Fair and Efficient Social Decision-Making

CSCI 699

Applications of Voting

Evi Micha

Participatory Budgeting

Participatory Budgeting

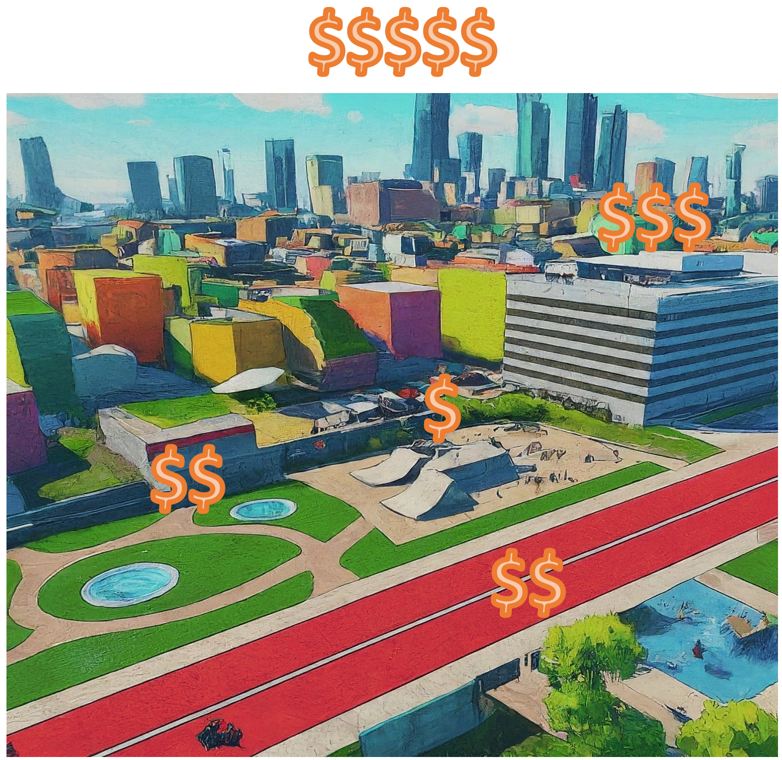

Participatory Budgeting

- Started in Porto Alegre, Brazil

- Currently, it is implemented in more than 1,500 cities around the world!
  - Madrid, 2021, EUR 50M
  - Paris, 2023, EUR 76M
  - Montreal, 2024, CAD 45M
  - Los Angeles, 2024, USD 8.5M

Participatory Budgeting in LA

Participatory Budgeting in LA

What is Participatory Budgeting?

Participatory budgeting (PB) is a democratic process in which community members decide how to spend part of a public budget.

PB started in Porto Alegre, Brazil in 1989 as an anti-poverty measure that helped reduce child mortality by nearly 20%. Since then, PB has spread to over 7,000 cities around the world, and has been used to decide budgets from states, counties, cities, housing authorities, schools, and other institutions. Click here to learn more about how Participatory Budgeting works.

The L.A. REPAIR Participatory Budgeting pilot program dedicates $8.5 million of City funding to nine underserved communities to decide how it should be spent on programs that serve the community.

$8.5M across 9 REPAIR Zones

1. DESIGN THE PROCESS — REPAIR Zone Committees develop the program Guidebook and plan for implementation.

2. IDEA COLLECTION — REPAIR Zone community members share their ideas for programs.

3. DEVELOP PROPOSALS — Nonprofits and community-based organizations transform ideas into program proposals.

4. VOTE — REPAIR Zone community members vote for the program proposals they want funded.

5. FUND & IMPLEMENT PROGRAMS — The City funds program proposals with the most votes and monitors implementation.

6. EVALUATE THE PROCESS — REPAIR Zone Committees and City reflect to make recommendations to the Mayor and City Council.

Participatory Budgeting in LA

Proposal A

Van Nuys Blvd. Mural Revitalization Project

ISSUE AREA:
Environment & Climate

ORGANIZATION:
Connectopod Learning, Inc.

BUDGET:
$353,925

TARGET POPULATION:
100 youth, 30,000 audience members

LOCATION:
Van Nuys Corridor

Connectopod proposes to engage youth in the restoration, creation, and documentation of local murals. First, local youth will restore four murals and create two new murals, led by Levi Ponce and local muralists who started the mural movement in Pacoima. Then, they will create a digital interactive map showing the revitalized Mural Mile and telling the individual background and meaning to each mural. Youth will learn interviewing and production skills, and document the process using a podcast and film centered around community conversations about these murals, the impact of murals in Los Angeles, and what Mural Mile means to the people of Pacoima. The goal is to embrace the history and culture around these murals, and build a sense of community pride.

Proposal B

Nurturing Health

ISSUE AREA:
Health & Wellbeing

ORGANIZATION:
El Nido Family Centers

BUDGET:
$315,772

TARGET POPULATION:
1,500 individuals

LOCATION:
El Nido's Pacoima Family Source Center and Parent Centers

This program aims to promote health and wellness within families. Key services include expanding the Pacoima Farmers Market with more fitness classes, nutrition workshops and healthy food demonstrations, a gardening program, and healthcare access. Also, the program will host a large Arleta community event, offer a yoga series, and organize nature excursions. This initiative will benefit over 1,500 residents, aiming to improve their health, nutrition, and community connection.

Basic Approval Model

- A set $N = \{1, \dots, n\}$ of voters
- $P = \{p_1, \dots, p_m\}$ of projects
- Each $p \in M$ has a cost $c(p) \geq 0$
  - For each $W \subseteq P$, $c(W) = \sum_{p \in W} c(p)$
- Budget limit $B \geq 0$
- The outcome $W \subseteq P$ is budget-feasible, i.e. $c(W) \leq B$
- Each $i \in N$ approves $A_i \subseteq P$
- Each $i \in N$ has utility $u_i(p)$ for each $p \in P$
  - Additive Utilities: For each $W \subseteq P$, $u_i(W) = \sum_{p \in W} u_i(p)$
  - For now, assume that $u_i(p)$ =1 for each $p \in A_i$, i.e. $u_i(W) = |W \cap A_i|$

Welfare Maximization

- Choose W that maximizes the utilitarian welfare (i.e. $\sum_{i \in N} u_i(w)$), subject to the budget constraint ($c(W) \leq B$) (Knapsack problem)
- Greedy Algorithm: Adds alternatives in order of approval score, skipping those that are unaffordable

$B = 52000$

Approval Votes	Cost	Greedy	Optimal
500	10,000	✖	✖
457	42,000	✖	
449	7,3000		✖
430	13,800		✖
398	9,500		✖
323	11,200		✖

Welfare Maximization

- Choose W that maximizes the utilitarian welfare (i.e. $\sum_{i \in N} u_i(w)$), subject to the budget constraint ($c(W) \leq B$) (Knapsack problem)
- Greedy Algorithm: Adds alternatives in order of approval score, skipping those that are unaffordable

Population: 12K

Population: 10K

Fairness

- Extended Justifies Representation in Committee Selection
  - W satisfies EJR if
    - For all $S \subseteq N$ and $\ell \in \{1, \dots, k\}$
    - If $|S| \geq \ell \cdot {}^{n}/_{k}$ (large) and $|\bigcap_{i \in S} A_i| \geq \ell$ (cohesive)
    - Then $u_i(W) \geq \ell$ for some $i \in S$
    - "If a group deserves ℓ candidates and has ℓ commonly approved candidates, then not every member should get less than ℓ utility"
  - PAV satisfies EJR
- Extended Justifies Representation in Participatory Budgeting
  - W satisfies EJR if
    - For all $S \subseteq N$ and $T \subseteq P$
    - If $c(T) \leq {}^{|S|}/_{n} \cdot B$ (affordable) and $T \subseteq \bigcap_{i \in S} A_i$ (cohesive)
    - Then $u_i(W) \geq |T|$ for some $i \in S$
    - "If a group can afford T that is approved by each $i \in S$, then not every member in S should get less than $|T|$ utility"
  - Does PAV satisfy EJR in Participatory Budgeting?

PAV

- Given a sequence $s = (1, {}^{1}/_{2}, {}^{1}/_{3}, \ldots)$, select a committee W that maximizes $\sum_{i \in N} s_{u_i(W)}$

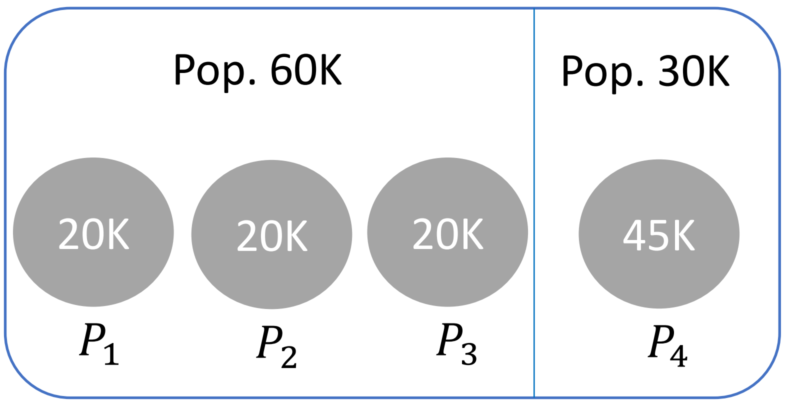

Budget 90K

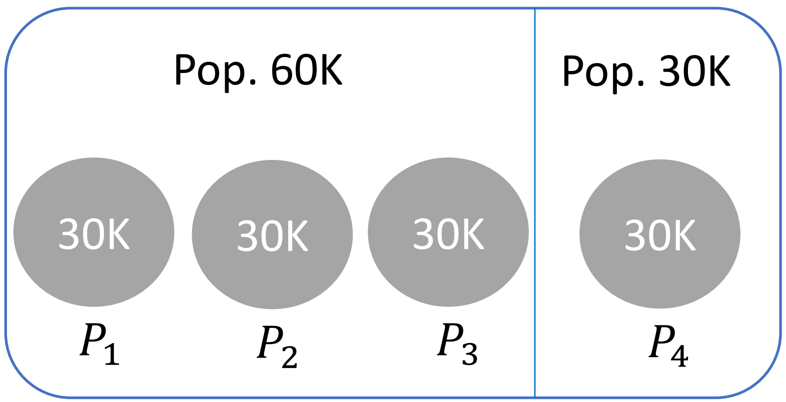

Budget 90K

PAV

- Given a sequence $s = (1, {}^1\!/_2, {}^1\!/_3, \ldots)$, select a committee W that maximizes $\sum_{i \in N} s_{u_i(W)}$

