LEARNING AND INFLUENCE IN NETWORKS: LECTURE I

BENJAMIN GOLUB

Abstract. This lecture covers classical models of repeated Bayesian learning in social networks.

1. Introduction

1.1. **Some motivating questions.**

(1) Why is there disagreement about factual matters?
   (a) Can we explain it rationally?
   (b) How does it depend on the network?
(2) Can decentralized learning (via talking or trading) diffuse/aggregate knowledge effectively?
   (a) How does this depend on the information people get?
   (b) How does this depend on the network?

1.2. **Overview of mini-course.** We start with Aumann's [1976] model of "interactive reasoning" and examine social learning models in that tradition, as well as sequential social learning. We then set out to address some of what the classical approach hasn't successfully analyzed: who is influential, what learning is like when it is imperfect, the causes of disagreement, and the rate of convergence. In the last lecture, we turn to recent work that aims to obtain models as tractable as the behavioral ones just discussed, but with Bayesian foundations.

Date Printed. October 15, 2018.
Department of Economics, Harvard University.

2. A general framework

•

A *social learning setting* consists of:

- N – a set of players;
- A – a common action space; e.g., $A = \{0, 1\}$ or $A = \mathbb{R}$;
- Θ – state space; e.g., $\Theta = \{\mathrm{H}, \mathrm{L}\}$;
- $u : A \times \Theta \to \mathbb{R}$ – a common utility function;
- S – private signal space;
- μ_i for each i – i's prior over $\Omega = \Theta \times S^N$.

The general idea is to start with an environment like this and then specify an extensive form: timing of moves and observation opportunities. We also specify a solution concept. We will study two examples in depth that make clear what this entails.

3. Martingales to Aumann

3.1. **Setting.** Take a social learning setting and suppose that the set N of players are located on a graph, G given by its *neighborhoods* $\mathcal{N}(i) \subseteq N$. The timing is as follows.

- At $t = 0$, receive signals s_i and take actions $a_i(0)$ [all simultaneously].
- For $t \geq 1$, observe $a_j(s)$ for all $s < t$, $j \in \mathcal{N}(n)$ and take actions $a_i(t)$ [all simultaneously].

Let $I_{i,t}$ be the information of individual i at time t (formally a σ-algebra on Ω).

For simplicity let's take

- $\Theta = \{0, 1\}$;
- $A = \mathbb{R}$ and a scoring-rule so that u is maximized by reporting $\mathbb{P}(\theta = 1 \mid I_{i,t})$;
- the solution concept as myopic best-response to all predecessors' strategies.

This is a generalization of the process of Geanakoplos and Polemarchakis [1982]. These sorts of models were later studied by Parikh and Krasucki [1990], Mueller-Frank [2013], and others.

3.2. **Basic questions.**

(1) Does each individual's belief converge? That is, does $\lim_t a_i(t)$ exist for each i?
(2) Do individuals come to agree? Is it the case that $\lim_t a_i(t) = \lim_t a_j(t)$ for all i, j?
(3) Does the eventual outcome aggregate information efficiently?

3.3. **Results.** On the first question, we have the following simple result.

Proposition 1. *$\overline{a}_i = \lim_t a_i(t)$ exists a.s. for each i.*

Proof. Each $(a_i(t)))_t$ is a bounded martingale with respect to the filtration $(I_{i,t})_t$. □

Now we consider the question of agreement.We say the graph is undirected if $i \in \mathcal{N}(j)$ implies $j \in \mathcal{N}(i)$. We say an undirected graph is connected if there is no nonempty proper subset $M \subseteq N$ so that $\mathcal{N}(M) \subseteq M$.

We will make the drastic simplification that Ω is finite.

