Game Theory

Repeated games

Xiang Sun

2021 Summer

A Motivating Example

- Consider the following Prisoners' Dilemma problem:

Player 2

Player 1		L_2	R_2
	L_1	$1, 1$	$5, 0$
	R_1	$0, 5$	$4, 4$

- If the game is played once, the unique Nash equilibrium is (L_1, L_2).
- What if the game is played more than once? Will the cooperative outcome (R_1, R_2) be achieved through repeated interactions (重复互动)?

Introduction

- Long-term (or repeated) interactions are very common.
- Examples:
  - Firms are engaged in competition over time.
  - Most employment relationships last for a long time.
  - Countries compete over tariffs years by years.
- In a long-term relationship, one must consider how his/her current behavior will influence others' behavior in the future, or how threats or promises about future behavior can affect current behavior.
- In these dynamic situations, one might care about "reputation", which is often used to describe how a person's past actions affect future beliefs and behavior.

Introduction

- We use repeated games (重复博弈) to study such interactions among players.
- In repeated games, we are interested in how repeated interactions among players would affect their behavior.
- Two types of repeated games:
    - finitely repeated games;
    - infinitely repeated games.
- The results predicted by these two types of games differ dramatically.

Example

- Consider the following repeated game (i.e., two-stage Prisoners' Dilemma game):
  - The two players play the simultaneous-move game twice;
  - Each player observes the outcome of the first play before the second game begins;
  - The payoff of each player in the whole game is simply the sum of two payoffs in both stages (i.e., no discounting).
- This game is an example of the two-stage imperfect information games that we have learned before.

Example (Cont.)

- We can use backwards induction to solve the game.
- In stage 2, the unique Nash equilibrium is (L_1, L_2), in which each player receives 1.
- In stage 1, the two players play the following equivalent game:

		Player 2	
		L_2	R_2
Player 1	L_1	$2, 2$	$6, 1$
	R_1	$1, 6$	$5, 5$

- Hence, (L_1, L_2) is the unique Nash equilibrium in stage 1.
- The subgame-perfect outcome: (L_1, L_2) is played in both periods.
- What is the subgame-perfect Nash equilibrium?

Finitely Repeated Games

- Let $G = \langle A_1, \ldots, A_n; u_1, \ldots, u_n \rangle$ denote a static game of complete information in which players 1 through n simultaneously choose actions a_1 through a_n from the action spaces A_1 through A_n, and the payoffs are $u_1(a_1, \ldots, a_n)$ through $u_n(a_1, \ldots, a_n)$.
- The game G is called the stage game (阶段博弈) of the repeated game.

Definition

Given a stage game G, let $G(T)$ denote the finitely repeated game in which G is played T times, with the outcomes of all preceding plays observed before the next play begins. The payoffs for $G(T)$ are simply the sum of the payoffs from the T stage games.

Subgame-perfect Nash equilibrium

Proposition

If the stage game G has a unique Nash equilibrium, then for any finite T, the repeated game $G(T)$ has a unique subgame-perfect outcome: the Nash equilibrium of G is played in every stage.

- In the Prisoners' Dilemma example, the unique outcome in each period is (L_1, L_2) regardless of how many times the game is played.
- The result in the above proposition can be extended even if G itself is a dynamic game of complete information.

Multiple Nash equilibria

- What if the stage game G has multiple Nash equilibria?
- Then there may be subgame-perfect outcomes of the repeated game $G(T)$ in which, for any $t < T$, the outcome of stage t is not a Nash equilibrium of G.
- Consider the following game:

Player 2

Player 1	L_2	M_2	R_2
L_1	$1, 1$	$5, 0$	$0, 0$
M_1	$0, 5$	$4, 4$	$0, 0$
R_1	$0, 0$	$0, 0$	$3, 3$

- There are two Nash equilibria: (L_1, L_2) and (R_1, R_2).

Multiple Nash equilibria (Cont.)

- Suppose the game is repeated twice.
- The outcome in stage 2 is either (L_1, L_2) or (R_1, R_2).
- Is it possible that the first-stage outcome is (M_1, M_2) in a subgame-perfect Nash equilibrium?
- Consider, for example, player i's strategy:
  - play M_i in the first stage;
  - play R_i if the first-stage outcome is (M_1, M_2); otherwise, play L_i.
- It can be verified that the strategy profile constitutes a subgame-perfect Nash equilibrium, in which the first-stage outcome is (M_1, M_2).
- Key: R_i serves as a reward and L_i serves as a punishment;
  $\Rightarrow$ enforce (M_1, M_2) to be implemented.