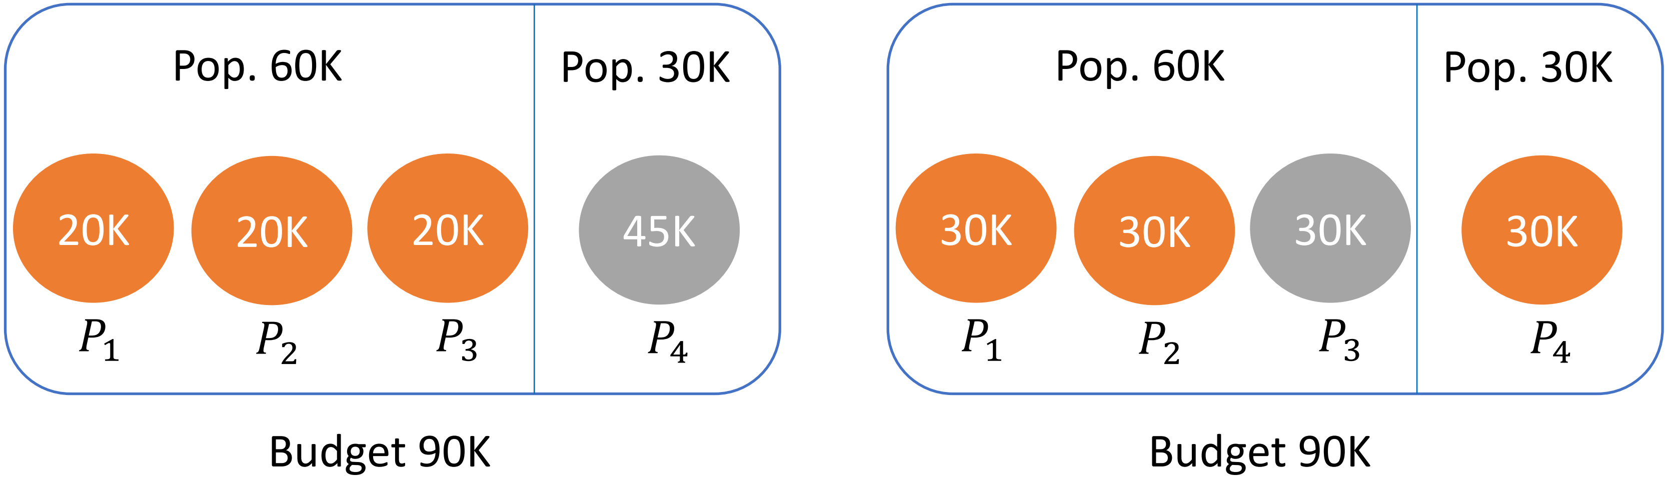

- PAV violates proportionality
- Theorem [Peters et al. '21]: Every voting rule that only depends on voters' utility functions and the collection of budget-feasible sets must fail proportionality

Method of Equal Shares

- Split the budget equally among the voters, i.e. each voter gets B/n
- Until the budget is exhausted:
  - For each project no funded yet, divide its cost as evenly as possible among its approvers
  - Fund an affordable alternative with the lowest max payment

	1	2	3	4
p_1	✖	✖	✖	
p_2		✖	✖	✖
p_3	✖			✖

p_1	p_2	p_3
6	10.5	8

$B = 20$

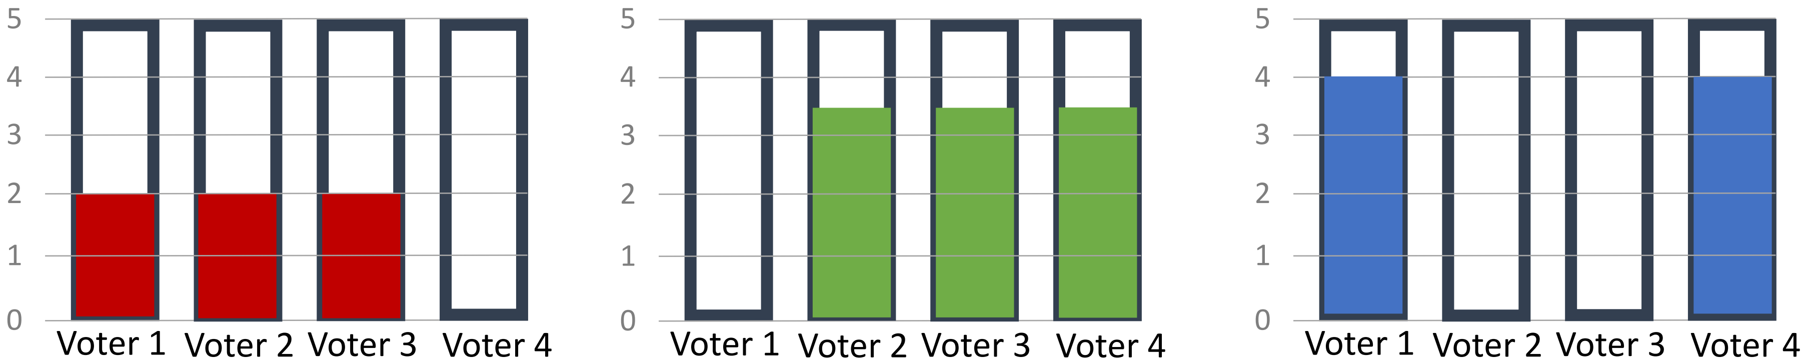

Method of Equal Shares

- Split the budget equally among the voters, i.e. each voter gets B/n
- Until the budget is exhausted:
  - For each project no funded yet, divide its cost as evenly as possible among its approvers
  - Fund an affordable alternative with the lowest max payment

	1	2	3	4
p_1	✖	✖	✖	
p_2		✖	✖	✖
p_3	✖			✖

p_1	p_2	p_3
6	10.5	8

$B = 20$

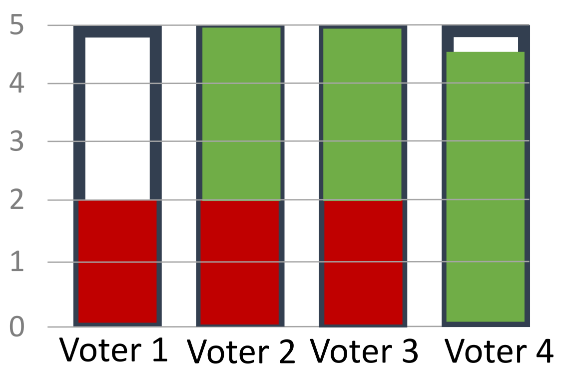

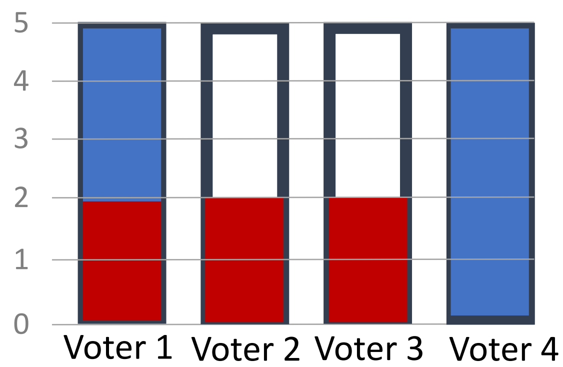

Method of Equal Shares

- Theorem [Peters et al. '21]: The Method of Equal Shares satisfies EJR

- In 2023, Method of Equal Shares was used in Aarua (Switzerland), Wieliczka (Poland) and Świecie (Poland)

Fairness

- Core in Committee Selection:
  - W satisfies core if
    - For all $S \subseteq N$ and $T \subseteq M$
    - If $|S| \geq |T| \cdot {}^{n}/_{k}$ (large)
    - Then $u_i(W) \geq u_i(T)$ for some $i \in S$
    - "If a group can afford T, then T should not be a (strict) Pareto improvement for the group"
- Core in Participatory Budgeting
  - W satisfies EJR if
    - For all $S \subseteq N$ and $T \subseteq P$
    - If $c(T) \leq {}^{|S|}/_{n} \cdot B$ (affordable)
    - Then $u_i(W) \geq u_i(T)$ for some $i \in S$
    - "If a group can afford T, then T should not be a (strict) Pareto improvement for the group"

Fairness

- The core can be empty
  - 3 voters, 3 projects with cost 2 each, budget is B=3
  - Whichever project we select, two of the voters can improve from (0,1)→(1,2)

	p_1	p_2	p_3
v_1	2	1	0
v_2	0	2	1
v_3	1	0	2

Reinforcement Learning with Human Feedback (RLHF)

AI Alignment

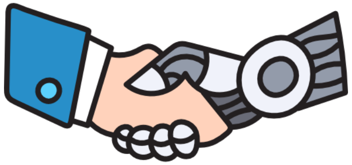

…AI alignment involves ensuring that an AI system's objectives match those of its designers…
(wikipedia)

Reinforcement Learning from Human Feedback (RLHF)

