From Consensus to Social Learning in Complex Networks

Ali Jadbabaie

Skirkanich Associate Professor of Innovation

Department of Electrical and Systems Engineering and GRASP Laboratory

UNIVERSITY of PENNSYLVANIA

Alireza Tahbaz-Salehi (Penn) and Victor Preciado (Penn)

Alvaro Sandroni (Kellogg School, Northwestern Univ)

US Army-commissioned study on Network Science by National Research Council (NAE+NAS)

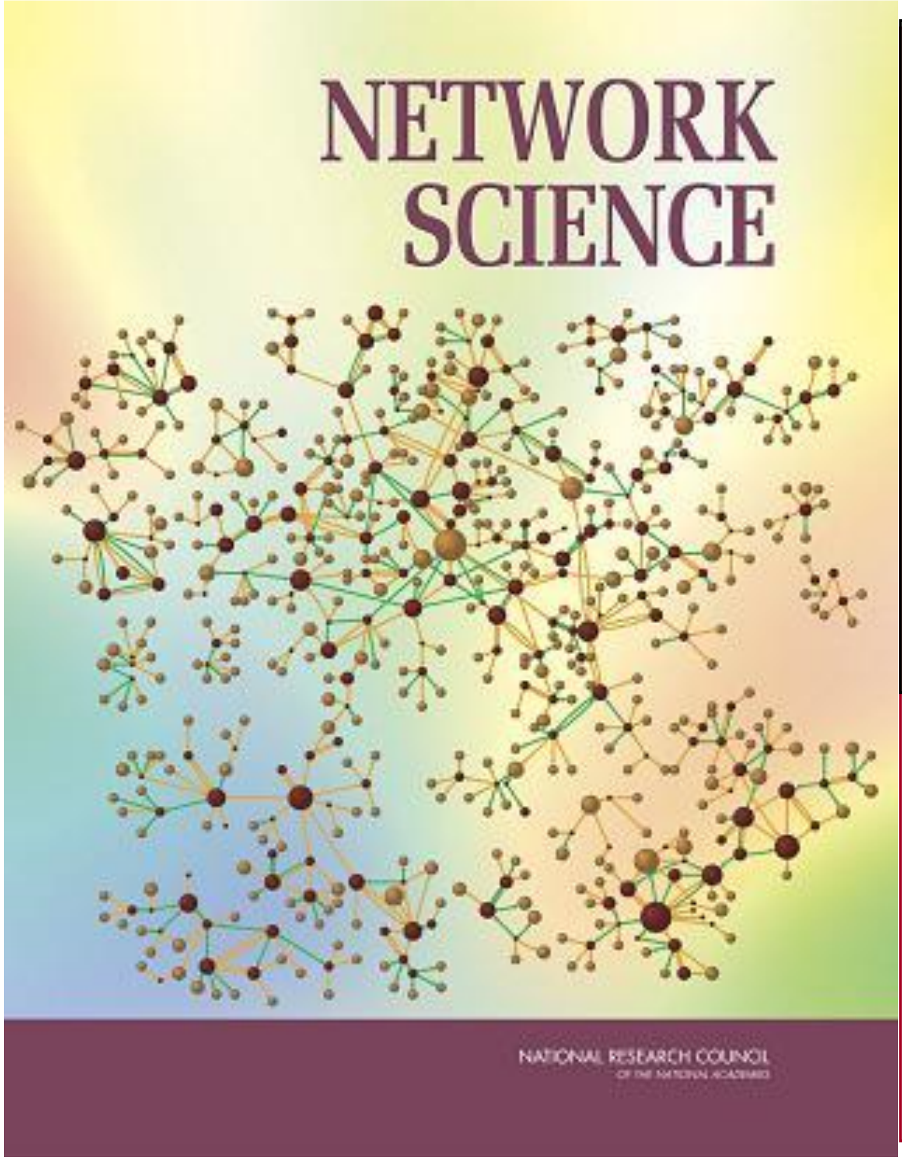

1. *Dynamics, spatial location, and information propagation in networks.*
2. *Modeling and analysis of very large networks.*
3. *Design and synthesis of networks.*
4. *Increasing the level of rigor and mathematical structure.*
5. *Abstracting common concepts across fields.*

6. *Better experiments and measurements of network structure.*
7. *Robustness and security of networks.*

Emergence of Consensus, synchronization, and flocking

VOLUME 75, NUMBER 6 PHYSICAL REVIEW LETTERS 7 AU

Novel Type of Phase Transition in a System of Self-Driv

Tamás Vicsek,[1,2] András Czirók,[1] Eshel Ben-Jacob,[3] In
[1]*Department of Atomic Physics, Eötvös University*

