Proposition Lemma

Lecture 3: Community Enforcement

Harry PEI
Department of Economics, Northwestern University

Mini Course at Oxford University

From Rich Information to Limited Information

Most of the repeated game models: Players' information is rich.

- Players can observe the entire history of actions (e.g., Fudenberg and Maskin 1986, Fudenberg and Levine 1989).
- 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 are unfamiliar with their partners (e.g., Maghribi traders in Medieval Europe, eCommerce platforms).
- People don't know who they are playing with (e.g., journal refereeing).

Can societies sustain good outcomes with limited information?

Model

- Time $t = 0, 1, 2, \ldots$
- $N \equiv 2n$ players, discount factor $\delta \in (0, 1)$.
- In each period, players are matched uniformly at random to play the prisoner's dilemma:

-	**Cooperate**	**Defect**
Cooperate	$1, 1$	$-l, 1+g$
Defect	$1+g, -l$	$0, 0$

with $g, l > 0$.

The matching process is independent across periods.

- **Monitoring structure:**

Each player only observes the action profile of his own matches.

He cannot observe the identity of his current/past opponents.

He cannot observe what happened in other matches.

How Can Players Sustain Cooperation?

Kandori (1992) proposes the following *contagion strategy*:

- Each player has two private states: c and d.
- The player plays C if his private state is c and plays D if his private state is d.

- All players' private states are c in period 0.
- For each player, his private state is c *if and only if* he hasn't observed anything other than (C, C) in his previous matches.

Lemma

For every $g, l > 0$ and $n \in \mathbb{N}$, there exists $\underline{\delta} \in (0, 1)$ such that when $\delta > \underline{\delta}$, all players using the contagion strategy is a Nash equilibrium.

How Can Players Sustain Cooperation?

Lemma

For every $g, l > 0$ and $n \in \mathbb{N}$, there exists $\underline{\delta} \in (0,1)$ such that when $\delta > \underline{\delta}$, all players using the contagion strategy is a Nash equilibrium.

Why is the contagion strategy a Nash equilibrium?

- We only need to verify players' incentives on the equilibrium path.

  All players play C on the equilibrium path.

- For any given player i, if he deviates to D, then
  - He obtains a one-period gain of g.
  - But he infects others in the community and spreads contagion.
  - Eventually, he will encounter someone who is in state d.

  When δ is large enough, his one-period gain is less than his long-term loss from spreading contagion.

How Can Players Sustain Cooperation?

Kandori (1992) proposes the following *contagion strategy*:

- Each player has two private states: c and d.
- The player plays C if his private state is c and plays D if his private state is d.

- All players' private states are c in period 0.
- For each player, his private state is c *if and only if* he hasn't observed anything other than (C, C) in his previous matches.

Question: Is the contagion strategy a sequential equilibrium?

- Not necessarily!
- Suppose you observe D in period 0, do you play D in period 1?

  You think that at most one person is infected.

  Playing C is attractive since it slows down contagion.

Theorem 1 in Kandori (1992)

Kandori (1992) shows that the contagion strategy is a sequential equilibrium when l is large enough relative to g and n.

Theorem 1 in Kandori (1992)

For every $g > 0$ and $n \in \mathbb{N}$, there exist $\underline{l} > 0$ and $\underline{\delta} \in (0, 1)$,

such that when $l > \underline{l}$ and $\delta > \underline{\delta}$,

all players using the contagion strategy is a sequential equilibrium.

Intuition: How to motivate players to play D in state d?

- When l is large relative to n and g,

  the loss from playing C while encountering someone playing D

  is much larger relative to the benefit from slowing down contagion.

Limitations of Kandori's result

Theorem 1 in Kandori (1992)

For every $g > 0$ and $n \in \mathbb{N}$, there exist $\underline{l} > 0$ and $\underline{\delta} \in (0, 1)$, such that when $l > \underline{l}$ and $\delta > \underline{\delta}$, all players using the contagion strategy is a sequential equilibrium.

It does not imply that players can cooperate in *all* prisoner's dilemma.

- l needs to be implausibly large as $n \to +\infty$.

