Game theory for wireless networks

Lecture 6

Outline of the lecture

- Learning in games
  - Fictitious play – existence of Nash eq.
  - No regreat learning – an introduction

- Correlated Nash equilibrium
- Types of games
  - Dummy, coordination games
  - Super-modular games
  - Potential Games

Fictitious play – recall from last lecture

- Player i: initial weight function

$$K_0^i : S^{-i} \to \mathcal{R}^+$$

- Game iteratively repeated → K updated:

$$K_t(s^{-i}) = K_{t-1}(s^{-i}) + \begin{cases} 1, & \text{if } s_{t-1}^{-i} = s^{-i} \\ 0, & \text{ow.} \end{cases}$$

- Given the frequency vector K → updates beliefs
  - The belief player i has at time t about its opponent to play s^{-i} at time t:

$$\gamma_t^i\left(s^{-i}\right) = \frac{K_t^i\left(s^{-i}\right)}{\sum_{\hat{s} \in S^{-i}} K_t^i\left(\hat{s}^{-i}\right)}$$

← Simple normalization

Fictitious play

- Given the updated belief γ_t^i

- Fictitious play: any rule $p_t^i(\gamma_t^i) \in BR^i(\gamma_t^i)$

- Not a unique fictitious play rule → there may be more than one best response to a particular assessment

Convergence properties for fictitious play

- Proposition:
  - (1) If s is a strict Nash equilibrium, and s is played at date t in the process of fictitious play, then s is played at all subsequent dates (Nash equilibria are absorbing for the process of fictitious play)
  - (2) Any pure-strategy steady state of fictitious play must be Nash equilibrium
- Definition: marginal empirical distribution of play:

$$d_t^j(s^j) = \frac{k_t(s^j) - k_0(s^j)}{t}$$

- Proposition:
  - If the empirical distributions over each player's choices converge, the strategy profile corresponding to the product of these distributions is a Nash eq.
- Proposition: the fictitious play process converges for a two person zero-sum game

No regret learning

- **No regret learning strategies:** probabilistic learning alg. which specify that players ***explore*** the space of actions by playing all actions with some non-zero probability, and ***exploit*** successful actions, by increasing the probability of employing those actions that generate high profits.
  - Learn mixed strategy equilibria
- No external regret algorithms
  - **External regret:** difference between the payoffs achieved by the strategies prescribed by the given algorithm, and the payoffs achieved by any other fixed sequence of decisions, in the worst case.
- No internal regret algorithms
  - **Internal regret:** difference between the payoffs achieved by the strategies prescribed by the given algorithm, and the payoffs that could be achieved by a remapped sequence of strategies.
    - Sequence is remapped, if there is a mapping f of the strategy space into itself, s.t. for each occurrence of strategy x (e.g. x=169), the mapped strategy y (e.g. y = 170) appears in the re-mapped sequence
  - No internal regret → no external regret

No External Regret Algorithm – an example alg.

- Freud & Schapire:
  - "Game theory, on-line prediction and boosting", Proc. of the 9th Annual Conference on Computational Learning Theory, pp. 325-332, ACM press, May 1996.
- No external regret via multiplicative updating
- ρ_i^t = cumulative payoff obtained through time t via strategy t+1:

$$\rho_i^t = \sum_{x=0}^{t} r_i^x$$

- The weight associated with strategy t+1, for $\beta \in [0,1)$:

$$w_i^{t+1} = \frac{(1+\beta)^{\rho_i^t}}{\sum_{j=1}^{S}(1+\beta)^{\rho_j^t}}$$

No external regret –cont.

- The weight calculation represents a measure of regret
- Need to know payoff that would be obtained for all possible strategies
- Naive player: knows only his payoff for the currently played strategy
  - Previous algorithm may be modified to work for naive players: [Auer, Chesa-Bianchi, Freud and Schapire, 1995]
    - Gambling in a rigged casino: the adversarial multi-armed bandit problem, proc. of the 36th Annual Symposium on Foundations of Computer Science, pp. 322-331. ACM Press, Nov. 1995.

No internal regret – an example alg.

- The regret for a player at time t: difference between the payoffs achieved using its strategy of choice, e.g. i, and the payoffs that could have been achieved had strategy j≠i been played instead:

$$R_{i\to j}^{t} = 1_i^t \left(r_j^t - r_i^t \right)$$

Indicator function: 1 if strategy i employed at t, 0 ow.

