Bayesian Learning in Social Networks

Asu Ozdaglar

Joint work with

Daron Acemoglu, Munther Dahleh, Ilan Lobel

Department of Electrical Engineering and Computer Science,

Department of Economics,

Operations Research Center

Massachusetts Institute of Technology

Motivation

- **Objective:** understand information aggregation in social networks.
- **Model:**
  - Dynamic game with unknown **state of the world**
  - Sequential decisions based on private signals and observation of past actions
  - Payoff conditional on underlying state (same for all agents)
- **Question:** Under what conditions do individuals make correct decisions (or learn the state) as the social network grows bigger ?

A Simple Motivating Model

- Model for Bayesian learning on a line **[Bikchandani, Hirschleifer, Welch (92), Banerjee (92)]**
- Two possible states of the world $\theta \in \{0, 1\}$, both equally likely
- A sequence of agents $(n = 1, 2, ...)$ making decisions $x_n \in \{0, 1\}$
- Agent n obtains utility 1 if $x_n = \theta$ and utility 0 otherwise
- Each agent has an iid private binary signals s_n, where $s_n = \theta$ with probability $> 1/2$
- Agent n knows his signal s_n and the decisions of previous agents $x_1, x_2, ..., x_{n-1}$
- Agent n chooses action 1 if

$$\mathbb{P}\left(\theta = 1 | s_n, x_1, x_2, ..., x_{n-1}\right) > \mathbb{P}\left(\theta = 0 | s_n, x_1, x_2, ..., x_{n-1}\right)$$

- If $s_1 = s_2 \neq \theta$, then all agents **herd** and $x_n \neq \theta$ for all agents,

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

Asymptotic Learning on a Line

- More general model studied by **[Smith and Sorensen (00)]**
- General signals s_n
- Private beliefs bounded if the resulting likelihood ratio is bounded away from 0 and ∞
- Private beliefs unbounded otherwise
- On the line there is asymptotic learning, $\lim_{n\to\infty} \mathbb{P}(x_n = \theta) = 1$, if private beliefs are unbounded
- No asymptotic learning if private beliefs are bounded

Social Networks

- Previous work considers situations where each individual observes all past actions. Thus no study of network topology
- In practice, most information obtained from an individual's social network; friends, neighbors, co-workers...
- How does network structure affect learning?
- How to model learning over networks?

Our Model

- Two possible states of the world $\theta \in \{0, 1\}$, both equally likely
- A sequence of agents $(n = 1, 2, ...)$ making decisions $x_n \in \{0, 1\}$. Agent n obtains utility 1 if $x_n = \theta$ and utility 0 otherwise
- Each agent has an iid private signal s_n in S. The signal is generated according to distribution $\mathbb{F}_\theta$, $\mathbb{F}_0$ and $\mathbb{F}_1$ absolutely continuous with respect to each other
- $(\mathbb{F}_0, \mathbb{F}_1)$ is the signal structure
- Agent n has a neighborhood $B(n) \subseteq \{1, 2, ..., n-1\}$ and observes the decisions x_k for all $k \in B(n)$. The set $B(n)$ is private information.
- The neighborhood $B(n)$ is generated according to an arbitrary distribution $\mathbb{Q}_n$
- $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ is the network topology and is common knowledge
- A social network consists of the signal structure and network topology
- **Asymptotic Learning:** Under what conditions does $\lim_{n \to \infty} \mathbb{P}(x_n = \theta) = 1$?

Network Topologies

- $\{\mathbb{Q}_n\}_{n\in\mathbb{N}}$ assigns probability 1 to neighborhood $\{1, 2..., n-1\}$ for each $n \in \mathbb{N}$—line
- $\{\mathbb{Q}_n\}_{n\in\mathbb{N}}$ assigns probability $1/n-1$ to each one of the subsets of size 1 of $\{1, 2..., n-1\}$ for each $n \in \mathbb{N}$—random sampling
- $\{\mathbb{Q}_n\}_{n\in\mathbb{N}}$ assigns probability 1 to neighborhood $\{n-1\}$ for each $n \in \mathbb{N}$
- $\{\mathbb{Q}_n\}_{n\in\mathbb{N}}$ assigns probability 1 to neighborhoods that are subsets of $\{1, 2, ..., K\}$ for each $n \in \mathbb{N}$ for some $K \in \mathbb{N}$—example of excessively influential agents

