Lecture 13: Reputation & Social Learning

Harry PEI
Department of Economics, Northwestern University

Spring Quarter, 2023

From Rich Information to Limited Information

Most of the models we have seen make an informational richness assumption:

- Players can observe the entire history of actions (e.g., Fudenberg and Maskin 1986, Fudenberg and Levine 1989, Smith and Sorensen 2000).
- Players can observe the entire history of some informative signals (e.g., Fudenberg, Levine and Maskin 1994, Fudenberg and Levine 1992).

In practice, people have limited info about others' past behaviors, e.g.,

- People do not recall events in the distant past.
- People do not know the exact sequence of actions/signals.
- People do not know the history of their trading partners (e.g., Maghribi traders in Medieval Europe, online dating apps).
- People do not know who they are playing with.

Can the society sustain good outcomes?

Outline

1. Reputation models with limited memories.
   - Can cooperation be sustained in *all* equilibria?
   - How do cooperative equilibria look like?

2. Repeated games played by a large community of players.
   - Can cooperation be sustained in *some* equilibrium?
   - How do cooperative equilibria look like?

Model

- Time: $t = 0, 1, 2, \ldots$
- Long-lived player 1 (e.g., seller) with discount factor $\delta \in (0,1)$

  vs a sequence of short-lived player 2s (e.g., buyers).
- P1 chooses $a_t \in A$, $\mathrm{P2}_t$ chooses $b_t \in B$, with A, B finite.
- Stage game payoffs: $u_1(a_t, b_t)$ and $u_2(a_t, b_t)$.

- P1 has two types:
  1. with prob $\pi_0 \in (0,1)$, a commitment type who plays his optimal pure commitment action a^*, which I assume is unique.
  2. with prob $1 - \pi_0$, a rational type who maximizes

$$(1-\delta) \sum_{t=0}^{\infty} \delta^t u_1(a_t, b_t).$$

- P1 observes all the actions taken in the past $h_1^t \equiv (a_s, b_s)_{s \leq t-1}$.

Modeling Innovation: P2s learn from previous P2s

Instead of observing the entire history of P1's actions,

- Each P2 only observes a bounded subset of P1's past actions.
- But they can observe previous P2s' actions.

P2_t observes

- P1's actions in the last $K \in \mathbb{N}$ periods $a_{\max\{0,t-K\}}, \ldots, a_{t-1}$.
- P2's actions in the last $M \in \mathbb{N} \cup \{+\infty\}$ periods $b_{\max\{0,t-M\}}, \ldots, b_{t-1}$,

Importantly, K is finite and $M \geq 1$.

Solution concept:

- Nash equilibrium for common properties of all equilibria.
- Perfect Bayesian equilibrium for existence of particular equilibrium.

Motivation

Application: Informal markets in developing countries:

- Consumers have limited access to seller's past records.
- Consumers can learn from other consumers' choices.

Cai, Chen and Fang (09), Zhang (10), Cai, De Janvry and Sadoulet (15).

Question: Can social learning provide adequate reputational incentives?

- Can patient P1 secure his optimal commitment payoff in all equilibria?

Takeaway: Observational learning can discourage reputation building.

- Observational learning enables consumers to imitate their predecessors.
- Limited observation of the seller's past actions rationalizes imitation.

Connections to Smith and Sorensen (2000)

My model resembles a social learning model when $M = +\infty$.

- $P2_t$'s private signal is $(a_{t-K}, ..., a_{t-1})$.
- $P2_t$ can observe the actions of their predecessors $(b_0, ..., b_{t-1})$.

There are two main differences:

- Learning about an exogenous payoff-relevant state

  vs Learning about the endogenous actions of a long-lived player.
- Players learn the state as $t \to +\infty$

  vs Long-lived player receives a high discounted average payoff.

Preview: There is no bad herd and the consumers' actions are informative.

Assumptions on Stage-Game Payoffs

Generic Stage-Game Payoffs

Players' stage-game payoff functions u_1 and u_2 satisfy:

1. *P2 has a unique best reply to every $a \in A$.*
2. *P1 has a unique best reply to every $b \in B$.*
3. *P1 has a unique (pure) Stackelberg action, denoted by a^*.*
   That is, $\arg\max_{a \in A} u_1(a, BR_2(a)) = \{a^*\}$.