- Cumulative regret:

$$R_{i\to j}^{T} = \sum_{t=0}^{T} R_{i\to j}^{t}$$

- Internal regret:

$$IR_{i\to j}^{T} = (R_{i\to j}^{T})^{+}$$

$$X^{+} = \max(X,0)$$

No internal regret – cont.

- Cumulative regret for playing all strategies but j:

$$IR_{S\to j}^{T}=\sum_{i=1}^{S} IR_{i\to j}^{T}$$

- Updating weights
  - If strategy i is played at time t,

$$w_j^{t+1}=\frac{1}{\mu}IR_{i\to j}^{t}$$

$$w_i^{t+1}=1-\sum_{j\neq i} w_j^{t+1}$$

  - μ = normalizing term:

$$\mu>\left(|S|-1\right)\max_{j\in S} IR_{i\to j}^{t}$$

No internal regret alg. – cont.

Some observations:

- Achieves no internal regret in the limit as $T \rightarrow \infty$
- Learning converges to the correlated Nash equilibrium
- Naive version also proposed
- Reference:
  - S. Hart & A. Mass Colell, “A simple adaptive procedure leading to correlated equilibrium”.

Correlated Nash equilibrium

- If players can engage in preplay communication, then go in separate rooms and choose their strategy independently
  - Might gain if they can build a signaling device → send signals to the separate rooms.
- Example of game with correlated equilibria

	L	R
U	5,1	0,0
D	4,4	1,5

  - Three Nash equilibria: (U,L), (D,R), and a mixed strategy eq. with equal probability on each pure strategy: payoff 2.5 for each player

Correlated Nash equilibria: cont

- If they can jointly observe a random variable, e.g. a coin flip
  - Player 1: U if heads, D if tails
  - Player 2: L if heads, R if tails

  → Payoff: (3,3)

  - More general: players can obtain any payoff vector in the convex hull of the set of Nash equilibrium payoffs
  - Cannot obtain any payoff outside the convex hull of the set of Nash equilibrium payoffs
- Can gain further if players receive different signals, but correlated
  - Build a signaling device that sends different, but correlated, signals to each of them
  - Device has three equally likely states: A, B, C
  - If A – player 1 completely informed, but it cannot distinguish between B and C
  - Player 2 – informed when C, cannot distinguish from A and B

Correlated Nash eq. – cont.

- A Nash eq. for the transformed game:
  - Player 1 plays U when A and D ow
  - Player 2: R when C, and L ow
  - Eq: (U,L), (D,L) and (D,R), occur with probability 1/3, payoff 3.33
  - Payoff outside the convex hull of the set of Nash equilibrium payoffs
- Note: signaling device may be interpreted as a recommendation on how to play
- Definition: A correlated eq. is any probability distribution p(.) over the pure strategies S1 x S2 x… Si, s.t. for every player i and every si, with p(si)>0,

$$\sum_{s_{-i} \in S_{-i}} p(s_{-i} \mid s_i) u_i(s_i, s_{-i}) \geq \sum_{s_{-i} \in S_{-i}} p(s_{-i} \mid s_i) u_i(s_i^{'}, s_{-i}) \quad \forall \ s_{i'} \in S_i$$

- Player I should not be able to gain by disobeying the recommendation to play si, if every other player obeys its recommendation

Types of games

Dummy game

- Unilateral deviations produce no change in the payoff of the deviating player.

a,b	c,b
a,d	c,d

- All pure strategy profiles are Nash eq.

$$u_i(s_i, s_{-i}) = u_i(s'_i, s_{-i})$$

Types of games: coordination games

Identical Interest Game

- For any choice of pure strategies, same payoffs for the players

a,a	b,b
c,c	d,d

$$u_i(s) = u_j(s), \quad \forall\, i, j \in N, \ \forall s \in S$$

- There exists some function *V*: $S \rightarrow \Re$ such that

$$u_i(s) = V(s) \ \ \forall i \in N, \ \forall\, s \in S$$

- All Maximizers of *V* are Nash eq.
- At least one Nash eq. must be Pareto Efficient

Types of games: supermodular games

- Games in which each player's marginal utility of increasing its strategy rises with increases in its rivals' strategies
- Super-modular games have pure strategy Nash equilibria

- To define supermodularity we need
  - Order structure on strategy spaces
  - Weak continuity requirements on payoffs
  - Supermodularity requirement

Order relation

- Let x and y denote two vectors in $\mathbf{R}^K$
- Let $x \geq y$ if $x_k \geq y_k$ for all $k = 1,2,\ldots,K$
- Let $x > y$ if $x \geq y$ and there exists some k such that $x_k > y_k$
- If a vector dominates another in one component, but is dominated in another component, vectors cannot be compared

