CUR 412: Game Theory and its Applications, Lecture 15

Prof. Ronaldo CARPIO

June 14, 2016

Announcements

- Homework 5 is due today.
- The final exam will be on June 20, 4-6 PM, in Boxue 507.

Review of Last Week

- In a *finitely* repeated Prisoner's Dilemma, the only NE and SPNE outcome is (D, D) in each period.
- However, in the *infinitely* repeated Prisoner's Dilemma, it is possible to have NE and SPNE where the outcome is (C, C) in each period, if the discount factor δ is high enough.
- For example, if both players play *GrimTrigger* and δ is high enough (depending on the payoffs), then an outcome of (C, C) in every period is a NE.
- The condition on δ means that players have to be sufficiently patient to value the long-term gain of (C, C) over the short-term gain of defecting (and then being punished by D forever).
- This shows that cooperation is possible even if both players are only motivated by self-interest (under certain conditions).

Review of Last Week

- Checking whether a given pair of strategies forms a NE is difficult, since there are an infinite number of ways to deviate.
- Checking whether a pair of strategies forms a SPNE is simpler, using the *one deviation* (or *one-shot deviation*) property.
- A *one-shot deviation* from a strategy s is:
  - deviating from s in the *first* period of the subgame, then *reverting* back to his strategy in s for the rest of the game.
- In a *finite* horizon game, a strategy profile is a SPNE if and only if it satisfies the one-deviation property.
- In an *infinite* horizon game where the discount factor is less than 1, a strategy profile is a SPNE if and only if it satisfies the one-deviation property.
- To check if a strategy pair is a SPNE, we need to check all possible one-shot deviations, and show that no player can get a higher discounted sum.

Tacit Collusion Among Firms

- One important application of repeated-game equilibria is the study of *tacit collusion* among firms.
- Usually, it is considered illegal for competitors to fix prices (i.e. agree not to compete on price, and keep prices high).
- If executives from rival firms are seen to meet and discuss prices, this is grounds for a lawsuit.
- However, even if they do not communicate, firms may collude, simply by observing the past history of prices.
- "Price wars" can be seen as times when firms seek to punish each other.

- Consider this Cournot duopoly problem with two firms.
- Each firm has cost of production $c_i(q_i) = 10q_i$.
- Market demand is given by $P = 100 - Q$.
- The two firms repeatedly play the Cournot duopoly game in time periods $t = 1, 2, 3, \ldots$ with discount factor δ.
- The NE of the Cournot game in a single period is $q_1 = q_2 = 30$, $p = 40$, and profits for each firm are 900.
- If there was a single monopolist firm, the optimal $q = 45$, $p = 55$, profits = 2025.
- Suppose in each period, each firm can choose to *Collude*, in which case they produce $q_i = 22.5$, half the monopoly quantity.
- Or, they can choose to *Defect*, in which case they maximize their own profit, given the quantity of the other firm.

- If both firms choose *Defect*, they choose the NE quantities ($q_i = 30$), which results in a profit of 900 for each firm.
- If both firms choose *Collude*, they choose half the monopoly quantities ($q_i = 22.5$), which results in a profit of 1012.5 for each firm.
- If one chooses *Collude* while the other chooses *Defect*, the firm that *Collude*s chooses $q_i = 22.5$, while the firm that *Defect*s chooses the best response to that, which is $q_j = 33.75$.
- The firm that chose *Collude* gets a profit of 759.375, while the firm that chose *Defect* gets a profit of 1139.06.
- Note that this situation is a Prisoner's Dilemma.
- We can then use the one-shot deviation principle to show that a given strategy profile (e.g. both players play Modified Grim Trigger) is a SPNE, for a certain value of the discount factor δ.

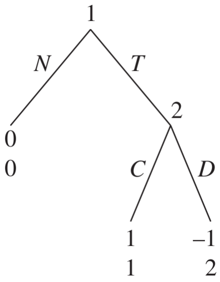

- The one-shot deviation principle is not limited to repeated Prisoner's Dilemma; it can be used in any repeated game situation to prove that a given strategy profile is or is not a SPNE.
- Consider this "trust game":
  1. Player 1 chooses whether to ask Player 2 to do something. He chooses Trust (T) or No Trust (N).
  2. Player 2 chooses to Cooperate (C) or Defect (D).
- Defecting is better for Player 2, at the expense of Player 1.

Trust Game

- The payoff matrix is:

	C	D
N	0,0	0,0
T	1,1	-1,2

- The only NE is (N, D).

