Lecture 11: Repeated Incomplete Information Games with Discounting

Harry PEI
Department of Economics, Northwestern University

Spring Quarter, 2021

Last Lecture: Shalev (1994)

Repeated games with two long-run players and without discounting.

- P1 privately observes a persistent state $\theta \in \Theta$
- P2's belief about the state is $\pi \in \Delta(\Theta)$.
- Players' stage-game payoffs $u_1(\theta, a_1, a_2)$ and $u_2(a_1, a_2)$.
- Player 1 maximizes $\lim_{T \to +\infty} \frac{1}{T} \sum_{t=1}^{T} u_1(\theta, a_{1,t}, a_{2,t})$.

  Player 2 maximizes $\lim_{T \to +\infty} \frac{1}{T} \sum_{t=1}^{T} u_2(a_{1,t}, a_{2,t})$.

Result: A payoff $((u_\theta)_\theta, v)$ is an equilibrium payoff if and only if there exists $\{\alpha^\theta\}_{\theta \in \Theta}$ with $\alpha^\theta \in \Delta(A_1 \times A_2)$ such that:

1. **Feasibility:** $u_\theta = u_1(\theta, \alpha^\theta)$ for every $\theta \in \Theta$. $v = \sum_{\theta \in \Theta} \pi(\theta) u_2(\alpha^\theta)$.
2. **IR:** $\sum_{\theta \in \Theta} p(\theta) u_\theta \geq \underline{u}(p)\ \forall\, p \in \Delta(\Theta)$. $u_2(\alpha^\theta) \geq \underline{v}\ \forall\, \theta \in \Theta$.
3. **IC:** For every $\theta, \theta' \in \Theta$, type θ weakly prefers α^θ to $\alpha^{\theta'}$.

Let $\mathscr{U} \subset \mathbb{R}^{|\Theta|}$ be the projection of this set on P1's payoff.

Today: Games with Discounting

- Time $t = 1, 2, \ldots$
- Two patient players: 1 and 2. Actions $a_1 \in A_1$ and $a_2 \in A_2$.
- P1 has private info about a persistent state $\theta \in \Theta$.

  P2's prior belief $\pi \in \Delta(\Theta)$.
- Stage-game payoffs $u_1(\theta, a_1, a_2)$ and $u_2(a_1, a_2)$.
- Both players can perfectly observe all the past actions.
- Players maximize:

$$\sum_{t=0}^{+\infty}(1-\delta_1)\delta_1^t u_1(\theta, a_{1,t}, a_{2,t}) \text{ and } \sum_{t=0}^{+\infty}(1-\delta_2)\delta_2^t u_2(a_{1,t}, a_{2,t}),$$

Cripps and Thomas (2003): Focus on P1's payoffs.

- What will happen when $\delta_1 \to 1$ and δ_2 is bounded away from 1?
- What will happen when both δ_1 and δ_2 are close to 1?

P2's Discount Factor is Bounded Away from 1

$\mathscr{U} \subset \mathbb{R}^{|\Theta|}$ is P1's equilibrium payoff set in a game w/o discounting.

Theorem: Necessary Condition for Equilibrium Payoff

Fix $\delta_2 \in (0,1)$ *and full support* π.

Then for every $\varepsilon > 0$, *there exists* $\underline{\delta}_1 \in (0,1)$ *such that for all* $\delta_1 \in (\underline{\delta}_1, 1)$,

if $\boldsymbol{u} \equiv (u_\theta)_{\theta \in \Theta}$ *is player* 1*'s equilibrium payoff, then*

$$\min_{\boldsymbol{u}^* \in \mathscr{U}} ||\boldsymbol{u}^* - \boldsymbol{u}|| < \varepsilon.$$

Proof Sketch

Theorem: Necessary Condition for Equilibrium Payoff

Fix $\delta_2 \in (0,1)$ and full support π.
Then for every $\varepsilon > 0$, there exists $\underline{\delta}_1 \in (0,1)$ such that for all $\delta_1 \in (\underline{\delta}_1, 1)$, if $\boldsymbol{u} \equiv (u_\theta)_{\theta\in\Theta}$ is player 1's equilibrium payoff, then

$$\min_{\boldsymbol{u}^* \in \mathscr{U}} ||\boldsymbol{u}^* - \boldsymbol{u}|| < \varepsilon.$$

Given any equilibrium $((\sigma_\theta)_{\theta\in\Theta}, \sigma_2)$ and $(a_1, a_2) \in A_1 \times A_2$, let

$$\alpha^\theta(a_1, a_2) \equiv \mathbb{E}^{(\sigma_\theta, \sigma_2)}\left[\sum_{t=0}^{+\infty}(1-\delta_1)\delta_1^t \mathbf{1}\{(a_{1,t}, a_{2,t}) = (a_1, a_2)\}\right].$$

Let $\alpha^\theta \in \Delta(A_1 \times A_2)$ be the allocation of type θ.

