Linear Social Learning in Networks with Rational Agents

Krishna Dasaratha   Kevin He

February **29**, 2020

How **Fast** Do Agents Learn in Different Networks?

- **Social learning**: info about unknown state dispersed among society of agents, agents act based on private signals and observations of social neighbors
- Social network affects the **rate** of private signal aggregation
- Speed matters: two networks both leading to correct long-run consensus can have very different welfare implications

Golub and Sadler (2016): "A significant gap in our knowledge concerns short-run dynamics and rates of learning in these models. [...] The complexity of Bayesian updating in a network makes this difficult, but even limited results would offer a valuable contribution to the literature."

Contributions and Key Results

Introduce **tractable model of rational sequential learning** that focuses on the role of social network on speed of learning

- rich signals, rich actions: Gaussian private signal, infer neighbors' beliefs perfectly from their actions
- strips away other sources of learning-rate inefficiency
- unique equilibrium of social-learning game has log-linear form

Highlight **network-based informational confounds**

- suppose 2 and 3 see 1, but 4 sees only 2 and 3
- 1's action confounds the info content of 2 and 3's behavior
- show how rational agents solve this signal-extraction problem

Generations network – observe subset of agents in previous gen

- derive how learning rate depends on network parameters
- info confounds in gen structure $\Rightarrow$ obstacle to fast learning
- extent of info loss: when network symmetric enough, social learning aggregates no more than 2 signals per gen asymptotically, even when gens are arbitrarily large

Related Literature

Rational sequential learning

- Banerjee (1992), Bikhchandani, Hirshleifer, and Welch (1992)
- Correct learning under mild conditions: Acemoglu, Dahleh, Lobel, and Ozdaglar (2011), Lobel and Sadler (2015). This paper: speed.

Obstructions to the efficient rate of learning

- Coarseness of the action space: Harel, Mossel, Strack, and Tamuz (2019 WP), Rosenberg and Vieille (2019), Hann-Caruthers, Martynov, and Tamuz (2018)
  - HMST's "rational groupthink": trapped in wrong consensus for a long time as small belief changes not reflected in actions
  - Rate of learning efficient if actions were rich
- Choosing from multiple info sources: Liang and Mu (2020)
- This paper: network-based obstructions to fast learning

Lobel, Acemoglu, Dahleh, and Ozdaglar (2009): compare two specific network structures with neighborhood size 1 — immediate past agent or one uniformly random past agent. This paper: arbitrary fixed networks.

Speed of learning under non-rational heuristics: Ellison and Fudenberg (1993), Golub and Jackson (2012). This paper: rational learning.

Outline

1. Setup and example of informational confound
2. Warm-up results
    2.1 Log-linearity of the equilibrium
    2.2 Signal-counting interpretation of equilibrium accuracy
    2.3 An IFF condition for long-run learning
3. The generations network
    3.1 Generation size and speed of learning
    3.2 Uniform bound on number of signals aggregated per generation
4. (if time) When does adding links improve accuracy?

Model and Notations

- Two equally likely states $\omega \in \{0, 1\}$
- Agents $i = 1, 2, 3, \ldots$ move in order, each acting once
  - i observes **private signal** $s_i \in \mathbb{R}$ and actions of **neighbors**, $N(i) \subseteq \{1, \ldots i-1\}$
  - picks **action** $a_i \in [0, 1]$ to maximize expectation of $-(a_i - \omega)^2$
- Signals are Gaussian and conditionally i.i.d. given state, $s_i \sim \mathcal{N}(1, \sigma^2)$ when $\omega = 1$ and $s_i \sim \mathcal{N}(-1, \sigma^2)$ when $\omega = 0$
- Neighborhoods define an **observation network** M, with $M_{i,j} = 1$ if $j \in N(i)$, $M_{i,j} = 0$ else. M is common knowledge.
- A **strategy** for i specifies i's play as a function of:
  1. observed actions from neighbors $N(i)$, and
  2. private signal s_i.
- Sequential nature of game $\Rightarrow$ there is a unique perfect-Bayesian **equilibrium** strategy profile

