New Perspectives on Social Choice

Eric Pacuit, University of Maryland

Lecture 12: Liquid Democracy and Deliberation

PHIL 808K

Suppose three voters with competences 80%, 75%, 70% choose between two alternatives. The probability that the simple majority is correct:

$$0.8 * 0.75 * 0.7 + 0.8 * 0.75 * 0.3 + 0.8 * 0.25 * 0.7 + 0.2 * 0.75 * 0.7 = 0.845$$

Suppose three voters with competences 80%, 75%, 70% choose between two alternatives. The probability that the simple majority is correct:

$$0.8 * 0.75 * 0.7 + 0.8 * 0.75 * 0.3 + 0.8 * 0.25 * 0.7 + 0.2 * 0.75 * 0.7 = 0.845$$

If the second voter delegates to the first (and most accurate!) voter, then the outcome under the majority rule depends solely on the ballot of that first voter, so the probability of being right goes down to 0.8.

Suppose three voters with competences 80%, 75%, 70% choose between two alternatives. The probability that the simple majority is correct:

$$0.8 * 0.75 * 0.7 + 0.8 * 0.75 * 0.3 + 0.8 * 0.25 * 0.7 + 0.2 * 0.75 * 0.7 = 0.845$$

If the second voter delegates to the first (and most accurate!) voter, then the outcome under the majority rule depends solely on the ballot of that first voter, so the probability of being right goes down to 0.8.

When the accuracies are 80%, 70%, 60%, we get the opposite effect:

$$0.8 * 0.7 * 0.6 + 0.8 * 0.7 * 0.4 + 0.8 * 0.3 * 0.6 + 0.2 * 0.7 * 0.6 = 0.788 < 0.8$$

Kahng, Mackenzie and Procaccia. *Liquid Democracy: An Algorithmic Perspective*. AAAI, 2018.

$G = (V, E, \vec{p})$ where $V = \{1, \ldots, n\}$ is a set of n voters, E are the edges in a directed social network where $(i, j) \in E$ means that i knows of voter j

Each $i \in V$ has a competence level p_i. The probability that i will vote *correctly*.

$\alpha \in [0, 1)$

i *approves* $j \in V$ if $(i, j) \in E$ and $p_j > p_i + \alpha$

$A_G(i) = \{j \in V \mid i \text{ approves } j\}$

Example

Consider the labeled graph $G_n = (V, E, \vec{p})$ over n vertices.
$E = \{(i, 1) \mid i \in V \setminus \{1\}\}$, i.e., G is a star with 1 at the center. Moreover, $p_1 = \frac{4}{5}$, $p_i = \frac{2}{3}$ for $i \in V \setminus \{1\}$, and $\alpha = \frac{1}{10}$.

Then, as n grows larger, $P_D(G_n)$ goes to 1 by the Condorcet Jury Theorem.

By contrast, all leaves approve the center, and a naive local delegation mechanism M would delegate all their votes. In that case, the decision depends only on the vote of the center, so $P_M(G_n) = \frac{4}{5}$ for all $n \in \mathbb{N}$ and $gain(M, G_n)$ converges to $-\frac{4}{5}$. So, M violates the *DNH* property.

Theorem For any $\alpha_0 \in [0, 1)$, there is no local mechanism that satisfies the *PG* and *DNH* properties.

"The main idea underlying Theorem 1 is that liquid democracy can correlate the votes to the point where the mistakes of a few popular voters tip the scales in the wrong direction."

D. Halpern, J. Halpern, A. Jadbabaie, E. Mossel, A. Procaccia, and M. Revel. *In Defense of Liquid Democracy*. manuscript, 2022.

Delegation Model

$M = (q, \varphi)$ where

- $q : [0,1] \to [0,1]$ is a function that maps a voter's competence to the probability that they delegate
- $\varphi : [0,1] \times [0,1] \to \mathbb{R}$ maps a pair of competencies to a weight.

Each voter i votes directly with probability $1 - q(p_i)$ and, conditioned on delegated with probability $q(p_i)$, delegates to voter $j = i$ with probability proportional to $\varphi(p_i, p_j)$.

Crucially, a voter does not need to "know" the competence of another voter to decide whether to delegate; rather, the delegation probabilities are merely influenced by competence in an abstract way captured by φ.

Upward Delegation Model

1. the probability that any voter i delegates is $q(p_i) = p$ and
2. the weight that any voter i puts on another voter j is $\varphi(p_i, p_j) = \mathbb{I}_{p_j - p_i > 0}$

This model captures that there might be some reluctance to delegate regardless of the voter's competence but assumes that voters act in the interest of society only delegating to voters that are more competent.

Upward Delegation Model

Generate a random graph as follows: add a single voter at a time in order of decreasing competence and allow the voter to either not delegate and create their own disconnected component, or delegate to the creator of any other component with probability proportional to p times the size of the component.

Upward Delegation Model

- For a voter, delegating to any voter in the previously created components is possible (since they have strictly higher competence) and would result in the votes being concentrated in the originator of that component by transitivity.

Upward Delegation Model

- For a voter, delegating to any voter in the previously created components is possible (since they have strictly higher competence) and would result in the votes being concentrated in the originator of that component by transitivity.
- This process is well-known in the study of random graphs.

Upward Delegation Model

- For a voter, delegating to any voter in the previously created components is possible (since they have strictly higher competence) and would result in the votes being concentrated in the originator of that component by transitivity.
- This process is well-known in the study of random graphs.
- It can be show that with high probability no component grows too large as long as $p < 1$.

Upward Delegation Model

- For a voter, delegating to any voter in the previously created components is possible (since they have strictly higher competence) and would result in the votes being concentrated in the originator of that component by transitivity.
- This process is well-known in the study of random graphs.
- It can be show that with high probability no component grows too large as long as $p < 1$.
- Also, there needs to be a constant improvement by continuity of the competence distribution, which ensures that a positive fraction of voters below a certain competence delegate to a positive fraction of voters with strictly higher competencies.

