Network-Induced Strategic Communication in Opinion Dynamics

Hassan Munif[1]

Joint work with: Anthony Couthures[2], Vineeth Varma[1], Samson Lasaulce[1], Tamer Başar[3]

[1]Université de Lorraine, CNRS, CRAN, F-54000 Nancy, France
[2]ICTEAM, UCLouvain, Louvain-la-Neuve, Belgium
[3]University of Illinois Urbana-Champaign, USA

5th November 2026 · Politecnico di Torino

UNIVERSITÉ DE LORRAINE

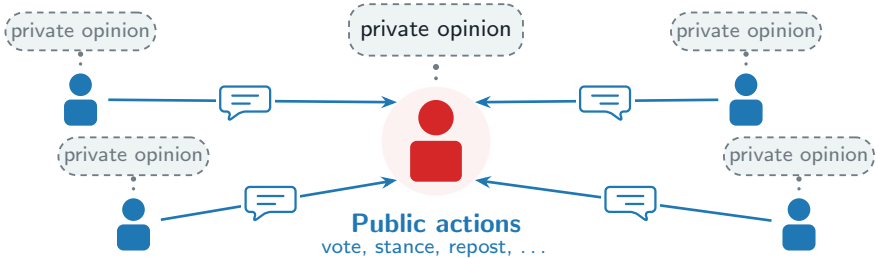
CRAN

How do opinions evolve on a network?



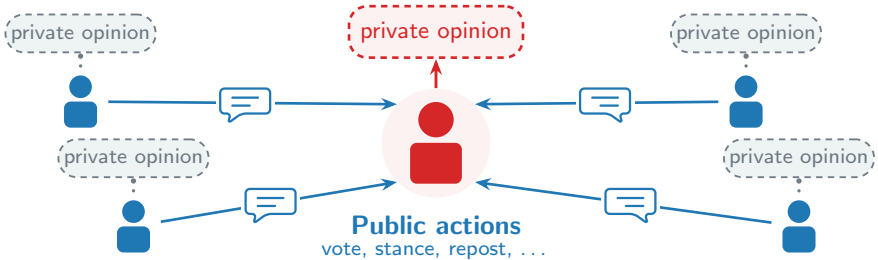
- Agents hold **private** opinions, which others cannot see.

How do opinions evolve on a network?



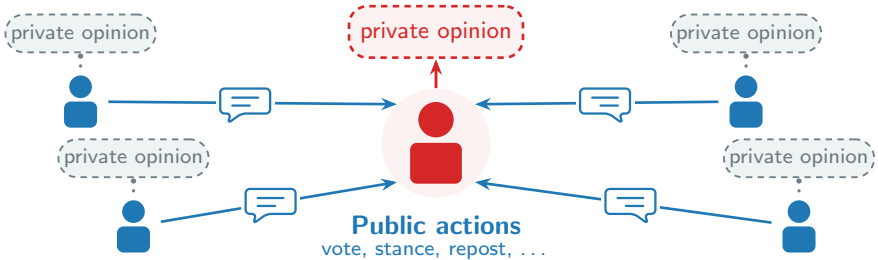
- Agents hold **private** opinions, which others cannot see.
- **Public actions** are the only thing others observe.

How do opinions evolve on a network?



- Agents hold **private** opinions, which others cannot see.
- **Public actions** are the only thing others observe.
- Whatever the mechanism of influence, opinions evolve from **what is public** (actions).

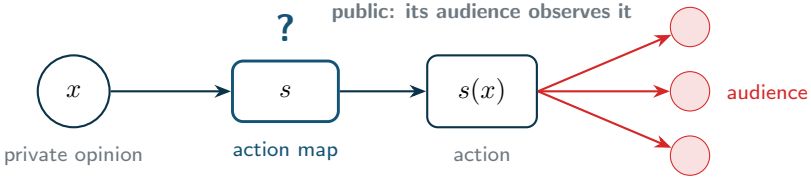
How do opinions evolve on a network?



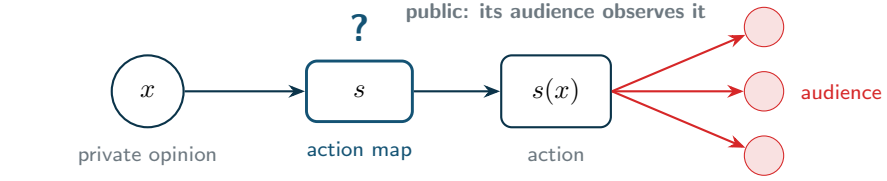
- Agents hold **private** opinions, which others cannot see.
- **Public actions** are the only thing others observe.
- Whatever the mechanism of influence, opinions evolve from **what is public** (actions).

Actions drive opinion evolution. . . but how do actions form?

Motivation: how does an opinion become an action?

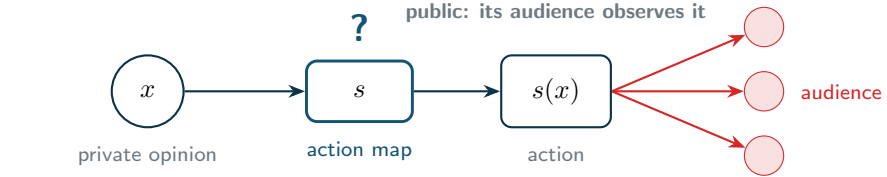


Motivation: how does an opinion become an action?



Same action, different opinions

Motivation: how does an opinion become an action?



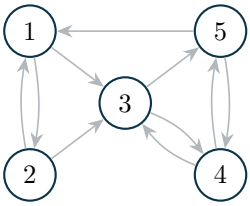


Same action, different opinions

Main question: how is the action map chosen for agents?

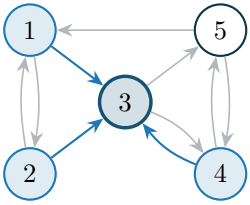
Communication graph

- **Network**: N agents in $\mathcal{V} = \{1, \ldots, N\}$, weighted digraph $\mathcal{G}$.
  Edge $j \to i$: j communicates to i.



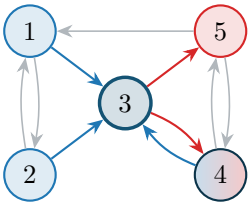
Communication graph

- **Network**: N agents in $\mathcal{V} = \{1, \dots, N\}$, weighted digraph $\mathcal{G}$.
  Edge $j \to i$: j communicates to i.
- **Sources** (in-neighborhood): $\mathcal{N}_i^{\mathrm{S}} = \{j : j \to i\}$
  e.g. $\mathcal{N}_3^{\mathrm{S}} = \{1, 2, 4\}$



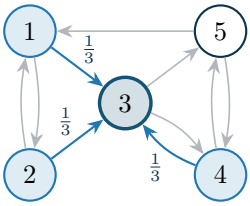
Communication graph

- **Network**: N agents in $\mathcal{V} = \{1, \ldots, N\}$, weighted digraph $\mathcal{G}$.
  Edge $j \to i$: j communicates to i.
- **Sources** (in-neighborhood): $\mathcal{N}_i^{\mathrm{S}} = \{j : j \to i\}$
  e.g. $\mathcal{N}_3^{\mathrm{S}} = \{1, 2, 4\}$
- **Audience** (out-neighborhood): $\mathcal{N}_i^{\mathrm{A}} = \{k : i \to k\}$
  e.g. $\mathcal{N}_3^{\mathrm{A}} = \{4, 5\}$



Communication graph

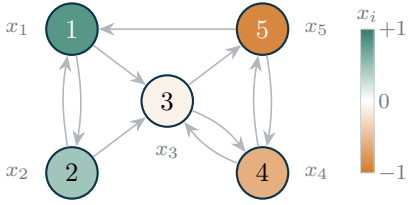
- **Network**: N agents in $\mathcal{V} = \{1, \dots, N\}$, weighted digraph $\mathcal{G}$.
  Edge $j \to i$: j communicates to i.
- **Sources** (in-neighborhood): $\mathcal{N}_i^{\mathrm{S}} = \{j : j \to i\}$
  e.g. $\mathcal{N}_3^{\mathrm{S}} = \{1, 2, 4\}$
- **Audience** (out-neighborhood): $\mathcal{N}_i^{\mathrm{A}} = \{k : i \to k\}$
  e.g. $\mathcal{N}_3^{\mathrm{A}} = \{4, 5\}$
- **Weight matrix $\boldsymbol{W}$**: row-stochastic, $W_{ii} = 0$.
  row i: how i weighs its sources; here $W_{ij} = 1/|\mathcal{N}_i^{\mathrm{S}}|$



$$\boldsymbol{W} = \begin{pmatrix} 0 & \frac{1}{2} & 0 & 0 & \frac{1}{2} \\ 1 & 0 & 0 & 0 & 0 \\ \frac{1}{3} & \frac{1}{3} & 0 & \frac{1}{3} & 0 \\ 0 & 0 & \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & 0 & \frac{1}{2} & \frac{1}{2} & 0 \end{pmatrix}$$

Communication graph

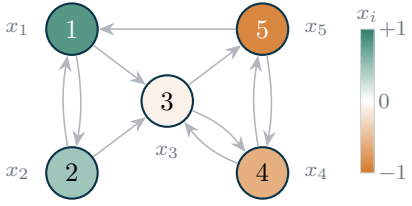
- **Network**: N agents in $\mathcal{V} = \{1, \dots, N\}$, weighted digraph $\mathcal{G}$.
  Edge $j \to i$: j communicates to i.
- **Sources** (in-neighborhood): $\mathcal{N}_i^{\mathrm{S}} = \{j : j \to i\}$
  e.g. $\mathcal{N}_3^{\mathrm{S}} = \{1, 2, 4\}$
- **Audience** (out-neighborhood): $\mathcal{N}_i^{\mathrm{A}} = \{k : i \to k\}$
  e.g. $\mathcal{N}_3^{\mathrm{A}} = \{4, 5\}$
- **Weight matrix** $\boldsymbol{W}$: row-stochastic, $W_{ii} = 0$.
  row i: how i weighs its sources; here $W_{ij} = 1/|\mathcal{N}_i^{\mathrm{S}}|$
- **Opinions**: each agent holds a private opinion $x_i \in [-1, 1]$.



$$\boldsymbol{W} = \begin{pmatrix} 0 & \frac{1}{2} & 0 & 0 & \frac{1}{2} \\ 1 & 0 & 0 & 0 & 0 \\ \frac{1}{3} & \frac{1}{3} & 0 & \frac{1}{3} & 0 \\ 0 & 0 & \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & 0 & \frac{1}{2} & \frac{1}{2} & 0 \end{pmatrix}$$

Communication graph

- **Network**: N agents in $\mathcal{V} = \{1, \dots, N\}$, weighted digraph $\mathcal{G}$.
  Edge $j \to i$: j communicates to i.
- **Sources** (in-neighborhood): $\mathcal{N}_i^{\mathrm{S}} = \{j : j \to i\}$
  e.g. $\mathcal{N}_3^{\mathrm{S}} = \{1, 2, 4\}$
- **Audience** (out-neighborhood): $\mathcal{N}_i^{\mathrm{A}} = \{k : i \to k\}$
  e.g. $\mathcal{N}_3^{\mathrm{A}} = \{4, 5\}$
- **Weight matrix $\boldsymbol{W}$**: row-stochastic, $W_{ii} = 0$.
  row i: how i weighs its sources; here $W_{ij} = 1/|\mathcal{N}_i^{\mathrm{S}}|$
- **Opinions**: each agent holds a private opinion $x_i \in [-1, 1]$.



$$\boldsymbol{W} = \begin{pmatrix} 0 & \frac{1}{2} & 0 & 0 & \frac{1}{2} \\ 1 & 0 & 0 & 0 & 0 \\ \frac{1}{3} & \frac{1}{3} & 0 & \frac{1}{3} & 0 \\ 0 & 0 & \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & 0 & \frac{1}{2} & \frac{1}{2} & 0 \end{pmatrix}$$

→ *Given the graph and the opinions: how do opinions evolve?*

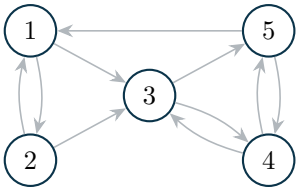
Dynamics of opinion evolution

Standard: averaging DeGroot 1974

$x_i(t+1) = \sum_j W_{ij}\, x_j(t)$

Simplified prediction: connected graph $\Rightarrow$ consensus

Richer models: (non-exhaustive)



every agent averages its sources

Dynamics of opinion evolution

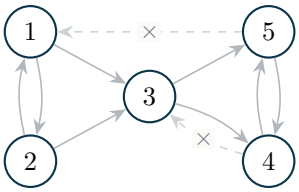
Standard: averaging DeGroot 1974

$x_i(t+1) = \sum_j W_{ij}\, x_j(t)$

Simplified prediction: connected graph $\Rightarrow$ consensus

Richer models: (non-exhaustive)

1. **Bounded confidence**: cut far edges
   Deffuant et al. 2000; Hegselmann & Krause 2002



far-apart opinions: the edge is cut

Dynamics of opinion evolution

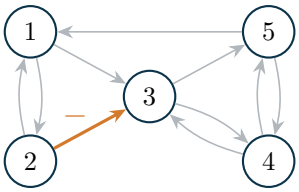
Standard: averaging DeGroot 1974

$x_i(t+1) = \sum_j W_{ij}\, x_j(t)$

Simplified prediction: connected graph $\Rightarrow$ consensus

Richer models: (non-exhaustive)

1. **Bounded confidence**: cut far edges
   Deffuant et al. 2000; Hegselmann & Krause 2002
2. **Antagonistic interactions**: negative edge weights ($W_{ij} < 0$)
   Altafini 2013



an antagonistic interaction: $W_{32} < 0$

Dynamics of opinion evolution

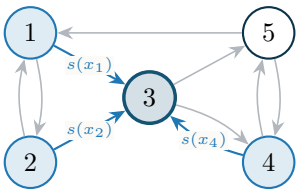
Standard: averaging DeGroot 1974

$x_i(t+1) = \sum_j W_{ij}\, x_j(t)$

Simplified prediction: connected graph $\Rightarrow$ consensus

Richer models: (non-exhaustive)

1. **Bounded confidence**: cut far edges
   Deffuant et al. 2000; Hegselmann & Krause 2002
2. **Antagonistic interactions**: negative edge weights ($W_{ij} < 0$)
   Altafini 2013
3. **Action maps**: share $s(x)$, not x
   Martins 2008; Bizyaeva et al. 2023; Couthures et al. 2025



3 observes the actions $s(x_1), s(x_2), s(x_4)$

Dynamics of opinion evolution

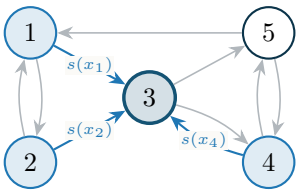
Standard: averaging DeGroot 1974

$x_i(t+1) = \sum_j W_{ij}\, x_j(t)$

Simplified prediction: connected graph $\Rightarrow$ consensus

Richer models: (non-exhaustive)

1. **Bounded confidence**: cut far edges
   Deffuant et al. 2000; Hegselmann & Krause 2002
2. **Antagonistic interactions**: negative edge weights ($W_{ij} < 0$)
   Altafini 2013
3. **Action maps**: share $s(x)$, not x
   Martins 2008; Bizyaeva et al. 2023; Couthures et al. 2025



3 observes the actions $s(x_1), s(x_2), s(x_4)$

... and many more, e.g.

