



HAL
open science

Cycles in synchronous iterative voting: general robustness and examples in Approval Voting

Benoît Kloeckner

► **To cite this version:**

Benoît Kloeckner. Cycles in synchronous iterative voting: general robustness and examples in Approval Voting. Social Choice and Welfare, inPress. hal-02434670v4

HAL Id: hal-02434670

<https://hal.science/hal-02434670v4>

Submitted on 9 Feb 2022

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.

Cycles in synchronous iterative voting: general robustness and examples in Approval Voting

Benoît R. Kloeckner,^{*†}

February 9, 2022

Abstract. We consider synchronous iterative voting, where voters are given the opportunity to strategically choose their ballots depending on the outcome deduced from the previous collective choices.

We propose two settings for synchronous iterative voting, one of classical flavor with a discrete space of states, and a more general continuous-space setting extending the first one. We give a general robustness result for cycles not relying on a tie-breaking rule, showing that they persist under small enough perturbations of the behavior of voters. Then we give examples in Approval Voting of electorates applying simple, sincere and consistent heuristics (namely Laslier’s Leader Rule or a modification of it) leading to cycles with bad outcomes, either not electing an existing Condorcet winner, or possibly electing a candidate ranked last by a majority of voters. Using the robustness result, it follows that those “bad cycles” persist even if only a (large enough) fraction of the electorate updates its choice of ballot at each iteration.

We complete these results with examples in other voting methods, including ranking methods satisfying the Condorcet criterion; an *in silico* experimental study of the rarity of preference profiles exhibiting bad cycles; and an example exhibiting chaotic behavior.

1 Introduction

Iterative Voting. In any voting system, the choice of a ballot by any voter is guided by her preferences between the candidates and the expected effect of each ballot she

^{*}ORCID 0000-0002-4966-7864, benoit.kloeckner@u-pec.fr

[†]Univ Paris Est Creteil, CNRS, LAMA, F-94010 Creteil, France – Univ Gustave Eiffel, LAMA, F-77447 Marne-la-Vallée, France

can cast. To estimate the effect of any ballot, one needs to know the ballots cast by the other voters, or at least to have some information on them. Assuming perfect information is neither realistic nor theoretically useful: in large electorates, changing one's ballot rarely changes 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. Given a model of voter's response to new information, one obtains a new set of equilibria, which we call *dynamical equilibria*: situations when the intents of voters stay unchanged after the outcome they induce is broadcast. In this article, the single word *equilibrium* will only be used to denote dynamical equilibria, never Nash equilibria. Natural questions are thus whether such dynamical equilibrium exist; whether such equilibrium is unique; what properties such equilibria have, in particular which candidates can be elected at equilibrium; and which other sets of intents converge to an equilibrium after successive adjustments. The field of *iterative voting* is the study of these questions, under various modeling assumptions.

Synchronized and asynchronous iterative voting. By far the most studied version of iterative voting is *asynchronous*, i.e. only some voters (most commonly a single one) adjust their intents given the available information, and information is updated before some other voter(s) make their adjustments, etc. While this model has both a theoretical interest and practical applications to some collective decision processes, it is ill-suited to model large-scale political election, where the number of times when the information is updated (poll publications) is much lesser than the number of voters.

In contrast, we consider in this article *synchronous* iterative voting, where all voters are given the information at time t , and simultaneously adjust their intents to build the new situation at time $t + 1$. It is important to observe that what makes iterative voting synchronous 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.

Synchronous iterative voting can be thought of as a model for political elections, where information is broadcasted through polls; we will thus name our dynamical model *Polling*

Dynamics; similar points of view have 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).

However this is not our sole motivation: considering synchronized iterative voting is also interesting from a theoretical perspective, when trying to define a “best” ballot for each voter. Indeed, 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 heuristic, which dictates for each possible expected outcome which ballot to cast, and we inductively apply these heuristics and update the expected outcome. At a fixed point of the Polling Dynamics, which we call a *dynamical equilibrium*, anticipating each other’s strategy result in the same situation, which is therefore stable. 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 [GHS19].

We shall be especially interested in 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. Approval Voting has the very interesting feature that several sincere ballots coexist for each order of preference, so that there exist heuristics that are both sincere, simple, and non-trivial; let us discuss this in more detail.

On sincerity, strategic voting and straightforwardness 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 others’ plans or that are too contrived cannot be applied uniformly, creating asymmetries between voters. Even in the absence of manipulation or strategic voting, sincere ballots 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 is thus *straightforwardness* [Gib73], i.e. the existence of strategies that yield an optimal outcome whatever choices are made by the other voters. Since straightforwardness cannot be hoped in general [Gib73, Sat75], one can try to determine whether certain systems often offer weak forms of straightforwardness (see again [Dur15]).

In Approval Voting, one can say that a ballot is *sincere* whenever all candidates present on the ballot are preferred to all candidates not present on it [BF07]. As soon as there are more than two candidates each voter has several sincere ballots, corresponding to the

various points 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 ballots. In other words, sincerity in the above sense does not imply straightforwardness. One of our goals is to give examples showing how very far from straightforward Approval Voting can be even when restricting to sincere ballots.

Is Approval Voting a Condorcet system “in practice”? While Approval Voting is known not to be a Condorcet system, it has been argued that it might be close to Condorcet “in practice”. We can e.g. read in [BF10]: “In particular, Condorcet candidates [...] almost always win under AV”, and some elections where this assertion can be checked are mentioned as evidence.

Several theoretical results could also be seen as providing a basis for the above claim. First, even though many other voting systems also satisfy these properties, let us mention that Brams and Sanver [BS03] showed that when a Condorcet winner exists, her election is a strong Nash equilibrium; and in [SS04] it is shown that only Condorcet Winners are elected at strong Nash equilibriums. Strong Nash equilibriums are said by Brams and Sanver to be “globally stable”, but one point not considered in either works above is whether they should be expected to be reached in practice.

A stronger argument originates in Laslier [Las09], where it is 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 heuristic, the “Leader Rule” (LR). 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.¹

Description of the main results. To introduce the Polling Dynamics, we propose both a classical-flavored *discrete space* setting, and an enlarged *continuous space* setting that allows to continuously perturb a situation (i.e. a given electorate with fixed heuristics), for example in order to assume that at each step of the dynamics, a small fraction of voters fail to adjust their ballots. This makes it possible to discuss how dynamical phenomena change under such perturbations.

Our first main result (Theorem A in Section 3) is a quite general robustness result, showing under a mild assumption that cycles or equilibriums in the discrete space setting persist under perturbation inside the continuous space setting. All the examples we will give benefit this robustness. Our main contribution here is in the framework for Polling Dynamics, the proof of Theorem A being rather straightforward once the definitions are set up.

Among the possible perturbations, one can model the situation where instead of having all voters change their ballots when informed of the last poll, a (small enough) fraction

¹A 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 less rigorous than Laslier’s, since the equilibrium assumption is kept implicit.

keeps their ballot as they were. This particular case brings us closer to asynchronous voting, but only to some extent since we need to stay close to the discrete, synchronous model.

The next main results are examples showing that in Approval Voting, the Polling Dynamics can exhibit a “bad cycle” where a disliked candidate can be elected, even in the presence of a Condorcet winner (Theorems [B](#) and [C](#) in Section 4). The main point shall not be the mere existence of cycles, a rather unsurprising phenomenon,² 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 heuristic at some iteration of the process; in our examples the assumed strategic behaviors is consistent, i.e. constant in time; they are also sincere, simple, and strategically sound (the first example uses the Leader Rule, the second a slight modification where voters refuse to vote for several of their least-preferred candidates in any circumstance). Two obvious limitations to the relevance of these existence results are that such “bad cycles” could be rare, and that it is unrealistic to assume all voters to update their choice of ballots at each poll publication. The former limitation will be addressed by numerical experiments (which partly confirm it, and partly moderate it, see below); the latter is addressed by our robustness result. Indeed, Theorem [A](#) shows that if instead of assuming that all voters update their choice of ballot at each iteration, we only assume that a large enough proportion of them do, then the same cycle of winners still occurs.

Regarding whether Approval Voting should often elects the Condorcet winner, our examples cast a doubt. Laslier’s results mentioned above mean that in every situation where a Condorcet winner exists but is not elected, voters (individually) have a strategic interest in changing their ballots. Our examples expose that this does *not* mean that following this heuristic shall lead to the election of a Condorcet winner: it may also lead to a periodic pattern where the Condorcet winner is absent. More generally, even simple heuristics make the availability of information a possible impediment to reaching an equilibrium and/or to elect an existing Condorcet winner.

This new issue with the Approval Voting system complement in particular the examples provided in [\[SDL06\]](#).

We also provide a few examples in other voting systems in Section 5. Together with the Approval Voting examples, they show how extreme an impact polling can have 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 need not induce synchronization nor loss of memory. Even with perfect unrigged polls, the sheer *number* of polls (e.g. its

²An example of cycle under the Leader Rule is given as exercise 8.4.4 in [\[Mei18\]](#), but there are no Condorcet winner.

parity) can decide the outcome of the election.

It could be expected that such bad cycles are in practice rare. In Section 4.3 we explore numerically this question, for Approval Voting, under various cultures modeling the construction of random voters' preferences. The main take-away is that bad cycles are rare when voters apply the Leader Rule, but can be more common, appearing in more than 15% of preference profiles, in its modification where voters give themselves a limit to the candidates they may approve of. Most striking is that the culture where bad cycles are rarest for the Leader Rule is the one where they are most common for its modification.

The continuous-state setting enables one to model a broad variety of voter's collective behavior. In Section 6, we consider an example where voters are reluctant to approve of a candidate who is ranked low in their preference order, but may do so if a very close-call makes this move susceptible to improve the outcome. This arbitration can surprisingly make the Continuous-State Polling Dynamics a *chaotic dynamical system*, making the prediction of the next winners from the observation of a sequence of previous expected winners practically impossible.

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; a more detailed review can be found in [Mei17]. The present work distinguishes itself from previous ones by considering synchronous iterative voting, previously seldomly studied: most commonly *iterative voting* is used to mean what we named here *asynchronous iterative voting*. Another divide which mostly parallels the previous one is between works considering at each step a set of possible moves by voters (e.g. best-responses, quite commonly considered in asynchronous voting) and, as we do here, deterministic heuristics. The reader can find an overview on strategic voting in [Mei18].

In the case of Plurality Voting, discussed in slightly more details below, very general convergence results have been obtained in [MLR14, Mei15, MPRJ17]. Many other voting rules — not Approval — have been considered in [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 to some not necessarily best or better responses, have been studied in [OMP⁺15] and [GLR⁺13].

Non-myopic strategies have been considered for example in [BHY14, AGP17, OLP⁺15], showing notably that voters learning from past information produce relatively good outcomes. Specific behaviors of voters, assumed to be subject to either *truth bias* or *lazy bias*, are studied in [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 equilibria 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 equilibria 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. This is the kind of modeling assumptions that our framework is made to support.

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 for Approval Voting presented here, and to François Durand for introducing me to the Social choice theory, for many discussions on Voting Systems, and for many suggestions that helped me improve this article. It also benefited from relevant comments provided by Adrien Fauré, Jean-François Laslier, Reshef Meir and anonymous referees who I warmly thank.

2 Polling dynamics with discrete or continuous space

This section mostly consists of notations and definitions; its main novelty is to propose a continuous space framework for iterative voting, which allows to model much more varied situations than the usual discrete space versions. It models a situation where voters are grouped by type, e.g. according to their preference order, but where not all voters of a given type will react in the same way to an information update. For example, we can assume that a fraction of voters of each type is not aware of an information update, and will thus keep their previous intended ballot for the next iteration. In another direction, if the win is a close call, some voter's type may strategically adjust their ballots accordingly; but how close the call should be for them to bother changing their intent may vary from individual to individual, and our continuous space framework will make it possible to model this.

Examples of uses of the framework developed here are given below, first for Approval voting and in Section 5 for several other voting systems.

2.1 Common notations

We start with the notation common to both the discrete space and continuous space frameworks; when possible, we use the notation from [BCE⁺16].

The set of *voter types* is assumed to be finite, of cardinal n , and denoted by N . It will be convenient to denote voter types by positive integers ($N = \{1, \dots, n\}$) in the theoretical part, but by upper-case letters from the end of the alphabet ($N = \{Z, Y, X, \dots\}$) in examples. Some of these letters will be used for maps below, but the context will prevent any confusion. In many examples below, voter types will coincide with preferences; however distinguishing gives more flexibility in the modelling, e.g. we can handle the case of different types of voters having the same preferences but using different heuristics.

Let A be a finite set of *candidates* (or *alternatives*). Its cardinal is denoted by m and assumed to be at least 2 (most usually at least 3). Candidates will be denoted by lower-case letters from the beginning of the alphabet ($A = \{a, b, c, \dots\}$). The set of weak orderings on A , i.e. rankings of candidates allowing ties, is denoted by $\mathcal{R}(A)$. The set of linear ordering on A , i.e. rankings of candidates without ties, is denoted by $\mathcal{L}(A)$.