Repeated Trust Game

- Now, consider the infinitely repeated version of this game.
- Suppose both players use strategies similar to *Grim Trigger*:
- Player 1:
  - Play T in the beginning;
  - If there has never been a deviation from (T, C), play T, otherwise play N forever.
- Player 2:
  - Play C in the beginning;
  - If there has never been a deviation from (T, C), play C, otherwise play D forever.

- If both players do not deviate, their payoff sequence is $(1, 1), (1, 1), \ldots$
- Both players get a discounted sum of 1.
- Player 1 will get 0 if he deviates, so he is playing a best response.
- Let's look at a one-shot deviation from Player 2.
- Suppose Player 2 plays D, then reverts to his strategy.
- The outcomes will be: $(T, D), (N, D), (N, D), (N, D), \ldots$
- Payoff sequence: $(-1, 2), (0, 0), (0, 0), (0, 0), \ldots$
- Discounted sum for Player 2 is $(1 - \delta)2$.
- Player 2 has no incentive to deviate if $1 \geq (1 - \delta)2$, or $\delta \geq \frac{1}{2}$.

Ch 14.8: NE Payoffs in an infinitely repeated PD

- Let's return to the subject of NE (not necessarily SPNE) in the repeated Prisoner's Dilemma.
- We have seen that there are NE that result in (C, C) every period, resulting in a discounted average of 2 for each player.
- For example: if both players play *Grim Trigger* or *Tit-for-Tat*.
- There are also NE that result in (D, D) every period, resulting in a discounted average of 1 for each player.
- For example: if both players play *Always Defect*.
- Are there any other discounted average payoffs possible in a NE?

Feasible Discounted Average Payoffs

- First, let's consider what discounted average payoffs are possible when some strategy pair is played, without considering if it is NE.
- The simplest case is if a player always plays the same action, C or D.
- This generates a discounted average equal to one of the 4 outcomes of the single-stage game: $(3,0),(0,3),(1,1)$, or $(2,2)$.
- For example, the strategy pair
  - Player 1 always plays C
  - Player 2 always plays D
- results in the discounted average payoff $(0,3)$.

Feasible Discounted Average Payoffs

- Now, consider the strategy pair
  - Player 1 always plays C
  - Player 2 alternates between C and D.
- The outcome will be $(C, C), (C, D), (C, C), (C, D), \ldots$
- The payoffs will be $(2, 2), (0, 3), (2, 2), (0, 3), \ldots$
- The (undiscounted) average payoff is $(1, \frac{5}{2})$.
- For $\delta < 1$, Player 1's discounted average must exceed 1, since more weight will be placed on the 2 in the first period.
- Similarly, for $\delta < 1$, Player 2's discounted average must be *less than* $\frac{5}{2}$, since more weight will be placed on 2 instead of 3.

Feasible Discounted Average Payoffs

- If δ is close to 1, then the discounted average is close to $(1, \frac{5}{2})$.
- Thus, we can achieve a payoff halfway between $(2, 2)$ and $(0, 3)$, by choosing a strategy pair that alternates between (C, C) and (C, D), if δ is close to 1.
- We can do the same thing with any combination of the 4 outcomes of the single-stage game, resulting in a payoff at the midpoint of any pair of outcomes.
- By using different frequencies (e.g play C 3 times, then D 1 time), we can achieve any weighted combination of the 4 original payoffs.

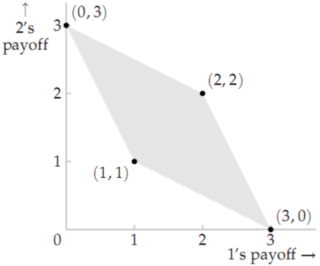

- We can generate different payoffs by *combining* different frequencies of the four outcomes of the strategic game.
- If the relative frequency of one outcome is higher than the others, then the discounted average will be closer to the payoff of that outcome.
- The set of all possible combinations is therefore the set of all *weighted averages* of the payoffs of the four outcomes.
- This is called the set of *feasible payoff profiles*.

Feasible Discounted Average Payoffs

- Now, we will show that some of these payoffs can be the result of NE, using strategies similar to *Grim Trigger*.
- First, in any NE, the payoff for each player must be at least 1, since either player can deviate to an *Always Defect* strategy which gives at least 1.
- Suppose (x_1, x_2) is a feasible payoff profile, where x_1, x_2 are at least as good as the payoff from (D, D).
- Then we can construct a finite sequence of outcomes $a^1, ...a^k$ with average payoff close to (x_1, x_2).
- Consider the strategy: play $a^1, ..., a^k$ in sequence repeatedly, and switch to D forever if the other player deviates from the sequence.
- Once punishment starts, the payoff for the other player is at most 1, so if δ is high enough, the other player will be deterred from deviating from the sequence.
- Therefore, for a high enough δ, this strategy pair is a NE.