The necessity of feasibility and IC are straightforward.

Proof Sketch

Theorem: Necessary Condition for Equilibrium Payoff

Fix $\delta_2 \in (0,1)$ and full support π.
Then for every $\varepsilon > 0$, there exists $\underline{\delta}_1 \in (0,1)$ such that for all $\delta_1 \in (\underline{\delta}_1, 1)$, if $\boldsymbol{u} \equiv (u_\theta)_{\theta\in\Theta}$ is player 1's equilibrium payoff, then

$$\min_{\boldsymbol{u}^*\in\mathscr{U}} ||\boldsymbol{u}^* - \boldsymbol{u}|| < \varepsilon.$$

IR-1: Generalize the Blackwell approachability theorem to discounted games where $\delta_1 \to 1$.

IR-2: Suppose P1 plays type θ's equilibrium strategy σ_θ.

- There exists $T \in \mathbb{N}$ s.t. $1 - \delta_2^T \approx 1$.
- P2's payoff conditional on θ is $\geq \underline{v} - \varepsilon$ if they are convinced that P1's strategy is close to σ_θ in the next T periods.
- There can be at most a bounded number of periods s.t. P2 believes that P1's strategy in the next T periods is far away from σ_θ.

Shalev's conditions are necessary, but not sufficient

Belonging to $\mathscr{U}$ is a necessary condition for P1's equilibrium payoff when δ_1 goes to 1, but it is in general not sufficient.

Example: Shalev's conditions are not sufficient

Suppose $\Theta = \{\theta\}$, $\theta, b, c > 0$, and $\delta_2 < \gamma^* \equiv \frac{c}{b+c}$.

-	T	N
H	$1-\theta, b$	$0,0$
L	$1, -c$	$0,0$

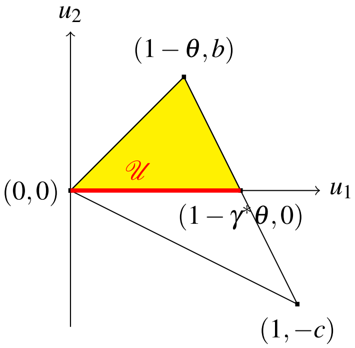

Example: Shalev's conditions are not sufficient

$\delta_2 < \gamma^* \equiv \frac{c}{b+c}$

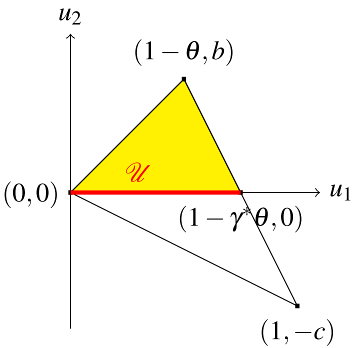

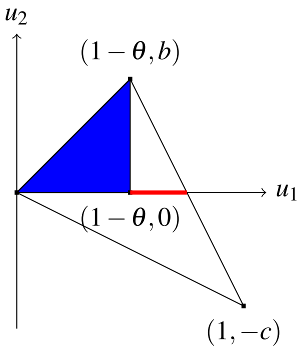

Intuition: P2's impatience introduces additional constraints on P1's equilibrium payoffs beyond feasibility, IR, and IC in Shalev (1994).

Both Players' Discount Factors are Close to 1

Let $\mathscr{U}$ be the set of $(u_\theta)_{\theta\in\Theta}$ s.t. there exist $\{\alpha^\theta\}_{\theta\in\Theta}$ and $v\in\mathbb{R}$ satisfying

1. **Feasibility:** $u_\theta = u_1(\theta,\alpha^\theta)$ for every $\theta\in\Theta$. $v=\sum_{\theta\in\Theta}\pi(\theta)u_2(\alpha^\theta)$.
2. **IR:** $\sum_{\theta\in\Theta}p(\theta)u_\theta \geq \underline{u}(p)\ \forall\ p\in\Delta(\Theta)$. $u_2(\alpha^\theta)\geq \underline{v}\ \forall\ \theta\in\Theta$.
3. **IC:** For every $\theta\neq\theta'$, type θ weakly prefers α^θ to $\alpha^{\theta'}$.

