6.254 : Game Theory with Engineering Applications
Lecture 13: Extensive Form Games

Asu Ozdaglar
MIT

March 18, 2010

Outline

- Extensive Form Games with Perfect Information
- One-stage Deviation Principle
- Applications
- Ultimatum Game
- Rubinstein-Stahl Bargaining Model

- Reading:
- Fudenberg and Tirole, Sections 4.1-4.4.

Introduction

- We have studied extensive form games which model sequential decision making.
- Equilibrium notion for extensive form games: Subgame Perfect (Nash) Equilibrium.
- It requires each player's strategy to be "optimal" not only at the start of the game, but also after every history.
- For finite horizon games, found by backward induction.
    - Backward induction refers to starting from the "last" subgames of a finite game, finding the best response strategy profiles or the Nash equilibria in the subgames, then assigning these strategies profiles and the associated payoffs to be subgames, and moving successively towards the beginning of the game.
- For finite/infinite horizon games, characterization in terms of one-stage deviation principle.

One-stage Deviation Principle

- Focus on multi-stage games with observed actions (or perfect information games).
- One-stage deviation principle is essentially the **principle of optimality** of dynamic programming.
- We first state it for finite horizon games.

Theorem (One-stage deviation principle)

For finite horizon multi-stage games with observed actions , s^* is a subgame perfect equilibrium if and only if for all i, t and h^t, we have

$$u_i(s_i^*, s_{-i}^* | h^t) \geq u_i(s_i, s_{-i}^* | h^t)$$

for all s_i satisfying

$$\begin{aligned} s_i(h^t) &\neq s_i^*(h^t), \\ s_{i|h^t}(h^{t+k}) &= s_{i|h^t}^*(h^{t+k}), \quad \textit{for all} \;\; k > 0, \;\textit{and all}\; h^{t+k} \in G(h^t). \end{aligned}$$

- Informally, s is a subgame perfect equilibrium (SPE) if and only if no player i can gain by deviating from s in a single stage and conforming to s thereafter.

One-stage Deviation Principle for Infinite Horizon Games

- The proof of one-stage deviation principle for finite horizon games relies on the idea that if a strategy satisfies the one stage deviation principle then that strategy cannot be improved upon by a finite number of deviations.
- This leaves open the possibility that a player may gain by an infinite sequence of deviations, which we exclude using the following condition.

Definition

Consider an extensive form game with an infinite horizon, denoted by G^∞. Let h denote an ∞-horizon history, i.e., $h = (a^0, a^1, a^2 \ldots)$, is an infinite sequence of actions. Let $h^t = (a^0, \ldots a^{t-1})$ be the restriction to first t periods. The game G^∞ is **continuous at infinity** *if for all players i, the payoff function u_i satisfies*

$$\sup_{h, \tilde{h} \text{ s.t. } h^t = \tilde{h}^t} |u_i(h) - u_i(\tilde{h})| \to 0 \quad as \quad t \to \infty.$$

One-stage Deviation Principle for Infinite Horizon Games

- The continuity at infinity condition is satisfied when the overall payoffs are a discounted sum of stage payoffs, i.e.,

$$u_i = \sum_{t=0}^{\infty} \delta_i^t g_i^t(a^t),$$

(where $g_i^t(a^t)$ are the stage payoffs, the positive scalar $\delta_i < 1$ is a discount factor), and the stage payoff functions are uniformly bounded, i.e., there exists some B such that $\max_{t,a^t} |g_i^t(a^t)| < B$.

Theorem

Consider an infinite-horizon game, G^∞, that is continuous at infinity. Then, the one stage deviation principle holds, i.e., the strategy profile s^* is an SPE if and only if for all i, h^t, and t, we have

$$u_i(s_i^*, s_{-i}^* | h^t) \leq u_i(s_i, s_{-i}^* | h^t),$$

for all s_i that satisfies $s_i(h^t) \neq s_i^*(h^t)$ and $s_i|h^t(h^{t+k}) = s_i^*|h^t(h^{t+k})$ for all $h^{t+k} \in G(h^t)$ and for all $k > 0$.

Examples: Value of Commitment

- Consider the entry deterrence game, but with a different timing as shown in the next figure.

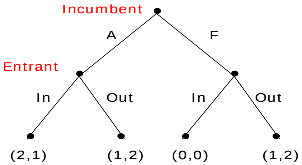

- Note: For consistency, first number is still the entrant's payoff.
- This implies that the incumbent can now commit to fighting (how could it do that?).
- It is straightforward to see that the unique SPE now involves the incumbent committing to fighting and the entrant not entering.
- This illustrates the **value of commitment**.

