Lecture 14: Limited Memories & Purification

Harry PEI
Department of Economics, Northwestern University

Spring Quarter, 2023

Overview

Last lecture: Reputation and social learning.

- Each short-run player observes a bounded number of the long-run player's action, and at least one short-run player's action.
- Proof: A mixed-strategy equilibrium in which learning is slow.

This lecture:

- Reputation effects with limited memory w/o social learning.
- Critique of mixed-strategy equilibria in extensive-form games with limited memories.

Primitives

- Long-lived player 1 *vs* short-lived player 2s.
- Player 1's action $x \in [0,1]$. Player 2's action $y \in [0,1]$.

- Player 1's stage-game payoff $u_1(x,y)$ satisfies:
  1. $u_1(x,y)$ is strictly decreasing in x,
  2. $u_1(x,y)$ is strictly increasing in y,
  3. $u_1(x,y)$ has strictly decreasing differences in (x,y).

- P2's stage-game payoff $u_2(x,y)$ has strictly increasing differences,

  P2 has a unique best reply to any of P1's mixed actions.
  - P2's best reply $y^* : \Delta[0,1] \to [0,1]$ is continuous and strictly increasing under FOSD.
  - Let us normalize $y^*(0) = 0$.

- P1's dominant action is 0. Unique stage-game NE is $(0,0)$.

Long-Run Player's Type

Player 1 has two types:

- With prob $\mu^* \in (0,1)$, P1 is a commitment type who plays $c > 0$.
- With prob $1-\mu^*$, P1 is rational and maximizes discounted payoff.

We assume that commitment to c is valuable:

$$u_1(c, y^*(c)) > u_1(0,0).$$

We also assume that δ is large enough such that:

$$\underbrace{u_1(c, y^*(c))}_{\text{P1's commitment payoff from } c} > (1-\delta)\underbrace{u_1(0, y^*(c))}_{\text{P1's payoff from deviating to 0}} + \delta \underbrace{u_1(0,0)}_{\text{P1's stage-game NE payoff}}.$$

Monitoring Structure

$P2_t$ only observes P1's action in the last K periods.

- $(a_{\max\{0,t-K\}}, \ldots, a_{t-1})$.

Important things to emphasize:

- Player 2 cannot observe the actions of previous short-run players.
- Player 2 cannot observe calendar time.

No social learning.

Player 2's prior assigns prob $(1-\delta)\delta^t$ to calendar time being t, and updates her belief according to Bayes rule after observing $(a_{\max\{0,t-K\}}, \ldots, a_{t-1})$.

Strategies & Stationary Perfect Bayesian equilibrium

The set of P2's histories $\mathcal{H} \equiv \bigcup_{k=0}^{K} [0, 1]^k$.

- Sequences of player 1's actions with length no more than K.

P2's strategy $\sigma_2 : \mathcal{H} \to [0, 1]$.

Focus on *stationary equilibria* s.t. P1's strategy depends only on $\mathcal{H}$.

Sufficient Statistics Result

Partition the set of length K histories into $K+1$ subsets:

- $h \in \mathcal{H}^K$ if the last K actions were c. Call them clean histories.
- For every $h \in \mathcal{H} \backslash \mathcal{H}^K$, let $I(h) \in \{0, 1, ..., K-1\}$ be the number of periods since the most recent action that is not c.

  Intuitively, if $I(h)$ is larger, then h is closer to a clean history.

  Let $\mathcal{H}^k$ be the set of histories with $I(h) = k$, for $k \in \{0, 1, ..., K-1\}$.

Proposition: Sufficient Statistics Result

In any stationary equilibrium, (i) players' actions after period K are measurable with respect to the above partition, and (ii) player 1 will only play 0 and c with positive probability.

This depends on the submodularity of P1's payoff.

Proof Sketch: Sufficient Statistics Result

Consider two histories h and h' s.t.

- P1's actions were the same in the last $K-1$ periods,

  P1's action K periods ago was not c under h and h'.

Lemma

P2's actions at h and h' are the same.
P1's continuation values at h and h' are the same.

Suppose by way of contradiction that $a_2(h) > a_2(h') \geq 0$.

