Learning and Dynamics in Networks

Asu Ozdaglar

Department of Electrical Engineering and Computer Science
Massachusetts Institute of Technology

American Control Conference
June, 2009

Networks and Strategic Interactions

- Networks running through almost every complex environment
  - social groups, markets, Web sites, ecosystems, supply chains, conflict

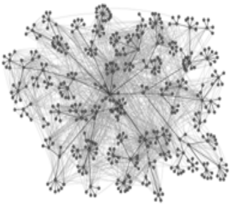

From Adamic

E-mail communication (mapped onto organizational hierarchy) at HP Labs

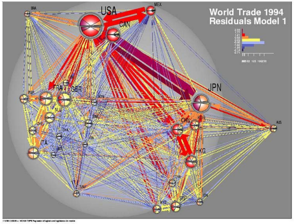

From Krempel and Plumber

A network representing international trade

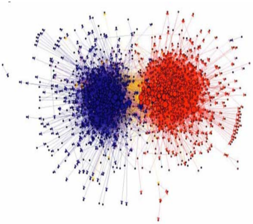

From Adamic

Network structure of political blogs prior to 2004 presidential elections

- Each individual's actions have consequences for outcomes of others
- Understanding interconnected systems requires reasoning about network structure as well as strategic behavior and feedback effects across individuals

Learning and Network Dynamics

- Belief formation crucial in social and economic networks
  - Formation of political opinions in voting
  - Learning about product quality
  - Information aggregation in financial and economic networks
  - Intentions in potential conflict situations
- How to model information dynamics in networks?
  - State of the system described by beliefs of individuals
  - Beliefs form and evolve over time based on private information, mutual information, and information exchange across individuals
- Central Question: Under what conditions (on network, interaction, and information structures) do these dynamics lead to efficient aggregation of disperse information?
- Similarity to cooperative engineering networks where there is aggregation of local information from decentralized sensors/agents
  - Same performance metrics: Accuracy and rate
  - Big new challenge: Strategic interactions

Roadmap

- Example explaining strategic interactions
- Distinction between different types of learning in different approaches
- A model of “consensus" learning
- A model of spread of misinformation and quantification of learning
- Bayesian learning over social networks (observational learning)
- Bayesian learning over social networks (communication learning)
- Learning, dynamics, and control over networks

A Motivating Example

- [Bikhchandani, Hirshleifer, Welch 92, Banerjee 92]
  - Agents arrive in town sequentially and choose to dine in an Indian or in a Chinese restaurant.
  - One restaurant is strictly better, underlying state $\theta \in \{Chinese, Indian\}$.
  - Agents have independent binary private signals.
  - Signals indicate the better option with probability $p > 1/2$.
  - Agents observe prior decisions, but not the signals of others.
- Realization: Assume $\theta = Indian$
  - Agent 1 arrives. Her signal indicates 'Chinese'.
  - She chooses to have a Chinese dinner.

A Motivating Example

- [Bikhchandani, Hirshleifer, Welch 92, Banerjee 92]
  - Agents arrive in town sequentially and choose to dine in an Indian or in a Chinese restaurant.
  - One restaurant is strictly better, underlying state $\theta \in \{Chinese, Indian\}$.
  - Agents have independent binary private signals.
  - Signals indicate the better option with probability $p > 1/2$.
  - Agents observe prior decisions, but not the signals of others.
- Realization: Assume $\theta = Indian$
  - Agent 1 arrives. Her signal indicates 'Chinese'.
  - She chooses to have a Chinese dinner.

A Motivating Example

- [Bikhchandani, Hirshleifer, Welch 92, Banerjee 92]
    - Agents arrive in town sequentially and choose to dine in an Indian or in a Chinese restaurant.
    - One restaurant is strictly better, underlying state $\theta \in \{Chinese, Indian\}$.
    - Agents have independent binary private signals.
    - Signals indicate the better option with probability $p > 1/2$.
    - Agents observe prior decisions, but not the signals of others.
- Realization: Assume $\theta = Indian$
    - Agent 1 arrives. Her signal indicates ‘Chinese’.
    - She chooses to have a Chinese dinner.

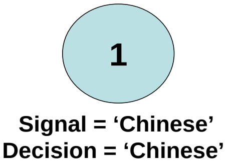

A Motivating Example

- [Bikhchandani, Hirshleifer, Welch 92, Banerjee 92]
  - Agents arrive in town sequentially and choose to dine in an Indian or in a Chinese restaurant.
  - One restaurant is strictly better, underlying state $\theta \in \{Chinese, Indian\}$.
  - Agents have independent binary private signals.
  - Signals indicate the better option with probability $p > 1/2$.
  - Agents observe prior decisions, but not the signals of others.
- Realization: Assume $\theta = Indian$
  - Agent 2 arrives. His signal indicates ‘Chinese’.
  - He also chooses to eat Chinese food.