An Example of Informational Confound

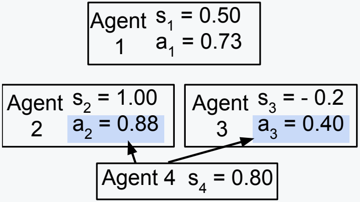

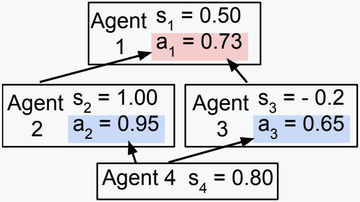

- 4 perfectly infers 2 and 3's signals from their actions
- 4's accuracy = 3 signals, fully incorporates info in s_2, s_3, and s_4

- a_1 influences both a_2 and a_3, but is unobserved by 4
- 4 cannot fully incorporate s_2 and s_3 without over-counting s_1
- optimal signal extraction: 4 puts "**2/3 as much weight**" on a_2 and a_3 as in other network
- 4's accuracy = "**3.67 signals**"
  - (to be formalized soon)

Log-Linearity of the Equilibrium

WLOG apply log-transformations and work with log-variables

- **log-signal**, $\tilde{s}_i := \ln\left(\frac{\mathbb{P}[\omega=1|s_i]}{\mathbb{P}[\omega=0|s_i]}\right)$, **log-actions**, $\tilde{a}_i := \ln\left(\frac{a_i}{1-a_i}\right)$
- these changes are 1-to-1, so there is a (unique) map from i's neighbors' log-actions and i's log-signal to i's eqm log-action
- next proposition says this map is linear

Proposition 1

For each agent i with $N(i) = \{j(1), ..., j(n_i)\}$, there exist constants $(\beta_{i,j(k)})_{k=1}^{n_i}$ s.t.

$$\tilde{a}_i^* = \tilde{s}_i + \sum_{k=1}^{n_i} \beta_{i,j(k)} \tilde{a}_{j(k)}^*.$$

The vector of coefficients $\vec{\beta}_{i,\cdot}$ is given by

$$\vec{\beta}_{i,\cdot} = 2\mathbb{E}[(\tilde{a}_{j(1)}^*, ..., \tilde{a}_{j(n_i)}^*) \mid \omega = 1] \cdot \mathrm{Cov}[\tilde{a}_{j(1)}^*, ..., \tilde{a}_{j(n_i)}^* \mid \omega = 1]^{-1}.$$

Discussion of Proposition 1

Proposition 1

For each agent i with $N(i) = \{j(1), ..., j(n_i)\}$, there exist constants $(\beta_{i,j(k)})_{k=1}^{n_i}$ s.t. $\tilde{a}_i^* = \tilde{s}_i + \sum_{k=1}^{n_i} \beta_{i,j(k)} \tilde{a}_{j(k)}^*$. The vector of coefficients $\vec{\beta}_{i,\cdot}$ is given by

$$\vec{\beta}_{i,\cdot} = 2\mathbb{E}[(\tilde{a}_{j(1)}^*, ..., \tilde{a}_{j(n_i)}^*) \mid \omega = 1] \cdot \mathrm{Cov}[\tilde{a}_{j(1)}^*, ..., \tilde{a}_{j(n_i)}^* \mid \omega = 1]^{-1}.$$

- For general private signal distributions, Bayesian updating in networks intractable as Golub and Sadler (2016) point out
- Gaussian info structure leads to **log-linear eqm** and **closed-form expression of linear weights** that solve signal-extraction problem: downweight neighbors' log-actions if they have higher equilibrium correlation conditional on ω
- $\vec{\beta}_{i,\cdot}$ depends on network M, but not on signal precision $1/\sigma^2$
- Can use Prop 1 inductively to express $\vec{\beta}_{i,\cdot}$ and how a_i^* incorporates predecessors' signals just as a function of M

Signal-Counting Interpretation of Eqm Accuracy

- Imagine i observes $n \in \mathbb{N}_+$ indep private signals, no other info
- Then $\tilde{a}_i \sim \mathcal{N}\left(\pm n \cdot \frac{2}{\sigma^2}, n \cdot \frac{4}{\sigma^2}\right)$

Definition

Social learning **aggregates** $r \in \mathbb{R}_+$ **signals by agent** i if the equilibrium log-action $\tilde{a}_i^*$ has the conditional distributions $\mathcal{N}\left(\pm r \cdot \frac{2}{\sigma^2}, r \cdot \frac{4}{\sigma^2}\right)$ in the two states.

