CUR 412: Game Theory and its Applications, Lecture 14

Prof. Ronaldo CARPIO

June 19, 2015

Announcements

- Homework #5 is due on Tuesday, June 23. You may put it in my mailbox on the 9th floor of Boxue, or in the envelope on my office door at 123 Qiusuo.
- The final will be on July 2, 8:00-10:00 AM, in Boxue 604.
- Final will be closed-book and covers the second half of the course. Old finals and solutions are on the course website.
- Bring a non-programmable calculator.

Review of Last Lecture

- A *repeated game* is a situation where a game is repeatedly played for several (possibly infinite) periods.
- We assume that agents value payoffs in the *present*, more than payoffs in the *future* by a constant *discount factor*, δ.
- δ is between 0 and 1; an agent with a lower δ is said to be more *impatient*, i.e. places lesser weight on the future.
- The overall value of a sequence of payoffs $w^1, w^2, w^3, \ldots$ is given by the *discounted average*:

$$(1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} w^t$$

- A strategy for a repeated game must specify an action for every possible history.

Review of Last Lecture

- We saw some examples of strategies: *Grim Trigger*, *Punish for k periods*, *Tit-for-Tat*.
- Strategies can be defined as a function mapping any history to an action.
- Or, we can represent them with a state diagram.
- To check whether a pair of strategies forms a NE in the infinitely repeated Prisoner's Dilemma, we must:
  - Calculate the discounted average of the payoff stream to both players when they *do not* deviate;
  - Show that there is no alternative strategy for either player that gives a higher discounted average.

Review of Last Lecture

- The value of δ, the discount factor, can affect whether a strategy pair can be a NE or not.
- Generally, δ must be higher than some level for a NE to exist.
- That is, the threat of future punishment can deter defection in the present only if players place a high enough weight on the future.
- Some known NE: when both players use *Grim Trigger* or *Tit-for-Tat*.

Ch 14.9: Subgame Perfect Equilibria and the One-Deviation Property

- We know that Nash equilibria of an extensive form game may include threats that are not credible.
- This is particularly important in repeated games, since the threat of punishment is the only way to deter a player from playing *Defect*.
- If punishment is not credible, then there will be no reason to expect anything but (D, D).
- The concept of subgame perfect equilibrium dealt with this issue by requiring that strategies also be Nash equilibria in all subgames.
- However, this is not easy to check in an infinite game.
- The good news is that we can take advantage of a result that gives a much simpler condition that is easier to check.

The One-Deviation Property

- Suppose we have a strategy profile s. s satisfies the **one-deviation property** if no player can increase his payoff in any subgame through a *one-shot deviation*:
  - 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.
- This is also called the *one-shot deviation principle*.

The One-Deviation Property

- The basic idea is that any change in strategy can be broken down into a sequence of one-period changes.
- If there is no one-period change that can increase payoffs, then there is no change in strategy that can increase payoffs.
- A full proof is somewhat technically advanced; however, if you are interested, you can search for "one-shot deviation principle".
- For repeated games, this means that in order to check whether some strategy profile s is a SPNE, we must check all possible one-shot deviations for all players.
- for every possible history, compare the payoffs to adhering to s, versus deviating for one period, then reverting back to s.
- Let's look at some of the strategies we've seen, and check if it is a SPNE when both players play them.

Grim Trigger

- *Grim Trigger* is a strategy that punishes defection, so it seems like it should be possible to be part of a SPNE.
- We will show that it is *not* a SPNE when both players play Grim Trigger.
- Suppose the previous history ended with (C, D).
- We don't specify how this occurred; in fact, if both players adhere to Grim Trigger, then the outcome will be (C, C) in every period. However, both players always have the choice of playing D, so it is one of the possible subgames.

Grim Trigger

- First, calculate the payoffs from *not* deviating. If both players adhere to *Grim Trigger*:
  - outcome path will be: $(D, C), (D, D), (D, D), \ldots$ (Player 1 will start punishing in the first period, Player 2 will start in the second period)
  - payoffs will be: $(3, 0), (1, 1), (1, 1), \ldots$
  - Player 1's discounted average:

$$(1-\delta)(3+\delta+\delta^2+\ldots) = (1-\delta)(3+\frac{\delta}{1-\delta}) = 3-2\delta$$

  - Player 2's discounted average:

$$(1-\delta)(0+\delta+\delta^2+\ldots) = \delta$$

Grim Trigger

- Now, suppose Player 2 deviates by choosing D in all periods (note that this is not a one-shot deviation).
  - outcome path: $(D, D), (D, D), (D, D), \ldots$
  - payoffs: $(1, 1), (1, 1), \ldots$
  - Both players' discounted average is 1.
- By our assumptions on $\delta, 1 > \delta$.
- Therefore, Player 2 has an incentive to deviate, given a history ending with (C, D), so this violates the requirement for a SPNE.