Monotone-Supermodular Payoffs (MSM)

There is a complete order on A, $\succ_A$, and a complete order on B, $\succ_B$

1. *$u_1(a,b)$ is strictly decreasing in a and is strictly increasing in b.*
2. *$u_2(a,b)$ has strictly increasing differences in (a,b).*
3. *The Stackelberg action a^* is not the lowest element of A.*

Leading Example: Product Choice Game

- Players' stage-game payoffs:

–	Trust	Not Trust
High Effort	**1**, 1	$-c_N, 0$
Low Effort	$1 + c_T, -1$	**0**, 0

with $c_T, c_N > 0$.

- The commitment type plays H in every period.
- P1's optimal commitment payoff is 1 and his minmax payoff is 0.

Reputation Failure Result

P2_t observes:

- P1's actions in the last $K \in \mathbb{N}$ periods $a_{\max\{0,t-K\}},\ldots,a_{t-1}$.
- P2's actions in the last $M \in \mathbb{N} \cup \{+\infty\}$ periods $b_{\max\{0,t-M\}},\ldots,b_{t-1}$.

where K is finite and $M \geq 1$.

Let a^* be the optimal commitment action and $b^* \equiv \text{BR}_2(a^*)$.

Let a' be the lowest action and $b' \equiv \text{BR}_2(a')$.

Reputation Failure Theorem

For every $K \in \mathbb{N}$, there exist $\overline{\pi}_0 \in (0,1)$ and $\underline{\delta} \in (0,1)$ such that such that for every $\pi_0 < \overline{\pi}_0$ and $\delta > \underline{\delta}$,

there exists an equilibrium s.t. rational type P1 receives payoff $u_1(a',b')$.

Comparison with Fudenberg and Levine (1989)

P2_t observes:

- P1's actions in the last $K \in \mathbb{N}$ periods $a_{\max\{0,t-K\}}, \ldots, a_{t-1}$.
- P2's actions in the last $M \in \mathbb{N} \cup \{+\infty\}$ periods $b_{\max\{0,t-M\}}, \ldots, b_{t-1}$.

where K is finite and $M \geq 1$.

Reputation Failure Theorem

For every $K \in \mathbb{N}$, there exist $\overline{\pi}_0 \in (0,1)$ and $\underline{\delta} \in (0,1)$ such that such that for every $\pi_0 < \overline{\pi}_0$ and $\delta > \underline{\delta}$,

there exists an equilibrium s.t. rational type P1 receives payoff $u_1(a', b')$.

Fudenberg and Levine (1989): $K = +\infty$ and M can be arbitrary.

- Suppose P2 observes all of P1's past actions, then for every $\pi_0 \in (0,1)$, P1's payoff is at least $u_1(a^*, b^*)$ in all equilibria as $\delta \to 1$.
- P1 can guarantee payoff $u_1(a^*, b^*)$ by imitating the commitment type.

The Role of $M \geq 1$

Counterexample when $M = 0$:

- If payoffs are strictly *supermodular*, i.e., $0 < c_T < c_N$, then P1 can secure a payoff strictly greater than 0 in all equilibria.

–	Trust	Not Trust
High Effort	$1, 1$	$-c_N, 0$
Low Effort	$1 + c_T, -1$	$0, 0$

Proposition: Social Learning is Essential for Reputation Failure

Suppose $0 < c_T < c_N$ *and* $(K, M) = (1, 0)$.

For every $\pi_0 > 0$, *there exists* $\underline{\delta} \in (0, 1)$, *such that when* $\delta > \underline{\delta}$,

the seller's payoff is at least $\delta - (1 - \delta)c_N$ *in every PBE.*

We will come back to this next week.

Proof of Theorem 1: Constructing Imitation Equilibria

–	Trust	Not Trust
High Effort	$1,1$	$-c_N,0$
Low Effort	$1+c_T,-1$	$0,0$

Construct a PBE (*imitation equilibrium*), s.t.

