Lecture 19: Infinitely Repeated Games

Mauricio Romero

Lecture 19: Infinitely Repeated Games

Introduction to Infinitely Repeated Games

Introduction to Infinitely Repeated Games

- One of the features of **finitely** repeated games was that if the stage game had a **unique** Nash equilibrium, then the only subgame perfect Nash equilibrium was the repetition of that unique stage game Nash equilibrium

- One of the features of **finitely** repeated games was that if the stage game had a **unique** Nash equilibrium, then the only subgame perfect Nash equilibrium was the repetition of that unique stage game Nash equilibrium

- This happened because there was a last period from which we could induct backwards (and there was a domino effect!)

- One of the features of **finitely** repeated games was that if the stage game had a **unique** Nash equilibrium, then the only subgame perfect Nash equilibrium was the repetition of that unique stage game Nash equilibrium

- This happened because there was a last period from which we could induct backwards (and there was a domino effect!)

- When the game is instead **infinitely** repeated, this argument no longer applies since there is no such thing as a last period

- Lets first define what an infinitely repeated game is

- Lets first define what an infinitely repeated game is

- We start with a stage game whose utilities are given by $u_1, u_2, \ldots, u_n$

- Lets first define what an infinitely repeated game is
- We start with a stage game whose utilities are given by $u_1, u_2, \ldots, u_n$
- Each player i has an action set A_i

- Lets first define what an infinitely repeated game is
- We start with a stage game whose utilities are given by $u_1, u_2, \ldots, u_n$
- Each player i has an action set A_i
- In each period $t = 0, 1, 2, \ldots$, players simultaneously choose an action $a_i \in A_i$ and the chosen action profile $(a_1, a_2, \ldots, a_n)$ is observed by all players

- Lets first define what an infinitely repeated game is
- We start with a stage game whose utilities are given by $u_1, u_2, \ldots, u_n$
- Each player i has an action set A_i
- In each period $t = 0, 1, 2, \ldots$, players simultaneously choose an action $a_i \in A_i$ and the chosen action profile $(a_1, a_2, \ldots, a_n)$ is observed by all players
- Then play moves to period $t + 1$ and the game continues in the same manner.

- It is impossible to draw the extensive form of this infinitely repeated game

- It is impossible to draw the extensive form of this infinitely repeated game
- Each information set of each player i associated with a finitely repeated game corresponded to a history of action profiles chosen in the past

- It is impossible to draw the extensive form of this infinitely repeated game
- Each information set of each player i associated with a finitely repeated game corresponded to a history of action profiles chosen in the past
- We can represent each information set of player i by a history:

$$h^0 = (\emptyset), h^1 = (a^0 := (a_1^0, \ldots, a_n^0)), \ldots, h^t = (a^0, a^1, \ldots, a^{t-1})$$

- It is impossible to draw the extensive form of this infinitely repeated game
- Each information set of each player i associated with a finitely repeated game corresponded to a history of action profiles chosen in the past
- We can represent each information set of player i by a history:

$$h^0 = (\emptyset), h^1 = (a^0 := (a_1^0, \ldots, a_n^0)), \ldots, h^t = (a^0, a^1, \ldots, a^{t-1})$$

- We denote the set of all histories at time t as H^t

Prisoner's Dilemma

	C_2	D_2
C_1	$1, 1$	$-1, 2$
D_1	$2, -1$	$0, 0$

- For example, if the stage game is the prisoner's dilemma, at period 1, there are 4 possible histories:

$$\{(C_1, C_2), (C_1, D_2), (D_1, C_2), (D_1, D_2)\} = H^1.$$

- For example, if the stage game is the prisoner's dilemma, at period 1, there are 4 possible histories:

$$\{(C_1, C_2), (C_1, D_2), (D_1, C_2), (D_1, D_2)\} = H^1.$$

- For time t, H^t consists of 4^t possible histories

- For example, if the stage game is the prisoner's dilemma, at period 1, there are 4 possible histories:

$$\{(C_1, C_2), (C_1, D_2), (D_1, C_2), (D_1, D_2)\} = H^1.$$

- For time t, H^t consists of 4^t possible histories

- This means that there is a **one-to-one** mapping between all possible histories and the information sets if we actually wrote out the whole extensive form game tree

