beta)$ avoiding I .

If $x(\mu) \in A$, α will receive the votes of all voters positioned in a half line starting at $(x(\alpha) + x(\beta))/2$ and containing $x(\mu)$, while other candidates receive votes only from voters positioned in A or in $I \cup B$. It follows that $\alpha' = \alpha$ and β' is the candidate positioned next to α , in the direction of $x(\mu)$ (possibly $\beta' = \mu$). In this case, we thus have $\alpha' = \alpha$ and $|x(\beta') - x(\mu)| < |x(\beta) - x(\mu)|$.

In all other cases, α' is the candidate positioned next to α , in the direction of μ (possibly $\alpha' = \mu$), therefore $|x(\alpha') - x(\mu)| < |x(\alpha) - x(\mu)|$. \square

Let \prec be the "lexicographic" order on outcomes, defined by $\alpha\beta \succ \alpha'\beta'$ whenever

$$|x(\alpha') - x(\mu)| < |x(\alpha) - x(\mu)|, \text{ or } \alpha' = \alpha \text{ and } |x(\beta') - x(\mu)| < |x(\beta) - x(\mu)|.$$

The previous lemma shows that along an orbit of PD, the outcome can only decrease in this order until μ becomes winner. Since there cannot be an infinite decreasing sequence of outcome, eventually μ becomes winner. Now, as is well-known, this is a stable situation: for all β , $\text{PD}(\mu\beta) = \mu\beta_0$ where β_0 is the candidate getting the better score in a duel against μ . Therefore, every state converges under PD to the state $\mu\beta_0$, and the proof of Theorem 5.2 is complete.

6 A continuous-state model

The above discrete-time, discrete space model is quite crude and thus can fail to convince that cycles could appear in practice. Discreteness of time is relevant, since polls occur at precise times, triggering some ballot adjustments: completely asynchronous dynamics,

where voters have the opportunity to change their vote one by one and the information is shared immediately after each vote change, is unrealistic whenever voters are numerous, in particular for political election in a large population.

Discreteness of space is more of a weakness of our 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 introduce 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 whatever the voter's strategies are in case of an almost tie. Moreover, in this continuous model the cycle are attractive: all states close enough to the cycle are attracted to it. In the next section, we will take advantage of this continuous setting to consider other possible models of voters' reactions to polls.

6.1 Continuous Polling Dynamics

A continuous state-space

We expand the setting of Section 2 to a continuous-space model. To each voter type Ω is associated a set $\mathcal{B}_\Omega \subset \mathcal{B}$ 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 $f_\Omega(\mathcal{O})$ of the strategy f_Ω of this type. This restriction is meant to reduce the dimension of the space of states.

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

$$\Delta(\mathcal{X}) := \left\{ P = (p_x)_{x \in \mathcal{X}} \in [0, 1]^{\mathcal{X}} \mid \sum_{x \in \mathcal{X}} p_x = 1 \right\}.$$

An element of $\Delta(\mathcal{X})$ can be interpreted as a probability vector on \mathcal{X} or simply as the proportions of a distribution of some quantity over the elements of \mathcal{X} . A *state* is now an element of

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

i.e. a state $s \in \mathcal{S}$ gives for each voter type and each admissible ballot the proportion of the voters of that type planning 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. A state $s = (s^\Omega)_{\Omega \in \mathcal{T}}$ is said to be *extreme* when for each Ω , in $P = s^\Omega$ exactly one of the proportion is 1 and the others are 0 (i.e., all voters in any given voter type will cast the same ballot). If $B \in \mathcal{B}_\Omega$, we will denote by $\mathbf{1}_B$ the element of $\Delta(\mathcal{B}_\Omega)$ that gives weight 1 to B and 0 to the other ballots. The set of extreme states is finite and can be identified with $\prod_{\Omega \in \mathcal{T}} \mathcal{B}_\Omega$, the element $(\mathbf{1}_{B_\Omega})_{\Omega \in \mathcal{T}} \in \mathcal{S}$ representing the situation where for each Ω , all voters of this type cast the ballot B_Ω .

Generalized strategies and Continuous Polling Dynamics

There is a natural mapping which associates to each state the election that would result from the casting of the planned ballots:

$$v : \mathcal{S} \rightarrow \mathcal{O}$$

$$s = \left(s^\Omega = (s_B^\Omega)_{B \in \mathcal{B}_\Omega} \right)_{\Omega \in \mathcal{T}} \mapsto \left(\sum_{\Omega} n_\Omega \cdot s_B^\Omega \right)_{B \in \mathcal{B}}$$

where s_B^Ω is considered zero whenever $B \notin \mathcal{B}_\Omega$. The mapping $AV \circ v : \mathcal{S} \rightarrow \mathcal{O}$ thus sends a state to the outcome that would result from it. Here, we will consider as set of outcomes \mathcal{O} the set of all detailed vote results; an outcome thus takes the form $(V_\alpha)_{\alpha \in \mathcal{C}}$ where V_α is the proportion of votes in favor of candidate α . Recall that we assumed a tie-breaking rule, e.g. by alphabetical order, so that from such an outcome the total ranking of all candidates, in particular the winner and runner-up, can be deduced. These extended outcomes enable voters decision to depend on the scores of the different candidates.

A *generalized strategy* for a voter type Ω is a map $g_\Omega : \Delta(\mathcal{B}_\Omega) \times \mathcal{O} \rightarrow \Delta(\mathcal{B}_\Omega)$ recording the respective proportions of voters of type Ω casting each admissible ballot in reaction to an outcome. With this definition, we allow that not all voters of a given type react in the same way to an outcome; some may for example be more reluctant to change their ballot (hence the $\Delta(\mathcal{B}_\Omega)$ factor on the starting set, serving as memory of the previously planned ballots by this type of voters). We do not permit the decisions of voters to depend on the ballots previously cast by other types than theirs: this information is considered unavailable. As in Section 2 we introduce a relation between preferences and strategy: a generalized strategy is *sincere* if it only gives positive proportion to sincere ballots for this type's preferences.

An *continuous electorate* is now the data of \mathcal{T} and for each voter type Ω of: a set of admissible ballots \mathcal{B}_Ω , a number of voters n_Ω , and a generalized strategy g_Ω .

The *continuous polling dynamics* (CPD) associated with a continuous electorate is the map $\Phi : \mathcal{S} \rightarrow \mathcal{S}$ defined by

$$\Phi(s) = \left(g_\Omega \left(s^\Omega; AV \circ v(s) \right) \right)_{\Omega \in \mathcal{T}}$$

In other words, from a state s and for each voter type Ω we extract the information s^Ω on ballots previously cast by the voters of this type, we compute the expected outcome $AV \circ v(s)$, and then apply the generalized strategies to build the next state. Note that we shifted where in the cycle

$$\dots \rightarrow \text{outcome} \rightarrow \text{voters' choice} \rightarrow \text{outcome} \rightarrow \dots$$

we record information: in the discrete setting, we recorded the successive outcomes; here we record the successive choices of voters, which usually provides richer information.

Embedding and perturbation of the Polling Dynamics

Given an electorate (in the sense of Section 2), we choose as suggested above $\mathcal{B}_\Omega = f_\Omega(\mathcal{O})$ and consider as generalized strategies the translation of the strategies in the continuous

model, i.e. $g_\Omega(P, o) = \mathbf{1}_{f_\Omega(o)}$ – in other words, all voters apply the strategy of their type given the expected outcome, irrespective of the information P on their previously cast ballots. This yields the following CPD:

$$\Phi_0(s) = \left(\mathbf{1}_{f_\Omega \circ AV \circ v(s)} \right)_{\Omega \in \mathcal{I}}.$$

To compute $\Phi_0(s)$, one thus first applies v to s to obtain the expected election, then applies AV to the result to obtain the expected outcome, then applies to this the strategies f_Ω to obtain the voters new choice of ballots. 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 (discrete) polling dynamics given by the electorate. In this section our goal is to study the dynamics of CPDs which are *close* to Φ_0 in a sense we now define.

