Beneficial strategic reasoning in iterative voting

Joint work with Umberto Grandi, Francesca Rossi,
K. Brent Venable and Toby Walsh

Andrea Loreggia
Department of Mathematics
University of Padova
23 January 2014

UNIVERSITÀ DEGLI STUDI DI PADOVA

Università degli Studi di Padova

Manipulation in elections is usually considered a bad thing, to be avoided or at least to be made computationally difficult to achieve.

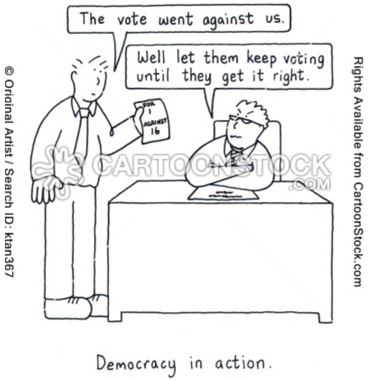

Can we get a better outcome with iterated manipulation of simple rules?

Practical Examples

Università degli Studi di Padova

In practice, iterative manipulation does occur:

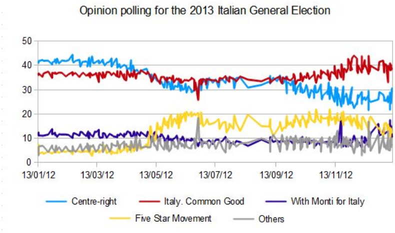

Iterative response to repeated polls

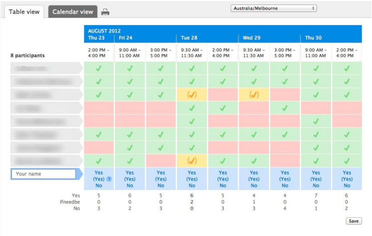

Approval voting with iterative manipulation

Image source: Wikipedia, Doodle.com

1. The setting:
   - Voting rules (in brief)
   - Iterative voting
   - Restricted manipulation: *M1* and *M2*
2. Theoretical evaluation
   - Convergence: Yes! (unknown for STV)
   - Axiomatic properties: transfer to iterative rules
3. Experimental evaluation
   - Condorcet efficiency: Increase!
   - Average position of the winner: Increase!
   - Results with impartial culture assumption, urn model, real-world data

