

Convergence and Quality of Iterative Voting Under Non-Scoring Rules

(Extended Abstract)

Aaron Koolyk
The Hebrew University
Jerusalem, Israel
aaron.koolyk@mail.huji.ac.il

Omer Lev
University of Toronto
Toronto, Canada
omerl@cs.toronto.edu

Jeffrey S. Rosenschein
The Hebrew University
Jerusalem, Israel
jeff@cs.huji.ac.il

ABSTRACT

Iterative voting is a social choice mechanism whereby voters are allowed to continually make strategic changes to their stated preferences until no further change is desired. We study the iterative voting framework for several common voting rules and show that, for these rules, an equilibrium may never be reached. We also consider several variations of iterative voting and show that with these variations equilibrium likewise may not be reached. Finally, we present an empirical analysis of the quality of candidates elected through iterative voting.

Keywords

Social Choice; Equilibrium; Iterative Voting

1. INTRODUCTION

The topic of voting, that is, how to aggregate diverse individual preferences into a collective decision, is of great importance in many automated agent scenarios; it has thus been the topic of much research in multiagent systems. One innovative voting model that was recently proposed is that of *iterative voting* [3]. Classic voting protocols terminate with the announcement of a winner following a single round of ballot submission; in iterative voting, on the other hand, as long as some voter wishes to change their vote, they may do so (when multiple voters wish to change their votes simultaneously, an arbitrary voter is chosen). The process terminates when no voter wishes to change their vote.

Iterative voting thus embraces the inevitable manipulability of voting, and considers agents' uniform ability to vote strategically as a collective opportunity. But in addition to being an intriguing method for reaching consensus, iterative voting has been proposed as a formal solution concept for voting games. Standard Nash equilibria are of limited utility in voting scenarios in which the group outcome is generally robust to changes in any single voter's action. The set of iterative voting equilibria, however, is a subset of Nash equilibria, and in particular those iterative voting equilibria reachable from the truthful profile could be considered a more natural (or meaningful) solution concept.

Appears in: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016)*, J. Thangarajah, K. Tuyls, C. Jonker, S. Marsella (eds.), May 9–13, 2016, Singapore.

Copyright © 2016, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

The most salient questions regarding iterative voting thus have two interpretations. Regarding iterative voting as a method for reaching consensus: does the process terminate? And does it arrive at “good” conclusions? In parallel, regarding iterative voting as a solution concept: do solutions exist? And what is their “price of stability/anarchy”?

2. CYCLICITY

We consider the convergence of several common non-scoring voting rules under the following iterative dynamics.

BR: An ordered pair of profiles is in the *best response* dynamic if the preferences of all voters but one are identical; the voter whose preferences change prefers the outcome of the second profile to that of the first profile; and of all possible changes to his preferences, the outcome under the second profile is preferred at least as much as the outcome under any other possible change.

TOP: In this dynamic, a changed vote necessarily assigns the candidate which the voter will make the winner the top spot in the new preference order. In many of the voting rules we consider (and in any weakly-monotone rule) this dynamic is a subset of the best-response dynamic, and, indeed, it generalizes the dynamic used in [3].

TB: This dynamic requires the new winner to be at the top of the new ballot, and the previous winner to be at the bottom. While in many scoring rules (e.g., plurality and veto) this is a subset of best response moves (and generalizes those used in [2, 5]), this is not true in general, and particularly in the voting rules we study in this work.

KT: This dynamic restricts best response to those moves with minimum Kendall-Tau distance from the previous vote.

SWAP: This restrictive dynamic, allows only changes of a single adjacent swap (a new vote is of a KT distance of one from the previous vote; a notion referred to as a ‘swap’ in the literature on bribery).

The following result establishes that neither best response iterative voting nor any of the four variations on iterative voting can be guaranteed to converge for a number of well-studied voting rules; see [1, 4].

THEOREM 1. *Maximin, Copeland, Bucklin, STV, Second Order Copeland, and Ranked Pairs with linear order tie-breaking do not converge under BR, TOP, TB, KT or SWAP.*

3. QUALITY

In this section, we present simulations of best-response-iterative voting for the six rules we have studied. Simulations were run for each rule with four candidates and both

ten voters and twenty five voters. 10,000 initial (truthful) profiles were sampled uniformly. Each profile evolved under best response dynamics and was run to completion (or detection of a cycle) 20 times. Each step was made by uniformly sampling a voter and a best-response move.

For all of the rules, cycles occur quite rarely; see Table 1. So although we have shown that all these rules *can* cycle, the frequency with which they do is very low.

