

Restricted Manipulation in Iterative Voting: Condorcet Efficiency and Borda Score

Umberto Grandi,¹ Andrea Loreggia,¹ Francesca Rossi,¹ K. Brent Venable² and Toby Walsh³

¹ University of Padova, Italy,

umberto.uni@gmail.com, andrea.loreggia@gmail.com, frossi@math.unipd.it

²Tulane University and IHMC, U.S., kvenabl@tulane.edu

³NICTA and UNSW, Australia, toby.walsh@nicta.com.au

Abstract

In collective decision making, where a voting rule is used to take a collective decision among a group of agents, manipulation by one or more agents is usually considered negative behavior to be avoided, or at least to be made computationally difficult for the agents to perform. However, there are scenarios in which a restricted form of manipulation can instead be beneficial. In this paper we consider the iterative version of several voting rules, where at each step one agent is allowed to manipulate by modifying his ballot according to a set of restricted manipulation moves which are computationally easy and require little information to be performed. We prove convergence of iterative voting rules when restricted manipulation is allowed, and we present experiments showing that most iterative voting rules have a higher Condorcet efficiency and a higher average position of the winner than their non-iterative version.

1 Introduction

In multi-agent systems, often agents need to take a collective decision. A voting rule can be used to decide which decision to take, mapping the agents' preferences over the possible candidate decisions into a winning decision for the collection of agents. In these kind of scenarios, it may be desirable that agents do not have any incentive to manipulate, that is, to misreport their preferences in order to influence the result of the voting rule in their favor. Indeed, manipulation is usually seen as bad behavior from an agent, to be avoided or at least to be made computationally difficult to accomplish. While we know that every voting rule is manipulable when no domain restriction is imposed on the agents' preferences (Gibbard, 1973; Satterthwaite, 1975), we can try to make sure that a voting rule is computationally difficult to manipulate for single agents or coalitions of agents. In these scenarios where agents' preferences are aggregated only once by a voting rule, the presence of manipulators is not fair to sincere agents, since they can be penalized.

In this paper we consider a different setting, in which instead manipulation is allowed in a fair way. More precisely,

agents express their preferences over the set of possible decisions and the voting rule selects the current winner, as in the usual case. However, this is just a temporary winner, since at this point a single agent may decide to manipulate, that is, to change her preference if by doing so the result changes in her favor. The process repeats with a new agent manipulating until we eventually reach a convergence state, i.e., a profile where no single agent can get a better result by manipulating. We call such a process *iterative voting*. In this scenario, manipulation can be seen as a way to achieve consensus, to give every agent a chance to vote strategically (a sort of fairness), and to account for inter-agent influence over time.

A practical example of this process is Doodle,¹ a very popular on-line system to select a time slot for a meeting by considering the preferences of the participants. In Doodle, each participant can approve as many time slots as she wants, and the winning time slot is the one with the largest number of approvals. At any point, each participant can modify her vote in order to get a better result, and this can go on for several steps. Depending on the voting rule, on the tie-breaking rule (to be used when there are several tied winners), and on the possible manipulation moves (that is, how agents are allowed to change their preferences in a single step), we may get convergence or not. We will say that the iterative version of a voting rule converges if it gets to a stable state no matter the initial profile.

Iterative voting has been the subject of numerous publications in recent years. Previous work has focused on iterating the plurality rule (Meir *et al.*, 2010), on the problem of convergence for several voting rules (Lev and Rosenschein, 2012), on the convergence of plurality decisions between multiple agents (Airiau and Endriss, 2009) and on the problem of analyzing epistemic conditions of strategic voting (Chopra *et al.*, 2004). Lev and Rosenschein (2012) showed that, if we allow agents to manipulate in any way they want (i.e., to provide their best response to the current profile), then the iterative version of most voting rules do not converge. Therefore, an interesting problem is to seek restrictions on the manipulation moves to guarantee convergence of the associated iterative rule. Restricted manipulation moves are good not only for convergence, but also because they can be easier to accomplish for the manipulating agent. In fact, contrar-

¹<http://doodle.com/>

ily to what we aim for in classical voting scenarios, here we do not want manipulation to be computationally difficult to achieve. It is actually desirable that the manipulation move be easy to compute while not requiring too much information for its computation.

