Game theory for wireless networks

Lecture 3

Dynamic games of complete information

- Extensive form games
  - In the examples studied so far, players choose their actions simultaneously
  - Can the strategic form game model the situation in which the order in which players move influences the outcome of the game?
  - Recall strategic (normal) form game characterized by three elements:
    - The set of players: {1, 2, …, I } (finite set)
    - The pure strategy space for each player i: S_i
    - Payoff (utility functions) for each profile of strategies: **s** = (s_1, …, s_I)

Example: Cournot vs. Stackelberg equilibrium

- Actions: choices of output levels: q1 and q2
- **Cournot:** both players choose their actions simultaneously, i.e., try to simultaneously maximize their utility functions
  - Example:

$$u_i(q_1,q_2)=[12-(q_1+q_2)]q_i \Rightarrow \begin{cases} \dfrac{\partial u_1(q_1,q_2)}{\partial q_1}=0 \Rightarrow r_1(q_2)=6-\dfrac{q_2}{2} \\ \dfrac{\partial u_2(q_1,q_2)}{\partial q_2}=0 \Rightarrow r_2(q_1)=6-\dfrac{q_1}{2} \end{cases}$$

- **Stackelberg:** player 1 chooses first, then player 2 observes the output q1, and consequently chooses q2
  - Is it the same equilibrium?
  - For which player this game is more advantageous?

Stackelberg equilibrium: cont.

- Player 2 sees q1, computes r2(q1) in the same fashion as before
- Player 1: knows that player 2 will maximize its utility based on q1, can compute r2(q1), and then maximize its utility by appropriately selecting q1.

$$u_i(q_1,q_2)=[12-(q_1+q_2)]q_i \Rightarrow \begin{cases} \dfrac{\partial u_2(q_1,q_2)}{\partial q_2}=0 \Rightarrow r_2(q_1)=6-\dfrac{q_1}{2} \\ u_1(q_1,q_2)=q_1\left(6-\dfrac{q_1}{2}\right) \\ \dfrac{\partial u_1(q_1)}{\partial q_1}=0 \Rightarrow q_1^*=6 \end{cases}$$

Stackelberg equilibrium: cont

- The resulting equilibrium point: $q^*_1=6$, $q^*_2=3$, with payoffs (18,9)
- Obtained by backward induction
- Cournot equilibrium: $q^C_1=4$, $q^C_2=4$, with payoffs (16,16)

- Leader has the advantage
- Other possible equilibria?
- Maybe, but not credible: would rely on empty threats from player 2, to maintain a different level q2

Extensive form games

- Stackelberg game (leader-follower) – example of a game in which players move sequentially and the order of the players' moves matters
- Multi-stage game
- Game of perfect information: exactly one player moves at a given stage, all the others have the one element choice: "do nothing"
- What is an extensive form game?

Extensive form games

The extensive form of a game contains the following inf:

- The set of players $i \in \mathbf{I}$
- The order of moves: Game tree
- The players' payoffs as a function of the previous moves
- What are the players choices when they move
- What each player knows when he makes its choice
- The probability distribution over any exogenous events

Exogenous events: moves by nature

Characterizing previous moves: definitions and notations

- Multi-stage game:
  - Players move simultaneously at stage k (do not know the actions of their opponents for stage k)
  - Know all the actions chosen at previous stages: 0,1,2,…, k-1.
  - Particular case: Stackelberg example (2 stage game) – at one stage just one player moves, the other one has action "do nothing".
- At stage k, i-th player chooses an action from the choice set $A_i\left(\mathbf{h}^k\right)$

$\mathbf{h}^k = \left(\mathbf{a}^0, \mathbf{a}^1, \ldots, \mathbf{a}^k\right)$ = the history at the end of stage k

$\mathbf{a}^k = \left(a_1^k, a_2^k, \ldots, a_I^k\right)$ = stage k strategy profile

Game begins at stage 0, with $\mathbf{h}^0 = \varnothing$

More notations

- H^k = the set of all stage k histories
- Z = the set of terminal histories

- Player's i payoff represented as

$$u_i : H^K \to R$$

- For many applications: the payoff is additive over the game stages (weighted sum)

- Single stage payoffs for user i at stage k: $g_i\left(a^k\right)$

Game tree representation

- Components of a game tree
  - Vertices (nodes) – represent a particular history
    - Usually represented as 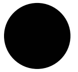
    - Exception: node with no history (first stage) 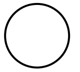
    - Labeled with the user id
  - Edges: correspond to the actions taken
    - Labeled with the action element
  - Simultaneous moves can also be modeled
  - When a node uncertain of past history (particular case is simultaneous moves, a dashed line unites its vertices for that stage. Sometimes this is represented also by encircling the vertices with an ellipse

Game tree representation – cont.

- Time progresses in one direction: typically from left to right, or top to bottom.