- absolutely stubborn agents, $W_{ii} = 1$ (Acemoglu et al. 2013)
- coevolution of opinions and actions (Aghbolagh et al. 2023)

Dynamics of opinion evolution

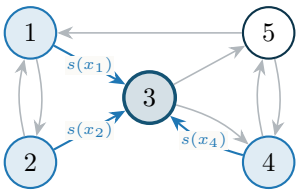
Standard: averaging DeGroot 1974

$x_i(t+1) = \sum_j W_{ij}\, x_j(t)$

Simplified prediction: connected graph $\Rightarrow$ consensus

Richer models: (non-exhaustive)

1. **Bounded confidence**: cut far edges
   Deffuant et al. 2000; Hegselmann & Krause 2002
2. **Antagonistic interactions**: negative edge weights ($W_{ij} < 0$)
   Altafini 2013
3. **Action maps**: share $s(x)$, not x
   Martins 2008; Bizyaeva et al. 2023; Couthures et al. 2025



3 observes the actions $s(x_1), s(x_2), s(x_4)$

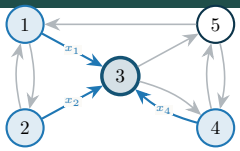
... and many more, e.g.

- absolutely stubborn agents, $W_{ii} = 1$ (Acemoglu et al. 2013)
- coevolution of opinions and actions (Aghbolagh et al. 2023)

Our focus, route 3: fixed graph, positive weights; the map s shapes what is exchanged.

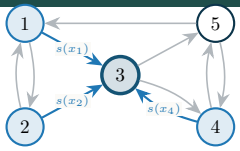
Opinion dynamics with an action map

Dynamics: $x_i(t+1) = \sum_{j \in \mathcal{N}_i^{\mathrm{S}}} W_{ij} \underbrace{x_j(t)}_{\text{what enters } i\text{'s update from } j}$



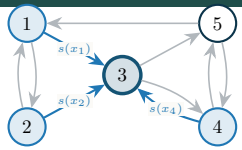
Opinion dynamics with an action map

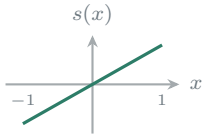
Dynamics: $x_i(t+1) = \sum_{j \in \mathcal{N}_i^{\mathrm{S}}} W_{ij} \underbrace{s\big(x_j(t)\big)}_{\text{what enters } i\text{'s update from } j}$



Opinion dynamics with an action map

Dynamics: $x_i(t+1) = \sum_{j \in \mathcal{N}_i^{\mathrm{S}}} W_{ij} \underbrace{s\big(x_j(t)\big)}_{\text{what enters } i\text{'s update from } j}$





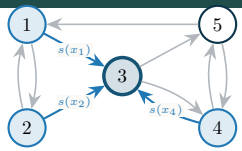
$s(x) = x$

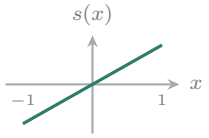
share the opinion as it is

DeGroot 1974

Opinion dynamics with an action map

Dynamics: $x_i(t+1) = \sum_{j \in \mathcal{N}_i^{\mathrm{S}}} W_{ij} \underbrace{s\big(x_j(t)\big)}_{\text{what enters } i\text{'s update from } j}$

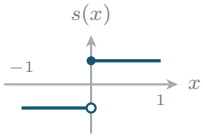




$s(x) = x$

share the opinion as it is

DeGroot 1974



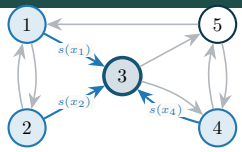
$s(x) = \mathrm{sign}(x)$

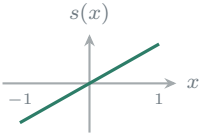
continuous opinions, discrete actions (CODA)

Martins 2008; Chowdhury et al. 2016; Ceragioli & Frasca 2018

Opinion dynamics with an action map

Dynamics: $x_i(t+1) = \sum_{j \in \mathcal{N}_i^{\mathrm{S}}} W_{ij} \underbrace{s\big(x_j(t)\big)}_{\text{what enters } i\text{'s update from } j}$

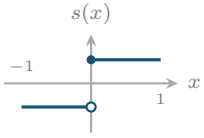




$s(x) = x$

share the opinion as it is

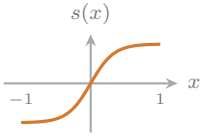
DeGroot 1974



$s(x) = \mathrm{sign}(x)$

continuous opinions, discrete actions (CODA)

Martins 2008; Chowdhury et al. 2016; Ceragioli & Frasca 2018



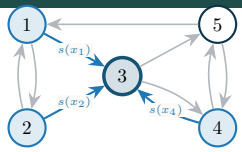
$s(x) = \tanh(kx)$

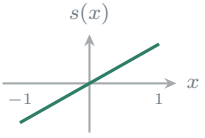
bounded, sigmoidal exchange

Bizyaeva, Franci & Leonard 2023; Couthures et al. 2025

Opinion dynamics with an action map

Dynamics: $x_i(t+1) = \sum_{j \in \mathcal{N}_i^{\mathrm{S}}} W_{ij}\, \underbrace{s\big(x_j(t)\big)}_{\text{what enters } i\text{'s update from } j}$

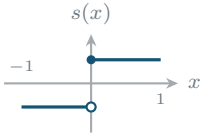




$s(x) = x$

share the opinion as it is

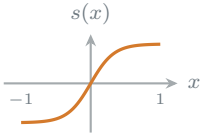
DeGroot 1974



$s(x) = \mathrm{sign}(x)$

continuous opinions, discrete actions (CODA)

Martins 2008; Chowdhury et al. 2016; Ceragioli & Frasca 2018



$s(x) = \tanh(kx)$

bounded, sigmoidal exchange

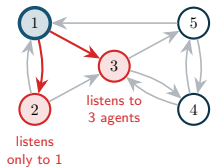
Bizyaeva, Franci & Leonard 2023; Couthures et al. 2025

Typically assumed: s is **fixed in advance**, and **common** to all agents.

Who chooses the action map?

Typically assumed: s is ~~**fixed in advance**~~, and ~~**common** to all agents~~.

But agents sit at **different places** in the network:

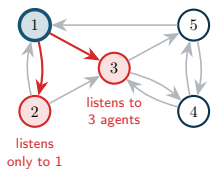


audience of agent 1

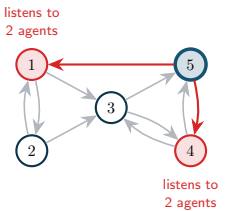
Who chooses the action map?

Typically assumed: s is ~~fixed in advance~~, and ~~common to all agents~~.

But agents sit at **different places** in the network:



audience of agent 1

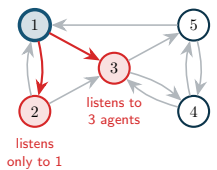


audience of agent 5

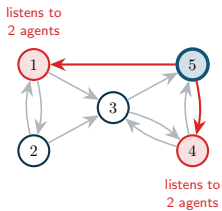
Who chooses the action map?

Typically assumed: s is ~~fixed in advance~~, and ~~common to all agents~~.

But agents sit at **different places** in the network:



audience of agent 1



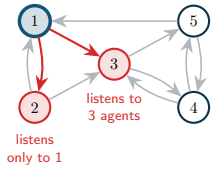
audience of agent 5

Agents 1 and 5 face different audiences. Should they act the same way?

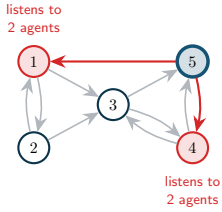
Who chooses the action map?

Typically assumed: s is ~~fixed in advance~~, and ~~common to all agents~~.

But agents sit at **different places** in the network:



audience of agent 1



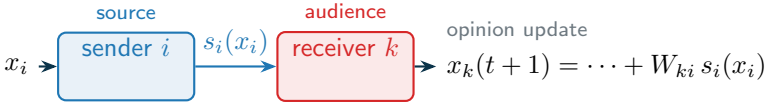
audience of agent 5

Agents 1 and 5 face different audiences. Should they act the same way?

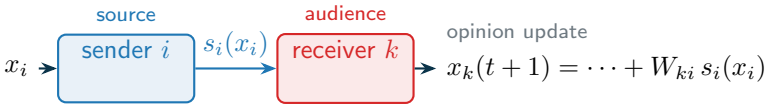
In this work, we look for a principle that motivates each agent's map:

$$s(x_j) \longrightarrow s_j(x_j).$$

Our approach: a strategic communication game

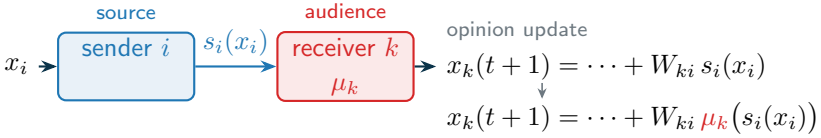


Our approach: a strategic communication game



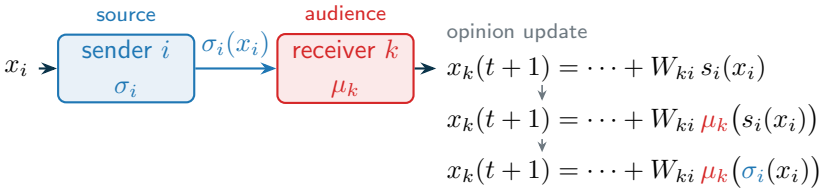
Hidden assumption: the receiver takes $s_i(x_i)$ into its update as it is.

Our approach: a strategic communication game



- **Decoder** μ_k: how k estimates i's opinion from what it observes.

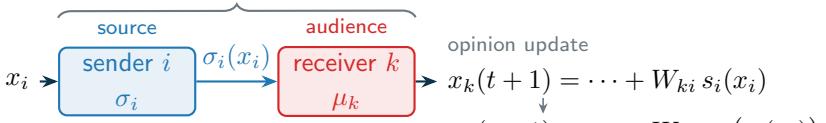
Our approach: a strategic communication game



- **Decoder** μ_k: how k estimates i's opinion from what it observes.
- **Encoder** σ_i: how i turns its opinion into an action.

Our approach: a strategic communication game

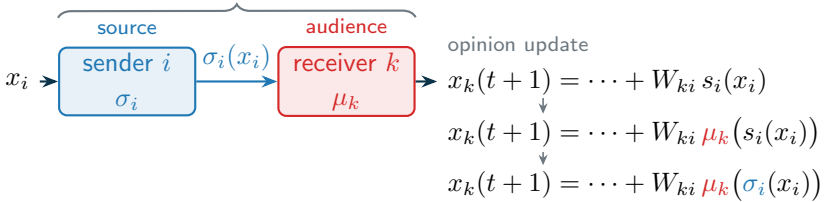
$s_i = \mu_k \circ \sigma_i$: **effective signaling function**



- **Decoder** μ_k: how k estimates i's opinion from what it observes.
- **Encoder** σ_i: how i turns its opinion into an action.
- **Effective signaling function** $s_i = \mu_k \circ \sigma_i$: what actually enters the dynamics.

Our approach: a strategic communication game

$s_i = \mu_k \circ \sigma_i$: **effective signaling function**

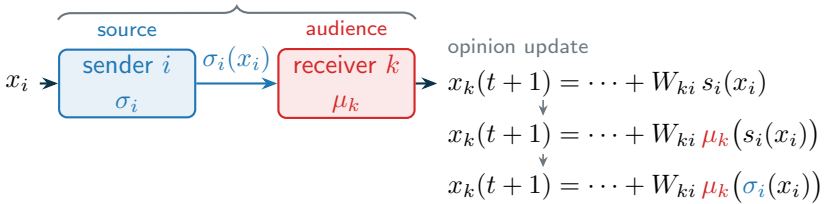


- **Decoder** μ_k: how k estimates i's opinion from what it observes.
- **Encoder** σ_i: how i turns its opinion into an action.
- **Effective signaling function** $s_i = \mu_k \circ \sigma_i$: what actually enters the dynamics.

$s_i = \mu_k \circ \sigma_i$ **is an equilibrium of a game.**

Our approach: a strategic communication game

$s_i = \mu_k \circ \sigma_i$: **effective signaling function**

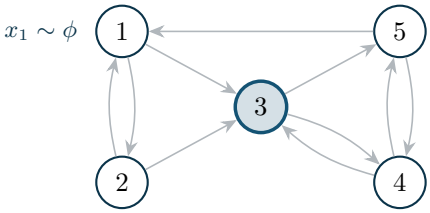


- **Decoder** μ_k: how k estimates i's opinion from what it observes.
- **Encoder** σ_i: how i turns its opinion into an action.
- **Effective signaling function** $s_i = \mu_k \circ \sigma_i$: what actually enters the dynamics.

