Bayesian Learning in Social Networks

Ilan Lobel

NYU

Simons Institute
October 2022

With Daron Acemoglu, Munther Dahleh and Asu Ozdaglar (MIT)
Review of Economic Studies 2011

The Starting Point of the Social Learning Literature

People often copy the actions of others

- Product going viral
- Meme stock trading

Herd behavior is an important economic phenomenon

- Think asset market bubble

But can it be rational?

- Two seminal papers (Bikhchandani, Hirshleifer, Welch 1992, Banerjee 1992) argued yes.

The Starting Point of the Social Learning Literature

People often copy the actions of others

- Product going viral
- Meme stock trading

Herd behavior is an important economic phenomenon

- Think asset market bubble

But can it be rational?

- Two seminal papers (Bikhchandani, Hirshleifer, Welch 1992, Banerjee 1992) argued yes.

The Starting Point of the Social Learning Literature

People often copy the actions of others

- Product going viral
- Meme stock trading

Herd behavior is an important economic phenomenon

- Think asset market bubble

But can it be rational?

- Two seminal papers (Bikhchandani, Hirshleifer, Welch 1992, Banerjee 1992) argued yes.

The BHW 1992 Model: An Example

Agents arrive in town sequentially and need to choose a restaurant:

- Chinese or Indian food?

One restaurant is better, but no one knows which one (equal priors).

Agents have independent private signals indicating where to go.

- Signal is correct with 70% probability.

Agents observe choices of others but not their signals.

Realization:

The BHW 1992 Model: An Example

Agents arrive in town sequentially and need to choose a restaurant:

- Chinese or Indian food?

One restaurant is better, but no one knows which one (equal priors).

Agents have independent private signals indicating where to go.

- Signal is correct with 70% probability.

Agents observe choices of others but not their signals.

Realization:

The BHW 1992 Model: An Example

Agents arrive in town sequentially and need to choose a restaurant:

- Chinese or Indian food?

One restaurant is better, but no one knows which one (equal priors).

Agents have independent private signals indicating where to go.

- Signal is correct with 70% probability.

Agents observe choices of others but not their signals.

Realization:

The BHW 1992 Model: An Example

Agents arrive in town sequentially and need to choose a restaurant:

- Chinese or Indian food?

One restaurant is better, but no one knows which one (equal priors).

Agents have independent private signals indicating where to go.

- Signal is correct with 70% probability.

Agents observe choices of others but not their signals.

Realization:

The BHW 1992 Model: An Example

Agents arrive in town sequentially and need to choose a restaurant:

- Chinese or Indian food?

One restaurant is better, but no one knows which one (equal priors).

Agents have independent private signals indicating where to go.

- Signal is correct with 70% probability.

Agents observe choices of others but not their signals.

Realization:

The BHW 1992 Model: An Example

Agents arrive in town sequentially and need to choose a restaurant:

- Chinese or Indian food?

One restaurant is better, but no one knows which one (equal priors).

Agents have independent private signals indicating where to go.

- Signal is correct with 70% probability.

Agents observe choices of others but not their signals.

Realization:

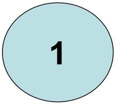

Signal = ‘Chinese’
Decision = ‘Chinese’

The BHW 1992 Model: An Example

Agents arrive in town sequentially and need to choose a restaurant:

- Chinese or Indian food?

One restaurant is better, but no one knows which one (equal priors).

Agents have independent private signals indicating where to go.

- Signal is correct with 70% probability.

Agents observe choices of others but not their signals.

Realization:

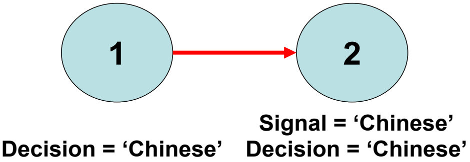

The BHW 1992 Model: An Example

Agents arrive in town sequentially and need to choose a restaurant:

- Chinese or Indian food?

One restaurant is better, but no one knows which one (equal priors).

Agents have independent private signals indicating where to go.

- Signal is correct with 70% probability.

Agents observe choices of others but not their signals.

Realization:

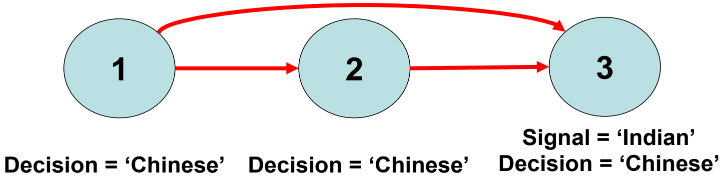

The BHW 1992 Model: An Example

Agents arrive in town sequentially and need to choose a restaurant:

- Chinese or Indian food?

One restaurant is better, but no one knows which one (equal priors).

Agents have independent private signals indicating where to go.

- Signal is correct with 70% probability.

Agents observe choices of others but not their signals.

Realization:

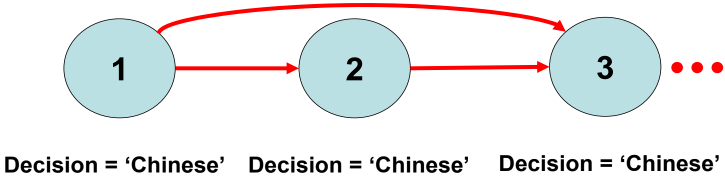

Asymptotic Learning

There are infinitely many independent signals about the state.

