Lecture 1: Commitment Payoff Theorem

Long-Run Short-Run Models with Perfect Monitoring

Harry PEI
Department of Economics, Northwestern University

Spring Quarter, 2023

What is this course about?

1. Repeated games and reputation effects.
   - What will happen when players can build reputations?
   - When will reputations work and when will they break down?
2. Bargaining under incomplete information.
   - Can people trade efficiently when they have private information?
3. Bayesian social learning (Ben will teach DeGroot learning).
   - When can observational learning aggregate information?
4. Sustaining cooperation with limited information.
   - Community enforcement.
   - Repeated games with limited memories.

- Instructor: Harry Pei.

  Email: harrydp@northwestern.edu

  Office: KGH 3391.

- Lectures: Tuesdays and Thursdays, 3:30-5:20 pm, KGH 3301.

- Proposal: Lectures on Thursdays 3:30-7:30 pm, KGH 3301.

- We have two guest lectures:

  Krishna Dasaratha on May 4th: Behavioral social learning.

  Daniel Clark on June 1st: Community enforcement.

  Zoom Meeting ID: 915 7004 8164.

- Office hour: by appointment.

Evaluation

Requirements:

- 1 hour presentation: a paper I suggested in class, or a paper you propose, or an on-going work of yours (theory/empirics).

  If you present other people's paper, you should read it carefully and critically, and provide *thoughtful* comments.

- Write up: solve a problem I mentioned in class, or solve a problem you come up by yourself, or write a literature review, or submit a research proposal, or submit an on-going work of yours.

Unsolicited advice:

- Grades don't matter in grad school.
- Go to seminars (theory, strategy, political econ, finance, macro).
- Try to find opportunities to present and to talk about your work.
- Don't do anything only for the sake of pleasing your advisors.

Rules for my Lectures

- **Please ask questions**.
- Let me know if I am going too slow or too fast.
- Interrupt me if:
  - I made a mistake (highly likely)
  - There is something unclear,
  - There is something you don't understand.
- Email me if you have suggestions.

What is a reputation?

Google: A widespread belief that someone or something has a particular habit or characteristic.

Two approaches to study reputations:

1. The habit view: Players convince their opponents that they will behave in a particular way (e.g., always cooperate, tit-for-tat).
2. The characteristic view: Players signal payoff-relevant characteristics over time (e.g., low production cost, high ability, high quality).

Similarity: Dynamic games with incomplete information, one informed player facing one/multiple uninformed opponent(s).

Difference: Nature of the informed player's private info.

We will start from the habit view and might move to the characteristic view.

Intellectual History: The Chainstore Paradox

- A monopolist has branches in $T \in \mathbb{N}$ locations, with T finite.

  He faces *one potential competitor in each location*.

- In period $s \in \{1, 2, ..., T\}$, the monopolist plays against the competitor in the s-th location.

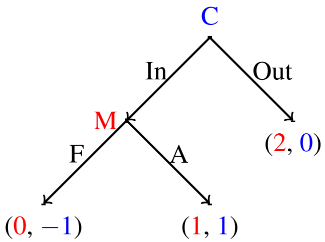

- Monopolist's total payoff is the sum of payoffs in T locations.
- Every competitor perfectly observes all actions chosen before.

The Chainstore Paradox

There is a unique subgame perfect equilibrium:

- Every competitor chooses *In* and monopolist chooses *Accommodate*.

What is wrong with this prediction?

- No matter how long the time horizon is, the monopolist never fights.
- Even if a competitor observes the monopolist fighting the past 1000 entrants, he still believes that he will be accommodated with prob 1.

Something is missing in complete information game repeated games.

Intellectual History: Commitment Type Models

How to fix this? *Gang of four*.

- Kreps and Wilson (1982), Milgrom and Roberts (1982).

Idea: Perturb the game with a small prob of commitment type.

