6.254 : Game Theory with Engineering Applications
Lecture 12: Extensive Form Games

Asu Ozdaglar
MIT

March 16, 2010

Outline

- Extensive Form Games with Perfect Information
- Backward Induction and Subgame Perfect Nash Equilibrium
- One-stage Deviation Principle
- Applications

- Reading:
- Fudenberg and Tirole, Chapter 3 (skim through Sections 3.4 and 3.6), and Sections 4.1-4.2.

Extensive Form Games

- We have studied strategic form games which are used to model one-shot games in which each player chooses his action once and for all simultaneously.
- In this lecture, we will study extensive form games which model multi-agent sequential decision making.
- Our focus will be on multi-stage games with observed actions where:
  - All previous actions are observed, i.e., each player is perfectly informed of all previous events.
  - Some players may move simultaneously at some stage k.
- Extensive form games can be conveniently represented by **game trees**.
- Additional component of the model, histories (i.e., sequences of action profiles).

Example 1 – Entry Deterrence Game:

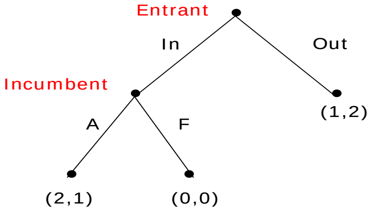

- There are two players.
- Player 1, the entrant, can choose to enter the market or stay out. Player 2, the incumbent, after observing the action of the entrant, chooses to accommodate him or fight with him.
- The payoffs for each of the action profiles (or histories) are given by the pair (x, y) at the leaves of the game tree: x denotes the payoff of player 1 (the entrant) and y denotes the payoff of player 2 (the incumbent).

Example 2 – Investment in Duopoly

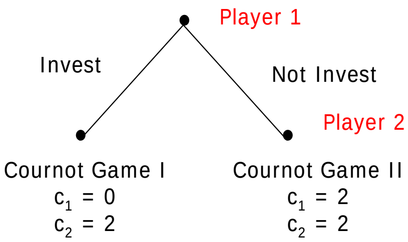

- There are two players in the market.
- Player 1 can choose to invest or not invest. After player 1 chooses his action, both players engage in a Cournot competition.
- If player 1 invests, then they will engage in a Cournot game with $c_1 = 0$ and $c_2 = 2$.
- Otherwise, they will engage in a Cournot game with $c_1 = c_2 = 2$.
- We can also assume that there is a fixed cost of f for player 1 to invest.

Extensive Form Game Model

- A set of players, $\mathcal{I} = \{1, \ldots, I\}$.
- Histories: A set H of sequences which can be finite or infinite.

$h^0 = \emptyset$	initial history
$a^0 = (a_1^0, \ldots, a_I^0)$	stage 0 action profile
$h^1 = a^0$	history after stage 0
$\vdots$	$\vdots$
$h^{k+1} = (a^0, a^1, \ldots, a^k)$	history after stage k

  - If the game has a finite number $(K + 1)$ of stages, then it is a finite horizon game.
  - Let $H^k = \{h^k\}$ be the set of all possible stage k histories.
  - Then H^{K+1} is the set of all possible *terminal histories*, and $H = \cup_{k=0}^{K+1} H^k$ is the set of all possible histories.

Extensive Form Game Model

- Pure strategies for player i is defined as a contingency plan for every possible history h^k.
  - Let $S_i(H^k) = \bigcup_{h^k \in H^k} S_i(h^k)$ be the set of actions available to player i at stage k.
  - Let $s_i^k : H^k \to S_i(H^k)$ such that $s_i(h^k) \in S_i(h^k)$.
  - Then the pure strategy of player i is the set of sequences $s_i = \{s_i^k\}_{k=0}^K$, i.e., a pure strategy of a player is a collection of maps from all possible histories into available actions.
  - Observe that the path of strategy profile s will be $a^0 = s^0(h^0)$, $a^1 = s^1(a^0)$, $a^2 = s^2(a^0, a^1)$, and so on.
- Preferences are defined on the outcome of the game H^{K+1}. We can represent the preferences of player i by a utility function $u_i : H^{K+1} \to \mathbb{R}$. As the strategy profile s determines the path $a^0, \ldots, a^k$ and hence h^{K+1}, we will denote $u_i(s)$ as the payoff to player i under strategy profile s.