- Since h and h' are not clean, P1 is known to be rational.
- $a_2(h)$ depends only on $a_1(h)$, $a_2(h')$ depends only on $a_1(h')$.

Hence, there exist a_1^* and a_1' with $a_1^* > a_1'$ s.t.

- P1 plays a_1^* with positive prob at h.
- P1 plays a_1' with positive prob at h'.

Proof Sketch: Sufficient Statistics Result

Suppose by way of contradiction that $a_2(h) > a_2(h') \geq 0$.

P1 plays a_1^* at h, and plays $a_1'(< a_1^*)$ at h'.

P1's incentive constraint at h:

$$(1-\delta)u_1(a_1^*, a_2(h)) + \delta V(h, a_1^*) \geq (1-\delta)u_1(a_1', a_2(h)) + \delta V(h, a_1').$$

P1's incentive constraint at h':

$$(1-\delta)u_1(a_1', a_2(h')) + \delta V(h', a_1') \geq (1-\delta)u_1(a_1^*, a_2(h')) + \delta V(h', a_1^*).$$

Since h and h' differ only in action K periods ago, we have $V(h, a_1^*) = V(h', a_1^*)$ and $V(h, a_1') = V(h', a_1')$. Hence,

$$u_1(a_1^*, a_2(h)) - u_1(a_1', a_2(h)) \geq \frac{\delta}{1-\delta}(V(h, a_1') - V(h, a_1^*))$$

$$= \frac{\delta}{1-\delta}(V(h', a_1') - V(h', a_1^*)) \geq u_1(a_1^*, a_2(h')) - u_1(a_1', a_2(h'))$$

Proof Sketch: Sufficient Statistics Result

Therefore, $a_2(h) > a_2(h')$, $a_1^* > a_1'$, but

$$u_1(a_1^*, a_2(h)) - u_1(a_1', a_2(h)) \geq u_1(a_1^*, a_2(h')) - u_1(a_1', a_2(h')).$$

This leads to a contradiction since u_1 has strictly decreasing differences.

Hence, P2's actions at h and h' must be the same.

P1's continuation values at h and h' must be the same as well.

- Why? P2's actions are the same, and P1's continuation value in the next period does not depend on whether the current history is h or h'.

Proof Sketch: Sufficient Statistics Result

We have just shown that at any two histories h and h' s.t.

- P1's actions were the same in the last $K - 1$ periods,

  P1's action K periods ago was not c under h and h'.

then

- P2's actions are the same at h and h'.
- P1's continuation values at the same at h and h'.

What about h and h' s.t.

- P1's actions were the same in the last $K - 2$ periods,

  P1's action $K - 1$ periods ago was not c under h and h'?

What can we say about $V(h, a_1)$ and $V(h', a_1)$?

Proof Sketch: Sufficient Statistics Result

Therefore, for any $k \in \{1, 2..., K\}$, if h and h' are such that:

- P1's actions were the same in the last $k-1$ periods,

  P1's action k periods ago was not c under h and h'.

At h and h', P1's continuation values and P2's actions are the same.

Recall the definition of $I(h)$.

- If $I(h) = I(h')$, then either $h = h'$ or

  there exists $k \in \{1, 2, ..., K\}$ such that P1's actions were the same in the last $k-1$ periods and their actions k periods ago were not c.

Hence, P2's action and P1's continuation value depend only on $I(h)$.

- This implies that P1 plays either 0 or c.
- Hence, P1's mixed actions can be ranked according to FOSD.
- If $a_2(h) = a_2(h')$ while h and h' are not clean,

  then P1's action at h and h' must be the same.

Characterize Stationary PBE

- μ_K: Prob P2's belief assigns to the commitment type at clean histories.
- β_k: Prob that rational-type P1 plays c at $\mathcal{H}^k$, for every $k \in \{0, 1, ..., K\}$.
- y_k: P2's action at $\mathcal{H}^k$, for every $k \in \{0, 1, ..., K\}$.