Preferences of each voter type will be modeled by an element of $\mathcal{R}(A)$; when disallowing ties, we will restrict to its subset $\mathcal{L}(A)$. A *preference profile* is a pair (P, w) where:

- $P = (\succsim_1, \dots, \succsim_n) \in \mathcal{R}(A)^n$ is a family of preferences, one for each voter type; the corresponding strict preferences are denoted by \succ_i ,
- $w = (w_1, \dots, w_n) \in \mathbb{R}_+^n$ is a family of *weights* $w_i > 0$, which represents the number of voters of each type. A usual situation is to take $w_i = 1$ for all i (or a positive integer, when grouping several voters in each type), but we allow more general weights for the continuous space framework.

A preference profile is represented by an array, each column of which represents a voter type, the column head carrying the weight (we shall often indicate the voter type on top for clarity) and the column listing the candidates in decreasing order of preference. With 3 candidates a, b, c an example of preference profile is:

1	2	3
100	101	102
a	b	c
b	a	b
c	c	a

meaning that $n = 3$ and e.g. that $a \succ_1 b$ but $b \succ_2 a$, that there are 102 voters of type 3, etc. For concision, an element of $\mathcal{L}(A)$ is denoted by listing the candidates in decreasing order of preference, e.g. $\succ := acb$ means $a \succ c \succ b$. Elements of $\mathcal{R}(A)$ are denoted similarly with tied elements between parentheses.

Let \mathcal{B} be a finite set of *ballots* that can be cast by voters. For example \mathcal{B} can be $\mathcal{L}(A)$, or for Approval Voting the set $\mathcal{P}(A)$ of subsets of A . An Approval Voting ballot $B \in \mathcal{P}(A)$ is said to be *sincere* for voters of type i whenever

$$\forall a, b \in A: (b \in B \text{ and } a \succ_i b) \implies a \in B$$

i.e. when every candidate strictly preferred to any approved candidate is also approved.

Let \mathcal{O} be a set representing the possible *outcomes* of the election. We can let $\mathcal{O} = A$ if we are only interested in the winner, but it is convenient to take a larger set \mathcal{O} to include all information that can be made available and used by voters to decide which ballot to cast. For example, for Approval Voting we can take $\mathcal{O} = [0, 1]^A$, an element $r = (r_a, r_b, \dots)$ of \mathcal{O} giving the proportion of voters approving of each candidate. We let $W : \mathcal{O} \rightarrow A$ be the function mapping a result to the corresponding winner (we shall always assume a tie-breaking rule, e.g. breaking ties in favor of the earlier candidate in the alphabetical order).

2.2 Polling dynamics with discrete space

We start with the usual case where in any given circumstance all voters of the same type cast the same ballot, which we call here the *discrete space* framework. A *ballot profile* is thus an element of \mathcal{B}^n and represents the ballots cast by each voter type.

We consider a *social choice function*, also called *voting rule* $f : \mathcal{B}^n \rightarrow A$ and assume the choice of \mathcal{O} and W makes it possible to decompose it as $f = W \circ g$ where $g : \mathcal{B}^n \rightarrow \mathcal{O}$ is called the *information function*. Observe that by assuming a tie-breaking rule for W , we only consider *resolute* social choice functions. For example, with the above choice of outcomes, the information function is given by

$$g(B_1, \dots, B_n) = \left(\frac{1}{\sum_{i \in N} w_i} \sum_{i \in N | \alpha \in B_i} w_i \right)_{\alpha \in A}$$

and the social choice function maps (B_1, \dots, B_n) to the candidates with the highest approval rating, ties broke by alphabetical order.

Each voter type i is assumed to have a *heuristic* $\sigma_i : \mathcal{O} \rightarrow \mathcal{B}$, representing the way voters of this type choose their next ballot given the information of the outcomes (which is itself determined from the previous ballot profile). Formally, the preferences of voter types do not appear in the Polling Dynamics to be defined in the next paragraph. Most usually, they do appear indirectly, through heuristics: more often than not, $\sigma_i(r)$ take the form $\sigma(\succsim_i, r)$ where σ is a map $\mathcal{R}(A) \times \mathcal{O} \rightarrow \mathcal{B}$ (or $\mathcal{L}(A) \times \mathcal{O} \rightarrow \mathcal{B}$ if ties in preferences are not allowed). An example of such a map σ is the Leader Rule of Laslier, which we mentioned earlier and recall below.

We obtain a *Polling Dynamics* on the set of ballot profiles

$$\begin{aligned} \varphi : \quad & \mathcal{B}^n \rightarrow \mathcal{B}^n \\ (B_1, \dots, B_n) & \mapsto (\sigma_1(g(B_1, \dots, B_n)), \dots, \sigma_n(g(B_1, \dots, B_n))) \end{aligned}$$

obtained by determining the outcome according to the ballot profile in argument and then applying each voter's heuristic to obtain a new ballot profile.

Alternatively, we can take as space the set of outcomes and consider the *shifted Polling Dynamics*

$$\begin{aligned}\psi &: \mathcal{O} \rightarrow \mathcal{O} \\ r &\mapsto g(\sigma_1(r), \dots, \sigma_n(r))\end{aligned}$$

mapping an outcome to a new one, after applying each voter's heuristic then determining the resulting outcome.

These two maps are strongly related: $\psi \circ g = g \circ \varphi$ (in dynamical systems theory we say they are *semi-conjugated*); φ seems to carry more information, but their dynamical study are actually equivalent because we assumed in the model that heuristics only depend on the outcome, not on the full details of the ballots.

By a *dynamical equilibrium* we mean a fixed point of either ψ or ϕ , i.e. either an outcome r such that $\psi(r) = r$, or a ballot profile (B_1, \dots, B_n) such that $\varphi(B_1, \dots, B_n) = (B_1, \dots, B_n)$. There is a one-to-one identification between these two points of view: when r is fixed by ψ , then $(\sigma_1(r), \dots, \sigma_n(r))$ is fixed by φ and sent back to r by g ; and when (B_1, \dots, B_n) is fixed by φ , $g(B_1, \dots, B_n)$ is fixed by ψ and sent back to (B_1, \dots, B_n) by the heuristics.

The Leader Rule and its modification. The Leader Rule is an example of heuristic for Approval Voting depending on tie-free preferences $\succsim \in \mathcal{L}(A)$; assuming the expected winner $\omega(r)$ and the expected runner-up $\rho(r)$ can both be deduced from the outcome r , it is defined by

$$\text{LR}(\succsim, r) = \{\alpha \in A \mid \alpha \succ \omega(r) \text{ or } \alpha = \omega(r) \succ \rho(r)\};$$

in words, all candidates preferred to the expected winner are approved, and the expected winner is approved if and only if she is preferred to the expected runner-up. When we say that voters (assumed to have tie-free preferences) apply the Leader Rule, we mean that their heuristics take the form $\sigma_i(r) = \text{LR}(\succsim_i, r)$.

When ties are allowed, i.e. $\succsim \in \mathcal{R}(A)$, there are several possible ways to extend the Leader Rule. One could use the very same formula as above, but it would have the consequence that when the outcome is the ranking $abcd$ and the preferences are $(ab)(cd)$, the resulting ballot is blank. But in this situation, casting the ballot $\{a, b\}$ may prevent a worst candidate to be elected, and is always preferable to a blank ballot. We will thus consider the following *Modified Leader Rule* for preferences with ties by

$$\text{MLR}(\succsim, r) = \{\alpha \in A \mid \alpha \succ \omega(r) \text{ or } \alpha \simeq \omega(r) \succ \rho(r) \text{ or } (\alpha \simeq \omega(r) \text{ and } \nexists \beta \succ \alpha)\}$$

in particular, compared to the above, when the expected winner is tied for top in the preferences, then she and all those tied with her are approved. Again, when we say that voters follow the Modified Leader Rule we mean that $\sigma_i(r) = \text{MLR}(\succsim_i, r)$.

2.3 Polling dynamics with continuous space

We now turn to Polling Dynamics with continuous space (CS), the goal of which is to allow more flexibility in our assumption on voters' heuristics. In particular, we want to be able to consider a continuum of behaviors in each voter type, for example accounting for variable levels of bias, be it for example truth-bias (favoring sincere ballots) or lazy-bias (aversion to change one's ballot in view of new information). Example 4.8 illustrates this framework.

The discrete space of ballot profiles \mathcal{B}^n thus has to be replaced by a more complicated object, encoding the proportion of each voter type casting each ballot. To this effect we use the *simplex* $\Delta(X)$ over a finite set X , defined as

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

The *full continuous space* is then $\Delta(\mathcal{B})^n$; an element of this space is a doubly indexed family $(p_B^i)_{i \in N, B \in \mathcal{B}}$ where p_B^i is the proportion of ballots B cast among voters of type i , and is called a *CS ballot profile*.

We now define a *CS information function* as a map $G : \Delta(\mathcal{B})^n \rightarrow \mathcal{O}$, with the *associated CS social choice function* $F = W \circ G : \Delta(\mathcal{B})^n \rightarrow A$. For example, with the above choice of outcomes, for Approval Voting the CS information function is given by

$$G((p_B^i)_{i,B}) = \left(\frac{1}{\sum_{i \in N} w_i} \sum_{\substack{i \in N \\ B \in \mathcal{B} \mid \alpha \in B}} p_B^i w_i \right)_{\alpha \in A}$$

where $p_B^i w_i$ represents the total number of voters of type i casting the ballot B .

The full continuous space often has an unnecessarily large dimension, since under most heuristics voters will only cast ballots among a small subset \mathcal{B}_i of \mathcal{B} (one subset for each voter type; for example, for Approval Voting we could restrict to ballots that are sincere with respect to the type's preferences). We thus define a more convenient *continuous space* by

$$\mathcal{P} = \prod_{i \in N} \Delta(\mathcal{B}_i) \subset \Delta(\mathcal{B})^n$$

the elements of which are *admissible CS ballot profiles*, or *states*. We still denote by G and F the restrictions of the CS information function and social choice function to this space. Observe that the full CS appears as a special case by taking $\mathcal{B}_i = \mathcal{B}$ for all i .

A *CS Polling Dynamics* is simply a map $\Phi : \mathcal{P} \rightarrow \mathcal{P}$. In general, it is not so much the dynamics of Φ that will be of interest, but the sequences of outcomes, which can be recovered as $G(\Phi^k(p_0))$ where k is the number of iterations and $p_0 = (p_B^i)_{i,B} \in \mathcal{P}$ is the initial ballots intended to be cast by the voters.

This very broad definition of a CS Polling Dynamics allows for much modeling flexibility, but the main interest is when such a map is deduced from some sort of heuristics of voters. We will not need this consideration in the robustness result which is very general, but let us give some definitions to show how this can be done.

We will assume that the choice of voters is based only on the predicted outcome and the ballots previously cast by voters of their type (this is enough to permit a fraction of voters not to adjust their choice, e.g. because they are not aware of the last poll). A *CS heuristic* for voters of type i is thus a map $\Sigma_i : \mathcal{O} \times \Delta(\mathcal{B}_i) \rightarrow \Delta(\mathcal{B}_i)$, and given a CS heuristic for each voter type the corresponding CS Polling Dynamics is given for all state $p = (p_B^i)_{i,B} \in \mathcal{P}$, setting $p^i = (p_B^i)_B$, by

$$\Phi(p) = \left(\Sigma_1(G(p), p^1), \Sigma_2(G(p), p^2), \dots, \Sigma_n(G(p), p^n) \right).$$

3 Robustness of tie-free cycles

In this section we want to prove that cycles of a discrete space Polling Dynamics that do not rely on the tie-breaking rule are robust, i.e. they persist under small enough perturbation. “Small enough” is a void concept in the discrete space setting, and we use CS Polling Dynamics to model the perturbations. The first step is to see that discrete space Polling Dynamics can always be realized as particular cases of CS Polling Dynamics.

3.1 Embedding the discrete space Polling Dynamics into the continuous space.

Observe that \mathcal{B}^n embeds naturally into the full CS, by identifying (B_1, \dots, B_n) with the element $\pi(B_1, \dots, B_n) = (p_B^i)_{i,B} \in \Delta(\mathcal{B})^n$ defined by

$$p_B^i = \begin{cases} 1 & \text{when } B = B_i \\ 0 & \text{otherwise} \end{cases}$$

Such an elements of the full CS, taking only the values 0 and 1, is called an *extreme state*. We shall say that the CS information function G *extends* the information function g whenever $G \circ \pi = g$, i.e. for all $(B_1, \dots, B_n) \in \mathcal{B}^n$, $G(\pi(B_1, \dots, B_n)) = g(B_1, \dots, B_n)$. Without this assumption we would not be encoding the same voting method in both settings.

Consider fixed the (discrete space) heuristics σ_i and the corresponding Polling Dynamics φ , and define the *associated* CS Polling Dynamics by

$$\Phi_0(p) = \pi(\sigma_1(G(p)), \dots, \sigma_n(G(p))),$$

i.e. starting from a CS ballot profile $p \in \Delta(\mathcal{B})^n$ we compute its outcome $G(p)$, then apply the heuristics to get new ballot choices for each voter type, then embed this into the full CS with the map π to obtain a new CS ballot profile. Note that Φ_0 is well-defined on the full CS, but all its values are extreme states.

