



HAL
open science

Iterative voting with partial preferences

Zoi Terzopoulou, Panagiotis Terzopoulos, Ulle Endriss

► **To cite this version:**

Zoi Terzopoulou, Panagiotis Terzopoulos, Ulle Endriss. Iterative voting with partial preferences. Artificial Intelligence, 2024, 332, 10.1016/j.artint.2024.104133 . hal-04739729

HAL Id: hal-04739729

<https://hal.science/hal-04739729v1>


Submitted on 16 Oct 2024

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.



Distributed under a Creative Commons Attribution - NonCommercial - NoDerivatives 4.0 International License



Iterative voting with partial preferences

Zoi Terzopoulou^{a,*}, Panagiotis Terzopoulos^b, Ulle Endriss^c

^a GATE, Saint-Etienne School of Economics, Jean Monnet University, Saint-Etienne, France

^b Deepcure Inc., Boston, MA, USA

^c Institute for Logic, Language and Computation, University of Amsterdam, Amsterdam, the Netherlands

ARTICLE INFO

Keywords:

Social choice theory
Iterative voting
Partial preferences

ABSTRACT

Voting platforms can offer participants the option to sequentially modify their preferences, whenever they have a reason to do so. But such iterative voting may never converge, meaning that a state where all agents are happy with their submitted preferences may never be reached. This problem has received increasing attention within the area of computational social choice. Yet, the relevant literature hinges on the rather stringent assumption that the agents are able to rank all alternatives they are presented with, i.e., that they hold preferences that are linear orders. We relax this assumption and investigate iterative voting under partial preferences. To that end, we define and study two families of rules that extend the well-known k -approval rules in the standard voting framework. Although we show that for none of these rules convergence is guaranteed in general, we also are able to identify natural conditions under which such guarantees can be given. Finally, we conduct simulation experiments to test the practical implications of our results.

1. Introduction

Models of collective decisions constitute the foundation of every democratic society. Nowadays, such models must also account for decisions by AI-powered agents made on our behalf rather than by human agents directly, and they must provide sufficient flexibility for decision processes that take place in online forums rather than physical environments. The primary means for collective decision making is *voting*, where agents express their preferences on a number of alternatives and the winners are selected based on the aggregation of these preferences [38].

We speak of *iterative voting* when voting takes place over a series of rounds in which the participating agents have the opportunity to repeatedly update the preferences they report. A well-known example is the `Doodle` tool for choosing a date for a meeting between several people, which allows users to modify their reported preferences over time. The analysis of iterative voting within the field of *computational social choice* was initiated over a decade ago by Meir et al. [21] and has been a hot topic ever since [30,29,11,5,27,20,26,28,18,9,35,14,22,37,12].

A central feature that all of the models of iterative voting studied in the literature have in common is the assumption that all agents hold and report preferences that are linear orders, i.e., complete and transitive rankings over the full set of alternatives. While this assumption is in line with the classical models of voting studied in social choice theory [38] and while this assumption undoubtedly has its place for the analysis of certain applications, it also clearly represents a significant deviation from reality when considering what kinds of preferences agents actually hold and when contemplating the design of innovative systems for collective

* Corresponding author.

E-mail address: zoi.terzopoulou@cnrs.fr (Z. Terzopoulou).

decision making that would allow agents to express richer forms of preference. Specifically, it is reasonable to assume that an agent may in fact be unable to compare some of the alternatives, thereby causing an intrinsic incompleteness in her preferences. For example, an agent may not have sufficient information to rank two alternatives, she may lack the computing power to fully evaluate that information even when it is available, or she may simply not be interested in part of the space of alternatives. Indeed, this phenomenon manifests itself in large parts of the `PreFlib` collection of real-world preference data [19]. Therefore, it is essential to consider voting models that simultaneously accommodate iteration and partial preferences, and this paper fills this gap. Note, however, that studying the concepts of incomplete information and iterative voting in parallel is not a totally novel idea. Dery et al. [6] explored iterative voting under incompleteness, but from the mechanism designer's rather than the agents' point of view. In our setting, in contrast, agents are assumed to hold intrinsically incomplete preferences, and there is no noise on the information that the mechanism designer receives. As far as the concrete voting rules being studied are concerned, most work on iterative voting to date has focused on the important family of *positional scoring rules*. Some exceptions exist, such as the work of Koolyk et al. [14], who also studied non-scoring rules—namely, Copeland, Maximin, Bucklin, STV, and ranked pairs—showing that convergence is not guaranteed under various non-standard heuristic that agents may employ when changing their preferences. Another exception is met in the paper of Endriss et al. [9], who showed that Copeland and Maximin (as well as positional scoring rules) always converge given that agents have limited information about each others' preferences and only know the winning alternative in every round. In this paper, we start from the basics and put positional scoring rules on the spotlight. These are simple voting rules that assign points to the alternatives depending on their position in the agents' preferences. Examples include the *plurality rule*, where an alternative receives a point whenever it is ranked first by an agent, and the *antiplurality rule*, where an alternative receives a point whenever it is *not* ranked last by an agent. Both belong to the family of *k-approval rules*, under which an alternative receives a point whenever it is ranked amongst an agent's top k alternatives. So plurality is 1-approval and antiplurality is $(m - 1)$ -approval, where m is the total number of alternatives. Interestingly, generalising the family of *k-approval rules* so as to fit a model of voting with partial preferences is not straightforward. To see why, observe that for the special case of linear preferences we can interpret the point allocation mechanism defining *k-approval* either as “approving the top k alternatives” in any given linear order or as “vetoing the bottom $(m - k)$ alternatives”.¹

But for partial orders, these two notions do not coincide any longer. To account for this fact, we introduce two families of voting rules for partial orders, the *k-approval rules* and the *k-veto rules*. Our analysis of iterative voting pertains to these two families.

More generally, the literature on voting for partial preferences is not as rich as the one for complete preferences. Yet, some concrete suggestions for voting rules tailored to partial preferences have been made, many of which focus on *k-truncated preferences* (that is, linear orders over a subset of the alternatives, assuming that the unranked alternatives are inferior to the ranked ones). Note that since truncated preferences are a special case of partial preferences, definitions corresponding to them do not directly apply in our framework. Baumeister et al. [3] look at campaigning problems for truncated preferences (e.g., the benefits that may stem from convincing the agents to rank more alternatives) and study the optimistic and pessimistic variants of the Borda rule (i.e., a specific positional scoring rule that we are not considering here); Emerson [8] focuses on the same generalisations of the Borda rule, exploring the strategic incentives each of them generates for the agents. Terzopoulou and Endriss [32] work with incomplete preferences that consist of pairwise comparisons between alternatives and introduce Kemeny-like rules that assign weights to these preferences according to their size. Also, Narodytska and Walsh [23] study strategic voting from a computational perspective, and examine scoring rules (along with STV and Copeland) for truncated preferences. A series of papers further away from this work are also interested in approximations of classical rules for *k-truncated preferences* and study the degree to which they agree with the original ones for linear orders [10,2,4,1].

Now, the basic scenario of iterative voting is the following. Initially, each agent reports a preference regarding the available alternatives. Then, after observing the reports of their peers, agents can choose to—*one-by-one*—modify their reported preferences in order to obtain a better outcome. This process is repeated *ad infinitum*—or, preferably, until we reach a point where no agent wants to make any further updates. In line with most previous literature, we assume that agents are *strategic* but *myopic*: they only attempt to improve the outcome in the immediately next round, without thinking further ahead. The main question we are interested in concerns the problem of convergence:

Do iterative voting processes with partial preferences reach stable states, where no agent has an incentive to (or is able to, given the types of moves the system permits her to make) change her reported preference?

Understanding when convergence can be guaranteed is crucial for the design and effective use of collective decision making tools. If convergence cannot be guaranteed, then it may be impossible to reach a final decision.

In a framework that allows for incompleteness, not only the sincere preferences of the agents but also their reported ones may constitute partial orders. As observed by Kruger and Terzopoulou [15], this grants a great degree of freedom on the part of agents who want to update their preferences. They can alter their preferences by (i) *adding*, (ii) *omitting*, or (iii) *flipping* pairwise comparisons.

¹ It is important to note that it is not the agent who approves or vetoes certain alternatives. Rather, this is how a given voting rule *interprets* the preference relation reported by the agent. Indeed, as is well understood in social choice theory, one should not confuse the reported preferences of an agent with the allocation of points to alternatives the voting rule translates that preference to. For example, it is not the case that, just because we use, say, the 3-approval rule to take a decision, every agent associates their three most preferred alternatives with a utility of 1 and all others with a utility of 0. We want to maintain this distinction between ordinal preferences on the one hand, and the way in which those ordinal preferences are translated into point allocations on the other—also when working with partial preferences. This is why we ask agents to report preference relations rather than to allocate points to alternatives directly.




Fig. 1. A strict partial preference represented as a directed acyclic graph.

So, on top of analysing the general convergence qualities of voting rules and comparing our results to those from the literature for the special case of linear preferences, we also investigate the impact of restricting the permitted types of preference updates on convergence. Notably, we find that none of the rules we analyse is guaranteed to always converge—a finding that stands in direct contrast with seminal convergence results by Meir et al. [21], Lev and Rosenschein [17], and Reyhani and Wilson [30] regarding plurality and antiplurality for linear preferences—but also that for all of the rules we study we can attain such an assurance under certain restrictions, which we characterise.

In the literature on iterative voting with linear preferences various other assumptions have been modified to facilitate convergence results, typically concerning behavioural characteristics of the agents, such as *truth-biasedness* [27,28], *farsightedness* [26], bounded computational power [11], *laziness* [7,28], and *incomplete information* [29,9]. These restrictions differ from the ones we consider in that they cannot be easily enforced by an external mechanism designer. On the contrary, enforcing our (combinations of) restrictions can easily be enforced on a voting platform, by simply disallowing the agents to perform certain changes in their reported preferences.

For every convergence result we obtain, we also establish the maximum number of rounds required for a stable state to emerge. Finally, we have generated synthetic preference data and conducted simulation experiments to check how often our theoretical results regarding the absence of guaranteed convergence for certain settings have negative implications in practice. Fortunately, we find that this happens only rarely.

The remainder of this paper is structured as follows. Section 2 introduces our model, while Sections 3 and 4 contain our theoretical results for the approval and the veto family of rules, respectively. Section 5 discusses our experiments and Section 6 concludes. The supplementary material for this paper [34] includes both the code required to run our experiments and the data generated during those experiments.

2. The model

In this section we present our basic voting model with partial preferences and also describe our iterative framework.

2.1. Voting preliminaries

We have a group of agents $N = \{1, \dots, n\}$, with $n \geq 2$, and a finite set of m alternatives $A = \{a, b, c, \dots\}$. Every agent i truthfully holds a *strict partial preference* \triangleright_i , which is an *irreflexive* and *transitive* binary relation over A : $(a, b) \in \triangleright_i$ means that agent i prefers alternative a over alternative b (or similarly, that a is superior to b and that b is inferior to a). We call \mathcal{D} the domain of all strict partial preferences. \mathcal{D} includes linear orders, i.e., preferences where all alternatives are ranked.²

We can graphically represent partial preferences as directed acyclic graphs with m nodes, where each node is labelled by a unique alternative and a directed edge from a to b means that a is more preferred than b (see Fig. 1). The *position* of an alternative in a partial preference refers to the node that the alternative labels in the graphical representation of the preference (i.e., the topological order of the corresponding directed acyclic graph). Given a partial preference \triangleright , we say that two alternatives are *isomorphically equivalent* if the nodes they label coincide under some isomorphism of the graph corresponding to \triangleright . For example, in Fig. 1 the alternatives a and c (as well as the alternatives b and d) are isomorphically equivalent.