An example of a restricted manipulation move is called *k*-pragmatists in Reijngoud and Endriss (2012): a *k*-pragmatist just needs to know the top *k* candidates in the collective candidate order, and will move the most preferred of those candidates to the top position of her preference. To compute this move, a *k*-pragmatist needs very little information (just the top *k* current candidates), and with this information it is computationally easy to perform the move. This move assures convergence of all positional scoring rules, Copeland, and Maximin, with linear tie-breaking. Note that each agent can apply this manipulation rule only once (since the top *k* candidates are always the same), and this is the main reason for convergence.

In this paper we introduce two restricted manipulation moves within the scenario of iterative voting and we analyze some of their theoretical and practical properties. Both manipulation moves we consider are polynomial to compute and require little information to be used. We show that convergence is guaranteed under both moves for those rules we consider, except for STV for which we only have experimental evidence of convergence. Moreover, we show that if a voting rule satisfies some axiomatic properties, such as Condorcet consistency or unanimity, then its iterative version will also satisfy the same properties as well. For voting rules that are not Condorcet consistent, we tested experimentally whether their Condorcet efficiency (that is, the probability to elect the Condorcet winner) improves by adopting the iterative version. Our experiments show that the Condorcet efficiency of voting rules improves when our restricted manipulation moves are used. Another parameter we used to test our manipulation moves is the average position of the winner in the initial profile. As in the previous case, we showed that this parameter increases when restricted manipulation is allowed.

The paper is organized as follows. In the first section we introduce the basic definitions of iterative voting and we define two new restricted manipulation moves. The second section contains theoretical results on convergence and preservation of axiomatic properties, and in the third section we present our experimental evaluation of restricted iterative voting. The last section contains our conclusions and points at some directions for future research.

2 Background Notions

In this section we recall the basic notions of social choice that we shall use in this paper, we present the setting of iterative voting, and we define a number of restrictions on the manipulation moves that agents can perform.

2.1 Voting Rules

Let \mathcal{X} be a finite set of m candidates and \mathcal{I} be a finite set of n individuals. We assume individuals have preferences over candidates in \mathcal{X} in the form of *strict linear orders*, i.e., transitive, anti-symmetric and complete binary relations. Individuals express their preferences in form of a ballot b_i (e.g., the

top candidate, a set of approved candidates, or the full linear order) and we call the choice of a ballot for each individual a profile $\mathbf{b} = (b_1, \dots, b_n)$. In this paper, we assume that individuals submit as a ballot for the election their full linear order, and we thus use the two notions of ballot and preference interchangeably.

A (non-resolute) *voting rule* F associates with every profile $\mathbf{b} = (b_1, \dots, b_n)$ a non-empty subset of winning candidates $F(\mathbf{b}) \in 2^{\mathcal{X}} \setminus \emptyset$. There is a wide collection of voting rules that have been defined in the literature (see, e.g., Brams and Fishburn, 2002) and here we focus on the following ones:

Positional scoring rules (PSR): Let (s_1, \dots, s_m) be a scoring vector such that $s_1 \geq \dots \geq s_m$ and $s_1 > s_m$. If a voter ranks candidate c at j -th position in her ballot, this gives s_j points to the candidate. The candidates with the highest score win. We focus on four particular PSR: *Plurality* with scoring vector $(1, 0, \dots, 0)$, *veto* with vector $(1, \dots, 1, 0)$, *2-approval* with vector $(1, 1, 0, \dots, 0)$, *3-approval* with vector $(1, 1, 1, 0, \dots, 0)$, and *Borda* with vector $(m-1, m-2, \dots, 0)$.

Copeland: The score of candidate c is the number of pairwise comparisons she wins (i.e., contests between c and another candidate a such that there is a majority of voters preferring c to a) minus the number of pairwise comparisons she loses. The candidates with the highest score win.

Maximin: The score of a candidate c is the smallest number of voters preferring it in any pairwise comparison. The candidates with the highest score win.

Single Transferable Vote (STV): At the first round the candidate that is ranked first by the fewest number of voters gets eliminated (ties are broken following a predetermined order of candidates). Votes initially given to the eliminated candidate are then transferred to the candidate that comes immediately after in the individual preferences. This process is iterated until one alternative is ranked first by a majority of voters.