Pretty simulations, few proofs

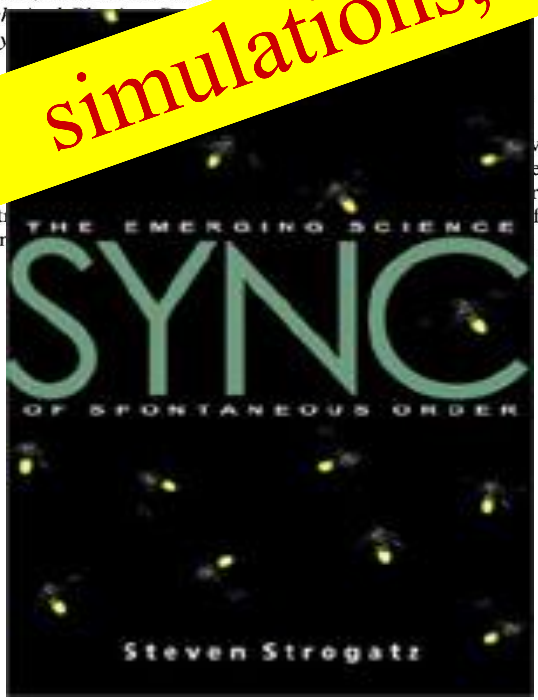

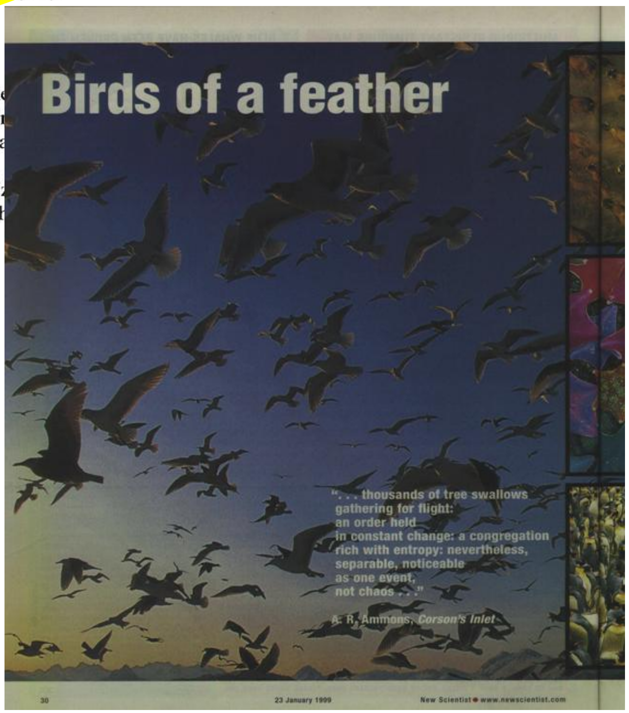

Spectacular Progress in nnderstanding Networked Systems

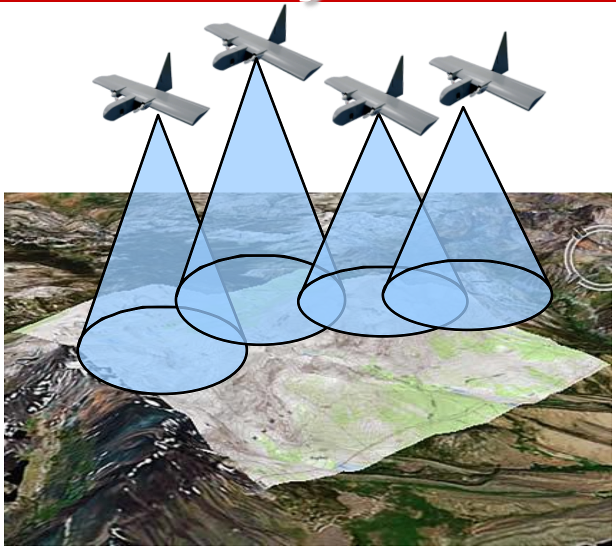

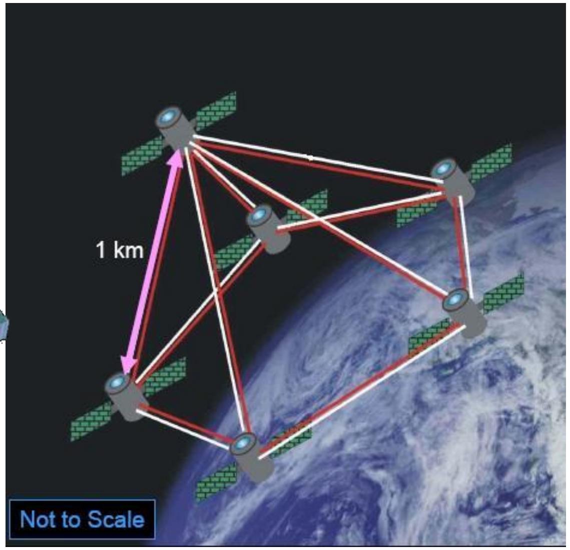

Understanding and predicting collective phenomena

Social and Economic Networks

- Epidemics and Pandemics
- Bubbles
- Bank Runs

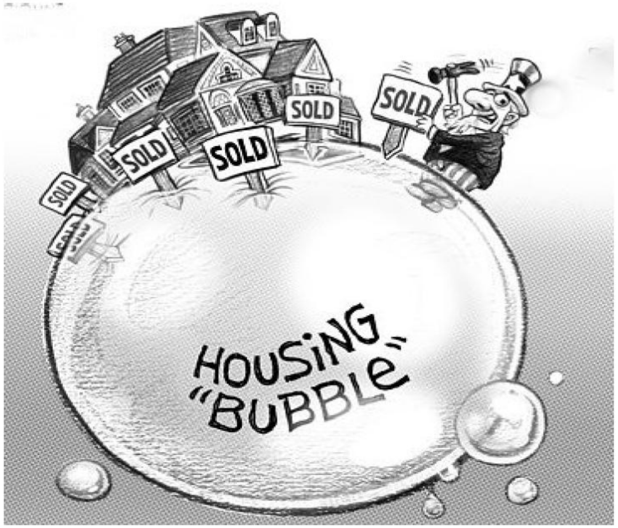

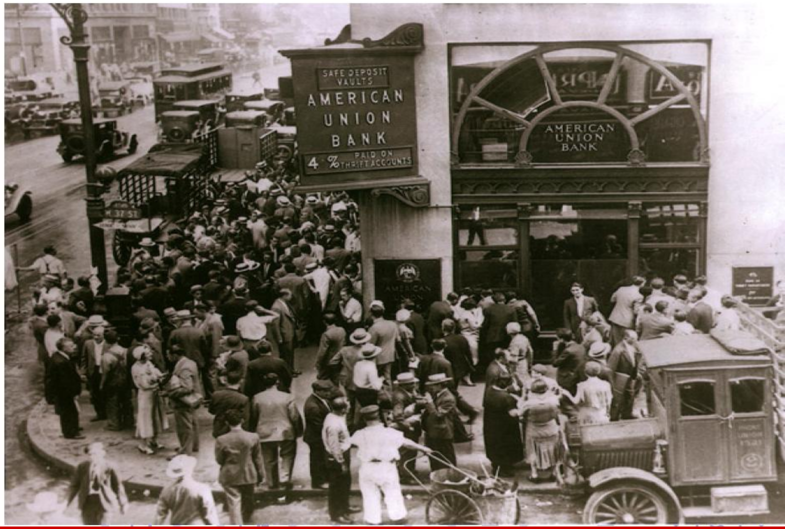

Example: Flocking and Motion Coordination

- How can a group of moving agents collectively decide on direction, based on nearest neighbor interaction?

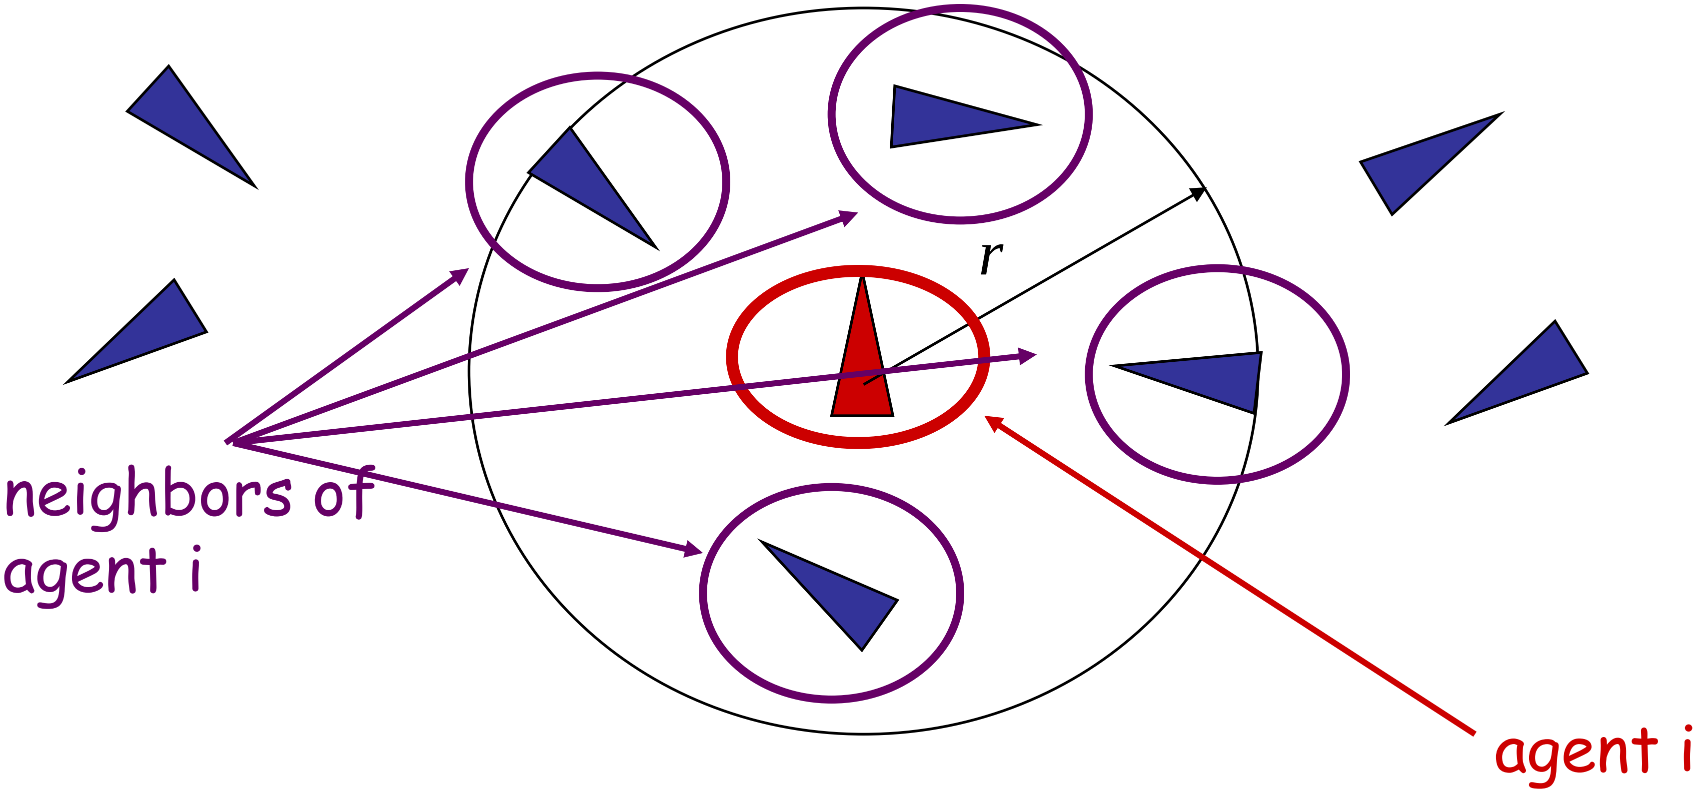

How does <u>global behavior (herding)</u> emerge from <u>local interactions</u>?

