Restricted Manipulation in Iterative Voting: Condorcet Efficiency and Borda Score

Umberto Grandi, Andrea Loreggia, Francesca Rossi[1],
K. Brent Venable[2], and Toby Walsh[3]

[1] University of Padova, umberto.uni@gmail.com, andrea.loreggia@gmail.com, frossi@math.unipd.it
[2] Tulane University and IHMC, kvenabl@tulane.edu
[3] NICTA and UNSW, 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 iterative restricted manipulation yields a positive increase in the Condorcet efficiency and Borda score for a number of standard voting rules.

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 at least choose a voting rule that is computationally difficult to manipulate for single agents or coalitions.

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, i.e., 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.

There are two prototypical situations in which iterative manipulation takes place. The first example is represented by the response of an electorate to a series of information polls about the result of a political election. At each step individuals may realize that their favorite candidate does not have chances to win and report a different preference in the subsequent poll. The second example is Doodle,[4] a very popular on-line system to select a time slot for a meeting. 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.

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), and on the convergence of plurality decisions between multiple agents (Airiau and Endriss, 2009). 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, contrarily 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. We then perform an experimental analysis of four restrictions on the set of manipulation strategies. For voting rules that are not Condorcet consistent, we test whether their Condorcet efficiency (that is, the probability to elect the Condorcet winner) improves by adopting the iterative versions. The second param-

[4] http://doodle.com/.

eter that we test is the Borda score of the winner in the truthful profile. Our results show that, with the exception of the Borda rule, both parameters never decrease in iteration, and a significant increase can be observed when the number of candidates is higher than the number of voters, as it is the case in a typical Doodle poll.

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 theory that we shall use in this paper, we present the setting of iterative voting, and we define three notions of restricted 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 p_i 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, which we also assume is a linear order over $\mathcal{X}$, and a profile $\mathbf{b} = (b_1, \dots, b_n)$ is defined by the choice of a ballot for each of the individuals. 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 in this paper we focus on the following definitions:

Positional scoring rules (PSR): Let $(s_1, \dots, s_m)$ be a scoring vector such that $s_1 > s_m$ and $s_1 \geq \dots \geq s_m$. 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 PSRs: *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). 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 Strategic Manipulation and Iterative Voting

A classical problem studied in voting theory is that of manipulation: do individuals have incentive to misreport their truthful 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 very 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 the form of restrictions on individual preferences or in the 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.

We consider a sequence of repeated elections with individuals manipulating one at a time at each step. The iteration process starts at $\mathbf{b}^0$ (which we shall refer to as the *truthful profile*) and continues to $\mathbf{b}^1, \dots, \mathbf{b}^k, \dots$. A turn function τ identifies one single individual i_k that is allowed to manipulate at step k, while all other individual ballots remain unchanged (e.g., τ follows the order in which individuals are given)[5]. The individual manipulator uses the *best response* strategy: she changes her full ballot by selecting the linear order which results in the best possible outcome based on her truthful preference. In case the result of the election cannot be changed to a better candidate for the manipulator, we say that the individual does not have incentive to manipulate. 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, and that this property is highly dependent on the tie-breaking rule under consideration. For instance, iterative plurality always converges with any tie-breaking

[5] Formally, a turn function takes as input the history of moves that have been played at step k, i.e., the sequence of profiles $\mathbf{b}^0 \dots \mathbf{b}^{k-1}$ 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).

rule, while the iterative version of PSRs and Maximin do not always converge. The following example shows that the iterative version of the Copeland rule does not converge even if we choose a linear tie-breaking rule:

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

2.3 Restricted Manipulation Moves

Given that convergence is not guaranteed when voters manipulate choosing their best response, an interesting problem is to devise suitable restrictions on the set of manipulation strategies available to the agents in order to obtain convergence. Initial work on this topic was done by Reijngoud and Endriss (2012). In this section we review their definition and we introduce two novel notions of restricted manipulation. 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 $\tau(k) = i$.

- *k-pragmatist* (Reijngoud and Endriss, 2012): 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.[6]
- *M1*: the manipulator i moves 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 $w = F(\mathbf{b}^k)$ in b_i^k to the top of her reported ballot b_i^{k+1}, among those that can become the new winner of the election.

