Voter Response to Iterated Poll Information

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

[joint work with Annemieke Reijngoud]

Outlook

What are the properties of the 'meta voting rules' we get if voters can repeatedly update their voting intentions before a rule is applied?

Two ideas:

- Voters may have only partial information (opinion polls)
- Strategic best responses (à la game theory) are but one option

Parameters to vary:

- Voting rule
- Poll information given to voters (full profile, scores, ...)
- Response policy (strategic, pragmatic, ...)

Parameters we won't vary here (but could in principle):

- Turn-taking policy (theorems: whatever / experiments: fixed)
- Tie-breaking rule (we use a fixed lexicographic order)

Basics

Finite set of *alternatives* $\mathcal{X}$. Finite set of *voters* $\mathcal{N} = \{1, \ldots, n\}$.

$\mathcal{L}(\mathcal{X})$ is the set of all *preferences* or *ballots* (strict linear orders on $\mathcal{X}$).

A *profile* $\boldsymbol{R} = (R_1, \ldots, R_n) \in \mathcal{L}(\mathcal{X})^n$ is a vector of ballots.

A (possibly irresolute) *voting rule* is a function $F : \mathcal{L}(X)^n \to 2^{\mathcal{X}} \setminus \{\emptyset\}$.

For this talk, we use a lexicographic *tie-breaking rule:* $a \succ b \succ c \succ \cdots$

Examples: Plurality / Borda / Copeland / ...

Poll Information Functions

An opinion poll reflects (usually partial) information of a given profile:

- score obtained under Plurality / Borda / Copeland / ...
- winner(s) / ranking of alternatives for a given voting rule
- number of occurrences of each distinct ballot
- (weighted) majority graph

A *poll information function* (PIF) is mapping any given profile into the chosen information structure $\mathcal{I}$ used for modelling polls:

$$\pi : \mathcal{L}(\mathcal{X})^n \to \mathcal{I}$$

Receiving signal $\pi(\boldsymbol{R})$ and knowing her own R_i, the *information set* of voter i is the set of profiles she considers possible:

$$\mathcal{W}_i^{\pi(\boldsymbol{R})} \quad = \quad \{\boldsymbol{R'} \in \mathcal{L}(\mathcal{X})^n \mid R'_i = R_i \text{ and } \pi(\boldsymbol{R'}) = \pi(\boldsymbol{R})\}$$

PIF π is *at least as informative* as PIF π' if always $\mathcal{W}_i^{\pi(\boldsymbol{R})} \subseteq \mathcal{W}_i^{\pi'(\boldsymbol{R})}$.

Informativeness

For example, for *Borda* we get this hierarchy of informativeness:

$$\text{Profile} \rightarrow \#\text{Ballot} \rightarrow \text{WMG} \begin{array}{c} \nearrow \\ \\ \searrow \end{array} \begin{array}{c} \text{Score} \rightarrow \text{Rank} \rightarrow \text{Winner} \\ \\ \text{MajorityGraph} \end{array} \begin{array}{c} \searrow \\ \\ \nearrow \end{array} \text{Zero}$$

<u>Related work:</u> There seem to be connections (yet to be explored) between our PIF's and work on the *compilation* of intermediate election results (Chevaleyre et al., 2009).

Y. Chevaleyre, J. Lang, N. Maudet, and G. Ravailly-Abadie. Compiling the Votes of a Subelectorate. Proc. IJCAI-2009.

Strategic Response to a Single Poll

Information sets (and therefore also PIF's) lead to an interesting model of strategic manipulation:

- Given your information set, can you misrepresent your preferences in such a way that in some possible world you get a better outcome and in no possible world you get a worse outcome?

Conitzer et al. (2011) use the same notion of manipulation.
Difference: they allow arbitrary information sets (not induced by PIF).

V. Conitzer, T. Walsh, and L. Xia. Dominating Manipulations in Voting with Partial Information. Proc. AAAI-2011.

Some Results

Some results (not the main topic today):

- *Gibbard-Satterthwaite* generalises to case where voters only know *π-poll*, provided voting rule is *strongly computable from π-images* (meaning: poll is enough to compute outcome for every possible choice of ballot for myself).
- Example for a positive result: For polls with *winner information*, *Antiplurality* is *immune* to manipulation (for $n \geqslant 2m-2$).
- There are (somewhat contrived) scenarios where *manipulation* is beneficial even when you have *no information* at all.

Repeated Response to Polls

Let F be a *resolute* voting rule (i.e., including a tie-breaking rule).