Theorem 1

Under any $\delta > \overline{\delta}$, memory length $K \in \mathbb{N}$, and prior belief $\mu^* > 0$. Every stationary PBE takes the following form:

- *The rational type P1 plays 0 for sure at clean histories.*
- *Prob of c increases with the index $0 < \beta_0 < \beta_1 < ... < \beta_{K-1} < \mu_K$.*
- *Trust increases with the index $y_0 < y_1 < ... < y_{K-1} < y_K$,*

Long-Run Player's Equilibrium Actions (in red)

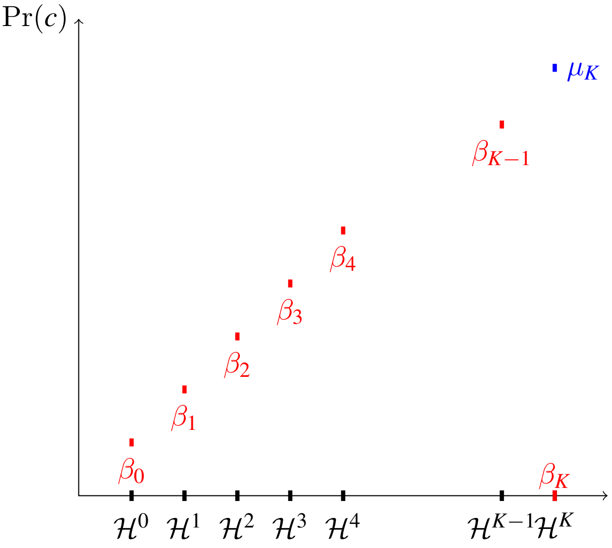

The Short-Run Players' Actions and Beliefs

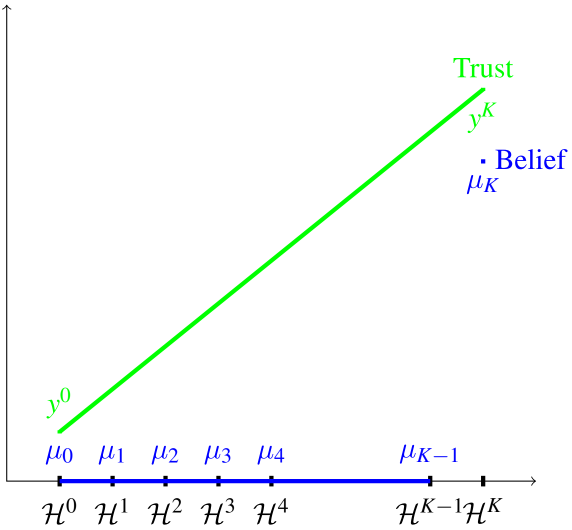

Characterize Stationary PBE

Takeaway from their characterization result:

1. P2's trust increases with the index $I(h)$.
2. P1 betrays for sure at clean histories.
3. At non-clean histories, P1's prob of playing c increases with $I(h)$.

Intuition: $u_1(a, b)$ has strictly decreasing differences.

- P1 has stronger incentive to betray when he is trusted more.

Caveat: Supermodularity/submodularity in the stage game usually do not imply much in the repeated game.

- Why? My action today affects your observation tomorrow.

  e.g., in the proof of the sufficient statistics result, we only use submodularity at specific pairs of histories.

Lower Bound on P1's Equilibrium Payoff

Theorem 2

For any $\varepsilon > 0$ *and* $\mu^* \in (0,1)$, *there exists* $K(\varepsilon, \mu^*) \in \mathbb{N}$ *such that for any* $K > K(\varepsilon, \mu^*)$, *P1's payoff at any history of any stationary PBE is at least*

$$(1-\delta^K) g_1(c,0) + \delta^K g_1(c, y^*(c)) - \varepsilon.$$

When K is large enough,

- as long as P1 is sufficiently patient, he can secure his commitment payoff $g_1(c, y^*(c))$ in *all* stationary PBEs.

This stands in contrast to Fudenberg and Levine (1989,1992):

- No meaningful payoff lower bound that applies to all histories.

Proof Sketch

The proof is different from the one in Fudenberg and Levine:

- Fudenberg and Levine's technique requires P2 observing everything her predecessors observed.
- This is not the case when memories are bounded.

