Lecture 15: The Efficiency of Social Learning

Harry PEI
Department of Economics, Northwestern University

Spring Quarter, 2022

Overview

The social learning results we have seen so far:

- focus on asymptotic outcomes,

  i.e., whether agents take the correct action as $t \to +\infty$.

However, asymptotic efficiency does not mean much for social welfare.

- Asymptotic efficiency $\not\Rightarrow$ Agents take the correct action in finite time.
- Welfare losses can be large even under asymptotic efficiency.

Today: Rosenberg and Vieille (2019).

- The efficiency of social learning.

Social Learning with Binary States, Actions, and Signals

- Time $t = 0, 1, \ldots$
- State $\omega \in \{0, 1\}$, equally likely.
- Action of agent t: $a_t \in \{0, 1\}$. Payoff is $\mathbf{1}\{a_t = \omega\}$.
- Agent t observes s_t and $\{a_0, \ldots, a_{t-1}\}$ and then chooses a_t.
- Agent t's private signal $s_t \sim G(\cdot|\omega)$, conditionally independent.
- Agent t's private belief is $q_t \equiv \mathbb{E}[\omega|s_t]$.
  - $\hookrightarrow$ Conditional distribution of private beliefs is F_ω.
  - $\hookrightarrow$ For every ω, F_ω has continuous and positive density f_ω.
  - $\hookrightarrow$ Let $F \equiv \frac{1}{2}F_0 + \frac{1}{2}F_1$ be the distribution of private beliefs.
- The log likelihood ratio (LLR) of private belief is $l_t \equiv \log \frac{q_t}{1-q_t}$.
- The public belief is π_t, with log likelihood ratio $\Pi_t \equiv \log \frac{\pi_t}{1-\pi_t}$.

Benchmark: Asymptotic Efficiency

Theorem: Smith and Sorensen (2000)

$\lim_{t\to+\infty} \Pr(a_t = 1|\omega = 1) = 1$ *if and only if* $1 \in co\Big(supp(F_1)\Big)$.

$\lim_{t\to+\infty} \Pr(a_t = 0|\omega = 0) = 1$ *if and only if* $0 \in co\Big(supp(F_0)\Big)$.

Actions are asymptotically correct *iff* private signals are unbounded.

However, it can be the case that

- actions are asymptotically efficient,
- but in expectation, agents are wrong infinitely many periods.

  e.g., $\Pr(a_t \neq \omega) = \frac{1}{t}$ and $\{a_t \neq \omega\} \subset \{a_{t-1} \neq \omega\}$.

Standards for Efficiency & Welfare

Efficiency criteria while taking the rate of learning into account.

1. The expected value of $W_\delta \equiv (1-\delta)\sum_{t=1}^{+\infty} \delta^t \mathbf{1}\{a_t \neq \omega\}$ for some $\delta \in (0,1)$. This is hard. The results are not clean.
2. Whether the expected value of $W \equiv \sum_{t=0}^{+\infty} \mathbf{1}\{a_t \neq \omega\}$ is finite (efficient) or infinite (inefficient).
3. Whether the expected value of $\tau \equiv \inf\{t \geq 0 | a_t = \omega\}$ is finite (efficient) or infinite (inefficient).

Both 2 and 3 are stronger than asymptotic efficiency.

- If $\lim_{t\to\infty} \mathbb{E}[\mathbf{1}\{a_t \neq \omega\}] > 0$, then $\mathbb{E}[W]$ is infinite.
- If $\lim_{t\to\infty} \mathbb{E}[\mathbf{1}\{a_t \neq \omega\}] > 0$, then agents herd on an incorrect action starting from period 0 with positive prob.
- e.g., suppose $\Pr(a_t \neq \omega) = \frac{1}{t}$ and $\{a_t \neq \omega\} \subset \{a_{t-1} \neq \omega\}$,

  then actions are asymptotically efficient, but $\mathbb{E}[W] = \mathbb{E}[\tau] = +\infty$.