We consider the supremum norm $|\cdot|$ on \mathcal{S} , i.e. given two states $s = (s_B^\Omega)$ and $\bar{s} = (\bar{s}_B^\Omega)$ we set

$$\|s - \bar{s}\| = \sup_{\Omega \in \mathcal{I}, B \in \mathcal{B}_\Omega} |s_B^\Omega - \bar{s}_B^\Omega|$$

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

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

We use D as a shortcut for $D_{\mathcal{S}}$, which is a metric on the set of all CPDs. 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.

6.2 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 exists an electorate on three candidates such that:*

- *there are a Condorcet winner and a consensual loser,*
- *each voter type is assigned a consistent, sincere and simple strategy,*
- *for any CPD Φ where a proportion $p \geq 85\%$ 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 voters keeping their previous ballot and the CPD being arbitrary when margins are lower than 4%), Φ has an attractive 2-cycle one of whose states elects the consensual loser.*

An example is presented in Figure 7.

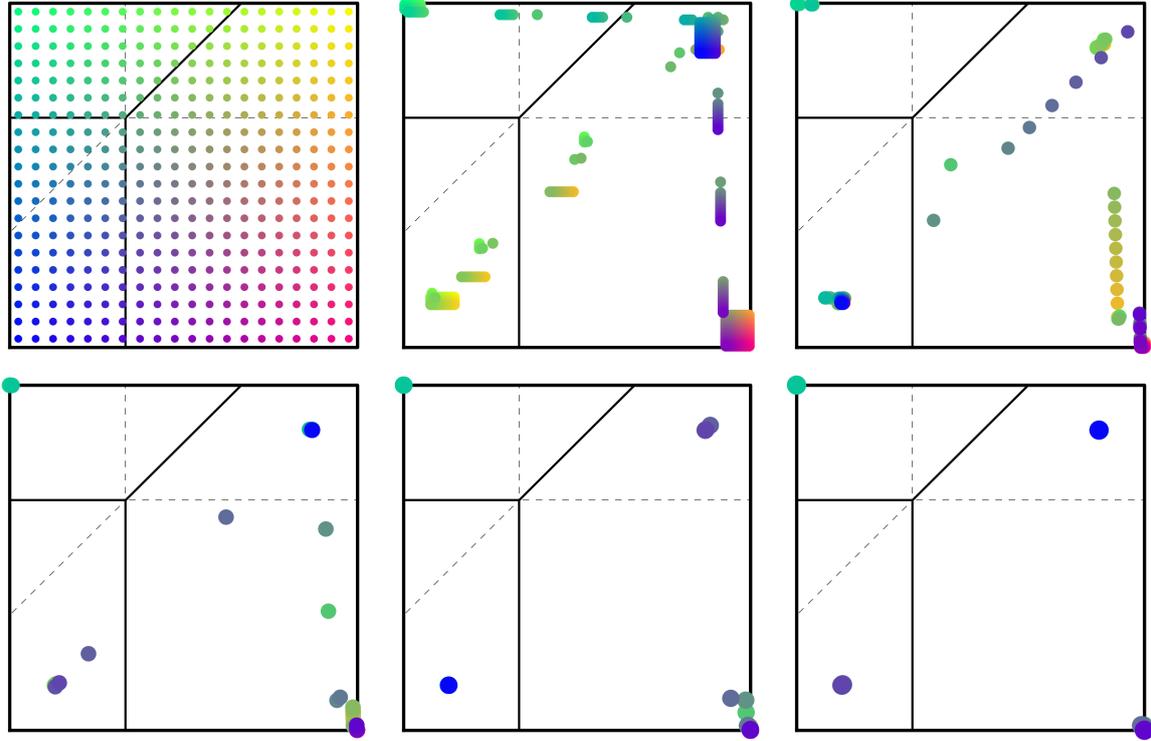


Figure 7: An example of a CPD verifying the hypotheses of Theorem C. Top-left: initial points of the plane; then from left to right then top to bottom, the first five iterations of a continuous CPD where, whenever margins are above 4%, 85% of voters apply the Leader Rule and 15% keep their ballot unchanged. Points are drawn with the color corresponding to initial position (darker above for better readability). After the first iteration, already most points are attracted to the periodic points near the four corners. The top-left and bottom-right corners are attracting fixed points, and there is an attractive orbit of period 2 near the other corners.

Proof. We consider the following four-types electorate

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

with the modified Leader Rule as in Section 3.2 and the corresponding sets of admissible ballots. Again, a is a Condorcet winner and c a consensual loser (now refused by 7/12th of the electorate: we traded some badness of c for more stability). Note that we can make a a Condorcet Winner in the strongest sense (see Section 2.3) by using the trick mentioned in Remark 3.4. Observe that

- $\Delta(\mathcal{B}_Y) = \Delta(\{\{a\}\})$ and $\Delta(\mathcal{B}_W)$ are both singletons and can be ignored in the product,
- $\Delta(\mathcal{B}_X) = \Delta(\{\{b\}, \{a, b\}\})$ can be identified with $[0, 1]$, denoting by x the proportion of voters of type X that cast the ballot $\{a, b\}$,
- $\Delta(\mathcal{B}_Z) = \Delta(\{\{a\}, \{a, b\}\})$ 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 the notation above, we thus have

$$s_{\{a,b\}}^Z = z \quad s_{\{a\}}^Z = 1 - z \quad s_{\{a,b\}}^X = x \quad s_{\{b\}}^X = 1 - x.$$

We can thus identify \mathcal{S} and $[0, 1]^2$ with coordinates (x, z) , see Figure 8, which contains the end of proof and which we now explain.

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 \mathcal{S} corresponding to the six possible outcomes (the lines themselves are attributed to one outcome according to the tie-breaking rule). 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 CPD with the property assumed in the third item, we have $\Phi(x, z) = ((1 - p)x, (1 - p)z)$ whenever $(x, z) \in A_1$ and $\Phi(x, z) = (p + (1 - p)x, p + (1 - p)z)$ whenever $(x, z) \in A_2$. One easily checks that $\Phi(A_1) \subset A_2$ and $\Phi(A_2) \subset A_1$. It follows that $\Phi^2(A_1) \subset A_1$, and since Φ^2 is a contraction on A_1 (of ratio $(1 - p)^2$), the Contraction Mapping Theorem ensures it has a fixed point $(x_1, z_1) \in A_1$. Then the orbit of (x_1, z_1) is a 2-cycle with one state inducing the outcome abc and the other inducing cab . Moreover any element of the open set $A_1 \cup A_2$ is attracted to this cycle exponentially fast: e.g. if $(x, z) \in A_1$, for all $n \in \mathbb{N}$ we have $|\Phi^{2n}(x, z) - (x_1, z_1)| \leq (1 - p)^{2n}|(x, z) - (x_1, z_1)|$ and $|\Phi^{2n+1}(x, z) - \Phi(x_1, z_1)| \leq (1 - p)^{2n+1}|(x, z) - (x_1, z_1)|$. \square

6.3 A general stability result

The above example was meant to be tangible and explicit, but the underlying phenomenon is quite general.