Voting Systems

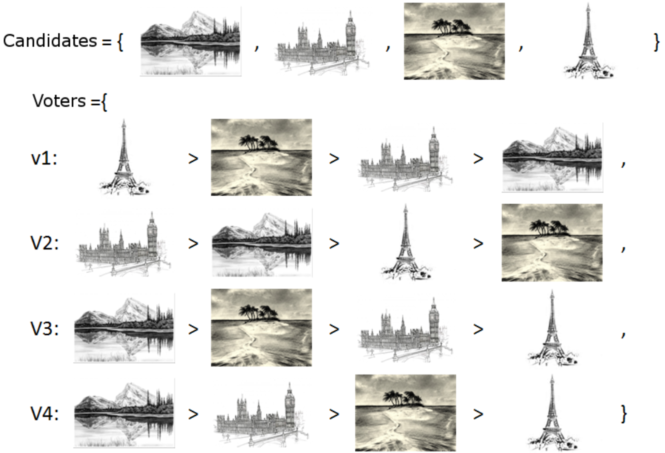

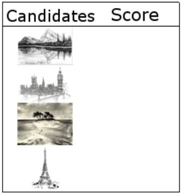

Voting Systems - Plurality

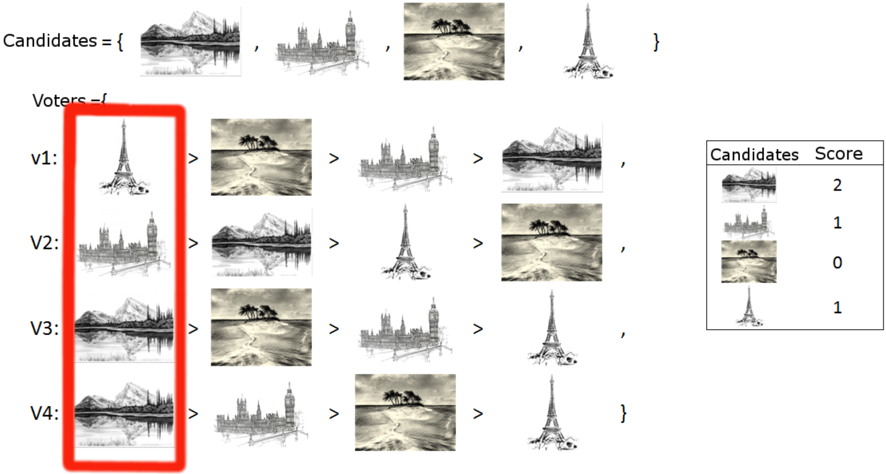

Voting Systems - Borda

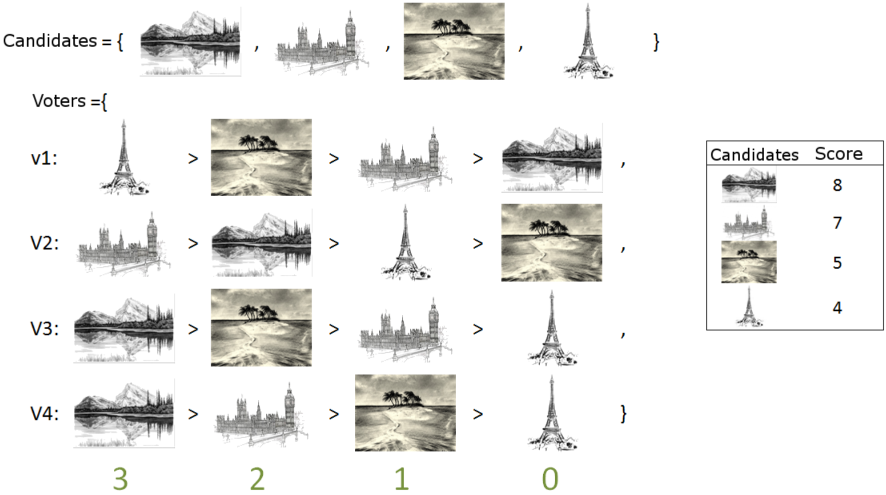

Candidates = {  ,  ,  ,  }

Voters ={

v1:  >  >  >  ,

V2:  >  >  >  ,

V3:  >  >  >  }

Candidates	Score
	2
	3
	0
	1

Formally: an election is $E = (C, V)$, such that:

- C is a set of m alternatives or candidates
- V is a set of n voters, they are represented by the list of their preferences over the candidates

R is a voting rule, it is used by the voting system to aggregate the preferences and choose the winner
A tie-breaking rule could be used to define the winner when ties happens. It can be deterministic, linear or random.

Different kind of voting rules:

- Positional scoring rules (they use scoring vector):
  - Plurality $v =< 1, 0, \ldots, 0 >$
  - Veto $v =< 1, 1, \ldots, 1, 0 >$
  - K-Approval $v =< \underbrace{1, 1, \ldots, 1}_{k}, 0, 0, \ldots, 0 >$
  - Borda $v =< m - 1, m - 2, \ldots, 0 >$
- Approval
- Single Transferable Voting
- Condorcet-consistent methods

  (They elect the Condorcet Winner, when it exists. Which is the candidate that defeats all the other in the pair-wise-competition):
  - Copeland
  - Maximin
  - Cup rule

