Performance Evaluation

Second Part
Lecture 6

Giovanni Neglia
INRIA – EPI Maestro
13 February 2012

Game Trees (Extensive form)

- Sequential play
  - players take turns in making choices
  - previous choices may be available to players
- Game represented as a tree
  - each non-leaf node represents a decision point for some player
  - edges represent available choices

Game Trees: simplified poker

- Rose and Colin put 1$ each in the pot and take a card (Ace or King)
- Colin may bet other 2$ or drop
- If Colin bets
  - Rose can put other 2$ and call (and the highest card wins)
  - or can fold (and Colin takes the money)
- If Colin drops
  - Rose takes all the money in the pot

Tree of the simplified poker

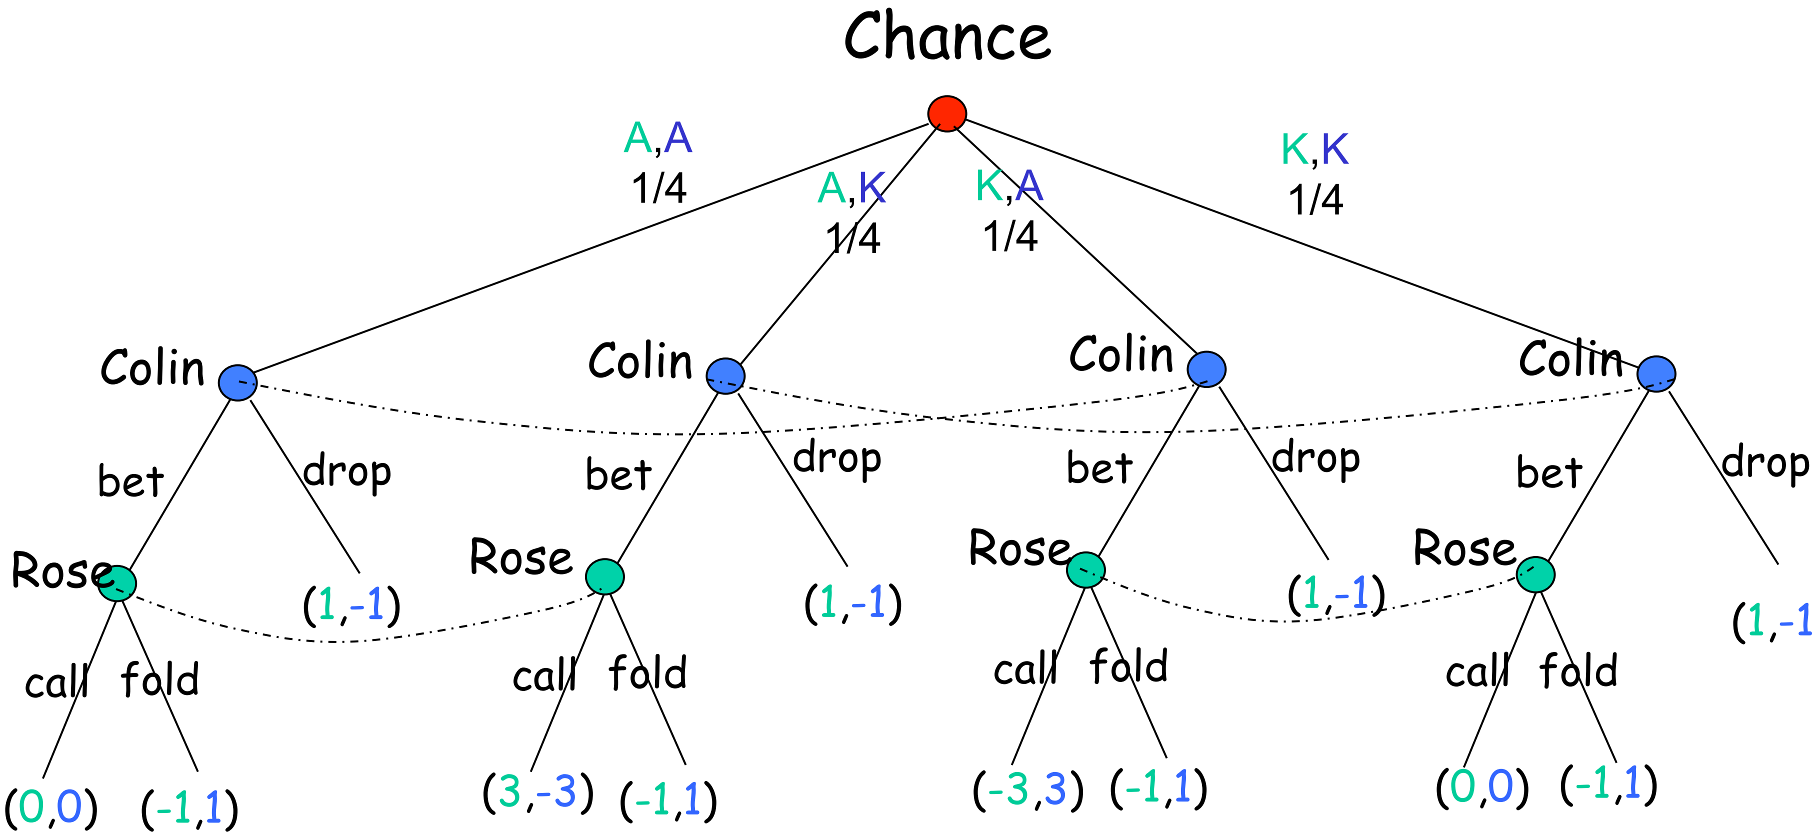

Arc joins states of a player in the same *information set*:

- when playing the player cannot distinguish these states
  - –the known sequence of past events is the same
  - –the set of future actions is the same

Game trees: more formal definition

1. each node is labeled by the player (including Chance) who makes a choice at that node
2. each branch leading by a node corresponds to a possible choice of the player at the node
3. each branch corresponding to a choice made by Chance is labeled with the corresponding probability
4. each leaf is labeled by players payoffs
5. nodes of each player are partitioned in information sets

Game trees and matrix games

- Each game tree can be converted in a matrix game!
- Connecting idea: strategy in game tree
  - it specifies a priori all the choices of the player in each situation
    - only need to specify for each information set
  - e.g. in simplified poker
    - for Colin 4 possible strategies
      - "always bet" (bb), "bet only if ace" (bd), "bet only if king" (db), "always drop" (dd)
    - for Rose 4 possible strategies
      - "always call" (cc), "call only if ace" (cf), "call only if king" (fc), "always fold" (ff)

Game trees and matrix games

- Each game tree can be converted in a matrix game!
- Once identified the strategies of every player…
- …use the expected payoffs of the game tree as payoffs of the matrix game