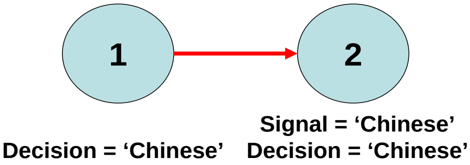

A Motivating Example

- [Bikhchandani, Hirshleifer, Welch 92, Banerjee 92]
  - Agents arrive in town sequentially and choose to dine in an Indian or in a Chinese restaurant.
  - One restaurant is strictly better, underlying state $\theta \in \{Chinese, Indian\}$.
  - Agents have independent binary private signals.
  - Signals indicate the better option with probability $p > 1/2$.
  - Agents observe prior decisions, but not the signals of others.
- Realization: Assume $\theta = Indian$
  - Agent 3 arrives. Her signal indicates 'Indian'.
  - She disregards her signal and copies the decisions of agents 1 and 2.

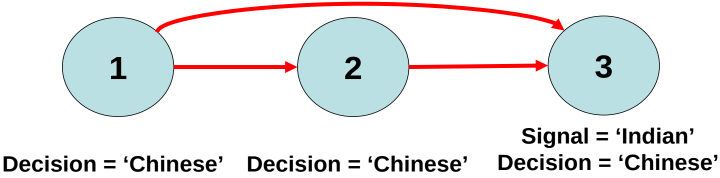

A Motivating Example

- [Bikhchandani, Hirshleifer, Welch 92, Banerjee 92]
  - Agents arrive in town sequentially and choose to dine in an Indian or in a Chinese restaurant.
  - One restaurant is strictly better, underlying state $\theta \in \{Chinese, Indian\}$.
  - Agents have independent binary private signals.
  - Signals indicate the better option with probability $p > 1/2$.
  - Agents observe prior decisions, but not the signals of others.
- Realization: Assume $\theta = Indian$
  - If the first two agents choose Chinese, everyone else selects Chinese.
  - People do not converge on the better restaurant.

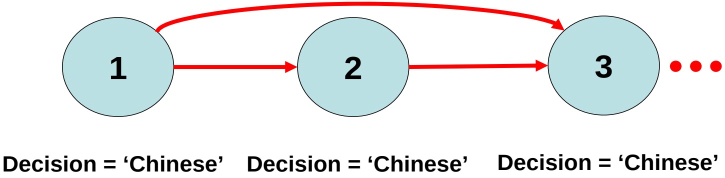

Take-away lessons

- *Game theoretic reasoning*
  - Agents $n \geq 3$ infer the signals of agents 1 and 2 from their behavior, because they conjecture the strategy used by these agents.
- *Game theoretic challenges to information aggregation*
  - Selfish behavior as a barrier to efficient information aggregation.
  - Social planner would have “collected" the signals of several agents by inducing them to act according to their signals.
  - Informational externality: Decisions I take reveal information useful for others, which does not feature in my decision making.
- *Game theoretic challenges to modeling*
  - How to analyze more realistic and complex networks with learning?

Modeling

- *How do agents act in reality?*
  - Bayesian rational learning: (as in the example)
    - Pro: Natural benchmark and often simple heuristics can replicate it
    - Con: Often complex
  - Non-Bayesian myopic learning: (rule-of-thumb)
    - Pro: Simple and often realistic
    - Con: Arbitrary rules-of-thumb, different performances from different rules, how to choose the right one?
- *What do agents observe?*
  - Observational learning: observe past actions (as in the example)
    - Most relevant for markets
  - Communication learning: communication of beliefs or estimates
    - Most relevant for friendship networks (such as Facebook)

A Benchmark Myopic Learning Model

- Beliefs updated by taking weighted averages of neighbors' beliefs [DeGroot 74], [Golub and Jackson 07]
- A finite set $\{1, \ldots, n\}$ of agents
- Interactions captured by an $n \times n$ nonnegative interaction matrix T
  - $T_{ij} > 0$ indicates the trust or weight that i puts on j
  - T is a stochastic matrix (row sum=1)
- There is an underlying state of the world $\theta \in \mathbb{R}$
- Each agent has initial belief $x_i(0)$; we assume $\theta = 1/n \sum_{i=1}^{n} x_i(0)$
- Each agent at time k updates his belief $x_i(k)$ according to

$$x_i(k+1) = \sum_{j=1}^{n} T_{ij} x_j(k)$$

  - Reasonable rule-of-thumb, but myopic
  - Update rule similar to consensus and optimization algorithms [Tsitsiklis 84], [Bertsekas, Tsitsiklis 95], [Jadbabaie, Lin, Morse 03], [Nedić, Ozdaglar 07], [Lobel, Ozdaglar 08]

Convergence and Learning

- Letting $x(k) = [x_1(k), \ldots, x_n(k)]$, the evolution of beliefs given by

$$x(k+1) = Tx(k) \qquad \text{for all } k \geq 0$$

- Under mild conditions (strong connectedness and aperiodicity of the directed graph induced by T), beliefs converge to a consensus equal to a weighted average of initial beliefs.
- Consider a sequence of networks (society) $\{T(n)\}$ and let $\bar{x}(n) \in \mathbb{R}$ be the limiting consensus belief in each $T(n)$.
- We say that asymptotic learning occurs if

$$\lim_{n\to\infty} |\bar{x}(n) - \theta| = 0$$

Influential Agents and Learning

- A set of agents B is called an influential family if the beliefs of all agents outside B affected by beliefs of B (in finitely many steps)