The hypotheses we introduced in this section are precisely what needs to be assumed to ensure that the continuous space framework extends the discrete space situation: Φ_0 produces the same sequences of outcome than φ when starting at an extreme state.

Lemma 3.1. *Assume that G extends g and that the heuristics only pick admissible ballots, i.e. $\sigma_i(r) \in \mathcal{B}_i$ for all voter type $i \in N$ and all outcome $r \in \mathcal{O}$. Then the associated CS Polling Dynamics takes its values in \mathcal{P} , and for all (B_1, \dots, B_n) the sequences of outcomes*

$$g(\varphi^k(B_1, \dots, B_n)) \quad \text{and} \quad G(\Phi_0^k(\pi(B_1, \dots, B_n)))$$

are the same.

Proof. Since heuristics only pick admissible ballots, for all $p \in \Delta(\mathcal{B})^n$ we have

$$(\sigma_1(G(p)), \dots, \sigma_n(G(p))) \in \prod_{i \in N} \mathcal{B}^i$$

so that $\Phi_0(p) \in \mathcal{P}$.

Since G extends g , for all (B_1, \dots, B_n) we have

$$\begin{aligned} \Phi_0(\pi(B_1, \dots, B_n)) &= \pi(\sigma_1(G(\pi(B_1, \dots, B_n))), \dots, \sigma_n(G(\pi(B_1, \dots, B_n)))) \\ &= \pi(\sigma_1(g(B_1, \dots, B_n)), \dots, \sigma_n(g(B_1, \dots, B_n))) \\ &= \pi(\varphi(B_1, \dots, B_n)) \end{aligned}$$

By induction, we deduce that $\Phi_0^k \circ \pi = \pi \circ \varphi^k$ for all positive integer k . Using again that G extends g , we get $G \circ \Phi_0^k \circ \pi = G \circ \pi \circ \varphi^k = g \circ \varphi^k$, as desired. \square

3.2 Perturbations and robustness.

The second step is to define what it means for two CS Polling Dynamics to be close one to the other. We consider the metric induced on the full CS and on \mathcal{P} by the supremum norm $\|\cdot\|$, i.e. given two CS ballot profiles $p = (p_B^i)_{i,B}$ and $\bar{p} = (\bar{p}_B^i)_{i,B}$ we set

$$\|p - \bar{p}\| = \sup_{i \in N, B \in \mathcal{B}} |p_B^i - \bar{p}_B^i|.$$

This induces the usual topology on the full CS and, by restriction, on \mathcal{P} . We denote by $\bar{B}(p, \varepsilon)$ the closed ball of radius $\varepsilon \geq 0$ and center $p \in \mathcal{P}$ with respect to this metric.

We then consider the uniform metric, defined for every pair Φ_1, Φ_2 of CS Polling Dynamics by

$$D(\Phi_1, \Phi_2) := \sup_{p \in \mathcal{P}} \|\Phi_1(p) - \Phi_2(p)\|.$$

Definition 3.2. We say that a ballot profile (B_1, \dots, B_n) is *tie-free* (implicitly, with respect to heuristics $(\sigma_i)_i$, an information function g and an extension G) whenever there exist a neighborhood U of $\pi(B_1, \dots, B_n)$ in \mathcal{P} such that for all $p \in U$:

$$W \circ G(p) = W \circ g(B_1, \dots, B_n) \quad \text{and} \quad \sigma_i(G(p)) = \sigma_i(g(B_1, \dots, B_n)) \quad \forall i \in N.$$

A ballot profile (B_1^0, \dots, B_n^0) belongs to a *tie-free cycle* for the Polling Dynamics φ whenever $\varphi^k(B_1^0, \dots, B_n^0) = (B_1^0, \dots, B_n^0)$ for some positive integer k (called a *period* of the cycle) and the successive ballot profiles $(B_1^j, \dots, B_n^j) = \varphi^j(B_1^0, \dots, B_n^0)$ are all tie-free.

Let us explain the rationale behind this definition. First, the condition $W \circ G(p) = W \circ g(B_1, \dots, B_n)$ means that the winner does not change if the considered CS ballot profile is close enough to $\pi(B_1, \dots, B_n)$. Second, all considered heuristics $\sigma : \mathcal{O} \rightarrow \mathcal{B}^n$ and their composition with the CS information function $\sigma \circ G : \mathcal{P} \rightarrow \mathcal{B}^n$ will be piecewise continuous, and thus piecewise constant (since they take value in a finite set). A state p at which $\sigma \circ G$ is not continuous means that either a small change in the state can change the outcomes radically, or a small change in the outcome can change the resulting ballot profile for this heuristic. This correspond to breaking a tie, either in the information function or in the heuristic. For example in Approval Voting, when applying the Leader Rule with the above continuous choice of outcomes $\mathcal{O} = [0, 1]^A$: when the first two candidates are tied, the tie-breaking rule embedded in W is taken into account to determine the expected winner; and when several candidates are tied for runner-up, a tie-breaking rule must implicitly be written in the heuristic itself.

Theorem A. *Assume that G extends g and that the heuristics $\sigma_1, \dots, \sigma_n$ only pick admissible ballots. If the Polling Dynamics φ has a tie-free cycle of period k , an element of which is denoted by (B_1^0, \dots, B_n^0) , then there exist $\varepsilon > 0$ with the following property: for all CS Polling Dynamics Φ such that $D(\Phi, \Phi_0) \leq \varepsilon$, for all $p^0 \in \mathcal{P}$ such that $\|p^0 - \pi(B_1^0, \dots, B_n^0)\| \leq \varepsilon$, for all positive integer j ,*

$$W \circ G(\Phi^j(p^0)) = W \circ g(\varphi^j(B_1^0, \dots, B_n^0)).$$

If moreover Φ is continuous, then there exist $p^0 \in \mathcal{P}$ such that $\|p^0 - \pi(B_1^0, \dots, B_n^0)\| \leq \varepsilon$ and $\Phi^k(p^0) = p^0$.

In words, if the discrete space dynamics has a tie-free cycle, then perturbing the dynamics and starting point in the CS setting does not change the (periodic) sequence of winners; and if the perturbed CS Polling Dynamics is continuous, then this sequence of winners is furthermore realized by a cycle in \mathcal{P} , which is a perturbation of the original cycle. We shall see that for some natural examples, the cycle moreover attracts a neighborhood of p^0 (see Examples 3.4 and 4.8, which are easily generalized).

Note that Theorem A makes no rationality assumption: it applies not only to best-response heuristics but allow arbitrary information-based heuristics, for example the pragmatist response policy of [RE12].

Proof. For each positive integer j , set $B^j := (B_1^j, \dots, B_n^j) := \varphi^j(B_1^0, \dots, B_n^0)$. Since B^j is tie-free, $\pi(B^j)$ admits a neighborhood U^j (and we can choose the sequence $(U^j)_j$ to be k -periodic) such that for all $p \in U^j$,

$$W \circ G(p) = W \circ g(B^j) \quad \text{and} \quad \sigma_i(G(p)) = \sigma_i(g(B^j)) \quad \forall i \in N.$$

By definition of a neighborhood, for each j there exist $\varepsilon^j > 0$ such that $\bar{B}(\pi(B^j), \varepsilon^j) \subset U^j$; let $\varepsilon = \min(\varepsilon^0, \varepsilon^1, \dots, \varepsilon^{k-1})$.

The second tie-free condition $\sigma_i(G(p)) = \sigma_i(g(B^j))$ for all $i \in N$ and the definition of Φ_0 ensures that for all $p \in U^j$, $\Phi_0(p) = \pi(B^{j+1})$. Let Φ be a CS Polling Dynamics such that $D(\Phi, \Phi_0) \leq \varepsilon$ and let $p^0 \in \mathcal{P}$ such that $\|p^0 - \pi(B_1^0, \dots, B_n^0)\| \leq \varepsilon$. Then

$\Phi_0(p^0) = \pi(B^1)$ and $\Phi(p^0) \in \bar{B}(\pi(B^1), \varepsilon) \subset U^1$. The same reasoning and an induction ensures that for all j , the state $p^j := \Phi^j(p^0)$ lies in $\bar{B}(\pi(B^j), \varepsilon) \subset U^j$. The first tie-free condition $W \circ G(p^j) = W \circ g(B^j)$ then gives the desired conclusion.

Assume further that Φ is continuous. Observe that \mathcal{P} is a polyhedron of \mathbb{R}^d for some d , and that balls $\bar{B}(s, \varepsilon)$ are thus convex sets, in particular homeomorphic to closed balls. For all $p \in \bar{B}(\pi(B^0), \varepsilon)$, we have $\Phi^{k-1}(p) \in U_{k-1}$ and thus $\Phi^k(p) \in \bar{B}(\pi(B^0), \varepsilon)$: Φ^k is a continuous map sending the topological ball $\bar{B}(\pi(B^0), \varepsilon)$ into itself. Brouwer's fixed point theorem ensures that there exist $p^0 \in \bar{B}(\pi(B^0), \varepsilon)$ such that $\Phi^k(p^0) = p^0$. \square

Remark 3.3. Theorem A applies to dynamical equilibria, by taking $k = 1$.

Example 3.4. Consider Plurality voting with three candidates $A = \{a, b, c\}$, three voter types $N = \{X, Y, Z\}$ and preference profile

X	Y	Z
4	2	3
a	b	c
bc	ac	b
	a	

since we use Plurality, $\mathcal{B} = A = \{a, b, c\}$; assume voters only consider the expected winner to choose how to vote, so that also $\mathcal{O} = A$ and W is the identity map. The voting rule f and the information function g coincide; both take the triple of ballots $(v_X, v_Y, v_Z) \in \{a, b, c\}^3$ and return the candidate with the most votes.

For the discrete polling dynamics, we consider the following heuristics: $\sigma_X(\alpha) = a$ and $\sigma_Y(\alpha) = b$ for all $\alpha \in \{a, b, c\}$ (voters of type X and Y always vote for their favorite); $\sigma_Z(a) = b$ and $\sigma_Z(\alpha) = c$ for $\alpha \in \{b, c\}$ (voters of type Z vote for their favorite unless their least favorite threatens to win, in which case they settle for their second choice). The shifted polling dynamics is then given by

$$\psi(a) = b, \quad \psi(b) = a, \quad \psi(c) = a$$

with a 2-cycle (a, b) . Let us now consider a natural perturbation of this dynamics. The full continuous space is 6-dimensional: we need two numbers for each voting types, representing the proportions voting for two of the candidates (the third proportion being deduced from the first two). The CS information function G compounds the votes in favor of each candidate, and returns the candidate with the most votes (tied broke by alphabetical order, say).

Assume $\mathcal{B}_X = \{a\}$, $\mathcal{B}_Y = \{b\}$ and $\mathcal{B}_Z = \{b, c\}$ (i.e. Voters of type X and Y still always vote for their favorite candidate, voters of type Z can vote for either of their two preferred candidates). Then the continuous space \mathcal{P} can be identified with $[0, 1]$, a number $z \in [0, 1]$ representing the proportion of voters of type Z voting for c , and the restriction of G to \mathcal{P} is then given by $G(z) = b$ when $z < \frac{1}{3}$ and $G(z) = a$ when $z \geq \frac{1}{3}$ (indeed a always receives 4 votes, while c receives $3z \leq 3$ and b receives $5 - 3z$). We

consider a parameter $t \in [0, 1]$ representing the reluctance of voters to change their votes and define a corresponding CS heuristic by

$$\Sigma_Z^t(\alpha, z) = \begin{cases} 1 - t(1 - z) & \text{when } \alpha \neq a \\ tz & \text{when } \alpha = a \end{cases}$$

(a proportion t of voters of type Z who should switch votes under the discrete heuristic σ_Z do not do so under Σ_Z^t). We denote by φ the discrete Polling Dynamics and by Φ_t the CS Polling dynamics driven by Σ_Z^t . The CS Polling Dynamics Φ_0 associated to φ is obtained as follows: given $z \in [0, 1]$, we compute the expected winner $G(z)$ and determine the behavior of voters with the discrete heuristics: when $z < \frac{1}{3}$, all voters of type Z shall vote for c in the next round so that $\Phi_0(z) = 1$ while when $z \geq \frac{1}{3}$, they all vote for b so that $\Phi_0(z) = 0$. Every starting point z is immediately mapped to an extreme state, as always, and in this particular case they all are attracted to the cycle (a, b) . For $t \in [0, 1]$, we get $\Phi_t(z) = 1 - t + tz$ when $z < \frac{1}{3}$ and $\Phi_t(z) = tz$ when $z \geq \frac{1}{3}$ (in particular the $t = 0$ matches Φ_0); the graph of Φ_t consist in two lines of slope t , as pictured in Figure 1. The map $t \mapsto \Phi_t$ is continuous in the uniform topology, so that Theorem A ensures that for t small enough and z_0 close enough to 0 or 1, the iteration of Φ_t starting at z_0 alternates between the intervals $[0, \frac{1}{3})$ and $[\frac{1}{3}, 1]$, producing the same cycle of outcomes than the discrete Polling Dynamics.