Example Network Topology

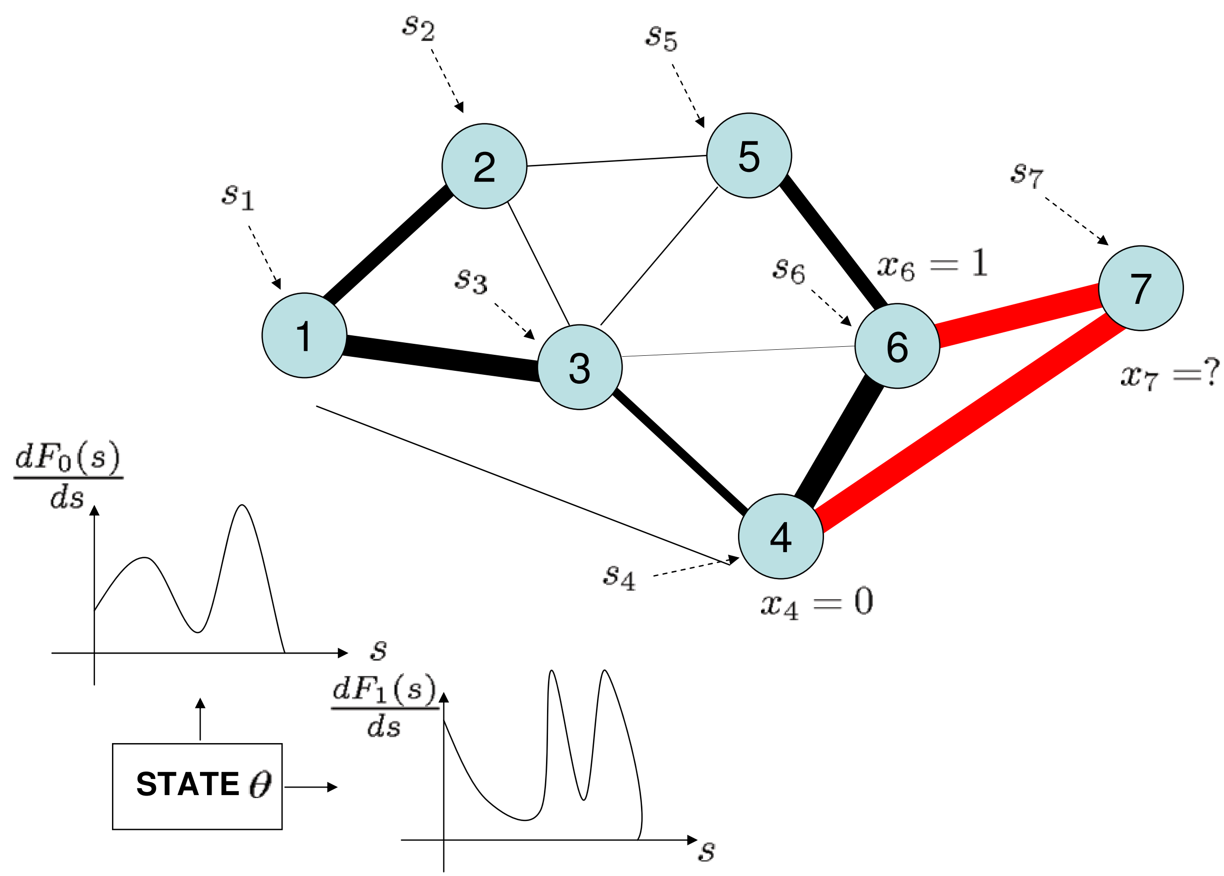

Related Literature

- **Bayesian Learning**
  - Banerjee (92), Bikhchandani, Hirshleifer and Welch (92), Smith and Sorensen (00)
  - Banerjee and Fudenberg (04), Smith and Sorensen (98), Gale and Kariv (03), Celen and Kariv (04)
- **Boundedly Rational Learning in Networks**
  - Ellison and Fudenberg (93, 95), Bala and Goyal (98, 01)
  - DeMarzo, Vayanos, Zwiebel (03), Golub and Jackson (07)
- **Decentralized Detection**
  - Cover (69), Papastavrou and Athans (90), Tay, Tsitsiklis and Win (06, 07).

