Computational Social Choice: Spring 2017

Ulle Endriss
Institute for Logic, Language and Computation
University of Amsterdam

Plan for Today

So far we have modelled voting as a *one-shot* event: voters declare their preferences and the voting rule computes a definitive outcome.

But reality often is more complex. Examples:

- committee members may hold several straw polls before deciding
- people using online tools such as Doodle can update their ballots

Today we are going to introduce the model of *iterative voting* to (approximately) capture such phenomena:

- ▶ All voters vote. They then inspect the outcome and one voter may decide to update her ballot (i.e., to "manipulate"). Repeat.

Main question today: will this always *converge*?

Formal Framework

A group of *voters* $N = \{1, \ldots, n\}$ choose from a set of *alternatives* X. Let $(\succ_1, \ldots, \succ_n) \in \mathcal{L}(X)^n$ denote the profile of *true preferences*.

Voting proceeds in rounds. $\boldsymbol{R}^t = (R_1^t, \ldots, R_n^t) \in \mathcal{L}(X)^n$ is the profile of *declared preferences* in round $t \geqslant 0$. Assume $R_i^0 = \succ_i$ for all $i \in N$.

We *break ties* using a fixed lexicographic order $\rhd \in \mathcal{L}(X)$. Thus, every voting rule F induces a resolute rule $F_\rhd : \boldsymbol{R} \mapsto \max_\rhd(F(\boldsymbol{R}))$.

After round t, voter $i \in N$ has a *better response* $R_i^\star \in \mathcal{L}(X)$ if:

$$F_\rhd(R_i^\star, \boldsymbol{R}_{-i}^t) \succ_i F_\rhd(R_i^t, \boldsymbol{R}_{-i}^t)$$

<u>Remark:</u> profile with no better responses = pure Nash equilibrium.

After each round, one voter with better responses implements one of them. The process stops when there are no more better responses.

We speak of *convergence* for the *voting rule* F and *update policy* P (= a class of better responses), if the process always stops eventually.

Example

Under the *plurality rule*, if voters can update their ballots using *arbitrary better responses*, we do not get convergence:

Alternatives $X = \{a, b, c, d\}$. Nine voters. Two each vote for a, b, c. The other three have these preferences:

$$d \succ_1 a \succ_1 b \succ_1 c \qquad d \succ_2 a \succ_2 b \succ_2 c \qquad c \succ_3 b \succ_3 a \succ_3 d$$

Then we may encounter a cycle (at each step we show the current scores of the four alternatives and the votes of our three voters):

$$\begin{matrix}[2,2,3,2] \\ (d,d,c)\end{matrix} \xrightarrow{1} \begin{matrix}[2,3,3,1] \\ (b,d,c)\end{matrix} \xrightarrow{2} \begin{matrix}[3,3,3,0] \\ (b,a,c)\end{matrix} \xrightarrow{3} \begin{matrix}[3,4,2,0] \\ (b,a,b)\end{matrix} \xrightarrow{1}$$

$$\begin{matrix}[3,3,3,0] \\ (c,a,b)\end{matrix} \xrightarrow{3} \begin{matrix}[3,2,4,0] \\ (c,a,c)\end{matrix} \xrightarrow{1} \begin{matrix}[3,3,3,0] \\ (b,a,c)\end{matrix} \xrightarrow{3} \cdots$$

[tie-breaking order: $a \rhd b \rhd c \rhd d$]

Best Responses

A *best response* is a better response $R_i^\star$ that cannot be topped:

$$F_\triangleright(R_i^\star, \boldsymbol{R}_{-i}^t) \quad \succ_i \quad F_\triangleright(R_i^t, \boldsymbol{R}_{-i}^t) \quad \text{and}$$
$$F_\triangleright(R_i', \boldsymbol{R}_{-i}^t) \quad \not\succ_i \quad F_\triangleright(R_i^\star, \boldsymbol{R}_{-i}^t) \quad \text{for all } R_i' \in \mathcal{L}(X)$$

There may be several best responses. It is often reasonable to assume that a voter will select a specific one. For instance:

- one that *minimises the swap distance* to the voter's previous ballot
- one that *maximises the margin of victory* for the new winner
- one that ranks the new winner at the top (*direct best response*)

Remark: For the plurality rule, the last two concepts coincide.