Proposition 2. *If G is undirected and connected, and the prior is common, for each i and j, a.s., $\overline{a}_i = \overline{a}_j$.*

Proof. Define $\overline{I}_i$ to be the limit of the increasing sequence of σ-algebras $(I_{i,t})_t$. From the facts that G is undirected and Ω is finite we deduce that if i and j are neighbors, then $\mathbb{E}[\theta = 1 \mid \overline{I}_i]$ and $\mathbb{E}[\theta = 1 \mid \overline{I}_j]$ are common knowledge at sufficiently late times. Then by Aumann's Theorem they must be the same. By connectedness this holds across the entire network. □

This leaves the question of whether aggregation is accurate.

Example 1. Take $N = \{1, 2\}$. Let s_1 and s_2 be independent, drawn uniformly at random from $S = \{0, 1\}$. Let $\theta = s_1 \text{ XOR } s_2$.

Note that $\mathbb{E}[\theta \mid s_i] = \frac{1}{2}$ for both values of s_i, so no information is revealed by an announcement.

Then $a_i(t) = \frac{1}{2}$ for all t but clearly information is not aggregated.

However, this example is nongeneric. If we think about the prior over $\Omega = \Theta \times S_1 \times S_2$ that gave rise to this, it is very uniform. If we were to perturb it a little bit, then announcements would become informative and information would get aggregated. Indeed, we have the following (loosely stated) result.

Claim 1. With a finite Ω and generic priors over Ω, for a nontrivial event E, posterior probabilities of E fully reveal all private information.

3.4. **Remarks and take-aways.**

- Social learning in networks ties in naturally with central ideas of economic theory, and the frameworks associated with them. E.g.:
  - beliefs are martingales;
  - agreeing to disagree.
- What is robust (and not)?
  - Convergence of individual beliefs: This is a quite robust feature as long as (i) individuals don't forget and (ii) the state in question is fixed.
    - Related to studies of market prices: Ostrovsky (ECMA 2012).
  - Agreement: This is fairly robust, certainly in the belief announcement framework. Raises some questions:
    - Technical: What if Ω weren't finite? To my knowledge this isn't fully resolved.
    - More substantively: is undirected important? Yes, I think it can't be dropped fully. Exercise: Find a counterexample.

        * Important extension: what can we conclude with a coarser action space? If i takes action a infinitely often and j takes action $a' \neq a$ infinitely often, then it is eventually common knowledge between them that i thinks a is weakly better while j thinks a' is better. I.e. it is eventually common knowledge that

$$\mathbb{E}[u(a,\theta) \mid \overline{I}_i] \geq \mathbb{E}[u(a,\theta) \mid \overline{I}_i] \qquad \text{and} \qquad \mathbb{E}[u(a,\theta) \mid \overline{I}_j] \leq \mathbb{E}[u(a,\theta) \mid \overline{I}_j].$$

          By same logic as Aumann's Agreeing to Disagree result, this forces equalities everywhere. So neighbors have to be indifferent between actions taken infinitely often.

    - Aggregation:
        * How do we assess the result?
            · Aggregation is perfect because people pay very careful attention to exactly what the report reveals about everything that's happened.
            · Result therefore insensitive to any aspect of network structure, etc.
        * Seems quite fragile. But do we learn anything from this fragility?
            · Genericity assumptions are deceptive when there is a dimensionality/cardinality mismatch.
            · "Smuggling information" is feasible in designed protocols but it can even happen "naturally".
            · For social learning questions to be interesting, actions should not reveal beliefs about *everything.*
    - Overall gripe: network doesn't matter very much beyond coarse/quantitative properties. Individuals' deductive abilities overwhelm everything.

References

R. J. Aumann. Agreeing to disagree. *Annals of Statistics*, 4:1236–1239, 1976.

J.D. Geanakoplos and H.M. Polemarchakis. We can't disagree forever. *Journal of Economic Theory*, 28:192–200, 1982.

M. Mueller-Frank. A general framework for rational learning in social networks. *Theoretical Economics*, 8(1):1–40, 2013.

R. Parikh and P. Krasucki. Communication, consensus, and knowledge. *Journal of Economic Theory*, 52:178–189, 1990.