Infinitely Repeated Games

What happens if the Prisoners' Dilemma game is played forever?

Present value

Definition

Let π_t be the payoff in stage t. Given the discount factor $\delta \in (0, 1)$, the present value (现值) of the infinite sequence of payoffs $\pi_1, \pi_2, \ldots$ is

$$\pi_1 + \delta\pi_2 + \delta^2\pi_3 + \cdots = \sum_{t=1}^{\infty} \delta^{t-1}\pi_t.$$

Infinitely Repeated Games

- Recall $G = \langle A_1, \ldots, A_n; u_1, \ldots, u_n \rangle$ is the stage game of repeated games.

Definition

Given a stage game G, let $G(\infty, \delta)$ denote the infinitely repeated game in which G is played forever and players share the discount factor δ.

- For each t, the outcomes of the $t-1$ preceding plays are observed before the t-th stage begins.
- Each player's payoff in $G(\infty, \delta)$ is the present value of the player's payoffs from the infinite sequence of stage games.

Infinitely Repeated Games

Consider the following infinitely repeated game of Prisoners' Dilemma:

- In stage 1, the two players play the stage game G and receive payoffs $\pi_{1,1}$ and $\pi_{2,1}$;
- In stage t, the players observe the actions chosen in the preceding $t-1$ stages, and then play G to receive $\pi_{1,t}$ and $\pi_{2,t}$;
- The payoff of the infinitely repeated game is the present value of the sequence of payoffs: $\sum_{t=1}^{\infty} \delta^{t-1}\pi_{i,t}$ for player $i = 1, 2$.

Strategies in Infinitely Repeated Games

- There are infinitely many strategies for the players.
- Some common strategies:
  1. noncooperative strategy:
     - play L_i in every stage.
  2. (grim) trigger strategy (触发策略):
     - play R_i in the first stage;
     - in stage t, if the outcome of all $t-1$ preceding stages has been (R_1, R_2), then play R_i; otherwise, play L_i.
  3. tit-for-tat (or tit for two tats) strategy (以牙还牙策略)
  4. carrot-and-stick strategy (or two-phase strategy) (胡萝卜加大棒策略)

Strategies in Infinitely Repeated Games

- We focus on the first two strategies.
- If both players adopt the noncooperative strategy, (L_1, L_2) is repeated forever.
- Using a trigger strategy, player *i* cooperates until someone fails to cooperate, which triggers a switch to noncooperation forever.
- If both players adopt the trigger strategy, then the outcome of the infinitely repeated game is (R_1, R_2) in every stage.
- Question: Is it a NE/SPNE in the infinitely repeated game where both players adopt the trigger strategy (i.e., cooperation is achieved)?

Nash Equilibria

Claim

Both players adopting the noncooperative strategy is a Nash equilibrium.

Proof.

- Assume player i plays L_i in every stage.
- Then player j's best response is also "to play L_j in every stage".

$\square$

Nash Equilibria

Claim

Both players adopting the trigger strategy is a Nash equilibrium if and only if $\delta \geq \frac{1}{4}$.

Nash Equilibria: Proof

- Assume player i has adopted the trigger strategy. We seek to show player j's best response is also to adopt the trigger strategy.
- It suffices to check when

$$\text{“follow trigger strategy”} \geq \text{“every deviations”}.$$

- 偏离触发策略的方式有很多，分为以下两类处理。
- Case 1: At the node where the outcome in a previous stage is not (R_1, R_2).

* Since player i plays L_i forever, player j's best response is also to play L_j forever.

Nash Equilibria: Proof (Cont.)

- Case 2: In the first stage or in a stage where all the preceding outcomes have been (R_1, R_2).
  - If player j follows the trigger strategy, then he should play R_j in this stage, and the outcome from this stage onwards will be (R_1, R_2) in every stage. Thus, player j's payoff from this stage onwards is
  $$4 + 4\delta + 4\delta^2 + \cdots = \sum_{t=1}^{\infty} 4 \times \delta^{t-1} = \frac{4}{1-\delta}.$$
  - If player j plays L_j in this stage (not follow the trigger strategy), player i still plays R_i in this stage but L_i forever from the next stage. And then player j will also play L_j from the next stage onwards, which is his optimal choice. This means player j's payoff from this stage onwards is
  $$5 + 1\delta + 1\delta^2 + \cdots = 5 + \sum_{t=1}^{\infty} \delta^t = 5 + \frac{\delta}{1-\delta}.$$

Nash Equilibria: Proof (Cont.)