An intuitive model (Vicsek' 1995)

The heading value updated (in discrete time) as a **weighted average** of the value of its neighbors: move one step along updated direction

$$\theta_i(k+1) = <\theta_i(k)>_r := \text{atan}\frac{(\sum_{j\in\mathcal{N}_i(k)} \sin\theta_j(k)) + \sin\theta_i(k)}{(\sum_{j\in\mathcal{N}_i(k)} \cos\theta_j(k)) + \cos\theta_i(k)}$$

Locally:

$$\theta_i(k+1) = \frac{1}{d_i(k)+1}\left(\sum_{j\in\mathcal{N}_i(k)} w_{ij}\theta_j(k) + w_{ii}\theta_i(k)\right)$$

Neighborhood relation depends on heading value, resulting in change in topology

MAIN QUESTION : ***When do all headings converge to the same value?***

A network which changes as a result of node dynamics

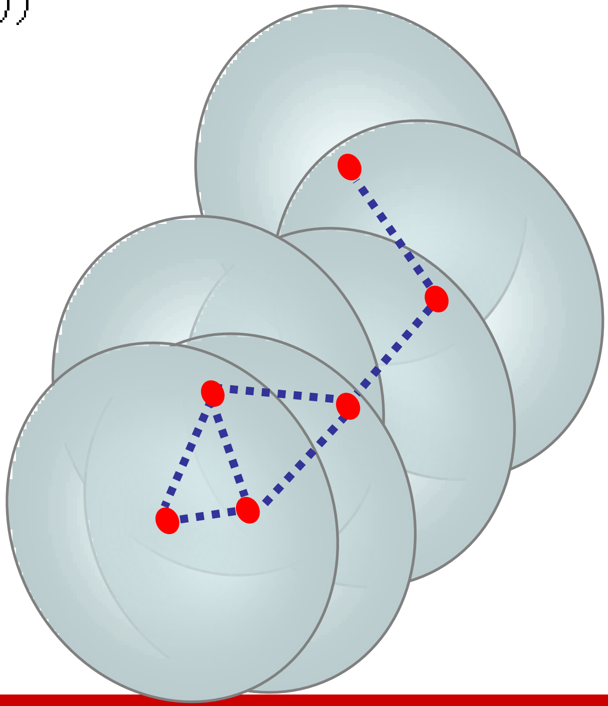

Flocking and opinion dynamics

- Bounded confidence opinion model (*Krause, 2000, Hendrix et al. 2008)*
  - Nodes update their opinions as a weighted average of the opinion value of their friends
  - Friends are those whose opinion is already close (e.g. within 1 unit)
  - When will there be fragmentation and when will there be convergence of opinions?
  - Node dynamics changes topology
  - Vicsek model in 1d
  - Special case: Gossiping: each node only talks to one neighbor at a time
  - Simulations informative but not enough

$$x_i(k) = \sum_{j \in \mathcal{N}_i} W_{ij}(k) x_j(k-1)$$

$$W_{ij}(k) \geq 0, \quad \sum_{j \in \mathcal{N}_i} W_{ij}(k) = 1.$$

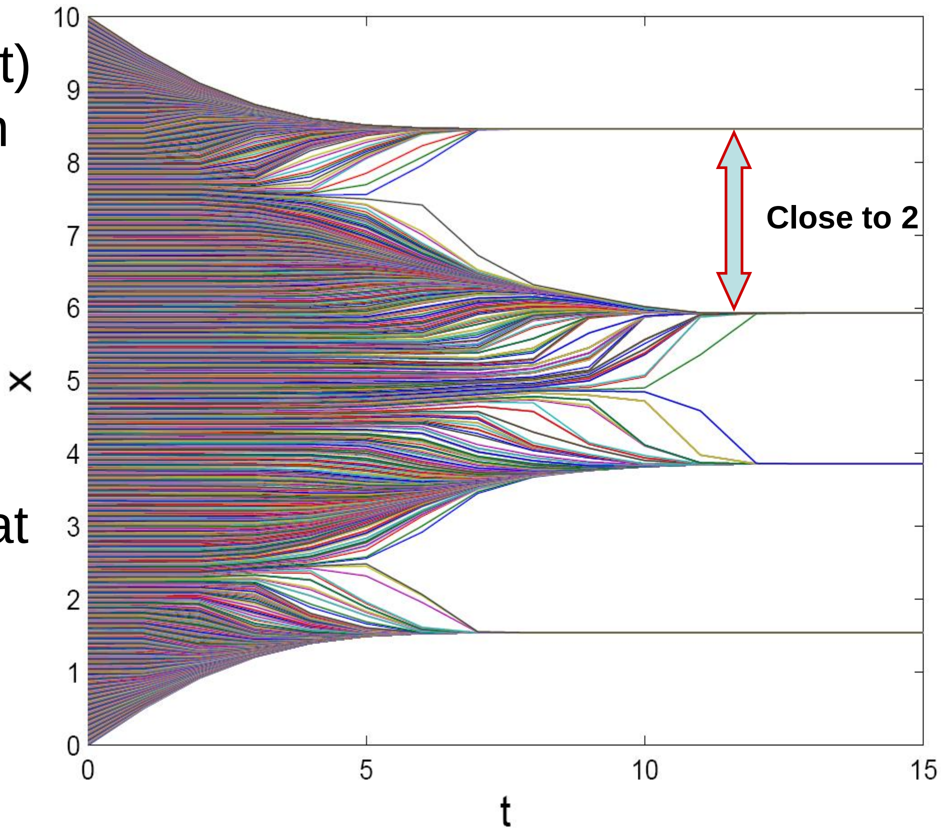

Consensus in changing networks

Theorem (Jadbabaie et al. 2003, Tsitsiklis'84): *If there is a sequence of bounded, non-overlapping time intervals T_k, such that over any interval of length T_k, the network of agents is "jointly connected", then all agents will asymptotically reach consensus.*

- **Special case: network is connected "once in a while"**
  - Similar result for continuous time, leader follower, ….

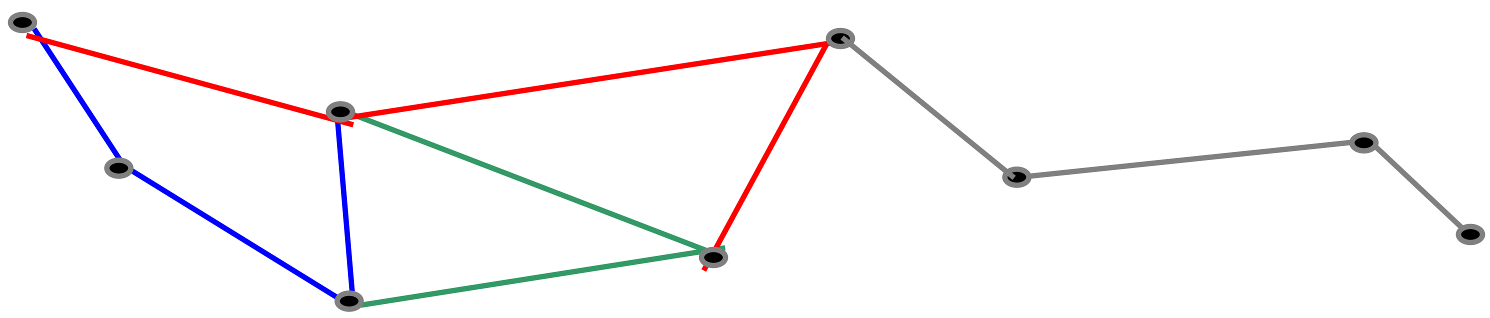

Consensus literature: an incomplete survey

Opinion Dynamics:
[DeGroot 1974, Chaterjee & Seneta 1974]

Parallel and Distributed Computation:
[Tsitsiklis 1984, Tsitsiklis *et al.* 1986]

Distributed Control and Optimization:

- distributed multi-agent optimization: [Nedić & Ozdaglar 2008]
- velocity alignment of kinematic agents: [Jadbabaie *et al.* 2003]
- continuous time dynamics: [Olfati-Saber & Murray 2004]
- directed networks: [Ren & Beard 2004, Lin & Francis 2004]
- nonlinear updates: [Moreau 2004, Lin *et al.* 2005]
- random networks: [Hatano & Mesbahi 2005, Wu 2006]
- time delays: [Angeli & Bliman 2006]
- quantized values: [Savkin 2004]

Consensus over random networks

In the real world, network communication links are random.
(*link failures, interference, physical obstacles, etc.*)

- What are the conditions for consensus when the weight matrices W_k are random?

Reaching consensus in probability: For all $x(0)$, $\epsilon > 0$, and $i, j \in \mathcal{V}$,

$$\mathbb{P}(|x_i(k) - x_j(k)| > \epsilon) \to 0$$

Reaching consensus almost surely: For all $x(0)$ and $i, j \in \mathcal{V}$,

$$|x_i(k) - x_j(k)| \to 0 \quad \textit{almost surely}$$

Consensus in Random Networks

The matrices W_k are independent and identically distributed.

- Edges of the graph maybe dependent at one time-instant.
- The graphs are independent from one time step to another.

The graphs could even be correlated so long as they are stationary-ergodic.

