Civic Participation III: Extreme Democracy

Ariel Procaccia

March 13, 2018

These notes are based off a presentation by Ariel Procaccia (CMU, Department of Computer Science) for the section on "Civic Participation" in the Mechanism Design for Social Good Reading Group. The notes are taken by members of the reading group with some figures and texts taken from the accompanying paper. Questions and comments from reading group members during the presentation are labeled as such. Please contact the reading group organizers with any questions or comments.

1 Introduction

"Marshall McLuhan once said that politics is solving today's problems with yesterday's tools."
– Pia Mancini.

The goal of this talk is to design some new tools using computer science.

The talk will have three parts:

- *Liquid Democracy:* Voting systems where voters can delegate votes.
- *Participatory Budgeting:* Systems for communities to collectively decide which projects to fund.
- *Virtual Democracy:* Automating moral decisionmaking using models for moral preferences and simulated voting.

Liquid Democracy

Some forms of government you might care about:

- *Monarchy.* A single person decides for everyone.
  Problem: Not particularly democratic.
- *Direct democracy.* One person, one vote on every issue.
  Advantage: Way more democratic.
  Problem: Doesn't scale particularly well.

- *Representative democracy.* People directly elect representatives, who then vote to make all decisions.
  Advantage: Fairly democratic, scales better.
  Problem: Representatives might not vote how their constituents want.

Liquid Democracy.

- Each person chooses to vote for themselves or delegate their vote to another individual.
- Delegations are transitive.
- The result is a DAG representing vote delegation. Sinks vote for everyone who delegated to them.

This can be thought of as lying between representative democracy and direct democracy.

Existing liquid democracy systems:

- LiquidFeedback: Used by the German Pirate Party since 2010.
- DemocracyOS: Used by the Argentinian Net Party since 2012. The Net Party sought to win seats in Argentinian parliament, and then the elected Net Party representatives would vote according to the outcomes of a system of liquid democracy held among party members. The Net Party did not ultimately win any seats.
- Flux: Used by the Australian Flux Party since 2016. Similar to DemocracyOS. The Flux party has failed to win any seats.

Modeling Liquid Democracy

Setup:

- Undirected graph $G = (V, E)$ (e.g. a social network on the set of voters)
- Two alternatives: one correct, one incorrect.
- Voters have competence level p_i: the probability that the voter can recognize the correct alternative.
- Objective: Maximize the probability of the aggregate vote selecting the correct alternative.

Delegation and voting:

- Voters can only delegate to neighbors in G.
- Assume for all voters i, j, voter i will only consider delegating to j if $p_j > p_i + \alpha$ for some $\alpha > 0$.
- The outcome is determined by the majority vote after delegation.

A bad example:

- n voters with competence .6

- 1 central voter with competence .8

The central voter gets all the votes, and we lose out on accuracy which would come from concentration.

Are all delegation mechanisms this bad?

Definition 1. *A delegation scheme is* local *if each voter's delegation decision is only a function of their local neighborhood in G.*

What are some properties that a reasonable delegation scheme should satisfy?

- *Do no harm:* For any $\epsilon > 0$, sufficiently large instances guarantee a loss of at most ϵ.
- *Positive gain:* There is some $\gamma > 0$ and an infinite family of instances where the scheme does at least γ better than direct democracy.

Theorem 1 (Local delegation is doomed to fail.)**.** *For any α such that delegation requires $p_j > p_i + \alpha$, there is no local scheme satisfying both do no harm and positive gain.*

Main problem/proof intuition: Local mechanisms can't prevent the case where a few vertices amass many votes.

A modest positive result: Capping the number of delegated votes at $O(\sqrt{\log n})$ (which is not a local mechanism) satisfies both reasonableness properties above.

Bigger takeaway: Non-local mechanisms for liquid democracy are a promising area for future work.

Participatory Budgeting

Basic idea: Cities hold public votes to determine the allocation of city funds.

- Porto Alegre: Entire budget (other than regular expenses like debt service) was allocated via participatory budgeting starting in 1989.
- Paris: 100 Million Euro allocated in one recent election (not entire budget).
- Madrid: 24 Million Euro allocated (not entire budget).
- New York: $\approx$ 31 Million Dollars allocated (not entire budget).

Theoretical Setup:

- n voters
- m alternatives
- Alternative a has cost c_a.
- There is a total budget B.
- Each voter has a utility $u_i(a)$ for each alternative.

- Objective: Find a budget-feasible set of alternatives maximizing total welfare.

Problem: Having each agent report their utility for each alternative is typically too complicated.

Alternative bidding languages ("input formats"):

- Ranking by value.
- Ranking by value for money ($u_i(a)/c_a$).
- Knapsack voting. Voters submit their optimal set subject to budget constraints. Used in practice.
- Threshold approval. (New to this work.) List an alternative if the voter's utility for it is above a certain threshold.

Assumption: Voters report truthfully.

How do we compare different rules?

Definition 2. *The* distortion *of a voting rule f on a vote profile $\boldsymbol{\sigma}$ is the worst-case ratio of optimal social welfare to f's social welfare over all utility profiles which are consistent with votes $\boldsymbol{\sigma}$. See [1].*

Intuition: How expressive is the voting language?

So how well do these alternative languages do, in terms of worst-case distortion?

- Ranking by value: $\Theta(\sqrt{m})$.
- Ranking by value for money: $\Theta(\sqrt{m})$.
- Knapsack voting: $\Theta(m)$.
- Threshold approval: $\Theta(\log^2 m)$.

Note that $O(m)$ can be achieved by picking a feasible subset uniformly at random. If you choose even just a single project uniformly at random, the expected value is exactly a $1/m$-fraction of the sum of values of all projects. As the value of the optimal solution is at most the sum of values of all projects, the ratio is therefore at most m.

How about on empirical distributions?

Method: Given empirical data from a participatory budgeting election,

- Reconstruct utilities by selecting utility functions uniformly at random which are consistent with the votes.
- Map these utilities into votes in each of the alternative languages.
- Compute the outcomes according to the choice rule which minimizes worst-case distortion.

- Measure resultant welfare.

Results:

- Empirical welfare loss due to distortion was markedly better than in the worst case.
- The four languages above had the same ranking as in the worst-case setup.

Conclusion: Threshold approval performs well.

Open question: In practice, how do you get people to actually report properly according to this rule?

Virtual Democracy

Only briefly covered.

Motivation: As more and more of our daily lives are automated, AI will more regularly face moral decisions. How should they decide?

Approach:

- Ask humans to make comparisons between alternatives (i.e. gather data).
- Learn a model for each person's moral preferences.
- Summarize society's model for moral preferences (i.e. across individuals).
- When presented with a new moral dilemma, predict people's votes and aggregate them into a decision.

References

[1] The distortion of cardinal preferences in voting, Procaccia, Ariel D and Rosenschein, Jeffrey S, International Workshop on Cooperative Information Agents, 317–331, 2006, Springer