- If agents fully shared what they knew, they'd figure out the state.
- The "Wisdom of Crowds" (Condorcet 1788, Galton 1906).

The wisdom goes away if people observe only actions of others.

- The probability of making a bad decision stays bounded away from zero as n grows (**failure of asymptotic learning**).

Asymptotic Learning

There are infinitely many independent signals about the state.

- If agents fully shared what they knew, they'd figure out the state.
- The "Wisdom of Crowds" (Condorcet 1788, Galton 1906).

The wisdom goes away if people observe only actions of others.

- The probability of making a bad decision stays bounded away from zero as n grows (**failure of asymptotic learning**).

Signals

Are these results driven by the signal structure?

Smith and Sorensen 2000: state $\theta \in \{0, 1\}$, signal s_n drawn from $\mathbb{F}_\theta$.

- Define private belief $p_n = P(\theta = 1|s_n)$.
- Let $\underline{p} = \inf_s P(\theta = 1|s)$ and $\bar{p} = \sup_s P(\theta = 1|s)$.
- If $\underline{p} > 0$ and $\bar{p} < 1$, then **private beliefs are bounded**.
- If $\underline{p} = 0$ and $\bar{p} = 1$, then **private beliefs are unbounded**.

Theorem

If private beliefs are bounded, asymptotic learning fails.
If private beliefs are unbounded, asymptotic learning succeeds.

Signals

Are these results driven by the signal structure?

Smith and Sorensen 2000: state $\theta \in \{0, 1\}$, signal s_n drawn from $\mathbb{F}_\theta$.

- Define private belief $p_n = P(\theta = 1 | s_n)$.
- Let $\underline{p} = \inf_s P(\theta = 1 | s)$ and $\bar{p} = \sup_s P(\theta = 1 | s)$.
- If $\underline{p} > 0$ and $\bar{p} < 1$, then **private beliefs are bounded**.
- If $\underline{p} = 0$ and $\bar{p} = 1$, then **private beliefs are unbounded**.

Theorem

If private beliefs are bounded, asymptotic learning fails.
If private beliefs are unbounded, asymptotic learning succeeds.

Signals

Are these results driven by the signal structure?

Smith and Sorensen 2000: state $\theta \in \{0, 1\}$, signal s_n drawn from $\mathbb{F}_\theta$.

- Define private belief $p_n = P(\theta = 1|s_n)$.
- Let $\underline{p} = \inf_s P(\theta = 1|s)$ and $\bar{p} = \sup_s P(\theta = 1|s)$.
- If $\underline{p} > 0$ and $\bar{p} < 1$, then **private beliefs are bounded**.
- If $\underline{p} = 0$ and $\bar{p} = 1$, then **private beliefs are unbounded**.

Theorem
If private beliefs are bounded, asymptotic learning fails.
If private beliefs are unbounded, asymptotic learning succeeds.

Private Beliefs

The private belief of agent n is

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

The signal structure has unbounded private beliefs if

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

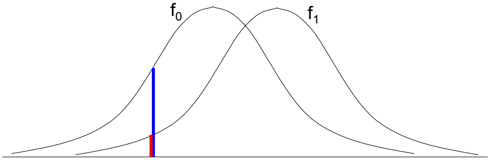

The Martingale Approach

Define the social belief $q_n = P(\theta = 1 | x_1, ..., x_n)$.

Since everyone observes all prior actions, $\{x_i\}$ defines a filtration.

The social beliefs $\{q_i\}$ are a martingale with respect to $\{x_i\}$.

This is a bounded martingale. By the martingale convergence theorem, almost all sample paths converge.

Sample paths of $\{q_i\}$ must converge to points where new private signals barely affect them.

With unbounded private beliefs, this means $\{q_i\}$ converges to $\{0, 1\}$.

Rationality implies learning since beliefs can't be fully wrong.

With bounded private beliefs, learning gets stuck away from $\{0, 1\}$.

Learning over Social Networks

Assumption so far: everyone observes the actions of **all** predecessors.

- At the heart of the proof technique.
- At the same time, it's an unrealistic assumption.

Can we study social learning if agents are embedded in a complex social network?

- A social network is more than a deterministic graph.
- Think complex random graph.
- People in the network only know their local neighborhood.
- They form beliefs on the underlying graph structure based on actions they observe.

Learning over Social Networks

Assumption so far: everyone observes the actions of **all** predecessors.

- At the heart of the proof technique.
- At the same time, it's an unrealistic assumption.

Can we study social learning if agents are embedded in a complex social network?

- A social network is more than a deterministic graph.
- Think complex random graph.
- People in the network only know their local neighborhood.
- They form beliefs on the underlying graph structure based on actions they observe.

Learning over Social Networks

Assumption so far: everyone observes the actions of **all** predecessors.

- At the heart of the proof technique.
- At the same time, it's an unrealistic assumption.

Can we study social learning if agents are embedded in a complex social network?

- A social network is more than a deterministic graph.
- Think complex random graph.
- People in the network only know their local neighborhood.
- They form beliefs on the underlying graph structure based on actions they observe.

Learning over Social Networks

Assumption so far: everyone observes the actions of **all** predecessors.

- At the heart of the proof technique.
- At the same time, it's an unrealistic assumption.

