Social Choice (and Mechanism Design) on Social Networks

Umberto Grandi

IRIT – University of Toulouse

13 May 2019

Outline of the tutorial

1. Social Choice and Social Networks (Umberto)
   - Quick introduction on social choice
   - Effects of social networks on collective choice
   - Social choice on networks
   - Opinion diffusion
2. Mechanism Design on Social Networks (Dengji)
   - Promotions via Social Networks
   - Mechanism Design Overview
   - Mechanism Design on Social Networks
   - Truthful Diffusion Mechanisms
   - The Generalization to Combinatorial Settings

For references and extra material on the first part consult:

Umberto Grandi. *Social Choice on Social Networks*. In U. Endriss (editor), Trends in Computational Social Choice, pp. 169-184, AI Access, 2017.

Voting works well until a paradox is found

Elections in the U.S. and in many other countries are decided using the plurality rule: the candidate who gets the most votes win.

Assume that the preferences of the individuals in Florida are as follows:

49%:	Bush $\succ$ Gore $\succ$ Nader
20%:	Gore $\succ$ Nader $\succ$ Bush
20%:	Gore $\succ$ Bush $\succ$ Nader
11%:	Nader $\succ$ Gore $\succ$ Bush

Bush results as the winner of the election, but:

- Gore wins against any other candidate in pairwise election.
- Nader supporters have an incentive to manipulate.

Computational Social Choice

Different social choice problems studied:

- Choosing a winner given individual preferences over candidates
- Allocate resources to users in an optimal way
- Finding a stable matching of students to schools

Different computational techniques used:

- Algorithm design to implement complex mechanisms
- Complexity theory to understand limitations
- Knowledge representation techniques to compactly model preferences
- Simulations and real-world data available on `Preflib.org`

Algorithmic Game Theory and Algorithmic Decision Theory are related research areas (resp. strategic interaction, and single-person decisions)

F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

Social choice I: aggregating individuals' views

The typical ingredients are a set of voters expressing their preferences or tastes over a set of alternatives:

4 friends 

$\succ$ 

$\succ$ 

4 friends 

$\succ$ 

$\succ$ 

1 friend 

$\succ$ 

$\succ$ 

Borda winner is 

STV winner is 

How to decide which rule to use? Typically checking its axiomatic properties, such as unanimity, resistance to clones, Condorcet consistency...

Social choice II: reconstructing the truth

In other applications there exists a ground truth which a set of individuals want to reconstruct, starting from their noisy estimates.

The classical result is Condorcet's jury theorem:

- two alternatives c and $\bar{c}$, with c the correct one
- each voter has an independent probability p to guess the correct alternative
- if $p > 1/2$ the probability that the majority vote is the correct alternative tends to 1 increasing the size of the electorate

We can also say that the majority rule is the maximum likelyhood estimator for the noise model described above.

Social networks in the pre-vote phase

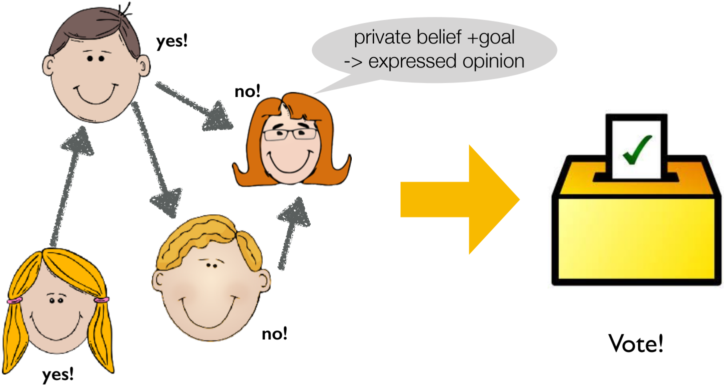

Mutual influence
(deliberation?)