Confidence-Based Delegation

1. the probability $q(p_i)$ that any voter i delegates is decreasing in p_i and
2. the weight that any voter i puts on another voter j is $\varphi(p_i, p_j) = 1$

This model captures that competence does not affect the probability of receiving delegations, only the probability of delegating.

General Continuous Delegation

1. the probability that any voter i delegates is $q(p_i) = p$ and
2. the weight that voter i places on voter j is $\varphi(p_i, p_j)$, where φ is continuous and increases in its second coordinate.

The delegation distribution is slightly skewed towards more competent voters.

For a set $[n]$ of n voters, let $\vec{p} = (p_1, \ldots, p_n)$ be the vector of competences for each voter.

A delegation graph $G_n = ([n], E)$ on n voters is a directed graph with voters as vertices and a directed edge $(i, j) \in E$ denoting that i delegates their vote to j.

The outdegree of a vertex in the delegation graph is at most 1 since each voter can delegate to at most one person.

Voters that do not delegate have no outgoing edges.

In a delegation graph G_n, the delegations received by a voter i, $\text{dels}_i(G_n)$, is the total number of people that (transitively) delegated to i in G_n (i.e., the total number of ancestors of i in G_n).

The weight of a voter i, $\text{weight}_i(G_n)$ is $\text{dels}_i(G_n)$ if i delegates, and 0 otherwise.

Let $\text{max-weight}(G_n) = \max_{i \in [n]} \text{weight}_i(G_n)$ and
$\text{total-weight}(G_n) = \sum_{i=1}^{n} \text{weight}_i(G_n)$

The tuple $(\vec{p}, G_n)$ is called a delegation instance on n voters.

Let $V_i = 1$ if voter i would vote correctly if i did vote and $V_i = 0$ otherwise.

Fixed competencies $\vec{p}$ induce a probability measure $\mathbb{P}_{\vec{p}}$ over the n possible binary votes V_i where $V_i \sim \mathrm{Bern}(p_i)$

The tuple $(\vec{p}, G_n)$ is called a delegation instance on n voters.

Let $V_i = 1$ if voter i would vote correctly if i did vote and $V_i = 0$ otherwise.

Fixed competencies $\vec{p}$ induce a probability measure $\mathbb{P}_{\vec{p}}$ over the n possible binary votes V_i where $V_i \sim \mathrm{Bern}(p_i)$

Let X_n^D be the number of correct votes under direct democracy, that is, $X_n^D = \sum_{i=1}^n V_i$.

Let $X_{G_n}^F$ be the number of correct votes under liquid democracy with delegation graph G_n, $X_{G_n}^F = \sum_{i=1}^n \mathrm{weight}_i(G_n) V_i$.

The tuple $(\vec{p}, G_n)$ is called a delegation instance on n voters.

Let $V_i = 1$ if voter i would vote correctly if i did vote and $V_i = 0$ otherwise.

Fixed competencies $\vec{p}$ induce a probability measure $\mathbb{P}_{\vec{p}}$ over the n possible binary votes V_i where $V_i \sim \mathrm{Bern}(p_i)$

Let X_n^D be the number of correct votes under direct democracy, that is, $X_n^D = \sum_{i=1}^n V_i$.

Let $X_{G_n}^F$ be the number of correct votes under liquid democracy with delegation graph G_n, $X_{G_n}^F = \sum_{i=1}^n \mathrm{weight}_i(G_n) V_i$.

The probability that direct democracy and liquid democracy are correct are $\mathbb{P}_{\vec{p}}[X_n^D > n/2]$ and $\mathbb{P}_{\vec{p}}[X_{G_n}^F > n/2]$, respectively.

Gain

The **gain** of an instance $(\vec{p}, G_n)$ is

$$\mathrm{gain}(\vec{p}, G_n) = \mathbb{P}_{\vec{p}}[X^F_{G_n} > n/2] - \mathbb{P}_{\vec{p}}[X^D_n > n/2]$$

Each voter's competence p_i is sampled i.i.d. from a fixed distribution $\mathcal{D}$ with support contained in $[0, 1]$. Delegations will be chosen according to a model $M = (q, \varphi)$.

A competence distribution $\mathcal{D}$, a model M, and a number n of voters induce a probability measure $\mathbb{P}_{\mathcal{D},M,n}$ over all instances of $(\vec{p}, G_n)$ of size n.

Definition 1 (Probabilistic do no harm). A model M satisfies probabilistic do no harm with respect to a class $\mathfrak{D}$ of distributions if, for all distributions $\mathcal{D} \in \mathfrak{D}$ and all $\epsilon, \delta > 0$, there exists $n_0 \in \mathbb{N}$ such that for all $n \geq n_0$,

$$\mathbb{P}_{\mathcal{D},M,n}[\mathrm{gain}(\vec{p}, G_n) \geq -\epsilon] > 1 - \delta$$

Definition 2 (Probabilistic positive gain). A model M satisfies probabilistic positive gain with respect to a class $\mathfrak{D}$ of distributions if, for all distributions $\mathcal{D} \in \mathfrak{D}$ and all $\epsilon, \delta > 0$, there exists $n_0 \in \mathbb{N}$ such that for all $n \geq n_0$,

$$\mathbb{P}_{\mathcal{D},M,n}[\mathrm{gain}(\vec{p}, G_n) \geq 1 - \epsilon] > 1 - \delta$$

Theorem 1 (Upward Delegation Model). For all $p \in (0,1)$, M_p^U satisfies probabilistic do no harm and probabilistic positive gain with respect to the class $\mathcal{D}^C$ of all continuous distributions.

Theorem 2 (Confidence-Based Delegation Model). All models M_q^C with monotonically decreasing q satisfy probabilistic do no harm and probabilistic positive gain with respect to the class $\mathcal{D}^C$ of all continuous distributions.

Theorem 3 (Continuous General Delegation Model). All models $M_{p,\varphi}^S$ with $p \in (0,1)$ and φ that is non-zero, continuous, and increasing in its second coordinate satisfy probabilistic do no harm and probabilistic positive gain with respect to the class $\mathcal{D}^C$ of all continuous distributions.