Can we study social learning if agents are embedded in a complex social network?

- A social network is more than a deterministic graph.
- Think complex random graph.
- People in the network only know their local neighborhood.
- They form beliefs on the underlying graph structure based on actions they observe.

Learning over Social Networks

Assumption so far: everyone observes the actions of **all** predecessors.

- At the heart of the proof technique.
- At the same time, it's an unrealistic assumption.

Can we study social learning if agents are embedded in a complex social network?

- A social network is more than a deterministic graph.
- Think complex random graph.
- People in the network only know their local neighborhood.
- They form beliefs on the underlying graph structure based on actions they observe.

Learning over Social Networks

Assumption so far: everyone observes the actions of **all** predecessors.

- At the heart of the proof technique.
- At the same time, it's an unrealistic assumption.

Can we study social learning if agents are embedded in a complex social network?

- A social network is more than a deterministic graph.
- Think complex random graph.
- People in the network only know their local neighborhood.
- They form beliefs on the underlying graph structure based on actions they observe.

Our Model — States, Decisions and Signals

State of the world:

- Two possible states $\theta \in \{0, 1\}$, both equally likely.

Decisions:

- A sequence of agents $(n \in \mathbb{N})$ making decisions $x_n \in \{0, 1\}$.
- Agent n obtains utility 1 if $x_n = \theta$ and utility 0 otherwise.

Signals:

- Each agent has an iid private signal s_n in an arbitrary space S.
- The signal is generated according to distribution $\mathbb{F}_\theta$. The pair $(\mathbb{F}_0, \mathbb{F}_1)$ is the signal structure.

Our Model — States, Decisions and Signals

State of the world:

- Two possible states $\theta \in \{0, 1\}$, both equally likely.

Decisions:

- A sequence of agents ($n \in \mathbb{N}$) making decisions $x_n \in \{0, 1\}$.
- Agent n obtains utility 1 if $x_n = \theta$ and utility 0 otherwise.

Signals:

- Each agent has an iid private signal s_n in an arbitrary space S.
- The signal is generated according to distribution $\mathbb{F}_\theta$. The pair $(\mathbb{F}_0, \mathbb{F}_1)$ is the signal structure.

Our Model — States, Decisions and Signals

State of the world:

- Two possible states $\theta \in \{0, 1\}$, both equally likely.

Decisions:

- A sequence of agents $(n \in \mathbb{N})$ making decisions $x_n \in \{0, 1\}$.
- Agent n obtains utility 1 if $x_n = \theta$ and utility 0 otherwise.

Signals:

- Each agent has an iid private signal s_n in an arbitrary space S.
- The signal is generated according to distribution $\mathbb{F}_\theta$. The pair $(\mathbb{F}_0, \mathbb{F}_1)$ is the signal structure.

Our Model — The Social Network

Neighborhoods:

- 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 neighborhood $B(n)$ is private information.
- The set $B(n)$ is generated according to a distribution $\mathbb{Q}_n$.
- The neighborhoods of the different agents are independent.
- $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ is the network topology and is common knowledge.

Private information:

- Agent n's information set is $\mathcal{I}_n = \{s_n, B(n), x_k \text{ for all } k \in B(n)\}$.

Social network = signal structure + network topology.

Our Model — The Social Network

Neighborhoods:

- 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 neighborhood $B(n)$ is private information.
- The set $B(n)$ is generated according to a distribution $\mathbb{Q}_n$.
- The neighborhoods of the different agents are independent.
- $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ is the network topology and is common knowledge.

Private information:

- Agent n's information set is $\mathcal{I}_n = \{s_n, B(n), x_k \text{ for all } k \in B(n)\}$.

Social network = signal structure + network topology.

Our Model — The Social Network

Neighborhoods:

- 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 neighborhood $B(n)$ is private information.
- The set $B(n)$ is generated according to a distribution $\mathbb{Q}_n$.
- The neighborhoods of the different agents are independent.
- $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ is the network topology and is common knowledge.

Private information:

- Agent n's information set is $\mathcal{I}_n = \{s_n, B(n), x_k \text{ for all } k \in B(n)\}$.

Social network = signal structure + network topology.

An Example of a Social Network

For every agent n, the signal $s_n \sim N(\theta, 1)$. For each agent $n > 1$,

$$B(n) = \begin{cases} \emptyset, & \text{with probability } 1/3; \\ \{n-1\}, & \text{with probability } 1/3; \\ \{1, ..., n-1\}, & \text{with probability } 1/3. \end{cases}$$

Realization:

An Example of a Social Network

For every agent n, the signal $s_n \sim N(\theta, 1)$. For each agent $n > 1$,

$$B(n) = \begin{cases} \emptyset, & \text{with probability } 1/3; \\ \{n-1\}, & \text{with probability } 1/3; \\ \{1, ..., n-1\}, & \text{with probability } 1/3. \end{cases}$$

Realization:

- The state is $\theta = 0$.
- All private signals s_n are iid Gaussian(0,1).

An Example of a Social Network

For every agent n, the signal $s_n \sim N(\theta, 1)$. For each agent $n > 1$,

$$B(n) = \begin{cases} \emptyset, & \text{with probability } 1/3; \\ \{n-1\}, & \text{with probability } 1/3; \\ \{1, ..., n-1\}, & \text{with probability } 1/3. \end{cases}$$

Realization:

- Agent 1 arrives.
- His signal is $s_1 = -0.4$ and his neighborhood is $B(1) = \emptyset$.
- He chooses action $x_1 = 0$.