With the exception of STV, all rules considered thus far are non-resolute, i.e., they associate a set of winning candidates with every profile of preferences. A tie-breaking rule is then used to eliminate ties in the outcome. In this paper we focus on *linear tie-breaking*: the set \mathcal{X} of candidates is ordered by $\prec_{\mathcal{X}}$, and in case of ties the alternative ranked highest by $\prec_{\mathcal{X}}$ is chosen as the unique outcome. Other forms of tie-breaking are possible, e.g., a random choice of a candidate from the winning set. The issue of tie-breaking has been shown to be crucial to ensure convergence of the iterative version of a voting rule (Lev and Rosenschein, 2012).

2.2 Iterative Voting

A classical problem studied in voting theory is that of manipulation: do individuals have incentive to misreport their preferences, in order to force a candidate they prefer as winner of the election? The celebrated Gibbard-Satterthwaite Theorem (Gibbard, 1973; Satterthwaite, 1975) showed that under arguably natural conditions all voting rules can be manipulated. Following this finding, a considerable amount of work has

been spent on devising conditions to avoid manipulation, e.g., in form of restrictive conditions on individual preferences, or in form of computational barriers that make the calculation of manipulation strategies too hard for the agents (Bartholdi and Orlin, 1991; Faliszewski and Procaccia, 2010).

In this paper, we take a different stance on manipulation: we consider the fact that individuals are allowed to change their preferences as a positive aspect of the voting process, that may eventually lead to a better result after a sufficient number of steps. Thus, we consider a sequence of repeated elections with individuals manipulating one at a time at each step, i.e., to modify her ballot in order to change the outcome of the election in her favor. The iteration process starts at \mathbf{b}^0 (which we shall refer to as the truthful profile) and continues to $\mathbf{b}^1, \dots, \mathbf{b}^K$. At each step only one individual $\tau(k)$ is allowed to manipulate, following a turn function τ (e.g., τ follows the order in which individuals are given), while all other individual ballots remain unchanged.² The iterative process converges if there exists a k_0 such that no individual has incentive to manipulate after k_0 steps of iteration.

The setting of iterative voting was first introduced and studied by Meir *et al.* (2010) for the case of the plurality rule, and expanded by Lev and Rosenschein (2012). In their work, the authors describe the iterated election process as a *voting game*, in which convergence of the iterative process corresponds to reaching a Nash equilibria of the game. They show that convergence is rarely guaranteed with most voting rules under consideration: for instance, the iterative version of PSRs and Maximin do not always converge, even with deterministic tie-breaking (i.e., not randomized). On the other hand, plurality always converges with any tie-breaking rule, as well as veto with linear tie-breaking. The tie-breaking rule is a crucial factor to obtain convergence. For instance, veto does not always converge if we choose a non-linear tie-breaking rule, and Copeland does not always converge even if we choose a linear tie-breaking rule, as we show in the following example:

Example 1. *Let there be two voters and three candidates, with $c \succ b \succ a$ as the tie-breaking order. The initial profile \mathbf{b}^0 has $a > b > c$ as the preference of the first voter and $c > a > b$ for the second voter. Candidate a wins in pairwise comparison with b , and no other candidate win any other pairwise comparison. Thus, the winner using the Copeland rule is a . Voter 2 has now an incentive to change her preferences to $c > b > a$, in which case by the tie-breaking order the winner is c , which is preferred by voter 2 in her truthful preference. Now voter 1 is unhappy, and changes her ballot to $b > a > c$ to force candidate b as the winner. This results in a further incentive for voter 2 to change her ballot to $c > a > b$, again forcing c as winner of the election. Finally, voter 1 changes again her ballot to $a > b > c$ to obtain a as winner of the election, moving back to the initial profile and creating a cycle of iterated manipulation.*

²Formally, a turn function takes as input the history of moves that have been played, i.e., the sequence of profiles $\mathbf{b}^0 \dots \mathbf{b}^K$ and outputs an individual in \mathcal{I} . A turn function is strongly related to the notion of scheduler in weakly acyclic games (Apt and Simon, 2012).