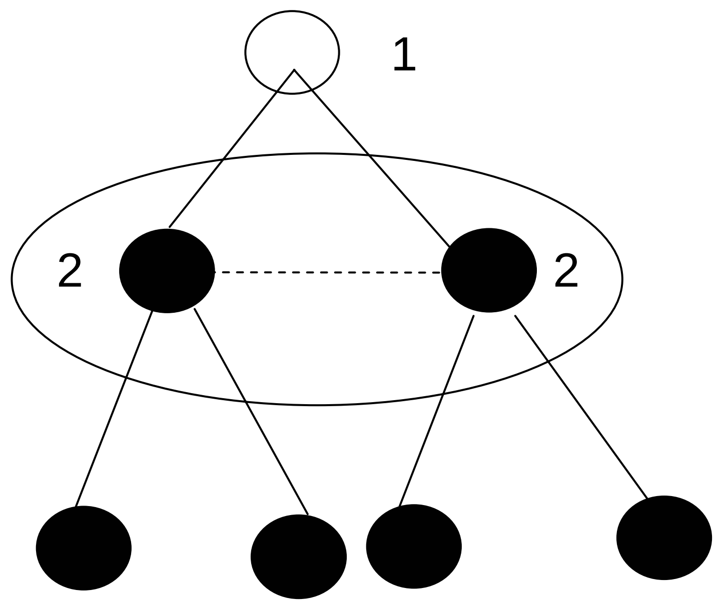

Stackelberg game representation

- Simplifying assumption: each player has only three possible output levels: 3, 4, 6