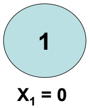

An Example of a Social Network

For every agent n, the signal $s_n \sim N(\theta, 1)$. For each agent $n > 1$,

$$B(n) = \begin{cases} \emptyset, & \text{with probability } 1/3; \\ \{n-1\}, & \text{with probability } 1/3; \\ \{1, \ldots, n-1\}, & \text{with probability } 1/3. \end{cases}$$

Realization:

- Agent 2 arrives.
- Her signal is $s_2 = -0.1$ and her neighborhood is $B(2) = \{1\}$.
- She chooses action $x_2 = 0$.

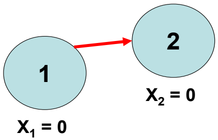

An Example of a Social Network

For every agent n, the signal $s_n \sim N(\theta, 1)$. For each agent $n > 1$,

$$B(n) = \begin{cases} \emptyset, & \text{with probability } 1/3; \\ \{n-1\}, & \text{with probability } 1/3; \\ \{1, ..., n-1\}, & \text{with probability } 1/3. \end{cases}$$

Realization:

- Agent 3 arrives.
- Her signal is $s_3 = 0.7$ and her neighborhood is $B(3) = \emptyset$.
- She chooses action $x_3 = 1$.

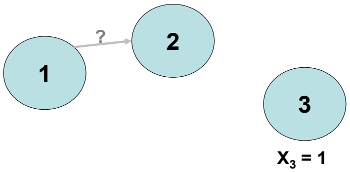

An Example of a Social Network

For every agent n, the signal $s_n \sim N(\theta, 1)$. For each agent $n > 1$,

$$B(n) = \begin{cases} \emptyset, & \text{with probability } 1/3; \\ \{n-1\}, & \text{with probability } 1/3; \\ \{1, ..., n-1\}, & \text{with probability } 1/3. \end{cases}$$

Realization:

- Agent 4 arrives.
- His signal is $s_4 = 0.4$ and his neighborhood is $B(4) = \{1, 2, 3\}$.
- Agent 4 must solve a complex estimation problem!

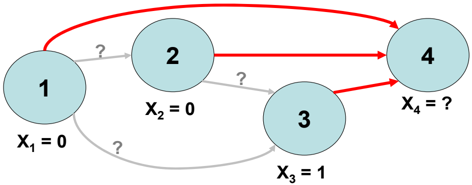

Solution Concept

A pure strategy σ_n for individual n is a mapping from $\mathcal{I}_n$ 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

Strategy profile σ^* is a pure-strategies Perfect Bayesian Equilibrium if

$$\sigma_n^*(\mathcal{I}_n) \in \arg\max_{y \in \{0,1\}} \mathbb{P}_{(y, \sigma^*_{-n})}(y = \theta \,|\mathcal{I}_n) \quad \text{for each } n \in \mathbb{N}.$$

A pure strategies PBE exists. We denote the set of PBEs by Σ^*.

Definition

We say that asymptotic learning occurs in equilibrium σ if

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

Solution Concept

A pure strategy σ_n for individual n is a mapping from $\mathcal{I}_n$ 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

Strategy profile σ^* is a pure-strategies Perfect Bayesian Equilibrium if

$$\sigma_n^*(\mathcal{I}_n) \in \arg\max_{y\in\{0,1\}} \mathbb{P}_{(y,\sigma_{-n}^*)}(y = \theta \,|\mathcal{I}_n) \quad \text{for each } n \in \mathbb{N}.$$

A pure strategies PBE exists. We denote the set of PBEs by Σ^*.

Definition

We say that asymptotic learning occurs in equilibrium σ if

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

Solution Concept

A pure strategy σ_n for individual n is a mapping from $\mathcal{I}_n$ 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

Strategy profile σ^* is a pure-strategies Perfect Bayesian Equilibrium if

$$\sigma_n^*(\mathcal{I}_n) \in \arg\max_{y \in \{0,1\}} \mathbb{P}_{(y, \sigma_{-n}^*)}(y = \theta \,|\mathcal{I}_n) \quad \text{for each } n \in \mathbb{N}.$$

A pure strategies PBE exists. We denote the set of PBEs by Σ^*.

Definition

We say that asymptotic learning occurs in equilibrium σ if

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

Bayesian Learning over Networks is Hard to Analyze

- No monotonicity.

Bayesian Learning over Networks is Hard to Analyze

- No monotonicity.

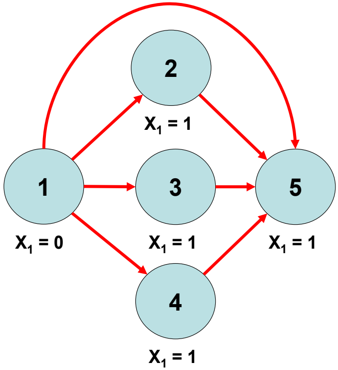

Bayesian Learning over Networks is Hard to Analyze

- No monotonicity.