Our Contributions

- We study sequential decision-making and information aggregation in social networks
- We establish decision rules used in perfect Bayesian equilibria
- When the signals lead to **unbounded private beliefs**:
  - We fully characterize the set of network topologies that lead to learning
- When the signals lead to **bounded private beliefs**:
  - We show most 'reasonable' networks do not lead to learning
  - We show learning is possible with stochastic network topologies

Perfect Bayesian Equilibria

- Agent n's information set is $I_n = \{s_n, B(n), x_k \text{ for all } k \in B(n)\}$
- A strategy for individual n is $\sigma_n : \mathcal{I}_n \to \{0, 1\}$
- A strategy profile is a sequence of strategies $\sigma = \{\sigma_n\}_{n \in \mathbb{N}}$.
  - A strategy profile σ induces a probability measure $\mathbb{P}_\sigma$ over $\{x_n\}_{n \in \mathbb{N}}$.

Definition: A strategy profile σ^* is a pure-strategy **Perfect Bayesian Equilibrium** if for each $n \in \mathbb{N}$

$$\sigma_n^*(I_n) \in \operatorname{argmax}_{y \in \{0,1\}} \mathbb{P}_{(y, \sigma_{-n}^*)}(y = \theta \mid I_n)$$

- A pure strategy PBE exists. Denote the set of PBEs by Σ^*.

Definition: Given a signal structure $(\mathbb{F}_0, \mathbb{F}_1)$ and a network topology $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$, we say that **asymptotic learning occurs in equilibrium** σ if x_n converges to θ in probability (according to measure $\mathbb{P}_\sigma$), that is,

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

Equilibrium Decision Rule

Lemma: The decision of agent n, $x_n = \sigma(I_n)$, satisfies

$$x_n = \begin{cases} 1, & \text{if } \mathbb{P}_\sigma(\theta = 1 \mid s_n) + \mathbb{P}_\sigma\big(\theta = 1 \mid B(n), x_k \text{ for all } k \in B(n)\big) > 1, \\ 0, & \text{if } \mathbb{P}_\sigma(\theta = 1 \mid s_n) + \mathbb{P}_\sigma\big(\theta = 1 \mid B(n), x_k \text{ for all } k \in B(n)\big) < 1, \end{cases}$$

and $x_n \in \{0, 1\}$ otherwise.

- The belief about the state decomposes into two parts:
  - the **Private Belief**: $\mathbb{P}_\sigma(\theta = 1 \mid s_n)$;
  - the **Social Belief**: $\mathbb{P}_\sigma(\theta = 1 \mid B(n), x_k \text{ for all } k \in B(n))$.

Private Beliefs

Lemma: The private belief of agent n is

$$p_n(s_n) = \mathbb{P}_\sigma(\theta = 1 | s_n) = \left(1 + \frac{d\mathbb{F}_0(s_n)}{d\mathbb{F}_1(s_n)}\right)^{-1}.$$

Definition: The signal structure has **bounded private beliefs** if there exists some $0 < m, M < \infty$ such that the Radon-Nikodym derivate $d\mathbb{F}_0/d\mathbb{F}_1$ satisfies

$$m < \frac{d\mathbb{F}_0}{d\mathbb{F}_1}(s) < M,$$

for almost all $s \in S$ under measure $(\mathbb{F}_0 + \mathbb{F}_1)/2$. The signal structure has **unbounded private beliefs** if

$$\inf_{s \in S} \frac{d\mathbb{F}_0}{d\mathbb{F}_1}(s) = 0 \quad \text{and} \quad \sup_{s \in S} \frac{d\mathbb{F}_0}{d\mathbb{F}_1}(s) = \infty.$$

- Bounded private beliefs $\Leftrightarrow$ bounded likelihood ratio
- 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**.

