

Extensive Form Games: Backward Induction and Imperfect Information Games

Lecture 10

Lecture Overview

- 1 Recap
- 2 Centipede Game
- 3 Imperfect-Information Extensive-Form Games
- 4 Perfect Recall

Introduction

- The normal form game representation does not incorporate any notion of sequence, or time, of the actions of the players
- The **extensive form** is an alternative representation that makes the temporal structure explicit.
- Two variants:
 - **perfect information** extensive-form games
 - a “game tree” consisting of choice nodes and terminal nodes
 - choice nodes labeled with players, and each outgoing edge labeled with an action for that player
 - terminal nodes labeled with utilities
 - **imperfect-information** extensive-form games
 - we'll get to this today

Pure Strategies

- Overall, a pure strategy for a player in a perfect-information game is a complete specification of which deterministic action to take at every node belonging to that player.

Definition


Let $G = (N, A, H, Z, \chi, \rho, \sigma, u)$ be a perfect-information extensive-form game. Then the pure strategies of player i consist of the cross product

$$\prod_{h \in H, \rho(h)=i} \chi(h)$$

- Using this definition, we recover the old definitions of mixed strategies, best response, Nash equilibrium, ...

Induced Normal Form

- we can “convert” an extensive-form game into normal form



	<i>CE</i>	<i>CF</i>	<i>DE</i>	<i>DF</i>
<i>AG</i>	3, 8	3, 8	8, 3	8, 3
<i>AH</i>	3, 8	3, 8	8, 3	8, 3
<i>BG</i>	5, 5	2, 10	5, 5	2, 10
<i>BH</i>	5, 5	1, 0	5, 5	1, 0

Subgame Perfection

- Define **subgame of G rooted at h** :
 - the restriction of G to the descendants of H .
- Define **set of subgames of G** :
 - subgames of G rooted at nodes in G
- s is a **subgame perfect equilibrium** of G iff for any subgame G' of G , the restriction of s to G' is a Nash equilibrium of G'
- Notes:
 - since G is its own subgame, every SPE is a NE.
 - this definition rules out “non-credible threats”

Computing Subgame Perfect Equilibria

Idea: Identify the equilibria in the bottom-most trees, and adopt these as one moves up the tree

```

function BACKWARDINDUCTION (node  $h$ ) returns  $u(h)$ 
if  $h \in Z$  then
   $\lfloor$  return  $u(h)$ 
 $best\_util \leftarrow -\infty$ 
forall  $a \in \chi(h)$  do
   $util\_at\_child \leftarrow$  BACKWARDINDUCTION( $\sigma(h, a)$ )
  if  $util\_at\_child_{\rho(h)} > best\_util_{\rho(h)}$  then
     $\lfloor$   $best\_util \leftarrow util\_at\_child$ 
return  $best\_util$ 
  
```

- $util_at_child$ is a vector denoting the utility for each player
- the procedure doesn't return an equilibrium strategy, but rather labels each node with a vector of real numbers.
 - This labeling can be seen as an extension of the game's utility function to the non-terminal nodes
 - The equilibrium strategies: take the best action at each node.

Computing Subgame Perfect Equilibria

Idea: Identify the equilibria in the bottom-most trees, and adopt these as one moves up the tree

```


function BACKWARDINDUCTION (node  $h$ ) returns  $u(h)$ 
if  $h \in Z$  then
   $\perp$  return  $u(h)$ 
 $best\_util \leftarrow -\infty$ 
forall  $a \in \chi(h)$  do
   $util\_at\_child \leftarrow$  BACKWARDINDUCTION( $\sigma(h, a)$ )
  if  $util\_at\_child_{\rho(h)} > best\_util_{\rho(h)}$  then
     $\perp$   $best\_util \leftarrow util\_at\_child$ 
return  $best\_util$ 
  
```

- For zero-sum games, BACKWARDINDUCTION has another name: the **minimax** algorithm.
 - Here it's enough to store one number per node.
 - It's possible to speed things up by **pruning** nodes that will never be reached in play: "alpha-beta pruning".