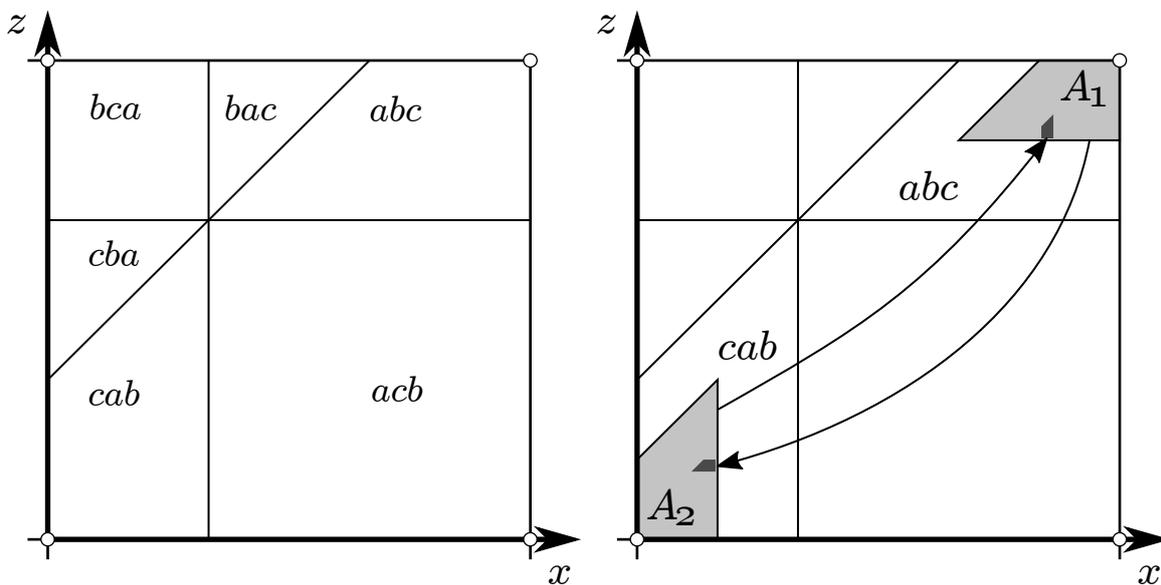


Figure 8: A continuous-space example. *Left:* the state space \mathcal{S} . Coordinates x, z represent respectively the proportion of voters of type X casting the ballot $\{a, b\}$ (the others casting $\{b\}$) and the proportion of voters of type Z casting the ballot $\{a, b\}$ (the others casting $\{a\}$). Regions are drawn showing how the outcome depends on x, z . The corners are the extreme states, corresponding to the four outcomes that are possible under the CDP Φ_0 : if all voters adjust their ballot after a poll, 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-left corner and the acb region to the lower-right corner. *Right:* on the same state space, any CPD where 85% of voters adjust their ballot according to the modified Leader Rule when candidates are separated by 4% margins will send the light-grey regions A_1 and A_2 into each other (images shown in dark grey), thus ensuring a 2-cycle with one state near the upper-right corner, one near the lower-left corner.

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

Consider first the continuous electorate corresponding to the CPD Φ_0 , i.e. assume temporarily that all voters apply their type's strategy at each poll. For each poll i along the cycle, the ballots B_i gives the candidates some scores $S_i(a), S_i(b), \dots$ expressed in percentage of votes received; we set

$$\varepsilon_0 = \min_{1 \leq i \leq k} \min_{\alpha \neq \beta \in \mathcal{C}} |S_i(\alpha) - S_i(\beta)|$$

the least margin separating two candidates according to any ballot vector B_i .

When $\varepsilon_0 > 0$ (i.e. when the tie-breaking rule is not needed in the considered cycle), $(0, \varepsilon_0]$ is non-empty and for all $\varepsilon \in (0, \varepsilon_0]$ we denote by A_i^ε the sub region, in the part of \mathcal{S} corresponding to B_i , where all candidates are separated by a fraction of the electorate at least ε . Then A_i^ε has non-empty interior (it is neither empty nor reduced to e.g. a subspace of dimension less than the dimension of \mathcal{S}).

Theorem D. *Assume the above setting, in particular that the Polling Dynamics has a cycle o_1, \dots, o_k , with each o_i surrounded by a region $A_i^{\varepsilon_0}$ in which o_i is the outcome with margins at least ε_0 between the candidates. If $\varepsilon_0 > 0$, for all $\varepsilon \in (0, \varepsilon_0)$ there is a $\delta > 0$ with the following property. Every CPD Φ such that, for each i :*

- Φ is continuous on each A_i^ε , and
- $D_{A_i^\varepsilon}(\Phi, \Phi_0) < \delta$

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 CPD that is close enough to Φ_0 near the cycle exhibits a similar cycle. In particular, if we only assume a large enough fraction of each voter's type adjust their ballot at each polls, and if we only assume the prescribed strategies away from almost-ties, we still have a cycle.

Proof. By definition of ε_0 , each A_i^ε is the intersection of \mathcal{S} with a polyhedron containing in its interior the extreme point corresponding to B_i , and Φ_0 maps A_i^ε 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 . \square

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. The cycle of states is attracting as soon as Φ is contracting on each A_i^ε ; this is for example the case when Φ

is a variant of a Polling Dynamics, where only a fraction of voters adjust their ballots upon receiving the result of the last poll, such as in Theorem C. More generally, Φ is contracting on each A_i^ε when in the absence of almost ties, voters of any given type tend to agree more on their next choice than they did on the previous one. In particular the discrete Polling Dynamics, where all voters of a type make the same choice given an expected outcome, exhibits a very extreme form of contractivity.

7 Reluctance to cooperate and chaos

Now we turn to other generalized strategies, motivated by the reluctance a voter is expected to have to approve of candidates that are low on her preferences; such an approbation can be seen as a *cooperation* with other voters sharing a strong dislike for the expected winner. Consider a voter driven by a strategy, say the Leader Rule, and planning to cast a certain ballot. When the information of a poll becomes available, if adding one or several names on her ballot could improve the outcome, there is a competition between the strategic reasoning and the impetus not to approve of these disliked candidates. It can also be tempting to remove names from the ballot if they seem not to be needed anymore to prevent a worse candidate to be elected.

We shall make the following assumptions:

Continuity: *all generalized strategies are continuous: a small change in the outcome o or the planned ballots s^Ω results in a small change in $g_\Omega(s^\Omega, o)$.*

Opportunity-driven strategies: *voters are less likely to include in (and more likely to exclude from) their ballot a name that is low in their preferences when the probability that this inclusion (or exclusion) would influence the outcome is small.*

Political weight of full outcomes: *voters aim primarily to elect a highly preferred candidate, but also try that candidates they dislike much have fewer votes and candidates they like much have more votes.*

The first assumption is meant to depart from the unrealistic discontinuities implied by a perfect application of strategies such as the Leader Rule: when two candidates are very close one to the other, which one is first should not change completely the choices of voters.

The second assumption, which is at the heart of this section, is to be understood as a trade-off between a reluctance to approve of a low candidate in one's preferences and the potential gain in the outcome. This trade-off can be estimated differently from voter to voter, even inside a given type –this is typically what the continuous model is designed to model. It is to be noted that the “opportunity-driven strategies” assumption is very different from the models above: far from almost ties, voters will be more numerous to restrict to ballots with few names, while near almost-ties they will be more numerous to instead apply a strategy such as the Leader Rule.

The last assumption backs the second and reflects that, to a certain extent, the final number of votes received will have political effects to be taken into account, e.g. a

disliked candidate may look legitimate if ranked high, even if not elected, and could be more likely to be a threat in next elections. In some countries, votes can also have consequences on party funding, and voters could thus be careful not to fund unnecessarily a candidate they dislike. This assumption is flexible, as one can put more or less weight to these considerations.

Let us now consider a simple, piecewise linear model following these requirements and have a look at its CPD.

7.1 A simple model based on the example of Theorem C

We consider again the following four-types electorate

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

(recall that a is a Condorcet winner and c a consensual loser). As above, the corresponding sets of admissible ballots makes \mathcal{S} a square with coordinates (x, z) where x, z are the proportions of voters of type X, Z respectively that “cooperate” by voting $\{a, b\}$.

We recall an outcome is a triple $o = (V_a, V_b, V_c)$, and we choose here the following generalized strategies (where we do not use the possibility of dependency on $s^\Omega = x$ or z):

$$g_Z(z, (V_a, V_b, V_c)) = C \circ S(V_a, V_b, V_c) \qquad g_X(x, (V_a, V_b, V_c)) = C \circ S(V_b, V_a, V_c)$$

where C, S will be defined below, to be interpreted as follows:

- S is a “safety function”, quantifying how unlikely it seems that cooperating would be useful to counter a threat by c (notice how V_a, V_b are exchanged in g_X , to take into account the preferences of this type). S will be in particular very small when both c is expected winner or close to be expected winner, *and* collaborating has a good chance to prevent her to win,
- C is a “collaboration function”, translating a level of safety into a proportion of collaborations (high safety resulting in low collaboration).

We take here a very simple safety function:

$$S(V_1, V_2, V_3) = \begin{cases} |V_2 - V_3| & \text{when } V_2 > V_1 \\ \frac{1}{2}|V_2 - V_3| + \frac{1}{2}|V_1 - V_3| & \text{otherwise.} \end{cases}$$

i.e. when the second-preferred candidates is ranked higher than the preferred one, the safety is the margin in her race with the least-preferred one; otherwise, safety is an average of the margins in races of the two preferred candidates against the last one. For example, if $V_1 \gg V_3 \simeq V_2$ then the safety is not too small (the least preferred candidate has little chance of winning) but not maximal, since there is an opportunity to ensure c finishes at the last rank. When $V_1 \simeq V_2$, the safety is close to both $|V_1 - V_3|$ and $|V_2 - V_3|$ (in particular, S is continuous).