- When agents use arbitrary strategy profile, conditional distributions of $\tilde{a}_i$ need not equal $\mathcal{N}\left(\pm r \cdot \frac{2}{\sigma^2}, r \cdot \frac{4}{\sigma^2}\right)$ for any r
- But, **equilibrium** log-actions always admit this kind of signal-counting interpretation

Proposition 2

There exist $(r_i)_{i \geq 1}$ so that social learning aggregates r_i signals by agent i. These $(r_i)_{i \geq 1}$ depend on the network M, but not on σ^2.

Condition for Long-Run Learning

Say society **learns completely in the long run** if equilibrium actions (a_i^*) converge to ω in probability.

Proposition 3

Society learns completely in the long if and only if $\lim_{i\to\infty}\left[\max_{j\in N(i)} j\right]=\infty.$

- If we consider the most recent neighbor of each agent, then this sequence of most-recent neighbors tends to ∞
- Analog of Acemoglu, Dahleh, Lobel, and Ozdaglar (2011)'s **expanding observations** property for deterministic network
- Mild and clearly necessary: else for some $C<\infty$, infinitely many i cannot access the signal of any $j>C$ except their own
- Rest of paper studies speed of learning by comparing $(r_i)_{i\geq 1}$ on different networks

The Maximal Generations Network

- $K \geq 1$ agents per generation
- Agents in gen t observe all agents in gen $t-1$

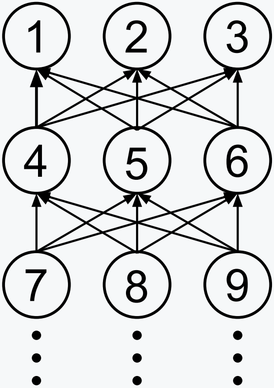

Proposition 4

In the maximal generations network:

- *Society learns completely in the long run with any K.*
- $r_i = i \cdot \frac{(2K-1)}{K^2} + o(i)$.
- *In the long run, social learning aggregates...*
  - *fewer signals per agent with larger K*
  - *fewer than 2 signals per generation with any K*
- *For any K and any agent i in generation $t \geq 3$, $r_i \leq K + 3t - 5$.*

Bounds on Signals Aggregated Per Generation

- Agents in generation t have observation paths of length $t-1$
- Can show in any network, this implies $r_i \geq t$
- Social learning must aggregate at least 1 signal per gen
- This lower-bound not too far from the actual learning rate:

$$r_i = i \cdot \frac{(2K-1)}{K^2} + o(i)$$

(No more than **2 signals** per gen in long-run, for any K)

$$r_i \leq K + 3t - 5, \qquad \text{for } t \geq 3$$

(No more than **3 signals** per gen **starting with gen 3**, for any K)

Slower Per-Agent Rate of Learning with Larger Gens

- If $K = 1$, every agent perfectly incorporates all past private signals $\Rightarrow$ fastest possible speed of social learning
- Slower learning with larger K holds numerically starting from agent $i = 16$ when comparing among $K \in \{2, 3, 4, 5\}$

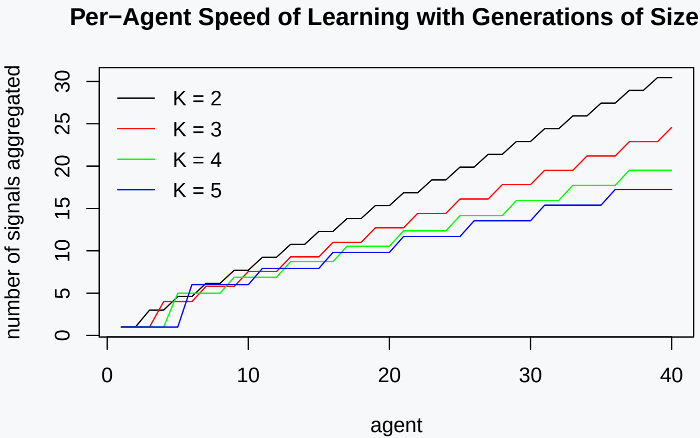

Slower Per-Agent Rate of Learning with Larger Gens

- Suppose we are interested in a "confidence threshold" metric of learning efficiency
  - e.g., how many agents does it take to become at least 90% confident in the true state more than 90% of the time?
- Societies with smaller generations reach any such threshold earlier when private signals are imprecise

Corollary 1

For every $0.5 < \bar{a} < 1$, $0 < \bar{p} < 1$, *and every pair of generation sizes* $K_1 < K_2$, *there exists a bound* $\tau > 0$ *on signal precision such that whenever* $0 < 1/\sigma^2 < \tau$, *the earliest* i *such that* $\mathbb{P}[a_i > \bar{a} \mid \omega = 1] > \bar{p}$ *is smaller with* K_1 *agents per generation than with* K_2 *agents per generation in the maximal generations network.*

Quiz: Which Network Leads to Faster Learning?