Standards for Efficiency & Welfare

Efficiency & welfare taking the rate of learning into account.

2. Whether the expected value of $W \equiv \sum_{t=0}^{+\infty} \mathbf{1}\{a_t \neq \omega\}$ is finite (efficient) or infinite (inefficient).
3. Whether the expected value of $\tau \equiv \inf\{t \geq 0 | a_t = \omega\}$ is finite (efficient) or infinite (inefficient).

$W \equiv \sum_{t=0}^{+\infty} \mathbf{1}\{a_t \neq \omega\} < +\infty \Rightarrow \tau \equiv \inf\{t \geq 0 | a_t = \omega\} < +\infty$.

- Why? Because $\tau \leq W$.

Main result: The two standards are equivalent.

- They also provide a characterization based on the primitives.

Benchmark: Efficiency When Signals are Public

Benchmark: Efficiency When Signals are Public

Suppose for every $t \in \mathbb{N}$, *agent* t *can observe* $\{s_0, \ldots, s_t\}$, *then* $\mathbb{E}[W] < +\infty$.

Proof: $\mathbb{E}[W] = \frac{1}{2}\sum_{t=0}^{+\infty} \Pr(a_t = 0|\omega = 1) + \frac{1}{2}\sum_{t=0}^{+\infty} \Pr(a_t = 1|\omega = 0)$.

- Agent t takes action 1 only if $\sum_{i=0}^{t} l_i \geq 0$.
- For any $\lambda > 0$, $\Pr(a_t = 1|\omega = 0) = \Pr(\sum_{i=0}^{t} l_i \geq 0|\omega = 0)$

  $= \Pr\left(\exp(\lambda \sum_{i=0}^{t} l_i) \geq 1 \middle| \omega = 0\right) \leq \mathbb{E}\left[\exp(\lambda \sum_{i=0}^{t} l_i) \middle| \omega = 0\right]$

  $= \left(\mathbb{E}[e^{\lambda l_i}|\omega = 0]\right)^{t+1}$.
- Therefore, $\Pr(a_t = 1|\omega = 0) \leq \left(\inf_{\lambda > 0} \mathbb{E}[e^{\lambda l_i}|\omega = 0]\right)^{t+1}$.
- We only need to show that $\inf_{\lambda > 0} \mathbb{E}[e^{\lambda l_i}|\omega = 0] < 1$.

Benchmark: Efficiency When Signals are Public

We only need to show that $\inf_{\lambda>0} \mathbb{E}[e^{\lambda l_i} | \omega = 0] < 1.$

- Recall that q_t is agent t's private belief, i.e., $l_t = \log \frac{q_t}{1-q_t}$, and $f_\omega(\cdot)$ is the pdf of distribution over private beliefs conditional on ω.
- Bayes Rule: For almost all $q \in [0,1]$, we have $\frac{f_1(q)}{f_0(q)} = \frac{q}{1-q}$.
- $\mathbb{E}[e^{\lambda l_i} | \omega = 0] = \mathbb{E}\left[\exp(\lambda \log \frac{q_i}{1-q_i}) \middle| \omega = 0\right] = \mathbb{E}\left[q_i^\lambda (1-q_i)^{-\lambda} \middle| \omega = 0\right]$

  $= \mathbb{E}\left[f_1(q)^\lambda f_0(q)^{-\lambda} \middle| \omega = 0\right] = \int_0^1 f_1(q)^\lambda f_0(q)^{1-\lambda} dq.$
- The value of $\int_0^1 f_1(q)^\lambda f_0(q)^{1-\lambda} dq$ is 1 when $\lambda = 0, 1$.

  $\mathbb{E}[e^{\lambda l_i} | \omega = 0]$ is a strictly convex function of λ since $\lambda \to e^\lambda$ is convex.
- Therefore, $\int_0^1 f_1(q)^\lambda f_0(q)^{1-\lambda} dq < 1$ when $\lambda \in (0,1)$

Main Result

Recall that $W \equiv \sum_{t=0}^{+\infty} \mathbf{1}\{a_t \neq \omega\}$ and $\tau \equiv \inf\{t \geq 0 | a_t = \omega\}$.