- Case 2 (Cont.): Playing the trigger strategy is optimal iff

$$\frac{4}{1-\delta} \geq 5 + \frac{\delta}{1-\delta} \Longleftrightarrow \delta \geq \frac{1}{4}.$$

- Summarizing Cases 1 and 2, the trigger strategies constitute a Nash equilibrium for the game iff $\delta \geq \frac{1}{4}$.

Subgame-perfect Nash Equilibrium

Claim

The trigger-strategy Nash equilibrium in the infinitely repeated Prisoners' Dilemma game is subgame perfect.

Subgame-perfect Nash Equilibrium: Proof

- In an infinitely repeated game, a subgame is characterized by its previous history. The subgames can be grouped as follows:
  1. Subgames whose previous histories are always a finite sequence of (R_1, R_2).
  2. Subgames whose previous histories contain other outcomes different from (R_1, R_2).
- For a subgame in Case (i), the players' strategies in such a subgame are again the trigger strategies, which is a Nash equilibrium for the whole game and thus for the subgame as well.
- For a subgame in Case (ii), the players' strategies are simply to repeat (L_1, L_2) all the time in the subgame, which is also a Nash equilibrium.

Subgame-perfect Nash Equilibrium: Proof (Cont.)

- We can also show directly that trigger strategies constitute a subgame-perfect Nash equilibrium.
- Alternatively, we can use an approach based on the following result:

One-deviation principle (单阶段偏离原则)

A strategy profile is a subgame-perfect Nash equilibrium if and only if, for each player i and for each subgame, no single deviation would raise player i's payoff in the subgame.

Collusion between Cournot Duopolists

- In the Cournot model, the unique Nash equilibrium involving each firm producing $q_c = \frac{a-c}{3}$, and earning a profit of $\pi_c = \frac{(a-c)^2}{9}$.
- If there is a monopolist, then the monopoly quantity is $q_m = \frac{a-c}{2}$ and profit is $\pi_m = \frac{(a-c)^2}{4}$.

* If the two firms can collude to produce $\frac{q_m}{2}$ each, then they jointly produce the monopoly quantity q_m. Each of them obtains a profit of $\frac{\pi_m}{2} = \frac{(a-c)^2}{8}$.

- If firm i produces $\frac{q_m}{2}$, then the best response for firm j is to produce $q_d = \frac{3(a-c)}{8}$. In this case, firm i's profit is $\frac{3(a-c)^2}{32}$, while firm j's profit is $\pi_d = \frac{9(a-c)^2}{64}$.

Collusion between Cournot Duopolists (Cont.)

- Consider the infinitely repeated game based on the Cournot stage game when both firms have the discount factor $0 < \delta < 1$.
- Trigger strategy:
  - in period 1, produce half of the monopoly quantity, $\frac{q_m}{2}$.
  - in period t, produce $\frac{q_m}{2}$ if both firms have produced $\frac{q_m}{2}$ in all preceding $t-1$ periods; otherwise, produce the Cournot quantity q_c.
- Here the cooperative output is $\frac{q_m}{2}$ and noncooperative output is q_c.

Firm 2

Firm 1	q_c	$\frac{q_m}{2}$	q_d
q_c	π_c, π_c		
$\frac{q_m}{2}$		$\frac{\pi_m}{2}, \frac{\pi_m}{2}$	$\pi_d, \frac{3(a-c)^2}{32}$
q_d		$\frac{3(a-c)^2}{32}, \pi_d$	

- Question: Is the collusive (cooperative) outcome sustained?

Trigger-strategy SPNE

Claim

For the infinitely repeated game with the Cournot stage game, both firms playing the trigger strategy is a subgame-perfect Nash equilibrium if and only if $\delta \geq \frac{9}{17}$.

Proof.

- Suppose firm *i* has adopted the trigger strategy, we need to show firm *j*'s best response is also to player the trigger strategy in any subgame.
- There are again two types of subgames to be checked.

