Game theory for wireless networks

Lecture 5

Outline of the lecture

- Dynamic games of incomplete information
  - Perfect Bayesian equilibrium
  - Sequential equilibrium
- Learning in games:
  - Fictitious play

Dynamic games of incomplete information

- Recall: static games of incomplete information
  - The game/payoffs depend on the type of players. A player knows its own type but it does not know the types of the other players.
  - Transform a game of incomplete information → game of imperfect information
    - Assign probabilities for the types of the players
    - Perceived as a move by nature
    - Represents the players' *apriori* belief on the types of other players
- What changes for the dynamic game?
  - Players have the chance of updating their beliefs based on the observed actions of the other players.

Subgame perfection for dynamic games of incomplete information?

- The concept of subgame perfection → harder to apply for games of incomplete information
  - Start of a period does not form a well-defined subgame
  - Formally: the only proper subgame of a game of incomplete information is the whole game → any Nash equilibrium is subgame perfect
  - To illustrate that: consider the following example of a game of imperfect information

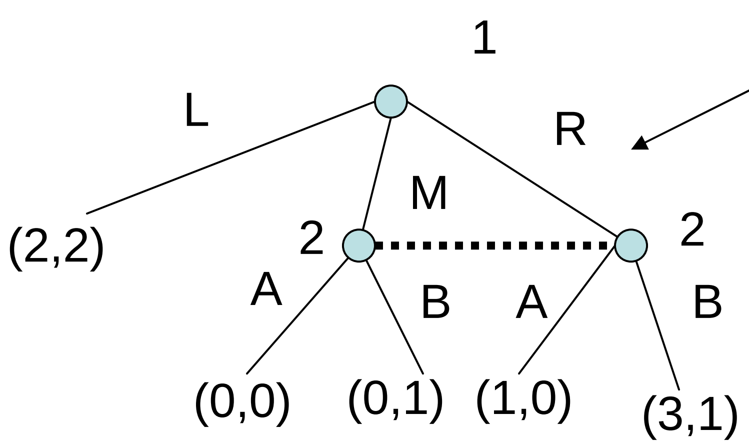

What are the Nash eq. for this game?

Subgame perfection

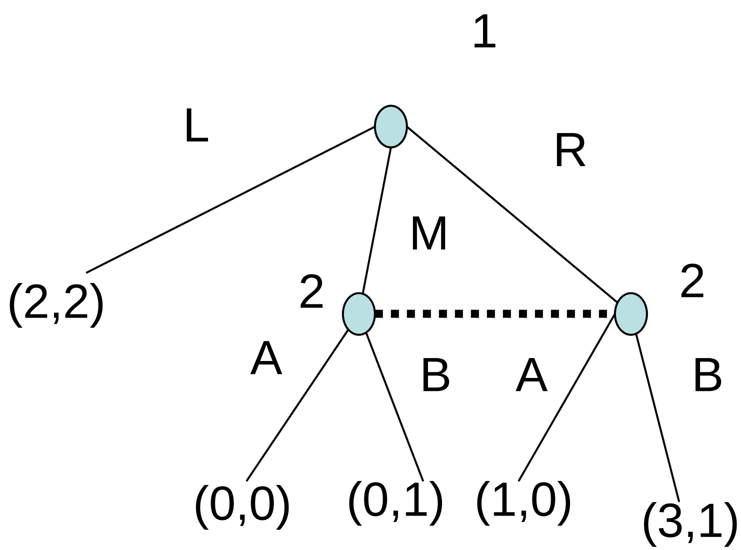

- Nash equilibria: (L,A) and (R,B)
- Both subgame perfect

However, (L,A) is a sub-optimal action → may be ruled out for the case of sequential equilibrium