- With probability $\varepsilon > 0$, the monopolist is *irrational*, doesn't care about payoffs, and mechanically fights in every period.
- With probability $1 - \varepsilon$, the monopolist is *rational*, maximizes the sum of his payoffs across periods.

Result: Gang of Four

Theorem: Gang of Four

For every $\varepsilon > 0$, there exists $T^* \in \mathbb{N}$ such that if $T \geq T^*$,

then on the equilibrium path of every sequential equilibrium,

- *The rational monopolist chooses F & each potential entrant chooses Out in all except for the last T^* periods*

Proof: Backward induction.

Takeaway: The option to build reputations can dramatically affect patient players' incentives and behaviors.

Proof: Take Home Exercise

Bonus Question: Figure out why sequential equilibrium

Proof Idea: Characterize the equilibrium via backward induction.

Equilibrium Behavior:

- In the first $T - t^*$ periods, rational incumbent plays F and entrant stays out. No learning takes place.
- In the last t^* periods, entrant enters with positive prob, rational incumbent mixes between F and A. Learning happens gradually.

  Probability of entry makes the rational incumbent indifferent, and rational incumbent's mixing probability makes the entrant indifferent.

Establish Uniqueness of Sequential Equilibrium Outcome: Pin down the entrants' on-path beliefs in the last few periods.

Robustness of the Gang of Four Insight?

Gang of four result requires:

- Finite horizon and backward induction.
- Particular stage-game payoff functions.
- Entrant can perfectly observe monopolist's action.
- Sequential equilibrium.

Robustness of the Gang of Four Insight?

Another concern: Does it rely on the specification of incomplete info?

- Let $G = (N, A, u)$ be an n-player normal form game.
- Let $\alpha^* \in \times_{i=1}^n \Delta(A_i)$ be a stage-game NE with payoff $\mathbf{w} \in \mathbb{R}^n$.

Folk Theorem under Incomplete Information: Fudenberg and Maskin (1986)

For any $\varepsilon > 0$ and any payoff vector $\mathbf{v} > \mathbf{w}$, there exists $T^* \in \mathbb{N}$ such that for any $T > T^*$, there exists a strategy profile $\{s_i\}_{i \in N}$ such that in the T-fold repetition of G with public randomization where each player i is rational with probability $1 - \varepsilon$ and is committed to s_i with probability ε, there is an equilibrium where players' average payoff is within ε of $\mathbf{v}$.

Not directly applicable to gang-of-four since it requires n long-run players.

- Takeaway: Predictions sensitive to the specification of incomplete info?

Proof: Fudenberg and Maskin's Folk Theorem

For every player i, define strategy s_i as:

- plays according to $\alpha \in \Delta(A)$ if everyone played according to α before, and plays α_i^* otherwise.

where α is a mixed action profile that gives players payoff $\mathbf{v}$.

Consider an auxiliary repeated game where:

- After any normal-type player deviates from α, they can only play α^*.

We show that every equilibrium in the auxiliary game satisfies:

- Each player i prefers to follow s_i except for the last few periods.

Proof: Fudenberg and Maskin's Folk Theorem

We show that every equilibrium in the auxiliary game satisfies:

- Each player i prefers to follow s_i except for the last few periods.

Suppose there are T^* periods left and I consider whether to deviate:

- **My loss from deviation:** When all opponents are committed, I will be punished for at least for T^* periods.
- **My gain from deviation:** I can gain in at most one period, since all players revert to static Nash afterwards.

When T^* is large relative to $\frac{1}{\varepsilon^{N-1}}$, each player prefers to follow s_i in all except for the last T^* periods.

Proof: Fudenberg and Maskin's Folk Theorem

Fix any equilibrium in the auxiliary game where:

- If any normal-type player deviates from α, he can only play α^*.

This remains an equilibrium in the original game.

Because α^* is a stage-game Nash, i.e., no player can do better when others play α^*.