Strategies in Extensive Form Games

Example:

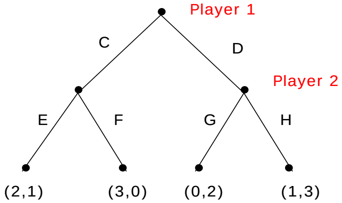

- Player 1's strategies: $s_1 : H^0 = \emptyset \rightarrow S_1 = \{C, D\}$; two possible strategies: C,D
- Player 2's strategies: $s^2 : H^1 = \{\{C\}, \{D\}\} \rightarrow S_2$; four possible strategies: which we can represent as EG, EH, FG and FH.
- If $s = (C, EG)$, then the outcome will be $\{C, E\}$. On the other hand, if the strategy is $s = (D, EG)$, the outcome will be $\{D, G\}$.

Strategies in Extensive Form Games (continued)

- Consider the following two-stage extensive form version of matching pennies.

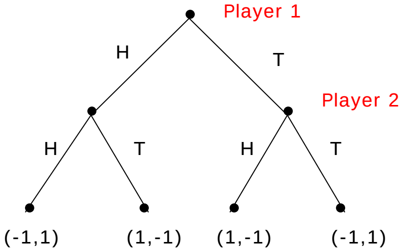

- How many strategies does player 2 have?

Strategies in Extensive Form Games (continued)

- Recall: strategy should be a *complete contingency plan.*
- Therefore: player 2 has four strategies:
  1. heads following heads, heads following tails HH;
  2. heads following heads, tails following tails HT;
  3. tails following heads, tails following tails TT;
  4. tails following heads, heads following tails TH.

Strategies in Extensive Form Games (continued)

- Therefore, from the extensive form game we can go to the strategic form representation. For example:

Player 1/Player 2	HH	HT	TT	TH
heads	$(-1, 1)$	$(-1, 1)$	$(1, -1)$	$(1, -1)$
tails	$(1, -1)$	$(-1, 1)$	$(-1, 1)$	$(1, -1)$

- So what will happen in this game?

Strategies in Extensive Form Games (continued)

- Can we go from strategic form representation to an extensive form representation as well?
- To do this, we need to introduce information sets:
  - Information sets model the information players have when they are choosing their actions.
  - They can be viewed as a generalization of the idea of a history.
  - The information sets, $h \in H$, partition the nodes of the game tree: the interpretation of the information set $h(x)$ containing node x is that the player who is choosing an action at x is uncertain if he is at x or at some other $x' \in h(x)$.
  - We require that if $x' \in h(x)$, then the same player moves at x and x' and also that $A(x) = A(x')$.
- A game has perfect information if all its information sets are singletons (i.e., all nodes are in their own information set).

Strategies in Extensive Form Games (continued)

- The following two extensive form games are representations of the simultaneous-move matching pennies.
- The loops represent the information sets of the players who move at that stage. These are imperfect information games.
- These games represent exactly the same strategic situation: each player chooses his action not knowing the choice of his opponent.
- Note: For consistency, first number is still player 1's payoff.

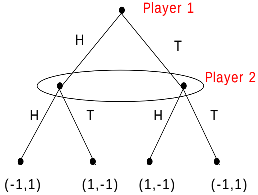

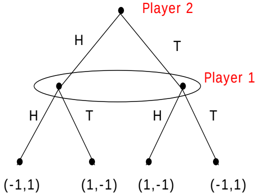

Entry Deterrence Game

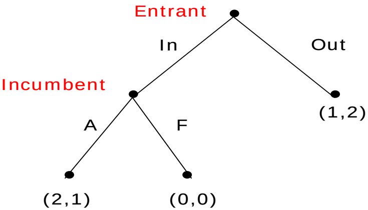

- Equivalent strategic form representation:

Entrant\Incumbent	Accommodate	Fight
In	$(2,1)$	$(0,0)$
Out	$(1,2)$	$(1,2)$

- Two pure Nash equilibria: (In,A) and (Out,F).

Are These Equilibria Reasonable?

