SF2972 GAME THEORY
Repeated games

Jörgen Weibull

February 2017

1 Introduction

Q1: Can repetition enable "better" outcomes than "static" equilibrium?

- Peace instead of war?
- Resolution of the tragedy of the commons?
- Collusion in oligopolistic markets?
- Keeping together criminal gangs?

Q2: Can repetition enable "worse" outcomes than "static" equilibrium?

- Better for one party but worse for another? Worse for all parties?

Key concepts

Threats and promises

Punishments and rewards

Credibility

- Credible threats ”cost nothing” but ”credible promises” may be costly!

Example 1.1 *Consider a repeated prisoners' dilemma protocol (in monetary gains):*

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

(a) Suppose this is played $T = 100$ times, each time as a simultaneous-move game, under perfect monitoring (of past moves), and that each player evaluates plays in terms of the sum of own monetary gains:

$$\Pi_i = \sum_{t=1}^{T} \pi_i \left(a \left(t \right) \right) \qquad i = 1, 2$$

where $a\left(t\right) \in \{c, d\}^2\ \forall t$. If $T = 100$, how would you play? What does the extensive form look like? What is a strategy? Subgame? Find all SPE! Is cooperation possible in SPE?

(b) Suppose everything is as in (a), except that now T is a geometrically distributed random variable. After each round, the game continues with probability $\delta \in (0,1)$ to the next round, with statistically independent draws each time. Then

$$\Pr[T=1] = 1-\delta, \ \Pr[T=2] = \delta(1-\delta), \ \Pr[T=3] = \delta^2(1-\delta), \ldots$$

How would you now play? What is a strategy? Subgame? Payoff functions? How define SPE? Find some SPE! Is cooperation possible in SPE? [Discounting?]

(c) Suppose everything is as in (b), except that the random variable T has a probability distribution with finite support, say $\Pr\left[T \leq 10^9\right] = 1$.

(d) Suppose everything as in (a),(b) or (c), except that now monitoring is imperfect. Two main cases: public monitoring (both players observe the same noisy signal about last round's play), private monitoring (each player observes a private noisy signal about last round's play)

Example 1.2 *Finitely repeated play of a coordination game with an added strictly dominated strategy:*

$$\begin{array}{cccc} & a & b & c \\ a & \mathbf{3,3} & \mathbf{0,0} & 8,0 \\ b & \mathbf{0,0} & \mathbf{1,1} & 0,0 \\ c & 0,8 & 0,0 & 7,7 \end{array}$$

Suppose each player adds up his or her period payoffs. Assume perfect monitoring.

Repeated play of (b, b) gives payoff 1 to each player in each round. Can this be obtained in SPE?

Repeated play of (a, a) gives payoff 3 to each player in each round. Can this be obtained in SPE?

Is it possible, in SPE, to obtain higher payoffs than $3T$ for each player?

2 Infinitely repeated games with discounting

- Simultaneous-move stage game $G = \langle N, A, \pi \rangle$, for

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

with each A_i is finite (or, more generally, compact)

- Terminology: $a_i \in A_i$ “actions”

- Time periods $t = 0, 1, 2, ...$

- Perfect monitoring: all actions observed after each period

- Write $\alpha_i \in \Delta(A_i)$ if α_i is a randomized action choice, a "mixed action", by player i

- Write $\mathbb{N}$ for the non-negative integers (that is, including zero)

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

   In the initial period $t = 0$: $H_0 = \{h_0\}$ (h_0 is the "null history")

   In any period $t > 0$: $h = \langle h_0, a(0), a(1), ..., a(t-1) \rangle \in H_t = H_0 \times A^t$

2. *Plays*: infinite sequences of action profiles

$$\tau = \langle a(0), a(1), ..., a(t), ... \rangle \in A^{\infty}$$

3. *Behavior strategies* $y_i : H \rightarrow \Delta(A_i)$

   (a) For any history $h \in H$: $y_i(h) = \alpha_i \in \Delta(A_i)$ is i's (local) randomization, in the next period, over his or her action set

   (b) Y_i denote the set of behavior strategies for player i, and let $Y = \times_{i \in N} Y_i$

4. Each behavior-strategy profile $y \in Y$, when used, recursively defines a play $\tau \in A^\infty$:

   (a) $a(0) \in A$ is the realization of $y(h_0) \in \boxdot(A)$

   (b) $a(1) \in A$ is the realization of $y(h_0, a(0)) \in \boxdot(A)$

   (c) $a(2) \in A$ is the realization of $y(h_0, a(0), a(1)) \in \boxdot(A)$ etc.