Cooperation is very fragile and is not robust to trembles.

- One defection causes cooperation to breakdown all together.

Ellison (1994)

Ellison (1994) sharpens Kandori's results in three steps:

1. He assumes that players have a public randomization device and shows that cooperation is feasible for all g and l.
2. He shows that the public randomization device is dispensable, i.e., patient players can sustain cooperation without it.
3. He shows that the equilibria he constructs are robust to trembles.

Community Enforcement with Public Randomization

- Time $t = 0, 1, 2, ...$
- $N \equiv 2n$ players, all share the same discount factor $\delta \in (0, 1)$.
- In each period, players are matched uniformly at random to play the prisoner's dilemma:

-	**Cooperate**	**Defect**
Cooperate	$1, 1$	$-l, 1+g$
Defect	$1+g, -l$	$0, 0$

with $g, l > 0$.

The matching process is independent across periods.

- By the end of each period, a random variable $q_t \sim U[0, 1]$.
- Each player only observes the action profile of his own matches and the entire history of public randomizations.

Contagion Strategy with Moderate Punishment

Ellison (1994) proposes the following modified version of Kandori's contagion strategy, parameterized by $\widehat{q} \in [0, 1]$, call it $\sigma_{\widehat{q}}$:

- Each player has two private states: c and d.
- The player plays C if his private state is c and plays D if his private state is d.

- All players' private states are c in period 0.
- For every $t \geq 1$, a player's period t state is c *if and only if* period $t-1$ action profile was (C, C) *or* $q_{t-1} \geq \widehat{q}$.

Intuition: When actions other than (C, C) occur, punish with prob $\widehat{q}$.

- After contagion has began, an amnesty with prob $1 - \widehat{q}$ in each period.
- Kandori (1992)'s contagion strategy corresponds to $\widehat{q} = 1$.

Why can moderating punishment help?

Two ICs are required to sustain cooperation in sequential equilibrium.

- Incentive to play C when their private state is c, which is stronger when $\widehat{q}$ is larger.
- Incentive to play D when their private state is d, which is stronger when $\widehat{q}$ is smaller.

  Why? The only benefit from playing C is to slow down contagion, and $\widehat{q} = 0$ kills contagion all together.

Does there exist a $\widehat{q} \in [0, 1]$ that can satisfy both constraints?

- Ellison's answer: Yes! For all δ large enough.

Proposition 1 in Ellison (1994)

Proposition 1 in Ellison (1994)

In the community enforcement game with public randomization.
For every g, l, n, there exists $\underline{\delta} \in (0, 1)$ such that when $\delta > \underline{\delta}$,
there exists a sequential equilibrium where (C, C) is always played on-path.

In fact, for every δ large enough, there exists a $\widehat{q} \in [0, 1]$ such that all players playing $\sigma_{\widehat{q}}$ is such a sequential equilibrium.

Thought experiment: Fix any $\widehat{q}$ and suppose others play $\sigma_{\widehat{q}}$,

- Suppose you think that k other players are in state d,

  does your incentive to play D increase or decrease in k?

Ellison (1994) shows that it is increasing in k.

- **Intuition:** When there are more infected players already, an additional infected player makes less difference.

Proof: Increasing Incentives to Defect

Let ω be a realization of the matching process (for everyone from 0 to $+\infty$).

Let $f(k, \delta, \widehat{q}, \omega)$ be player 1's continuation value when

- He plays D until observing $q \geq \widehat{q}$, i.e., until the amnesty.
- Other players use strategy $\sigma_{\widehat{q}}$.
- k of the other players are in private state d.
- The realized matching process is ω.