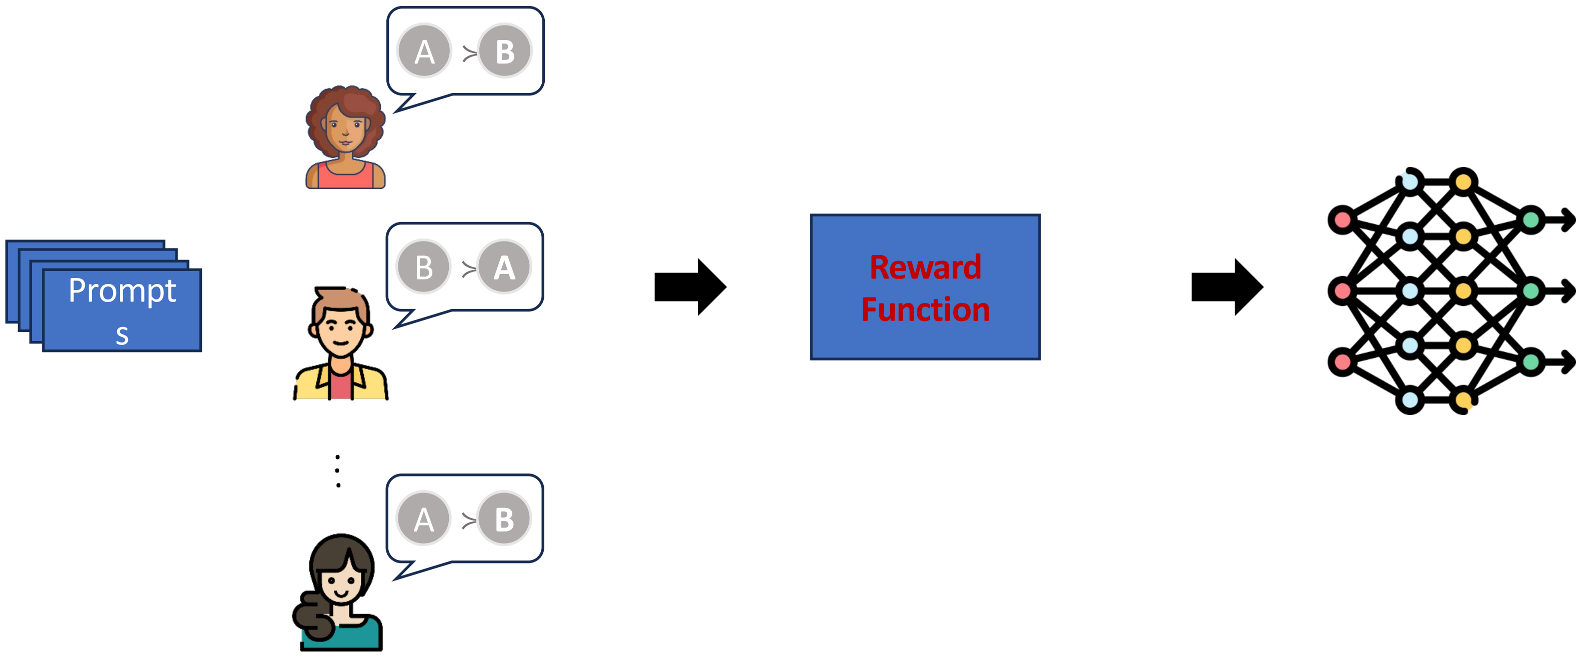

Reinforcement Learning from Human Feedback (RLHF)

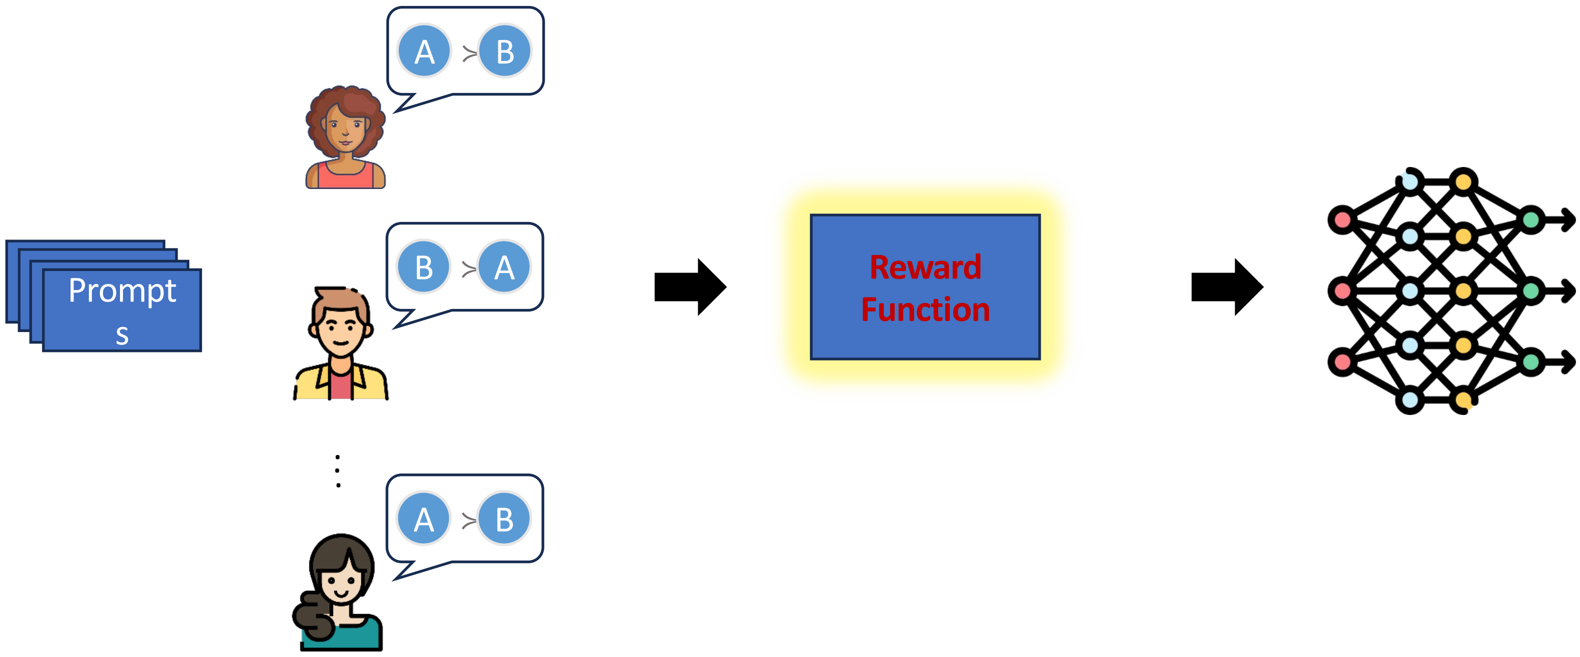

Random Utility Model

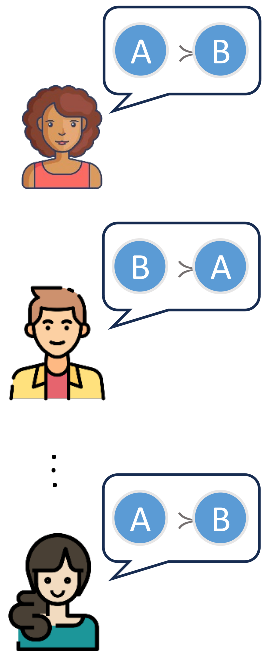

BTL Model

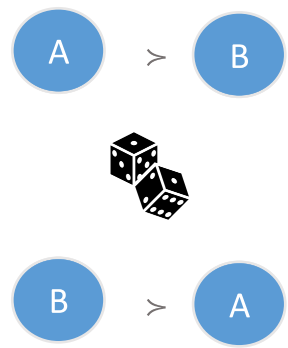

$$\frac{e^{r_\theta(A)}}{e^{r_\theta(A)} + e^{r_\theta(B)}}$$

$$\frac{e^{r_\theta(B)}}{e^{r_\theta(A)} + e^{r_\theta(B)}}$$

Profile of Ordinal Preferences

$$inf_\theta \; L(\theta;\pi) = inf_\theta \; \sum_{A\neq B} n_{A\succ B}(\pi) \cdot ln(1 + e^{r_\theta(B) - r_\theta(A)})$$