From Delegation to Deliberation

$$\mathit{Influence} = \mathit{Delegation}^{-1}$$

C. List (2018). *Democratic Deliberation and Social Choice: A Review*. the Oxford Handbook of Deliberative Democracy.

- The consensus hypothesis: Deliberation tends to generate a consensus

  “rather than aggregating or filtering preferences, the political system should be set up with a view to changing them by public debate and confrontation....[T]here would not be any need for an aggregation mechanism, since a rational discussion would tend to produce unanimous preferences.” (Elster, 1986, p. 112)

- The consensus hypothesis: Deliberation tends to generate a consensus

  "rather than aggregating or filtering preferences, the political system should be set up with a view to changing them by public debate and confrontation....[T]here would not be any need for an aggregation mechanism, since a rational discussion would tend to produce unanimous preferences." (Elster, 1986, p. 112)

- The no-helpful-change hypothesis: Deliberation does not helpfully reduce preference diversity

  "[P]ublic deliberation on a pending item seldom seems to change anyone's mind. . . [D]ue to the network [structure of individual opinions], the effects of deliberative persuasion are typically latent, indirect, delayed, or disguised." (Mackie, 2006, p. 279).

While the consensus hypothesis was too optimistic, the “no-change” hypothesis is too pessimistic. Experience suggests that deliberation sometimes changes people’s minds, and there is some social-scientific evidence that deliberation promotes reflection and learning, and generates more considered preferences...

Arrow's Universal Domain Conditiona

Universal Domain: The domain of the social welfare (choice) function is the set of *all profiles*

Arrow's Universal Domain Conditiona

Universal Domain: The domain of the social welfare (choice) function is the set of *all profiles*

Epistemic Rationale: "If we do not wish to require any prior knowledge of the tastes of individuals before specifying our social welfare function, that function will have to be defined for every logically possible set of individual orderings." (Arrow, 1963, pg. 24)

Universal Domain

1. Voters are free to choose any ranking.
2. The voters' choices are independent.

	x_1			x_2			y		
s_1	a	b	c	c	b	a	b	c	a
s_2	a	c	b	b	c	a	c	b	a
s_3	c	a	b	b	a	c	c	b	a
s_4	b	c	a	a	c	b	b	c	a

Universal Domain

1. Voters are free to choose any ranking.
2. The voters' choices are independent.

	x_1			x_2			y		
s_1	a	b	c	c	b	a	b	c	a
s_2	a	c	b	b	c	a	c	b	a
s_3	c	a	b	b	a	c	c	b	a
s_4	b	c	a	a	c	b	b	c	a

Universal Domain

1. Voters are free to choose any ranking.
2. The voters' choices are independent.

	x_1			x_2			y		
s_1	a	b	c	c	b	a	b	c	a
s_2	a	c	b	b	c	a	c	b	a
s_3	c	a	b	b	a	c	c	b	a
s_4	b	c	a	a	c	b	b	c	a
s_5		$\dots$		c	a	b		$\dots$	

Universal Domain

1. Voters are free to choose any ranking.
2. The voters' choices are independent.

	x_1			x_2			y		
s_1	a	b	c	c	b	a	b	c	a
s_2	a	c	b	b	c	a	c	b	a
s_3	c	a	b	b	a	c	c	b	a
s_4	b	c	a	a	c	b	b	c	a
s_5	c	a	b	c	a	b	c	a	b

Universal Domain

1. Voters are free to choose any ranking.
2. The voters' choices are independent.

	x_1			x_2			y		
s_1	a	b	c	c	b	a	b	c	a
s_2	a	c	b	b	c	a	c	b	a
s_3	c	a	b	b	a	c	c	b	a
s_4	b	c	a	a	c	b	b	c	a
s_5	c	a	b	c	a	b	c	a	b

Universal Domain

1. Voters are free to choose any ranking.
2. The voters' choices are independent.

	x_1			x_2			y		
s_1	a	b	c	c	b	a	b	c	a
s_2	a	c	b	b	c	a	c	b	a
s_3	c	a	b	b	a	c	c	b	a
s_4	b	c	a	a	c	b	b	c	a
s_5	c	a	b	c	a	b	c	a	b
s_6	a	b	c	c	a	b	???		

Universal Domain

1. Voters are free to choose any ranking.
2. The voters' choices are independent.

	x_1			x_2			y		
s_1	a	b	c	c	b	a	b	c	a
s_2	a	c	b	b	c	a	c	b	a
s_3	c	a	b	b	a	c	c	b	a
s_4	b	c	a	a	c	b	b	c	a
s_5	c	a	b	c	a	b	c	a	b
s_6	a	b	c	c	a	b	b c		

Universal Domain

1. Voters are free to choose any ranking.
2. The voters' choices are independent.

	x_1			x_2			y		
s_1	a	b	c	c	b	a	b	c	a
s_2	a	c	b	b	c	a	c	b	a
s_3	c	a	b	b	a	c	c	b	a
s_4	b	c	a	a	c	b	b	c	a
s_5	c	a	b	c	a	b	c	a	b
s_6	a	b	c	c	a	b	$b\,c, c\,a$		

Universal Domain

1. Voters are free to choose any ranking.
2. The voters' choices are independent.

	x_1			x_2			y		
s_1	a	b	c	c	b	a	b	c	a
s_2	a	c	b	b	c	a	c	b	a
s_3	c	a	b	b	a	c	c	b	a
s_4	b	c	a	a	c	b	b	c	a
s_5	c	a	b	c	a	b	c	a	b
s_6	a	b	c	c	a	b	b	c	a

