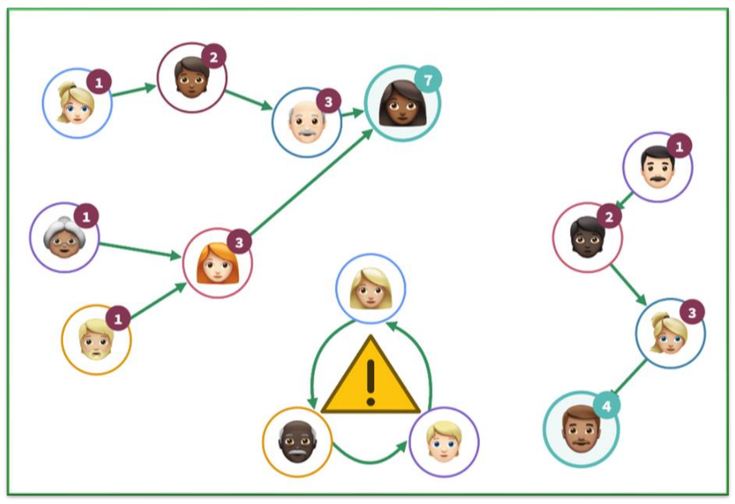
Cycles in Liquid Democracy

A Game-Theoretic Justification

Markus Brill *Oxford* · Rachael Colley *Glasgow* · Anne-Marie George *Oslo* · Grzegorz Lisowski *Groningen* · Georgios Papasotiropoulos *Warsaw* · Ulrike Schmidt-Kraepelin *Eindhoven*



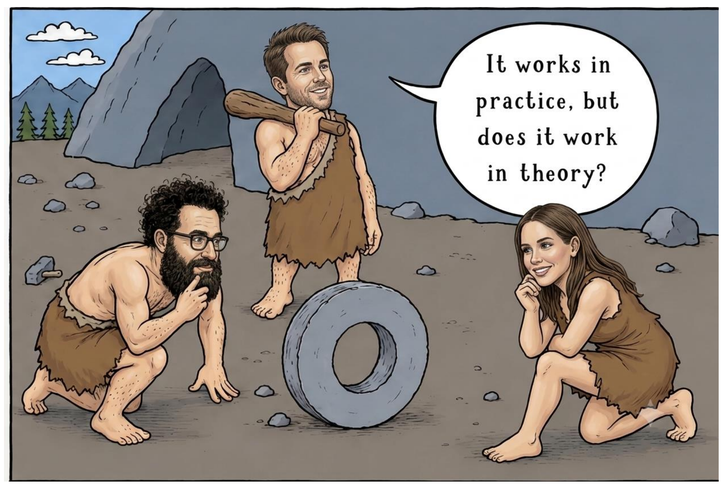
Theory vs Practice:

Are cycles a bug or a feature?

Most of the **academic literature** treats delegation cycles as a problem, as voting weight is lost.

Practitioners argue that delegation cycles are unproblematic and are **formed intentionally**.

Platforms allow voters to set **default delegations** that remain fixed across **multiple decisions**.

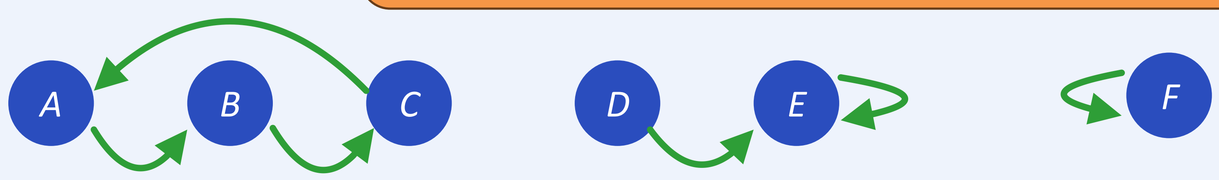


Our **default delegation model** captures strategic incentives of participants on liquid democracy platforms. We provide a **game-theoretic analysis** of how voters choose delegates.

The Default Delegation Model

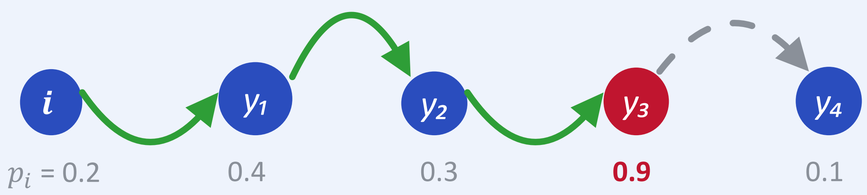
Default delegations

Each voter nominates a **backup delegate** (possibly themselves), to be used whenever they don't cast a vote themselves. In the resulting **delegation graph**, every voter has **out-degree 1**.



Probabilistic participation

To model **uncertainty** about participation, we assume that each voter i casts a vote with probability p_i. This yields a probability distribution over i's **ultimate delegate**.