- $P2_t$'s strategy depends only on (a_{t-1}, b_{t-1}).
- P1's period t action depends only on (a_{t-1}, b_{t-1}) and his *reputation*.

Strategy profile:

- When $(a_{t-1}, b_{t-1}) = \varnothing$ or (L, N), P2 plays N and the rational type P1 mixes s.t. the unconditional prob of H is $1/2$.
- When $(a_{t-1}, b_{t-1}) = (H, N)$, P2 plays T with prob $\frac{1-\delta}{\delta} c_N$ and the rational type P1 mixes s.t. the unconditional prob of H is $1/2$.
- When $(a_{t-1}, b_{t-1}) = (L, T)$, P2 plays T with prob $1 - \frac{1-\delta}{\delta} c_T$ and the rational type P1 mixes s.t. the unconditional prob of H is $1/2$.
- When $(a_{t-1}, b_{t-1}) = (H, T)$, P2 plays T for sure and P1 plays H.

Proof of Theorem 1: Constructing Imitation Equilibria

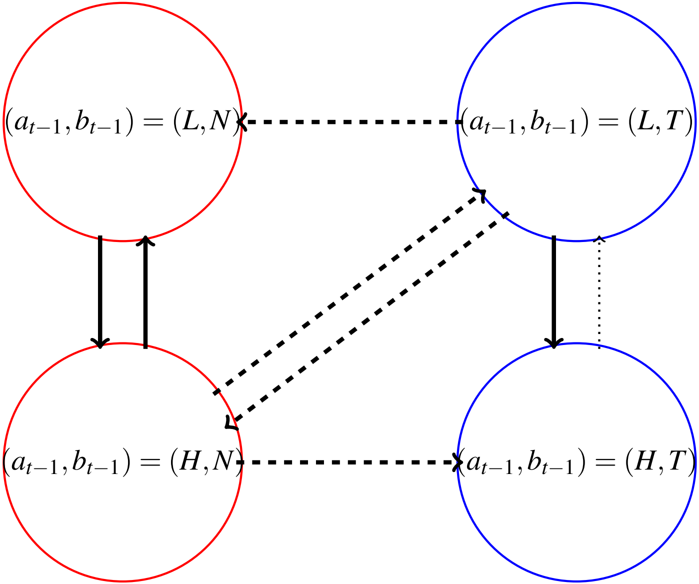

Interpretation of Imitation Equilibria

- When $(a_{t-1}, b_{t-1}) = \varnothing$ or (L, N), P2 plays N and the rational type P1 mixes s.t. the unconditional prob of H is $q \in (0, 1/2]$.
- When $(a_{t-1}, b_{t-1}) = (H, N)$, P2 plays T with prob $\frac{1-\delta}{\delta} c_N$ and the rational type P1 mixes s.t. the unconditional prob of H is $1/2$.
- When $(a_{t-1}, b_{t-1}) = (L, T)$, P2 plays T with prob $1 - \frac{1-\delta}{\delta} c_T$ and the rational type P1 mixes s.t. the unconditional prob of H is $1/2$.
- When $(a_{t-1}, b_{t-1}) = (H, T)$, P2 plays T for sure and P1 plays H.

Interpretation: Buyer plays N in period 0.

- Every buyer imitates the previous buyer with prob ≈ 1.
- When the seller plays H, buyers' actions are $(N, N, ..., N, T, T, ...)$
- The seller's undiscounted average payoff from playing H in every period is 1 but his discounted average payoff is 0.
- Buyers never herd on N and their actions are informative.

Driving Forces Behind Imitation Equilibria

Reputation failure is caused by imitation, i.e., choosing $b_t = b_{t-1}$.

- Imitation is feasible when $M \geq 1$.
- Imitation can be rationalized when K is finite no matter how large M is.

Imitation is *not* rational when $K = +\infty$.

- Let $\pi_t \equiv \Pr(\text{commitment type})$.
- If buyer t plays N but $a_t = H$, then $\pi_{t+1} \geq 2\pi_t$.
- If each buyer imitates her predecessor by playing N, the seller's reputation will exceed $1/2$, after which it is not rational to imitate.