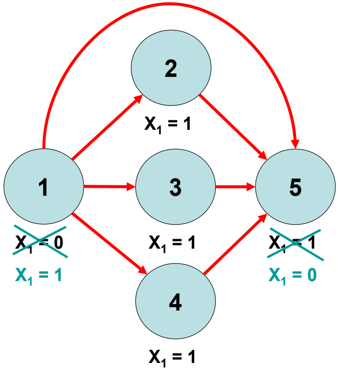

Bayesian Learning over Networks is Hard to Analyze

- No monotonicity.

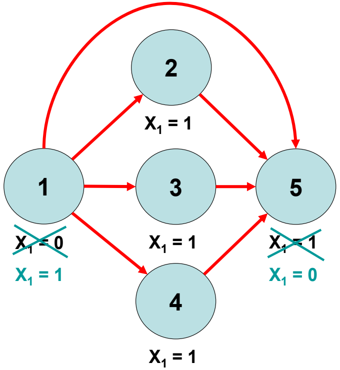

- Less can be more.

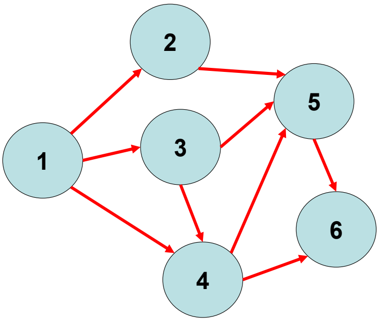

Bayesian Learning over Networks is Hard to Analyze

- No monotonicity.

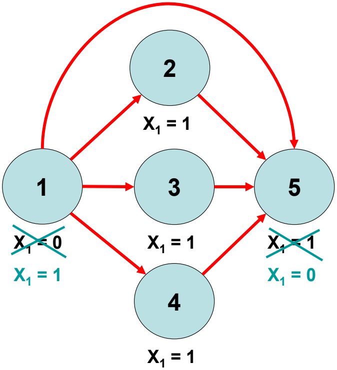

- Less can be more.

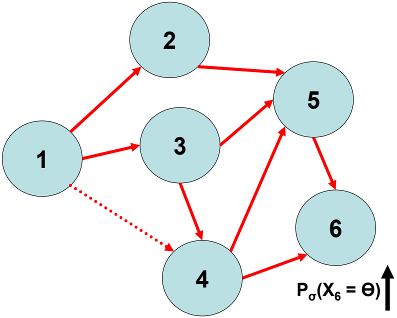

Our Approach: An Improvement Function

Consider an agent n observing only the action of agent b: $B(n) = \{b\}$.

In equilibrium, it must be the case that

$$\mathbb{P}_\sigma(x_n = \theta | B(n) = \{b\}) \geq \mathbb{P}_\sigma(x_b = \theta)$$

since agent n can copy agent b.

Can we make this inequality strict?

If yes, can we use this improvement function as a Lyapunov function?

Our Approach: An Improvement Function

Consider an agent n observing only the action of agent b: $B(n) = \{b\}$.

In equilibrium, it must be the case that

$$\mathbb{P}_\sigma(x_n = \theta | B(n) = \{b\}) \geq \mathbb{P}_\sigma(x_b = \theta)$$

since agent n can copy agent b.

Can we make this inequality strict?

If yes, can we use this improvement function as a Lyapunov function?

Our Approach: An Improvement Function

Consider an agent n observing only the action of agent b: $B(n) = \{b\}$.

In equilibrium, it must be the case that

$$\mathbb{P}_\sigma(x_n = \theta | B(n) = \{b\}) \geq \mathbb{P}_\sigma(x_b = \theta)$$

since agent n can copy agent b.

Can we make this inequality strict?

If yes, can we use this improvement function as a Lyapunov function?

Observing a Single Decision

Lemma

If $B(n) = \{b\}$, then agent n's equilibrium decision is based on 2 thresholds L_σ^b and U_σ^b:

$$x_n = \begin{cases} 0, & p_n < L_\sigma^b; \\ x_b, & p_n \in (L_\sigma^b, U_\sigma^b); \\ 1, & p_n > U_\sigma^b. \end{cases}$$

- Strict improvement if there is there is a chance $x_n \neq x_b$.
- Therefore, strict improvement if private signals are unbounded.
- The thresholds L_σ^b and U_σ^b are functions of $\mathbb{P}_\sigma(x_b = \theta | \theta = 0)$ and $\mathbb{P}_\sigma(x_b = \theta | \theta = 1)$.
- For a Lyapunov function, we need a uniform strict improvement for all values where $\mathbb{P}_\sigma(x_b = \theta | \theta = 0) + \mathbb{P}_\sigma(x_b = \theta | \theta = 1) = k$.

Observing a Single Decision

Lemma

If $B(n) = \{b\}$, then agent n's equilibrium decision is based on 2 thresholds L_σ^b and U_σ^b:

$$x_n = \begin{cases} 0, & p_n < L_\sigma^b; \\ x_b, & p_n \in (L_\sigma^b, U_\sigma^b); \\ 1, & p_n > U_\sigma^b. \end{cases}$$

- Strict improvement if there is there is a chance $x_n \neq x_b$.
- Therefore, strict improvement if private signals are unbounded.
- The thresholds L_σ^b and U_σ^b are functions of $\mathbb{P}_\sigma(x_b = \theta | \theta = 0)$ and $\mathbb{P}_\sigma(x_b = \theta | \theta = 1)$.
- For a Lyapunov function, we need a uniform strict improvement for all values where $\mathbb{P}_\sigma(x_b = \theta | \theta = 0) + \mathbb{P}_\sigma(x_b = \theta | \theta = 1) = k$.