Perfect Bayesian eq. for multi-stage games

- **The Basic Signaling Game**
  - Simplest type of game: 2 players
    - Player 1: leader (sender)
    - Player 2: follower (receiver)
    - Player 1 has private inf. about its type $\theta \in \Theta$, action $a_1 \in A_1$
    - Player 2: its type is common knowledge, action $a_2 \in A_2$
    - Space of mixed actions: $\mathcal{A}_1$ and $\mathcal{A}_2$, with elements α_1 and α_2
    - Utility of player i: $u_i(\alpha_1, \alpha_2, \theta)$
    - Player 2 – prior belief about player's 1 type: *p* → common knowledge
    - Strategy for player 1: probability distribution $\sigma_1(.| \theta)$ over actions $a_1 \in A_1$ for each type θ
    - Strategy for player 2: probability distribution $\sigma_2(.| \theta)$ over actions $a_2 \in A_2$ for each $a_1 \in A_1$

Basic signaling game: payoffs

- Payoff for player 1, given its type θ:

$$u_1(\sigma_1, \sigma_2, \theta) = \sum_{a_1} \sum_{a_2} \sigma_1(a_1 \mid \theta) \sigma_2(a_2 \mid a_1) u_1(a_1, a_2, \theta)$$

- Player 2's (ex ante – beforehand payoff) to strategy $\sigma_2(.|a_1)$, when player 1 plays $\sigma_1(.|\, \theta)$:

$$\sum_{\theta} p(\theta) \left[\sum_{a_1} \sum_{a_2} \sigma_1(a_1 \mid \theta) \sigma_2(a_2 \mid a_1) u_2(a_1, a_2, \theta) \right]$$

- Should we make the decision based on the above computed payoff?

Posterior beliefs

- Player 2, observes the action of player 2 → must update its belief on θ, and its choice of action → posterior distribution over Θ: $\mu(.|a_1)$
- How to compute $\mu(.|a_1)$?
  - Player 1 actions may depend on its type
  - Let $\sigma_1^*(.|\theta)$ to be player 1 strategy
  - Know p(.), $\sigma_1^*(.|\theta)$ and observe a_1: use Bayes rule
- Extension of subgame-perfect eq. → perfect Bayesian eq.
  - Player 2 max. its payoff conditional on a_1:

$$\sum_{\theta} \mu(\theta \mid a_1) u_2\big(a_1, \sigma_2(. \mid a_1), \theta\big) = \sum_{\theta} \sum_{a_2} \mu(\theta \mid a_1) \sigma_2(a_2 \mid a_1) u_2\big(a_1, a_2, \theta\big)$$

Perfect Bayesian Equilibrium (PBE)

- **Definition:** A perfect Bayesian eq. of a signaling game is a strategy profile σ* and posterior beliefs $\mu(.|a_1)$, s.t.:

- (P1): $\forall\theta,\ \sigma_1^*(.|\theta) \in \arg\max_{\alpha_1} u_1(\alpha_1, \sigma_2^*, \theta)$

- (P2): $\forall a_1,\ \sigma_2^*(.|a_1) \in \arg\max_{\alpha_2} \sum_{\theta} \mu(\theta|a_1) u_2(a_1, \alpha_2, \theta)$

- (B) $\mu(\theta|a_1) = \dfrac{p(\theta)\sigma_1^*(a_1|\theta)}{\sum_{\theta'\in\Theta} p(\theta')\sigma_1^*(a_1|\theta')}, \quad \text{if} \quad \sum_{\theta'\in\Theta} p(\theta')\sigma_1^*(a_1|\theta') > 0$

μ(.|a1), is any prob. distr. on Θ, if $\sum_{\theta'\in\Theta} p(\theta')\sigma_1^*(a_1|\theta') = 0$

Backward induction solution?

- Backward induction was used in games with perfect information
- PBE: strategies are optimal given the beliefs and the beliefs are obtained from equilibrium strategies and observed actions
  – Circularity → PBE cannot be determined by backward induction

An example of a signaling game

- **Two period reputation game**
  - 2 firms on the market

  Period 1: both firms of the market; only firm 1 action a1
  - Actions for firm 1: prey or accommodate
    - If prey: firm 2 gets P2<0
    - If accommodate: firm 2 gets D2>0
  - Type of firm 1: “sane” or “crazy”
    - “crazy” – always prey
    - “sane”
      - If accommodates: payoff for 1: D1>0
      - If preys: payoff for 1: P1< D1 → prefers to accommodate
      - If 2 exits→ M1>D1 (monopoly)

Two period reputation game: cont.

