Econ 421
Fall, 2023
Li, Hao
UBC

Lecture 6. Infinite-horizon Games

1. Infinite horizon

Potentially infinite repeated strategic interactions among players.

- An identical or similar *stage game* is repeated, which can be a simultaneous-move game, as well as an extensive-form game.
- Histories of previous stages are all observed, allowing players to condition their play in current stage on the past.

Applications of multi-stage games include:

- Alternating-offer bargaining.
- War of attrition.
- Collusion in repeated oligopoly.

We generalize extensive-form games by allowing infinite histories.

- A terminal history may be an infinite history.
- Evaluating an infinite terminal history: for each i, any payoff u received in t periods in the future is worth $\delta_i^t u$ today, where the discount factor $\delta_i \in (0, 1)$.
- Player i prefers one infinite sequence of payoffs $(u_i^0, \ldots, u_i^t, \ldots)$ to another $(\tilde{u}_i^0, \ldots, \tilde{u}_i^t, \ldots)$ if and only if

$$\sum_{t=0}^{\infty} \delta_i^t u_i^t > \sum_{t=0}^{\infty} \delta_i^t \tilde{u}_i^t.$$

Example (Rubinstein's alternating-offer bargaining game). There are two players who bargain over a dollar. The game begins in period 0 with player 1 proposing a way to split the dollar. If 2 accepts, game ends with the dollar split as proposed by player 1. If 2 rejects, game proceeds to period 1 in which 2 makes a counter proposal. Game ends if player 1 accepts it, otherwise player 1 gets to make a proposal in period 2, and so on. If the game ends in period t with one player proposing $x_t \in [0, 1]$ to the other player and the latter accepting it, then the former's payoff is $\delta^t(1 - x_t)$ and the latter's payoff is $\delta^t x_t$, where δ is the common discount factor. If the game never ends, each player gets 0.

Formally, any history $h^t \in H$ is one of the following:

- $h^t = \emptyset$ or $h^t = (x_0, R, x_1, R, \ldots, x_t, R)$ (x_s is offer in s and R is *Reject*), with $P(h^t) = 1$ if t is odd and $P(h^t) = 2$ if t is even;

- $h^t = (x_0, R, x_1, R, \ldots, x_t)$, with $P(h^t) = 1$ if t is odd, $P(h^t) = 2$ if t is even;

- $h^t = (x_0, R, x_1, R, \ldots, x_t, A)$ (A is *Accept*), with payoffs $\delta^t x_t$ to player 1 and $\delta^t(1 - x_t)$ to player 2 if t is odd, and $\delta^t(1 - x_t)$ to player 1 and $\delta^t x_t$ to player 2 if t is even;

- $h^t = (x_0, R, x_1, R, \ldots,)$, with payoff 0 to both player.

Example (Repeated Cournot duopoly). Suppose that two Cournot duopolists with the same constant marginal cost c (and no fixed cost) face an inverse demand function given by $P(Q) = a - Q$ with $a > c$. Consider this Cournot duopoly game being played for an infinite number of periods. Each firm evaluates an infinite sequence of profits $(\pi_t)_{t=0}^{\infty}$ according the geometric sum $\sum_{t=0}^{\infty} \delta^t \pi_t$, where $\delta \in (0, 1)$ is the common discount factor.

Formally, any $h^t \in H$ except $\emptyset$ is a sequence of pairs of quantities from period 0 to period $t-1$, $((q_1^0, q_2^0), \ldots, (q_1^{t-1}, q_2^{t-1}))$, such that

- $P(h^t) = \{\text{Firm 1}, \text{Firm 2}\}$;
- set of actions for each firm is $A_1(h^t) = A_2(h^t) = [0, \infty)$;
- set of terminal histories is an infinite sequence of pairs of quantities, $((q_1^0, q_2^0), \ldots, (q_1^{t-1}, q_2^{t-1}), \ldots)$, with payoff $\sum_{t=0}^{\infty} \delta^t q_j^t (a - q_1^t - q_2^t - c)$ to Firm j, $j = 1, 2$.

SPE applies without change.

- With infinite terminal histories, backward induction does not work, but we have the following extension.

