Learning over Complex Social Networks

Munther A. Dahleh,

Laboratory for Information and Decision Systems

EECS/MIT

SYSID 2009, Saint-Malo, France

Collaborations

- Daron Acemoglu: Economics, MIT
- Ilan Lobel: Business School, NYU
- Asuman Ozdalgar: EECS, MIT

Drug Prescription

The drugs your physician prescribes may well depend on the behavior of an opinion leader in his or her social network in addition to your doctor's own knowledge of or familiarity with those products.

Smoking

Whether a person quits smoking is largely shaped by social pressures, and people tend to quit smoking in groups. If a spouse quits smoking, the other spouse is 67% less likely to smoke. If a friend quits, a person is 36% less likely to still light up. Siblings who quit made it 25% less likely that their brothers and sisters would still smoke.

Social Networks and Politics

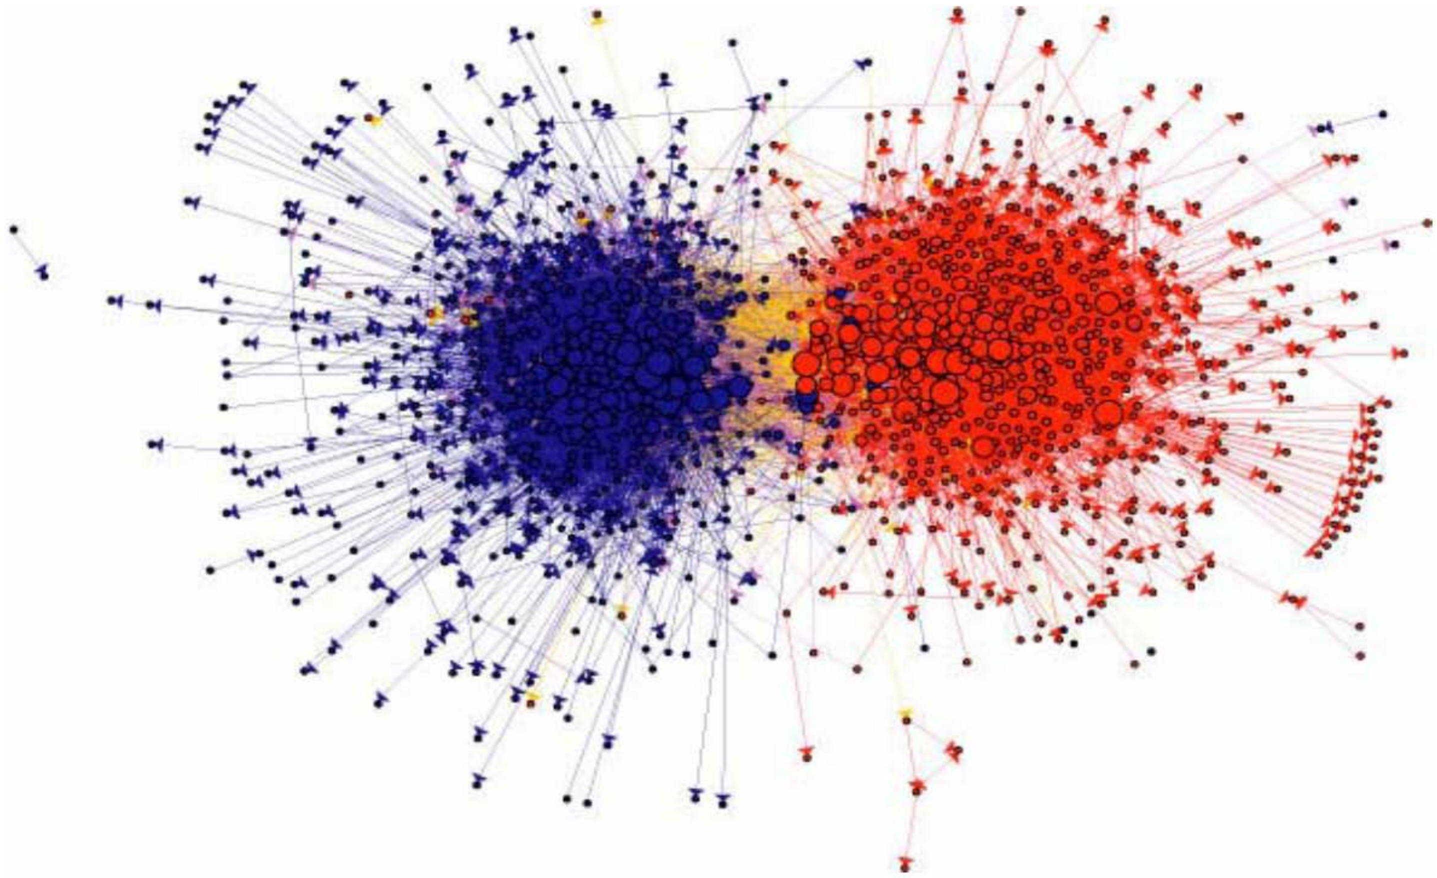

Network structure of political blogs prior to 2004 presidential elections

The Tipping Point: M. Gladwell

The Tipping Point is that magic moment when an idea, trend, or social behavior crosses a threshold, tips, and spreads like wildfire. Just as a single sick person can start an epidemic of the flu, so too can a small but precisely targeted push cause fashion trend, the popularity of a new product, or a drop of crime rate.

Objective

Develop Models that can capture the impact of a social network on learning and decision making

Who's Buying the Newest Phone and Why?

