REPEATED GAMES

Jörgen Weibull

April 13, 2010

Q1: Can repetition induce cooperation?

- Peace and war
- Oligopolistic collusion
- Cooperation in the tragedy of the commons

Q2: Can a game be repeated?

- Game protocols vs. games.

- Preferences

Key concepts: Threats and promises, punishments and rewards

References

1. Advanced information on the 2005 Economics Nobel Memorial Prize in Economics (Aumann and Schelling)

2. Aumann R.J. (1959): "Acceptable points in general cooperative n-person games", in R. D. Luce and A. W. Tucker (eds.), *Contributions to the Theory of Games IV, Annals of Mathematics Study* 40, 287-324, Princeton University Press.

3. Friedman J. (1971): "A non-cooperative equilibrium for supergames", *Review of Economic Studies* 38, 1-12.

4. Benoit J.-P. and V. Krishna (1985): "Finitely repeated games", *Econometrica* 53, 905-922.

5. Fudenberg D. and E. Maskin (1986): “Folk theorems for repeated games with discounting and incomplete information”, *Econometrica* 54, 533-554.

6. Abreu D., P. Dutta and L. Smith (1994): “The Folk theorem for repeated games: a NEU condition”, *Econometrica* 62, 939-948.

7. Mailath G. and L. Samuelson (2006): *Repeated Games and Reputations*, Oxford University Press.

Disposition of lecture

1. Examples
2. Infinitely repeated games with discounting
3. Solution concepts
4. The one-shot deviation principle
5. Folk theorems
6. Renegotation-proofness

1 Examples

Example 1.1 *One repetition can be enough! Play the following game-protocol twice:*

	a	b	c
a	$\mathbf{3,3}$	$\mathbf{0,0}$	$8,0$
b	$\mathbf{0,0}$	$\mathbf{1,1}$	$0,0$
c	$0,8$	$0,0$	$7,7$

Assume risk-neutral and additive preferences over plays

A coordination game to which we have added a strictly dominated strategy, c

Assume that moves are observed after first period

Defines a finite extensive-form game. How many subgames?

Find a SPE that obtains payoff 10 to each player!

Example 1.2 *Repeated Prisoners' Dilemma protocol:*

	c	d
c	$3,3$	$0,4$
d	$4,0$	$1,1$

Assume risk-neutral and additive preferences with discounting:

$$\Pi_i = \sum_{t=0}^{T} \delta^t \pi_i(s_t) \qquad i = 1,2$$

for $\delta \in (0,1)$.

(a) Find all SPE for each $T < +\infty$

(b) Find all SPE for $T = +\infty$

Example 1.3 *Repeated Bertrand duopoly*

$$D(p) = \max\{0, 1-p\}$$

$$\pi_1(p_1, p_2) = \begin{cases} (1-p_1)(p_1-c) & \text{if } p_1 < p_2 \\ \frac{1}{2}(1-p_1)(p_1-c) & \text{if } p_1 = p_2 \\ 0 & \text{if } p_1 > p_2 \end{cases}$$

$$\pi_2(p_1, p_2) = \begin{cases} (1-p_2)(p_2-c) & \text{if } p_2 < p_1 \\ \frac{1}{2}(1-p_2)(p_2-c) & \text{if } p_2 = p_1 \\ 0 & \text{if } p_2 > p_1 \end{cases}$$

($p_i, c \in [0, 1]$)

Unique NE in one-shot interaction:

$$p_1^* = p_2^* = c \quad \pi_1^* = \pi_2^* = 0$$

Optimal cartel price = monopoly price:

$$\max_{p} (1-p)(p-c)$$
$$\Rightarrow \quad p^m = \frac{c+1}{2}$$

$$\Pi^m = \max_{p_1, p_2} (\pi_1 + \pi_2) = \frac{(1-c)^2}{4} > 0$$

Assume risk-neutrality, additivity and discounting $\delta \in (0,1)$

Can optimal cartel profits be maintained under finite repetition?

Under infinite repetition?

- Is the infinitely repeated game the relevant model?

