Social and Economic Networks

Learning

Xiang Sun

2019 Fall

Outline

1. Observational learning
   - Herding
   - Herding and sequential decision-making
   - Learning from neighbors
2. DeGroot model
   - Convergence
   - Consensus in beliefs
   - Wise learning
   - Social influence

Reference

- Jackson, Chapter 8
- Easley and Kleinberg, Chapter 16
- MIT Open Course Networks, Fall 2009, Lecture 22–23
- MIT Open Course Networks, Spring 2018, Lecture 22–23

Learning

- People are influenced by others:
  - the opinions they hold,
  - the products they buy,
  - the political positions they support,
  - the activities they pursue,
  - the technologies they use,
  - and many other things.
- It may be rational for an individual to imitate the choices of others even if the individual's own information suggests an alternative choice.
  - Choosing a restaurant in an unfamiliar town.

Learning (Cont.)

Social networks play a central role in the sharing of information and the formation of opinions.

- Providing information about scientific research and results.
- Advising friends on which movies to see.
- Relaying information about the abilities and profit of a potential new employee in a firm.
- Debating the relative merits of politicians.

Learning (Cont.)

Given the role of social networks in the formation of opinions and beliefs, and the subsequent shaping of behaviors, it is critical that we have a thorough understanding of this how the structure of social networks affects learning:

- Whether individuals in a society come to hold a common belief or remain divided in opinions.
- Which individuals have the most influence over the beliefs in a society.
- How quickly individuals learn.
- Whether initially diverse information scattered throughout the society can be aggregated in an accurate manner.

Two kinds of learning

- (Bayesian) Observational learning
  - Individuals observe actions and results experienced by their neighbors and the information in a sophisticated manner (Bayesian update).
  - It provides conditions under which individuals come to act similarly over time.
- Communication learning
  - Individuals exchange information with their neighbors over time and then update by taking some weighted average of what they hear.
  - Non-Bayesian, myopic, rule of thumb.
  - Tractable, and allows us to incorporate rich network structures.
  - DeGroot model (1974).

1. Observational learning
   - Herding
   - Herding and sequential decision-making
   - Learning from neighbors

2. DeGroot model
   - Convergence
   - Consensus in beliefs
   - Wise learning
   - Social influence

Example

- Milgram, Bickman, and Berkowitz in the 1960s.
- There are several groups of people ranging in size from just one person to as many as fifteen people.
- In each round, each group of people stand on a street corner and stare up into the sky.
- They then observed how many passersby stopped and also looked up at the sky.
    - With only one person looking up, very few passersby stopped.
    - If five people were staring up into the sky, then more passersby stopped, but most still ignored them.
    - Finally, with fifteen people looking up, they found that 45% of passersby stopped and also stared up into the sky.

Example (Cont.)

- One interpretation: A social force for conformity grows stronger as the group conforming to the activity becomes larger.
- Another possible explanation: A possible mechanism gives rise to the conformity observed in this kind of situation.
  - Initially the passersby saw no reason to look up (they had no private or public information that suggested it was necessary).
  - But with more and more people looking up, future passersby may have rationally decided that there was good reason to also look up (since perhaps those looking up knew something that the passersby didn't know).

$\Rightarrow$ Information cascades may be at least part of the explanation for many types of imitation in social settings.

Herding

- Herding or information cascade:
  - people make decisions sequentially,
  - later people observe the actions of earlier people and infer something about what the earlier people know.
- Individuals in a cascade are imitating the behavior of others, but it is not mindless imitation. Rather, it is the result of drawing rational inferences from limited information.

Simple model

- Consider a group of people (numbered 1, 2, 3, ...) who will sequentially make decisions—that is, individual 1 will decide first, then individual 2 will decide, and so on.
- Each individual make a decision: accepting or rejecting some option:
    - whether to adopt a new technology, wear a new fashion, eat in a new restaurant, commit a crime, vote for a particular political candidate, or choose one route to a common destination rather than an alternative route.

Simple model: State

- At the start of everything, before any individual has made a decision, we assume that the world is randomly placed into one of two possible states:
    - it is either placed in a state in which the option is a good idea,
    - or a state in which the option is actually a bad idea.
- G represents the state where the option is a good idea.
- B represents the state where the option is a bad idea.
- Prior probability: $\text{Prob}(G) = p$ and $\text{Prob}(B) = 1 - p$.