Observing a Single Decision

Lemma

If $B(n) = \{b\}$, then agent n's equilibrium decision is based on 2 thresholds L_σ^b and U_σ^b:

$$x_n = \begin{cases} 0, & p_n < L_\sigma^b; \\ x_b, & p_n \in (L_\sigma^b, U_\sigma^b); \\ 1, & p_n > U_\sigma^b. \end{cases}$$

- Strict improvement if there is there is a chance $x_n \neq x_b$.
- Therefore, strict improvement if private signals are unbounded.
- The thresholds L_σ^b and U_σ^b are functions of $\mathbb{P}_\sigma(x_b = \theta | \theta = 0)$ and $\mathbb{P}_\sigma(x_b = \theta | \theta = 1)$.
- For a Lyapunov function, we need a uniform strict improvement for all values where $\mathbb{P}_\sigma(x_b = \theta | \theta = 0) + \mathbb{P}_\sigma(x_b = \theta | \theta = 1) = k$.

Observing a Single Decision

Lemma

If $B(n) = \{b\}$, then agent n's equilibrium decision is based on 2 thresholds L_σ^b and U_σ^b:

$$x_n = \begin{cases} 0, & p_n < L_\sigma^b; \\ x_b, & p_n \in (L_\sigma^b, U_\sigma^b); \\ 1, & p_n > U_\sigma^b. \end{cases}$$

- Strict improvement if there is there is a chance $x_n \neq x_b$.
- Therefore, strict improvement if private signals are unbounded.
- The thresholds L_σ^b and U_σ^b are functions of $\mathbb{P}_\sigma(x_b = \theta | \theta = 0)$ and $\mathbb{P}_\sigma(x_b = \theta | \theta = 1)$.
- For a Lyapunov function, we need a uniform strict improvement for all values where $\mathbb{P}_\sigma(x_b = \theta | \theta = 0) + \mathbb{P}_\sigma(x_b = \theta | \theta = 1) = k$.

Observing a Single Decision

Lemma

If $B(n) = \{b\}$, then agent n's equilibrium decision is based on 2 thresholds L_σ^b and U_σ^b:

$$x_n = \begin{cases} 0, & p_n < L_\sigma^b; \\ x_b, & p_n \in (L_\sigma^b, U_\sigma^b); \\ 1, & p_n > U_\sigma^b. \end{cases}$$

- Strict improvement if there is there is a chance $x_n \neq x_b$.
- Therefore, strict improvement if private signals are unbounded.
- The thresholds L_σ^b and U_σ^b are functions of $\mathbb{P}_\sigma(x_b = \theta | \theta = 0)$ and $\mathbb{P}_\sigma(x_b = \theta | \theta = 1)$.
- For a Lyapunov function, we need a uniform strict improvement for all values where $\mathbb{P}_\sigma(x_b = \theta | \theta = 0) + \mathbb{P}_\sigma(x_b = \theta | \theta = 1) = k$.

The Lyapunov Function

Lemma

If private beliefs are unbounded, there exists a function $\mathcal{Z}$ such that

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

where

$$\mathcal{Z}(\alpha) > \alpha \ \textit{for all} \ \ \alpha < 1.$$

- Such a $\mathcal{Z}$ does not exist if private beliefs are bounded.

Corollary

If agents are in a line, $B(n) = \{n-1\}$, asymptotic learning happens if and only if private beliefs are unbounded.

The Lyapunov Function

Lemma

If private beliefs are unbounded, there exists a function $\mathcal{Z}$ such that

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

where

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

- Such a $\mathcal{Z}$ does not exist if private beliefs are bounded.

Corollary

If agents are in a line, $B(n) = \{n-1\}$, asymptotic learning happens if and only if private beliefs are unbounded.

The Lyapunov Function

Lemma

If private beliefs are unbounded, there exists a function $\mathcal{Z}$ such that

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

where

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

- Such a $\mathcal{Z}$ does not exist if private beliefs are bounded.

Corollary

If agents are in a line, $B(n) = \{n-1\}$, asymptotic learning happens if and only if private beliefs are unbounded.

Generalizing to Deterministic Networks

Suppose $b \in B(n)$. Then,

$$\mathbb{P}_\sigma(x_n = \theta | b \in B(n)) \geq \mathcal{Z}(\mathbb{P}_\sigma(x_b = \theta))$$

since agent n has the following heuristic available:

- Ignore all decisions from $B(n) \setminus \{b\}$;
- Choose optimally based on (s_n, x_b).

With complex neighborhoods, it's impossible to characterize equilibrium strategies.

- But we can still lower bound the quality of decisions!

If "lines" exist in the network for all agents, we can prove asymptotic learning under unbounded private beliefs.

Generalizing to Deterministic Networks

Suppose $b \in B(n)$. Then,

$$\mathbb{P}_\sigma(x_n = \theta | b \in B(n)) \geq \mathcal{Z}(\mathbb{P}_\sigma(x_b = \theta))$$

since agent n has the following heuristic available:

- Ignore all decisions from $B(n) \setminus \{b\}$;
- Choose optimally based on (s_n, x_b).

With complex neighborhoods, it's impossible to characterize equilibrium strategies.

