ECO512
Eric Maskin
emaskin@princeton.edu

Repeated Games

Text: Mailath and Samuelson – Repeated Games and Reputation

Assessment: Problem sets plus term paper (due end of August)
└ 2 for the whole semester └ original research └ 20th August

Why are repeated games important?

Humans different from other species because they cooperate (even genetic foreigners cooperate).

Cooperation sometimes conflicts with self interest.

e.g. Prisoner's Dilemma.

	C	D
c	2, 2	-1, 3
d	3, -1	1, 1

(d, D) is equilibrium, (c, C) is efficient

Repeated games can explain cooperation among rational genetic foreigners.

e.g. Repeated Prisoner's Dilemma

Play the above game infinitely and players maximise Σ
– to avoid infinite payoffs we can:
* discount (later)
* compare sequences of payoffs according to the overtaking criterion.

Overtaking Criterion (OC)

$V = (v^1, v^2, v^3, \dots)$
$W = (w^1, w^2, w^3, \dots)$

we say $V \underset{oc}{>} W$ if $\sum_{t=1}^{T} v^t > \sum_{t=1}^{T} w^t$ for all $T \geq T^*$

i.e. V is eventually better than W.

Consider the trigger strategy or grim strategy:
"play C as long as nobody has deviated from C previously, otherwise play D".

If both players play the trigger strategy we have NE using the overtaking criterion.

If no deviation, players get: $2+2+2+\dots$

If a player deviates I get at most: $3+1+1+1+\dots$

After 3 periods the sums corresponding to the trigger strategy is better/larger than the sum corresponding to the deviation.

Careful: OC does not form a total order

Right now: Discounting criterion more prevelant in research

But, OC is more about getting that people make finite approximations to infinite games.

This is *not* the only NE in a repeated prisoner's dilemma. Many other equilibria are possible:

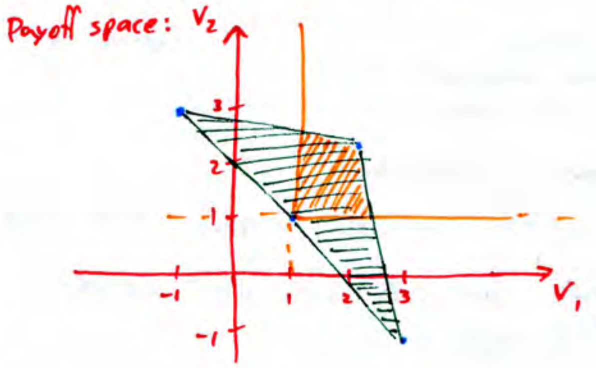

The average payoff:

$$\lim_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} v_i^t$$

~~hence it must~~ must be in the green region

If the players play rationally, however, they can guarantee themselves payoffs of at least 1, thus any payoff in the orange region can be achieved.

General Definitions

$g: S_1 \times \dots \times S_n \to \mathbb{R}^n$ is a stage game (map of strategies to outcomes)

S_i is player i's one-shot strategy space.

$g_i(s_1, \dots, s_n)$ is player i's payoff when strategies $s_1, \dots, s_n$ are used.

$$V^* = \text{convex hull } \{ g(s_1, \dots, s_n) \mid (s_1, \dots, s_n) \in S_1 \times \dots \times S_n \}$$

$$\underline{v}_i = \min_{\tilde{s}_{-i}} \max_{s_i} g_i(s_i, \tilde{s}_{-i})$$

allows for mixed strategies.

The Folk Theorem for repeated games

→ known since 50's, didn't make it into the literature until the 1970s.

Theorem: If $(v_1, \dots, v_n) \in V^*$ and $v_i > \underline{v}_i$, then there exists a NE of the repeated game giving rise to average payoffs $(v_1, \dots, v_n)$.

↳ using the overtaking criterion

Proof: Obvious ∎

↳ infinite punishment of deviating player.

Let us now consider: Complicated Prisoner's dilemma

	C	D	P
C	2,2	-1,3	-1,-1
D	3,-1	1,1	0,-2
P	-1,-1	-2,0	-1,-1

$\underline{V_i} = 0$ now.
so by Folk theorem, we get more equilibrium payoffs.
e.g. $(\frac{1}{2}, \frac{1}{2})$ is a NE
However it is <u>not</u> a <u>SPNE</u>

So how do we get to a perfect folk theorem?
↳ for SPNE

The Perfect Folk Theorem for repeated games

If $(v_1, \ldots, v_n) \in V^*$ and $v_i > \underline{v_i}$ then there exists a subgame perfect Nash equilibrium of the repeated game giving rise to average payoffs $(v_1, \ldots, v_n)$.