Social Beliefs Need Not Be Monotone

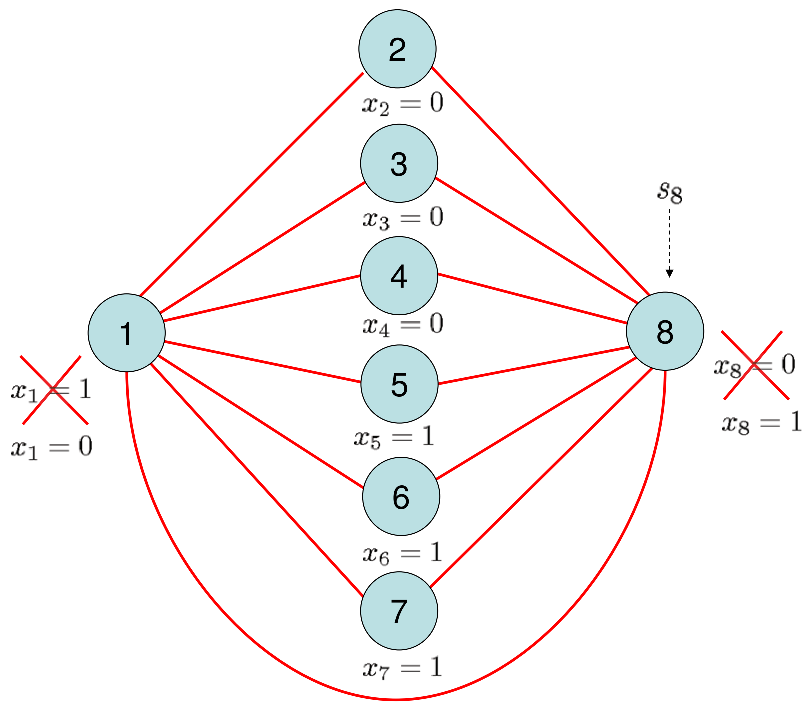

- There exist signal structures $(\mathbb{F}_0, \mathbb{F}_1)$ such that for all equilibria σ,

$$\mathbb{P}_\sigma (\theta = 1 | x_1 = \ldots = x_4 = 0, x_5 = \ldots = x_7 = 1) >$$
$$\mathbb{P}_\sigma (\theta = 1 | x_2 = \ldots = x_4 = 0, x_1 = x_5 = \ldots = x_7 = 1)$$

- Need a strategy of analysis not relying on monotonicity

Properties of Network Topology

Definition: A network topology $\{\mathbb{Q}_n\}_{n\in\mathbb{N}}$ has **expanding observations** if for all K,

$$\lim_{n\to\infty} \mathbb{Q}_n \left(\max_{b\in B(n)} b < K \right) = 0.$$

Otherwise, it has **nonexpanding observations**

- Expanding observations do not imply connected graph
- Nonexpanding observations equivalently : There exists some K, $\epsilon > 0$ and an infinite subset $\mathcal{N} \in \mathbb{N}$ such that

$$\mathbb{Q}_n \left(\max_{b\in B(n)} b < K \right) \geq \epsilon \quad \text{for all} \quad n \in \mathcal{N}.$$

- A finite group of agents is **excessively influential** if there exists an infinite number of agents who, with probability uniformly bounded away from 0, observe only the actions of a subset of this group.
  - For example, a group is excessively influential if it is the source of *all information* for an infinitely large component of the network
- Nonexpanding observations $\Leftrightarrow$ excessively influential agents

Main Results

Theorem 1: Assume that the network topology $\{\mathbb{Q}_n\}_{n\in\mathbb{N}}$ has nonexpanding observations. Then, there exists no equilibrium $\sigma \in \Sigma^*$ with asymptotic learning.

Theorem 2: Assume that the signal structure $(\mathbb{F}_0, \mathbb{F}_1)$ has unbounded private beliefs and the network topology $\{\mathbb{Q}_n\}_{n\in\mathbb{N}}$ has expanding observations. Then, asymptotic learning occurs in every equilibrium $\sigma \in \Sigma^*$.