2.3 Restricted Manipulation Moves

The convergence of the iterated version of a voting rule can be obtained by restricting the set of manipulation strategies available to the agents. We now list a number of restrictions that have been studied in the literature, and we add two new definitions to this list. Let \mathbf{b}^k be the current profile at step k , \mathbf{b}^0 be the initial (truthful) profile, and F be a voting rule. Assume that $\tau(k) = i$.

Best response (no restriction): the manipulator i changes her full ballot by selecting the linear order which results in the best possible outcome for her truthful preference b_i^0 (Lev and Rosenschein, 2012).

k-pragmatist: the manipulator i moves to the top of her reported ballot the most preferred candidate following b_i^0 among those that scored in the top k positions (Reijngoud and Endriss, 2012).

M1: the manipulator i moves once the second-best candidate in b_i^0 to the top of her reported ballot b_i^{k+1} , unless the current winner $w = F(\mathbf{b}^k)$ is already her best or second-best candidate in b_i^0 .

M2: the manipulator i moves the most preferred candidate in b_i^0 which is above $w = F(\mathbf{b}^k)$ to the top of her reported ballot b_i^{k+1} , among those that become the new winner of the election after manipulation took place.

Different restrictions on manipulation moves induce different iterated versions of a voting rule:

Definition 1. *Let F be a voting rule and M a restriction on manipulation moves. $F^{M,\tau}$ associates with every profile \mathbf{b} the outcome of the iteration of F using turn function τ and manipulation moves in M if this converges, and \uparrow otherwise.*

In the following we shall omit the superscript τ from the notation when this will be clear from the context.

We follow three criteria in our choice of a restriction on the set of manipulation moves: (i) the convergence of the iterated voting rule associated with the restriction, (ii) the information to be provided to voters for computing their strategy³, and (iii) the computational complexity of computing the manipulation move at every step. An ideal restriction always guarantees convergence, requires as little information as possible, and is easy to compute.

As we pointed out at the end of the previous section, convergence is not guaranteed in most cases if the set of manipulation moves is not restricted (i.e., using best responses). Reijngoud and Endriss (2012) show convergence for PSRs using the k -pragmatist restriction, and we shall investigate convergence results for *M1* and *M2* in the following section. Let us move to the other two parameters: on the one hand, *M1* requires as little information as possible to be computed, i.e., only the winner of the current election, and is also very easy to compute. On the other hand, computing the best response requires an agent to have full knowledge of a profile, and may be computationally difficult to compute (see, e.g., Bartholdi

³This parameter is called *poll information function* by Reijngoud and Endriss (2012).

and Orlin, 1991). The k -pragmatist restriction has good properties: it is easy to compute, and the information required to compute the best strategy is just the set of the candidates ranked in the top k positions. $M2$ also requires little information for the agents: the vector of scores of candidates in case of scoring rules, the majority graph for Copeland and Maximin. In the case of STV the full profile is instead required. Moreover, from the point of view of the manipulator, $M2$ is computationally easy (i.e., polynomial) to perform.

3 Convergence and Axiomatic Properties

In this section we prove that the iterated version of PSR, Maximin and Copeland converge when using our two new restricted manipulation moves. We also analyze, for a number of axiomatic properties, the behavior of the iterated version of a voting rule with respect to properties of the non-iterated version. An *iteration step* is an effective change of ballot by one of the voters, moving from \mathbf{b}^j to a different profile \mathbf{b}^{j+1} .

Theorem 2. F^{M1} converges for every voting rule F .

Proof. The proof of this statement is straightforward from our definitions. The iteration process starts at the truthful profile \mathbf{b}_0 , and each agent is then allowed to switch the top candidate with the one in second position. Thus, the iteration process stops after at most n steps. \square

Theorem 3. F^{M2} converges if F is a PSR, the Copeland rule or the Maximin rule.

Proof. The winner of an election using a PSR, Copeland or Maximin is defined as the candidate maximizing a certain score (or with maximal score and higher rank in the tie-breaking order). Since the maximal score of a candidate is bounded, it is sufficient to show that the score of the winner increases at every iteration step (or, in case the score remains constant that the position of the winner in the tie-breaking order increases) to show that the iterative process converges.