Convergence of Best Responses under Plurality

Theorem 1 (Meir et al., 2010) *Iterative voting restricted to arbitrary best responses converges for the plurality rule.*

Note that we assume lexicographic tie-breaking and starting from the true profile. Not all of the literature makes the same assumptions.

To be precise, Meir et al. (2010) showed this for *direct best responses* only, and this also is all we are going to prove here.

Reijngoud (2011) and Brânzei et al. (2013) later strengthened the result of Meir et al. (2010) to obtain the one stated above.

R. Meir, M. Polukarov, J.S. Rosenschein, and N.R. Jennings. Convergence to Equilibria in Plurality Voting. Proc. AAAI-2010.

A. Reijngoud. Voter Response to Iterated Poll Information. MoL, 2011.

S. Brânzei, I. Caragiannis, J. Morgenstern, and A.D. Procaccia. How Bad is Selfish Voting? Proc. AAAI-2013.

Proof

Response = a voter transfers 1 point from one alternative to another.
Direct = the receiving alternative is the new winner.

So can distinguish two *types* of direct best responses:

- transfer of 1 point from the old winner to the new winner
- transfer of 1 point from an old non-winner to the new winner

Let $W^t \subseteq X$ be the set of alternatives that, for *some* sequence of direct best responses, win in *some* round $t' \geqslant t$.

On the next slide, we prove the following lemma:

> *If the update after round t involves the transfer of 1 point away from non-winner x, then $x \notin W^{t'}$ for all rounds $t' > t$.*

But this proves the theorem, because there can be no infinite sequence of *consecutive* updates involving point transfers *away from the winner:* every voter can update *at most $m-1$ times* in such a sequence. ✓

Proof of the Lemma

Consider an update after round t involving the transfer of 1 point from non-winner $x_i^t \neq F_\triangleright(x_i^t, \boldsymbol{R}_{-i}^t)$ to new winner $x_i^{t+1} = F_\triangleright(x_i^{t+1}, \boldsymbol{R}_{-i}^t)$.

Claim: $x_i^t \notin W^{t'}$ for all rounds $t' > t$, i.e., x_i^t can never win again.

Proof: Let $s(x, \boldsymbol{R})$ denote the score of x under $\boldsymbol{R}$ and define the *refined score* of x in round t as follows:

$$s_\triangleright^t(x) \quad = \quad s(x, \boldsymbol{R}^t) + \frac{1}{m} \cdot \#\{y \in X \mid x \triangleright y\}$$

Let $\alpha := s_\triangleright^t(F_\triangleright(\boldsymbol{R}^t))$ be the winning score in round t.

Now show that for every round $t' > t$, these *two invariants* hold:

(i) $s_\triangleright^{t'}(x_i^t) < \alpha - 1$ $\quad$ (ii) $s_\triangleright^{t'}(x) \geqslant \alpha$ for at least two $x \in X$

True for $t' = t + 1$: (i) x_i^t didn't win in round t and then lost 1 point. (ii) The old winner has α points and the new winner has more.

Also for all $t' > t$: x_i^t can't get promoted, as it wouldn't win (gaining just 1 point). Old winner has score $\geqslant \alpha$. New winner even higher. ✓

Other Rules?

Lev and Rosenschein (2012) and Reyhani and Wilson (2012) proved a very similar result for *antiplurality* (independently from each other).

Those working on this topic believe that plurality and antiplurality are the only voting rules for which such results are attainable, but to date this conjecture has not been made entirely precise.

Next: Various ways of altering our assumptions to still get results . . .

O. Lev, J.S. Rosenschein. Convergence of Iterative Voting. Proc. AAMAS-2012.

R. Reyhani and M.C. Wilson. Best Reply Dynamics for Scoring Rules. Proc. ECAI-2012.

Convergence for Arbitrary Better Responses

Are there voting rules that converge for arbitrary better responses?
Yes: every dictatorship does. *What about more interesting rules?*

At least for a nonstandard rule, there exists a positive result:

The *direct kingmaker* is the voting rule under which voter 1 selects one of the other voters who then selects the winning alternative.