Driving Forces Behind Imitation Equilibria

Why is imitation rational even when $M = +\infty$?

Let $\pi_t \equiv \Pr(\text{commitment type})$.

- If buyer t plays N but $a_t = H$, then $\pi_{t+1} \geq 2\pi_t$.
- Each buyer observes at most K of the seller's actions, after which the seller's reputation is multiplied by 2^K.
- Buyers may observe *all* previous buyers' actions, however,

$$\frac{\Pr\Big((b_0,...,b_{t-1}) = (N,N,...,N)\Big|\text{commitment type}\Big)}{\Pr\Big((b_0,...,b_{t-1}) = (N,N,...,N)\Big|\text{rational type}\Big)} \leq 1.$$

A long sequence of N encourage the buyers to play N.

Differences between $K = +\infty$ and $M = +\infty$

When $M = +\infty$, P2 observes an unbounded sequence of informative signals about P1's actions.

Why can't we apply Fudenberg and Levine (1992)'s result and show that P1 can secure his commitment payoff?

Connections with Canonical Reputation Results

- Gossner (2011) shows that in canonical reputation models, for every equilibrium strategy profile (σ_1, σ_2), we have

$$\mathbb{E}^{(a^*,\sigma_2)}\left[\sum_{t=0}^{\infty} d\Big(y_t(\cdot|a^*)\Big\|y_t(\cdot)\Big)\right] \leq -\log \pi_0,$$

  where y_t is the distribution of $P2_t$'s signal about a_t.

  The above inequality applies when K is finite and $M = +\infty$ by taking y_t as the distribution of $(b_{t+1}, ..., b_{t+K})$.

- In canonical reputation models, $d\big(y_t(\cdot|a^*)\big\|y_t(\cdot)\big)$ is bounded above 0 when $P2_t$ does not have a strict incentive to play b^*.
- In the constructive proof of Theorem 1, $d\big(y_t(\cdot|a^*)\big\|y_t(\cdot)\big) > 0$ when $P2_t$ does not have a strict incentive to play b^*.

  However, $d\big(y_t(\cdot|a^*)\big\|y_t(\cdot)\big)$ vanishes to 0 as $\delta \to 1$.

What about other equilibria?

Theorem 1: Reputation fails in one equilibrium.

How general is this lesson? What about other equilibria?

- For example, grim-trigger equilibria.

I proceed in two steps in order to address this question:

1. I propose a refinement that selects low-payoff equilibria,
2. I examine the common properties of all equilibria.

For simplicity, I will state all results in the product choice game.

–	**Trust**	**Not Trust**
High Effort	$1, 1$	$-c_N, 0$
Low Effort	$1 + c_T, -1$	$0, 0$

Refinement

Equilibrium refinement:

1. **Punishment after bad outcome:** $b_t = N$ if $(a_{t-1}, b_{t-1}) = (L, N)$.
2. **No initial trust:** $b_0 = N$.

Proposition: Equilibrium Refinement

For every $(\delta, \pi_0) \in (0,1)^2$, player 1's payoff equals 0 in every PBE that satisfies punishment after bad outcome and no initial trust.

Why? $b_0 = N$ implies that L is P1's best reply in period 0.

- Hence, period 1 history is (L, N) with positive prob, after which P2 plays N in period 1.
- This implies that L is P1's best reply in period 1...

  ... ...
- Iteratively apply this argument, we know that playing L in every period is P1's best reply, from which his payoff is 0.

Refinement

Equilibrium refinement:

1. **Punishment after bad outcome:** $b_t = N$ if $(a_{t-1}, b_{t-1}) = (L, N)$.
2. **No initial trust:** $b_0 = N$.

Compare imitation equilibria to grim-trigger equilibria:

- Punishment after bad outcome is satisfied by both equilibria.
- The key difference is the prob that P1 is trusted in period 0.

