

Introduction to Game Theory

Lecture Note 4: Extensive-Form Games and Subgame Perfect Equilibrium


Haifeng Huang

University of California, Merced

Spring 2020

Extensive games with perfect information

- What we have studied so far are strategic-form games, where players simultaneously choose an action (or a mixed strategy) once and for all. Now we study extensive-form games (extensive games; dynamic games), where players move sequentially.
- An example: A challenger decides whether or not to enter (an electoral race); if the challenger enters, the incumbent decides to fight or acquiesce.



Concepts


- Some concepts:
 - ▷ The **empty history** (\emptyset): the start of the game.
 - ▷ A **terminal history**: a sequence of actions that specifies what may happen in the game from the start of the game to an action that ends the game.
 - ▷ A **subhistory**, or **history** (h), of a finite sequence of actions (a^1, a^2, \dots, a^k) refers to a sequence of the form (a^1, a^2, \dots, a^m) , where $1 \leq m \leq k$, or the empty history.
 - ▷ A **proper subhistory** is a subhistory that is not equal to the entire sequence (i.e., $m < k$).
- So in the entry game,
 - ▷ Terminal histories: $(in, acquiesce)$, $(in, fight)$, and out .
 - ▷ Subhistories of the sequence $(in, acquiesce)$: \emptyset , $(in, acquiesce)$, and in .

Definition

- **Perfect information:** each player is perfectly informed of the *history* of what has happened so far, up to the point where it is her turn to move.
- An **extensive game with perfect information** consists of
 - ① A set of players
 - ② A set of terminal histories
 - ③ A player function that assigns a player to every sequence that is a proper subhistory of some terminal history (i.e., to every node in each terminal history)
 - ⇒ Player function in the entry game: $P(\emptyset) = \text{Challenger}$ and $P(in) = \text{Incumbent}$.
 - ④ For each player, preferences over the set of terminal histories


Strategy

- A **strategy** of player i in an extensive game with perfect information specifies what action i takes for each history after which it is her turn to move; i.e., it is a **plan of action** for *all contingencies*.
- In the following game, player 1's strategies are C and D , while player 2's strategies are EG , EH , FG , and FH .



Strategy

- In the following game, player 2's strategies are E and F , while 1's strategies are CG , CH , DG , and DH .




- Even if player 1's plan is to choose D after the start of the game, a strategy needs to specify what she will do after history CE !

Nash equilibrium in extensive games

- Let s^* denote a strategy profile, and $O(s^*)$ denote a terminal history generated by s^* . The strategy profile s^* in an extensive game with perfect information is a Nash equilibrium if, for every player i and every strategy r_i of player i , $O(s^*)$ is at least as good for i as the terminal history $O(r_i, s_{-i}^*)$. In other words, $u_i(O(s^*)) \geq u_i(O(r_i, s_{-i}^*)), \forall r_i$.
- How to obtain NEs of an extensive game? **Convert the extensive game into a strategic form game**, by renaming the strategies in the extensive form as actions in the strategic form and making the payoffs to a terminal history generated by a strategy profile as the payoffs to a action profile.

Nash equilibrium in extensive games: example

The following extensive game



is converted as below with NE: (CH, F), (DG, E), (DH, E).

	E	F
CG	1, 2	3, 1
CH	0, 0	3, 1
DG	2, 0	2, 0
DH	2, 0	2, 0

NE not good enough for extensive games


- There is something unsatisfactory about the Nash equilibrium concept in extensive games. It ignores the sequential structure of the game and treats strategies as choices made once and for all.
 - ▷ For example, in the previous game (CH, F) is a NE. But when it's player 1's turn to move after history E, will she really choose H over G?
 - ▷ Take a simpler example. In the entry game, what are the NEs? (In, A) and (Out, F) . But is the incumbent's threat to fight credible?
- We need a refinement of the Nash equilibrium for extensive games, which is called subgame perfect equilibrium.