$s_i = \mu_k \circ \sigma_i$ **is an equilibrium of a game.**

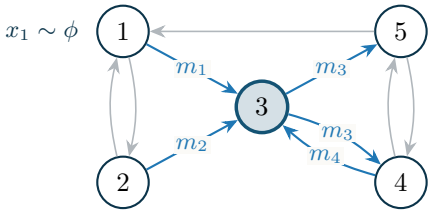
→ *Let's formalize the game played by the agents.*

One round of communication



- Opinions $x_i \sim \phi$ i.i.d.; the prior ϕ is common knowledge

One round of communication

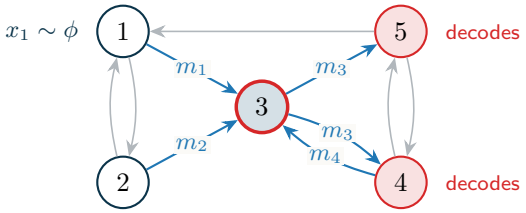


- Opinions $x_i \sim \phi$ i.i.d.; the prior ϕ is common knowledge

1 · Encode: $m_i = \sigma_i(x_i) \in [-1, 1]$

one public message for the whole audience

One round of communication



- Opinions $x_i \sim \phi$ i.i.d.; the prior ϕ is common knowledge

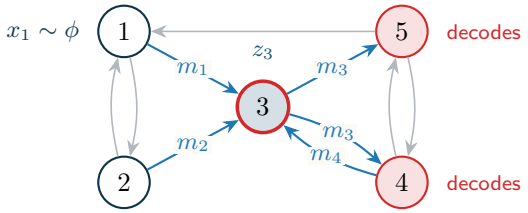
1 · Encode: $m_i = \sigma_i(x_i) \in [-1, 1]$

one public message for the whole audience

2 · Decode: $\widehat{x}_j = \mu_j(m_j)$

opinions are independent, so all receivers of j form the same estimate

One round of communication



- Opinions $x_i \sim \phi$ i.i.d.; the prior ϕ is common knowledge

1 · Encode: $m_i = \sigma_i(x_i) \in [-1, 1]$

one public message for the whole audience

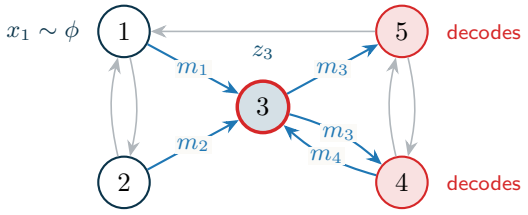
2 · Decode: $\widehat{x}_j = \mu_j(m_j)$

opinions are independent, so all receivers of j form the same estimate

3 · Update: $z_i = \zeta_i\big(x_i, \widehat{\boldsymbol{x}}\big) = x_i(t+1)$

the new opinion from the current one and the estimates

One round of communication



- Opinions $x_i \sim \phi$ i.i.d.; the prior ϕ is common knowledge

1 · Encode: $m_i = \sigma_i(x_i) \in [-1, 1]$

one public message for the whole audience

2 · Decode: $\widehat{x}_j = \mu_j(m_j)$

opinions are independent, so all receivers of j form the same estimate

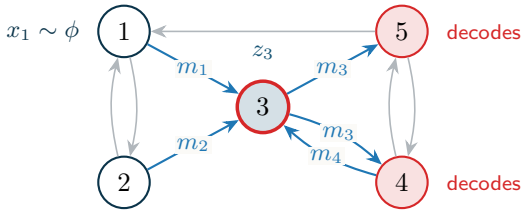
3 · Update: $z_i = \zeta_i\big(x_i, \widehat{\boldsymbol{x}}\big) = x_i(t+1)$

the new opinion from the current one and the estimates

- **Cheap talk**[1]: model of strategic communication
- Messages are costless, non-binding and unverifiable

[1] Crawford & Sobel (1982), Strategic information transmission, *Econometrica*.

One round of communication



- Opinions $x_i \sim \phi$ i.i.d.; the prior ϕ is common knowledge

1 · Encode: $m_i = \sigma_i(x_i) \in [-1, 1]$

one public message for the whole audience

2 · Decode: $\widehat{x}_j = \mu_j(m_j)$

opinions are independent, so all receivers of j form the same estimate

3 · Update: $z_i = \zeta_i\big(x_i, \widehat{\boldsymbol{x}}\big) = x_i(t+1)$

the new opinion from the current one and the estimates

- **Cheap talk**[1]: model of strategic communication
- Messages are costless, non-binding and unverifiable

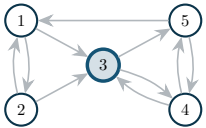
→ *Each agent chooses what to say, how to estimate, and how to update. What is the goal of each agent?*

[1] Crawford & Sobel (1982), Strategic information transmission, *Econometrica*.

Agent's cost function: three drives

Reminder: z_i is the opinion update, $z_i = x_i(t+1)$

$$c_i(\boldsymbol{z}, \boldsymbol{x}) =$$

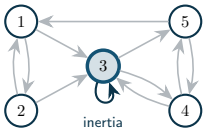


Agent's cost function: three drives

Reminder: z_i is the opinion update, $z_i = x_i(t+1)$

$$c_i(\boldsymbol{z}, \boldsymbol{x}) = \underbrace{\alpha_i \left(z_i - x_i\right)^2}_{\textbf{inertia}}$$

- $\alpha_i \in [0, 1)$: **inertia** (self-persistence)

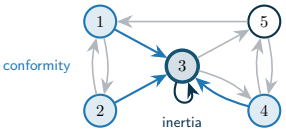


Agent's cost function: three drives

Reminder: z_i is the opinion update, $z_i = x_i(t+1)$

$$c_i(\boldsymbol{z}, \boldsymbol{x}) = \underbrace{\alpha_i \, (z_i - x_i)^2}_{\textbf{inertia}} + \underbrace{(1-\alpha_i)\Big(z_i - \sum_{j \in \mathcal{N}_i^{\mathrm{S}}} W_{ij} \, x_j\Big)^2}_{\textbf{social conformity}}$$

- $\alpha_i \in [0, 1)$: **inertia** (self-persistence)
- **conformity**: pull toward the sources (peer pressure)

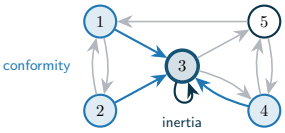


Agent's cost function: three drives

Reminder: z_i is the opinion update, $z_i = x_i(t+1)$

$$c_i(\boldsymbol{z}, \boldsymbol{x}) = \underbrace{\underbrace{\alpha_i \, (z_i - x_i)^2}_{\textbf{inertia}} + \underbrace{(1-\alpha_i)\Big(z_i - \sum_{j \in \mathcal{N}_i^{\mathrm{S}}} W_{ij} \, x_j\Big)^2}_{\textbf{social conformity}}}_{\text{(receiver)}}$$

- $\alpha_i \in [0, 1)$: **inertia** (self-persistence)
- **conformity**: pull toward the sources (peer pressure)



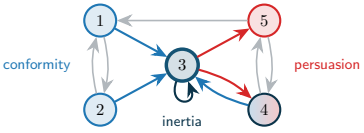
Agent's cost function: three drives

Reminder: z_i is the opinion update, $z_i = x_i(t+1)$

$$c_i(\boldsymbol{z}, \boldsymbol{x}) = \underbrace{\underbrace{\alpha_i \left(z_i - x_i\right)^2}_{\text{inertia}} + \underbrace{(1-\alpha_i)\Big(z_i - \sum_{j \in \mathcal{N}_i^{\mathrm{S}}} W_{ij}\, x_j\Big)^2}_{\text{social conformity}}}_{\text{(receiver)}} + \underbrace{\beta_i \sum_{k \in \mathcal{N}_i^{\mathrm{A}}} \left(z_k - x_i\right)^2}_{\text{persuasion}}$$

(sender)

- $\alpha_i \in [0, 1)$: **inertia** (self-persistence)
- **conformity**: pull toward the sources (peer pressure)
- $\beta_i \in \{0, 1\}$: **persuasion parameter**, whether the agent wants to persuade its audience



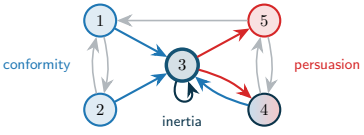
Agent's cost function: three drives

Reminder: z_i is the opinion update, $z_i = x_i(t+1)$

$$c_i(\boldsymbol{z}, \boldsymbol{x}) = \underbrace{\underbrace{\alpha_i \left(z_i - x_i\right)^2}_{\text{inertia}} + \underbrace{(1-\alpha_i)\Big(z_i - \sum_{j \in \mathcal{N}_i^{\mathrm{S}}} W_{ij}\, x_j\Big)^2}_{\text{social conformity}}}_{\text{(receiver)}} + \underbrace{\beta_i \sum_{k \in \mathcal{N}_i^{\mathrm{A}}} \left(z_k - x_i\right)^2}_{\text{persuasion}}$$

(sender)

- $\alpha_i \in [0, 1)$: **inertia** (self-persistence)
- **conformity**: pull toward the sources (peer pressure)
- $\beta_i \in \{0, 1\}$: **persuasion parameter**, whether the agent wants to persuade its audience
- $\beta_i = 0$: **non-persuasive sender**; $\beta_i = 1$: **persuasive sender**



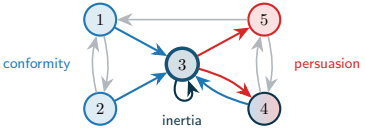
Agent's cost function: three drives

Reminder: z_i is the opinion update, $z_i = x_i(t+1)$

$$J_i(\boldsymbol{\sigma}, \boldsymbol{\mu}, \boldsymbol{\zeta}) = \mathbb{E}\Big[\underbrace{\alpha_i \,(z_i - x_i)^2}_{\text{inertia}} + \underbrace{(1-\alpha_i)\Big(z_i - \sum_{j \in \mathcal{N}_i^{\mathrm{S}}} W_{ij}\, x_j\Big)^2}_{\text{social conformity}} + \underbrace{\beta_i \sum_{k \in \mathcal{N}_i^{\mathrm{A}}} (z_k - x_i)^2}_{\text{persuasion}}\Big]$$

(receiver) (sender)

- $\alpha_i \in [0,1)$: **inertia** (self-persistence)
- **conformity**: pull toward the sources (peer pressure)
- $\beta_i \in \{0,1\}$: **persuasion parameter**, whether the agent wants to persuade its audience
- $\beta_i = 0$: **non-persuasive sender**; $\beta_i = 1$: **persuasive sender**



The strategic communication game

A stage game of incomplete information

- **Players:** the agents $i \in \mathcal{V}$, on the graph $\mathcal{G}$.

The strategic communication game

A stage game of incomplete information

- **Players:** the agents $i \in \mathcal{V}$, on the graph $\mathcal{G}$.
- **Private states:** $x_i \sim \phi$ i.i.d. on $[-1, 1]$; ϕ common knowledge.

The strategic communication game

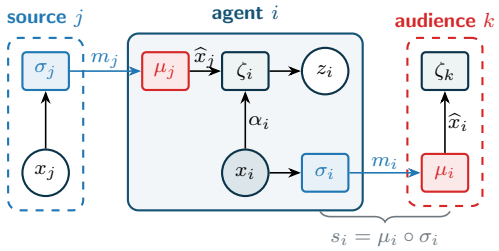
A stage game of incomplete information

- **Players:** the agents $i \in \mathcal{V}$, on the graph $\mathcal{G}$.
- **Private states:** $x_i \sim \phi$ i.i.d. on $[-1, 1]$; ϕ common knowledge.
- **Strategies** (pure):

  encoder $\sigma_i : [-1, 1] \to \mathcal{M} = [-1, 1]$

  decoder $\mu_j : \mathcal{M} \to [-1, 1]$, for every source j

  opinion update rule $\zeta_i : \left(x_i, \widehat{\boldsymbol{x}}\right) \mapsto z_i = x_i(t+1)$

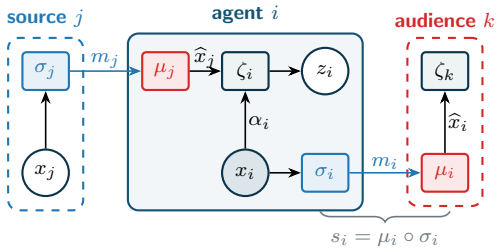


top row: i as a receiver · bottom row: i as a sender

The strategic communication game

A stage game of incomplete information

- **Players:** the agents $i \in \mathcal{V}$, on the graph $\mathcal{G}$.
- **Private states:** $x_i \sim \phi$ i.i.d. on $[-1, 1]$; ϕ common knowledge.
- **Strategies** (pure):
  encoder $\sigma_i : [-1, 1] \to \mathcal{M} = [-1, 1]$
  decoder $\mu_j : \mathcal{M} \to [-1, 1]$, for every source j
  opinion update rule $\zeta_i : \big(x_i, \widehat{\boldsymbol{x}}\big) \mapsto z_i = x_i(t+1)$
- **Costs:** $J_i(\boldsymbol{\sigma}, \boldsymbol{\mu}, \boldsymbol{\zeta}) = \mathbb{E}_{\boldsymbol{x}}\big[c_i(\boldsymbol{z}, \boldsymbol{x})\big]$.

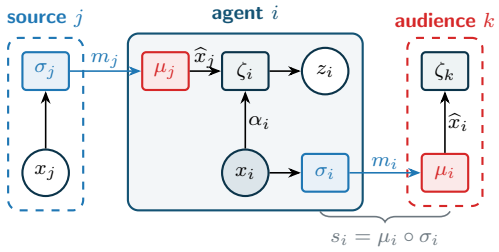


top row: i as a receiver · bottom row: i as a sender

The strategic communication game

A stage game of incomplete information

- **Players:** the agents $i \in \mathcal{V}$, on the graph $\mathcal{G}$.
- **Private states:** $x_i \sim \phi$ i.i.d. on $[-1, 1]$; ϕ common knowledge.
- **Strategies** (pure):
  encoder $\sigma_i : [-1, 1] \to \mathcal{M} = [-1, 1]$
  decoder $\mu_j : \mathcal{M} \to [-1, 1]$, for every source j
  opinion update rule $\zeta_i : \big(x_i, \widehat{\boldsymbol{x}}\big) \mapsto z_i = x_i(t+1)$
- **Costs:** $J_i(\boldsymbol{\sigma}, \boldsymbol{\mu}, \boldsymbol{\zeta}) = \mathbb{E}_{\boldsymbol{x}}\big[c_i(\boldsymbol{z}, \boldsymbol{x})\big]$.
- **Timing:** encode, decode, update; simultaneous moves in each phase.

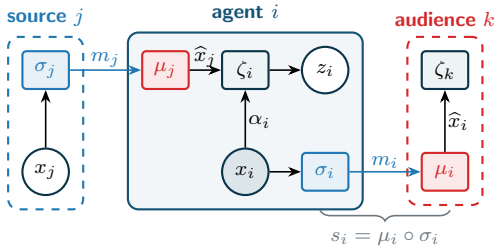


top row: i as a receiver · bottom row: i as a sender

The strategic communication game

A stage game of incomplete information

- **Players:** the agents $i \in \mathcal{V}$, on the graph $\mathcal{G}$.
- **Private states:** $x_i \sim \phi$ i.i.d. on $[-1, 1]$; ϕ common knowledge.
- **Strategies** (pure):
  encoder $\sigma_i : [-1, 1] \to \mathcal{M} = [-1, 1]$
  decoder $\mu_j : \mathcal{M} \to [-1, 1]$, for every source j
  opinion update rule $\zeta_i : \big(x_i, \widehat{\boldsymbol{x}}\big) \mapsto z_i = x_i(t+1)$
- **Costs:** $J_i(\boldsymbol{\sigma}, \boldsymbol{\mu}, \boldsymbol{\zeta}) = \mathbb{E}_{\boldsymbol{x}}\big[c_i(\boldsymbol{z}, \boldsymbol{x})\big]$.
- **Timing:** encode, decode, update; simultaneous moves in each phase.



top row: i as a receiver · bottom row: i as a sender

→ *Solution concept: what does it mean for the agents to play well?*

Solution concept: perfect Bayesian equilibrium

Definition (PBE). $(\boldsymbol{\sigma}^*, \boldsymbol{\mu}^*, \boldsymbol{\zeta}^*)$ is a PBE if, for every agent i: (a refinement of Nash equilibrium)

1. **Best-response encoding:** send the best message, given the estimate it will induce.
   for every encoder σ_i, $J_i\big(\sigma_i^*, \boldsymbol{\sigma}_{-i}^*, \boldsymbol{\mu}^*, \boldsymbol{\zeta}^*\big) \leq J_i\big(\sigma_i, \boldsymbol{\sigma}_{-i}^*, \boldsymbol{\mu}^*, \boldsymbol{\zeta}^*\big)$.