Lemma

For every $k' > k$, we have

$$f(k, \delta, \widehat{q}, \omega) - f(k+1, \delta, \widehat{q}, \omega) \geq f(k', \delta, \widehat{q}, \omega) - f(k'+1, \delta, \widehat{q}, \omega).$$

Proof: Increasing Incentives to Defect

Lemma

For every $k' > k$, we have

$$f(k, \delta, \widehat{q}, \omega) - f(k+1, \delta, \widehat{q}, \omega) \geq f(k', \delta, \widehat{q}, \omega) - f(k'+1, \delta, \widehat{q}, \omega).$$

Let us compare $f(k, \delta, \widehat{q}, \omega)$ to $f(k+1, \delta, \widehat{q}, \omega)$:

- They are the same after period t if $q_s \geq \widehat{q}$ for some $s < t$.
- If $q_s < \widehat{q}$ for all $s < t$, then player 1's period t stage-game payoffs are different only when he is matched with someone
  1. who will not be infected before t by the first k players,
  2. who will be infected before t by the $k+1$th player,

  in which case his payoff is reduced by $1 + g$.
- The red set is independent of k while the blue set shrinks with k.
- The intersection of these two sets is smaller when k increases.

Proof: Expression for the Payoff Difference

Lemma

$\frac{f(k,\delta,\widehat{q},\omega)-f(k+1,\delta,\widehat{q},\omega)}{1-\delta}$ *depends on* $\widehat{q}$ *and* δ *only through* $\delta\widehat{q}$.

Let us compare $f(k,\delta,\widehat{q},\omega)$ to $f(k+1,\delta,\widehat{q},\omega)$:

- They are the same after period t if $q_s \geq \widehat{q}$ for some $s < t$.
- If $q_s < \widehat{q}$ for all $s < t$, then player 1's period t stage-game payoffs are different only when he is matched with someone:
  1. who will not be infected before t by the first k players,
  2. who will be infected before t by the $k+1$th player.

Hence,

$$f(k,\delta,\widehat{q},\omega) - f(k+1,\delta,\widehat{q},\omega) = (1-\delta)\sum_{t=0}^{+\infty} \delta^t \widehat{q}^t (1+g)\mathbf{1}\{...\}$$

where "..." stands for the event that "you encounter someone that belongs to both the red and the blue set".

Proof of the Ellison Theorem

Proposition

For every δ large enough, there exists $\widehat{q} \in [0, 1]$ such that all players using strategy $\sigma_{\widehat{q}}$ is a sequential equilibrium.

Choose $\widehat{q}$ such that players are indifferent between C and D when $k = 0$:

$$(1 - \delta)g = \delta \cdot \widehat{q} \cdot \mathbb{E}_\omega[f(0, \delta, \widehat{q}, \omega) - f(1, \delta, \widehat{q}, \omega)]. \quad (1)$$

A player's incentive to play D after he is infected:

- If his current partner is infected, then playing D and playing C leads to the same continuation value, yet playing D leads to a benefit l.
- If his current partner is not infected, then the payoff difference between playing D and playing C is:

  $$(1 - \delta)g - \delta \cdot \widehat{q} \cdot \mathbb{E}_{k,\omega}[f(k, \delta, \widehat{q}, \omega) - f(k + 1, \delta, \widehat{q}, \omega)],$$

  which is positive given (1) and $f(k, ..) - f(k + 1, ...)$ is decreasing in k.

Proof of the Ellison Theorem

Choose $\widehat{q}$ such that players are indifferent between C and D when $k = 0$:

$$(1-\delta)g = \delta \cdot \widehat{q} \cdot \mathbb{E}_\omega[f(0,\delta,\widehat{q},\omega) - f(1,\delta,\widehat{q},\omega)]. \tag{2}$$

Question: Does there exist such a $\widehat{q}$?

Yes! Why? Let $\widehat{q} = 1$.

- By continuity, there exists $\widehat{\delta} \in (0,1)$ such that when $\delta > \widehat{\delta}$, inequality (2) is true when $\widehat{q} = 1$.

Since $\frac{f(k,\delta,\widehat{q},\omega) - f(k+1,\delta,\widehat{q},\omega)}{1-\delta}$ depends on $\widehat{q}$ and δ only through $\delta\widehat{q}$, for every $\delta > \widehat{\delta}$, we can set $\widehat{q} = \widehat{\delta}/\delta$ which satisfies (2).

Remove the Public Randomization Device

Ellison's construction relies on a public randomization device.