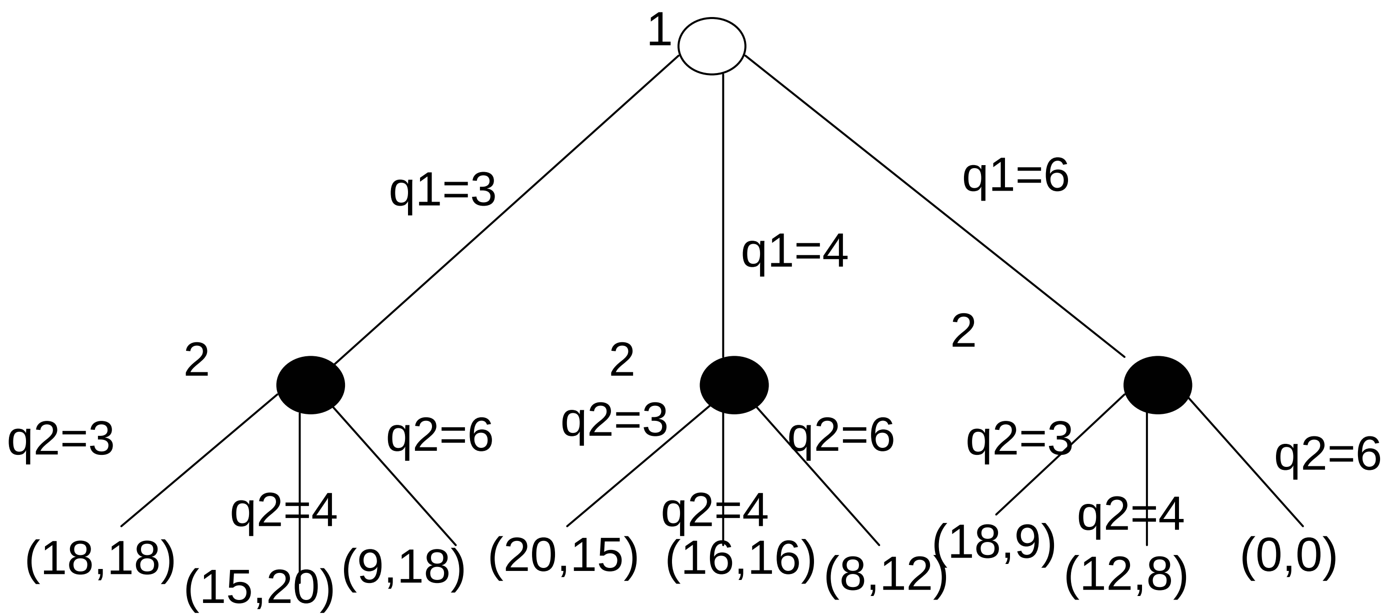

Cournot equivalent

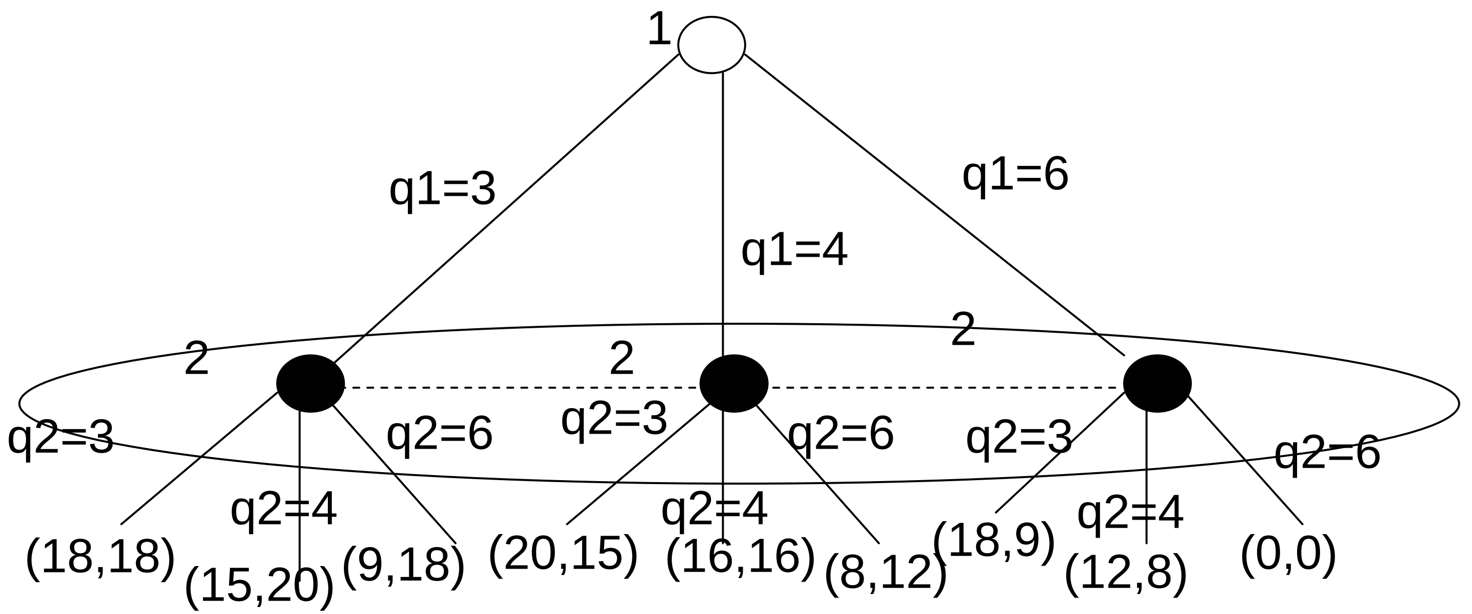

- Is this representation unique?

Strategies and equilibria for extensive form games

- H_i = player i information sets
- $A_i = \bigcup_{h_i \in H_i} A(h_i)$ = set of all actions for player i
- A pure strategy is a map

$$s_i : H_i \to A_i, \quad s_i(h_i) \in A(h_i), \quad \forall h_i \in H_i$$

- Player i pure strategy space, is the space of all si
- The number of player's i pure strategy:

$$\# S_i = \prod_{h_i \in H_i} \# A(h_i)$$

- Path of s = information sets that are reached with positive probability
- Pure strategy Nash equilibrium: a strategy profile s*, such that each player's i strategy (s_i*) maximizes his expected payoff, given the strategies of his opponents (s_{-i}*)

Stackelberg game example

- How many pure strategies for players?