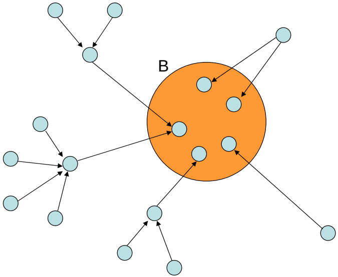

- With a sequence of (uniformly) influential agents, no asymptotic learning
  - Lack of doubly stochasticity of T
  - Interpretation: Information of influential agents overrepresented
- Distressing result since influential families (e.g., media, local leaders) common in practice

Towards a Richer Model

- Too myopic and mechanical: If communicating with same people over and over again (deterministically), some recognition that this information already been incorporated.
- No notion of misinformation or extreme views that can spread in the network.
- No analysis of what happens in terms of quantification of learning without doubly stochasticity

A Model of Misinformation

- Misinformation over networks [Acemoglu, Ozdaglar, ParandehGheibi 09]
- Finite set $\mathcal{N} = \{1, \ldots, n\}$ of agents, each with initial belief $x_i(0)$.
- Time continuous: each agent recognized according to iid Poisson processes.
- $x_i(k)$: belief of agent i after k^{th} communication.
- Conditional on being recognized, agent i meets agent j with probability p_{ij}:
  - With probability β_{ij}, the two agents agree and exchange information

$$x_i(k+1) = x_j(k+1) = (x_i(k) + x_j(k))/2.$$

  - With probability γ_{ij}, disagreement and no exchange of information.
  - With probability α_{ij}, i is influenced by j

$$x_i(k+1) = \epsilon x_i(k) + (1-\epsilon)x_j(k)$$

  for some $\epsilon > 0$ small. Agent j's belief remains unchanged.
- We say that j is a forceful agent if $\alpha_{ij} > 0$ for some i.

Evolution of Beliefs

- Letting $x(k) = [x_1(k), \ldots, x_n(k)]$, evolution of beliefs written as

$$x(k+1) = W(k)x(k),$$

where $W(k)$ is a random matrix given by

$$W(k) = \begin{cases} A_{ij} \equiv I - \frac{(e_i - e_j)(e_i - e_j)'}{2} & \text{with probability } p_{ij}\beta_{ij}/n, \\ J_{ij} \equiv I - (1-\epsilon)\, e_i(e_i - e_j)' & \text{with probability } p_{ij}\alpha_{ij}/n, \\ I & \text{with probability } p_{ij}\gamma_{ij}/n. \end{cases}$$

- The matrix $W(k)$ is a (row) stochastic matrix for all k, and is iid over all k, hence

$$E[W(k)] = \tilde{W} \qquad \text{for all } k \geq 0.$$

- We refer to the matrix $\tilde{W}$ as the mean interaction matrix.

Social Network and Influence Matrices

- Using the belief update model, we can decompose $\tilde{W}$ as:

$$\begin{aligned}\tilde{W} &= \frac{1}{n}\sum_{i,j} p_{ij}\Big[\beta_{ij}A_{ij} + \alpha_{ij}J_{ij} + \gamma_{ij}I\Big] \\ &= \frac{1}{n}\sum_{i,j} p_{ij}\Big[(1-\gamma_{ij})A_{ij} + \gamma_{ij}I\Big] + \frac{1}{n}\sum_{i,j} p_{ij}\alpha_{ij}\big[J_{ij} - A_{ij}\big] \\ &= T + D.\end{aligned}$$

- Matrix T represents the underlying social interactions: social network matrix
- Matrix D represents the influence structure in the society: influence matrix
- Decomposition of $\tilde{W}$ into a doubly stochastic and a remainder component
- Social network graph: the undirected (and weighted) graph $(\mathcal{N}, \mathcal{A})$, where $\mathcal{A} = \{\{i,j\} \mid T_{ij} > 0\}$, and the edge $\{i,j\}$ weight given by $T_{ij} = T_{ji}$
- Interaction dynamics nonsymmetric version of gossip algorithms [Boyd, Ghosh, Prabhakar, Shah 03]

Assumptions

Assumption (Connectivity and Interaction)

(i) *The graph* $(\mathcal{N}, \mathcal{E})$*, where* $\mathcal{E} = \{(i,j) \mid p_{ij} > 0\}$*, is strongly connected.*

(ii) *We have*

$$\beta_{ij} + \alpha_{ij} > 0 \qquad \text{for all } (i,j) \in \mathcal{E}.$$

- Positive probability that even forceful agents obtain information from the other agents in the society.
- Captures the idea that “no man is an island"

Convergence to Consensus

Theorem

The beliefs $\{x_i(k)\}$, $i \in \mathcal{N}$ *converge to a consensus belief, i.e., there exists a random variable* $\bar{x}$ *such that*

$$\lim_{k\to\infty} x_i(k) = \bar{x} \qquad \textit{for all } i \textit{ with probability one}.$$

Moreover, there exists a probability vector $\bar{\pi}$ *with* $\lim_{k\to\infty} \tilde{W}^k = e\bar{\pi}'$, *such that*

$$E[\bar{x}] = \sum_{i=1}^{n} \bar{\pi}_i x_i(0) = \bar{\pi}' x(0).$$

- Convergence to consensus guaranteed; but with forceful agents, consensus belief is a random variable.
- We are interested in providing an upper bound on

$$E\left[\bar{x} - \frac{1}{n}\sum_{i\in\mathcal{N}} x_i(0)\right] = \sum_{i\in\mathcal{N}} \left(\bar{\pi}_i - \frac{1}{n}\right) x_i(0).$$

- $\bar{\pi}$: consensus distribution, and $\bar{\pi}_i - \frac{1}{n}$: excess influence of agent i

Global Bounds on Consensus Distribution

Theorem

Let π denote the consensus distribution. Then,

$$\left\|\pi - \frac{1}{n}e\right\|_2 \leq \frac{1}{1-\lambda_2}\frac{\sum_{i,j} p_{ij}\alpha_{ij}}{n},$$

where λ_2 is the second largest eigenvalue of the social network matrix T.