Solution concept: perfect Bayesian equilibrium

Definition (PBE). $(\boldsymbol{\sigma}^*, \boldsymbol{\mu}^*, \boldsymbol{\zeta}^*)$ is a PBE if, for every agent i: (a refinement of Nash equilibrium)

1. **Best-response encoding:** send the best message, given the estimate it will induce.
   for every encoder σ_i, $J_i\big(\sigma_i^*, \boldsymbol{\sigma}_{-i}^*, \boldsymbol{\mu}^*, \boldsymbol{\zeta}^*\big) \le J_i\big(\sigma_i, \boldsymbol{\sigma}_{-i}^*, \boldsymbol{\mu}^*, \boldsymbol{\zeta}^*\big)$.
2. **Bayes consistency:** update the belief on x_i using Bayes' rule, and estimate the sender's opinion: $\mu_i^*(m) = \mathbb{E}\big[x_i \mid \sigma_i^*(x_i) = m\big]$.

Solution concept: perfect Bayesian equilibrium

Definition (PBE). $(\boldsymbol{\sigma}^*, \boldsymbol{\mu}^*, \boldsymbol{\zeta}^*)$ is a PBE if, for every agent i: (a refinement of Nash equilibrium)

1. **Best-response encoding:** send the best message, given the estimate it will induce.
   for every encoder σ_i, $J_i\big(\sigma_i^*, \boldsymbol{\sigma}_{-i}^*, \boldsymbol{\mu}^*, \boldsymbol{\zeta}^*\big) \le J_i\big(\sigma_i, \boldsymbol{\sigma}_{-i}^*, \boldsymbol{\mu}^*, \boldsymbol{\zeta}^*\big)$.
2. **Bayes consistency:** update the belief on x_i using Bayes' rule, and estimate the sender's opinion:
   $\mu_i^*(m) = \mathbb{E}\big[x_i \mid \sigma_i^*(x_i) = m\big]$.
3. **Best-response update:** update the opinion given the estimates of its sources.
   $\zeta_i^*(x_i, \widehat{\boldsymbol{x}}) \in \arg\min_{z_i} \mathbb{E}\big[c_i \mid x_i, \widehat{\boldsymbol{x}}\big]$.

Solution concept: perfect Bayesian equilibrium

Definition (PBE). $(\boldsymbol{\sigma}^*, \boldsymbol{\mu}^*, \boldsymbol{\zeta}^*)$ is a PBE if, for every agent i: (a refinement of Nash equilibrium)

1. **Best-response encoding:** send the best message, given the estimate it will induce.
   for every encoder σ_i, $J_i\big(\sigma_i^*, \boldsymbol{\sigma}_{-i}^*, \boldsymbol{\mu}^*, \boldsymbol{\zeta}^*\big) \le J_i\big(\sigma_i, \boldsymbol{\sigma}_{-i}^*, \boldsymbol{\mu}^*, \boldsymbol{\zeta}^*\big)$.
2. **Bayes consistency:** update the belief on x_i using Bayes' rule, and estimate the sender's opinion:
   $\mu_i^*(m) = \mathbb{E}\big[x_i \mid \sigma_i^*(x_i) = m\big]$.
3. **Best-response update:** update the opinion given the estimates of its sources.
   $\zeta_i^*(x_i, \widehat{\boldsymbol{x}}) \in \arg\min_{z_i} \mathbb{E}\big[c_i \mid x_i, \widehat{\boldsymbol{x}}\big]$.

An equilibrium fixes both what is communicated and what it means.

Solution concept: perfect Bayesian equilibrium

Definition (PBE). $(\boldsymbol{\sigma}^*, \boldsymbol{\mu}^*, \boldsymbol{\zeta}^*)$ is a PBE if, for every agent i: (a refinement of Nash equilibrium)

1. **Best-response encoding:** send the best message, given the estimate it will induce.
   for every encoder σ_i, $J_i\big(\sigma_i^*, \boldsymbol{\sigma}_{-i}^*, \boldsymbol{\mu}^*, \boldsymbol{\zeta}^*\big) \le J_i\big(\sigma_i, \boldsymbol{\sigma}_{-i}^*, \boldsymbol{\mu}^*, \boldsymbol{\zeta}^*\big)$.
2. **Bayes consistency:** update the belief on x_i using Bayes' rule, and estimate the sender's opinion:
   $\mu_i^*(m) = \mathbb{E}\big[x_i \mid \sigma_i^*(x_i) = m\big]$.
3. **Best-response update:** update the opinion given the estimates of its sources.
   $\zeta_i^*(x_i, \widehat{\boldsymbol{x}}) \in \arg\min_{z_i} \mathbb{E}\big[c_i \mid x_i, \widehat{\boldsymbol{x}}\big]$.

An equilibrium fixes both what is communicated and what it means.

→ *Solve backwards: the update is chosen last, so we solve it first.*

Best-response opinion update

- Receiver's cost $c_i = \alpha_i\,(z_i - x_i)^2 + (1-\alpha_i)\Big(z_i - \sum_{j\in\mathcal{N}_i^{\mathrm{S}}} W_{ij}\,x_j\Big)^2$ x_j unknown $\rightarrow$ estimate $\mu_j(m_j)$

Proposition (optimal update). Given the decoded estimates $\mu_j(m_j)$, the unique optimal update is

$$z_i^* = \alpha_i\,x_i + (1-\alpha_i)\sum_{j\in\mathcal{N}_i^{\mathrm{S}}} W_{ij}\,\mu_j(m_j).$$

Best-response opinion update

- Receiver's cost $c_i = \alpha_i\,(z_i - x_i)^2 + (1-\alpha_i)\Big(z_i - \sum_{j\in\mathcal{N}_i^{\mathrm{S}}} W_{ij}\,x_j\Big)^2$ x_j unknown $\to$ estimate $\mu_j(m_j)$

Proposition (optimal update). Given the decoded estimates $\mu_j(m_j)$, the unique optimal update is

$$z_i^* = \alpha_i\,x_i + (1-\alpha_i)\sum_{j\in\mathcal{N}_i^{\mathrm{S}}} W_{ij}\,\mu_j(m_j).$$

- This gives the opinion update dynamic

$$x_i(t+1) = \alpha_i\,x_i(t) + (1-\alpha_i)\sum_{j\in\mathcal{N}_i^{\mathrm{S}}} W_{ij}\,\mu_j\big(m_j(t)\big).$$

Best-response opinion update

- Receiver's cost $c_i = \alpha_i\,(z_i - x_i)^2 + (1-\alpha_i)\Big(z_i - \sum_{j\in\mathcal{N}_i^{\mathrm{S}}} W_{ij}\,x_j\Big)^2$ x_j unknown $\to$ estimate $\mu_j(m_j)$

Proposition (optimal update). Given the decoded estimates $\mu_j(m_j)$, the unique optimal update is

$$z_i^* = \alpha_i\,x_i + (1-\alpha_i)\sum_{j\in\mathcal{N}_i^{\mathrm{S}}} W_{ij}\,\mu_j(m_j).$$

- This gives the opinion update dynamic

$$x_i(t+1) = \alpha_i\,x_i(t) + (1-\alpha_i)\sum_{j\in\mathcal{N}_i^{\mathrm{S}}} W_{ij}\,\mu_j\big(m_j(t)\big).$$

- Source j enters only through its estimate $\mu_j(m_j) = \mathbb{E}[x_j \mid m_j]$.

Best-response opinion update

- Receiver's cost $c_i = \alpha_i\,(z_i - x_i)^2 + (1-\alpha_i)\Big(z_i - \sum_{j\in\mathcal{N}_i^{\mathrm{S}}} W_{ij}\,x_j\Big)^2$ x_j unknown $\to$ estimate $\mu_j(m_j)$

Proposition (optimal update). Given the decoded estimates $\mu_j(m_j)$, the unique optimal update is

$$z_i^* = \alpha_i\,x_i + (1-\alpha_i)\sum_{j\in\mathcal{N}_i^{\mathrm{S}}} W_{ij}\,\mu_j(m_j).$$

- This gives the opinion update dynamic

$$x_i(t+1) = \alpha_i\,x_i(t) + (1-\alpha_i)\sum_{j\in\mathcal{N}_i^{\mathrm{S}}} W_{ij}\,\mu_j\big(m_j(t)\big).$$

- Source j enters only through its estimate $\mu_j(m_j) = \mathbb{E}[x_j \mid m_j]$.
- Next: the encoder σ_i, which fixes that estimate.

Best-response opinion update

- Receiver's cost $c_i = \alpha_i\,(z_i - x_i)^2 + (1-\alpha_i)\Big(z_i - \sum_{j\in\mathcal{N}_i^{\mathrm{S}}} W_{ij}\,x_j\Big)^2$ x_j unknown $\to$ estimate $\mu_j(m_j)$

Proposition (optimal update). Given the decoded estimates $\mu_j(m_j)$, the unique optimal update is

$$z_i^* = \alpha_i\,x_i + (1-\alpha_i)\sum_{j\in\mathcal{N}_i^{\mathrm{S}}} W_{ij}\,\mu_j(m_j).$$

- This gives the opinion update dynamic

$$x_i(t+1) = \alpha_i\,x_i(t) + (1-\alpha_i)\sum_{j\in\mathcal{N}_i^{\mathrm{S}}} W_{ij}\,\mu_j\big(m_j(t)\big).$$

- Source j enters only through its estimate $\mu_j(m_j) = \mathbb{E}[x_j \mid m_j]$.
- Next: the encoder σ_i, which fixes that estimate.

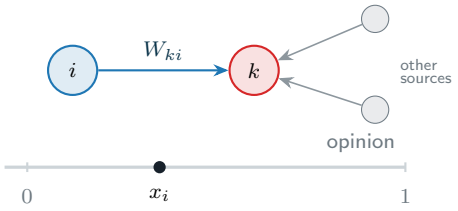
What estimate does i want its audience to form? We zoom in on one edge $i \to k$.

Zooming in: one edge interaction $i \to k$

1- Receiver k: separate the terms depending on m_i

$$z_k = \alpha\, x_k + \widetilde{W}_{ki}\, \widehat{x}_i + (1-\alpha) \sum_{j \neq i} W_{kj}\, \widehat{x}_j$$

$\widehat{x}_i = \mu_i(m_i)$; $\widetilde{W}_{ki} = (1-\alpha)\, W_{ki}$: the **susceptibility** of k to i



Zooming in: one edge interaction $i \to k$

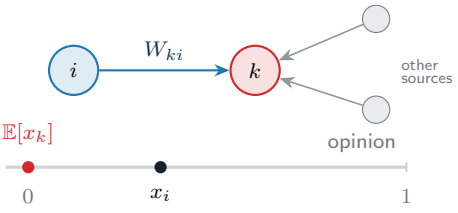
1- Receiver k: separate the terms depending on m_i

$$z_k = \alpha\, x_k + \widetilde{W}_{ki}\, \widehat{x}_i + (1-\alpha) \sum_{j \neq i} W_{kj}\, \widehat{x}_j$$

$\widehat{x}_i = \mu_i(m_i)$; $\widetilde{W}_{ki} = (1-\alpha)\, W_{ki}$: the **susceptibility** of k to i

2- Sender i: reminder: persuasion cost $\beta\,(z_k - x_i)^2$

$$\mathbb{E}\big[(z_k - x_i)^2 \mid x_i\big] = \big(\widetilde{W}_{ki}\, \widehat{x}_i - x_i\big)^2 + \text{const}$$



Zooming in: one edge interaction $i \to k$

1- Receiver k: separate the terms depending on m_i

$$z_k = \alpha\, x_k + \widetilde{W}_{ki}\, \widehat{x}_i + (1-\alpha) \sum_{j \neq i} W_{kj}\, \widehat{x}_j$$

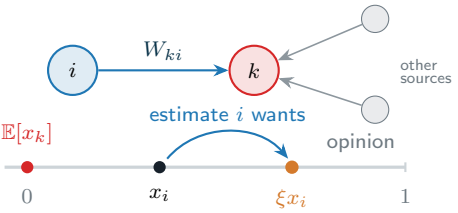
$\widehat{x}_i = \mu_i(m_i)$; $\widetilde{W}_{ki} = (1-\alpha)\, W_{ki}$: the **susceptibility** of k to i

2- Sender i: reminder: persuasion cost $\beta\,(z_k - x_i)^2$

$$\mathbb{E}\big[(z_k - x_i)^2 \mid x_i\big] = \big(\widetilde{W}_{ki}\, \widehat{x}_i - x_i\big)^2 + \text{const}$$

3- Sender's target estimate: $\widehat{x}_i = \dfrac{x_i}{\widetilde{W}_{ki}}$

We write $\widehat{x}_i = \xi\, x_i$, where $\xi = \dfrac{1}{(1-\alpha)\, W_{ki}} \geq 1$ is the **exaggeration factor**.



Zooming in: one edge interaction $i \to k$

1- Receiver k: separate the terms depending on m_i

$$z_k = \alpha\, x_k + \widetilde{W}_{ki}\, \widehat{x}_i + (1-\alpha) \sum_{j \neq i} W_{kj}\, \widehat{x}_j$$

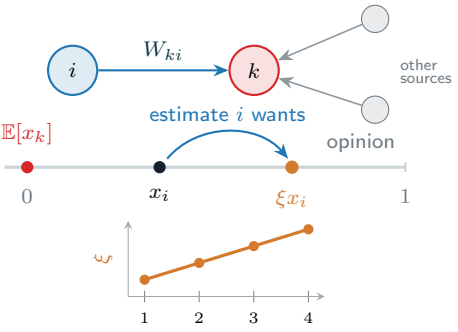
$\widehat{x}_i = \mu_i(m_i)$; $\widetilde{W}_{ki} = (1-\alpha)\, W_{ki}$: the **susceptibility** of k to i

2- Sender i: reminder: persuasion cost $\beta\,(z_k - x_i)^2$

$$\mathbb{E}\big[(z_k - x_i)^2 \mid x_i\big] = \big(\widetilde{W}_{ki}\, \widehat{x}_i - x_i\big)^2 + \text{const}$$

3- Sender's target estimate: $\widehat{x}_i = \dfrac{x_i}{\widetilde{W}_{ki}}$

We write $\widehat{x}_i = \xi\, x_i$, where $\xi = \dfrac{1}{(1-\alpha)\, W_{ki}} \geq 1$ is the **exaggeration factor**.



n sources for agent k ($W_{ki} = 1/n$, $\alpha = \frac{1}{2}$)

The smaller the weight the receiver gives the sender, the stronger the incentive to exaggerate.