- **Proposition** (One-shot deviation principle). A strategy profile s^* is a subgame perfect equilibrium in a finite-horizon multi-stage game, or an infinite-horizon game with discounting, if and only if, for each player i and each subgame $\Gamma(h^t)$ with $i \in P(h^t)$, conditional on reaching h^t there is no action $a_i \in A_i(h^t)$ such that player i obtains a strictly higher payoff in $\Gamma(h^t)$ by deviating from $s_i^*(h^t)$ to a_i in period t and then reverting back to s_i^* afterwards.

Basic idea behind one-shot principle.

- In words, one-shot principle says that a strategy profile is SPE if and only if there is no player and no subgame such that the player has a strictly better strategy in the subgame that differs from the player's proposed strategy in one single move.

- The "only if" part is immediate.

- The "if" part implies that there is no need to check multi-shot deviations, and follows from Principle of Optimality that there is no profitable multi-shot deviation if there is none in one shot.

Proof. "Only if" follows from definition of SPE.

For "if," fix a strategy profile s^*. Suppose that there is no profitable one-shot deviation from s_i^* against s_{-i}^* for any i but that s^* is not a SPE. Then, there exist a player i, a subgame $\Gamma(h^t)$, and a strategy s_i such that s_i is a strictly better response than s_i^* to s_{-i}^* in $\Gamma(h^t)$. Let τ be the largest $t' > t$ such that there is a history $h^{t'}$ following h^t with $s_i(h^{t'}) \neq s_i^*(h^{t'})$.

- (i) Suppose τ is finite. Define a new strategy $s_i^{\tau-1}$ in $\Gamma(h^t)$ that agrees with s_i from period t to $\tau-1$ and with s_i^* afterwards. Then $s_i^{\tau-1}$ is also is a strictly better response than s_i^* to s_{-i}^* in $\Gamma(h^t)$. The same is true if we define $s_i^{\tau-2}$ that agrees with s_i up to period $\tau-2$ and with s_i^* afterwards. A contradiction eventually.

- (ii) Suppose τ is infinite. Due to discounting, there is $\tilde{\tau}$ sufficiently large, such that the strategy $\tilde{s}_i$ in $\Gamma(h^t)$ given by s_i up to period $\tilde{\tau}$ and by s_i^* afterwards is also a strictly better response than s_i^* to s_{-i}^* in $\Gamma(h^t)$. Apply (i).

2. Alternating-offer bargaining

Construct SPE that is stationary (equilibrium offer is constant regardless of who makes the offer when) and efficient (equilibrium offer is accepted with probability one).

- Equilibrium offer satisfies

$$x_* = \delta(1 - x_*),$$

which gives

$$x_* = \frac{\delta}{1+\delta}.$$

- Equilibrium is then given by: always offer x_* as proposer, and accept any offer $x \geq x_*$ and reject all $x < x_*$.

- Verify the equilibrium by one-shot deviation principle.

  - Proposer strictly prefers having x_* accepted to accepting x_* in next period.

  - Responder is indifferent between accepting x_* and rejecting it and having x_* accepted in next period.

Equilibrium properties.

- Efficiency: there is no delay in equilibrium.
- First-mover advantage: player 1's equilibrium payoff is greater than player 2's equilibrium payoff.
  - Advantage arises from the delay in making a counter offer.
  - It disappears when δ goes to 1.

3. Repeated prisoner's dilemma

Consider a repeated Prisoner's Dilemma game, with common discount factor $\delta \in (0, 1)$ and stage game given by

	D	C
D	$1, 1$	$3, 0$
C	$0, 3$	$2, 2$

Finitely repeated.

- Backward induction implies a unique SPE.
- In the equilibrium, the outcome is (D, D) in each period.
- There is no possibility to support the Pareto efficient outcome (C, C) in a subgame perfect equilibrium, even though repeated plays allow players to condition their players in the current on the history of the past players.

Infinitely repeated.

- The Pareto efficient stage outcome (C, C) can be sustained in every period in a SPE if players are sufficiently patient.

- Consider *trigger strategy*: start with C; continue with it so long as the last period's outcome is (C, C); otherwise permanently switch to D regardless of history.