- For example, if the stage game is the prisoner's dilemma, at period 1, there are 4 possible histories:

$$\{(C_1, C_2), (C_1, D_2), (D_1, C_2), (D_1, D_2)\} = H^1.$$

- For time t, H^t consists of 4^t possible histories

- This means that there is a **one-to-one** mapping between all possible histories and the information sets if we actually wrote out the whole extensive form game tree

- As a result, we can think of each $h^t \in H^t$ as representing a particular information set for each player i in each time t

- What is a strategy in an infinitely repeated game?

- What is a strategy in an infinitely repeated game?

- It is simply a prescription of what player i would do at every information set or history

- What is a strategy in an infinitely repeated game?
- It is simply a prescription of what player i would do at every information set or history
- Therefore, it is a function that describes:

$$s_i : \bigcup_{t \geq 0} H^t \to A_i.$$

- What is a strategy in an infinitely repeated game?
- It is simply a prescription of what player i would do at every information set or history
- Therefore, it is a function that describes:

$$s_i : \bigcup_{t \geq 0} H^t \to A_i.$$

- Intuitively, s_i describes exactly what player i would do at every possible history h^t, where $s_i(h^t)$ describes what player i would do at history h^t

- For example in the infinitely repeated prisoner's dilemma, the strategy $s_i(h^t) = C_i$ for all h^t and all t is the strategy in which player i always plays C_i regardless of the history

- For example in the infinitely repeated prisoner's dilemma, the strategy $s_i(h^t) = C_i$ for all h^t and all t is the strategy in which player i always plays C_i regardless of the history

- There can be more complicated strategies such as the following:

$$s_i(h^t) = \begin{cases} C_i & \text{if } t = 0 \text{ or } h^t = (C, C, \dots, C), \\ D_i & \text{otherwise.} \end{cases}$$

- For example in the infinitely repeated prisoner's dilemma, the strategy $s_i(h^t) = C_i$ for all h^t and all t is the strategy in which player i always plays C_i regardless of the history

- There can be more complicated strategies such as the following:

$$s_i(h^t) = \begin{cases} C_i & \text{if } t = 0 \text{ or } h^t = (C, C, \ldots, C), \\ D_i & \text{otherwise.} \end{cases}$$

- The above is called a **grim trigger strategy**

- How are payoffs determined in the repeated game?

- How are payoffs determined in the repeated game?

- Suppose the *strategies* $s_1, \ldots, s_n$ are played which lead to the infinite sequence of action profiles:

$$a^0, a^1, \ldots, a^t, a^{t+1}, \ldots.$$

- How are payoffs determined in the repeated game?
- Suppose the *strategies* $s_1, \ldots, s_n$ are played which lead to the infinite sequence of action profiles:
$$a^0, a^1, \ldots, a^t, a^{t+1}, \ldots.$$
- Then the payoff of player i in this repeated game is given by:
$$\sum_{t=0}^{\infty} \delta^t u_i(a^t).$$

- How are payoffs determined in the repeated game?
- Suppose the *strategies* $s_1, \ldots, s_n$ are played which lead to the infinite sequence of action profiles:
$$a^0, a^1, \ldots, a^t, a^{t+1}, \ldots.$$
- Then the payoff of player i in this repeated game is given by:
$$\sum_{t=0}^{\infty} \delta^t u_i(a^t).$$
- Intuitively, the contribution to payoff of time t action profile a^t is discounted by δ^t

- It may be unreasonable to think about an infinitely repeated game

- It may be unreasonable to think about an infinitely repeated game
- However the discount factor instead could be interpreted by the probability of the game/relationship ending at any point in time.