Under the k-*pragmatist* restriction voters are allowed to move to the top of their reported ballot the individual which they prefer amongst the top k candidates ranked by F. *M1* allows only a very simple move: switch the first and second candidate in the manipulator's ballot unless she is already satisfied, i.e., in case the current winner is ranked first or second in her truthful ballot. *M2* is more complex: the manipulator selects those candidates that she prefers to the current winner in her current ballot at step k; then, starting from the most preferred one in the truthful ballot, she tries to put such candidate on the top of her current ballot and computes the outcome of the election; the

[6] Note that all voting rules considered in this paper can be easily extended to output a ranking of the candidates rather than just a single winner.

first candidate which succeeds in becoming the new winner of the manipulated election is the one chosen for the top position of her reported ballot.

While the choice of these restrictions may at first seem arbitrary, we believe they represent three basic prototypes of simple manipulation strategies for agents with bounded computational capabilities and limited access to information. Indeed, when evaluating restrictions on the set of manipulation moves we followed three criteria: (i) the convergence of the iterative voting rule associated with the restriction, (ii) the information to be provided to voters for computing their strategy, 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 computationally easy to compute. 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. 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 candidates which are ranked in the top k positions. *M2* also requires little information: the candidates' final score in case of scoring rules, the majority graph for Copeland and Maximin. Instead, in the case of STV the full profile is required. From the point of view of the manipulator, *M2* is computationally easy (i.e., polynomial) to perform.

We conclude by defining the iterative version of a voting rule depending on the different assumptions we can make on the set of manipulation moves:

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.*

3 Convergence and Axiomatic Properties

In this section we prove that the iterative versions of PSR, Maximin and Copeland converge when using our two new restrictions on the manipulation moves. We also analyze, for a number of axiomatic properties, the behavior of the iterative version of a voting rule with respect to the properties of the non-iterative version.

Theorem 1. *$F^{M1,\tau}$ converges for every voting rule F and turn function τ.*

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 allowed to switch the top candidate with the one in second position. The iteration process stops after at most n steps.

Theorem 2. *$F^{M2,\tau}$ converges for every turn function τ 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) remain unchanged when ballots are manipulated using *M2*. Thus, the Copeland score and the Maximin score of a new winner must by higher than that of c (or the new winner must be placed higher in the tie-breaking order).

This proof generalizes to show the convergence of the *M2*-iterative version of any voting rule which outputs as winners those candidates maximizing a given notion of score. 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, which is 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 iterative version. We refer to the literature for an explanation of these properties (see, e.g., Taylor, 2005). Let us call F_t^M the iterative version of voting rule F after t iteration steps (we omit the superscript τ, indicating that these results hold for every turn function). We say that a restricted manipulation move M *preserves* a given axiom if whenever a voting rule F satisfies the axiom then also F_t^M does satisfy it for all t.

Theorem 3. *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.

Theorem 4. *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*.

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 iterative 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 four restrictions on the set of manipulation moves, namely 2 and 3-*pragmatists*, *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 observe the variation of the Borda score of the winner, i.e., we compare the average position of the current winner in the initial truthful profile at convergence with the value of the same parameter in the initial profile. We focus on four voting rules: plurality, STV, Borda, 2 and 3-approval. Our findings show that both parameters never decrease by allowing iterated restricted manipulation, and that a substantial increase can be observed in case the number of candidates is higher than the number of voters (e.g., in a Doodle poll). We conclude the section by reporting on some initial experiments with real-world datasets.

4.1 Experimental setting

We generated profiles using the Polya-Eggenberger urn model (see, e.g., Berg, 1985). Individual ballots are extracted from an urn initially containing all $m!$ possible ballots, i.e., all linear orders over m candidates, and each time we draw a vote from the urn we put it back with a additional copies of the same vote. In this way we generate profiles with correlated preferences and we control the correlation ratio with the parameter a. In our experiment we tested three different settings: the *impartial culture assumption* (IC) when $a = 0$, the UM10 with 10%-correlation when $a = \frac{m!}{9}$, and the UM50 with 50%-correlation when $a = m!$. We generated 10.000 profiles for each experiment, restricting to profiles with a Condorcet winner when testing the Condorcet efficiency.