- Use one-shot principle to verify equilibrium.
  - In any subgame at $h^t \neq \{(C, C), \ldots, (C, C)\}$, there is no profitable one-shot deviation from D.
  - In any subgame at $h^t = \{(C, C), \ldots, (C, C)\}$, there is no profitable one-shot deviation from C if

$$2 + \delta \cdot 2 + \delta^2 \cdot 2 + \ldots \geq 3 + \delta \cdot 1 + \delta^2 \cdot 1 + \ldots,$$

which is equivalent to $\delta \geq \frac{1}{2}$.

- Generalized Prisoner's Dilemma, with $d > l$ and $h > c > d$:

	D	C
D	d, d	h, l
C	l, h	c, c

  - To support (C, C) as equilibrium outcome, we need

$$c + \delta \cdot c + \delta^2 \cdot c + \ldots \geq h + \delta \cdot d + \delta^2 \cdot d + \ldots,$$

    which is equivalent to $(c - d)\delta/(1 - \delta) \geq h - c$.

  - Present value of future losses is greater than the one time gain from deviation.

- Tit-for-tat.
  - Consider the following: start with C; in every subsequent period, play the action chosen by the opponent.
  - Use one-deviation principle to show that a pair of tit-for-tat strategies is not a subgame perfect equilibrium except for a single value of δ.
  - The legend of tit-for-tat, and an explanation.

4. Repeated Cournot duopoly

Benchmarks.

- The collusive outcome is each firm producing $\frac{1}{4}(a-c)$, with profit $\frac{1}{8}(a-c)^2$.
- The unique NE outcome in the stage game is each firm producing $\frac{1}{3}(a-c)$, with profit $\frac{1}{9}(a-c)^2$.

Nash-trigger strategies.

- There is a critical value of $\hat{\delta}$, such that for all $\delta \geq \hat{\delta}$, there is SPE that supports collusive outcome, using a strategy that triggers a permanent switch to NE after any deviation from collusion.

- For $\delta < \hat{\delta}$, Nash-trigger strategy cannot support collusion, but can support an outcome more profitable than NE.

- Can firms do even better than most profitable outcome supported by Nash-trigger strategy when $\delta < \hat{\delta}$?

Answer is yes, but need to a find a strategy that is more punishing than permanent switch to NE and is yet credible.

- Consider carrot-and-stick strategy: for some large fraction x to be specified, play the collusive quantity $\frac{1}{4}(a-c)$ if last period outcome is either collusive quantity pair $\left(\frac{1}{4}(a-c), \frac{1}{4}(a-c)\right)$, or $(x(a-c), x(a-c))$; otherwise, play $x(a-c)$.

- Carrot is collusive quantity $\frac{1}{4}(a-c)$; stick is punishment quantity $x(a-c)$, which lasts just one period and is supported by restarting.

- Find necessary and sufficient conditions on x using the one-shot deviation principle.

  - For all subgames at the beginning of period t, after either $\left(\frac{1}{4}(a-c), \frac{1}{4}(a-c)\right)$ or $(x(a-c), x(a-c))$ in previous period,

    $$\frac{1}{8(1-\delta)} \geq \frac{9}{64} + \delta x(1-2x) + \frac{\delta^2}{8(1-\delta)},$$

    which is equivalent to

    $$\delta(4x-1)^2 \geq \frac{1}{8}.$$

– For all other subgames at the beginning of period t,

$$x(1-2x)+\frac{\delta}{8(1-\delta)} \geq \frac{(1-x)^2}{4}+\delta x(1-2x)+\frac{\delta^2}{8(1-\delta)},$$

which is equivalent to

$$\delta(4x-1)^2 \geq 2(3x-1)^2.$$

- As long as $\delta \geq \frac{9}{32}$, there exists x that satisfies both inequalities.
  - Collusion can be supported by carrot-and-stick strategy as SPE outcome, so long as $\delta \geq \frac{9}{32}$.
  - This cutoff value of δ is lower than $\hat{\delta}$, the lowest discount factor that allows collusion to be supported by Nash-trigger strategy as SPE outcome.
- For $\delta < \frac{9}{32}$, carrot-and-stick strategy cannot support collusion, but can support an outcome more profitable than NE.