Key step of the proof:

Lemma

For any $\eta \in (0,1)$, $\exists\, \overline{K}(\eta)$ such that when $K > \overline{K}(\eta)$, $\mu(K) > 1-\eta$.

Intuition: Suppose $\mu(K) \leq 1-\eta$. Since $\beta(k) < \mu(K)$ for every $k \leq K-1$, the rational type reaches the clean history with prob no more than $(1-\eta)^K$.

- If K is large, then $(1-\eta)^K$ is small,

  which implies that P2's belief at clean histories assigns a high prob to the commitment type. This contradicts $\mu(K) \leq 1-\eta$.

Dynamic Games with Limited Memories

Some of the results rely on constructing mixed-strategy equilibria:

- Liu (2011), Liu and Skrzypacz (2014), Pei (2023).

Mixed-strategy equilibria are also used to show folk theorems in repeated games with private monitoring.

- Most notably, the literature on belief-free equilibria, e.g.,

  Ely and Välimäki (2002), Ely, Hörner and Olzewski (2005).

How robust are these mixed-strategy equilibria?

Background: Mixed Strategies & Harsanyi Purification

Battle of sexes game:

–	L	R
T	2, 1	0, 0
B	0, 0	1, 2

There exists one mixed-strategy Nash equilibrium:

- P1 chooses $\frac{2}{3}T + \frac{1}{3}B$. P2 chooses $\frac{1}{3}L + \frac{2}{3}R$.

How to pin down players' mixing probabilities?

- P1's mixing prob is pinned down by P2's indifference condition.
- P2's mixing prob is pinned down by P1's indifference condition.

This logic sounds weird.

Background: Mixed Strategies & Harsanyi Purification

Harsanyi: What if players privately observe private payoff shocks?

–	L	R
T	2, 1	0, 0
B	0, 0	1, 2

Suppose P1's payoff from T equals his payoff in the matrix plus ξ.

- ξ is drawn from a continuous distribution $U[0, \varepsilon]$.

  Only P1 observes the realization of ξ.

Suppose P2's payoff from L equals his payoff in the matrix plus η.

- η is drawn from a continuous distribution $U[0, \varepsilon]$.

  Only P2 observes the realization of η.

P1 chooses T when ξ is high and chooses B when ξ is low.

- From P2's perspective, he faces a distribution of P1's actions.

Background: Mixed Strategies & Harsanyi Purification

Harsanyi: What if players' observe some private payoff perturbations?

–	L	R
T	2, 1	0, 0
B	0, 0	1, 2

Suppose P1's payoff from T equals his payoff in the matrix plus ξ.

- ξ is drawn from a continuous distribution $U[0, \varepsilon]$.

  Only P1 observes the realization of ξ.

Suppose P2's payoff from L equals his payoff in the matrix plus η.

- η is drawn from a continuous distribution $U[0, \varepsilon]$.

  Only P2 observes the realization of η.

P2 chooses L when η is high and chooses R when η is low.

- From P1's perspective, he faces a distribution of P2's actions.

Background: Mixed Strategies & Harsanyi Purification

Let's solve for the Bayes Nash equilibrium in the perturbed game:

–	L	R
T	$2+\xi, 1+\eta$	$\xi, 0$
B	$0, \eta$	$1, 2$

Let $\xi^*, \eta^* \in [0, \varepsilon]$ be such that:

- P1 chooses *B iff* $\xi < \xi^* \Rightarrow$ P1 chooses *B* with prob $\frac{\xi^*}{\varepsilon}$.
- P2 chooses *R iff* $\eta < \eta^* \Rightarrow$ P2 chooses *R* with prob $\frac{\eta^*}{\varepsilon}$.