A *profile* is a vector $\triangleright = (\triangleright_1, \dots, \triangleright_n) \in \mathcal{D}^n$ consisting of the preferences of all agents in the group, and $(\triangleright_{-i}, \triangleright_i)$ is the new profile where all agents besides i report the same preferences as in \triangleright , while agent i reports the preference \triangleright_i . This setting clearly extends the standard voting framework where all agents are assumed to hold linear orders over the set of alternatives.

A *voting rule* is a function that maps every profile $\triangleright \in \mathcal{D}^n$ to a non-empty set of *winners* in A (that is, a subset of A). A *positional scoring rule* F_s is associated with a scoring function s that assigns a score $s_{\triangleright_i}(x) \in \mathbb{R}$ to every alternative x depending on its position in an agent's preference \triangleright_i ; if the alternatives x and y are isomorphically equivalent in a preference \triangleright_i , then it must hold that $s_{\triangleright_i}(x) = s_{\triangleright_i}(y)$ [15]. Then, the alternatives with the highest summed score over all agents' preferences are selected as winners. Formally, we have the following:

$$F_s(\triangleright) = \operatorname{argmax}_{x \in A} \sum_{i \in N} s_{\triangleright_i}(x)$$

Later on, given a profile \triangleright , we will capture the scores of the m alternatives in A (in alphabetical order) with a vector in \mathbb{R}^m , showing the winning score in bold.

Here, we consider two families of positional scoring rules in the general context of partial preferences—the *approval* family and the *veto* family—that extend classical rules from standard voting with linear orders. Approval-type rules treat some of the most preferred alternatives of an agent as her approved ones, while veto-type rules consider some of the least preferred alternatives of the

² Note that in our setting truthful preferences *can* be complete, i.e., linear orders. Indeed, for several of the examples in this paper agents originally hold linear preferences and report partial ones only during the process of iteration.




Fig. 2. The scores depicted below the preference correspond to the following rules: 1-approval (top left), 1-veto (top right), 2-approval and 2-veto (middle), 5-approval (bottom left), 5-veto (bottom right).

agent as those she vetoes. This does not mean that the agent herself approves or vetoes the given alternatives—she is not asked to do so, and possibly she is not even able to do so (recall that we only assume that the agents form pairwise comparisons between some of the alternatives in A , which does not imply that they can themselves determine which alternatives to approve). Yet, voting rules often employ the idea of approval and veto scores as a tool that simplifies the selection of the winners. Approval-type and veto-type rules follow this general approach.

If we make the restrictive assumption that the agents hold complete preferences (i.e., linear orders over A), a classical k -approval voting rule considers approved by an agent all alternatives that are inferior to *at most* $k - 1$ other alternatives in the agent’s preference (these can be thought of as the k “top” alternatives in the linear order). In that setting, it is equivalent for the rule to consider vetoed the $m - k$ alternatives at the “bottom” of an agent’s preference (i.e., the alternatives that are inferior to at least k other alternatives, or equivalently, superior to at most $k - 1$ other alternatives).

But how would a k -approval rule deal with incomparable alternatives? Preserving the spirit of the complete case, let us say that an alternative x is *amongst the top k levels* in an agent’s preference if x is inferior to at most $k - 1$ other alternatives in that preference; analogously, an alternative x is *amongst the bottom k levels* in a preference if x is superior to at most $k - 1$ other alternatives in that preference.

We are now ready to provide our main definitions, of the approval and the veto class of rules for partial preferences.

Definition 1. Given $\triangleright \in D$ and $x \in A$, for $k \in \{1, \dots, m\}$, the **k -approval** rule assigns to x score 1 if x is amongst the top k levels in \triangleright , and score 0 otherwise.

Definition 2. Given $\triangleright \in D$ and $x \in A$, the **k -veto** rule assigns to x score 0 if x is amongst the bottom k levels in \triangleright , and score 1 otherwise.

When 1-approval and $(m - 1)$ -veto are applied to profiles of linear orders, they correspond to the plurality rule, while 1-veto and $(m - 1)$ -approval reduce to the antiplurality rule. For our general partial preferences, the following hold: under 1-approval, an alternative x is approved if and only if there exists no alternative that is superior to x ; under $(m - 1)$ -approval, an alternative x is approved if and only if there exists some alternative that is not superior to x ; under 1-veto, an alternative is approved if and only if there exists some alternative that is inferior to x ; and under $(m - 1)$ -veto, an alternative is approved if and only if all other alternatives are inferior to x .

Moreover, for the special case of profiles of linear orders, k -approval and $(m - k)$ -veto coincide (this is not true in general for partial preferences; they entail that *precisely* k alternatives are approved (while for general partial preferences, k -approval implies that *at least* k alternatives are approved)—see Fig. 2 for examples.

In order to ensure a unique winning alternative, we always enforce a lexicographic tie-breaking rule that, given a voting outcome $S \subseteq A$, selects the alternative a when $a \in S$, the alternative b when $b \in S$ and $a \notin S$, and so on.

2.2. Iterative processes

Given an arbitrary voting rule F , a process of iterative voting takes place in rounds. In every round, a profile is given and a winner is determined through F . We assume the following:

- ◊ In the very beginning (round 1), all agents are sincere. (assumption 1)

Assumption 1 is reasonable, given that, except for certain pathological cases involving contrived voting rules, agents with no information about the preferences of their peers have no incentive to submit an untruthful preference [29].³ As Meir et al. [21] also write, “it is often the case that the initial state is truthful, as agents know that they can reconsider and vote differently, if they are not happy with the current outcome”. In the literature of iterative voting with linear preferences, convergence results are often given for arbitrary initial states [21,18]. We will later see that even if we restrict the starting preference profiles to sincere ones, several of these results fail under partial preferences (while some of our convergence results also trivially hold for arbitrary initial profiles). Note also that studying convergence from sincere initial states is frequently done in existing literature on iterative voting [29,11,20].

Subsequently, an agent i with truthful preference \succ_i may change her reported preference from $\succ^t \in \mathcal{D}$ in round t to $\succ^{t+1} \in \mathcal{D}$ in round $t + 1$ only if

$$F(\succ_{-i}^t, \succ^{t+1}) \succ_i F(\succ_{-i}^t, \succ^t),$$

where \succ_{-i}^t is the profile of the rest of the group in round t . Importantly, the preferences submitted by the agents in every round have to be admissible, which we interpret as belonging to \mathcal{D} ; in particular, they have to be transitive. We also make the following assumption:

- ◊ When more than one agent can profit by changing their reported preference, any single one of them may do so. We impose no further constraints on who this agent might be. (assumption 2)

Since we allow for incomplete preferences, an agent i may modify her reported preference $\succ^t \in \mathcal{D}$ to a new one $\succ^{t+1} \in \mathcal{D}$ in three fundamental ways, or any combination thereof [15]:

- *Omission* (removing pairwise comparisons from \succ^t);
- *Addition* (adding pairwise comparisons to \succ^t);
- *Flipping* (swapping pairwise comparisons in \succ^t).

Not all omission, addition, or flipping moves are admissible though. For instance, an agent cannot modify her preference from $\{(a, b), (b, c), (a, c)\}$ to $\{(a, b), (b, c)\}$, since the latter violates transitivity and thus is not a valid partial order. In a voting platform (for instance an online poll), the mechanism designer can guarantee that the agents will not be able to employ certain changes on their reported preferences simply by forbidding the corresponding actions. For example, if an agent has already reported that she deems a better than b , the system can save this preference and not allow the agent to delete or hide it in future rounds. Checking for which (combinations of) restrictions we obtain well-behaved voting patterns is thus important for practical purposes.

The iterative process based on a voting rule F has achieved *convergence* in profile \succ (then called a *stable state*) when no agent has an incentive (or is able) to further change her preference. We also say that the rule F *guarantees* convergence if we always reach a stable state, starting from any truthful profile, for any number of alternatives m , of agents n , and for any order in which the voters may change their preferences. Assumption 2 implies that when we say that an iterative process is guaranteed to reach a stable state, this will include all possible orders in which the agents might act.

Amongst all alternatives that an agent might be able to make win, she will choose her most preferred one (i.e., the agents always choose *best responses*).⁴ After an agent has identified a best response, we additionally assume the following, informally described:

- ◊ An agent will not alter her reported preference more than necessary to achieve a desired outcome: the agents’ moves are *minimally diverging*. (assumption 3)

Assumption 3, based on the idea that *changing one’s preference is costly*, is necessary (yet not always sufficient) for our convergence results. Specifically, flipping one’s preference on a pair of alternatives will be taken to cost more than just removing or just adding the relevant pairwise comparison. This can be externally imposed in application contexts. For example, in an online poll the designer can force the participants to delete an already submitted ranking between two alternatives before they are able to report a different ranking regarding the same alternatives.⁵ Formally, for $\succ, \succ' \in \mathcal{D}$ and $(x, y) \in A^2$, we define $C_{xy}(\succ, \succ')$ to be the *cost* of moving between the two preferences with respect to the given pair (x, y) :

$$C_{xy}(\succ, \succ') = \begin{cases} 0 & \text{if } \succ, \succ' \text{ fully agree on } (x, y); \\ 1 & \text{if } \succ, \succ' \text{ fully disagree on } (x, y); \\ 1/2 & \text{otherwise,} \end{cases}$$

where we say that the two preferences *fully agree* on (x, y) if they both rank x above (or below) y , or if they both deem x incomparable to y ; and they *fully disagree* on (x, y) if one of them ranks x above y , while the other one ranks y above x .

³ Of course, in real-world decision-making scenarios it rarely will be the case that agents really have no information at all. For example, in political elections polls provide some background information. In small-scale online settings, such as choosing a meeting time with `Doodle`, the agents involved may be close personal acquaintances of one another and have at least some idea about the preferences of their peers.

⁴ If the agent cannot compare some of her favourite alternatives that can win, she will nondeterministically select one of them.

⁵ Also, the agents themselves may deem flipping more costly, since lying in that fashion intuitively is more extreme in social contexts.

Table 1

Summary of convergence results for the approval family of rules. “✓” indicates that convergence is guaranteed, while “✗” indicates that an example against convergence exists. When only omission or only addition is allowed, convergence is trivially achieved in a finite number of rounds.

	1-approval	k -approval	$(m - 1)$ -approval
omm	✓trivial	✓trivial	✓trivial
add	✓trivial	✓trivial	✓trivial
flip	✓Prop.2	✗Prop.5	✗Prop.7
omm + add	✗Prop.1	✗Prop.4	✓Prop.9
omm + flip	✓Prop.2	✓Prop.6	✓Prop.8
add + flip	✓Prop.3	✗Prop.5	✗Prop.7
omm + add + flip	✗Prop.1	✗Prop.4	✗Prop.7

For the total cost $C(\triangleright, \triangleright')$ of moving from \triangleright to \triangleright' , we sum over the scores $C_{xy}(\triangleright, \triangleright')$ for all pairs $(x, y) \in A^2$:

$$C(\triangleright, \triangleright') = \sum_{(x,y) \in A^2} C_{xy}(\triangleright, \triangleright')$$

To make an alternative a win in round $t + 1$, an agent will choose among the appropriate preferences the one with the lowest cost from \triangleright'_t .

