Lecture 14 - Opinion Dynamics, Financial Networks

Models for real networks
ECE 5260 – ORIE 5735

Anna Scaglione

Cornell Tech

March 16, 2026

Content

Cornell University

Review and overview

Cornell University

- We introduced some abstract models for network dynamics
  - Random walk
  - Consensus (synchronous and asynchronous)
  - General non-linear models
  - Affine linear models
  - Positive dynamics and flow networks
- We also introduced an application of non-linear modeling in the study of epidemic models
- Today we consider opinion dynamics in social networks

Review and overview

Cornell University

- We introduced some abstract models for network dynamics
  - Random walk
  - Consensus (synchronous and asynchronous)
  - General non-linear models
  - Affine linear models
  - Positive dynamics and flow networks
- We also introduced an application of non-linear modeling in the study of epidemic models
- Today we consider opinion dynamics in social networks

Review and overview

Cornell University

- We introduced some abstract models for network dynamics
  - Random walk
  - Consensus (synchronous and asynchronous)
  - General non-linear models
  - Affine linear models
  - Positive dynamics and flow networks
- We also introduced an application of non-linear modeling in the study of epidemic models
- Today we consider opinion dynamics in social networks

Review and overview

Cornell University

- We introduced some abstract models for network dynamics
  - Random walk
  - Consensus (synchronous and asynchronous)
  - General non-linear models
  - Affine linear models
  - Positive dynamics and flow networks
- We also introduced an application of non-linear modeling in the study of epidemic models
- Today we consider opinion dynamics in social networks

Review and overview

Cornell University

- We introduced some abstract models for network dynamics
  - Random walk
  - Consensus (synchronous and asynchronous)
  - General non-linear models
  - Affine linear models
  - Positive dynamics and flow networks
- We also introduced an application of non-linear modeling in the study of epidemic models
- Today we consider opinion dynamics in social networks

Review and overview

Cornell University

- We introduced some abstract models for network dynamics
  - Random walk
  - Consensus (synchronous and asynchronous)
  - General non-linear models
  - Affine linear models
  - Positive dynamics and flow networks
- We also introduced an application of non-linear modeling in the study of epidemic models
- Today we consider opinion dynamics in social networks

Outline

Cornell University

Outline

Cornell University

1 Review
- Stochastic and linear models
- General non-linar models
- Linear affine models

Stochastic and linear models

Cornell University

- **Random walk**: At each t:
  1. $\boldsymbol{x}^{\text{rw}}(t) = \boldsymbol{e}_{v_t}(t)$, where v_t is the vertex chosen at random and $\boldsymbol{e}_v(t)$ is the $v$thcoordinate vector
  2. The probability mass of the vertex being chosen is the vector such that $\boldsymbol{\pi}(t+1) = \boldsymbol{W}\boldsymbol{\pi}(t)$, where $\boldsymbol{W}$ is the transition probability matrix ($\boldsymbol{W} = \boldsymbol{I} - \boldsymbol{L}^{\text{rw}}$ for the uniform random walk)
  3. $\lambda_2(\boldsymbol{W}) < 1$ implies that $\boldsymbol{\pi}(t) = \boldsymbol{\pi}^* + O(\lambda_2^t(\boldsymbol{W}))$
- **Consensus Gossiping**: At each t:
  1. The state is updated **synchrounously** $\boldsymbol{x}(t+1) = \boldsymbol{W}\boldsymbol{x}(t)$,
  2. The $\|\boldsymbol{x}(t) - x^{\text{ave}}\mathbf{1}\| = O(\lambda_2^t(\boldsymbol{W}))$
  3. The state is updated **asynchrounously** $\boldsymbol{x}(t+1) = \boldsymbol{W}(t)\boldsymbol{x}(t)$,
  4. The $\mathbb{E}[\|\boldsymbol{x}(t) - x^{\text{ave}}\mathbf{1}\|] = O(\lambda_2^t(\overline{\boldsymbol{W}}))$

Stochastic and linear models

Cornell University

- **Random walk**: At each t:
  1. $\boldsymbol{x}^{\text{rw}}(t) = \boldsymbol{e}_{v_t}(t)$, where v_t is the vertex chosen at random and $\boldsymbol{e}_v(t)$ is the $v$thcoordinate vector
  2. The probability mass of the vertex being chosen is the vector such that $\boldsymbol{\pi}(t+1) = \boldsymbol{W}\boldsymbol{\pi}(t)$, where $\boldsymbol{W}$ is the transition probability matrix ($\boldsymbol{W} = \boldsymbol{I} - \boldsymbol{L}^{\text{rw}}$ for the uniform random walk)
  3. $\lambda_2(\boldsymbol{W}) < 1$ implies that $\boldsymbol{\pi}(t) = \boldsymbol{\pi}^* + O(\lambda_2^t(\boldsymbol{W}))$

- **Consensus Gossiping**: At each t:
  1. The state is updated **synchrounously** $\boldsymbol{x}(t+1) = \boldsymbol{W}\boldsymbol{x}(t)$,
  2. The $\|\boldsymbol{x}(t) - x^{\text{ave}}\mathbf{1}\| = O(\lambda_2^t(\boldsymbol{W}))$
  3. The state is updated **asynchrounously** $\boldsymbol{x}(t+1) = \boldsymbol{W}(t)\boldsymbol{x}(t)$,
  4. The $\mathbb{E}[\|\boldsymbol{x}(t) - x^{\text{ave}}\mathbf{1}\|] = O(\lambda_2^t(\overline{\boldsymbol{W}}))$

Stochastic and linear models

- **Random walk**: At each t:
  1. $\boldsymbol{x}^{\text{rw}}(t) = \boldsymbol{e}_{v_t}(t)$, where v_t is the vertex chosen at random and $\boldsymbol{e}_v(t)$ is the $v$thcoordinate vector
  2. The probability mass of the vertex being chosen is the vector such that $\boldsymbol{\pi}(t+1) = \boldsymbol{W}\boldsymbol{\pi}(t)$, where $\boldsymbol{W}$ is the transition probability matrix ($\boldsymbol{W} = \boldsymbol{I} - \boldsymbol{L}^{\text{rw}}$ for the uniform random walk)
  3. $\lambda_2(\boldsymbol{W}) < 1$ implies that $\boldsymbol{\pi}(t) = \boldsymbol{\pi}^* + O(\lambda_2^t(\boldsymbol{W}))$
- **Consensus Gossiping**: At each t:
  1. The state is updated **synchrounously** $x(t+1) = \boldsymbol{W}\boldsymbol{x}(t)$,
  2. The $\|\boldsymbol{x}(t) - x^{\text{ave}}\mathbf{1}\| = O(\lambda_2^t(\boldsymbol{W}))$
  3. The state is updated **asynchrounously** $x(t+1) = \boldsymbol{W}(t)\boldsymbol{x}(t)$,
  4. The $\mathbb{E}[\|\boldsymbol{x}(t) - x^{\text{ave}}\mathbf{1}\|] = O(\lambda_2^t(\overline{\boldsymbol{W}}))$

Stochastic and linear models

Cornell University

- **Random walk**: At each t:
  1. $\boldsymbol{x}^{\text{rw}}(t) = \boldsymbol{e}_{v_t}(t)$, where v_t is the vertex chosen at random and $\boldsymbol{e}_v(t)$ is the $v$thcoordinate vector
  2. The probability mass of the vertex being chosen is the vector such that $\boldsymbol{\pi}(t+1) = \boldsymbol{W}\boldsymbol{\pi}(t)$, where $\boldsymbol{W}$ is the transition probability matrix ($\boldsymbol{W} = \boldsymbol{I} - \boldsymbol{L}^{\text{rw}}$ for the uniform random walk)
  3. $\lambda_2(\boldsymbol{W}) < 1$ implies that $\boldsymbol{\pi}(t) = \boldsymbol{\pi}^* + O(\lambda_2^t(\boldsymbol{W}))$
- **Consensus Gossiping**: At each t:
  1. The state is updated **synchrounously** $x(t+1) = \boldsymbol{W}\boldsymbol{x}(t)$,
  2. The $\|\boldsymbol{x}(t) - x^{\text{ave}}\mathbf{1}\| = O(\lambda_2^t(\boldsymbol{W}))$
  3. The state is updated **asynchrounously** $\boldsymbol{x}(t+1) = \boldsymbol{W}(t)\boldsymbol{x}(t)$,
  4. The $\mathbb{E}[\|\boldsymbol{x}(t) - x^{\text{ave}}\mathbf{1}\|] = O(\lambda_2^t(\overline{\boldsymbol{W}}))$

Stochastic and linear models

Cornell University

- **Random walk**: At each t:
  1. $\boldsymbol{x}^{\text{rw}}(t) = \boldsymbol{e}_{v_t}(t)$, where v_t is the vertex chosen at random and $\boldsymbol{e}_v(t)$ is the $v$thcoordinate vector
  2. The probability mass of the vertex being chosen is the vector such that $\boldsymbol{\pi}(t+1) = \boldsymbol{W}\boldsymbol{\pi}(t)$, where $\boldsymbol{W}$ is the transition probability matrix ($\boldsymbol{W} = \boldsymbol{I} - \boldsymbol{L}^{\text{rw}}$ for the uniform random walk)
  3. $\lambda_2(\boldsymbol{W}) < 1$ implies that $\boldsymbol{\pi}(t) = \boldsymbol{\pi}^* + O(\lambda_2^t(\boldsymbol{W}))$

- **Consensus Gossiping**: At each t:
  1. The state is updated **synchrounously** $\boldsymbol{x}(t+1) = \boldsymbol{W}\boldsymbol{x}(t)$,
  2. The $\|\boldsymbol{x}(t) - x^{\text{ave}}\mathbf{1}\| = O(\lambda_2^t(\boldsymbol{W}))$
  3. The state is updated **asynchrounously** $\boldsymbol{x}(t+1) = \boldsymbol{W}(t)\boldsymbol{x}(t)$,
  4. The $\mathbb{E}[\|\boldsymbol{x}(t) - x^{\text{ave}}\mathbf{1}\|] = O(\lambda_2^t(\overline{\boldsymbol{W}}))$

Stochastic and linear models

Cornell University

- **Random walk**: At each t:
  1. $\boldsymbol{x}^{\text{rw}}(t) = \boldsymbol{e}_{v_t}(t)$, where v_t is the vertex chosen at random and $\boldsymbol{e}_v(t)$ is the $v$thcoordinate vector
  2. The probability mass of the vertex being chosen is the vector such that $\boldsymbol{\pi}(t+1) = \boldsymbol{W}\boldsymbol{\pi}(t)$, where $\boldsymbol{W}$ is the transition probability matrix ($\boldsymbol{W} = \boldsymbol{I} - \boldsymbol{L}^{\text{rw}}$ for the uniform random walk)
  3. $\lambda_2(\boldsymbol{W}) < 1$ implies that $\boldsymbol{\pi}(t) = \boldsymbol{\pi}^* + O(\lambda_2^t(\boldsymbol{W}))$
- **Consensus Gossiping**: At each t:
  1. The state is updated **synchrounously** $\boldsymbol{x}(t+1) = \boldsymbol{W}\boldsymbol{x}(t)$,
  2. The $\|\boldsymbol{x}(t) - x^{\text{ave}}\mathbf{1}\| = O(\lambda_2^t(\boldsymbol{W}))$
  3. The state is updated **asynchrounously** $\boldsymbol{x}(t+1) = \boldsymbol{W}(t)\boldsymbol{x}(t)$,
  4. The $\mathbb{E}[\|\boldsymbol{x}(t) - x^{\text{ave}}\mathbf{1}\|] = O(\lambda_2^t(\overline{\boldsymbol{W}}))$