Theorem: Rosenberg and Vieille (2019)

The following three statements are equivalent:

1. $\mathbb{E}[W|\omega = 0] < +\infty$.
2. $\mathbb{E}[\tau|\omega = 0] < +\infty$.
3. $\int_0^1 \frac{1}{F(q)} dq < +\infty$.

Similarly, the following three statements are equivalent:

1. $\mathbb{E}[W|\omega = 1] < +\infty$.
2. $\mathbb{E}[\tau|\omega = 1] < +\infty$.
3. $\int_0^1 \frac{1}{1-F(q)} dq < +\infty$.

Implications

Since $\mathbb{E}[W] < +\infty$ if and only if $\mathbb{E}[W|\omega = 0] < +\infty$ and $\mathbb{E}[W|\omega = 1] < +\infty$,

- $\mathbb{E}[W] < +\infty$ if and only if $\mathbb{E}[\tau] < +\infty$
- $\mathbb{E}[\tau] < +\infty$ if and only if $\int_0^1 \frac{1}{F(q)} dq < +\infty$ and $\int_0^1 \frac{1}{1-F(q)} dq < +\infty$.

My view of their main contribution:

- Characterize efficiency via two inequalities:
  $\int_0^1 \frac{1}{F(q)} dq < +\infty$ and $\int_0^1 \frac{1}{1-F(q)} dq < +\infty$.
- These inequalities are satisfied *only if* private beliefs are unbounded.
- But unbounded private belief is *not sufficient*.
- Example: $s \sim N(-1,1)$ if $\omega = 0$ and $s \sim N(1,1)$ if $\omega = 1$.

  One can check that private beliefs are unbounded.

  However, $F(q) = \Pr(s \leq \log \frac{q}{1-q}) = \Phi(\log \frac{q}{1-q})$ where Φ is the cdf for $N(0,2)$, and we know that $\int_0^1 \frac{1}{\Phi(\log \frac{q}{1-q})} dq = +\infty$.

Implications

Since $\mathbb{E}[W] < +\infty$ if and only if $\mathbb{E}[W|\omega = 0] < +\infty$ and $\mathbb{E}[W|\omega = 1] < +\infty$,

- $\mathbb{E}[W] < +\infty$ if and only if $\mathbb{E}[\tau] < +\infty$
- $\mathbb{E}[\tau] < +\infty$ if and only if $\int_0^1 \frac{1}{F(q)} dq < +\infty$ and $\int_0^1 \frac{1}{1-F(q)} dq < +\infty$.

My issue with their characterization result:

- What is the *economic* interpretation of $\int_0^1 \frac{1}{F(q)} dq < +\infty$ and $\int_0^1 \frac{1}{1-F(q)} dq < +\infty$?

Implications

Since $\mathbb{E}[W] < +\infty$ if and only if $\mathbb{E}[W|\omega = 0] < +\infty$ and $\mathbb{E}[W|\omega = 1] < +\infty$,

- $\mathbb{E}[W] < +\infty$ if and only if $\mathbb{E}[\tau] < +\infty$
- $\mathbb{E}[W] < +\infty$ if and only if $\int_0^1 \frac{1}{F(q)} dq < +\infty$ and $\int_0^1 \frac{1}{1-F(q)} dq < +\infty$.

They also show that $\int_0^1 \frac{1}{F(q)} dq < +\infty$ and $\int_0^1 \frac{1}{1-F(q)} dq < +\infty$ if and only if $\mathbb{E}[W] < +\infty$ in a model where agent t only observes s_t and a_{t-1}.

- In terms of their efficiency standards, observing one player's past action is equivalent to observing all players' past actions.

Intuition:

- The last player's action determines the sign of $\log \frac{\pi_t}{1-\pi_t}$.
- This together with calendar time provides a fairly good summary of the info contained in players' past actions.