Let

$$T^* \equiv \left\lceil \max_i \frac{\overline{v}_i - (1-\varepsilon^{n-1})\underline{v}_i}{\varepsilon^{n-1}(v_i - w_i)} \right\rceil .$$

Take $T \gg T^*$, players' average payoffs in these equilibria are close to $\mathbf{v}$.

Lectures 1 and 2: Fudenberg and Levine (1989, 1992)

Extend the gang of four insights to

- environments with an infinite horizon.
- general stage game payoffs.
- imperfect monitoring.
- weaker solution concepts.
- not sensitive to the details of incomplete info.

I will present all results in games with an infinite horizon.

- Their results also apply to games with long but finite horizon.

Infinitely Repeated Game with One Long-Run Player

- Time: $t = 0, 1, 2, \ldots$
- Long-lived player 1 (P1) *vs* a sequence of short-lived player 2s (P2).

  (alternative interpretation: P2 is a continuum of small players)

- Players simultaneously choose their actions $a_1 \in A_1$ and $a_2 \in A_2$.

  Actions in period t: $a_{1,t} \in A_1$ and $a_{2,t} \in A_2$.

- Stage-game payoffs: $u_1(a_{1,t}, a_{2,t})$, $u_2(a_{1,t}, a_{2,t})$.

  P1's *discounted average payoff*: $\sum_{t=0}^{\infty}(1-\delta)\delta^t u_1(a_{1,t}, a_{2,t})$.

- Public signal in period t: $y_t \in Y$,

  which is distributed according to $\rho(\cdot|a_{1,t}, a_{2,t}) \in \Delta(Y)$.

Introducing Commitment Types

P1 has a perfectly persistent type $\omega \in \Omega \equiv \{\omega^r\} \bigcup \Omega^m$.

1. ω^r denotes the *rational type*, who can flexibly choose his actions in order to maximize his discounted average payoff.

2. Each $\alpha_1^* \in \Omega^m \subset \Delta(A_1)$ represents a *commitment type*,

   who does not care about payoffs and plays α_1^* in every period.

P2's prior belief: $\pi \in \Delta(\Omega)$.

What can players observe?

- Player 1's history: $h_1^t \in \mathcal{H}_1^t \equiv \Omega \times \{A_1 \times Y\}^t$.
- Player 2's history: $h_2^t \in \mathcal{H}_2^t \equiv Y^t$.

Assumptions: A_1, A_2, Y and Ω^m are finite, π has full support.

Commitment Payoff Theorem: Perfect Monitoring

Let's make two simplifying assumptions:

1. Perfect monitoring: $Y = A_1 \times A_2$ and $\rho(a_1, a_2 | a_1, a_2) = 1$.
2. There exists a commitment type that plays a pure action $a_1^* \in A_1$.

For every commitment action $a_1^* \in \Omega^m$, P1's commitment payoff from a_1^*:

$$v_1^*(a_1^*) \equiv \min_{a_2 \in \text{BR}_2(a_1^*)} u_1(a_1^*, a_2).$$

Let $\underline{u}_1$ be P1's lowest stage-game payoff.

Commitment Payoff Theorem: Fudenberg and Levine (1989)

For every $\varepsilon > 0$, there exists $T \in \mathbb{N}$,

such that when π assigns prob more than ε to commitment type $a_1^* \in \Omega^m$,

rational P1's payoff in any Bayes Nash equilibrium is at least:

$$(1 - \delta^T)\underline{u}_1 + \delta^T v_1^*(a_1^*).$$

Commitment Payoff Theorem: $\delta \to 1$ Limit

Commitment Payoff Theorem: Fudenberg and Levine (1989)

For every $\varepsilon > 0$, there exists $T \in \mathbb{N}$,

such that when π assigns prob more than ε to commitment type $a_1^* \in \Omega^m$,

rational P1's payoff in any Bayes Nash equilibrium is at least:

$$(1-\delta^T)\underline{u}_1 + \delta^T v_1^*(a_1^*).$$

What happens when the informed player is patient, i.e., $\delta \to 1$?