- Grants an amnesty after each period with probability $1 - \widehat{q}$.
- Moderate the punishment to provide players incentives to punish.

Can we moderate the punishment without any public randomization?

- Yes, when δ is close enough to 1.
- Ellison introduces a cool trick to do this.

Ellison's Trick: How to Lower the Discount Factor

Theorem: Lowering the Discount Factor

Let $G(\delta)$ be any repeated complete info game.

Suppose there exists a non-empty interval (δ_0, δ_1) such that for every $\delta \in (\delta_0, \delta_1)$, $G(\delta)$ has an equilibrium $s^*(\delta)$ with outcome $\alpha \in \Delta(A)$.

Then there exists $\underline{\delta} < 1$ such that for every $\delta^* \in (\underline{\delta}, 1)$, there also exists a strategy profile $s^{**}(\delta^*)$ which is an equilibrium in $G(\delta^*)$ and implements α.

There exists $\underline{\delta} \in (0, 1)$ such that for every $\delta > \underline{\delta}$, there exists $N(\delta) \in \mathbb{N}$ such that $\delta^{N(\delta)} \in (\delta_0, \delta_1)$.

- Treat the entire repeated game as $N(\delta)$ separate repeated games.
- Repeated game 1 is played in period $0, N(\delta), 2N(\delta), \ldots$
- Repeated game 2 is played in period $1, N(\delta) + 1, 2N(\delta) + 1, \ldots$

Ellison's Theorem without Public Randomization

Proposition 4 in Ellison (1994)

In the community enforcement game without public randomization.

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

there exists a sequential equilibrium where (C, C) is always played on-path.

Strategy $\sigma_{\widehat{q}}$ with $\widehat{q} = 1$ does not require public randomization.

Recall that for every $k' > k$ and every ω, we have

$$f(k,\delta,1,\omega) - f(k+1,\delta,1,\omega) \geq f(k',\delta,1,\omega) - f(k'+1,\delta,1,\omega).$$

Therefore, we know that for every $k' > k$

$$\mathbb{E}_\omega[f(k,\delta,1,\omega) - f(k+1,\delta,1,\omega)] > \mathbb{E}_\omega[f(k',\delta,1,\omega) - f(k'+1,\delta,1,\omega)].$$

Ellison's Theorem without Public Randomization

Recall that fix $\widehat{q} = 1$, there exists $\widehat{\delta} \in (0, 1)$ such that

$$(1 - \widehat{\delta})g = \widehat{\delta} \cdot \mathbb{E}_\omega[f(0, \widehat{\delta}, 1, \omega) - f(1, \widehat{\delta}, 1, \omega)],$$

i.e., indifferent between C and D when discount factor is $\widehat{\delta}$, no other player is infected, and all players use strategy σ_1.

Since for every $k' > 0$,

$$\mathbb{E}_\omega[f(0, \delta, 1, \omega) - f(1, \delta, 1, \omega)] > \mathbb{E}_\omega[f(k', \delta, 1, \omega) - f(k' + 1, \delta, 1, \omega)],$$

there exists an open set of discount factors $(\widehat{\delta}, \widehat{\delta} + \varepsilon)$ such that every player

- prefers C to D when no other player is infected,
- prefers D to C when at least one other player is infected.

Applying the Ellison's trick, we know that (C, C) can be attained in sequential equilibrium for δ large enough even w/o public randomization.

Robustness to Trembles

Recall that a major critique of Kandori's construction is that the equilibrium is not robust to small trembles.

Ellison examines two types of trembles.

1. Independent trembles:
    - Each agent is forced to play D with prob $\varepsilon > 0$ in each period, and trembles are independent across players and across periods.
2. Correlated trembles:
    - Each agent is a commitment type with prob ε, and commitment types always play D. Players' types are independent of each other.

Robustness to Independent Trembles

Independent trembles: Each agent is forced to play D with prob $\varepsilon > 0$ in each period, and trembles are independent across players and across periods.