Simple model: State (Cont.)

- The state of the world is determined by some initial random event that the individuals can't observe.
- The indivoduals will try to use what they observe to make inferences about this state.
- Example:
    - the world is either in a state where the new restaurant is good or a state where it is bad;
    - the individuals in the model know that it was randomly placed in one of these two states, and they're trying to figure out which.

Simple model: Payoff

- Each individual receives a payoff based on her decision to accept or reject the option.
- If the individual chooses to reject the option, she receives a payoff of 0.
- The payoff for accepting depends on whether the option is a good idea or a bad idea:
  - If the option is a good idea, then the payoff is $v_g > 0$.
  - If the option is a bad idea, then the payoff is $v_b < 0$.
- We will also assume that the expected payoff from accepting in the absence of other information is equal to 0; in other words,

$$v_g p + v_b(1 - p) = 0.$$

* Before an individual gets any additional information, the expected payoff from accepting is the same as the payoff from rejecting.

Simple model: Signal

- Before any decisions are made, each individual gets a private signal that provides information about whether accepting is a good idea or a bad idea.
  - a review of the restaurants.
- There are two possible signals:
  - a high signal (denoted H), suggesting that accepting is a good idea;
  - a low signal (denoted L), suggesting that accepting is a bad idea.
- If accepting is in fact a good idea, then high signals are more frequent than low signals: $\text{Prob}(H \mid G) = q > \frac{1}{2}$.
- If accepting the option is a bad idea, then low signals are more frequent: $\text{Prob}(L \mid B) = q > \frac{1}{2}$.

Individual decisions

- Suppose that a person gets a high signal H.
- This shifts their expected payoff from $v_g \operatorname{Prob}(G) + v_b \operatorname{Prob}(B) = 0$ to
$$v_g \operatorname{Prob}(G \mid H) + v_b \operatorname{Prob}(B \mid H).$$
- Bayes' rule implies:

$$\begin{aligned}\operatorname{Prob}(G \mid H) &= \frac{\operatorname{Prob}(G) \cdot \operatorname{Prob}(H \mid G)}{\operatorname{Prob}(H)} \\ &= \frac{\operatorname{Prob}(G) \cdot \operatorname{Prob}(H \mid G)}{\operatorname{Prob}(G) \cdot \operatorname{Prob}(H \mid G) + \operatorname{Prob}(B) \cdot \operatorname{Prob}(H \mid B)} \\ &= \frac{pq}{pq + (1-p)(1-q)} > \frac{pq}{pq + (1-p)q} = p.\end{aligned}$$

Individual decisions (Cont.)

- As a result, the expected payoff shifts from 0 to a positive number, and so they should accept the option.
- Interpretation: A high signal is more likely to occur if the option is good than if it is bad, so if an individual observes a high signal they raise their estimate of the probability that the option is good.
- A completely analogous calculation shows that if the individual receives a low signal, they should reject the option.

Multiple signals

When the individuals get a sequence S of independently generated signals consisting of a high signals and b low signals, interleaved in some fashion, how do they act?

- The posterior probability $\text{Prob}(G \mid S)$ is greater than the prior $\text{Prob}(G)$ when $a > b$;
- The posterior $\text{Prob}(G \mid S)$ is less than the prior $\text{Prob}(G)$ when $a < b$;
- The two probabilities $\text{Prob}(G \mid S)$ and $\text{Prob}(G)$ are equal when $a = b$.

Multiple signals (Cont.)

- Bayes' rule imples

$$\text{Prob}(G \mid S) = \frac{\text{Prob}(G) \cdot \text{Prob}(S \mid G)}{\text{Prob}(S)},$$

where S is a sequence with a high signals and b low signals.

- The signals are generated independently.

$\Rightarrow$ $\text{Prob}(S \mid G) = q^a(1-q)^b$.

-

$$\begin{aligned}\text{Prob}(S) &= \text{Prob}(G)\,\text{Prob}(S \mid G) + \text{Prob}(B)\,\text{Prob}(S \mid B) \\ &= pq^a(1-q)^b + (1-p)(1-q)^a q^b.\end{aligned}$$

Multiple signals (Cont.)

- If $a > b$, then $(1-p)(1-q)^a q^b < (1-p)q^a(1-q)^b$.