Finally, we choose a simple collaboration function extrapolating linearly between the value 1 (all voters collaborate) when the safety vanishes, and the value 0 (no voter collaborates) when the safety is large enough:

$$C(t) = (1 - 5t)_+ := \max(0, 1 - 5t)$$

where as indicated $(\cdot)_+$ is the positive part, and where the coefficient 5 quantifies the reluctance to collaborate: the higher, the less likely X and Z voters will collaborate. 5 is rather moderate: $1 - 5t$ reaches 0 only at $t = 0.2$, i.e. when the safety margins reaches a staggeringly high 20% of voters.

Since a state (x, z) results in the outcome $(V_a, V_b, V_c) = (3x + 4, 3z + x, 5)$, these choices yield the CPD Φ defined by:

$$\Phi(x, z) = \left((1 - 5|3x - 1|)_+, (1 - 2.5|3x - 1| - 2.5|3z + x - 5|)_+ \right)$$

when $2x + 4 \geq 3z$ and

$$\Phi(x, z) = \left((1 - 2.5|3z + x - 5| - 2.5|3x - 1|)_+, (1 - 5|3z + x - 5|)_+ \right)$$

when $2x + 4 \leq 3z$. While Φ may look like an innocent map, drawing an orbit of this map reveals an interesting pattern (Figure 9). It turns out *all* orbits yield pretty much the same image.

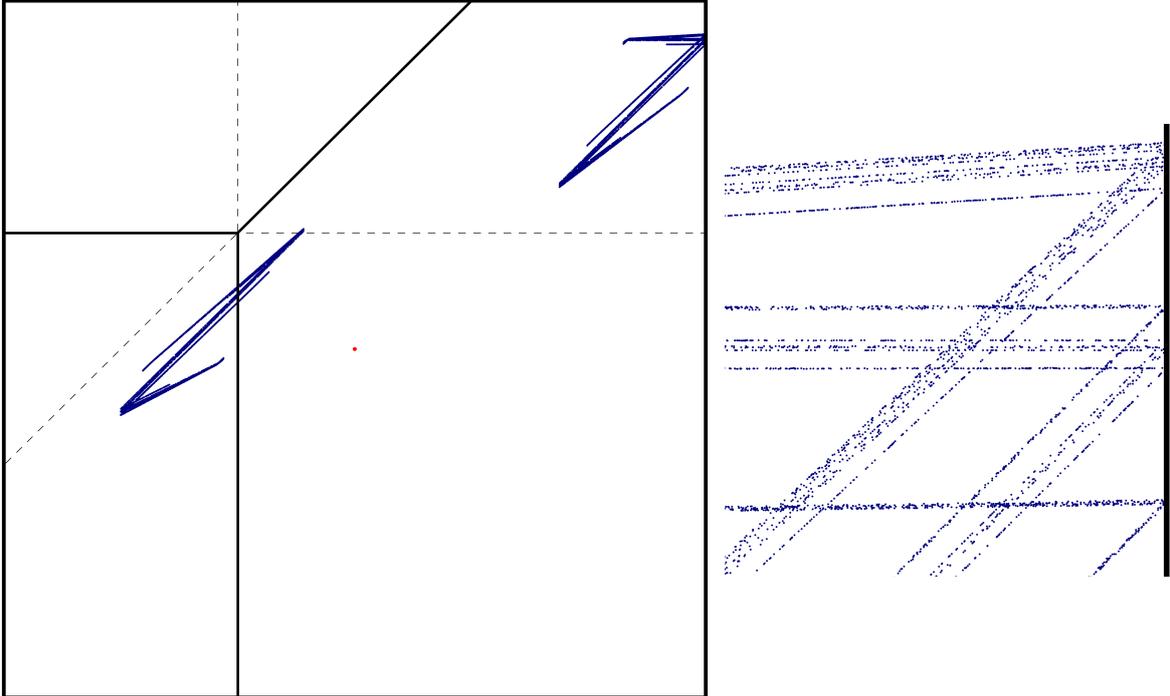


Figure 9: One orbit of the CPD Φ . Left: full picture. Right: a zoom to the upper-right part showing a linear/fractal structure.

In the next section, we will introduce tools (well known in the mathematical community) to assert the chaotic nature of this dynamics, and more importantly of the sequence of elected candidates along any orbit. While we will not *prove* chaos here, we will show compelling numerical evidence. We will also construct a simpler example where chaos follows from well-known results in the dynamical system literature. This is to be understood as a rather theoretical study; in practice there is very little difference between a bad cycle and a chaotic attractor. However it is worthwhile to realize that in a continuous-state model, orbits can have much more varied behavior than converging to an equilibrium or a periodic orbit.

7.2 Asserting chaos: a primer on entropy

In dynamical systems, *chaos* is not a formally defined word, but refers to a number of properties all seeking to translate the idea that orbits behave in an unpredictable fashion. The most prominent one is *entropy*, which comes in a number of versions. We start by giving a brief overview of entropies for maps, before introducing the variant that will be most relevant to us.

The *topological entropy* $h_{\text{top}}(f)$ of a map f is a non-negative number, defined in brief as the exponential rate of growth with $n \in \mathbb{N}$ of the number of length n orbit segments $(x, f(x), f^2(x), \dots, f^{n-1}(x))$, after identifying orbit segments that stay “close” one to another. A precise definition is for example given in [KH95], together with proofs of the claims stated below. The frontier of chaos is usually drawn between exponential and sub-exponential growth; with respect to entropy, one thus says that a map is chaotic whenever its topological entropy is non-zero.

When the map moreover has a invariant probability measure μ , i.e. a law of random variable X such that $f(X)$ has the same law as X , one can define the *Kolmogorov-Sinai entropy*⁵ $h_{\text{KS}}(f, \mu)$ of the pair (f, μ) . It is similar to the topological entropy, except that orbit segments are in a way weighted using μ : if many distinguishable orbit segments exist, but very few occur most of the time when we choose x randomly with law μ , then the Kolmogorov-Sinai entropy is rather small. When f is continuous and acts on a compact space, these two entropies are related by the *variational principle*:

$$h_{\text{top}}(f) = \sup_{\mu} h_{\text{KS}}(f, \mu)$$

where the least upper bound is taken over all f -invariant probability measures. What is most important to us is that for every invariant probability measure, $h_{\text{top}}(f) \geq h_{\text{KS}}(f, \mu)$, so that proving $h_{\text{KS}}(f, \mu) > 0$ is stronger than proving $h_{\text{top}}(f) > 0$.

Now, when $f = \Phi$ is a CPD, it may happen that the topological entropy is positive while the winner of the election is constant, or ultimately constant, on each orbit. In this case, the chaos of the map is irrelevant to the core feature of the election. To extract this feature from orbits, we introduce some notation. For each candidate α , let $W_{\alpha} \subset \mathcal{S}$ be the set of states leading to the election of α , formally

$$W_{\alpha} = \{s \in \mathcal{S} \mid \alpha \text{ wins in } \text{AV} \circ v(s)\}.$$

⁵Also called *metric entropy* or *measure-theoretic entropy*.

The set $\{W_\alpha : \alpha \in \mathcal{C}\}$ is a partition of \mathcal{S} , and given $s \in \mathcal{S}$ we denote by $w_s \in \mathcal{C}$ the winner of s , i.e. the candidate such that $s \in W_{w_s}$.

The set \mathcal{C} of candidates will serve as an alphabet, and sequence of winners will be written as words. Formally a *finite word* of length n is an expression $\alpha_1\alpha_2\cdots\alpha_n$, i.e. a n -tuple written without coma and parentheses. The set of words of length n is thus denoted by \mathcal{C}^n , and the set of finite words is $\mathcal{C}^* := \cup_{n \in \mathbb{N}} \mathcal{C}^n$. An *infinite word* ω is an infinite expression $\alpha_1\alpha_2\cdots$ (i.e. a sequence, written without comma and parentheses) and the set of infinite words is denoted by $\mathcal{C}^\mathbb{N}$. For all $k < m \in \mathbb{N}$ we set $\omega_{k:m} = \alpha_k\alpha_{k+1}\alpha_{k+2}\cdots\alpha_m$; each $\omega_{k:m}$ is called a *subword* of ω (of length $m - k + 1$).