- It may be unreasonable to think about an infinitely repeated game
- However the discount factor instead could be interpreted by the probability of the game/relationship ending at any point in time.
- Thus, an infinitely repeated game does not necessarily represent a scenario in which there are an infinite number of periods, but rather a relationship which ends in finite time with probability one, but in which the time at which the relationship ends is uncertain

- Lets see some examples of how to compute payoffs in the repeated game

- Lets see some examples of how to compute payoffs in the repeated game
- Consider first the strategy profile in which $s_i(h^t) = C_i$ for all $i = 1, 2$ and all h^t.

- Lets see some examples of how to compute payoffs in the repeated game
- Consider first the strategy profile in which $s_i(h^t) = C_i$ for all $i = 1, 2$ and all h^t.
- In this case, the payoff of player 1 in this repeated game is given by:

$$\sum_{t=0}^{\infty} \delta^t = \frac{1}{1-\delta}$$

- Lets see some examples of how to compute payoffs in the repeated game
- Consider first the strategy profile in which $s_i(h^t) = C_i$ for all $i = 1, 2$ and all h^t.
- In this case, the payoff of player 1 in this repeated game is given by:

$$\sum_{t=0}^{\infty} \delta^t = \frac{1}{1-\delta}$$

- What about in the grim trigger strategy profile?

- Lets see some examples of how to compute payoffs in the repeated game
- Consider first the strategy profile in which $s_i(h^t) = C_i$ for all $i = 1, 2$ and all h^t.
- In this case, the payoff of player 1 in this repeated game is given by:

$$\sum_{t=0}^{\infty} \delta^t = \frac{1}{1-\delta}$$

- What about in the grim trigger strategy profile?
- In that case, if all players play the grim trigger strategy profile, the sequence of actions that arise is again $(C, C, \ldots)$

- Lets see some examples of how to compute payoffs in the repeated game
- Consider first the strategy profile in which $s_i(h^t) = C_i$ for all $i = 1, 2$ and all h^t.
- In this case, the payoff of player 1 in this repeated game is given by:

$$\sum_{t=0}^{\infty} \delta^t = \frac{1}{1-\delta}$$

- What about in the grim trigger strategy profile?
- In that case, if all players play the grim trigger strategy profile, the sequence of actions that arise is again $(C, C, \ldots)$
- Thus the payoffs of all players is again $\frac{1}{1-\delta}$.

- How about a more complicated strategy profile?

- How about a more complicated strategy profile?
- Suppose that $s_i(h^0) = (C_1, D_2)$ and the strategy profile says to do exactly what the opponent did in the previous period

- How about a more complicated strategy profile?

- Suppose that $s_i(h^0) = (C_1, D_2)$ and the strategy profile says to do exactly what the opponent did in the previous period

- Then if both players play these strategies, then the sequence of actions that arise is:
$$(C_1, D_2), (D_1, C_2), (C_1, D_2), \ldots$$

- How about a more complicated strategy profile?

- Suppose that $s_i(h^0) = (C_1, D_2)$ and the strategy profile says to do exactly what the opponent did in the previous period

- Then if both players play these strategies, then the sequence of actions that arise is:
$$(C_1, D_2), (D_1, C_2), (C_1, D_2), \ldots$$

- Then the payoff to player 1 in this game is given by:
$$\sum_{t=0}^{\infty} \delta^{2t}(-1) + \delta^{2t+1} \cdot 2 = \frac{-1}{1-\delta^2} + \frac{2\delta}{1-\delta^2} = \frac{2\delta - 1}{1-\delta^2}.$$

- What is a subgame perfect Nash equilibrium in an infinitely repeated game?

- What is a subgame perfect Nash equilibrium in an infinitely repeated game?
- It is exactly the same idea as in the finitely repeated game or more generally extensive form games

- What is a subgame perfect Nash equilibrium in an infinitely repeated game?
- It is exactly the same idea as in the finitely repeated game or more generally extensive form games
- That is a strategy profile $s = (s_1, \ldots, s_n)$ is a subgame perfect game Nash equilibrium if and only if s is a Nash equilibrium in every subgame of the repeated game.

