EC3324 Autumn Lecture #06
Repeated Games

- Reading
  - Osborne, Chapters 14, 15
- Learning outcomes
  - Understand the concept of a repeated game
  - Find subgame perfect Nash equilibria in finitely and infinitely repeated games

Repeated Interaction

- So far, we studied **one-shot** games, i.e. played only once
- Economic interaction, however, is often **repeated**
  - Think, e.g., of two firms competing in a market or consumers who buy repeatedly from the same firm
- The same applies for other fields where game theory can be applied
  - e.g. politics, biology, anthropology, interaction with your neighbour
- If interaction is repeated => opportunity for **reciprocity**
    - This is a good reason why we see cooperation in our everyday life, for example we don't dump garbage in our neighbours' garden (dominant strategy in the one shot game!)
- Intuitively, repeated interaction allows for cooperative outcomes, based on threats of **punishments** or promises of rewards *in the future*
  - Neighbour will dump garbage in *your* garden…
- This is more plausible if threats and promises are credible
  - Hence we are primarily interested in **subgame perfect equilibria**

Repeated Games

- Such repeated interaction can be modeled as **repeated games:** the (same) **stage game** is played several times
- In each **period,** the stage game is played
- Stage game can be in extensive form or strategic form
- If the stage game is in strategic form, the repeated game is an extensive game with simultaneous moves (this is the case we consider here, but the logic applies to repeated sequential games as well)
- Fundamental distinction:
  - **Finitely** repeated games: the game is played a known, finite number of times
  - **Infinitely** repeated games: the game is played infinitely often, or rather, there is a probabilistic end, reflected in the discount factor

Repeated Games

- **Formally:** A repeated game is determined by a stage game G and the number of repetitions (periods) T.
- a^k action profile in period k, then $(a^1,a^2,\ldots,a^T)$ is a terminal history
- the set of available actions to player i after each history is A_i (the available actions are always those of the stage game)
- preferences are given by the (discounted) sum of payoffs over all periods
- in period t, the **strategy** of a player conditions on the history up to this point $(a^1,a^2,\ldots,a^{t-1})$
- The game is **finitely repeated** if T is finite and **infinitely repeated** otherwise

Finitely repeated games

- For finitely repeated games, we usually compute the payoffs as the sum of the payoffs in the individual stages
- We might also, if the (economic) application suggests it, **discount** later payoffs (as we will do for infinitely repeated games)
- But for finitely repeated games, discounting is not necessary to derive interesting results (and has no impact for discount factors very close to 1)

How to solve a finitely repeated game?

- The game has a known end, so a **finite game tree**
- But we know how to solve these games, more precisely, how to find a **SPNE**:
  - By **backward induction**: start at the end and move backwards
- What does this imply?

Example: T-times repeated Prisoner's dilemma

Normal form of the stage game:

	Defect	Cooperate
Defect	0,0	7,-2
Cooperate	-2,7	5,5

- Repeat this game T times, say 10, against your neighbour.
- *What would you do*?
- Is your strategy rational? Is it part of SPNE?

Example: solving the finitely repeated PD

- Repeating the game would let us hope to get to cooperation, e.g. by strategies like "**Tit-for-Tat**": start by cooperating, and then cooperate if the other did last period, but defect if the other defected last period.
- But what happens in the **last** period?
- Clearly, whatever the history, Defect is a **dominant** strategy
- So in **any** period T subgame, the **only** equilibrium is (D,D)
- Now do backward induction
  - knowing that we will play (D,D) in period T **irrespective** of the history, in period T-1, no way to change opponents action in period T, so also in period T-1 (D,D) in any SPNE
- Continuing, the **only** SPNE is (D,D) in each period, in any subgame
  - ((D,D) in each period is even the only **Nash-equilibrium outcome**)

When does finite repetition help?

Generalization:

- If the stage game has a **unique** Nash-equilibrium, the **only** SPNE is to **repeat** this equilibrium in each period

  (and even if we deviated in one period, i.e. in other subgames, we would still play this equilibrium in all future periods)

Even more generally:

- pick an equilibrium of the stage game. Then it is **always one** SPNE to **repeat** this equilibrium in every period
- This is somewhat depressing: repeating a game seems not to help

But even finite repetition can help:

- The crucial point above is the word “unique”: if the stage game has **more** than one equilibrium, we can do **better** than repeating a **stage-game equilibrium**

Example: modified Prisoner's dilemma, stage game

	Defect	Cooperate	Partial
Defect	0,0	7,-2	3,-1
Cooperate	-2,7	5,5	0,6
Partial	-1,3	6,0	3,3

- (D,D) is still an equilibrium, but (P,P) is as well
- So we have a **second** SPNE: play (P,P) in any period
- But we can do **much** better:
- Note that the **punishment** of playing (D,D) instead of (P,P) is bigger than the **temptation** to defect if the other cooperates
  - u(P,P) – u(D,D) = 3 > 2 = u(D,C) – u(C,C)
- Thus we get a SPNE as follows: play C in periods 1,…,T-1 and P in period T if the play is (C,C) (hence **not** a stage-game equilibrium) up to that point, otherwise play (D,D) for the rest of the game (details in seminar)
- This is SPNE, because both (D,D) and (P,P) are equilibria of the stage game
  - hence threat of D and promise of P are credible
- There are lots of other SPNE for T large enough