	Number of Cycles	Initial winner re-selected
Maximin 10	180 (0.09%)	28,244 (14.12%)
Maximin 25	55 (0.03%)	19,036 (9.52%)
Copeland 10	1132 (0.57%)	34,641 (17.32%)
Copeland 25	107 (0.05%)	23,540 (11.77%)
Bucklin 10	545 (0.27%)	33,932 (16.97%)
Bucklin 25	191 (0.10%)	34,834 (17.42%)
STV 10	79 (0.04%)	10,518 (5.26%)
STV 25	25 (0.01%)	13,777 (6.89%)
SOC 10	943 (0.47%)	33,715 (16.86%)
SOC 25	110 (0.06%)	22,607 (11.30%)
Ranked Pairs 10	226 (0.11%)	26,493 (12.25%)
Ranked Pairs 25	40 (0.02%)	17,967 (8.98%)

Table 1: Statistics of the 200,000 Paths.

The rarity of cycles suggest that perhaps iterative voting could be considered even with these rules and best response dynamics. In the rare case of a cycle, the situation could be deferred to a cycle-breaking rule. The election could be run again, or a profile within the cycle (or an outcome among the cycle’s outcomes) could be randomly chosen, for instance.

With regard to the quality of the outcomes, first note that quite often iterative voting leads to the original outcome. Many such instances are the result of original profiles which are non-manipulable. But many are also the result of manipulations, whose equilibrium reverted to the original winner; see Table 1 with respect to the latter.

Next, we assess the change in Borda score and Condorcet efficiency. The change between the Borda score of the original winner (in the original profile) and the Borda score of the equilibrium winner (in the original/truthful profile) is shown in Table 2.

With regard to Condorcet efficiency, we consider, for each of the 10,000 profiles (of which a Condorcet Winner existed in 4764 and 8413, respectively), whether the CW is selected by the voting rule, and whether it is selected in equilibrium. The latter is presented in Table 2 in terms of efficiency (out of 10,000) after aggregating equilibria over non-cycling paths.

Of the two rules that are not Condorcet consistent, Bucklin and STV actually improve their efficiency under iterative voting. These two rules also fared well under Borda criteria, suggesting that iterative Bucklin and iterative STV could be considered improvements on their static counterparts.

4. CONCLUSIONS

We have shown that for a number of common non-scoring voting rules, iterative voting under best response dynamics does not always converge. Even after restricting the dynamics to allow voters only limited changes to their ballots—

	Percent increase of avg. BS	Initial winner is CW	Equilibrium winner is CW
Maximin 10	-0.84	4764	0.47
Maximin 25	-0.34	8413	0.82
Copeland 10	-0.48	4764	0.47
Copeland 25	-0.10	8413	0.81
Bucklin 10	2.10	3717	0.45
Bucklin 25	2.19	4461	0.58
STV 10	0.40	4610	0.47
STV 25	0.32	7795	0.83
SOC 10	-0.46	4764	0.47
SOC 25	-0.17	8413	0.81
Ranked Pairs 10	-0.40	4764	0.47
Ranked Pairs 25	-0.35	8413	0.82

Table 2: Borda Score and Condorcet Efficiency at Equilibrium

whether by constraining the positioning of affected candidates, or by prioritizing minor ballot changes—they still do not always converge.

On the other hand, we have shown empirically that cycles occur quite infrequently for all these rules. Given the sparsity of cycles and the possibility of detecting and overcoming them, it would seem that best response iterative voting could be considered even for rules that do not always converge. The quality of the outcome of iterative voting is never much worse, and sometimes better, than non-iterative voting.

Acknowledgments

This research has been partly funded by Microsoft Research through its PhD Scholarship Program, and by Israel Science Foundation grant #1227/12.

REFERENCES

- [1] A. Koolyk. Iterative voting: Convergence, cyclicity, and quality. Master’s thesis, Hebrew University, 2015.
- [2] O. Lev and J. S. Rosenschein. Convergence of iterative voting. In *The Eleventh International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 611–618, Valencia, Spain, June 2012.
- [3] R. Meir, M. Polukarov, J. S. Rosenschein, and N. Jennings. Convergence to equilibria of plurality voting. In *The Twenty-Fourth National Conference on Artificial Intelligence*, pages 823–828, Atlanta, Georgia, July 2010.
- [4] S. Obraztsova, E. Markakis, M. Polukarov, Z. Rabinovich, and N. R. Jennings. On the convergence of iterative voting: How restrictive should restricted dynamics be? In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 993–999, Austin, Texas, January 2015.
- [5] R. Reyhani and M. C. Wilson. Best reply dynamics for scoring rules. In *The 20th European Conference on Artificial Intelligence (ECAI 2012)*, pages 138–144, Montpellier, France, August 2012.