For the construction *with* public randomization:

- For every g, l, n and $\eta > 0$, there exist $\bar{\varepsilon} > 0$ and $\underline{\delta} \in (0, 1)$, such that when $\delta > \underline{\delta}$ and $\varepsilon < \bar{\varepsilon}$, there exists a sequential equilibrium in which (C, C) is played with probability more than $1 - \eta$

For the construction *without* public randomization:

- Cooperation breaks down as $t \to +\infty$, but players' payoffs are close to 1 in the double limit where $\lim_{\varepsilon \to 0} \lim_{\delta \to 1} \dots$

Ellison's constructions are somewhat robust to small independent trembles.

- What about large trembles or general noisy monitoring?

What about Correlated Trembles?

Correlated trembles: Each agent is a commitment type with prob ε, and commitment types always play D.

Why are correlated trembles different from independent trembles?

- Independent trembles: After an amnesty, you know that everyone will play C with prob ≈ 1.
- Correlated trembles: Your belief about the number of commitment types depends on your private history.

  e.g., you may have no incentive to play C even after an amnesty if you believe that many players are commitment types.

Community Enforcement with Incomplete Information

Consider a large population of players playing the prisoner's dilemma:

- A fraction of the population are bad types who always play D,

  e.g., each player is normal w.p. $1-\varepsilon$ and is bad w.p. ε.

- In each period, players are randomly matched and can only observe the actions in their own match.

Two key findings:

- Sugaya and Wolitzky (2020): Anti-folk theorem w/o communication.
- Sugaya and Wolitzky (2021): Folk theorem when players can communicate via cheap talk messages.

A General Anonymous Repeated Game with Bad Types

- Discrete time $t = 0, 1, 2, ...$
- N players with discount factor δ.
- Each player's action set A, with $a_t \in A^N$ the action profile at t.

- Player i's type $\theta_i \in \{R, B\}$, with type B taking a^* in every period.
- Type distribution $p \in \Delta\left(\{R, B\}^N\right)$.

- Player i's private signal $y_{i,t} \sim F(\cdot|(a_\tau, y_\tau)_{\tau=0}^{t-1}, a_t)$.
- Public randomization device $\xi_t \sim U[0, 1]$.
- Player i's private history in period t consists of θ_i and $(a_{i,\tau}, y_{i,\tau}, \xi_\tau)_{\tau=0}^{t-1}$.

- Players' stage-game payoffs $(u_1, ..., u_N) : A^N \to [0, 1]^N$.

Symmetry Assumptions

Assumption: Symmetric Type Distribution

$p(\theta_1, ..., \theta_n)$ depends only on the number of bad types in $(\theta_1, ..., \theta_n)$.

Assumption: Symmetric Payoff Function

Fix $i, j \in \{1, 2, ..., N\}$. We have $u_i(a_i, a_{-i}) = u_j(a'_j, a'_{-j})$ if

- $a_i = a'_j$,
- *the number of other players playing each action is the same under a_{-i} and under a'_{-j}.*

Prisoner's Dilemma with Uniform Random Matching

Leading example: $N = 2n$ players are uniformly matched into pairs in each period to play the prisoner's dilemma.

- Payoffs are symmetric since matching is uniform and anonymous.

  Each opponent's action matters for your payoff with prob $\frac{1}{N-1}$.
- The private signal $y_{i,t}$ is the action profile in agent i's match, i.e., agent i perfect observes each opponent's action with prob $\frac{1}{N-1}$.
- The type distribution is symmetric when types are i.i.d. and each player is bad with probability ε.

Analysis

- Focus on symmetric equilibrium.
  - Given the symmetry assumptions and the presence of public randomization, this is without loss if the focus is on $\sum U_i/N$.

- Let $\mathcal{B}_n$ be the event that there are n bad players, with $p_n \equiv \Pr(\mathcal{B}_n)$.

- Let $q_n \equiv \Pr\left(n \text{ out of } N-1 \text{ other players are bad} \middle| \text{player i is rational}\right)$.
- Let $q_n^- \equiv q_{n-1}$. Let $q_N \equiv 0$ and $q_0^- \equiv 0$.

  Both $q \equiv (q_0, ..., q_N)$ and $q^- \equiv (q_0^-, ..., q_N^-)$ are prob distributions.