Deterministic Topologies

- In a deterministic network, π *is an information path* of agent n if for each i, $\pi_i \in B(\pi_{i+1})$ and the last element of π is n. The *information depth L(n)* is the number of elements in the maximal $\pi(n)$.

Corollary: Assume that the signal structure $(\mathbb{F}_0, \mathbb{F}_1)$ has unbounded private beliefs and that the network topology is deterministic. Then, asymptotic learning occurs for all equilibria if and only if $\{L(n)\}_{n \in \mathbb{N}}$ goes to infinity.

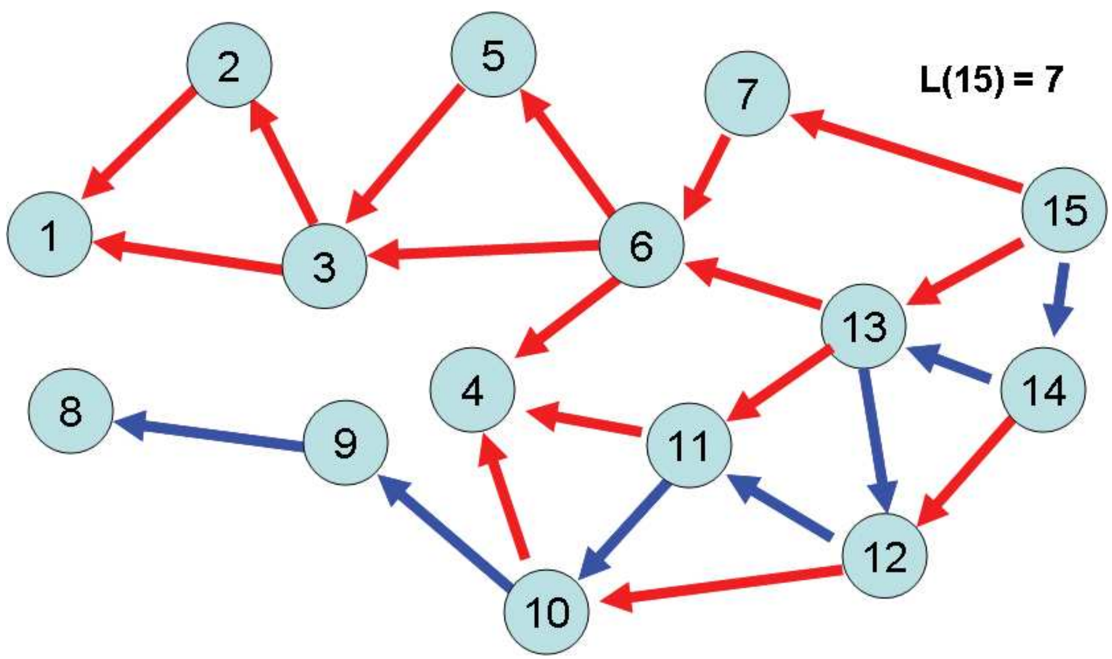

Proof Idea of Theorem 1

- Since nonexpanding observations, there exists some K, $\epsilon > 0$ and an infinite subset $\mathcal{N} \subset \mathbb{N}$ such that

$$\mathbb{Q}_n \left(\max_{b \in B(n)} b < K \right) \geq \epsilon \text{ for all } n \in \mathcal{N}.$$

- Then, for any $n \in \mathcal{N}$ and any equilibrium σ,

$$\begin{aligned} \mathbb{P}_\sigma(x_n = \theta) &= \mathbb{P}_\sigma \left(x_n = \theta \,\middle|\, \max_{b \in B(n)} b < K \right) \mathbb{Q}_n \left(\max_{b \in B(n)} b < K \right) \\ &\quad + \mathbb{P}_\sigma \left(x_n = \theta \,\middle|\, \max_{b \in B(n)} b \geq K \right) \mathbb{Q}_n \left(\max_{b \in B(n)} b \geq K \right) \\ &\leq 1 - \epsilon + \epsilon \mathbb{P}_\sigma \left(x_n = \theta \,\middle|\, \max_{b \in B(n)} b < K \right) \end{aligned}$$