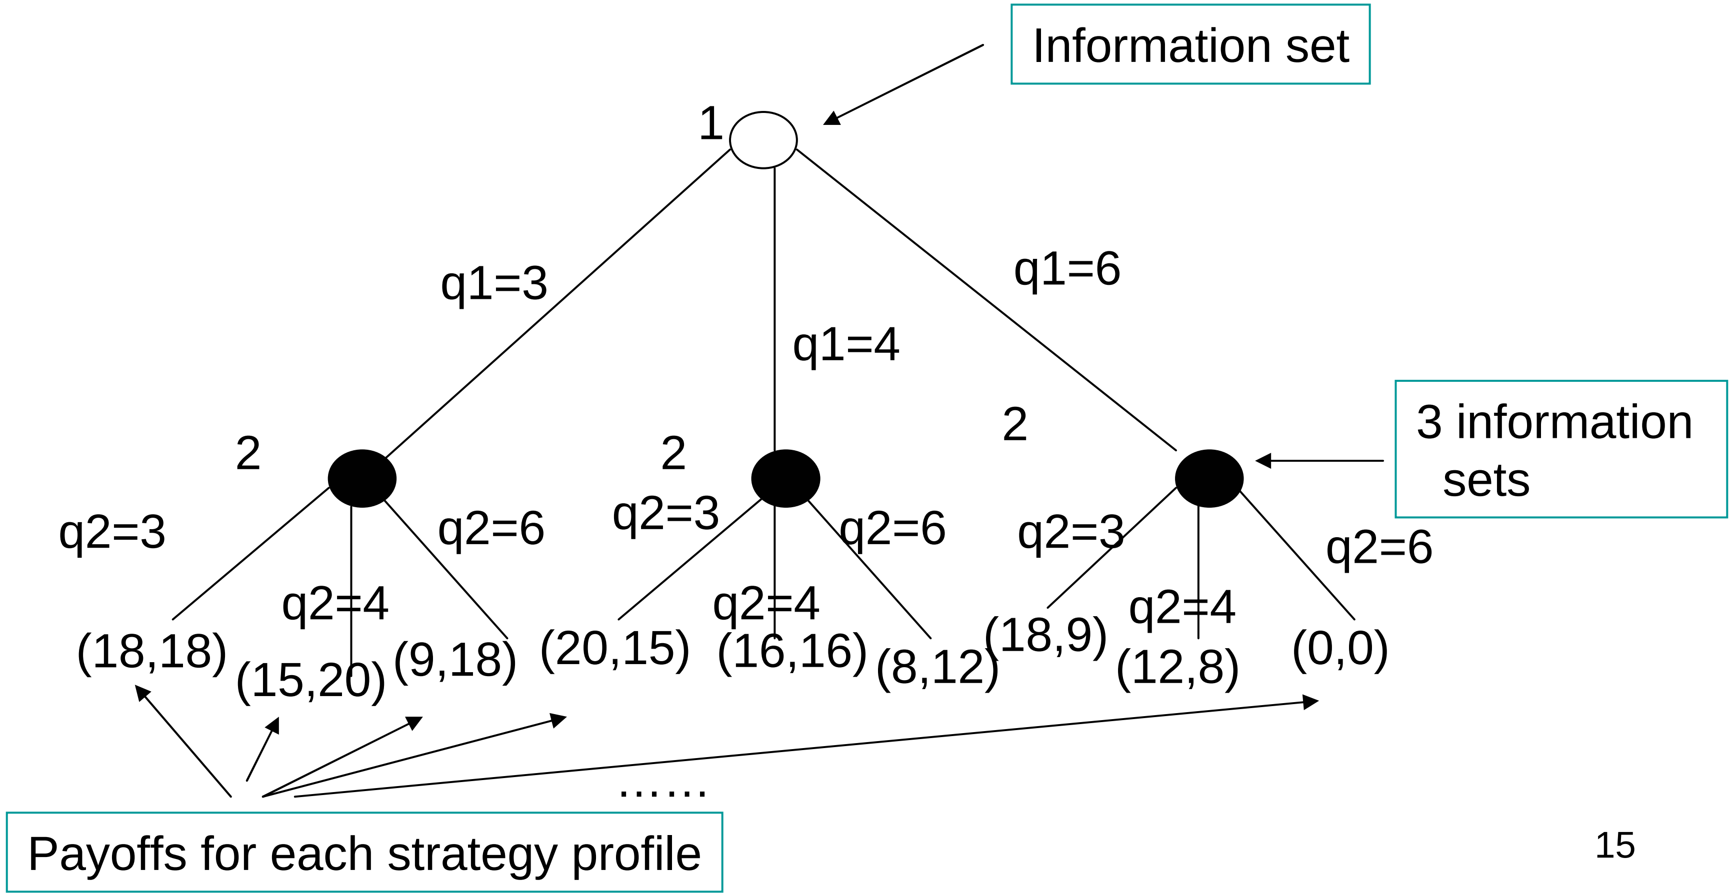

Mixed strategies

- Behavior strategies
- A behavior strategy specifies a probability distribution over actions at each h_i, and the probability distributions at different information sets are independent.

- Nash eq. in behavior strategies = profile such that no player can increase its expected payoff by using a different behavior strategy.
- Mixed strategies and behavior strategies are equivalent for games of perfect recall
  - No player ever forgets any information he knew
  - Library analogy for mixed and behavior strategies (ex. In book)

Equivalence with games in strategic form

- Basic idea: represent every strategy and their corresponding payoff:
- Consider an example game
  - Player 1: {A,B}
  - Player 2: {(C,C’),(C,D’),(D,C’),(D,D’)}