Here we can compute exactly how large a perturbation we can afford. For all $t \in [0, \frac{1}{2})$, there is a 2-cycle consisting of the points $\frac{t}{1+t} < \frac{1}{3}$ and $\frac{1}{1+t} > \frac{2}{3}$. Since the slope of the graph of Φ_t is less than one, this 2-cycle attracts (exponentially fast) all nearby orbits. At $t = \frac{1}{2}$, one can check that orbits converge to an almost-cycle $(\frac{1}{3}, \frac{2}{3})$; this is not a true cycle since $\frac{1}{3}$ is sent to $\frac{1}{6}$ because of the tie-breaking rule, but points $z < \frac{1}{3}$ close to $\frac{1}{3}$ are sent near $\frac{2}{3}$. One can check further that for $t > \frac{1}{2}$, no orbit realizes the sequence $ababa\dots$ of winners (but $aaa\dots$ and $bbb\dots$ are not realized either, the winner still alternate between a and b , but not as regularly as for $t < \frac{1}{2}$).

Remark 3.5. As mentioned by one of the anonymous referees, it seems Theorem A should apply to asynchronous iterative voting. It would need to change the definitions and notation to adjust to this case; let us briefly describe how one could proceed.

The usual discrete model for asynchronous iterative voting yields an oriented graph, with an arrow from one ballot profile to another when an admissible move takes the former to the latter. This can be modeled as a set-valued function on the set of ballot profiles. The continuous space dynamics could thus be any set-valued function on the CS space we defined here. The discrete model would embed in that CS space (and give a particular example of CS dynamics) as follows: given a state, first we apply the CS information function G , and then we look at all admissible moves given this result. We can either assume all voters of a given type apply the chosen move, or only a (non-zero) “quantum” of them (rounding if necessary in such a way that every paths ultimately reflects a paths occurring in the original discrete model). In both cases we get a set-valued map on the CS space. Perturbations can be measured by the uniform Hausdorff distance, i.e. when Φ and Φ' are set-valued maps, $D(\Phi, \Phi')$ would be the least ε such

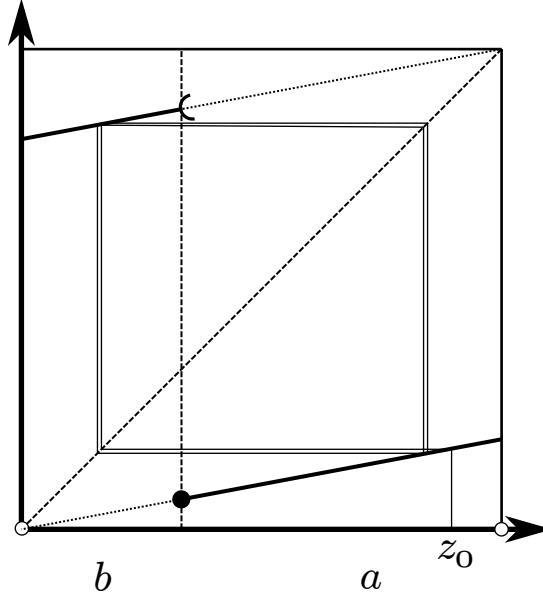


Figure 1: Graph of Φ_t and an orbit. Dotted: construction lines. Dashed: diagonal and delimitation between the two possible outcomes.

that for all state p , every element of $\Phi(p)$ is at distance at most ε from some element of $\Phi'(p)$, and every element of $\Phi'(p)$ is at distance at most ε from some element of $\Phi(p)$.

Then a variant of Theorem A can be expected, with basically the same proof: for each cycle in the discrete model that stays away from the frontiers between candidates' winning regions, if we consider a small enough perturbation we shall find a nearby CS cycle with the same sequence of winners.

4 Bad cycles in Approval Voting

Let us recall some definitions relative to what can be considered “good” or “bad” collective choices among the candidates.

Given a preference profile, 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 she exists she 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. Similarly 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.

Last, we will use a stronger notion than Condorcet loser: a candidate is said to be an *absolute majority loser* whenever she is a Condorcet loser and there is a strict majority of the electorate that ranks her last (possibly tied with others) in their preferences. An absolute majority loser rarely exists.

4.1 First example: non-convergence of the Leader Rule

In this Section all heuristics will be given by the Leader Rule, which only relies on the identity of the expected winner and runner-up, not on the expected scores of candidates (its definition is recalled in Section 2.2). We can thus consider a simple set of outcomes, $\mathcal{O} = \{\omega\rho: \omega \neq \rho \in A\}$ the set of order pairs of distinct candidates, and the winner map is defined by $W(\omega\rho) = \omega$; ρ represents the runner-up. Preferences of voter types are without tie (i.e. in $\mathcal{L}(A)$) as needed for the Leader Rule. The information function g is obtained by ranking candidates by decreasing order of approval numbers ($r_\alpha = \sum_{i|\alpha \in B_i} w_i$), breaking ties by alphabetical order, and then selecting the two first candidates, keeping their order. The CS information function G is defined similarly, with approval numbers

$$r_\alpha = \sum_{\substack{i \in N \\ B|\alpha \in B}} p_B^i w_i,$$

making it an extension of g . The admissible ballots of type i are the sincere ballots with respect to the preference order \succsim_i ; the LR ensures that heuristics only pick admissible ballots. Our goal is to prove the following result.

Theorem B. *Using Approval Voting, there exists a preference profile for 4 candidates such that:*

- *there is a Condorcet winner,*
- *each voter has preferences without ties,*
- *assuming voters follow the Leader Rule, the Polling Dynamics has a cycle along which the Condorcet winner is never elected. Moreover, a majority of the initial ballot profiles lead to this cycle, among those who are both sincere and expressive (i.e. approving a non-empty, strict subset of candidates).*

Proof. We set $A = \{a, b, c, d\}$ and consider the following preference profile with 7 types of voters:

T	U	V	W	X	Y	Z
100	1000	1001	1002	1004	1008	1016
a	b	c	d	b	c	d
b	a	a	a	c	d	b
c	c	d	b	a	a	a
d	d	b	c	d	b	c

and heuristics given by the Leader Rule.

The voters of type U, V, W like a but each prefers one of b, c, d better, while the voters of type 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 $b \succ c \succ d \succ b$ with a close to tie with each of b, c, d . Meanwhile, voters of type T prefers a to any other candidate, 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 types U to Z are chosen, for the sake of fanciness, to exclude any perfect tie (different sums of distinct powers of 2 never agree).

Figure 2 represents the shifted Polling Dynamics ψ in the form of a graph. We only give the details of the computations along the cycle, others are similar.

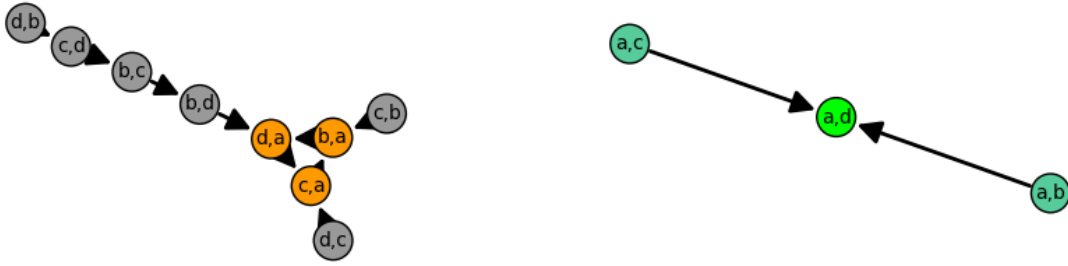


Figure 2: The shifted Polling Dynamics of the first example. Outcomes where the Condorcet winner is elected are shown in green, light green for the dynamical equilibrium. The periodic outcomes not electing the Condorcet winner are shown in orange.

Consider the outcome ba . Under the Leader Rule, it leads to the following ballot profile and results:

$$\begin{array}{llllll}
 B_T = \{a\} & B_U = \{b\} & B_X = \{b\} & r_a = 3111 & r_b = 3020 & \\
 & B_V = \{c, a, d\} & B_Y = \{c, d, a\} & & r_c = 2009 & \\
 & B_W = \{d, a\} & B_Z = \{d, b\} & & r_d = 4027 &
 \end{array}$$

so that $\psi(ba) = da$ - i.e. a stays second, while the previously unthreatening d comes in first position. The strategic adjustments triggered by the outcome da are as follows:

$$\begin{array}{llllll}
 B_T = \{a, b, c\} & B_U = \{b, a, c\} & B_X = \{b, c, a\} & r_a = 3105 & r_b = 2104 & \\
 & B_V = \{c, a\} & B_Y = \{c, d\} & & r_c = 4113 & \\
 & B_W = \{d\} & B_Z = \{d\} & & r_d = 3026 &
 \end{array}$$

so that $\psi^2(ba) = \psi(da) = ca$. The corresponding strategic adjustments are then:

$$\begin{array}{llllll}
 B_T = \{ab\} & B_U = \{b, a\} & B_X = \{b, c\} & r_a : 3118 & r_b : 4122 & \\
 & B_V = \{c\} & B_Y = \{c\} & & r_c : 3013 & \\
 & B_W = \{d, a, b\} & B_Z = \{d, b, a\} & & r_d : 2018 &
 \end{array}$$

so that $\psi^3(ba) = ba$.

The graph shows that more than two third of the outcomes lead to the bad cycle. A computer assisted enumeration of all possible sincere, expressive ballot profiles shows that 1353 out of 2187, i.e. more than 60 %, lead to the bad cycle. \square

Remark 4.1. It is worthwhile to observe (although a bit tedious to check) that in this example, *all* outcomes can be obtained from a certain sincere ballot profile. For example, *dc* is obtained if all voters cast the ballot with only their preferred candidates. That all outcomes are covered is hardly surprising: there are only 12 of them, but $3^7 = 2187$ combinations of sincere ballots (excluding abstentions).

Remark 4.2. In the cases of $m > 4$ 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 outcomes where one of the dummy candidates wins will all be sent by φ to an outcome where *a* wins, 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 50 voters of type *S* with preference order *becda*. 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 4.3. As pointed out by one of the anonymous referees, Theorem B does not hold for 3 candidates. Assume there is a Condorcet Winner *a* and two other candidates *b, c*. Since preferences have no ties, there are 6 possible preferences, and we can reduce to six voter types; let $T_{\alpha\beta}$ denotes the voters ranking α first and β second, and let it also denotes the number 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,

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

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 *ab* or *ac*.

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 *ab* or *ac* is an equilibrium, attracting the other one of these two outcomes.

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 ca, ba 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.

Remark 4.4. If we do not ask for a Condorcet winner to exist, a cycle can be produced with $m = 3$ candidates, e.g. with the preference profile

X	Y	Z
10	11	12
a	b	c
b	c	a
c	a	b

where starting with voters approving only their top candidate, we obtain a cycle of outcomes cb, bc, ab .

From Theorem A we deduce that the cycle persists when we perturb the Polling Dynamics in the CS setting (to be meaningful, here we change back the set of outcomes to $\mathcal{O} = [0, 1]^A$).

Corollary 4.5. *Consider a preference profile as given by Theorem B. Let Φ_0 be the induced CS Polling Dynamics on the continuous space \mathcal{P} defined by making admissible exactly the sincere ballots. Then there exist $\varepsilon > 0$ and an open subset U of \mathcal{P} such that for all CS Polling Dynamics Φ with $D(\Phi, \Phi_0) \leq \varepsilon$ and all $p^0 \in U$, the sequence of winners*

$$W \circ G(\Phi^j(p^0)) \quad j = 0, 1, 2, \dots$$

is periodic and does not contain the Condorcet winner.

4.2 Second example: the possible election of an absolute majority loser

Our 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 candidates, and it exhibits a cycle where an absolute majority loser could get elected. We conserve most the setting of the previous Section: \mathcal{O} , W , g and G are as above. The only change is that we allow for ties in preferences and more varied heuristics.

Theorem C. *Using Approval Voting, there exists a preference profile on 3 candidates with ties in preferences allowed and sincere heuristics such that:*

- there are a Condorcet winner and an absolute majority loser,
- the Polling Dynamics has a 2-cycle, one of whose ballot profiles elects the absolute majority loser,
- there is an equilibrium not electing the Condorcet winner,
- moreover only one of four sincere, expressive ballot profiles avoid the above bad cycle and equilibrium.

Proof. We consider the following preference profile:

Z	Y	X	W
3	1	3	5
a	a	b	c
b	bc	a	ab
c		c	

with each voter type $i \in N$ using the Modified Leader Rule. In particular 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 4 to 3 (voters of type W abstaining) and c with a score of 7 to 5. Moreover c is a Condorcet loser, loosing to b by 6 to 5; actually c is a worst candidate of 7 out of the 12 voters, making it an absolute majority loser.