Proof Sketch

Theorem: Rosenberg and Vieille (2019)

The following three statements are equivalent:

1. $\mathbb{E}[W|\omega = 0] < +\infty$.
2. $\mathbb{E}[\tau|\omega = 0] < +\infty$.
3. $\int_0^1 \frac{1}{F(q)} dq < +\infty$.

$1 \Rightarrow 2$ is straightforward.

- Why? Because $\tau \leq W$.

In what follows, we show that $2 \Rightarrow 1$ and $2 \Leftrightarrow 3$.

Sorensen's Overturning Principle

Recall that π_t is the public belief in period t.

Overturning Principle

For every $t \in \mathbb{N}$, if $a_t = 1$, then $\pi_{t+1} \geq 1/2$; if $a_t = 0$, then $\pi_{t+1} \leq 1/2$.

Proof: By the law of iterated expectations,

$$\pi_{t+1} = \mathbb{E}[\omega | a_0, \dots, a_t] = \mathbb{E}\left[\underbrace{\mathbb{E}[\omega | a_0, \dots, a_t, s_t]}_{\equiv p_t} \middle| a_0, \dots, a_t\right] = \mathbb{E}[p_t | a_0, \dots, a_t].$$

The conclusion follows since $a_t = 1$ iff $p_t \geq 1/2$ and $a_t = 0$ iff $p_t \leq 1/2$.

Choice is Never Wrong with Positive Probability

Focus on unbounded signals where $F(q) > 0$ for all $q > 0$

- otherwise, agents' private beliefs are bounded away from 0, so all three conditions are violated.

Let σ be the first period s.t. agents take the wrong action.

Choice is Never Wrong with Positive Probability

There exists $c > 0$ such that $\Pr(\sigma = +\infty | \omega = 0, \pi_0) \geq c$ for every $\pi_0 \leq 1/2$.

The proof uses the optional stopping theorem.

- σ is a stopping time w.r.t. $\{\mathscr{F}_t\}_{t\in\mathbb{N}}$ if $\{\sigma = t\} \in \mathscr{F}_t$ for every $t \in \mathbb{N}$.

Optional Stopping Theorem

Suppose $\{X_t\}_{t\in\mathbb{N}}$ is a supermartingale and σ is a stopping time, then the process $\{X_{\sigma\wedge t}\}_{t\in\mathbb{N}}$ is a supermartingale.

Choice is Never Wrong with Positive Probability

Let σ be the first period s.t. agents take the wrong action.

Since $X_t \equiv \frac{\pi_t}{1-\pi_t}$ is a supermartingale conditional on $\omega = 0$, and $\sigma + 1$ is a stopping time, $X_{\sigma+1\wedge t}$ is also a supermartingale conditional on ω_0, we have:

$$\mathbb{E}[X_{\sigma+1\wedge t}|\omega = 0] \leq X_1 = \frac{\pi_1}{1-\pi_1}.$$

- By definition $a_\sigma = 1$, so $\pi_{\sigma+1} \geq \frac{1}{2}$ and $X_{\sigma+1} \geq 1$ if $\sigma < +\infty$.
- Markov's inequality implies that $\Pr(\sigma < +\infty|\omega = 0) \leq \frac{\pi_1}{1-\pi_1}$.
- When $\pi_0 \leq 1/2$, there exists $\alpha > 0$ s.t. $\pi_1 < 1/2 - \alpha$ when $a_0 = 0$.
- Hence $\Pr(\sigma < +\infty|\omega = 0, \pi_0)$ is bounded away from 1.

Proof Sketch for $2 \Rightarrow 1$

Theorem: Rosenberg and Vieille (2019)

If $\mathbb{E}[\tau|\omega = 0] < +\infty$, *then* $\mathbb{E}[W|\omega = 0] < +\infty$.

What do agents' actions look like?