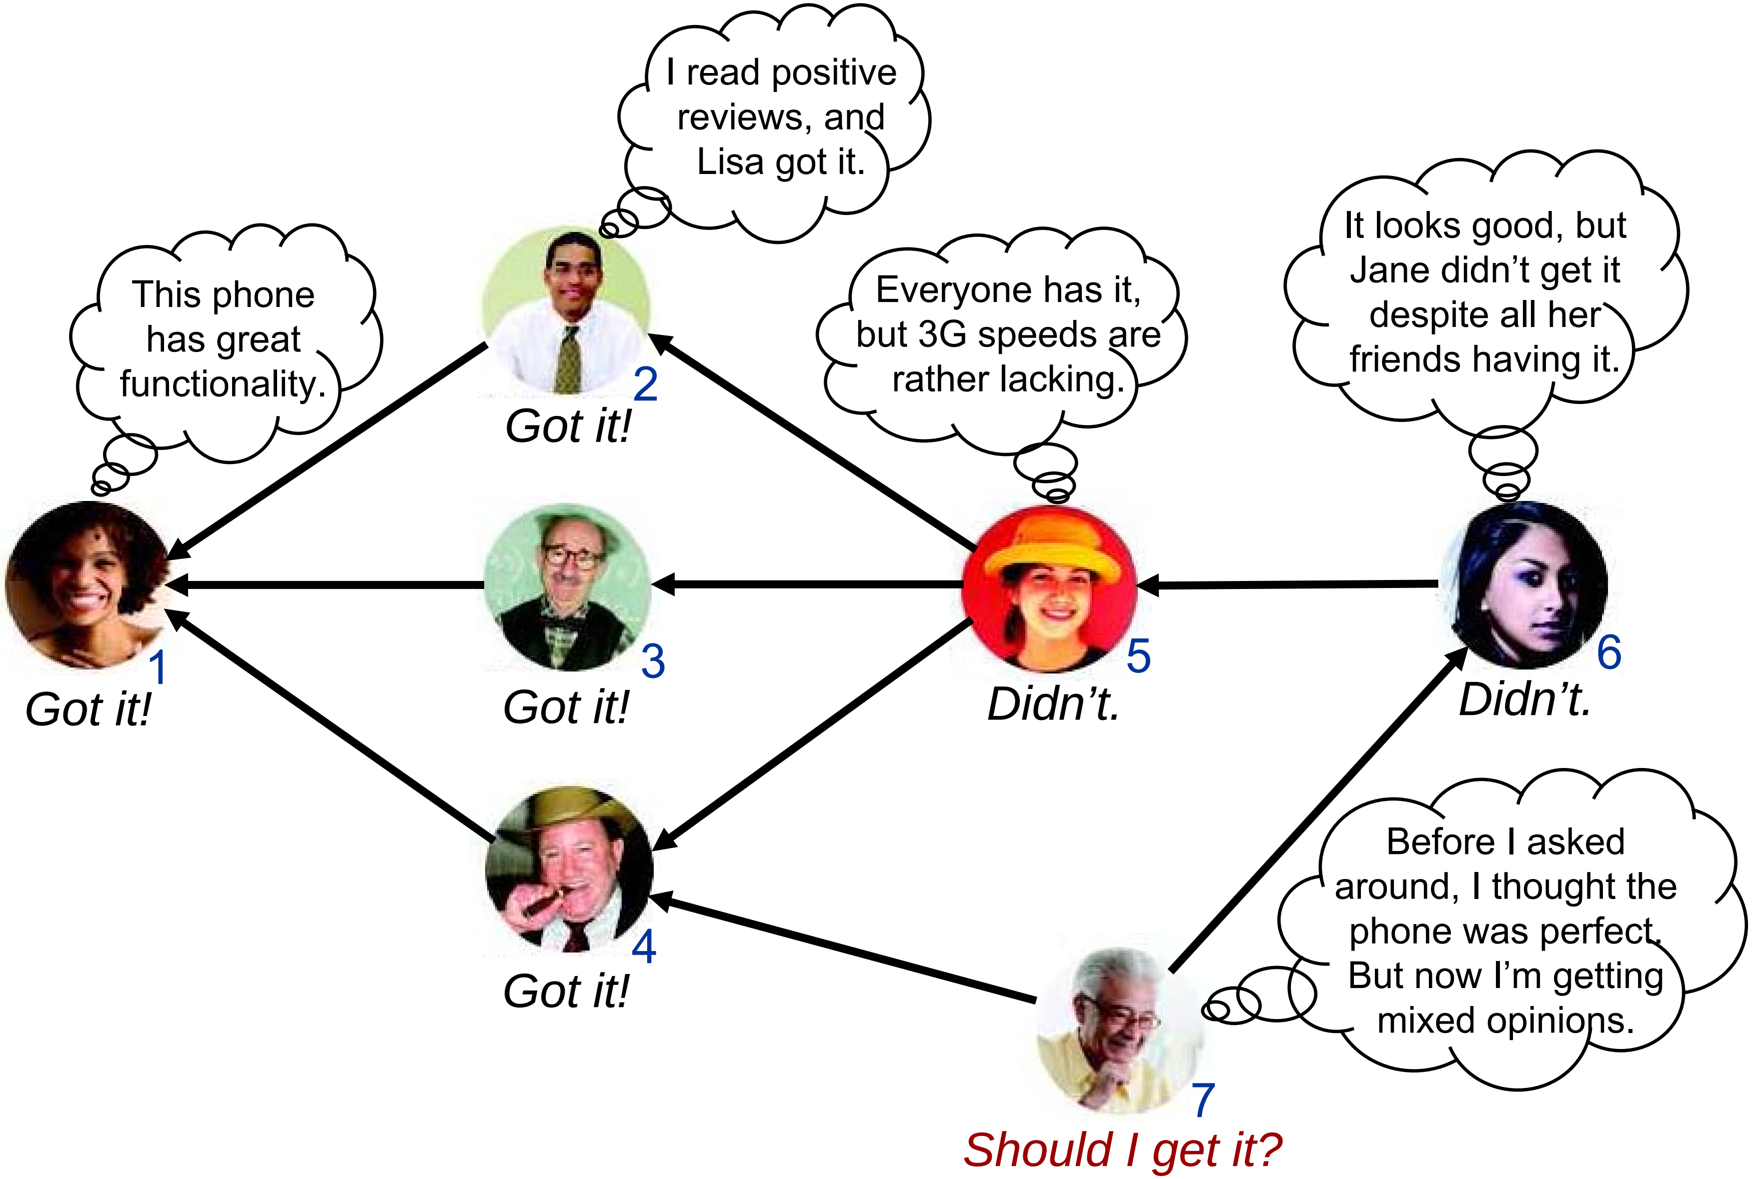

Formulation: Two States

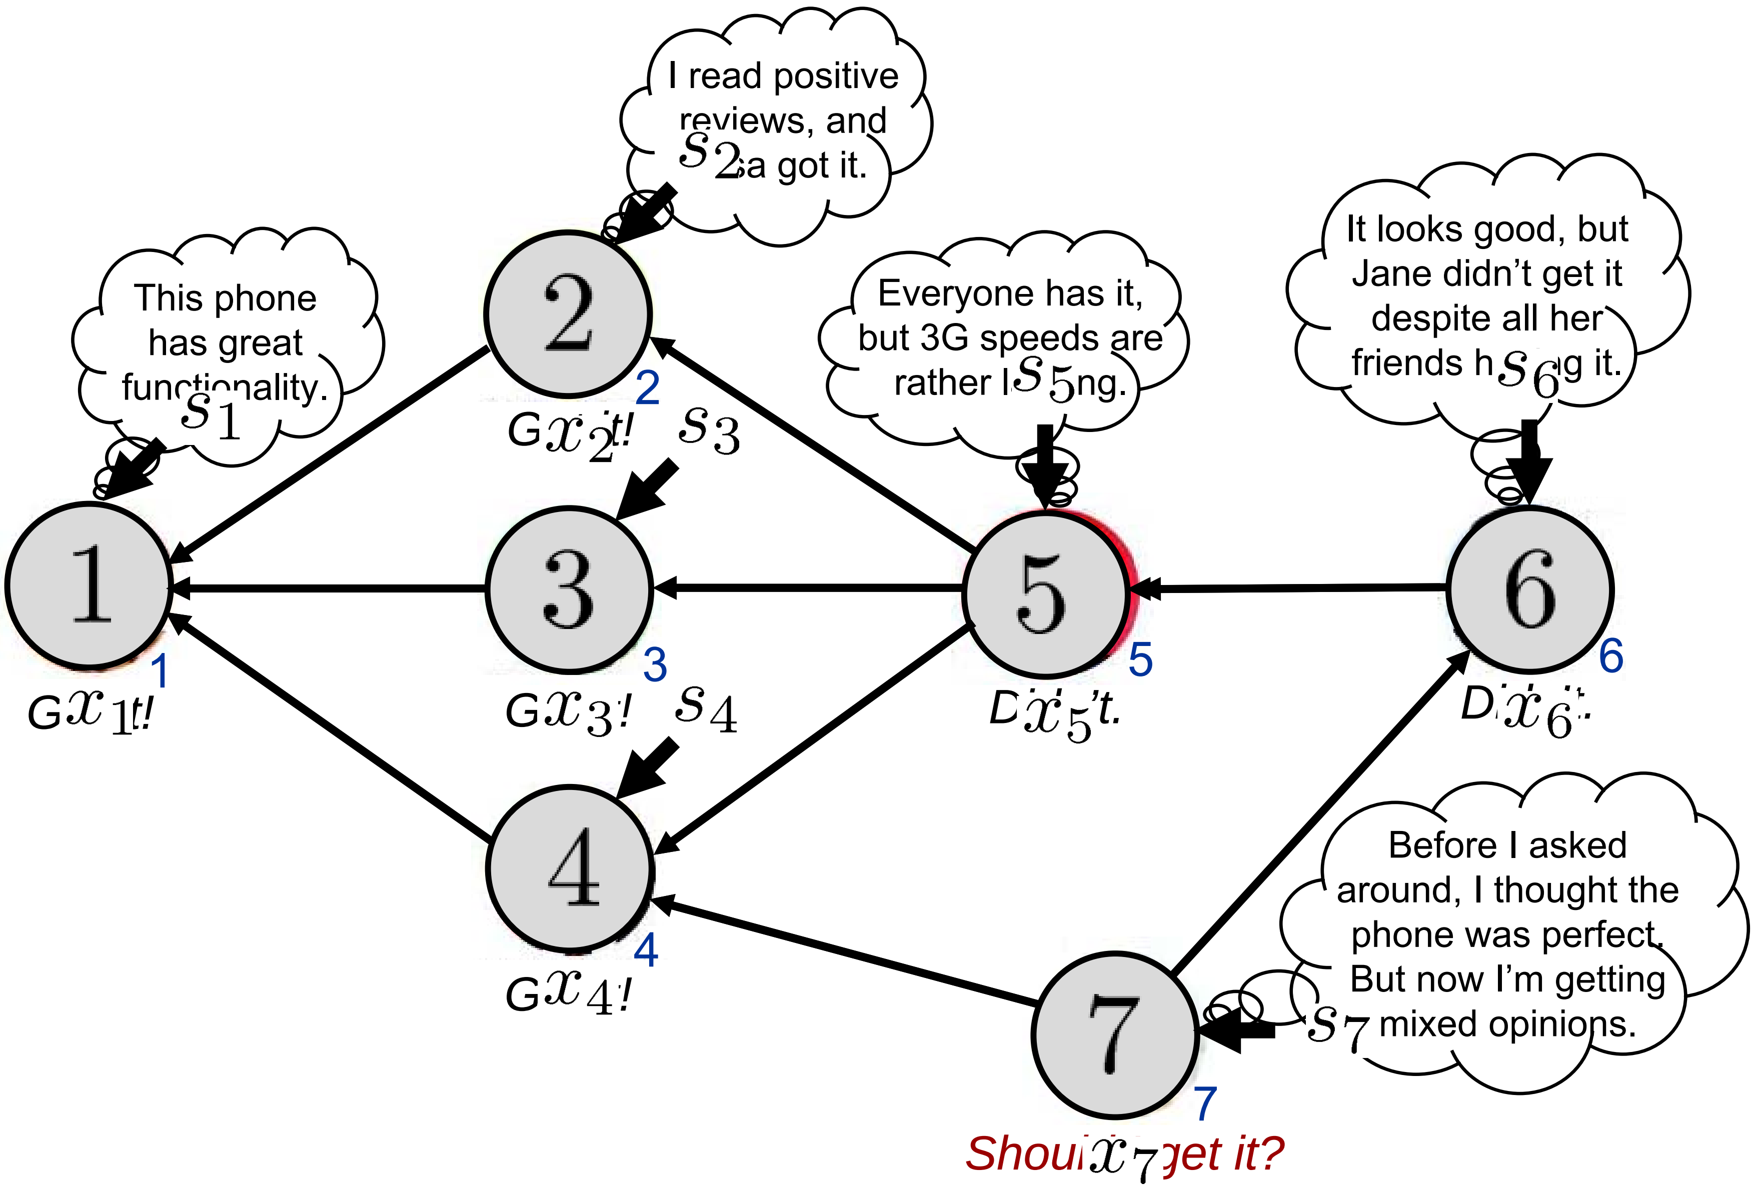

General Setup

- Two possible states of the world $\theta \in \{0,1\}$ both equally likely.
- A sequence of agents $(n = 1, 2, ...)$ making binary decisions x_n.
- Agent *n* obtains utility 1 if $x_n = \theta$ and utility 0 otherwise.

- Each agent has an iid private signal $s_n \in [0,1]$. The signal is sampled from a cumulative density F_θ .

- The neighborhood: $\Omega_n \subset \{1, \ldots n-1\}$

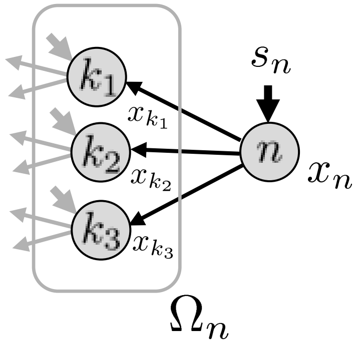

- Information: $I_n = \{s_n, \Omega_n, x_k \text{ for all } k \in \Omega_n\}$

Rationality

- Rational Choice: Given information set I_n agent n chooses

$$\sigma_n(I_n) \in \arg \max_{y \in \{0,1\}} \mathbf{P}(\theta = y | I_n) .$$

- Strategy profile: $\sigma = \{\sigma_n\}$

- Asymptotic Learning: Under what conditions does

$$\lim_{n \to \infty} \mathbf{P}_\sigma(x_n = \theta) = 1$$

Equilibrium Decision Rule

- The belief about the state decomposes into two parts
  - Private Belief, $p_n(s_n) = \mathbf{P}_\sigma(\theta = 1 | s_n) = \left(1 + \frac{dF_0(s_n)}{dF_1(s_n)}\right)^{-1}$
  - Social belief, $\mathbf{P}_\sigma(\theta = 1 | \Omega_n, x_k \text{ for all } k \in \Omega_n)$

