Repeated Games and the Folk Theorem

Lecture 9

Lecture Overview

1. Recap
2. Repeated Games
3. Infinitely Repeated Games
4. Folk Theorem

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
- $I = (I_1, \ldots, I_n)$, 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 $h \in I_{i,j}$ and $h' \in I_{i,j}$.

Normal-form games

- We can represent any normal form game.

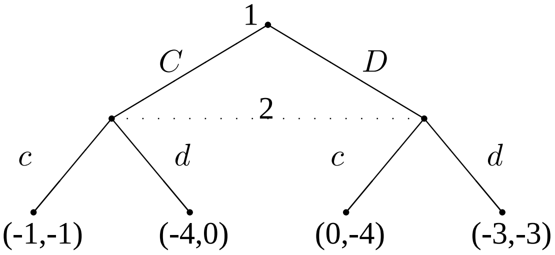

- Note that it would also be the same if we put player 2 at the root node.

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

Perfect Recall: mixed and behavioral strategies coincide

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 $h_0, a_0, h_1, a_1, h_2, \ldots, h_n, a_n, h$ from the root of the game to h (where the h_j are decision nodes and the a_j are actions) and any path $h_0, a'_0, h'_1, a'_1, h'_2, \ldots, h'_m, a'_m, h'$ 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.

Lecture Overview

1. Recap
2. Repeated Games
3. Infinitely Repeated Games
4. Folk Theorem

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	$0, -4$	$-3, -3$

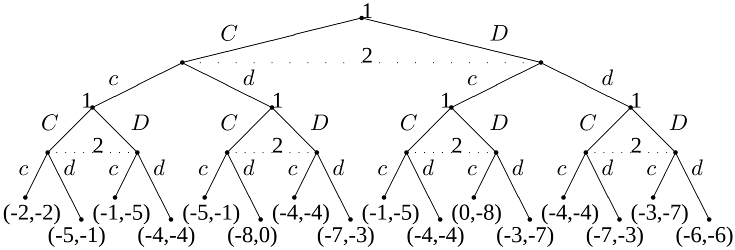

Example

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

$\Rightarrow$

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

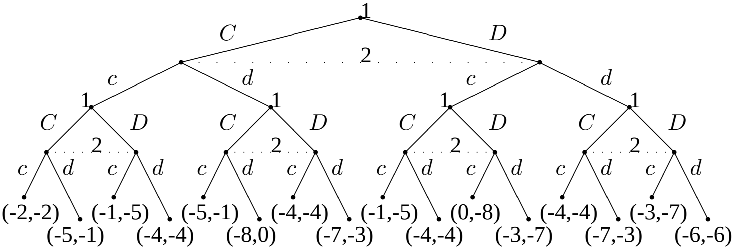

Play repeated prisoner's dilemma with one or more partners.
Repeat the game five times.

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 in behavioral strategies (called a stationary strategy)
- 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

1. Recap
2. Repeated Games
3. Infinitely Repeated Games
4. Folk Theorem

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

$$\lim_{k \to \infty} \sum_{j=1}^{k} \frac{r_j}{k}.$$

Discounted reward

Definition

Given an infinite sequence of payoffs $r_1, r_2, \ldots$ for player i and discount factor β with $0 \leq \beta \leq 1$, i's future discounted reward 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.

Lecture Overview

1. Recap
2. Repeated Games
3. Infinitely Repeated Games
4. Folk Theorem

Definitions

- Consider any n-player game $G = (N, A, u)$ and any payoff vector $r = (r_1, r_2, \ldots, r_n)$.
- Let $v_i = \min_{s_{-i} \in S_{-i}} \max_{s_i \in S_i} u_i(s_{-i}, s_i)$.
  - i's minmax value: the amount of utility i can get when $-i$ play a minmax strategy against him

Definition

A payoff profile r is enforceable if $r_i \geq v_i$.

Definition

A payoff profile r is feasible if there exist rational, non-negative values α_a such that for all i, we can express r_i as $\sum_{a \in A} \alpha u_i(a)$, with $\sum_{a \in A} \alpha_a = 1$.

- a payoff profile is feasible if it is a convex, rational combination of the outcomes in G.

Folk Theorem

Theorem (Folk Theorem)

Consider any n-player game G and any payoff vector $(r_1, r_2, \dots, r_n)$.

1. *If r is the payoff in any Nash equilibrium of the infinitely repeated G with average rewards, then for each player i, r_i is enforceable.*
2. *If r is both feasible and enforceable, then r is the payoff in some Nash equilibrium of the infinitely repeated G with average rewards.*

Folk Theorem (Part 1)

Payoff in Nash $\rightarrow$ enforceable

Part 1: Suppose r is not enforceable, i.e. $r_i < v_i$ for some i. Then consider a deviation of this player i to $b_i(s_{-i}(h))$ for any history h of the repeated game, where b_i is any best-response action in the stage game and $s_{-i}(h)$ is the equilibrium strategy of other players given the current history h. By definition of a minmax strategy, player i will receive a payoff of at least v_i in every stage game if he adopts this strategy, and so i's average reward is also at least v_i. Thus i cannot receive the payoff $r_i < v_i$ in any Nash equilibrium.

Folk Theorem (Part 2)

Feasible and enforceable $\rightarrow$ Nash

Part 2: Since r is a feasible payoff profile, we can write it as $r_i = \sum_{a \in A} \left(\frac{\beta_a}{\gamma}\right) u_i(a)$, where β_a and γ are non-negative integers.[1] Since the combination was convex, we have $\gamma = \sum_{a \in A} \beta_a$. We're going to construct a strategy profile that will cycle through all outcomes $a \in A$ of G with cycles of length γ, each cycle repeating action a exactly β_a times. Let (a^t) be such a sequence of outcomes. Let's define a strategy s_i of player i to be a trigger version of playing (a^t): if nobody deviates, then s_i plays a_i^t in period t. However, if there was a period t' in which some player $j \neq i$ deviated, then s_i will play $(p_{-j})_i$, where (p_{-j}) is a solution to the minimization problem in the definition of v_j.

[1] Recall that α_a were required to be rational. So we can take γ to be their common denominator.

Folk Theorem (Part 2)

Feasible and enforceable $\rightarrow$ Nash

First observe that if everybody plays according to s_i, then, by construction, player i receives average payoff of r_i (look at averages over periods of length γ). Second, this strategy profile is a Nash equilibrium. Suppose everybody plays according to s_i, and player j deviates at some point. Then, forever after, player j will receive his $\min\max$ payoff $v_j \leq r_j$, rendering the deviation unprofitable.