One edge is a cheap-talk game

Both care about the estimate $\widehat{x}_i$, but have different preferences:

Receiver k does not know x_i

wants a **correct** estimate:

$$\min_{\widehat{x}} \mathbb{E}\left[(\widehat{x} - x_i)^2 \mid m_i\right] \;\Rightarrow\; \widehat{x}_i = \mathbb{E}[x_i \mid m_i]$$

Sender i knows x_i, chooses m_i

wants k's estimate to be ξx_i:

$$\min_{m_i} \left(\widehat{x}_i - \xi\, x_i\right)^2$$

One edge is a cheap-talk game

Both care about the estimate $\widehat{x}_i$, but have different preferences:

Receiver k does not know x_i

wants a **correct** estimate:

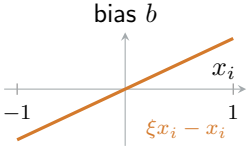
$$\min_{\widehat{x}} \mathbb{E}\big[(\widehat{x} - x_i)^2 \mid m_i\big] \;\Rightarrow\; \widehat{x}_i = \mathbb{E}[x_i \mid m_i]$$

Sender i knows x_i, chooses m_i

wants k's estimate to be ξx_i:

$$\min_{m_i} \big(\widehat{x}_i - \xi\, x_i\big)^2$$

Bias: the gap between the two targets, $b(x_i) = \xi x_i - x_i$.



One edge is a cheap-talk game

Both care about the estimate $\widehat{x}_i$, but have different preferences:

Receiver k does not know x_i

wants a **correct** estimate:

$$\min_{\widehat{x}} \mathbb{E}\big[(\widehat{x} - x_i)^2 \mid m_i\big] \;\Rightarrow\; \widehat{x}_i = \mathbb{E}[x_i \mid m_i]$$

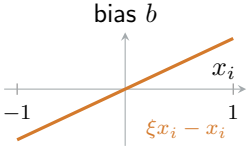
Sender i knows x_i, chooses m_i

wants k's estimate to be ξx_i:

$$\min_{m_i} \big(\widehat{x}_i - \xi\, x_i\big)^2$$

Bias: the gap between the two targets, $b(x_i) = \xi x_i - x_i$.

The **cheap-talk** setting of Crawford & Sobel (1982): informed sender, uninformed receiver, conflicting goals, costless messages. *Can reliable communication hold?*



One edge is a cheap-talk game

Both care about the estimate $\widehat{x}_i$, but have different preferences:

Receiver k does not know x_i

wants a **correct** estimate:

$$\min_{\widehat{x}} \mathbb{E}\big[(\widehat{x} - x_i)^2 \mid m_i\big] \;\Rightarrow\; \widehat{x}_i = \mathbb{E}[x_i \mid m_i]$$

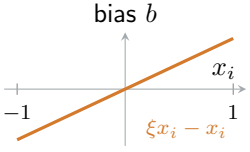
Sender i knows x_i, chooses m_i

wants k's estimate to be ξx_i:

$$\min_{m_i} \big(\widehat{x}_i - \xi\, x_i\big)^2$$

Bias: the gap between the two targets, $b(x_i) = \xi x_i - x_i$.

The **cheap-talk** setting of Crawford & Sobel (1982): informed sender, uninformed receiver, conflicting goals, costless messages. *Can reliable communication hold?*



One edge = the cheap-talk game $\Gamma(\xi)$.

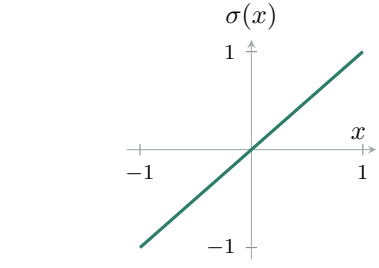
Two easy cases: a non-persuasive sender, or a naive receiver

Which encodings σ are equilibria of the sender–receiver game $\Gamma(\xi)$?

Two easy cases: a non-persuasive sender, or a naive receiver

Which encodings σ are equilibria of the sender–receiver game $\Gamma(\xi)$?

- **Non-persuasive sender**, $\beta = 0$: the sender is indifferent to the encoding, so any encoding is an equilibrium, including $\sigma(x) = x$ (DeGroot).

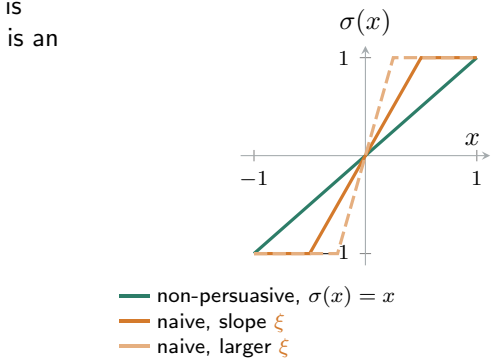


Two easy cases: a non-persuasive sender, or a naive receiver

Which encodings σ are equilibria of the sender–receiver game $\Gamma(\xi)$?

- **Non-persuasive sender**, $\beta = 0$: the sender is indifferent to the encoding, so any encoding is an equilibrium, including $\sigma(x) = x$ (DeGroot).
- **Naive receiver**, $\mu(m) = m$: the sender communicates ξx, clipped to $[-1, 1]$:

$$\sigma(x) = \mathrm{sat}(\xi\, x).$$

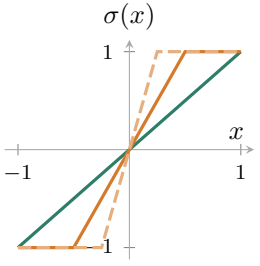


Two easy cases: a non-persuasive sender, or a naive receiver

Which encodings σ are equilibria of the sender–receiver game $\Gamma(\xi)$?

- **Non-persuasive sender**, $\beta = 0$: the sender is indifferent to the encoding, so any encoding is an equilibrium, including $\sigma(x) = x$ (DeGroot).
- **Naive receiver**, $\mu(m) = m$: the sender communicates ξx, clipped to $[-1, 1]$:

$$\sigma(x) = \mathrm{sat}(\xi\, x).$$



The saturated maps of nonlinear opinion dynamics, with slope ξ.

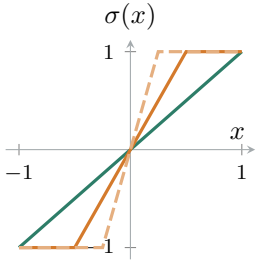
Bizyaeva, Franci & Leonard 2023; Couthures et al. 2025

Two easy cases: a non-persuasive sender, or a naive receiver

Which encodings σ are equilibria of the sender–receiver game $\Gamma(\xi)$?

- **Non-persuasive sender**, $\beta = 0$: the sender is indifferent to the encoding, so any encoding is an equilibrium, including $\sigma(x) = x$ (DeGroot).
- **Naive receiver**, $\mu(m) = m$: the sender communicates ξx, clipped to $[-1, 1]$:

$$\sigma(x) = \mathrm{sat}(\xi\, x).$$



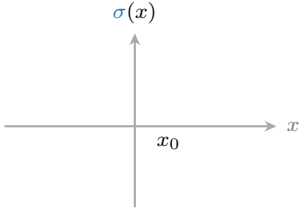
The saturated maps of nonlinear opinion dynamics, with slope ξ.

Bizyaeva, Franci & Leonard 2023; Couthures et al. 2025

$\rightarrow$ *A naive receiver gets exploited. What if the receiver is Bayesian?*

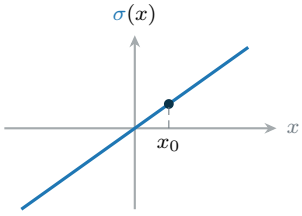
Bayesian receiver: truthful encoding is not an equilibrium

- **Bayesian** receiver: $\mu(m) = \mathbb{E}\big[x \mid \sigma(x) = m\big]$



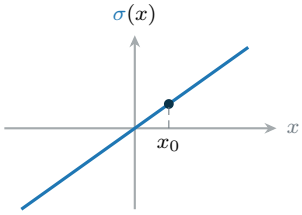
Bayesian receiver: truthful encoding is not an equilibrium

- **Bayesian** receiver: $\mu(m) = \mathbb{E}\big[x \mid \sigma(x) = m\big]$
- **Injective** σ (each message comes from one opinion)
  σ invertible $\Longrightarrow$ exact estimate: $\mu(\sigma(x)) = x$, e.g. truthful, $\sigma(x) = x$



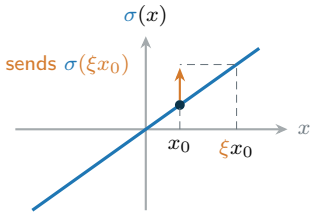
Bayesian receiver: truthful encoding is not an equilibrium

- **Bayesian** receiver: $\mu(m) = \mathbb{E}\big[x \mid \sigma(x) = m\big]$
- **Injective** σ (each message comes from one opinion)
  σ invertible $\Longrightarrow$ exact estimate: $\mu(\sigma(x)) = x$, e.g. truthful, $\sigma(x) = x$
- So for an opinion x_0, the sender can induce any estimate.



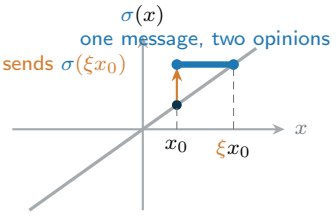
Bayesian receiver: truthful encoding is not an equilibrium

- **Bayesian** receiver: $\mu(m) = \mathbb{E}\big[x \mid \sigma(x) = m\big]$
- **Injective** σ (each message comes from one opinion)
  σ invertible $\Longrightarrow$ exact estimate: $\mu(\sigma(x)) = x$, e.g. truthful, $\sigma(x) = x$
- So for an opinion x_0, the sender can induce any estimate.
- Its preferred estimate is ξx_0, so the sender has an incentive to deviate: it sends $\sigma(\xi x_0)$ $\Rightarrow$ **the sender deviates**



Bayesian receiver: truthful encoding is not an equilibrium

- **Bayesian** receiver: $\mu(m) = \mathbb{E}\big[x \mid \sigma(x) = m\big]$
- **Injective** σ (each message comes from one opinion)
  σ invertible $\Longrightarrow$ exact estimate: $\mu(\sigma(x)) = x$, e.g. truthful, $\sigma(x) = x$
- So for an opinion x_0, the sender can induce any estimate.
- Its preferred estimate is ξx_0, so the sender has an incentive to deviate: it sends $\sigma(\xi x_0)$ $\Rightarrow$ **the sender deviates**



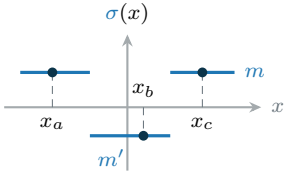
At equilibrium, the encoding is not injective.

↓

some message is induced by different opinions

Which opinions induce the same message? Intervals

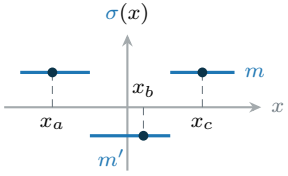
- **Single crossing of the sender's cost**
  $c(y,x) = (y - \xi x)^2$, writing $y = \widehat{x}$:
  $\dfrac{\partial^2 c}{\partial x\, \partial y} = -2\xi < 0$, so for $y < y'$, $c(y,x) - c(y',x)$ is increasing in x
  $\Rightarrow$ if the sender at x prefers a higher (lower) estimate, this holds at any higher (lower) opinion



Which opinions induce the same message? Intervals

- **Single crossing of the sender's cost**
  $c(y,x) = (y - \xi x)^2$, writing $y = \widehat{x}$:
  $\dfrac{\partial^2 c}{\partial x\, \partial y} = -2\xi < 0$, so for $y < y'$, $c(y,x) - c(y',x)$ is increasing in x
  $\Rightarrow$ if the sender at x prefers a higher (lower) estimate, this holds at any higher (lower) opinion

- Suppose $x_a < x_b < x_c$, with x_a and x_c inducing m and x_b inducing m'

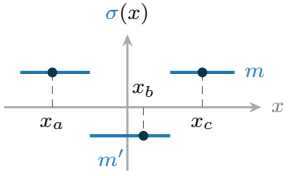


Which opinions induce the same message? Intervals

- **Single crossing of the sender's cost**
  $c(y,x) = (y - \xi x)^2$, writing $y = \widehat{x}$:
  $\dfrac{\partial^2 c}{\partial x\, \partial y} = -2\xi < 0$, so for $y < y'$, $c(y,x) - c(y',x)$ is increasing in x
  $\Rightarrow$ if the sender at x prefers a higher (lower) estimate, this holds at any higher (lower) opinion

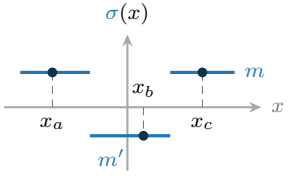
- Suppose $x_a < x_b < x_c$, with x_a and x_c inducing m and x_b inducing m'

- If the sender at x_b preferred m', so would the sender at x_a or at x_c $\Rightarrow$ **contradiction**



Which opinions induce the same message? Intervals

- **Single crossing of the sender's cost**
  $c(y,x) = (y - \xi x)^2$, writing $y = \widehat{x}$:
  $\dfrac{\partial^2 c}{\partial x\, \partial y} = -2\xi < 0$, so for $y < y'$, $c(y,x) - c(y',x)$ is increasing in x
  $\Rightarrow$ if the sender at x prefers a higher (lower) estimate, this holds at any higher (lower) opinion

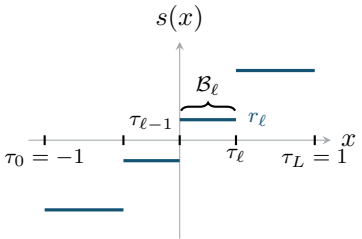


- Suppose $x_a < x_b < x_c$, with x_a and x_c inducing m and x_b inducing m'
- If the sender at x_b preferred m', so would the sender at x_a or at x_c $\Rightarrow$ **contradiction**

Opinions inducing the same message form an interval $\rightarrow$ quantizer

Main result: equilibrium encoders are quantizers

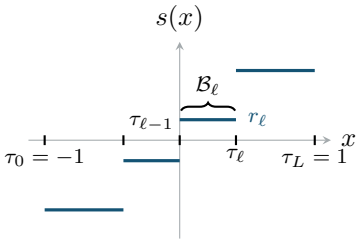
Theorem. In every equilibrium of the communication game $\Gamma(\xi)$, the effective signaling function $s = \mu \circ \sigma$ corresponds to an **interval quantizer**.



a 4-message equilibrium code

Main result: equilibrium encoders are quantizers

Theorem. In every equilibrium of the communication game $\Gamma(\xi)$, the effective signaling function $s = \mu \circ \sigma$ corresponds to an **interval quantizer**.

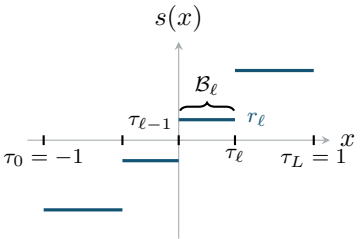


a 4-message equilibrium code

- $[-1, 1]$ is cut into finitely many intervals $\mathcal{B}_1, \ldots, \mathcal{B}_L$, with $\mathcal{B}_\ell = [\tau_{\ell-1}, \tau_\ell]$

Main result: equilibrium encoders are quantizers

Theorem. In every equilibrium of the communication game $\Gamma(\xi)$, the effective signaling function $s = \mu \circ \sigma$ corresponds to an **interval quantizer**.