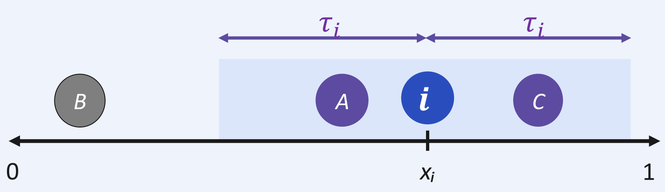
Probability that voter y_ℓ is the ultimate delegate of voter $i = y_0$:

$$p_{y_\ell} \times \prod_{r=0}^{\ell-1} (1 - p_{y_r})$$

(there is no ult. del. with prob. $\prod_{r=0}^{k} (1 - p_{y_r})$)

Voter utility

We assume that voters are positioned on the interval $[0, 1]$. Voters want their ultimate delegate **as close as possible** to their own position, and have a **tolerance τ**.



$$u_i = \begin{cases} \tau_i - \text{dist}(i, y) & \text{if } y \text{ is the ult. del. of } i \\ 0 & \text{if } i \text{ has no ult. del.} \end{cases}$$

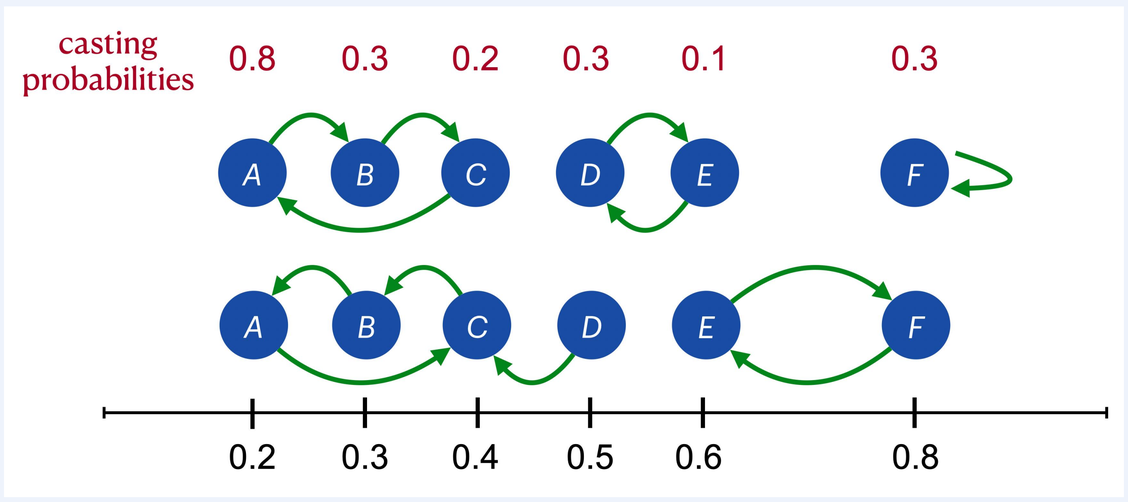
(u_i could also be normalized)

Results

Existence of Nash Equilibria

A delegation profile is a **Nash equilibrium** (NE) if no voter gains by unilaterally changing their delegate.

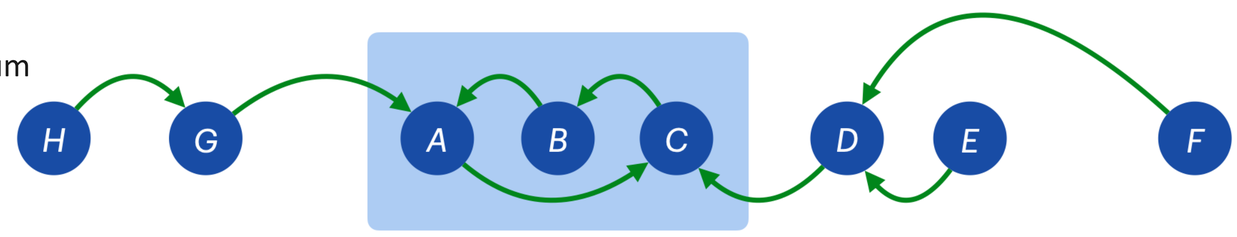
Equilibria need **not** be **unique** (example on the right, with $\tau = \frac{1}{4}$) and need **not exist** generally, yet **best-response dynamics** found an NE in all 20,000 randomly generated instances we ran.



Structure of Nash Equilibria

Assuming mutual acceptance, every delegation graph in equilibrium contains **exactly one cycle** per non-trivial weakly connected component, with trees feeding into it.

Cycles aren't an accident — they're the rule.



Quality of Nash Equilibria

Social welfare of delegation profile $\mathbf{d}$: $SW(\mathbf{d}) = \sum_i u_i(\mathbf{d})$.

Price of Anarchy (PoA) measures how far away an NE can be from the social optimum, in terms of social welfare:

$$\text{PoA} = \frac{SW(\mathbf{d_{SW}})}{SW(\mathbf{d_{NE}})}$$

$\mathbf{d_{SW}}$: social-welfare maximizing profile

$\mathbf{d_{NE}}$: Nash equilibrium profile

In the worst case, the PoA can be **arbitrarily large** — but never more than $\frac{1}{p_{\min}}$.

In experiments, the PoA is **close to 1**: Nash equilibria reach near-optimal welfare.

Future Directions

- Computational complexity
- More general metric spaces
- Other utility functions
- Other aspects of liquid democracy
- Real-world data???