MS&E 246: Lecture 10
Repeated games

Ramesh Johari

What is a repeated game?

A *repeated game* is:

A dynamic game constructed by playing the same game over and over.

It is a dynamic game of imperfect information.

This lecture

- Finitely repeated games
- Infinitely repeated games
  - Trigger strategies
  - The folk theorem

Stage game

At each stage, the same game is played: the *stage game* G.

Assume:

- G is a simultaneous move game
- In G, player i has:
  - Action set A_i
  - Payoff $P_i(a_i, \mathbf{a}_{-i})$

Finitely repeated games

$G(K)$: G is repeated K times

Information sets:
All players observe outcome of each stage.

What are:
strategies? payoffs? equilibria?

History and strategies

Period t history h_t:

$h_t = (\mathbf{a}(0), \ldots, \mathbf{a}(t\text{-}1))$ where

$\mathbf{a}(\tau)$ = action profile played at stage τ

Strategy s_i:

Choice of stage t action $s_i(h_t) \in A_i$

for each history h_t

i.e. $a_i(t) = s_i(h_t)$

Payoffs

Assume payoff = *sum of stage game payoffs*

$$\Pi_i(\mathbf{s}) = \sum_{t=0}^{K-1} P_i(s_1(h_t), \ldots, s_N(h_t))$$

Example: Prisoner's dilemma

Recall the Prisoner's dilemma:

		Player 1	
		defect	cooperate
Player 2	defect	(1,1)	(4,0)
	cooperate	(0,4)	(2,2)

Example: Prisoner's dilemma

Two volunteers

Five rounds

No communication allowed!

Round	1	2	3	4	5	Total
Player 1	1	1	1	1	1	5
Player 2	1	1	1	1	1	5

SPNE

Suppose $\mathbf{a}^{\text{NE}}$ is a stage game NE.

Any such NE gives a SPNE:

Player i plays a_i^{NE} at every stage, regardless of history.

Question: Are there any other SPNE?

SPNE

How do we find SPNE of $G(K)$?

Observe:

Subgame starting after history h_t is identical to $G(K - t)$

SPNE: Unique stage game NE

Suppose G has a unique NE $\mathbf{a}^{\text{NE}}$

Then regardless of period K history h_K, last stage has unique NE $\mathbf{a}^{\text{NE}}$

$\Rightarrow$ At SPNE, $s_i(h_K) = a_i^{\text{NE}}$

SPNE: Backward induction

At stage K - 1, *given* $\mathbf{s}_{-i}(\cdot)$, player i chooses $s_i(h_{K-1})$ to maximize:

$$\underbrace{P_i(s_i(h_{K-1}), \mathbf{s}_{-i}(h_{K-1}))}_{\textit{payoff at stage } K-1} + \underbrace{P_i(\mathbf{s}(h_K))}_{\textit{payoff at stage } K}$$

SPNE: Backward induction

At stage K - 1, *given* $\mathbf{s}_{-i}(\cdot)$, player i chooses $s_i(h_{K-1})$ to maximize:

$$\underbrace{P_i(s_i(h_{K-1}), \mathbf{s}_{-i}(h_{K-1}))}_{\textit{payoff at stage } K-1} + \underbrace{P_i(\mathbf{a}^{NE})}_{\textit{payoff at stage } K}$$

We know: at last stage, $\mathbf{a}^{NE}$ is played.

SPNE: Backward induction

At stage K - 1, *given* $\mathbf{s}_{-i}(\cdot)$, player i chooses $s_i(h_{K-1})$ to maximize:

$$\underbrace{P_i(s_i(h_{K-1}), \mathbf{s}_{-i}(h_{K-1}))}_{\textit{payoff at stage } K-1}$$

$\Rightarrow$ *Stage game NE again!*

SPNE: Conclusion

Theorem:

If stage game has unique NE $\mathbf{a}^{\text{NE}}$, then finitely repeated game has unique SPNE:

$s_i(h_t) = a_i^{\text{NE}}$ for all h_t

Example: Prisoner's dilemma

Moral: "Cooperate" should never be played.