5. Each player's *preferences over plays* is assumed to be representable by the *Bernoulli function*

$$v_i(\tau) = (1-\delta) \sum_{t=0}^{\infty} \delta^t \pi_i[a(t)]$$

   for some common *discount factor* $\delta \in (0,1)$

   This is the *normalized present value* of the stream of stage-game payoffs.

6. *Payoff functions* $u_i : Y \longrightarrow \mathbb{R}$ are defined as the *normalized expected present value* of the payoff stream:

$$u_i(y) = (1 - \delta) \cdot \mathbb{E}_y \left[\sum_{t=0}^{\infty} \delta^t \pi_i [a(t)] \right]$$

This defines an *infinitely repeated game with discounting*, Γ^δ

Remark: The assumption that preferences over plays take this simple additive form (over one's own per-period payoffs) is a very strong assumption

3 Solution concepts

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

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

- Just as in the case of finite extensive-form games, a behavior-strategy profile is a NE if and only if it is sequentially rational on its own path.

- *Continuation strategies*: given any history $h \in H$, the restriction of a behavior-strategy profile y to the subset of histories that begin with h:

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

- Recall that under perfect monitoring every history is the root of a subgame

Definition 3.2 *A behavior-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 *Unconditional play of any NE of the stage game G in each period, can be supported in SPE in Γ^δ, for any δ and for any time horizon $T \leq +\infty$*

Remark 3.2 *Unconditional play of any given sequence of NE of the stage game G can also be supported in SPE*

4 The one-shot deviation principle

In dynamic programming: this principle is called *unimprovability*

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

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

Such a deviation from a strategy profile $y \in Y$ *is* **profitable** *if*

$$u_{i|h^*}\left(y_i', y_{-i}\right) > u_{i|h^*}(y)$$

- Nash equilibria have no profitable one-shot deviations on their paths, but may have profitable one-shot deviations off their paths

- But not so for subgame perfect equilibria:

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 sketch:

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

2. not SPE $\Rightarrow$ $\exists$ profitable one-shot deviation by "payoff continuity at infinity" (in class)

Example 4.1 *Reconsider the Prisoners' dilemma and use the one-shot deviation principle to test well-known strategy profiles for SPE, given some* $\delta \in (0, 1)$*: grim trigger, tit-for-tat, all D etc.*

5 Folk theorems

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

Q: In infinitely repeated games with discounting and perfect monitoring, what payoff vectors (normalized expected present value of stream of stage-game payoffs) can be supported in SPE?

A: For sufficiently patient players (high $\delta < 1$): any *feasible* and *individually rational* payoff vector *in the stage game*

- Why called "folk theorems"?

- Early versions: NE instead of SPE, limit average payoffs (no discounting) instead of present values under discounting

5.1 The Nash-threat folk theorem

- Suppose that each action set A_i be compact (not necessarily finite), write $A = \times_{i \in N} A_i$ and let each stage-game payoff function $\pi_i : A \to \mathbb{R}$ be continuous

- Then any payoff vector in the stage game that strictly Pareto dominates some stage-game NE can be supported in SPE if the players are sufficiently patient:

Theorem 5.1 (Friedman, 1971) *Assume that $v = \pi(\hat{a}) > \pi(a^*)$ for some $\hat{a} \in A$ and some NE $a^* \in A$ 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: Let $y \in Y$ in Γ^δ be defined by $y(h_0) = \hat{a} \in A$, $y(h) = \hat{a}$ for all $h \in H$ in which all players took actions $\hat{a}$ in all preceding periods. For other $h \in H$: $y(h) = a^*$

1. **On the path** of y: No profitable one-shot deviation for player i iff

$$(1-\delta) \cdot M_i + \delta \cdot \pi_i(a^*) \leq \pi_i(\hat{a}) \tag{1}$$

where $M_i = \max_{a_i \in A_i} \pi_i(a_i, \hat{a}_{-i})$ (and note that $M_i \geq \pi_i(\hat{a}) > \pi_i(a^*)$)

   (a) Inequality (1) holds iff

$$\delta \geq \bar{\delta}_i = \frac{M_i - \pi_i(\hat{a})}{M_i - \pi_i(a^*)}$$

   (b) Let $\bar{\delta} = \max_{i \in N} \bar{\delta}_i$. Then $\bar{\delta} < 1$.