Number of voters in π that prefer A to B

Heterogeneous Prefences

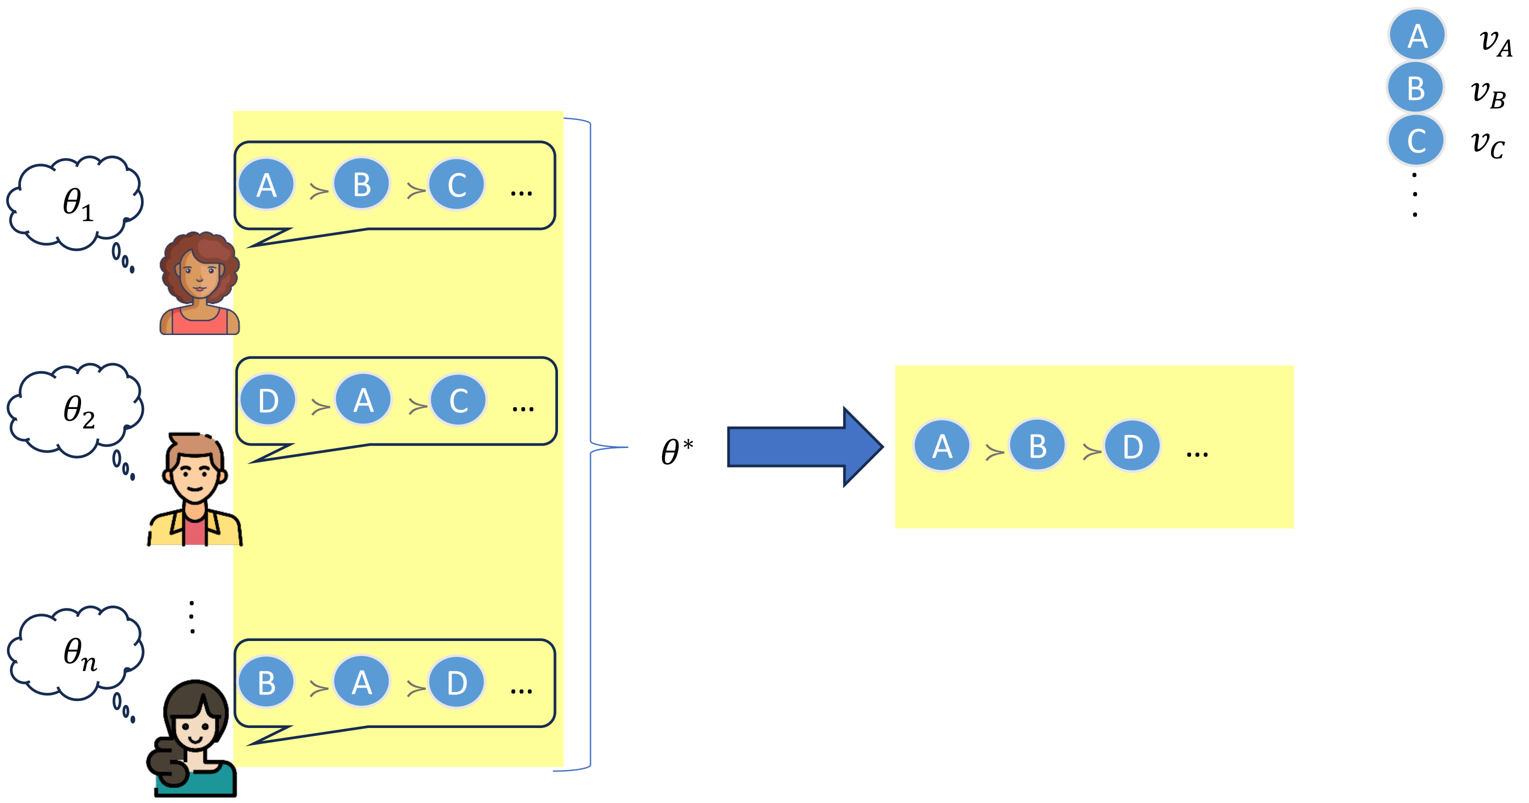

Linear Model: $\boldsymbol{r_\theta(v) = \langle \theta, v \rangle}$

Linear Social Choice

A $v_A = [20,0,0]$

B $v_B = [0,20,0]$

C $v_C = [0,10,10]$

D $v_D = [0,0,1]$

E $v_E = [1,0,0]$

$\theta = [\theta_1, \theta_2, \theta_3]$ ➡ A $\succ$ B $\succ$ C $\succ$ D $\succ$ E

Linear Social Choice

A $v_A = [20,0,0]$

B $v_B = [0,20,0]$

C $v_C = [0,10,10]$

D $v_D = [0,0,1]$

E $v_E = [1,0,0]$

$\theta = [\theta_1, \theta_2, \theta_3]$ ➡ A ≻ B ≻ C ≻ D ≻ E

$\theta_1 > \theta_2$

Linear Social Choice

A $v_A = [20,0,0]$

B $v_B = [0,20,0]$

C $v_C = [0,10,10]$

D $v_D = [0,0,1]$

E $v_E = [1,0,0]$

$\theta = [\theta_1, \theta_2, \theta_3]$ ➡ A $\succ$ B $\succ$ C $\succ$ D $\succ$ E

$\theta_1 > \theta_2$

$\theta_2 > \theta_3$

Linear Social Choice

A $v_A = [20,0,0]$

B $v_B = [0,20,0]$

C $v_C = [0,10,10]$

D $v_D = [0,0,1]$

E $v_E = [1,0,0]$

$\theta = [\theta_1, \theta_2, \theta_3]$ ➡ A $\succ$ B $\succ$ C $\succ$ D $\succ$ E

$\theta_1 > \theta_2$

$\theta_2 > \theta_3$

$\theta_3 > \theta_1$

Linear Social Choice

A $v_A = [20,0,0]$

B $v_B = [0,20,0]$

C $v_C = [0,10,10]$

D $v_D = [0,0,1]$

E $v_E = [1,0,0]$

$\theta = [\theta_1, \theta_2, \theta_3]$ ➡

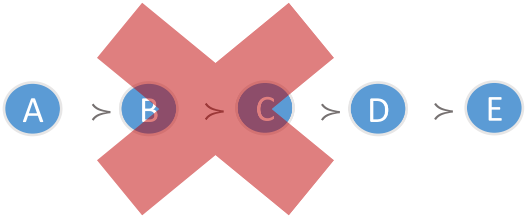

Linear Rank Aggregation Rules

Axiomatic Approach

Goals:

- What axioms are satisfied by aggregation methods used by existing RLHF algorithms?
- Are there alternative aggregation methods that offer stronger axiomatic guarantees?

- **Pareto Optimality:** A linear rank aggregation rule f satisfies Pareto optimality if, whenever every voter prefers candidate a over candidate b, then candidate a is ranked higher than candidate b in the output ranking

- **Pairwise Majority Consistency (PMC):** A ranking σ is called a PMC ranking for profile π if for all a, b ∈ C, $a \succ_\sigma b$ if and only if a majority of voters rank $a \succ b$. A linear rank aggregation rule satisfies PMC if, when a PMC ranking σ exists for the input profile π and σ is feasible, then $f(\pi) = \sigma$

Loss-Based Rules

A loss function $\ell: \mathbb{R} \rightarrow \mathbb{R}$

$$\inf_{\theta} \; L(\theta; \pi, \ell) = \inf_{\theta} \; \sum_{a \neq b} n_{a \succ b}(\pi) \cdot \ell(r_\theta(b) - r_\theta(a))$$

BTL model: $\ell(x) = \ln(1 + e^x)$

Theorem (informal): If a linear rank aggregation rule f optimizes a loss function that is either nondecreasing and weakly convex, or strictly convex then f ***fails PO and PMC***

A Social Choice Based Rule

- **Leximax Copeland subject to PO**

σ_1	σ_2	…	σ_n
1	2		
2	1		
3	3		
⋮	⋮		
m	$m - 1$		

Copeland →

σ^*
2
1
3
⋮
$m - 1$

→

σ'

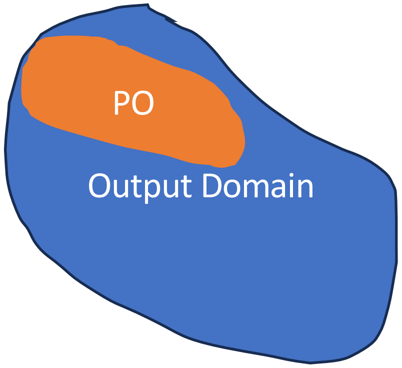

A Social Choice Based Rule

- **Leximax Copeland subject to PO**

σ_1	σ_2	…	σ_n
1	2		
2	1		
3	3		
⋮	⋮		
m	$m-1$		

Copeland →

σ^*
2
1
3
⋮
$m-1$

→

σ'

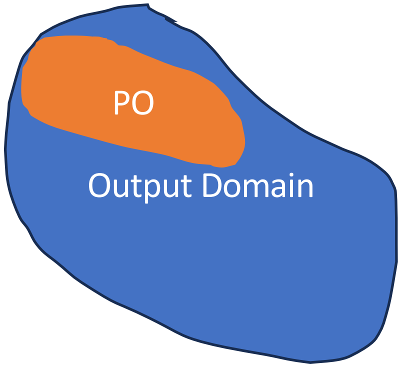

A Social Choice Based Rule

- **Leximax Copeland subject to PO**

σ_1	σ_2	…	σ_n
1	2		
2	1		
3	3		
⋮	⋮		
m	$m - 1$		

Copeland →

σ^*
2
1
3
⋮
$m-1$

→

σ'