Theorem

The agents reach consensus almost surely if and only if $|\lambda_2(\mathbb{E}W_k)| < 1$, *that is, the expected graph contains a rooted spanning tree.*

A. Tahbaz-Salehi and A. Jadbabaie,
A necessary and sufficient condition for consensus over random networks
IEEE Transactions on Automatic Control, April 2008.

Also Hatano & Mesbahi 2006; Wu 2006; Picci & Taylor 2007; Fagnani & Zampieri 2008; Porfiri & Stilwell 2007

What about consensus value?

A random graph sequence means a random consensus value x^*.

What is its distribution?

Open problem, but ...

- $\mathbb{E}x^* = x(0)^T \mathbf{v}_1(\mathbb{E}W_k)$
- $\text{var}(x^*) = [x(0) \otimes x(0)]^T \mathbf{v}_1(\mathbb{E}[W_k \otimes W_k]) - [x(0)^T \mathbf{v}_1(\mathbb{E}W_k)]^2$
- Degenerate distribution iff all matrices have same left eigenvector.
- Large-scale behavior

Can we say more?

Almost Surely

[Tahbaz-Salehi and Jadbabaie, TAC, 2009 (To Appear)]
[Preciado, Tahbaz-Salehi, and Jadbabaie 2009]

Switching Erdos-Renyi graphs

- Consider a network with *n* nodes and a vector of initial values, **x**(0)
- Repeated local averaging using a *switching and directed* graph $\mathcal{G}(n,p)$
- In each time step, $\mathcal{G}_k(n,p)$ is a realization of a random graph where edges appear with probability, $\Pr(a_{ij}=1)=p$, independently of each other

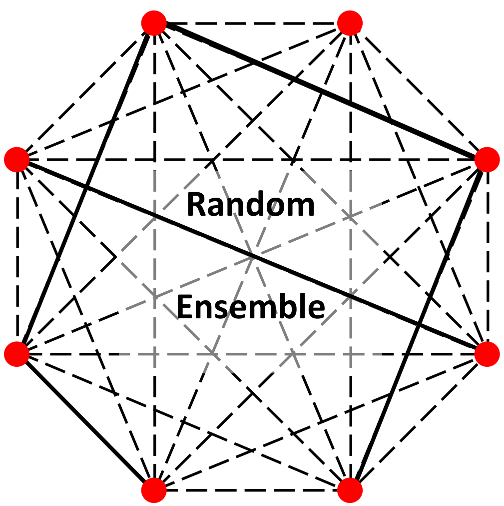

Consensus dynamics

$$\mathbf{x}(k+1) = W_k\,\mathbf{x}(k)$$

$$W_k = (D_k + I_n)^{-1}(A_k + I_n)$$

Stationary behavior

$$\mathbf{x}(k) = U_k\mathbf{x}(0), \text{ with } U_k = W_{k-1}W_{k-2}\ldots W_0,$$

$$\lim_{k\to\infty} U_k = \mathbf{1}\mathbf{v}^T, \text{ where } \mathbf{v} \text{ is a random vector,}$$

$$x^* = \lim_{k\to\infty} x_i(k) \text{ is a random variable.}$$

We can find a close form expression for the mean & variance of x*

Mean and variance in E-R graphs

- Remember, for any IID graph sequence

$$\begin{aligned} \mathbb{E}x^* &= \mathbf{x}(0)^T \mathbf{v}_1(\mathbb{E}W_k) \\ var(x^*) &= [\mathbf{x}(0)\otimes\mathbf{x}(0)]^T \ \mathbf{v}_1(\mathbb{E}[W_k\otimes W_k]) - \left[\mathbf{x}(0)^T \ \mathbf{v}_1(\mathbb{E}W_k)\right]^2 \end{aligned}$$

- Expected weight matrix is symmetric! Therefore mean is just average of initial condiitons!

$$\begin{aligned} \mathbb{E}W_k &= \mathbb{E}w_{ij}\mathbf{1}_n\mathbf{1}_n^T + (\mathbb{E}w_{ii} - \mathbb{E}w_{ij}) I_n \\ &= \frac{1 - f_1(p,n)}{n-1}\mathbf{1}_n\mathbf{1}_n^T - \frac{1 - n\ f_1(p,n)}{n-1} I_n \end{aligned}$$

- Computing the variance of x^* is more complicated
  - Involves the Perron vector of the matrix $\boldsymbol{E}[W_k \otimes W_k]$.
    (it is not Kronecker product of two eigenvectors!)
- Can derive a closed form expression of the left eigenvector of $\boldsymbol{E}[W_k \otimes W_k]$ for any network size n, link probability p, and initial condition $x(0)$.

What does $E[W_k \otimes W_k]$ look like?

- $\mathbf{E}[W_k \otimes W_k]$ is **Not** $\mathbf{E}[W_k] \otimes \mathbf{E}[W_k]$, but it almost is!
- Only n of the n^2 entries are different!

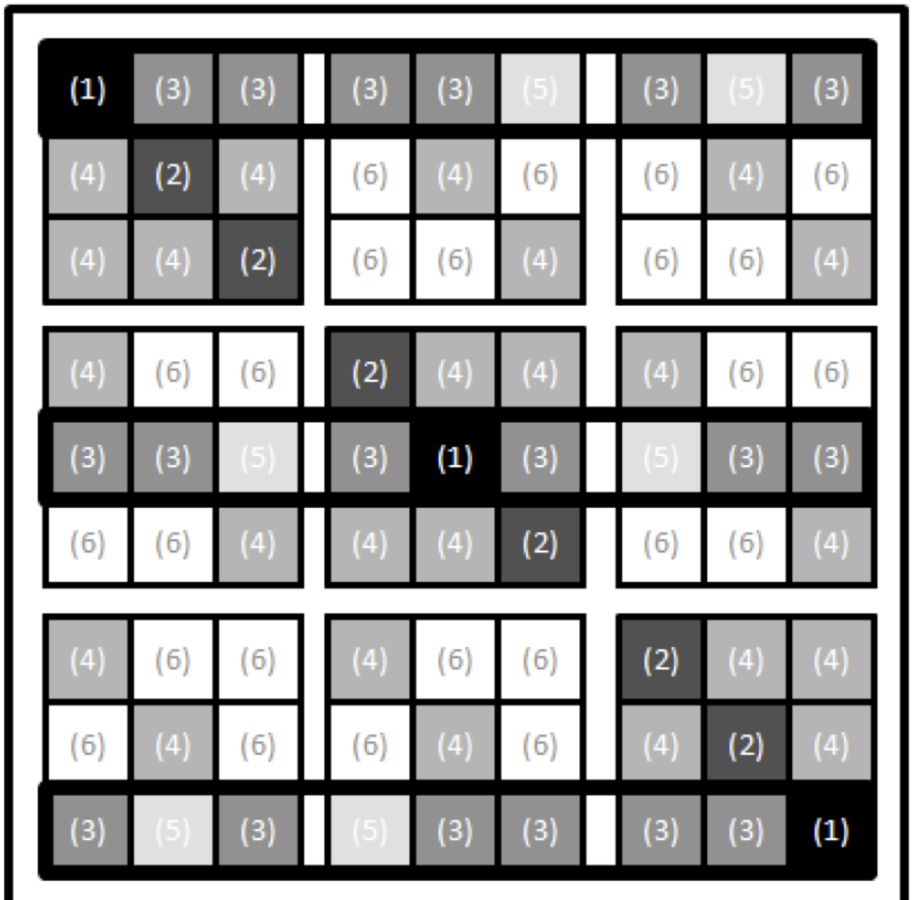

$$(1)\ \mathbb{E}(w_{ii}w_{ii}) = q^{n-1}H(p,n),$$

$$(2)\ \mathbb{E}(w_{ii}w_{jj}) = \left(\frac{1-q^n}{np}\right)^2,$$

$$(3)\ \mathbb{E}(w_{ii}w_{is}) = \mathbb{E}(w_{ij}w_{ij}) = \mathbb{E}(w_{ij}w_{ii}) = \frac{q-q^n(q+npH(p,n))}{n(n-1)pq},$$

$$(4)\ \mathbb{E}(w_{ii}w_{ri}) = \mathbb{E}(w_{ii}w_{rs}) = \frac{(1-q^n)(np-1+q^n)}{n^2p^2(n-1)},$$

$$(5)\ \mathbb{E}(w_{ij}w_{is}) = \frac{q(np-3)+q^n(3-3p+2np\ H(p,n))}{pqn(n-1)(n-2)},$$

$$(6)\ \mathbb{E}(w_{ij}w_{ji}) = \mathbb{E}(w_{ij}w_{js}) = \mathbb{E}(w_{ij}w_{ri})$$
$$= \mathbb{E}(w_{ij}w_{rj}) = \mathbb{E}(w_{ij}w_{rs}) = \left(\frac{np-1+q^n}{n(n-1)p}\right)^2,$$

q=1-p and *H(p,n)* can be written in terms of a hypergeometric function

A surprising result

- Theorem: $$\mathrm{var}(x^*) = \frac{1-\rho}{\delta}\sum_{i=1}^{n}\left[x_i(0)-\bar{x}(0)\right]^2,$$

$$\rho(p,n) \triangleq \frac{p(n-1)}{p(n-2)+1-(1-p)^n}, \qquad \delta(p,n) \triangleq \frac{n^2}{n+(n^2-n)\,\rho(p,n)}$$

- No Kronecker products or hyper-geometric functions
- As network size grows, variance of consensus value goes to zero!
- What about other random graph models?