Trigger-strategy SPNE: Proof

- Case 1: if a quantity other than $\frac{q_m}{2}$ has been chosen by any firm before the current period, then firm i chooses q_c from this period onwards.
  * The best response for firm j is also to choose q_c from this period onwards. Thus, playing the trigger strategy is optimal in this subgame.
- Case 2: in period t, if the outcomes of all previous periods are $\left(\frac{q_m}{2}, \frac{q_m}{2}\right)$.
  * Firm j's payoff from this period onwards if it follows the trigger strategy is

$$\frac{\pi_m}{2} + \frac{\pi_m}{2}\delta + \frac{\pi_m}{2}\delta^2 + \cdots = \frac{\pi_m}{2(1-\delta)}.$$

Trigger-strategy SPNE: Proof (Cont.)

* If firm j deviates from the trigger strategy by choosing a quantity other than $\frac{q_m}{2}$, then firm i produces $\frac{q_m}{2}$ in this period, but q_c from period $t+1$ onwards. Thus, it is optimal for firm j to produce q_d in this period and q_c from period $t+1$ onwards. Thus, firm j's present value of the payoffs from period t onwards is

$$\pi_d + \pi_c\delta + \pi_c\delta^2 + \cdots = \pi_d + \frac{\delta}{1-\delta}\pi_c.$$

* Therefore, trigger strategy is the best response for firm j to firm i's trigger strategy iff

$$\frac{\pi_m}{2(1-\delta)} \geq \pi_d + \frac{\delta}{1-\delta}\pi_c \iff \delta \geq \frac{\pi_d - \frac{\pi_m}{2}}{\pi_d - \pi_c} = \frac{9}{17}.$$

Two-phase strategy

- What happens if players are less patient, i.e., $\delta < \frac{9}{17}$? Are there any other strategies that can support the collusive outcome as a subgame-perfect Nash equilibrium?
- Consider the two-phase (or carrot-and-stick) strategy:
  - in the first period, produce half of the monopoly quantity $\frac{q_m}{2}$;
  - in period t, produce $\frac{q_m}{2}$ if both firms produce $\frac{q_m}{2}$ or both firms produce x in period $t-1$; otherwise, produce x.
- This strategy involve a (one-period) punishment phase in which the firm produces x and a (potentially infinite) collusive phase in which the firm produces $\frac{q_m}{2}$.
- Such a strategy punishes
  - a firm for deviating from the collusive phase;
  - a firm for deviating from the punishment phase.

Two-phase strategy SPNE

- If both firms produce x, the profit of each firm is denoted by $\pi(x) = (a - 2x - c)x$, where $x \leq \frac{a-c}{2}$.
- If firm i produces x, the best response of firm j is to produce $q_{dp} = \frac{a-x-c}{2}$ and the corresponding profit is denoted by $\pi_{dp}(x) = \frac{(a-x-c)^2}{4}$.
- There are two types of subgames:
    - collusive subgames: the outcome of previous period is either $(\frac{q_m}{2}, \frac{q_m}{2})$ or (x, x);
    - punishment subgames: the outcome of previous period is neither $(\frac{q_m}{2}, \frac{q_m}{2})$ nor (x, x).

Two-phase strategy SPNE

- To show both firms adopting the two-phase strategy is a subgame-perfect Nash equilibrium, we use the one-deviation principle.
- Suppose firm i has adopted the two-phase strategy.
- In collusive subgames, if firm j also adopts the two-phase strategy, its payoff is

$$(1 + \delta + \delta^2 + \cdots)\tfrac{1}{2}\pi_m = \left(1 + \delta + \tfrac{\delta^2}{1-\delta}\right)\tfrac{1}{2}\pi_m.$$

* If firm j deviates in this period only, then firm i still chooses $\frac{q_m}{2}$ in this period but x in the next period. Then firm j would choose q_d in this period and x in the next period. The payoff from deviation is

$$\pi_d + \delta\pi(x) + \delta^2\tfrac{1}{2}\pi_m + \delta^3\tfrac{1}{2}\pi_m + \cdots = \pi_d + \delta\pi(x) + \tfrac{\delta^2}{1-\delta}\tfrac{1}{2}\pi_m.$$

Two-phase strategy SPNE: Proof

* Thus, choosing the two-phase strategy is optimal iff

$$(1+\delta)\tfrac{1}{2}\pi_m \geq \pi_d + \delta\pi(x). \tag{1}$$

- In punishment subgames, it is optimal to choose the two-phase strategy for firm j iff

$$\pi(x) + \delta\tfrac{1}{2}\pi_m \geq \pi_{dp}(x) + \delta\pi(x). \tag{2}$$

- Both firms adopting the two-phase strategy is a subgame-perfect Nash equilibrium iff (1) and (2) hold.