Single-Peaked Profiles

1	1	1
A	B	C
B	C	A
C	A	B

Single-Peaked Profiles

1	1	1
A	B	C
B	C	A
C	A	B

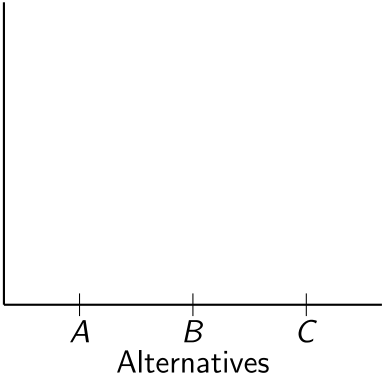

Single-Peaked Profiles

1	1	1
A	B	C
B	C	A
C	A	B

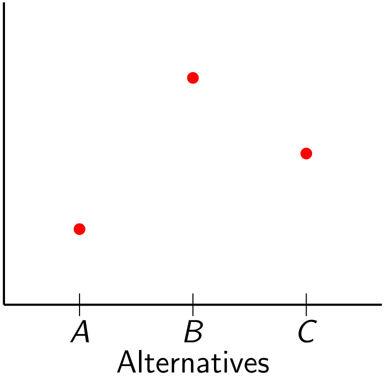

Single-Peaked Profiles

1	1	1
A	B	C
B	C	A
C	A	B

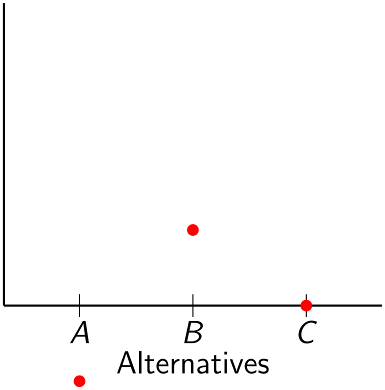

Single-Peaked Profiles

1	1	1
A	B	C
B	C	A
C	A	B

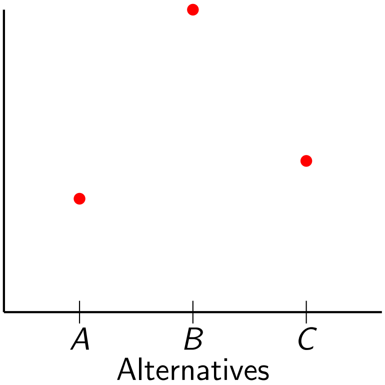

Single-Peaked Profiles

1	1	1
A	B	C
B	C	A
C	A	B

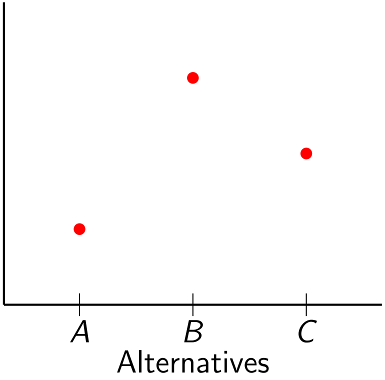

Single-Peaked Profiles

1	1	1
A	B	C
B	C	A
C	A	B

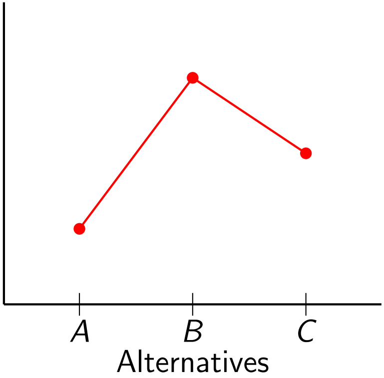

Single-Peaked Profiles

1	1	1
A	B	C
B	C	A
C	A	B

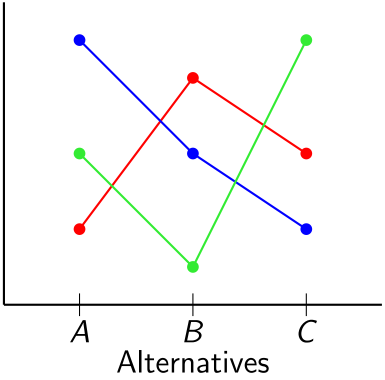

Single-Peaked Profiles

1	1	1
A	B	C
B	C	A
C	A	B

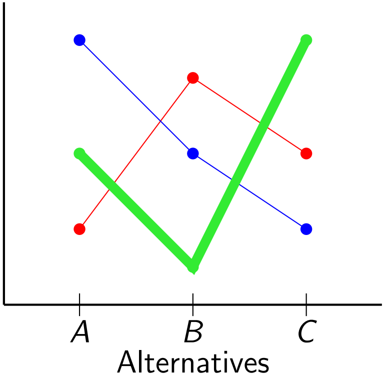

Single-Peaked Profiles

1	1	1
A	B	C
B	C	B
C	A	A

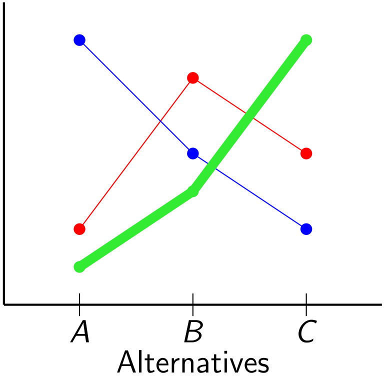

Single-Peaked Profiles

1	1	1
C	C	B
A	B	C
B	A	A

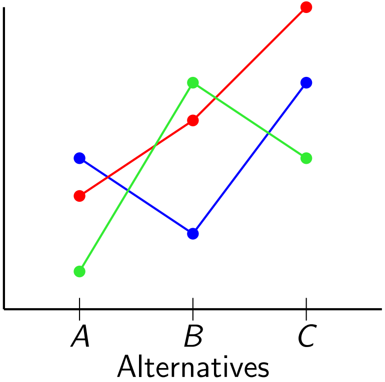

Single-Peaked Profiles

1	1	1
C	C	B
A	B	C
B	A	A

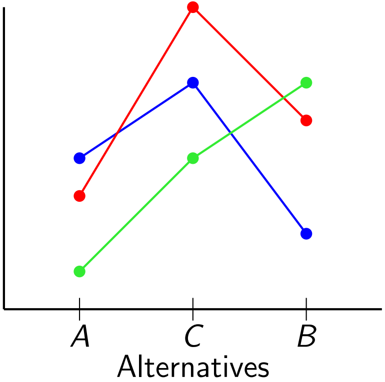

Single-Peaked Profiles

1	1	1
A	B	C
B	C	A
C	A	B

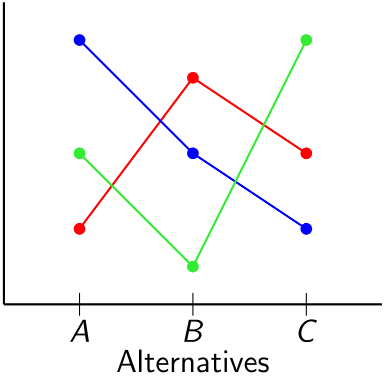

Single-Peaked Profiles

1	1	1
A	B	C
B	C	A
C	A	B

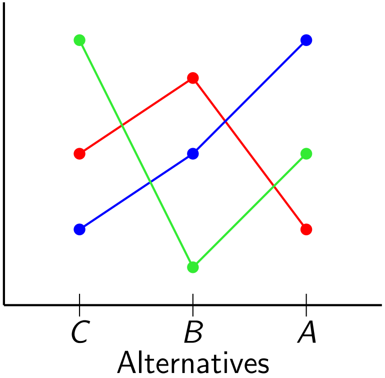

Single-Peaked Profiles

1	1	1
A	B	C
B	C	A
C	A	B

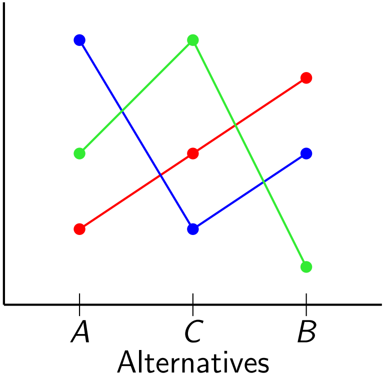

Single-Peaked Profiles

1	1	1
A	B	C
B	C	A
C	A	B

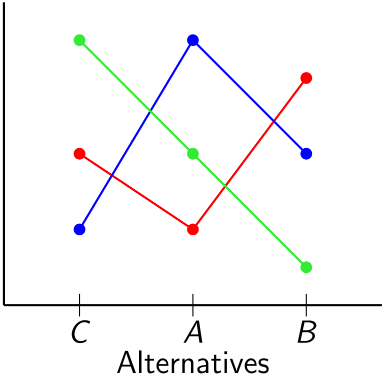

D. Black. *On the rationale of group decision-making*. Journal of Political Economy, 56:1, pgs. 23 - 34, 1948.

Single-Peakedness: the preferences of group members are said to be single-peaked if the alternatives under consideration can be represented as points on a line and each of the utility functions representing preferences over these alternatives has a maximum at some point on the line and slopes away from this maximum on either side.

Single-Peakedness: the preferences of group members are said to be single-peaked if the alternatives under consideration can be represented as points on a line and each of the utility functions representing preferences over these alternatives has a maximum at some point on the line and slopes away from this maximum on either side.

Theorem. If there is an odd number of voters that display single-peaked preferences, then a Condorcet winner exists.

Domain Restrictions

W. Gaertner (2001). *Domain Conditions in Social Choice Theory*. Cambridge University Press, 2001.

E. Kalai, E. Muller, and M. Satterthwaite (1979). *Social welfare functions when preferences are convex and continuous: Impossibility results*. Public Choice, 34, pp. 87 - 97.

E. Elkind (2018). *Restricted Preference Domains in Social Choice: Two Perspectives*. International Symposium on Algorithmic Game Theory, SAGT 2018, pp. 12 - 18.

E. Elkind, M. Lackner and D. Peters. *Preference Restrictions in Computational Social Choice: Recent Progress*. Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-16).

D. Miller. *Deliberative Democracy and Social Choice*. Political Studies, 40, pgs. 54 - 67, 1992.