- Period 2 - player 2 selects action a2: stay or exit
  - If exits, it gets a 0 payoff, and player 1 gets M1> D1
- Assumptions:
  - Player 1 knows his type
  - Player 2 believes that player 1 is sane with probability p
  - δ = discount factor between the two periods
- Building reputation for the sane player
  - Player 1 may try to convince player 2 that he is crazy, to get M1>D1 in the second period of the game.

Taxonomy of PBE

- **Separating equilibrium:** the two types of player 1, choose two different actions in period 1
  - Firm 2 has complete information for the second period

$$\mu(\theta = \text{sane} \mid a_1 = \text{accomodate}) = 1$$
$$\mu(\theta = \text{crazy} \mid a_1 = \text{prey}) = 1$$

- **Pooling equilibrium:** the two types of player 1, choose the same action in period 1

$$\mu(\theta = \text{sane} \mid a_1 = \text{prey}) = p$$

- **Hybrid (semi-separating equilibria):** the sane type may randomize between preying and accommodating

$$\mu(\theta = \text{sane} \mid a_1 = \text{prey}) \in (0, p)$$
$$\mu(\theta = \text{sane} \mid a_1 = \text{accommodate}) = 1$$

What type of equilibrium? Existence.

- Separating eq. existence: sufficient and necessary condition

$$D_1(1+\delta) \geq P_1 + \delta M_1 \Leftrightarrow \delta(M_1 - D_1) \leq (D_1 - P_1)$$

- Pooling eq. → to enforce exit for player 2. Condition:

$$pD_2 + (1-p)P_2 \leq 0$$

- If the above two conditions do not hold → hybrid PBE

Note: uniqueness of the eq. in this case, is due to the fact that the “crazy” type is assumed to always prey.

Multi-stage games with observed actions and incomplete information

- Each player i has type θ_i, and types are independent

$$p(\theta) = \prod_{i=1}^{I} p_i(\theta_i)$$

- At each period t (t=0,1,2,… T), players choose their actions simultaneously, and the actions are revealed at the end of the period
- Players' action set at a date t is type independent
- Behavior strategy: $\sigma_i(a_i \mid h^t, \theta)$
- Payoffs $u_i(h^{T+1}, \theta)$

- Subgame perfection → BNE not only for the whole game, but also for the "continuation game" starting at period t after all possible histories h^t
  - Continuation games → proper subgames?
    - No. They do not stem from a singleton information set

Continuation games → true games

- Need to specify the players' beliefs at the start of each continuation game.
- **Definition:** A perfect Bayesian equilibrium is a (σ,μ) that satisfies (P) and (B(i) – B(iv)).

B(i) Posterior beliefs are independent, and all types of player i have the same belief.

- For all θ, t and h^t:

$$\mu_i(\theta_{-i} \mid \theta_i, h^t) = \prod_{j \neq i} \mu_i(\theta_j \mid h^t)$$

- even unexpected events will not change the independence assumption for the type of the opponents

Perfect Bayesian equilibrium: cont

- B(ii) Beliefs are updated according to Bayes' rule:
  - For all i,j, h^t, and a_j^t, If there exist

$$\hat{\theta}_j,\ \text{s.t. } \mu_i\left(\theta_j \mid h^t\right) > 0,\ \sigma_{\mathrm{j}}\left(a_j^t \mid h^t, \theta_j\right) > 0,\ \text{then}$$

$$\mu_i\left(\theta_j \mid h^t, a^t\right) = \frac{\mu_i\left(\theta_j \mid h^t\right)\sigma_j\left(a_j^t \mid h^t, \theta_j\right)}{\sum_{\hat{\theta}_j} \mu_i\left(\hat{\theta}_j \mid h^t\right)\sigma_j\left(a_j^t \mid h^t, \hat{\theta}_j\right)}$$

- B(iii) Don't signal what you don't know
  - For all i,j, h^t, and a^t and $\hat{a}^t$

$$\mu_i\left(\theta_j \mid (h^t, a^t)\right) = \mu_i\left(\theta_j \mid (h^t, \hat{a}^t)\right),\ \text{if } a_j^t = \hat{a}_j^t$$

Perfect Bayesian eq. – cont.