Let us start with PSR. Recall that the score of a candidate c under PSR is $\sum_i s_i$ where s_i is the score given by the position of c in ballot b_i . Using $M2$, the manipulator moves to the top a candidate which lies above the current winner c . Thus, the position – and hence the score – of c remains unchanged, and the new winner must have a strictly higher score (or a better position in the tie-breaking order) than the previous one. The case of Copeland and Maximin can be solved in a similar fashion: it is sufficient to observe that the relative position of the current winner c with all other candidates (and thus also its score) remains unchanged when ballots are manipulated using $M2$. Thus, the Copeland score and the Maximin score of a new winner must be higher than that of c (or the new winner must be placed higher in the tie-breaking order). \square

While currently we do not have a proof of convergence for STV, we observed experimentally that its iteration always terminates on profiles with a Condorcet winner when a suitable turn function described in the following section is used.

Voting rules are traditionally studied using axiomatic properties, and we can inquire whether these properties extend from a voting rule to its iterated version. We refer to the literature for an explanation of these properties (see, e.g., Taylor,

2005). Let us call F_k^{M} the iterated version of voting rule F after k iteration steps. We say that a restricted manipulation move M *preserves* a given axiom if whenever a voting rule F satisfies the axiom then also F_k^{M} does satisfy it for all k .

Theorem 4. $M1$ and $M2$ preserve unanimity.

Proof. Assume that the iteration process starts at a unanimous profile \mathbf{b} in which candidate c is at top position of all individual preferences. If F is unanimous, then $F(\mathbf{b}) = c$, and no individual has incentives to manipulate either using $M1$ or $M2$. Thus, iteration stops at step 1 and $F_t^{M1}(\mathbf{b}) = c$ and $F_t^{M2}(\mathbf{b}) = c$, satisfying the axiom of unanimity. \square

Theorem 5. $M1$ and $M2$ preserve Condorcet consistency.

Proof. Let c be the Condorcet winner of a profile \mathbf{b} . If F is Condorcet-consistent then $F(\mathbf{b}) = c$. As previously observed, when individuals manipulate using either $M1$ or $M2$ the relative position of the current winner with all other candidates does not change, since the manipulation only involves candidates that lie above the current winner in the individual preferences. Thus c remains the Condorcet winner in all iteration steps \mathbf{b}^k . Since $F_k^{M1}(\mathbf{b}) = F(\mathbf{b}^k)$ and F is Condorcet-consistent, we have that $F_k^{M1}(\mathbf{b}) = c$ and thus F_k^{M1} is Condorcet consistent. Similarly for $M2$. \square

Other properties that transfer from a voting rule to its iterative version are neutrality and anonymity (supposing the turn function satisfies an appropriate version of neutrality and anonymity). The Pareto-condition does not transfer to the iterated version, as can be shown by adapting an example by Reijngoud and Endriss (2012, Example 3).

4 Experimental Evaluation of Restricted Manipulation Moves

In this section we evaluate our two restricted manipulation moves $M1$ and $M2$ under two important aspects. First, we measure whether the restricted iterative version of a voting rule has a higher Condorcet efficiency than the initial voting rule, i.e., whether the probability that a Condorcet winner (if it exists) gets elected is higher for the iterative rather than non-iterative rule. Second, we investigate whether the average position of the winner (i.e., its Borda score) in the truthful profile increases by allowing iterated restricted manipulation using $M1$ or $M2$. We compare our findings with the k -pragmatist restriction for $k = 2, 3$.

Our results are obtained using a program implemented in Java ver.1.6.0. The software generates profiles with uniform distribution, i.e., we make use of the *impartial culture assumption*. This assumption has received much criticism in recent years (see, e.g., Regenwetter *et al.*, 2006). However, it remains the most common assumption used in social choice theory, and thus represents the obvious starting point for our empirical evaluation. Our test set contains 10.000 profiles for both experiments, restricting to profiles with a Condorcet winner in the first case. We set the number of candidates to 5 as done by Reijngoud and Endriss (2012), and varied the number of voters from 20 to 100.