$\Rightarrow$ $\text{Prob}(S) < pq^a(1-q)^b + (1-p)q^a(1-q)^b = q^a(1-q)^b$.

$\Rightarrow$ $\text{Prob}(G \mid S) > \text{Prob}(G)$.

Sequential decision-making

- When a given person decides whether to accept or reject the option, they have access to their own private signal and also the accept/reject decisions of all earlier people.
- However, they do not see the actual private signals of any of these earlier people.
- Person 1 will follow his own private signal.
- Person 2 will know that person 1's decision reveals their private signal, and so it's as though person 2 gets two signals.
    - If these signals are the same, person 2's decision is easy.
    - If they are different, then as we saw before, person 2 will be indifferent between accepting and rejecting.
    - Here we will assume she follows her own private signal. Thus, either way, person 2 is following her own signal.

Sequential decision-making (Cont.)

- Person 3 knows that person 1 and person 2 both acted on their private signals, so it is as though person 3 has received three independent signals (the two he infers, and his own).
- From the previous argument, we know that person 3 will follow the majority signal (high or low) in choosing whether to accept or reject.
- If person 1 and person 2 made opposite decisions (i.e. they received opposite signals), then
  - person 3 will use his own signal as the tie-breaker;
  - future people will know that person 3's decision was based on his own signal, and so they can use this information in their own decisions.

Herding

- On the other hand, if person 1 and person 2 made the same decision (i.e. had the same signal), then
  - person 3 will follow this regardless of what his own signal says;
  - future people will know that person 3's decision conveys no information about his signal, and future people will all be in the same position as person 3.
  - In this case, a cascade has begun. That is, we are in a situation where no individual's decision can be influenced by his own signal. No matter what they see, every individual from person 3 on will make the same decision that 1 and 2 made.

Sequential decision-making (Cont.)

- Let's now consider how this process unfolds through future people (person N) beyond person 3.
- Suppose that person N knows that everyone before her has followed their own signal—that is, suppose the accept/reject decisions of these earlier people exactly coincide with whether they received a high or low signal, and person N knows this.
- If the number of acceptances among the people before N is equal to the number of rejections, then N's signal will be the tie-breaker, and so N will follow her own signal.

Sequential decision-making (Cont.)

- If the number of acceptances among the people before N differs from the number of rejections by one, then
  - either N's private signal will make her indifferent,
  - or it will reinforce the majority signal.
- Either way, N will follow her private signal (since we assume a person follows their own signal in the case of indifference).

Herding (Cont.)

- If the number of acceptances among the people before N differs from the number of rejections by two or more, then however N's private signal turns out, it won't outweigh this earlier majority.
- As a result, N will follow the earlier majority and ignore her own signal.
- In this case, the people numbered $N+1$, $N+2$, and onward will know that person N ignored her own signal (whereas we've assumed that all earlier people were known to have followed their private signals).
- So they will each be in exactly the same position as N.
- That is, each of them too will ignore their own signals and follow the majority, and hence a cascade has begun.

Herding (Cont.)

- As long as the number of acceptances differs from the number of rejections by at most one, each person in sequence is simply following their own private signal in deciding what to do.
- But once the number of acceptances differs from the number of rejections by two or more, a cascade takes over, and everyone simply follows the majority decision forever.

Herding (Cont.)

- It is very hard for this difference to remain in such a narrow interval (between -1 and $+1$) forever.
- Prob(herding) $\geq$ Prob(three people in a row get the same signal).
- We divide the first N people into blocks of three consecutive people each.
- The people in any one block will receive identical signals with $q^3 + (1-q)^3$.
- The probability that none of these blocks consists of identical signals is therefore

$$[1 - q^3 - (1-q)^3]^{\frac{N}{3}}.$$

- As N goes to infinity this quantity goes to 0.

Bayesian learning

- Individuals observe actions and results experienced by their neighbors and the information in a sophisticated manner.
- Conclusion: If agents can observe each other's actions and outcomes over time, and all agents have the same preferences and face the same form of uncertainty, then they end up with similar payoffs over time.
- Idea: an agent who is doing significantly worse than a neighbor must come to realize this over time, and will eventually change actions and come to do as well as the neighbor.
- This then implies that all connected agents must end up with the same limiting payoffs.