- Proof using perturbation theory of Markov Chains
  - View $\tilde{W}$ as a perturbation of matrix T by the influence matrix D
- λ_2 related to mixing time of a Markov Chain
  - When the spectral gap $(1 - \lambda_2)$ is large, we say that the Markov Chain induced by T is fast-mixing
- In fast-mixing graphs, forceful agents will themselves be influenced by others (since $\beta_{ij} + \alpha_{ij} > 0$ for all i, j)
  - Beliefs of forceful agents moderated by the society before they spread

Location of Forceful Agents

- Previous bound does not depend on the location of the forceful agents

Example: Consider 6 agents connected with social network graph induced by T and two different misinformation scenarios:

- forceful link over a bottleneck and forceful link inside a cluster

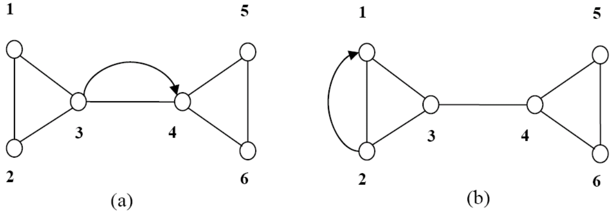

The stationary distribution for each case is given by

$$\pi_a = \frac{1}{6}(1.25, 1.25, 1.25, 0.75, 0.75, 0.75)', \quad \pi_b = \frac{1}{6}(0.82, 1.18, 1, 1, 1, 1)'.$$

Location of Forceful Agents

- Previous bound does not depend on the location of the forceful agents

Example: Consider 6 agents connected with social network graph induced by T and two different misinformation scenarios:

- forceful link over a bottleneck and forceful link inside a cluster

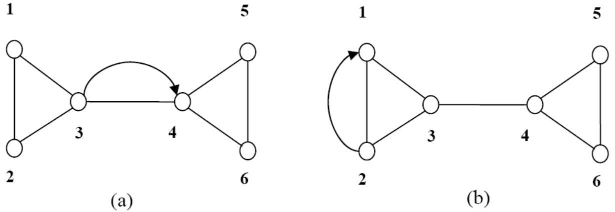

The stationary distribution for each case is given by

$$\pi_a = \frac{1}{6}(1.25, 1.25, 1.25, 0.75, 0.75, 0.75)', \quad \pi_b = \frac{1}{6}(0.82, 1.18, 1, 1, 1, 1)'.$$

Exact Characterization with Mean First Passage Times

Theorem

Let $\bar{\pi}$ denote the consensus distribution. For every agent k,

$$\bar{\pi}_k - \frac{1}{n} = \sum_{i,j} \frac{p_{ij}\alpha_{ij}}{2n^2}((1-2\epsilon)\bar{\pi}_i + \bar{\pi}_j)\big(m_{ik} - m_{jk}\big) \quad \text{for all } k,$$

where m_{ij} is the mean first passage time from state i to state j of a Markov chain $(X_t, t = 0, 1, 2, \ldots)$ with transition matrix T, i.e.,

$$m_{ij} = \mathbb{E}[T_j \mid X_0 = i],$$

where $T_i = \inf\{t \geq 0 \mid X_t = i\}$.

- Excess influence of each agent depends on the relative distance of that agent to the forceful and the influenced agent
  - Explains the insensitivity of the agents in the right cluster in the previous example.

Information Bottlenecks – Relative Min-Cuts

Theorem

Let $\bar{\pi}$ be the consensus distribution. For all k, we have

$$\left|\bar{\pi}_k - \frac{1}{n}\right| \leq \sum_{i,j} \frac{3p_{ij}\alpha_{ij}}{2n}\left(\frac{\log n}{\rho_{ij}}\right),$$

where ρ_{ij} is the minimum normalized relative cut value between i and j of the Markov chain induced by the social network matrix T, i.e.,

$$\rho_{ij} = \inf_{S \subset \mathcal{N}} \left\{ \frac{\sum_{h\in S}\sum_{l\in S^c} T_{hl}}{|S|} \mid i \in S, j \notin S \right\}.$$

- Proof relies on bounding the mean commute time using Max flow-Min cut Theorem.