Subgame

- Let Γ be an extensive game with perfect information, with player function P . For any nonterminal history h of Γ , the **subgame** $\Gamma(h)$ following the history h is the extensive game that starts after history h .
- The subgame following the empty history \emptyset is the entire game itself. Every other subgame of an extensive game is called a **proper subgame**.

Subgame: example

The following game has 3 subgames: the whole game, the game following history C, and the game following history D.




Subgame perfect Nash equilibrium (SPNE)

A **subgame perfect Nash equilibrium** is a strategy profile s^* with the property that in no subgame following history h can any player i do better by choosing a strategy different from s_i^* , given that every other player j adheres to s_j^* . I.e., $u_i(O_h(s^*)) \geq u_i(O_h(r_i, s_{-i}^*)), \forall r_i, i$.


Subgame perfect Nash equilibrium (SPNE): example

In the following extensive game, (CH, F) and (DH, E) are NE but not SPNE; (DG, E) is SPNE.



Relationship between NE and SPNE

- Every SPNE is a NE; the reverse is not true.
- A strategy profile that induces a NE in every subgame is a SPNE.
- In the entry game, how many subgames are there? Identify the NE(s) and the SPNE(s).




Finding SPNE in finite horizon games: backward induction

- An extensive game is **finite** if it has finite horizon (i.e., the length of the longest terminal history is finite), *and* the number of terminal histories is finite.
- **Length** of a subgame: the length of the longest history in the subgame
- **Backward induction:**
 - ① Find the optimal actions of the players who move in the subgames of length 1
 - ② Taken the actions found in step 1 as given, find the optimal actions of the players who move first in the subgames of length 2
 - ③ Continue this procedure, until you reach the start of the game


Backward induction: example

Solve for the SPNE in the following game by backward induction




Backward induction: another example

- What about the following game?



Backward induction: another example


- What about the following game?



- (C,EG), (D,EG), (C,EH), (D,FG), (C,FH), and (D,FH)

Practice with backward induction: roll call voting

In the US Senate, roll calls are held in alphabetical order. Suppose Party D has (only) three senators, each of whom prefers a unpopular bill to pass. However, conditional on the outcome of the vote, each also prefers to be seen as independent from the party line and not to be on the record supporting the bill. What will be the equilibrium of the following roll call (Spenkuch, Montagnes, and Magleby 2018)? Two “yea” votes are needed for passage.




Some general results

- Proposition: *The set of SPNE of a finite horizon extensive game with perfect information is equal to the set of strategy profiles isolated by the procedure of backward induction.*
- **Zermelo's Theorem:** Every finite extensive game with perfect information has a SPNE. Moreover, if no player has the same payoffs at any two terminal nodes, then there is a unique SPNE that can be derived from backward induction.
- In other words, the game Go actually has a solution! It's just we have not found it, even in the age of super computers.

The ultimatum game

Two people need to decide how to divide a dollar. Person 1 offers person 2 an amount x , and keeps the rest $(1 - x)$. If person 2 accepts the offer, then they respectively receive $1 - x$ and x . If 2 rejects the offer, then neither person receives any of the dollar. Each person cares only about the money she receives. What's the SPNE?




Solving the ultimatum game

- Person 2's optimal choice: accept if $x > 0$; accept or reject if $x = 0$.
- If person 2 is to accept all offers (including $x = 0$), then person 1's best strategy is to offer $x = 0$; if person 2 is to reject the offer $x = 0$, then person 1 has no best strategy, b/c there is no minimum x that is greater than 0.
- Therefore the unique SPNE is person 1's strategy is to offer 0, and person 2's strategy is to accept whatever that is offered.
- If person 2 is allowed to make a counter offer, the game becomes a bargaining and there will be a different solution.
- Experiments on this game has different results across different cultures.

Agenda control: an application of the ultimatum game

- There is a status quo policy y_0 . A committee proposes to change the policy, and afterwards the legislature decides to accept the committee's proposal or to reject it. If the proposal is rejected, the status quo policy is maintained. The legislature's ideal policy is 0, while the committee's ideal policy is $y_c > 0$. Both players have single peaked preferences, i.e., the further the final policy is from their ideal points, the worse their payoff. What policy will the committee propose in a SPNE?