The turn function used in our experiments associates to each voter i a dissatisfaction index $d_i(k)$, which increases by one point for each iteration step in which the individual has an incentive to manipulate but is not allowed to do so by the turn function. At iteration step k the individual that has the highest dissatisfaction index is allowed to move (in the first step, and in case of ties, the turn follows the initial order in which voters are given). We also performed experiments using a sequential turn function which assigns turns depending on the order in which individuals are given (i.e., individuals with lower index has higher priority), obtaining similar results to those shown below (with the notable exception of STV, which does not converge using the sequential turn function).

The observed average number of steps to obtain convergence for iteration with moves $M1$, 2 and 3 -pragmatist grows linearly with respect to the number of voters. Instead, $M2$ reaches convergence after a moderately small number of steps, and the iteration takes place in between 10-30% of the cases. In Table 1, we report the average and maximal number of steps after which $M2$ converges using plurality, STV and Borda, among those profiles in which iteration actually takes place (over 10.000 profiles with 50 voters).

	# profiles with iteration	average # steps	maximal # steps
Plurality	2902	11.8	27
STV	1173	1.7	7
Borda	1961	8.1	31
2-Approval	2395	9.1	17

Table 1: Convergence figures for $M2$.

4.1 Condorcet efficiency

Figure 1 compares, for several voting rules, the Condorcet efficiency of the respective iterative version using restricted manipulation moves $M1$, $M2$, 2 -pragmatist and 3 -pragmatist. The number of voters is set to $n = 50$.

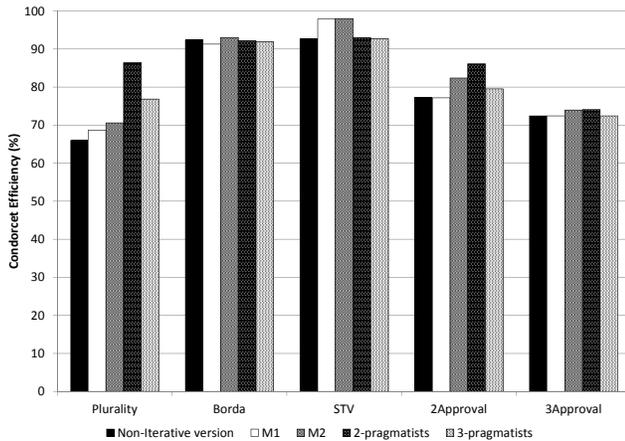


Figure 1: Condorcet efficiency of iterated voting rules with 50 voters and 5 candidates.

Except for the case of the Borda rule, the Condorcet efficiency of the iterated version of a voting rule improves with respect to the non-iterated version, and the growth is higher when voters manipulate the election using $M2$ rather than $M1$. The same behavior was observed varying the test set as well as performing similar experiment varying the number of candidates and voters. A plausible reason for this behavior is the difference in range of candidates that can be helped by the two manipulation moves. While $M1$ may help a Condorcet winner being elected only if it was ranked second by some of the individuals, $M2$ may help a candidate even if it was ranked lower. Let us also stress that while the increase in Condorcet efficiency using $M1$ is minimal, it is still surprising that such a simple move can result in a better performance than the original version of the voting rule. The 2 -pragmatist restriction perform quite well with the plurality rule and 2 -approval, while for the other rules our restriction $M2$ results in a better performance.

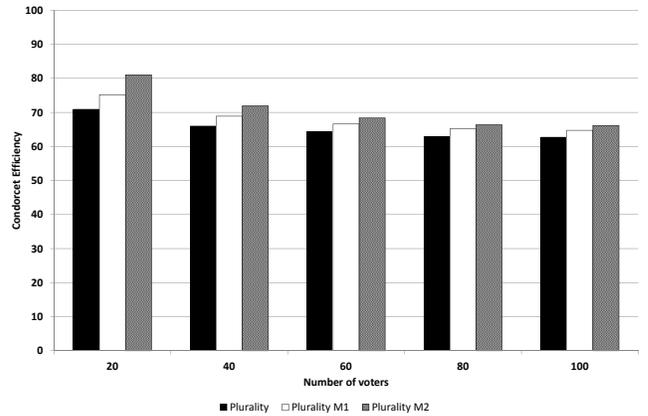


Figure 2: Condorcet efficiency for the plurality rule.