Lastly, let us write $s_t(x) = \sum_{i \in N} s_{\triangleright'_t}^i(x)$ for the total score that alternative x collects in round t . Suppose that an agent changes the winner from w_t to w_{t+1} . For convenience, we categorise her possible moves into two (exhaustive, not necessarily disjoint) types:

- *Type 1:* $s_{t+1}(w_t) <_\ell s_t(w_t)$. That is, agent i decreases the score of w_t to make some other alternative the winner.
- *Type 2:* $s_t(w_{t+1}) <_\ell s_{t+1}(w_{t+1})$. That is, agent i increases the score of w_{t+1} to make it the winner.

In the above definition, given two alternatives $x, y \in A$, we write $s_t(x) <_\ell s_{t'}(y)$ when either (i) the score of x in round t is strictly smaller than the score of y in round t' , or (ii) the score of x in round t is the same as the score of y in round t' , but the lexicographic tie-breaking rule ranks y above x . For simplicity, we say that the score of x is smaller than the score of y .

3. The approval family

This section explores convergence for the approval family. Table 1 summarises the relevant results and indicates where their proofs can be found.

3.1. 1-approval


We know that the plurality rule—the counterpart of 1-approval for linear orders—is guaranteed to converge when the agents are only allowed to report complete preferences [21]. This result does not survive in richer settings of incompleteness. In fact, there exists a counterexample for convergence of the iterative 1-approval process, where the agents only alter a single pairwise comparison in their sincere preferences.

Proposition 1. *If omission and addition are allowed, then the iterative 1-approval process is not guaranteed to converge (independently of whether flipping is allowed).*

Proof. Consider a sincere preference profile as follows:


- | | |
|--|--|
| Agent 1 : $c \triangleright d \triangleright a \triangleright b$ | Agent 5 : $d \triangleright c \triangleright b \triangleright a$ |
| Agent 2 : $b \triangleright d \triangleright c \triangleright a$ | Agent 6 : $d \triangleright b \triangleright c \triangleright a$ |
| Agent 3 : $a \triangleright d \triangleright b \triangleright c$ | Agent 7 : $c \triangleright b \triangleright d \triangleright a$ |
| Agent 4 : $b \triangleright c \triangleright d \triangleright a$ | Agent 8 : $a \triangleright b \triangleright c \triangleright d$ |

Every time an agent makes a move, she will be altering the position of the alternative ranked second in her sincere preference. For example, consider agent 1 with sincere preference $c \triangleright d \triangleright a \triangleright b$: We write “ $1d \uparrow$ ” when agent 1 moves alternative d on top by omission to make d win instead of a , and “ $1d \downarrow$ ” when she removes alternative d from the top by addition and submits her sincere preference to make c win, as illustrated below (note that the dashed arrows illustrate the direction of the manipulation move from one preference to another, while the solid arrows are part of the preference graph—this convention is adopted throughout the paper).



In the sincere preference of agent 1, alternative c gets 1 point and all other alternatives get 0 points; in the manipulated preference, alternatives c and d get 1 point each and all other alternatives still get 0 points.

Starting with the score-vector on the top-left corner of the figure below, the first seven agents will participate in a cycle.



A counterexample for convergence thus exists. \square

We can, though, obtain convergence if we restrict the ways the agents are allowed to change their preferences. Our proofs (and certain proofs in later parts of this paper too) are akin to those of Reyhani and Wilson [30] for the plurality and antiplurality rules under linear preferences, showing that the set of potentially winning alternatives decreases along a path of best responses. Some of the ideas we develop—for example the fact that after a flipping move the non-winning alternative loses its chance to become a winner in the future—also featured in the original work of Meir et al. [21].

We state a few results in this section for the $(m - 1)$ -veto in parallel with 1-approval (and will call upon them later on, in the veto section), whenever the reasoning for the two rules is analogous. Proposition 2 will rely on the assumptions of transitivity and of minimal divergence, while this will not be the case for Proposition 3.⁶ We will often use Lemma 1, which is a consequence of the definitions of our rules.

Lemma 1. *The first manipulation move of every agent under the iterative 1-approval process or the iterative $(m - 1)$ -veto process will be of type 2.*

Proof. Let us prove the claim for the 1-approval rule (the reasoning is analogous for $(m - 1)$ -veto). Suppose that alternative x is the winner in round t , where agent i holds her truthful preference \succ_i^t . Further suppose that agent i can make an alternative y such that $y \succ_i^t x$ the winner in round $t + 1$. By definition of 1-approval, if $s_{\succ_i^t}(x) = 1$, it would be impossible to have that $y \succ_i^t x$. This means that $s_{\succ_i^t}(x) = 0$, so agent i cannot decrease the score of the winner in round $t + 1$. Necessarily, her move will be of type 2. \square

Proposition 2. *If addition is not allowed, then the iterative 1-approval process is guaranteed to converge.*

Proof. We will show by induction that if addition is not allowed, then an arbitrary agent i with sincere preference \succ_i will also never employ flipping. She will be changing her submitted preference only by omission, making a new alternative the winner each time. Since this can happen for at most finitely many rounds, convergence follows.

⁶ Both in Proposition 2 and in Proposition 3, we prove that the agents will only perform moves of type 2, and convergence is a logical consequence of this (in the former case, because omission moves can be only done finitely many times; in the latter case, because making a new alternative the winner can be only done finitely many times). But the underlying reasons for our proofs' validity are very different (linked to either the absence of the ability or the absence of the interest to make the relevant manipulation moves, respectively). On the one hand, when addition is not allowed, the agents would only want to perform a move of type 1 by flipping, but all relevant moves would violate transitivity and would thus not be admissible. On the other hand, when omission is not allowed, the agents' incentives will always concern moves of type 2, and their desired actions will only employ flipping.

Our induction basis is established by Lemma 1: We know that a move of type 2 will be done by omission, since the minimal divergence assumption prioritises omission over flipping. Suppose now that the statement holds for the k^{th} time agent i changes her preference. We will show that it also holds for the $(k + 1)^{\text{th}}$ time she changes her preference. Say this happens in round t .

Case 1: Agent i 's move in round t is of type 2 (and not of type 1). Then, omission will be employed as for round 1, and a new alternative will be placed on top.

Case 2: Agent i 's move in round t is of type 1, making alternative y the new winner, instead of a previous winner x . Since addition is not allowed, the only way that x can have its score decreased is when x is a top alternative in \triangleright_i^t and agent i flips some pairwise preference $x \triangleright_i^t z$ to $z \triangleright_i^{t+1} x$. Note that $y \triangleright_i x$ and $x \triangleright_i z$ (because up until round t the agent has only omitted pairwise preferences), and $y \triangleright_i z$ by transitivity. Since in round t alternative z is not on top for agent i , $y \triangleright_i^t z$ holds because the induction hypothesis tells us that the only pairwise comparisons of the form $x_1 \triangleright x_2$ that have been omitted are such that x_2 is forced to win, and we are not allowing addition. Thus, since $x \triangleright_i^t z$ exists, removing $y \triangleright_i^t z$ would have been pointless. Hence, $y \triangleright_i^{t+1} z$ holds. But $y \triangleright_i^{t+1} z$ and $z \triangleright_i^{t+1} x$ imply $y \triangleright_i^{t+1} x$, by transitivity. This is impossible, since x is a top alternative in round t and addition is not allowed. \square

Proposition 3. *If omission is not allowed, then the iterative 1-approval process and the iterative $(m-1)$ -veto process are guaranteed to converge.*

Proof. We will prove the following claim: Every time an agent changes her preference, she promotes a winning alternative that is different from those she has promoted before. Because the set of alternatives A is finite, convergence follows.

Suppose that agent i is the first to change her reported (sincere) preference \triangleright_i in round 1. By Lemma 1, her move will be of type 2, aiming at increasing the score of some alternative y . For the 1-approval rule, this move will be done by flipping: Agent i will obtain a new pairwise preference with $y \triangleright_i^2 x$, where x was a top alternative in \triangleright_i that was not winning. For the $(m - 1)$ -veto rule, the move may also be done by flipping (which will concern alternatives x and y as described for 1-approval), or it may be done by addition (which means that y was a top alternative in \triangleright_i). We show the following:

(*) Every alternative y' that agent i sincerely places above y (i.e., $y' \triangleright_i y$) has no chance of winning in later rounds unless the score of the winner decreases.

Condition (*) is clearly true for alternative x , which loses one point after i 's move; it is also true for every non-top alternative z such that $z \triangleright_i y$ (if such z exists), since z could not become a winner in round 2 (we know that because agent i submitted a best response), and the score of z remained the same after i 's move. Lastly, suppose that there exists a top alternative r such that $r \triangleright_i y$ and r did not have its score decreased after agent i 's move: that is, it is not the case that $y \triangleright_i^2 r$. As omission is not allowed, it should also hold that $r \triangleright_i^2 y$, which is impossible since our rules must assign y score 1 in round 2. This concludes the proof of Condition (*).

Now, as is the case for agent i , every agent j 's first move will be of type 2 and condition (*) will also hold for j . So the winning score will keep increasing. From (*) it follows that when an agent modifies her reported preference for the second time, she will have no incentive to use a type-1 move to decrease the score of the winner, because no better alternative will be able to win. This means that each agent who performs a second move will promote another alternative than the one she promoted in her first move, increasing its score and making it the winner. Continuing with this reasoning, the same will be true for everyone who performs a third move, and so on. \square

Corollary 1. *The iterative 1-approval process is guaranteed to converge if and only if omission is not allowed or addition is not allowed.*

The proofs of Proposition 2 and Proposition 3 show that when omission (but not addition) is allowed or when addition and flipping (but not omission) are allowed, a stable state will be reached by the iterative 1-approval process after at most nm rounds: Every time an agent changes her preferences, she moves a new alternative on top to make it win. When only addition is allowed, every sincere profile is a stable state.

3.2. k -approval, for $1 < k < m - 1$

By adjusting the non-convergence example for the 1-approval process under omission and addition, we construct one for k -approval, where $1 < k < m - 1$.

Proposition 4. *If omission and addition are allowed, then the iterative k -approval process is not guaranteed to converge (independently of whether flipping is allowed).*

Proof. Fix some $k \in \{2, \dots, m - 2\}$, and consider alternatives e_1, \dots, e_{k-1} . In the profile provided in the proof of Proposition 1, place all new alternatives immediately after the top alternative of each agent. For example, for $k = 3$ the preference of agent 1 becomes $c \triangleright e_1 \triangleright e_2 \triangleright d \triangleright a \triangleright b$. If needed, add some more agents to guarantee that the new alternatives have lower score than all other alternatives in the profile. \square

Moreover, if we forbid omission, we find an existing non-convergence example for the iterative k -approval process where flipping is employed in the work of Lev and Rosenschein [18]: proof of Theorem 4, Part 1. Proposition 5 is thus true.

Proposition 5. *If flipping but not omission is allowed, then the iterative k -approval process is not guaranteed to converge (independently of whether addition is allowed).*

The conditions under which convergence of k -approval is guaranteed are more elaborate than those related to 1-approval. Particularly, in Proposition 6 we need to precisely describe the combinations of all allowed types of moves, as opposed to simply excluding addition (Proposition 2) or omission (Proposition 3).