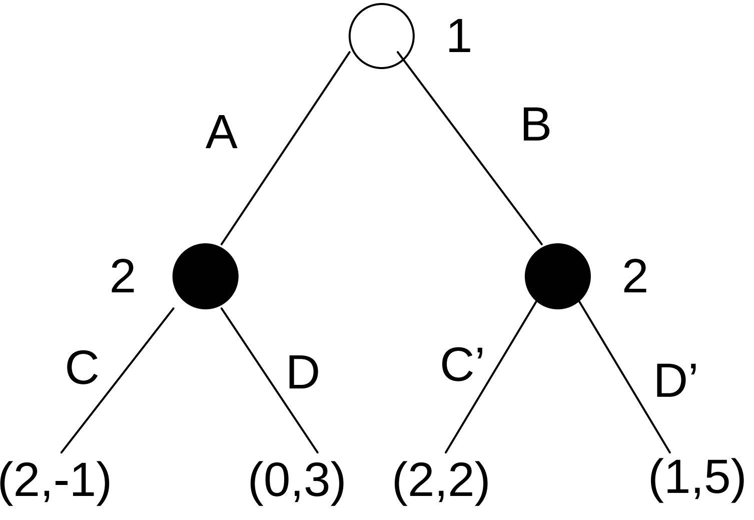

	C,C’	C,D’	D,C’	D,D’
A	2,-1	2,-1	0,3	0,3
B	2,2	1,5	2,2	1,5

Another example

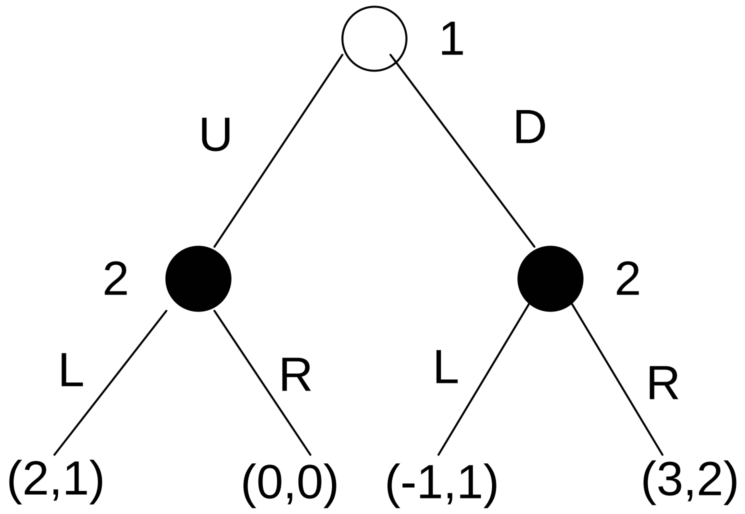

Find strategic form equivalence
Find Nash equilibria

Nash equilibria

- If the extensive form is finite → the corresponding strategic form is finite → Nash theorem guarantees the existence of a mixed-strategy equilibrium
- Iterative strict dominance notion extends for extensive form games as well
  - Weaker notion: a player cannot strictly prefer one action over another at an information set that is not reached, given its opponent's play
- Theorem (Zermelo, Kuhn): **A finite game of perfect information has a pure strategy Nash equilibrium.**
  - Proof based on many player generalization of backward induction in dynamic programming
  - Idea: the game finite → has a set of penultimate nodes → players moving at these nodes chooses strateg that leads to terminal node with max. payoff. Players that have successors the penultimate nodes, choose actions that max. their payoffs, given the choice of the penultimate nodes. Etc… Roll back to the three → resulting strategy is a Nash eq.

Nash equilibria comments

- Previous theorem may not hold if the hypothesis are weakened
  - Infinite games
    - Node with an infinite number of successors: continuum of actions
    - Path with an infinite number of nodes: multi-stage games with an infinite number of stages
  - Not perfect information (i.e., some of the information sets are not singletons)

Backward induction and subgame perfection

- Find Nash eq. by backward induction:
  - Reason backwards on what each rational player would play
  - Assumption: Starting at any decision point in the game, a player’s strategy (from that point on) is a best response to the strategies of other players → Sequential Rationality
  - Subgame perfect Nash equilibrium → key concept for the backward induction technique

Subgame

- Definition: A proper subgame G of an extensive form game T, consists of a single node and all of its successors in T, with the property that if $x' \in G$, and $x'' \in h(x')$, then $x'' \in G$. The information sets and payoffs are inherited from the original game.

Simplified: A subgame is determined by assuming that h has already happened, and selecting the game from that point onward.