Let $\widehat{\mathscr{U}}$ be the set of $(u_\theta)_{\theta\in\Theta}$ s.t. there exist $\{\alpha^\theta\}_{\theta\in\Theta}$ and $v\in\mathbb{R}$ satisfying

1. **Feasibility:** $u_\theta = u_1(\theta,\alpha^\theta)$ for every $\theta\in\Theta$. $v=\sum_{\theta\in\Theta}\pi(\theta)u_2(\alpha^\theta)$.
2. **Strict IR:** $\sum_{\theta\in\Theta}p(\theta)u_\theta > \underline{u}(p)\ \forall\ p\in\Delta(\Theta)$. $u_2(\alpha^\theta) > \underline{v}\ \forall\ \theta\in\Theta$.
3. **Strict IC:** For every $\theta\neq\theta'$, type θ strictly prefers α^θ to $\alpha^{\theta'}$.

Both Players' Discount Factors are Close to 1

Theorem: Sufficient Condition for Equilibrium Payoff

For every $\boldsymbol{u} \in \widehat{\mathscr{U}}$ *and* $\varepsilon > 0$, *there exists* $\underline{\delta} \in (0,1)$ *such that*

whenever $1 > \delta_1, \delta_2 > \underline{\delta}$,

there is an equilibrium s.t. P1's payoff is within an ε*-neighborhood of* $\boldsymbol{u}$.

When both players' discount factors are close to 1, every payoff that is *strictly* IR and IC can be approximately attained in the discounted game.

- What are the connections between $\mathscr{U}$ and $\widehat{\mathscr{U}}$?
- Is either the ε approximation or the strict IR/IC conditions redundant?

Follow-up works

Hörner and Lovo (2009)

- Characterize *belief-free equilibrium payoffs* when $\delta_1 = \delta_2 \to 1$.
- Allows for two-sided private information and interdependent values.

Hörner, Lovo and Tomala (2011) generalize it to three or more players.

Peski (2014)

- Focus on private value games. Allows for two-sided private info.
- Characterize the equilibrium payoff set when $\delta_1 = \delta_2 \to 1$.

An open question: What if one of the players is not very patient?

- Pei (2021 TE): Monotone-supermodular games.
- Pei (Working Paper): Provide strategic foundations for the sender's commitment in Bayesian persuasion models.

Challenge: Figure out what the additional constraints are.

Example

Time: $t = 0, 1, 2, \ldots$

A long-lived P1, discount factor $\delta \in (0,1)$ *vs* a sequence of myopic P2s.

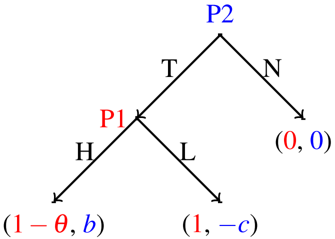

in which $b > 0$ and $c > 0$, let $\gamma^* \equiv \frac{c}{b+c}$

The terminal node in each period is perfectly observed.

P1 has perfectly persistent private information about θ:

- $\theta \in \Theta \equiv \{\theta_1, \ldots, \theta_m\} \subset [0,1)$, with $0 < \theta_1 < \theta_2 < \ldots < \theta_m < 1$.
- P2's full support prior $\pi \in \Delta(\Theta)$.

Result: P1's Highest Equilibrium Payoff

For every $j \in \{1,2,...,m\}$, let

$$v_j^* \equiv \underbrace{(1-\gamma^*\theta_j)}_{\text{Type } \theta_j\text{'s Stackelberg payoff}} \times \underbrace{\frac{1-\theta_1}{1-\gamma^*\theta_1}}_{\text{incomplete information multiplier } (<1)} .$$

Theorem 1: Highest Equilibrium Payoff

For every $\varepsilon > 0$, there exists $\underline{\delta} \in (0,1)$ s.t. when $\delta > \underline{\delta}$,

1. *$\exists$ sequential equilibrium s.t. P1's payoff is within ε of $(v_1^*,...,v_m^*)$.*
2. *$\nexists$ BNE s.t. type θ_1's payoff is more than v_1^*.*

   $\nexists$ BNE and $j \in \{2,...,m\}$, s.t. type θ_j's payoff is more than $v_j^* + \varepsilon$.

Lessons from Theorem 1

$$v_j^* \equiv \underbrace{(1-\gamma^*\theta_j)}_{\text{Type } \theta_j\text{'s Stackelberg payoff}} \times \underbrace{\frac{1-\theta_1}{1-\gamma^*\theta_1}}_{\text{incomplete information multiplier } (<1)} .$$

1. Type θ_j's highest equilibrium payoff only depends on:
   (a) His own cost of playing H: θ_j.
   (b) The *lowest* cost in the support of P2's prior belief.