Other random graph models

- Erdos-Renyi graphs are easier: no correlation between entries of adjacency matrix
- E-R graphs have Poisson degree distributions, whereas many real large networks have heavy tailed degree distributions
- Note: degree distribution does not uniquely identify the topology (not even close!)
- Fixing degree distribution (and higher degree correlations) will fix the moments of the Laplacian spectrum!

Degree distribution≠ Topology

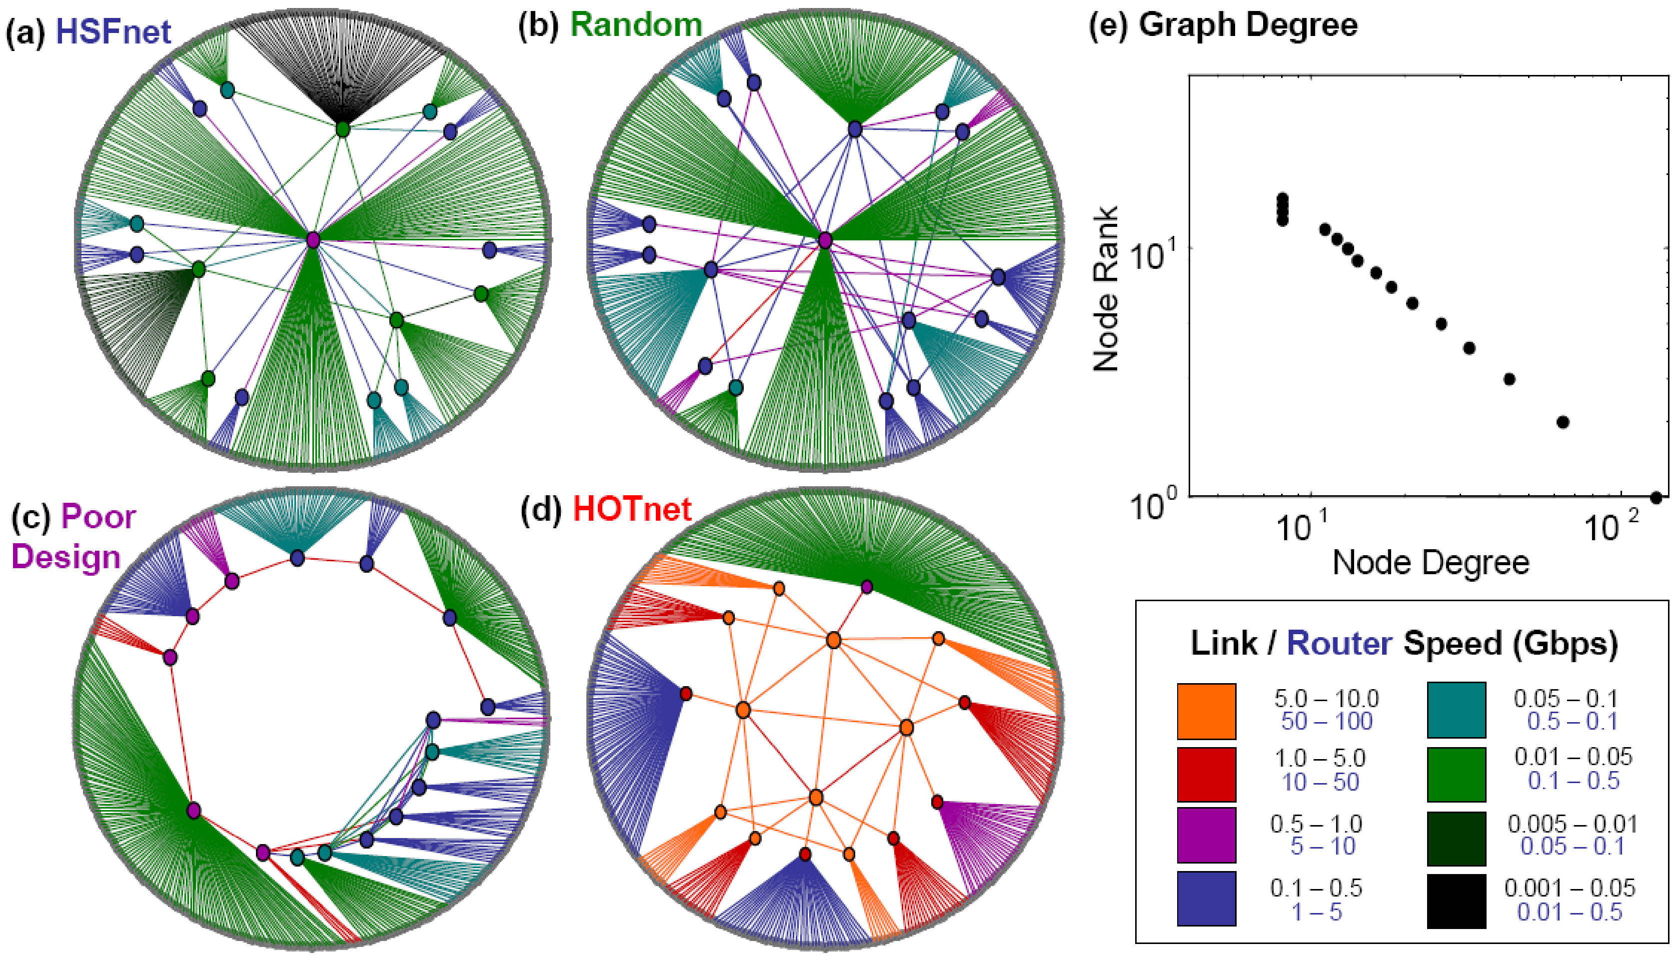

4 graphs with exact # of nodes, links and degree distribution but VERY different topologies Li, Alderson, Willinger & Doyle 06

Random graph models with prescribed degree distribution

- Degree distributions Do tell us something though
- Moments of spectra of graph Laplacians!

- Generalized static models [Chung and Lu, 2003]:
  - Random graph with a prescribed expected degree sequence
  - We can impose an expected degree w_i on the i-th node

$$(w_1, w_2, ..., w_n) \quad \rho = (\Sigma_i w_i)^{-1}$$

$$\mathbb{P}\,(i \sim j) = \rho w_i w_j$$

$$\mathbb{E}[d_i] = w_i$$

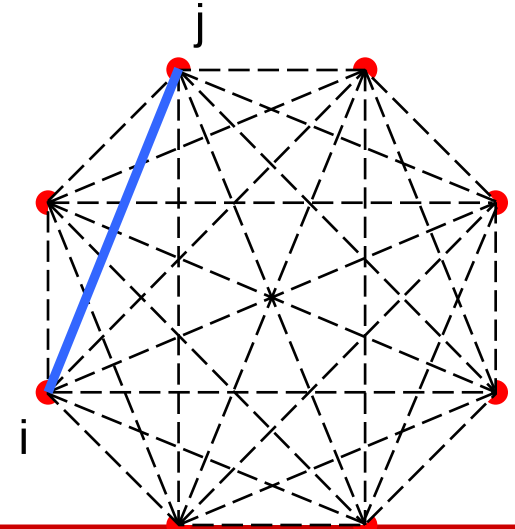

A CLT for eigenvalues of Chung-Lu Graph Laplacians

- **Main Problem:** adjacency entries are correlated

- **Numerical Experiment:** Represent the histogram of eigenvalues for several realizations of this random graph