- $1, 1, ..., 1, 0, 0, ...0, 1, 1, ..., 1, 0, ...$
- $0, 0, ...0, 1, 1, ..., 1, 0, 0, ..., 0, 1, 0...$

A sequence of wrong actions followed by a sequence of correct actions...

$\mathbb{E}[W]$ = expected number of wrong sequences $\times$ expected length of each wrong sequence.

- How many wrong sequences are there?
- How long is each wrong sequence?

Proof Sketch for $2 \Rightarrow 1$

How many wrong sequences are there?

Choice is Never Wrong with Positive Probability

Let σ be the first period s.t. agents make wrong choices. There exists $c > 0$ such that $\Pr(\sigma = +\infty | \omega = 0, \pi_0) \geq c$ for every $\pi_0 \leq 1/2$.

Overturning Principle

For every $t \in \mathbb{N}$, if $a_t = 1$, then $\pi_{t+1} \geq 1/2$; if $a_t = 0$, then $\pi_{t+1} \leq 1/2$.

If $a_t = 0$, then $\pi_{t+1} \leq 1/2$, so with prob $c > 0$ the sequence of correct actions starting from period $t + 1$ never ends.

- The prob that the kth wrong sequence happens $\leq (1 - c)^{k-1}$.

  Hence, the expected number of wrong sequences is finite.

Proof Sketch for 2 $\Rightarrow$ 1

How long is each wrong sequence?

Since we have assumed that $\mathbb{E}[\tau|\omega = 0] < +\infty$,

- The expected length of each wrong sequence is finite.
- W = expected number of wrong sequences $\times$ their expected length.

  We know that $\mathbb{E}[W|\omega = 0]$ is also finite.

Tricky part: $\mathbb{E}[\tau|\omega = 0] < +\infty$ only implies that the expected length is finite under the prior belief.

- **Observation**: When a wrong sequence starts, the public belief $\leq 1/2$.
- The authors show a monotonicity result, that if $\mathbb{E}[\tau|\omega = 0] < +\infty$ under prior $1/2$, then $\mathbb{E}[\tau|\omega = 0] < +\infty$ under any prior less than $1/2$.

Proof Sketch for 2 ⇔ 3

Theorem: Rosenberg and Vieille (2019)

$\mathbb{E}[\tau|\omega = 0] < +\infty$ *if and only if* $\int_0^1 \frac{1}{F(q)} dq < +\infty$.

Recall that τ is the first time agents take the correct action, so

Summation by parts:

$$\mathbb{E}[\tau|\omega = 0] = \sum_{t=1}^{+\infty} \Pr(\tau \geq t|\omega = 0).$$

Since $a_t = 1$ is the wrong action when $\omega = 0$, we have

$$\Pr(\tau \geq t|\omega = 0) = \Pr(a_1 = ... = a_{t-1} = 1|\omega = 0).$$

Proof Sketch for $2 \Leftrightarrow 3$

Recall that τ is the first time that agents take the correct action, so

$$\Pr(\tau \geq t | \omega = 0) = \Pr(a_1 = \ldots = a_{t-1} = 1 | \omega = 0).$$

Let π_t^* be the public belief when $a_1 = \ldots = a_{t-1} = 1$. Bayes rule:

$$\Pr(a_1 = \ldots = a_{t-1} = 1 | \omega = 0) = \Pr(a_1 = \ldots = a_{t-1} = 1 | \omega = 1) \cdot \frac{1 - \pi_t^*}{\pi_t^*}.$$

The RHS is bounded above and below by some linear function of $1 - \pi_t^*$.