Strategic Actions - Borda manipulation

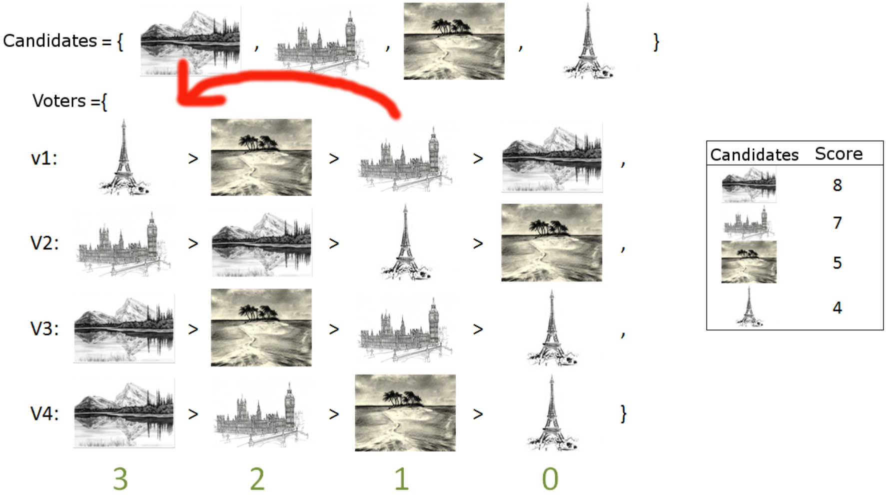

Università degli Studi di Padova

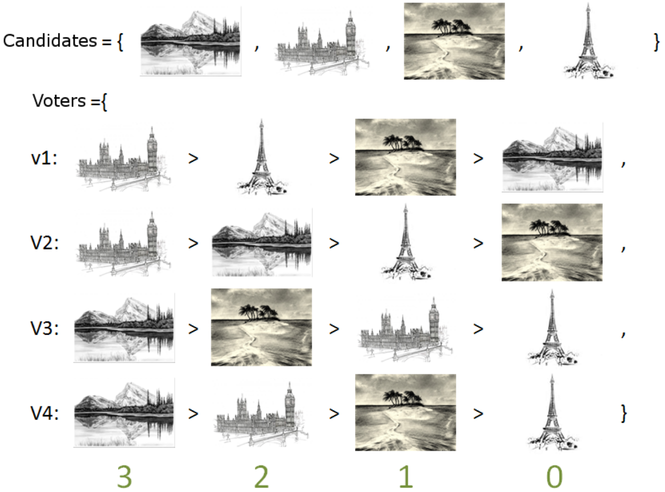

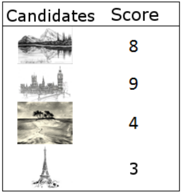

Manipulation occurs whenever a voter changes her ballot in her favour:
Is there any chance to avoid manipulation?

Theorem [Gibbard-Satterthwaite]

Given a voting rule F, one of the following facts must be true: (i) there is a candidate that never wins (ii) F is a dictatorship, (iii) F can be manipulated.

Needless to say, all voting rules presented are manipulable...

A. Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 1973.
M. A. Sattertwaithe, Strategy-proofness and Arrows conditions... *JET*, 1975.

Strategic manipulation in elections defines a voting game:

- Strategies are linear orders: individuals can change their preferences to obtain a better outcome
- The outcome is the result of the voting rule
- Utilities are defined by the truthful preferences of individuals

Strategic manipulation in elections defines a voting game:

- Strategies are linear orders: individuals can change their preferences to obtain a better outcome
- The outcome is the result of the voting rule
- Utilities are defined by the truthful preferences of individuals

Definition

Given a set of manipulation moves M, a voting rule F (and a turn function) the iterated voting rule F^M associates with every profile $\mathbf{b}$ the outcome of convergent iteration of manipulation moves in M (or $\uparrow$ if it does not converge).

Unrestricted manipulation does not always converge! But if it does, it converges to a Nash equilibrium of the voting game associated to F.

R. Meir Et Al. Convergence to equilibria in plurality voting. AAAI-2010.

O. Lev and J. S. Rosenschein. Convergence of iterative voting. AAMAS-2012.

Iterative Voting - Borda

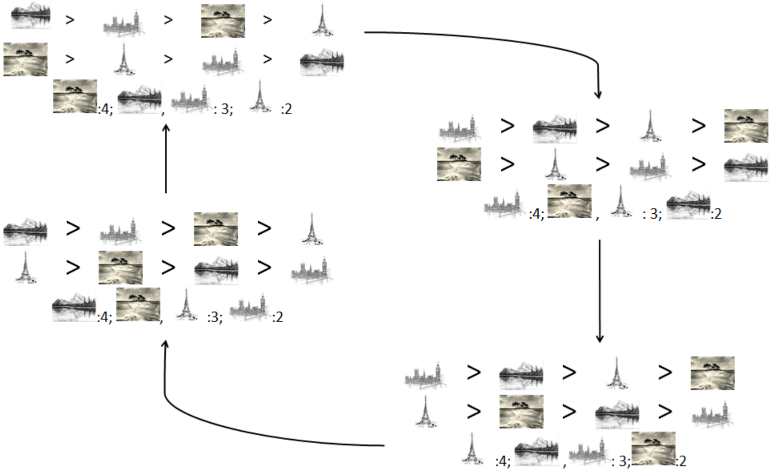

Voting Rule	Tie-breaking	Agent weight	Result
Plurality	Deterministic	1	C
Plurality	Random	1	C
Borda	Any	any	NC
Veto	Linear	unweighted	C
K-Approval, $k \geq 3$	Linear	any	NC
Maximin	Deterministic	any	NC