- But we can still lower bound the quality of decisions!

If "lines" exist in the network for all agents, we can prove asymptotic learning under unbounded private beliefs.

Generalizing to Deterministic Networks

Suppose $b \in B(n)$. Then,

$$\mathbb{P}_\sigma(x_n = \theta | b \in B(n)) \geq \mathcal{Z}(\mathbb{P}_\sigma(x_b = \theta))$$

since agent n has the following heuristic available:

- Ignore all decisions from $B(n) \setminus \{b\}$;
- Choose optimally based on (s_n, x_b).

With complex neighborhoods, it's impossible to characterize equilibrium strategies.

- But we can still lower bound the quality of decisions!

If "lines" exist in the network for all agents, we can prove asymptotic learning under unbounded private beliefs.

Generalizing to Deterministic Networks

Suppose $b \in B(n)$. Then,

$$\mathbb{P}_\sigma(x_n = \theta | b \in B(n)) \geq \mathcal{Z}(\mathbb{P}_\sigma(x_b = \theta))$$

since agent n has the following heuristic available:

- Ignore all decisions from $B(n) \setminus \{b\}$;
- Choose optimally based on (s_n, x_b).

With complex neighborhoods, it's impossible to characterize equilibrium strategies.

- But we can still lower bound the quality of decisions!

If "lines" exist in the network for all agents, we can prove asymptotic learning under unbounded private beliefs.

Deterministic Networks

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 cardinality of the maximal $\pi(n)$.

If $\lim_{n\to\infty} L(n) = \infty$, then all agents have long information paths.

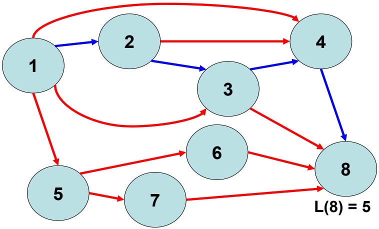

If $\liminf_{n\to\infty} L(n) < \infty$, then some agents don't have long information paths.

Expanding Observations

Definition

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

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

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.

Expanding observations $\Leftrightarrow$ no excessively influential agents.

Theorem

Assume that the network topology does not have expanding observations. Then, there exists no equilibrium with asymptotic learning.

Expanding Observations

Definition

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

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

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.

Expanding observations $\Leftrightarrow$ no excessively influential agents.

Theorem

Assume that the network topology does not have expanding observations. Then, there exists no equilibrium with asymptotic learning.

Expanding Observations

Definition

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

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

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.

Expanding observations $\Leftrightarrow$ no excessively influential agents.

Theorem

Assume that the network topology does not have expanding observations. Then, there exists no equilibrium with asymptotic learning.

Expanding Observations

Definition

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

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

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.

Expanding observations $\Leftrightarrow$ no excessively influential agents.

Theorem

Assume that the network topology does not have expanding observations. Then, there exists no equilibrium with asymptotic learning.

Learning under Unbounded Private Beliefs

Theorem

Assume that the signal structure has unbounded private beliefs and the network topology has expanding observations. Then, asymptotic learning occurs in every equilibrium.

- Under the unbounded private beliefs assumption, expanding observations characterizes asymptotic learning.

Learning under Unbounded Private Beliefs

Theorem

Assume that the signal structure has unbounded private beliefs and the network topology has expanding observations. Then, asymptotic learning occurs in every equilibrium.

- Under the unbounded private beliefs assumption, expanding observations characterizes asymptotic learning.

Learning under Unbounded Private Beliefs

Theorem

Assume that the signal structure has unbounded private beliefs and the network topology has expanding observations. Then, asymptotic learning occurs in every equilibrium.

- Under the unbounded private beliefs assumption, expanding observations characterizes asymptotic learning.

Influential vs. Excessively Influential

Consider the network topology $B(n) = \{1, n-1\}$.

Myopic models saying asymptotic learning does not happen in such networks because of the influence of agent 1.

In a Bayesian model, influential, but not excessively influential, individuals do not prevent learning.

- Intuition: the weight given to the information of influential individuals is reduced according to Bayes rule.

Learning is very robust to network structure under unbounded private beliefs.

Influential vs. Excessively Influential

Consider the network topology $B(n) = \{1, n-1\}$.

Myopic models saying asymptotic learning does not happen in such networks because of the influence of agent 1.

In a Bayesian model, influential, but not excessively influential, individuals do not prevent learning.

- Intuition: the weight given to the information of influential individuals is reduced according to Bayes rule.

Learning is very robust to network structure under unbounded private beliefs.

Influential vs. Excessively Influential

Consider the network topology $B(n) = \{1, n-1\}$.

Myopic models saying asymptotic learning does not happen in such networks because of the influence of agent 1.

In a Bayesian model, influential, but not excessively influential, individuals do not prevent learning.

- Intuition: the weight given to the information of influential individuals is reduced according to Bayes rule.

Learning is very robust to network structure under unbounded private beliefs.

Learning Despite Bounded Private Beliefs

Theorem

Assume the network topology satisfies the following three conditions:

- *expanding observations;*
- *"uninformed" agents:* $\sum_{n=1}^{\infty} \mathbb{P}(B(n) = \emptyset) = \infty$;
- *information aggregators:* $\mathbb{P}(B(n) = \{1, ..., n-1\}) \geq \epsilon \;\; \forall n$.