- Strategy profile σ is a perfect Bayesian equilibrium if and only if:

$$\mathbf{P}_\sigma(\theta = 1 | s_n) + \mathbf{P}_\sigma(\theta = 1 | \Omega_n, x_k \text{ for all } k \in \Omega_n) > 1 \Longrightarrow \sigma_n(I_n) = 1;$$

$$\mathbf{P}_\sigma(\theta = 1 | s_n) + \mathbf{P}_\sigma(\theta = 1 | \Omega_n, x_k \text{ for all } k \in \Omega_n) < 1 \Longrightarrow \sigma_n(I_n) = 0.$$

Selfish vs. Engineered Response: Star Topology (Cover)

- Hypothesis testing:

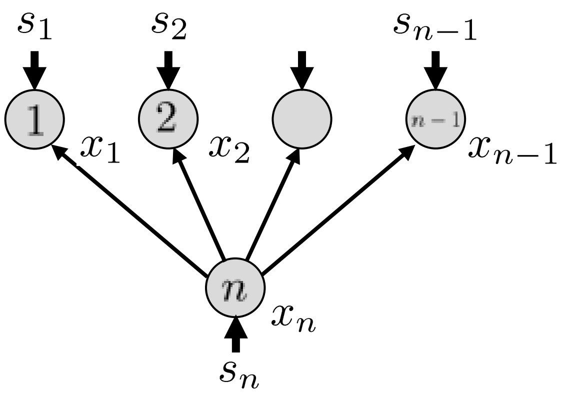

- If nodes communicate their observations: $L(\cdot) = \text{Likelihood Ratio}$

$$\mathbf{P}\{\frac{1}{n}\sum_i L(s_i) - \mathbf{E}L(S) > d\} \leq e^{ng(d)}$$

- What if nodes communicate only their decisions: $x_i = f(s_i)$

$$\mathbf{P}\{\frac{1}{n}\sum_i L(x_i) - \mathbf{E}L(x) > d\} \leq e^{ng_1(d)}$$

Selfishness and Herding Phenomenon: [Banerjee (92), BHW 92]

- Setup:
  - Full network: $\Omega_n = \{1, 2, ..., n-1\}$
  - $s_i = \theta$ with probability 0.8

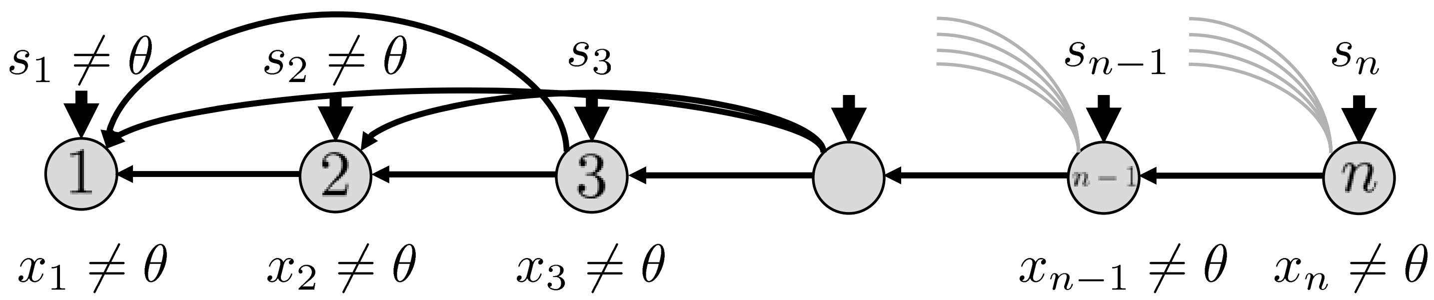

ABSENCE OF COLLECTIVE WISDOM

Line Network [AP90, TTW 07]

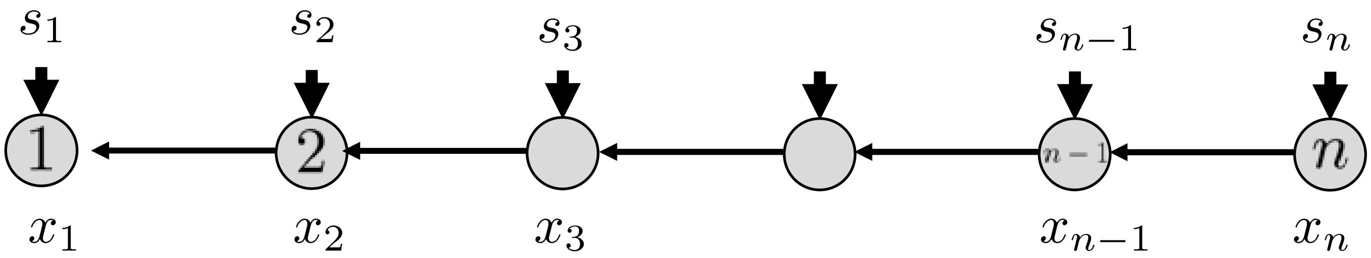

Line Network [AP90, TTW 07]

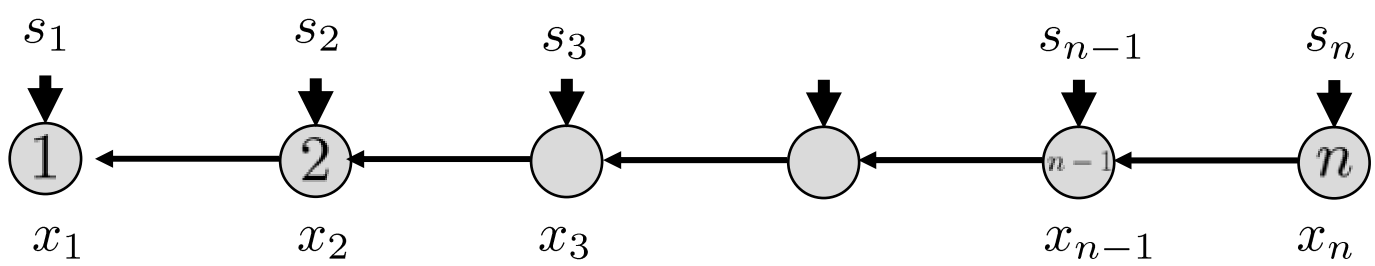

- When the set of observed decisions $\Omega_n = \{n-1\}$, for any equilibrium σ, the decision is a threshold policy:

$$x_n = \begin{cases} 0, & p_n(s_n) < L_{n-1}; \\ x_{n-1}, & p_n(s_n) \in (L_{n-1}, U_{n-1}); \\ 1, & p_n(s_n) > U_{n-1}. \end{cases}$$