Example:

$x = (2,1,5,3)$, $y = (3,5,7,3)$ $y \geq x$ $y > x$

$x = (2,0,8,1)$, $y = (1,5,9,1)$ Cannot be compared

"Meet" and "join" relations

"meet of x and y" $x \wedge y \equiv (\min\{x_1, y_1\}, \min\{x_2, y_2\}, \ldots, \min\{x_K, y_K\})$

"join of x and y" $x \vee y \equiv \{\max(x_1, y_1), \max(x_2, y_2), \ldots, \max(x_K, y_K)\}$

S is a sublattice of $\mathbf{R}^{m_i}$ if s, $s^* \in S$ then $s \wedge s^* \in S$ and $s \vee s^* \in S$.

Sublattice property – Every bounded sublattice has a greatest and a least element.

Increasing Differences Property

Definition

$u_i(s_i, s_{-i})$ has increasing differences in (s_i, s_{-i}) if, for all

$s_i, \tilde{s}_i \in S_i$ and $s_{-i}, \tilde{s}_{-i} \in S_{-i}$ such that $s_i \geq \tilde{s}_i$ and $s_{-i} \geq \tilde{s}_{-i}$

$$u_i\left(s_i, s_{-i}\right) - u_i\left(\tilde{s}_i, s_{-i}\right) \geq u_i\left(s_i, \tilde{s}_{-i}\right) - u_i\left(\tilde{s}_i, \tilde{s}_{-i}\right)$$

Intuitive explanation: an increase in the strategies of *i*’s rivals Increases the value of playing a high strategy for player *i*.

Supermodular Games

Definition:

$u_i(s_i, s_{-i})$ is supermodular in s_i if for each s_{-i}

$$u_i\left(s_i, s_{-i}\right) + u_i\left(\tilde{s}_i, s_{-i}\right) \leq u_i\left(s_i \wedge \tilde{s}_i, s_{-i}\right) + u_i\left(s_i \vee \tilde{s}_i, s_{-i}\right) \;\forall s_i, \tilde{s}_i \in S_i$$

The above relation is satisfied with equality if S_i is single-dimensional

Definition:

A game is supermodular if the following conditions are met for all i

S_i is a sublattice of $\mathbf{R}^m$

u_i has increasing differences in (s_i, s_{-i})

u_i is supermodular in s_i

Continuous payoff functions

- If u_i is twice continuously differentiable, u_i is a supermodular game iff, for any s_k and s_l (components of s)

$$\frac{\partial^2 u_i}{\partial s_l \partial s_k} \geq 0$$

- Example: Bertrand game
  - Oligopoly with demand function:

$$D_i(p_i, p_{-i}) = a_i - b_i p_i + \sum_{j \neq i} d_{ij} p_j, \; b_i > 0, \; d_{ij} > 0$$

  - And utility:

$$u_i(p_i, p_{-i}) = (p_i - c_i) D_i(p_i, p_{-i})$$

  - Supermodular game: $\dfrac{\partial^2 u_i}{\partial p_i \partial p_j} > 0, \quad \forall i, j \neq i$

Nash eq. for supermodular games

➢ (Topkis) A supermodular game for which each S_i is compact and each u_i is upper semi-continuous in s_i for each s_{-I}, then the set of pure strategy Nash eq. is non-empty and contains greatest and least elements $\bar{s}, \underline{s}$,respectively

Note on upper semi-continuity:

A function *f:X → R*
- upper semi-continuous if, for every a in R, the preimage of [a,∞) is closed.
- lower semi-continuous if, for every a in R, the preimage of (-∞,a] is closed

f:X → R, the image of x is f(x). The preimage of y is $f^{-1}(y) = \{x \mid f(x) = y\}$ *, for all x whose image is y*

Best response for supermodular games

- (Milgrom and Roberts) A best response dynamic for supermodular game with compact action spaces and upper semi continuous objective functions converges to a region bounded by the greatest and least elements in the set of Nash eq.

- If the Nash eq. is unique, then the best response dynamic converges to the Nash eq..