Modified Grim Trigger

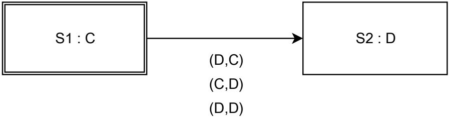

- However, there is a modified version of Grim Trigger that is a SPNE when both players play it.
- *Grim Trigger* moves into the punishment state if the *other* player plays D.
- In contrast, *Modified Grim Trigger* moves into the punishment state if *either* player plays D.

Modified Grim Trigger

- Now, let's use the One-Deviation Property to show that it is a SPNE when both players play *Modified Grim Trigger*.
- We need to check all possible possible histories for the case where a one-shot deviation increases a player's payoff.
- This is vastly simplified because the strategy distinguishes between two types of histories:
  - D has never been played by either player (therefore, play C)
  - D has been played by either player (therefore, play D)
- For each of these types of histories, we will check the conditions that satisfy the one-deviation princple.

Case 1: D has not been played

- Case 1: Suppose D has never been played by either player.
- This is the case if the history is empty (i.e. at the beginning of the game), or if only (C, C) has been played in all periods.
- First, let's calculate the payoffs from *not* deviating.
  - outcome path: $(C, C), (C, C), \ldots$
  - payoffs: $(2, 2), (2, 2), \ldots$
  - Both players' discounted average is 2.

Case 1: D has not been played

- Now, consider the payoffs from a one-shot deviation. Suppose Player 1 does the one-shot deviation and plays D in the first period, then reverts back to *Modified Grim Trigger*.
  - outcome path: $(D, C), (D, D), (D, D), \ldots$
  - payoffs: $(3, 0), (1, 1), (1, 1), \ldots$
  - Player 1's discounted average is $3 - 2\delta$, same as shown above for Grim Trigger.
- Therefore, this one-shot deviation will give a higher payoff if $3 - 2\delta > 2$, or if $\delta < \frac{1}{2}$.
- The same is true if Player 2 does the one-shot deviation.

Case 2: D has been played

- Case 2: D has been played in the past, by either player.
- Then, *Modified Grim Trigger* will switch to the punishment state. If the players do *not* deviate:
  - outcome path: $(D, D), (D, D), ...$
  - payoffs: $(1, 1), (1, 1), ...$
  - Both players' discounted average is 1.

Case 2: D has been played

- Suppose Player 1 does a one-shot deviation by playing C.
  - outcome path: $(C, D), (D, D), (D, D), ...$
  - payoffs: $(0, 3), (1, 1), (1, 1), ...$
  - Player 1's discounted average is $(1 - \delta)(\delta + \delta^2 + ...) = \delta$
- This one-shot deviation will not give a higher payoff for any value of δ.
- The same is true if it is Player 2 that considers a one-shot deviation.

Modified Grim Trigger

- The requirements for the two cases we've seen are:
  - Case 1: D has never been played. A one-shot deviation is profitable if $\delta < \frac{1}{2}$.
  - Case 2: D has been played in the past, by either player. A one-shot deviation is never profitable.
- Combining the requirements for both cases, a profitable (i.e. higher payoff) one-shot deviation does not exist if $\delta \geq \frac{1}{2}$, and therefore the strategy profile is a SPNE.

Tit for Tat

- Recall: $Tit - for - Tat$ plays C at the beginning of the game, then plays whatever the other player did in the previous period.
- Therefore, in terms of determining the outcome, we need only look at the previous round's actions. There are four possibilities:
  1. History ends in (C, D)
  2. History ends in (D, C)
  3. History ends in (C, C)
  4. History ends in (D, D)
- Let's examine each case in turn. We will find the conditions on δ that ensure a profitable one-shot deviation does not exist.
- We need to check both players' possible one-shot deviations.