Bala-Goyal model

- n players in an undirected connected network g.
- Choose action A or B in each period $t \in \{1, 2, \ldots\}$.
- In each period agent gets a payoff based on choice:
  - action A results in a payoff of 1.
  - action B results in a payoff of 2 with probability p and 0 with probability $1 - p$.
- p is unknown taking on finite set of values.

Bala-Goyal model (Cont.)

- Players also observe neighbors' choices.
- Each player maximizes discounted stream of payoffs

$$\mathbf{E}\left[\sum_{t=1}^{\infty} \delta^t \cdot \pi_{it}\right],$$

where $\delta \in (0, 1)$ is a discount parameter and π_{it} is the payoff that i receives at time t.

Bala-Goyal model: Challenges

- Seeing that a neighbor chooses an action B might indicate that the individual's neighbors have had good outcomes from B in the past.
- Beyond simply seeing actions and outcomes, an individual can make inferences about outcomes of indirect neighbors by observing the action choices of neighbors.

Bala-Goyal model: Result

Proposition

If p is not exactly $\frac{1}{2}$, then with probability 1 there is a time such that all agents lay just one action (and all play the same action) from that time onward.

Proof

- Suppose contrary.
- Some agent plays B infinitely often.
- That agent will converge to true belief p by the law of large numbers.
- In order for agent to play B infinitely often, it must be that $p > \frac{1}{2}$, otherwise agent would stop playing B.

Proof (Cont.)

- With probability 1, all agents who see B played infinitely often converge to a belief that B pays 2 with probability $p > \frac{1}{2}$.
- Neighbors of agent must play B, after some time, and so forth.
- All agents must play B from some time on.

Play the right action?

- The fact that all agents end up choosing the same action does not imply that they end up with the same limiting beliefs, nor does it imply that they end up choosing the "right" action.
- If B is the right action then play the right action if converge to it, but might not.

* Each player starts with a low belief.

- If A is the right action, then must converge to right action.

Conclusions

- Consensus action chosen.
- Not necessarily consensus belief.
- Speed of convergence?

Limitations

- Homogeneity of actions and payoffs across players.
- What if heterogeneity?
- Repeated actions over time.
- Stationarity.
- Networks are not playing role here.

DeGroot model

- Repeated communication.
- Information comes only once.
- See how information disseminates.
- Who has influence, convergence speed, network structure impact …

Bounded rationality

- Repeatedly average beliefs of self with neighbors.
- Non-Bayesian if weights do not adjust over time.
- Can under-weight neighbors (just as in experiments).

DeGroot model (Cont.)

- Individuals $\{1, 2, \ldots, n\}$.
- Individuals in a society start with initial opinions on a subject.
- Let these be represented by an n-dimensional vector of probabilities, $p(0) = \big(p_1(0), p_2(0), \cdots, p_n(0)\big)$.
- Each $p_i(0)$ lies in $[0, 1]$, and might be thought of as the probability that a given statement is true, or the quality of a given product, or the likelihood that the individual might engage in a given activity, etc.

DeGroot model: Updating

- The interaction patterns are captured through a possibly weighted and directed $n \times n$ nonnegative matrix T (social influence matrix).
- The interpretation of T_{ij} is that it represents the weight or trust that agent i places on the current belief of agent j in forming his or her belief for the next period.

* T_{ij}: agent j's impact on agent i.

- T: a (row) stochastic matrix, so that its entries across each row sum to one.
- Updating

$$p_i(t) = \sum_j T_{ij} \cdot p_j(t-1).$$

DeGroot model: Updating (Cont.)

- Updating
$$p_i(t) = \sum_j T_{ij} \cdot p_j(t-1).$$
-
$$p_i(t) = \sum_j T_{ij} \cdot p_j(t-1) = \big(T \cdot p(t-1)\big)_i.$$
-
$$p(t) = \begin{pmatrix} p_1(t) \\ p_2(t) \\ \vdots \\ p_n(t) \end{pmatrix} = \begin{pmatrix} \big(T \cdot p(t-1)\big)_1 \\ \big(T \cdot p(t-1)\big)_2 \\ \vdots \\ \big(T \cdot p(t-1)\big)_n \end{pmatrix} = T \cdot p(t-1) = T^t \cdot p(0).$$

DeGroot model: Illustration 1

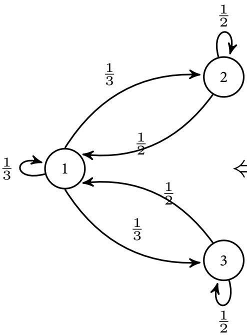

$$\iff T = \begin{bmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{2} & \frac{1}{2} & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} \end{bmatrix}$$

DeGroot model: Illustration 1 (Cont.)