Our results are obtained using a program implemented in Java ver.1.6.0. We model two prototypical examples of iterative voting: in the *electoral simulation* we set the number of candidates to $m = 5$ and the number of voters to $n = 500$ to model situations in which a large population needs to decide on a small set of candidates. In the *Doodle simulation* we set the number of candidates to $m = 25$ and the number of voters to $n = 10$ to model a small group of people deciding over a number of time slots. In both cases we performed additional experiments varying the number of voters and candidates but keeping their ratio fixed without observing significant variation in the results.

The turn function used in our experiments associates with 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 initial experiments using a sequential turn function which assigns turns depending on the order in which individuals are given obtaining similar results to those shown below (with the exception of STV, which does not converge using the sequential turn function).

It is interesting to observe that the higher the correlation in the profile the smaller the number of profiles in which iteration takes place (with the notable exception of the

Borda rule). In Figure 1 it is shown, for the Doodle simulation, the percentage of profiles in which iteration takes place for the three different correlation ratios considered. In the case of plurality, convergence is reached after an average of 3 steps and a maximal of 9 steps. The figures for the other voting rules are similar.

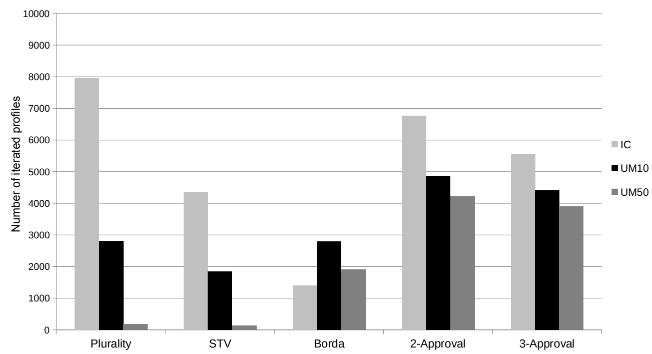

Fig. 1. Number of profiles with iteration compared to the correlation ratio.

4.2 Condorcet efficiency

Figures 2 and 3 compare the four restrictions on manipulation moves with respect to the Condorcet efficiency of the iterative version of the five voting rules under consideration, respectively for the Doodle simulation and the electoral simulation. In both experiments the correlation ratio is set at 10%.

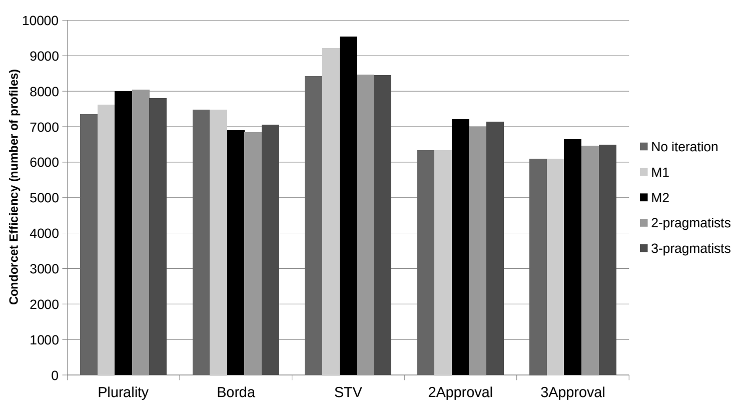

Fig. 2. Doodle experiment with UM10: Condorcet efficiency.

Except for the case of the Borda rule, the Condorcet efficiency of the iterative version of a voting rule always improves with respect to the non-iterative version, and the growth is consistently higher when voters manipulate the election using *M2* rather than *M1*. 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 performs quite well with the plurality rule in both experiments. STV has the highest performance of all voting rules considered thus far with respect to Condorcet efficiency and this performance is amplified by the use of iterated manipulation, resulting in the election of a Condorcet winner in almost 95 percent of the cases. As remarked earlier, we observed convergence in all profiles considered. The increase in Condorcet efficiency is more noticeable in the Doodle simulation rather than in the electoral situation. Thus, when the number of individuals is considerably higher than the number of alternatives the iterative process leads to a minimal increase in Condorcet efficiency.