Theorem (One-stage deviation principle)

s is a subgame perfect Nash equilibrium (SPNE) if and only if at every time t, and every history and every player i, player i cannot profit by deviating just at time t and following the strategy s_i' from time $t+1$ on

- This is extremely useful since we only need to check that s_i is optimal against all possible one-stage deviations rather than having to check that it is optimal against all s_i'.

- This is extremely useful since we only need to check that s_i is optimal against all possible one-stage deviations rather than having to check that it is optimal against all s_i'.

- We will now put this into practice to analyze subgame perfect Nash equilibria of infinitely repeated games

- Lets go back to the infinitely repeated prisoner's dilemma

- Lets go back to the infinitely repeated prisoner's dilemma
- What is an example of a subgame perfect Nash equilibrium?

- Lets go back to the infinitely repeated prisoner's dilemma

- What is an example of a subgame perfect Nash equilibrium?

- One kind of equilibrium should be straightforward: each player plays D_1 and D_2 always at all information sets

- Lets go back to the infinitely repeated prisoner's dilemma
- What is an example of a subgame perfect Nash equilibrium?
- One kind of equilibrium should be straightforward: each player plays D_1 and D_2 always at all information sets
- Why is this a SPNE?

- Lets go back to the infinitely repeated prisoner's dilemma
- What is an example of a subgame perfect Nash equilibrium?
- One kind of equilibrium should be straightforward: each player plays D_1 and D_2 always at all information sets
- Why is this a SPNE?
- We can use the one-stage deviation principle

Prisoner's Dilemma

	C_2	D_2
C_1	$1, 1$	$-1, 2$
D_1	$2, -1$	$0, 0$

- Under this strategy profile s_1^*, s_2^*, for all histories h^t,

$$V_1(s_1^*, s_2^* \mid h^t) = V_2(s_1^*, s_2^* \mid h^t) = 0.$$

- Under this strategy profile s_1^*, s_2^*, for all histories h^t,

$$V_1(s_1^*, s_2^* \mid h^t) = V_2(s_1^*, s_2^* \mid h^t) = 0.$$

- Thus, for all histories h^t,

$$\underbrace{u_i(D_i, D_{-i})}_{0} + \delta \underbrace{V_i(s_1^*, s_2^* \mid h^t)}_{0} > \underbrace{u_i(C_i, D_{-i})}_{-1} + \delta \underbrace{V_i(s_1^*, s_2^* \mid h^t)}_{0}$$

- Under this strategy profile s_1^*, s_2^*, for all histories h^t,

$$V_1(s_1^*, s_2^* \mid h^t) = V_2(s_1^*, s_2^* \mid h^t) = 0.$$

- Thus, for all histories h^t,

$$\underbrace{u_i(D_i, D_{-i})}_{0} + \delta \underbrace{V_i(s_1^*, s_2^* \mid h^t)}_{0} > \underbrace{u_i(C_i, D_{-i})}_{-1} + \delta \underbrace{V_i(s_1^*, s_2^* \mid h^t)}_{0}$$

- Thus, (s_1^*, s_2^*) is a SPNE

In fact this is not specific to the prisoner's dilemma as we show below:

Theorem

Let a^* be a Nash equilibrium of the stage game. Then the strategy profile s^* in which all players i play a_i^* at all information sets is a SPNE for any $\delta \in [0, 1)$.

- What other kinds of SPNE are there?

- What other kinds of SPNE are there?

- In finitely repeated games, this was the only SPNE with prisoner's dilemma since the stage game had a unique Nash equilibrium

- What other kinds of SPNE are there?