- Let f give the best estimate of the state given a finite set of iid signals

$$\mathbb{P}_\sigma \left(x_n = \theta \,\middle|\, \max_{b \in B(n)} b < K \right) \leq \mathbb{P}\left(f(s_1, s_2, ..., s_{K-1}, s_n) = \theta \right) < 1$$

- The result follows

Proof of Theorem 2: Roadmap

- Characterization of equilibrium strategies when observing a single agent
- **Strong improvement principle** when observing one agent
- **Generalized strong improvement principle**
- Asymptotic learning with unbounded private beliefs and expanding observations

Observing a Single Decision

- Given σ and n, let us define Y_n^σ and N_n^σ as

$$Y_n^\sigma = \mathbb{P}_\sigma(x_n = 1 \mid \theta = 1), \qquad N_n^\sigma = \mathbb{P}_\sigma(x_n = 0 \mid \theta = 0).$$

- The unconditional probability of a correct decision is

$$\frac{1}{2}(Y_n^\sigma + N_n^\sigma) = \mathbb{P}_\sigma(x_n = \theta)$$

- We also define the *thresholds* L_n^σ and U_n^σ in terms of these probabilities:

$$L_n^\sigma = \frac{1 - N_n^\sigma}{1 - N_n^\sigma + Y_n^\sigma}, \qquad U_n^\sigma = \frac{N_n^\sigma}{N_n^\sigma + 1 - Y_n^\sigma}.$$

Proposition: Let $B(n) = \{b\}$ for agent n. Agent n's decision x_n in $\sigma \in \Sigma^*$ satisfies

$$x_n = \begin{cases} 0, & \text{if } p_n < L_b^\sigma \\ x_b, & \text{if } p_n \in (L_b^\sigma, U_b^\sigma) \\ 1, & \text{if } p_n > U_b^\sigma. \end{cases}$$

Observing a Single Decision (continued)

- Let the conditional distribution of private belief p be

$$\mathbb{G}_j(r) = \mathbb{P}(p \leq r \mid \theta = j)$$

- Let $\underline{\beta}$ and $\overline{\beta}$ be the lower and upper support of private beliefs
- Equilibrium decisions:

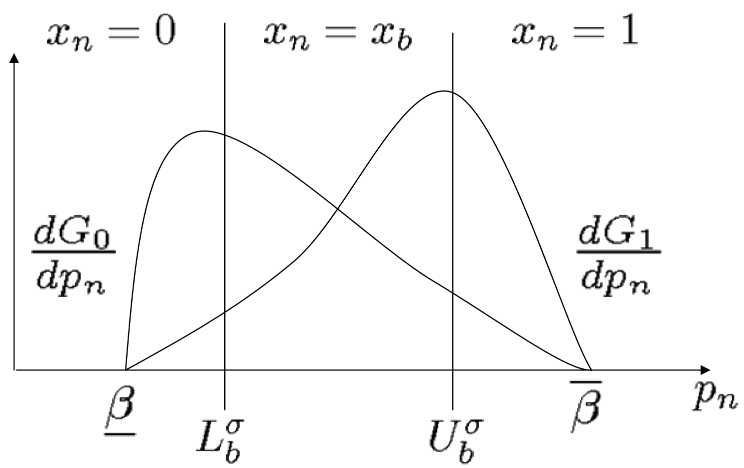

Strong Improvement Principle

- Agent n has the option of copying the action of any agent in his neighborhood:

$$\mathbb{P}_\sigma(x_n = \theta \mid B(n) = \mathfrak{B}) \geq \max_{b \in \mathfrak{B}} \mathbb{P}_\sigma(x_b = \theta).$$

  - Similar to the *welfare improvement principle* in Banerjee and Fudenberg (04) and Smith and Sorensen (98), and *imitation principle* in Gale and Kariv (03)

- Using the equilibrium decision rule and the properties of private beliefs, we establish a **strict gain** of agent n over agent b.