   The multiplier $\frac{1-\theta_1}{1-\gamma^*\theta_1}$ converges to 1 as $\theta_1 \downarrow 0$.

2. Type θ_1's payoff is no more than his highest equilibrium payoff in the repeated complete information game.

3. Types θ_2 to θ_m can strictly benefit from incomplete information.

$$v_j^* > \underbrace{1-\theta_j}_{\text{highest payoff under complete info}}, \quad \text{for all } j \neq 1.$$

Understand the Formula $v_j^* = (1-\gamma^*\theta_j)\cdot\frac{1-\theta_1}{1-\gamma^*\theta_1}$

v_j^* is the value of the following constrained optimization problem:

$$\max_{\mathscr{Y}\in\Delta\{N,H,L\}}\left\{(1-\theta_j)\underbrace{\mathscr{Y}(H)}_{\text{prob of terminal node } H}+\underbrace{\mathscr{Y}(L)}_{\text{prob of terminal node } L}\right\},$$

subject to:

$$(1-\theta_1)\mathscr{Y}(H)+\mathscr{Y}(L)\leq 1-\theta_1,$$

and

$$\frac{\mathscr{Y}(H)}{\mathscr{Y}(L)}\geq\frac{\gamma^*}{1-\gamma^*}.$$

Two Type Example

Each $\mathscr{Y}$ is captured by a point in the yellow set.

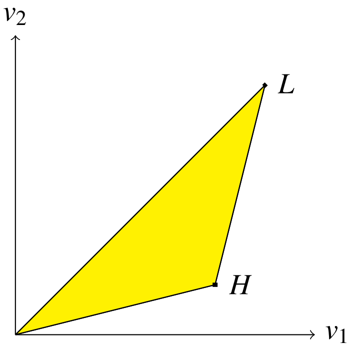

Constraint 1

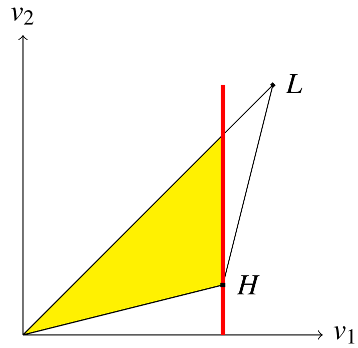

$$(1-\theta_1)\mathscr{Y}(H)+\mathscr{Y}(L)\leq 1-\theta_1$$

Constraint 2

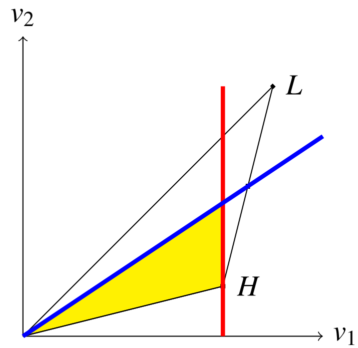

$$\frac{\mathscr{Y}(H)}{\mathscr{Y}(L)} \geq \frac{\gamma^*}{1-\gamma^*}$$

Proof of Theorem 1: Overview

For any equilibrium $\sigma \equiv \Big((\sigma_\theta)_{\theta\in\Theta}, \sigma_2\Big)$, and any $\theta_j \in \{\theta_1, ..., \theta_m\}$:

- Let $\mathscr{Y}^j \in \Delta\{N, H, L\}$ be the *discounted average frequency of terminal nodes* induced by $(\sigma_{\theta_j}, \sigma_2)$.

$$\mathscr{Y}^j(y) \equiv \mathbb{E}^{(\sigma_{\theta_j}, \sigma_2)}\Big[\sum_{t=0}^{\infty}(1-\delta)\delta^t \mathbf{1}\{y_t = y\}\Big] \text{ for every } y \in \{N, H, L\}.$$

We know that:

1. Type θ_j's equilibrium payoff equals his expected payoff from $\mathscr{Y}^j$.
2. Type θ_1's expected payoff from $\mathscr{Y}^j \leq$ type θ_1's equilibrium payoff.

The proof consists of two steps.

- For every $j \in \{1, ..., m\}$, $\mathscr{Y}^j$ must satisfy the two constraints
$\Rightarrow$ Type θ_j's payoff cannot exceed v_j^*.
- Construct an equilibrium that approximately attains $(v_1^*, ..., v_m^*)$.

Proof: Type θ_j's Payoff $\leq v_j^*$

Lemma 1

Type θ_1's payoff in any BNE is no more than $1-\theta_1$.