Outline

Cornell University

1 Review
- Stochastic and linear models
- General non-linar models
- Linear affine models

General non-linar models

- Non-linear network dynamics are more often studied in the form of ODEs. The general form we introduced is as follows:

$$\dot{x}_i = f(x_i) + \sum_{j \in \mathcal{N}_i} w_{ij} g(x_i, x_j)$$

- These systems can also be written in vector form:

$$\dot{\boldsymbol{x}} = \boldsymbol{f}(\boldsymbol{x}), \quad \text{where} \quad [\boldsymbol{f}(\boldsymbol{x})]_i = f(x_i) + \sum_{j \in \mathcal{N}_i} w_{ij} g(x_i, x_j)$$

- We also indicated how, even if there is no general solution, it is useful to gain insight about the stability of the stationary points where $\dot{x}_i(t) = 0$ for all $i \in \mathcal{V}$:

$$f(x_i^*) + \sum_{j \in \mathcal{N}_i} g(x_i^*, x_j^*) = 0, \quad \forall i \in \mathcal{V} \quad \text{or}$$

General non-linar models

Cornell University

- Non-linear network dynamics are more often studied in the form of ODEs. The general form we introduced is as follows:

$$\dot{x}_i = f(x_i) + \sum_{j \in \mathcal{N}_i} w_{ij} g(x_i, x_j)$$

- These systems can also be written in vector form:

$$\dot{\boldsymbol{x}} = \boldsymbol{f}(\boldsymbol{x}), \quad \text{where} \;\; [\boldsymbol{f}(\boldsymbol{x})]_i = f(x_i) + \sum_{j \in \mathcal{N}_i} w_{ij} g(x_i, x_j)$$

- We also indicated how, even if there is no general solution, it is useful to gain insight about the stability of the stationary points where $\dot{x}_i(t) = 0$ for all $i \in \mathcal{V}$:

$$f(x_i^*) + \sum_{j \in \mathcal{N}_i} g(x_i^*, x_j^*) = 0, \quad \forall i \in \mathcal{V} \;\; \text{or}$$

Local stability

- The starting point is a fixed point $\boldsymbol{x}^*$ for the equations, defining:

$$\alpha_i = \left.\frac{\partial f}{\partial x}\right|_{x=x_i^*}, \ \beta_{ij} = \left.\frac{\partial g\,(u,v)}{\partial u}\right|_{u=x_i^*,v=x_j^*}, \ \gamma_{ij} = \left.\frac{\partial g\,(u,v)}{\partial v}\right|_{u=x_i^*,v=x_j^*}$$

$$\Rightarrow \dot{\boldsymbol{\epsilon}} \equiv \boldsymbol{M}\boldsymbol{\epsilon}, \quad \text{with} \quad [\boldsymbol{M}]_{ij} = \left[\alpha_i \delta(i-j) + \sum_j \beta_{ij} A_{ij}\right] + \gamma_{ij} A_{ij}$$

and $\delta(i)$ is the Kronecker delta.

- Note that $\boldsymbol{M} = \alpha \boldsymbol{I} + \beta \boldsymbol{L}$ in the special case $\boldsymbol{x}^* = x^* \boldsymbol{1}$, $\left.\frac{\partial f}{\partial x}\right|_{x=x^*} = \alpha$ and if $\left.\frac{\partial g(u,v)}{\partial v}\right|_{u=x^*,v=x^*} = -\left.\frac{\partial g(u,v)}{\partial u}\right|_{u=x^*,v=x^*} = \beta$

- The solution depends on the EVD of $\boldsymbol{M} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^{-1}$ and $\boldsymbol{\varepsilon}(0) = \boldsymbol{U}^{-1}\boldsymbol{\epsilon}(0)$, and $\boldsymbol{x}^*$ is stable if all $\lambda_k < 1$:

$$\boldsymbol{x}(t) \approx \boldsymbol{x}^* + \sum_{k=1}^{N} \varepsilon_k(0) e^{\lambda_k t} \boldsymbol{u}_k$$

Local stability

Cornell University

- The starting point is a fixed point $\boldsymbol{x}^*$ for the equations, defining:

$$\alpha_i = \left.\frac{\partial f}{\partial x}\right|_{x=x_i^*}, \ \beta_{ij} = \left.\frac{\partial g\,(u,v)}{\partial u}\right|_{u=x_i^*, v=x_j^*}, \ \gamma_{ij} = \left.\frac{\partial g\,(u,v)}{\partial v}\right|_{u=x_i^*, v=x_j^*}$$

$$\Rightarrow \dot{\boldsymbol{\epsilon}} \equiv \boldsymbol{M}\boldsymbol{\epsilon}, \quad \text{with} \quad [\boldsymbol{M}]_{ij} = \left[\alpha_i \delta(i-j) + \sum_j \beta_{ij} A_{ij}\right] + \gamma_{ij} A_{ij}$$

and $\delta(i)$ is the Kronecker delta.

- Note that $\boldsymbol{M} = \alpha \boldsymbol{I} + \beta \boldsymbol{L}$ in the special case $\boldsymbol{x}^* = x^* \mathbf{1}$, $\left.\frac{\partial f}{\partial x}\right|_{x=x^*} = \alpha$ and if $\left.\frac{\partial g(u,v)}{\partial v}\right|_{u=x^*, v=x^*} = -\left.\frac{\partial g(u,v)}{\partial u}\right|_{u=x^*, v=x^*} = \beta$

- The solution depends on the EVD of $\boldsymbol{M} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^{-1}$ and $\boldsymbol{\varepsilon}(0) = \boldsymbol{U}^{-1}\boldsymbol{\epsilon}(0)$, and $\boldsymbol{x}^*$ is stable if all $\lambda_k < 1$:

$$\boldsymbol{x}(t) \approx \boldsymbol{x}^* + \sum_{k=1}^{N} \varepsilon_k(0) e^{\lambda_k t} \boldsymbol{u}_k$$

Local stability

Cornell University

- The starting point is a fixed point $\boldsymbol{x}^*$ for the equations, defining:

$$\alpha_i = \left.\frac{\partial f}{\partial x}\right|_{x=x_i^*}, \ \beta_{ij} = \left.\frac{\partial g\,(u,v)}{\partial u}\right|_{u=x_i^*, v=x_j^*}, \ \gamma_{ij} = \left.\frac{\partial g\,(u,v)}{\partial v}\right|_{u=x_i^*, v=x_j^*}$$

$$\Rightarrow \dot{\boldsymbol{\epsilon}} \equiv \boldsymbol{M}\boldsymbol{\epsilon}, \quad \text{with} \quad [\boldsymbol{M}]_{ij} = \left[\alpha_i \delta(i-j) + \sum_j \beta_{ij} A_{ij}\right] + \gamma_{ij} A_{ij}$$

  and $\delta(i)$ is the Kronecker delta.

- Note that $\boldsymbol{M} = \alpha \boldsymbol{I} + \beta \boldsymbol{L}$ in the special case $\boldsymbol{x}^* = x^* \mathbf{1}$, $\left.\frac{\partial f}{\partial x}\right|_{x=x^*} = \alpha$ and if $\left.\frac{\partial g(u,v)}{\partial v}\right|_{u=x^*, v=x^*} = -\left.\frac{\partial g(u,v)}{\partial u}\right|_{u=x^*, v=x^*} = \beta$

- The solution depends on the EVD of $\boldsymbol{M} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^{-1}$ and $\boldsymbol{\varepsilon}(0) = \boldsymbol{U}^{-1}\boldsymbol{\epsilon}(0)$, and $\boldsymbol{x}^*$ is stable if all $\lambda_k < 1$:

$$\boldsymbol{x}(t) \approx \boldsymbol{x}^* + \sum_{k=1}^{N} \varepsilon_k(0) e^{\lambda_k t} \boldsymbol{u}_k$$

Outline

Cornell University

Affine linear dynamics

Cornell University

Linear affine equations in discrete and ODE form respectively are

$$\boldsymbol{x}(t+1) = \boldsymbol{W}\boldsymbol{x}(t) + \boldsymbol{b}, \quad \dot{\boldsymbol{x}}(t) = \boldsymbol{M}\boldsymbol{x}(t) + \boldsymbol{b}$$

Affine discrete dynamics

If and only if $\boldsymbol{W}$ has all $\lambda_k(\boldsymbol{W})$ <1 the affine discrete dynamics $\boldsymbol{x}$ converge to:

$$\boldsymbol{x}^* = (\boldsymbol{I} - \boldsymbol{W})^{-1}\boldsymbol{b}$$

If $w_{ij} \geq 0$ then $\boldsymbol{W}$ has to be **sub-stochastic**, i.e. $\boldsymbol{1}^\top \boldsymbol{W} < 1$.

Affine continuous dynamics

If and only if $\boldsymbol{M}$ is Hurwitz if all its eigenvalues have negative real part, i.e. $\mathbb{R}\{\lambda_k\} < 0 \quad \forall k = 1, N$ the continuous dynamics converge to:

$$\boldsymbol{x}^* = -\boldsymbol{M}^{-1}\boldsymbol{b}$$

If $\boldsymbol{M}$ is metzler and $\boldsymbol{b} \geq 0$ then all $\boldsymbol{x}(t)$ including $\boldsymbol{x}^*$ are non-negative.

Affine linear dynamics

Cornell University

Linear affine equations in discrete and ODE form respectively are

$$\boldsymbol{x}(t+1) = \boldsymbol{W}\boldsymbol{x}(t) + \boldsymbol{b}, \quad \dot{\boldsymbol{x}}(t) = \boldsymbol{M}\boldsymbol{x}(t) + \boldsymbol{b}$$

Affine discrete dynamics

If and only if $\boldsymbol{W}$ has all $\lambda_k(\boldsymbol{W})$ <1 the affine discrete dynamics $\boldsymbol{x}$ converge to:

$$\boldsymbol{x}^* = (\boldsymbol{I} - \boldsymbol{W})^{-1}\boldsymbol{b}$$

If $w_{ij} \geq 0$ then $\boldsymbol{W}$ has to be **sub-stochastic**, i.e. $\mathbf{1}^\top \boldsymbol{W} < 1$.

Affine continuous dynamics

If and only if $\boldsymbol{M}$ is Hurwitz if all its eigenvalues have negative real part, i.e. $\mathbb{R}\{\lambda_k\} < 0 \quad \forall k = 1, N$ the continuous dynamics converge to:

$$\boldsymbol{x}^* = -\boldsymbol{M}^{-1}\boldsymbol{b}$$

If $\boldsymbol{M}$ is metzler and $\boldsymbol{b} \geq 0$ then all $\boldsymbol{x}(t)$ including $\boldsymbol{x}^*$ are non-negative.