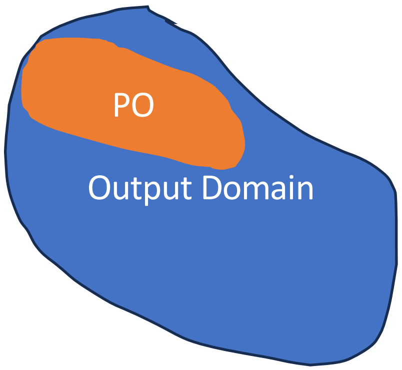

A Social Choice Based Rule

- **Leximax Copeland subject to PO**

σ_1	σ_2	…	σ_n
1	2		
2	1		
3	3		
$\vdots$	$\vdots$		
m	$m - 1$		

Copeland →

σ^*
2
~~1~~
3
$\vdots$
$m - 1$

→

σ'
1

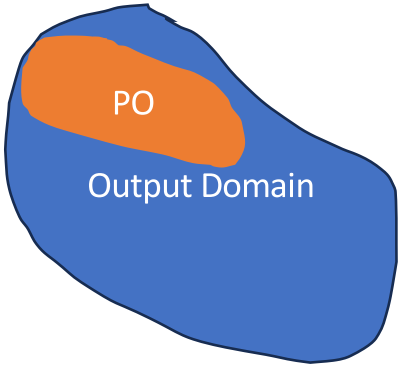

A Social Choice Based Rule

- **Leximax Copeland subject to PO**

σ_1	σ_2	…	σ_n
1	2		
2	1		
3	3		
⋮	⋮		
m	$m-1$		

Copeland →

σ^*
2
~~1~~
3
⋮
$m-1$

→

σ'
1

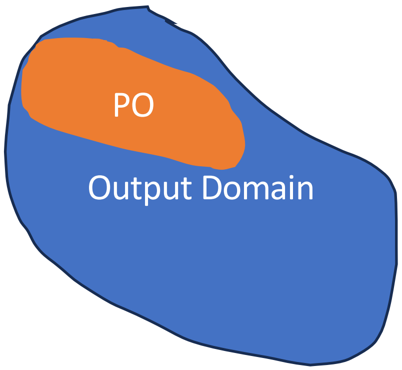

A Social Choice Based Rule

- **Leximax Copeland subject to PO**

σ_1	σ_2	…	σ_n
1	2		
2	1		
3	3		
⋮	⋮		
m	$m-1$		

Copeland →

σ^*
2
~~1~~
3
⋮
$m-1$

→

σ'
1

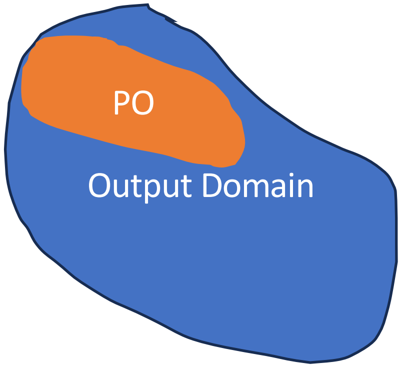

A Social Choice Based Rule

- **Leximax Copeland subject to PO**

σ_1	σ_2	…	σ_n
1	2		
2	1		
3	3		
⋮	⋮		
m	$m-1$		

Copeland →

σ^*
2
~~1~~
~~3~~
⋮
$m-1$

→

σ'
1
3

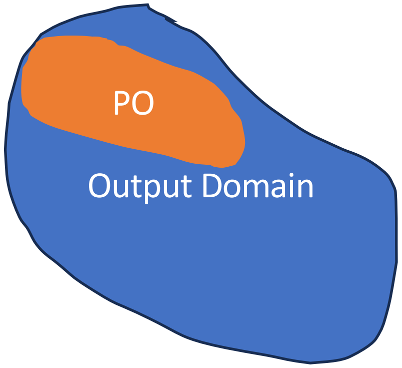

A Social Choice Based Rule

- **Leximax Copeland subject to PO**

σ_1	σ_2	…	σ_n
1	2		
2	1		
3	3		
⋮	⋮		
m	$m - 1$		

Copeland →

σ^*
2
~~1~~
~~3~~
⋮
$m - 1$

→

σ'
1
3

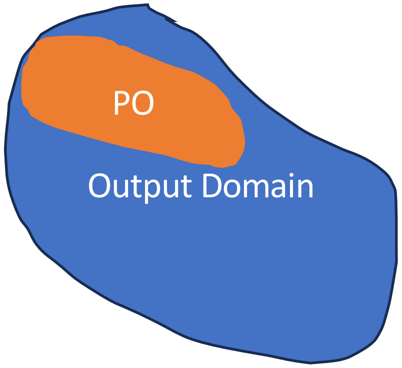

A Social Choice Based Rule

- **Leximax Copeland subject to PO**

σ_1	σ_2	…	σ_n
1	2		
2	1		
3	3		
⋮	⋮		
m	$m-1$		

Copeland →

σ^*
~~2~~
~~1~~
~~3~~
⋮
$m-1$

→

σ'
1
3
2

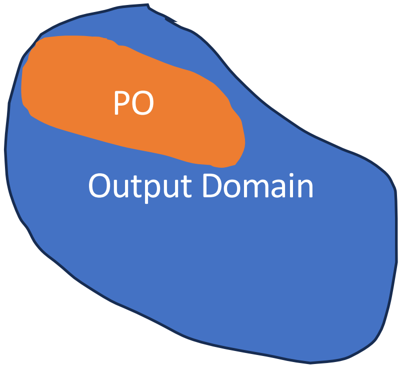

A Social Choice Based Rule

- **Leximax Copeland subject to PO**

σ_1	σ_2	…	σ_n
1	2		
2	1		
3	3		
⋮	⋮		
m	$m-1$		

Copeland →

σ^*
~~2~~
~~1~~
~~3~~
⋮
~~$m-1$~~

→

σ'
1
3
2
m

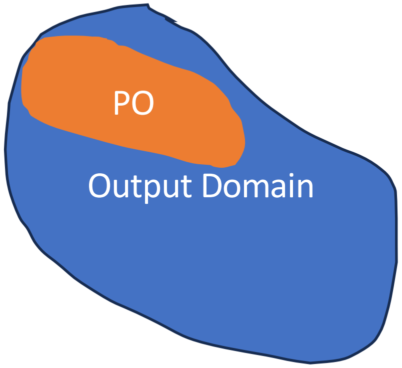

A Social Choice Based Rule

- **Theorem:** Leximax Copeland subject to PO **satisfies**
  a) PO
  b) PMC

A Social Choice Based Rule

- **Theorem:** Leximax Copeland subject to PO **satisfies**
  a) PO
  b) PMC
  c) majority consistency
  d) winner monotonicity

- **Majority Consistency:** A linear rank aggregation rule f satisfies majority consistency if when a candidate a is ranked first by a majority of voters in the input profile, a is ranked first in the output ranking

- **Winner Monotonicity:** A linear rank aggregation rule f satisfies winner monotonicity if, when a candidate a is ranked first in the output ranking, elevating a in any voter's preference does not cause a to lose their top position in the updated aggregate ranking

A Social Choice Based Rule

- **Theorem:** Leximax Copeland subject to PO **satisfies**
  a) PO
  b) PMC
  c) majority consistency
  d) winner monotonicity

  and can be implemented in polynomial time by solving $O(m^2)$ small linear programs

- **Majority Consistency:** A linear rank aggregation rule f satisfies majority consistency if when a candidate a is ranked first by a majority of voters in the input profile, a is ranked first in the output ranking

- **Winner Monotonicity:** A linear rank aggregation rule f satisfies winner monotonicity if, when a candidate a is ranked first in the output ranking, elevating a in any voter's preference does not cause a to lose their top position in the updated aggregate ranking