Examples: Stackleberg Model of Competition

- Consider a variant of the Cournot model where player 1 chooses its quantity q_1 first, and player 2 chooses its quantity q_2 after observing q_1. Here, player 1 is the Stackleberg leader.
- Suppose again that both firms have marginal cost c and the inverse demand function is given by $P(Q) = \alpha - \beta Q$, where $Q = q_1 + q_2$, where $\alpha > c$.
- This is a dynamic game, so we should look for SPE. How to do this?
- Backward induction—this is not a finite game, but all we have seen so far applies to infinite games as well.
- Look at a subgame indexed by player 1 quantity choice, q_1. Then player 2's maximization problem is essentially the same as before

$$\begin{aligned}
\max_{q_2 \geq 0} \pi_2(q_1, q_2) &= [P(Q) - c]\, q_2 \\
&= [\alpha - \beta(q_1 + q_2) - c]\, q_2.
\end{aligned}$$

Stackleberg Competition (continued)

- This gives best response

$$q_2 = \frac{\alpha - c - \beta q_1}{2\beta}.$$

- Now the difference is that player 1 will choose q_1 recognizing that player 2 will respond with the above best response function.
- Player 1 is the Stackleberg leader and player 2 is the follower.
- This means player 1's problem is

$$\begin{aligned} \text{maximize}_{q_1 \geq 0} \quad & \pi_1(q_1, q_2) = [P(Q) - c]\, q_1 \\ \text{subject to} \quad & q_2 = \frac{\alpha - c - \beta q_1}{2\beta}. \end{aligned}$$

Or

$$\max_{q_1 \geq 0} \left[\alpha - \beta \left(q_1 + \frac{\alpha - c - \beta q_1}{2\beta} \right) - c \right] q_1.$$

Stackleberg Competition (continued)

- The first-order condition is

$$\left[\alpha - \beta\left(q_1 + \frac{\alpha - c - \beta q_1}{2\beta}\right) - c\right] - \frac{\beta}{2} q_1 = 0,$$

  which gives

$$q_1^S = \frac{\alpha - c}{2\beta}.$$

- And thus

$$q_2^S = \frac{\alpha - c}{4\beta} < q_1^S$$

- Why lower output for the follower?
- Total output is

$$Q^S = q_1^S + q_2^S = \frac{3\left(\alpha - c\right)}{4\beta},$$

  which is greater than Cournot output. Why?

Multiplicity of Subgame Perfect Equilibria

- Question: What happens if in some subgame more than one action is optimal? Consider all optimal actions and trace back implications of each in all of the longer subgames.
  To illustrate this, consider the following game.

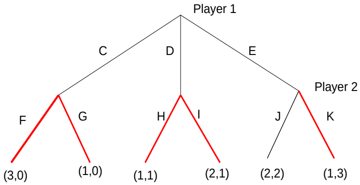

- Player 2's optimal strategies in this game are given by:

FHK FIK
GHK GIK

Multiplicity of Subgame Perfect Equilibria

- Now consider player 1's optimal strategies for every combination of optimal actions for player 2:

2's opt. strt.		1's BR	6 SPE's
			(C, FHK)
FHK	$\longrightarrow$	C	(C, FIK)
FIK	$\longrightarrow$	C	(C, GHK)
GHK	$\longrightarrow$	C,D,E	(D, GHK)
GIK	$\longrightarrow$	D	(E, GHK)
			(D, GIK)

Bargaining Problems

- We next study bargaining problems, which can be naturally modeled as an extensive game.
- We start by studying the ultimatum game, which is a simple game that is the basis of a richer model.
- Two people use the following procedure to split c dollars:
  - 1 offers 2 some amount $x \leq c$
  - if 2 accepts the outcome is: $(c - x, x)$
  - if 2 rejects the outcome is: $(0, 0)$
  - Note: each person cares about the amount of money he receives and we assume that x can be any scalar, not necessarily integral.
- Question: What is an SPE for this game?
- Let us use an extensive game model for the negotiation process:

SPE of the Ultimatum Game

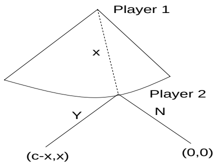

- It is a finite horizon game, so we can use backward induction to find the SPE of this game.
- There is a different possible subgame for each value of x, so we need to find the optimal action of player 2 for each such subgame:
  - if $x > 0 \longrightarrow$ Yes
  - $x = 0 \longrightarrow$ indifferent between Yes and No
- How many different optimal strategies does player 2 have?