In the case of Plurality a significant increase can be obtained with both $M1$ and $M2$. In Figure 2 we show the trend in Condorcet efficiency when voters vary from $n = 20$ to $n = 100$. It can be observed that the increase is higher for smaller numbers of voters and stabilizes at around $n = 60$. The same behavior can be observed in Figure 3 for the case of STV.

STV has the highest performance of all voting rules considered thus far. STV has already a high Condorcet efficiency, but this is amplified by the use of manipulation moves. In Figure 3 we show that the efficiency can be augmented to more than 95 percent. As remarked earlier, we observed convergence in all profiles considered.

4.2 Borda score

In a second experiment we measured the average position of the winner (i.e. its Borda score) in the initial truthful profile, starting with the winner of the non-iterative version and comparing it with that of the winner of the iterated version after convergence. Given a candidate c , let p_i be the position of c in the initial preference b_i^0 of voter i . We compute the average position of c as $\frac{1}{n} \sum_{i=1}^n p_i$.

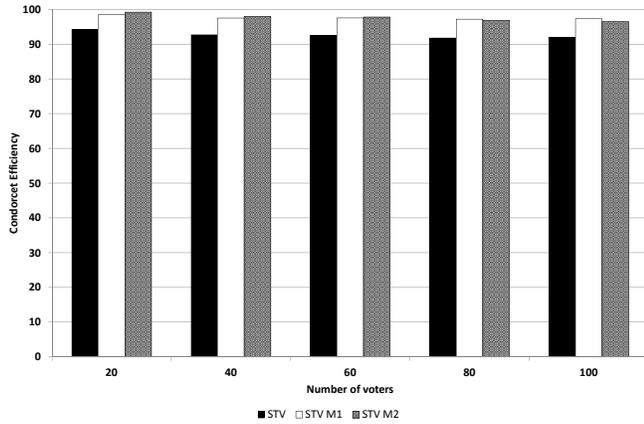


Figure 3: Condorcet efficiency for iterated STV.

In Figure 4 we compare the winner’s average position for the non-iterative version of some of the voting rules considered thus far. Observe that the voting rule with the highest average position of the winner is Borda, which by definition elects those candidates with the highest average position.

In the iterative setting we observed that the average position of the winner improves for all voting rules except for Borda. While a significant improvement can be obtained by using restricted manipulation moves $M2$, 2 and 3-*pragmatists*, only minor changes are observed when using $M1$. For instance, in Figure 5 we show the Borda score of iterated Plurality with move $M1$, $M2$, 2-*pragmatist* (2P) and 3-*pragmatist* (3P) comparing it with its non-iterated version. Similar graphs can be produced for the other voting rules.

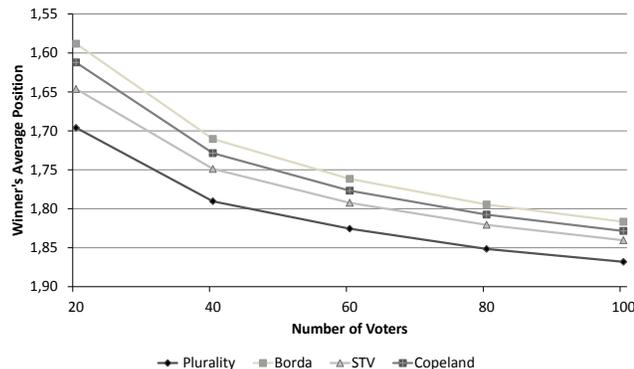


Figure 4: Winner average position (non-iterative case).

5 Conclusions and Future Work

This paper studies the iteration of classical voting rules allowing individuals to manipulate the outcome of the election using a restricted set of manipulation moves.

We provided two new definitions of manipulation moves $M1$ and $M2$ and showed that they lead to convergence for all voting rules considered (cf. Theorem 2 and 3) except for

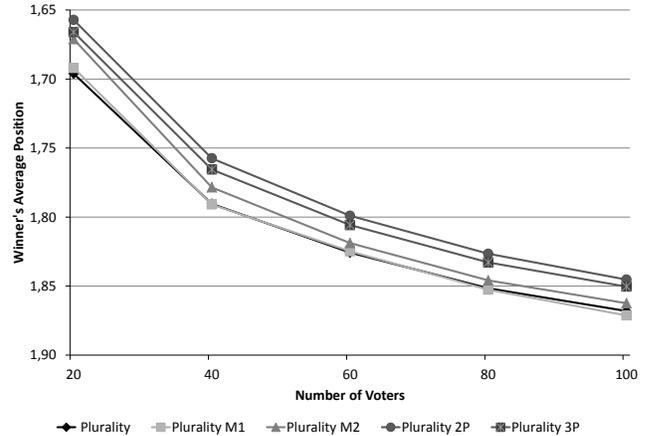


Figure 5: Winner average position for iterated plurality.