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

$$B_Z = \{a, b\} \quad B_Y = \{a\} \quad B_X = \{b, a\} \quad B_W = \{c\}$$

$$r_a = 7 \quad r_b = 6 \quad r_c = 5$$

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 heuristics they choose to vote only for their favorite candidate:

$$B_Z = \{a\} \quad B_Y = \{a\} \quad B_X = \{b\} \quad B_W = \{c\}$$

$$r_a = 4 \quad r_b = 3 \quad r_c = 5.$$

The second poll thus results in a 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 same ballots being cast as in the first poll, so we get a 2-cycle, in which the worst candidate is elected in one of the outcomes.

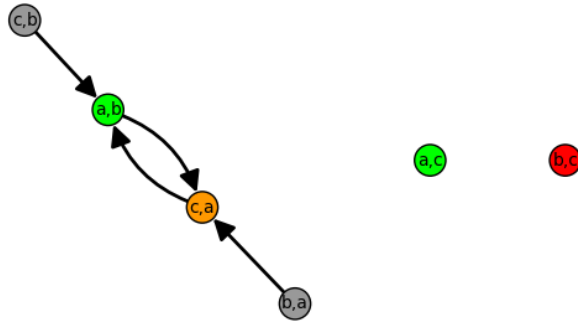


Figure 3: The shifted Polling Dynamics of the second example. Light green outcomes are accessible after an arbitrary large number of iterations and elect the Condorcet winner; orange (periodic) and red (fixed point) outcomes are accessible after an arbitrarily large number of iterations but do not elect the Condorcet winner. The orange outcome is arguably the most problematic, as it elects the absolute majority loser c .

Figure 3 represents the shifted Polling Dynamics ψ in graph form.

We see that of 6 outcomes, 4 lead to the cycle that can elect either the Condorcet winner a or the absolute majority loser c depending on whether the number of polls conducted before the election is odd or even, while the other 2 are equilibria, one electing the Condorcet winner a and the other the Condorcet runner-up b . There are 4 sincere, expressive ballot profiles (two possible choices for each of Z and X), one for each of the equilibria and each of the 2-cycle outcomes. \square

Remark 4.6. 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 heuristic 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 C also holds with the stronger definition of Condorcet winner (see beginning of Section 4).

Remark 4.7. An 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 of 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 .

Again, Theorem A ensures that this bad cycle is robust. Let us make this more concrete by considering an explicit perturbation.

Example 4.8. We now give a CS setting extending the discrete-space example in the proof of Theorem C. Let the sets of admissible ballots be

$$\mathcal{B}_Z = \{\{a\}, \{a, b\}\}, \quad \mathcal{B}_Y = \{\{a\}\}, \quad \mathcal{B}_X = \{\{b\}, \{a, b\}\} \quad \text{and} \quad \mathcal{B}_W = \{\{c\}\}$$

and consider outcomes giving the proportion of votes received by each candidates: $\mathcal{O} = [0, 1]^A$.

Observe that since the simplex over a singleton is a singleton, and the simplex over a two-element set is an interval, \mathcal{P} can be identified with the square $[0, 1]^2$, with coordinates (x, z) where x (respectively z) represents the proportion of voters of type X (respectively Z) casting the ballot $\{a, b\}$. The leftmost part of Figure 4 represents \mathcal{P} . 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); they are concurrent at the point where all three candidates are tied, and delimit six areas, one for each possible outcome (all of which are possible under the chosen preference profile and restriction of ballots, as we can see in the figure). We denote by A_{abc} the region where the outcome is abc (boundary segments are attributed according to the tie-breaking rule), and similarly for the other five regions.

Given $\delta \in [0, 1]$, we let \mathcal{T}_δ be the set of $(x, z) \in \mathcal{P}$ such that $(r_a, r_b, r_c) := V(x, z)$ avoids ties by a margin at least δ , i.e. $|r_\alpha - r_\beta| \geq \delta$ for all $\alpha \neq \beta \in A$. We consider any perturbation $(\Phi_\varepsilon)_{\varepsilon \in [0, 1]}$ of the embedded discrete space Polling Dynamics such that:

$$\Phi_\varepsilon(x, z) = (1 - \varepsilon)\Phi_0(x, z) + \varepsilon(x, z) \quad \forall (x, z) \in \mathcal{T}_\delta$$

which models the situation where after each poll where no two candidates are δ -close to be tied, a proportion ε of voters keep their ballots unchanged and the remaining $(1 - \varepsilon)$ apply the heuristics of their types defined in the discrete setting (here the LR for types X and Z , types Y and W always voting only for their top candidate).

We will show that taking $\delta = .04$ and $\varepsilon = .15$ is small enough for the conclusion of Theorem A to apply. Figure 5 shows the first five iterations of such a perturbation for $\delta = .04$ and $\varepsilon = .15$, and we can observe a bad 2-cycle with points in A_{abc} and A_{cab} .

The region $A_1 \subset A_{abc}$ 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 $A_2 \subset A_{cab}$ 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. We have $\Phi_{.15}(x, z) = (.15x, .15z)$ whenever $(x, z) \in A_1$ and $\Phi_{.15}(x, z) = (.85 + .15x, .85 + .15z)$ whenever $(x, z) \in A_2$. One easily checks that $\Phi_{.15}(A_1) \subset A_2$ and $\Phi_{.15}(A_2) \subset A_1$. It follows that $\Phi_{.15}^2(A_1) \subset A_1$, and since $\Phi_{.15}^2$ is a contraction on A_1 (of ratio $.15^2$, much small than 1), 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.

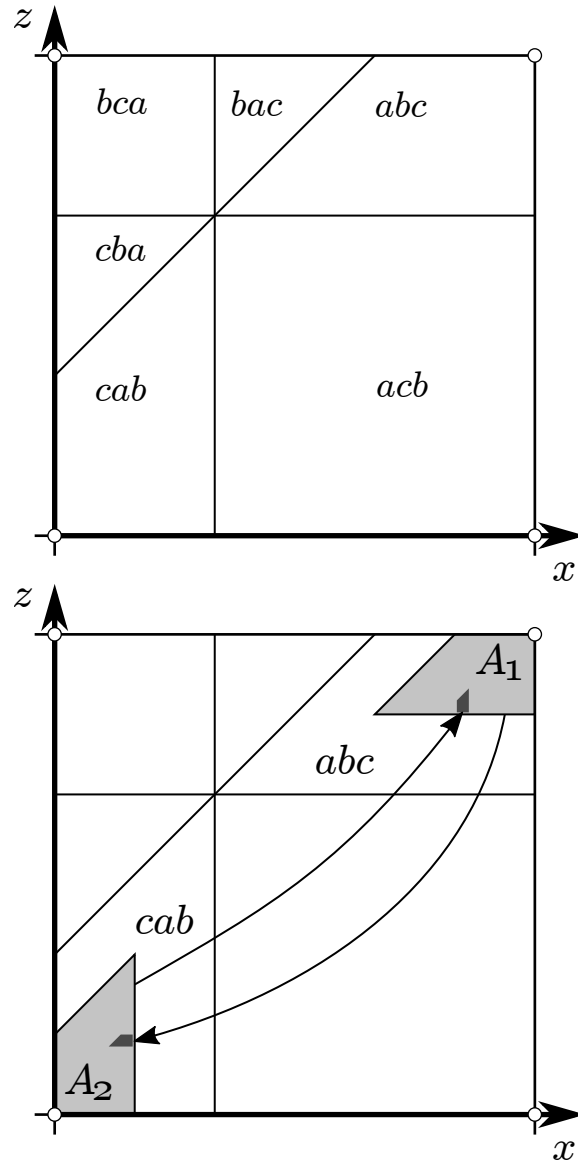


Figure 4: A continuous-space example. *Left:* the state space \mathcal{P} . The corners are the extreme states, corresponding to the four outcomes that are attainable from a ballot profile under the discrete-space Polling Dynamics. *Right:* on the same state space, any CS Polling Dynamics where 85% of 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.




Figure 5: An example of a CS Polling Dynamics obtained by perturbation of the Polling Dynamics used in the proof of Theorem C, showing robustness of the bad 2-cycle. 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.

4.3 Numerical study

The examples of the previous section are interesting from a theoretical point of view, showing that the existence of equilibria 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 preference profiles where a Condorcet winner exists. We continue to follow the setup of Sections 4.1 and 4.2, using either the Leader Rule or the modified Leader Rule. This leaves us to model the preference profile, and for this we compare several “cultures”.

Our results complement for example the findings of [LLV00] and [Las10]. Compared to the latter, for each preference profile we determine the existence of a bad cycle anywhere in the space of outcomes, and of arbitrary length.

Main experimental setup We fix the number of candidates to $m = 6$, and the number of voter types to $n = 20$, $N = \{1, \dots, 20\}$, a compromise between complexity of the preference profile and computational power limitations. The influence of these numbers is tested at the end of the section.

Each voter type is assigned a number of voters uniformly drawn in $[0, 1]$ and is assigned a preference modeled by a linear order on $\bar{A} = A \cup \{\text{limit}\}$; when we use the Leader Rule, *limit* is simply ignored, but when we use its modified version, the candidates ranked beyond *limit* are considered tied for the last rank and thus never put on this voters' ballots. For example, the order $a \succ b \succ \text{limit} \succ c \succ d \succ e \succ f$ corresponds to the preference $abcdef$ when the LR is used, and to $ab(cdef)$ when its modified version is used.

We consider the following ways of constructing preferences:

- *Impartial Culture*: each voter type draws uniformly an element $\mathcal{L}(\bar{A})$,
- *d-dimensional culture*: each candidate and each voter types is given a uniform random position in $[0, 1]^d$ where each coordinate represents a “political axis”; the preferences of a voter type is obtained by sorting the candidates by increasing distance to the voter type, using the ℓ^1 metric

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

that sums the levels of disagreement along the various axes; *limit* is then inserted by drawing a random “distance” to the voter type, with the same law than the distance between two uniform points in $[0, 1]^d$

We ran experiments for each combination of a heuristic among the Leader Rule and 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 preference profiles and recorded whether a Condorcet winner exists, and if yes whether the Polling Dynamics has a “bad cycle” (or bad equilibrium), i.e. where some of the ballot profiles 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. The first take-away is that in practice

The Leader Rule is very effective in electing the Condorcet winner.

Only the Impartial Culture witnesses slightly larger odds of not electing her while structured cultures very rarely produce bad cycles, the linear ($d = 1$) culture having produced none in 100 000 attempts (this is no accident, see Section 4.4).

On the contrary, when voters apply the Modified Leader Rule and thus, with the kind of preferences used here, refuse to approve of candidates 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 culture, with 15% of preference profiles leading to a bad

Culture \ Heuristic	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 preference profiles where a Condorcet winner exists; in normal size, the proportion of preference profiles with a bad cycle or equilibrium, among preference profiles having a Condorcet winner (both rounded). The 95% confidence Wilson score interval gives a deviation of less than 0.23 percentage point for all values; the deviation is even less than 0.04 percentage point for the value 0.3% and below.

cycle in the Polling Dynamics. The situation that is most favorable to the Leader Rule, is also the worst one for its modification!

Imagine that voters of a given type are aware of a Condorcet winner quite low on their preferences, and that they know that when everyone uses the Leader Rule, the election of the Condorcet winner is very likely. By not applying the Leader Rule and instead refuse to approve candidates below a threshold, they can get significant odds to obtain the election of a candidate they prefer to the Condorcet winner:

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 least not below a certain threshold.

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 considers that each voter assumes imperfect recording of ballots and think as if she would be the only one to change her ballot, trying to optimize the expected outcome. Our numerical results show that whenever voters expect other voters of the same type to also change their votes strategically, the Leader Rule might not be their best heuristic.

This experiment also gave the opportunity to find especially problematic Polling Dynamics, see Figures 6 and 7 where green outcomes correspond to the Condorcet winner being elected, orange outcomes are part of a cycle but do not elect the Condorcet winner, and red outcomes are equilibria not electing the Condorcet winner.

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. 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 equilibria more likely. The main take away, valid in all considered conditions except the 1-dimensional culture with the LR, is that:

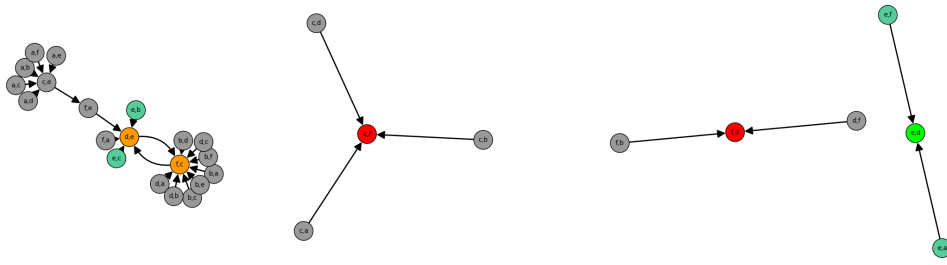


Figure 6: A bad cycle and two equilibria not electing the Condorcet winner under the modified Leader Rule, produced with the Impartial Culture. Only 3 outcomes out of 30 end up electing the Condorcet winner, and there are outcomes where the Condorcet winner is elected, but that lie in the basin of attraction of the bad cycle (this never happens with the Leader Rule).