- What about finitely repeated but with uncertain duration?

- Interpretation of the discount factor as a continuation probability

- A partial relief: the discontinuity at infinity disappears if the n-player stage-game has multiple Nash equilibria, the payoffs of which span a subset of dimension n. Theorem 3.7 in Benoit and Krishna (1985)

- Analysis of the case of a decreasing discount factor: Bernheim and Dasgupta (1995), *Journal of Economic Theory*

2 Infinitely repeated games with discounting

- Simultaneous-move Euclidean stage-game protocol $G = (N, A, \pi)$, for

$$N = \{1, ..., n\} \quad A = \times_{i=1}^{n} A_i \quad \pi : A \to \mathbb{R}^n$$

with each $A_i \subset \mathbb{R}^{m_i}$ compact and $\pi_i : R^m \to \mathbb{R}$ continuous. Hence $\exists M \in \mathbb{R}$ s.t. $|\pi_i(a)| < M$ for all i and $a \in A$

- Terminology: $a_i \in A_i$ "actions": could be pure or mixed strategies

- Time periods $t \in \mathbb{N}$

- Perfect monitoring: actions observed after each period (very strong assumption if actions are mixed strategies!)

1. *Histories* $H = \cup_{t \in \mathbb{N}} H_t$:

   In the initial period $t = 0$: $H_0 = \{\theta\}$ (θ is the "null history")

   In any period $t > 0$: $\langle a(0), a(1), ..., a(t-1) \rangle \in H_t = A^t$

2. *Plays*: (infinite) sequences of action profiles $\tau = \langle a(0), a(1), ..., a(t), ... \rangle \in A^\infty$

3. *Behavior strategies* $y_i : H \rightarrow A_i$. Let Y_i be the i'th player's set of behavior strategies: For any history $h \in H$, $y_i(h) \in A_i$ is i's subsequent action

4. Each behavior-strategy profile y recursively defines a play:

$$a(0) = (y_1(\theta), ..., y_n(\theta))$$

$$a(1) = (y_1(\theta, a(0)), ..., y_n(\theta, a(0)))$$

$$a(2) = (y_1(\theta, a(0), a(1)), ..., y_n(\theta, a(0), a(1)))...$$

5. Each player's preferences over plays is assumed to allow additive representation

$$u_i(y) = (1-\delta) \sum_{t=0}^{+\infty} \delta^t \pi_i [a(t)]$$

for some *discount factor* $\delta \in (0, 1)$, the same for all players i.

$u_i(\tau)$ is the *average* (or normalized) *discounted value* of the payoff stream associated with play τ

6. This defines an infinite extensive-form game Γ^δ, an *infinitely repeated game with discounting*, with *stage game* G

- History-contingent *continuation strategies*, given any history $h \in H$: the restriction of the function y to the subset of histories that begin with h. Notation:

$$y_{|h} = \left(y_{1|h}, .., y_{n|h}\right)$$

3 Solution concepts

We will use *subgame perfect equilibrium*

Definition 3.1 *A strategy profile* y *is a* **NE** *of* Γ^δ *if*

$$u_i\left(y\right) \geq u_i\left(y_i', y_{-i}\right) \quad \forall i \in N, y_i' \in Y_i$$

Definition 3.2 *A strategy profile* y *is a* **SPE** *of* Γ^δ *if*

$$u_i\left(y_{|h}\right) \geq u_i\left(y_{i|h}', y_{-i|h}\right) \quad \forall i \in N, y_i' \in Y_i, h \in H$$

Remark 3.1 *A strategy profile* y *is a NE of* Γ^δ *iff it prescribes sequential rationality on its own path (as in finite extensive-form games)*

Remark 3.2 *Unconditional play of a NE of the stage game G, in each period, can be supported in SPE in Γ^δ, for any δ.*

Remark 3.3 *Unconditional play of any given sequence of NE of the stage game G, one NE assigned to each period (irrespective of history), can be supported in SPE in Γ^δ, for any δ.*

4 The one-shot deviation principle

- In dynamic programming: "unimprovability"