Affine linear dynamics

Cornell University

Linear affine equations in discrete and ODE form respectively are

$$\boldsymbol{x}(t+1) = \boldsymbol{W}\boldsymbol{x}(t) + \boldsymbol{b}, \quad \dot{\boldsymbol{x}}(t) = \boldsymbol{M}\boldsymbol{x}(t) + \boldsymbol{b}$$

Affine discrete dynamics

If and only if $\boldsymbol{W}$ has all $\lambda_k(\boldsymbol{W})$ <1 the affine discrete dynamics $\boldsymbol{x}$ converge to:

$$\boldsymbol{x}^* = (\boldsymbol{I} - \boldsymbol{W})^{-1}\boldsymbol{b}$$

If $w_{ij} \geq 0$ then $\boldsymbol{W}$ has to be **sub-stochastic**, i.e. $\mathbf{1}^\top \boldsymbol{W} < 1$.

Affine continuous dynamics

If and only if $\boldsymbol{M}$ is Hurwitz if all its eigenvalues have negative real part, i.e. $\mathbb{R}\{\lambda_k\} < 0 \ \ \forall k = 1, N$ the continuous dynamics converge to:

$$\boldsymbol{x}^* = -\boldsymbol{M}^{-1}\boldsymbol{b}$$

If $\boldsymbol{M}$ is metzler and $\boldsymbol{b} \geq 0$ then all $\boldsymbol{x}(t)$ including $\boldsymbol{x}^*$ are non-negative.

Outline

Cornell University

Why Study Social Network Dynamics?

Cornell University

Many collective phenomena emerge from interactions on networks:

- How opinions spread on social media
- How misinformation becomes viral
- How financial bubbles take place
- How innovations diffuse through societies

Key question:

How do individual beliefs evolve when people repeatedly influence each other?

We will model this as a **dynamical system on a network**.

Why Study Social Network Dynamics?

Many collective phenomena emerge from interactions on networks:

- How opinions spread on social media
- How misinformation becomes viral
- How financial bubbles take place
- How innovations diffuse through societies

Key question:

How do individual beliefs evolve when people repeatedly influence each other?

We will model this as a **dynamical system on a network**.

Why Study Social Network Dynamics?

Many collective phenomena emerge from interactions on networks:

- How opinions spread on social media
- How misinformation becomes viral
- How financial bubbles take place
- How innovations diffuse through societies

Key question:

How do individual beliefs evolve when people repeatedly influence each other?

We will model this as a **dynamical system on a network**.

Understanding social network systems

- *Opinion dynamics* model agents' influencing each others beliefs

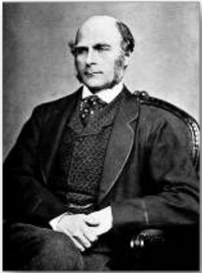

- Early interest in them: Nicolas Marquis de Condorcet, Francis Galton famous paper "Vox populi"sought to prove mathematically that *crowds were wiser* than each individual

Learning under social pressure

Cornell University

- Galton averaged many answers about the weight of an ox, noting the average was very close to the true value
- **Modern opinion dynamics** recognize the critical difference between voting by individuals whose preference are based on their individual information and individuals who observe the votes of others

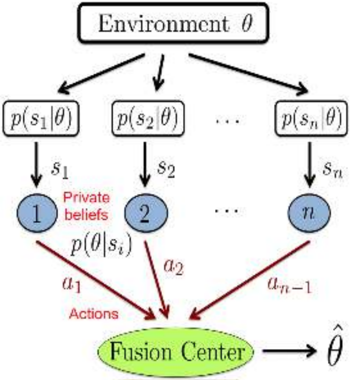

Learning under social pressure

- Galton averaged many answers about the weight of an ox, noting the average was very close to the true value
- **Modern opinion dynamics** recognize the critical difference between voting by individuals whose preference are based on their individual information and individuals who observe the votes of others

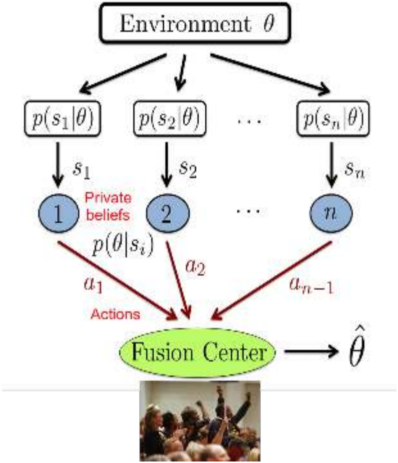

Outline

Cornell University

2 Opinion dynamics modeling
- Models for opinion diffusion
- Random walk based models

Opinion Dynamics

Cornell University

The social systems equations

Describe through a deterministic or a probabilistic dynamical model how a set of agents combine their private beliefs with observations (direct or indirect) of the peers' beliefs. The goal is understanding the *emergent behavior* of society.

- Agent $i = 1, \ldots, n$ holds an initial opinion $x_i(0)$ on a topic θ. For example, for agent i it can be $x_i(0) = p(\theta|s_i)$ or $x_i(0) = \mathbb{E}[\theta|s_i]$
- There is a network process (social activities) through which these opinions are shared and updated

agents beliefs at time t $\boldsymbol{x}(t) = (x_1(t), \ldots, x_n(t))$

continuous time $\mapsto$ $\dot{\boldsymbol{x}}(t) = \boldsymbol{f}(\boldsymbol{x}(t))$, discrete time $\mapsto$ $\boldsymbol{x}(t+1) = \boldsymbol{f}(\boldsymbol{x}(t))$

Opinion Dynamics

The social systems equations

Describe through a deterministic or a probabilistic dynamical model how a set of agents combine their private beliefs with observations (direct or indirect) of the peers' beliefs. The goal is understanding the *emergent behavior* of society.

- Agent $i = 1, \ldots, n$ holds an initial opinion $x_i(0)$ on a topic θ. For example, for agent i it can be $x_i(0) = p(\theta|s_i)$ or $x_i(0) = \mathbb{E}[\theta|s_i]$
- There is a network process (social activities) through which these opinions are shared and updated

agents beliefs at time t $\boldsymbol{x}(t) = (x_1(t), \ldots, x_n(t))$

continuous time $\mapsto$ $\dot{\boldsymbol{x}}(t) = \boldsymbol{f}(\boldsymbol{x}(t))$, discrete time $\mapsto$ $\boldsymbol{x}(t+1) = \boldsymbol{f}(\boldsymbol{x}(t))$

Opinion Dynamics

The social systems equations

Describe through a deterministic or a probabilistic dynamical model how a set of agents combine their private beliefs with observations (direct or indirect) of the peers' beliefs. The goal is understanding the *emergent behavior* of society.

- Agent $i = 1, \ldots, n$ holds an initial opinion $x_i(0)$ on a topic θ. For example, for agent i it can be $x_i(0) = p(\theta|s_i)$ or $x_i(0) = \mathbb{E}[\theta|s_i]$
- There is a network process (social activities) through which these opinions are shared and updated

agents beliefs at time t $\boldsymbol{x}(t) = (x_1(t), \ldots, x_n(t))$

continuous time $\mapsto$ $\dot{\boldsymbol{x}}(t) = \boldsymbol{f}(\boldsymbol{x}(t))$, discrete time $\mapsto$ $\boldsymbol{x}(t+1) = \boldsymbol{f}(\boldsymbol{x}(t))$

Social Network Model

- The social network is modeled as a directed weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \boldsymbol{W})$
- The weight w_{ij} captures the "trust" node i places in node j.
- **The opinion** of agent i on the kth topic θ_k after the t interactions is $x_i(t;k)$

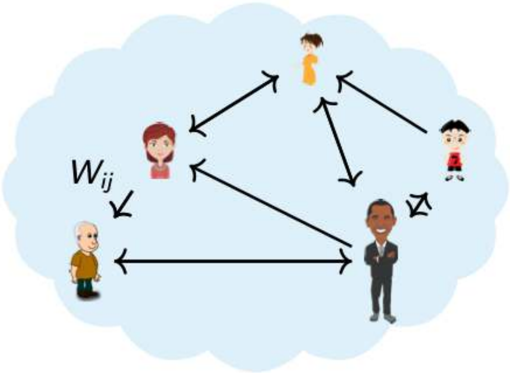

Social Network Model

- The nodes $\mathcal{V} = \{1, ..., n\}$ — are set of agents (or individuals).
- $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ — set of edges (or friendship), with $|\mathcal{E}| \ll n^2$.
- Each edge has a weight W_{ij}. $\boldsymbol{W} \in \mathbb{R}_+^{n \times n}$ — weighted adjacency matrix, with $\boldsymbol{W}\mathbf{1} = \mathbf{1}$ (each row adds up to one).
- We focus on analyzing **steady-state opinions** on each topic k, $\lim_{t \to +\infty} \boldsymbol{x}(t; k)$.

Social Network Model

- The nodes $\mathcal{V} = \{1, ..., n\}$ — are set of agents (or individuals).
- $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ — set of edges (or friendship), with $|\mathcal{E}| \ll n^2$.
- Each edge has a weight W_{ij}. $\boldsymbol{W} \in \mathbb{R}_+^{n \times n}$ — weighted adjacency matrix, with $\boldsymbol{W}\mathbf{1} = \mathbf{1}$ (each row adds up to one).
- We focus on analyzing **steady-state opinions** on each topic k, $\lim_{t \to +\infty} \boldsymbol{x}(t; k)$.