Since type θ_1's expected payoff from $\mathscr{Y}^j \leq$ type θ_1's equilibrium payoff, this lemma implies that $\mathscr{Y}^j$ must satisfy the first constraint.

Lemma 2

For every $\varepsilon > 0$, there exists $\underline{\delta} \in (0,1)$ such that when $\delta > \underline{\delta}$, for every BNE and for every $j \in \{1,2,\ldots,m\}$,

$$\frac{\mathscr{Y}^j(H)}{\mathscr{Y}^j(L)} \geq \frac{\gamma^* - \varepsilon}{1-(\gamma^* - \varepsilon)}.$$

Apply Gossner's learning argument to type θ_j's equilibrium strategy.

- Conditional on P2 plays T, σ_{θ_j} plays H with prob at least γ^* in all except for a bounded number of periods.

Proof of Lemma 1: Type θ_1's Payoff $\leq 1 - \theta_1$

Induction on the number of types in the support of P2's posterior.

- One type: Direct implication of Fudenberg, Kreps and Maskin (90).
- Suppose lowest-cost type θ's payoff is no more than $1 - \theta$ when there are $\leq n - 1$ types, what happens when there are n types?

Let θ be the lowest-cost type. Partition on-path histories into 3 subsets:

1. P2 plays N with prob 1,
2. P2 plays T with positive prob and type θ plays H with positive prob,
3. P2 plays T with positive prob and type θ plays H with zero prob.

The following strategy is type θ's best reply to σ_2:

- Until reaching a Class 3 history, plays σ_θ at Class 1 histories and plays H for sure at Class 2 histories.

Proof of Lemma 1: Type θ_1's Payoff $\leq 1 - \theta_1$

Three classes of histories:

1. P2 plays N with prob 1,
2. P2 plays T with positive prob and type θ plays H with positive prob,
3. P2 plays T with positive prob and type θ plays H with zero prob.

Consider the following best reply of type θ to σ_2:

- Until reaching a Class 3 history, plays σ_θ at Class 1 histories and plays H for sure at Class 2 histories.

Class 1 histories: Type θ's stage-game payoff is 0.

Class 2 histories: Type θ's stage-game payoff is at most $1 - \theta$.

Proof: Type θ_1's Payoff $\leq 1 - \theta_1$

What about type θ's continuation value at Class 3 history?

3. P2 plays T with positive prob and type θ plays H with 0 prob.

Class 3 history h^t**:** $\exists$ another type that plays H with positive prob.

- After observing H at h^t, at most $n-1$ types in the support of P2's posterior belief.
- Induction hypothesis $\Rightarrow$ there exists some type $\theta'(>\theta)$ whose continuation value at h^t is at most $1-\theta'$.
- If type θ' imitate type θ's strategy starting from h^t, then type θ' receives continuation value is at least:

$$\text{Type } \theta\text{'s Continuation Value at } h^t + (\theta - \theta') \quad (\leq 1 - \theta')$$

$\Rightarrow$ Type θ's continuation value at h^t is at most $1-\theta$.

Proof: Type θ_1's Payoff $\leq 1 - \theta_1$

Suppose type θ plays the following best reply to σ_2:

- Until reaching a Class 3 history, plays σ_θ at Class 1 histories and plays H for sure at Class 2 histories.

Then

- His stage-game payoff $\leq 1 - \theta$ before play reaches a Class 3 history.
- His continuation value is $\leq 1 - \theta$ after reaching a Class 3 history.

Type θ's payoff is no more than $1 - \theta$ at a history where:

- He is the lowest cost type in the support of P2's posterior.
- There are at most n types in the support of P2's posterior.

Construct equilibria that approximately attain v^*

How to approximately attain payoff $v^* \equiv (v_1^*, \ldots, v_m^*)$?

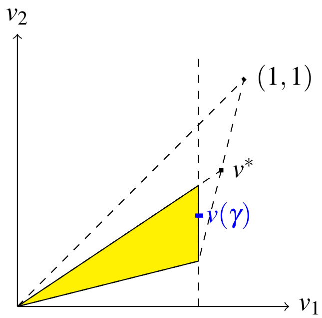

Lemma 1

For every $\eta \in (0,1)$ and $v(\gamma)$, there exists $\underline{\delta} \in (0,1)$, such that for every $\delta > \underline{\delta}$ and $\pi_0(\theta_1) \geq \eta$, there exists an equilibrium in which P1's payoff is $v(\gamma)$.

Overview: Equilibrium Construction

Types $\theta_2 \sim \theta_m$ adopt the same strategy. Type θ_1 plays differently.