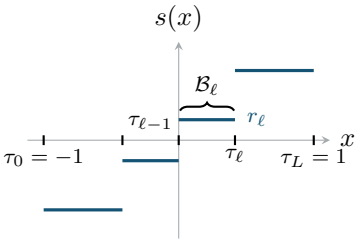


a 4-message equilibrium code

- $[-1, 1]$ is cut into finitely many intervals $\mathcal{B}_1, \ldots, \mathcal{B}_L$, with $\mathcal{B}_\ell = [\tau_{\ell-1}, \tau_\ell]$
- Each interval induces a different message: $\sigma(x) = m_\ell$ on $\mathcal{B}_\ell$

Main result: equilibrium encoders are quantizers

Theorem. In every equilibrium of the communication game $\Gamma(\xi)$, the effective signaling function $s = \mu \circ \sigma$ corresponds to an **interval quantizer**.

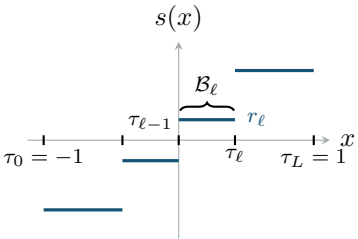


a 4-message equilibrium code

- $[-1, 1]$ is cut into finitely many intervals $\mathcal{B}_1, \ldots, \mathcal{B}_L$, with $\mathcal{B}_\ell = [\tau_{\ell-1}, \tau_\ell]$
- Each interval induces a different message: $\sigma(x) = m_\ell$ on $\mathcal{B}_\ell$
- Each message induces a different estimate: $\mu(m_\ell) = r_\ell = \mathbb{E}[x \mid x \in \mathcal{B}_\ell]$

Main result: equilibrium encoders are quantizers

Theorem. In every equilibrium of the communication game $\Gamma(\xi)$, the effective signaling function $s = \mu \circ \sigma$ corresponds to an **interval quantizer**.



a 4-message equilibrium code

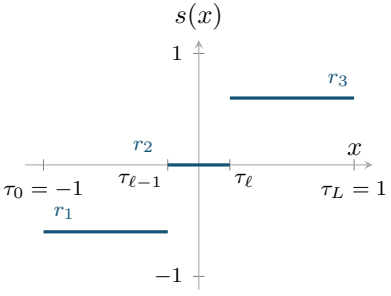
- $[-1, 1]$ is cut into finitely many intervals $\mathcal{B}_1, \ldots, \mathcal{B}_L$, with $\mathcal{B}_\ell = [\tau_{\ell-1}, \tau_\ell]$
- Each interval induces a different message: $\sigma(x) = m_\ell$ on $\mathcal{B}_\ell$
- Each message induces a different estimate: $\mu(m_\ell) = r_\ell = \mathbb{E}[x \mid x \in \mathcal{B}_\ell]$

$\rightarrow$ *How to characterize this quantizer?*

Characterizing the equilibrium quantizer

- **Reminder:** cost $(y - \xi x)^2$; each estimate is a mean, $r_\ell = \mathbb{E}[x \mid x \in \mathcal{B}_\ell]$

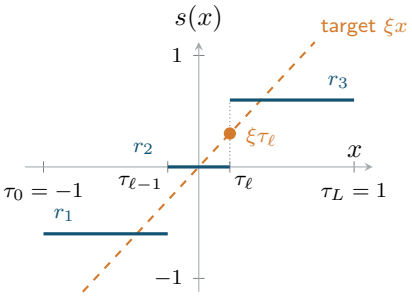
Characterizing the equilibrium quantizer



a 3-message equilibrium code

- **Reminder:** cost $(y - \xi x)^2$; each estimate is a mean, $r_\ell = \mathbb{E}[x \mid x \in \mathcal{B}_\ell]$
- **Example:** $x \sim \mathcal{U}[-1, 1]$: $r_\ell = \frac{1}{2}\left(\tau_{\ell-1} + \tau_\ell\right)$

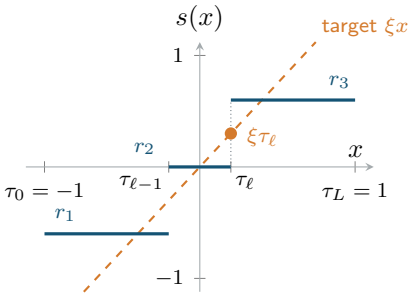
Characterizing the equilibrium quantizer



a 3-message equilibrium code

- **Reminder:** cost $(y - \xi x)^2$; each estimate is a mean, $r_\ell = \mathbb{E}[x \mid x \in \mathcal{B}_\ell]$
- **Example:** $x \sim \mathcal{U}[-1, 1]$: $r_\ell = \frac{1}{2}\big(\tau_{\ell-1} + \tau_\ell\big)$
- **At a threshold τ_ℓ, the sender is indifferent** between r_ℓ and $r_{\ell+1}$:
  $(r_\ell - \xi\tau_\ell)^2 = (r_{\ell+1} - \xi\tau_\ell)^2 \iff \xi\,\tau_\ell = \frac{1}{2}\big(r_\ell + r_{\ell+1}\big)$:
  the target is the midpoint

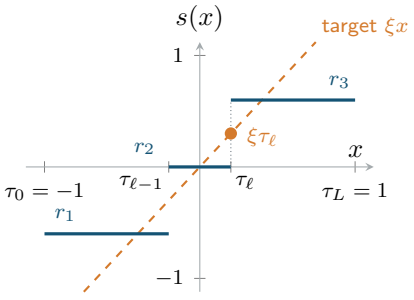
Characterizing the equilibrium quantizer



a 3-message equilibrium code

- **Reminder:** cost $(y - \xi x)^2$; each estimate is a mean, $r_\ell = \mathbb{E}[x \mid x \in \mathcal{B}_\ell]$
- **Example:** $x \sim \mathcal{U}[-1, 1]$: $r_\ell = \frac{1}{2}\big(\tau_{\ell-1} + \tau_\ell\big)$
- **At a threshold τ_ℓ, the sender is indifferent** between r_ℓ and $r_{\ell+1}$:
  $(r_\ell - \xi\tau_\ell)^2 = (r_{\ell+1} - \xi\tau_\ell)^2 \iff \xi\,\tau_\ell = \frac{1}{2}\big(r_\ell + r_{\ell+1}\big)$:
  the target is the midpoint
- r_ℓ **depends on $\tau_{\ell-1}$ and τ_ℓ** $\Rightarrow$ a recurrence for the thresholds:
  $\frac{1}{2}\big(\tau_{\ell-1} + \tau_\ell\big) + \frac{1}{2}\big(\tau_\ell + \tau_{\ell+1}\big) = 2\xi\,\tau_\ell$,
  with $\tau_0 = -1$, $\tau_L = 1$

Characterizing the equilibrium quantizer



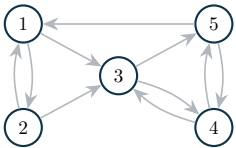
a 3-message equilibrium code

- **Reminder:** cost $(y - \xi x)^2$; each estimate is a mean, $r_\ell = \mathbb{E}[x \mid x \in \mathcal{B}_\ell]$
- **Example:** $x \sim \mathcal{U}[-1, 1]$: $r_\ell = \frac{1}{2}\big(\tau_{\ell-1} + \tau_\ell\big)$
- **At a threshold τ_ℓ, the sender is indifferent** between r_ℓ and $r_{\ell+1}$:
  $(r_\ell - \xi\tau_\ell)^2 = (r_{\ell+1} - \xi\tau_\ell)^2 \iff \xi\,\tau_\ell = \frac{1}{2}\big(r_\ell + r_{\ell+1}\big)$:
  the target is the midpoint
- r_ℓ **depends on $\tau_{\ell-1}$ and τ_ℓ** $\Rightarrow$ a recurrence for the thresholds:
  $\frac{1}{2}\big(\tau_{\ell-1} + \tau_\ell\big) + \frac{1}{2}\big(\tau_\ell + \tau_{\ell+1}\big) = 2\xi\,\tau_\ell$,
  with $\tau_0 = -1$, $\tau_L = 1$

$\rightarrow$ *The encoding depends on the zoomed-in edge through ξ only. What happens at the network level?*

The network game decouples

- In the network, each agent sends to several receivers and hears several senders.



network game

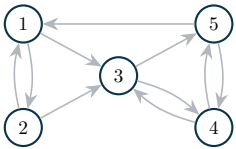
The network game decouples

- In the network, each agent sends to several receivers and hears several senders.

Theorem (decoupling). A profile $(\boldsymbol{\sigma}, \boldsymbol{\mu}, \boldsymbol{\zeta})$ is a PBE of the network game $\iff$ each sender i plays an equilibrium of its own game $\Gamma(\xi_i)$, with ξ_i set by its audience's weights.

Proof idea: reducing the sender's cost by the law of iterated expectations, the other senders' encodings do not matter in expectation:

$$\sum_{k \in \mathcal{N}_i^{\mathrm{A}}} \left(z_k - x_i\right)^2 \;\longrightarrow\; \sum_{k \in \mathcal{N}_i^{\mathrm{A}}} \left(\widetilde{W}_{ki}\, \widehat{x}_i - x_i\right)^2 + \text{const}$$



network game

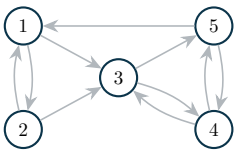
The network game decouples

- In the network, each agent sends to several receivers and hears several senders.

Theorem (decoupling). A profile $(\boldsymbol{\sigma}, \boldsymbol{\mu}, \boldsymbol{\zeta})$ is a PBE of the network game $\iff$ each sender i plays an equilibrium of its own game $\Gamma(\xi_i)$, with ξ_i set by its audience's weights.

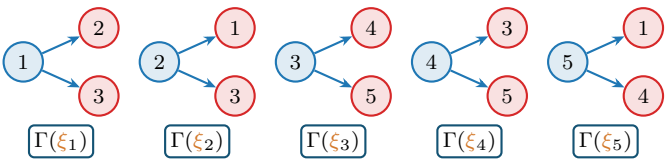
Proof idea: reducing the sender's cost by the law of iterated expectations, the other senders' encodings do not matter in expectation:

$$\textstyle\sum_{k \in \mathcal{N}_i^{\mathrm{A}}} \left(z_k - x_i\right)^2 \;\longrightarrow\; \sum_{k \in \mathcal{N}_i^{\mathrm{A}}} \left(\widetilde{W}_{ki}\, \widehat{x}_i - x_i\right)^2 + \text{const}$$



network game

$\iff$ PBE



one sender game per agent

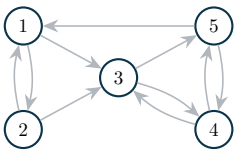
The network game decouples

- In the network, each agent sends to several receivers and hears several senders.

Theorem (decoupling). A profile $(\boldsymbol{\sigma}, \boldsymbol{\mu}, \boldsymbol{\zeta})$ is a PBE of the network game $\iff$ each sender i plays an equilibrium of its own game $\Gamma(\xi_i)$, with ξ_i set by its audience's weights.

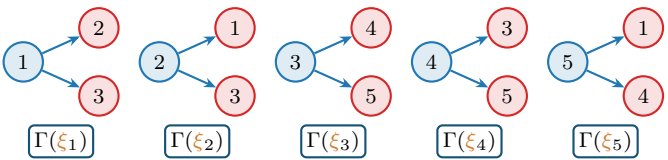
Proof idea: reducing the sender's cost by the law of iterated expectations, the other senders' encodings do not matter in expectation:

$$\textstyle\sum_{k \in \mathcal{N}_i^{\mathrm{A}}} \left(z_k - x_i\right)^2 \;\longrightarrow\; \sum_{k \in \mathcal{N}_i^{\mathrm{A}}} \left(\widetilde{W}_{ki}\, \widehat{x}_i - x_i\right)^2 + \text{const}$$



network game

$\iff$ PBE



one sender game per agent

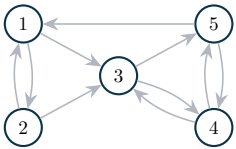
The network game decouples

- In the network, each agent sends to several receivers and hears several senders.

Theorem (decoupling). A profile $(\boldsymbol{\sigma}, \boldsymbol{\mu}, \boldsymbol{\zeta})$ is a PBE of the network game $\iff$ each sender i plays an equilibrium of its own game $\Gamma(\xi_i)$, with ξ_i set by its audience's weights.

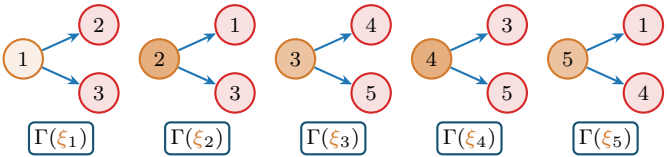
Proof idea: reducing the sender's cost by the law of iterated expectations, the other senders' encodings do not matter in expectation:

$$\sum_{k \in \mathcal{N}_i^{\mathrm{A}}} \left(z_k - x_i\right)^2 \longrightarrow \sum_{k \in \mathcal{N}_i^{\mathrm{A}}} \left(\widetilde{W}_{ki}\, \widehat{x}_i - x_i\right)^2 + \text{const}$$



network game

$\iff$ PBE



one sender game per agent

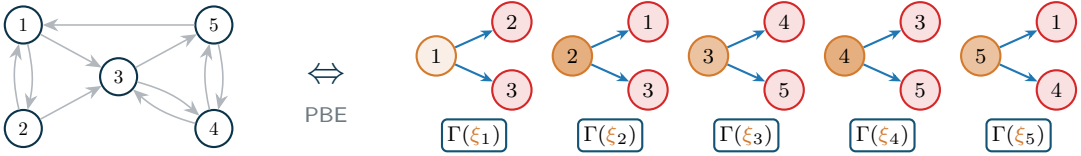
The network game decouples

- In the network, each agent sends to several receivers and hears several senders.

Theorem (decoupling). A profile $(\boldsymbol{\sigma}, \boldsymbol{\mu}, \boldsymbol{\zeta})$ is a PBE of the network game $\iff$ each sender i plays an equilibrium of its own game $\Gamma(\xi_i)$, with ξ_i set by its audience's weights.