Proposition

In every equilibrium that satisfies punishment after bad outcome and P2 plays T with prob ε in period 0, P1's payoff is no more than

$$\varepsilon\Big\{\gamma(\pi_0,\delta)(1+c_T)+(1-\gamma(\pi_0,\delta))\Big\},$$

where $\gamma(\pi_0,\delta) \to 0$ as $\delta \to 1$ and $\pi_0 \to 0$.

Properties of All Nash Equilibria

P2s herd on $b \in B$ at h^t if they play b for every $h^s \succeq h^t$.

Proposition: No Bad Herd

For every $(\delta, \pi_0) \in (0,1)^2$ and in every Nash equilibrium, P2s cannot herd on N at any h^t that satisfies $\pi(h^t) > 0$.

Proof idea:

- If P2s herd on N at h^t, then the rational type P1 has no intertemporal incentive and hence, strictly prefers L.
- Upon observing H, P2 knows that P1 is committed and has a strict incentive to play T. This leads to a contradiction.

Properties of All Nash Equilibria

Proposition: High Asymptotic Payoff

Suppose $M = +\infty$. For every $(\delta, \pi_0) \in (0,1)^2$ and in every NE,

$$\liminf_{t\to\infty} \frac{1}{t}\mathbb{E}^{(H,\sigma_2)}\left[\sum_{s=0}^{t-1} u_1(a_s, b_s)\right] \geq \frac{K}{K+1} - \frac{1}{K+1}c_N.$$

Rough intuition: If $a_t = H$ for every $t \in \mathbb{N}$, then for every $t \in \mathbb{N}$,

- *either* P2's future actions are informative about a_t,

  or rational P1 plays L in equilibrium, so the next K P2s believe that P1 is committed after observing H.

Our Theorem: P1 receives his minmax payoff 0 in some equilibria.

$\Rightarrow$ P1's payoff from imitating commitment type $\leq$ his minmax payoff.

When $\frac{K}{K+1} - \frac{1}{K+1}c_N > 0$, the only plausible explanation for our theorem:

- If P1 plays H, then he can eventually guarantee a positive asymptotic payoff, but the rate at which play reaches this phase goes to 0 as $\delta \to 1$.

Reputation Result with Contemporaneous Information

For this part, suppose $P2_t$ observes

- all previous P2's actions $b_0, \ldots, b_{t-1}$
- P1's actions in the last K periods $a_{\max\{0,t-K\}}, \ldots, a_{t-1}$
- a private signal $s_t \sim f(\cdot|a_t) \in \Delta(S)$ of P1's current action.

  Let $\mathbf{f} \equiv \{f(\cdot|a)\}_{a \in A}$, with $f(s|a)$ the prob of s under a.

  We also assume that S is a countable set.

Unbounded Informativeness

$\mathbf{f}$ *is unboundedly informative about* a^* *if for every* $M > 0$*, there exists* $s \in S$*, such that* $f(s|a^*) > M \max_{a \neq a^*} f(s|a)$.

- When S is finite, unbounded informativeness requires the existence of $s^* \in S$ such that $f(s^*|a) > 0$ if and only if $a = a^*$.
- When S is infinite, allows $f(\cdot|a)$ to have full support for all a.

Monotone Likelihood Ratio Property (MLRP)

For this part, suppose $P2_t$ observes

- all previous P2's actions $b_0, \ldots, b_{t-1}$
- P1's actions in the last K periods $a_{\max\{0,t-K\}}, \ldots, a_{t-1}$
- a private signal $s_t \sim f(\cdot|a_t) \in \Delta(S)$ of P1's current action.

  Let $\mathbf{f} \equiv \{f(\cdot|a)\}_{a \in A}$, with $f(s|a)$ the prob of s under a.

  We also assume that S is a countable set.

Monotone Likelihood Ratio Property (MLRP)

$\mathbf{f} \equiv \{f(\cdot|a)\}_{a \in A}$ satisfies MLRP if there exists a complete order on S, $\succ_S$, such that $f(s|a)f(s'|a') \geq f(s|a')f(s'|a)$ for every $a \succ_A a'$ and $s \succ_S s'$.

If $|A| = 2$, then any $\mathbf{f}$ satisfies MLRP.

Result

f being unboundedly informative is sufficient and almost necessary.