Type ξ^* player 1's indifference between *T* and *B* requires that

$$\left(1 - \frac{\eta^*}{\varepsilon}\right)(2 + \xi^*) = \frac{\eta^*}{\varepsilon}.$$

$$\Rightarrow \quad \varepsilon(2 + \xi^*) = \eta^*(3 + \xi^*).$$

Background: Mixed Strategies & Harsanyi Purification

Let's solve for the equilibrium in the perturbed game:

–	L	R
T	$2+\xi, 1+\eta$	$\xi, 0$
B	$0, \eta$	$1, 2$

Let $\xi^*, \eta^* \in [0, \varepsilon]$ be such that:

- P1 chooses B *iff* $\xi < \xi^* \Rightarrow$ P1 chooses B with prob $\frac{\xi^*}{\varepsilon}$.
- P2 chooses R *iff* $\eta < \eta^* \Rightarrow$ P2 chooses R with prob $\frac{\eta^*}{\varepsilon}$.

Type η^* player 2's indifference between L and R requires that:

$$\left(1 - \frac{\xi^*}{\varepsilon}\right)(1 + \eta^*) = 2\frac{\xi^*}{\varepsilon}.$$

$$\Rightarrow \quad \varepsilon(1 + \eta^*) = \xi^*(3 + \eta^*).$$

Background: Mixed Strategies & Harsanyi Purification

Let's solve for the equilibrium in the perturbed game:

–	L	R
T	$2+\xi, 1+\eta$	$\xi, 0$
B	$0, \eta$	$1, 2$

These two equations lead to a unique positive solution (ξ^*, η^*):

$$\varepsilon(1+\eta^*) = \xi^*(3+\eta^*) \quad \varepsilon(2+\xi^*) = \eta^*(3+\xi^*).$$

As $\varepsilon \to 0$, $\xi^*/\varepsilon \to 1/3$ and $\eta^*/\varepsilon \to 2/3$.

- As $\varepsilon \to 0$, players' action distribution in the perturbed game converge to the mixed-strategy equilibrium in the unperturbed game.

Harsanyi Purification Theorem

- N player finite normal-form game $\mathcal{G} \equiv (I, A, u)$, with $A \equiv \prod_{i=1}^{n} A_i$.
- Perturbation for player i: $\eta_i \in \mathbb{R}^{|A_i|}$ drawn according to continuous distribution μ_i s.t. player i's payoff from a is $u_i(a) + \eta_i(a)$.
- Only player i observes η_i, η_i is independent of η_j for every $i \neq j$.

Purification Theorem

For every regular Nash equilibrium $\sigma \in \Delta(A)$ *of* $\mathcal{G}$ *and for every* $\{\mu_1^k, ..., \mu_n^k\}_{k \in \mathbb{N}}$ *with* $\lim_{k \to +\infty} \mu_i^k = 0$ *for every i.*

For every $\varepsilon > 0$ *there exists* $\bar{k} \in \mathbb{N}$ *such that for every* $k > \bar{k}$*, the perturbed game* $\{\mathcal{G}, \mu_1^k, ..., \mu_n^k\}$ *has a Bayes Nash equilibrium s.t.*

- *each player has a strict incentive almost surely,*
- *this BNE induces a distribution that is within* ε *of* σ*.*

Harsanyi Purification Theorem

Purification Theorem

For every regular Nash equilibrium $\sigma \in \Delta(A)$ *of* $\mathcal{G}$ *and for every* $\{\mu_1^k, ..., \mu_n^k\}_{k\in\mathbb{N}}$ *with* $\lim_{k\to+\infty} \mu_i^k = 0$ *for every i.*

For every $\varepsilon > 0$ *there exists* $\bar{k} \in \mathbb{N}$ *such that for every* $k > \bar{k}$*, the perturbed game* $\{\mathcal{G}, \mu_1^k, ..., \mu_n^k\}$ *has a Bayes Nash equilibrium s.t.*

- *each player has a strict incentive almost surely,*
- *this BNE induces a distribution that is within* ε *of* σ*.*

Regular Nash equilibrium is a refinement of Nash equilibrium.

- In generic finite normal-form games, every Nash equilibrium is a regular Nash equilibrium (Van Damme 1991).