Table: C: converge - NC: non-converge

[Meir, Polukarov, Rosenschein, Jennings - 2010][Lev, Rosenschein - 2012]

Manipulation moves studied in the literature:

- **Best response** (no restriction): choose the ballot that changes the outcome of the election in the best way.
- **k-pragmatist**: put in first position your favourite candidate among the top k in the outcome of the voting rule.

A. Reijngoud and U. Endriss. Voter response to iterated poll information. AAMAS-2012.

Manipulation moves studied in the literature:

- **Best response** (no restriction): choose the ballot that changes the outcome of the election in the best way.
- **k-pragmatist**: put in first position your favourite candidate among the top k in the outcome of the voting rule.

How to evaluate a restriction on manipulation moves?

Convergence	**Computation**	**Information**
Guaranteed	Not costly	Low
(small number of steps)	(not NP-hard!)	(top candidate, scores..)

- Restricted Best Response (RBR): the new winner has a score greater or equal than the winner at previous step
- Bounded WPs: bounded number of worsening flips, that is when a voter moves a more preferred candidate below a less preferred candidate

Rule	Restriction	Tie-break	Result
Veto		NON-Linear	NC
Copeland		Any	NC
Copeland	RBR	NON-Linear	NC
Copeland	RBR	Linear	C
STV		Linear	NC
Approval		Linear	NC
Maximin		Linear	C
Cup Rule		Linear	NC
Cup Rule	bounded WPs	Linear	C

Table: C: converge - NC: non-converge

Scoring protocols and Condorcet efficiency:

- M1: give a second chance to the candidate who is ranked in the second position in the truthful preference, by moving her to the top
- M2: move to the first position the candidate that can become a winner and it is more preferred than the current winner

With M1 and M2 little information is needed, e.g. for scoring protocols just the current winner and the final scores.

Iteration starts at $\mathbf{b}^0$ (truthful) and continues to $\mathbf{b}^1, \ldots, \mathbf{b}^k$ until convergence

M1

Move to the top the second-best candidate in b_i^0 (truthful), unless the current winner $w = F(\mathbf{b}^k)$ is already her best or second-best candidate in b_i^0 (truthful)

$$\begin{array}{c} a \succ b \succ c \succ d \\ c \succ b \succ a \succ d \\ d \succ b \succ c \succ a \\ \hline \text{Plurality: } a \end{array} \quad \rightarrow \quad \begin{array}{c} a \succ b \succ c \succ d \\ b \succ c \succ a \succ d \\ d \succ b \succ c \succ a \\ \hline \text{Plurality: } a \end{array} \quad \rightarrow \quad \begin{array}{c} a \succ b \succ c \succ d \\ b \succ c \succ a \succ d \\ b \succ d \succ c \succ a \\ \hline \text{Plurality: } b \end{array}$$

Minimal computation cost, minimal information required.

M2

Move to the top the best candidate in b_i^0 *(truthful) which is above* $w = F(\mathbf{b}^k)$ *in* b_i^k *(reported), among those that can become the new winner of the election*

$$\begin{array}{c} \mathbf{a} \succ b \succ c \succ d \\ b \succ c \succ \mathbf{a} \succ d \\ d \succ \mathbf{a} \succ b \succ c \\ c \succ d \succ b \succ \mathbf{a} \\ \hline \text{Plurality: } a \end{array} \rightarrow \begin{array}{c} a \succ b \succ \mathbf{c} \succ d \\ c \succ b \succ a \succ d \\ d \succ a \succ b \succ \mathbf{c} \\ \mathbf{c} \succ d \succ b \succ a \\ \hline \text{Plurality: } c \end{array} \rightarrow \begin{array}{c} a \succ b \succ c \succ d \\ c \succ b \succ a \succ d \\ a \succ d \succ b \succ c \\ c \succ d \succ b \succ a \\ \hline \text{Plurality: } a \end{array}$$

Low computation cost, low information required (score, majority graph).