- The total variation distance between q and q^- is:

$$\Delta \equiv \max_{\mathcal{N} \subset \{0,1,\ldots,N\}} \left| \sum_{n \in \mathcal{N}} (q_n - q_n^-) \right|.$$

Analysis

Interpretations of the two distributions q and q^-:

- Let $q_n \equiv \Pr\left(n \text{ out of } N-1 \text{ other players are bad} \middle| \text{player i is rational}\right)$.
- Let $q_n^- \equiv q_{n-1}$.

Suppose the rational type's equilibrium strategy is *not* a^* in every period.

- If I am rational and play my equilibrium strategy, then q is my belief about the total number of people playing a^* in every period.
- If I am rational but I deviate to a^* in every period, then q^- is my belief about the total number of people playing a^* in every period.

- Therefore, Δ measures the *detectability* of a rational type's deviation to the bad type's strategy.

Lower Bound on Rational Type's Payoff

Let $U_i(\theta)$ be player 1's equilibrium payoff conditional on type profile θ.

Let

$$u_n^R \equiv \mathbb{E}[U_i(\theta)|\theta_i = R, \mathcal{B}_n] \text{ and } u_n^B \equiv \mathbb{E}[U_i(\theta)|\theta_i = B, \mathcal{B}_n].$$

Lemma

In every equilibrium of the repeated game, we have

$$\sum_{n=0}^{N-1} q_n u_n^R \geq \sum_{n=0}^{N-1} q_n u_n^B - \Delta.$$

What is the rational type's expected payoff when he plays his equilibrium strategy?

- $\sum_{n=0}^{N-1} q_n u_n^R$.

Lower Bound on Rational Type's Payoff

Let $U_i(\theta)$ be player 1's equilibrium payoff conditional on type profile θ.

Let

$$u_n^R \equiv \mathbb{E}[U_i(\theta)|\theta_i = R, \mathcal{B}_n] \text{ and } u_n^B \equiv \mathbb{E}[U_i(\theta)|\theta_i = B, \mathcal{B}_n].$$

Lemma

In any equilibrium,

$$\sum_{n=0}^{N-1} q_n u_n^R \geq \sum_{n=0}^{N-1} q_n u_n^B - \Delta.$$

What is the rational type's expected payoff when he deviates and plays a^* in every period?

- $\sum_{n=0}^{N-1} q_n u_{n+1}^B = \sum_{n=0}^{N} q_n^- u_n^B$.

  (comes directly from $q_n^- = q_{n-1}$)

Proof: Lower Bound on Payoff

Let

$$u_n^R \equiv \mathbb{E}[U_i(\theta)|\theta_i = R, \mathcal{B}_n] \text{ and } u_n^B \equiv \mathbb{E}[U_i(\theta)|\theta_i = B, \mathcal{B}_n].$$

Lemma

In any equilibrium,

$$\sum_{n=0}^{N-1} q_n u_n^R \geq \sum_{n=0}^{N-1} q_n u_n^B - \Delta.$$

Rational type's payoff from deviating to a^* in every period is given by $\sum_{n=0}^{N-1} q_n u_{n+1}^B = \sum_{n=0}^{N} q_n^- u_n^B$. Therefore,

$$\sum_{n=0}^{N-1} q_n u_{n+1}^B = \sum_{n=0}^{N-1} q_n u_n^B - \sum_{n=0}^{N} (q_n - q_n^-) u_n^B \geq \sum_{n=0}^{N-1} q_n u_n^B - \Delta$$

The blue term is no more than his equilibrium payoff $\sum_{n=0}^{N-1} q_n u_n^R$.

Pairwise Dominant Action

This lemma is useful in games where a^* is a pairwise dominant action:

Assumption: Pairwise Dominance

Action $a^* \in A$ is a pairwise dominant action if there exists $c > 0$ such that for every $a \neq a^*$ and $a_{-ij} \in A^{N-2}$, we have

$$u_i(a_i = a^*, a_j = a, a_{-ij}) - u_j(a_j = a, a_i = a^*, a_{-ij}) > c.$$

This neither implies nor is implied by a^* being a dominant action.