Classes of Opinion Dynamics models

- Alignment of social behavior:
  - Consensus models
    - Herding behavior (Economics)
    - Fad or trend behavior (Social Psychology)
    - Bandwagon effect (political science)
- Disagreement, polarization
  - Social influence models for “mavens” or “stubborn” agents
  - Bounded confidence models

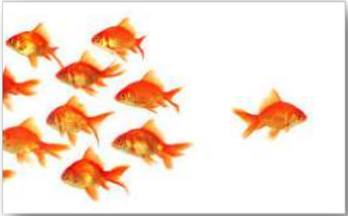

"Vox populi"

Cornell University

- The kth topic θ_k is a Gaussian zero mean random variable with variance σ_o ($\theta_k \sim \mathcal{N}(0,\sigma_o^2)$)
- The private signals signal are $s_i = \theta_k + w_i$ where $w_i \sim \mathcal{N}(0, \sigma_i^2)$. For simplicity $\sigma_i = 1$
- Agent i initial opinion is $x_i(0) = s_i$
- Communications are in a sequence

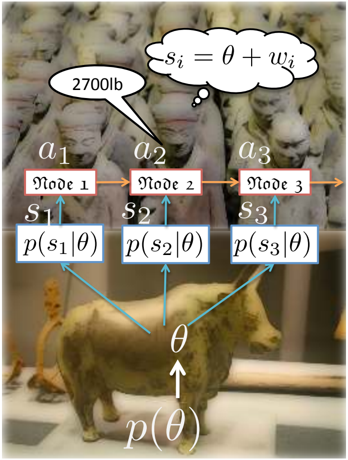

"Vox populi"

Cornell University

- The kth topic θ_k is a Gaussian zero mean random variable with variance σ_o ($\theta_k \sim \mathcal{N}(0,\sigma_o^2)$)
- The private signals signal are $s_i = \theta_k + w_i$ where $w_i \sim \mathcal{N}(0, \sigma_i^2)$. For simplicity $\sigma_i = 1$
- Agent i initial opinion is $x_i(0) = s_i$
- Communications are in a sequence

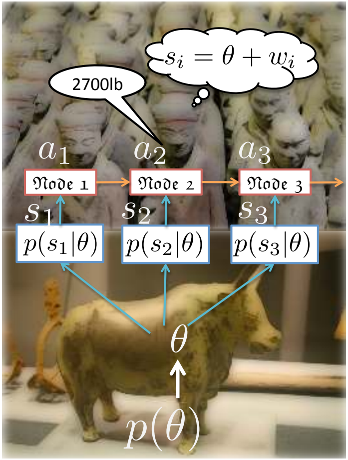

Optimum weights

Cornell University

- Let us ignore the index k, of the topic. If the private signals signal are $s_i = \theta + w_i$ where $w_i \sim \mathcal{N}(0,1)$ and $\theta \sim \mathcal{N}(0,\sigma_o^2)$, then $s_i \sim\sim \mathcal{N}(0,\sigma_o^2+1)$.
- Let $x(t)$ be the estimate of θ that the tth node in the sequence make (node-t opinion): this estimate $x(t)$ is a function of its own private information s_t and the estimate the previous node in the sequence declared $x(t-1)$:

$$x(t) = G(s_t, x(t-1))$$

Optimum weights

Cornell University

- Which function? A reasonable choice for $G(s_t, x(t-1))$ that it is the minimum mean squared error estimate (MMSE) of θ
- For this Gaussian case, the function $G(s_t, x(t-1))$ MMSE estimate is the conditional expectation:

$$x(t) = \mathbb{E}[\theta | s_t, x(t-1)]$$

- As it turns out:

$$x(t) = \mathbb{E}[\theta | s_t, x(t-1)] = (1 - W_{t,t-1})s_t + W_{t,t-1}x(t-1;k)$$

where $W_{i,j} = \frac{t-1}{t}$ for $(i,j) = (t,t-1)$, else $W_{i,j} = 0$

Optimal sequential update

Cornell University

- This implies that optimum update to minimize the error is similar to a consensus update, if we assume that s_t is what the nodes want to average:

$$x(t) = (1 - W_{t,t-1})s_t + W_{t,t-1}x(t-1).$$

- The opinion converges to the true value in the mean square sense, because the variance decreases to zero as $O(1/t)$

$$x(t) \sim \mathcal{N}\left(\theta, \frac{\sigma_o^2}{t}\right)$$

Optimal sequential update

Cornell University

- This implies that optimum update to minimize the error is similar to a consensus update, if we assume that s_t is what the nodes want to average:

$$x(t) = (1 - W_{t,t-1})s_t + W_{t,t-1}x(t-1).$$

- The opinion converges to the true value in the mean square sense, because the variance decreases to zero as $O(1/t)$

$$x(t) \sim \mathcal{N}\left(\theta, \frac{\sigma_o^2}{t}\right)$$

Opinion Dynamics as Network Consensus

Cornell University

Consider a network of individuals exchanging opinions.

- Node i represents a person
- $x_i(t)$ represents the opinion of person i at time t
- Edges represent influence relationships

A simple opinion update model is

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

Interpretation:

- W_{ij} measures how much person i trusts person j
- Each person updates their belief as a weighted average of neighbors

This is exactly the **DeGroot opinion dynamics model**.

Opinion Dynamics as Network Consensus

Consider a network of individuals exchanging opinions.

- Node i represents a person
- $x_i(t)$ represents the opinion of person i at time t
- Edges represent influence relationships

A simple opinion update model is

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

Interpretation:

- W_{ij} measures how much person i trusts person j
- Each person updates their belief as a weighted average of neighbors

This is exactly the **DeGroot opinion dynamics model**.

Opinion Dynamics as Network Consensus

Consider a network of individuals exchanging opinions.

- Node i represents a person
- $x_i(t)$ represents the opinion of person i at time t
- Edges represent influence relationships

A simple opinion update model is

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

Interpretation:

- W_{ij} measures how much person i trusts person j
- Each person updates their belief as a weighted average of neighbors

This is exactly the **DeGroot opinion dynamics model**.

Opinion Dynamics as Network Consensus

Consider a network of individuals exchanging opinions.

- Node i represents a person
- $x_i(t)$ represents the opinion of person i at time t
- Edges represent influence relationships

A simple opinion update model is

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

Interpretation:

- W_{ij} measures how much person i trusts person j
- Each person updates their belief as a weighted average of neighbors

This is exactly the **DeGroot opinion dynamics model**.

DeGroot Opinion Dynamics

Cornell University

- DeGroot introduced this model for opinion dynamics in the 70s. Topics of discussions indexed by $k = 1, ..., K$, each iteration is analogous to average consensus

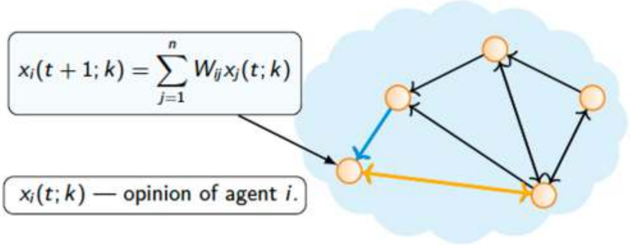

- Agents takes a weighted average of neighbors' opinions (recall $\boldsymbol{W}\mathbf{1} = \mathbf{1}$).

DeGroot Opinion Dynamics

- The DeGroot model is identical to Average Consensus Gossiping (ACG) we already have studied, we just used $\boldsymbol{W}$ as the weight matrix instead of $\boldsymbol{W}$.
- Its convergence is geometric (these are linear dynamics, $\lambda_1(\boldsymbol{W}) = 1$) and the rate is the second largest eigenvalue $\lambda_2(\boldsymbol{W}) < 1$.

Fact: *if G is **strongly connected** and $\boldsymbol{W}\mathbf{1} = \mathbf{1}$, $\lambda_1(\boldsymbol{W}) = 1$, $\lambda_2(\boldsymbol{W}) < 1$*

$$\lim_{t\to\infty} \boldsymbol{x}(t;k) = c(k)^{\star}\mathbf{1} \Longrightarrow \textit{consensus is reached.}$$

DeGroot Opinion Dynamics

Cornell University

- The DeGroot model is identical to Average Consensus Gossiping (ACG) we already have studied, we just used $\boldsymbol{W}$ as the weight matrix instead of $\boldsymbol{W}$.
- Its convergence is geometric (these are linear dynamics, $\lambda_1(\boldsymbol{W}) = 1$) and the rate is the second largest eigenvalue $\lambda_2(\boldsymbol{W}) < 1$.

Fact: *if G is **strongly connected** and $\boldsymbol{W}\mathbf{1} = \mathbf{1}$, $\lambda_1(\boldsymbol{W}) = 1$, $\lambda_2(\boldsymbol{W}) < 1$*

$$\lim_{t\to\infty} \boldsymbol{x}(t;k) = c(k)^{\star}\mathbf{1} \Longrightarrow \textit{consensus is reached.}$$

DeGroot Opinion Dynamics

Cornell University

- The DeGroot model is identical to Average Consensus Gossiping (ACG) we already have studied, we just used $\boldsymbol{W}$ as the weight matrix instead of $\boldsymbol{W}$.
- Its convergence is geometric (these are linear dynamics, $\lambda_1(\boldsymbol{W}) = 1$) and the rate is the second largest eigenvalue $\lambda_2(\boldsymbol{W}) < 1$.

Fact: *if G is* ***strongly connected*** *and* $\boldsymbol{W}\mathbf{1} = \mathbf{1}$, $\lambda_1(\boldsymbol{W}) = 1$, $\lambda_2(\boldsymbol{W}) < 1$

$$\lim_{t \to \infty} \boldsymbol{x}(t;k) = c(k)^{\star}\mathbf{1} \Longrightarrow \textit{consensus is reached.}$$

Long term influence Influence consensus opinion

Recall that the DeGroot model $\boldsymbol{x}(t+1) = \boldsymbol{W}\boldsymbol{x}(t)$ has generally a directed network. If the network is strongly connected,

$$\lim_{t\to\infty} \boldsymbol{x}(t) = \mathbf{1}\,\overbrace{\boldsymbol{\pi}^T\boldsymbol{x}(0)}^{c}, \quad \text{where} \quad \boldsymbol{\pi}^T\boldsymbol{W} = \boldsymbol{\pi}^T$$

Hence $c = \sum_{n=1}^{N} \pi_i x_i(0)$ (the larger π_i the larger the influence of that agent). Interpretation:

- π_i measures the long-run influence of agent i
- agents with larger π_i have more impact on the final opinion

The vector π is the **stationary distribution of a random walk on the network** and is related to the eigenvector centrality and page-rank algorithm

Long term influence Influence consensus opinion

Recall that the DeGroot model $\boldsymbol{x}(t+1) = \boldsymbol{W}\boldsymbol{x}(t)$ has generally a directed network. If the network is strongly connected,

$$\lim_{t\to\infty} \boldsymbol{x}(t) = \mathbf{1}\, \overbrace{\boldsymbol{\pi}^T \boldsymbol{x}(0)}^{c}, \quad \text{where} \quad \boldsymbol{\pi}^T \boldsymbol{W} = \boldsymbol{\pi}^T$$

Hence $c = \sum_{n=1}^{N} \pi_i x_i(0)$ (the larger π_i the larger the influence of that agent). Interpretation:

- π_i measures the long-run influence of agent i
- agents with larger π_i have more impact on the final opinion

The vector π is the **stationary distribution of a random walk on the network** and is related to the eigenvector centrality and page-rank algorithm

Long term influence Influence consensus opinion

Recall that the DeGroot model $\boldsymbol{x}(t+1) = \boldsymbol{W}\boldsymbol{x}(t)$ has generally a directed network. If the network is strongly connected,

$$\lim_{t\to\infty} \boldsymbol{x}(t) = \mathbf{1}\, \overbrace{\boldsymbol{\pi}^T \boldsymbol{x}(0)}^{c}, \quad \text{where} \quad \boldsymbol{\pi}^T \boldsymbol{W} = \boldsymbol{\pi}^T$$

Hence $c = \sum_{n=1}^{N} \pi_i x_i(0)$ (the larger π_i the larger the influence of that agent). Interpretation:

- π_i measures the long-run influence of agent i
- agents with larger π_i have more impact on the final opinion

The vector π is the **stationary distribution of a random walk on the network** and is related to the eigenvector centrality and page-rank algorithm

Evolution of the opinions

- The more meshed the network is, the lower is $0 \leq \lambda_2 \leq 1$. The algebraic connectivity of the graph is $1 - \lambda_2(\boldsymbol{W})$.