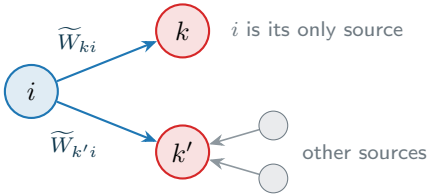
Proof idea: reducing the sender's cost by the law of iterated expectations, the other senders' encodings do not matter in expectation:

$$\sum_{k \in \mathcal{N}_i^{\mathrm{A}}} \left(z_k - x_i\right)^2 \longrightarrow \sum_{k \in \mathcal{N}_i^{\mathrm{A}}} \left(\widetilde{W}_{ki}\, \widehat{x}_i - x_i\right)^2 + \text{const}$$



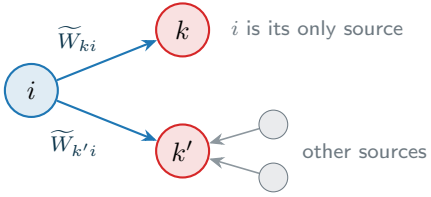
→ *Multiple receivers, one message: how does the exaggeration ξ_i change?*

Two receivers, one public message

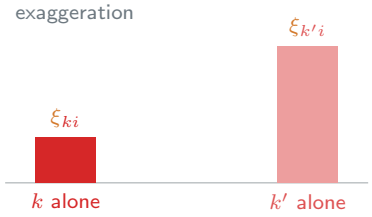


Sender's expected cost:
$\sum_{\ell=k,k'} \left(\widetilde{W}_{\ell i}\, \widehat{x}_i - x_i\right)^2 + \text{const.}$

Two receivers, one public message

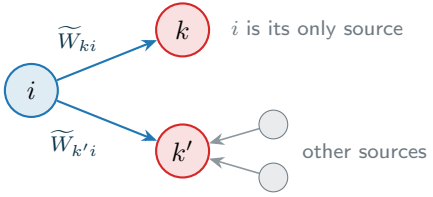


Sender's expected cost:
$\sum_{\ell=k,k'} \left(\widetilde{W}_{\ell i}\, \widehat{x}_i - x_i\right)^2 + \text{const.}$

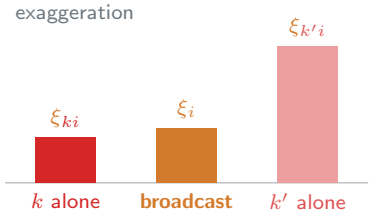


- **Gossip**, one message per receiver: $\xi_{ki} = 1/\widetilde{W}_{ki}$.

Two receivers, one public message

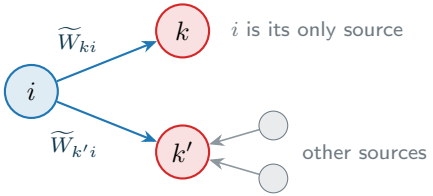


Sender's expected cost:
$\sum_{\ell=k,k'} \left(\widetilde{W}_{\ell i}\, \widehat{x}_i - x_i\right)^2 + \text{const.}$

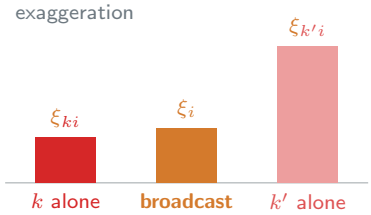


- **Gossip**, one message per receiver: $\xi_{ki} = 1/\widetilde{W}_{ki}$.
- **Broadcast**, one public message: $\xi_i = \dfrac{\sum_k \widetilde{W}_{ki}}{\sum_k \widetilde{W}_{ki}^2}$.

Two receivers, one public message



Sender's expected cost:
$\sum_{\ell=k,k'} \left(\widetilde{W}_{\ell i}\, \widehat{x}_i - x_i\right)^2 + \text{const.}$

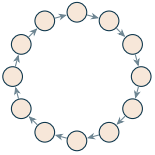


- **Gossip**, one message per receiver: $\xi_{ki} = 1/\widetilde{W}_{ki}$.
- **Broadcast**, one public message: $\xi_i = \dfrac{\sum_k \widetilde{W}_{ki}}{\sum_k \widetilde{W}_{ki}^2}$.

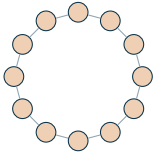
Broadcast: a weighted compromise between the one-edge interactions.

Exaggeration on standard graphs

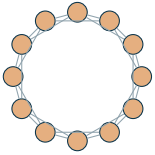
$N = 12$, $\alpha_i = \frac{1}{2}$, equal weights over neighbors; node shade and label: ξ_i



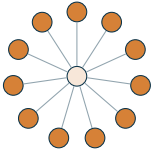
directed cycle
$\xi = 2$



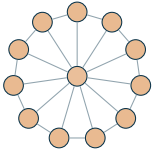
cycle
$\xi = 4$



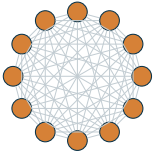
4-regular ring
$\xi = 8$



star
hub $\xi = 2$, leaves $\xi = 22$



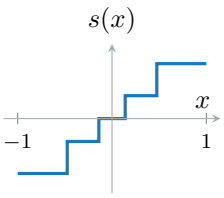
wheel
hub $\xi = 6$, rim $\xi \approx 6.6$



complete
$\xi = 22$

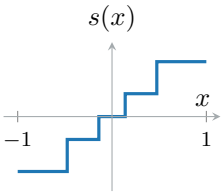
High ξ: effectively binary public expression

$\xi = 1.1$

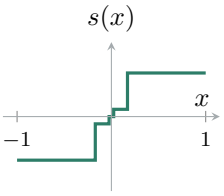


High ξ: effectively binary public expression

$\xi = 1.1$

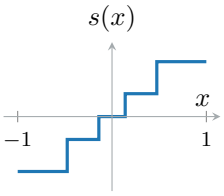


$\xi = 2$

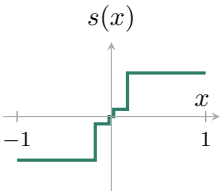


High ξ: effectively binary public expression

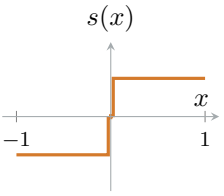
$\xi = 1.1$



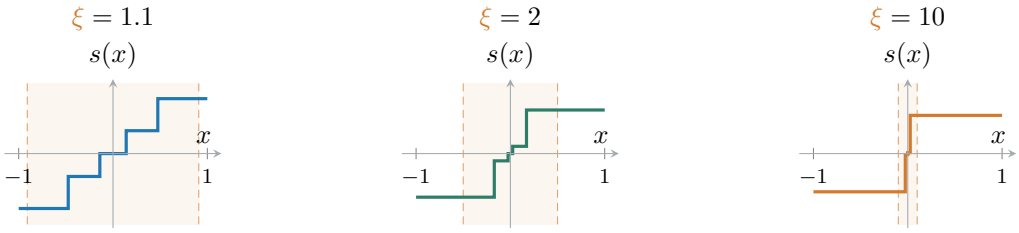
$\xi = 2$



$\xi = 10$



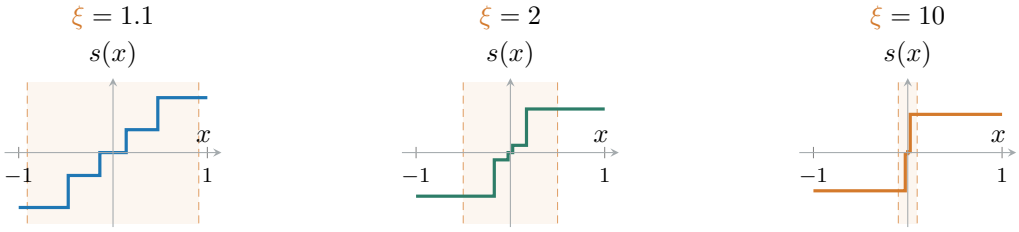
High ξ: effectively binary public expression



dashed: $|x| = 1/\xi$

At equilibrium, every interior threshold lies in $[-1/\xi,\ 1/\xi]$, for any prior and any number of messages.

High ξ: effectively binary public expression

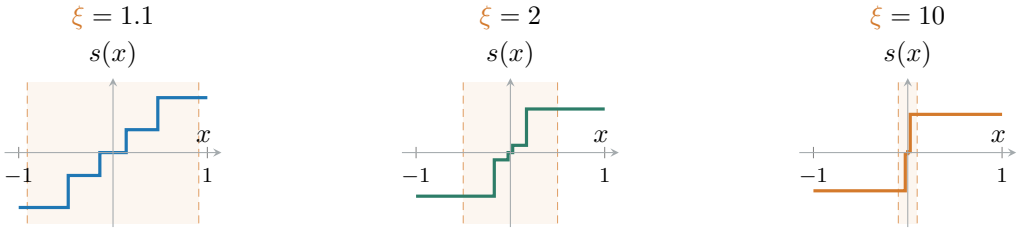


dashed: $|x| = 1/\xi$

At equilibrium, every interior threshold lies in $[-1/\xi,\, 1/\xi]$, for any prior and any number of messages.

The larger ξ, the closer the encoding to $\mathrm{sign}(x)$ $\Rightarrow$ public expression becomes effectively binary.

High ξ: effectively binary public expression



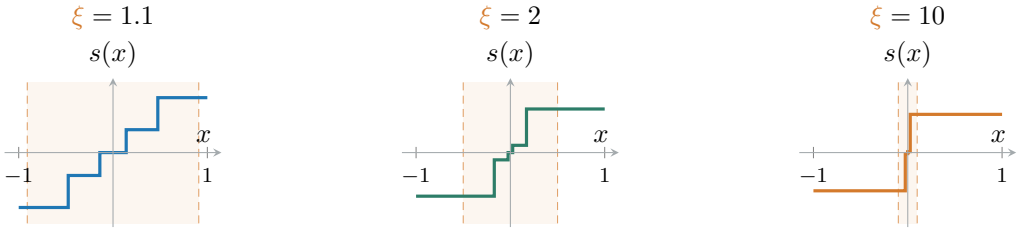
dashed: $|x| = 1/\xi$

At equilibrium, every interior threshold lies in $[-1/\xi, 1/\xi]$, for any prior and any number of messages.

The larger ξ, the closer the encoding to $\mathrm{sign}(x)$ $\Rightarrow$ public expression becomes effectively binary.

This is similar to CODA models (Martins 2008; Chowdhury et al. 2016).

High ξ: effectively binary public expression



dashed: $|x| = 1/\xi$

At equilibrium, every interior threshold lies in $[-1/\xi, 1/\xi]$, for any prior and any number of messages.

The larger ξ, the closer the encoding to $\mathrm{sign}(x)$ $\Rightarrow$ public expression becomes effectively binary.

This is similar to CODA models (Martins 2008; Chowdhury et al. 2016).

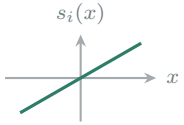
$\rightarrow$ *Next: brief analysis of the dynamics*

Back to the dynamics

- Collective opinion update: $\boldsymbol{x}(t+1) \coloneqq \boldsymbol{z}^*(t) = \boldsymbol{D}_\alpha\, \boldsymbol{x}(t) + (\boldsymbol{I} - \boldsymbol{D}_\alpha)\, \boldsymbol{W}\, \boldsymbol{s}\big(\boldsymbol{x}(t)\big), \quad \boldsymbol{D}_\alpha = \mathrm{diag}(\boldsymbol{\alpha})$

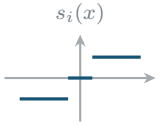
Back to the dynamics

- Collective opinion update: $\boldsymbol{x}(t+1) \coloneqq \boldsymbol{z}^*(t) = \boldsymbol{D}_\alpha\, \boldsymbol{x}(t) + (\boldsymbol{I} - \boldsymbol{D}_\alpha)\, \boldsymbol{W}\, \boldsymbol{s}\big(\boldsymbol{x}(t)\big), \quad \boldsymbol{D}_\alpha = \mathrm{diag}(\boldsymbol{\alpha})$
- Each agent uses its effective signaling function $s_i = \mu_i \circ \sigma_i$:



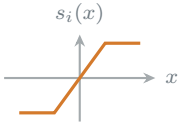
non-persuasive: $s_i(x) = x$

DeGroot



Bayesian: $s_i = \mu_i \circ \sigma_i$

strategic CODA



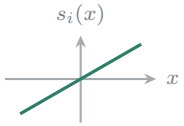
naive receiver:

$s_i(x) = \mathrm{sat}(\xi_i x)$

saturated OD

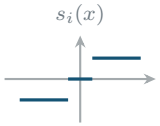
Back to the dynamics

- Collective opinion update: $\boldsymbol{x}(t+1) \coloneqq \boldsymbol{z}^*(t) = \boldsymbol{D}_\alpha\, \boldsymbol{x}(t) + (\boldsymbol{I} - \boldsymbol{D}_\alpha)\, \boldsymbol{W}\, \boldsymbol{s}\big(\boldsymbol{x}(t)\big), \quad \boldsymbol{D}_\alpha = \mathrm{diag}(\boldsymbol{\alpha})$
- Each agent uses its effective signaling function $s_i = \mu_i \circ \sigma_i$:



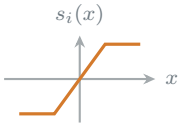
non-persuasive: $s_i(x) = x$

DeGroot



Bayesian: $s_i = \mu_i \circ \sigma_i$

strategic CODA



naive receiver:

$s_i(x) = \mathrm{sat}(\xi_i x)$

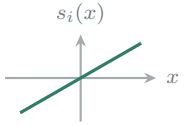
saturated OD

Assumption: **Myopic agents:** each s_i is computed once, against the initial prior, and reused in every period.

Conventions form slowly, opinions move fast.

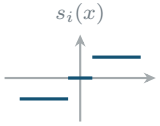
Back to the dynamics

- Collective opinion update: $\boldsymbol{x}(t+1) \coloneqq \boldsymbol{z}^*(t) = \boldsymbol{D}_\alpha\, \boldsymbol{x}(t) + (\boldsymbol{I} - \boldsymbol{D}_\alpha)\, \boldsymbol{W}\, \boldsymbol{s}\big(\boldsymbol{x}(t)\big), \quad \boldsymbol{D}_\alpha = \mathrm{diag}(\boldsymbol{\alpha})$
- Each agent uses its effective signaling function $s_i = \mu_i \circ \sigma_i$:



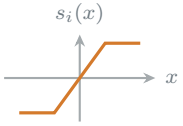
non-persuasive: $s_i(x) = x$

DeGroot



Bayesian: $s_i = \mu_i \circ \sigma_i$

strategic CODA



naive receiver:

$s_i(x) = \mathrm{sat}(\xi_i x)$

saturated OD

Assumption: **Myopic agents:** each s_i is computed once, against the initial prior, and reused in every period.

Conventions form slowly, opinions move fast.

→ *We focus on the Bayesian case: strategic CODA*

Strategic CODA: clusters, but no extremism

$$\boldsymbol{x}(t+1) = \boldsymbol{D}_\alpha \, \boldsymbol{x}(t) + (\boldsymbol{I} - \boldsymbol{D}_\alpha) \, \boldsymbol{W} \, \boldsymbol{s}\big(\boldsymbol{x}(t)\big) =: \boldsymbol{f}\big(\boldsymbol{x}(t)\big)$$

Strategic CODA: clusters, but no extremism

$$\boldsymbol{x}(t+1) = \boldsymbol{D}_\alpha \boldsymbol{x}(t) + (\boldsymbol{I} - \boldsymbol{D}_\alpha)\boldsymbol{W}\,\boldsymbol{s}\big(\boldsymbol{x}(t)\big) =: \boldsymbol{f}\big(\boldsymbol{x}(t)\big)$$

1- Steady states exist

$\boldsymbol{x}^*$ is steady $\iff$ $\boldsymbol{x}^* = \boldsymbol{W}\boldsymbol{s}(\boldsymbol{x}^*)$: each opinion is the weighted average of its sources' estimates.