Case 1: Previous outcome was (C, D)

- Suppose the history ends with (C, D).
- If both players do *not* deviate:
  - outcome path: $(D, C), (C, D), (D, C), ...$ (both players will alternate between C and D forever)
  - payoffs: $(3, 0), (0, 3), (3, 0), ...$
  - Player 1's discounted average:

$$(1-\delta)(3+0+3\delta^2+0+3\delta^4+...) = (1-\delta)\frac{3}{1-\delta^2} = \frac{3}{1+\delta}$$

  - Player 2's discounted average:

$$(1-\delta)(0+3\delta+0+3\delta^3+0+...) = (1-\delta)\frac{3\delta}{1-\delta^2} = \frac{3\delta}{1+\delta}$$

Case 1: Previous outcome was (C, D)

- Suppose Player 1 does a one-shot deviation by playing C in the first period.
  - outcome path: $(C, C), (C, C), \ldots$
  - payoffs: $(2, 2), (2, 2), \ldots$
  - Player 1's discounted average is 2.
- Therefore, this one-shot deviation will give a higher payoff if $2 > \frac{3}{1+\delta}$, or if $\delta > \frac{1}{2}$.
- Now, suppose Player 2 does a one-shot deviation by playing D in the first period:
  - outcome path: $(D, D), (D, D), \ldots$
  - payoffs: $(1, 1), (1, 1), \ldots$
  - Player 2's discounted average is 1.
- Therefore, this one-shot deviation will give a higher payoff if $1 > \frac{3\delta}{1+\delta}$, or if $\delta < \frac{1}{2}$.
- Combining the two conditions, a profitable one-shot deviation does not exist if $\delta = \frac{1}{2}$.

Case 2: Previous outcome was (D, C)

- Suppose the previous outcome was (D, C).
- This is the same situation as the previous case, with the players reversed.
- Therefore, the result is the same: a profitable one-shot deviation does not exist if $\delta = \frac{1}{2}$.

Case 3: Previous outcome was (C, C)

- Suppose the previous outcome was (C, C).
- If both players do *not* deviate:
  - outcome path: $(C, C), (C, C)...$
  - payoffs: $(2, 2), (2, 2), (2, 2), ...$
  - Both players' discounted average: 2

Case 3: Previous outcome was (C, C)

- Suppose Player 1 does a one-shot deviation by playing D in the first round.
  - outcome path: $(D, C), (C, D), (D, C), (C, D), \ldots$
  - payoffs: $(3, 0), (0, 3), (3, 0), \ldots$
  - Player 1's discounted average: $\frac{3}{1+\delta}$ (shown in the first type of history).
- The one-shot deviation will be profitable if $\frac{3}{1+\delta} > 2$, or if $\delta < \frac{1}{2}$.
- The same is true if it is Player 2 that is considering a one-shot deviation.
- Therefore, a profitable one-shot deviation does not exist if $\delta \geq \frac{1}{2}$.

Case 4: Previous outcome was (D, D)

- Suppose the previous outcome was (D, D).
- If both players do *not* deviate:
  - outcome path: $(D, D), (D, D)...$
  - payoffs: $(1, 1), (1, 1), (1, 1), ...$
  - Both players' discounted average: 1

Case 4: Previous outcome was (D, D)

- Suppose Player 1 does a one-shot deviation by playing C in the first round.
  - outcome path: $(C, D), (D, C), (C, D), \ldots$
  - payoffs: $(0, 3), (3, 0), (0, 3), (3, 0), \ldots$
  - Player 1's discounted average: $\frac{3\delta}{1+\delta}$ (see Case 1 for the formula for the discounted avg)
- The one-shot deviation will be profitable if $\frac{3\delta}{1+\delta} > 1$, or if $\delta > \frac{1}{2}$.
- The same is true if it is Player 2 that is considering a one-shot deviation.
- A profitable one-shot deviation will not exist if $\delta \leq \frac{1}{2}$.

Tit for Tat

- Summarizing the results for each case:
  - Case 1 (previous outcome was (C, D)): no profitable deviation exists if $\delta = \frac{1}{2}$.
  - Case 2 (previous outcome was (D, C)): no profitable deviation exists if $\delta = \frac{1}{2}$.
  - Case 3 (previous outcome was (C, C)): no profitable deviation exists if $\delta \geq \frac{1}{2}$.
  - Case 4 (previous outcome was (D, D)): no profitable deviation exists if $\delta \leq \frac{1}{2}$.
- Combining all the conditions, no profitable deviation exists if and only if $\delta = \frac{1}{2}$.
- This result depends on the particular payoffs of the game in each period.
- If we change the payoffs, the range of δ that supports a SPNE may change.

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.

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.

Announcements

- Homework #5 is due on Tuesday, June 23. You may put it in my mailbox on the 9th floor of Boxue, or in the envelope on my office door at 123 Qiusuo.
- The final will be on July 2, 8:00-10:00 AM, in Boxue 604.
- Final will be closed-book and covers the second half of the course. Old finals and solutions are on the course website.
- Bring a non-programmable calculator.