Subgame example

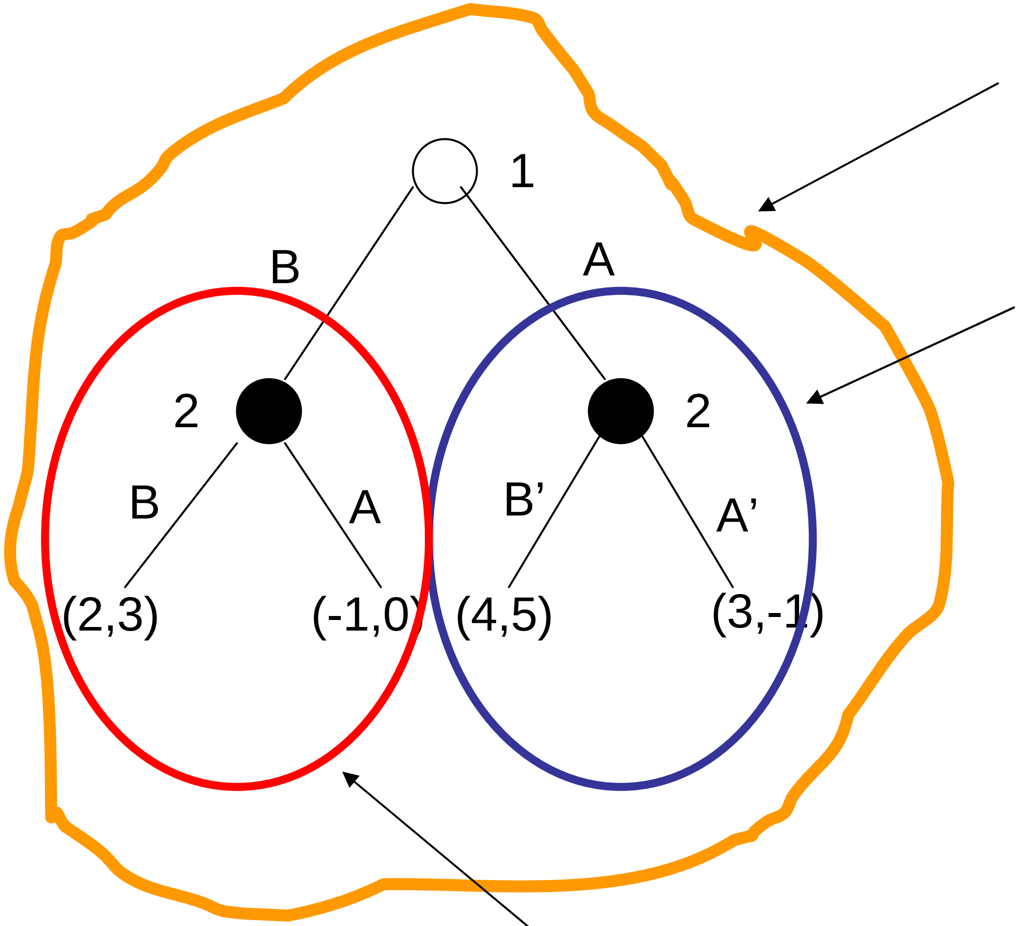

Original game is its own subgame

Subgame given A

Subgame given B

Another subgame example

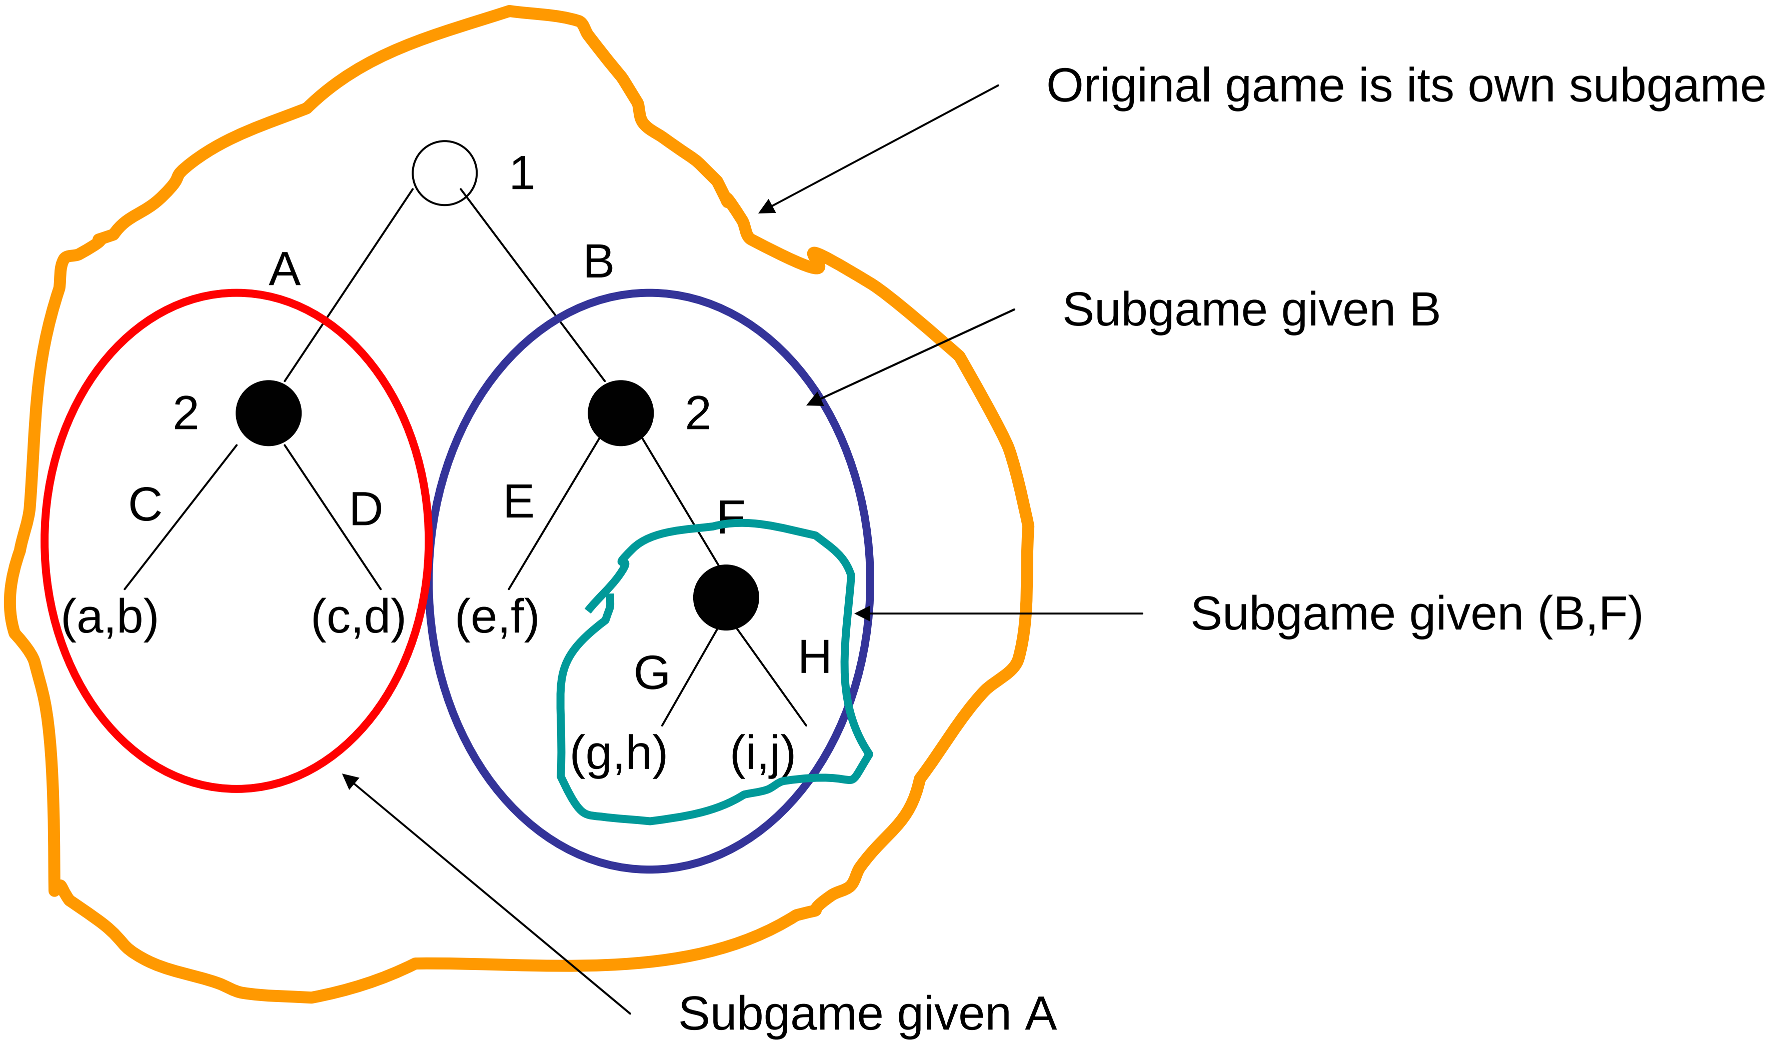

Subgame perfect equilibrium

- **Definition:** A behavior-strategy profile σ of an extensive-form game is a subgame-perfect equilibrium, if the restriction of σ to G is a Nash equilibrium of G for every proper subgame G.

Backward induction algorithm

- Identify all terminal subgames
- Determine the Nash eq. for these subgames
- Modify the original game tree by replacing the terminal subgames with the Nash equilibrium payoffs
- Repeat until the three is reduced to one stage game, and then determine the Nash equilibrium.