2. **Off the path** of y: the stage-game NE a^* is prescribed in each period after any such history h, so there is no profitable one-shot deviation

5.2 Example: Cournot duopoly

- Two identical firms, producing the same good, for which the demand function is

$$D(p) = 100 - p$$

in each time period $t = 1, 2, ...$

- No fixed costs and a constant marginal production cost of $c \geq 0$ per unit

- Each firm i independently decides on its on output, $q_i(t)$, in each period $t = 0, 1, 2, ...$

- The resulting market price in period t:

$$p(t) = 100 - [q_1(t) + q_2(t)]$$

- Profits in period t:

$$\pi_i [q(t)] = (100 - [q_1(t) + q_2(t)] - c) \cdot q_i(t)$$

- Perfect monitoring: past outputs are observed (or, equivalently, past prices are observed)

- The stage game G has a unique NE:

$$q_1 = q_2 = q^* = \frac{100 - c}{3}$$

- Let $Q^* = 2q^*$. This industry output exceeds monopoly industry output $\hat{Q}$:

$$\hat{Q} = \frac{1}{2}(100 - c) < \frac{2}{3}(100 - c) = Q^*$$

- Equilibrium industry profit fall short of monopoly industry profit:

$$\Pi^* = 2\left(\frac{100-c}{3}\right)^2 < \left(\frac{100-c}{2}\right)^2 = \hat{\Pi}$$

- Note that the sum of profits is a function of the sum of outputs:

$$\pi_1 + \pi_2 = (100 - (q_1+q_2) - c)\cdot(q_1+q_2)$$

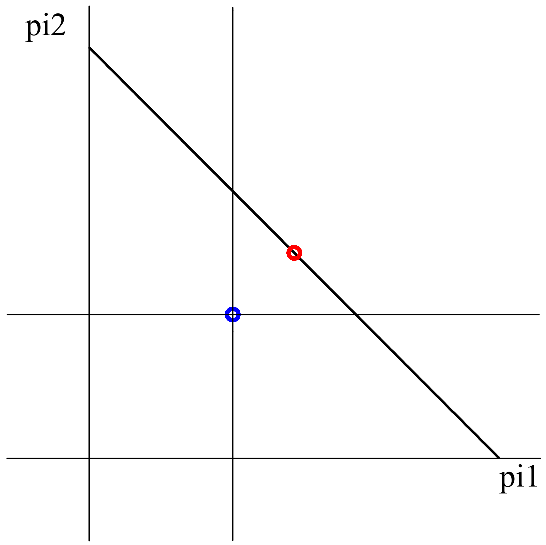

- Suppose infinitely repeated with discount factor δ (for example $\delta = e^{-r\Delta}$ where r is the interest rate and Δ the period length)

- Consider the following pure (behavior) strategy, s_i^*: start out with some quantity $\hat{q}_i \in (0, 100)$, and supply this output in all future periods, as long as no deviation from these output levels, $\hat{q} = (\hat{q}_1, \hat{q}_2)$ has been observed. If a deviation occurs: play the static Cournot equilibrium, q^*, in all future periods.

- No profitable one-shot deviations in any history containing a deviation from $\hat{q}$. The strategy pair $\left(s_1^*, s_2^*\right)$ is thus a SPE iff

$$\pi_i(\hat{q}) \geq (1-\delta) \cdot \max_{q_i \in [0,100]} \pi_i(q_i, \hat{q}_{-i}) + \delta \cdot \pi_i(q^*) \quad \text{for } i = 1, 2$$

- Possible to support also other outcomes in SPE? Lower than static Cournot profits for one firm, or even for both firms?

6 General folk theorems

Definition 6.1 *An action profile* $a = (a_1, .., a_n) \in A$ *is a* **minmax action-profile against player** i *if*

$$a_{-i} \in A^0_{-i} = \arg\min_{a_{-i}} \left(\max_{a_i} \pi_i (a_i, a_{-i}) \right)$$

- It is as if the others gang up to jointly punish i and i, knowing their punishment (a_{-i}) defends her/himself as best she/he can

Definition 6.2 *Player* i*'s* **minmax value***:*

$$v^0_i = \min_{a_{-i}} \left(\max_{a_i} \pi_i (a_i, a_{-i}) \right)$$

Definition 6.3 *A payoff vector* $v \in \mathbb{R}^n$ *is* **strictly individually rational** *if* $v > v^0$.

- In some games the resulting minmax value can be (much) lower if the punishers use mixed strategies

- Reconsider the Prisoner's dilemma, the matching-pennies game, a 2x2 coordination game

$$\begin{array}{ccc} & C & D \\ C & 3,3 & 1,4 \\ D & 4,1 & 2,2 \end{array} \qquad \begin{array}{ccc} & H & T \\ H & 1,-1 & -1,1 \\ T & -1,1 & 1,-1 \end{array} \qquad \begin{array}{ccc} & A & B \\ A & 2,2 & 0,0 \\ B & 0,0 & 1,1 \end{array}$$

- What are the minmax vectors under pure/mixed minmaxing?

