Lecture 5: The DeGroot Learning Model

Alexander Wolitzky

MIT

6.207/14.15: Networks, Spring 2022

DeGroot Learning Model: Preview

The DeGroot learning model (introduced in 1974 by statistician Morris DeGroot) is a simple, important model of how people in a network update their opinions over time and eventually reach a group consensus

Basic idea: each period, every agent in the network updates her opinion by taking a weighted average of her own opinion and her neighbors' opinions, with constant, time-invariant weights.

Given all the agents' initial opinions and the matrix describing what (constant) weight each agent puts on each other agent's opinion, we will be able to compute the dynamics of everyone's opinions and the long-run group consensus belief.

Among other interesting results, we'll see that an agent's **influence** on the group's long-run consensus belief is determined by her eigenvector centrality in the weight matrix.

Does Repeated Averaging Make Sense?

Before getting into the math, let's step back and ask what such a model of repeated averaging can and cannot capture.

First of all, understanding **social learning**—how groups of people aggregate their information and form beliefs and preferences—is a key question in social science, with countless applications.

- Which candidate becomes popular and wins an election?
- What are the prices for shares in various companies on the stock market?
- Which restaurants, books, or movies become popular?

Good models of social learning can guide empirical analysis, decisions, and policy on these topics. So the general problem is quite important. We will spend substantial time in this course on understanding social learning and information aggregation.

Does Repeated Averaging Make Sense? (cntd.)

There are two broad approaches to modeling social learning.

Bayesian social learning: Agents are Bayesian statisticians. Agents are trying to learn some unknown state of the world (e.g., a company's long-run profitability), they start with a prior belief, and they update it using Bayes' rule in light of information they receive from the environment or from other agents.

- These models are in some sense the "gold standard" for understanding rational learning. However, they assume a lot of rationality on the part of agents and they can be very complicated outside of simple examples.
- Great models of how people *should* learn, whether they're great models of how people *do* learn is more nuanced.
- We will study these models (and assess their strengths and weaknesses) later in the course after we introduce Bayesian game theory.

Does Repeated Averaging Make Sense? (cntd.)

Non-Bayesian/"rule of thumb" learning: Agents do something simpler, more "heuristic" than Bayesian learning.

You can perhaps think of these models as

- a rough-and-ready approximation of Bayesian learning when that's too complicated, or
- a more accurate model for real people who aren't so rational, or
- a description of updating something that isn't quite a "belief" in the proper probabilistic sense (e.g., how people update their "opinion" on an issue, or their tastes for one product or another).

These various interpretations may make more or less sense in different contexts.

Does Repeated Averaging Make Sense? (cntd.)

The DeGroot learning model is the simplest and best-known model of non-Bayesian social learning in a network.

- The DeGroot repeated averaging procedure is typically not well-founded in terms of rational social learning. (The agents in the model aren't "good Bayesians.")
- Nonetheless, the model is still useful and important for the above reasons.
- It is also useful as a complement to the Bayesian models we'll see later in the course, which arguably have stronger foundations but are often less tractable.

DeGroot Model

- Finite set of agents $N = \{1, \ldots, n\}$.
- Discrete time: $t = 0, 1, 2, \ldots$
- Interactions capture by an $n \times n$ non-negative **updating matrix** A.
  - $A_{ij} \geq 0$ indicates the level of weight or trust that i puts on j
  - A is a **row-stochastic matrix**: $\sum_{j=1}^{n} A_{ij} = 1$ for all i.
- Each agent i has **initial belief** (or "opinion"; not a belief in the Bayesian sense) $x_i(0) \in [0, 1]$.
- From period t to $t+1$, agent i updates her belief by linear averaging according to A:

$$x_i(t+1) = \sum_{j=1}^{n} A_{ij} x_j(t).$$

Opinion Dynamics

How does the vector of beliefs $x(t)$ evolve over time?

We have

$$x_i(t) = \sum_{j=1}^{n} A_{ij} x_j(t-1) \text{ for all } i \in N \text{ and } t > 0.$$

In matrix notation, this says

$$x(t) = Ax(t-1).$$

Iterating, we obtain

$$x(t) = A^t x(0).$$

Thus, belief vector $x(t)$ evolves as a Markov chain with transition matrix A.

- This is the same as the "viruses" we used to interpret eigenvector centrality in Lecture 3, with the difference that here we take the dynamics literally. Also, because A is row-stochastic the mass of viruses at each node always stays between 0 and 1.