Proposition 6. *If omission and flipping are allowed, but addition is not, then the iterative k -approval process is guaranteed to converge, for all $1 < k < m - 1$.*

Proof. Suppose that a cycle exists, aiming for a contradiction. The agents that participate in it will not employ omission, since such moves cannot be repeated when addition is not allowed. Then, all moves in the cycle will employ flipping. Moreover, we claim that all moves in the cycle will be of type 1. To prove the claim, first note that flipping in some round t of the cycle will decrease the score of some alternative a in order to make another alternative win. Then if $a \neq w_t$, the new winner could simply be obtained by omission (because of minimally diverging moves). So $a = w_t$, which is the definition of type 1.

Thus, in every round t of the cycle the winner w_t has its score decreased (precisely, $s_{t+1}(w_t) <_\ell s_t(w_t)$), and some alternative has its score increased by one point. Let us now prove that the following holds:

$$s_{t+1}(w_{t+1}) <_\ell s_t(w_t)$$

The reasoning is as follows: We know that w_t is the winner of round t that must have its score strictly decreased in order for w_{t+1} to be the winner of round $t + 1$, and this is done by flipping. If $s_t(w_{t+1}) = s_t(w_t)$, it means that w_t wins because it is lexicographically prioritized to w_{t+1} , but has the same score as it in round t . But then, from the minimal divergence assumption, the voter should instead have decreased the score of w_t by omission without changing the score of w_{t+1} . This is impossible since we have excluded omission. So it must be the case that $s_t(w_{t+1}) < s_t(w_t)$. If now w_t is lexicographically inferior to w_{t+1} , again the voter should instead have decreased the score of w_t by omission without changing the score of w_{t+1} , which we have excluded. So the only option is that $s_t(w_{t+1}) < s_t(w_t)$ and w_t is lexicographically prioritised to w_{t+1} . Then, $s_t(w_{t+1}) + 1 <_\ell s_t(w_t)$. Finally, because the score of an alternative under the k -approval voting rule cannot gain more than one point with a single move, we have that $s_{t+1}(w_{t+1}) \leq s_t(w_{t+1}) + 1$. This leads us to the conclusion: $s_{t+1}(w_{t+1}) <_\ell s_t(w_t)$.

Moreover, there must exist at least one alternative y that has its score increased an infinite number of times (since the cycle includes infinitely many steps, but the alternatives are finitely many). Then, alternative y must also have its score decreased an infinite number of times (otherwise it would be impossible for the score of the winner to always decrease). More specifically, every round t of the cycle where y has its score decreased must be followed by some round $t' > t$ where y has its score increased. Consider such rounds t and t' , where y does not have its score altered in rounds between t and t' . For all $x \in A \setminus \{y\}$, the following holds:

$$\begin{aligned} s_{t+1}(x) &\leq_\ell s_{t+1}(w_{t+1}) <_\ell s_t(w_t) \\ s_t(w_t) &= s_t(y) = s_{t+1}(y) + 1 \end{aligned}$$

Alternative y may have its score increased from round t' to round $t' + 1$ only if $y \neq w_{t'}$. Then, we have the following:

$$s_{t'}(x) \leq_\ell s_{t'}(w_{t'}) \leq s_t(w_{t+1}) <_\ell s_t(w_t) = s_{t'}(y) + 1$$

Also, it holds that $s_{t'+1}(x) \leq_\ell s_{t'}(x)$, for all $x \in A \setminus \{y\}$. This means that if y has its score increased in round t' (that is, $s_{t'+1}(y) = s_{t'}(y) + 1$), y will be the winner in round $t' + 1$; but the agent can achieve this simply by increasing the score of y by one point, without changing the score of other alternatives. Given minimally diverging moves, this should happen by omission, which implies a contradiction. \square

Keeping in mind that if only omission or only addition is allowed, then convergence is straightforward, we state Corollary 2.

Corollary 2. *The iterative k -approval process is guaranteed to converge, for all $1 < k < m - 1$, if and only if omission and addition are not allowed simultaneously, and flipping is allowed only when omission is allowed as well.*

Interestingly, under the assumption of minimally diverging moves, allowing more freedom to the agents—that is, omission on top of flipping—favours convergence.

When flipping and omission (but not addition) are allowed, iterative k -approval will reach a stable state after at most n^2m^2 rounds: Every agent performs omission at most m times, each of which with the purpose of increasing the score of a new alternative; between every two of these nm steps, at most nm flipping rounds materialise, where the score of the winner decreases and the score of some non-winner increases.

When only omission or addition is allowed, convergence occurs after at most nm rounds: Every agent’s move consists of changing the score of a new alternative.

3.3. $(m - 1)$ -approval


Probably not surprisingly, we also find a counterexample for convergence concerning the $(m - 1)$ -approval process.

Proposition 7. *If flipping but not omission is allowed, then the iterative $(m - 1)$ -approval process is not guaranteed to converge (independently of whether addition is allowed as well). The same holds if omission, addition, and flipping are all allowed.*

Proof. Consider a sincere preference profile as follows:


- Agent 1 : $c \triangleright a \triangleright b \triangleright d \triangleright e$
- Agent 2 : $d \triangleright c \triangleright b \triangleright a \triangleright e$
- Agent 3 : $a \triangleright d \triangleright b \triangleright c \triangleright e$
- Agent 4 : $b \triangleright a \triangleright c \triangleright d \triangleright e$
- Agent 5 : $b \triangleright d \triangleright a \triangleright c \triangleright e$
- Agent 6 : $b \triangleright c \triangleright d \triangleright a \triangleright e$
- Agent 7 : $a \triangleright c \triangleright d \triangleright e \triangleright b$

Every time an agent makes a move, she will be altering the position of the alternative ranked second in her sincere preference. For example, consider agent 1 with sincere preference $c \triangleright a \triangleright b \triangleright d \triangleright e$. The first time that agent 1 performs a manipulation act, we write “ $1a\downarrow$ ” when she moves alternative a at the bottom by flipping and submits the (minimally diverging) preference $c \triangleright b \triangleright d \triangleright e \triangleright a$ to prevent a be the winner instead of c . Then, we write “ $1a\uparrow$ ” when she removes alternative a from the bottom to make a win instead of b , submitting the preference $c \triangleright b \triangleright d \triangleright a \triangleright e$. Finally, when agent 1 needs to decrease the score of a again (for which we also write “ $1a\downarrow$ ”), she will submit the preference $c \triangleright b \triangleright d \triangleright e \triangleright a$. In the cycle, she will keep alternating between the two insincere preferences (see below). Note that the notation $kx\downarrow$ is used for a move of agent k that decreases the score of alternative x , without corresponding to unique starting and ending preferences.



For all preferences of agent 1, all alternatives get 1 point each except for the bottom alternative that gets 0 points.

The first six agents are going to participate in a cycle, starting from the score-vector on the top-left corner of the figure below. Note that the cycle only starts from the second iteration of the described sequence of deviations.



It is not difficult to convince ourselves that when all three types of omission, addition, and flipping are allowed, the above counterexample continues to constitute a non-convergence case (where, for example, agent 1 will first employ flipping and then will alternate between increasing the score of a by omitting the pair (e, a) and decreasing the score of a by addition). \square

The non-convergence result for $(m - 1)$ -approval when only flipping is allowed (even when the agents all hold linear orders) contrasts with the fact that the antiplurality rule in standard voting always converges [18]. This happens because in the latter framework it is assumed that when an agent changes her reported preference, she always ranks the current winner in the lowest position, while this is not the case under our minimal divergence assumption. Note that this contrast does not appear between 1-approval and the plurality rule, since we have shown that the former rule also converges under partial preferences when only flipping is allowed. Hence, in the more general setting of this paper, we observe a new difference between (an extension of) plurality and (an extension) of antiplurality.

Once more, when addition is not permitted, then allowing for omission on top of flipping is key to convergence (Proposition 8).

Table 2
Summary of convergence results for the veto family of rules. “✓” indicates that convergence is guaranteed, while “✗” indicates that an example against convergence exists. When only omission or only addition is allowed, convergence is trivially achieved in a finite number of rounds.

	1-veto	k -veto	$(m - 1)$ -veto
omm	✓trivial	✓trivial	✓trivial
add	✓trivial	✓trivial	✓trivial
flip	✗ Prop.11	✗ Prop.13	✓ Prop.17
omm + add	✗ Prop.10	✗ Prop.13	✓ Prop.16
omm + flip	✓ Prop.12	✓ Prop.14	✓ Prop.17
add + flip	✗ Prop.11	✗ Prop.13	✓ Prop.18
omm + add + flip	✗ Prop.10	✗ Prop.13	✗ Prop.15

Proposition 8. *If omission and flipping are allowed, but addition is not, then the iterative $(m - 1)$ -approval process is guaranteed to converge.*

Proof. The proof is identical to that of Proposition 6. □

When flipping and omission are allowed, but addition is not, then convergence will be achieved after at most n^2m^2 rounds, as we argued is the case for k -approval.

Furthermore, when omission and addition are both allowed but flipping is not, convergence will be reached in at most n rounds: Every agent may only add a pairwise preference once, to reduce the score of an alternative by placing it at the bottom of a linear order (and when only omission or only addition is allowed, convergence will be reached after at most 0 and n rounds, respectively).

Proposition 9. *If omission and addition are allowed, but flipping is not, then the iterative $(m - 1)$ -approval process is guaranteed to converge.*

Combining our observations, we now know what the conditions characterising convergence for the $(m - 1)$ -approval rule are.

Corollary 3. *The iterative $(m - 1)$ -approval process is guaranteed to converge if and only if flipping is not allowed, unless omission but not addition is allowed as well.*

4. The veto family

This section examines convergence for the veto family. Table 2 summarises the relevant results and indicates where their proofs can be found.

4.1. 1-veto


For the 1-veto rule, we find a similar (in terms of the score-changes involved) counterexample as for $(m - 1)$ -approval. However, different types of manipulation moves play a role now (specifically, omission and addition instead of flipping).

Proposition 10. *If omission and addition are allowed, then the iterative 1-veto process is not guaranteed to converge (independently of whether flipping is allowed).*

Proof. Consider a sincere preference profile as follows:


- Agent 1 : $c \triangleright a \triangleright b \triangleright d$
- Agent 2 : $d \triangleright c \triangleright b \triangleright a$
- Agent 3 : $a \triangleright d \triangleright b \triangleright c$
- Agent 4 : $b \triangleright a \triangleright c \triangleright d$
- Agent 5 : $b \triangleright d \triangleright a \triangleright c$
- Agent 6 : $b \triangleright c \triangleright d \triangleright a$
- Agent 7 : $a \triangleright c \triangleright d \triangleright b$
- Agent 8 : $a \triangleright c \triangleright d \triangleright b$
- Agent 9 : $a \triangleright c \triangleright d \triangleright b$

Every time an agent makes a move, she alters the position of the alternative ranked second in her sincere preference. For example, consider agent 1 with the sincere preference $c \triangleright a \triangleright b \triangleright d$. The first time that agent 1 performs a manipulation act, we write “ $1a\downarrow$ ” when she moves alternative a to the bottom by omission and submits the new preference $\{(c, a), (c, b), (b, d), (c, d)\}$, to make c win instead of a . Then, we write “ $1a\uparrow$ ” when she removes a from the bottom by adding the pair (a, d) , to make a win instead of b . When she needs to decrease the score of a again, agent 1 will just omit the pair (a, d) from her preference. In the cycle, she will alternate between the two insincere preferences (see below for an illustration). Note that the notation $kx\downarrow$ is again used for a move of agent k that decreases the score of alternative x , although the specific preferences corresponding to the move may vary.



