Repeated Games and Cooperation

Seminar VIII

Marco Pelliccia

Outline

- The problem of cooperation
- Finitely-repeated prisoner's dilemma
- Infinitely-repeated games and cooperation
- Folk theorems
- Cooperation in finitely-repeated games
- Social preferences

Prisoner's dilemma

- How to sustain cooperation in the society?
- Recall the **prisoner's dilemma**, which is the canonical game for understanding incentives for defecting instead of cooperating.

		Player 2	
		Cooperate	Defect
Player 1	Cooperate	$1, 1$	$-1, 2$
	Defect	$2, -1$	$0, 0$

- Recall that the strategy (D, D) is the unique NE. In fact, D strictly dominates C and thus $(D; D)$ is the dominant equilibrium.
- In society, we have many situations of this form, but we often observe some amount of cooperation.

Prisoner's dilemma

- How to sustain cooperation in the society?
- Recall the **prisoner's dilemma**, which is the canonical game for understanding incentives for defecting instead of cooperating.

		Player 2	
		Cooperate	Defect
Player 1	Cooperate	$1, 1$	$-1, 2$
	Defect	$2, -1$	$0, 0$

- Recall that the strategy (D, D) is the unique NE. In fact, D strictly dominates C and thus $(D; D)$ is the dominant equilibrium.
- In society, we have many situations of this form, but we often observe some amount of cooperation.
- Why?

Repeated Games

- In many strategic situations, players interact repeatedly over time.
- Perhaps repetition of the same game might foster cooperation.
- By **repeated games** we refer to a situation in which the same **stage game** (one-shot game) is played at each date for some duration of T periods (also called "supergames").
- Key new concept: **discounting**.
- The future payoffs are discounted and are thus less valuable.

Discounting

- We model time preferences by assuming that future payoffs are discounted proportionately (exponentially) at some rate $\delta \in [0, 1)$ called **discount rate**.
- For example, in a two-period game with stage payoffs given by u^1 and u^2, overall payoffs will be

$$U = u^1 + \delta u^2$$

- With the interest rate interpretation, we would have

$$\delta = \frac{1}{1 + r}$$

where r is the interest rate.

Mathematical Model

- Imagine that N players are playing a strategic form game $G = \{N, (A_i)_{i \in N}, (u_i)_{i \in N}\}$ for $T > 1$ periods. At each period, the outcomes of all past periods are observed by all players.
- T can be either finite or $T = \infty$.
- A_i is the set of actions at each stage, and

$$u_i : A \longrightarrow \mathbb{R}$$

  where $A = A_1 \times ... \times A_N$.
- That is, $u_i(a_i^t, a_{-i}^t)$ is the state payoff of player i when action profile $a^t = (a_i^t, a_{-i}^t)$ is played.

Mathematical Model (cont.)

- We use the notation $\mathbf{a} = \{a^t\}_{t=0}^{T}$ to denote the sequence of action profiles (pure strategies).
- The payoff to player i in the repeated game

$$U(\mathbf{a}) = \sum_{t=0}^{T} \delta^t u_i(a_i^t, a_{-i}^t)$$

  where $\delta \in [0, 1)$.
- We denote the T-period repeated game with discount factor δ by $G^T(\delta)$.

Finitely-Repeated Prisoner's dilemma

- Recall

		Player 2	
		Cooperate	Defect
Player 1	Cooperate	$1, 1$	$-1, 2$
	Defect	$2, -1$	$0, 0$

- What happens if this game was played $T < \infty$ times?
- We use the equilibrium notion of **subgame perfect Nash equilibrium** (SPE). Recall that SPE $\iff$ backward induction.
- Start in the last period, at time T. What will happen?

Finitely-Repeated Prisoner's dilemma (cont.)

- In the last period, "defect" is a dominant strategy regardless of the history of the game. So the subgame starting at T has a dominant strategy equilibrium (D, D).
- Move to stage $T - 1$. Again, for the same argument above, the dominant strategy at $T - 1$ is (D, D).

Finitely-Repeated Prisoner's dilemma (cont.)

- In the last period, "defect" is a dominant strategy regardless of the history of the game. So the subgame starting at T has a dominant strategy equilibrium (D, D).
- Move to stage $T - 1$. Again, for the same argument above, the dominant strategy at $T - 1$ is (D, D).
- There exists a unique SPE: (D, D) at each date.

Finitely-Repeated Prisoner's dilemma (cont.)

- In the last period, "defect" is a dominant strategy regardless of the history of the game. So the subgame starting at T has a dominant strategy equilibrium (D, D).
- Move to stage $T - 1$. Again, for the same argument above, the dominant strategy at $T - 1$ is (D, D).
- There exists a unique SPE: (D, D) at each date.
- In fact this is a special case of a more general result.

Equilibria of Finitely-Repeated Games

Theorem

Consider a repeated game $G^T(\delta)$ for $T < \infty$. Suppose that the stage game G has a unique pure strategy equilibrium a^*. Then, G^T has a unique SPE. In this unique SPE, $a^t = a^*$ for each $t = 0, 1, ..., T$ regardless of history.

Proof: The proof follows exactly the same logic as the Prisoner's dilemma example. By backward induction, at date T, we will have $a^T = a^*$. given this, then we have $a^{T-1} = a^*$, and continuing inductively $a^t = a^*$ for each $t = 0, 1, ..., T$ regardless of history.