Example

Suppose $n = 3$ and the updating matrix is given by

$$A = \begin{pmatrix} 1/3 & 1/3 & 1/3 \\ 1/2 & 1/2 & 0 \\ 0 & 1/4 & 3/4 \end{pmatrix}.$$

- Agent 1 puts equal weight on everyone's opinion.
- Agent 2 puts $1/2$ weight each on her own opinion and agent 1's opinion.
- Agent 3 puts $3/4$ weight on her own opinion and puts $1/4$ weight on agent 1's opinion.

Example (cntd.)

Suppose the vector of initial opinions is

$$x(0) = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}.$$

- Agent 1 starts with opinion 1; agents 2 and 3 start with opinion 0.

After 1 period of updating, beliefs are given by

$$\begin{aligned} x(1) &= Ax(0) \\ &= \begin{pmatrix} 1/3 & 1/3 & 1/3 \\ 1/2 & 1/2 & 0 \\ 0 & 1/4 & 3/4 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \\ &= \begin{pmatrix} 1/3 \\ 1/2 \\ 0 \end{pmatrix}. \end{aligned}$$

Example (cntd.)

After 2 periods of updating, beliefs are given by

$$\begin{aligned} x(2) &= Ax(1) = A^2 x(0) \\ &= \begin{pmatrix} 1/3 & 1/3 & 1/3 \\ 1/2 & 1/2 & 0 \\ 0 & 1/4 & 3/4 \end{pmatrix} \begin{pmatrix} 1/3 \\ 1/2 \\ 0 \end{pmatrix} \\ &= \begin{pmatrix} 5/18 \\ 5/12 \\ 1/8 \end{pmatrix}. \end{aligned}$$

Example (cntd.)

To calculate the belief dynamics, we keep left-multiplying by the matrix A.

In the long-run, beliefs converge to

$$\begin{aligned} x^* &= \lim_{t \to \infty} A^t x(0) \\ &= \begin{pmatrix} 3/11 \\ 3/11 \\ 3/11 \end{pmatrix}. \end{aligned}$$

- We'll see how this is computed later in the lecture.

Two things to note:

1. Everyone's belief converged (no cycling).
2. Everyone's belief converged to the same thing (the group reached **consensus**).

Do Beliefs Always Converge in the Long Run?

In the example, the belief vector $x(t)$ converged in the long-run (no cycling), and the limiting belief vector was constant (consensus).

This happens for "typical" updating matrices A, but not for all of them.

An Example with Cycles

Suppose the updating matrix is

$$A = \begin{pmatrix} 0 & 1/2 & 1/2 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}.$$

- Agent 1 puts equal weight on the other agents' opinions and 0 weight on her own opinion; the other agents put all weight on agent 1.

Suppose the vector of initial opinions is

$$x(0) = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}.$$

Example with Cycles (cntd.)

Then

$$x(1) = \begin{pmatrix} 0 & 1/2 & 1/2 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix}.$$

And

$$x(2) = \begin{pmatrix} 0 & 1/2 & 1/2 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}.$$

But this brings us back to $x(0)$. Hence, $x(t)$ will alternate between $x(0)$ and $x(1)$ forever—beliefs will not converge.

Aperiodic Matrices

Mathematically, the problem here is that powers of the matrix A cycle: that is, A is **periodic**.

$$
\begin{aligned}
A &= \begin{pmatrix} 0 & 1/2 & 1/2 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix} \\
A^2 &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1/2 & 1/2 \\ 0 & 1/2 & 1/2 \end{pmatrix} \\
A^3 &= \begin{pmatrix} 0 & 1/2 & 1/2 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix} = A.
\end{aligned}
$$

Aperiodic Matrices (cntd.)

In general, a matrix is **aperiodic** if the greatest common divisor of all directed cycle lengths is 1 (where "directed cycles" are defined in the network where there is a directed link from i to j iff $A_{ij} > 0$).

- The condition holds in the first example above, but not in the second example as in that case the length of all directed cycles are even.

A simple sufficient condition for a matrix to be aperiodic is that there exists some agent i such that $A_{ii} > 0$.

- That is, if anyone puts positive weight on her own opinion then the matrix is aperiodic.

Important (and intuitive) fact from Markov chain theory: if the matrix A is aperiodic, then $\lim_{t\to\infty} A^t$ converges.

- Another consequence of the Perron-Frobenius theorem we encountered in Lecture 3.

Strongly Connected Matrices

Even if $\lim_{t\to\infty} A^t$ converges, agents can have different long-run beliefs if the matrix is not strongly connected (i.e., for some i and j, there is no directed path from i to j).

Trivial example:

$$A = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

- Each agent sticks with her initial belief forever, so if $x_1(0) \neq x_2(0)$ there is no long-run consensus.