In all preferences above, alternative d gets 0 points and alternatives c and b get 1 point each. What changes is that alternative a alternates between getting 1 and 0 points.

The first six agents are going to participate in a cycle, starting from the score-vector on the top-left corner of the figure below. Note that the cycle only starts from the second iteration of the described sequence of deviations.



A counterexample for convergence thus exists. \square


When flipping but not omission is allowed, a similar counterexample to that in the proof of Proposition 10 can be constructed, in which only flipping is employed. We provide Proposition 11.

Proposition 11. *If flipping but not omission is allowed, then the iterative 1-veto process is not guaranteed to converge (independently of whether addition is allowed).*

Proof. Consider a sincere preference profile as follows:


- Agent 1 : $c \succ a \succ b \succ d \succ e$ Agent 6 : $b \succ c \succ d \succ a \succ e$
- Agent 2 : $d \succ c \succ b \succ a \succ e$ Agent 7 : $a \succ c \succ d \succ b \succ e$
- Agent 3 : $a \succ d \succ b \succ c \succ e$ Agent 8 : $a \succ c \succ d \succ b \succ e$
- Agent 4 : $b \succ a \succ c \succ d \succ e$ Agent 9 : $a \succ c \succ d \succ e \succ b$
- Agent 5 : $b \succ d \succ a \succ c \succ e$

Every time an agent makes a move, she alters the position of the alternative ranked second in her sincere preference. For example, consider agent 1 with the sincere preference $c \succ a \succ b \succ d \succ e$. The first time that agent 1 performs a manipulation act, we write “ $1a\downarrow$ ” when she moves alternative a to the bottom by flipping and submits the new preference $c \succ b \succ d \succ e \succ a$, to make c win instead of a . Then, we write “ $1a\uparrow$ ” when she removes a from the bottom again by flipping and submits the new preference $c \succ b \succ d \succ a \succ e$, to make a win instead of b . When she needs to decrease the score of a again, agent 1 will submit the preference $c \succ b \succ d \succ e \succ a$. In the cycle, she will alternate between the two insincere preferences, where a is either ranked last or second to last (see below for an illustration). The notation $kx\downarrow$ is again used for a move of agent k that decreases the score of alternative x , without corresponding to unique starting and ending preferences.



In all preferences above, alternatives $b, c,$ and d get 1 point each. What changes is that alternatives a and e alternate between getting 1 and 0 points.

The first six agents are going to participate in a cycle, starting from the score-vector at the top-left corner of the figure below. Note that the cycle only starts from the second iteration of the described sequence of deviations.



A counterexample for convergence thus exists. \square

Again, we will study whether convergence can be guaranteed under further restrictions on the agents' moves.

Proposition 12. *If flipping and omission (but not addition) are allowed, then the iterative 1-veto process is guaranteed to convergence.*

Proof. We will show by induction that if flipping and omission (but not addition) are allowed, then an arbitrary agent i with sincere preference \triangleright_i will never employ flipping. She will be changing her submitted preference only by omission, making a new alternative the winner each time. Since this can happen for at most finitely many rounds, convergence follows.

For our induction basis, note that the first manipulation move of every agent under 1-veto will be by omission, and of type 1: the only possible incentive for the agent will constitute in decreasing the score of an alternative in order to make a different alternative win, which will be achieved by omission from the minimal divergence assumption.

Suppose now that the statement holds for the k^{th} time agent i changes her preference. We will show that it also holds for the $(k + 1)^{\text{th}}$ time she changes her preference. Say this happens in round t .

Case 1: Agent i 's move in round t is of type 1 (and not of type 2). Then, omission will be employed for the same reason as for round 1.

Case 2: Agent i 's move in round t is of type 2, increasing the score of alternative y to make it the new winner, instead of a previous winner x . Since addition is not allowed, the only way that y can have its score increased is when y is a bottom alternative in \triangleright_i^t and agent i flips some pairwise preference $z \triangleright_i^t y$ to $y \triangleright_i^{t+1} z$. Note that $y \triangleright_i x$ and $z \triangleright_i y$ (because up until round t the agent has only omitted pairwise preferences), and $z \triangleright_i x$ by transitivity. Since in round t alternative z is not on the bottom for agent i , $z \triangleright_i^t x$ holds because the induction hypothesis tells us that the only pairwise comparisons of the form $x_1 \triangleright x_2$ that have been omitted are such that x_1 is forced to lose, and we are not allowing addition. Hence, $z \triangleright_i^{t+1} x$ also holds because the agents perform minimal changes on their submitted preferences. But $z \triangleright_i^{t+1} x$ and $y \triangleright_i^{t+1} z$ imply $y \triangleright_i^{t+1} x$, by transitivity. This is impossible, since y is a bottom alternative in round t and addition is not allowed. \square

Knowing that if only omission or only addition is allowed, then convergence will always be achieved, we obtain Corollary 4.

Corollary 4. *The iterative 1-veto process is guaranteed to converge if and only if omission and addition are not allowed simultaneously, and flipping is allowed only when omission is allowed as well.*

Convergence for 1-veto can be reached in at most nm rounds when flipping and omission (but not addition) are allowed or when only omission is allowed: Each agent will change her preference at most m times, to prevent a new alternative from being the winner. When only addition is allowed, all sincere profiles constitute stable states.

4.2. k -veto, for $1 < k < m - 1$

The analysis for k -veto draws inspiration from both our work regarding 1-veto and that regarding k -approval. First, all counterexamples for 1-veto can be reconstructed with some additional dummy alternatives, so that the relevant k comes into play.

Proposition 13. *If omission and addition are allowed simultaneously (independently of whether flipping is allowed), or if flipping is allowed without omission (independently of whether addition is allowed), then the iterative k -veto process is not guaranteed to converge.*

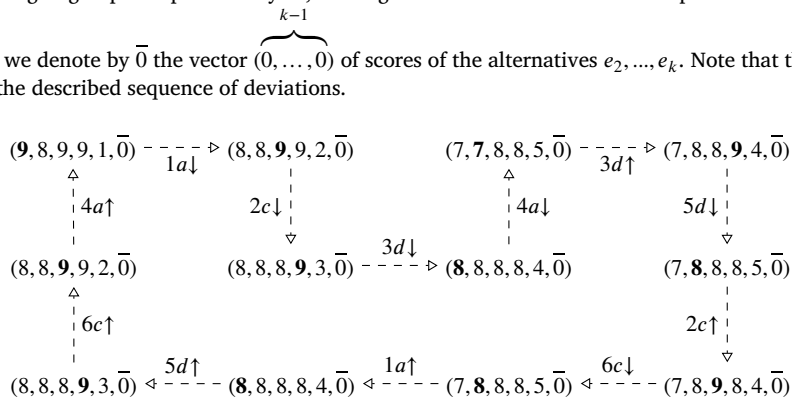
Proof. Let us show how the counterexample of Proposition 11 is adjusted for a k -veto rule, proving that when flipping is allowed without omission, convergence is not guaranteed (the counterexample of Proposition 10 can be adjusted in a directly analogous way). Consider a sincere preference profile as follows:

- Agent 1 : $c \triangleright a \triangleright b \triangleright d \triangleright e_1 \triangleright \dots \triangleright e_k$
- Agent 2 : $d \triangleright c \triangleright b \triangleright a \triangleright e_1 \triangleright \dots \triangleright e_k$
- Agent 3 : $a \triangleright d \triangleright b \triangleright c \triangleright e_1 \triangleright \dots \triangleright e_k$
- Agent 4 : $b \triangleright a \triangleright c \triangleright d \triangleright e_1 \triangleright \dots \triangleright e_k$
- Agent 5 : $b \triangleright d \triangleright a \triangleright c \triangleright e_1 \triangleright \dots \triangleright e_k$
- Agent 6 : $b \triangleright c \triangleright d \triangleright a \triangleright e_1 \triangleright \dots \triangleright e_k$
- Agent 7 : $a \triangleright c \triangleright d \triangleright b \triangleright e_1 \triangleright \dots \triangleright e_k$
- Agent 8 : $a \triangleright c \triangleright d \triangleright b \triangleright e_1 \triangleright \dots \triangleright e_k$
- Agent 9 : $a \triangleright c \triangleright d \triangleright e_1 \triangleright \dots \triangleright e_k \triangleright b$

Every time an agent makes a move, she alters the position of the alternative ranked second in her sincere preference. For example, consider agent 1 with the sincere preference $c \triangleright a \triangleright b \triangleright d \triangleright e_1 \triangleright \dots \triangleright e_k$. The first time that agent 1 performs a manipulation act, we write “ $1a\downarrow$ ” when she moves alternative a lower by flipping and submits the new preference $c \triangleright b \triangleright d \triangleright e_1 \triangleright a \triangleright \dots \triangleright e_k$, to make c win instead of a . Then, we write “ $1a\uparrow$ ” when she moves a higher again by flipping and submits the new preference $c \triangleright b \triangleright d \triangleright a \triangleright e_1 \triangleright \dots \triangleright e_k$, to make a win instead of b . When she needs to decrease the score of a again, agent 1 will return to the preference $c \triangleright b \triangleright d \triangleright e_1 \triangleright a \triangleright \dots \triangleright e_k$. In the cycle, she will alternate between the two insincere preferences.

The first six agents are going to participate in a cycle, starting from the score-vector at the top-left corner of the figure below. For

reasons of space, below we denote by $\bar{0}$ the vector $(0, \dots, 0)$ of scores of the alternatives e_2, \dots, e_k . Note that the cycle only starts from the second iteration of the described sequence of deviations.



A counterexample for convergence thus exists. □

Concerning the positive results, if only omission or only addition is allowed, then convergence easily follows in at most mn steps. Moreover, we again find that allowing omission together with flipping is advantageous to convergence. The relevant proof is close to the proof of Proposition 6 regarding k -approval, although several of the technical tricks employed in the two cases are different—importantly, both proofs make explicit use of the minimal divergence assumption.

Proposition 14. *If flipping and omission (but not addition) are allowed, then the iterative k -veto process is guaranteed to converge.*

Proof. Suppose that a cycle exists, aiming for a contradiction. The agents that participate in it will not employ omission, since such moves cannot be repeated when addition is not allowed. Then, all moves in the cycle will employ flipping. Moreover, we claim that all moves in the cycle will be of type 2. To prove the claim, first note flipping in some round t of the cycle will increase the score of some alternative a and decrease the score of some alternative b , in order to make a the winner. Then if $a \neq w_{t+1}$, the new winner could simply be obtained via decreasing the score of b by omission (because of minimally diverging moves). So $a = w_{t+1}$, which is the definition of type 2.

Let i be an agent that makes her first move in the cycle, say in round t , increasing the score of some x_1 to make it the winner, and decreasing the score of some y_1 . For the scores to change appropriately, it must be the case that $s_{\triangleright_i}(x_1) = s_t(x_1) = 0$ and $s_{\triangleright_i}(y_1) = s_t(y_1) = 1$.

Since $x_1 \triangleright_i y_1$ and agent i has an incentive for the manipulation move, we know that $y_1 \neq w_t$. And since (from the minimal divergence assumption) x_1 and y_1 will be the only alternatives of which the scores change in round t , it holds that:

$$s_{t+1}(w_{t+1}) >_\ell s_{t+1}(w_t) = s_t(w_t)$$

So the score of the winner strictly increases in round t . Moreover, the alternative y_1 cannot be the winner in a future round unless the score of the winner decreases.