Definition 6.4 *The set of* **feasible payoff vectors** *in the stage game G is the convex hull of the direct payoff image of the action space:*

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

- Why is convexification natural?

Definition 6.5 *The set of* **feasible and strictly individually rational** *payoff vectors in the stage game G:*

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

- Reconsider the above examples!

6.1 Two-player games

- Assume $n = 2$, $A = A_1 \times A_2$ compact and $\pi : A \to \mathbb{R}^2$ continuous

Definition 6.6 *A* **mutual minmax profile** *in G is an action profile $\left(a_1^0, a_2^0\right) \in A$ such that $a_1^0 \in A_1$ is a minmax action against 2 and $a_2^0 \in A_2$ a minmax action against 1.*

- Note that $\pi\left(a_1^0, a_2^0\right) \leq v^0$ (since a player's minmax action is not necessarily a best-reply to the other's minmax action)

- Main result: Any payoff vector in the stage game that strictly Pareto dominates the minmax payoff vector can be supported in SPE if the players are sufficiently patient. Proof: Threat of temporary mutual minmaxing.

Theorem 6.1 (Fudenberg and Maskin, 1986) *Let* $n = 2$*, and suppose* $\hat{a} \in A$ *is such that* $\pi(\hat{a}) > v^0$*. There exists a* $\bar{\delta} \in (0, 1)$ *such that play of* $\hat{a} \in A$ *in each period is supported by a SPE in* Γ^{δ}*, for any* $\delta \in \left[\bar{\delta}, 1\right)$*.*

Proof sketch:

Given $\hat{a} \in A$ such that $\pi(\hat{a}) > v^0$, consider a behavior-strategy profile $y = (y_1, y_2)$ in the repeated game, with "penalty duration" L:

1. Start by playing $\hat{a} = (\hat{a}_1, \hat{a}_2)$, and play $\hat{a}$ if $\hat{a}$ was always played so far

2. Also play $\hat{a}$ if sometime in the past the mutual minmax profile a^0 was played for L consecutive periods after which no other action pair than $\hat{a}$ was ever played

3. For all other histories: play a^0

- L has to be long enough to deter deviations in phases 1 and 2, but short enough to deter deviation in phase 3. Such an L always exists!

- Use the one-shot deviation principle!

  – One-shot deviations in phases 1&2 unprofitable iff

$$\max_{a_i \in A_i} \pi_i\left(a_i, \hat{a}_{-i}\right) - \pi_i\left(\hat{a}\right) < \left(\delta + \delta^2 + .. + \delta^L\right)\left[\pi_i\left(\hat{a}\right) - \pi_i\left(a^0\right)\right]$$

  – One-shot deviations in phase 3 unprofitable iff

$$v_i^0 - \pi_i\left(a^0\right) \leq \delta^L \cdot \left[\pi_i\left(\hat{a}\right) - \pi_i\left(a^0\right)\right]$$

- Draw picture in class

- Can this theorem explain why two rational persons stand in a street beating each other with a stick?

- Reconsider the Cournot duopoly example!

6.2 Games with more than two players

- For $n > 2$ there may exist no mutual minmax action-profile:

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

- A player can unilaterally deviate from minmaxing of another player, and obtain a payoff 1, instead of the minmax value 0

- The proof for $n = 2$ cannot be generalized. Not only that, the claim is not valid for generally valid for $n > 2$!

Definition 6.7 *Two players in G, say i and j, have* **equivalent** *payoff functions if $\pi_j = \alpha\pi_i + \beta$ for some $\alpha > 0$ and $\beta \in \mathbb{R}$.*

- The so-called *NEU* condition, or *Non-Equivalent-Utilities* condition: no pair of players have equivalent payoffs functions

- Assume that $A = \times_{i=1}^{n} A_i$ is compact and $\pi : A \to \mathbb{R}^n$ is continuous

Theorem 6.2 (Abreu, Dutta and Smith, 1994) *Assume G satisfies NEU. Suppose $\hat{a} \in A$ is such that $\pi(\hat{a}) > v^0$. Then there exists a $\bar{\delta} \in (0,1)$ such that play of $\hat{a} \in A$ in each period is supported by a SPE in Γ^{δ}, for every $\delta \in \left[\bar{\delta}, 1\right)$.*

- See Abreu, Dutta and Smith (1994) and/or Mailath & Samuelson (2006)

7 Concluding comment

- Note the *neutrality* of the folk theorems: they do *not* say that repetition will necessarily lead to cooperation, only that it *enables* cooperation *if* players are sufficiently patient

- Interesting implications of the folk theorem also for "bad" outcomes