100 nodes

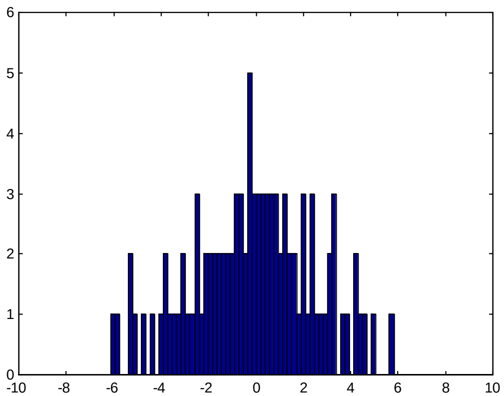

500 nodes

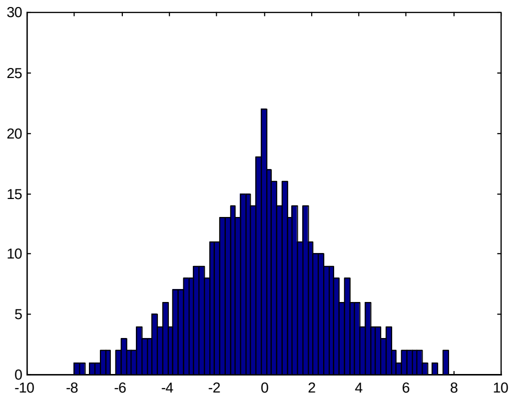

1000 nodes

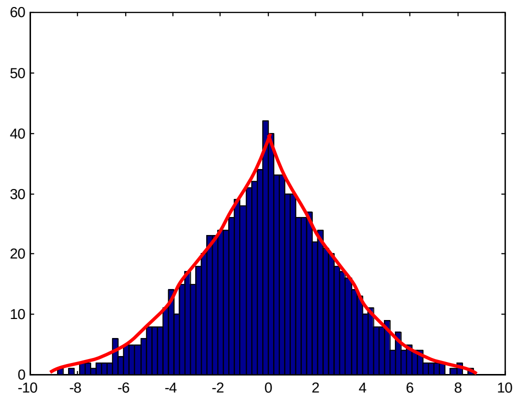

- **Limiting Spectral Density:** Analytical expression only possible for very particular cases.

Contribution: Estimation of the shape of the bulk for a given expected degree sequence, $(w_1,\ldots,w_n)$.

Symbolic Polynomials for Expected Spectral Moments for large graphs

- Our symbolic expressions are in terms of $W_k = \sum_{i=1}^{n} w_i^{\,k}$

$$\mathbb{E}[m_2(\tilde{A}_n)] = (1+o(1))\frac{1}{n}\left(W_1^2\right),$$
$$\mathbb{E}[m_4(\tilde{A}_n)] = (1+o(1))\frac{1}{n}\left(2W_1^2 W_2\right),$$
$$\mathbb{E}[m_6(\tilde{A}_n)] = (1+o(1))\frac{1}{n}\left(2W_1^3 W_3 + 3W_1^2 W_2^2\right),$$
$$\mathbb{E}[m_8(\tilde{A}_n)] = (1+o(1))\frac{1}{n}(2W_1^4 W_4 + 8W_1^3 W_2 W_3 + 4W_1^2 W_2^3)$$

- Numerical verification: 500 nodes random power-law, β=2.5

2s-th order	Analytical Expectation	Numerical Realization	Relative Error
$m_2(\tilde{A}_n)$	2.8088e+004	2.8024e+004	0.23 %
$m_4(\tilde{A}_n)$	1.6363e+009	1.6237e+009	0.77 %
$m_6(\tilde{A}_n)$	1.2075e+014	1.1870e+014	1.69 %
$m_8(\tilde{A}_n)$	1.0040e+019	9.7485e+018	2.90 %
$m_{10}(\tilde{A}_n)$	8.9736e+023	8.5881e+023	4.30 %

Only one typical realization!

Other related work

- Closed form solutions for moments of adjacency/Laplacian matrices of random geometric graphs (all moments in 1d, first 3 moments in 2d)

(Preciado & J., CDC 09)

- Given the moments, we can estimated the shape of eigenvalue distribution, and estimate the spread

- Can predict synchronizability, speed of convergence once we know the spread of eigenvalues *(Preciado & J., CDC 09)*

When is consensus a good thing?

Do consensus algorithms aggregate information correctly?

Sometimes.

- Computing the maximum likelihood estimator [Boyd, Xiao, and Lall 2006]
- Learning in large networks [Golub and Jackson 2008]

In many scenarios agreement is not sufficient.
Agents need to agree on the “right” value: learning.

Consensus and Naïve Social learning

- Need to make sure update converges to the correct value

Agents initially receive a noisy signal about the true state of the world.
Update their beliefs as a weighted average of the neighbors' beliefs.
In a connected network, people reach asymptotic consensus.
What is this value if the size of the network grows?

If no agent is overly influential, then the consensus value converges to the true state of the world in probability, that is, everybody learns the true state.

Wisdom of crowds

B. Golub and M. O. Jackson.
Naïve Learning in Social Networks: Convergence, Influence, and the Wisdom of Crowds
Unpublished Manuscript, December 2008.

Social learning

- There is a (pay-off relevant) true state of the world, among countably many (eg quality of a product, suitability of a political candidate, …)
- We start from a prior distribution, would like to update the distribution (or *belief on the true state)* with more observations
- Ideally we use Bayes rule to do the information aggregation
- Works well when there is one agent (*Blackwell, Dubin'1962), becomes hard when more than 2!*

Social Learning

Bayesian learning over social networks:
[Banerjee 1998]
[Smith and Sørensen 1998]
[Acemoglu, Dahleh, Lobel, and Ozdaglar 2008]

Rule-of-thumb learning over networks (DeGroot's Model):
[DeMarzo, Vayanos, and Zwiebel 2003]
[Acemoglu, Nedić, and Ozdaglar 2008]
[Golub and Jackson 2008]
[Acemoglu, Ozdaglar, and ParandehGheibi 2009]

Non-Bayesian learning:
[Ellison and Fudenberg 1993, 1995]
[Bala and Goyal 1998, 2001]

Bayesian learning

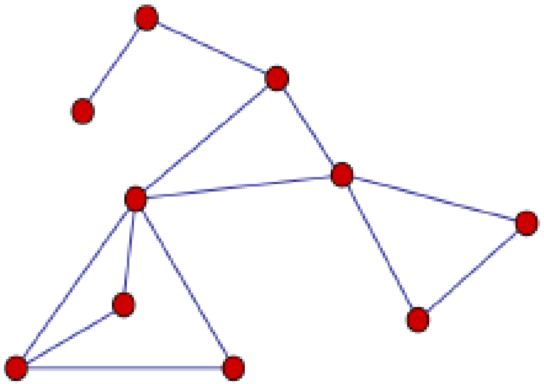

$$\mu_{i,t}(\theta) = \mathbb{P}\left[\theta = \theta^* \,|\, \mathcal{F}_{i,t}\right]$$

where

$$\mathcal{F}_{i,t} = \sigma\left(s_1^i, \ldots, s_t^i, \{\mu_{j,k} : j \in \mathcal{N}_i, k \le t\}\right)$$

is the information available to agent i up to time t.

Agents need to make rational deductions about everybody's beliefs based on only observing neighbors' beliefs:

Problem with Bayesian Social learning

1. Incomplete network information
2. Incomplete information about other agents' signal structures
3. Higher order beliefs matter ▸ Example
4. The source of each piece of information is not immediately clear