Infinitely-Repeated Games

- Consider the **infinitely-repeated game** G^{∞}.
- Now, $\mathbf{a} = \{a^t\}_{t=0}^{\infty}$ denotes the infinite sequence of action profiles.
- The payoff of player i is then

$$U(\mathbf{a}) = (1 - \delta) \sum_{t=0}^{\infty} \delta^t u_i(a_i^t, as_{-i}^t)$$

- This summation is well defined since $\delta < 1$.
- The term in front is introduced as a normalization, so that utility remains bounded even when $\delta \to 1$.

Trigger Strategies

- A **trigger strategy** essentially threatens other players with a “worse” punishment or action if they deviate from an implicitly agreed action profile.
- A **non-forgiving** trigger strategy (or *grim* strategy) s would involve the punishment *forever* after a single deviation.

Cooperation with Trigger Strategies in the Repeated Prisoner's dilemma

- Recall

		Player 2	
		Cooperate	Defect
Player 1	Cooperate	$1, 1$	$-1, 2$
	Defect	$2, -1$	$0, 0$

- Suppose both the players use the following grim trigger strategy s^*:
  - Play C in every period unless someone has ever played D in the past.
  - Play D forever if someone has played D in the past.
- We next show that s^* is an SPE if $\delta \geq 1/2$.

Cooperation with Trigger Strategies in the Repeated Prisoner's dilemma (cont.)

- Step 1: cooperation is best response to cooperation.
  - Suppose that so far both the players played C all the time. Then given s^* played by the other player, the payoffs to cooperation and defection are

    Payoff from C: $(1-\delta)(1+\delta+\delta^2+\cdots) = (1-\delta)\frac{1}{1-\delta} = 1$.

    Payoff from D: $(1-\delta)(2+0+0+\cdots) = 2(1-\delta)$.

- Cooperation is better if $2(1-\delta) \geq 1$, or alternatively said, deviation to D is not profitable if $\delta \geq 1/2$.

Cooperation with Trigger Strategies in the Repeated Prisoner's dilemma (cont.)

- Step 2: defection is best response to defection.
- Suppose that there has been some D in the past. Then, according to s^*, the other player will always play D. Against this, D is a best response.
- This argument is true in every subgame, so s^* is a SPE.

Multiplicity of Equilibria

- Cooperation is an equilibrium, but so are many other strategy profiles.
- Multiplicity of equilibria endemic in repeated games.
- Note that this multiplicity only occurs at $T = \infty$.
- For any finite T, the prisoner's dilemma has a unique SPE.

Repetition Can Lead to Bad Outcomes

- Suppose the following example.

		Player 2		
		A	B	C
	A	$2,2$	$2,1$	$0,0$
Player 1	B	$1,2$	$1,1$	$-1,0$
	C	$0,0$	$0,-1$	$-1,-1$

- The action *A* strictly dominates both *B* and *C* for both the players. The unique NE of the stage game is (A, A).
- If $\delta \geq 1/2$, this game has an SPE in which (B, B) is played in every period.
- It is supported by the trigger strategy: Play *B* in every period unless someone deviates, and play *C* if there is any deviation.

Repetition Can Lead to Bad Outcomes

- Suppose the following example.

		Player 2		
		A	B	C
Player 1	A	$2, 2$	$2, 1$	$0, 0$
	B	$1, 2$	$1, 1$	$-1, 0$
	C	$0, 0$	$0, -1$	$-1, -1$

- The action *A* strictly dominates both *B* and *C* for both the players. The unique NE of the stage game is (A, A).
- If $\delta \geq 1/2$, this game has an SPE in which (B, B) is played in every period.
- It is supported by the trigger strategy: Play *B* in every period unless someone deviates, and play *C* if there is any deviation.

Folk Theorem

- There is a “folk” theorem stating that one can support cooperation in repeated prisoner’s dilemma games and other “non-one stage” equilibrium outcomes in infinitely-repeated games with sufficiently high discount factors.

Cooperation in Finitely-Repeated Games

- We saw that finitely-repeated games with unique stage equilibrium do not allow cooperation or any other outcome than the repetition of this unique equilibrium.
- This is no longer the case when there are multiple equilibria in the stage game.
- Consider the following example,

		Player 2		
		A	B	C
Player 1	A	$3, 3$	$0, 4$	$-2, 0$
	B	$4, 0$	$1, 1$	$-2, 0$
	C	$0, -2$	$0, -2$	$-1, -1$

- The stage game has two pure NE, (B, B) and (C, C). The most cooperative outcome, (A, A), is not an equilibrium.
- In the twice repeated version of this game, we can support (A, A) in the first period!

Cooperation in Finitely-Repeated Games (cont.)

- Use the threat of switching to (C, C) in order to support (A, A) in the first period, and (B, B) in the second.
- Assume for simplicity $\delta = 1$ (no discounting).
- If we support (A, A) in the first period and (B, B), then each player obtains 4.
- If a player deviates and plays B in the first period, then in the second period the opponent plays C, and thus her best response would be to play C as well giving her -1. Thus, the total payoff from defecting will be 3. $\Rightarrow$ The deviation is not profitable!

How Do People Play in Repeated Games?

- In lab experiments, there is more cooperation in prisoner's dilemma games than predicted by theory.
- Cooperation increases as the game is repeated, even if there is only finite rounds of repetition.
- The results may differ because of **social preferences**.

How Do People Play in Repeated Games?

- Types of social preferences:
  - **Altruism:** individuals receive utility from being nice to others.
  - **Fairness:** individuals receive utility from being fair to others.
  - **Vindictiveness:** individuals like to punish those deviating from "fairness" or other accepted norms of behaviour.
- All of these types of social preferences seem to play some role in experimental results.