Theorem: Reputation Effects & Unbounded Informativeness

Suppose **f** *satisfies MLRP.*

1. *If* **f** *is unboundedly informative about* a^*, *then for every* $\pi_0 > 0$ *and* $\varepsilon > 0$, *there exists* $\delta^* \in (0,1)$ *such that when* $\delta > \delta^*$, *player* 1*'s payoff is at least* $u_1(a^*,b^*) - \varepsilon$ *in every equilibrium.*
2. *If there exists* $\varepsilon > 0$ *such that* $f(s|a') \geq \varepsilon f(s|a^*)$ *for every* $s \in S$, *then for every* $K \in \mathbb{N}$, *there exists* $\overline{\pi}_0 \in (0,1)$ *such that for every* $\pi_0 < \overline{\pi}_0$ *and* $\delta > \underline{\delta}(u_1,u_2)$, *there exists an equilibrium in which player* 1*'s payoff is* $u_1(a',b')$.

Unbounded Informativeness Secures Stackelberg Payoff

Reminiscent of BHW and Smith-Sørensen: In social learning models,

- myopic players' actions are asymptotically efficient *if and only if* their private signals are unboundedly informative.

However, P2 asymptotically learns about P1's type is neither necessary nor sufficient for P1 to receive his Stackelberg payoff.

1. High asymptotic payoff $\not\Rightarrow$ High discounted average payoff.
2. P2 cares about P1's endogenous actions but not his type.

Proof Sketch of the Commitment Payoff Theorem

For now, let us focus on the case in which S is finite.

- There exists $s^* \in S$ such that $f(s^*|a) > 0$ iff $a = a^*$.

Suppose rational P1 deviates and plays a^* in every period.

1. Is b_t responsive to s_t? If not responsive, why?
2. If b_t is responsive to s_t, then is b_t informative about P1's type?
3. A remaining step that deals with P2's heterogenous private beliefs. (*I won't talk about it for the interest of time*)

Is b_t responsive to s_t? Not necessarily

Two reasons for why b_t is irresponsive to s_t:

1. Prior prob of a^* is too low, unwilling to play b^* no matter what.
2. Prior prob of a^* is too high, willing to play b^* no matter what.

When $\mathbf{f}$ is unboundedly informative,

- the first reason is ruled out, since P2 plays b^* after observing s^*.

  (a more sophisticated argument is needed when S is infinite)
- the second reason: P1's stage-game payoff is $u_1(a^*, b^*)$.

Suppose b_t is responsive to s_t, is b_t responsive to P1's type?

If $\mathbf{f}$ is unboundedly informative and $|A| = 2$, then

- b_t is responsive to $s_t \Rightarrow b_t$ is responsive to P1's type.

P2 is willing to play b^* after observing s *iff* $\frac{f(s|a^*)}{f(s|a')}$ is above some cutoff.

$\Rightarrow$ $\Pr(b_t = b^*|a^*) - \Pr(b_t = b^*|a') \geq 0.$

Since P2 plays b^* following s^* and $\mathbf{f}$ is unboundedly informative about a^*.

- $\exists$ $C > 0$ such that

$$\Pr(b_t = b^*|a^*) - \Pr(b_t = b^*|a') \geq C(1 - \Pr(b_t = b^*|a^*)).$$

As long as P2 plays b^* with prob less than $1 - \varepsilon$, the informativeness of b_t about P1's type is bounded above 0, which does not depend on δ.

Suppose b_t is responsive to s_t, is b_t responsive to P1's type?

If **f** is unboundedly informative and $|A| \geq 3$, then

- b_t is responsive to s_t $\not\Rightarrow$ b_t is responsive to P1's type.
- It could be the case that $\Pr(b_t = b^*) < 1$ while

$$\Pr(b_t = b^* | a_t = a^*) = \Pr(b_t = b^* | a_t \neq a^*).$$

Example:

-	b^*	b'
$\overline{a}$	$1,4$	$-2,0$
a^*	$2,1$	$-1,0$
$\underline{a}$	$3,-2$	$0,0$

Let $S \equiv \{\overline{s}, s^*, \underline{s}\}$, with $f(s^*|a^*) = 2/3, f(\underline{s}|a^*) = 1/3, f(\overline{s}|\overline{a}) = 1,$ $f(\overline{s}|\underline{a}) = 1/3$, and $f(\underline{s}|\underline{a}) = 2/3$.