Liquid Democracy

Democracy

Direct Democracy

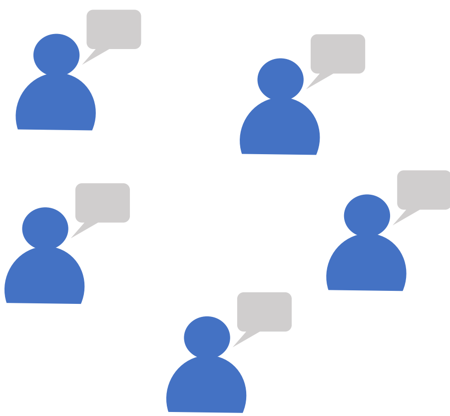

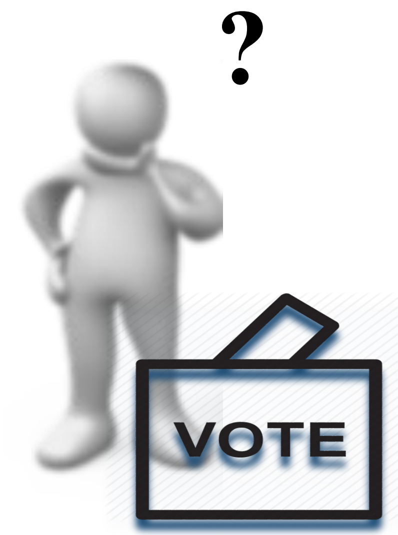

Democracy

Direct Democracy

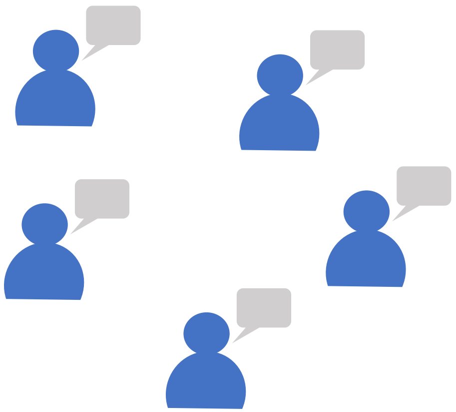

Representative Democracy

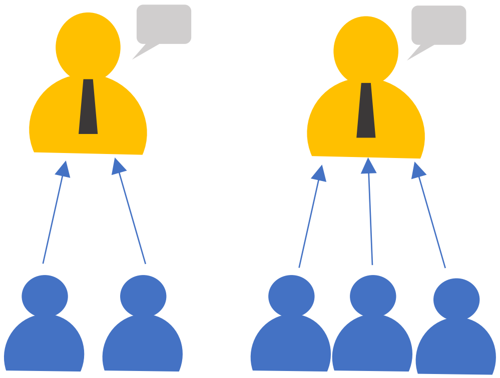

Democracy

Direct Democracy

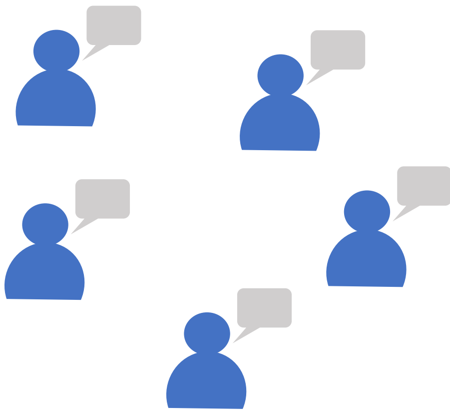

Liquid Democracy

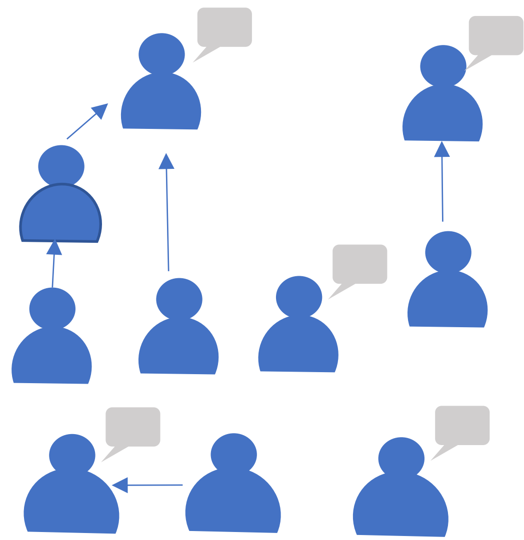

Representative Democracy

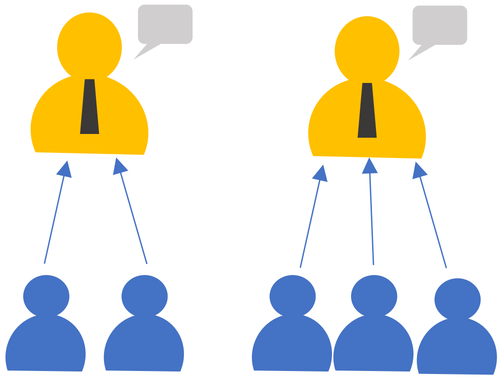

Liquid Democracy

- Voters consist a social network with one-way connections

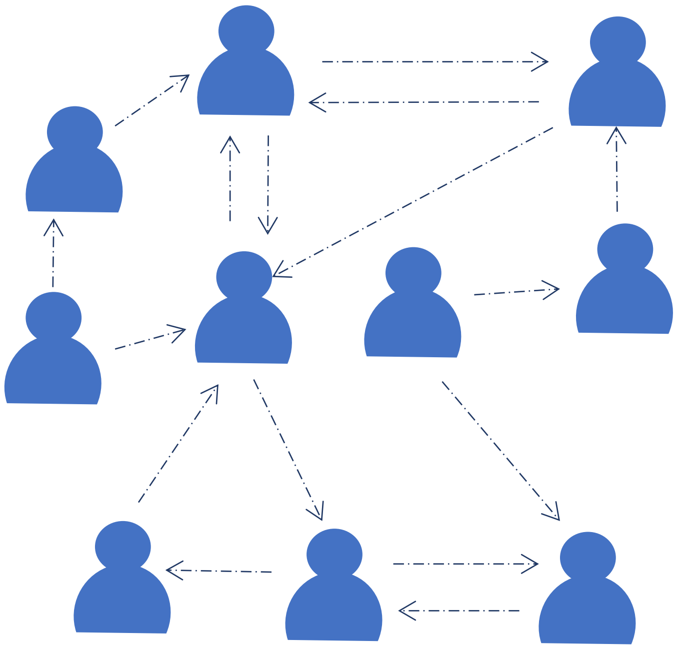

Liquid Democracy

- Voters consist a social network with one-way connections
- Every voter can either cast her vote or delegate her right to vote to another citizen

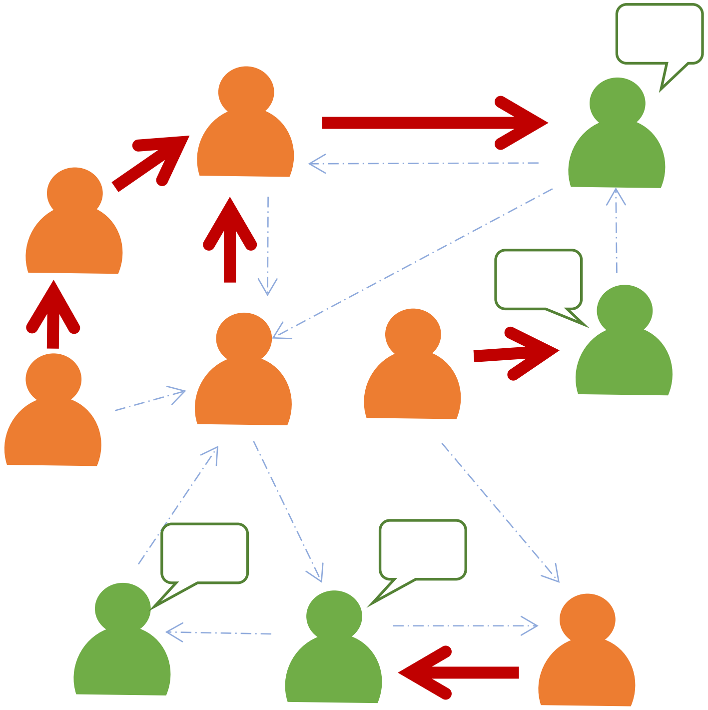

Liquid Democracy

- Voters consist a social network with one-way connections
- Every voter can either cast her vote or delegate her right to vote to another citizen
- Delegations are transitive

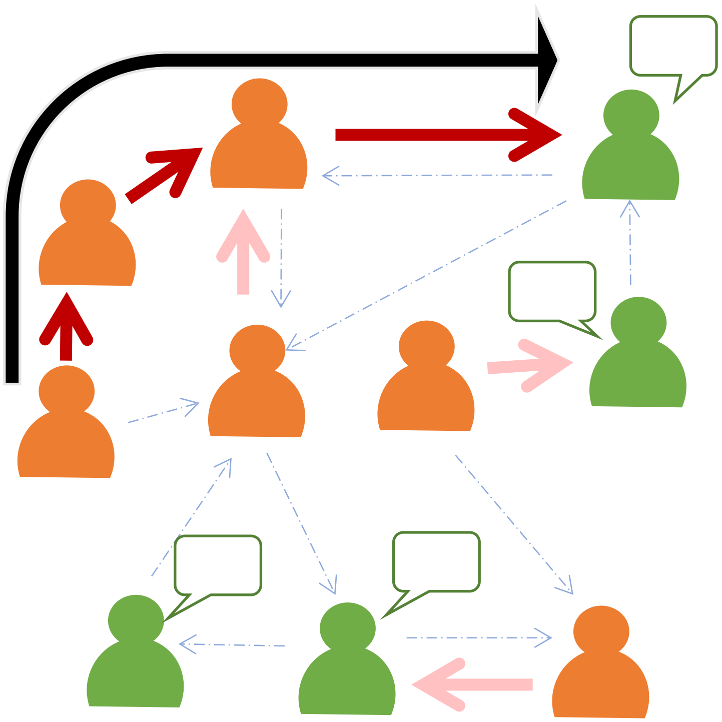

Liquid Democracy

- Voters consist a social network with one-way connections
- Every voter can either cast her vote or delegate her right to vote to another citizen
- Delegations are transitive
- Every voter has a weight equal to the number of the votes that have been delegated to her, plus her right to vote

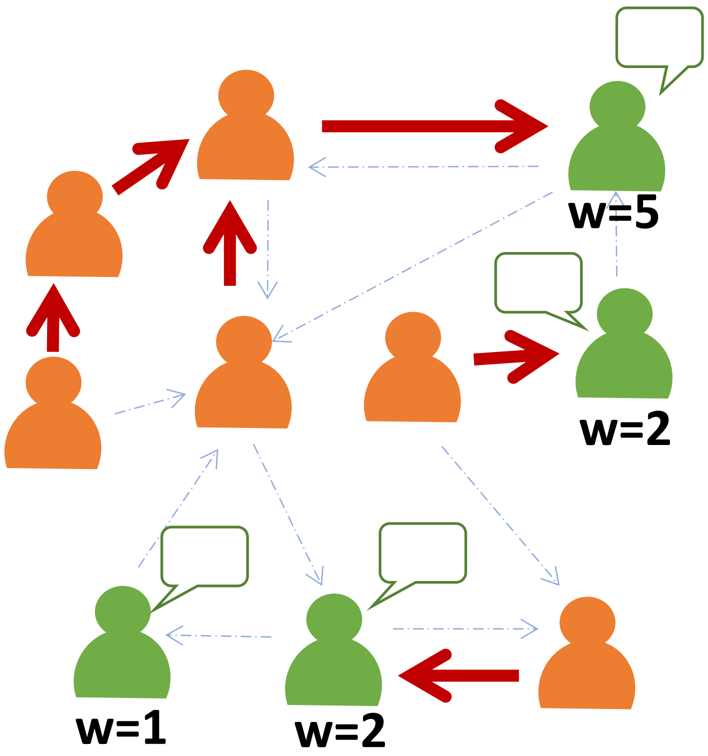

Liquid Democracy

- Voters consist a social network with one-way connections
- Every voter can either cast her vote or delegate her right to vote to another citizen
- Delegations are transitive
- Every voter has a weight equal to the number of the votes that have been delegated to her, plus her right to vote
- In a binary election, the outcome is usually decided by weighted majority

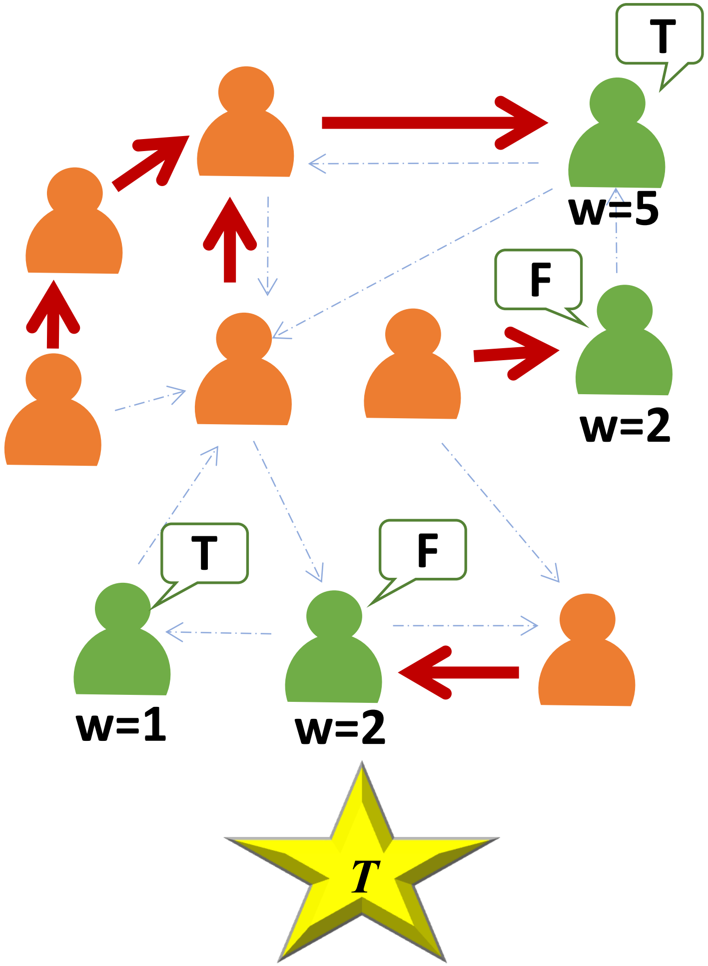

Liquid Democracy Applications

Germany, 2010

Argentina, 2012

Australia, 2016

Model

- Elections with two alternatives, T and F, such that $T \succ F$
- Each voter i has **competency level** p_i, which is the probability that i casts a vote for T
- The social network is represented by a directed graph, $G(V, E)$, where agents are the nodes and an edge (i, j) indicates that i "follows" j
- Optimistic View: i ***approves*** j if $(i, j) \in E$ and $p_j > p_i + \alpha$