- In finitely repeated games, this was the only SPNE with prisoner's dilemma since the stage game had a unique Nash equilibrium

- When the repeated game is infinitely repeated, this is no longer true

- Consider for example the grim trigger strategy profile that we discussed earlier. Each player plays the following strategy:

$$s_i^*(h^t) = \begin{cases} C_i & \text{if } h^t = (C, C, \ldots, C) \\ D_i & \text{if } h^t \neq (C, C, \ldots, C). \end{cases}$$

- Consider for example the grim trigger strategy profile that we discussed earlier. Each player plays the following strategy:

$$s_i^*(h^t) = \begin{cases} C_i & \text{if } h^t = (C, C, \dots, C) \\ D_i & \text{if } h^t \neq (C, C, \dots, C). \end{cases}$$

- We will show that if δ is sufficiently high, so that the players are sufficiently patient, the strategy profile of grim trigger strategies is indeed a SPNE

- Consider for example the grim trigger strategy profile that we discussed earlier. Each player plays the following strategy:

$$s_i^*(h^t) = \begin{cases} C_i & \text{if } h^t = (C, C, \dots, C) \\ D_i & \text{if } h^t \neq (C, C, \dots, C). \end{cases}$$

- We will show that if δ is sufficiently high, so that the players are sufficiently patient, the strategy profile of grim trigger strategies is indeed a SPNE

- The *equilibrium path of play* for this SPNE is for players to play C in every period

- How do we show that the above is indeed an SPNE?

- How do we show that the above is indeed an SPNE?

- We use the one-stage deviation principle again

- How do we show that the above is indeed an SPNE?
- We use the one-stage deviation principle again
- We need to check the one-stage deviation principle at every history h^t.

Case 1:

- Suppose first that $h^t \neq (C, C, \ldots, C)$

Case 1:

- Suppose first that $h^t \neq (C, C, \ldots, C)$
- Players are each suppose to play D_i

Case 1:

- Suppose first that $h^t \neq (C, C, \ldots, C)$
- Players are each suppose to play D_i
- Thus, we need to check that

$$u_i(D_i, D_{-i}) + \delta V_i(s^* \mid (h^t, D))$$
$$\geq u_i(C_i, D_{-i}) + \delta V_i(s^* \mid (h^t, (C_i, D_{-i})))$$

Case 1:

- Suppose first that $h^t \neq (C, C, \ldots, C)$
- Players are each suppose to play D_i
- Thus, we need to check that

$$
\begin{aligned}
& u_i(D_i, D_{-i}) + \delta V_i(s^* \mid (h^t, D)) \\
& \geq u_i(C_i, D_{-i}) + \delta V_i(s^* \mid (h^t, (C_i, D_{-i})))
\end{aligned}
$$

- But since $h^t \neq (C, C, \ldots, C)$,
  $V_i(s^* \mid (h^t, D)) = V_i(s^* \mid (h^t, (C_i, D_{-i}))) = u_i(D_i, D_{-i})$.

Case 1:

- Suppose first that $h^t \neq (C, C, \ldots, C)$
- Players are each suppose to play D_i
- Thus, we need to check that

$$u_i(D_i, D_{-i}) + \delta V_i(s^* \mid (h^t, D)) \\ \geq u_i(C_i, D_{-i}) + \delta V_i(s^* \mid (h^t, (C_i, D_{-i})))$$

- But since $h^t \neq (C, C, \ldots, C)$,
  $V_i(s^* \mid (h^t, D)) = V_i(s^* \mid (h^t, (C_i, D_{-i}))) = u_i(D_i, D_{-i})$.
- So the above inequality is satisfied if and only if

$$u_i(D_i, D_{-i}) \geq u_i(C_i, D_{-i}).$$

Case 1:

- Suppose first that $h^t \neq (C, C, \ldots, C)$
- Players are each suppose to play D_i
- Thus, we need to check that

$$u_i(D_i, D_{-i}) + \delta V_i(s^* \mid (h^t, D)) \\ \geq u_i(C_i, D_{-i}) + \delta V_i(s^* \mid (h^t, (C_i, D_{-i})))$$

