Repeated Games

ISCI 330 Lecture 16

March 13, 2007

Lecture Overview

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.

Example

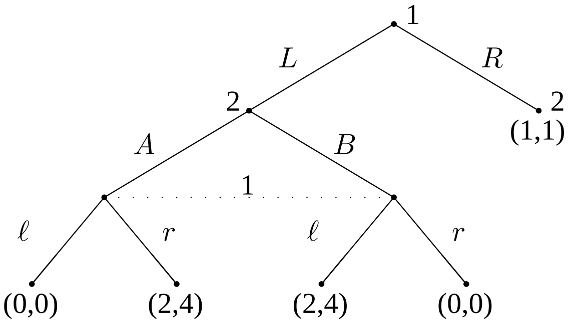

- What are the equivalence classes for each player?
- The pure strategies for each player are a choice of an action in each equivalence class.

Normal-form games

- We can represent any normal form game.

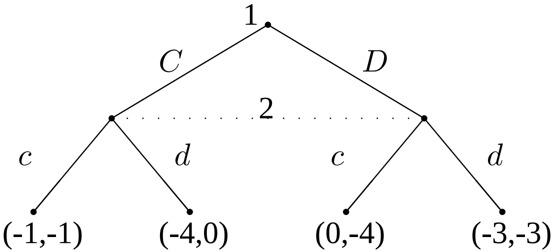

  - Note that it would also be the same if we put player 2 at the root node.

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 are now able both to convert NF games to EF, and EF games to NF.

Lecture Overview

Introduction

- Play the same normal-form game over and over
  - each round is called a "stage game"
- Questions we'll need to answer:
  - what will agents be able to observe about others' play?
  - how much will agents be able to remember about what has happened?
  - what is an agent's utility for the whole game?
- Some of these questions will have different answers for finitely- and infinitely-repeated games.

Finitely Repeated Games

- Everything is straightforward if we repeat a game a finite number of times
- we can write the whole thing as an extensive-form game with imperfect information
  - at each round players don't know what the others have done; afterwards they do
  - overall payoff function is additive: sum of payoffs in stage games

Example

	C	D
C	$-1, -1$	$-4, 0$
D	$0, -4$	$-3, -3$

$\Rightarrow$

	C	D
C	$-1, -1$	$-4, 0$
D	$0, -4$	$-3, -3$

Example

	C	D
C	$-1, -1$	$-4, 0$
D	$0, -4$	$-3, -3$

$\Rightarrow$

	C	D
C	$-1, -1$	$-4, 0$
D	$0, -4$	$-3, -3$

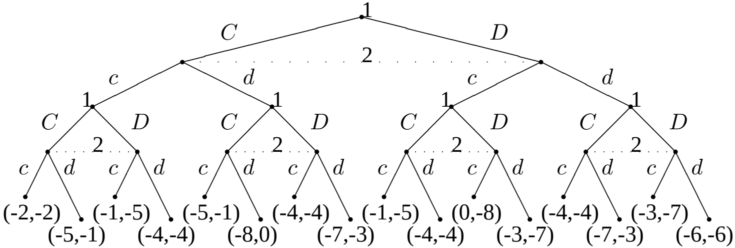

Notes

- Observe that the strategy space is much richer than it was in the NF setting
- Repeating a Nash strategy in each stage game will be an equilibrium (called a stationary strategy)
  - however, there can also be other equilibria
- In general strategies adopted can depend on actions played so far
- We can apply backward induction in these games when the normal form game has a dominant strategy.

Lecture Overview

Infinitely Repeated Games

- Consider an infinitely repeated game in extensive form:
  - an infinite tree!
- Thus, payoffs cannot be attached to terminal nodes, nor can they be defined as the sum of the payoffs in the stage games (which in general will be infinite).

Definition

Given an infinite sequence of payoffs $r_1, r_2, \ldots$ for player i, the average reward of i is $\lim_{k\to\infty} \Sigma_{j=1}^{k} r_j / k$.

Discounted reward

Definition

Given an infinite sequence of payoffs $r_1, r_2, \ldots$ for player i and a discount factor β with $0 \leq \beta \leq 1$, the future discounted rewards of i is $\sum_{j=1}^{\infty} \beta^j r_j$.

- Interpreting the discount factor:
  1. the agent cares more about his well-being in the near term than in the long term
  2. the agent cares about the future just as much as the present, but with probability $1 - \beta$ the game will end in any given round.
- The analysis of the game is the same under both perspectives.

Strategy Space

- What is a pure-strategy in an infinitely-repeated game?

Strategy Space

- What is a pure-strategy in an infinitely-repeated game?
  - a choice of action at every decision point
  - here, that means an action at every stage game
  - ...which is an infinite number of actions!
- Some famous strategies (repeated PD):
  - Tit-for-tat: Start out cooperating. If the opponent defected, defect in the next round. Then go back to cooperation.
  - Trigger: Start out cooperating. If the opponent ever defects, defect forever.

Nash Equilibria

- With an infinite number of equilibria, what can we say about Nash equilibria?
  - we won't be able to construct an induced normal form and then appeal to Nash's theorem to say that an equilibrium exists
  - Nash's theorem only applies to finite games
- Furthermore, with an infinite number of strategies, there could be an infinite number of pure-strategy equilibria!
- It turns out we can characterize a set of payoffs that are achievable under equilibrium, without having to enumerate the equilibria.