- P1's payoff lower bound $\to v_1^*(a_1^*)$.
- Patient P1 receives at least his commitment payoff from a_1^*.

The payoff lower bound does not depend on the details of the type space.

- It only requires commitment type a_1^* to occur with positive prob.

Proof: Overview

Commitment Payoff Theorem: Fudenberg and Levine (1989)

For every $\varepsilon > 0$, there exists $T \in \mathbb{N}$,

such that when π assigns prob more than ε to commitment type $a_1^* \in \Omega^m$,

rational P1's payoff in any Bayes Nash Equilibrium is at least:

$$(1-\delta^T)\underline{u}_1 + \delta^T v_1^*(a_1^*).$$

Fix the parameters (π, δ). For every Bayes Nash Equilibrium (σ_1, σ_2),

- Consider rational-type P1's payoff

  if he deviates from σ_1 and mechanically plays a_1^* in every period.

- Let this payoff be U_1^*.
- By definition, rational P1's equilibrium payoff $\geq U_1^*$.

Proof: P1's payoff if he deviates and plays a_1^*

In every period,

- *either* P2's action is supported in $\mathrm{BR}_2(a_1^*)$.

  or P2 has an incentive to play actions outside $\mathrm{BR}_2(a_1^*)$.

In the 1st case, P1's stage-game payoff $\geq v_1^*(a_1^*)$.

In the 2nd case, there exists $\gamma > 0$ such that:

- P2 believes that a_1^* is played with prob less than $1-\gamma$ in that period.

  Such γ *depends only* on players' stage-game payoff functions.
- After P2 observes P1 plays a_1^* in that period, Bayes Rule suggests that:

$$\text{Posterior Prob of Type } a_1^* = \frac{(\text{Prior Prob of Type } a_1^*) \cdot \Pr(a_1^* | \text{type } a_1^*)}{\text{unconditional prob of } a_1^*}$$

$$\geq \frac{\text{Prior Prob of Type } a_1^*}{1-\gamma}.$$

- This can happen in at most $T \equiv \lceil \log \varepsilon / \log(1-\gamma) \rceil$ periods.

Proof: Wrap up

What is rational P1's payoff if he deviates and plays a_1^* in every period?

In periods where P2's action is supported in $\mathrm{BR}_2(a_1^*)$.

- P1's stage game payoff $\geq v_1^*(a_1^*)$.

In periods where P2's action is *not* supported in $\mathrm{BR}_2(a_1^*)$.

- P1 may receive low stage-game payoff,
- But there can be at most $T \equiv \lceil \log \varepsilon / \log(1-\gamma) \rceil$ such periods.

Lower bound on rational P1's payoff from playing a_1^* in every period:

$$(1-\delta^T)\underline{u}_1 + \delta^T v_1^*(a_1^*).$$

This is also a lower bound for the rational-type P1's equilibrium payoff.

How to Interpret Commitment Type?

Commitment type(s) capture the intuition that:

- Once we observe a player behaving in certain ways for a long time, we tend to believe that they will behave similarly in the future.
- This logic is missing in complete information game models.
- Commitment type is a modeling device that can capture this logic.

The proof captures this logic:

- Either P2 believes that P1 will play a_1^* and best replies to a_1^*.
- Or P2 does not believe that P1 will play a_1^*,

  but after observing P1 plays a_1^*, she will be *surprised*

  and the probability she assigns to commitment type a_1^* increases.