- But since $h^t \neq (C, C, \ldots, C)$, $V_i(s^* \mid (h^t, D)) = V_i(s^* \mid (h^t, (C_i, D_{-i}))) = u_i(D_i, D_{-i})$.
- So the above inequality is satisfied if and only if

$$u_i(D_i, D_{-i}) \geq u_i(C_i, D_{-i}).$$

- But this is satisfied since D is a Nash equilibrium of the stage game

Case 2:

- Suppose instead that $h^t = (C, C, \ldots, C)$

Case 2:

- Suppose instead that $h^t = (C, C, \ldots, C)$
- Players are both supposed to play C_i

Case 2:

- Suppose instead that $h^t = (C, C, \ldots, C)$
- Players are both supposed to play C_i
- Thus, we need to check that

$$u_i(C_i, C_{-i}) + \delta V_i(s^* \mid (h^t, C)) \\ \geq u_i(D_i, C_{-i}) + \delta V_i(s^* \mid (h^t, (D_i, C_{-i}))).$$

Case 2:

- Suppose instead that $h^t = (C, C, \ldots, C)$
- Players are both supposed to play C_i
- Thus, we need to check that

$$u_i(C_i, C_{-i}) + \delta V_i(s^* \mid (h^t, C)) \\ \geq u_i(D_i, C_{-i}) + \delta V_i(s^* \mid (h^t, (D_i, C_{-i}))).$$

- In this case,

$$V_i(s^* \mid (h^t, C)) = u_i(C_i, C_{-i}) \\ = 1, V_i(s^* \mid (h^t, (D_i, C_{-i}))) = u_i(D) = 0.$$

Case 2:

- Suppose instead that $h^t = (C, C, \dots, C)$
- Players are both supposed to play C_i
- Thus, we need to check that

$$u_i(C_i, C_{-i}) + \delta V_i(s^* \mid (h^t, C)) \\ \geq u_i(D_i, C_{-i}) + \delta V_i(s^* \mid (h^t, (D_i, C_{-i}))).$$

- In this case,

$$V_i(s^* \mid (h^t, C)) = u_i(C_i, C_{-i}) \\ = 1, V_i(s^* \mid (h^t, (D_i, C_{-i}))) = u_i(D) = 0.$$

- Therefore, the above is satisfied if and only if

$$1 + \delta \geq 2 \iff \delta \geq 1/2.$$

Case 2:

- Suppose instead that $h^t = (C, C, \ldots, C)$
- Players are both supposed to play C_i
- Thus, we need to check that

$$u_i(C_i, C_{-i}) + \delta V_i(s^* \mid (h^t, C)) \\ \geq u_i(D_i, C_{-i}) + \delta V_i(s^* \mid (h^t, (D_i, C_{-i}))).$$

- In this case,

$$V_i(s^* \mid (h^t, C)) = u_i(C_i, C_{-i}) \\ = 1, V_i(s^* \mid (h^t, (D_i, C_{-i}))) = u_i(D) = 0.$$

- Therefore, the above is satisfied if and only if

$$1 + \delta \geq 2 \iff \delta \geq 1/2.$$

- Thus the grim trigger strategy profile s^* is a SPNE if and only if $\delta \geq 1/2$.

- The above findings that SPNE may involve the repetition of action profile that is not a stage game NE is not specific to just the infinitely repeated prisoner's dilemma as the following theorem demonstrates.

Theorem (Folk theorem)

Suppose that a^* is a Nash equilibrium of the stage game. Suppose that $\hat{a}$ is an action profile of the Nash equilibrium such that

$$u_1(\hat{a}) > u_1(a^*), \ldots, u_n(\hat{a}) > u_n(a^*).$$

Then there is some $\delta^* < 1$ such that whenever $\delta > \delta^*$, there is a SPNE in which on the equilibrium path of play, all players play $\hat{a}$ in every period.