Two-phase strategy SPNE: Proof (Cont.)

- The two conditions (1) and (2) can be rewritten as

$$\delta\left(\frac{1}{2}\pi_m - \pi(x)\right) \geq \pi_d - \frac{1}{2}\pi_m, \tag{3}$$

$$\delta\left(\frac{1}{2}\pi_m - \pi(x)\right) \geq \pi_{dp}(x) - \pi(x). \tag{4}$$

- Intuitions: The gain this period from deviating must not exceed the discounted value of the loss next period from punishment.

Two-phase strategy SPNE (Cont.)

- Consider the case $\delta = \frac{1}{2} < \frac{9}{17}$.
- Condition (3) is satisfied iff $\frac{x}{a-c} \leq \frac{1}{8}$ or $\frac{x}{a-c} \geq \frac{3}{8}$.
- Condition (4) is satisfied iff $\frac{3}{10} \leq \frac{x}{a-c} \leq \frac{1}{2}$.
- Thus, two-phase strategies constitute a subgame-perfect Nash equilibrium in the game iff $\frac{3}{8}(a-c) \leq x \leq \frac{1}{2}(a-c)$.

Subgame-perfect Nash Equilibrium

- In the Prisoners' Dilemma example, the cooperative outcome, which cannot be achieved in stage game or in any finitely repeated game, can be sustained if the stage game is played forever.
- The condition is that the discount factor is sufficiently large (or players are sufficiently patient).

Folk theorem (无名氏定理)

Cooperative equilibria which do not exist in static games can be achieved in repeated games.

Feasible Payoff

Definition

The payoffs $(x_1, \ldots, x_n)$ are feasible in the stage game G if they are a convex combination (i.e., a weighted average, where the weights are all nonnegative and sum to one) of the pure-strategy payoffs of G.

- In the Prisoners' Dilemma example, all pure-strategy payoffs $(1, 1)$, $(0, 5)$, $(4, 4)$ and $(5, 0)$ are feasible.
- The payoffs $(2.5, 2.5)$ are also feasible, which can be achieved if player i adopts the mixed-strategy $\frac{1}{2}L_i + \frac{1}{2}R_i$ for $i = 1, 2$.
- All feasible payoffs are depicted in the shaded region of Figure 1.

Feasible Payoff

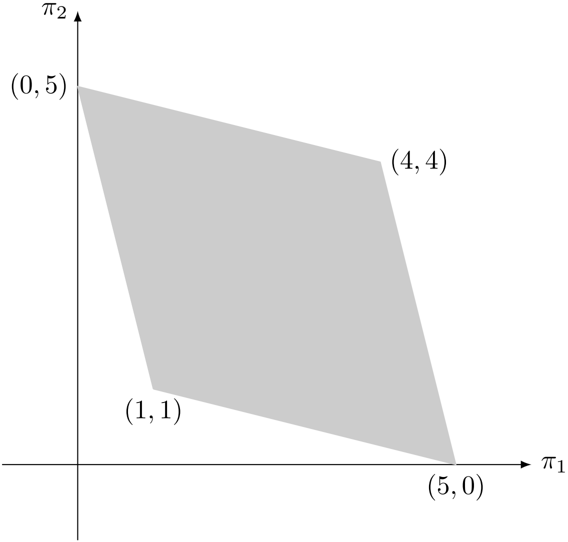

Average Payoff

Definition

Given the discount factor δ, the average payoff of the infinite sequence of payoffs $\pi_1, \pi_2, \ldots$ is

$$(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\pi_t.$$

- Both present value and average payoff can present a player's payoff in an infinitely repeated game.
- Average payoff is directly comparable to the payoffs from the stage game.

Friedman Theorem

Theorem (Friedman 1971)

Let G be a finite, static game of complete information. Let $(e_1, \ldots, e_n)$ denote the payoffs from a Nash equilibrium of G, and let $(x_1, \ldots, x_n)$ denote any feasible payoffs from G, where $x_i > e_i$ for each player i. If the discount factor δ is sufficiently close to one, then there exists a subgame-perfect Nash equilibrium in the infinitely repeated game $G(\infty, \delta)$ that achieves $(x_1, \ldots, x_n)$ as the average payoff.

- Friedman theorem is part of the Folk theorem.
- Fudenberg and Maskin (1986) have shown that the above result can be extended if the equilibrium payoffs are replaced by reservation payoffs.

Friedman Theorem