Keeps track of the following state variables:

- P1's reputation: Prob of $\theta = \theta_1$, denoted by $\eta(h^t)$.
- P1's continuation value.

Three-phase equilibrium:

- One learning phase: P2 plays T all the time & slowly learns θ.
- Two absorbing phases: Learning about θ stops.

Phase transition happens with positive prob at h^t:

either P2 believes that $\theta = \theta_1$ occurs with prob 1.

or convex weight of L in P1's continuation value $< 1 - \delta$.

Players' Strategies in the Learning Phase

- P2 plays T in every period.
- Let $\eta^* \equiv \gamma^* \eta(h^0)$ and let

$$\Delta(\cdot) \equiv \eta(\cdot) - \eta^*.$$

- Each type of P1's mixed action at h^t is pinned down by:

$$\Delta(h^t, L) = (1 - \lambda\gamma^*)\Delta(h^t),$$

and

$$\Delta(h^t, H) = \min\left\{1 - \eta^*, \Big(1 + \lambda(1 - \gamma^*)\Big)\Delta(h^t)\right\}.$$

- $\lambda > 0$ measures the speed of learning.

Why this Belief Updating Formula?

1. Conditional on remaining in the learning phase, P1's reputation depends only on the number of times H and L have been played.

2. Respect P2's incentive constraint at each learning phase history:
   - Relative speed of reputation increases is low enough. H is played with probability at least γ^*.

3. For every $\widetilde{\gamma} > \gamma^*$, there exists $\lambda > 0$ small enough s.t.

$$\Big(1 + \lambda(1 - \gamma^*)\Big)^{\widetilde{\gamma}}\Big(1 - \lambda\gamma^*\Big)^{1-\widetilde{\gamma}} > 1.$$

Intuition Behind Belief Updating Formula

For every $\widetilde{\gamma} > \gamma^*$, there exists λ small enough s.t.

$$\Big(1+\lambda(1-\gamma^*)\Big)^{\widetilde{\gamma}}\Big(1-\lambda\gamma^*\Big)^{1-\widetilde{\gamma}} > 1.$$

Recall that we want to attain payoff $v(\gamma)$ for some $\gamma > \gamma^*$.

- P1's continuation value increases if frequency of $H > \gamma$.
- P1's continuation value decreases if frequency of $H < \gamma$.

Issue: P1's continuation value explodes.

Solution: If P1 plays H too frequently, then his reputation reaches 1 and play reaches the absorbing phase.

Red formula:

- We can find $\widetilde{\gamma} \in (\gamma^*, \gamma)$ and $\lambda > 0$ such that if P1 plays H with frequency above $\widetilde{\gamma}$, then P1's reputation increases.

First Absorbing Phase: P1's Reputation Reaches 1

Type θ_j's continuation payoff is $v_1(h^t)\frac{1-\theta_j}{1-\theta_1}$ for every j.

- randomize between terminal nodes N and L.

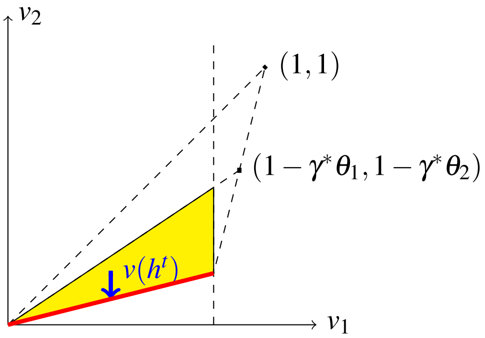

Type $\theta_j \neq \theta_1$ has no incentive to reach this phase because the continuation payoff is always above the red line (due to 2nd absorbing phase).

Second Absorbing Phase: P1 (nearly) depletes L

Transits to 2nd absorbing phase with positive probability if

- Weight of L in $v(h^t)$ is $< 1 - \delta$.

When this convex weight $= 0$,

- Transition happens with probability 1.
- Deliver $v(h^t)$ by randomizing between N and H.

When this convex weight $\in (0, 1 - \delta)$,

- Technical complication without public randomization.