- Hint: the committee's proposal depends on the location of y_0 .

The holdup game

- Person 1 decides to exert high effort ($E = H$) or low effort ($E = L$), $H > L$, in producing a pie. H leads to a big pie with size C_H , and L leads to a small pie with size C_L . Person 2 then chooses to offer person 1 with X percent of the pie and keep the rest for himself. Person 1 can accept or reject the offer. If she rejects, neither gets any of the pie. Person 2's payoff is what he receives. Person 1's payoff is the share of the pie she gets minus the cost of her effort. What's the SPNE?

The holdup game

- Person 1 decides to exert high effort ($E = H$) or low effort ($E = L$), $H > L$, in producing a pie. H leads to a big pie with size C_H , and L leads to a small pie with size C_L . Person 2 then chooses to offer person 1 with X percent of the pie and keep the rest for himself. Person 1 can accept or reject the offer. If she rejects, neither gets any of the pie. Person 2's payoff is what he receives. Person 1's payoff is the share of the pie she gets minus the cost of her effort. What's the SPNE?
- After person 1's first move, the game is the same as the ultimatum game, and that subgame has a unique SPNE, in which person 2 offers person 1 zero. Foreseeing this, person 1 chooses low effort.
- Concerns about one party's bargaining power leads to failure in cooperation in producing a larger pie.

Rotten kid theorem

- A child takes an action a that affects both his income $c(a)$ and his parent's income $p(a)$. For all actions of the child $p(a) > c(a)$. The child cares only about his own income while the parent cares about both her income and the child's income. Specifically, $u_c = c(a)$, and $u_p = \min\{c(a), p(a)\}$. After the child's action, the parent can transfer some of her money to the child.
- What kind of action will the child take? Does he take an action to maximize his income, his parent's income, or the total family income?

Rotten kid theorem: solution

- The parent's transfer t will be $\frac{p(a)-c(a)}{2}$. And so the child's final income will be $\frac{p(a)+c(a)}{2}$. That is, a child will take an action to maximize the total family income, even though he only cares about himself.
- I.e., if a family head is caring enough about all other family members (benevolent dictator), then other family members are motivated to maximize family income, even if their own welfare depends solely on their own individual income. A well-known result about family economics and perhaps authoritarianism, due to Becker (1974).

Vote buying among pirates

Five pirates (P_1 , P_2 , P_3 , P_4 and P_5) must split a treasure consisting of 100 (indivisible) gold pieces. They use the following procedure: first, pirate P_1 proposes a division of the 100 gold pieces, then all pirates vote for or against; if it passes, the proposed division is implemented, if not, pirate P_1 is beheaded. Afterwards pirate P_2 proposes, and the same process is repeated with P_3 , P_4 and P_5 until either agreement is reached, or there is only one pirate left. In order for a proposal to pass, it must meet majority approval. Pirates cannot abstain, and if they are indifferent they vote against, since a beheading is always fun. Ties are broken in favor of the proposal. How many gold pieces does each pirate end up with? Does any pirate get beheaded?

Another model of vote buying/interest group lobbying

Two interest groups compete in bribing k legislators, k odd, to vote for their respective favorite bills. Interest group X values bill x at $V_x > 0$, and values bill y at zero. Interest group Y values bill y at $V_y > 0$, and values bill x at zero. X moves first and decides how much to pay to each legislator, and Y moves second. A legislator will vote for the bill favored by the group that offers her more payment. If the two payments are equal, she votes for y (a simplifying tie-breaking rule). Each interest group wants to win and minimize the payment. What is the SPNE?

A model of vote buying/interest group lobbying

- Hint: Since k is odd, let $\mu = \frac{k+1}{2}$ be the number that constitutes bare majority of the legislature. Y moves second, so to get bill y passed it only needs to pay the same amount as what X has paid the μ legislators that have received the lowest payments.
- This is the second mover's advantage.