Suppose initially everyone votes truthfully (not unreasonable, given that zero information often implies strategy-proofness).

Then iterate this protocol:

(1) Apply *PIF* to current profile and broadcast result.

(2) Pick one voter i (*turn taking*).

(3) Voter i decides whether/how to *update* given her *response policy*, her true *preference*, her current *ballot*, and her *information set*.

We make no assumptions regarding the *turn-taking policy*, except that voters who want to update have priority.

In some cases, at some point nobody will update anymore (*termination*).

We continue till termination or for some fixed number of *rounds* t.

Response Policies

Possible response policies:

- A *truth-teller* always reports her true preference.
- Suppose the PIF in use provides at least ranking information. A *k-pragmatist* moves her favourite alternative amongst the k front-runners to the top and leaves the rest of the ballot as it is.
- A *strategist* chooses a best response given her information set: no other response gives a strictly better outcome for some possible profile and is at least as good for all possible profiles.

Related work: Meir et al. (2010) use the strategist's response policy, except that they always select one specific best response (and they assume full information, namely score information for plurality).

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

Iterated Voting Games

An *iterated voting game* $G = \langle F, \pi, \boldsymbol{\delta} \rangle$ is defined by a resolute voting rule F, a PIF π, and response policies $\boldsymbol{\delta} = (\delta_1, \ldots, \delta_n)$.

Every iterated voting game $G = \langle F, \pi, \boldsymbol{\delta} \rangle$ and number of rounds to be played $t \in \mathbb{N}$ *induce* a new (irresolute) voting rule F^t:

$$F^t(\boldsymbol{R}) \;=\; \left\{ x \in \mathcal{X} \;\middle|\; \begin{array}{l} \textit{with } \boldsymbol{R} \textit{ being the (truthful) initial profile, for} \\ \textit{some turn-taking policy, } x \textit{ wins after } t \textit{ rounds} \end{array} \right\}$$

Remark: Could also do this for specific turn-taking policies.

A game G *terminates* in round t if no voter wishes to change her ballot anymore. If G always terminates after $\leqslant t$ rounds, we write $F^\star$ for F^t.

Termination Results

Theorem 1 *Let F be a positional scoring rule (with lexicographic tie-breaking). If polls give ranking information and all voters are k-pragmatists (for any k) or truth-tellers, then the corresponding iterated voting game terminates.*

Proof: The main insight is that the set of the k top-ranked alternatives never changes and that thus each pragmatist updates at most once. ✓

Also works for *Copeland*, *maximin*, and *Bucklin*.

Remark: Meir et al. (2010) and Lev and Rosenschein (2012) prove similar results for *strategists* rather than pragmatists.

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

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

Transfer Results: Unanimity, but not Pareto

▶ Which properties of F *transfer* to F^t (and $F^\star$, when well-defined)?

Theorem 2 *For all response policies and all polls, unanimity transfers.*

Proof: Immediate. ✓

But beware: *Pareto efficiency* does not transfer. Example:

- Voting rule: elect lexi-first amongst Pareto optimal alternatives
- PIF: winner information only
- Response policy: strategic
- Truthful profile $\boldsymbol{R}$: `voter 1` $= b \succ c \succ a$, `voter 2` $= c \succ a \succ b$
- Note that $F^0(\boldsymbol{R}) = b$. But $F^t(\boldsymbol{R}) = a$ as soon as `voter 2` has updated, even though a is Pareto-dominated by c.

Transfer Results: Condorcet Consistency

Theorem 3 *For all truth-tellers and k-pragmatists (for any k) and polls providing rank information, Condorcet consistency transfers.*

Proof: By induction. Base case: CW wins first round (as F is CC). Suppose CW won previous round. Then every k-pragmatist will consider CW. For any competitor, more than half of the voters will prefer CW, so CW wins again. ✓

Experimental Results: Condorcet Efficiency

The *Condorcet efficiency* of a voting rule is its probability to elect the Condorcet winner when it exists (wrt. some distribution over profiles).

▶ How does the Condorcet efficiency of $F^t/F^\star$ depend on that of F?