C. List, R. Luskin, J. Fishkin and I. McLean. *Deliberation, Single-Peakedness, and the Possibility of Meaningful Democracy: Evidence from Deliberative Polls*. Journal of Politics, 75(1), pgs. 80 - 95, 2013.

The *meta-consensus* hypothesis: Deliberation tends to generate a meta-consensus, which is associated with “single-peaked” preferences.

1. Group deliberation leads people to identify a common (semantic) issue dimension in terms of which to conceptualize the decision problem at stake [such as socio-economic left versus right, secular versus religious, or urban versus rural].

1. Group deliberation leads people to identify a common (semantic) issue dimension in terms of which to conceptualize the decision problem at stake [such as socio-economic left versus right, secular versus religious, or urban versus rural].

2. For a given such issue dimension, group deliberation leads people to agree on how the decision options are aligned from left to right with respect to that issue dimension; so people determine which (geometric) structuring dimension best represents the given (semantic) issue dimension.

1. Group deliberation leads people to identify a common (semantic) issue dimension in terms of which to conceptualize the decision problem at stake [such as socio-economic left versus right, secular versus religious, or urban versus rural].

2. For a given such issue dimension, group deliberation leads people to agree on how the decision options are aligned from left to right with respect to that issue dimension; so people determine which (geometric) structuring dimension best represents the given (semantic) issue dimension.

3. Once a (semantic) issue dimension and a corresponding (geometric) structuring dimension have been identified as relevant, group deliberation leads each individual to determine a most preferred position (his or her "peak") on that dimension, with decreasing preference as options get increasingly distant from the most preferred position.

Data from Deliberative Polls support the hypothesis just discussed. Deliberative Polls enable us to compare the participants' preferences before and after deliberation....The finding was that post-deliberation preferences tended to be closer to single-peaked than predeliberation preferences...The increases in proximity to single-peakedness were greater for low-salience issues..., on which people's opinions were presumably less entrenched, than for highsalience issues..., on which people were presumably more opinionated.

C. List, R. Luskin, J. Fishkin and I. McLean. *Deliberation, Single-Peakedness, and the Possibility of Meaningful Democracy: Evidence from Deliberative Polls*. Journal of Politics, 75(1), pgs. 80 - 95, 2013.

It remains an open question whether deliberation would also generate single-peaked preferences on issues beyond those covered in the study, and how scalable the mechanism is, i.e., whether the effect could occur in larger groups or in the electorate as a whole...