Theorem

F^{M1} *converges for every voting rule.*

Proof idea: *M1* can be applied only once by each individual.

Theorem

F^{M2} *converges for PSR, Copeland and Maximin.*

Proof idea: the score of the winner increases at every step, or remains the same and the candidate moves up in the tie-breaking order.

The result generalises to every rule electing the candidate that maximises a notion of score.

Axiomatic properties are preserved at every step of the iteration:

Theorem

M1 and M2 preserve unanimity.

If we start from a unanimous profile, the winner is always the top preferred candidate at every step of the iteration.

Theorem

M1 and M2 preserve Condorcet consistency.

Same for anonymity and neutrality. Pareto-condition does not transfer.

Condorcet Efficiency: Impartial Culture Assumption

Università degli Studi di Padova

For Plurality better 2P and 3P, for all others *M2* is better.
Positive performance of *M1*, even if little changes.
50 voters, 10.000 profiles, 5 alternatives

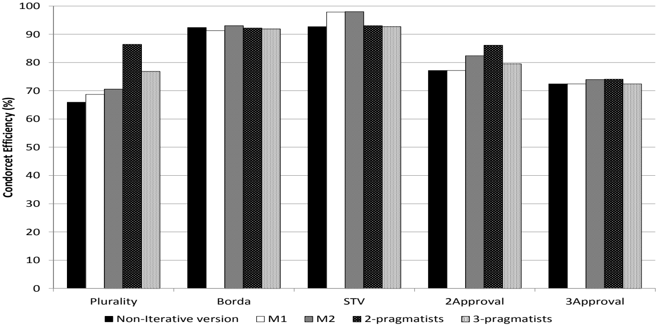

One further motivation for iterated manipulation is that the Condorcet winner may be extracted without having to ask for the full profile.

But: is it more costly to iterate or to ask for the full profile?

	# 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

Profiles are 50×5, maximal number of iterations is 27: good for Plurality!
Iteration takes place between 10% and 30% of the cases:
Not very costly, given the increase in Condorcet efficiency!

Urn model allows for correlation in the preferences of individuals
Iteration decreases with the increase of correlation (25 candidates, 10 voters):

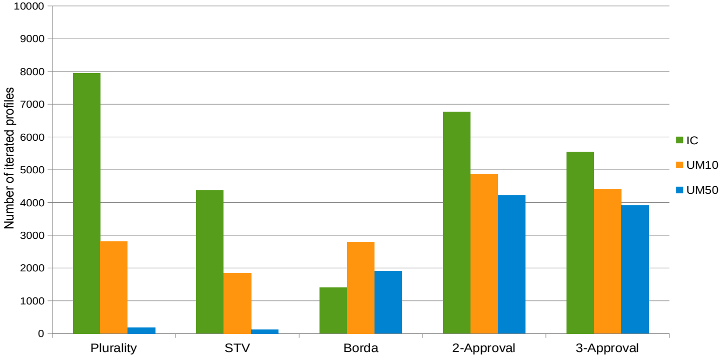

IC: no correlation. UM10: 10% copies in the urn. UM50: 50% copies in the urn.

Condorcet Efficiency: Urn Model

Good results for low number of voters and high number of candidates
Modelling a classic Doodle poll (25 time slots, 10 voters)

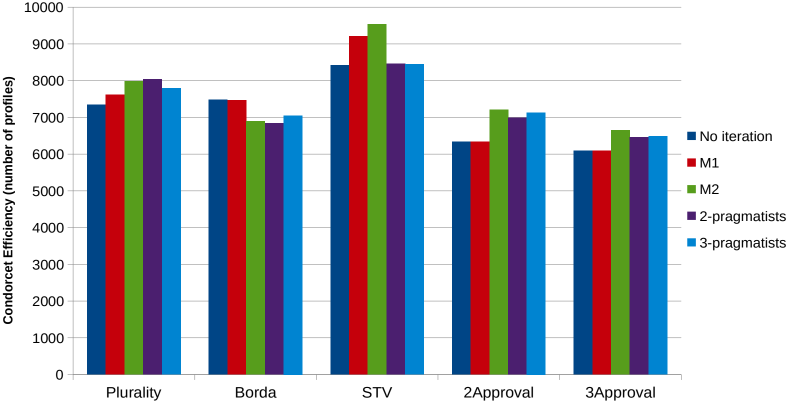

Average Position of the Winner (aka Borda score)

How much preferred is the winner in average?