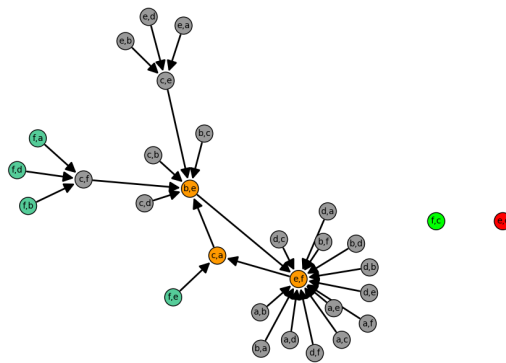


Figure 7: An example in the one-dimensional culture with the modified Leader Rule: the bad cycle attracts all outcomes but two, including most outcomes electing the Condorcet winner.

A larger number of candidates seems to make bad cycles and equilibria 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.

4.4 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 equilibria and that all equilibria elect ω . Here we prove that in this particular culture, we moreover have convergence of the Polling Dynamics.

Definition 4.9. We say that a preference profile *can be modeled by a one-dimensional culture* when there exist a *positional mapping*

$$x : A \cup N \rightarrow \mathbb{R}$$

Cul. & Strat.	C & T	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 preference profiles (except for 6 candidates and 20 voter types, where we reported the previous results with larger samples), yielding typically 4 times larger confidence intervals than samples of size 100,000 and leaving the expected deviations under the percentage point. In small, the proportion of preference profiles where a Condorcet winner exists; in normal size, the proportion of preference profiles with a bad cycle or equilibrium, among profiles having a Condorcet winner (both rounded).

with the following property: the preferences \succsim of any voter type $i \in N$ is given by

$$\forall \alpha \neq \beta \in A, \quad \alpha \succsim_i \beta \quad \text{if and only if} \quad |x(\alpha) - x(i)| < |x(\beta) - x(i)|.$$

For simplicity, we assume further that $|x(\alpha) - x(i)| \neq |x(\beta) - x(i)|$ for all $\alpha \neq \beta$ and all i , 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 4.10. *Under Approval Voting, if the preference profile can be modeled by a one-dimensional culture and voters apply the Leader Rule, then the Polling Dynamics*

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 4.10. We use outcomes reduced to the winner and runner-up, we assume that the preference profile can be modeled by a one-dimensional culture and we fix a positional mapping x . Denoting by $w = \sum_{i \in N} w_i$ the total number of voters, a *median* is a value $m \in \mathbb{R}$ such that

$$\sum_{i: x(i) \leq m} w_i \geq \frac{w}{2} \quad \text{and} \quad \sum_{i: x(i) \geq m} w_i \geq \frac{w}{2}$$

i.e., at least half the voters lie on the left of m , and at least half have lie 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: $m = x(i_m)$, where the voter type i_m is called the *median type*. By assumption, there is a single candidate μ whose position is closest to $x(i_m)$.

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

Lemma 4.11. *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 $x(i_m)$ and thus contains at least half the voters. Since there are no possible ties, μ dominates α . \square

Lemma 4.12. *Let $\alpha\beta$ be an outcome such that $\alpha \neq \mu$ and let $\alpha'\beta' = \psi(\alpha\beta)$ be the next outcome in the (shifted) Polling Dynamics. 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 H_α with extremity $x(\alpha)$ avoiding I , and the open half-line H_β with extremity $x(\beta)$ avoiding I .

If $x(\mu) \in H_\alpha$, 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 H_α or in $I \cup H_\beta$. 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 $>$ denote the ‘‘geo-lexicographic’’ order on outcomes, defined by $\alpha\beta > \alpha'\beta'$ 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 ψ , the outcome can only decrease in this order until μ becomes winner. Since there cannot be an infinite decreasing sequence

of outcomes, eventually μ becomes winner. Now, as is well-known, this is a stable situation: for all β , $\psi(\mu\beta) = \mu\beta_0$ where β_0 is the candidate getting the better score in a duel against μ . Therefore, every outcome converges under ψ to the outcome $\mu\beta_0$ (and every ballot profile of B^n converges under φ to the ballot profile induced by the outcome $\mu\beta_0$ and the Leader Rule). The proof of Theorem 4.10 is complete.

5 Other Voting rules

In order to show how our framework applies to general voting rule, let us consider three further examples: one for Plurality voting, where contrary to Example 3.4 heuristics use quantitative information in the outcome rather than only the ranking; one for Instant Run-Off voting (IRV); one for the *condorcified* version of IRV (or, for that matter, of any ranked voting).

There is a paucity of examples in the literature concerning *synchronous* iterative voting; we can still mention Example 1 in [CPP04] (whose heuristics imply that voters know the number of voters of the same type), and Example 7 in [BF07] (which needs a change of heuristic along the cycle).

5.1 A plurality cycle

We consider Plurality voting, with outcomes giving the total number of votes of each candidate, with $n = 4$ voter types and $m = 3$ candidates, $A = \{a, b, c\}$. We thus set $\mathcal{O} = [0, w]^A$ where $w = \sum_{i \in N} w_i$ is the total number of voters, an element of which is written as a triple $r = (r_a, r_b, r_c)$. Ballots bear a single name, i.e. $\mathcal{B} = A$, and the information function g sends a family $(B_i)_{i \in N} \in A^N$ of ballots to the outcome

$$g((B_i)_{i \in N}) = \left(\sum_{i \in N, B_i = \alpha} w_i \right)_{\alpha \in A}.$$

The winner map sends (r_a, r_b, r_c) to the candidate with the most votes, with ties broken in favor of the earlier candidate in the alphabetical order.

We consider the preference profile

X	Y	Z
10	13	14
a	b	c
b	a	ab
c	c	

and assume the following heuristics:

- voters of type Z always vote for c ,

- voters of type X vote as follows. If their least-preferred candidate c is in a close contest for winning with their second-preferred candidate b , more precisely if

$$|r_c - r_b| \leq .03(r_a + r_b + r_c) \quad \text{and} \quad \min(r_b, r_c) > r_a + .03(r_a + r_b + r_c)$$

then they vote for b ; in all other cases they vote sincerely, for a ,

- voters of type Y vote as voters of type X , with a and b switched.

These heuristics allow voters of type X and Y to vote strategically whenever their preferred choice is far from winning, but their second-preferred choice needs support to avoid the win of their least-preferred candidate. The factor $.03$ is chosen because it is the typical uncertainty in polls on about a thousand voters. With the above preference profile, $.03(r_a + r_b + r_c) = 1.11$.

If voters first report their preferred candidates, we obtain the outcome

$$r_0 = (10, 13, 14)$$

and we are in the situation where c wins in close contest with b . The above heuristics lead voters of type X to vote for b , while others remain sincere, leading to the new outcome $r_1 = \psi(r_0) = (0, 23, 14)$ where b wins by a large margin. The above heuristics lead every voter to vote sincerely, so that $\psi(r_1) = r_0$, and we have a two-cycle (an outcome of which elects the absolute majority loser). By Theorem A, this bad cycle persists in a small enough CS perturbation.

5.2 An IRV cycle

We now consider IRV and give an example with $n = 3$ and $m = 4$. Ballots and outcomes are linear rankings of candidates: $\mathcal{B} = \mathcal{O} = \mathcal{L}(A)$. The information function g is defined as follows. Given an argument $(B_i)_{i \in N} \in \mathcal{B}^N$, we first rank candidates according to the total weight of ballots where they are ranked top (again, ties broke in alphabetical order). The last ranked candidate is eliminated, and we rank the remaining candidates according to the total weight of ballots where they are ranked top *among candidates that have not been eliminated*. The last ranked candidate is again eliminated. The resulting outcome is denoted with the last standing candidate first, then the other by reversed order of elimination. The winner map sends a ranking $r \in \mathcal{O} = \mathcal{L}(A)$ to its top candidate.

We consider the preference profile

X	Y	Z
10	11	12
b	c	d
a	a	a
c	b	b
d	d	c

so that the Condorcet order coincides with the alphabetical order, in particular a is the Condorcet winner. Let us define simple heuristics that make use of an expected outcome to try to improve it. Denote by $\alpha\beta\gamma\delta$ the preferences of the considered type i of voters (e.g. for $i = X$, $\alpha = b$, $\beta = a$, $\gamma = c$ and $\delta = d$). Given an outcome r , if the expected winner is neither α nor β and β was eliminated before α then $\sigma_i(r) = \beta\alpha\gamma\delta$; for every other outcome r , $\sigma_i(r) = \alpha\beta\gamma\delta$. In other words, voters strategically invert their two first preferred candidates when the top one had a better run but still could not prevent a worse candidate to be elected, giving their second choice a chance to do better.

If voters start voting sincerely, we obtain the outcome $cdba$ with the Condorcet winner eliminated with no votes at all in the first round, and ballots of type X voters ultimately transferred to candidate c against d . The above heuristics have voters of type X and Z invert their two most preferred candidates and vote $abcd$, $adbc$ respectively. The outcome is then $acbd$. Then all voters resume voting sincerely, and we have a 2-cycle one of whose ballot profiles elects the third candidate in the Condorcet order. Observe that every outcome where the Condorcet winner is elected leads voters to cast their sincere ballot, thus leading to this bad cycle.

According to Theorem A, this cycle is again robust under perturbation in a CS setting.

5.3 A cycle in condorcified ranked voting

Consider a ranked voting rule, i.e. one with $\mathcal{B} = \mathcal{L}(A)$. It has been proved by Durand, Mathieu and Noirie [DMN14, DMN16] and independently by Green-Armytage, Tideman and Cosman [GATC16] that under mild conditions *condorcification* of a voting system (i.e. electing the Condorcet winner if she exist and applying the given rule otherwise) cannot increase manipulability and often reduces it. We can therefore ask whether a condorcified voting rule is less susceptible to the presence of bad cycles when a Condorcet winner is present. Unsurprisingly, the answer is negative when voters who do not like the Condorcet winner strategically choose to rank her lower in their ballot

Let us consider a condorcified voting rule, i.e. for all $(B_i)_{i \in N} \in \mathcal{B}^N$ for which their is a Condorcet winner γ , $W \circ g((B_i)_{i \in N}) = \gamma$. Assume moreover that \mathcal{O} contains the information whether or not the winner was a Condorcet winner (with respect to the *ballot* profile, since the preference profile is not available information). Consider the preference profile

Z	Y	X	W
150	102	101	100
a	b	c	b
b	c	a	a
c	a	b	c

for which a is the sincere Condorcet winner. Assume the following heuristics for voter of any type i , whose preferences are denoted as above by $\alpha\beta\gamma$: if β is elected as Condorcet winner, then they cast the ballot $\alpha\gamma\beta$, in an attempt to give α a shot. Otherwise, they vote sincerely.

If at first all voters vote sincerely, in the first outcome r_0 the winner is a , declared a Condorcet winner. This leads voters of type X to next cast the ballot cba and voters of type W to cast the ballot bca . This makes b the Condorcet winner of the new ballot profile. The above heuristics then lead all voters to vote sincerely, except voters of type Z who cast the ballot acb ; this makes again a the Condorcet winner, and we have a 2-cycle, one of whose ballot profiles elects a candidate that is not the Condorcet winner of the preference profile. Moreover this cycle attracts all outcomes where a is a Condorcet winner of the ballot profile.

Again, Theorem A ensures that this cycle is robust under perturbation in a CS setting.

6 Chaos

We finish with an example illustrating the flexibility of the CS setting for the Polling Dynamics. The starting point is to observe that the embedding Φ_0 of a discrete space Polling Dynamics in a CS setting cannot be continuous (unless it is a constant map). Indeed, it takes a finite number of values, and the state space \mathcal{P} is connected. There must exist some CS ballot profiles near which arbitrarily close CS ballot profiles are sent to very different images by Φ_0 . In practice, assuming the outcomes carry continuous information (e.g. shares of votes for each candidate), a heuristic like the Leader Rule would probably not be applied blindly by all voters when the expected winner and runner-off are in a close-call contest, or when the runner-up is in a close-call contest with the next candidate. Different voters will have different confidence in the available information and a small change in the expected outcome should result in a small change in the ballots cast, i.e. the CS Polling Dynamics should be continuous.

We now give a relatively simple example with $m = 3$, based on a reluctance to add one's second-preferred candidate unless it seems likely to improve the outcome, and we observe that a complicated dynamics emerges. The take-away is that

Even simple CS Polling Dynamics can exhibit a chaotic behavior, where the sequence of winners is impossible to predict reliably from the observation of arbitrarily many of its first terms.

By looking at the influence of a small number of core parameters of the model, we shall see that chaos is neither universal nor restricted to exceptional parameters.

A simple continuous CS Polling Dynamics We consider the four-types preference profile used in the proof of Theorem C:

Z	Y	X	W
3	1	3	5
a	a	b	c
b	bc	a	ab
c		c	

for which a is a Condorcet winner and c an absolute majority loser. As above, the corresponding sets of admissible ballots makes \mathcal{P} a square with coordinates (x, z) where x, z are the proportions of voters of type X, Z respectively voting $\{a, b\}$. This choice of ballot can be seen either as a form of cooperation between the two types of voters, or as the result of risk aversion (which are complementary views, not opposed ones).