Given any state s , its *winners word* is the word $\omega(s) = w_s w_{\Phi(s)} w_{\Phi^2(s)} \cdots$. What we want to address is how difficult it would be to guess a length ℓ segment of winners

$$\omega(s)_{k+1:k+\ell} = w_{\Phi^{k+1}(s)} w_{\Phi^{k+2}(s)} \cdots w_{\Phi^{k+\ell}(s)}.$$

We shall measure this with a variant of entropy for words. As above, two versions (topological and Kolmogorov-Sinai) can be defined.

Definition 7.1. Let $\omega = \alpha_1\alpha_2\cdots$ be an infinite word. For all $\ell \in \mathbb{N}$, let $S^\ell(\omega)$ be the number of distinct subwords of length ℓ of ω . We define the *topological entropy* of the word ω as

$$h_{\text{top}}(\omega) = \lim_{\ell \rightarrow \infty} \frac{1}{\ell} \log S^\ell(\omega).$$

For all $n, \ell \in \mathbb{N}$ with $\ell \leq n$ and all finite word θ of length ℓ , we denote by $S_n^\theta(\omega)$ the number of times the subword θ appears in $\omega_{1:n}$, we set

$$P_n^\theta(\omega) := \frac{S_n^\theta(\omega)}{n - \ell + 1}$$

the proportion of length ℓ subwords of $\omega_{1:n}$ equal to θ , and by $P_n^\ell(\omega)$ the “probability vector” $(P_n^\theta(\omega))_{\theta \in \mathcal{C}^\ell}$. Finally we define the *Kolmogorov-Sinai entropy* of the word ω as

$$h_{\text{KS}}(\omega) = \lim_{\ell \rightarrow \infty} \frac{1}{\ell} \limsup_{n \rightarrow \infty} H(P_n^\ell(\omega))$$

where

$$H(p_1, p_2, \dots, p_k) = \sum_{i=1}^k -p_i \log p_i \quad \text{with the convention } 0 \log 0 = 0.$$

Let us point out a few observations to explain these definitions.

The limit in the definition of $h_{\text{top}}(\omega)$ exists by Fekete’s lemma since the sequence $(\log S^\ell)_\ell$ is subadditive (a subword of length $\ell + k$ is the concatenation of a subword of length ℓ and a subword of length k , so that $S^{\ell+k} \leq S^\ell \times S^k$). The topological entropy of a word corresponds to the topological entropy of a particular dynamical system constructed from ω called a “subshift”, but we will not use this. Observe that $h_{\text{top}}(\omega) \in [0, \log |\mathcal{C}|]$ since there are in total only $|\mathcal{C}|^\ell$ words of length ℓ . The interpretation of $h_{\text{top}}(\omega)$ is that it takes the value h when the number of subwords of length ℓ in ω is roughly of the order of $e^{h\ell}$ (or subexponential if $h = 0$).

In the definition of $h_{\text{KS}}(\omega)$ the limit exists again by subadditivity, which is more subtle than above and one of the reasons to choose the function H . When $(p_i)_{1 \leq i \leq N}$ runs over all “probability vectors” of dimension N (i.e. $p_i \in [0, 1]$ and $\sum_i p_i = 1$), the expression $H(p_1, \dots, p_N)$ runs over $[0, \log N]$, and it attains its maximum $\log N$ at a single point, namely $(\frac{1}{N}, \frac{1}{N}, \dots, \frac{1}{N})$ (this is a simple application of the concavity of $x \mapsto -x \log x$). Moreover $H(p_1, \dots, p_N, 0, 0, \dots, 0) = H(p_1, \dots, p_N)$, and it follows that $h_{\text{KS}}(\omega) \in [0, h_{\text{top}}(\omega)]$, with equality when all subwords of a given length ℓ that appear in ω do so with asymptotically the same abundance. When some subwords appear much more often than others, the Kolmogorov-Sinai entropy is lower.

The value of an entropy is to be interpreted as the “uncertainty” of a random guess, with the case an object drawn among N uniformly having uncertainty $\log N$; this makes uncertainty linear in the number of objects to be guessed in the sense that guessing ℓ independent objects each uniformly drawn among N has uncertainty $\ell \log N$. More generally, one defines the uncertainty of a choice made according to a probability vector (p_1, \dots, p_N) as $H(p_1, \dots, p_N)$, and this definition enjoys many natural properties (see e.g. [Wal82] Theorem 4.1).

To guess a segment $\omega_{k+1:k+\ell}$, one can take all subwords of ω of length ℓ and choose one randomly. If this choice is made uniformly, the uncertainty is of the order of $\ell \cdot h_{\text{top}}(\omega)$. If this choice is made with each subword given a probability proportional to its abundance in $\omega_{1:n}$ for large n , uncertainty is down to approximately $\ell \cdot h_{\text{KS}}(\omega)$.

If one knows $\omega_{k:m}$ and wants to guess the next letters $\alpha_{m+1} \cdots \alpha_{m+\ell}$, one would rather consider all subwords of ω of length $m - k + 1 + \ell$ starting with $\omega_{k:m}$, and guess among the words obtained by removing this $\omega_{k:m}$ prefix from them. If this choice is made uniformly, the uncertainty is

$$\log \frac{S^{m-k+1+\ell}(\omega)}{S^{m-k+1}(\omega)} \sim \ell \cdot h_{\text{top}}(\omega) \quad \text{when } \ell \rightarrow \infty.$$

If the choice is weighted by the respective abundance of the subwords of length $m - k + 1 + \ell$, then the uncertainty is asymptotically $\ell \cdot h$ where h depends on $\omega_{k:m}$: some cases may give much more information on the sequel than others. Then $h_{\text{KS}}(\omega)$ is an average of these values h , weighted according to the abundance of $\omega_{k:m}$ (rare subwords weight less in this average).

Definition 7.1 has two features that matter in our situation. First, its relevance: the entropy (either topological or, better, Kolmogorov-Sinai) of a winners word $\omega(s)$ gives a quantification of the chaos, as an uncertainty akin to randomness, in the sequence of winners starting from the state s . While $\omega(s)$ depends on s , in some situations we will observe that $h_{\text{KS}}(\omega(s))$ is independent of s . In any case, the existence of a state with $h_{\text{KS}}(\omega(s)) > 0$ means that at least in some starting intents of the voters, the Continuous Polling Dynamics results in an unpredictable result. Also note that for all s , we have $h_{\text{top}}(\Phi) \geq h_{\text{KS}}(\omega(s))$, and the CPD Φ is chaotic as soon as there is one starting state with positive entropy. Second, it is easy to estimate numerically these quantities. For h_{KS} , one simply plots the sequence $\ell \mapsto H(P_n^\ell(\omega))$ for a fixed large n . If for ℓ large enough the points are well-aligned, the slope of the interpolating line can be taken as an estimation of the Kolmogorov-Sinai entropy of $\omega(s)$ (see eg. Figure 10). Taking various

s then adds robustness to the analysis. With this tool at hand, let us return to the Example of Section 7.1.

7.3 Exploration of chaos for some CPD

If we take the example of Section 7.1, the word of the successive winners starting from the center of the space \mathcal{S} is (32 first letters shown, capital c for better readability):

$$\omega(.5, .5) = aaaCaCaaaCaCaCaaaCaCaCaCaCaaaCaC \dots$$

There is some visible structure in this word, for example there are never two successive c no four a in a row; apart from the first three letters, the word seems constructed only with the blocs $\beta = ca$ and $\gamma = caaa$. However, the succession of these blocks looks somewhat random:

$$\omega(.5, .5) = aaa\gamma\beta\gamma\gamma\beta\gamma\gamma\gamma\beta\gamma\gamma\beta\gamma\gamma\gamma\beta\beta\beta\gamma \dots$$

To quantify the randomness of $\omega(.5, .5)$, we proceed as explained in the previous section and plot the sequence $\ell \mapsto H(P_{220}^\ell(\omega))$, obtaining the graph shown in Figure 10. We observe an extremely good alignment from $\ell = 4$ onward, with slope $\simeq 0.229$. Every starting state s yields very similar results, a strong numerical indication that the map Φ is chaotic, with chaotic sequences of winners. To check whether all states have orbits accumulating on the attractor pictured in Figure 7, we have drawn the first iterates of Φ in Figure 11. All orbits seem to accumulate to the attractor, which appears in fact connected (a fact that can be confirmed by topological arguments), made of a very thin curve between two arrow-shaped parts, which are exchanged by Φ .