- Define $N_{n-1} = \mathbf{P}_\sigma(x_{n-1} = \theta | \theta = 0)$ and $Y_{n-1} = \mathbf{P}_\sigma(x_{n-1} = \theta | \theta = 1)$ then

$$U_{n-1} = \frac{N_{n-1}}{Y_{n-1} + N_{n-1} - 1} \qquad L_{n-1} = \frac{1 - N_{n-1}}{Y_{n-1} + N_{n-1} - 1}$$

Line Topology: Decision Rule

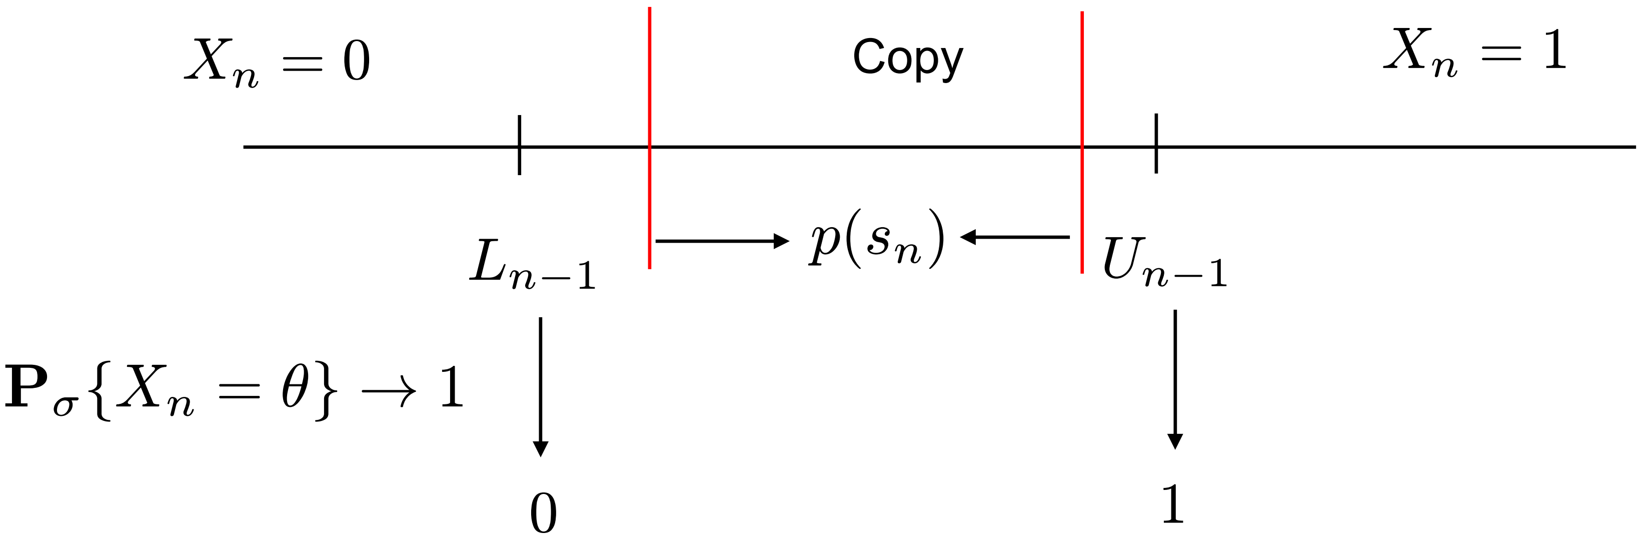

CONTRADICTION

Private Beliefs

- Private Beliefs $p_n(s_n) = \mathbf{P}_\sigma(\theta = 1 | s_n) = \left(1 + \frac{dF_0(s_n)}{dF_1(s_n)}\right)^{-1}.$

- Definition: The private beliefs are called unbounded if

$$\sup_{s \in S} \frac{dF_0(s)}{dF_1(s)} = \infty \quad \text{and} \quad \inf_{s \in S} \frac{dF_0(s)}{dF_1(s)} = 0$$

- If the private beliefs are unbounded, then there exist some agents with beliefs arbitrarily close to 0 and other agents with beliefs arbitrarily close to 1.

- Discrete example: $s_i = \theta$ with probablity 0.8?

The Line Network: One Step Improvement

When observing a single decision, agent n has the option to copy the decision he is observing. Therefore,

$$\mathbf{P}_\sigma(x_n = \theta | \Omega_n = \{n-1\}) \geq \mathbf{P}_\sigma(x_{n-1} = \theta).$$

The Line Network: A Lyapunov Function

- Unbounded Private Beliefs

- There exists a function $\mathcal{Z}$ such that:

$$\mathcal{Z}(\alpha) > \alpha \quad \text{for all} \quad \alpha < 1$$

- If agent n is observing a single other decision $(\Omega_n = \{n-1\})$, then

$$\mathbf{P}_\sigma(x_n = \theta | \Omega_n = \{n-1\}) \geq \mathcal{Z}(\mathbf{P}_\sigma(x_{n-1} = \theta))$$

The Line Network

- Proposition: If every agent can only observe the previous decision, i.e.,$\Omega_n = \{n-1\}$, then in any equilibrium σ , $\lim_{n\to\infty} \mathbf{P}_\sigma(x_n = \theta)$ exists and:
  - if the private beliefs are unbounded, then

$$\lim_{n\to\infty} \mathbf{P}_\sigma(x_n = \theta) = 1$$

  - if the private beliefs are bounded, then

$$\lim_{n\to\infty} \mathbf{P}_\sigma(x_n = \theta) < 1$$

- For the bounded case, we show there is a positive probability everyone will choose action 1 irrespective of the state (Herding).

- The proof follows from the existence of the Lyapunov function.

Full Network [SS 00]

- If all previous decisions are observable, i.e., $\Omega_n = \{1, ..., n-1\}$, then asymptotic learning occurs if and only if private beliefs are unbounded.

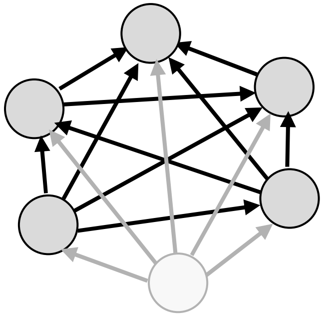

- Proof is based on a martingale convergence argument, but does not extend to general network topologies.

Are Unbounded Beliefs Enough?

- Asymptotic learning occurs if and only if the private beliefs are unbounded in the following two topologies:
  - line Network: $\Omega_n = \{n-1\}$
  - full Network: $\Omega_n = \{1, 2, ..., n-1\}$
- In what other kind of networks do unbounded beliefs characterize learning ?
- Obviously, not all possible networks. Example: $\Omega_n = \emptyset$ for all n.

General Case: Random Network

.