Lecture Overview


- 1 Recap
- 2 Centipede Game
- 3 Imperfect-Information Extensive-Form Games
- 4 Perfect Recall

Centipede Game



- Play this as a fun game...

Backward Induction



- What happens when we use this procedure on Centipede?
 - In the only equilibrium, player 1 goes down in the first move.
 - However, this outcome is Pareto-dominated by all but one other outcome.
- Two considerations:
 - practical: human subjects don't go down right away
 - theoretical: what should you do as player 2 if player 1 doesn't go down?
 - SPE analysis says to go down. However, that same analysis says that P1 would already have gone down. How do you update your beliefs upon observation of a measure zero event?
 - but if player 1 knows that you'll do something else, it is rational for him not to go down anymore... a paradox
 - there's a whole literature on this question

Lecture Overview

- 1 Recap
- 2 Centipede Game
- 3 Imperfect-Information Extensive-Form Games**
- 4 Perfect Recall

Intro

- Up to this point, in our discussion of extensive-form games we have allowed players to specify the action that they would take at every choice node of the game.
- This implies that players know the node they are in and all the prior choices, including those of other agents.
- We may want to model agents needing to act with partial or no knowledge of the actions taken by others, or even themselves.
- This is possible using **imperfect information** extensive-form games.
 - each player's choice nodes are partitioned into **information sets**
 - if two choice nodes are in the same information set then the agent cannot distinguish between them.


Formal definition

Definition

An **imperfect-information game** (in extensive form) is a tuple $(N, A, H, Z, \chi, \rho, \sigma, u, I)$, where


- $(N, A, H, Z, \chi, \rho, \sigma, u)$ is a perfect-information extensive-form game, and
- $I = (I_1, \dots, I_n)$, where $I_i = (I_{i,1}, \dots, I_{i,k_i})$ is an equivalence relation on (that is, a partition of) $\{h \in H : \rho(h) = i\}$ with the property that $\chi(h) = \chi(h')$ and $\rho(h) = \rho(h')$ whenever there exists a j for which $h \in I_{i,j}$ and $h' \in I_{i,j}$.

Example



- What are the equivalence classes for each player?
- What are the pure strategies for each player?


Example



- What are the equivalence classes for each player?
- What are the pure strategies for each player?
 - choice of an action in each **equivalence class**.
- Formally, the pure strategies of player i consist of the cross product $\times_{I_{i,j} \in I_i} \chi(I_{i,j})$.

Normal-form games

- We can represent any normal form game.



- Note that it would also be the same if we put player 2 at the root node.


Induced Normal Form

- Same as before: enumerate pure strategies for all agents
- Mixed strategies are just mixtures over the pure strategies as before.
- Nash equilibria are also preserved.
- Note that we've now defined both mapping from NF games to IIEF and a mapping from IIEF to NF.
 - what happens if we apply each mapping in turn?
 - we might not end up with the same game, but we do get one with the same strategy spaces and equilibria.

Randomized Strategies


- It turns out there are two meaningfully different kinds of randomized strategies in imperfect information extensive form games
 - mixed strategies
 - behavioral strategies
- **Mixed strategy**: randomize over pure strategies
- **Behavioral strategy**: independent coin toss every time an information set is encountered

Randomized strategies example




- Give an example of a behavioral strategy:

Randomized strategies example



- Give an example of a behavioral strategy:
 - A with probability $.5$ and G with probability $.3$
- Give an example of a mixed strategy that is not a behavioral strategy:


Randomized strategies example



- Give an example of a behavioral strategy:
 - A with probability $.5$ and G with probability $.3$
- Give an example of a mixed strategy that is not a behavioral strategy:
 - $(.6(A, G), .4(B, H))$ (why not?)
- In this game every behavioral strategy **corresponds to** a mixed strategy...

Games of imperfect recall


Imagine that player 1 sends two proxies to the game with the same strategies. When one arrives, he doesn't know if the other has arrived before him, or if he's the first one.



- What is the space of pure strategies in this game?