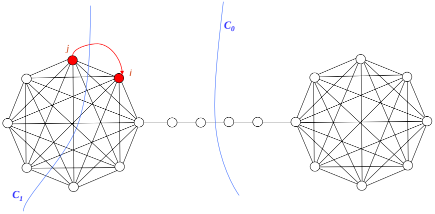

Shortcomings of This Model

- Even though non-deterministic interaction pattern, information still replicated.
- Where do these rules come from?
- Line of Attack: Develop Bayesian models
- Imagine the Chinese-Indian restaurant example, but with two realistic features:
    - Social network structure (every agent does not observe the full past)
    - Heterogeneity of preferences

Bayesian Learning over Networks– without heterogeneity

- Learning over general networks [Acemoglu, Dahleh, Lobel, Ozdaglar 08]
- 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 $\mathbb{F}_\theta$ (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 $\mathbb{Q}_n$ (independently for all n) (network topology)
  - The sequence $\{\mathbb{Q}_n\}_{n\in\mathbb{N}}$ is common knowledge.
- Asymptotic Learning: Under what conditions does $\lim_{n\to\infty} \mathbb{P}(x_n = \theta) = 1$?

An Example of a Social Network

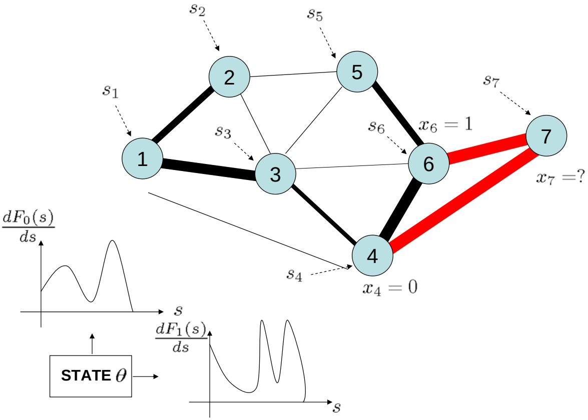

Perfect Bayesian Equilibria

- 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 all n

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

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

Definition

We say that asymptotic learning occurs in equilibrium σ if x_n converges to θ in probability,

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

Some Difficulties of Bayesian Learning

- No following the crowds

Some Difficulties of Bayesian Learning

- No following the crowds

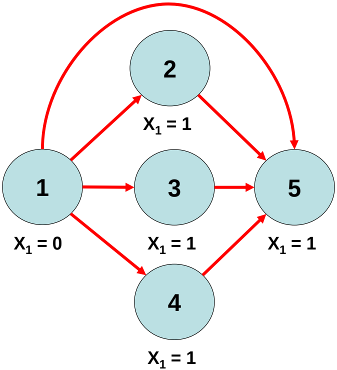

Some Difficulties of Bayesian Learning

- No following the crowds

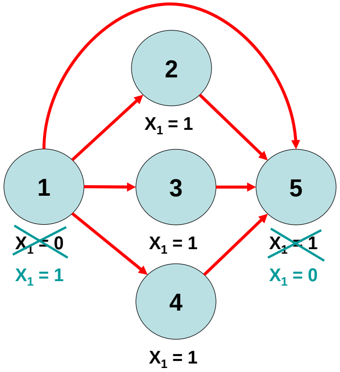

Some Difficulties of Bayesian Learning

- No following the crowds

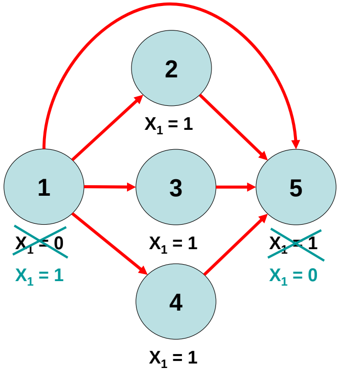

- Less can be more

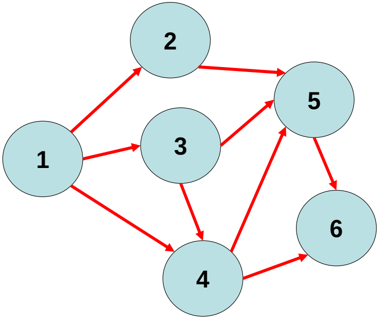

Some Difficulties of Bayesian Learning

- No following the crowds

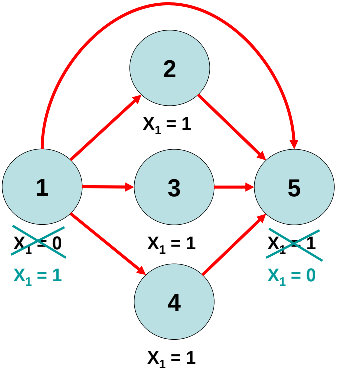

- Less can be more.

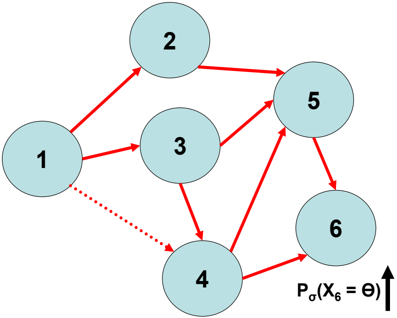

Equilibrium Decision Rule

Lemma

The decision of agent n, $x_n = \sigma(\mathcal{I}_n)$, 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.

- Implication: 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

- 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 unbounded private beliefs if

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

- If the private beliefs are unbounded, then there exist agents with beliefs arbitrarily strong in both directions.
  - Gaussian signals yield unbounded beliefs; discrete signals yield bounded beliefs.

Properties of Network Topology

Definition

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

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

- Excessive influence:
  - 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.
  - Expanding observations $\Leftrightarrow$ no excessively influential agents.