Definition of Regular Equilibrium

Fix a strategy profile $a^* \equiv (a_1^*, ..., a_n^*)$. For every $\boldsymbol{p} \in \Delta(A)$, let

$$F_i^k(\boldsymbol{p}|a^*) \equiv p_i^k \left\{ u_i(\boldsymbol{p}_{-i}, a_i^k) - u_i(\boldsymbol{p}_{-i}, a_i^*) \right\} \text{ for every } i \in I \text{ and } a_i^k \neq a_i^*,$$

$$F_i(\boldsymbol{p}|a^*) = \sum_{k=1}^{|A_i|} p_i^k - 1$$

Let

$$J(\boldsymbol{p}^*, a^*) \equiv \left. \frac{\partial \mathbf{F}(\boldsymbol{p}|a^*)}{\partial \boldsymbol{p}} \right|_{\boldsymbol{p}=\boldsymbol{p}^*}$$

An equilibrium $\boldsymbol{p}^*$ is *regular* if and only if there exists $a^* \in A$ such that $J(\boldsymbol{p}^*, a^*)$ is not a singular matrix.

From Normal Form to Extensive Form

Can we generalize this insight to extensive-form games?

Bhaskar: Overlapping Generation Repeated Games

- Time $t = 0, 1, 2, \ldots$ One agent born in each period.
- The agent who is born is period t:
  - Receives $K \in \mathbb{N}$ units of endowment in period t.
  - He shares $a_t \in \{0, 1, \ldots, K\}$ with his predecessor.
  - He may receive transfers from his successor in period $t + 1$.
- Agent t's payoff is $u(a_t, a_{t+1})$, which is strictly increasing in a_{t+1} and is strictly decreasing in a_t.
- The efficient level of transfer:

$$k^* \in \arg \max_{k \in \{0,1,\ldots,K\}} \Big\{ u(K - k, k) + u(k, K - k) \Big\}.$$

- We assume that k^* is strictly positive.
- We also assume that $u(K - k^*, k^*) > u(K, 0)$.

Benchmark: Unbounded Memories

Benchmark: Sustaining Cooperation with Perfect Information

If every agent can perfectly observe all previous agents' actions, then there exists a pure-strategy equilibrium that sustains transfer level k^*.

Proof: Grim-trigger strategies.

Introducing Limited Memories

For every $j > i$, agent j's info about agent i's action is a partition $B_{j,i}$ of A_i.

- **Assumption 1:** For every $k > j > i$, $B_{k,i}$ is weakly coarser than $B_{j,i}$.
- **Assumption 2:** There are infinitely many agents whose actions are no longer visible after a certain period.

Can we sustain cooperation via pure-strategy equilibria?

Theorem

Suppose $u(a_t, a_{t+1}) \neq u(a'_t, a'_{t+1})$ *for any* $(a_t, a_{t+1}) \neq (a'_t, a'_{t+1})$.

All agents choose 0 *in every pure-strategy equilibrium.*

Proof

Theorem

Suppose $u(a_t, a_{t+1}) \neq u(a'_t, a'_{t+1})$ for any $(a_t, a_{t+1}) \neq (a'_t, a'_{t+1})$.
All agents choose 0 in every pure-strategy equilibrium.

Suppose agent i's action is not visible starting from period $j+1$.

- Will agent j's action depend on agent i's action?
- Suppose h_j and h'_j differ only in agent i's action, and agent j plays a at h_j and plays a' at h'_j. Then agent j's incentive constraints imply that:

$$u(a, a_{j+1}(h_j, a)) \geq u(a', a_{j+1}(h_j, a'))$$

and

$$u(a', a_{j+1}(h'_j, a')) \geq u(a, a_{j+1}(h'_j, a)).$$

Not true since $a_{j+1}(h_j, a) = a_{j+1}(h'_j, a)$ and $a_{j+1}(h_j, a') = a_{j+1}(h'_j, a')$.

Proof

Theorem

Suppose $u(a_t, a_{t+1}) \neq u(a'_t, a'_{t+1})$ *for any* $(a_t, a_{t+1}) \neq (a'_t, a'_{t+1})$.
All players choose 0 *in every pure-strategy equilibrium.*

Suppose agent i's action is not visible starting from period $j+1$.

- Agent j's action does not depend on agent i's action.
- Can agent $j-1$'s action depend on agent i's action?
- Can agent $i+1$'s action depend on agent i's action?