- Find two counterexamples to convince yourself.

In the prisoner's dilemma game with uniform random matching:

- D is a pairwise dominant action since

$$\frac{x+1}{N-1}(1+g) \geq \frac{x}{N-1} - l \cdot \frac{N-1-x}{N-1} + \underbrace{\min\{g,l\}}_{\equiv c},$$

where x is the number of people playing C other than i and j.

Upper Bound on Rational Type's Payoff

Fix an equilibrium. When the rational type plays his equilibrium strategy,

- let γ_n be the occupation measure with which he plays actions other than a^* conditional on there are n bad types in the population.

Recall that

$$u_n^R \equiv \mathbb{E}[U_i(\theta)|\theta_i = R, \mathcal{B}_n] \text{ and } u_n^B \equiv \mathbb{E}[U_i(\theta)|\theta_i = B, \mathcal{B}_n].$$

Lemma

If a^* is a pairwise dominant action, then $u_n^B \geq u_n^R + \gamma_n c$ for every n.

This follows from the definition of pairwise dominant actions.

Lower Bound on the Occupation Measure of a^*

Combining the two lemmas:

Lemma

In any equilibrium, $\sum_{n=0}^{N-1} q_n u_n^R \geq \sum_{n=0}^{N-1} q_n u_n^B - \Delta$.

Lemma

If a^* *is a pairwise dominant action, then* $u_n^B \geq u_n^R + \gamma_n c$ *for every* n.

we obtain the following inequality:

$$\Delta \geq \sum_{n=0}^{N-1} q_n (u_n^B - u_n^R) \geq c \cdot \sum_{n=0}^{N-1} q_n \gamma_n.$$

The expected occupation measure of actions other than a^*, $\sum_{n=1}^{N-1} q_n \gamma_n$, is no more than $\frac{\Delta}{c}$, i.e., the expected occupation measure of a^* is at least $1 - \frac{\Delta}{c}$.

Anti-Folk Theorem

Recall that the expected occupation measure of actions other than a^*, $\sum_{n=1}^{N-1} q_n \gamma_n$, is no more than $\frac{\Delta}{c}$.

If $\Delta \to 0$, then:

- In every equilibrium, the rational type plays a^* in almost all periods.
- Social welfare is close to the case in which everyone is bad.

This leads to an anti-folk theorem, i.e., all payoffs are close to $U(a^*)$.

When is it the case that $\Delta \to 0$ as $N \to +\infty$?

Leading example: Each player is bad with prob ε, and players' types are independently drawn from the same distribution.

Fix $\varepsilon > 0$.

- $q_n = \binom{N-1}{n}(1-\varepsilon)^{N-n}\varepsilon^n$.
- $q_n^- = q_{n-1} = \binom{N-1}{n-1}(1-\varepsilon)^{N-n+1}\varepsilon^{n-1}$.

Since q_n is single-peaked in n, the total variation distance is

$$\Delta = q_0 + (q_1 - q_0) + ... + (q_k - q_{k-1}) = q_k$$

where $q_k \equiv \max_{n\in\{0,1,..,N\}} q_n$.

As $N \to +\infty$, $\max_{n\in\{0,..,N\}} \binom{N-1}{n}(1-\varepsilon)^{N-n}\varepsilon^n \to 0$.

Therefore, $\Delta \to 0$ as $N \to +\infty$.

Conclusion: Anti-Folk Theorem under Incomplete Info

Sugaya and Wolitzky (2020)'s result implies that:

- In a repeated prisoner's dilemma with uniform random matching and each player is a bad type who always defects with prob ε,

  all equilibrium payoffs converge to the minmax value as $N \to +\infty$.

Hence, it is impossible to sustain cooperation in large populations.

Sugaya and Wolitzky (2021) focus on this specific setting.

- Theorem 1 in Sugaya and Wolitzky (2021): Extend the anti-folk theorem to when players can observe their partners' identities.
- As $(1 - \delta)N \to +\infty$, every NE payoff is close to 0.