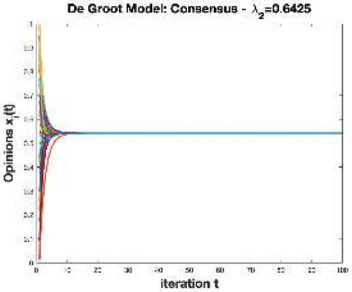

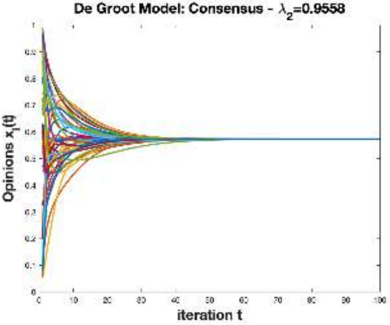

- A variant of the model is the randomized consensus

Randomizing the opinion exchange

Cornell University

- Pairs of neighbors wake up at random, exchange and then update their opinions to the average of their current ones, i.e.:

$$x_i(t+1;k) = x_j(t+1;k) = 0.5x_i(t;k) + 0.5x_j(t;k)$$

Note the trust matrix is switching randomly (I.e. $\boldsymbol{W}(t)\mathbf{1} = \mathbf{1}$).

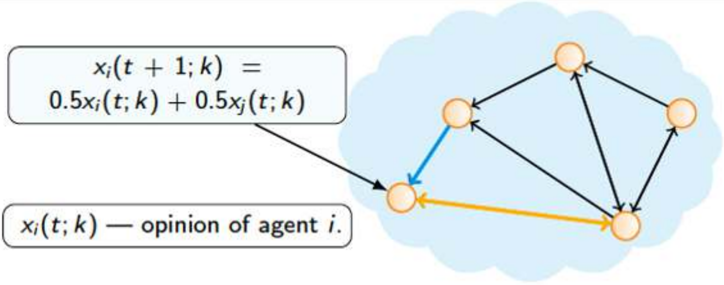

Convergence of Randomized Interaction

- If the process $\boldsymbol{W}(t)$ is stationary (e.g. the pair choice is i.i.d. vs t) in expectation the dynamics look exactly the same. Let $\overline{\boldsymbol{x}}(t;k) = \mathbb{E}[\boldsymbol{x}(t;k)|\theta_k]$ and $\overline{\boldsymbol{W}} = \mathbb{E}[\boldsymbol{W}(t)]$

$$\overline{\boldsymbol{x}}(t+1;k) = \overline{\boldsymbol{W}}\,\overline{\boldsymbol{x}}(t;k)$$

- In this case if the graph is strongly connected and $\overline{W}\mathbf{1} = \mathbf{1}$ we have almost sure convergence to consensus

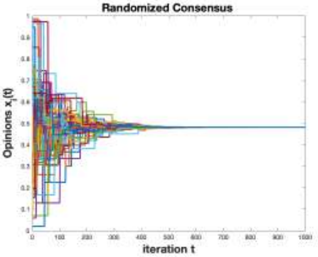

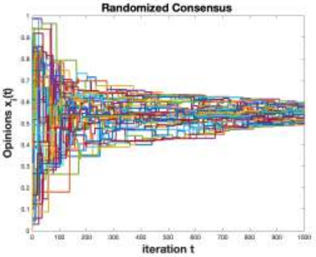

Convergence of Randomized Interaction

Cornell University

- If the process $\boldsymbol{W}(t)$ is stationary (e.g. the pair choice is i.i.d. vs t) in expectation the dynamics look exactly the same. Let $\overline{\boldsymbol{x}}(t;k) = \mathbb{E}[\boldsymbol{x}(t;k)|\theta_k]$ and $\overline{\boldsymbol{W}} = \mathbb{E}[\boldsymbol{W}(t)]$

$$\overline{\boldsymbol{x}}(t+1;k) = \overline{\boldsymbol{W}}\,\overline{\boldsymbol{x}}(t;k)$$

- In this case if the graph is strongly connected and $\overline{\boldsymbol{W}}\mathbf{1} = \mathbf{1}$ we have almost sure convergence to consensus

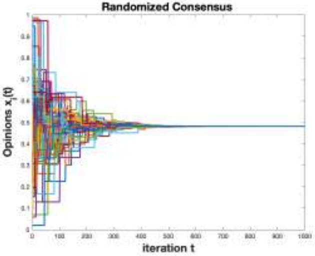

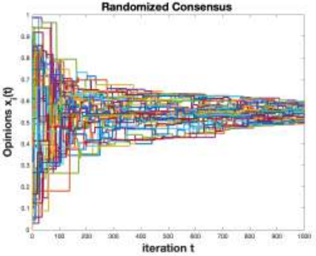

Disagreement and polarization

Cornell University

- Consensus is more the exception than the rule (polarization)

Social network data / opinion:

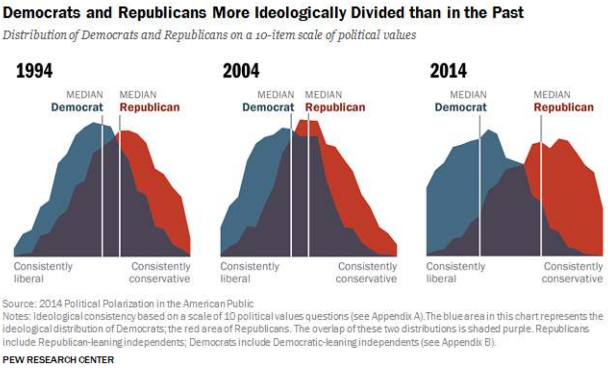

- **Polarization** is common in social networks' opinions (democrats vs. republicans)

DeGroot Model with Stubborn Agents

Cornell University

- Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \boldsymbol{W})$ to include S stubborn agents with $S \ll n$.
- **Stubborn agents** are agents whose opinions do not change, e.g., the leaders, guiding the rest towards different actions. Mathematically, we have:

$$x_i(t+1;k) = x_i(t;k), \ \forall\ t, \text{ if agent } i \text{ is stubborn.}$$

- *DeGroot Opinion Dynamics* — let us partition $\boldsymbol{x}(t;k) = (\boldsymbol{z}(t;k), \boldsymbol{y}(t;k))$

$$\begin{pmatrix} \boldsymbol{z}(t+1;k) \\ \boldsymbol{y}(t+1;k) \end{pmatrix} = \underbrace{\begin{pmatrix} \boldsymbol{I} & \boldsymbol{0} \\ \boldsymbol{B} & \boldsymbol{D} \end{pmatrix}}_{\text{weighted adjacency matrix } \boldsymbol{W}} \begin{pmatrix} \boldsymbol{z}(t;k) \\ \boldsymbol{y}(t;k) \end{pmatrix}$$

DeGroot Model with Stubborn Agents

Cornell University

- Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \boldsymbol{W})$ to include S stubborn agents with $S \ll n$.
- **Stubborn agents** are agents whose opinions do not change, e.g., the leaders, guiding the rest towards different actions. Mathematically, we have:

$$x_i(t+1;k) = x_i(t;k), \ \forall\ t, \text{ if agent } i \text{ is stubborn.}$$

- *DeGroot Opinion Dynamics* — let us partition $\boldsymbol{x}(t;k) = (\boldsymbol{z}(t;k), \boldsymbol{y}(t;k))$

$$\begin{pmatrix} \boldsymbol{z}(t+1;k) \\ \boldsymbol{y}(t+1;k) \end{pmatrix} = \underbrace{\begin{pmatrix} \boldsymbol{I} & \boldsymbol{0} \\ \boldsymbol{B} & \boldsymbol{D} \end{pmatrix}}_{\text{weighted adjacency matrix } \boldsymbol{W}} \begin{pmatrix} \boldsymbol{z}(t;k) \\ \boldsymbol{y}(t;k) \end{pmatrix}$$

DeGroot Model with Stubborn Agents

- Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \boldsymbol{W})$ to include S stubborn agents with $S \ll n$.
- **Stubborn agents** are agents whose opinions do not change, e.g., the leaders, guiding the rest towards different actions. Mathematically, we have:

$$x_i(t+1;k) = x_i(t;k), \ \forall \ t, \ \text{if agent } i \text{ is stubborn.}$$

- *DeGroot Opinion Dynamics* — let us partition $\boldsymbol{x}(t;k) = (\boldsymbol{z}(t;k), \boldsymbol{y}(t;k))$

$$\begin{pmatrix} \boldsymbol{z}(t+1;k) \\ \boldsymbol{y}(t+1;k) \end{pmatrix} = \underbrace{\begin{pmatrix} \boldsymbol{I} & \boldsymbol{0} \\ \boldsymbol{B} & \boldsymbol{D} \end{pmatrix}}_{\text{weighted adjacency matrix } \boldsymbol{W}} \begin{pmatrix} \boldsymbol{z}(t;k) \\ \boldsymbol{y}(t;k) \end{pmatrix}$$

DeGroot Model with Stubborn Agents

Cornell University

$$\begin{pmatrix} \boldsymbol{z}(t+1;k) \\ \boldsymbol{y}(t+1;k) \end{pmatrix} = \underbrace{\begin{pmatrix} \boldsymbol{I} & \boldsymbol{0} \\ \boldsymbol{B} & \boldsymbol{D} \end{pmatrix}}_{\text{weighted adjacency matrix } \boldsymbol{W}} \begin{pmatrix} \boldsymbol{z}(t;k) \\ \boldsymbol{y}(t;k) \end{pmatrix}$$

If we write each block explicitely:

$$\boldsymbol{z}(t+1;k) = \boldsymbol{z}(t;k)$$

$$\boldsymbol{y}(t+1;k) = \boldsymbol{D}\boldsymbol{y}(t;k) + \boldsymbol{B}\boldsymbol{z}(t;k)$$ is an affine mapping!

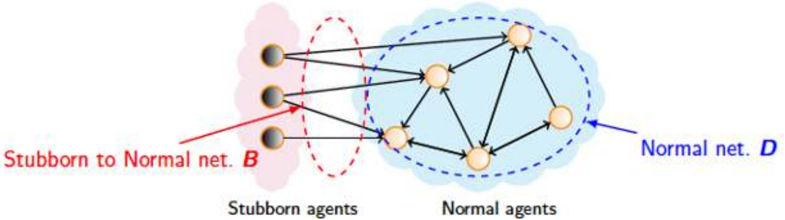

DeGroot Model with Stubborn Agents

Cornell University

$$\begin{pmatrix} \boldsymbol{z}(t+1;k) \\ \boldsymbol{y}(t+1;k) \end{pmatrix} = \underbrace{\begin{pmatrix} \boldsymbol{I} & \boldsymbol{0} \\ \boldsymbol{B} & \boldsymbol{D} \end{pmatrix}}_{\text{weighted adjacency matrix } \boldsymbol{W}} \begin{pmatrix} \boldsymbol{z}(t;k) \\ \boldsymbol{y}(t;k) \end{pmatrix}$$

If we write each block explicitely:

$$\begin{aligned} \boldsymbol{z}(t+1;k) &= \boldsymbol{z}(t;k) \\ \boldsymbol{y}(t+1;k) &= \boldsymbol{D}\boldsymbol{y}(t;k) + \boldsymbol{B}\boldsymbol{z}(t;k) \end{aligned}$$

is an affine mapping!