Axelrod's tournament (1980):

Winning strategy was "tit for tat":

Cooperate if and only if
your opponent did so at the last stage

SPNE: Multiple stage game NE

Note:

If multiple NE exist for stage game NE,

there may exist SPNE where actions are played that appear in *no stage game NE*

(See Gibbons, 2.3.A)

Infinitely repeated games

- History, strategy definitions same as finitely repeated games

- Payoffs:

  Sum might not be finite!

Discounting

Define payoff as:

$$\Pi_i(\mathbf{s}) = (1 - \delta) \sum_{t=0}^{\infty} \delta^t P_i(s_1(h_t), \ldots, s_N(h_t))$$

i.e., *discounted sum* of stage game payoffs

This game is denoted $G(\delta, \infty)$

(*Note:* (1 - δ) is a normalization)

Discounting

Two interpretations:

1. Future payoffs worth less than today's payoffs

2. Total # of stages is a geometric random variable

Folk theorems

- Major problem with infinitely repeated games:

  If players are patient enough,
  SPNE can achieve "any" reasonable payoffs.

Prisoner's dilemma

Consider the following strategies, (s_1, s_2):

1. Play C at first stage.
2. If h_t = ((C,C), ..., (C,C)),
   then play C at stage t.

   Otherwise play D.

i.e., punish the other player for defecting

Prisoner's dilemma

Note: $G(\delta, \infty)$ is *stationary*

Case 1: Consider any subgame where at least one player has defected in h_t.

Then (D,D) played forever.

This is NE for subgame,

since (D,D) is stage game NE.

Prisoner's dilemma

Step 2: Suppose h_t = ((C,C), ..., (C,C)).

Player 1's options:

(a) Follow s_1 $\Rightarrow$ play C forever

(b) Deviate at time t $\Rightarrow$ play D forever

Prisoner's dilemma

Given s_2:

Playing C forever gives payoff:

$(1-\delta)(P_1(C,C) + \delta P_1(C,C) + \ldots) = P_1(C,C)$

Playing D forever gives payoff:

$(1-\delta)(P_1(D,C) + \delta P_1(D,D) + \ldots)$

$= (1-\delta) P_1(D,C) + \delta P_1(D,D)$

Prisoner's dilemma

So cooperate if and only if:

P_1(C,C) $\geq$ (1 - δ) P_1(D,C) + δ P_1(D,D)

Note: if P_1(C,C) > P_1(D,D),
then this is always true for δ close to 1

Conclude:
If δ close to 1, then (s_1, s_2) is an SPNE

Prisoner's dilemma

In our game:

Need $2 \geq (1 - \delta)\, 4 + \delta \quad \Rightarrow \quad \delta \geq 2/3$

So cooperation can be sustained if time horizon is *finite but uncertain*.

Trigger strategies

In a *(Nash) trigger strategy* for player i :

1. Play a_i at first stage.
2. If h_t = ($\mathbf{a}$, ..., $\mathbf{a}$),
   then play a_i at stage t.
   Otherwise play a_i^{NE}.

Trigger strategies

If $\mathbf{a}$ Pareto dominates $\mathbf{a}^{NE}$,

trigger strategies will be an SPNE for large enough δ

Formally: need

$P_i(\mathbf{a}) > (1 - \delta)\, P_i(a_i', \mathbf{a}_{-i}) + \delta\, P_i(\mathbf{a}^{NE})$

for all players i and actions a_i'.

Achievable payoffs

Achievable payoffs:

T = Convex hull of $\{ (P_1(\mathbf{a}), P_2(\mathbf{a})) : a_i \in S_i \}$

e.g., in Prisoner's Dilemma:

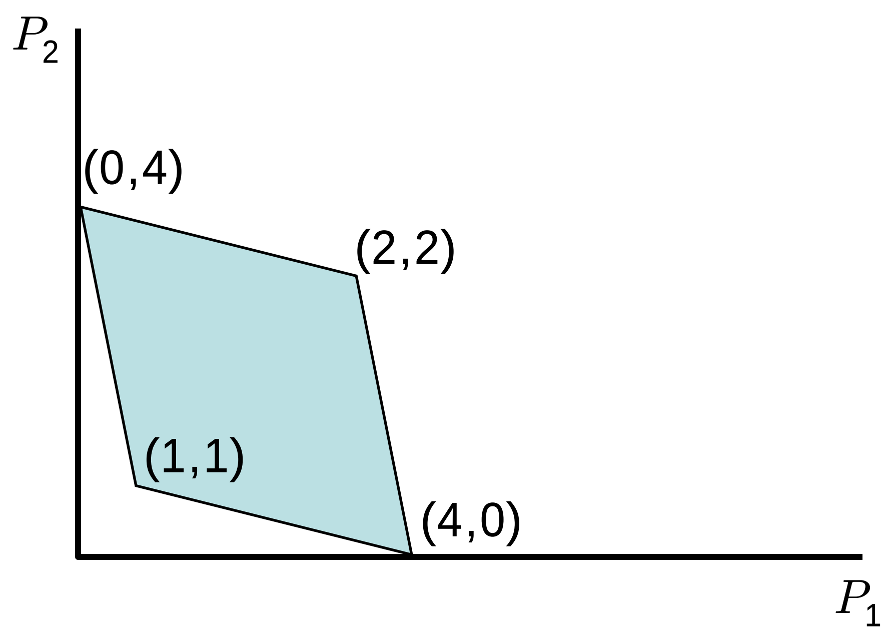

Achievable payoffs and SPNE

A key result in repeated games:

Any "reasonable" achievable payoff can be realized in an SPNE of the repeated game, if players are patient enough.

Simple proof: generalize prisoner's dilemma.

Randomization

- To generalize, suppose before stage t all players observe i.i.d. uniform r.v. U_t
- History:
  $h_t = (\mathbf{a}(0), \ldots, \mathbf{a}(t\text{-}1), U_0, \ldots, U_t)$
- Players can use U_t to *coordinate* strategies at stage t

Randomization

E.g., suppose players want to achieve

$\mathbf{P} = \alpha\, \mathbf{P}(\mathbf{a}) + (1 - \alpha)\, \mathbf{P}(\mathbf{a}')$

If $U_t \leq \alpha$: Player i plays a_i

If $U_t > \alpha$: Player i plays a_i'

We'll call this the $\mathbf{P}$-*achieving action* for i.

(Uniquely defined for all $\mathbf{P} \in T$.)

Randomization

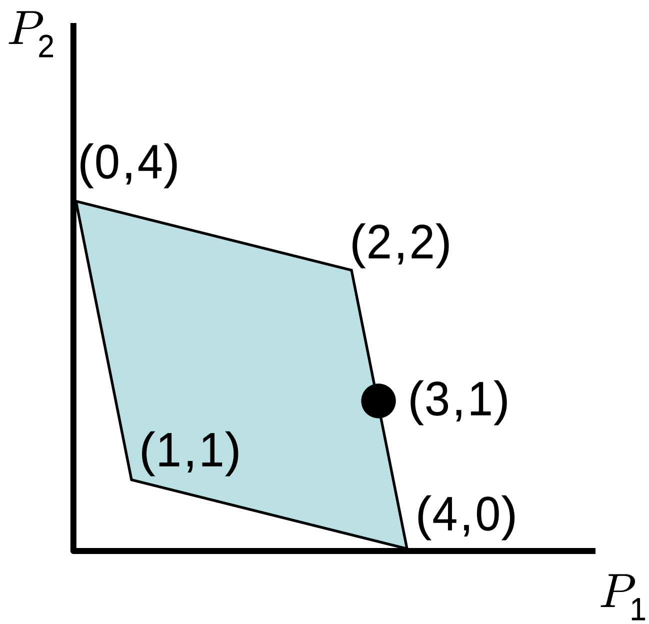

E.g., Prisoner's Dilemma

Let $\mathbf{P}$ = (3,1).