It remains an open question whether deliberation would also generate single-peaked preferences on issues beyond those covered in the study, and how scalable the mechanism is, i.e., whether the effect could occur in larger groups or in the electorate as a whole...

DELIBERATION Day—a new national holiday. It will be held one week before major national elections. Registered voters will be called together in neighborhood meeting places, in small groups of 15, and larger groups of 500, to discuss the central issues raised by the campaign. Each deliberator will be paid $150 for the day's work of citizenship, on condition that he or she shows up at the polls the next week. All other work, except the most essential, will be prohibited by law.

Details follow.

B. Ackerman and J. S. Fishkin (2002). *Deliberation Day*. Journal of Political Philosophy 10, pp. 129-152.

V. Ottonelli and D. Porello (2013). *On the elusive notion of meta-agreement*. Politics, Philosophy & Economics 12(1), pp. 68 - 92.

D. Porello (2016). *Single-peakedness and semantic dimensions of preferences*. Logic Journal of the IGPL, 24(4), pp. 570 - 583.

Modeling Deliberation

C. List (2011). *Group Communication and the Transformation of Judgments: An Impossibility Result*. Journal of Political Philosophy 19, pp. 1 - 27.

J. Perote-Pena and A. Piggins (2015). *A Model of Deliberative and Aggregative Democracy*. Economics and Philosophy 31, pp. 93-121.

Game Theoretic Issues

D. Landa and A. Meirowitz (2009). *Game Theory, Information, and Deliberative Democracy*. American Journal of Political Science 53, pp. 427 - 444.

T. Feddersen (1998). *Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts under Strategic Voting*. American Political Science Review, 92, pp. 23 - 35.

Suppose a 12-member jury (initially in conditions without deliberation) has to reach a verdict in a criminal trial and uses the unanimity rule, where a "guilty" verdict is reached if and only if all jurors vote for "guilty".

Suppose further that each juror has received some independent and private information, a binary signal of the form "guilty" or "not guilty", which is fallible but correlated with the truth; say, it has a 70% chance of being correct.

However, suppose that each juror is committed to the principle "convict if and only if the defendant's guilt is beyond reasonable doubt", understood as a probability of guilt above 99%. Surprisingly, the jurors may then lack an incentive to vote truthfully.

Suppose I am one of the 12 jurors, and suppose, for the sake of argument, the others will vote truthfully. Should I then vote truthfully too?

- My vote will make a difference only if everyone else votes for "guilty": If some others vote for "not guilty", then the outcome will be a "not guilty" verdict, no matter how I vote. If everyone else votes for "guilty", on the other hand, then my vote will be pivotal; it will determine whether we reach unanimity.

Suppose I am one of the 12 jurors, and suppose, for the sake of argument, the others will vote truthfully. Should I then vote truthfully too?

- My vote will make a difference only if everyone else votes for "guilty": If some others vote for "not guilty", then the outcome will be a "not guilty" verdict, no matter how I vote. If everyone else votes for "guilty", on the other hand, then my vote will be pivotal; it will determine whether we reach unanimity.
- Now, if the others vote truthfully, there is only a small chance that they are all wrong: all 11 signals would have to be incorrect, a chance of (30%): less than two in a million and well below the threshold of reasonable doubt.

Suppose I am one of the 12 jurors, and suppose, for the sake of argument, the others will vote truthfully. Should I then vote truthfully too?

- My vote will make a difference only if everyone else votes for "guilty": If some others vote for "not guilty", then the outcome will be a "not guilty" verdict, no matter how I vote. If everyone else votes for "guilty", on the other hand, then my vote will be pivotal; it will determine whether we reach unanimity.
- Now, if the others vote truthfully, there is only a small chance that they are all wrong: all 11 signals would have to be incorrect, a chance of (30%): less than two in a million and well below the threshold of reasonable doubt.
- So, whether or not my own signal supports a "guilty" verdict, I should vote for "guilty" under the present assumptions. Even if my own signal suggests "not guilty", it is more likely that this signal is wrong than that the others' signals are all wrong.

The example shows that even when all jurors have the same goal – namely to convict if and only if the defendant's guilt is beyond reasonable doubt – voting truthfully is not generally a dominant strategy.

The bottom line is that if we model deliberation as nothing more than an opportunity for voters to share private information before voting—a form of "straw polling"—then the participants' incentives for and against truth-telling depend on several factors:

1. whether there is a consensus on the underlying goal (if so, deliberation induces truth-telling; if not, it doesn't generally do so);
2. whether, in the absence of a consensus, there is uncertainty about voters' biases (if so, truth-telling is sometimes rational; if not, it may not be); and
3. the voting rule (if majority rule, then deliberation sometimes induces truth-telling; if it is the unanimity rule, it may not)

S. Rafiee Rad and O. Roy (2021). *Deliberation, Single-Peakedness, and Coherent Aggregation*. American Political Science Review, 115(2), pp. 629 - 648.

- Each round starts with one of the participants announcing her current preferences and the others updating their opinions accordingly. The next participant then announces her updated preferences, at which point the others once again update their preferences.

- This goes on until all participants have announced their preferences, and then a new round starts, following the same procedure. The order of the participants is randomly reassigned in each round.

- The preference updates follow a distance-minimization rule, weighed by possible biases that the participants may have towards their own opinions.

Let P be a preference profile. After the ith member announces her preference ranking, the new profile P' where for each j, P'_j is the ranking that minimizes

$$\sqrt{\mathrm{rel}_i d(P_i, P'_j)^2 + \mathrm{rel}_j d(P_j, P'_j)^2}$$

where d is a distance measure on rankings, and $\mathrm{rel}_i \in [0,1]$ with $\mathrm{rel}_j = 1 - \mathrm{rel}_i$ representing the bias each voter has towards her own opinion.