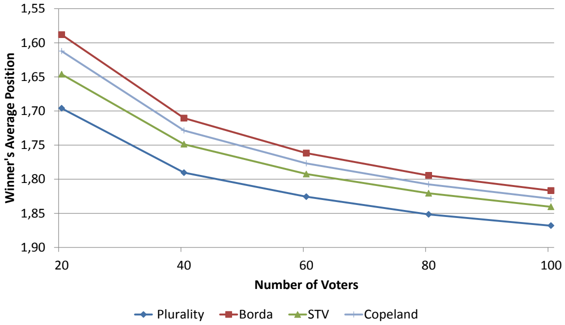

Recall that Borda elects the candidate with the highest "average position"

For all voting rules (except for Borda) the position of the winner shows an increase by allowing iterated restricted manipulation:

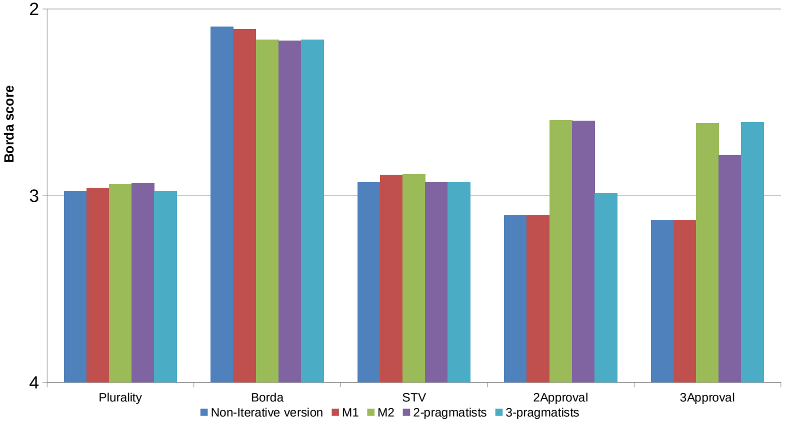

Data constructed from www.preflib.org, Netflix dataset
Profiles are too often unanimous, iteration takes place in $< 0.1\%$ of profiles! 10 candidates, 100 voters: too little iteration!

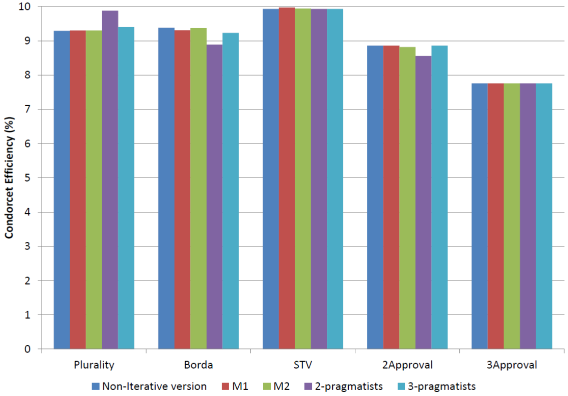

Same experiment with Skate dataset: too much correlation, too little iteration

Conclusions and Future Work

Università degli Studi di Padova

We introduced two new restricted manipulation moves which are easy to compute and need small amount of information, and we evaluated:

- Convergence of restricted iterative voting
- Condorcet efficiency
- Average position of the winner (Borda score)
- Number of iteration steps

Restricted manipulation in iterative voting increases the Condorcet efficiency and the average position of the winner in a limited number of steps.

Lots of future questions:

- More realistic distribution of preferences: real-world data from elections
- Other ideas for restricted manipulation move?
- Other parameters to evaluate performance of iteration?

Andrea Loreggia
Department of Mathematics
University of Padova
loreggia@math.unipd.it
`http://www.math.unipd.it/~loreggia/`