Naïve vs. Rational learning

Naïve learning

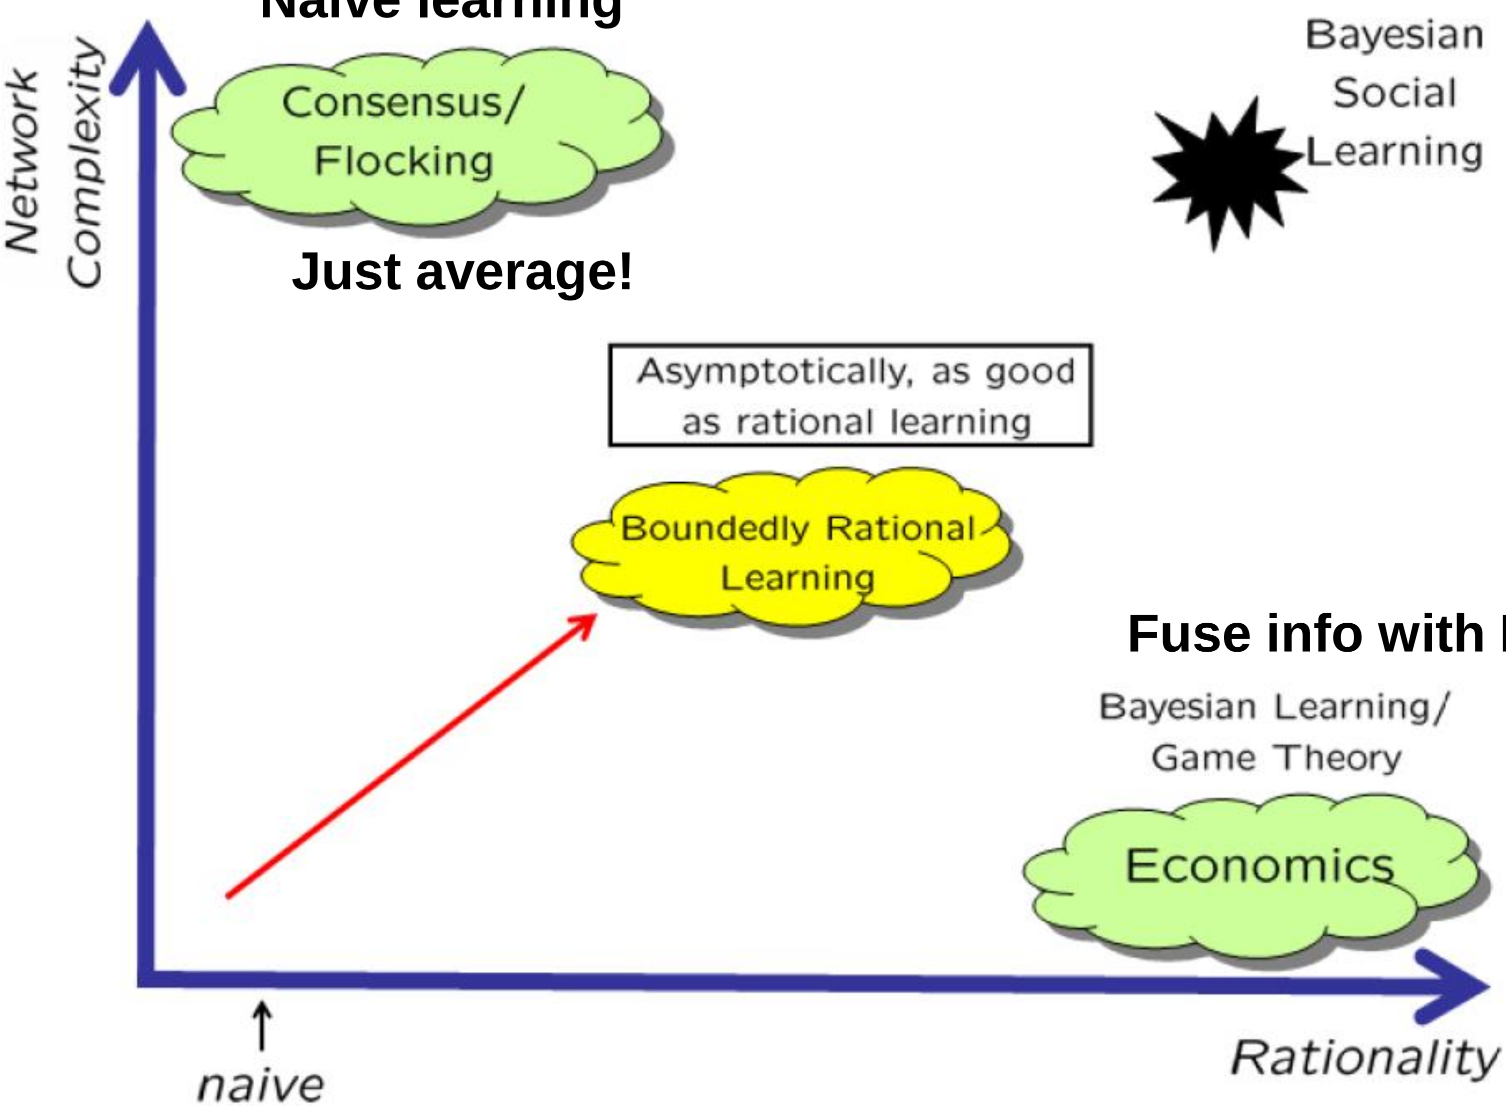

Locally Rational, Globally Naïve: Bayesian learning under peer pressure

Need a local and computationally tractable update, which hopefully delivers asymptotic social learning.

Agent i is

- Bayesian when it comes to her observation
- non-Bayesian when incorporating others information

[Tahbaz-Salehi, Sandroni, and Jadbabaie 2009]

Model Description

$\mathcal{N} = \{1, 2, \ldots, n\}$	individuals in the society
$G = (\mathcal{N}, \mathcal{E})$	social network
Θ	finite parameter space
$\theta^* \in \Theta$	the unobservable true state of the world
$s_t = (s_t^1, \ldots, s_t^n)$	s_t^i is the signal observed by agent i at time t
$S = S_1 \times S_2 \times \cdots \times S_n$	signal space
$\ell(s\|\theta)$	the likelihood function (prob. of observing s if the true state is θ)
$\ell_i(s^i\|\theta)$	the marginal likelihood function

Model Description

$\mu_{i,t}(\theta)$	time t beliefs of agent i (a probability measure on Θ)
$\mu_{i,0}(\theta)$	agent i's prior belief
$\mathbb{P}^* = \otimes_{t=1}^{\infty} \ell(\cdot \mid \theta^*)$	the true probability measure

Agent i's time t forecasts of the next observation profile:

$$m_{i,t}(s_{t+1}) = \int_{\Theta} \ell(s_{t+1} \mid \theta) d\mu_{i,t}(\theta)$$

What do we mean by learning?

Definition

The Forecasts of agent i are eventually correct on a path $\{s_t\}_{t=1}^{\infty}$ if, along that path,

$$m_{i,t}(\cdot) \to \ell_i(\cdot|\theta^*) \quad \text{as} \quad t \to \infty.$$

Definition

Agent i asymptotically learns the true parameter θ^* on a path $\{s_t\}_{t=1}^{\infty}$ if, along that path,

$$\mu_{i,t}(\theta^*) \to 1 \quad \text{as} \quad t \to \infty.$$

- Asymptotic learning, in this setup, is stronger.
- Depends on the information structure.

Belief Update Rule

$$\mu_{i,t+1}(\theta) = a_{ii} BU(\mu_{i,t}; s^i_{t+1})(\theta) + \sum_{j \in \mathcal{N}_i} a_{ij} \mu_{j,t}(\theta)$$

where

$$BU(\mu_{i,t}; s^i_{t+1})(\theta) = \mu_{i,t}(\theta) \frac{\ell_i(s^i_{t+1}|\theta)}{m_{i,t}(s^i_{t+1})}$$

$$a_{ij} \geq 0 \quad , \quad \sum_{j \in \mathcal{N}_i} a_{ij} = 1$$

- Individuals rationally update the beliefs after observing the signal
- exhibit a bias towards the average belief in the neighborhood

Why this update?

$$\mu_{i,t+1}(\theta) = a_{ii}\mu_{i,t}(\theta)\frac{\ell_i(s_{t+1}^i|\theta)}{m_{i,t}(s_{t+1}^i)} + \sum_{i \neq j} a_{ij}\mu_{j,t}(\theta) \qquad \forall\theta \in \Theta$$

- Does not require knowledge about the network.
- Does not require deduction about the beliefs of others.
- Does not require knowledge about other agents' signallings.
- The update is local and tractable.
- If the signals are uninformative, reduces to the consensus update.
- Reduces to the benchmark Bayesian case if agents assign weight zero to the beliefs of their neighbors. [Blackwell and Dubins 1962]

Eventually correct forecasts

$$\mu_{i,t+1}(\theta) = a_{ii}\mu_{i,t}(\theta)\frac{\ell_i(s^i_{t+1}|\theta)}{m_{i,t}(s^i_{t+1})} + \sum_{i\neq j} a_{ij}\mu_{j,t}(\theta) \qquad \forall\theta\in\Theta$$

Theorem

Suppose that

1. the social network is strongly connected,
2. $a_{ii} > 0$ for all $i \in \mathcal{N}$,
3. there exists an agent i such that $\mu_{i,0}(\theta^*) > 0$.

Then the forecasts of all agents are eventually correct $\mathbb{P}^*$-almost surely, that is, $m_{i,t}(\cdot) \to \ell_i(\cdot|\theta^*)$.

Why strong connectivity?

What if the network has a directed spanning tree but is not strongly connected?