STV. We also showed that a number of axiomatic properties, such as unanimity and Condorcet consistency, are preserved in the iteration process. We evaluated the performance of our restricted manipulation moves with respect to the Condorcet efficiency of the iterated version of a voting rule as well as the Borda score of the winner in the initial truthful profile. We performed initial experiments using the impartial culture assumption. We showed that restricted manipulation in iterative voting yields a positive increase in Condorcet efficiency, and that, predictably, the best performance is obtained when more information is given to agents (cf. the case of STV with $M2$). The average position of the winner also grows with iteration, leading to elected candidates which are in average preferred to that of the initial non-iterated procedure.

While for some manipulation moves it can be questionable whether the gain in Condorcet efficiency and average position of the winner is worth the increase in computational cost deriving from iteration, our manipulation move $M2$ results in one of the best performances while keeping the number of iteration steps considerably low (cf. Table 1). For instance, by iterating 2-approvals with $M2$ it is possible to extract the Condorcet winner by asking a considerably lower amount of information from individuals than the full profile. Manipulation moves such as 2 and 3-*pragmatists* and $M1$ require instead a higher number of steps for convergence.

This work gives rise to a number of interesting directions to be explored in future research. First, different restrictions on manipulation moves may be considered, and their performance should be compared with that of existing definitions. We tested a move similar to $M2$, which did not restrict the choice of a candidate to those who become the new winners of the iterated election, obtaining a performance comparable to that of $M2$. Restricted manipulation moves may also be evaluated using other parameters, and we plan to test them on more realistic distributions of profiles of preferences (such as the urn model), or by exploiting data extracted from Internet-based polling services like Doodle.

References

- Stéphane Airiau and Ulle Endriss. Iterated majority voting. In *Proceedings of the 1st International Conference on Algorithmic Decision Theory (ADT-2009)*, 2009.
- Krzysztof R. Apt and Sunil Simon. A classification of weakly acyclic games. In *5th International Symposium on Algorithmic Decision Theory (SAGT-2012)*, 2012.
- John J. Bartholdi and James B. Orlin. Single transferable vote resists strategic voting. *Social Choice and Welfare*, 8:341–354, 1991.
- Steven J. Brams and Peter C. Fishburn. Voting procedures. In Kenneth Arrow, Amartya Sen, and Kotaro Suzumura, editors, *Handbook of Social Choice and Welfare*. Elsevier, 2002.
- Samir Chopra, Eric Pacuit, and Rohit Parikh. Knowledge-theoretic properties of strategic voting. In *Proceedings of the 9th European Conference on Logics in Artificial Intelligence (JELIA 2004)*, 2004.
- Piotr Faliszewski and Ariel D. Procaccia. AI’s war on manipulation: Are we winning? *AI Magazine*, 31(4):53–64, 2010.
- Allan Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, 1973.
- Omer Lev and Jeffrey S. Rosenschein. Convergence of iterative voting. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2012)*, 2012.
- Reshef Meir, Maria Polukarov, Jeffrey S. Rosenschein, and Nicholas R. Jennings. Convergence to equilibria in plurality voting. In *Proceedings of the Twenty-fourth conference on Artificial Intelligence (AAAI-2010)*, 2010.
- M. Regenwetter, B. Grofman, A. Marley, and I. Tsetlin. *Behavioral social choice*. Cambridge University Press, 2006.
- Annemieke Reijngoud and Ulle Endriss. Voter response to iterated poll information. In *Proceedings of the 11th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2012)*, June 2012.
- Mark Allen Satterthwaite. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187 – 217, 1975.
- Alan D. Taylor. *Social choice and the mathematics of manipulation*. Cambridge University Press, 2005.