- The equilibrium (Out,F) is sustained by a **noncredible threat** of the monopolist.
- 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*.
- For infinite horizon games, characterization in terms of **one-stage deviation principle.**

Subgames

- To define subgame perfection formally, we first define the idea of a subgame.
  - Informally, a subgame is a portion of a game that can be analyzed as a game in its own right.
- Recall that a game G is represented by a game tree. Denote the set of nodes of G by V_G.
- Recall that history h^k denotes the play of a game after k stages. In a perfect information game, each node $x \in V_G$ corresponds to a unique history h^k and vice versa. This is not necessarily the case in imperfect information games.

Subgames (continued)

Definition (Subgames)

A subgame G' of an extensive form game G consists of a single node and all its successors in G, with the property that if $x' in V_{G'}$ and $x'' \in h(x')$, then $x'' \in V_{G'}$. The information sets and payoffs of the subgame are inherited from the original game.

- The definition requires that all successors of a node is in the subgame and that the subgame does not "chop up" any information set.
  - This ensures that a subgame can be analyzed in its own right.
  - This implies that a subgame starts with a node x with a singleton information set, i.e., $h(x) = x$.
- In **perfect information games**, subgames coincide with nodes or stage k histories h^k of the game. In this case, we use the notation h^k or $G(h^k)$ to denote the subgame.
- A restriction of a strategy s to subgame G', $s_{|G'}$ is the action profile implied by s in the subgame G'.

Subgames: Examples

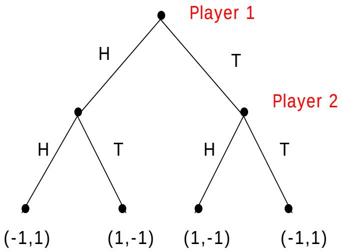

- Recall the two-stage extensive-form version of the matching pennies game
- In this game, there are two proper subgames and the game itself which is also a subgame, and thus a total of three subgames.

Subgame Perfect Equilibrium

Definition

(Subgame Perfect Equilibrium) *A strategy profile* s^* *is a Subgame Perfect Nash equilibrium (SPE) in game* G *if for any subgame* G' *of* G, $s^*_{|G'}$ *is a Nash equilibrium of* G'.

- Loosely speaking, subgame perfection will remove noncredible threats, since these will not be Nash equilibria in the appropriate subgames.
- In the entry deterrence game, following entry, F is not a best response, and thus not a Nash equilibrium of the corresponding subgame. Therefore, (Out,F) is not a SPE.
- How to find SPE? One could find all of the Nash equilibria, for example as in the entry deterrence game, then eliminate those that are not subgame perfect.
- But there are more economical ways of doing it.

Backward Induction

- Backward induction refers to starting from the last subgames of a finite game, then 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.

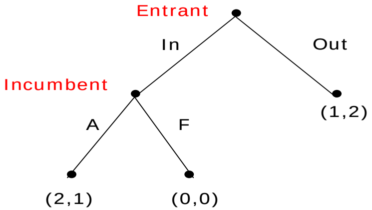

Backward Induction (continued)

Theorem

Backward induction gives the entire set of SPE.

Proof: backward induction makes sure that in the restriction of the strategy profile in question to any subgame is a Nash equilibrium.

- Backward induction is straightforward for games with perfect information and finite horizon.
- For imperfect information games, backward induction proceeds similarly: we identify the subgames starting from the leaves of the game tree and replace it with one of the Nash equilibrium payoffs in the subgame.
- For infinite horizon games: we will rely on a useful characterization of the subgame perfect equilibria given by the "one stage deviation principle."

Existence of Subgame Perfect Equilibria

Theorem

Every finite perfect information extensive form game G has a pure strategy SPE.

Proof: Start from the end by backward induction and at each step one strategy is best response.

Theorem

Every finite extensive form game G has a SPE.

Proof: Same argument as the previous theorem, except that some subgames need not have perfect information and may have mixed strategy equilibria.

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., $h = (a^0, a^1, a^2 ...)$, is an infinite sequence of actions. Let $h^t = (a^0, ...a^{t-1})$ 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*

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

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.

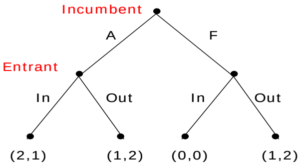

- Note: For consistency, first number is still the entrant's payoff.
- This implies that the incumbent can now commit to fighting (how could it do that?).
- It is straightforward to see that the unique SPE now involves the incumbent committing to fighting and the entrant not entering.
- This illustrates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

6.254 Game Theory with Engineering Applications
Spring 2010