- B(iv) All players have to have the same belief about the type of another player
  - Imposed because of the req. of eq. analysis: players have the same belief about each other's strategies.
  - For all θ_k, and h^t

$$\mu_i\left(\theta_k \mid h^t\right) = \mu_k\left(\theta_k \mid h^t\right) = \mu\left(\theta_k \mid h^t\right), \text{ for } i \neq j \neq k$$

- (P) For each player i, type θ_i, alternative strategy $\sigma'_{i,}$, and history h^t

$$u_i\left(\sigma \mid h^t, \theta_i, \mu(. \mid h^t)\right) \geq u_i\left(\sigma'_i, \sigma_{-i} \mid h^t, \theta_i, \mu(. \mid h^t)\right)$$

Sequential equilibrium

- We saw already that the requirement that the players' strategies form a Nash equilibrium is too weak → formally the only proper subgame for the games of incomplete, or imperfect information is the whole game.
- Recall the initial example

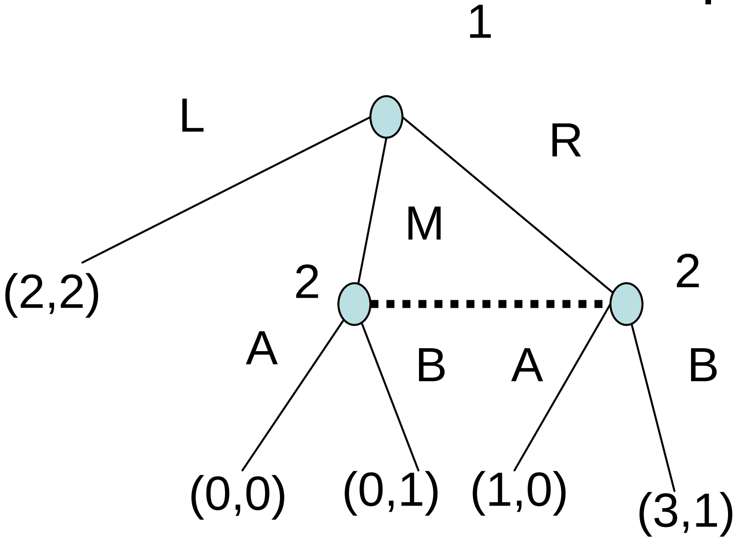

- Nash equilibria: (L,A) and (R,B)
- Both subgame perfect

Is (LA) equilibrium plausible?

Sequential equilibrium: cont.

- (L,A) is not plausible
  - Whatever player's 2 beliefs on player's 1 move (M or R), he must chose B if he has an opportunity to move.

- Need to generalize the previous condition (P) → given the system of beliefs, no player can gain by deviating at any information set.
- (s) An assessment (σ,μ) is *sequentially rational* if, for any information set h, and alternative strategy $\sigma'_{i(h)}$,

$$u_{i(h)}\left(\sigma \mid h, \mu(h)\right) \geq u_{i(h)}\left(\sigma'_{i(h)}, \sigma_{-i(h)} \mid h, \mu(h)\right)$$

Sequential eq. cont.

- Consistency condition on beliefs is also introduced

- (C) An assessment (σ,μ) is consistent if

$$(\sigma, \mu) = \lim_{n \to +\infty} (\sigma^n, \mu^n)$$

For some sequence $(\sigma^n, \mu^n) \in \Psi^0$ ← The set of all assessments

- **Definition:** A sequential equilibrium is an assessment (σ,μ) that satisfies (S) and (C)

Existence: For any finite extensive-form game, there exist at least one sequential equilibrium.