The Set of Limiting Equilibrium Payoffs V^*

Let $V^* \subset \mathbb{R}^m$ be such that for every $(v_1, v_2, \ldots, v_m) \in V^*$, there exist $\mathscr{Y}_1, \ldots, \mathscr{Y}_m \in \Delta\{N, H, L\}$, such that:

$$\mathbb{E}_{\mathscr{Y}_j}[u_1(\theta_j, y)] = v_j \text{ for every } j \in \{1, 2, \ldots, m\},$$

$$\frac{\mathscr{Y}_j(H)}{\mathscr{Y}_j(L)} \geq \frac{\gamma^*}{1-\gamma^*} \text{ for every } j \in \{1, 2, \ldots, m\},$$

$$\mathbb{E}_{\mathscr{Y}_j}[u_1(\theta_j, y)] \geq \mathbb{E}_{\mathscr{Y}_k}[u_1(\theta_j, y)] \text{ for every } j, k \in \{1, 2, \ldots, m\},$$

$$\mathbb{E}_{\mathscr{Y}_1}[u_1(\theta_1, y)] \leq 1 - \theta_1.$$

Theorem 1': Limiting Equilibrium Payoff Set

For every $v \in int(V^*)$, there exists $\underline{\delta} \in (0,1)$ such that when $\delta > \underline{\delta}$, there exists a sequential equilibrium s.t. P1's payoff is v.

For every $v \in ext(V^*)$, there exists $\underline{\delta} \in (0,1)$ such that when $\delta > \underline{\delta}$, there exists no BNE s.t. P1's payoff is v.

Next Step: Connect this model to reputation models

Fudenberg and Levine's approach to study repeated incomplete info games:

- Assume the behavior of one type,
- Study the common properties of P1's equilibrium payoff.

Let's approach the problem from a different perspective:

- All types' behaviors are endogenous.
- Characterize patient P1's highest equilibrium payoff.
- Focus on P1 optimal equilibria and study the common properties of P1's behavior in those equilibria.

Advantages/disadvantages of the rational type approach:

- Proof is constructive, better understanding of behavior, which commitment behaviors are more reasonable.
- Only have results for a class of games, cannot refine equilibria when P1's actions are identified.

Equilibrium Behavior

How does P1 behave in equilibria that approximately attain v^*?

Type θ player 1's strategy $\sigma_\theta : \mathscr{H} \to \Delta(A_1)$.

Player 2's strategy $\sigma_2 : \mathscr{H} \to \Delta(A_2)$.

Under strategy profile $\left((\sigma_\theta)_{\theta\in\Theta}, \sigma_2\right)$,

- σ_θ is *stationary* if $\sigma_\theta(h) = \sigma_\theta(h')$ for every h, h' that

  occur with positive probability under $(\sigma_\theta, \sigma_2)$,
  and $\sigma_2(h)$, $\sigma_2(h')$ attach positive prob to T.

- σ_θ is *completely mixed* if $\sigma_\theta(h)$ is nontrivially mixed for every h that

  occurs with positive probability under $(\sigma_\theta, \sigma_2)$,
  and $\sigma_2(h)$ attaches positive prob to T.

Equilibrium Behavior

How does P1 behave in equilibria that approximately attain v^*?

Recall that $\Theta \equiv \{\theta_1, \theta_2, \ldots, \theta_m\}$.

Theorem 2 (Nonstationary Equilibrium Behavior)

When $m \geq 2$, for every small enough $\varepsilon > 0$, there exists $\overline{\delta}$, s.t. if $\delta > \overline{\delta}$, in any BNE that attains payoff within ε of $(v_1^*, \ldots, v_m^*)$.

1. *no type of P1 uses stationary strategies or completely mixed strategies*
2. *no type of P1 has a completely mixed equilibrium best reply.*

This is also true when $\theta_1 = 0$.

- No matter how low P1's cost is,

  his equilibrium behavior must depend nontrivially on past play.

Proof of Theorem 2

Suppose toward a contradiction that there exists a BNE s.t.

1. P1's payoff is within ε of v^*.
2. Some type θ_j plays a non-trivially mixed action at every history.

$\Rightarrow$ Playing L at every history is type θ_j's best reply.

Playing H at every history is type θ_j's best reply.

$\Rightarrow$ $\forall\, k > j$, type θ_k plays L w.p. 1 at every *on-path* history.

$\forall\, i < j$, type θ_i plays H w.p. 1 at every *on-path* history.

However, none of these stationary pure strategies are consistent with P1's equilibrium payoff is approximately v^*.

Proof of Theorem 2, Continued...

Suppose some type θ_k plays L at every on-path history,

- P2 will learn type θ_k's strategy in bounded number of periods, after which they will play N.
- Type θ_k's payoff is close to 0 as $\delta \to 1$.
- This leads to a contradiction.

Suppose types θ_1 to θ_{j-1} play H at every on-path history, with $j \geq 2$.

- Type θ_j's long-term payoff cannot exceed $(1-\delta)+\delta(1-\theta_j)$. Separated from all lower types after playing L for one period.
- However, $(1-\delta)+\delta(1-\theta_j)$ is strictly less than v_j^* when δ is large.
- This leads to a contradiction.