Note that here the notion of a "true initial profile" is not well-defined anymore, so we drop this assumption.

Proposition 2 (Meir et al., 2016) *Iterative voting with arbitrary better responses converges for the direct kingmaker.*

Proof: Voters other than 1 only switch to more preferred alternatives. Once they stop, voter 1 can at most move through all of them once. ✓

R. Meir, M. Polukarov, J.S. Rosenschein, and N.R. Jennings (pers. comm., 2016).

Convergence for Pragmatic Voters

Computing responses is difficult. A *k-pragmatist* will only use a better response that amounts to moving her favourite alternative amongst the k front-runners to the top and leaving the rest of the ballot as it is.

Exercise: *What else do you call a k-pragmatist with $k = |X|$?*

Proposition 3 (Reijngoud and Endriss, 2012) *Iterative voting amongst k-pragmatists converges for any positional scoring rule.*

Proof: Convergence follows from the insight that for a PSR the set of the k front-runners never changes under this update policy: they never lose points and the other alternatives never gain points. ✓

A. Reijngoud and U. Endriss. Voter Response to Iterated Poll Information. Proc. AAMAS-2012.

Convergence under Incomplete Information

We know that restricting access to information can reduce manipulation. *Maybe this can help here as well?* Yes (but not much known to date).

Recall: *Copeland* rule = max(# pairwise wins − # pairwise losses)

Suppose a voter has *incomplete information* (e.g., just who wins). Suppose she only updates if her move is a *better response* for some profile she considers possible and *not a worse response* for any of them.

Weaker notion of convergence: we have reached a *stable outcome* once voters may still update but the winner won't change anymore.

Theorem 4 (Endriss et al., 2016) *When voters are given only winner information, iterative Copeland converges to a stable outcome.*

Proof: Omitted.

U. Endriss, S. Obraztsova, M. Polukarov, and J.S. Rosenschein. Strategic Voting with Incomplete Information. Proc. IJCAI-2016.

Social Benefits of Iteration

Using computer simulations, several papers exhibit results showing that iterated manipulation can be socially beneficial. Examples:

- *Condorcet efficiency* (frequency of electing Condorcet winners) may increase (for rules that are not Condorcet extensions).
- Borda scores of winners may increase, i.e., *average voter regret* may decrease (for rules other than Borda).

Upon reflection, this is not so surprising: particularly for low-info rules such as plurality, iteration allows you to make your ballot more relevant.

A. Reijngoud and U. Endriss. Voter Response to Iterated Poll Information. Proc. AAMAS-2012.

U. Grandi, A. Loreggia, F. Rossi, K.B. Venable, and T. Walsh. Restricted Manipulation in Iterative Voting: Condorcet Efficiency and Borda Score. Proc. ADT-2013.

R. Meir, O. Lev, and J.S. Rosenschein. A Local-Dominance Theory of Voting Equilibria. Proc. EC-2014.

Other Research Directions

Other questions that have been investigated in the literature:

- Think of iterated $F_{\triangleright}$ with update policy P as a new voting rule. What axiomatic properties of F *transfer* to this new rule?
- Are there other *update policies* that may be more attractive in terms of their *computational simplicity* or *cognitively plausibility*?
- What if the *initial profile* need not be the truthful one?
- What about *random tie-breaking* (or other tie-breaking rules)?
- What about *simultaneous* (but uncoordinated) updates?
- What about *coordinated* updates (i.e., coalitional manipulation)?
- What about the *frequency of termination* in case convergence cannot be guaranteed in general ($\hookrightarrow$ simulation)?

None of these questions have so far been answered exhaustively.

Summary

This has been an introduction to the topic of iterative voting, a natural model for certain types of interactive voting scenarios.

Our main concern has been *convergence*. Elusive, but:

- better responses converge for some artificial rules ("kingmaker")
- best responses converge (only) under plurality and antiplurality
- better results for simpler update policies (e.g., "pragmatism")
- better results also under incomplete information

Framework with many parameters. Not yet fully understood.

Interesting to try an understand the "*meta voting rules*" defined by a voting rule and an update policy: transfer results, outcome quality, . . .

What next? Voting in combinatorial domains. Then logic and SCT.