Parameters of experiments on following slides:

- 50 voters and 5 alternatives
- profiles generated using impartial culture assumption
- turn-taking: fixed order in which voters are offered to update
- games run until termination (with one exception)
- results shown are averaged over 10,000 trials

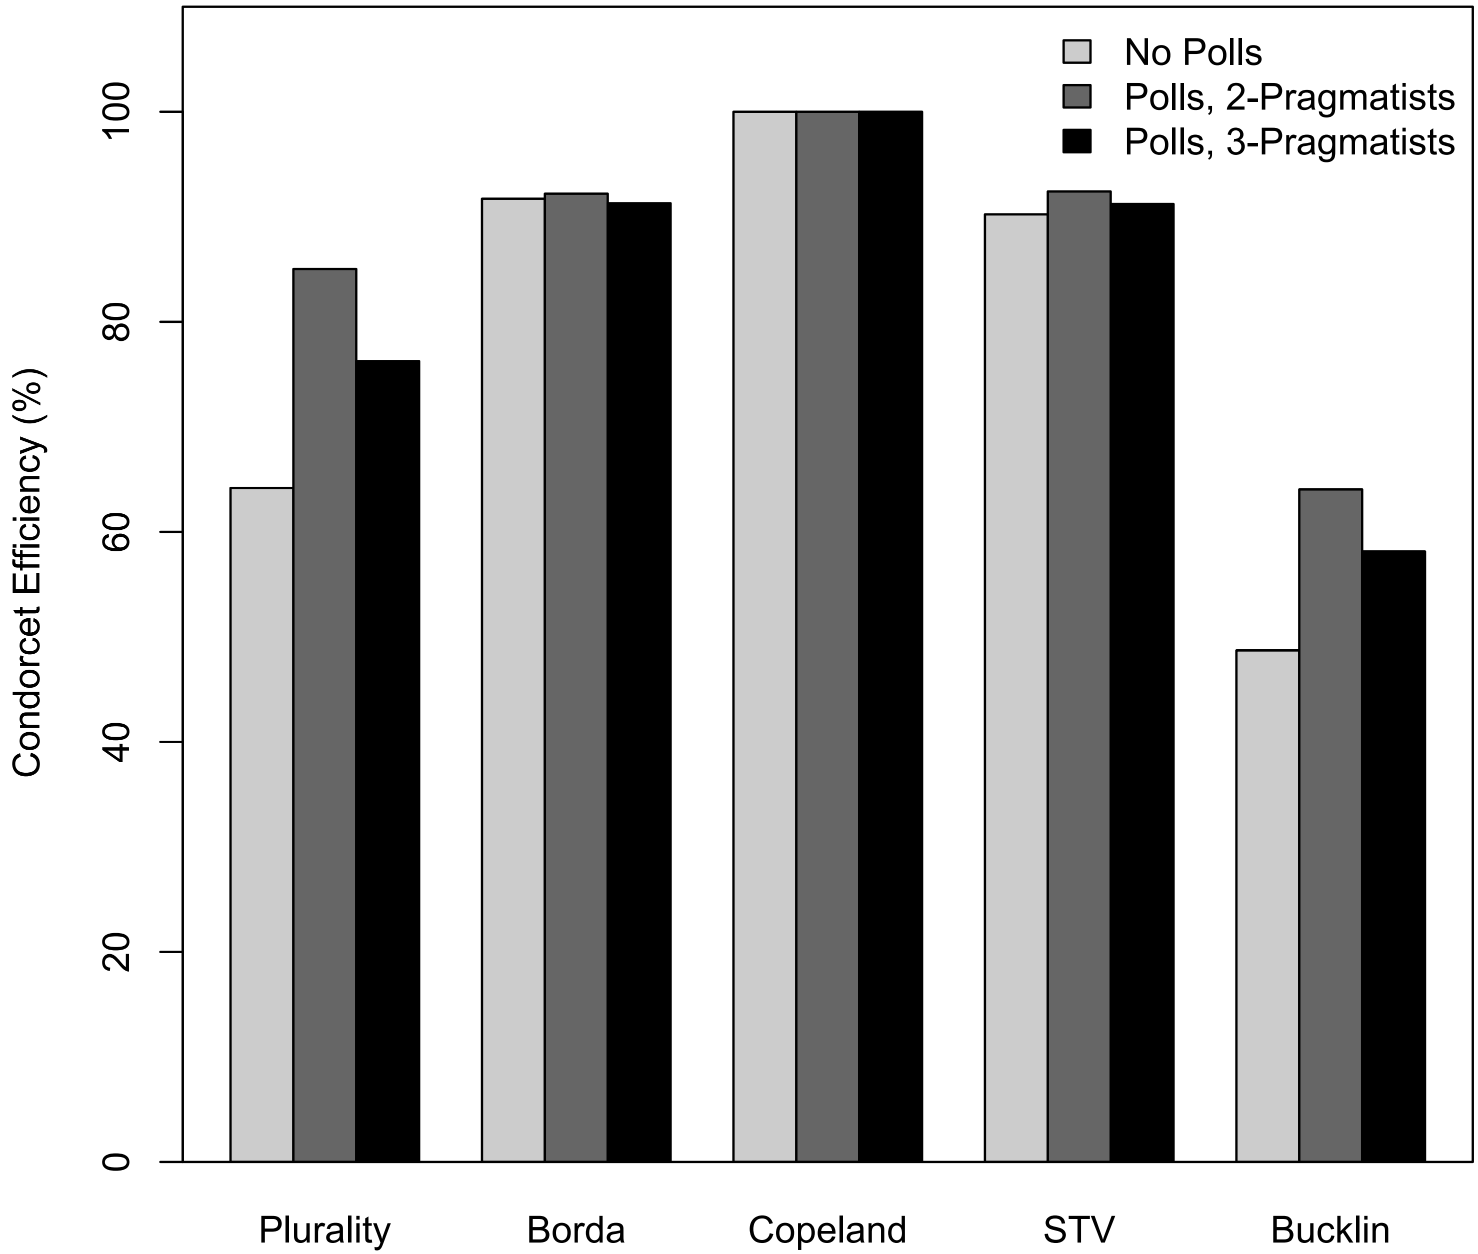
No Polls
Polls, 2-Pragmatists
Polls, 3-Pragmatists
Condorcet Efficiency (%)
0
20
40
60
80
100
Plurality
Borda
Copeland
STV
Bucklin