Social networks as parts of the voting mechanism

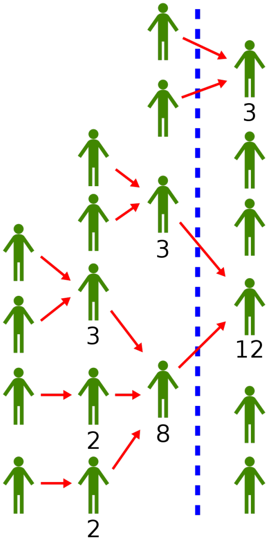

Part I:

Social network effects on collective choice

The majority illusion

Consider a two-candidate election (full node vs. empty node) with voters connected on a social network:

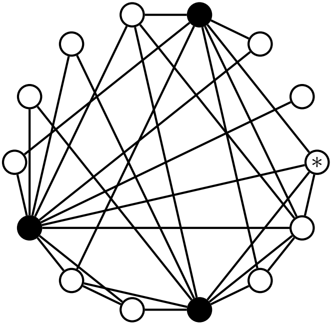

There is a clear majority of 3 vs 11 in favour of empty nodes. If we asked voters, based on their neighbourhood, how do they think the election will go:

- Take voter $*$: she sees one empty node and three full, so will reply that the full node will win
- The same for all 11 empty votes, resulting in a poll reporting a victory of full nodes with 11 vs 3!

Noisy votes

We are in the truth-tracking perspective with two candidates, c and $\bar{c}$:

Theorem [Conitzer, 2012]

If the probability that a voter estimates the correct alternative is independent from the probability of being influenced by her neighbours, then the best mechanism to recover the ground truth ignores the network.

Proof. Assume first that $Prob(P|c) = \prod_{i \in N} f_i(p_i, P_{N(i)}|c)$, where $P_{N(i)}$ is the profile restricted to i's neighbours. Suppose now that $f_i(p_i, P_{N(i)}|c) = g_i(p_i|c) \cdot h_i(p_i, P_{N(i)})$. Functions h_i do not depend on the current alternative c, so to maximise the latter figure one need only look at functions g_i.

Related work

1. The independent conversation model (Conitzer 2013, Tsang et al 2015, Procaccia et al 2015) assumes that a voter is influenced by the majoritarian opinion on a set of discussions with other voters. It has been refined assuming a tendency to be more easily convinced of true opinion.

2. In iterative voting a set of voters respond to the result of the previous election until a convergence result is (or not) found. Tsang and Larson 2016 and Sina et al 2015 studied iterations when the information available to voters is filtered by a social network.

3. A social network can be used to restrict the possible communication or interaction between voters, and then study the effects on coalition formation, voting equilibria, group activity selection...

For precise references consult the "Trends" chapter cited at the beginning.

Part II:
Social choice on networks

Liquid democracy

The principle: voters can directly vote (0/1) on the issue at stake, or delegate their vote to others who can in turn delegate their vote or vote directly:

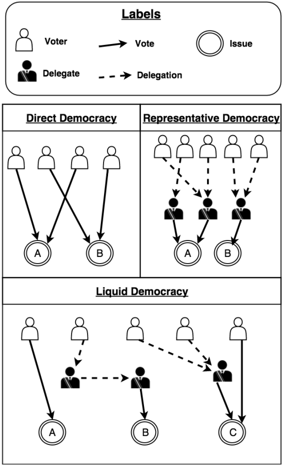

Picture from Medium.com

Liquid democracy

Multiple ways of computing the result given a delegation graph:

- The weight of a voting individual is the number of individuals delegating (directly or by transitive delegation): this is the implementation in the *Liquid Feedback* software, and studied in the setting of multiple interconnected issues, cycles of delegations, and as a truth tracking mechanism (Christoff and Grossi 2017, Bloembergen et al, 2019)
- Spectral ranking techniques such as Page Rank or Katz index can be used to compute the weights of voting individuals (Boldi et al 2009, 2011)
- Voters can have partial orderings over alternatives, and refine them by delegating part of their orderings to other voters (Brill and Talmon, 2017)

Ratings and recommendations

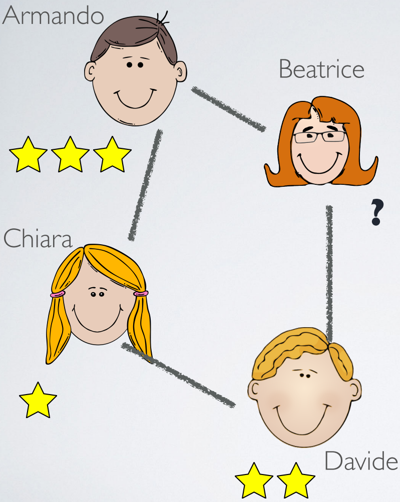

How good is
the restaurant
for Beatrice?