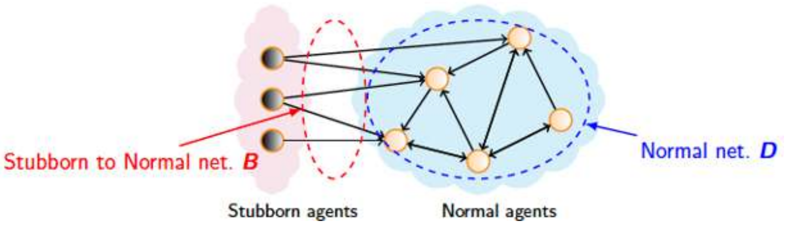

DeGroot Model with Stubborn Agents

Cornell University

- Let $\boldsymbol{x}(t;k) = (\boldsymbol{z}(t;k); \boldsymbol{y}(t;k))$ such that $\boldsymbol{y}(t;k)$ corresponds to the normal agents.

Steady-state: *if sub-graph $\mathcal{V} \setminus \mathcal{S}]$ is* ***strongly connected*** *(+other conditions we discussed for affine dynamics):*

$$\boldsymbol{y}^k := \lim_{t\to\infty} \boldsymbol{y}(t;k) = (\boldsymbol{I} - \boldsymbol{D})^{-1} \boldsymbol{B} \boldsymbol{z}^k.$$

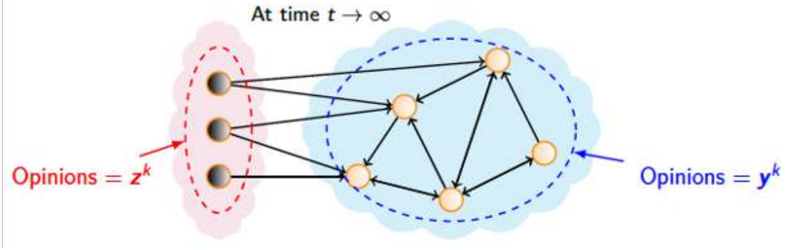

DeGroot Model with Stubborn Agents

- Let $\boldsymbol{x}(t;k) = (\boldsymbol{z}(t;k); \boldsymbol{y}(t;k))$ such that $\boldsymbol{y}(t;k)$ corresponds to the normal agents.

Steady-state: *if sub-graph $\mathcal{V} \setminus \mathcal{S}]$ is* ***strongly connected*** *(+other conditions we discussed for affine dynamics):*

$$\boldsymbol{y}^k := \lim_{t\to\infty} \boldsymbol{y}(t;k) = (\boldsymbol{I} - \boldsymbol{D})^{-1} \boldsymbol{B} \boldsymbol{z}^k.$$

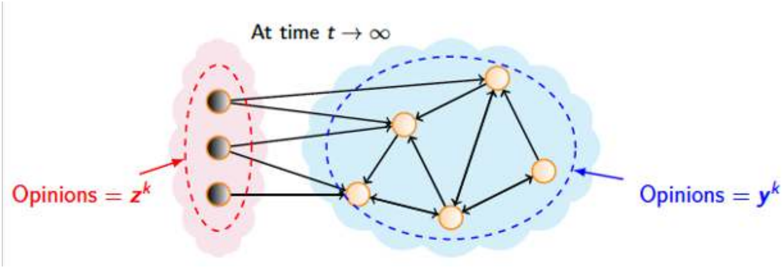

Evolution of the opinion

Cornell University

- The opinions $\boldsymbol{y}^k$ is an output of the linear transformation '$(\boldsymbol{I} - \boldsymbol{D})^{-1}\boldsymbol{B}$.
- The observations $\{\boldsymbol{y}^k, \boldsymbol{z}^k\}_{k=1}^K$ is at most rank $\min\{S, K\} \Longrightarrow$ rank of data $= S =$ number of stubborn agents.
- In general there is no consensus. It occurs if there is only one stubborn agent, or all stubborn agents have the same opinion.

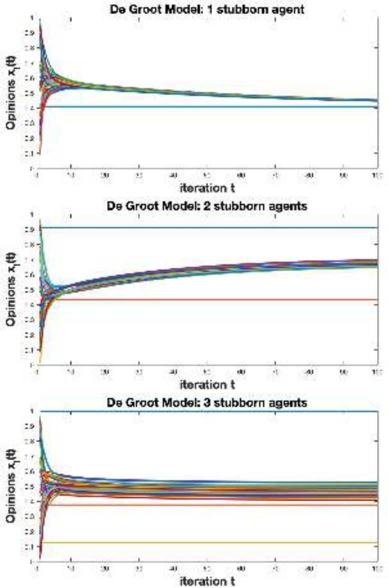

Evolution of the opinion

Cornell University

- The opinions $\boldsymbol{y}^k$ is an output of the linear transformation '$(\boldsymbol{I}-\boldsymbol{D})^{-1}\boldsymbol{B}$.
- The observations $\{\boldsymbol{y}^k, \boldsymbol{z}^k\}_{k=1}^{K}$ is at most rank $\min\{S, K\} \Longrightarrow$ rank of data $= S =$ number of stubborn agents.
- In general there is no consensus. It occurs if there is only one stubborn agent, or all stubborn agents have the same opinion.

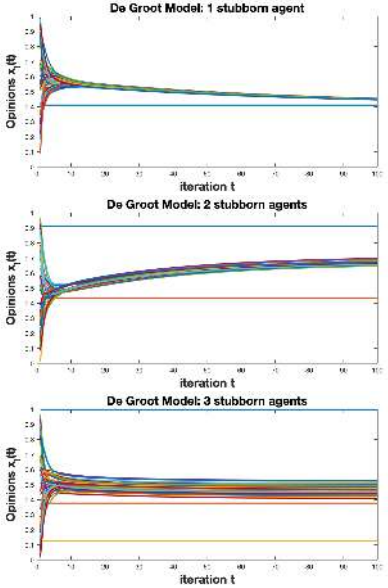

Evolution of the opinion

- The opinions $\boldsymbol{y}^k$ is an output of the linear transformation '$(\boldsymbol{I} - \boldsymbol{D})^{-1}\boldsymbol{B}$.
- The observations $\{\boldsymbol{y}^k, \boldsymbol{z}^k\}_{k=1}^K$ is at most rank $\min\{S, K\} \Longrightarrow$ rank of data $= S =$ number of stubborn agents.
- In general there is no consensus. It occurs if there is only one stubborn agent, or all stubborn agents have the same opinion.

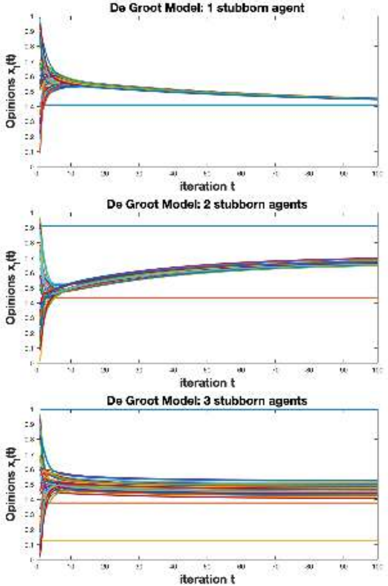

Stubborn agents and affine systems

Cornell University

- The model is a **discrete affine linear update** model we studied in the last lecture.

$$\boldsymbol{x}(t+1) = \boldsymbol{W}\boldsymbol{x}(t) + \boldsymbol{b}$$

- Here we started from non-affine dynamics, but ended up with an analogous equation for the **non-stubborn agents dynamics**:

$$\boldsymbol{y}(t+1;k) = \overbrace{\boldsymbol{D}}^{\text{equivalent to } \boldsymbol{W}} \boldsymbol{y}(t;k) + \overbrace{\boldsymbol{B}\,\boldsymbol{z}(t;k)}^{\text{equivalent to } \boldsymbol{b}}$$

- The steady state solution is consistent with $\boldsymbol{x}^* = (\boldsymbol{I} - \boldsymbol{W})^{-1}\boldsymbol{b}$ which we said is the limit of affine linear systems

$$\boldsymbol{y}^k := \lim_{t\to\infty} \boldsymbol{y}(t;k) = (\boldsymbol{I} - \boldsymbol{D})^{-1}\boldsymbol{B}\boldsymbol{z}^k$$

Stubborn agents and affine systems

Cornell University

- The model is a **discrete affine linear update** model we studied in the last lecture.

$$\boldsymbol{x}(t+1) = \boldsymbol{W}\boldsymbol{x}(t) + \boldsymbol{b}$$

- Here we started from non-affine dynamics, but ended up with an analogous equation for the **non-stubborn agents dynamics**:

$$\boldsymbol{y}(t+1;k) = \overbrace{\boldsymbol{D}}^{\text{equivalent to } \boldsymbol{W}} \ \boldsymbol{y}(t;k) + \overbrace{\boldsymbol{B}\ \boldsymbol{z}(t;k)}^{\text{equivalent to } \boldsymbol{b}}$$

- The steady state solution is consistent with $\boldsymbol{x}^* = (\boldsymbol{I} - \boldsymbol{W})^{-1}\boldsymbol{b}$ which we said is the limit of affine linear systems

$$\boldsymbol{y}^k := \lim_{t\to\infty} \boldsymbol{y}(t;k) = (\boldsymbol{I} - \boldsymbol{D})^{-1}\boldsymbol{B}\boldsymbol{z}^k$$

Stubborn agents and affine systems

Cornell University

- The model is a **discrete affine linear update** model we studied in the last lecture.

$$\boldsymbol{x}(t+1) = \boldsymbol{W}\boldsymbol{x}(t) + \boldsymbol{b}$$

- Here we started from non-affine dynamics, but ended up with an analogous equation for the **non-stubborn agents dynamics**:

$$\boldsymbol{y}(t+1;k) = \overbrace{\boldsymbol{D}}^{\text{equivalent to } \boldsymbol{W}} \boldsymbol{y}(t;k) + \overbrace{\boldsymbol{B}\ \boldsymbol{z}(t;k)}^{\text{equivalent to } \boldsymbol{b}}$$

- The steady state solution is consistent with $\boldsymbol{x}^* = (\boldsymbol{I} - \boldsymbol{W})^{-1}\boldsymbol{b}$ which we said is the limit of affine linear systems

$$\boldsymbol{y}^k := \lim_{t\to\infty} \boldsymbol{y}(t;k) = (\boldsymbol{I} - \boldsymbol{D})^{-1}\boldsymbol{B}\boldsymbol{z}^k$$

Bounded Confidence and Polarization

Cornell University

- Are stubborn agents necessary to create social division? R. Hegselmann and U. Krause introduced a model that shows polarization can emerge on its own
- $d_{ij}(t;k)$ is the opinion distance of agents i and j
- Agents "talk" iff $d_{ij}(t;k) < \tau \rightarrow$ update is non linear