Now suppose that i makes her second move in the cycle in round t' , increasing the score of some x_2 to make it the winner, and decreasing the score of some $y_2 \neq y_1$. Note that agent i could have made x_2 win with her first move in round t too, because if this weren't the case, the following inequalities would imply that x_2 could not win in round t' either:

$$s_t(x_2) + 1 <_\ell s_t(w_t) <_\ell s_k(w_k) \text{ for all } t < k < t'$$

Since i could make x_2 win with her first move but she did not do so, we know that $x_1 \triangleright_i x_2$. Also $s_{\triangleright_i^{t'}}(w_{t'}) = 0$, because $x_2 \triangleright_i w_{t'}$ (so that there exists an incentive to make the move) and $x_1 \triangleright_i x_2$ (otherwise x_2 should have been made the winner with the first move of agent i), so $s_{\triangleright_i}(x_1) = 0$ means that $s_{\triangleright_i}(x_2) = 0$ and $s_{\triangleright_i}(w_{t'}) = 0$.

Note that $s_{\triangleright_i^{t'}}(y_2) = 1$ must hold for the manipulation move in round t' that decreases the score of y_2 to be possible, so $y_2 \neq w_{t'}$. As earlier, it is implied that the score of the winner strictly increases in round t' :

$$s_{t'+1}(w_{t'+1}) >_\ell s_{t'+1}(w_{t'}) = s_{t'}(w_{t'})$$

Moreover, y_1 and y_2 cannot be the winners in a future round unless the score of the winner decreases. Continuing with the same reasoning, we see that the score of the winner never decreases, while every time an agent makes a new move in the cycle, she must decrease the score of a different alternative. Since the alternatives are finitely many but the cycle contains infinitely many rounds, we reached a contradiction. \square

We know from the proof of Proposition 14 that in every sequence of flipping steps, all moves will be of type 2, constantly making new alternatives win. Such a sequence can take at most mn steps. Furthermore, an agent i may make use of omission in some round t in order to decrease the score of an alternative y and make a different alternative win. Suppose that there exists some alternative x such that $x \triangleright_i^t y$ (if such an alternative does not exist, then it will apparently not be possible to change the score of y by flipping in the future). Note that for a decrease in the score of y to be realisable, there must also exist some suitable alternative z such that $y \triangleright_i^t z$ (and thus $x \triangleright_i^t z$ by transitivity). After the agent's omission move that will be minimally diverging, it will still be the case that $x \triangleright_i^{t+1} y$ and $x \triangleright_i^{t+1} z$, but not that $y \triangleright_i^{t+1} z$. Since we assumed that addition is not allowed, it will not be possible for the agent to flip the pairwise preference between x and y to change the score of y in the future. Hence, the score of every alternative can be changed via omission at most n times. In total, when flipping and omission are allowed, convergence of the iterative k -veto process will be reached in at most $2mn$ steps.

Corollary 5. *The iterative k -veto process is guaranteed to converge, for all $1 < k < m - 1$, if and only if omission and addition are not allowed simultaneously, and flipping is allowed only when omission is allowed as well.*

4.3. $(m - 1)$ -veto


For $(m - 1)$ -veto, we identify a counterexample that in terms of scores is similar to the one for 1-approval, but in which the agents employ all three types of addition, omission, and flipping (as opposed to 1-approval, where only omission and addition were used).

Proposition 15. *When addition, omission, and flipping are all allowed, then the iterative $(m - 1)$ -veto process is not guaranteed to converge.*

Proof. Consider a sincere preference profile as follows (where with the notation $(c \triangleright d \triangleright a)$, b we aim at abbreviating the preference $\{(c, d), (c, a), (d, a)\}$):


Agent 1 : $(c \triangleright d \triangleright a), b$	Agent 5 : $(d \triangleright c \triangleright b), a$
Agent 2 : $(b \triangleright d \triangleright c), a$	Agent 6 : $(d \triangleright b \triangleright c), a$
Agent 3 : $(a \triangleright d \triangleright b), c$	Agent 7 : $(c \triangleright b \triangleright d), a$
Agent 4 : $(b \triangleright c \triangleright d), a$	Agent 8 : $(a \triangleright b \triangleright c), d$

Every time an agent makes a move, she alters the position of the alternative ranked second in her sincere preference. For example, consider agent 1 with sincere preference $(c \triangleright d \triangleright a), b$. In the first manipulation act of agent 1, we write “ $1d\uparrow$ ” when she moves alternative d on the top by flipping and addition and submits the preference $\{(d, c), (c, a), (d, a), (d, b)\}$ to make d win instead of a . We write “ $1d\downarrow$ ” when she removes alternative d from the top by omission and submits the preference $\{(d, c), (c, a), (d, a)\}$ to make c win. When agent 1 needs to increase the score of d again, she will simply add back the pairwise preference (d, b) . In the cycle, agent 1 will alternate between the two insincere preferences, as depicted below.



When agent 1 places alternative d on top, then d gets score 1 and all other alternatives get score 0; in the other preferences of agent 1, all alternatives get score 0.

The first seven agents are going to participate in a cycle, starting from the score-vector at the top-left corner of the figure below. Note that the cycle only starts from the second iteration of the described sequence of deviations.



A counterexample for convergence thus exists. \square

Strikingly, the iterative $(m - 1)$ -veto process is not guaranteed to converge *only* if omission, addition, and flipping are all allowed.

Proposition 16. *If flipping is not allowed, then the iterative $(m - 1)$ -veto process is guaranteed to converge.*

Proof. The $(m - 1)$ -veto rule assigns score 0 to all alternatives except for a single case: score 1 is assigned to the most preferred alternative in a linear order. Since flipping is not allowed, an agent may have an incentive to change her preference once, by addition, increasing the score of some alternative she finds inferior to no other alternative. Subsequently, no score can be altered by addition; also, omission can only decrease the score of the most preferred alternative, which is not beneficial to the agent. Convergence follows. \square

Proposition 17. *If addition is not allowed, then the iterative $(m - 1)$ -veto process is guaranteed to converge.*

Proof. A convergence proof can be obtained analogously to the proof of Proposition 3. Again, we will start with a proof of the following claim: Every time an agent changes her preference, she promotes a winning alternative that is different from those she has promoted before. Because the set of alternatives A is finite, convergence follows.

Suppose that agent i is the first to change her reported (sincere) preference \triangleright_i in round 1. Since addition is not allowed, agent i must sincerely hold a linear preference order and make a move of type 2 by flipping (note that for $(m - 1)$ -veto, omission may only decrease the score of the top alternative in a linear preference, which does not constitute a manipulation incentive for an agent). So agent i 's move will aim at increasing the score of some alternative y : she will report a new pairwise preference with $y \triangleright_i^2 x$, where x was the unique top alternative in \triangleright_i and was not winning. We show that:

(*) Every alternative y' that agent i sincerely places above y (i.e., $y' \triangleright_i y$) has no chance of winning in later rounds unless the score of the winner decreases.

Condition (*) is clearly true for alternative x , which loses one point after i 's move; it is also true for every non-top alternative $z \neq x$ such that $z \triangleright_i y$ (if such z exists), since z could not become a winner in round 2 (we know that because agent i submitted a best response), and the score of z remained the same after i 's move.

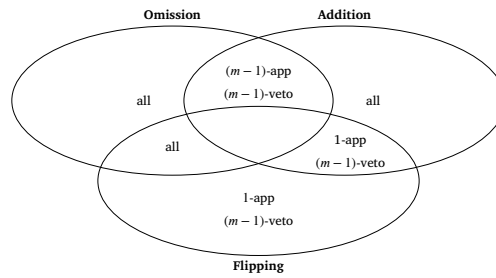


Fig. 3. Convergence of voting rules in the approval and the veto family. For addition (omission) only, all rules converge (and convergence clearly holds even if iteration starts for an arbitrary—not necessarily sincere—preference profile). For omission and addition, $(m - 1)$ -approval / veto converges because flipping is needed for an agent to want to make a move twice. For flipping and omission, convergence is guaranteed for all rules because omission can do the same job as flipping, but is preferable due to being cheaper, while it cannot be repeated. For flipping and addition, “plurality-like” rules converge because the score of the winning alternative will constantly increase, reducing the set of future winners. Notably, $(m - 1)$ -veto converges under any condition where other rules converge too.

Now, as is the case for agent i , every agent j 's first move will be of type 2 and condition $(*)$ will also hold for j . So the winning score will keep increasing. From $(*)$ it follows that when an agent modifies her reported preference for the second time, she will have no incentive to use a type-1 move to decrease the score of the winner, because no better alternative will be able to win. This means that each agent who performs a second move will promote another alternative than the one she promoted in her first move, increasing its score and making it the winner. Continuing with this reasoning, the same will be true for everyone who performs a third move, and so on. \square

Proposition 18. *If omission is not allowed, then the iterative $(m - 1)$ -veto process is guaranteed to converge.*

Proof. Convergence is implied by Proposition 3. \square

Corollary 6. *The iterative $(m - 1)$ -veto process is guaranteed to converge if and only if omission, addition, and flipping are not simultaneously allowed.*

If only omission is allowed, convergence is immediate (holds after 0 rounds) for $(m - 1)$ -veto, and if addition but no flipping is allowed, convergence requires at most n rounds. Finally, if flipping is allowed (possibly together with addition or omission), the proof of Proposition 3 shows that we will have convergence after at most mn rounds.

Fig. 3 provides a graphical summary of our findings regarding all rules.

5. Simulations

An important limitation of our theoretical results of the previous sections is that they only speak about guarantees for convergence; said differently, they speak about convergence that is achieved starting from *all* possible profiles. Analogous guarantees are reflected by results on the strategyproofness of aggregation rules, requiring that agents *never* have an incentive to misrepresent their preferences. When guarantees cannot be given, we do not yet provide any information about how *often* convergence is achieved. Furthermore, when a stable state is reached, we only know the *maximum* number of rounds needed to arrive to it—in practice, convergence can be much faster.

In this section, we complement our previous analytical results with observations that we have attained through computer simulations in Python. As we will see, convergence is essentially always achieved by the voting rules of our interest, even if we allow all possible types of manipulation (still, bearing in mind the minimal divergence assumption). This means that all stable states constitute Nash equilibria, where no agent can unilaterally benefit from changing her reported preference. For both the approval family and the veto family (particularly for $k \in \{1, m - 1\}$), we explore:

- the frequency of non-manipulable sincere profiles;
- the frequency of manipulable profiles where at the end of the iterative process, when convergence is achieved, the winner is the same as the one before manipulation;
- the number of rounds needed for a stable profile to be reached;
- the quality of the winner, in social terms, in sincere vs. in equilibrium profiles.

The first item above is independent of iterative voting. A number of previous works have asked the same question for voting with complete preferences, but the corresponding problem under incomplete preferences is significantly underexplored. Note that whenever we refer to incomplete preferences in this section, we still mean partial (i.e., transitive) ones.

For profiles of linear orders, Kelly [13] conducted an experimental analysis for resolute voting rules satisfying various axiomatic properties (such as anonymity, neutrality, and a version of monotonicity) and found that the Borda rule's performance is amongst the very best ones (although the rule has been often criticised in terms of its susceptibility to manipulation). Kelly's experiments

relied on the *impartial culture* (IC) assumption (implying a uniformly random distribution of the preference profiles), with at most five alternatives and three agents.