The vote buying model: analysis (1)

- Using backward induction, we first analyze Y 's optimal strategy. Let m_x be the total amount of money that interest group X has paid to the μ legislators that have received the lowest payments from X .
- If $m_x > V_y$, then it is not worthwhile for Y to buy off the bare majority (or any majority). Y 's optimal strategy is then to pay zero to every legislator.
- If $m_x < V_y$, then Y should match whatever X has paid to the μ members of the “cheapest” bare majority, and have bill y passed.
- If $m_x = V_y$, the two above strategies are both best responses to X 's strategies.

The vote buying model: analysis (2)

- Foreseeing what Y will do, X should come up with a way to deter Y , or give up if Y cannot be deterred.
- Y is willing to pay each of μ legislator an amount up to $\frac{V_y}{\mu}$ to get y passed. So to deter Y , X needs to pay at least $\frac{V_y}{\mu}$ to all k legislators. Why?
- So if V_x is less than $\frac{kV_y}{\mu}$, the fight is not worthwhile for X , and X should pay each legislator zero; i.e., its strategy is $(x_1, x_2, \dots, x_k) = (0, 0, \dots, 0)$.
- If V_x is greater than $\frac{kV_y}{\mu}$, X should pay each of the k legislators an amount $\geq \frac{V_y}{\mu}$. There is no smallest number that is $> \frac{V_y}{\mu}$, so no SPNE in which X pays every legislator an amount $> \frac{V_y}{\mu}$. There is a SPNE in which X pays each legislator $\frac{V_y}{\mu}$, and Y pays zero (Y is indifferent).

The vote buying model: conclusion

- When $V_x \neq \frac{kV_y}{\mu}$, there is a unique SPNE, in which Y 's strategy is to
 - ▷ match X 's payment to the "cheapest" μ legislators after a history in which $m_x < V_y$, and
 - ▷ make no payment to any legislator after a history in which $m_x \geq V_y$.and X 's strategy is to make no payment if $V_x < \frac{kV_y}{\mu}$, and pay each legislator $\frac{V_y}{\mu}$ if $V_x > \frac{kV_y}{\mu}$.
- If $V_x = \frac{kV_y}{\mu}$, there is an equilibrium in which X pays zero (and Y pays zero), and an equilibrium in which X pays each legislator $\frac{V_y}{\mu}$, and Y pays zero (If Y will match X , then X should pay zero).
- In equilibrium, Y always pays zero!

Oligopolistic competition: the Stackelberg model

- Two firms produce the same product. The unit cost of production is c . Let q_i be firm i 's output, $Q = \sum_{i=1}^2 q_i$, then the market price P is $P(Q) = \alpha - Q$, where α is a constant.
- In the Cournot model the two firms move simultaneously. Here firm 1 moves first, and firm 2 moves next.
- Using backward induction, we first analyze firm 2's decision:

$$q_2 = b_2(q_1) = \frac{\alpha - q_1 - c}{2} \quad (1)$$

- Then firm 1 takes firm 2's best response as given, and its optimal strategy is to maximize

$$\begin{aligned} \pi_1 &= (\alpha - (q_1 + q_2))q_1 - cq_1 \\ &= \left(\alpha - \left(q_1 + \frac{\alpha - q_1 - c}{2}\right)\right)q_1 - cq_1 \end{aligned} \quad (2)$$

The Stackelberg model (2)

- Differentiating equation (2) with regard to q_1 and using the first order condition, we know

$$q_1^* = \frac{1}{2}(\alpha - c).$$


And then from equation (1) we get

$$q_2^* = \frac{1}{4}(\alpha - c).$$

- Firm 1 increases its output (in fact, also profit) than in the Cournot model, while firm 2 decreases its output.

Extension 1: allowing for simultaneous moves in extensive games


- A combination of extensive games with perfect information and normal form games.
- Difference from extensive games w/o simultaneous moves: i is a member of player function $P(h)$, rather than $P(h) = i$
- Example: a variant of Battle of Sexes



- Player 1's strategies: (Ballet, A), (Ballet, B), (Soccer, A), (Soccer, B); 2's strategies: A, B.