- The neighborhood $\Omega_n \subset \{1, \ldots n-1\}$ is generated according to an arbitrary distribution $\mathbf{Q_n}$.

- The sequence of independent distributions $\{\mathbf{Q_n}, \mathbf{n} \in \mathbf{N}\}$ is called the network topology and is common knowledge.

- Agent n Information set:

$$I_n = \{s_n, \Omega_n, x_k \text{ for all } k \in \Omega_n\}$$

- Asymptotic Learning: Under what conditions does

$$\lim_{n\to\infty} \mathbf{P}_\sigma(x_n = \theta) = 1$$

The World According to Agent 7

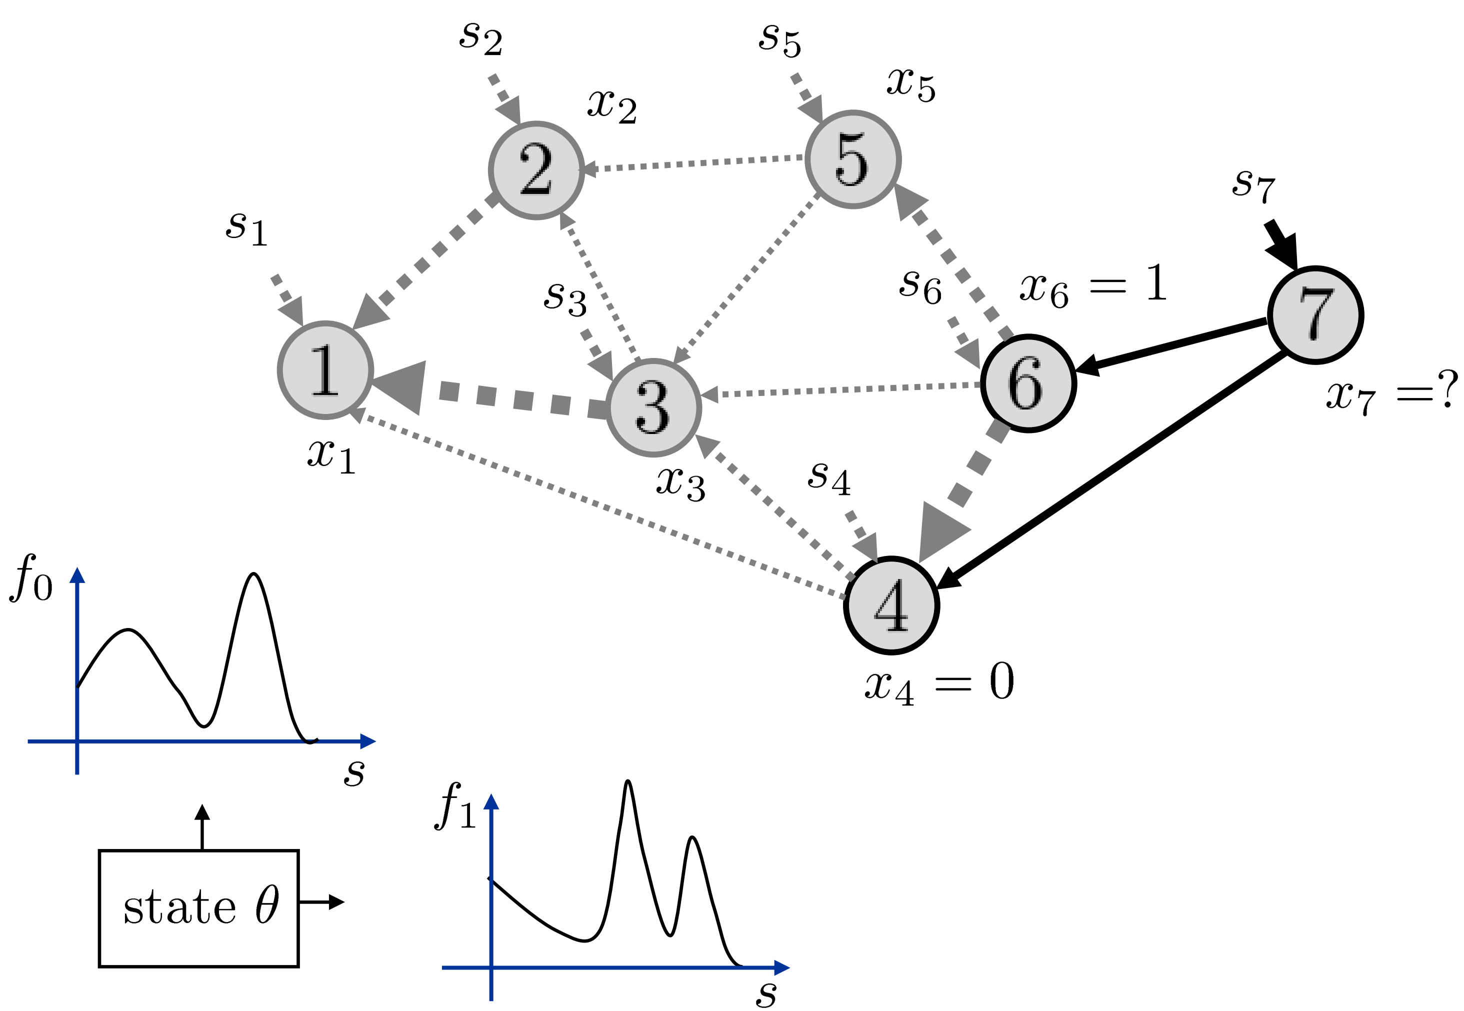

Infinitely Many Observations Are Not Enough

- We can construct examples where agents have unbounded beliefs, observe arbitrarily many decisions and still do not learn the state

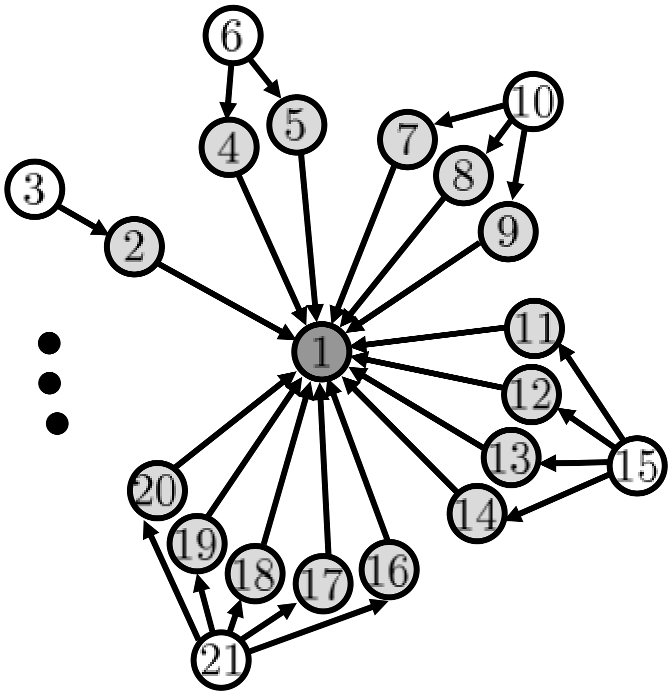

Expanding Observations

- Definition: A network topology $\{\mathbf{Q}_n\}_{n\in\mathbf{N}}$ is said to have expanding observations if for all $\epsilon > 0$, and all $K \in \mathbf{R}$, there exists some N such that for all $n \geq N$