Then asymptotic learning occurs in all equilibria.

- Uniformed agents act based on their signals.
- Aggregators infer the state (proof via martingale).
- Information paths spread information about the true state.

Learning Despite Bounded Private Beliefs

Theorem

Assume the network topology satisfies the following three conditions:

- *expanding observations;*
- *"uninformed" agents:* $\sum_{n=1}^{\infty} \mathbb{P}(B(n) = \emptyset) = \infty$;
- *information aggregators:* $\mathbb{P}(B(n) = \{1, ..., n-1\}) \geq \epsilon \ \forall n$.

Then asymptotic learning occurs in all equilibria.

- Uniformed agents act based on their signals.
- Aggregators infer the state (proof via martingale).
- Information paths spread information about the true state.

Learning Despite Bounded Private Beliefs

Theorem

Assume the network topology satisfies the following three conditions:

- *expanding observations;*
- *"uninformed" agents:* $\sum_{n=1}^{\infty} \mathbb{P}(B(n) = \emptyset) = \infty$*;*
- *information aggregators:* $\mathbb{P}(B(n) = \{1, ..., n-1\}) \geq \epsilon \ \forall n$.

Then asymptotic learning occurs in all equilibria.

- Uniformed agents act based on their signals.
- Aggregators infer the state (proof via martingale).
- Information paths spread information about the true state.

Learning Despite Bounded Private Beliefs

Theorem

Assume the network topology satisfies the following three conditions:

- *expanding observations;*
- *"uninformed" agents:* $\sum_{n=1}^{\infty} \mathbb{P}(B(n) = \emptyset) = \infty$*;*
- *information aggregators:* $\mathbb{P}(B(n) = \{1, ..., n-1\}) \geq \epsilon \;\; \forall n.$

Then asymptotic learning occurs in all equilibria.

- Uniformed agents act based on their signals.
- Aggregators infer the state (proof via martingale).
- Information paths spread information about the true state.

Lack of Learning Without Aggregators

Theorem

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

$$\lim_{n \to \infty} \max_{b \in B(n)} b = \infty \quad a.s., \tag{1}$$

then asymptotic learning does not occur in any equilibrium.

- Implication: With bounded beliefs, learning requires aggregators.
- Caveat: Eq. (1) is stronger than expanding observations.

Lack of Learning Without Aggregators

Theorem

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

$$\lim_{n\to\infty} \max_{b\in B(n)} b = \infty \quad a.s., \tag{1}$$

then asymptotic learning does not occur in any equilibrium.

- Implication: With bounded beliefs, learning requires aggregators.
- Caveat: Eq. (1) is stronger than expanding observations.

Lack of Learning Without Aggregators

Theorem

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

$$\lim_{n \to \infty} \max_{b \in B(n)} b = \infty \quad a.s., \tag{1}$$

then asymptotic learning does not occur in any equilibrium.

- Implication: With bounded beliefs, learning requires aggregators.
- Caveat: Eq. (1) is stronger than expanding observations.

Two Follow-Up Papers (with Evan Sadler)

General networks

- We drop the independent neighborhoods assumption.
- A equilibrium failure worse than lack of asymptotic learning emerges (lack of information diffusion).
- Agents need to know who to look at for the improvement heuristic to perform well.

Diverse preferences

- Martingale-style aggregation is positively affected.
- But the improvement heuristic is negatively affected.

Two Follow-Up Papers (with Evan Sadler)

General networks

- We drop the independent neighborhoods assumption.
- A equilibrium failure worse than lack of asymptotic learning emerges (lack of information diffusion).
- Agents need to know who to look at for the improvement heuristic to perform well.

Diverse preferences

- Martingale-style aggregation is positively affected.
- But the improvement heuristic is negatively affected.

Concluding Thoughts

Prior state-of-the-art method was based on martingale convergence.

- Requires a filtration... far from ideal.

Our paper proposed an alternative approach: improvement heuristic.

- It fully characterizes the unbounded private beliefs case.
- It made some progress on the bounded beliefs case, but a general characterization is still an open problem.

Bayesian learning in networks is a rich and important problem and several amazing papers have been written after our work.

- Mossel, Sly and Tamuz (Econometrica 2015).
- Dasaratha and He (EC 2021 Best Paper Award).

Concluding Thoughts

Prior state-of-the-art method was based on martingale convergence.

- Requires a filtration... far from ideal.

Our paper proposed an alternative approach: improvement heuristic.

- It fully characterizes the unbounded private beliefs case.
- It made some progress on the bounded beliefs case, but a general characterization is still an open problem.

Bayesian learning in networks is a rich and important problem and several amazing papers have been written after our work.

- Mossel, Sly and Tamuz (Econometrica 2015).
- Dasaratha and He (EC 2021 Best Paper Award).

Concluding Thoughts

Prior state-of-the-art method was based on martingale convergence.

- Requires a filtration... far from ideal.

Our paper proposed an alternative approach: improvement heuristic.

- It fully characterizes the unbounded private beliefs case.
- It made some progress on the bounded beliefs case, but a general characterization is still an open problem.

Bayesian learning in networks is a rich and important problem and several amazing papers have been written after our work.

- Mossel, Sly and Tamuz (Econometrica 2015).
- Dasaratha and He (EC 2021 Best Paper Award).