Liquid vs Direct Democracy

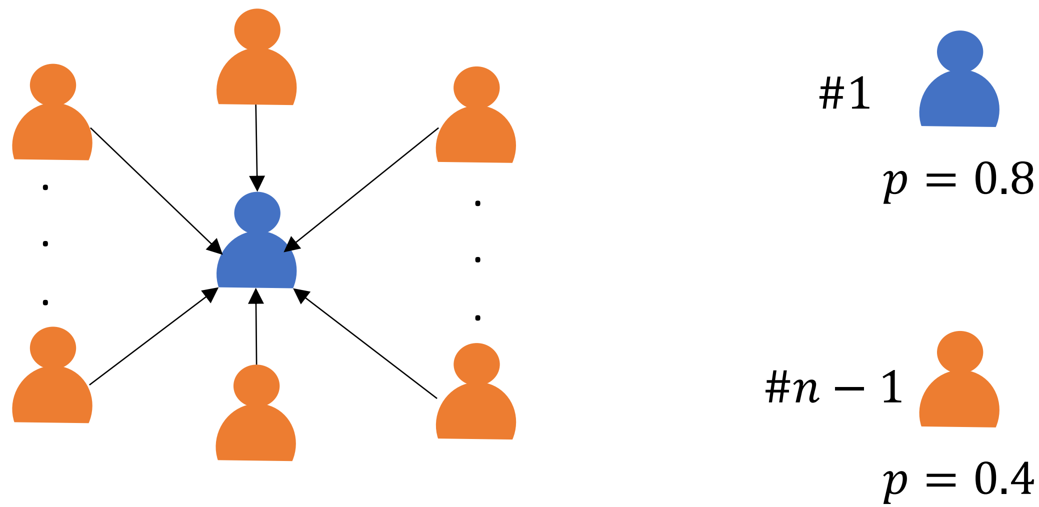

- Direct Democracy: As n grows , due to the Condorcet Jury Theorem, the probability that the majority votes for the correct alternative goes to 0
- Liquid Democracy: If all leaves delegate to the hub, the probability of electing the correct alternative is 0.8

Liquid vs Direct Democracy

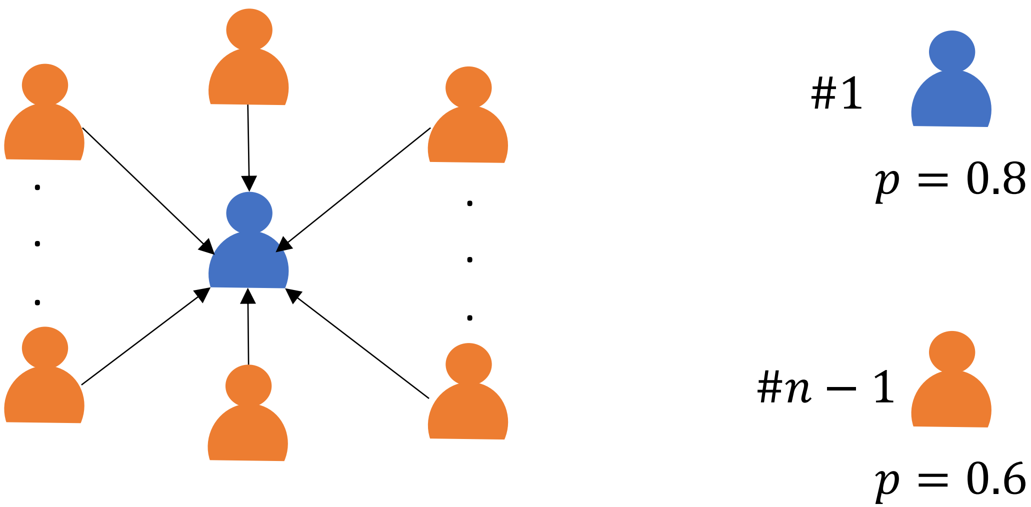

- Direct Democracy: As n grows , due to the Condorcet Jury Theorem, the probability that the majority votes for the correct alternative goes to 1
- Liquid Democracy: If all leaves delegate to the hub, the probability of electing the correct alternative is 0.8

Delegation Mechanisms

- Given $G(V, E)$ and vector $\boldsymbol{p}$, a mechanism M decides (possibly at random) the delegations
  - $P_M(G, \boldsymbol{p})$ is the probability that the outcome is T, under M
    - Apply M on $(G(V, E), \boldsymbol{p})$
    - Sample an instance from the distribution
    - Each sink has weight equal to the number of vertices with directed paths to it
    - Each sink i votes for the correct alternative with probability p_i
    - The winner is determined by weighted majority
  - A mechanism is ***local*** when the decision for every agent *i* depends on her neighborhood

Local Delegation Mechanisms

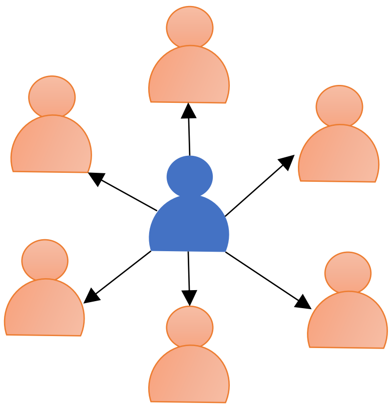

No delegations

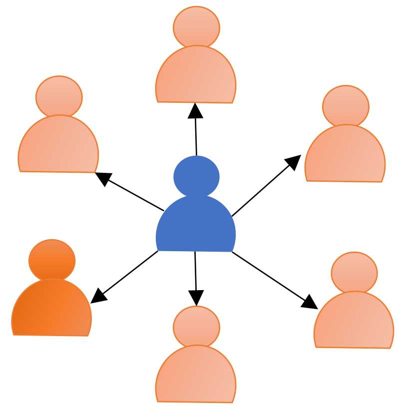

Delegate to the neighbor with the largest probability to vote for the correct alternative

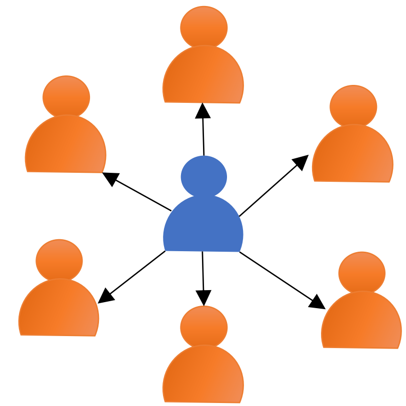

Delegate uniformly at random

Can we design a "good" delegation mechanism?

Local Delegation Mechanisms

- Gain of a mechanism M: $gain(M, G, \boldsymbol{p}) = P_M(G, \boldsymbol{p})$- $P_D(G, \boldsymbol{p})$ where D is direct democracy
- Do no harm: For every $\epsilon > 0$, there exists $n_0 \in \mathbb{Z}$ such that for all $G(V, E)$ with $|V| \geq n_o$ and $\boldsymbol{p}$, $gain(M, G, \boldsymbol{p}) \geq -\epsilon$
- Positive Gain: There exist $\gamma > 0$ and an instance $(G, \boldsymbol{p})$ such that $gain(M, G, \boldsymbol{p}) \geq \gamma$
- Theorem [Kahng et. al, 18]: For any $\alpha \in [0,1)$ there is no local delegation mechanism that satisfies do no harm and positive gain

Local Delegation Mechanisms

- Direct delegation Mechanism (D): No voter delegates
- Full delegation Mechanism (F): Every voter delegates to more informative voters whenever possible

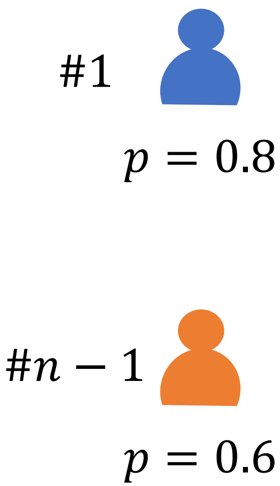

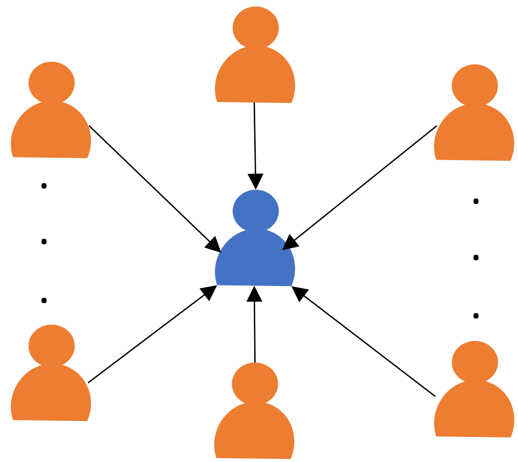

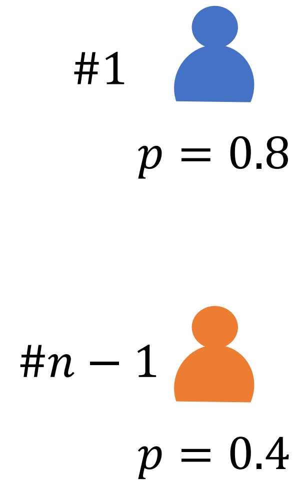

- Theorem [M and Caragiannis, 2019]: For any $\alpha \in [0,1)$ and $\delta > 0$, there exists an instance $(G, \boldsymbol{p})$ such that either $P_D(G,\boldsymbol{p})$- $P_M(G,\boldsymbol{p}) \geq \frac{1}{2} - \alpha - \delta$ or $P_F(G,\boldsymbol{p})$- $P_M(G,\boldsymbol{p}) \geq \frac{1}{2} - \alpha - \delta$