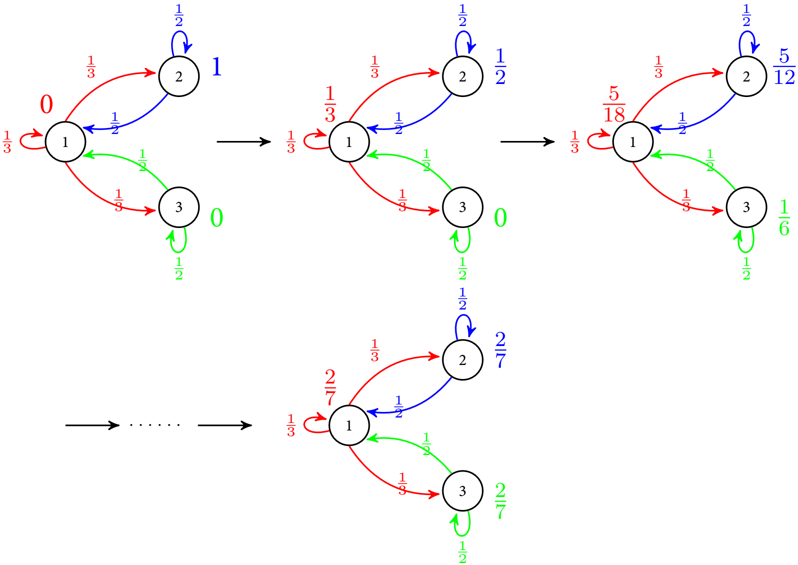

DeGroot model: Illustration 2

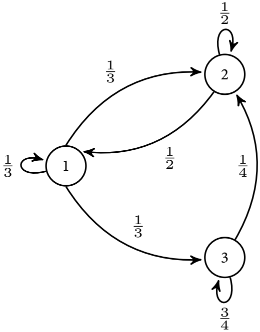

$$\iff T = \begin{bmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{2} & \frac{1}{2} & 0 \\ 0 & \frac{1}{4} & \frac{3}{4} \end{bmatrix}$$

DeGroot model: Illustration 2 (Cont.)

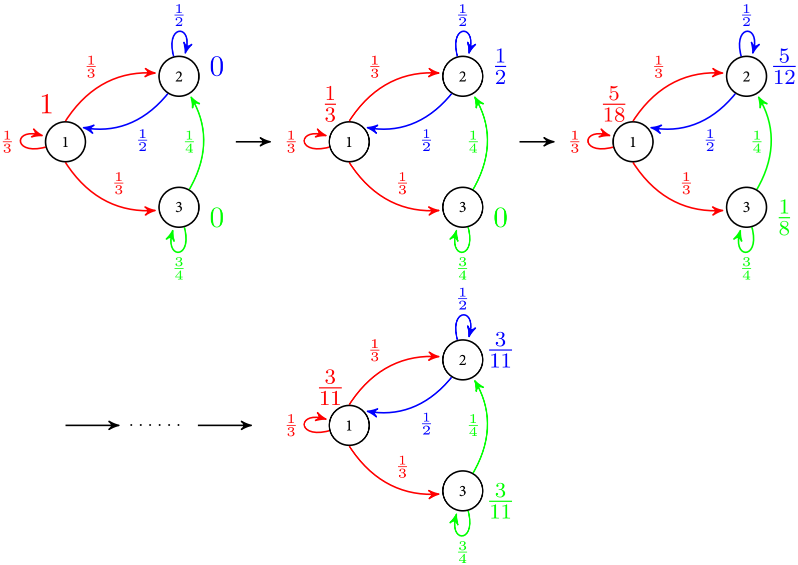

DeGroot model: Illustration 3

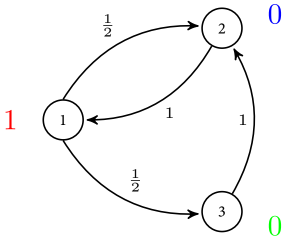

$$\Longleftrightarrow T = \begin{bmatrix} 0 & \frac{1}{2} & \frac{1}{2} \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$$