$$\mathbf{Q}_n\left(\max_{b\in\Omega_n} b < K\right) < \epsilon$$

- Conversely

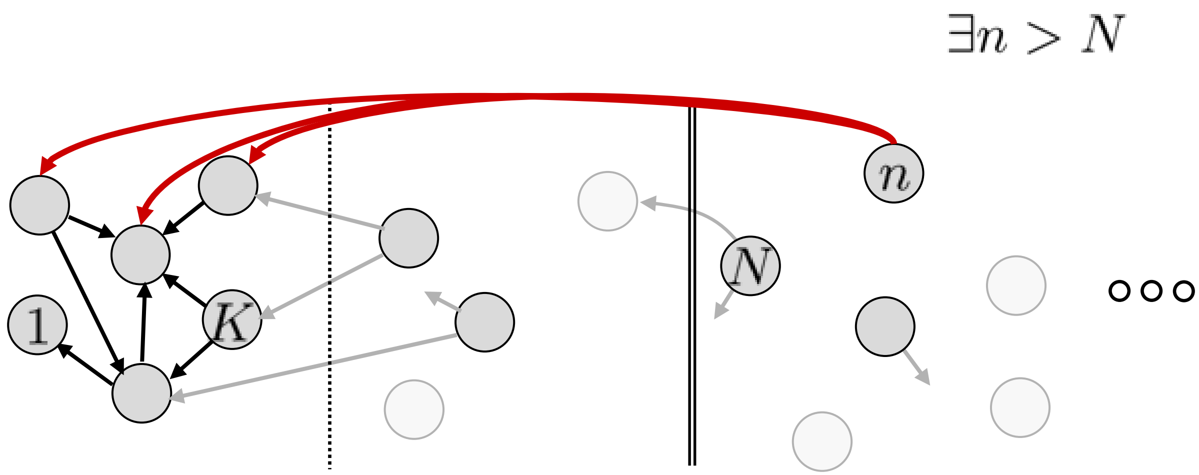

ABSENCE OF EXCESSIVELY INFLUENTIAL AGENTS

Influential References

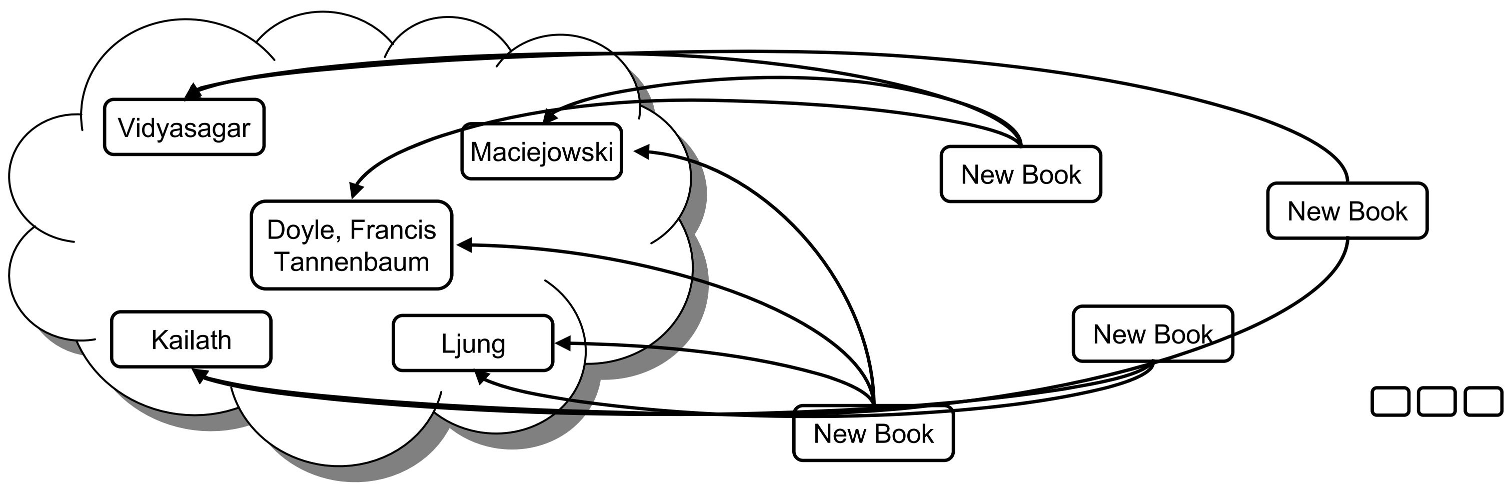

No “learning”!

Main Results

- Theorem (negative): If the topology does not have expanding observations, then asymptotic learning does not occur.

- Theorem (positive): If the private beliefs are unbounded and the topology has expanding observations, then asymptotic learning occurs.

Deterministic Networks: Information Trail

- Suppose the network topology is deterministic. For every pair such (m, n) that $m \in \Omega_n$, consider it to be an n, m-directed edge.
- Let $L(n)$ be the length of the longest path starting at node n. If private beliefs are unbounded, then asymptotic learning occurs if and only if

$$\lim_{n\to\infty} L(n) = \infty$$

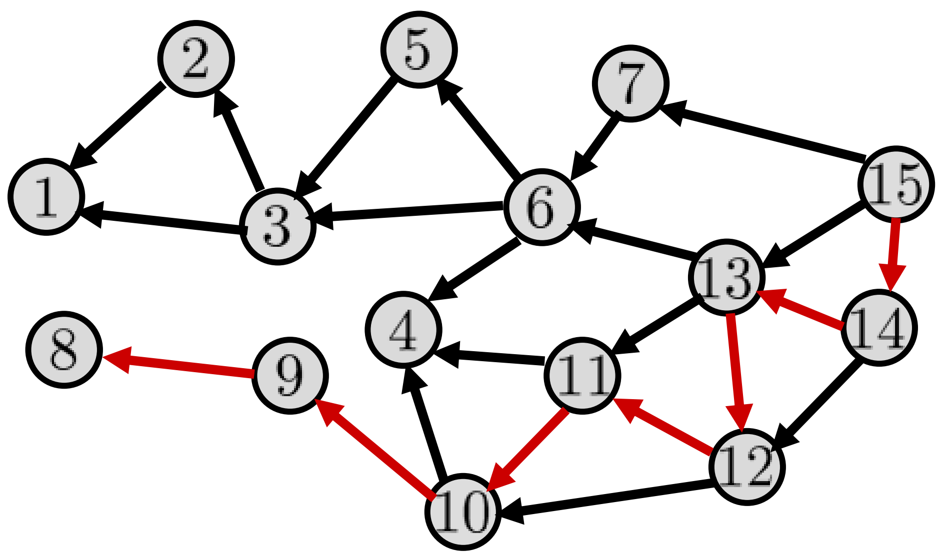

$L(15) = 7$

Deterministic Networks: Examples

- Full topology: $\Omega_n = \{1, 2, ..., n-1\}$

- Line topology: $\Omega_n = \{n-1\}$

Examples: A Random Sample

- Suppose each agent observes a sample $C > 0$ of randomly drawn (uniformly) decisions from the past. If the private beliefs are unbounded, then asymptotic learning occurs.