FIGURE 2. Proportion of Initial Cyclic Profiles Still Present after Deliberation (*y*-axis) on Strict Rankings as Bias Increases (*x*-axis; *n* = 51)

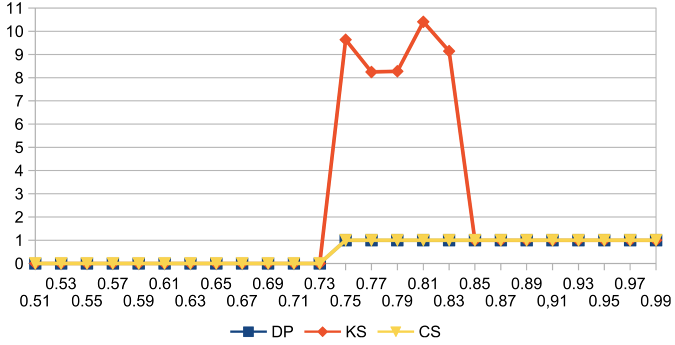

FIGURE 7. **Average Proximity to Single-Plateauedness (*y*-axis) as Bias Increases (*x*-axis; *n* = 51)**

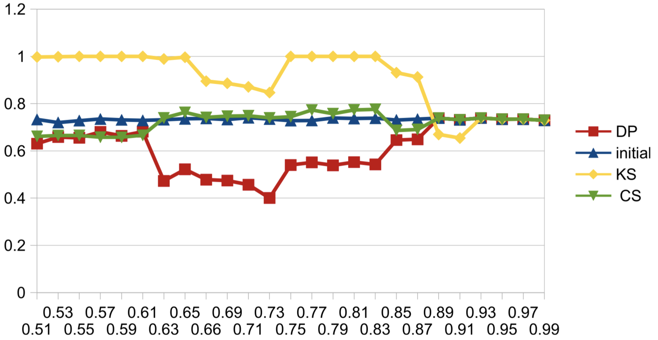

Strategic Voting under Uncertainty

S. Chopra, E. Pacuit and R. Parikh. *Knowledge-theoretic Properties of Strategic Voting*. JELIA 2004.

Strategic Voting under Uncertainty

S. Chopra, E. Pacuit and R. Parikh. *Knowledge-theoretic Properties of Strategic Voting*. JELIA 2004.

H. van Ditmarsch, J. Lang and A. Saffidine. *Strategic voting and the logic of knowledge*. TARK, 2013.

By default, voters will vote *sincerely*.

By default, voters will vote *sincerely*.

Given imperfect information about how the other voters will choose, when will a voter strategize?

Do voters assume that the other voters are voting sincerely or do they allow for the possibility of strategizing?

Protocol 1

If the current winner is X, then agent i will switch its vote to some candidate Y provided

1. Y is one of the top two candidates as indicated by a poll
2. Y is preferred to the other top candidate

4	3	2
A	B	C
C	C	A
B	A	B

4	3	2
A	B	C
C	C	A
B	A	B

Stage	A	B	C	Winner
1	4	3	2	A

4	3	2
A	B	C
C	C	A
B	A	B

Stage	A	B	C	Winner
1	4	3	2	A

4	3	2
A	B	C
C	C	A
B	A	B

Stage	A	B	C	Winner
1	4	3	2	A
2	6	3	0	A

4	3	2
A	B	C
C	C	B
B	A	A

4	3	2
A	B	C
C	C	B
B	A	A

Stage	A	B	C	Winner
1	4	3	2	A

4	3	2
A	B	C
C	C	B
B	A	A

Stage	A	B	C	Winner
1	4	3	2	A
2	4	5	0	B

4	3	2
A	B	C
C	C	A
B	A	B

4	3	2
A	B	C
C	C	B
B	A	A

C is the Condorcet winner (and the Borda winner)

Protocol 2

If the current winner is X, then agent i will switch its vote to some candidate Y provided

1. i prefers Y to X, and
2. the current total for Y plus agent i's votes for Y is greater than the current total for X.

40	30	15	8	7
A	B	C	D	C
D	A	B	A	A
B	C	D	B	B
C	D	A	C	D

40	30	15	8	7
A	B	C	D	C
D	A	B	A	A
B	C	D	B	B
C	D	A	C	D

Stage	40	30	15	8	7	A	B	C	D	Winner
1	A	B	C	D	C	40	30	22	8	A

40	30	15	8	7
A	B	C	D	C
D	A	B	A	A
B	C	D	B	B
C	D	A	C	D

Stage	40	30	15	8	7	A	B	C	D	Winner
1	A	B	C	D	C	40	30	22	8	A
2	A	B	B	D	C	40	45	7	8	B

40	30	15	8	7
A	B	C	D	C
D	A	B	A	A
B	C	D	B	B
C	D	A	C	D

Stage	40	30	15	8	7	A	B	C	D	Winner
1	A	B	C	D	C	40	30	22	8	A
2	A	B	B	D	C	40	45	7	8	B

40	30	15	8	7
A	B	C	D	C
D	A	B	A	A
B	C	D	B	B
C	D	A	C	D

Stage	40	30	15	8	7	A	B	C	D	Winner
1	A	B	C	D	C	40	30	22	8	A
2	A	B	B	D	C	40	45	7	8	B

40	30	15	8	7
A	B	C	D	C
D	A	B	A	A
B	C	D	B	B
C	D	A	C	D

Stage	40	30	15	8	7	A	B	C	D	Winner
1	A	B	C	D	C	40	30	22	8	A
2	A	B	B	D	C	40	45	7	8	B

40	30	15	8	7
A	B	C	D	C
D	A	B	A	A
B	C	D	B	B
C	D	A	C	D

Stage	40	30	15	8	7	A	B	C	D	Winner
1	A	B	C	D	C	40	30	22	8	A
2	A	B	B	D	C	40	45	7	8	B
3	D	B	B	A	A	15	45	0	40	B

40	30	15	8	7
A	B	C	D	C
D	A	B	A	A
B	C	D	B	B
C	D	A	C	D

Stage	40	30	15	8	7	A	B	C	D	Winner
1	A	B	C	D	C	40	30	22	8	A
2	A	B	B	D	C	40	45	7	8	B
3	D	B	B	A	A	15	45	0	40	B
4										

40	30	15	8	7
A	B	C	D	C
D	A	B	A	A
B	C	D	B	B
C	D	A	C	D

Stage	40	30	15	8	7	A	B	C	D	Winner
1	A	B	C	D	C	40	30	22	8	A
2	A	B	B	D	C	40	45	7	8	B
3	D	B	B	A	A	15	45	0	40	B
4	A	B	B	D	A	47	45	0	8	A

Protocol 3

If the current winner is X, then agent i will switch its vote to some candidate Y provided

1. i prefers Y to X, and
2. the some past votes for Y plus agent i's votes for Y is greater than the current total for X.