Learning Theorem – with Unbounded Beliefs

Theorem

Assume unbounded private beliefs and expanding observations. Then, asymptotic learning occurs in every equilibrium $\sigma \in \Sigma^*$.

- Implication: Influential, but not excessively influential, individuals do not prevent learning.
  - This contrasts with results in models of myopic learning.
  - Intuition: The weight given to the information of influential individuals is adjusted in Bayesian updating.

Proof of Theorem – A 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

Proposition

Let $B(n) = \{b\}$ for some agent n. There exists L_b^σ and U_b^σ such that agent n's decision x_n in $\sigma \in \Sigma^*$ satisfies

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

- Let $\mathbb{G}_j(r) = \mathbb{P}(p \leq r \mid \theta = j)$ be the conditional distribution of the private belief with $\underline{\beta}$ and $\overline{\beta}$ denoting the lower and upper support

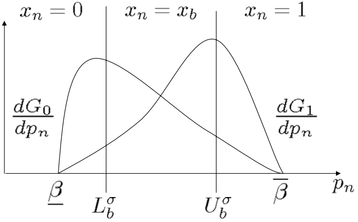

Strong Improvement Principle

- Agent n has the option of copying the action of his neighbor b:

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

- 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

- With multiple agents, learning no worse than observing just one of them.
- Equilibrium strategy is better than the following heuristic:
  - Discard all decisions except the one from the most informed neighbor.
  - Use equilibrium decision rule for this new information set.

Proposition (Generalized Strong Improvement Principle)

For any $n \in \mathbb{N}$, *any set* $\mathfrak{B} \subseteq \{1, ..., n-1\}$ *and any* $\sigma \in \Sigma^*$,

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

Proof of Theorem:

- 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 with Bounded Beliefs

Theorem

Assume that the signal structure has bounded private beliefs. If 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 in any equilibrium $\sigma \in \Sigma^*$.

- Implication: With bounded beliefs, no learning from observing neighbors or sampling the past.

Learning with Bounded Beliefs

- Theorem: There exist random network topologies for which learning occurs in all equilibria.

Example

Let the network topology be

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

Asymptotic learning occurs in all equilibria $\sigma \in \Sigma^*$ for any signal structure $(\mathbb{F}_0, \mathbb{F}_1)$.

- Result contrasts with prior literature.
- Proof Idea:
  - Social beliefs form a martingale.
  - Martingale convergence implies almost sure convergence of actions.
  - The rate of contrary actions gives away the state.

Diversity and Learning

- So far, all agents have the same preferences.
  - They all prefer to take action $= \theta$, and with the same intensity.
- In realistic situations, not only diversity of opinions, but also diversity of preferences.
- How does diversity of preferences affect social learning?
- Naive conjecture: diversity will introduce additional noise and make learning harder or impossible.
- Our Result: in the line topology, diversity always facilitates learning.

Model with Heterogeneous Preferences

- Assume $B(n) = \{1, ..., n-1\}$ [Acemoglu, Dahleh, Lobel, Ozdaglar 09]
- Let agent n have private preference t_n independently drawn from some $\mathbb{H}$.
- The payoff of agent n given by:

$$u_n(x_n, t_n, \theta) = \begin{cases} I(\theta = 1) + 1 - t_n & \text{if } x_n = 1 \\ I(\theta = 0) + t_n & \text{if } x_n = 0 \end{cases}$$

- Theorem: With unbounded preferences, i.e., $[0, 1] \subseteq supp(\mathbb{H})$, asymptotic learning occurs in all equilibria $\sigma \in \Sigma^*$ for any signal structure $(\mathbb{F}_0, \mathbb{F}_1)$.
  - Heterogeneity pulls learning in opposite directions:
    - Actions of others are less informative (direct effect)
    - Each agent uses more of his own signal in making decisions and, therefore, there is more information in the history of past actions (indirect effect)
  - Indirect effect dominates the direct effect! (relies on martingale convergence for the social belief sequence)

Extensions

- Correlated neighborhoods
  - Expanding observations not a sufficient condition
  - Encompasses random graph models
- Diversity of preferences with general network topologies
- Rate of learning
  - Presented by Ilan Lobel on Thursday
- Previous model based on observational learning
- In practice, belief formation also depends on communication with friends, neighbors, and media sources
  - What was captured by the myopic models
- Next, a learning model with communication and observation.
  - Much more of effect of network structure

A Model of Communication Learning

- Effect of communication on learning [Acemoglu, Bimpikis, Ozdaglar 09]
- Two possible states of the world, $\theta \in \{0, 1\}$
- A set $\mathcal{N} = \{1, \ldots, n\}$ of agents and a friendship network given

Stage 1: Network Formation

- Additional link formation is costly, c_{ij}^n : cost incurred by i to link with j
- Induces the communication network $G^n = (\mathcal{N}, \mathcal{E}^n)$

Stage 2: Information Exchange (over the communication network G^n)

- Each agent receives an iid private signal, $s_i \sim \mathbb{F}_\theta$
- Agents receive all information acquired by their direct neighbors
- At each time period t they can choose:
  (1) irreversible action 0 (2) irreversible action 1 (3) wait