Learning in games

- Why learning?
  - For introspection, the rules of the game, rationality of the players, payoff functions – all common knowledge
  - Another problem: for multiple equilibria, how players come to expect the same equilibrium?

- Applicability
- Repeated games
- Teach opponent to play a best response to a particular action, by repeating it over and over again

Example of sophisticated learning

- How would you play this game, if you were player 1?

	L	R
U	1,0	3,2
D	2,1	4,0

Sophisticated learning?

- Most learning theory → models for which the incentive is small to alter the future play of the opponents.

  - Examples:

    - large anonymous population: population size large compared to the discount factor

    - Players locked in their choices and discount factor small compared to maximum speed at which the system can possibly adjust

Common models for learning

- **Fictitious play**
  - Players observe only their own matches and play a best response to the historical frequency of play
- **Partial best-response**
  - A fixed portion of users switches each period from its current action to a BR to the aggregate statistics from the previous period
- **Replicator Dynamics**
  - The fraction of the population using a given strategy, grows proportionally to that strategy's current payoff.

One type of learning: Cournot adjustment

- Unique Nash eq. is at the intersection of the reaction curves

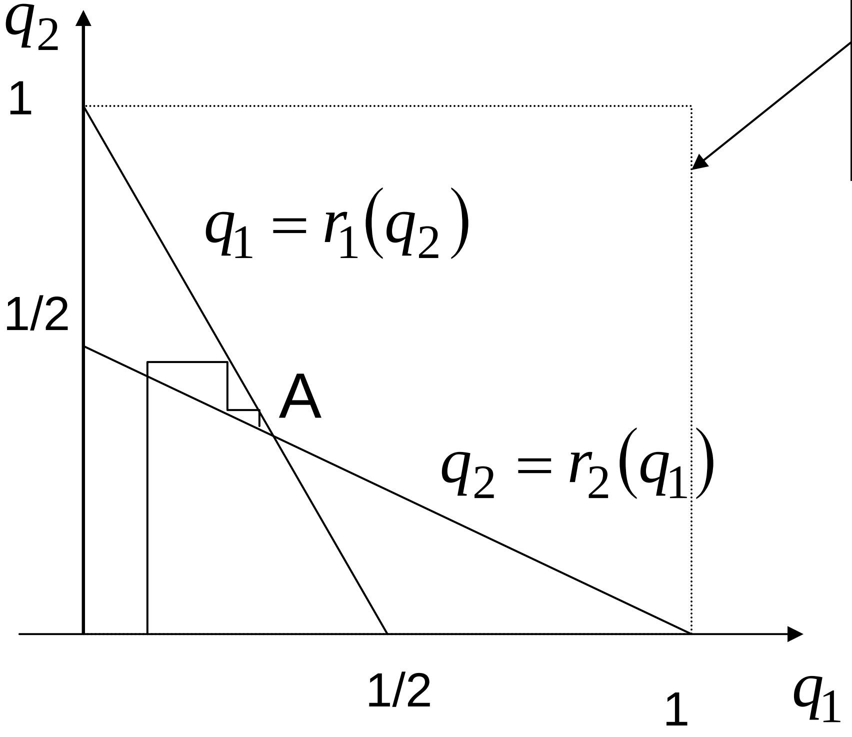

The process converges to Nash equilibrium from any Starting point →eq. globally stable

Fictitious play

- Repeated game
- Stationary assumption
- Each player: belief of opponents “strategy” by looking at what happened
- Player then plays best response (BR) according the his belief
- Belief: a prediction of the opponent action distribution, i.e. the degree to which player i believes player j will play a certain action.
- Players choose their actions in each period , s.t. to maximize their expected payoff, with respect to their belief for the current period.

Updating beliefs

- Player i: initial weight function

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

- 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