Games of imperfect recall


Imagine that player 1 sends two proxies to the game with the same strategies. When one arrives, he doesn't know if the other has arrived before him, or if he's the first one.



- What is the space of pure strategies in this game?
 - 1: (L, R) ; 2: (U, D)

Games of imperfect recall


Imagine that player 1 sends two proxies to the game with the same strategies. When one arrives, he doesn't know if the other has arrived before him, or if he's the first one.



- What is the space of pure strategies in this game?
 - 1: (L, R) ; 2: (U, D)
- What is the mixed strategy equilibrium?

Games of imperfect recall


Imagine that player 1 sends two proxies to the game with the same strategies. When one arrives, he doesn't know if the other has arrived before him, or if he's the first one.



- What is the space of pure strategies in this game?
 - 1: (L, R) ; 2: (U, D)
- What is the mixed strategy equilibrium?
 - Observe that D is dominant for 2. R, D is better for 1 than L, D , so R, D is an equilibrium.


Games of imperfect recall

Imagine that player 1 sends two proxies to the game with the same strategies. When one arrives, he doesn't know if the other has arrived before him, or if he's the first one.




- What is the space of pure strategies in this game?
 - 1: (L, R) ; 2: (U, D)
- What is the mixed strategy equilibrium?
 - Observe that D is dominant for 2. R, D is better for 1 than L, D , so R, D is an equilibrium.

Games of imperfect recall



- What is an equilibrium in behavioral strategies?

Games of imperfect recall



- What is an equilibrium in behavioral strategies?
 - again, D strongly dominant for 2
 - if 1 uses the behavioural strategy $(p, 1 - p)$, his expected utility is $1 * p2 + 100 * p(1 - p) + 2 * (1 - p)$
 - simplifies to $-99p^2 + 98p + 2$
 - maximum at $p = 98/198$
 - thus equilibrium is $(98/198, 100/198), (0, 1)$
- Thus, we can have behavioral strategies that are different from mixed strategies.

Lecture Overview

- 1 Recap
- 2 Centipede Game
- 3 Imperfect-Information Extensive-Form Games
- 4 Perfect Recall**

Perfect Recall: mixed and behavioral strategies coincide

No player forgets anything he knew about moves made so far.

Definition

Player i has **perfect recall** in an imperfect-information game G if for any two nodes h, h' that are in the same information set for player i , for any path $h_0, a_0, h_1, a_1, h_2, \dots, h_n, a_n, h$ from the root of the game to h (where the h_j are decision nodes and the a_j are actions) and any path $h_0, a'_0, h'_1, a'_1, h'_2, \dots, h'_m, a'_m, h'$ from the root to h' it must be the case that:

- 1 $n = m$
- 2 For all $0 \leq j \leq n$, h_j and h'_j are in the same equivalence class for player i .
- 3 For all $0 \leq j \leq n$, if $\rho(h_j) = i$ (that is, h_j is a decision node of player i), then $a_j = a'_j$.

G is a game of perfect recall if every player has perfect recall in it.

Perfect Recall

Clearly, every perfect-information game is a game of perfect recall.

Theorem (Kuhn, 1953)

*In a game of perfect recall, any mixed strategy of a given agent **can be replaced by an equivalent behavioral strategy**, and any behavioral strategy can be replaced by an equivalent mixed strategy. Here two strategies are equivalent in the sense that they induce the same probabilities on outcomes, for any fixed strategy profile (mixed or behavioral) of the remaining agents.*

Corollary

In games of perfect recall the set of Nash equilibria does not change if we restrict ourselves to behavioral strategies.

Computing Equilibria of Games of Perfect Recall

How can we find an equilibrium of an imperfect information extensive form game?

- One idea: convert to normal form, and use techniques described earlier.
 - Problem: exponential blowup in game size.
- Alternative (at least for perfect recall): **sequence form**
 - for zero-sum games, computing equilibrium is polynomial in the size of the extensive form game
 - exponentially faster than the LP formulation we saw before
 - for general-sum games, can compute equilibrium in time exponential in the size of the extensive form game
 - again, exponentially faster than converting to normal form