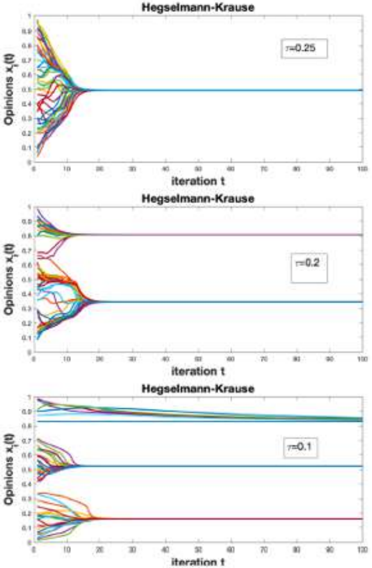

Bounded Confidence and Polarization

Cornell University

- Are stubborn agents necessary to create social division? R. Hegselmann and U. Krause introduced a model that shows polarization can emerge on its own
- $d_{ij}(t;k)$ is the opinion distance of agents i and j
- Agents "talk" iff $d_{ij}(t;k) < \tau \rightarrow$ update is non linear

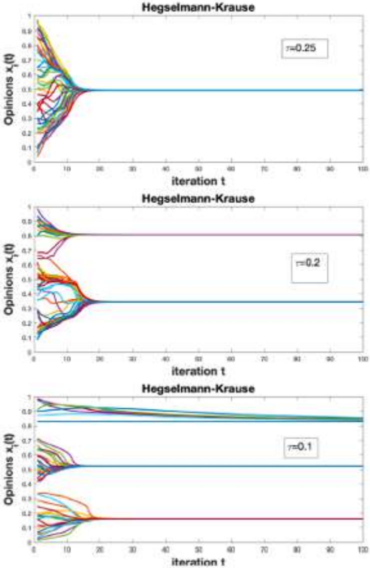

Alternative: Bounded Confidence

Cornell University

- With $u(x)$ the unit step $u(x) = 1, x > 0$ and $u(x) = 0$ else, the HK model update is a non-linear discrete dynamic:

$$x_i(t+1;k) = \sum_{j=1}^{n} W_{ij} u(\tau - d_{ij}(t;k)) x_j(t;k)$$

$$\boldsymbol{x}(t+1;k) = \boldsymbol{W}(\boldsymbol{x}(t;k))\boldsymbol{x}(t;k)$$

- The randomized variant is due to G. Deffuant

$$x_i(t+1;k) = x_i(t;k) + 0.5u(\tau - d_{ij}(t;k))(x_j(t;k) - x_i(t;k))$$

$$x_j(t+1;k) = x_j(t;k) + 0.5u(\tau - d_{ij}(t;k))(x_i(t;k) - x_j(t;k))$$

Alternative: Bounded Confidence

- With $u(x)$ the unit step $u(x) = 1, x > 0$ and $u(x) = 0$ else, the HK model update is a non-linear discrete dynamic:

$$x_i(t+1;k) = \sum_{j=1}^{n} W_{ij} u(\tau - d_{ij}(t;k)) x_j(t;k)$$

$$\boldsymbol{x}(t+1;k) = \boldsymbol{W}(\boldsymbol{x}(t;k))\boldsymbol{x}(t;k)$$

- The randomized variant is due to G. Deffuant

$$x_i(t+1;k) = x_i(t;k) + 0.5u(\tau - d_{ij}(t;k))(x_j(t;k) - x_i(t;k))$$

$$x_j(t+1;k) = x_j(t;k) + 0.5u(\tau - d_{ij}(t;k))(x_i(t;k) - x_j(t;k))$$

Alternative: Bounded Confidence

Cornell University

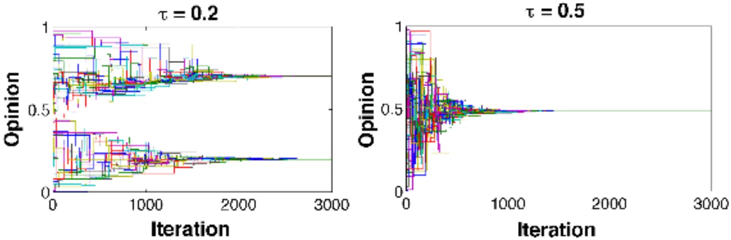

Phase transition to consensus

Let $\overline{W}_{ij}$ probability of the interaction and $d_{ij}(t;k) = |x_i(t;k) - x_j(t;k)|$. A **necessary** condition for the system to converge almost surely to consensus:

$$\tau > \overline{d}_k := \sum_{ij} \overline{W}_{ij} d_{ij}(0;k)$$

Phase transition from consensus to polarization

Let $d_{ij}(t;k) = |x_i(t;k) - x_j(t;k)|$. A phase transition from a population of polarized opinions to consensus. A **necessary** condition for the system to converge almost surely to consensus

$$\tau > \bar{d}_k := \sum_{ij} \overline{W}_{ij} d_{ij}(0;k)$$

$\overline{W}_{ij}$ probability of the interaction.

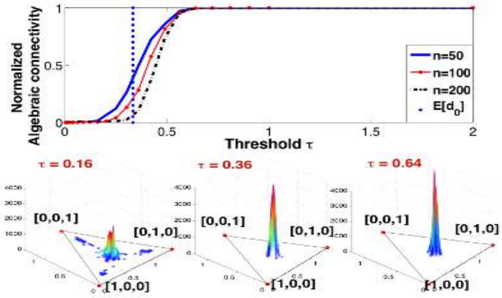

The θ_k in this case has three possible outcomes: the belief is a point in the 3-D simplex. The frequency with which consensus is observed decreases as the threshold decreases.

Outline

Cornell University

2 Opinion dynamics modeling
- Models for opinion diffusion
- Random walk based models

Bounded opinions: Voter model

Cornell University

- Agents have a prior belief $x_i(0;k) = p(\theta_k|s_i)$ on a topic θ_k
- Agents can only observe actions $a_i(t;k)$ (typically binary, i.e. $\in \{1,0\}$) and update their beliefs
- The fact that the observations are quantized makes a significant difference

Variants of the voter model

- Rational agents with bounded opinions - Bikhchandani, Hirshleifer and Welch (BHW)
- Binary opinion dynamics - Peter Clifford and Aidan Sudbury '73 is inspired by Ising model for spin alignment

Both of them lead to consensus or herding

Bounded opinions: Voter model

Cornell University

- Agents have a prior belief $x_i(0;k) = p(\theta_k|s_i)$ on a topic θ_k
- Agents can only observe actions $a_i(t;k)$ (typically binary, i.e. $\in \{1,0\}$) and update their beliefs
- The fact that the observations are quantized makes a significant difference

Variants of the voter model

- Rational agents with bounded opinions - Bikhchandani, Hirshleifer and Welch (BHW)
- Binary opinion dynamics - Peter Clifford and Aidan Sudbury '73 is inspired by Ising model for spin alignment

Both of them lead to consensus or herding

Bounded opinions: Voter model

- Agents have a prior belief $x_i(0;k) = p(\theta_k|s_i)$ on a topic θ_k
- Agents can only observe actions $a_i(t;k)$ (typically binary, i.e. $\in \{1,0\}$) and update their beliefs
- The fact that the observations are quantized makes a significant difference

Variants of the voter model

- Rational agents with bounded opinions - Bikhchandani, Hirshleifer and Welch (BHW)
- Binary opinion dynamics - Peter Clifford and Aidan Sudbury '73 is inspired by Ising model for spin alignment

Both of them lead to consensus or herding

Bounded opinions: Voter model

Cornell University

- Agents have a prior belief $x_i(0;k) = p(\theta_k|s_i)$ on a topic θ_k
- Agents can only observe actions $a_i(t;k)$ (typically binary, i.e. $\in \{1,0\}$) and update their beliefs
- The fact that the observations are quantized makes a significant difference

Variants of the voter model

- Rational agents with bounded opinions - Bikhchandani, Hirshleifer and Welch (BHW)
- Binary opinion dynamics - Peter Clifford and Aidan Sudbury '73 is inspired by Ising model for spin alignment

Both of them lead to consensus or herding

Bounded opinions: Voter model

Cornell University

- Agents have a prior belief $x_i(0;k) = p(\theta_k|s_i)$ on a topic θ_k
- Agents can only observe actions $a_i(t;k)$ (typically binary, i.e. $\in \{1,0\}$) and update their beliefs
- The fact that the observations are quantized makes a significant difference

Variants of the voter model

- Rational agents with bounded opinions - Bikhchandani, Hirshleifer and Welch (BHW)
- Binary opinion dynamics - Peter Clifford and Aidan Sudbury '73 is inspired by Ising model for spin alignment

Both of them lead to consensus or herding

Bounded opinions: Voter model

Cornell University

- Agents have a prior belief $x_i(0; k) = p(\theta_k | s_i)$ on a topic θ_k
- Agents can only observe actions $a_i(t; k)$ (typically binary, i.e. $\in \{1, 0\}$) and update their beliefs
- The fact that the observations are quantized makes a significant difference

Variants of the voter model

- Rational agents with bounded opinions - Bikhchandani, Hirshleifer and Welch (BHW)
- Binary opinion dynamics - Peter Clifford and Aidan Sudbury '73 is inspired by Ising model for spin alignment

Both of them lead to consensus or herding

Herding of rational agents

Cornell University

- Bikhchandani, Hirshleifer and Welch (BHW): θ is randomly determined between the values of 1 ("good state") and 0 ("bad state")
- The agents gain c if $a_i = \theta = 1$ they loose c if $a_i = 1$ and $\theta = 0$ and $a_i = 0$ yields no gain or cost

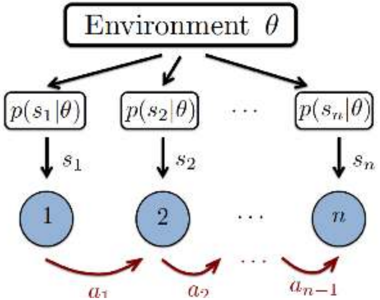

The updates are in sequence. The tth agent to update observed s_t and the actions of the previous agents to decides

$$a(t;k) = \text{argmax}_{\theta \in \{0,1\}} P(\theta | s_t, a(t-1;k), \ldots, a(1;k)) \rightarrow \text{rational choice}$$

Herding of rational agents

Cornell University

- Bikhchandani, Hirshleifer and Welch (BHW): θ is randomly determined between the values of 1 ("good state") and 0 ("bad state")
- The agents gain c if $a_i = \theta = 1$ they loose c if $a_i = 1$ and $\theta = 0$ and $a_i = 0$ yields no gain or cost

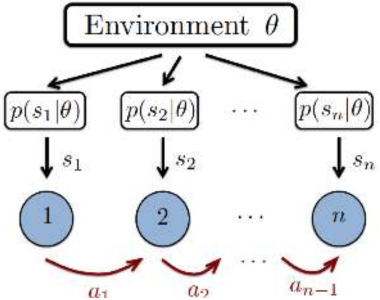

The updates are in sequence. The tth agent to update observed s_t and the actions of the previous agents to decides

$$a(t;k) = \text{argmax}_{\theta\in\{0,1\}} P(\theta|s_t, a(t-1;k), \ldots, a(1;k)) \rightarrow \text{rational choice}$$

Herding of rational agents

Cornell University

Information cascade

The update rule (for cond. independent s_i) for the likelihood ratio is a martingale. In finite iterations the agents *herd* and make the same decision irrespective of their private belief, which has finite probability of being wrong!