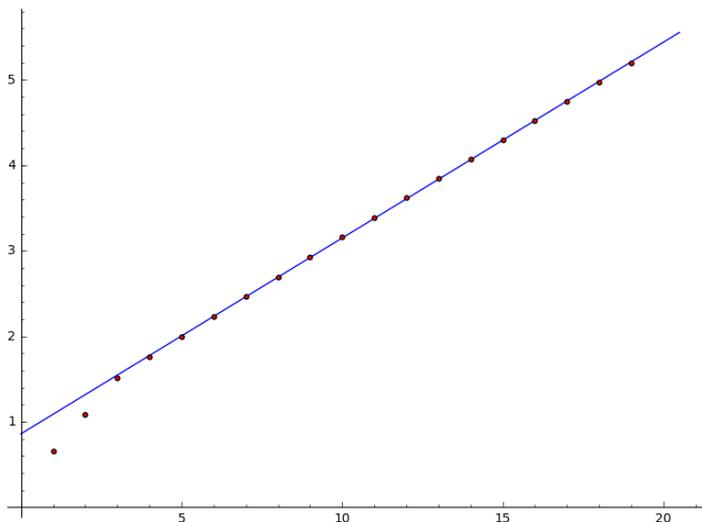


Figure 10: Plot of the entropy-estimating sequence $\ell \mapsto H(P_{220}^\ell(\omega))$. The slope of the marked line is $\simeq 0.2291$

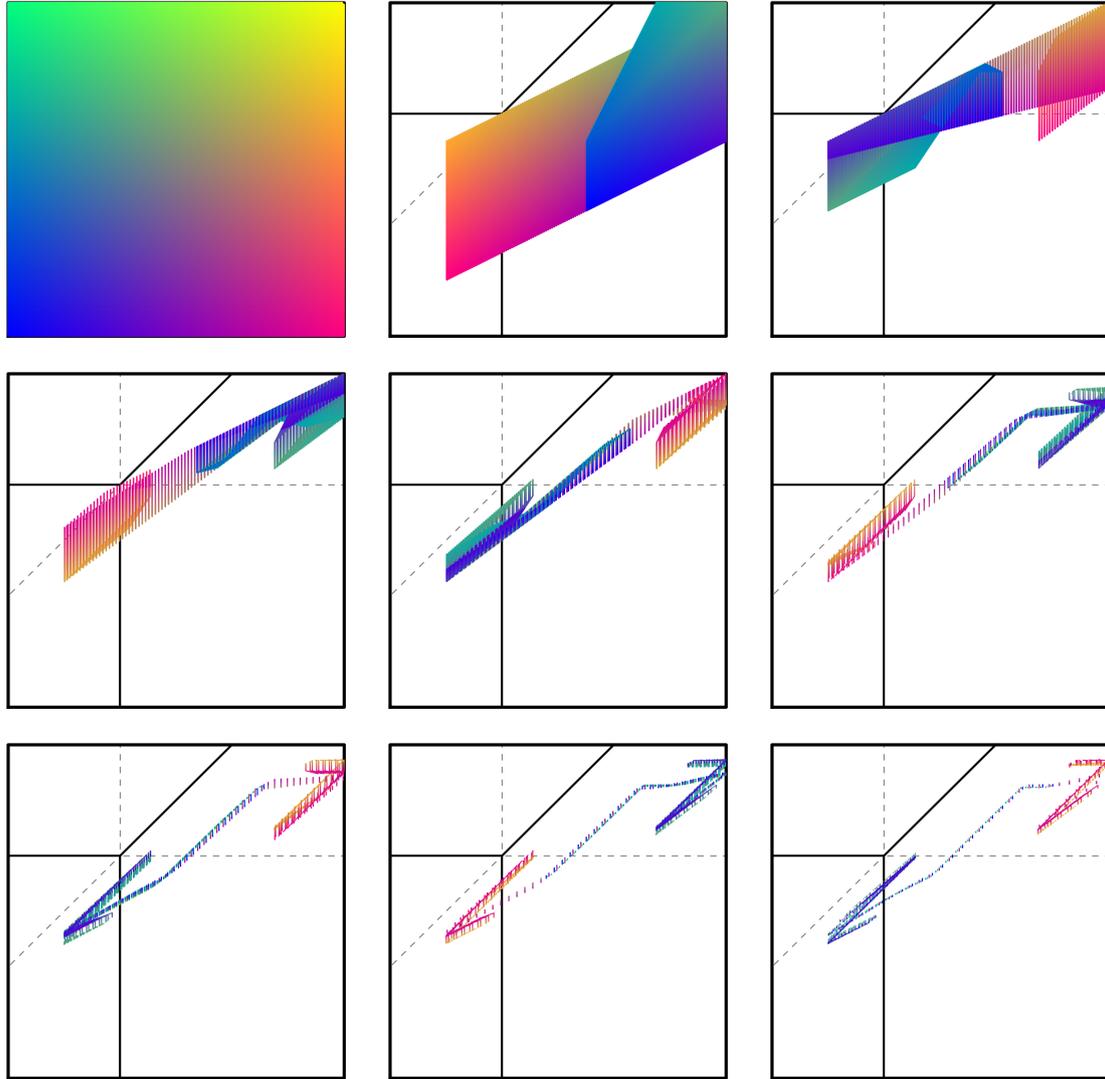


Figure 11: From left to right then top to bottom, 200×200 points of the state space regularly spaced and their 8 first iterates by Φ (colors preserved along each orbit).

While this is only a specific example, it was not difficult to construct: every addition made to the example of Theorem C has been chosen as the simplest I could think of satisfying the hypotheses of Continuity, Opportunity-driven strategies, and Political weight of full outcomes. To dispel the possible feeling that this is an accidental discovery, let us first look at the same example but with randomly chosen populations for all four voter types. The following examples complements the previous one to form a representative sample of what a quick experimental exploration revealed.

Example 7.2. Take

$$n_X = 0.05, \quad n_Y = 0.02, \quad n_Z = 0.8, \quad n_W = 0.8$$

Then a is constantly elected. Observe that it always gets at least .82 votes while c receives only .8, so the only possible contender would be b , but she would need almost all of Z voters to cooperate and most X voters not to cooperate, a behavior prevented in our model. This is a non-chaotic, actually very stable example.

Example 7.3. Take now

$$n_X = 0.6, \quad n_Y = 0.08, \quad n_Z = 0.56, \quad n_W = 0.82$$

Then applying the same method as above, we get a higher entropy estimation:

$$h_{\text{KS}}(\omega(.5, .5)) \simeq 0.36$$

but the word itself looks a bit different, with long sequences of consecutive a :

$$\omega(.5, .5) = aaaCaaaCaCaaaaaCaCaaaaaaaaaCaaaaaaaaaaaaaaaaCaCaCaCaCa \dots$$

This could be example of an “intermittent” behavior, with slow regions where Φ is relatively tame but from which all orbits eventually escape to enter strongly chaotic regions, constantly alternating between two behaviors: predictable and chaotic. See one orbit in Figure 12 (left); other orbits produce very similar pictures.

Example 7.4. Changing only slightly the previous example,

$$n_X = 0.6, \quad n_Y = 0.08, \quad n_Z = 0.56, \quad n_W = 0.81$$

a radical change in behavior appears: entropy seems to vanish, with $H(P_{220}^\ell)$ plateauing abruptly from $\ell = 10$ onward. When looking at subwords of $\omega(.5, .5)$ of length 10 and more, one observes that there are exactly 22 of them, no matter the length. Closer observation then reveals that $\omega(.5, .5)$ is the concatenation of copies of the length 22 word $aaacacaaacacacacaaaaaa$. Other starting states yield similar results, indicating that Φ has here an attracting periodic orbit, of period 22 (see Figure 12, right). At the time scale of an election, it would appear chaotic, but in the longer run it is not.