Local Delegation Mechanisms

- High level idea of the proof:
  - Case I: Delegate with probability $< \frac{1}{2}$

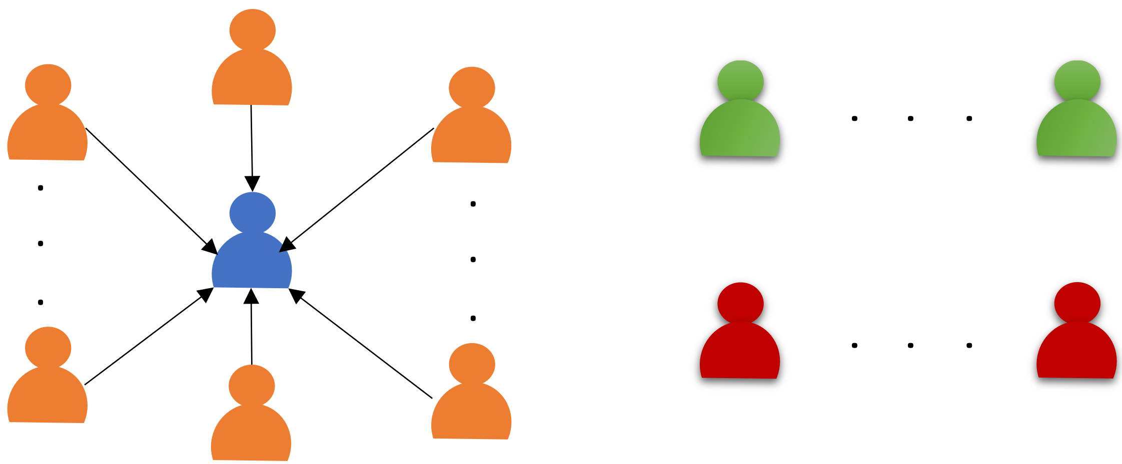

$\#1$ — $\frac{1}{2} + \alpha + \frac{\delta}{2}$

$\#k$ — $1/2$

$\# \frac{n+1}{2} - k - 1$ — 1

$\# \frac{n-1}{2}$ — 0

Local Delegation Mechanisms

- High level idea of the proof:
  - Case II: Delegate with probability $\geq \frac{1}{2}$

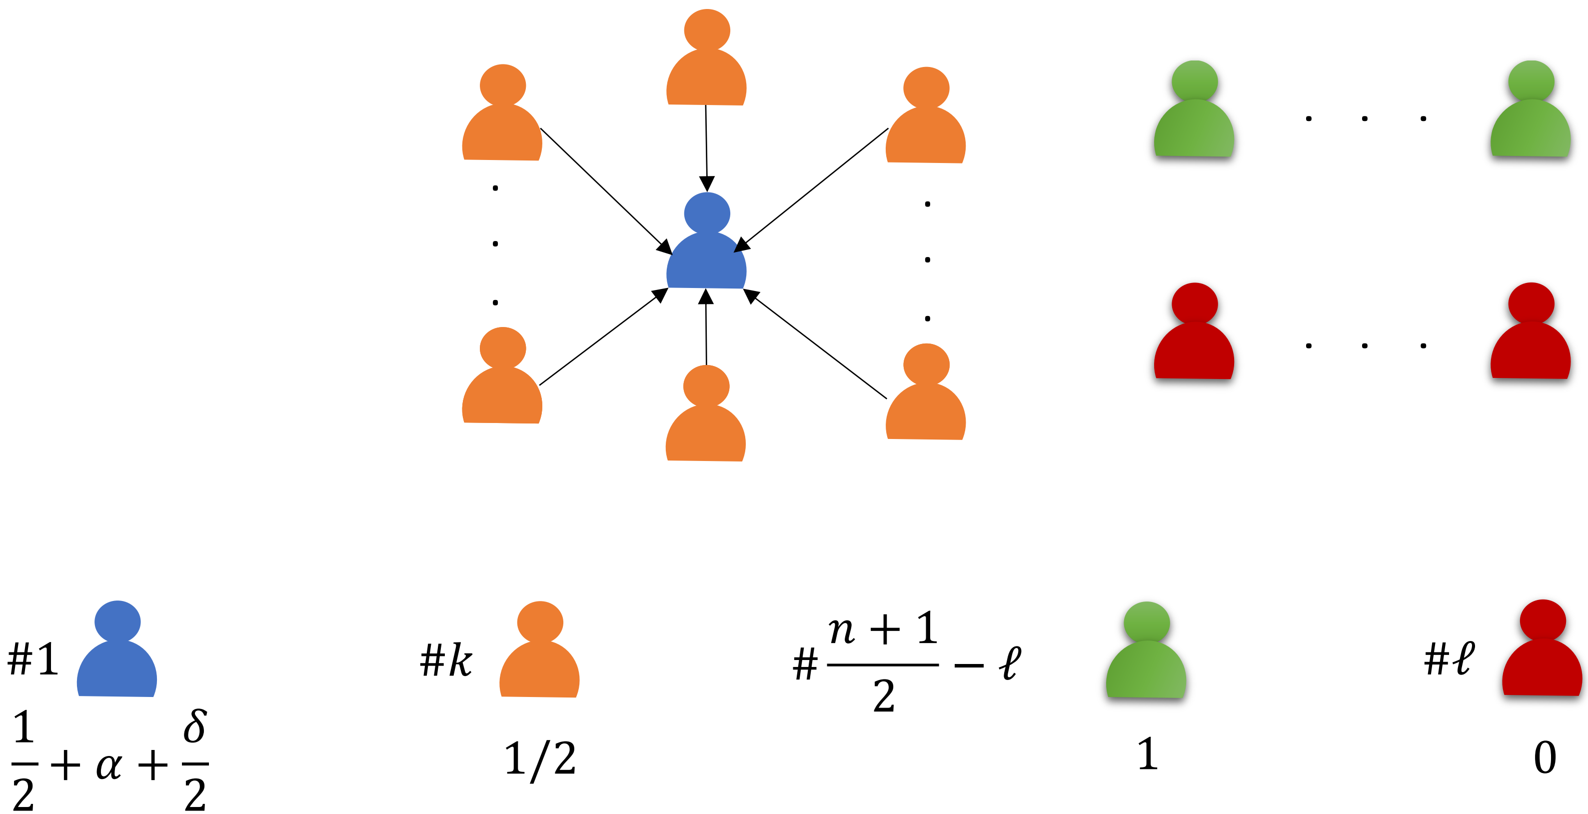

Delegations to Less Informed Agents

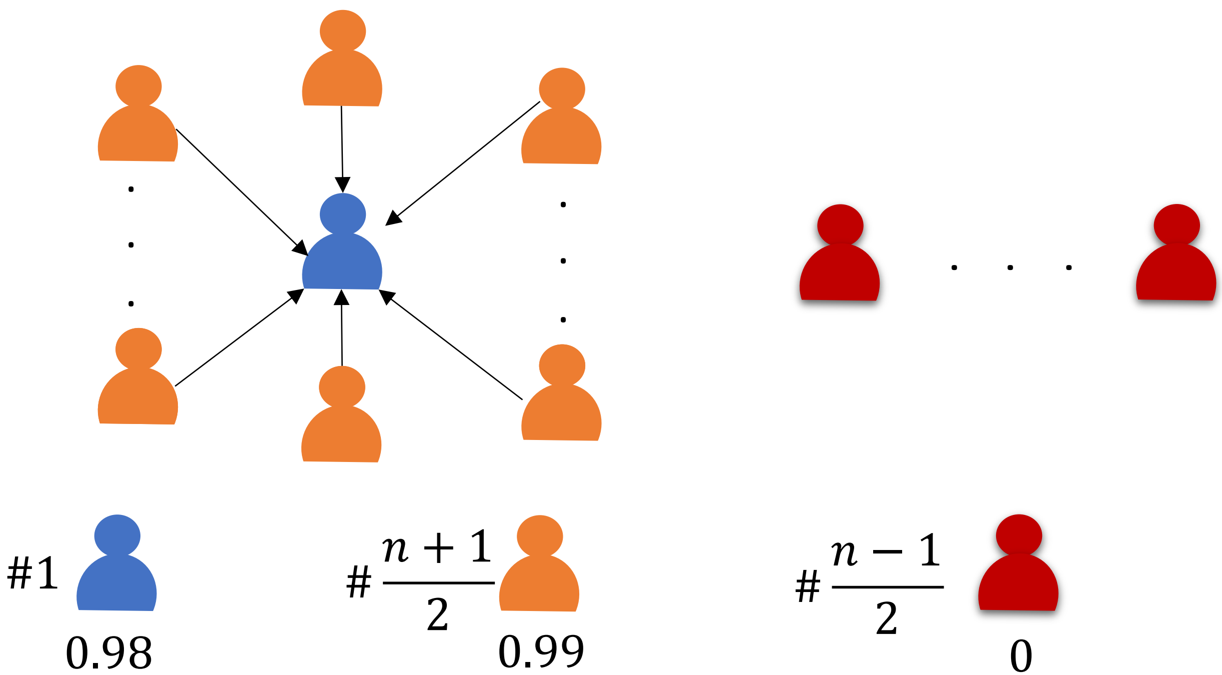

- $P_D(G, \boldsymbol{p})$ goes to 0, as n grows
- $P_F(G, \boldsymbol{p}) = 0.98$