Solution of the example


- The subgame following the history *soccer* has two Nash equilibria: (A, A) and (B, B).
- If the outcome of that subgame is (A, A), then player 1 will choose soccer in her initial move; if the subgame outcome is (B, B), she will choose ballet in her initial move.



- Hence, two SPNE: ((Soccer, A), A) and ((Ballet, B), B)

An exercise


- What are each player's strategies in the following game?
What are the NEs in the two subgames?



- What are the SPNEs?

An exercise


- What are each player's strategies in the following game?
What are the NEs in the two subgames?



- What are the SPNEs?
- (ACE, CE), (ACF, CF), (ADF, DF), (BDE, DE)

Illustration: Agenda manipulation

- Binary agenda: committee members simultaneously vote whether to choose or eliminate the first alternative; if eliminate, then move on to vote on the second alternative; and so on.




- Let's assume voters do not use weakly dominated strategies (sophisticated voting).

Agenda manipulation

- Assuming sophisticated voting, the Condorcet winner, if it exists, always wins regardless of the voting agenda.
- If a Condorcet winner does not exist, then the winner depends on the voting agenda, even with sophisticated voting.
- E.g., 3 members and 3 alternatives. A: $x \succ y \succ z$; B: $y \succ z \succ x$; C: $z \succ x \succ y$. Any alternative can win, given an appropriate voting agenda.
- Another example. 3 members, 5 alternatives. A: $x \succ y \succ v \succ w \succ z$; B: $z \succ x \succ v \succ w \succ y$; C: $y \succ z \succ w \succ v \succ x$.
- Design a binary voting agenda in which z is the result of sophisticated voting.

Agenda manipulation in committees (3)


- The following agenda produces z as the winner.



- Exchanging x and z produces x as the winner; exchanging y and z produces y as the winner.
- The agenda-setting power is very consequential!

Extension 2: Allowing for exogenous uncertainty

- Now we allow random exogenous events to occur during the course of the extensive game.
- In the game below, two players: 1 and 2; c stands for chance.



- Backward induction: 2 chooses C; if 1 chooses B, she'll get 3 with prob $\frac{1}{2}$, and 0 with prob $\frac{1}{2}$. So 1's expected payoff is 1.5 from B, better than 1 from A. So 1 chooses B.

An infinite horizon example: sequential duel (1)


- Two people alternatively decide whether or not to shoot the other person; each has infinite bullets so this is infinite horizon game. Each person i 's shots hit the target with prob p_i . Each cares only about her own chance of survival.
- Cannot use backward induction since the horizon is infinite.
- The strategy pair in which neither person ever shoots is SPNE.
 - ⇒ Each person survives. No outcome is better for either of them.
- The strategy pair in which both people always shoot is SPNE too. Why?

Sequential duel (2)

- Suppose person 2 shoots whenever it is her turn to move. Can player 1 benefit from deviating (at some point) from the strategy of always shooting?
- Denote person 1's probability of survival when she follows the strategy of always shooting as π_1 .
- Suppose person 1 deviates to not shooting (only) at the start of the game. This reduces her chance of survival, since it reduces her chance of eliminating person 2, who will always shoot her. More precisely, it reduces her survival probability from π_1 to $(1 - p_2)\pi_1$.
- Similarly, person 1 should not deviate after any history that ends with person 2 shooting and missing.
- The same logic holds for person 2.

Controversy on backward induction: the centipede game


- Two players move alternatively. Each can continue (C) or stop (S). The game ends after k periods (say, $k = 100$).



- What is the SPNE?

The centipede game

- The backward induction method would lead every player to stop whenever it's her turn to move. The outcome is that the first player stops immediately at the start of the game, yielding a payoff of $(2,0)$ to the two players.



- But the result is unappealing intuitively. If the players can continue for some steps, each of them can reap a much higher payoff. It's rational to be a little irrational!
- Key: *common knowledge* of the rationality assumption.