  (1) Yes for all $x \geq 0$
  (2) Yes if $x > 0$ and No if $x = 0$

SPE of the Ultimatum Game

- Trace back the implications of each of player 2's optimal strategies, i.e., consider player 1's optimal strategy for each of these strategies:
  - For (1): player 1's optimal offer is $x = 0$
  - For (2): player 1's optimal offer is:

    $x = 0 \rightarrow 0$

    $x > 0 \rightarrow c - x \quad \max_{x>0}(c - x) \Rightarrow$ no optimal solution

    $\Rightarrow$ no offer of player 1 is optimal!
  - Unique SPE:
    - 1 offers 0
    - person 2 accepts all offers
    - Outcome: $(0, y) \Rightarrow$ 1 gets all the pie

Remarks:

- One-sided outcome $\longrightarrow$ one-sided structure of the game should allow 2 to make a counter offer after rejection then it is more like bargaining
- This SPE is not supported by experimental evidence (cultural effect come into play and behavior exhibits some sort of concern for fair outcomes or reciprocity).

SPE of the Ultimatum Game

Exercises:

1. What if the amount of money available is in multiples of a cent? Then there are 2 SPE's instead of 1:
    - Player 1 offers 0, and player 2 says Yes to all offers
    - Player 1 offers 1 cent, and player 2 says Yes to all offers except 0
2. Show that for every $\bar{x} \in [0, c]$, there exists a NE in which 1 offers $\bar{x}$. Find 2's optimal strategy.

Bargaining as an Extensive Game

- In the ultimatum game, player 2 is powerless. His only alternative to accepting is to reject which results in him getting no pie.
- Let us extend the model to give player 2 more power:
- We assume that $c = 1$. Moreover, let $x = (x_1, x_2)$ with $x_1 + x_2 = 1$ denote the allocation in the first part and $y = (y_1, y_2)$ with $y_1 + y_2 = 1$ denote the allocation in the second part.

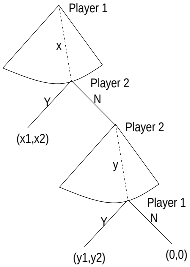

Bargaining as an Extensive Game

- The second part of the game is an ultimatum game in which 2 moves first. This has a unique subgame perfect equilibrium given by:
  - 2 offers nothing to 1
  - 1 accepts all offers
- We note that in every SPE, 2 obtains all the pie
- Last Mover's Advantage: Similar result with alternating offers. In every SPE, the player who makes the offer in the last period obtains all the pie.
- In our model so far, players indifferent about timing of an agreement. In real life however, bargaining takes time and time is valuable. Players preferences should reflect the fact that they have bias towards earlier agreements.

Finite Horizon Game with Alternating Offers

- Players alternate proposals, future discounted using the constant discount factor $0 < \delta_i < 1$ at each period.

Two Periods:

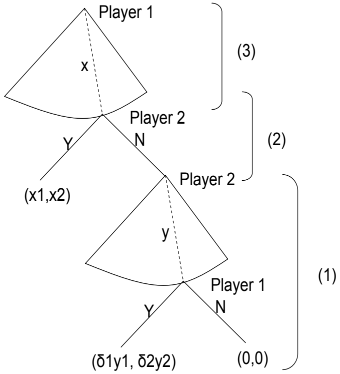

We find the SPE by backward induction:

Finite Horizon Game with Alternating Offers

1. In (1) (the last "ultimatum game"), the unique SPE: player 2 offers $(0, 1)$ player 1 accepts all proposals. The outcome is $(0, \delta_2)$.
2. In (2):

   $N \rightarrow (0, \delta_2)$
   $Y \rightarrow (x_1, x_2)$

   2 strategies

   - Y if $x_2 \geq \delta_2$
     N if $x_2 < \delta_2$
   - Y if $x_2 > \delta_2$
     N if $x_2 \leq \delta_2$
3. In (3): Players 1's optimal strategy is $(1 - \delta_2, \delta_2)$

Hence, the unique SPE of this game is:

- Player 1's initial proposal $(1 - \delta_2, \delta_2)$.
- Player 2 accepts all proposals where $x_2 \geq \delta_2$ and rejects all $x_2 < \delta_2$.
- Player 2 proposes $(0, 1)$ after any history in which he rejects a proposal of player 1.
- Player 1 accepts all proposals of player 2 (after a history in which 2 rejects 1's opening proposal).