Nitzan [24] also studied the Borda rule, as well as the plurality rule, for profiles of linear orders with $n \leq 90$. For both these rules, Nitzan showed that the probability of them being vulnerable to individual strategic manipulation is always greater than 0.165, 0.284, and 0.379 for three, four, and five alternatives, respectively.

In an iterative setting, Grandi et al. [11] considered positional scoring rules for profiles of linear orders too, and conducted simulations showing that iterative manipulation under certain restrictive types of moves (that are easy to compute from the agents' perspective) in general yields a positive increase in the Condorcet efficiency of the rules (i.e., in the proportion of profiles where an alternative that wins a majority competition over all other alternatives is elected). A positive increase was also observed with respect to the group's social welfare, assuming that an agent i 's cardinal utility given an alternative a is the Borda score that a gets in i 's sincere preference. Grandi et al. generated profiles using the IC assumption and versions of the urn model (one of which we will explain below). Similar observations were obtained by Reijngoud and Endriss [29].

Meir et al. [20] proposed the *local-dominance theory* of voting equilibria, suggesting a type of restriction for the agents' manipulation moves that constitutes a subset of their best responses. Using $n \in \{10, 20, 50\}$ and $m \in \{3, \dots, 8\}$, they showed that the iterative plurality process quickly converges to an equilibrium starting from sincere profiles, and that the quality of the winner (again based on Borda utilities) generally improves compared to sincere voting for the IC assumption, but not for the urn model.

The IC assumption, which is employed by ample papers on voting simulations, is investigated in depth by Nurmi [25]. Broadly speaking, this assumption is useful when one thinks of the space of all possible profiles and is interested in finding the probability that a certain profile satisfies a given property—by sampling sufficiently many profiles, the desired probability is approximated. Moreover, the IC assumption is known to provide some “worst-case” guarantees for the observations obtained through simulations; for instance, it maximises the probability of cyclic collective preferences to occur under the pairwise majority rule [36]. Yet, people's votes in the real world cannot be expected to follow a uniform distribution.

In order to capture a certain kind of correlation between the preferences in a group, the literature has considered variations of the Polya-Eggenberger *urn model*—the one employed by Meir et al. [20] works as follows: For a given number of alternatives m and a parameter ℓ , we start with an urn containing all possible preferences that an agent may hold (for linear orders, there are $m!$ such preferences in total—for partial orders, there are many more). We sample ℓ preferences from the urn uniformly at random, which will be the “leading” preferences. We make enough copies of each one of those preferences and place all of them back in the urn, so that in the end each leading preference constitutes $\frac{1}{\ell+1}$ the urn. Each agent then chooses a random preference from the urn.⁷

For every combination of $n \in \{10, 20, 50\}$ and $m \in \{3, 4, 5\}$, we generate 200 profiles under the IC assumption and the aforementioned urn model with $\ell = 2$ (we call this the 2urn model).⁸ We also explicitly generate profiles of linear orders, since in our theoretical results we commonly found that cycles were formed when iteration started from sincere profiles of complete preferences—exploring such profiles through extensive simulations will give us a better idea about the role that completeness plays for convergence.

5.1. Non-manipulable sincere profiles

A profile is *manipulable* under an aggregation rule if some agent has an incentive to change her preference and steer the outcome towards a winner that she prefers. Fig. 4 depicts our experimental results concerning the (quite high) frequency of sincere profiles that are *not* manipulable, for $m = 4$ (we did not observe any significant differences for $m = 5$).

Note that the difference with respect to the frequency of manipulation between 1-approval and $(m - 1)$ -veto is insignificant under complete profiles, as the two rules will initially coincide and differ only in an apparently negligible manner during manipulation. We confirm two natural hypotheses: that manipulable profiles are rarer under the 2urn distribution than under the IC one,⁹ as well as that larger numbers of alternatives are detrimental to manipulability.

Starting with incomplete profiles, we see an ordering with respect to the frequency of manipulability under the different rules for both the 2urn and the IC models—from less frequently manipulable to more frequently manipulable, the rules are: 1-approval, 1-veto, $(m - 1)$ -veto, and $(m - 1)$ -approval. Moreover, the difference between 1-approval and $(m - 1)$ -approval (and similarly between 1-veto and $(m - 1)$ -veto) is larger under incomplete initial profiles than under complete ones—possibly because there is a higher chance for ties when $k = m - 1$ under incomplete profiles.

⁷ The same process is followed both for linear and for partial preferences. Importantly, in this paper we have not paid particular attention to the different kinds of partial preferences that exist and their influence to convergence—applying the urn model, we have only focused on the degree of concordance between the preferences of the agents in a group.

⁸ However, five alternatives seem to be one too many for our program to terminate under the 1-veto rule, so we skip that case (this does not prohibit us from comparing our experimental observations with our analytical results, since cycles were already observed under the 1-veto rule for four alternatives, in the proof of Proposition 10). Interestingly, this problem does not arise under the 1-approval rule, nor under the $(m - 1)$ -approval rule, which brings to the surface further computational differences between the approval and the veto families. The difficulty linked to the iterative 1-veto process stems from the fact that the manipulation moves of the agents under that rule will often be very costly, which will make it harder, and slower, for our program to compute them. This can be seen if we consider an agent i that reports a preference $e \succ_i d \succ_i c \succ_i b \succ_i a$ in round t , and suppose that all alternatives have the same total score in that round (so a wins, because of the lexicographic tie-breaking rule). Then, agent i must reduce the score of three alternatives, b, c , and d , in order to make her favourite alternative win. Under the plurality rule on the other hand, such situations do not arise.

⁹ Surprisingly, 2urn-drawn profiles are more manipulable than IC-drawn ones for $(m - 1)$ -approval. This contrast will further stand out later, when we talk about convergence speed—we will address it then.




Fig. 4. Frequency of non-manipulable sincere profiles, incomplete ('incom') and complete ('com') ones, for $n \in \{10, 20, 50\}$.




Fig. 5. Frequency of winner-invariant profiles, incomplete ('incom') and complete ('com') ones, amongst all those that are manipulable, for $n \in \{10, 20, 50\}$.

Complete initial profiles increase manipulability for $(m - 1)$ -veto and $(m - 1)$ -approval. On the one hand this can be expected, since agents that rank more alternatives will have more incentives for manipulation by finding a different alternative preferable to the current winner. On the other hand, 1-veto produces fewer ties amongst the winners in the complete case, which seems to have a strong effect in preventing manipulation.

We also observe that the number of non-manipulable profiles increases with the size of the group, across rules, distributions, and numbers of alternatives: Since the rules we investigate assign the same score to many alternatives in a given preference (often to all of them), they can be expected to produce many ties in small groups, implying that several agents may be able to make their favourite alternatives win.

5.2. Winner invariance

Fig. 5 illustrates our results on winner invariance for $m = 4$, conditioning on manipulable profiles. We checked for other numbers of alternatives too (e.g., $m = 3$ and $m = 5$), and our observations were not affected.




Fig. 6. Speed of convergence (for incomplete, manipulable profiles, 2urn-drawn).

When manipulation cannot be avoided in a given profile, it would be comforting to know that allowing the agents to manipulate in rounds has no final impact on the winner. We explore how often such a scenario arises—where the winner in equilibrium coincides with the sincere winner—bringing out positive news in various cases.

Under all rules and distributions, winner invariance is in general achieved more frequently as the group grows larger and for complete sincere profiles: Since we saw that large groups encourage non-manipulability while complete profiles damage it, we cannot deduce any direct connection between winner invariance and manipulability.

For complete initial profiles, around 50% of the iteration processes end up not changing the winner, and plurality-like rules (i.e., 1-approval and $(m - 1)$ -veto) perform better than antiplurality-like ones (i.e., 1-veto and $(m - 1)$ -approval); this may be due to ties, which are expected to be more evident for antiplurality-like rules, and which may make agents pivotal triggering more manipulation moves. For incomplete initial profiles and smaller groups of agents, we detect that profiles drawn from the IC model are more frequently winner-invariant than those drawn from the 2urn model. This observation is rather curious—further research could shed more light on the reasons behind it.

5.3. Speed of convergence

Probably the most important conclusion of our simulations is that convergence is *always* achieved in practice, even if sometimes it takes a while. In the plots below (all concerning manipulable profiles with four alternatives and ten agents), we only depict up to fifty manipulation rounds. The number of alternatives does not significantly affect our observations, while larger groups—reasonably—delay convergence.¹⁰

Figs. 6 and 7 highlight the radical, negative effect of completeness on convergence speed; a similar—albeit weaker—effect of the IC assumption follows from Figs. 8 and 9 (with an exception regarding $(m - 1)$ -approval, which we discuss).

Plurality-like rules (i.e., 1-approval and $(m - 1)$ -veto) generally reach convergence faster than the rest—frequently, they converge after only a few rounds. Indeed, antiplurality-like rules are more sensitive to manipulation by definition: they enable the agents to manipulate by decreasing the score of the winning alternative, which means that some alternatives that have no chance of winning in the initial profile may become potential winners in future rounds (this phenomenon never emerges under the plurality-like rules, where the agents always increase the scores of alternatives by manipulating).

In the cases where the initial, sincere profiles of the agents are complete, 1-veto clearly is the slowest rule. At the same time, $(m - 1)$ -approval often performs quite well under the same circumstances—in particular, for complete and IC-drawn profiles, $(m - 1)$ -approval is amongst the fastest rules. Here is a possible interpretation for this: Agents can manipulate $(m - 1)$ -approval either by increasing the score of an attractive alternative, or by decreasing the score of an appalling one—they will start with the latter type of move in the initial, sincere profile. But only one alternative can have score 0 under the $(m - 1)$ -approval rule for complete profiles. So, a decrease in score of an alternative a will imply the increase in score of another alternative b ranked below a by the manipulator, who will only decrease the score of a if she is not in danger of making b win. In IC profiles, ties between the alternatives are more common than in 2urn profiles (because in the IC model, loosely speaking, alternatives are uniformly distributed across positions in the agents' preferences). Therefore, the probability that b will win is larger under the IC, disincentivising agents from manipulation.

¹⁰ So in large groups, it is less common to find profiles that are manipulable, but those that are manipulable take longer to converge.

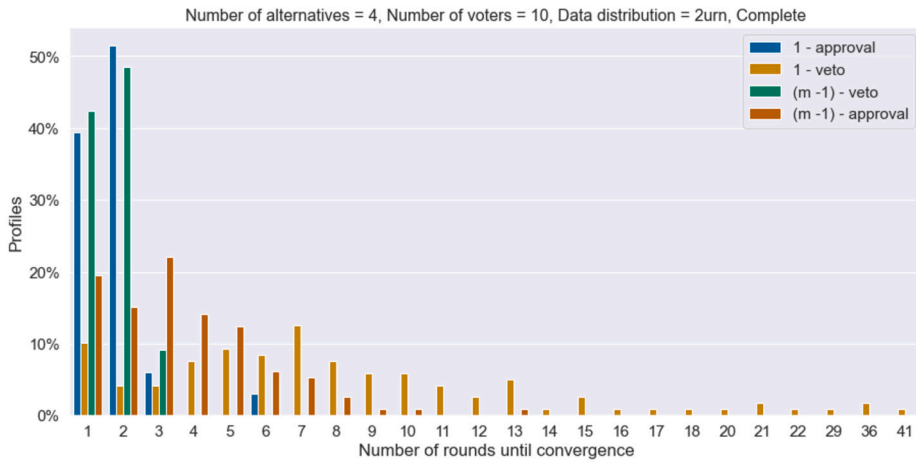


Fig. 7. Speed of convergence (for complete, manipulable profiles, 2urn-drawn).