Personalised ratings
are **resistant**
to bribery!

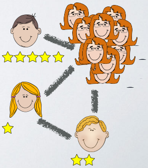

Grandi and Turrini, Personalised Ratings (IJCAI 2016)

Part III:
Opinion diffusion

Diffusion of expressed opinions

Let us focus on one aspect of the initial picture:

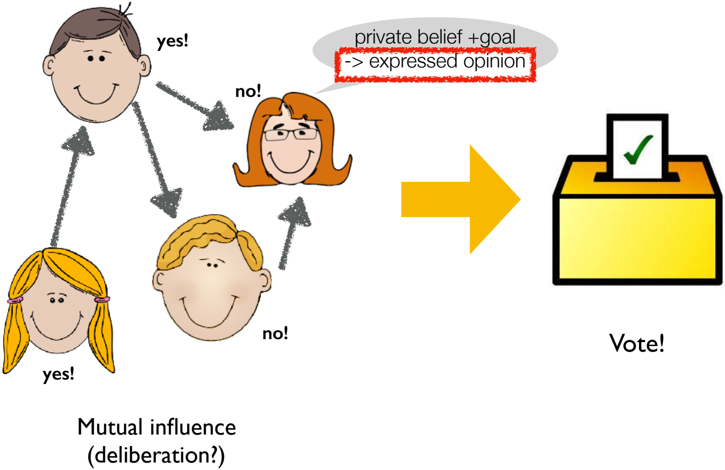

Social influence as aggregation

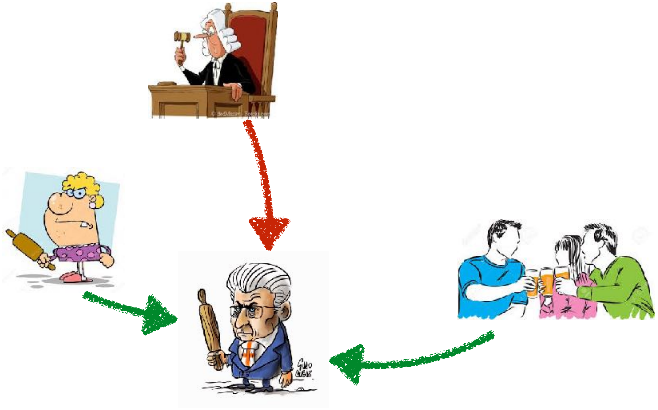

Are Salvini and Di Maio
fit to govern?

Social influence as aggregation

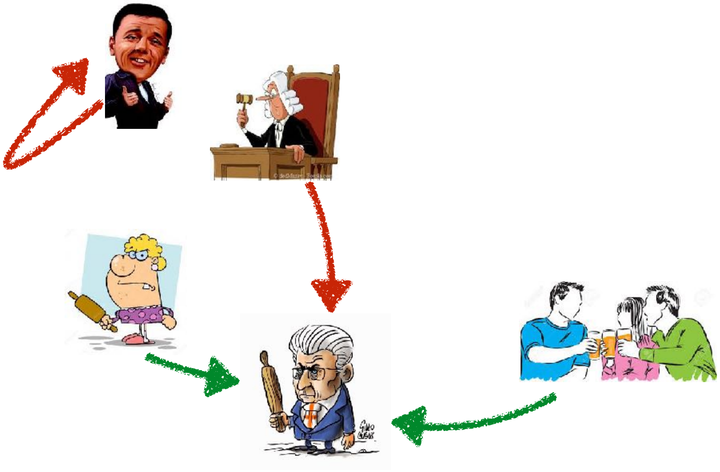

Are Salvini and Di Maio
fit to govern?

Social influence as aggregation

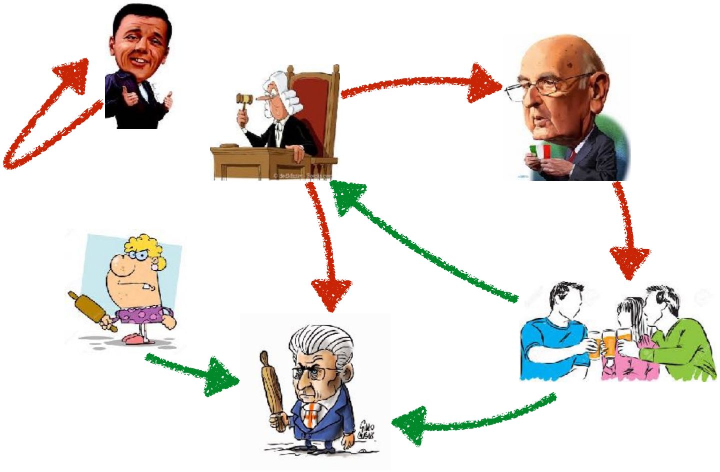

Are Salvini and Di Maio
fit to govern?