$$p(0) = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \to \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} \to \begin{pmatrix} \frac{1}{2} \\ 0 \\ 1 \end{pmatrix} \to \begin{pmatrix} \frac{3}{4} \\ \frac{1}{2} \\ 0 \end{pmatrix} \to \begin{pmatrix} \frac{1}{4} \\ \frac{3}{4} \\ \frac{1}{2} \end{pmatrix} \to \cdots \to \begin{pmatrix} \frac{2}{5} \\ \frac{2}{5} \\ \frac{2}{5} \end{pmatrix}$$

DeGroot model: Illustration 4

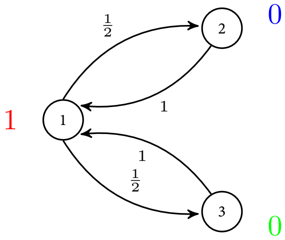

$$\iff T = \begin{bmatrix} 0 & \frac{1}{2} & \frac{1}{2} \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}$$

$$p(0) = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \to p(1) = \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix} \to \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \to \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix} \to \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \to \cdots$$

Convergence

- T converges if $\lim_{t \to \infty} T^t \cdot p(0)$ exists for all $p(0)$.
- T is aperiodic (非周期) if the greatest common divisor of its cycle lengths is one.
- Left: aperiodic; Right: periodic.

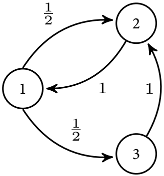

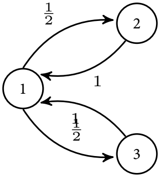

Convergence result

- Suppose the network is strongly connected— there is a directed path from any node to any other node.

* It is equivalent to assume the adjacency matrix T to be irreducible (不可约).

Result

T is strongly connected/irreducible, then T is convergent if and only if it is aperiodic.

- Result: T is strongly connected/irreducible, then T is convergent if and only if $\lim_{t\to\infty} T^t = (1, 1, \cdots, 1)^\intercal \cdot s$, where s is the unique left eigenvector of T associated with the eigenvalue 1.

Proof: Sufficiency

- Definition: T is primitive (素矩阵) if $T_{ij}^{t} > 0$ for all i and j after some t.
- Perkins (1961): If T is strongly connected and (row) stochastic, then it is aperiodic if and only if it is primitive.
- Meyer (2000): If T is strongly connected and primitive, then $\lim_{t\to\infty} T^t = (1, 1, \cdots, 1)^\intercal \cdot s$, where s is the unique left eigenvector of T associated with the eigenvalue 1.
- So strongly connectness and aperiodicity imply convergence.

Proof: Necessity

- Claim: If T is strongly connected, row-stochastic and convergent, then it is primitive.
- Since T is row-stochastic, Perron-Frobenius theorem implies that 1 is an eigenvalue of T and $1 \geq |\lambda|$ for any other eigenvalues λ of T.
- Let $S = \lim_{t \to \infty} T^t$.
- Then $ST = \lim_{t \to \infty} T^t T = S$.
- So each row is a left eigenvector of T with eigenvalue 1.

Proof: Necessity (Cont.)

- Since T is irreducible and nonnegative, Perron-Frobenius theorem implies that the eigenspaces associated with 1 is one-dimensional and T is a positive eigenvector associated with 1.
- Thus, each row of S can be taken to be positive.
- Since S is all positive, T is primitive.
- T is primitive then Perron-Frobenius theorem implies the eigenvector is unique, and all rows of S are the same s.

Convergence

- Aperiodicity is easy to satisfy.
- Have some agent weight him or herself.
  * If T is strongly connected and $T_{ii} > 0$ for some i, then T is aperiodic, and hence T is convergent.
- Have at least one communicating dyad and a transitive triple.
  * T is aperiodic if the greatest common divisor of its cycle lengths is one.

Consensus

- Beyond knowing whether or not beliefs converge, we are also interested in characterizing:
  - what beliefs converge to when they converge,
  - which agents have substantial influence in the society,
  - when it is that a consensus is reached.
- Agents reaches a consensus (共识) under T for an initial vector of beliefs $p(0)$ if $\lim_{t\to\infty} p_i(t) = \lim_{t\to\infty} p_j(t)$ for each i and j.