This examples incites us to plot very long orbits on top of our entropy estimation, to rule out an attracting periodic orbit of length significantly larger than the maximal length of subwords used in the entropy estimation. Rigorously proving positive entropy would need more sophisticated mathematical tools (one would search for a “horseshoe”).

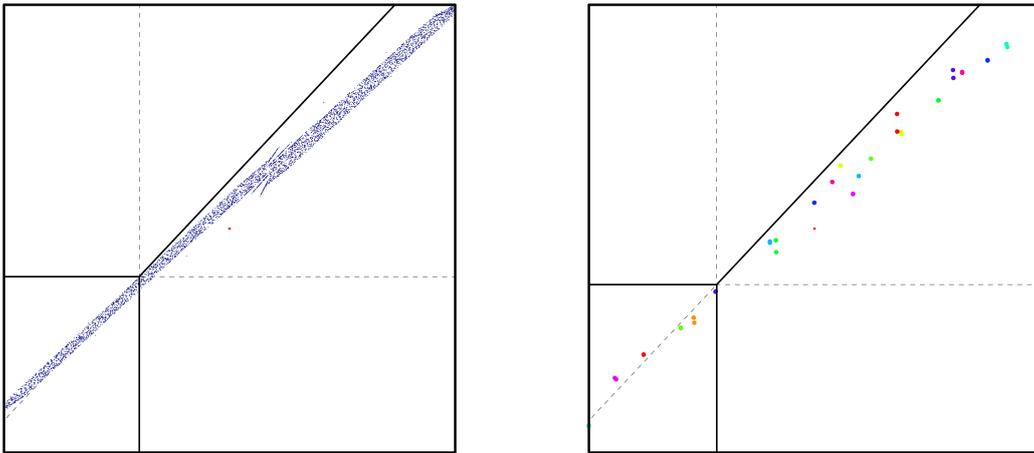


Figure 12: Comparison of the orbits of $s = (.5, .5)$ for two close populations: example 7.3 on the left, example 7.4 on the right, with colors of period 22. Both runs have 5000 points (larger points are used on the right for readability).

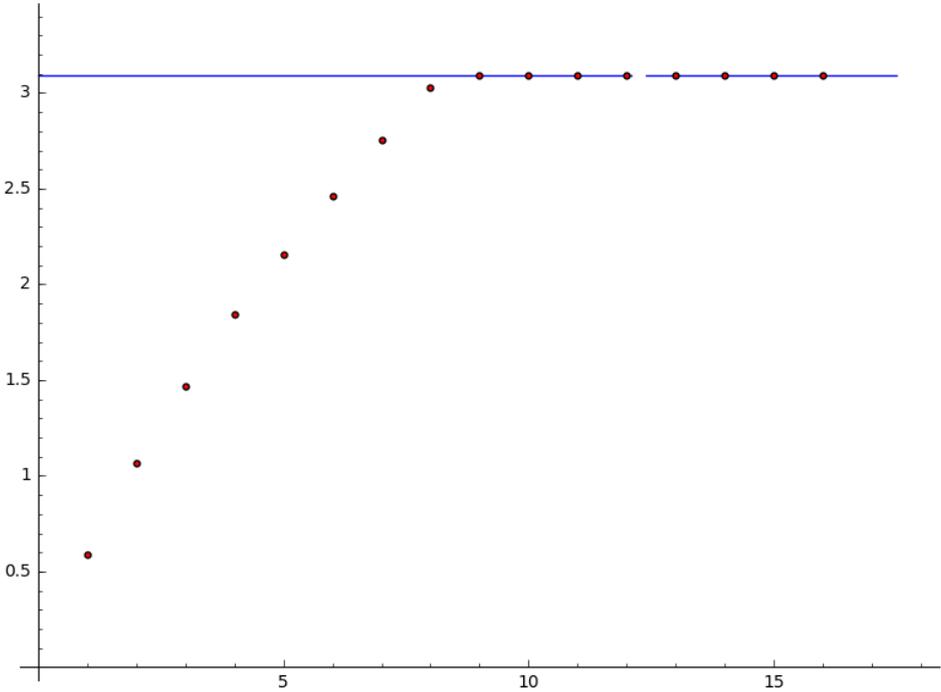


Figure 13: Plot of the entropy-estimating sequence $\ell \mapsto H(P_{20}^\ell(\omega))$ for example 7.4

Last, one can wonder what happens if we change the model. Let us give only one example, far from exhausting the possibilities.

Example 7.5. We consider the case studied in Section 7.1 with as single change a different collaboration function

$$C(t) = \frac{1}{1 + 45t}$$

which is decreasing convex, takes the value .1 at $t = 0.2$ where the previous reached 0, and has much stronger slope at 0 thus yielding less collaboration when safety is small and more collaboration when safety is large. We see again a chaotic behavior (entropy of $\omega(s)$ estimated at 0.42), but with a quite different attractor, see Figure 14.

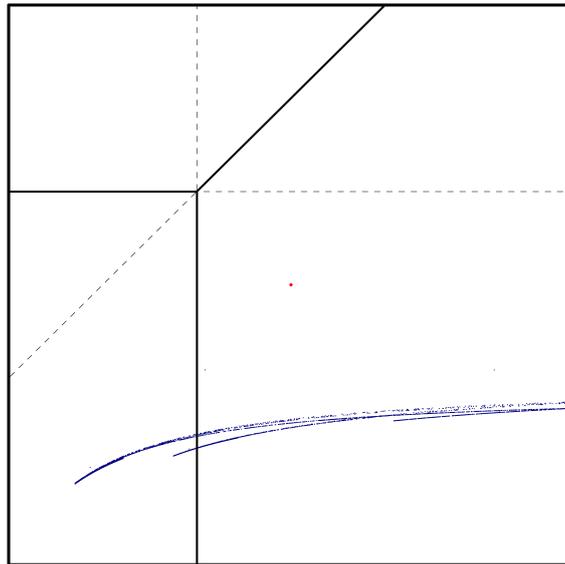


Figure 14: Orbit of $s = (.5, .5)$ for an alternative collaboration function (other starting states yield a very close picture).

7.4 A provably chaotic example

To conclude this section, let us give a very simple example where chaos can be rigorously proven without effort, only relying on the existing mathematical literature.

We consider again three candidates $\mathcal{C} = \{a, b, c\}$ and three voter types:

$$Z : abc \quad 2 \qquad Y : b(ac) \quad 3.5 \qquad X : c(ab) \quad 4.5$$

Voters of type Y always cast the ballot $\{b\}$, and voters of type X always cast the ballot $\{c\}$. Here b is a Condorcet winner, while c and a are consensual losers. Let z be the

proportion of voters of type Z “collaborating”, i.e. casting the ballot $\{a, b\}$, the $1 - z$ remaining fraction casting the ballot $\{a\}$. The state space is reduced to an interval, $\mathcal{S} = [0, 1]$ with coordinate z . We take the same model as in Section 7.1 for the behavior of voters of type Z , but with a higher reluctance to cooperate:

$$g_Z(z, (V_a, V_b, V_c)) = C \circ S(V_a, V_b, V_c) \quad C(t) = (1 + 10t)_+, \quad S(V_a, V_b, V_c) = |V_b - V_c|$$

where S is written that way to simplify, since candidate a will be last for all values of (V_a, V_b, V_c) that can be obtained for any z . Given z , the outcome is given by $V_b = (3.5 + 2z)/10$ and $V_c = 0.45$, so that $|V_b - V_c| = \frac{1}{10}|2z - 1|$; this yields the following CPD:

$$\Phi(z) = \begin{cases} 2z & \text{when } z \leq \frac{1}{2} \\ 2 - 2z & \text{when } z \geq \frac{1}{2} \end{cases}$$

This is the well-known *tent map*, which is an archetype of chaotic map. In particular, it has entropy $h_{\text{top}}(\Phi) = \log 2$ (see e.g. [KH95] page 499); moreover, the *Lebesgue measure* Leb is invariant for Φ , ergodic and of maximal entropy: $h_{\text{KS}}(\Phi, \text{Leb}) = \log 2$. As a consequence, for almost all z (in the sense of Lebesgue measure), the winners word also have $h_{\text{KS}}(\omega(z)) = \log 2$, and the orbit of z is everywhere dense in $[0, 1]$. Using other invariant measures, for example “Bernoulli measures”, one can show that for any value $h \in [0, 2]$ there are uncountably many z such that $h_{\text{KS}}(\omega(z)) = h$, but all these z form together a negligible set (let us recall that a subset of \mathbb{R} is *negligible* when for all ε , there exist a countable family of intervals whose union contains the subset, and whose sum of lengths is lesser than ε ; this is to be interpreted as any continuous random variable taking value outside the subset with probability 1).