Strategies to sustain this equilibrium:

2 Phases (or n+1 phases)

- Normal Phase — same as before, ensure average payoffs are $(v_1, \ldots, v_n)$.
- Phase where i is punished — i is punished by min-max-ing long enough to wipe out whatever gain he got by deviating and then return to normal phase.
  ↳ relative to the deviation
  could be by not punishing

Comment <u>OBSERVABILITY ISSUE</u>

If players are punishing by mixed strategy and prefer one pure strategy to another in their mixed strategy, a player needs to know whether the punishing player is really using a mixed strategy and not just a pure strategy. Mixed strategies could be public if randomising devices are used and they are observable ex-post.

↳ e.g. toss coin H/T, decide on strategy and play, other player observes outcome and looks at the coin to check you acted correctly.

→ This could occur through an agreement.
A Nash equilibrium is a self enforcing agreement between players.

The rest of the course will consider refinements of the Perfect folk theorem above.

Finitely Repeated Games

Finitely repeated Prisoners dilemma has a unique SPNE where players play (d, D) in every period.

There is a discontinuity at $T=\infty$, where we then have an infinite number of equilibria.

└ this behaviour is not ~~unique~~ exclusive to the prisoner's dilemma.

⇓

Theorem Suppose that g has a unique NE. Then in any finitely repeated version of g there is a unique SPNE where the one-shot equilibrium is played in every period.

Proof: Obvious. ∎

We will see that this discontinuity is not such a big problem... for example with games of imperfect information or games with multiple equilibria.

Discount Factor Criterion

Single discount factor $\delta \in (0,1)$ — Could get similar results for different δ's.

Players maximize $v^1 + \delta v^2 + \delta^2 v^3 + \ldots = \frac{v}{1-\delta}$ if v is average payoff.

How does discounting change the perfect folk theorem?

└ Proof no longer works — see Fudenberg and Maskin (1986).

Two-Person Perfect Folk Theorem for discounting

Consider $(v_1, v_2) \in V^*$ where $v_i > 0 \; \forall i$ (normalize $(\underline{v}_1, \underline{v}_2) = (0,0)$)

$\exists \delta^* < 1$ s.t. $\forall \delta > \delta^*$ $\exists$ SPNE of a repeated game with discount factor δ such that discounted average payoffs are (v_1, v_2), or total payoffs are $\left(\frac{v_1}{1-\delta}, \frac{v_2}{1-\delta}\right)$.

Proof: Two Phases: normal and punishment.

Normal phase is same as always, except discounting — close enough.

Punishment phase: Let d be the maximum possible gain from a 1-st deviation.

T period punishment such that players mutual minmax minimax eachother

$(T+1)v_i > v_i + d$ (0 when players minimax eachother)

└ if this holds, it will hold for δ close to 1.

Notes: Hitting rational V_i's is not an issue, but with irrational ~~V~~ V_i's it can be a problem. see Fudenberg & Maskin (1991) JET — Week 1·3

"On the dispensibility of public randomization in repeated games"

if we can randomize so that in expectation we reach the V_i's we are fine - but this is disepensible, see paper.

Problem with proof: There needs to exist mutual minimaxing - works for 2 players but <u>not</u> with 3 or more players.

But worse still: the theorem no longer holds.

example

1,1,1	0,0,0
0,0,0	0,0,0

0,0,0	0,0,0
0,0,0	1,1,1

minimax payoff for each player is 0.

should be able to get $(\frac{1}{8}, \frac{1}{8}, \frac{1}{8})$ by folk theorem, but we can't.

Claim $\forall \delta < 1$ $\nexists$ SPNE in which average payoffs are less than $\frac{1}{4}$.

Proof Let $V_* = \inf\{V_i \mid \exists \text{SPE with avg. payoff } V_i\}$. WTS $V_* \geq \frac{1}{4}$.

Look at an equilibrium which gives us $V_*(+\varepsilon)$ (ignore ε).

Consider 1st period. Let a_i be the prob. i chooses the 1st action. (a_1, a_2, a_3) and suppose $a_2 \geq \frac{1}{2}$, $a_3 \geq \frac{1}{2}$ so that $a_2 \cdot a_3 \geq \frac{1}{4}$

In the repeated game, any player can set:

$$\frac{1}{4} + \frac{\delta V_*}{1-\delta} \leq \frac{V_*}{1-\delta} \Rightarrow V_* \geq \frac{1}{4} \quad //$$

To get a folk theorem for a game where ~~players~~ we have more than 2 players we need a further assumption — full dimensionality:

$$V^* \subseteq \mathbb{R}^n$$

this set should have dimension n.