4	3	3
A	B	C
B	C	A
C	A	B

Stage	4	3	3	A	B	C	Winner
1	A	B	C	4	3	3	A

4	3	3
A	B	C
B	C	A
C	A	B

Stage	4	3	3	A	B	C	Winner
1	A	B	C	4	3	3	A
2	A	C	C	4	0	6	C

4	3	3
A	B	C
B	C	A
C	A	B

Stage	4	3	3	A	B	C	Winner
1	A	B	C	4	3	3	A
2	A	C	C	4	0	6	C
3	??	C	C			6	C

4	3	3
A	B	C
B	C	A
C	A	B

Stage	4	3	3	A	B	C	Winner
1	A	B	C	4	3	3	A
2	A	C	C	4	0	6	C
3	B	C	C	0	3	6	C

4	3	3
A	B	C
B	C	A
C	A	B

Stage	4	3	3	A	B	C	Winner
1	A	B	C	4	3	3	A
2	A	C	C	4	0	6	C
3	B	C	C	0	3	6	C
4	B	B	C	0	7	3	B

4	3	3
A	B	C
B	C	A
C	A	B

Stage	4	3	3	A	B	C	Winner
1	A	B	C	4	3	3	A
2	A	C	C	4	0	6	C
3	B	C	C	0	3	6	C
4	B	B	C	0	7	3	B
5	B	B	A	3	7	0	B

4	3	3
A	B	C
B	C	A
C	A	B

Stage	4	3	3	A	B	C	Winner
1	A	B	C	4	3	3	A
2	A	C	C	4	0	6	C
3	B	C	C	0	3	6	C
4	B	B	C	0	7	3	B
5	B	B	A	3	7	0	B
6	A	B	A	7	3	0	A

4	3	3
A	B	C
B	C	A
C	A	B

Stage	4	3	3	A	B	C	Winner
1	A	B	C	4	3	3	A
2	A	C	C	4	0	6	C
3	B	C	C	0	3	6	C
4	B	B	C	0	7	3	B
5	B	B	A	3	7	0	B
6	A	B	A	7	3	0	A
7	A	C	C	4	0	6	C

Stage	4	3	3	A	B	C	Winner
1	A	B	C	4	3	3	A
2	A	C	C	4	0	6	C
3	B	C	C	0	3	6	C
4	B	B	C	0	7	3	B
5	B	B	A	3	7	0	B
6	A	B	A	7	3	0	A
7	A	C	C	4	0	6	C
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮

Uncertainty About the Profiles

How should we represent the a voter's uncertainty about how the other voters' preferences?

- Quantitative Information: Probability measure over space of profiles
- Qualitative Information: Information partition on the space of profiles

Poll information function: $\pi : O(X)^V \to \mathcal{I}$, where $\mathcal{I}$ is some property of the profiles. For example,

Poll information function: $\pi : O(X)^V \to \mathcal{I}$, where $\mathcal{I}$ is some property of the profiles. For example,

- Profile: p is the identity function
- Ballot: p returns a vector recording how often each ballot occurs in the input profile
- (Weighted) Majority Graph: p returns the weighted majority graph
- Score: p returns the score for each candidate
- Rank: p returns the ranking of each candidate
- Winner: p returns the winner
- Zero: p returns 0

Let $\pi : O(X)^V \rightarrow \mathcal{I}$ be a poll information function and $\boldsymbol{P}$ a profile. The information cell generated by $\boldsymbol{P}$ is:

$$I_i(\boldsymbol{P}) = \{\boldsymbol{P}' \in O(X)^V \mid P'_i = R_i \text{ and } \pi(\boldsymbol{P}) = \pi(\boldsymbol{P}')\}$$

Ulle Endriss, Svetlana Obraztsova, Maria Polukarov, and Jeffrey S. Rosenschein. *Strategic Voting with Incomplete Information*. In Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI-2016), July 2016.

Iterated Voting

Allow voters the opportunity to *change* their votes in response to certain "poll" information.

- Will this process *converge*?

Iterated Voting

Allow voters the opportunity to *change* their votes in response to certain "poll" information.

- Will this process *converge*?
- Cf. the work on convergence to Nash equilibria in the Game Theory literature
- Cf. the work on *deliberative democracy*

O. Lev and J. Rosenschein. *Convergence of Iterative Voting*. AAMAS, 2012.

A. Reijngoud and U. Endriss. *Voter Response to Iterated Poll Information*. AAMAS, 2012.

A **voter response strategy**:

- Truth-teller: always voter truthfully
- Strategist: Computes best response to poll information and uses (any) one of them.
- k-Pragmatist: Does not compute a best response, but rather moves her favorite of the currently k-highest ranked candidates to the top of her ballot.

Given a social choice function F, F^t is the social choice function induced by starting with truthful preferences, the voting game is played for t rounds.

Given a social choice function F, F^t is the social choice function induced by starting with truthful preferences, the voting game is played for t rounds.

- Which properties of the social choice function are also satisfied by F^t? (Unanimity, Condorcet consistency)
- Run a simulation and measure the Condorcet efficiency of different voting rules.

M. Brill (2019). *Interactive Democracy: New Challenges for Social Choice Theory*. In J.-F. Laslier, H. Moulin, R. Sanver, and W. S. Zwicker, editors, Future of Economic Design, pp. 59–66.

U. Grandi (2017). *Social choice and social networks*. In U. Endriss, editor, Trends in Computational Social Choice, Ch. 9, pp. 169 - 184.

H. Landermore (2020). *Open Democracy: Reinventing Popular Rule for the Twenty-First Century*. Princeton University Press.