What will agent i choose when no agent's action depends on his?
What will agent $i-1$ choose when agent i chooses 0 no matter what?

Sustaining Cooperation via Mixed Strategies

Theorem

Suppose $u(K-k^*, k^*) > u(K, 0)$ and agent $t+1$ observes the action of agent t. There exists a mixed strategy equilibrium that sustains transfer level k^*.

Consider an equilibrium in which

- Agent 0 chooses k^*.
- Agent i chooses k^* if agent $i-1$ chooses k^*, and mixes between 0 and k^* if agent $i-1$ chooses any other action.

Agent i is indifferent between 0 and k^*,

- Their mixing probability depends on agent $i-1$'s action, and is chosen in order to make agent $i-1$ indifferent.
- Agent i's incentive to mix is provided by agent $i+1$'s mixed action.

Are These Mixed-Strategy Equilibria Robust?

Suppose agents' payoffs are perturbed so that agent t's payoff from (a_t, a_{t+1}) equals:

$$u(a_t, a_{t+1}) + \eta_t(a_t, a_{t+1}),$$

where $\boldsymbol{\eta}_t \equiv \{\eta_t(a_t, a_{t+1})\}_{(a_t, a_{t+1}) \in A^2}$ is drawn according to continuous distribution F^n, and the support of $\eta_t(a_t, a_{t+1})$ is close to 0.

- η_t is independent of η_s for every $t \neq s$,
- η_t is independent of the history of play.

Theorem

When the support of the perturbations are close to 0*, all agents choose* 0 *in every equilibrium of every perturbed game.*

Proof Sketch

Theorem

All agents choose 0 in every equilibrium of every perturbed game.

Suppose agent i's action is not visible starting from period $j+1$.

- Will agent j's action depend on agent i's action?
- Agent j strictly prefers a to a' at h_j iff:

$$u(a, a_{j+1}(h_j, a)) + \eta_j(a, a_{j+1}(h_j, a)) > u(a', a_{j+1}(h_j, a')) + \eta_j(a', a_{j+1}(h_j, a')).$$

or equivalently,

$$\eta_j(a, a_{j+1}(h_j, a)) - \eta_j(a', a_{j+1}(h_j, a')) > u(a', a_{j+1}(h_j, a')) - u(a, a_{j+1}(h_j, a)).$$

Proof Sketch

Theorem

All agents choose 0 in every equilibrium of every perturbed game.

Suppose agent i's action is not visible starting from period $j+1$.

- Suppose h'_j and h_j differ only in agent i's action
- Agent j strictly prefers a to a' at h'_j iff:

$$u(a, a_{j+1}(h'_j, a)) + \eta_j(a, a_{j+1}(h'_j, a)) > u(a', a_{j+1}(h'_j, a')) + \eta_j(a', a_{j+1}(h'_j, a')).$$

or equivalently,

$$\eta_j(a, a_{j+1}(h'_j, a)) - \eta_j(a', a_{j+1}(h'_j, a')) > u(a', a_{j+1}(h'_j, a')) - u(a, a_{j+1}(h'_j, a)).$$

Proof Sketch

Theorem

All agents choose 0 *in every equilibrium of every perturbed game.*

Suppose agent i's action is not visible starting from period $j+1$.

- Suppose h_j' and h_j differ only in agent i's action.

  Then $a_{j+1}(h_j, a) = a_{j+1}(h_j', a)$ and $a_{j+1}(h_j, a') = a_{j+1}(h_j', a')$.
- Agent j strictly prefers a to a' at h_j iff:

  $$\eta_j(a, a_{j+1}(h_j, a)) - \eta_j(a', a_{j+1}(h_j, a')) > u(a', a_{j+1}(h_j, a')) - u(a, a_{j+1}(h_j, a)).$$
- Agent j strictly prefers a to a' at h_j' iff:

  $$\eta_j(a, a_{j+1}(h_j', a)) - \eta_j(a', a_{j+1}(h_j', a')) > u(a', a_{j+1}(h_j', a')) - u(a, a_{j+1}(h_j', a)).$$
- Prob that agent j chooses each action is independent of agent i's action.