Proposition: (Strong Improvement Principle) Let $B(n) = \{b\}$ for some n and $\sigma \in \Sigma^*$ be an equilibrium. There exists a continuous, increasing function $\mathcal{Z} : [1/2, 1] \to [1/2, 1]$ with $\mathcal{Z}(\alpha) \geq \alpha$ such that

$$\mathbb{P}_\sigma(x_n = \theta \mid B(n) = \{b\}) \geq \mathcal{Z}\left(\mathbb{P}_\sigma(x_b = \theta)\right).$$

If the private beliefs are unbounded, then:

- $\mathcal{Z}(\alpha) > \alpha$ for all $\alpha < 1$
- $\alpha = 1$ is the unique fixed point of $\mathcal{Z}(\alpha)$

Generalized Strong Improvement Principle

- When multiple agents in the neighborhood, learning no worse than observing just one of them:

Proposition (Generalized Strong Improvement Principle) For any $n \in \mathbb{N}$, any set $\mathfrak{B} \subseteq \{1, ..., n-1\}$ and any equilibrium $\sigma \in \mathcal{S}$, we have

$$\mathbb{P}_\sigma \left(x_n = \theta \mid B(n) = \mathfrak{B} \right) \geq \mathcal{Z} \left(\max_{b \in \mathfrak{B}} \mathbb{P}_\sigma (x_b = \theta) \right).$$

Proof of Theorem 2

- Under expanding observations, one can construct a sequence of agents along which the generalized strong improvement principle applies
- Unbounded private beliefs imply that along this sequence $\mathcal{Z}(\alpha)$ strictly increases
- Until unique fixed point $\alpha = 1$, corresponding to **asymptotic learning**

No Learning under Bounded Beliefs

Theorem 3: If the private beliefs are bounded and the network topology satisfies one of the following conditions,

(a) $B(n) = \{1, ..., n-1\}$ for all n,

(b) $|B(n)| \leq 1$ for all n,

(c) there exists some constant M such that $|B(n)| \leq M$ for all n and

$$\lim_{n\to\infty} \max_{b\in B(n)} b = \infty \text{ with probability 1,}$$

then asymptotic learning does not occur.

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

Proof Idea - Theorem 3(c):

- Asymptotic learning implies social beliefs converge to 0 or 1 almost surely
- But with bounded beliefs, this implies individuals decide on the basis of social belief alone
- Then, positive probability of mistake–contradiction

Learning under Bounded Beliefs

Theorem 4: There exist network topologies where asymptotic learning occurs for any signal structure $(\mathbb{F}_0, \mathbb{F}_1)$.

- In the paper, characterization of a class of network topologies for which asymptotic learning occurs with bounded beliefs

Example: For all n,

$$B(n) = \begin{cases} \{1, ..., n-1\}, & \text{with probability } 1 - r(n); \\ \emptyset, & \text{with probability } r(n), \end{cases}$$

for some sequence $\{r(n)\}$ where $\lim_{n\to\infty} r(n) = 0$ and $\sum_{n=1}^{\infty} r(n) = \infty$.

In this case, asymptotic learning occurs for an arbitrary signal structure $(\mathbb{F}_0, \mathbb{F}_1)$ and at any equilibrium.

Proof Idea

- Individuals with empty neighborhood must act according to their private beliefs
- If they are identified by a marker, then simply apply weak law of large numbers
- For the stochastic network topology, we prove that eventually all agents with $B(n) = \{1, ..., n-1\}$ converge on a decision using martingale convergence.
- Eventually, everyone can identify the agents with $B(n) = \emptyset$ and extract true state from them using weak law of large numbers.

Summary

- When does asymptotic learning occur ?

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

- No asymptotic learning with unbounded beliefs due to excessively influential agents
- If there is a group of agents who are "influential", but not excessively so (for example, overrepresented in the information sets of others), this does not prevent asymptotic learning with unbounded beliefs $\Rightarrow$ *contrast with myopic learning*

Future Directions

- How does the rate of learning with unbounded beliefs depend on network topology?
- With bounded beliefs, how does the structure of the social network affect probability of wrong asymptotic beliefs?
- Learning in social networks with repeated actions and observations
- How does network structure interact with learning when underlying state is changing?
- Heterogeneous preferences