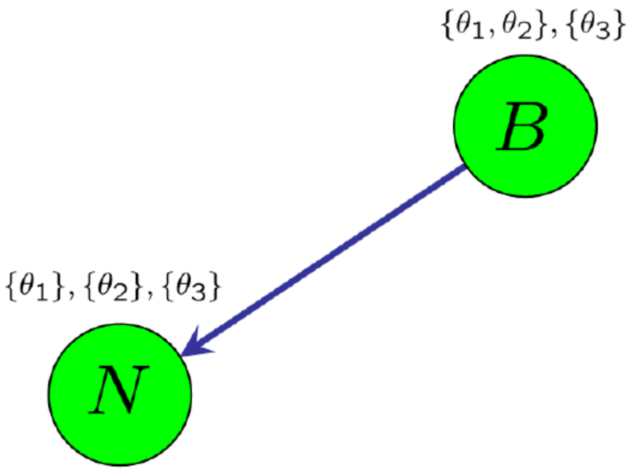

- $\mathcal{N} = \{B, N\}$
- $\Theta = \{\theta_1, \theta_2, \theta_3\}$
- $\theta^* = \theta_2$

$$\mu_{N,t+1}(\theta) = \lambda \mu_{N,t}(\theta) \frac{\ell_N(s_{t+1}^N | \theta)}{m_{N,t}(s_{t+1}^N)} + (1-\lambda)\mu_{B,t}(\theta) \qquad \forall \theta \in \Theta$$

- No convergence if different people interpret signals differently
- N is misled by listening to the less informed agent B

Example

In any strongly connected social network, forecasts of all agents are correct on almost all sample paths.

- $\mathcal{N} = \{1, 2\}$
- $\Theta = \{\theta_1, \theta_2, \theta_3\}$
- $\theta^* = \theta_2$

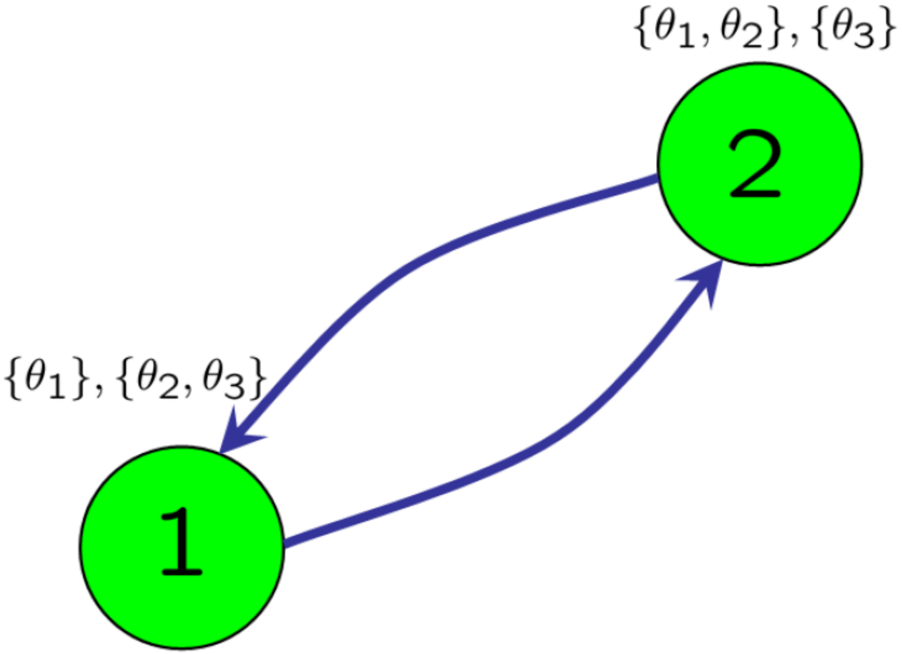

One can actually learn from others

Convergence of beliefs and consensus on correct value!

Theorem

Under the assumptions of the pervious theorem, the beliefs of all agents converge with $\mathbb{P}^*$-probability one.

$$\mu_{t+1}(\theta) = A\mu_t(\theta) + \mathrm{diag}\left(a_{ii}[\frac{\ell_i(s^i_{t+1}|\theta)}{m_{i,t}(s^i_{t+1})} - 1]\right)_{i=1,\dots,n} \mu_t(\theta)$$

Corollary

Under the assumptions of the theorem, all agents have asymptotically equal beliefs $\mathbb{P}^*$-almost surely.

Consensus!

Learning from others

All agents have asymptotically equal forecasts. Therefore,

- Each agent can correctly forecast every other agent's signals.

$$\forall i, j \in \mathcal{N} \quad \int_\Theta \ell_j(\cdot|\theta) d\mu_{i,t}(\theta) \longrightarrow \ell_j(\cdot|\theta^*) \qquad \mathbb{P}^* - \text{a.s.}$$

Local information of any agent is revealed to every other agent.

- This does not mean that the agents can forecast the joint distributions. They can only forecast the marginals correctly.
- To be expected: only marginals appear in the belief update scheme.

Information Aggregation

Theorem

Suppose that

(i) *social network is strongly connected,*

(ii) *all agents have strictly positive self-confidence,*

(iii) *there exists an agent with strictly positive prior on* θ^*.

Then,

1. *Every agent can eventually forecast the signals of every other agent correctly with* $\mathbb{P}^*$*-probability one.*
2. *If there exists a state* $\theta \in \Theta$ *and an agent* i *such that* $\ell_i(s^i|\theta) \neq \ell_i(s^i|\theta^*)$ *for some* $s^i \in S_i$*, then* $\mu_j(\theta) \to 0$ *for all* $j \in \mathcal{N}$.

Example

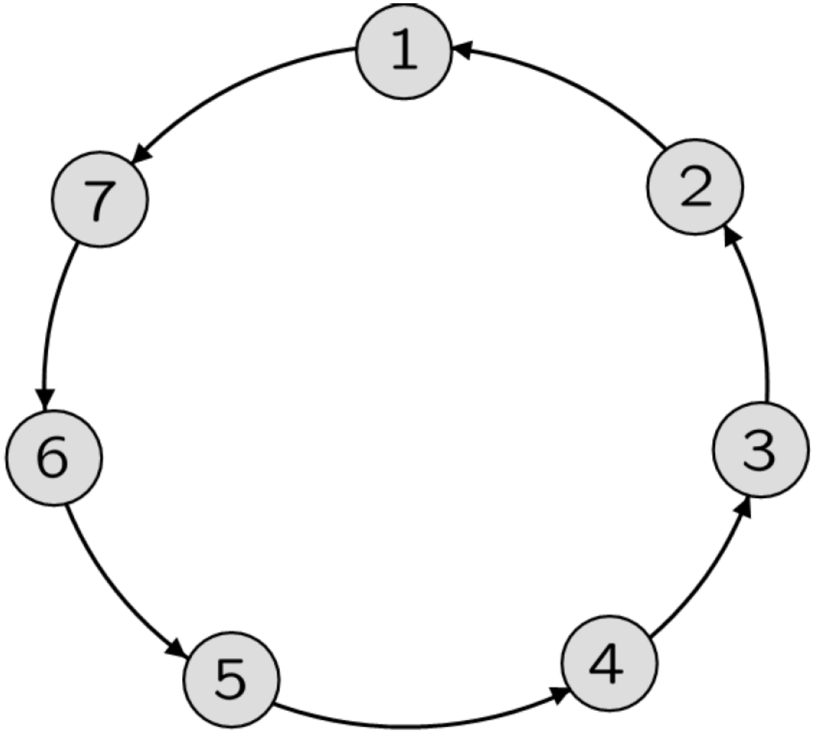

$$\Theta = \{\theta_1, \theta_2, \dots, \theta_7\}$$

$$\theta^* = \theta_1$$

$$S_i = \{H, T\}$$

$$\ell_i(H|\theta) = \begin{cases} \frac{i}{i+1} & \text{if } \theta = \theta_i \\ \frac{1}{i+1} & \text{otherwise} \end{cases}$$

Local information of every agent is revealed to every other agent.

Summary

How information is aggregated over networks?

- from local information to inference about global uncertainties

Bayesian learning over networks:

- optimal but not tractable

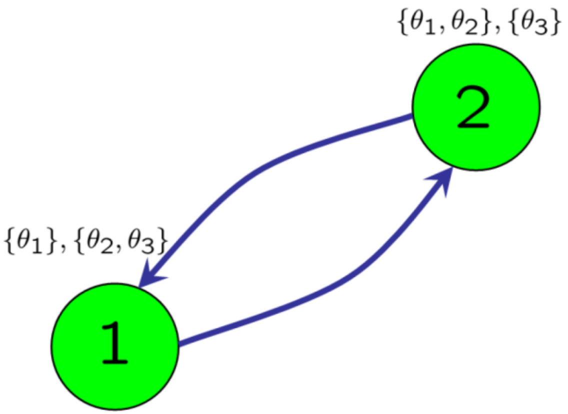

Extends to changing graphs under some conditions on weights

No need to be Bayesian for asymptotic learning:

- a non-Bayesian model of learning which is local and tractable.
- asymptotically optimal, under independent observations.

Acknowledgements

Research supported by:
DARPA/DSO SToMP
ARO/MURI W911NF-05-1-0381
ONR/YIP N00014-04-1-0467
ONR MURI N000140810747
NSF-ECS-0347285
AFOSR: Complex Networks Program