An outcome is a triple $r = (r_a, r_b, r_c) \in [0, 1]^3$ giving the shares of votes obtained by each candidate, and we choose the following CS heuristics (here the resulting distribution of ballots does not depend upon the value x or z in $\Delta(\mathcal{B}_i)$ that is by definition part of their argument):

$$\sigma_Z((r_a, r_b, r_c), z) = C_Z \circ S(r_a, r_b, r_c) \quad \sigma_X(x, (r_a, r_b, r_c), x) = C_X \circ S(r_b, r_a, r_c)$$

where C_Z, C_X, 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 r_a, r_b are exchanged in σ_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:

$$S(r_1, r_2, r_3) = \begin{cases} |r_2 - r_3| & \text{when } r_2 > r_1 \\ \frac{1}{2}|r_2 - r_3| + \frac{1}{2}|r_1 - r_3| & \text{otherwise.} \end{cases}$$

i.e. when the second-preferred candidate 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 $r_1 \gg r_3 \simeq r_2$ then the safety is not too small (the least preferred candidate has little chance of winning) but not maximal, meaning some voters will prefer to collaborate in order to ensure c finishes at the last rank. When $r_1 \simeq r_2$, the safety is close to both $|r_1 - r_3|$ and $|r_2 - r_3|$; when c is far above or far below, one’s vote seems unlikely to change the outcome and the incentive to collaborate is small.

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:

$$C_i(t) = (1 - c_i t)_+ := \max(0, 1 - c_i t)$$

where as indicated $(\cdot)_+$ is the positive part, and where the coefficient c_i ($i = Z, X$) quantifies the risk tolerance: the higher, the less likely X and Z voters will collaborate (low values of c_i thus correspond to high aversion to risk). For now, we take $c_Z = c_X = 5$, a rather moderate value: $1 - 5t$ 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 $(r_a, r_b, r_c) = (3x + 4, 3z + x, 5)$, these choices yield the CS Polling Dynamics Φ defined by:

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

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

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

when $2x + 4 \leq 3z$. While Φ may look like an innocent map, drawing an orbit of this map reveals an interesting pattern (Figure 8). It turns out *all* orbits yield pretty much the same image, typical of a *chaotic attractor*.

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, of which we shall only give one flavor (more information is available in numerous books, e.g. [KH95]). It is to be understood that “chaotic” situations are those of positive entropy, zero entropy being the sign for a relatively “tame” system.

While we will not *prove* chaos here, we will show compelling numerical evidence that some entropy is positive, meaning that an orbit have the same level of complexity as a random sequence of (skewed) coin tosses. Note that the fact that the map Φ is chaotic is not the most important point: what matters is that the sequence of winners along many orbits are chaotic, which here results from the attractor intersecting regions electing different candidates.

Numerical quantification of the chaos Given any state $s \in \mathcal{P}$, its *winners word* is the word $\omega(s)$ whose k -th letter is $W(G(\Phi^k(s)))$, the winner after k iterations. In the example above the word of the successive winners starting from the center of the space \mathcal{P} is (32 first letters shown, capital c for better readability):

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

(the choice of $(.5, .5)$ is arbitrary, but experimentally one checks that it has no bearing on the results exposed here). 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:

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

To quantify the randomness of $\omega(.5, .5)$, we use the following definition.

Definition 6.1. Let $\omega = \alpha_1\alpha_2 \dots$ be an infinite word on the alphabet A , i.e. a sequence of elements of A (written without parentheses or comma). For all $k < j \in \mathbb{N}$ we set $\omega_{k:j} = \alpha_k\alpha_{k+1}\alpha_{k+2} \dots \alpha_j$; each $\omega_{k:j}$ is called a *subword* of ω (of length $j - k + 1$).




Figure 8: One orbit of the CS Polling Dynamics Φ . Left: full picture. Right: a zoom to the upper-right part showing a linear/fractal structure.

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

$$P_k^\theta(\omega) := \frac{S_k^\theta(\omega)}{k - \ell + 1}$$

the proportion of length ℓ subwords of $\omega_{1:k}$ equal to θ , and by $P_k^\ell(\omega)$ the “probability vector” $(P_k^\theta(\omega))_{\theta \in A^\ell}$. Finally the *Kolmogorov-Sinai entropy* of the word ω is defined as

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

where

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

Remark 6.2. The limit in the definition exists by Fekete’s lemma, the needed subadditivity being of the perks of the function H . The entropy $h_{\text{KS}}(\omega)$ is never larger than $\log m$, with equality when all possible subwords of length ℓ appear in ω with the same asymptotic frequency.

The value of $h_{\text{KS}}(\omega)$ is to be interpreted as the “uncertainty” of a random guess. The reference case is that of an object uniformly drawn among N , having uncertainty

$\log N$; Observe that 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 $\log(N^\ell) = \ell \log N$. 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). The uncertainty of guessing a length- ℓ subword of ω , when ℓ is large and we look far on the right of the word, is asymptotically of the order of $\ell \cdot h_{\text{KS}}(\omega)$.

The sequence $\ell \mapsto H(P_{2^{20}}^\ell(\omega))$ is plotted in Figure 9. We observe an extremely good alignment from $\ell = 4$ onward, with slope $\simeq 0.229$, a strong indication that $h_{\text{KS}}(\omega(.5, .5))$ is close to this value. 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 5, we have drawn the first iterates of Φ in Figure 10. 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 .

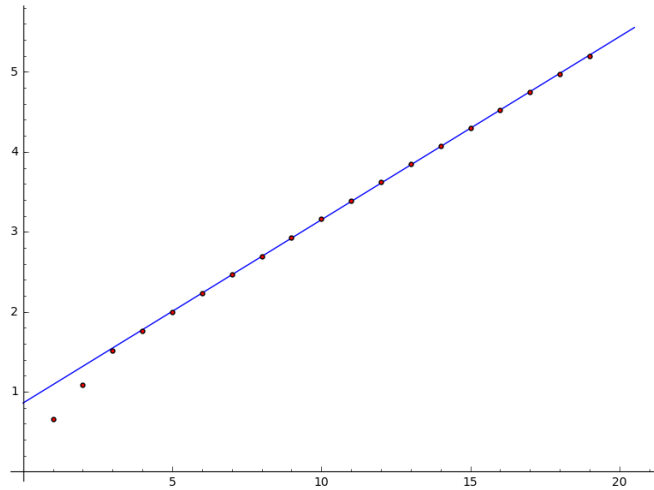


Figure 9: Plot of the entropy-estimating sequence $\ell \mapsto H(P_{2^{20}}^\ell(\omega))$. The slope of the marked line is $\simeq 0.2291$

While this is only a specific example, it was constructed from basic principles and it is possible that chaos be a quite common feature of CS Polling Dynamics. In order to start testing this hypothesis, we look at what happens when we let c_X and c_Z vary, see Figure 11. We observe that entropy is mostly concentrated in an L shaped region, where both voter types have not too high a risk aversion, and either type has not too high risk tolerance. However when Z has high risk-aversion, moderate-high values of c_X still result in positive entropy. This asymmetry must come from the asymmetry of the preference profile introduced by voters of type Y .

Looking at sequences of winners with various parameters show that in the upper-right square where entropy vanishes (high risk tolerance for both voter types), several different

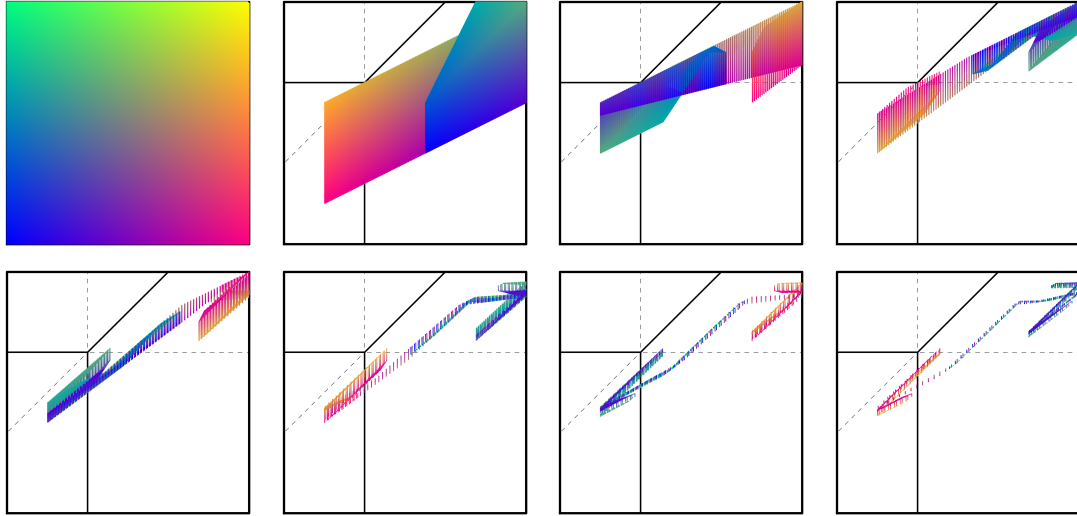


Figure 10: From left to right then top to bottom, 200×200 points of the state space regularly spaced and their 7 first iterates by Φ (colors preserved along each orbit).

dynamics can happen: c can be constantly elected or we can observe a periodic pattern involving a and c as winners. When X has high risk aversion, we only observed constant election of a . When Z has high risk aversion, for increasing values of c_X , we observe successively: constant election of a , periodic patterns involving a and b , chaotic patterns involving a and b , constant election of b .

Let us end with the influence of the preference profile. We look at the basic example ($c_X = c_Z = 5$.) with randomly chosen weights for all four voter types. The following examples complements the main one to form a representative sample of what a quick experimental exploration revealed for this model.

Example 6.3. With $w_X = 0.05$, $w_Y = 0.02$, $w_Z = 0.8$ and $w_W = 0.8$, 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 6.4. Take now $w_X = 0.6$, $w_Y = 0.08$, $w_Z = 0.56$ and $w_W = 0.82$. 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 :

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

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

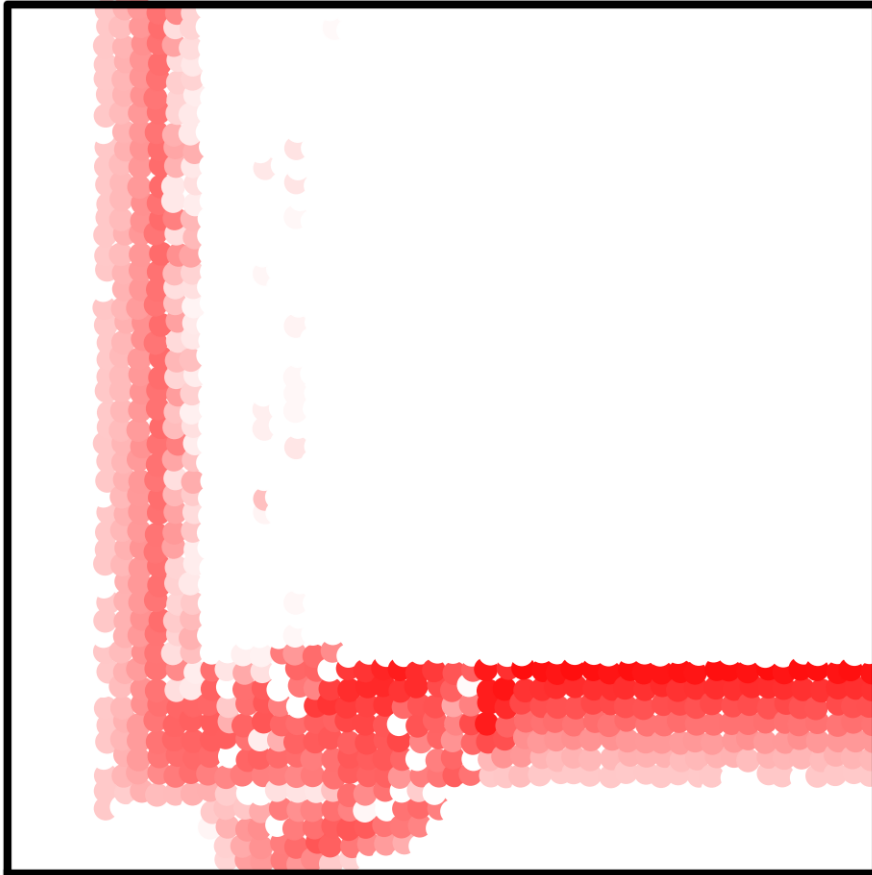


Figure 11: Entropy estimated for various values of c_X (abscissa) and c_Z (ordinate), from 1. to 30., mostly regularly spaced with a random perturbation. Each dot is one value of the pair (c_X, c_Z) , color indicating the estimated entropy (linear regression made on subwords of length 4 to 11 in $\omega(.5, .5)$): white is entropy 0 while pure red is entropy $\ln(2)$ or above.

regions, constantly alternating between two behaviors: predictable and chaotic. See one orbit in Figure 12 (left); other orbits produce very similar pictures.