The literature: qualitative/quantitative opinions

Most models of opinion diffusion are based on quantitative opinions in [0,1]:

- De Groot (1974) and Lehrer-Wagner (1981): individuals take the weighted average of the opinions of their neighbours
- First recent study involving logical constraints by Friedkin et al. (2016)
- Epidemics models: SIR models, cascades, ising spin...

Much less work exists on discrete opinons:

- Threshold models by Granovetter and Schelling (1978): 0/1 yes/no opinions, updated if the proportion of neighbours with the opposite opinion raises above a certain threshold
- Voter models (Holley and Ligget, 1975, Clifford and Sudbury, 1973): a random individual takes the opinion of random neighbour

Stability, not consensus

Much theoretical work aimed at characterising conditions to reach consensus

We view opinion diffusion as a pre-processing step before voting takes place.
Interesting questions: Will the process terminate? On what "kind" of profile (aligned, polarised, unanimous)? What voting rule should we use then?

Opinion diffusion as aggregation

Opinions can be more complex than single 0/1 views or parameters in [0,1]:

- Multi-issue binary views with constraints (AAMAS-15,-17-19)
- Preferences as linear orders over candidates (IJCAI-16)
- Belief bases as sets of propositional formulas (Schwind et al. 2015, 2016)

Opinion diffusion as aggregation

Opinions can be more complex than single 0/1 views or parameters in [0,1]:

- Multi-issue binary views with constraints (AAMAS-15,-17-19)
- Preferences as linear orders over candidates (IJCAI-16)
- Belief bases as sets of propositional formulas (Schwind et al. 2015, 2016)

How are individuals updating on complex opinions?
A simple idea is to look at the opinion of one's influencers and use:

- Aggregation rules from judgement/binary aggregation (constraints!)
- Voting rules from preference aggregation (transitivity!)
- Belief merging techniques

The architecture of a discrete time iterated diffusion process - Part I

In virtually all settings there are common features:

- A finite set of individuals $\mathcal{N} = \{1, \dots, n\}$
- A directed graph $E \subseteq \mathcal{N} \times \mathcal{N}$ representing the trust network
- Individual opinions (unspecified format for now) that we shall denote as B_i

Some further notation: $\mathit{Inf}(i) = \{j \mid (i, j) \in E\}$ is the set of influencers of individual i on E. Profile of opinions are $\boldsymbol{B} = (B_1, \dots, B_n)$.

The architecture of a discrete time iterated diffusion process - Part I

In virtually all settings there are common features:

- A finite set of individuals $\mathcal{N} = \{1, \ldots, n\}$
- A directed graph $E \subseteq \mathcal{N} \times \mathcal{N}$ representing the trust network
- Individual opinions (unspecified format for now) that we shall denote as B_i

Some further notation: $\mathit{Inf}(i) = \{j \mid (i, j) \in E\}$ is the set of influencers of individual i on E. Profile of opinions are $\boldsymbol{B} = (B_1, \ldots, B_n)$.

An aggregation function for individual opinion updates

Each individual $i \in \mathcal{N}$ is provided with a suitably defined F_i that merge a set of opinions into a single one. The updated opinion of i is $F_i(B_i, \boldsymbol{B}\restriction_{\mathit{Inf}(i)})$.

Examples: F_i is the majority rule, a belief merging operator...

The architecture of a discrete time iterated diffusion process - Part II

Opinion diffusion process

Let $turn : \mathbb{N} \to 2^{\mathcal{N}}$ *indicate at each point in time the set of agents updating. Let* B_i^t *be the opinion of agent* i *at time* $t \in \mathbb{N}$*, and:*

$$B_i^{t+1} = \begin{cases} F_i(B_i, \boldsymbol{B}^t\restriction_{Inf(i)}) & \text{if } i \in turn(t) \\ B_i^t & \text{otherwise.} \end{cases}$$

If $turn(t) = \mathcal{N}$ the process is called synchronous, if $turn$ selects one individual uniformly at random the process is called asynchronous

Disclaimer: when opinions are on multiple issues or preferences we will also specify at each point in time the issue on which the update is performed.