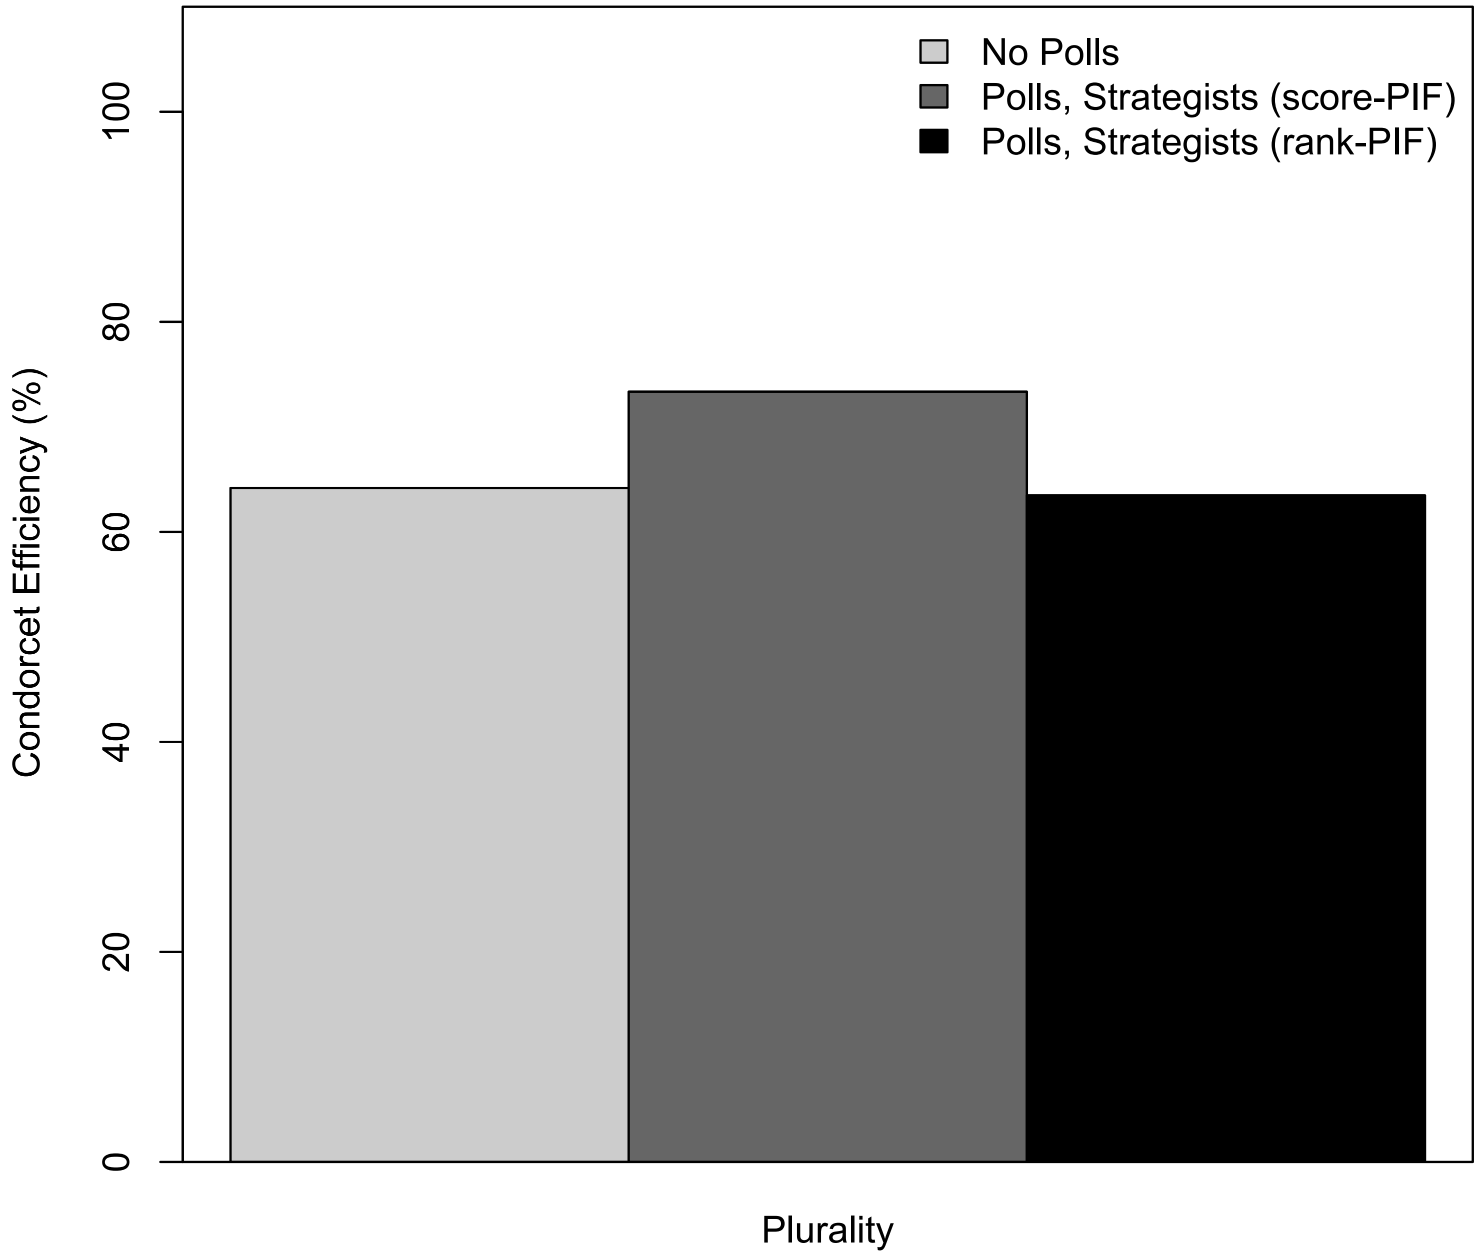
No Polls
Polls, Strategists (score-PIF)
Polls, Strategists (rank-PIF)
Condorcet Efficiency (%)
0
20
40
60
80
100
Plurality

Summary

Presented a simple model for *iterated voting* under *partial information:*

- What we call 'polls' encode natural partial view of a profile (e.g., who wins, scores, majority graph, ...).
- One voter at a time responds in view of her information set and her response policy (e.g., k-pragmatism).
- Understand resulting 'meta voting rule': termination, transfer of properties, reinforcement of properties (e.g., Condorcet efficiency).

All results are from our AAMAS paper. Annemieke's thesis has further results, e.g., on *Condorcet loser efficiency* and on *approval voting*.

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

A. Reijngoud. Voter Response to Iterated Poll Information. Master's thesis. ILLC, University of Amsterdam, 2011.

Points for Discussion

- Most work (unlike ours) assumes that voters have full information. I think *partial information* models deserve much more attention (not just for iterated voting). Our PIF's are just one way.

- What other *response policies* are reasonable?

  Problems with game-theoretical approach ('stategists'):

  - the usual bounded rationality arguments
  - inconsistent to insist on optimising locally but to completely ignore effects across rounds ($\rightsquigarrow$ Stéphane's talk)

- What is the impact of the *impartial culture assumption* on results such as ours on Condorcet efficiency?