Finite Horizon Game with Alternating Offers

The outcome of the game is:

- Player 1 proposes $(1 - \delta_2, \delta_2)$.
- Player 2 accepts.
- Resulting payoff : $(1 - \delta_2, \delta_2)$.

Desirability of an earlier agreement yields a positive payoff for player 1.

- For 3 periods, similar analysis using backward induction.
- Iterating, we get Stahl's Bargaining Model.

Stahl's Bargaining Model: Finite Horizon

# *Periods*	player 1 gets
2 periods	$1-\delta_2$
3 periods	$1-\delta_2+\delta_1\delta_2$
5 periods	$1-\delta_2+\delta_1\delta_2(1-\delta_2)+\delta_1\delta_2$
$2k$ periods	$1-\delta_2\left[\frac{1-(\delta_1\delta_2)^k}{1-(\delta_1\delta_2)}\right]$
$2k+1$ periods	$1-\delta_2\left[\frac{1-(\delta_1\delta_2)^k}{1-(\delta_1\delta_2)}\right]+(\delta_1\delta_2)^k$

Taking the limit as $k \rightarrow \infty$, we see that player 1 gets $x_1^* = \frac{1-\delta_2}{1-\delta_1\delta_2}$ at the SPE.

Rubinstein's Infinite Horizon Bargaining Model

- Instead of two players alternating offers for a period of time, there is no deadline, they can alternate offers forever.
- There are two types of terminal histories:

$$(x^1, N, x^2, N.....x^t, N...) \quad \rightarrow \quad \text{every offer rejected}$$
$$(x^1, N, x^2, N.....x^t, Y)$$

- This game does not have a finite horizon, so we cannot use backward induction to find the SPE.
- We will instead guess a strategy profile and verify that it forms an SPE using the one-stage deviation principle.
- The strategy of a player in this game involves:
  1. Offer in period 1
  2. Response to history (x^1, N, x^2)
  3. counteroffer for history (x^1, N, x^2, N)
- Each player faces the same subgame in all periods: The absolute payoffs are different but the preferences are the same, because all options are discounted by the same factor

Rubinstein's Infinite Horizon Bargaining Model

- Therefore, we focus on stationary policies in which each player always make the same proposal and always accepts the same set of proposals.
- We define

$$x_1^* = \frac{1-\delta_2}{1-\delta_1\delta_2}, \qquad x_2^* = \frac{\delta_2(1-\delta_1)}{1-\delta_1\delta_2},$$
$$y_1^* = \frac{1-\delta_1}{1-\delta_1\delta_2}, \qquad y_2^* = \frac{\delta_1(1-\delta_2)}{1-\delta_1\delta_2}.$$

- We consider the following strategy profile (s_1^*, s_2^*):
  - player 1 proposes x^* and accepts y if and only if $y_1 \geq y_1^*$.
  - player 2 proposes y^* and accepts x if and only if $x_2 \geq x_2^*$.
- We verify that this strategy profile is an SPE using one-stage deviation principle.

Rubinstein's Infinite Horizon Bargaining Model

- First note that this game has 2 types of subgames:
  1. One in which first move is an offer:
  2. One in which first move is a response to an offer:
- For the first type of subgame: Suppose offer made by player 1
  - Fix 2's strategy at s_2^*
  - if player 1 adopts $s_1^* \Rightarrow 2$ accepts, player 1 gets x_1^*
  - if 1 offers $> x_2^*$, 2 accepts leading to a lower payoff than x_1^* for player 1.
  - if 1 offers $< x_2^*$, 2 rejects, offers y^*, player 1 accepts, leading to a payoff of $\delta_1 y_1^*$. Since $\delta_1 y_1^* < x_1^*$, player 1 is better off using s_1^*.
- For the second type of subgame: Suppose player 1 is responding
  - Fix 2's strategy at s_2^*
  - Denote by (y_1, y_2) the offer to which player 1 is responding
  - if player 1 adopts s_1^*, he accepts the offer iff $y_1 > y_1^*$
  - if player 1 rejects some offer $y_1 \geq y_1^*$, player 1 will get $\delta_1 x_1^* = y_1^*$, thus he cannot increase his payoff by deviating.

Hence s^* is an SPE (in fact the *unique SPE*, check FT, Section 4.4.2 to verify).