- P2 plays b^* if $s_t \in \{s^*, \overline{s}\}$ and plays b' if $s_t = \underline{s}$.
- P1 plays $0.5 \circ a^* + 0.25 \circ \overline{a} + 0.25 \circ \underline{a}$.

Suppose b_t is responsive to s_t, is b_t responsive to P1's type?

If **f** is unboundedly informative and **f** satisfies MLRP, then

- b_t is responsive to s_t $\Rightarrow$ b_t is responsive to P1's type.

For every $\alpha \in \Delta(A)$ and $\beta : S \to \Delta(B)$, let

- $\gamma(\alpha, \beta) \in \Delta(B)$ be the distribution over b induced by (α, β)

Proposition: Social Learning is Essential for Reputation Failure

When **f** *is unboundedly informative about a^* and satisfies MLRP, there exists $C > 0$ such that for every $\alpha \in \Delta(A)$ with $a^* \in supp(\alpha)$, and every $\beta : S \to \Delta(B)$ that best replies against α. If $\gamma(\alpha, \beta)[b^*] < 1 - \varepsilon$, then*

$$d\Big(\gamma(\alpha, \beta) \Big\| \gamma(a^*, \beta)\Big) > C\varepsilon^2.$$

Connections with Canonical Reputation Results

Connections between my result and existing results:

- Gossner shows that

$$\mathbb{E}^{(a^*,\sigma_2)}\left[\sum_{t=0}^{\infty} d\Big(y_t(\cdot|a^*)\Big|\Big|y_t(\cdot)\Big)\right] \leq -\log \pi_0.$$

- In canonical reputation models, $d\big(y_t(\cdot|a^*)\big\|y_t(\cdot)\big)$ is bounded above 0 as long as P2 does not have a strict incentive to play b^*.
- In the model with contemporary information, there exists increasing function $g(\varepsilon)$ with $g(0) = 0$ s.t.

  $d\big(y_t(\cdot|a^*)\big\|y_t(\cdot)\big) \geq g(\varepsilon)$ when $\Pr(b_t = b^*) < 1 - \varepsilon$.
- Take any $\varepsilon > 0$, the expected number of periods s.t. P2 plays b^* with prob $< 1 - \varepsilon$ is finite and is independent of δ.

Conclusion

Patient player's incentive to build reputations when uninformed players

- have bounded observation of the patient player's past actions,
- can learn from other uninformed players' actions.

Reputation failure caused by the uninformed players' imitation.

- It takes a long time for the patient player to receive a high payoff.
- Imitation limits actions' informativeness and slows down learning.

When $M = +\infty$, returns from reputation hinges on the speed of learning.

- Without contemporaneous info, social learning is too slow.
- With unboundedly informative contemporaneous signals, the speed of social learning is bounded above zero.

Related Literature

1. Social learning and imitation: Banerjee (92), Bikhchandani, et al (92), Smith and Sørensen (00), Gale and Kariv (03), Hann-Caruthers, et al (18), Harel, et al (21), Kartik et al. (21).

   Difference: Speed of learning instead of asymptotic outcome.

2. Efficiency of social learning: Rosenberg and Vieille (19).

   Difference: Learn about an exogenous state *vs* endogenous actions.

3. Reputation effects: Fudenberg and Levine (89,92), Gossner (11).

   Difference: Informativeness of signal depends on δ.

4. Reputation with limited memory: Liu (11), Liu and Skrzypacz (14).

   Difference: No social learning in their models and no reputation result.

5. Bad reputation: Ely and Välimäki (03), Ely, Fudenberg and Levine (08)

   Difference: Learning cannot stop & high asymptotic payoffs.

Next Lecture

Reputation model with limited memory w/o social learning.

Critique of mixed-strategy equilibria in extensive-form games.