Stage 1: Forming the communication network

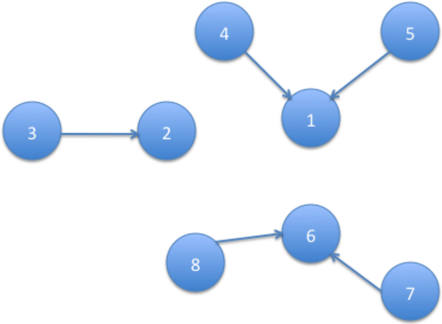

Stage 1: Forming the communication network

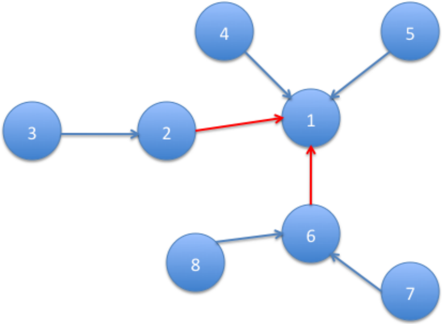

Stage 2: Information Exchange

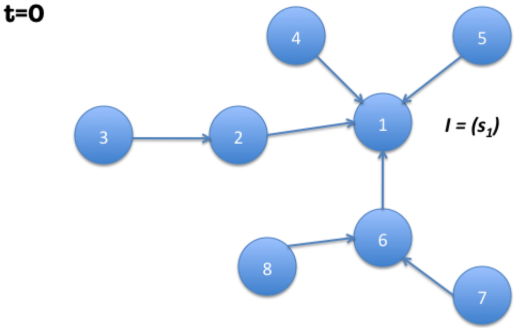

Stage 2: Information Exchange

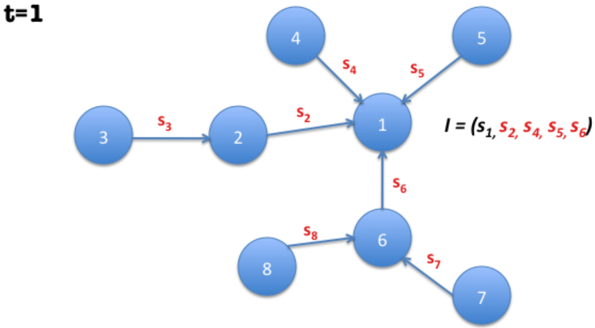

Stage 2: Information Exchange

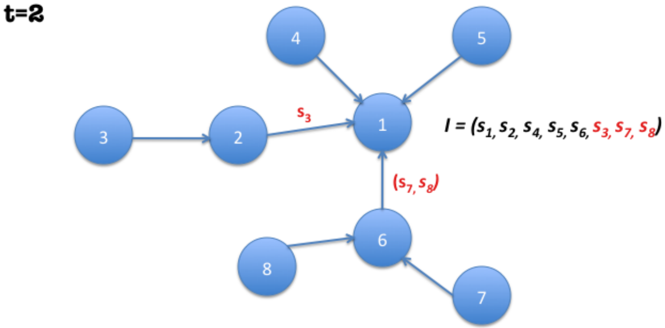

Model

- This talk: Focus on stage 2
- Agent i's payoff is given by

$$u_i(\mathbf{x_i^n}, \theta) = \begin{cases} \delta^\tau \pi & \text{if } x_{i,\tau}^n = \theta \text{ and } x_{i,t}^n = \text{ ``wait" for } t < \tau \\ 0 & \text{otherwise} \end{cases}$$

  - $\mathbf{x_i^n} = [x_{i,t}^n]_{t \geq 0}$: sequence of agent i's decisions, $x_{i,t}^n \in \{0, 1, \text{``}wait\text{''}\}$
  - δ: discount factor ($\delta < 1$)
  - τ: time when action is taken (agent collects information up to τ)
  - π: payoff - normalized to 1
- Assumption: Communication between agents is not strategic
  - Agents cannot manipulate the information they send to neighbors
  - Results extend to ϵ-equilibrium with strategic communication!
- Let $B_{i,t}^n = \{j \neq i \mid \exists$ a directed path from j to i with at most t links in $G^n\}$
  - All agents that are at most t links away from i in G^n
- Agent i's information set at time t:

$$I_{i,t}^n = \{s_i, G^n, s_j \text{ for all } j \in B_{i,t}^n\}$$

Equilibrium and Learning

- Given a sequence of communication networks $\{G^n\}$ (society):
  - Strategy for agent i at time t is $\sigma^n_{i,t} : \mathcal{I}^n_{i,t} \rightarrow \{\text{“wait"}, 0, 1\}$

Definition

A strategy profile $\sigma^{n,*}$ is a Perfect-Bayesian Equilibrium if for all i and t,

$$\sigma^{n,*}_{i,t} \in \arg \max_{y \in \{\text{“wait”},0,1\}} E_{(y,\sigma^{n,*}_{-i,t})} \left(u_i(\mathbf{x^n_i}, \theta) | I^n_{i,t} \right) .$$

- Let

$$M^n_{i,t} = \begin{cases} 1 & \text{if } x_{i,\tau} = \theta \text{ for some } \tau \leq t \\ 0 & \text{otherwise} \end{cases}$$

