Extensive Form Games: Backward Induction and Imperfect Information Games

CPSC 532A Lecture 10

October 12, 2006

Lecture Overview

Recap

Backward Induction

Imperfect-Information Extensive-Form Games

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
  - imperfect-information extensive-form games

Definition

A (finite) perfect-information game (in extensive form) is a tuple $G = (N, A, H, Z, \chi, \rho, \sigma, u)$, where

- N is a set of n players
- $A = (A_1, \ldots, A_n)$ is a set of actions for each player
- H is a set of non-terminal choice nodes
- Z is a set of terminal nodes, disjoint from H
- $\chi : H \rightarrow 2^A$ is the action function
  - assigns to each choice node a set of possible actions
- $\rho : H \rightarrow N$ is the player function
  - assigns to each non-terminal node a player $i \in N$ who chooses an action at that node

Definition

A (finite) perfect-information game (in extensive form) is a tuple $G = (N, A, H, Z, \chi, \rho, \sigma, u)$, where

- $\chi : H \rightarrow 2^A$ is the action function
  - assigns to each choice node a set of possible actions
- $\rho : H \rightarrow N$ is the player function
  - assigns to each non-terminal node a player $i \in N$ who chooses an action at that node
- $\sigma : H \times A \rightarrow H \cup Z$ is the successor function
  - maps a choice node and an action to a new choice node or terminal node such that for all $h_1, h_2 \in H$ and $a_1, a_2 \in A$, if $\sigma(h_1, a_1) = \sigma(h_2, a_2)$ then $h_1 = h_2$ and $a_1 = a_2$
- $u = (u_1, \ldots, u_n)$, where $u_i : Z \rightarrow \mathbb{R}$ is a utility function for player i on the terminal nodes Z

Note: the choice nodes form a tree, so we can identify a node with its history.

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

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

Nash Equilibria

Given our new definition of pure strategy, we are able to reuse our old definitions of:

- mixed strategies
- best response
- Nash equilibrium

Theorem

Every perfect information game in extensive form has a PSNE

This is easy to see, since the players move sequentially.

Induced Normal Form

- In fact, the connection to the normal form is even tighter
  - we can "convert" an extensive-form game into normal form

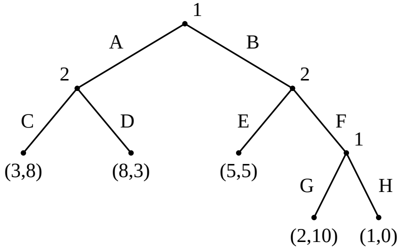

	CE	CF	DE	DF
AG	$3,8$	$3,8$	$8,3$	$8,3$
AH	$3,8$	$3,8$	$8,3$	$8,3$
BG	$5,5$	$2,10$	$5,5$	$2,10$
BH	$5,5$	$1,0$	$5,5$	$1,0$

Subgame Perfection

- Define subgame of G rooted at h:
  - the restriction of G to the descendents 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"

Lecture Overview

Recap

Backward Induction

Imperfect-Information Extensive-Form Games

Centipede Game

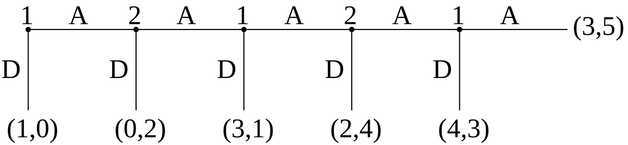

- Play this as a fun game...

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 ∈ Z then return u(h) {h is a terminal node}
    best_util ← −∞
    for all a ∈ χ(h) do {all actions available at node h}
        util_at_child ←BackwardInduction(σ(h, a))
        if util_at_child_ρ(h) > best_util_ρ(h) then
            best_util ← util_at_child
        end if
    end for
    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.

Backward Induction

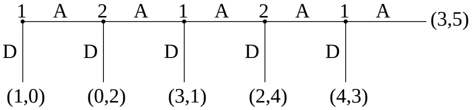

- What happens when we use this procedure on Centipede?
  - In the only equilibrium, player 1 goes down in the first move.
  - However, this outcome is Pareto-dominated by all but one other outcome.
- Two considerations:
  - practical: human subjects don't go down right away
  - theoretical: what should 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

Recap

Backward Induction

Imperfect-Information Extensive-Form Games

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 games (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
- $I = (I_1, \ldots, I_n)$, where $I_i = (I_{i,1}, \ldots, I_{i,k_i})$ 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 $h \in Ii, j$ and $h' \in Ii, j$.

Example

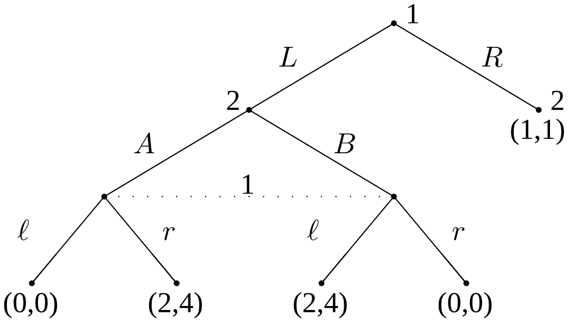

- What are the equivalence classes for each player?
- What are the pure strategies for each player?