Information Cascade Model Setup

- It is a sequential decision-making process where the tth agent choose between two actions $a(t)$, A and B, corresponding to two possible states of the world, $\theta = H = 1$ (high) and $\theta = L = 0$ (low).
- Each agent receives an independent private signal $s(t)$ which can be either $h = 1$ (favoring H) or $l = 0$ (favoring L). For simplicity let us assume a symmetric private signal structure:

$$P(s_t = h|\theta = H) = P(s = l|\theta = L) = p > 0.5.$$

Key point

Each agent observes all past actions but not the private signals of previous agents.

Decision Rule and Information Cascades

- Each agent follows the Maximum A Posteriori (MAP) rule, choosing A for a given history$=(a(t-1),\ldots,a(1))$ and s_t if:

$$P(H|\text{history}_t, s_t) > P(L|\text{history}_t, s_t).$$

- Using Bayes' theorem, the belief update based on past actions and the private signal can be expressed in terms of the log-likelihood ratio (LLR):

$$\Lambda_t = \log\frac{P(H|\text{history}_t)}{P(L|\text{history}_t)} + \log\frac{P(s_t|H)}{P(s_t|L)}. \tag{1}$$

- The agent selects A if $\Lambda_n > 0$, otherwise B is chosen.

First Two Agents Follow Their Signals

The first agent has no prior observations and follows their private signal:

$$a_1 = \begin{cases} A, & \text{if } s_1 = h, \\ B, & \text{if } s_1 = l. \end{cases}$$

The second agent observes a_1 and updates their belief, still considering their own signal.

Third Agent and the Onset of Herding

Cornell University

If the first two agents choose the same action, say A, the third agent updates their belief as:

$$\Lambda_3 = \log \frac{P(H|A,A)}{P(L|A,A)} + \log \frac{P(s_3|H)}{P(s_3|L)}. \tag{2}$$

If $P(H|A,A)$ is high enough, then even if $s_3 = l$, we may have $\Lambda_3 > 0$, leading to the choice of A. This initiates a cascade where subsequent agents ignore their private signals.

Information Loss and Herding

Once an information cascade starts, all subsequent agents choose the same action irrespective of their private signals. This results in potential inefficiency, as incorrect cascades can emerge purely due to early misalignment of private signals.

Voter model

Cornell University

- Introduced by Peter Clifford and Aidan Sudbury analyzed first by Richard A Holley and Thomas M Liggett
- In interacting social agents copy at random the action of a neighbor
- The next action $a_i(t;k)$ is equivalently a Bernoulli random variable:

$$a_i(t;k) \sim \mathcal{B}(x_i(t;k)), \qquad x_i(t+1;k) = \frac{1}{|\mathcal{N}_i|}\sum_{j\in\mathcal{N}_i}^{n} a_j(t;k)$$

Voter model at $t \to \infty$

Consensus: Reached over connected finite graphs and $1-2$D lattices. All actions are almost surely the same, i.e. $t \to \infty$ $\boldsymbol{a}(t;k) = a^\infty$ a.s.. The final collective action can be $a^\infty = 0$ or $a^\infty = 1$. Given $x_i(0;k)$:

$$P(a^\infty = 1) = \frac{1}{n}\mathbf{1}\mathbf{1}^T\boldsymbol{x}(0;k).$$

Voter model

Cornell University

- Introduced by Peter Clifford and Aidan Sudbury analyzed first by Richard A Holley and Thomas M Liggett
- In interacting social agents copy at random the action of a neighbor
- The next action $a_i(t;k)$ is equivalently a Bernoulli random variable:

$$a_i(t;k) \sim \mathcal{B}(x_i(t;k)), \quad x_i(t+1;k) = \frac{1}{|\mathcal{N}_i|} \sum_{j \in \mathcal{N}_i}^{n} a_j(t;k)$$

Voter model at $t \to \infty$

Consensus: Reached over connected finite graphs and $1-2$D lattices. All actions are almost surely the same, i.e. $t \to \infty$ $\boldsymbol{a}(t;k) = a^\infty$ a.s.. The final collective action can be $a^\infty = 0$ or $a^\infty = 1$. Given $x_i(0;k)$:

$$P(a^\infty = 1) = \frac{1}{n} \mathbf{1}\mathbf{1}^T \boldsymbol{x}(0;k).$$

Voter model

Cornell University

- Introduced by Peter Clifford and Aidan Sudbury analyzed first by Richard A Holley and Thomas M Liggett
- In interacting social agents copy at random the action of a neighbor
- The next action $a_i(t;k)$ is equivalently a Bernoulli random variable:

$$a_i(t;k) \sim \mathcal{B}(x_i(t;k)), \qquad x_i(t+1;k) = \frac{1}{|\mathcal{N}_i|}\sum_{j\in\mathcal{N}_i}^{n} a_j(t;k)$$

Voter model at $t \to \infty$

Consensus: Reached over connected finite graphs and $1-2$D lattices. All actions are almost surely the same, i.e. $t \to \infty$ $\boldsymbol{a}(t;k) = a^{\infty}$ a.s.. The final collective action can be $a^{\infty}=0$ or $a^{\infty}=1$. Given $x_i(0;k)$:

$$P(a^{\infty}=1) = \frac{1}{n}\mathbf{1}\mathbf{1}^T\,\boldsymbol{x}(0;k).$$

Voter model

- Introduced by Peter Clifford and Aidan Sudbury analyzed first by Richard A Holley and Thomas M Liggett
- In interacting social agents copy at random the action of a neighbor
- The next action $a_i(t;k)$ is equivalently a Bernoulli random variable:

$$a_i(t;k) \sim \mathcal{B}(x_i(t;k)), \qquad x_i(t+1;k) = \frac{1}{|\mathcal{N}_i|}\sum_{j \in \mathcal{N}_i}^{n} a_j(t;k)$$

Voter model at $t \to \infty$

Consensus: Reached over connected finite graphs and $1-2$D lattices. All actions are almost surely the same, i.e. $t \to \infty$ $\boldsymbol{a}(t;k) = a^{\infty}$ a.s.. The final collective action can be $a^{\infty} = 0$ or $a^{\infty} = 1$. Given $x_i(0;k)$:

$$P(a^{\infty} = 1) = \frac{1}{n}\mathbf{1}\mathbf{1}^T \, \boldsymbol{x}(0;k).$$

Analysis of the voter model

Cornell University

- The key analysis methods is the dual approach: it tracks the origin of an opinion at node i at time T back in time, by viewing it as a random *walk*
- The state of each agent i at time T can be traced back to an opinion jumping from a neighbor to another
- If the origins of opinions of all agents at some in time point coalesce, the walks will continue together

Analysis of the voter model

Cornell University

- The key analysis methods is the dual approach: it tracks the origin of an opinion at node i at time T back in time, by viewing it as a random *walk*
- The state of each agent i at time T can be traced back to an opinion jumping from a neighbor to another
- If the origins of opinions of all agents at some in time point coalesce, the walks will continue together

Analysis of the voter model

Cornell University

- The key analysis methods is the dual approach: it tracks the origin of an opinion at node i at time T back in time, by viewing it as a random *walk*
- The state of each agent i at time T can be traced back to an opinion jumping from a neighbor to another
- If the origins of opinions of all agents at some in time point coalesce, the walks will continue together

Analysis of the voter model

- The number of steps backward when each of the random walks coalesces is the consensus convergence time

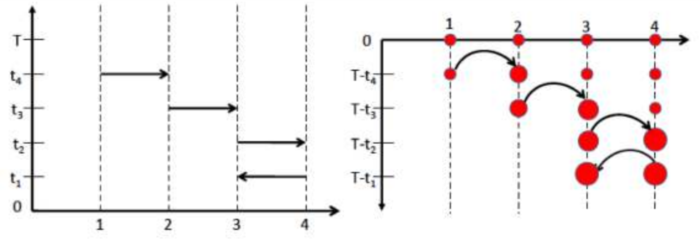

(a) (b)

Voter model with stubborn agents

Cornell University

- No consensus. The network continues to randomly change votes
- Mobilia in 2003 performed a mean field study.

Analysis

If every non-stubborn node in the network has at least one directed path to a stubborn agent votes $\boldsymbol{a}(t;k)$ converges in distribution, i.e. for $t \to +\infty$ $\boldsymbol{a}(t;k) \sim \mathcal{B}(\boldsymbol{x}(+\infty))$.

Voter model with stubborn agents

Cornell University

- No consensus. The network continues to randomly change votes
- Mobilia in 2003 performed a mean field study.

Analysis

If every non-stubborn node in the network has at least one directed path to a stubborn agent votes $a(t;k)$ converges in distribution, i.e. for $t \to +\infty$ $\boldsymbol{a}(t;k) \sim \mathcal{B}(\boldsymbol{x}(+\infty))$.

Voter model with stubborn agents

- No consensus. The network continues to randomly change votes
- Mobilia in 2003 performed a mean field study.

Analysis

If every non-stubborn node in the network has at least one directed path to a stubborn agent votes $\boldsymbol{a}(t;k)$ converges in distribution, i.e. for $t \to +\infty$ $\boldsymbol{a}(t;k) \sim \mathcal{B}(\boldsymbol{x}(+\infty))$.

Stubborn agents position

Cornell University

- Placement of stubborn nodes makes a huge difference on the variance

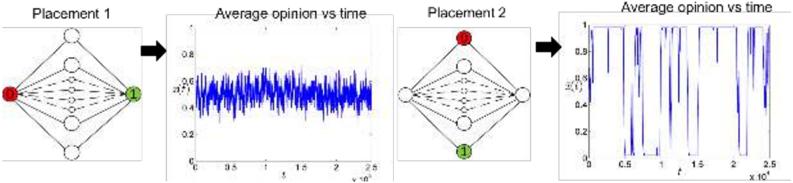

Optimal placement

The problem of maximizing social bias is a sub-modular problem → near optimal solution is greedy.

Stubborn agents position

Cornell University

- Placement of stubborn nodes makes a huge difference on the variance

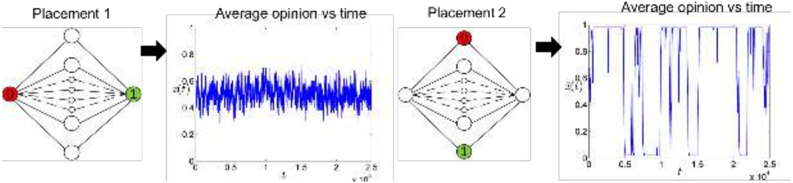

Optimal placement

The problem of maximizing social bias is a sub-modular problem $\rightarrow$ near optimal solution is greedy.

Outline

1 Review

2 Opinion dynamics modeling

3 Conclusions

Concluding remarks

Cornell University

- In this lecture we discussed social networks models for opinion dynamics
- These models were example of linear, non-linear and affine dynamics that we introduced in general before
- We will introduced models for financial contagion in the next lecture.