$\mathbf{P}$-achieving actions:

Player 1 plays C if $U_t \leq$ ½ and D if U_t > ½

Player 2 plays C if $U_t \leq$ ½ and C if U_t > ½

Randomization and triggering

So now suppose $\mathbf{P} \in T$ and:

$P_i > P_i(\mathbf{a}^{\text{NE}})$ for all i

Trigger strategy:

Punish forever (by playing a_i^{NE}) if opponent deviates from $\mathbf{P}$-achieving action

Randomization and triggering

Both players using this trigger strategy is again an SPNE for large enough δ.

Formally: need

$$(1 - \delta)\, P_i(p_i, \mathbf{p}_{-i}) + \delta\, P_i$$
$$> (1 - \delta)\, P_i(a_i', \mathbf{p}_{-i}) + \delta\, P_i(\mathbf{a}^{NE})$$

for all players i and actions a_i'.

(Here p is $\mathbf{P}$-achieving action for player i, and $\mathbf{p}_{-i}$ is $\mathbf{P}$-achieving action vector for all other players.)

Randomization and triggering

Both players using this trigger strategy is again an SPNE for large enough δ.

Formally: need

$$(1 - \delta)\ P_i(p_i,\ \mathbf{p}_{-i}) + \delta\ P_i$$
$$> (1 - \delta)\ P_i(a_i',\ \mathbf{p}_{-i}) + \delta\ P_i(\mathbf{a}^{\mathrm{NE}})$$

for all players i and actions a_i'.

(At time t:

LHS is payoff if player i does not deviate after seeing U_t;

RHS is payoff if player i deviates to a_i' after seeing U_t)

Folk theorem

Theorem (Friedman, 1971):

Fix a Nash equilibrium $\mathbf{a}^{\mathrm{NE}}$, and

$\mathbf{P} \in T$ such that

$P_i > P_i(\mathbf{a}^{\mathrm{NE}})$ for all i

Then for large enough δ,

there exists an SPNE $\mathbf{s}$ such that:

$\Pi_i(\mathbf{s}) = P_i$

Minimax payoffs

What is the *minimum* payoff Player 1 can guarantee himself?

$$\min_{a_2 \in A_2} \left\{ \quad \max_{a_1 \in A_1} P_1(a_1, a_2) \quad \right\}$$

Minimax payoffs

What is the *minimum* payoff Player 1 can guarantee himself?

$$\min_{a_2 \in A_2} \left\{ \underbrace{\max_{a_1 \in A_1} P_1(a_1, a_2)} \right\}$$

Given a_2, this is the highest payoff player 1 can get...

Minimax payoffs

What is the *minimum* payoff Player 1 can guarantee himself?

$$\underbrace{\min_{a_2 \in A_2} \left\{ \max_{a_1 \in A_1} P_1(a_1, a_2) \right\}}$$

...so Player 1 can guarantee himself this payoff if he knows how Player 2 is punishing him

Minimax payoffs

What is the *minimum* payoff Player 1 can guarantee himself?

$$\min_{a_2 \in A_2} \left\{ \max_{a_1 \in A_1} P_1(a_1, a_2) \right\}$$

This is m_1, the *minimax value* of Player 1.

Generalization

Theorem (Fudenberg and Maskin, 1986):

Folk theorem holds for all $\mathbf{P}$ such that $P_i > m_i$ for all i

(Technical note:

This result requires that dimension of T = # of players)

Finite vs. infinite

Theorem (Benoit and Krishna, 1985):

Assume: for each i, we can find two NE $\mathbf{a}^{\text{NE}}$, $\underline{\mathbf{a}}^{\text{NE}}$ such that $P_i(\mathbf{a}^{\text{NE}}) > P_i(\underline{\mathbf{a}}^{\text{NE}})$

Then as $K \to \infty$,
set of SPNE payoffs of $G(K)$
approaches $\{ \mathbf{P} \in T : P_i > m_i \}$

(Same technical note as Fudenberg-Maskin applies)

Finite vs. infinite

In the unique Prisoner's Dilemma NE,

only one NE exists

$\Rightarrow$ Benoit-Krishna result fails

Note at Prisoner's Dilemma NE,
each player gets minimax value.

Summary

Repeated games are a simple way to model interaction over time.

(1) In general, too many SPNE $\Rightarrow$ not very good predictive model

(2) However, can gain insight from *structure* of SPNE strategies