Example

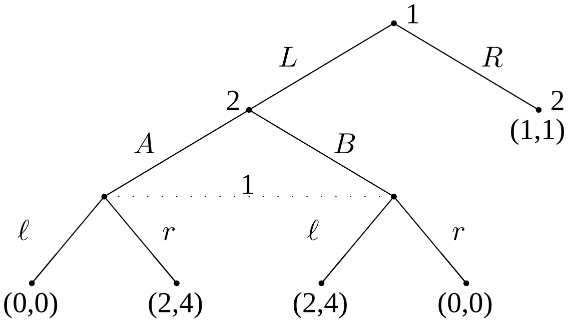

- What are the equivalence classes for each player?
- What are the pure strategies for each player?
  - choice of an action in each equivalence class.
- Formally, the pure strategies of player i consist of the cross product $\times_{I_{i,j} \in I_i} \chi(I_{i,j})$.

Normal-form games

- We can represent any normal form game.

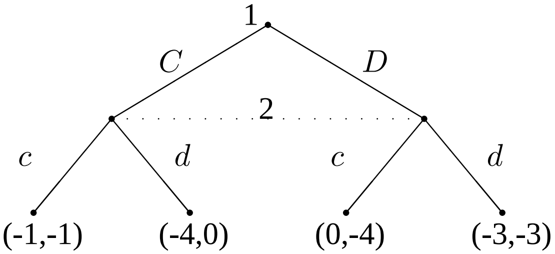

  - Note that it would also be the same if we put player 2 at the root node.

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

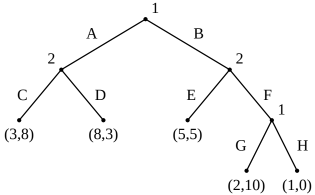

- Give an example of a behavioral strategy:

Randomized strategies example

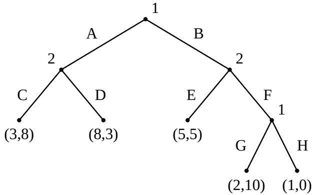

- Give an example of a behavioral strategy:
  - A with probability $.5$ and G with probability $.3$
- Give an example of a mixed strategy that is not a behavioral strategy:

Randomized strategies example

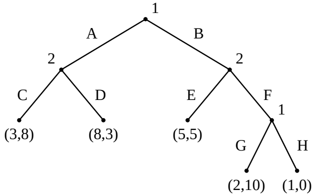

- Give an example of a behavioral strategy:
  - A with probability $.5$ and G with probability $.3$
- Give an example of a mixed strategy that is not a behavioral strategy:
  - $(.6(A, G), .4(B, H))$ (why not?)
- In this game every behavioral strategy corresponds to a mixed strategy...

Games of imperfect recall

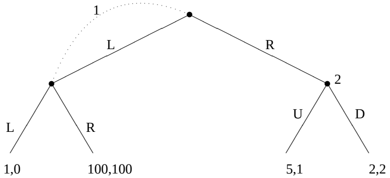

- What is the space of pure strategies in this game?

Games of imperfect recall

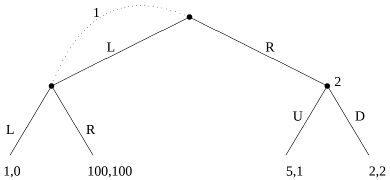

- What is the space of pure strategies in this game?
  - 1: (L, R); 2: (U, D)

Games of imperfect recall

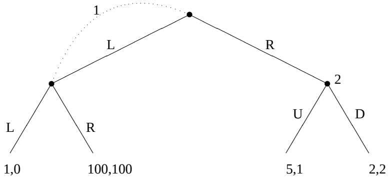

- What is the space of pure strategies in this game?
  - 1: (L, R); 2: (U, D)
- What is the mixed strategy equilibrium?

Games of imperfect recall

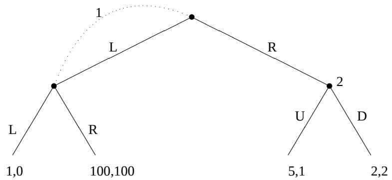

- What is the space of pure strategies in this game?
  - 1: (L, R); 2: (U, D)
- What is the mixed strategy equilibrium?
  - Observe that D is dominant for 2. R, D is better for 1 than L, D, so R, D is an equilibrium.

Games of imperfect recall

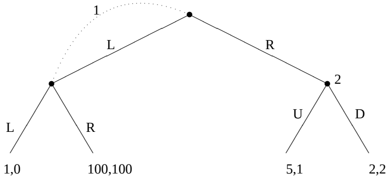

- What is the space of pure strategies in this game?
  - 1: (L, R); 2: (U, D)
- What is the mixed strategy equilibrium?
  - Observe that D is dominant for 2. R, D is better for 1 than L, D, so R, D is an equilibrium.

Games of imperfect recall

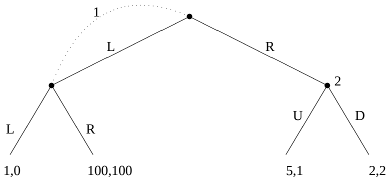

- What is an equilibrium in behavioral strategies?

Games of imperfect recall

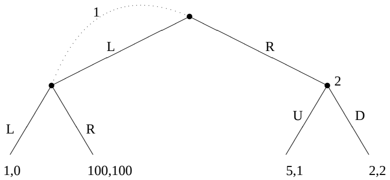

- What is an equilibrium in behavioral strategies?
  - again, D strongly dominant for 2
  - if 1 uses the behavioural strategy $(p, 1-p)$, his expected utility is $1 * p2 + 100 * p(1-p) + 2 * (1-p)$
  - simplifies to $-99p^2 + 98p + 2$
  - maximum at $p = 98/198$
  - thus equilibrium is $(98/198, 100/198),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