Consensus/convergence and aperiodicity

- Theorem: Agents reaches a consensus for every initial vector of beliefs under T if and only if T is aperiodic.
- Necessity: Consensus $\Rightarrow$ convergence $\Rightarrow$ aperiodicity.
- Sufficiency:

$$p(\infty) = \lim_{t\to\infty} T^{\top} p(0) = \begin{pmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{pmatrix} \begin{pmatrix} s_1 & s_2 & \cdots & s_n \end{pmatrix} \begin{pmatrix} p_1(0) \\ p_2(0) \\ \vdots \\ p_n(0) \end{pmatrix}$$

$$= \begin{pmatrix} s_1p_1(0) + \cdots + s_np_n(0) \\ s_1p_1(0) + \cdots + s_np_n(0) \\ \vdots \\ s_1p_1(0) + \cdots + s_np_n(0) \end{pmatrix} = \begin{pmatrix} s \cdot p(0) \\ s \cdot p(0) \\ \vdots \\ s \cdot p(0) \end{pmatrix}.$$

Consensus in beliefs

- The agents reach a consensus whenever T converges.
- The limit belief $p_i(\infty)$ is $s \cdot p(0)$, where s is the left eigenvector of T associated with eigenvalue 1.
- The belief converges to (normalized) eigenvector weighted (s) sum of original beliefs $p(0)$.

Wise learning

- Consensus is not necessarily a good thing.
- In the herding example, there is consensus, but this could lead to the wrong outcome.
- We would like to consensus to be at

$$p(\infty) = \frac{1}{n} \sum_{i=1}^{n} p_i(0) = \theta,$$

so that individuals learn the underlying state.
- If this happens, we say that the society is wise.

Wise learning (Cont.)

Result

The society is wise if and only if T is doubly stochastic.

- Intuition: Otherwise, there is no balance in the network, so some agents are influential; their opinion is listened to more than they listen to other people's opinion.

1. Observational learning
   - Herding
   - Herding and sequential decision-making
   - Learning from neighbors

2. DeGroot model
   - Convergence
   - Consensus in beliefs
   - Wise learning
   - Social influence

Limiting beliefs

- Limiting beliefs would be weighted averages of the initial beliefs.
- The relative weights would be the influences that the various agents have on the final consensus beliefs.
- $p_i(\infty) = s \cdot p(0)$.
- $s_i = \sum_j s_j T_{ji}$.

$\Rightarrow$ High influence from being paid attention to by people with high influence.

- Related to eigenvector centrality (left eigenvector centrality).

Influential agents

- A set of agents B is called an influential family if the beliefs of all agents outside B is affected by beliefs of B (in finitely many steps).
- The presence of influential agents implies no asymptotic learning:
  - The presence of influential agents is the same thing as 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.

Stubborn agents

- An agent who places high weight on self will maintain belief while others converge to that agent's belief.
- Groups that are highly introspective will have substantial influence.

Equal weights

- Suppose equally weight connections.
- Suppose also that $T_{ij} > 0$ if and only if $T_{ji} > 0$.
- d_i is i's degree.
- So, $T_{ij} = \frac{1}{d_i}$ for each i and j that i has a (directed) link to.

$\Rightarrow$ Weight friends equally.

Equal weights (Cont.)

- Let $D = \sum_k d_k$.
- Claim: $s_i = \frac{d_i}{D}$ for each i.
- Verify: $s_i = \sum_j s_j T_{ji} = \sum_{j:\ ji \in g} \frac{1}{d_j} \frac{d_j}{D} = \frac{d_i}{D}$ (degree centrality).

- 周围的社会结构对于一个简单的决策会发挥决定性的作用："德格鲁特学习"模型，其最终的结果取决于每个节点的初始估计数和特征向量中心度。
- 两种偏差：回声，重复计算（更广泛）。
- 如果满足几项关键条件，德格鲁特学习模型也能带来非常准确的结果：多样性的观点，不能有系统性偏差，交流网络具有良好的平衡性（每个人的特征向量中心度相对于其他人的中心度之和而言足够小）。
- 选择性关注：网络结构失衡时，人们对小团体过度关注，形成了"个别人物法则"。

* 互联网时代信息极大丰饶，但严谨的新闻调查的激励被削弱，假新闻泛滥成灾。