□

Behavioral Prediction in Binary Action Games

How does P1 behave in equilibria that approximately attain v^*?

Theorem 3 (Equilibrium Action Frequencies)

When $m \geq 2$, for every small enough $\varepsilon > 0$, there exists $\overline{\delta}$, s.t. if $\delta > \overline{\delta}$, in every BNE $\Big((\sigma_\theta)_{\theta\in\Theta}, \sigma_2\Big)$ that attains payoff within ε of $(v_1^*, \ldots, v_m^*)$.

1. *For every $\theta \neq \theta_m$, and for every best reply $\widehat{\sigma}_\theta$ of type θ against σ_2,*

$$\frac{\mathbb{E}^{(\widehat{\sigma}_\theta,\sigma_2)}\left[\sum_{t=0}^{\infty}(1-\delta)\delta^t \mathbf{1}\{y_t = H\}\right]}{\mathbb{E}^{(\widehat{\sigma}_\theta,\sigma_2)}\left[\sum_{t=0}^{\infty}(1-\delta)\delta^t \mathbf{1}\{y_t = L\}\right]} \geq \frac{\gamma^* - \varepsilon}{1-(\gamma^* - \varepsilon)}.$$

Behavioral Prediction in Binary Action Games

How does P1 behave in equilibria that approximately attain v^*?

Theorem 3 (Equilibrium Action Frequencies)

When $m \geq 2$, for every small enough $\varepsilon > 0$, there exists $\overline{\delta}$, s.t. if $\delta > \overline{\delta}$, in every BNE $\Big((\sigma_\theta)_{\theta \in \Theta}, \sigma_2\Big)$ that attains payoff within ε of $(v_1^*, \ldots, v_m^*)$.

2. *For every $\theta \neq \theta_1$, and for every best reply $\widehat{\sigma}_\theta$ of type θ against σ_2,*

$$\frac{\mathbb{E}^{(\widehat{\sigma}_\theta, \sigma_2)}\left[\sum_{t=0}^{\infty}(1-\delta)\delta^t \mathbf{1}\{y_t = H\}\right]}{\mathbb{E}^{(\widehat{\sigma}_\theta, \sigma_2)}\left[\sum_{t=0}^{\infty}(1-\delta)\delta^t \mathbf{1}\{y_t = L\}\right]} \leq \frac{\gamma^* + \varepsilon}{1-(\gamma^* + \varepsilon)}.$$

Behavioral Prediction in Binary Action Games

How does P1 behave in equilibria that approximately attain v^*?

The two statements of Theorem 3:

1. For every best reply of type $\theta \neq \theta_m$, frequency ratio between H and L is more than $\frac{\gamma^* - \varepsilon}{1-(\gamma^* - \varepsilon)}$.
2. For every best reply of type $\theta \neq \theta_1$, frequency ratio between H and L is less than $\frac{\gamma^* + \varepsilon}{1-(\gamma^* + \varepsilon)}$.

Implications:

1. Pin down the action frequencies of all types except for θ_1 and θ_m.
2. Applies to all pure strategy best replies.

Repeated Communication Games

Repeated communication games with private lying cost.

- Sender has persistent private info about her lying cost.
- Sender private observes i.i.d. state $\omega_t \in \Omega$.
- Sender sends message $m_t \in \Omega$ to the receiver.
- Receiver takes an action $a_t \in A$.
- Period t receiver observes $\{a_s, m_s, \omega_s\}_{s=0}^{t-1}$ and m_t.

Sender's stage-game payoff is $u_s(\omega_t, a_t) - C \cdot \mathbf{1}\{m_t \neq \omega_t\}$, where $C \in \{C_1, ..., C_n\}$ is the sender's persistent private info.

Receiver's stage-game payoff $u_r(\omega_t, a_t)$.

Another Application: Repeated Communication

Pei (2021): Stage-game payoffs follows from the leading example in KG.

- Characterize every type of patient sender's highest equilibrium payoff.

  Conditions s.t. highest payoff $\approx$ Bayesian persuasion payoff.

  A microfoundation for the commitment assumption in Bayesian persuasion games in models without any commitment.

- No rational type uses the optimal disclosure policy in every period, no matter how large the lying cost is.

  Stands in contrast to Mathevet et al. (2019)

- The possibility of having a high lying cost can hurt some type of sender who has a low lying cost.

  A novel outside option effect.

Next Lecture: Social Learning

- Banerjee (1992)
- Bikhchandani, Hirshleifer and Welch (1992)
- Smith and Sørensen (2000)