Perfect Folk Theorem for discounting - general n-person game

If V^* has full dimensionality and $(v_1, \ldots, v_n) \in V^*$, $v_i > 0$ $\forall i$ $\exists \delta^* < 1$ s.t. $\forall \delta > \delta^*$ $\exists$ SPNE of the repeated game with discount δ and discounted average payoffs $(v_1, \ldots, v_n)$.

For this week's materials see:

Fudenberg – Maskin (1986)

Rubinstein (1977)

Aumann and Shapley (1988)

Perfect folk theorem for discounting – general n-person game

Assume V^* has dimension n (i.e. full dimensionality). For any $(v_1, \ldots, v_n) \in V^*$ s.t. $v_i > 0\ \forall i$, there exists $\delta^* < 1$ such that $\forall \delta > \delta^*$ $\exists$ SPE of the infinitely repeated game with discount factor δ in which discounted average payoffs are $(v_1, \ldots, v_n)$.

Proof: We will write proof for $n=3$, since notation is easier.

Fix $(\hat{v}_1, \hat{v}_2, \hat{v}_3) \in \text{int } V^*$, and choose $\varepsilon > 0$ such that

$$(\hat{v}_1 + \varepsilon, \hat{v}_2 + \varepsilon, \hat{v}_3 + \varepsilon) \in V^*$$

and $0 < \hat{v}_i < (\hat{v}_i + \varepsilon) < v_i$. The payoff we want to get is (v_1, v_2, v_3).

Phase I: Choose actions that result in payoffs (v_1, v_2, v_3).

If player i deviates in phase I, go to:

Phase II^i: Punish for T periods, such that $\begin{cases} \hat{v}_i + d_i < (1+T)\hat{v}_i \\ v_i + d_i < (1+T) v_i \end{cases}$ choose T so the inequalities holds.

Phase III^i: Choose actions that result in $(\hat{v}_i, \hat{v}_{-i} + \varepsilon)$.

$$d_i = \max_{s_i'} \left[g_i(s_i', s_{-i}) - g_i(s_i, s_{-i}) \right]$$

δ is fixed by ensuring in phase II ~~punishment~~ punishers do not want to deviate.

see: Fudenberg and Maskin (1986), Mailath and Samuelson

To ensure punisher does not use pure strategy and randomizes correctly in phase II, could give higher ε's in phase III ~~for the~~ depending on the pure strategy played in phase II, so that punisher is indifferent between his actions. — <u>interesting idea</u>

Full dimensionality not necessary — see Mailath and Samuelson
weaker condition exists — paper by Abreu, Smith...

Consider a finitely repeated game, where each period players know there is a probability δ that the game continues and $1-\delta$ that it ends in that period.

These conditions ensure that the expected utilities of the players are the same as in an infinitely-repeated game with discount factor δ.

There are other ways in which infinitely repeated games are similar to finitely repeated games. For example when the stage-game has more than one equilibrium.

We have the following **THEOREM:**

THEOREM: Consider the stage game that has multiple NE, and let $w_i = \min_{(s_1,\dots,s_n)\in NE} g_i(s_1,\dots,s_n)$.

Suppose for all i there exists a Nash equilibrium of the stage game $(s_1^i,\dots,s_n^i)$ such that $g_i(s_1^i,\dots,s_n^i) > w_i$ and $\dim V^* = n$. $(*)$

Then for all $\varepsilon > 0$, and all $(v_1,\dots,v_n)\in V^*$, $v_i > 0$, there exists a T^* s.t. for all $T > T^*$ $\exists$SPE of T-period finitely-repeated game in which player i's average payoff is within ε of v_i.

Note: Difference between infinitely and finitely repeated game is simply this ε.

Proof: Benoit-Krishna (1985), we shall prove something simpler, but no full dimensionality.

Theorem: Assume $(*)$ holds. Choose $(v_1,\dots,v_n)\in V^*$, $\varepsilon > 0$ s.t. $v_i > w_i$ for all i. Then $\exists T^*$ s.t. $\forall T > T^*$ $\exists$SPE of the T-period game with average payoff within ε of v_i.

Proof: If noone deviates, in the final $\hat{T}$ periods switch around the $(s_1^i,\dots,s_n^i)$'s. these are NE

Let player i's average payoffs be $v_i^* (> w_i)$ (in this $\hat{T}$ periods).

Let d_i = i's max gain from deviating in 1-shot game.

Choose $\hat{T}$ so that $\hat{T}(v_i^* - w_i) > d_i$

Proposed equilibrium: Before the final $\hat{T}$ periods, play strategies that — not necessarily Nash equilib — get within ε of $(v_1,\dots,v_n)$. Choose T^* big enough so that within $T^* - \hat{T}$ periods we can get within ε of v and that the $\hat{T}$ periods is small enough so as not to remove us from the ε neighbourhood of v.