However, if the matrix is strongly connected, then if beliefs do converge, they must converge to a vector where everyone has the same belief.

- Otherwise, every agent with the lowest limit belief puts weight only on agents with weakly higher limit beliefs, and some agent with the lowest limit belief puts strictly positive weight on some agent with a strictly higher limit belief, which is impossible.

Long-Run Consensus

Putting together what we've said about aperiodic and strongly connected matrices, we have:

Theorem

Consider the directed network defined by A, where there is a link from i to j iff $A_{ij} > 0$. If this network is strongly connected and aperiodic, then $\lim_{t\to\infty} A^t = A^*$ exists. Moreover, the rows of the matrix A^* are all the same: for every vector of initial beliefs x, we have $(A^*x)_i = (A^*x)_j$ for all agents i and j.

Note: the long-run updating matrix A^* is determined by the 1-shot updating matrix A, independent of the initial beliefs.

- This implies that each agent gets the same weight in forming the long-run consensus, regardless of the initial beliefs.
- Since the long-run consensus is given by A^*x, agent i's weight is the i^{th} element of (any one of the) rows of the matrix A^*.
- We now show how to compute each agent's weight. This also lets us compute the long-run beliefs A^*x.

Long-Run Influence

Agent i's weight in determining the long-run consensus is called her **long-run influence** s_i.

Formally, say that a column vector $s \in [0,1]^n$ is a **long-run influence vector** if $\sum_i s_i = 1$ and

$$x_1^* = s' \cdot x(0) = \sum_{i=1}^{n} s_i x_i(0)$$

for every vector of initial beliefs $x(0)$.

- (Since x_i^* is the same for all i, could have written x_i^* for any i instead of x_1^*.)

Long-Run Influence (cntd.)

If s is a long-run influence vector, then we have

$$s' \cdot x(1) = s' \cdot (Ax(0)) = s' \cdot x(0)$$

(as both equal x_1^*).

Since this must hold for every vector $x(0)$, we have

$$s'A = s', \text{ or equivalently } A's = s.$$

That is, s is a (actually, the) **unit eigenvector of** A'.

This is precisely the vector of **eigenvector centralities** of A' (as defined in Lecture 2).

- Punchline: in DeGroot learning, an agent's long-run influence equals her eigenvector centrality.

Computing Long-Run Influence and Beliefs

We now have a recipe for computing long-run influence and beliefs.

1. First compute the long-run influence vector s (i.e., the unit eigenvector of the updating matrix A').
2. Then compute the long-run belief as $s' \cdot x(0)$.

Note: to compute the long-run beliefs given updating matrix A and initial beliefs $x(0)$, we start by **ignoring** the initial beliefs and instead computing the influence vector s (which doesn't depend on initial beliefs), and **then** calculate the long-beliefs as an average of the initial beliefs with weights given by s.

Let's see how this works in an example.

Example

Recall our first example:

$$A = \begin{pmatrix} 1/3 & 1/3 & 1/3 \\ 1/2 & 1/2 & 0 \\ 0 & 1/4 & 3/4 \end{pmatrix}.$$

The long-run influence vector s is given by the system of equations:

$$\begin{aligned} s_1 &= \frac{1}{3}s_1 + \frac{1}{2}s_2 \\ s_2 &= \frac{1}{3}s_1 + \frac{1}{2}s_2 + \frac{1}{4}s_3 \\ s_3 &= \frac{1}{3}s_1 + \frac{3}{4}s_3. \end{aligned}$$

Example (cntd.)

Solving this system of equations gives

$$
\begin{aligned}
s_1 &= \frac{3}{11} \\
s_2 &= \frac{4}{11} \\
s_3 &= \frac{4}{11}.
\end{aligned}
$$

For example, if the vector of initial beliefs is given by

$$x(0) = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix},$$

then the long-run consensus is given by

$$x^* = s' \cdot x(0) = \frac{3}{11}.$$

Wisdom of the Crowd

Another question: with DeGroot learning, when are large networks "wise"?

- Suppose each initial opinion $x_i(0)$ is an independent, unbiased estimate of some true state of the world θ.
- When does the consensus long-run belief converges in probability to θ as the network grows large?

- Recall that the long-run consensus equals $\sum_{i=1}^{n} s_i x_i(0)$.
- By the law of large numbers, this sum converges in probability to θ if and only if

$$\lim_{n\to\infty} \max_{i\leq n} s_i(n) = 0.$$

- That is, the network is wise if and only if each node's long-run influence is small.

Equal Long-Run Influence

Under what conditions is each node's long-run influence small?