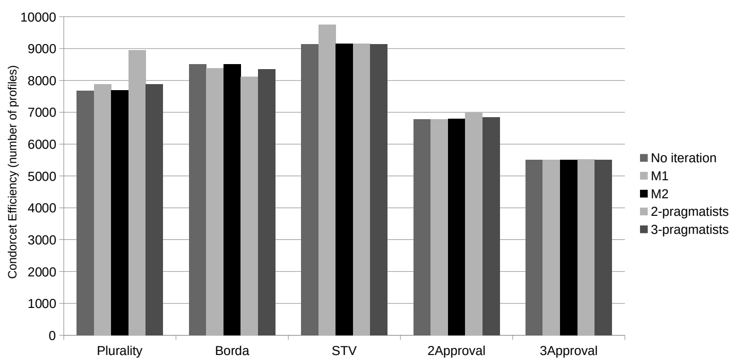

Fig. 3. Electoral experiment with UM10: Condorcet efficiency.

We also run the same two experiments with different correlation ratios: Using the IC assumption the increase in Condorcet efficiency is more significant, while with the UM50 assumption the results are much less perturbed by iteration. This should not come as a surprise, given that the amount of profiles in which iteration takes place decrease rapidly with the growth of the correlation ratio.

4.3 Borda score

The second parameter we used to assess the performance of restricted manipulation moves is known in the literature as the Borda score. Given a candidate c, let p_i be the position of c in the initial preference b_i^0 of voter i (from bottom to top, i.e., if a candidate is ranked first she gets $m-1$ points, while if she is ranked last she gets 0 points). We compute the Borda score of c as $\sum_{i=1}^{n} p_i$.

For each voting rule and each restriction on the set of manipulation moves we compared the score of the winner of the non-iterative version with that of the winner of the iterative version after convergence. Since the Borda rule elects by definition those candidates with the highest Borda score, we did not evaluate the iterative version of Borda with respect to this parameter. Our results showed that in both the Doodle and the electoral simulation with UM10 the Borda score increases minimally if we allow for iterated restricted manipulation, resulting in a chart similar to that in Figure 3. The best results are in this case obtained by using *M2* and 2-*pragmatists* restriction with 2 and 3-approval. As in the previous section, by decreasing the correlation of the generated preferences we obtain a more significant increase in the Borda score after iteration.

4.4 Real-world Datasets

We performed initial experiments using data from *Preflib* (Mattei and Walsh, 2013), a library of preference datasets collected from various sources. In order to mimic the original preference distribution, we generated 10.000 profile with 5 candidates and 500 voters drawing votes with impartial culture assumption from two original datasets: the Netflix Prize Data (Bennett and Lanning, 2007) and the Skating Data. What we observed is that preferences contained in such datasets are quite correlated, with iteration taking place in just a handful of profiles (in the order of 5–10 per 10.000 profiles). Experiments run on data from political elections may have a chance to lead to more significant results, once our setting has been adapted to the case of partial orders over candidates, as required by the electoral datasets available at present state.

5 Conclusions and Future Work

This paper studies the effect of iteration on classical voting rules by 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 1 and 2) except for STV. We showed that a number of axiomatic properties, such as unanimity and Condorcet consistency, are preserved in the iteration process. We evaluated experimentally the performance of our restricted manipulation moves with respect to the Condorcet efficiency of the iterative version of a voting rule as well as the Borda score of the winner in the initial truthful profile. We performed two simulations based on prototypical examples of iterated manipulation: the first simulation with the number of candidates smaller but comparable to the number of voters to model scheduling with Doodle, and the second with the number of voters much higher than that of candidates to model iterated polls before a political election. With the exception of the Borda rule, we showed that restricted manipulation in iterative voting yields a positive increase in both the Condorcet efficiency and Borda score and that the best performance is obtained when the number of candidates is higher than the number of individuals.

In future work we plan to analyze different versions of manipulation moves, and to compare their performance with that of the existing definitions. By adapting the framework defined in this paper to account for preferences expressed as partial orders it will also be possible to exploit preference data on real-world elections, to assess with more accuracy the effects of iterated restricted manipulation on more realistic distributions of preferences.

Bibliography

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.

James Bennett and Stan Lanning. The Netflix prize. In *Proceedings of the KDD Cup and workshop*, 2007.

Sven Berg. Paradox of voting under an urn model: The effect of homogeneity. *Public Choice*, 47(2):377–387, 1985.

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.

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.

Nicholas Mattei and Toby Walsh. Preflib: A library of preference data (www.preflib.org). 2013.

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.

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.