Definition 4.1 *A* **one-shot deviation** *from $y_i \in Y_i$ is a strategy $y_i' \neq y_i$ that agrees with y_i at all histories but one: $\exists h' \in H$ such that*

$$y_i'(h) = y_i(h) \quad \forall h \neq h'$$

Definition 4.2 *A one-shot deviation y_i from y_i (at h') is* **profitable** *if*

$$u_i\left(y_{i|h'}', y_{-i|h'}\right) > u_i\left(y_{|h'}\right)$$

Proposition 4.1 (One-shot deviation principle) *A strategy profile y is a SPE of Γ^δ if and only if $\nexists$ profitable one-shot deviation.*

Proof:

1. SPE $\Rightarrow$ no profitable one-shot deviation.

2. not SPE $\Rightarrow$ $\exists$ profitable one-shot deviation by “payoff continuity at infinity”

5 Folk theorems

- We focus on infinitely repeated games, as opposed to the limit of finitely repeated games, see Benoit and Krishna (1985)

- Aumann (1959), Friedman (1971), Aumann and Shapley (1976), Rubinstein (1979), Fudenberg and Maskin (1986), Abreu, Dutta and Smith (1994)

Q: What payoff outcomes (average discounted payoff vectors) can be supported in SPE of the infinitely repeated game?

A: "For sufficiently patient players: any *feasible* and *individually rational* payoff vector" (The Folk theorem)

5.1 The Nash-threat folk theorem

- The idea: Payoff vectors that can be achieved in the stage game and that Pareto dominates some stage-game Nash equilibrium can be supported in SPE by threat of reversion forever to that Nash equilibrium, granted the players are sufficiently patient.

Theorem 5.1 (Friedman) *Assume that $v = \pi(a) > \pi(a^*)$ for some $a \in A$ and some Nash equilibrium $a^* \in A$ in G. There exists a $\bar{\delta} \in (0, 1)$ such that v is a SPE payoff outcome in Γ^{δ}, for every $\delta \in \left[\bar{\delta}, 1\right)$.*

Proof: Consider the following behavior-strategy profile $y \in Y$ in Γ^{δ}: $y(\theta) = a$ and $y(h) = a$ for all histories in which all players took actions a in all preceding periods. For all other histories, let $y(h) = a^*$.

1. On the path of y: there is no profitable for player i iff

$$(1-\delta) M+\delta \pi_i (a^*) \leq \pi_i (a)$$

   This holds for all $\delta < 1$ sufficiently close to 1 (since by hypothesis $\pi_i (a^*) < \pi_i (a)\ \forall i$)

2. Off the path of y: then the stage-game NE a^* is prescribed each period, and so no profitable deviation exists

5.2 The general folk theorem for two-player games

Definition 5.1 *An action* $\bar{a}_i \in A_i$ *is a* **minmax action** *in* G **against player** $j \neq i$ *if*

$$\bar{a}_i \in \arg\min_{a_i \in A_i} \left(\max_{a_j \in Aj} \pi_j(a) \right)$$

Definition 5.2 *Player* j*'s* **minmax value***:*

$$v_j^0 = \min_{a_i \in A_i} \left(\max_{a_j \in Aj} \pi_j(a) \right)$$

Definition 5.3 *A payoff vector* $v \in \mathbb{R}^2$ *is (strictly)* **individually rational** *if* $v > \left(v_1^0, v_2^0\right)$.

Definition 5.4 *A* **mutual minmax profile** *in G is a an action profile $\bar{a} \in A$ such that $\bar{a}_1$ is a minmax action against player 2 and $\bar{a}_2$ is a minmax action against player 1*

- Note that $\pi(\bar{a}) \leq v^0$ (since the players do not necessarily best-reply to the other's minmax action)

- Note the difference between deterministic and randomized minmaxing

- Examples: PD, MP

Definition 5.5 *The set of* **feasible payoff vectors** *in G: the convex hull of the direct payoff image of the action space:*

$$V = co[\pi(A)] \subset \mathbb{R}^n$$

- Why is convexification natural?