Different conditions can be given. A simple approach is to ask when all nodes have **equal** long-run influence, in that $s_i = \frac{1}{n}$ for all i.

It turns out that all nodes have equal long-run influence if and only if the network is **doubly stochastic.**

Stochastic Matrices

We've assumed A is **row stochastic matrix**: $\sum_j A_{ij} = 1$ for all i.

- This says $x_i(t+1)$ is a weighted average of the $x_j(t)$'s (including $x_i(t)$).

A matrix A is called **doubly stochastic** if all rows *and columns* sum to 1:

$$\sum_j A_{ij} = 1 \text{ for all } i, \text{ and } \sum_i A_{ij} = 1 \text{ for all } j.$$

- In DeGroot learning, this means that each agent gets the same average weight in others' updating rules.

Equal Long-Run Influence (cntd.)

Theorem

In DeGroot learning, everyone has equal long-run influence if and only if the influence matrix A is doubly stochastic.

We skip the proof.

Simple example: suppose $n = 2$, so A is a 2x2 matrix.

- In this case, "doubly stochastic" simply means that the weight 1 puts on 2's opinion equals the weight that 2 puts on 1's opinion.
- The theorem says that $s_1 = s_2 = \frac{1}{2}$ if and only if this condition holds, which makes sense.

Generalizations of DeGroot

Since the DeGroot model is so simple, it can be extended in various directions to capture other possible aspects of rule-of-thumb social learning. We'll mention just a couple possibilities.

Time-varying weight on own belief (DeMarzo, Vayanos, and Zwiebel, 2000)*:* If agents realize they are learning over time (or wish to avoid "double-counting" their neighbors' beliefs), they might put more weight on their own belief. A natural model of this is

$$x_i(t) = (1 - \lambda_t) x_i(t-1) + \lambda_t \sum_{j=1}^{n} A_{ij} x_j(t-1),$$

where λ_t decreases over time. (λ_t constant is standard DeGroot.)

Generalizations of DeGroot (cntd.)

Ignoring people with distant beliefs (Krause, 2000): Perhaps only people with nearby beliefs are worth listening to. A natural model of this is

$$x_i(t) = \frac{\sum_{j \in N_i : |x_i(t-1) - x_j(t-1)| < d} x_j(t-1)}{|\{j \in N_i : |x_i(t-1) - x_j(t-1)| < d\}|}.$$

In these models, under some conditions the group can be partitioned into sets of agents who form distinct consensuses. See Jackson Ch. 8 if curious.

Generalizations of DeGroot (cntd.)

Initially uninformed agents (Banerjee, Breza, Chandrasekhar, and Mobius, 2021): Perhaps only a subset of agents $S \subset N$ initially have any opinion at all, and other agents develop opinions (according to the usual weighted averaging) only when at least one of their neighbors develops an opinion (or is in the initial "seed set" S).

- Now each agent $i \in S$ gets extra influence, as she pushes her beliefs onto all agents in $N \backslash S$ who are closer to i than to any other agent in S.
- The set of agents in $N \backslash S$ who are closer to i than to anyone else in S is called the **Voronoi set** of i.

In this model, the group always eventually reaches a single consensus as in standard DeGroot. However, roughly speaking, the influence of each node $i \in S$ on the consensus is now determined by the size of her Voronoi set. This is the number of agents that agent i "persuades" before the agents in the network start averaging their opinions.

Remark: Aggregation vs. Diffusion

DeGroot learning is a model of opinion **aggregation**: everyone starts with an opinion, we ask how they get aggregated into a social consensus.

In the coming weeks, we will spend considerable time on models of **diffusion** (of information, but also new products, diseases, etc.): only a few people start with a piece of information, we ask how it spreads through society.

The Banerjee et al model is a hybrid: "having an opinion" diffuses, while people who already have opinions aggregate them a la DeGroot.

It's useful to understand the difference between aggregation and diffusion, while also recognizing that in the real-world social learning has elements of both.

DeGroot Learning: Summary

- DeGroot learning is a simple and tractable model of imitative, non-Bayesian social learning.
- If the updating matrix is strongly connected and aperiodic, the group reaches consensus in the long-run.
- The long-run consensus is computed by first computing each individual's long-run influence, which equals her eigenvector centrality.
- Even though DeGroot learning is very simple and somewhat naive, it satisfies a kind of "wisdom of the crowds" if each node is equally inflential. For example, this is the case if the updating matrix is doubly stochastic.
- The DeGroot model is amendable to various extensions that can sometimes be more realistic. However, typically these extensions only partially preserve the DeGroot model's strong results on consensus and influence.