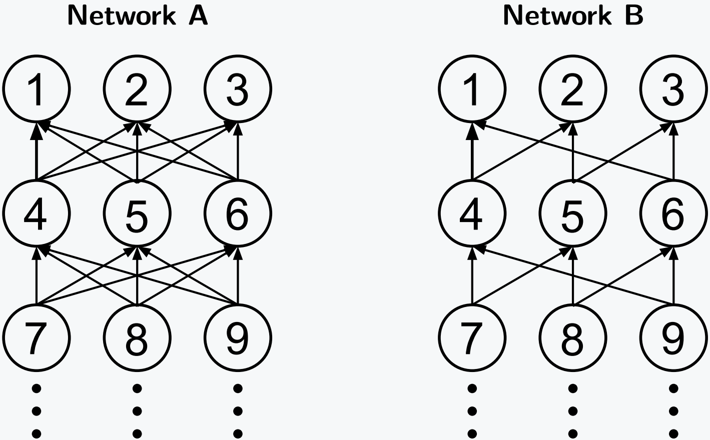

- **Network A** is the maximal generations network with $K = 3$
- **Network B** puts agents in each gen into 3 slots, $k \in \{1, 2, 3\}$. $k = 1$ sees 1 and 2, $k = 2$ sees 2 and 3, $k = 3$ sees 3 and 1.
- Unclear as **Network B** has fewer social obs, but also less info confounding. Need: rate of learning on more general networks.

Generations Network with Partial Observations

- Generations network with K agents per gen
- $\Psi_k \subseteq \{1, ..., K\}$, **inter-generational observation set**, define which gen $t-1$ slots are observed by a gen t agent in slot k
- Maximal generations network is the case of $\Psi_k = \{1, ..., K\}$

Definition

Consider the directed graph on $\{1, ..., K\}$ induced by $(\Psi_k)_k$ where $k_1 \to k_2$ if and only if $k_2 \in \Psi_{k_1}$. This induced graph is **aperiodic at *k*** if there exists at least one cycle that starts at k, and no integer larger than 1 divides the length of every cycle that starts at k. The observation sets are **weakly aperiodic** if every slot is either aperiodic or has an in-degree of 0.

Satisfied if $k \in \Psi_k$ for all k, for example.

Generations Network with Partial Observations

Definition

The observation sets are **strongly regular** if all agents observe $d \geq 1$ neighbors and all pairs of agents in the same generation share d_s common neighbors. That is, for all $i_1 \neq i_2$ in same generation $t \geq 2$, $|N(i_1)| = d$ and $|N(i_1) \cap N(i_2)| = d_s$.

A condition about the symmetry of $(\Psi_k)_k$. Adapts usual notion of strongly regular networks to generations networks.

For example, **Network B** from the "quiz" is strongly regular with $d = 2$, $d_s = 1$.

Speed of Learning with Partial Observations

Theorem 1

Suppose $(\Psi_k)_k$ are weakly aperiodic and strongly regular. Then

$$r_i = \left(1 + \frac{d^2 - d}{d^2 - d + d_s}\right)\frac{i}{K} + o(i).$$

- Exact asymptotic rate of learning for a broader class of generations networks
- Term in parenthesis increases in d and decreases in d_s — more obs speeds up rate of learning per gen but more confounding slows it down, all else equal
- Maximal gen network has the worst rate of learning, among all weakly aperiodic, strongly regular gen networks with same d
  - Because actions very confounded in maximal gen network
- But Theorem 1 shows asymptotic bound of 2 signals per gen applies to **all** such networks, strengthening Prop 4

Quiz: Which Network Leads to Faster Learning?