Infinitely repeated games

- For infinitely repeated games, we usually compute the payoffs as the **discounted** sum of the payoffs in the individual stages
  - otherwise, payoffs often become infinite and we cannot distinguish between payoffs for different strategies
- i.e. let u_{it} be the payoff to player i in period t, then the total payoff is

$$u_{i0} + \delta\, u_{i1} + \ldots + \delta^{t}\, u_{it} + \ldots$$

- The discount factor $\delta < 1$ reflects two aspects:
  - **impatience** (consumption now is better than consumption tomorrow)
  - **uncertain continuation**, i.e. the game reaches a further stage only with some probability $\delta < 1$
    - $1 - \delta$ is then the probability of a breakdown in communication
  - the discount factor can reflect any combination of both

How to solve an infinitely repeated game?

- By definition an infinitely repeated game has **no known end**
- Hence we cannot apply **backward induction**
- But we can still apply the criterion for **subgame perfection**
  - note that each period starts **new** subgames
  - so a SPNE has to be an equilibrium for each game starting in any period
  - note that the subgames are **essentially identical** (except for the history):
  - in period t, the payoff for that period t is not discounted, the payoff for the futute, from period $t+1$ onwards, is discounted by δ and the game has infinite horizon
  - so the decision problem a player faces in period t is essentially the same as in period 1 (or any other period)

How to solve an infinitely repeated game?

- To see whether a strategy combination is a SPNE, we need to check whether a player has an incentive to deviate in any subgame
- Consider the repeated prisoner's dilemma:
- **Claim:** the following is a SPNE if $\delta > 2/7$:
  – play C as long as **both** players have chosen C in **all** previous periods
  – play D otherwise (variant of "**grim trigger**" strategy)
- **Proof:**
  – Consider first a subgame after somebody has chosen D in any previous period
  – Assume player *j* sticks to the above strategy, hence chooses D in **all further periods**
  – then player *i* cannot do better than choosing D as well (optimal in this period and no chance to influence behavior in **future** periods)
  – so the above strategy yields an equilibrium in any subgame following (at least) one D

- if we are in a subgame following only (C,C) in **all** previous periods, the payoff to *i* from **deviating** (choosing D even though *j* chooses C) is
- $7 + \delta\, 0 + \ldots + \delta^t\, 0 + \ldots = 7$
- but the payoff from sticking to C (assuming *j* chooses the suggested strategy) is
- $5 + \delta\, 5 + \ldots + \delta^t\, 5 + \ldots = 5 / (1 - \delta)$
- hence deviating does not pay if
- $5 / (1 - \delta) > 7 \leftrightarrow 5 > 7 - 7\delta \leftrightarrow 7\delta > 2 \leftrightarrow \delta > 2/7$
- Thus if $\delta > 2/7$ there is no incentive to deviate in a subgame following mutual cooperation in all preceding periods
- So in no subgame do we have an incentive to deviate
- Hence the above strategy yields a SPNE

- Thus **infinitely** repeated SPNE play can be **dramatically** different than **finitely** repeated play

  (e.g. always C instead of always D in Prisoner's dilemma)

- In infinitely repeated game, there is always a **future**, so promises and threats are **not empty**

- The above equilibrium is **not** the **only** equilibrium: we can establish (lots of) other paths of play as SPNE ("Folk-Theorem")

- **Note:** It is sufficient to **punish** for a **finite** number of periods ("**forgiving trigger**")
  - Such strategies are crucial if errors of a player can cause a particular trigger punishment by the opponent

Experimental Results on Repeated Games

- Typically, players do **some** steps of backward induction, but not many
- Hence, there is often a high degree of cooperation until the second or third to last period and then a sharp decrease: "**endgame effect"**
- Thus, repetition often **increases cooperation**, e.g. for Prisoner's dilemma or trust games cooperation rates are generally higher than for one-shot games
- But **negative reciprocity** can also lead to worse outcomes, e.g. **public good games** often start with intermediate contribution rates (similar to one-shot), but then decrease over time, often to 0
- "Infinitely repeated experiments" with random stopping do not differ much from finitely repeated games, but remove endgame effect, because no defined end
- **Simulations (Axelrod, 1984):** Tit-for-Tat outperforms other strategies, though not theoretically robust

Problem set #06

1. Consider the modified prisoner's dilemma game and the suggested SPNE (play C in periods 1,…,T-1 and P in period T if the play is (C,C) up to that point, otherwise play (D,D) for the rest of the game)
   a) show that it does not pay to deviate in period T
   b) show that it does not pay to deviate in period T-1
   c) show that it does not pay to deviate in any period $t < T$-1
   d) find other SPNE's (there are a lot!)
   e) replace the 7 for (D,C) with 10. Find the minimum number of repetitions such that there is a SPNE where (C,C) is played in the first period.
2. Consider Bertrand competition with demand $D = 20$-p, constant marginal cost of 0 and two firms. Let prices be constrained to integers.
   a) Show that the stage game has two Nash-equilibria.
   b) Now consider the game to be repeated T times. Show that there is a SPNE where both firms set the monopoly price $p^m = 10$ up to $t = T$-3.
3. Osborne Ex 429.1
4. (Osborne Ex 454.3)