Fig. 8. Speed of convergence (for incomplete, manipulable profiles, IC-drawn).




Fig. 9. Speed of convergence (for complete, manipulable profiles, IC-drawn).




Fig. 10. Improvement in social welfare from the sincere to the equilibrium winner (for manipulable, complete, IC-drawn profiles, where the winner changes at the end of iteration).

5.4. Quality of outcomes under iterative voting

Is iteration profitable for the agents, in terms of social welfare, when it changes the winner? To answer this question we first have to define social welfare, which is not straightforward for incomplete preferences. For complete preferences, social welfare is routinely calculated using Borda scores. The idea is that the Borda score that an alternative gets from an agent's sincere preference encodes how much the agent is *relieved* by seeing that alternative win, or equivalently how much she does not *regret* seeing it win.¹¹

However, there is no unique way to extend Borda scores for incomplete preferences [8,33]. We could examine versions of *pessimistic* and *optimistic* social welfare, corresponding to the relevant Borda generalisations, and implying that the agents are driven only by regret or only by relief, respectively. To keep things simple, here we will only exhibit results with respect to the *averaged* version of the Borda scores, meaning that our agents' utility is calculated by subtracting the regret they feel from their relief. Formally, the utility that an agent i receives from a winning alternative a corresponds to the number of alternatives that i truthfully ranks below a minus the number of alternatives that i ranks above a . By summing up the utilities of all agents, we obtain the social welfare of the group. The *quality of iteration* then is the improvement in social welfare from the winner in the sincere profile to the winner in the equilibrium profile.

There are several conclusions that we draw from our experiments regarding quality of iteration. Firstly, a substantial percentage of all manipulable profiles (amongst those that change the winning alternative only) do not have any impact on the agents' social welfare. This means that the group as a whole is often not affected by the consequences of iterative manipulation. When an effect on the quality of the winner arises, it takes both positive and negative values; these values are somewhat symmetric, but tend to weigh more on the positive side. Fig. 10 (for IC-drawn, complete profiles with four alternatives and ten voters) is representative of our results.

We have also observed that incomplete sincere profiles imply a wider range of values for the quality of iteration; the same holds for larger groups of agents and larger sets of alternatives. We do not include any relevant plot to demonstrate these effects here because the extreme length of their corresponding ranges makes the plots difficult to read. We have not detected any particular pattern with respect to the different rules.

6. Conclusion

In this paper we have developed an original model of iterative voting for agents with partial preferences. Since voting rules that allow for partial preferences are scarce, we have generalised the well-known k -approval rules, traditionally studied in the context of iterative voting with linear preferences, by defining two natural families of approval-based and veto-based rules (these families include rules that extend the popular plurality and antiplurality rules, from complete to partial preferences). We have illuminated the convergence properties of all rules in the two families for different types of moves that the agents may perform under partial preferences, thus drawing a fine-grained picture (illustrated in Fig. 3). Knowing exactly under which restrictions on the moves of the agents convergence is guaranteed, and being able to impose those restrictions externally, can contribute to the design of safe-to-use voting platforms. We have also provided theoretical upper bounds on the number of rounds needed for a stable state to be reached, whenever convergence of a rule is guaranteed. To test the more practical implications of our results, we have lastly conducted simulation experiments. Summarising our findings, we can now confidently say that iterative manipulation processes with partial

¹¹ Alternative ways to evaluate an outcome for a group of agents without inducing cardinal scores from their ordinal preferences exist—for example, one can calculate the *Copeland score* of the winner (i.e., the number of alternatives to which the winner is deemed superior for a majority of the agents).

preferences are not so alarming as they may have seemed initially. First, very often they will not even start. Second, even if they start, they will eventually stop. Third, when they stop, they will probably do not have any impact on the winning alternative. Fourth, even if they do have an impact on the winning alternative, this will likely be a positive one for the group. All these observations agree with the ones made in the special case where the agents are only allowed to report complete preferences—recall, e.g., the work of Meir et al. [20].

Evidently, our work raises plenty of follow-up questions: Is it computationally easy for the agents to determine their minimally diverging moves? Are there further sensible restrictions that we can impose to guarantee convergence? To what extent does the degree of incompleteness of the agents' preferences affect the number of rounds required for convergence? What other rules for partial preferences would be of interest in an iterative setting, and what if the agents can also make simultaneous updates?¹² Finally, questions we explored in this paper in the context of voting can also be asked within other aggregation frameworks, such as judgment aggregation (Terzopoulou and Endriss [31] have made a start for the complete case). The concepts of omission, addition, and flipping make sense in judgment aggregation too, and iteration might give rise to intriguing observations for various aggregation processes.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

We have shared the link to our data and code

Acknowledgements

We would like to thank the anonymous *Artificial Intelligence* reviewers for their helpful feedback on our work.

References

- [1] M. Ayadi, N. Amor, J. Lang, Approximating voting rules from truncated ballots, *Auton. Agents Multi-Agent Syst.* 36 (2022) 1–29.
- [2] M. Ayadi, N. Amor, J. Lang, D. Peters, Single transferable vote: incomplete knowledge and communication issues, in: *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 2019, pp. 1288–1296.
- [3] D. Baumeister, P. Faliszewski, J. Lang, J. Rothe, Campaigns for lazy voters: truncated ballots, in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2012, pp. 577–584.
- [4] M. Bentert, P. Skowron, Comparing election methods where each voter ranks only few candidates, in: *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 2020, pp. 2218–2225.
- [5] S. Brânzei, I. Caragiannis, J. Morgenstern, A.D. Procaccia, How bad is selfish voting?, in: *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI)*, 2013, pp. 138–144.
- [6] L. Dery, S. Obraztsova, Z. Rabinovich, M. Kalech, Lie on the fly: strategic voting in an iterative preference elicitation process, *Group Decis. Negot.* 28 (2019) 1077–1107.
- [7] Y. Desmedt, E. Elkind, Equilibria of plurality voting with abstentions, in: *Proceedings of the 11th ACM Conference on Electronic Commerce (EC)*, 2010, pp. 347–356.
- [8] P. Emerson, The original Borda count and partial voting, *Soc. Choice Welf.* 40 (2013) 353–358.
- [9] U. Endriss, S. Obraztsova, M. Polukarov, J.S. Rosenschein, Strategic voting with incomplete information, in: *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, 2016, pp. 236–242.
- [10] Y. Filmus, J. Oren, Efficient voting via the top-k elicitation scheme: a probabilistic approach, in: *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*, 2014, pp. 295–312.
- [11] U. Grandi, A. Loreggia, F. Rossi, K.B. Venable, T. Walsh, Restricted manipulation in iterative voting: Condorcet efficiency and Borda score, in: *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT)*, 2013, pp. 181–192.
- [12] J. Kavner, L. Xia, Strategic behavior is bliss: iterative voting improves social welfare, in: *Proceedings of the 34th Conference on Advances in Neural Information Processing Systems (NeurIPS)*, 2021, pp. 19021–19032.
- [13] J.S. Kelly, Almost all social choice rules are highly manipulable, but a few aren't, *Soc. Choice Welf.* 10 (1993) 161–175.
- [14] A. Koolyk, T. Strangway, O. Lev, J.S. Rosenschein, Convergence and quality of iterative voting under non-scoring rules, in: *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 2017, pp. 273–279.
- [15] J. Kruger, Z. Terzopoulou, Strategic manipulation with incomplete preferences: possibilities and impossibilities for positional scoring rules, in: *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2020, p. 645,653.
- [16] J.F. Laslier, The leader rule: a model of strategic approval voting in a large electorate, *J. Theor. Polit.* 21 (2009) 113–136.
- [17] O. Lev, J.S. Rosenschein, Convergence of iterative voting, in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2012, pp. 611–618.
- [18] O. Lev, J.S. Rosenschein, Convergence of iterative scoring rules, *J. Artif. Intell. Res.* 57 (2016) 573–591.
- [19] N. Mattei, T. Walsh, Preflib: a library of preference data, in: *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT)*, 2013, pp. 259–270.
- [20] R. Meir, O. Lev, J.S. Rosenschein, A local-dominance theory of voting equilibria, in: *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*, 2014, pp. 313–330.
- [21] R. Meir, M. Polukarov, J.S. Rosenschein, N.R. Jennings, Convergence to equilibria in plurality voting, in: *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*, 2010, pp. 823–828.

¹² Laslier [16] studied simultaneous moves in approval voting.

- [22] R. Meir, M. Polukarov, J.S. Rosenschein, N.R. Jennings, Iterative voting and acyclic games, *Artif. Intell.* 252 (2017) 100–122.
- [23] N. Narodytka, T. Walsh, The computational impact of partial votes on strategic voting, in: *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI)*, 2014, pp. 657–662.
- [24] S. Nitzan, The vulnerability of point-voting schemes to preference variation and strategic manipulation, *Public Choice* 47 (1985) 349–370.
- [25] H. Nurmi, An assessment of voting system simulations, *Public Choice* 73 (1992) 459–487.
- [26] S. Obratsova, O. Lev, M. Polukarov, Z. Rabinovich, J.S. Rosenschein, Farsighted voting dynamics, in: *Proceedings of the 1st IJCAI Workshop on Algorithmic Decision Theory*, 2015.
- [27] S. Obratsova, E. Markakis, D.R. Thompson, Plurality voting with truth-biased agents, in: *Proceedings of the 6th International Symposium on Algorithmic Game Theory (SAGT)*, 2013, pp. 26–37.
- [28] Z. Rabinovich, S. Obratsova, O. Lev, E. Markakis, J.S. Rosenschein, Analysis of equilibria in iterative voting schemes, in: *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, 2015, pp. 1007–1013.
- [29] A. Reijngoud, U. Endriss, Voter response to iterated poll information, in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2012, pp. 635–644.
- [30] R. Reyhani, M.C. Wilson, Best reply dynamics for scoring rules, in: *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI)*, 2012, pp. 672–677.
- [31] Z. Terzopoulou, U. Endriss, Modelling iterative judgment aggregation, in: *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, 2018, pp. 1234–1241.
- [32] Z. Terzopoulou, U. Endriss, Aggregating incomplete pairwise preferences by weight, in: *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 2019, pp. 595–601.
- [33] Z. Terzopoulou, U. Endriss, The Borda class: an axiomatic study of the Borda rule on top-truncated preferences, *J. Math. Econ.* 92 (2021) 31–40.
- [34] Z. Terzopoulou, P. Terzopoulos, U. Endriss, Code and data for “Iterative Voting with Partial Preferences”, Available at, https://github.com/PanagiotisTerzopoulos/iterative_voting_with_partial_preferences, 2021.
- [35] A. Tsang, K. Larson, The echo chamber: strategic voting and homophily in social networks, in: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2016, pp. 368–375.
- [36] I. Tsetlin, M. Regenwetter, B. Grofman, The impartial culture maximizes the probability of majority cycles, *Soc. Choice Welf.* 21 (2003) 387–398.
- [37] A. Wilczynski, Poll-confident voters in iterative voting, in: *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 2019, pp. 2205–2212.
- [38] W.S. Zwicker, Introduction to the theory of voting, in: F. Brandt, V. Conitzer, U. Endriss, J. Lang, A.D. Procaccia (Eds.), *Handbook of Computational Social Choice*, Cambridge University Press, 2016, pp. 23–56, chapter 2.