Example backward induction algorithm

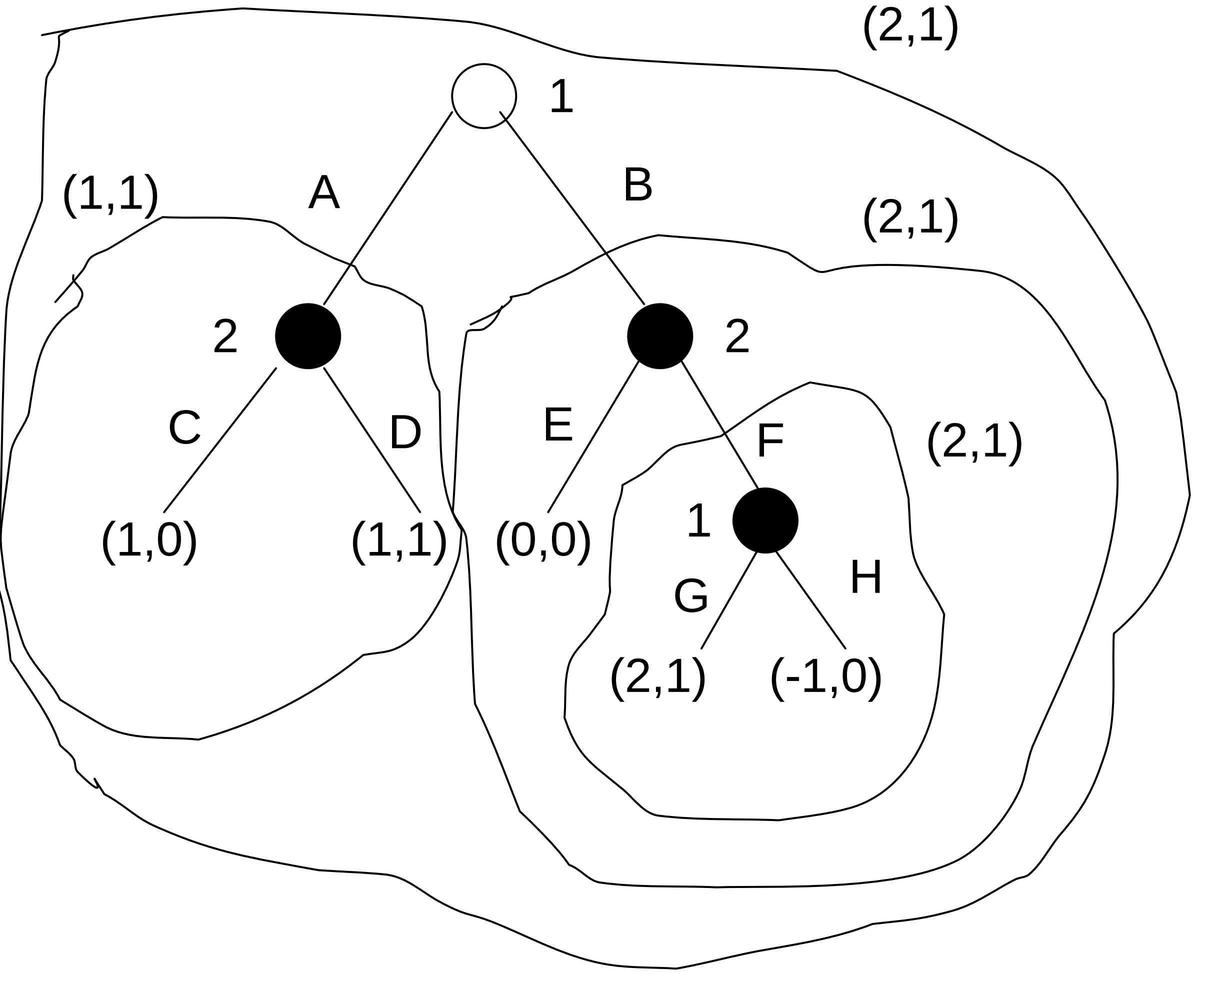

{B,F,G} = Nash equilibrium

Payoff is (2,1)

Subgame perfection and one stage deviation principle

- **Finite horizon games:** In a finite multistage game with observed actions, a strategy profile s is subgame perfect, iff no player can gain by deviating from s in a single stage and then conforming to s thereafter.
- **Infinite horizon games:** In an infinite multistage game with observed actions that is continuous at infinity (events in distant future are relatively unimportant), profile s is subgame perfect, iff there is no player i and strategy $s\hat{}_l$ that agrees with si except at a single t and h^t, such that $s\hat{}_l$ is a better response to s_{-l} , than si, conditioned on history h^t being reached.

Repeated games

- Introduces new equilibria: players may condition their actions on the way their opponents play in previous periods.
- Example: prisoner's dilemma
- Payoffs depend only on current actions ($g_i(a^t)$ shown in matrix)
- Players discount future payoffs with a common discount factor δ
- Questions: how the eq. payoffs vary with the horizon T, and the discount factor ?

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

Repeated games-cont.

- Cummulative payoff
- Normalized – to make it comparable for different time horizons T
- The utility of a sequence

$$\left\{a^0, a^1, \ldots, a^T\right\}$$

$$\frac{1-\delta}{1-\delta^{T+1}} \sum_{t-0}^{T} \delta^t g_i\left(a^t\right) \longleftarrow \text{Average discounted payoff}$$

Prisoners' dilemma example

- Game played only once: equilibrium - both defect
- Game repeated a finite # of times
  - Both defect remains a sub-game perfect equilibrium
  - If horizon is infinite, and $\delta > ½$ → new subgame perfect eq.:
    - "Cooperate in the first period and continue to cooperate as long as no player has ever defected. If any player has ever defected, then defect for the rest of the game".
  - Two classes of subgames: A – no player has defected
  - B – defect i has occurred
- If player i conforms to A for all stages of the game : payoff is 1
- If deviates at time t: its normalized payoff is

$$(1-\delta)(1+\delta+...+\delta^{t-1}+2\delta^{t}+0+0+...)=1-\delta^{t}(2\delta-1)$$

New equilibrium for Prisoners' dilemma

- In every subgame, no player can gain by deviating once from the specified strategy and then conforming → one stage deviation principle holds → the strategies form a subgame perfect eq.
- Depending on the size of the discount factor, there can exist many other perfect equilibria
- Next time: we will discuss folk theorem
  - Repeated play with patient players not only makes cooperation possible (more efficient payoffs) but it leads to a large set of other equilibrium outcomes.

Homework –study this problem + new hw assignment

- Consider a duopoly game between firms A and B, in which they both have the choice of selecting a high or a low price for the good they produce. They make the choice of price simultaneously every year, and the game is played repeatedly, year after year. If the firms both select the high price, they both earn $80.000 per year. If one sets a high price, and the other sets a low price, the low price firm earns $100.000, and the high price firm earns $30.000. If they both set a low price, they will both earn $50.000.
- 1) Construct the payoff matrix of the one shot game and find the Nash eq. for this case.
- 2) Assume that the firms decide to play this game for a period of 3 years. Determine the game tree. What are the total equilibrium payoffs for this scenario?
- 3) Assume now that the firms are going to play this game forever, and also assume that an inflation rate factor of 25% apply (a discount factor $\delta = 0.8$). If both firms play the "Grim-Trigger" strategy (both price high, until one of them defects and they both price low for the rest of the game). Will the firm have incentives to cooperate for this scenario?
- 4) Suppose that the firms play the above game repeatedly, but unexpectedly, the market for their product disappears at the end of year 3. What is their profit at the end of year 3? Compare with the scenario in 2)

- Find solution at http://www.ece.stevens-tech.edu/~ccomanic/ee800c.html (Homework 2 solution)