Example 6.5. Changing only slightly the previous example with $w_X = 0.6$, $w_Y = 0.08$, $w_Z = 0.56$ and $w_W = 0.81$, a radical change in long-term behavior appears: entropy seems to vanish, with $H(P_{20}^\ell)$ 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’’).

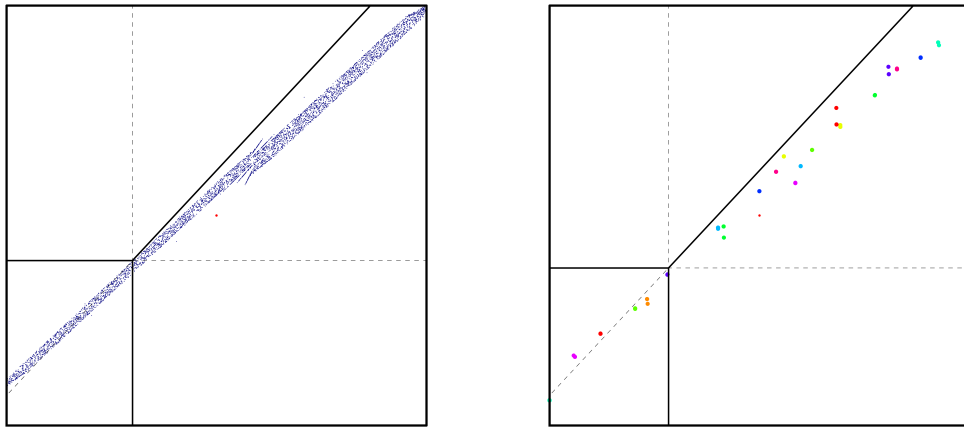


Figure 12: Comparison of the orbits of $s = (.5, .5)$ for two close populations: Example 6.4 on the left, Example 6.5 on the right, with colors of period 22. Both runs have 5000 points (larger points are used on the right for readability).

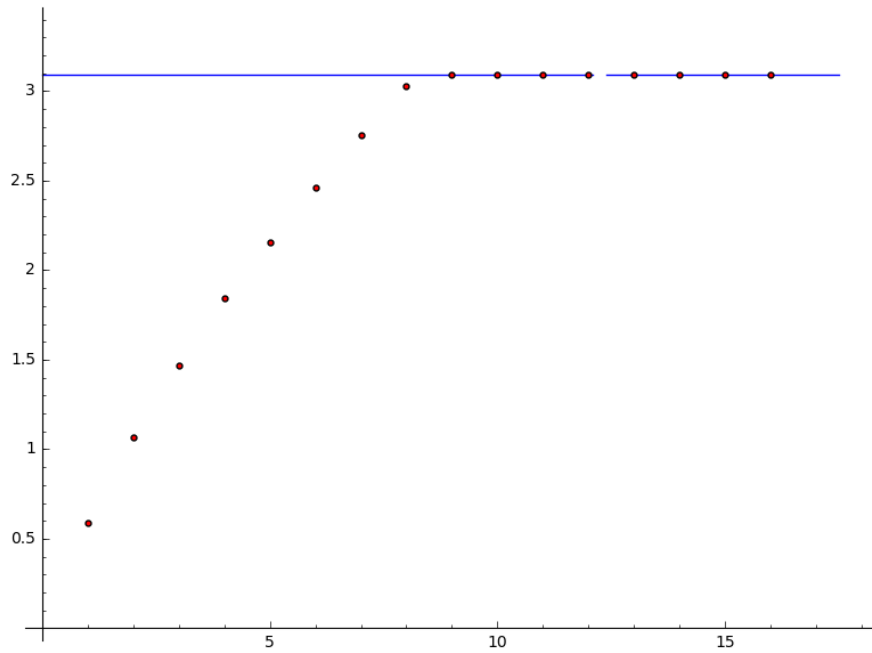


Figure 13: Plot of the entropy-estimating sequence $\ell \mapsto H(P_{220}^\ell(\omega))$ for Example 6.5

7 Conclusion

We have explained why *synchronized* iterated voting is relevant both to model political elections preceded by polls, and to iterative anticipation of other voters’ strategies. We proposed a framework for synchronized iterated voting, producing a “Polling Dynamics”, in two flavors: one with a discrete space of states, the other with a continuous space allowing more flexible modeling. We showed that the former can be seen as a particular case of the latter.

Our first main result is that cycles of outcomes are robust under perturbation, so that finding a cycle in the discrete-space setting ensures any small enough variation of the model in the continuous-space setting still exhibit the same cycle.

Then we produced two examples in Approval voting, where voters apply *simple, consistent, sincere* heuristics but cycles with sub-optimal or outright bad outcomes appear. The first example somewhat mitigate an important result of Laslier: while the equilibria of the Polling Dynamics following from his Leader Rule, as he proved, elect the Condorcet winner when she exist, they may fail to attract most of the possible outcomes.

We performed *in silico* experiments to assert the prevalence of these electoral conundrums. They are very rare for the Leader Rule, and even impossible when the culture is unidimensional. However a slight relaxation of the Leader Rule makes precisely the unidimensional culture the worst one among those tested, with up to more than 15% of electorate producing a bad cycle.

We thus showed that under Approval Voting, not only convergence to equilibrium may

not happen, but cycles can lead individually sound heuristics to result collectively in the worst possible outcome. We then gave example of bad cycles for other voting systems, in particular Condorcet systems, showing that these issues are not at all specific to Approval Voting.

Last, we considered a simplistic example of continuous-space Polling Dynamics ensuring continuity, i.e. small changes in the expected outcome leads to small changes in the ballots cast. It turned out this model has a chaotic behavior, and we conjectured that chaos is not uncommon at all for continuous-space Polling Dynamics. We supported this conjecture by looking at the influence of some parameters of the model, showing that chaos can be observed in a non-negligible range. The way in which a particular model of voters behavior and a particular preference profile result in either constant, periodic or chaotic patterns is only illustrated here, and would deserve a full study that is way beyond the scope of the present work.

References

- [AGP17] Stéphane Airiau, Umberto Grandi, and Filippo Studzinski Perotto, *Learning agents for iterative voting*, International Conference on Algorithmic Decision Theory, Springer, 2017, pp. 139–152. [1](#)
- [BCE⁺16] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia (eds.), *Handbook of computational social choice*, Cambridge University Press, 2016. [2.1](#)
- [BCMP13] Simina Brânzei, Ioannis Caragiannis, Jamie Morgenstern, and Ariel D Procaccia, *How bad is selfish voting?*, Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence, 2013, pp. 138–144. [1](#)
- [BF07] Steven Brams and Peter C Fishburn, *Approval voting*, Springer Science & Business Media, 2007. [1](#), [1](#), [5](#)
- [BF10] Steven J Brams and Peter C Fishburn, *Going from theory to practice: the mixed success of Approval Voting*, Handbook on Approval Voting, Springer, 2010, pp. 19–37. [1](#)
- [BHY14] Clark Bowman, Jonathan K Hodge, and Ada Yu, *The potential of iterative voting to solve the separability problem in referendum elections*, Theory and decision **77** (2014), no. 1, 111–124. [1](#)
- [Bla86] Duncan Black, *The theory of committees and elections*, Springer, 1986, Originally published 1958. [4.4](#)
- [BS03] Steven J. Brams and M. Remzi Sanver, *Voter sovereignty and election outcomes*, <http://dev.gtcenter.org/Archive/Conf03/Sovereignty.pdf>, 2003. [1](#)

- [CPP04] Samir Chopra, Eric Pacuit, and Rohit Parikh, *Knowledge-theoretic properties of strategic voting*, European Workshop on Logics in Artificial Intelligence, Springer, 2004, pp. 18–30. [1](#), [5](#)
- [DMN14] François Durand, Fabien Mathieu, and Ludovic Noirie, *Making most voting systems meet the condorcet criterion reduces their manipulability*, <https://hal.inria.fr/hal-01009134/file/condorcification.pdf>, 2014. [5.3](#)
- [DMN16] ———, *Can a condorcet rule have a low coalitional manipulability?*, ECAI 2016, IOS Press, 2016, pp. 707–715. [5.3](#)
- [Dur15] François Durand, *Vers des modes de scrutin moins manipulable*, Ph.D. thesis, Université Pierre et Marie Curie, 2015. [1](#)
- [GATC16] James Green-Armytage, T Nicolaus Tideman, and Rafael Cosman, *Statistical evaluation of voting rules*, Social Choice and Welfare **46** (2016), no. 1, 183–212. [5.3](#)
- [GHRS19] Umberto Grandi, Daniel Hughes, Francesca Rossi, and Arkadii Slinko, *Gibbard–satterthwaite games for k -approval voting rules*, Mathematical Social Sciences **99** (2019), 24–35. [1](#)
- [Gib73] Allan Gibbard, *Manipulation of voting schemes: a general result*, Econometrica: journal of the Econometric Society (1973), 587–601. [1](#)
- [GLR⁺13] Umberto Grandi, Andrea Loreggia, Francesca Rossi, Kristen Brent Venable, and Toby Walsh, *Restricted manipulation in iterative voting: Condorcet efficiency and borda score*, International Conference on Algorithmic Decision-Theoretic Theory, Springer, 2013. [1](#)
- [KH95] Anatole Katok and Boris 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. [6](#)
- [KSLR17] Aaron Koolyk, Tyrone Strangway, Omer Lev, and Jeffrey S. Rosenschein, *Convergence and quality of iterative voting under non-scoring rules*, Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, 2017. [1](#)
- [Las09] Jean-François Laslier, *The Leader Rule: a model of strategic approval voting in a large electorate*, Journal of Theoretical Politics **21** (2009), no. 1, 113–136. [1](#), [1](#), [4.3](#), [4.4](#)
- [Las10] ———, *In silico voting experiments*, Handbook on approval voting, Springer, 2010, pp. 311–335. [4.3](#)

- [LLV00] Dominique Lepelley, Ahmed Louichi, and Fabrice Valognes, *Computer simulations of voting systems*, *Advances in Complex Systems* **3** (2000), no. 01n04, 181–194. [4.3](#)
- [LR16] Omer Lev and Jeffrey S. Rosenschein, *Convergence of iterative scoring rules*, *J. Artif. Intell. Res.* **57** (2016), 573–591. [1](#)
- [Mei15] Reshef Meir, *Plurality voting under uncertainty*, *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI Press, 2015, pp. 2103–2109. [1](#)
- [Mei17] ———, *Iterative voting*, *Trends in computational social choice* (2017), 69–86. [1](#)
- [Mei18] ———, *Strategic voting*, *Synthesis Lectures on Artificial Intelligence and Machine Learning* **13** (2018), no. 1, 1–167. [2](#), [1](#)
- [MLR14] Reshef Meir, Omer Lev, and Jeffrey S. Rosenschein, *A local-dominance theory of voting equilibria*, *Proceedings of the Fifteenth ACM Conference on Economics and Computation (New York, NY, USA)*, ACM, 2014, pp. 313–330. [1](#)
- [MPRJ17] Reshef Meir, Maria Polukarov, Jeffrey S. Rosenschein, and Nicholas R. Jennings, *Iterative voting and acyclic games*, *Artificial Intelligence* **252** (2017), 100–122. [1](#)
- [OLP⁺15] Svetlana Obraztsova, Omer Lev, Maria Polukarov, Zinovi Rabinovich, and Jeffrey S. Rosenschein, *Farsighted voting dynamics*, *1st Workshop on Algorithmic Game Theory at IJCAI, 2015*. [1](#)
- [OMP⁺15] Svetlana Obraztsova, Evangelos Markakis, Maria Polukarov, Zinovi Rabinovich, and Nicholas R Jennings, *On the convergence of iterative voting: how restrictive should restricted dynamics be?*, *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI Press, 2015, pp. 993–999. [1](#)
- [RE12] Annemieke Reijngoud and Ulle Endriss, *Voter response to iterated poll information*, *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2, International Foundation for Autonomous Agents and Multiagent Systems*, 2012, pp. 635–644. [1](#), [3.2](#)
- [ROL⁺15] Zinovi Rabinovich, Svetlana Obraztsova, Omer Lev, Evangelos Markakis, and Jeffrey S. Rosenschein, *Analysis of equilibria in iterative voting schemes*, *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI Press, 2015. [1](#)
- [Sat75] Mark Allen Satterthwaite, *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, 187–217. [1](#)

- [SDL06] Francesco De Sinopoli, Bhaskar Dutta, and Jean-François Laslier, *Approval voting: three examples*, *Int. J. Game Theory* **35** (2006), 27–38. [1](#)
- [SS04] Murat R Sertel and M Remzi Sanver, *Strong equilibrium outcomes of voting games are the generalized condorcet winners*, *Social Choice and Welfare* **22** (2004), no. 2, 331–347. [1](#)
- [Wal82] Peter Walters, *An introduction to ergodic theory*, Graduate Texts in Mathematics, vol. 79, Springer-Verlag, New York-Berlin, 1982. MR 648108 [6.2](#)