Folk Theorem for Infinitely Repeated PD

- Proposition 435.1:
  - For any $0 < \delta < 1$, the discounted average payoff for player i in a NE is at least $u_i(D, D)$.
  - Let (x_1, x_2) be a feasible pair of payoffs for which $x_i > u_i(D, D)$ for both players. If δ is high enough, there is a NE where the discounted average is (x_1, x_2).
  - For any δ, there is a NE in which the discounted average of player i is $u_i(D, D)$.
- This is called a "folk theorem" since it was well known before anyone published it.

Subgame Perfect Folk Theorem for Infinitely Repeated PD

- Proposition 447.1:
  - For any $0 < \delta < 1$, the discounted average payoff for player i in a SPNE is at least $u_i(D, D)$.
  - Let (x_1, x_2) be a feasible pair of payoffs for which $x_i > u_i(D, D)$ for both players. If δ is high enough, there is a SPNE where the discounted average is (x_1, x_2).
  - For any δ, there is a SPNE in which the discounted average of player i is $u_i(D, D)$.

A Review of Some Key Ideas

- Recall the original goal of the course: we want to be able to predict the outcomes of strategic situtations.
- We have seen a variety of techniques and ideas that are used in order to try to answer this question.
- Let's review some of the topics we've covered and see some of the ideas behind them.

Idea 1: Play the Optimal Action

- The simplest idea is to simply see if there is an action that is always the optimal choice in any situation.
- If it exists, a rational player will always play it.
- This is a *strictly dominant* action.
- If all players have a strictly dominant action, then we can predict that will be the outcome.
- However, most games don't have a strictly dominant action.

Idea 2: Eliminate Sub-optimal Actions

- If we can't find a clearly optimal action, we can try to get rid of actions that are clearly sub-optimal.
- A *strictly dominated* action is never a best response to any situation, so a rational player will never play it.
- Therefore, it will not be part of any kind of equilibrium (NE, SPNE, etc).
- If we make the assumption that players know that other players are rational, then we can repeatedly eliminate dominated actions.
- The set of outcomes that remain are said to be *rationalizable*.

Idea 3: Find an Equilibrium of the System

- What if we are left with multiple outcomes even after eliminating all the dominated actions? (e.g. BoS)
- Instead of trying to directly find the "best" outcome, we can try a less ambitious goal: find a *stable point*, or equilibrium, of the system.
- An equilibrium is a point such that the system stays there, once it reaches that point.
- For example, when analyzing a market with a supply and demand curve, the equilibrium is the intersection.
- If the price diverges, there will be excess demand or supply, driving the price back to equilibrium.

Idea 3: Find an Equilibrium of the System

- Nash equilibrium is a specific type of equilibrium where we define "moving away" as due to the action of a *single* player.
- We can imagine different conditions for equilibrium; for example, if we consider the possibility of multiple players cooperating to change the outcome.
- This is the subject of *coalitional game theory*.
- Note that this sidesteps many important questions:
  - How do players discover the equilibrium outcomes?
  - How long does it take to reach equilibrium?
  - What if there are many equilibria; which ones are "better" than others?

Idea 4: Allow Combinations of Actions

- A mixed strategy is a way of combining actions, by creating a weighted average of actions.
- We have seen that a mixed strategy can dominate pure strategies, even if the pure strategies don't have a dominant action.
- Allowing combinations of actions also guarantees the existence of an equilibrium: this is Nash's proof that a NE exists in a finite game with mixed strategies.

Idea 5: Predict What Happens in Any Situation

- In extensive-form games, we use backwards induction to find a solution.
- This is equivalent to predicting the outcome for all possible situations (e.g. subgames), and assuming that this outcome will hold.
- The subgame perfect concept requires that equilibrium conditions must hold at every step of the game.
- However, as we've seen in the Centipede Game, it becomes harder to predict what happens many steps in the future.

Idea 6: Beliefs as Probabilities

- When there is uncertainty (about other players, or about the environment), one way to handle this is to use probability distributions to model beliefs.
- Bayes' Rule gives us a way to calculate the correct probabilities of any event, given knowledge of other events that may influence the outcome.
- We can expand our definition of equilibrium to include beliefs as well as strategies; this gives us our concept of weak sequential equilibrium.

Idea 7: There's Always a Future

- In the finitely repeated Prisoner's Dilemma, we've seen that the only NE and SPNE result in (D, D) every period.
- It is only possible to deter defection and maintain cooperation in the *infinitely* repeated game, because there is always the possibility of future punishment.
- This also requires that players have a sufficiently high discount factor; otherwise, the threat of future punishment is outweighed by the prospect of immediate gain.

- Thank you!