Definition 5.6 *The set of* **feasible and strictly individually rational** *payoff vectors in G:*

$$V^* = \left\{ v \in V : v > v^0 \right\}$$

- Examples: the PD, MP, Cournot duopoly.

Theorem 5.2 (Fudenberg and Maskin) *Assume that $n = 2$. For every $v \in V^*$ there exists a $\bar{\delta} \in (0, 1)$ such that v can be arbitrarily well approximated by a SPE payoff outcome in Γ^{δ}, for every $\delta \in \left[\bar{\delta}, 1\right)$.*

Proof sketch: Suppose that $v = \pi(a)$ for some action profile $a \in A$ and let $\bar{a} \in A$ be a mutual minmax action profile.* Let M be an upper bound

*More generally, any payoff vector $v \in V^*$ can be arbitrarily well approximated in the repeated game by a suitably chosen sequence of action profiles (such that the discounted average payoff comes close to v).

on all π_1 and π_2. There exists an integer $L > 0$ s.t.

$$L \cdot \min_{i \in N} \left[\pi_i \left(a^* \right) - \pi_i \left(\bar{a} \right) \right] > M - \min_{i \in N} \pi_i \left(a^* \right)$$

1. In case of defection from play of a: mutual minmaxing for L periods.

2. In case of defection from prescribed mutual minmaxing: re-start L periods of minmaxing

3. For all other histories: prescribe play of some stage-game NE

For $\delta \in (0, 1)$ sufficiently close to 1, this deters all deviations. We have to check:

(a) no profitable one-shot deviation from proposed path

(b) no profitable one-shot deviation from punishment phase. **End of proof sketch**

- Note the neutrality of the result: it does **not** say that repetition will lead to cooperation (better outcomes than repeated stage-game NE), only that it makes cooperation **possible**, if players are sufficiently patient!

- What if one player is patient and the other impatient? - A very subtle point. See Lehrer and Pauzner (1999), *Econometrica*.

- What if there is imperfect monitoring? - See Mailath and Samuelson (2006)

5.3 General folk theorem for n-player games

- For $n > 2$ there may not exist any mutual minmax action-profile. Example: player 1 chooses row, player 2 chooses column and player 3 chooses tri-matrix

$$\begin{array}{ccc} & L & R \\ T & 1,1,1 & 0,0,0 \\ B & 0,0,0 & 0,0,0 \end{array} \qquad \begin{array}{ccc} & L & R \\ T & 0,0,0 & 0,0,0 \\ B & 0,0,0 & 1,1,1 \end{array}$$

  Here $v^0 = (0,0,0)$ but $\nexists$ action profile that keeps all three players' payoffs non-positive

- Hence, proof cannot be generalized. Not only that: the claim can be show to be false for $n > 2$!

- Fudenberg-Maskin (1986) impose a "full dimensionality" condition (of direct payoff image of action space), and "reward" those who participate in punishment of a player i by returning, after punishment of i, to a SPE continuation play that is not so good for i but relatively good for the others

- More general result:

Definition 5.7 *Two players in G, say i and j, have* **equivalent** *payoff functions if $\pi_j = \alpha\pi_i + \beta$ for some $\alpha > 0$ and $\beta \in \mathbb{R}$.*

Theorem 5.3 (Abreu, Dutta and Smith) *Assume that no pair of players have equivalent payoff functions. For every $v \in V^*$ there exists a $\bar{\delta} \in (0,1)$ such that v is a SPE payoff outcome in Γ^{δ}, for every $\delta \in \left[\bar{\delta}, 1\right)$.*

- Proof idea: same as in Fudenberg-Maskin, "reward" those who participate in punishment of a player i by returning, after punishment of i, to a SPE continuation play that is not so good for i but relatively good for the punishers.

6 Renegotiation proofness

In class

1. Benoit J.-P. and V. Krishna (1993): "Renegotiation in finitely repeated games", *Econometrica* 61, 303-323.

2. van Damme E. (1988): "The impossibility of stable renegotiation", *Economics Letters* 26, 321-324.

Next time: discussion of selected research articles

THE END