Existence: $\boldsymbol{f}$ is order-preserving (Tarski).

Strategic CODA: clusters, but no extremism

$$\boldsymbol{x}(t+1) = \boldsymbol{D}_\alpha\, \boldsymbol{x}(t) + (\boldsymbol{I} - \boldsymbol{D}_\alpha)\, \boldsymbol{W}\, \boldsymbol{s}\big(\boldsymbol{x}(t)\big) =: \boldsymbol{f}\big(\boldsymbol{x}(t)\big)$$

1- Steady states exist

$\boldsymbol{x}^*$ is steady $\iff$ $\boldsymbol{x}^* = \boldsymbol{W}\boldsymbol{s}(\boldsymbol{x}^*)$: each opinion is the weighted average of its sources' estimates.

Existence: $\boldsymbol{f}$ is order-preserving (Tarski).

2- No absolute extremism

- The extreme message is estimated as its interval's mean: $r_L < 1$
- So every estimate lies in $[r_1, r_L] \subset (-1, 1)$
- Hence every steady state too: $\pm\mathbf{1}$ is never steady



Strategic CODA: clusters, but no extremism

$$\boldsymbol{x}(t+1) = \boldsymbol{D}_\alpha\, \boldsymbol{x}(t) + (\boldsymbol{I} - \boldsymbol{D}_\alpha)\, \boldsymbol{W}\, \boldsymbol{s}\big(\boldsymbol{x}(t)\big) =: \boldsymbol{f}\big(\boldsymbol{x}(t)\big)$$

1- Steady states exist

$\boldsymbol{x}^*$ is steady $\iff$ $\boldsymbol{x}^* = \boldsymbol{W}\boldsymbol{s}(\boldsymbol{x}^*)$: each opinion is the weighted average of its sources' estimates.

Existence: $\boldsymbol{f}$ is order-preserving (Tarski).

2- No absolute extremism

- The extreme message is estimated as its interval's mean: $r_L < 1$
- So every estimate lies in $[r_1, r_L] \subset (-1, 1)$
- Hence every steady state too: $\pm\mathbf{1}$ is never steady



Clusters and polarization can persist; absolute extremism cannot.

Strategic CODA: clusters, but no extremism

$$\boldsymbol{x}(t+1) = \boldsymbol{D}_\alpha\, \boldsymbol{x}(t) + (\boldsymbol{I} - \boldsymbol{D}_\alpha)\, \boldsymbol{W}\, \boldsymbol{s}\big(\boldsymbol{x}(t)\big) =: \boldsymbol{f}\big(\boldsymbol{x}(t)\big)$$

1- Steady states exist

$\boldsymbol{x}^*$ is steady $\iff$ $\boldsymbol{x}^* = \boldsymbol{W}\boldsymbol{s}(\boldsymbol{x}^*)$: each opinion is the weighted average of its sources' estimates.

Existence: $\boldsymbol{f}$ is order-preserving (Tarski).

2- No absolute extremism

- The extreme message is estimated as its interval's mean: $r_L < 1$
- So every estimate lies in $[r_1, r_L] \subset (-1, 1)$
- Hence every steady state too: $\pm\mathbf{1}$ is never steady



Clusters and polarization can persist; absolute extremism cannot.

→ *Steady states exist, but does the dynamics actually reach one?*

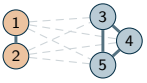
Convergence $\iff$ received public inputs settle

Proposition. For $0 \leq \alpha_i < 1$: $\boldsymbol{x}(t)$ converges $\iff$ $\boldsymbol{W}\boldsymbol{s}(\boldsymbol{x}(t))$ is eventually constant.

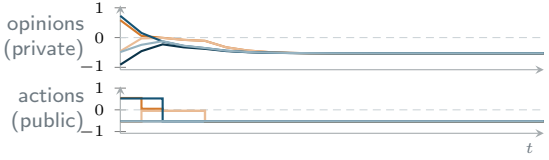
Convergence $\iff$ received public inputs settle

Proposition. For $0 \leq \alpha_i < 1$: $\boldsymbol{x}(t)$ converges $\iff$ $\boldsymbol{W}\boldsymbol{s}(\boldsymbol{x}(t))$ is eventually constant.

Two groups $\{1, 2\}$ and $\{3, 4, 5\}$; illustration: $L = 5$, same $\boldsymbol{x}(0)$; $\alpha_i = 0.5$ (high) or 0.05 (low).



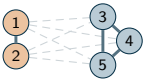
high inertia: the public profile settles



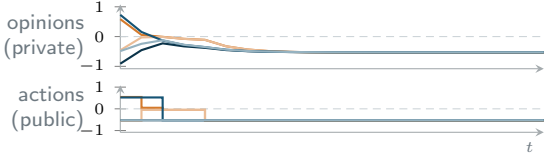
Convergence $\iff$ received public inputs settle

Proposition. For $0 \leq \alpha_i < 1$: $\boldsymbol{x}(t)$ converges $\iff$ $\boldsymbol{W}\boldsymbol{s}(\boldsymbol{x}(t))$ is eventually constant.

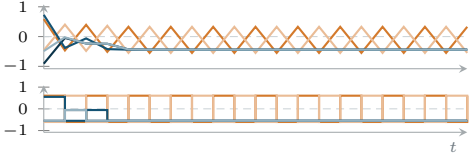
Two groups $\{1, 2\}$ and $\{3, 4, 5\}$; illustration: $L = 5$, same $\boldsymbol{x}(0)$; $\alpha_i = 0.5$ (high) or 0.05 (low).



high inertia: the public profile settles



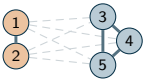
low inertia: it keeps switching



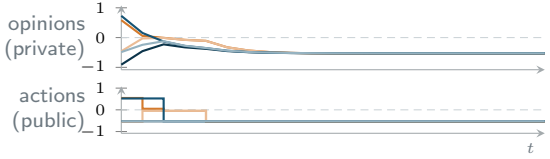
Convergence $\iff$ received public inputs settle

Proposition. For $0 \leq \alpha_i < 1$: $\boldsymbol{x}(t)$ converges $\iff$ $\boldsymbol{W}\boldsymbol{s}(\boldsymbol{x}(t))$ is eventually constant.

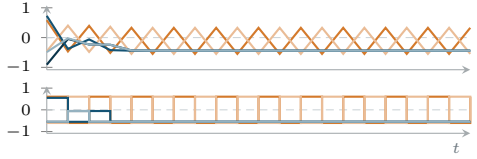
Two groups $\{1, 2\}$ and $\{3, 4, 5\}$; illustration: $L = 5$, same $\boldsymbol{x}(0)$; $\alpha_i = 0.5$ (high) or 0.05 (low).



high inertia: the public profile settles

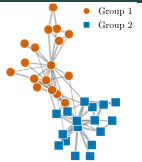


low inertia: it keeps switching



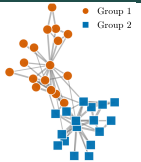
A group agrees in private only once it agrees in public.

Clusters persist: simulation on Zachary's karate club

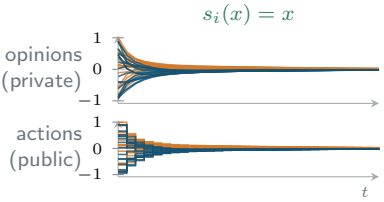


- Zachary's karate club (1977): 34 members, two factions
  $\alpha_i = \frac{1}{2}$, $L = 5$, $x_i(0) \sim \mathcal{U}[-1, 1]$

Clusters persist: simulation on Zachary's karate club

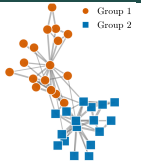


- Zachary's karate club (1977): 34 members, two factions
  $\alpha_i = \frac{1}{2}$, $L = 5$, $x_i(0) \sim \mathcal{U}[-1, 1]$

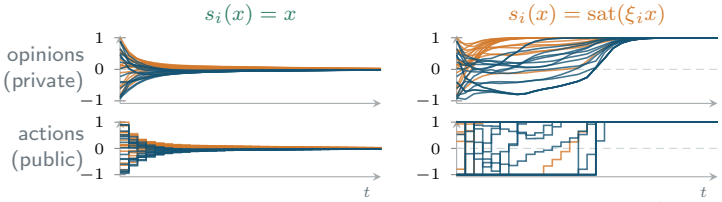


interior consensus

Clusters persist: simulation on Zachary's karate club



- Zachary's karate club (1977): 34 members, two factions
  $\alpha_i = \frac{1}{2}$, $L = 5$, $x_i(0) \sim \mathcal{U}[-1, 1]$

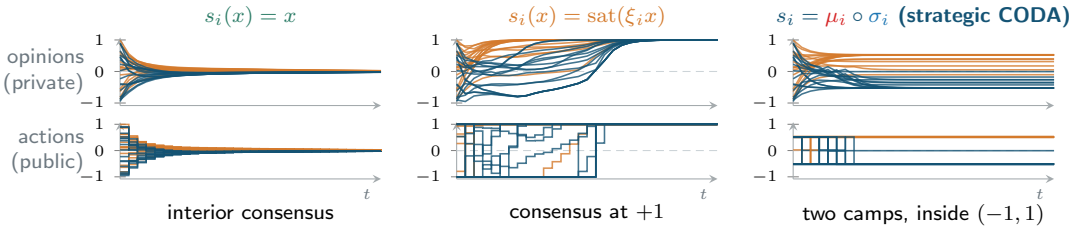


interior consensus | consensus at $+1$

Clusters persist: simulation on Zachary's karate club



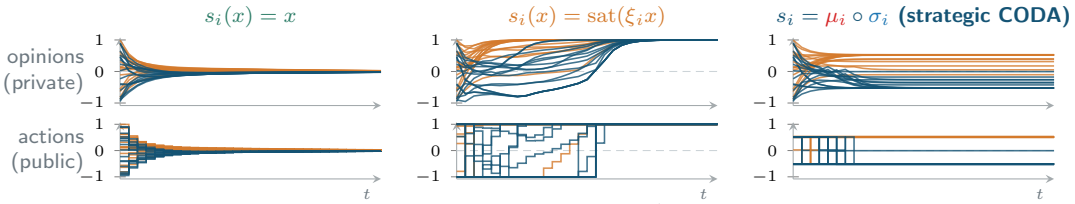
- Zachary's karate club (1977): 34 members, two factions
  $\alpha_i = \frac{1}{2}$, $L = 5$, $x_i(0) \sim \mathcal{U}[-1, 1]$



Clusters persist: simulation on Zachary's karate club

- Zachary's karate club (1977): 34 members, two factions
  $\alpha_i = \frac{1}{2}$, $L = 5$, $x_i(0) \sim \mathcal{U}[-1, 1]$





Same network and start: the encoding decides the outcome.

Summary

1. Agents are both **senders** and **receivers** of information.

Summary

1. Agents are both **senders** and **receivers** of information.
2. As senders, agents can be **persuasive**: they want to move their audience's opinion.

Summary

1. Agents are both **senders** and **receivers** of information.
2. As senders, agents can be **persuasive**: they want to move their audience's opinion.
3. As receivers, agents can be **naive** or **Bayesian**.

Summary

1. Agents are both **senders** and **receivers** of information.
2. As senders, agents can be **persuasive**: they want to move their audience's opinion.
3. As receivers, agents can be **naive** or **Bayesian**.
4. Each agent encodes its opinion with σ_i, its audience decodes the message with μ_i
   $\Rightarrow$ **effective signaling map** $s_i = \mu_i \circ \sigma_i$.

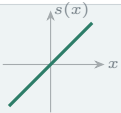
Summary

1. Agents are both **senders** and **receivers** of information.
2. As senders, agents can be **persuasive**: they want to move their audience's opinion.
3. As receivers, agents can be **naive** or **Bayesian**.
4. Each agent encodes its opinion with σ_i, its audience decodes the message with μ_i
   $\Rightarrow$ **effective signaling map** $s_i = \mu_i \circ \sigma_i$.
5. At equilibrium, effective signaling maps are **quantizers**: CODA-like signaling, where public actions are discrete.

The full picture

Naive receivers $\mu(m) = m$ **Bayesian receivers** $\mu(m) = \mathbb{E}[x \mid m]$

Non-persuasive



$s_i(x) = x$ $\beta_i = 0$: non-persuasive sender, truth-telling

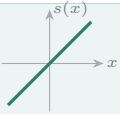
linear averaging (DeGroot 1974)

Persuasive

The full picture

Naive receivers $\mu(m) = m$ **Bayesian receivers** $\mu(m) = \mathbb{E}[x \mid m]$

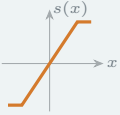
Non-persuasive



$s_i(x) = x$ $\beta_i = 0$: non-persuasive sender, truth-telling

linear averaging (DeGroot 1974)

Persuasive



$s_i(x) = \mathrm{sat}(\xi_i x)$

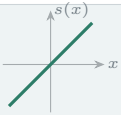
target $\xi_i x$, clipped to $[-1, 1]$

saturated models (Bizyaeva et al. 2023; Couthures et al. 2025)

The full picture

Naive receivers $\mu(m) = m$ | **Bayesian receivers** $\mu(m) = \mathbb{E}[x \mid m]$

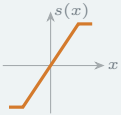
Non-persuasive



$s_i(x) = x$ $\beta_i = 0$: non-persuasive sender, truth-telling

linear averaging (DeGroot 1974)

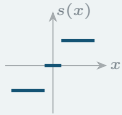
Persuasive



$s_i(x) = \mathrm{sat}(\xi_i x)$

target $\xi_i x$, clipped to $[-1, 1]$

saturated models (Bizyaeva et al. 2023; Couthures et al. 2025)



s_i = equilibrium quantizer

every equilibrium is an interval quantizer

CODA models (Martins 2008) → this work: *strategic CODA*

Perspectives

Same model, open questions

- **The dual role of the network.**
  The same graph drives agreement and generates the communication rules.
- **Network design.**
  Shape the graph to reach a target exaggeration profile, hence a desired public vocabulary.

Perspectives

Same model, open questions

- **The dual role of the network.**
  The same graph drives agreement and generates the communication rules.
- **Network design.**
  Shape the graph to reach a target exaggeration profile, hence a desired public vocabulary.

Extensions

- **Full gossip model.**
  Targeted messages, one per receiver.
- **Repeated game.**
  Relaxing the myopic agents assumption: receivers update their beliefs as the opinions evolve.

Thanks for your attention.

Questions?