Potential games

Exact potential game:

- A game for which you can construct a single-dimensional function

  $P: S \rightarrow \Re$ whose change in value is exactly equal to the utility change in value of the deviating player.

$$P(s_i, s_{-i}) - P(s'_i, s_{-i}) = u(s_i, s_{-i}) - u(s'_i, s_{-i})$$

Coordination games: potential games Combination of coordination and dummy games: potential games Most of congestion games: potential games

- A game $G = \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ is an exact potential game *iff* there exist functions $\{c_i\}_{i \in N}$ and $\{d_i\}_{i \in N}$ such that
  - $u_i = c_i + d_i$
  - $\langle N, \{S_i\}_{i \in N}, \{c_i\}_{i \in N} \rangle$ is a coordination game
  - $\langle N, \{S_i\}_{i \in N}, \{d_i\}_{i \in N} \rangle$ is a dummy game

Nash equilibrium

- The Nash equilibria of an exact potential game are the same as the ones for its constituent coordination game.

- Proof sketch:
  - unilateral deviation for a dummy game yields the same payoff → adding a dummy game *D* to any other game will preserve the Nash eq.
  - Any potential game can be decomposed as the sum of a coordination game and a dummy game
  - Nash eq. for the potential game must be the same as for the constituent coordination game

Nash eq. More properties

- For an exact potential game, the maximizers of the potential function are Nash eq. for the game.
- For a finite exact potential game, i.e., finite strategy space, finite player set, the game has at least one pure-strategy Nash equilibrium.
- All repeated games where each stage is the same finite Exact Potential Game (EPG) and all players are myopic, converge to the Nash equilibria of the stage game, if play follows a better response dynamic.
  - Definition: Better response dynamic

    At each stage, one player $i \in N$ is permitted to deviate from s_i to some randomly selected action $s'_i \in S_i$ iff it is an improvement deviation (the action improves utility).
- All repeated myopic games with the stage game an EPG, converge to the Nash eq. of the stage game, if play follows a best response dynamic.

How to recognize potential games?

- Find potential function
  - Using def. $P(s_i, s_{-i}) - P(s'_i, s_{-i}) = u(s_i, s_{-i}) - u(s'_i, s_{-i})$
  - For continuous payoff:

$$\frac{\partial P}{\partial s_i} = \frac{\partial u_i}{\partial s_i}, \quad \forall\, i \in N$$

$$\frac{\partial^2 P}{\partial s_i \partial s_j} = \frac{\partial^2 u_i}{\partial s_i \partial s_j} = \frac{\partial^2 u_j}{\partial s_i \partial s_j}, \quad \forall \mathrm{i}, \mathrm{j} \in \mathrm{N}$$

How to recognize potential games?

Bilateral Symmetric Interaction Game

Strategic form game for which a player's objective function is a sum of bilateral symmetric interaction (BSI) terms.
A BSI term is defined as

$$w_{ij} : S_i \times S_j \to \mathcal{R}\,, \text{ s.t. } w_{ij}(s_i, s_j) = w_{ji}(s_j, s_i) \;\; \forall (s_i, s_j) \in S_i \times S_j$$

The objective function: $u_i(s) = \sum_{j \neq i} w_{ij}\left(s_i, s_j\right) - h_i(a_i)$

Potential function for this game can be defined as:

$$\sum_{i \in N} \sum_{j=1}^{i-1} w_{ij}(s_i, s_j) - \sum_{i \in N} h_i(s_i)$$

How to recognize potential games?

- Coordination game
- Coordination + dummy games
- Congestion games

Homework

1) Assume that there are two players in a game. Player 1 has private information about his type, and can choose to send a message to player 1 to reveal his type. Player 1 can choose to send “1” or “0” on a signaling channel. His choices are

- “1” if he is type A, and “0” if he is type “B”
- “1” for both cases
- “0” for both cases

Player’s 2 apriori belief on 1’s type is ½ for both types.

After receiving the signal, player 2 chooses an action: a1, a2, a3, and both players get payoffs according to the following payoff matrix:

Homework: cont.

	a1	a2	a3
A	2,3	0,2	-1,0
B	1,0	2,2	0,3

- a) Consider the case of error free signaling channel
  - Show that there is no BNE (Bayesian Nash eq.) for which player 1 signals its type, and that the BNE is player 2 choose a2 no matter what player 1 says.
- b) Consider now that the channel occasionally (with probability ½) flips the “1” into “0” (“0” never gets flipped)
  - Show that a BNE equilibrium is: player 1 signals “1” if of type A and “0” if of type B, and player 2 plays a1 if “1” and “a2” if “0”.

Homework: cont

2) Build a potential function for the prisoner Dilemma game.

	C	D
C	1,1	-1,2
D	2,-1	0,0