Termination of diffusion on classes of graphs

Two forms of termination of the iterative process can be investigated:

Asymptotic termination

A diffusion model asymptotically terminates *on a class of graphs* $\mathcal{E} \subseteq 2^{N^2}$ *if for each graph* $E \in \mathcal{E}$ *and for each initial profile of opinions* $\boldsymbol{B}^0$ *we have*

$$\lim_{t \to +\infty} \mathbb{P}[\boldsymbol{B}^{t+1} \neq \boldsymbol{B}^t] \quad = \quad 0.$$

In asynchronous models equivalent to being *absorbing* Markov chain.

Universal termination

A diffusion model universally terminates *on a class of graphs* $\mathcal{E}$ *if there does not exist an infinite sequence of effective updates (ie. such that* $\boldsymbol{B}^{t+1} \neq \boldsymbol{B}^t$*).*

Typically hard to guarantee.

Convergence

Call a profile $\boldsymbol{B}^t$ *stable* if $F_i(\boldsymbol{B}^t) = B_i^t$ for all i, and a *termination profile* for $\boldsymbol{B}^0$ any stable profile reachable from $\boldsymbol{B}^0$.

What happens when the process terminates?

- Diffusion converges to unique profile if termination profiles coincide
- Diffusion converges to consensus if termination profiles are unanimous
- Other notions are of course possible...

Multi-issue binary views

An influence network between Ann, Bob and Jesse:

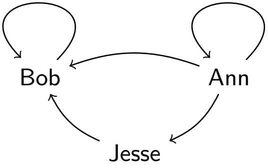

The three agents need to decide whether to approve the building of a swimming pool (first issue) and a tennis court (second issue) in the residence where they live. Here are their initial opinions and their evolution following propositional opinion diffusion with each agent syncronously using the majority rule:

Initial opinions	Profile $\boldsymbol{B}^1$	Profile $\boldsymbol{B}^2$
$B_A^0 = (0, 1)$	$B_A^1 = (0, 1)$	$B_A^2 = (0, 1)$
$B_B^0 = (0, 0)$	$B_B^1 = (0, 0)$	$B_B^2 = (0, 1)$
$B_J^0 = (1, 0)$	$B_J^1 = (0, 1)$	$B_J^2 = (0, 1)$

General termination result

A directed-acyclic graph (DAG) with loops is a directed graph that does not contain cycles involving more than one node.

Theorem [Grandi et al, AAMAS-2015]

If F_i satisfies ballot-monotonicity for all i, then synchronous POD universally terminates on the class of DAG with loops in at most diam$(E) + 1$ steps.

Proof. Start from the sources and propagate opinions.

Observations:

- The proof is a polynomial algorithm to compute the termination profile
- The theorem is not easy to strenghten: take the example of a circle
- The theorem works for any aggregator F_i, even those that do not treat issues independently

UG, E. Lorini and L. Perrussel. Propositional Opinion Diffusion. In *Proceedings of AAMAS-2015*.

Further work on propositional opinion diffusion

Necessary and sufficient conditions for universal termination of synch. POD:

- when F_i are independent, monotonic and responsive, and G is serial
- in terms of winning/losing coalitions of F_i interlocking on G

Z. Christoff and D. Grossi. Stability in Binary Opinion Diffusion. In *Proceedings of LORI-2017.*

Manipulating the result of opinion diffusion at convergence is computationally hard (for undirected networks):

- by bribing some vertices to change their opinion
- by controlling network links
- by controlling the update sequence (turn)

R. Bredereck and E. Elkind. Manipulating Opinion Diffusion in Social Networks. In *Proceedings of IJCAI-2017.*

The influence of a Condorcet cycle

An influence network with 4 agents and 3 alternatives. The preferences 1, 2, and 3 form a *Condorcet cycle*: the majority relation of their preferences is cyclic