- Since agents are never wrong with positive prob, $\Pr(a_1 = \ldots = a_{t-1} = 1 | \omega = 1) > c > 0$, so

$$\Pr(a_1 = \ldots = a_{t-1} = 1 | \omega = 1) \cdot \frac{1 - \pi_t^*}{\pi_t^*} \geq c(1 - \pi_t^*).$$

- Since $\pi_t^* \geq \frac{1}{2}$ when $a_1 = \ldots = a_{t-1} = 1$, so

$$\Pr(a_1 = \ldots = a_{t-1} = 1 | \omega = 1) \cdot \frac{1 - \pi_t^*}{\pi_t^*} \leq 2(1 - \pi_t^*).$$

Proof Sketch for 2 $\Leftrightarrow$ 3

What we have done so far:

$$\mathbb{E}[\tau|\omega = 0] = \sum_{t=1}^{+\infty} \Pr(\tau \geq t|\omega = 0) = \sum_{t=1}^{+\infty} \Pr(a_1 = \ldots = a_{t-1} = 1|\omega = 0).$$

and

$$c(1-\pi^*) \leq \Pr(a_1 = \ldots = a_{t-1} = 1|\omega = 0) \leq 2(1-\pi_t^*).$$

Hence, $\sum_{t=1}^{+\infty} \Pr(a_1 = \ldots = a_{t-1} = 1|\omega = 0)$ is finite $\Leftrightarrow$ $\sum_{t=1}^{+\infty}(1-\pi_t^*)$ is finite.

Proof Sketch for $2 \Leftrightarrow 3$

When signals are unbounded, $\pi_t^* \to 1$ as $t \to +\infty$, since agents will take action 0 if they receive a sufficiently low signal.

- But how fast will π_t^* converge to 1?

If $a_1 = \ldots = a_{t-1} = 1$, then agent t takes action 1 iff his private belief $\geq 1 - \pi_t^*$, so the prob that he takes action 1 in state ω is $1 - F_\omega(1 - \pi_t^*)$.

Let $\Pi_t^* \equiv \log \frac{\pi_t^*}{1-\pi_t^*}$. We have

$$\Pi_{t+1}^* - \Pi_t^* = \log \frac{1 - F_1(1 - \pi_t^*)}{1 - F_0(1 - \pi_t^*)}.$$

As $\pi_t^* \to 1$,

$$\frac{F_1(1 - \pi_t^*)}{F_0(1 - \pi_t^*)} \to 0,$$

and given that $F = \frac{F_1 + F_0}{2} \approx \frac{F_0}{2}$, we have:

$$\Pi_{t+1}^* - \Pi_t^* = \log \frac{1 - F_1(1 - \pi_t^*)}{1 - F_0(1 - \pi_t^*)} \approx F_0(1 - \pi_t^*) \approx 2F(1 - \pi_t^*).$$

Proof Sketch for $2 \Leftrightarrow 3$

Recall that the object we need to bound is $\sum_{t=1}^{+\infty}(1-\pi_t^*)$.

Since $\Pi_t^* \equiv \log \frac{\pi_t^*}{1-\pi_t^*}$, we know that as $\pi_t^* \to 1$,

$$1-\pi_t^* = \frac{1}{1+e^{\Pi_t^*}} \approx e^{-\Pi_t^*}.$$

We approximate $\sum_{t=1}^{+\infty} e^{-\Pi_t^*}$ by

$$\int_0^{+\infty} e^{-\Pi^*(t)} dt$$

where $\Pi^*(0)=0$, $\Pi^*(t) = \log \frac{\pi^*(t)}{1-\pi^*(t)}$, and $d\Pi^*(t)/dt = 2F(1-\pi^*(t))$.

Changing variables, we obtain:

$$\int_0^{+\infty} e^{-\Pi^*(t)} dt = \int_0^1 \frac{1}{2F(q)} dq.$$

Kevin's Lecture: Steady State Learning Models

An old question: Why do people play equilibrium?

- Equilibrium is neither necessary nor sufficient for rationality.

Two schools of thoughts:

- Epistemic conditions for equilibrium.
- Equilibrium arise in the long run when players can learn.

Fudenberg and Levine (1993):

- A population of players, with some birth and death rate.
- Players are randomly matched to play some game.
- Each player observes what his opponents played in the past, and tries to infer the distribution over players' actions.

What are the properties of the steady state distribution over action profiles?

- When is this distribution a Nash equilibrium?