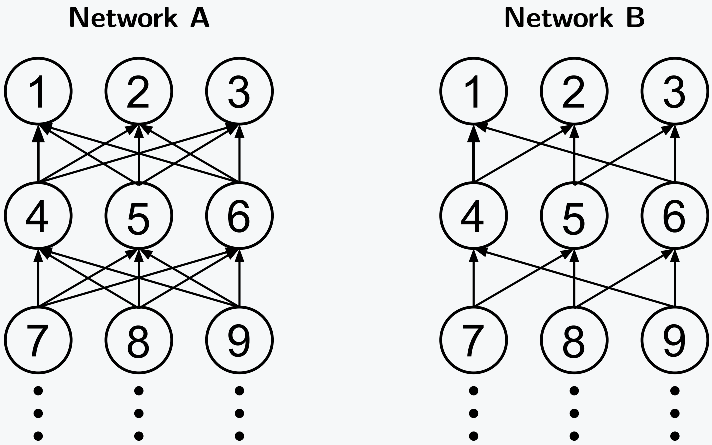

- Applying Theorem 1, rates of learning are the same in **Network A** ($d = 3, d_s = 3$) and **Network B** ($d = 2, d_s = 1$)!
- Extra social obs exactly cancel out reduced info content of each obs

Social Planner's Benchmark

Definition

$(\Psi_k)_k$ are **strongly connected** if for every $1 \leq k_1 \leq k_2 \leq K$, there exist t_1, t_2 so that $t_1 K + k_1$ is connected to $t_2 K + k_2$ in M.

Proposition 5

Suppose $(\Psi_k)_k$ are strongly connected and aperiodic. There is a log-linear strategy profile such that, for every $K_0 < K$, there exists a corresponding T so that for all $t \geq T$ and $1 \leq k \leq K$, the action of agent $(t-1)K + k$ is more accurate[1] *than $(t-1)K_0$ signals.*

- Slow learning of Thm 1 not intrinsic limitation of gen networks
- For K large, individuals only manage to aggregate an unboundedly small fraction of their private signals in eqm
- But a social planner can aggregate close to all signals

[1] i's action **more accurate than r signals** if it is more likely to lean towards the correct state than the action of someone who observes r indep signals.

Thank you!

When Does Adding Links Improve Accuracy?

- Consider general M, not necessarily with generations structure
- For two networks M and $M^\bullet$, write $M^\bullet \geq M$ if $M^\bullet$ can be generated from M by adding links
- Proposition 3 implies adding links improves long-run outcome
- But comparison of **Network A** and **Network B** from the "quiz" cautions effect on finite-time accuracy can be subtle
- Key issue is **intransitivities**: sequences of links $i_n \to i_{n-1}, i_{n-1} \to i_{n-2}, ..., i_2 \to i_1$ such that $i_n \not\to i_1$

Proposition 6

Suppose $M^\bullet \geq M$ and both networks are transitive at i. Then r_i is weakly higher on $M^\bullet$ than on M.

New Intransitivities Can Harm Accuracy

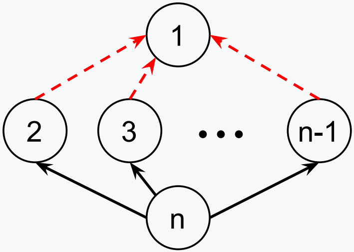

- M is the network with black links, transitive for n
- For $k \in \{2, ..., n-1\}$, adding $k-1$ red links from $2, ..., k$ to 1 creates a new network $M_k^\bullet$, not transitive for n
- In M, $r_n = n - 1$
- In $M_k^\bullet$, $r_n^{(k)} = 4 \cdot \frac{k-1}{k} + n - k$
- Adding 4 or more red links to M strictly harms n's welfare

Links Can Hurt Without Creating New Confounds

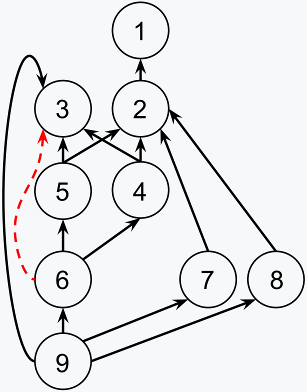

- If i already has intransitivities in M, adding links may decrease i's accuracy even without creating new intransitivities
- Adding red link strictly decreases r_9
- New link causes 6 to change eqm play in a way that creates negative info externality for 9
  - 6 puts more weight on $\tilde{a}_4$ and $\tilde{a}_5$, since she can now subtract off some confound using $\tilde{a}_3$
  - 6's action now contains worse over-counting $\tilde{s}_1$ and $\tilde{s}_2$, harms 9 who has a different signal-extraction problem than 6