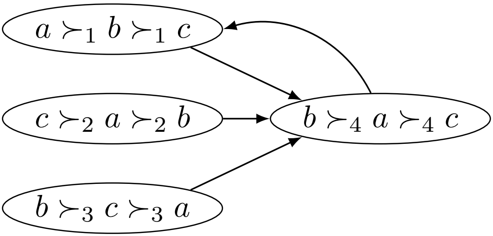

A possible branching of asynchronous pairwise preference diffusion (PPD):

- agent 4 updates on ab, moving to $a \succ_4 b \succ_4 c$
- no further updates possible: ac is no longer adjacent in $\succ_4$

The influence of a Condorcet cycle

An influence network with 4 agents and 3 alternatives. The preferences 1, 2, and 3 form a *Condorcet cycle*: the majority relation of their preferences is cyclic

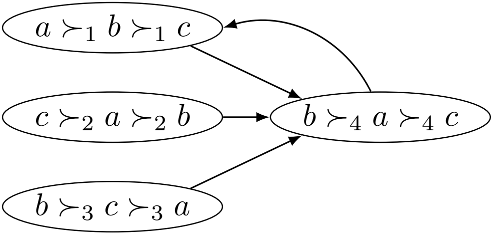

A possible branching of asynchronous pairwise preference diffusion (PPD):

- agent 4 updates on ab, moving to $a \succ_4 b \succ_4 c$
- no further updates possible: ac is no longer adjacent in $\succ_4$

A possible branching of synchronous PPD:

- agents 1 and 4 update repeatedly on pair ab
- an infinite update sequence starts

One interesting result and an open problem

Formalising the argument that mutual influence leads to aligned profiles:

Convergence to aligned profiles

If the sources of a DAG are aligned (single-peaked, single-crossing, Sen's restriction) then under mild conditions termination profiles are also aligned.

One interesting result and an open problem

Formalising the argument that mutual influence leads to aligned profiles:

Convergence to aligned profiles

If the sources of a DAG are aligned (single-peaked, single-crossing, Sen's restriction) then under mild conditions termination profiles are also aligned.

An open problem in opinion diffusion with constraints:

1. We show that asymptotic termination is guaranteed on all graphs (even cyclic ones) though under restrictive conditions (basically that no Condorcet cycle can ever occur). Can we relax this assumption?

M. Brill, E. Elkind, U. Endriss, and UG. Pairwise Diffusion of Preference Rankings in Social Networks. In *Proceedings of IJCAI-2016*.

S. Botan, U. Grandi, L. Perrussel. Multi-issue Opinion Diffusion under Constraints. In *Proceedings of AAMAS-2019*.

Constrained collective choices

Four individuals are deciding to build a skyscraper (S), a hospital (H), or a new road (R). Law says that if S and H are built then R also should be built.

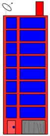

(Hosp and SkyS) implies Road

Voter 1: Y N N

Voter 2: N N Y

Voter 3: Y Y Y

Voter 4: N N N

Constrained collective choices

Four individuals are deciding to build a skyscraper (S), a hospital (H), or a new road (R). Law says that if S and H are built then R also should be built.

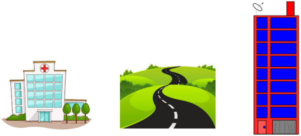

(Hosp and SkyS) implies Road

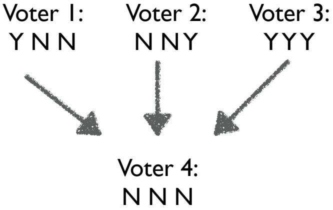

What can happen:

- If voter 4 asks her influencers on 3 issues at the time then voter 4 faces an inconsistent issue-by-issue majority (Y N Y)
- If voter 4 asks on 2 issues the result can be (Y N N) or (N N Y)
- Same result can be reached by asking two 1-issue questions in sequence

Come to our talk and poster of Multi-Issue Opinion Diffusion under Constraints, on Thursday, May 16, 10:30-12:00

Conclusions

In the first part of this tutorial we have seen:

1. An introduction to (computational) social choice: how to aggregate tastes and preferences, how to find a ground truth
2. What are the effects of social networks on social choice mechanisms: the majority illusion, and noisy votes
3. Can we devise mechanism that take into account networks? The case of liquid democracy and personalised recommendations
4. How to model the diffusion of complex opinions? The case of binary issues, preferences, and constrained binary opinions

In the second part of the tutorial...