Switch around in the final $\hat{T}$ periods, if i deviates, go to NE where i gets w_i for the rest of the game.

no problems in punishers carrying out punishment since they are playing NE and this is self enforcing.

Pricefixing Anti-trust cases, where colluding firms make use of multiplicity of equilibria

---

Uncertainty about payoffs of other players

You could be almost certain, but even a little bit of uncertainty gives us the previous results.

Literature: Kreps-Milgrom-Roberts-Wilson
Fudenberg-Maskin (1986)

The Model – uncertainty about payoffs of players

Let $\varepsilon>0$. With prob. $1-\varepsilon$, player i's payoffs are described by g_i. These payoffs are termed "rational" or "sane". With probability ε, player i's payoffs are different, termed "crazy" or "insane".

Theorem (Nash threat) Folk theorem for finitely repeated games with imperfect info
Let $(s_1^*, \dots, s_n^*)$ be NE of g and let $(V_1^*, \dots, V_n^*) = g(s_1^*, \dots, s_n^*)$.
For all $(v_1, \dots, v_n) \in V^*$ with $v_i > V_i^*$ (also true is $V_i^* = 0$, i.e. min max payoff), there exists T^* such that for all $T > T^*$ there is a T-period sequential game in which with probability $1-\varepsilon$ i's payoff is g_i and a sequential equilibrium (No longer Nash) of that game in which i's average payoff is within ε of v_i.

Note the double use of ε.
Note also that we choose the game (by fixing ε.) as well as the strategies.

Experimental evidence supports this type of crazy player existing.

Note: Type does not change during game.

Proof: let $(s_1, \dots, s_n)$ (not necessarily NE) be such that $g(s_1, \dots, s_n) \to (v_1, \dots, v_n)$ (within ε).
Crazy type of player i:
- Plays s_i as long as $(s_1, \dots, s_n)$ has not been deviated from in the past.
- Play s_i^* otherwise

Let $v_i^{max} = \max_{s \in S} g_i(s)$

$V_i^{min} = \min_{s \in S} g_i(s)$

choose $$T^* > \frac{V_i^{max} - (1-\varepsilon^{n-1})V_i^{min} + \varepsilon^{n-1} V_i^*}{\varepsilon^{n-1}(V_i - V_i^*)}$$

Construct a sequential equilibrium as follows:
- If anyone deviates from $(s_1, \dots, s_n)$ play $(s_1^*, \dots, s_n^*)$ thereafter
- If i is rational, i will play s_i except in the last T^* periods as long as no deviation from $(s_1, \dots, s_n)$.

Showing that this is equilibrium

If T periods are left, if rational player i deviates from s_i, he can at most get $V_i^{max} + (T-1)V_i^*$ (upper bound) in the rest of the game.

If he plays s_i until someone deviates and then plays s_i^*; when at least one other player is rational (happens with prob. $(1-\varepsilon)^{n-1}$), then i's payoff is at least $V_i^{min} + (T-1)V_i^*$ (lower bound)

when all other players are crazy, he gets Tv_i. occurs with prob. ε^{n-1}

Thus, the rational player's expected payoff from not deviating is greater if:
$$\varepsilon^{n-1} T v_i + (1-\varepsilon^{n-1})(V_i^{min} + (T-1)V_i^*) > V_i^{max} + (T-1)V_i^*$$
$\hookrightarrow$ This holds if $T > T^*$ if T^* is as defined above.

Why do the crazy types make a difference?

Not as discontinuous as first appears, since as $\varepsilon \to 0$, $T^* \to \infty$.

End finitely repeated games for now.

Remark:

Reputation in game theory is about convincing others your payoffs are crazy.

Renegotiation

How do we get to Nash equilibrium? Learning, evolution, some sort of dynamic process.

or NEGOTIATION

suppose there are 2 players and many equilibria in a game

If there is no outside enforcement, the only thing they could agree on is a Nash equilibrium

↳ A Nash equilibrium is a self enforcing aggreement.

Example

	C	D
c	2,2	-1,3
d	3,-1	0,0

Triger strategy:

* Play C unless someone deviated, thereafter play D.

BUT players could renegotiate and go back to C

However if this is anticipated, a player will deviate then renegotiate, then deviate etc.

Thus The triger strategy equilibrium is not renegotiation proof.

↳ we will examine ~~g~~ games which are renegotiation proof. next time.

Literature: Mailath-Samuelson - renegotiation in infinitely repeated games

Farrell-Maskin - renegotiation

Abreu - Pearce - Stacchetti - renegotiation

Also: Repeated games with imperfect observability, next week.