Examples: Binomial Sample (Erdős–Rényi)

- Suppose all links in the network are independent, and for two constants A and B we have $\mathbf{Q}_n(m \in \Omega_n) = \frac{A}{n^B}$,
  - If the beliefs are unbounded and $B < 1$, asymptotic learning occurs
  - If $B \geq 1$, asymptotic learning does not occur

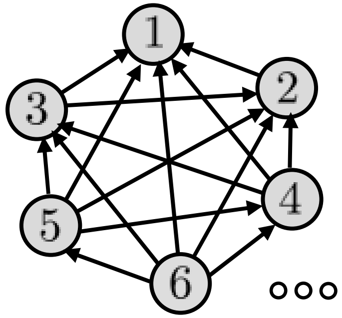

Back to Bounded Beliefs

- If the private beliefs are bounded, there exists some constant M such that $|\Omega_n| \leq M$ for all n, and

$$\lim_{n \to \infty} \max_{b \in \Omega_n} b = \infty \text{ with probability } 1$$

then asymptotic learning does not occur.

- Implication: No learning with random sampling and bounded beliefs.

Learning under Bounded Beliefs

THE VALUE OF INDEPENDENT THINKERS

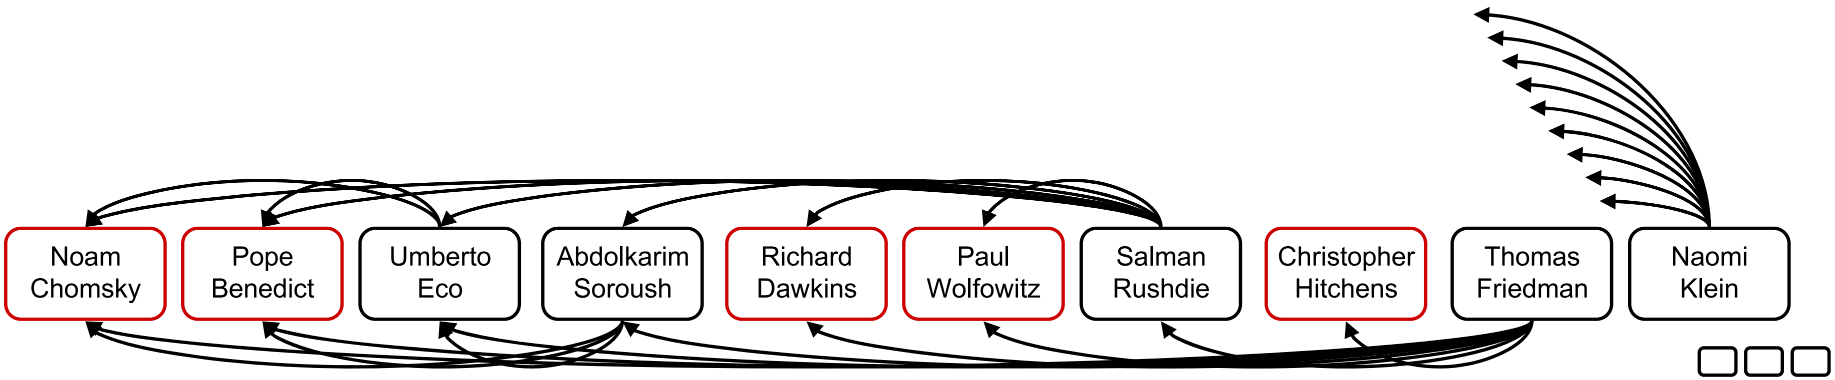

Learning under Bounded Beliefs

- There exist network topologies where asymptotic learning occurs for any signal structure

- Example:

$$\Omega_n = \begin{cases} \{1, ..., n-1\}, & \text{with probability } 1 - \frac{1}{n}; \\ \emptyset, & \text{with probability } \frac{1}{n}, \end{cases}$$

- There is a large number of agents that make a decision only based on their observations.

Summary

	Unbounded Beliefs	Bounded Beliefs
Expanding Observations	YES	USUALLY NO, SOMETIMES YES
Other Topologies	NO	NO

Preferential Attachment: Correlated Neigborhoods

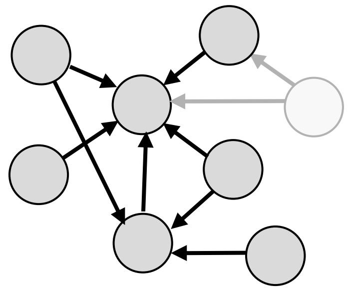

- Correlated graphs present another challenge
- Preferential Attachment: Probability of connection is proportional to the degree of each agent
- Expanding observations: only necessary for learning
- Sufficient conditions have been derived
  - Include above model under unbounded beliefs
- Open research problem

Rate of Convergence

- How does the network impact the speed of learning?

- For $\Omega_n = \{n-1\}$, then

$$\mathbf{P}_\sigma(x_n \neq \theta) = \Theta(\frac{1}{n})$$

- For the uniform sampling of one agent, i.e., $\mathbf{Q}_n(\Omega_n = \{k\}) = \frac{1}{n-1}$, then

$$\mathbf{P}_\sigma(x_n \neq \theta) = \Theta(\frac{1}{\log(n)})$$

Conclusions

- Just scratching the surface….
- Presented a simple model of information aggregation
  - Private signal
  - Network topology
- More complex models for sequential decision making
  - Dependent neighbors
  - Heterogeneous preferences
  - Multi-class agents
  - Cyclic decisions
- Rationality
- Topology measures: depth, diameter, conductance
  - Expanding observations
- Robustness

Related Literature

- Social Learning
  - Banerjee (92), Bikhchandani, Hirshleifer and Welch (92),
  - Smith and Sorensen (98, 00).

- Myopic / Boundedly Rational Learning in Networks
  - Bala and Goyal (98),
  - Kempe, Kleinberg and Tardos (03, 05).

- Decentralized Detection
  - Cover (69),
  - Athans and Papastavrou (90),
  - Tay, Tsitsiklis and Win (06, 07)