Refinement for Repeated Complete Info Games

Fudenberg, Kreps and Maskin (1990): Folk theorem under complete information

The set of P2's mixed strategy best replies:

$$\mathcal{A}_2 \equiv \{\alpha_2 \in \Delta(A_2) | \alpha_2 \text{ best replies against some } \alpha_1 \in \Delta(A_1)\}$$

Patient P1's lowest equilibrium payoff:

$$v^{\min} \equiv \min_{\alpha_2 \in \mathcal{A}_2} \max_{a_1 \in A_1} u_1(a_1, \alpha_2).$$

P1's highest equilibrium payoff:

$$v^{\max} \equiv \max_{\{(\alpha_1,\alpha_2) s.t. \alpha_2 \in \mathrm{BR}_2(\alpha_1)\}} \min_{a_1 \in \mathrm{supp}(\alpha_1)} u_1(a_1, \alpha_2).$$

In many games of interest, the option to build a reputation selects a subset of high payoffs for P1. Sometimes, it selects P1's highest equilibrium payoff.

Product Choice Game (Mailath and Samuelson 2001)

A firm (P1) and a sequence of consumers (P2s).

$-$	T	N
H	$2, 1$	$-1, 0$
L	$3, -1$	$0, 0$

Repeated complete information game:

- P1's payoff can be anything within $[0, 2]$.

Positive prob of commitment type that mechanically plays H.

- Rational firm guarantees payoff ≈ 2 in every BNE.

Some Common Misunderstandings

1. Can rational P1 convince P2s that he is a commitment type?

   Not with high prob on the equilibrium path! Belief is a martingale.

   Example: Think about a pooling equilibrium.

2. Will the rational-type P1 build a reputation?

   Not necessarily in the infinite horizon game. He may find it strictly optimal to separate from the commitment type in period 0.

3. Does it say much about the short-run players' welfare?

   No. Because rational-type P1's behavior cannot be pinned down.

Predictions on P1's Behavior?

Suppose there is a commitment type that plays P1's optimal pure commitment action a_1^* in every period, then

- What's the frequency with which the rational-type P1 plays a_1^*?

$$X^{(\sigma_1,\sigma_2)}(a_1^*) \equiv \mathbb{E}^{(\sigma_1,\sigma_2)}\left[\sum_{t=0}^{\infty}(1-\delta)\delta^t \mathbf{1}\{a_{1,t}=a_1^*\}\right]$$

Li and Pei (2021): In many games of interest, any action frequency that is compatible with

- P1 receiving payoff at least $v_1(a_1^*)$,
- P2's myopic incentives

can arise in some equilibria of the reputation game.

Li and Pei (2021)'s Theorem

Assumptions on stage-game payoffs:

- P1 has a unique optimal commitment action a_1^* and $\mathrm{BR}_2(a_1^*) = \{a_2^*\}$.
- $a_1^* \notin \mathrm{BR}_1(a_2^*)$.
- $u_1(a_1^*, a_2^*) > v^{\min} \equiv \min_{\alpha_2 \in \mathcal{A}_2} \max_{a_1 \in A_1} u_1(a_1, \alpha_2)$.

Let

$$F^*(u_1, u_2) \equiv \min_{(\alpha_1', \alpha_1'', a_2', a_2'', q) \in \Delta(A_1) \times \Delta(A_1) \times A_2 \times A_2 \times [0,1]} \left\{ q\alpha_1'(a_1^*) + (1-q)\alpha_1''(a_1^*) \right\},$$

subject to $a_2' \in \mathrm{BR}_2(\alpha_1')$, $a_2'' \in \mathrm{BR}_2(\alpha_1'')$, and

$$qu_1(\alpha_1', a_2') + (1-q)u_1(\alpha_1'', a_2'') \geq u_1(a_1^*, a_2^*).$$

Theorem: When δ is close enough to 1, rational-type P1's discounted frequency of playing a_1^* can be anything between $F^*(u_1, u_2)$ and 1.

What about long finite horizon? Sharper predictions?

Next Lecture

Commitment payoff theorem with imperfect monitoring:

- the public signals are noisy,
- commitment payoff from mixed commitment actions.

Papers to read:

- Fudenberg and Levine (1992), Gossner (2011).
- Kalai and Lehrer (1993), Sorin (1999).