Definition

We say that asymptotic learning occurs in society $\{G^n\}$ if for every $\epsilon > 0$

$$\lim_{n\to\infty} \lim_{t\to\infty} P_{\sigma^{n,*}} \left(\left[\frac{1}{n} \sum_{i=1}^n \left(1 - M^n_{i,t}\right) \right] > \epsilon \right) = 0$$

Agent Decision Rule

Lemma

Let $\sigma^{n,*}$ be an equilibrium and $I_{i,t}^{n}$ be an information set of agent i at time t. Then, the decision of agent i, $x_{i,t}^{n} = \sigma_{i,t}^{n,*}(I_{i,t}^{n})$ satisfies

$$x_{i,t}^{n} = \begin{cases} 0, & \text{if } \log L(s_i) + \sum_{j \in B_{i,t}^{n}} \log L(s_j) \leq -\log A_{i,t}^{n,*}, \\ 1, & \text{if } \log L(s_i) + \sum_{j \in B_{i,t}^{n}} \log L(s_j) \geq \log A_{i,t}^{n,*}, \\ \text{"wait"}, & \text{otherwise}, \end{cases}$$

where $L(s_i) = \frac{dP_\sigma(s_i \mid \theta=1)}{dP_\sigma(s_i \mid \theta=0)}$ is the likelihood ratio of signal s_i, and $A_{i,t}^{n,*} = \frac{p_{i,t}^{n,*}}{1-p_{i,t}^{n,*}}$, is a time-dependent parameter.

- $p_{i,t}^{n,*}$: belief threshold that depends on time and graph structure
- For this talk:
  - Focus on binary private signals $s_i \in \{0,1\}$
  - Assume $L(1) = \frac{\beta}{1-\beta}$ and $L(0) = \frac{1-\beta}{\beta}$ for some $\beta > 1/2$.

Minimum Observation Radius

Lemma

The decision of agent i, $x_{i,t}^n = \sigma_{i,t}^{n,*}(I_{i,t}^n)$ *satisfies*

$$x_{i,t}^n(I_{i,t}^n) = \begin{cases} 0, & \text{if } k_{i,0}^t - k_{i,1}^t \geq \log A_{i,t}^{n,*} \cdot \left(\log \frac{\beta}{1-\beta}\right)^{-1}, \\ 1, & \text{if } k_{i,1}^t - k_{i,0}^t \geq \log A_{i,t}^{n,*} \cdot \left(\log \frac{\beta}{1-\beta}\right)^{-1}, \\ \text{“wait"}, & \text{otherwise,} \end{cases}$$

where $k_{i,1}^t$ $(k_{i,0}^t)$ *denotes the number of 1's (0's) agent i has observed up to time t.*

Definition

We define the minimum observation radius of agent i, denoted by d_i^n, as

$$d_i^n = \arg\min_t \left\{ \left|B_{i,t}^n\right| \,\middle|\, \left|B_{i,t}^n\right| \geq \log A_{i,t}^{n,*} \cdot \left(\log \frac{\beta}{1-\beta}\right)^{-1} \right\}.$$

- Agent i receives at least $|B_{i,d_i^n}^n|$ signals before she takes an irreversible action
- $B_{i,d_i^n}^n$: Minimum observation neighborhood of agent i

A Learning Theorem

Definition

For any integer $k > 0$, we define the *k*-radius set, denoted by V_k^n, as

$$V_k^n = \{j \in \mathcal{N} \mid \left|B^n_{j,d_j^n}\right| \leq k\}$$

- Set of agents with "finite minimum observation neighborhood"
- Note that any agent i in the k-radius set has positive probability of taking the wrong action.

Theorem

Asymptotic learning occurs in society $\{G^n\}$ if and only if

$$\lim_{k\to\infty} \lim_{n\to\infty} \frac{\left|V_k^n\right|}{n} = 0$$

- A "large" number of agents with finite obs. neighborhoods precludes learning.

Information Hubs and Random Graph Models

- A node i is an information hub if it has an infinite number of direct neighbors,
$$\lim_{n \to \infty} |B_1^n(i)| = \infty$$
- Asymptotic learning occurs if and only if for all but a negligible fraction of agents, the shortest path to a hub is shorter than minimum observation radius.

Proposition

Asymptotic Learning occurs for

(a) *Complete and Star Graphs*

(b) *Power Law Graphs with exponent* $\gamma \leq 2$ *(with high probability)*

- *Intuition: The average degree is infinite - there exist many hubs.*

Asymptotic Learning fails for

(a) *Bounded Degree Graphs, e.g. expanders*

(b) *Preferential Attachment Graphs (with high probability)*

- *Intuition: Edges form with probability proportional to degree, but there exist many low degree nodes.*

Networks, Dynamics, and Learning

- This talk: A review of the emerging field of theoretical models of social learning in networks
  - Modeling strategic interactions between individuals
  - Characterizing effects of network structure
  - Game theory and stochastic dynamic analysis
- Literature so far focuses on modeling and understanding dynamics
- Next step: Control over networks
  - How can misinformation be contained?
  - Which networks are robust and resilient?
  - How can information exchange be facilitated?
- Mechanism Design approach (design of game forms) meets control theory over networks
- Large area of research at the intersection of Networks, Control Theory, Economics, Computer Science, Operations Research, Sociology,. . .