8 Conclusion

We have considered the synchronized iterated Approval Voting, assuming voters apply *simple, consistent, sincere* strategies. This “Polling 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, or as a (crude) model of actual polls before an election.

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 consensual loser,
- iii.* these bad cycles are quite uncommon under the Leader Rule, but can represent up to 15% (or more with enough candidates) of simulated electorate having a

Condorcet Winner when ties in preferences are allowed, in particular showing that in certain circumstances some voters' best course of action is *not* to follow the Leader Rule.

This shows that under Approval Voting, not only convergence to equilibrium may not happen, but cycles can turn individually sound strategies into sub-optimal collective strategies, possibly leading to the worst possible outcome.

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. To define these perturbations, we introduced the notion of *Continuous Polling Dynamics* (CPD), which opens new directions of study. We illustrated the possibilities of CPD by looking at the tip of the iceberg of more general models for voters' response to polls, and observed that even simple, innocent-looking models easily produce chaotic dynamics.

There are many directions in which these results can be extended. First, one would be interested in studying other voting methods; the sincerity hypothesis will have to be relaxed in situation where it excludes strategic voting. One could also deepen the analysis of prevalence of bad cycle to more "cultures", i.e. models of voter's preferences. Then, the Continuous Polling Dynamics can be studied for various models of Voter's strategies, trying to better understand what features of a model lead to predictable dynamics, and which lead to chaos.

It would be interesting to see if in actual elections, some light could be cast by the Continuous Polling Dynamics formalism on the sequences of poll results; a particular difficulty would be to distinguish the effect of poll's broadcasting from the other events that may have influenced voters. A recent possible example of poll's effect is the first round of the 2017 French presidential election where two candidates, Benoît Hamon and Jean-Luc Mélenchon, competed for similar parts of the electorates and a positive reinforcement could explain part of the decline of Hamon in favor of Mélenchon.

Similar dynamics could be used in other many-players games beyond the study of voting systems. The Continuous Polling Dynamics could be adapted to model generalized strategies, taking for example into account psychological factors in Voter's response to new information.

References

- [AGP17] S. AIRIAU, U. GRANDI & F. S. PEROTTO – “Learning agents for iterative voting”, in *International Conference on Algorithmic Decision Theory*, Springer, 2017, p. 139–152. [1.5](#)
- [BCMP13] S. BRÂNZEI, I. CARAGIANNIS, J. MORGENSTERN & A. D. PROCACCIA – “How bad is selfish voting?”, in *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*, 2013, p. 138–144. [1.5](#)
- [BF07] S. BRAMS & P. C. FISHBURN – *Approval voting*, Springer Science & Business Media, 2007. [1.2](#), [1.4](#)

- [BHY14] C. BOWMAN, J. K. HODGE & A. YU – “The potential of iterative voting to solve the separability problem in referendum elections”, *Theory and decision* **77** (2014), no. 1, p. 111–124. [1.5](#)
- [Bla86] D. BLACK – *The theory of committees and elections*, Springer, 1986, Originally published 1958. [5](#)
- [BS03] S. J. BRAMS & M. R. SANVER – “Voter sovereignty and election outcomes”, <http://dev.gtcenter.org/Archive/Conf03/Sovereignty.pdf>, 2003. [1.3](#)
- [CPP04] S. CHOPRA, E. PACUIT & R. PARIKH – “Knowledge-theoretic properties of strategic voting”, in *European Workshop on Logics in Artificial Intelligence*, Springer, 2004, p. 18–30. [1.1](#)
- [Dur15] F. DURAND – “Vers des modes de scrutin moins manipulable”, Thèse, Université Pierre et Marie Curie, 2015. [1.2](#)
- [End13] U. ENDRISS – “Sincerity and manipulation under approval voting”, *Theory and Decision* **74** (2013), no. 3, p. 335–355. [1.2](#)
- [GHR19] U. GRANDI, D. HUGHES, F. ROSSI & A. SLINKO – “Gibbard–satterthwaite games for k-approval voting rules”, *Mathematical Social Sciences* **99** (2019), p. 24–35. [1.1](#)
- [Gib73] A. GIBBARD – “Manipulation of voting schemes: a general result”, *Econometrica: journal of the Econometric Society* (1973), p. 587–601. [1.2](#)
- [GLR⁺13] U. GRANDI, A. LOREGGIA, F. ROSSI, K. B. VENABLE & T. WALSH – “Restricted manipulation in iterative voting: Condorcet efficiency and borda score”, in *International Conference on Algorithmic Decision Theory*, Springer, 2013. [1.5](#)
- [KH95] A. KATOK & B. HASSELBLATT – *Introduction to the modern theory of dynamical systems*, Encyclopedia of Mathematics and its Applications, vol. 54, Cambridge University Press, Cambridge, 1995, With a supplementary chapter by Katok and Leonardo Mendoza. [7.2](#), [7.4](#)
- [KSLR17] A. KOOLYK, T. STRANGWAY, O. LEV & J. S. ROSENSCHEIN – “Convergence and quality of iterative voting under non-scoring rules”, in *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, 2017. [1.5](#)
- [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. [1.1](#), [1.3](#), [2.1](#), [3.2](#), [4.2](#), [5](#)
- [LR16] O. LEV & J. S. ROSENSCHEIN – “Convergence of iterative scoring rules”, *J. Artif. Intell. Res.* **57** (2016), p. 573–591. [1.5](#)

- [Mei15] R. MEIR – “Plurality voting under uncertainty”, in *Proceedings of the Twenty-Ninth AAI Conference on Artificial Intelligence*, AAAI Press, 2015, p. 2103–2109. [1.5](#)
- [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), ACM, 2014, p. 313–330. [1.5](#)
- [MPRJ17] R. MEIR, M. POLUKAROV, J. S. ROSENSCHEIN & N. R. JENNINGS – “Iterative voting and acyclic games”, *Artificial Intelligence* **252** (2017), p. 100–122. [1.5](#)
- [OLP⁺15] S. OBRAZTSOVA, O. LEV, M. POLUKAROV, Z. RABINOVICH & J. S. ROSENSCHEIN – “Farsighted voting dynamics”, in *1st Workshop on Algorithmic Game Theory at IJCAI*, 2015. [1.5](#)
- [OMP⁺15] S. OBRAZTSOVA, E. MARKAKIS, M. POLUKAROV, Z. RABINOVICH & N. R. JENNINGS – “On the convergence of iterative voting: how restrictive should restricted dynamics be?”, in *Proceedings of the Twenty-Ninth AAI Conference on Artificial Intelligence*, AAAI Press, 2015, p. 993–999. [1.5](#)
- [RE12] A. REIJNGOUD & U. ENDRISS – “Voter response to iterated poll information”, in *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, International Foundation for Autonomous Agents and Multiagent Systems, 2012, p. 635–644. [1.1](#), [1.5](#)
- [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 AAI Conference on Artificial Intelligence*, AAAI Press, 2015. [1.5](#)
- [Sat75] M. A. SATTERTHWAITTE – “Strategy-proofness and arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions”, *Journal of economic theory* **10** (1975), no. 2, p. 187–217. [1.2](#)
- [SDL06] F. D. SINOPOLI, B. DUTTA & J.-F. LASLIER – “Approval voting: three examples”, *Int. J. Game Theory* **35** (2006), p. 27–38. [1.1](#)
- [SS04] M. R. SERTEL & M. R. SANVER – “Strong equilibrium outcomes of voting games are the generalized condorcet winners”, *Social Choice and Welfare* **22** (2004), no. 2, p. 331–347. [1.3](#)
- [Wal82] P. WALTERS – *An introduction to ergodic theory*, Graduate Texts in Mathematics, vol. 79, Springer-Verlag, New York-Berlin, 1982. [7.2](#)