Game trees and matrix games

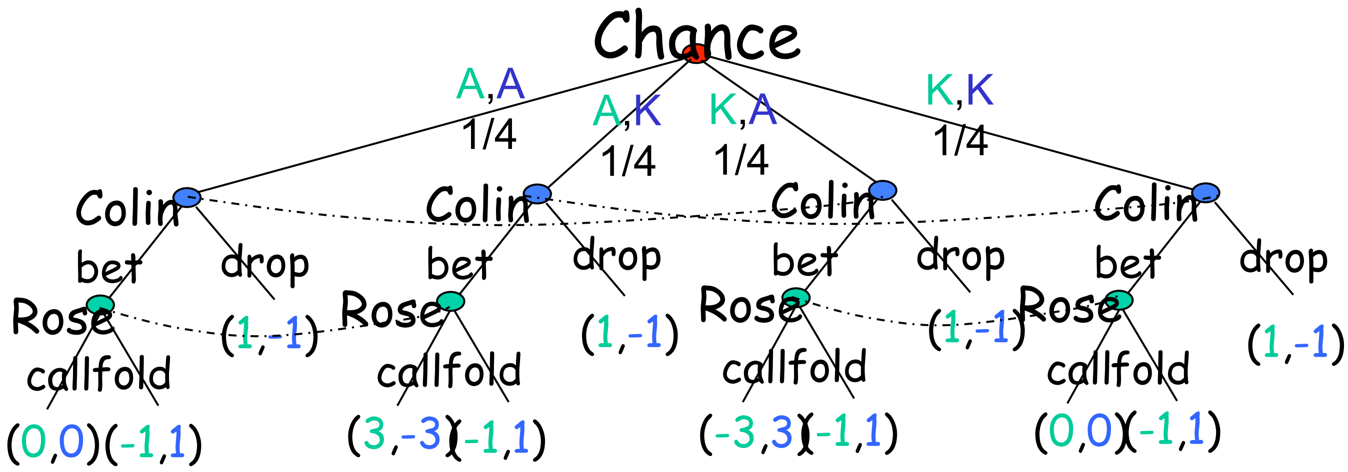

Study this game

Game trees and matrix games

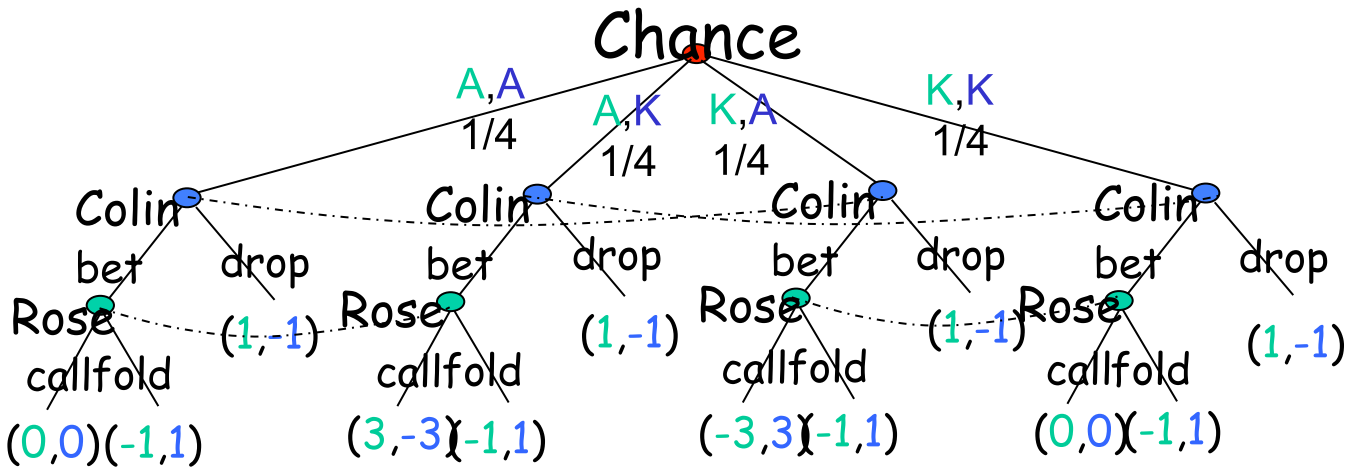

Colin

Rose

	bb	bd	db	dd
cc	0	-1/4	5/4	1
cf	1/4	1/4	1	1
fc	-5/4	-1/2	1/4	1
ff	-1	0	0	1

Study this game

Game trees and matrix games

- Each game tree can be converted in a matrix game!
- Problem: this approach does not scale with the size of the tree
  - exponential growth in the number of strategies
    - consider how many strategies are available in chess to White and to Black for their respective first move
- Try to study directly the game tree

Game trees with perfect information

- Definition
  1. no nodes are labeled by Chance
  2. all information sets consist of a single node
- Test: which among the following is a game with perfect information and why?
  - poker
  - tic, tac, toe
  - rock, scissor, paper
    - honestly and dishonestly played…
  - chess
  - guess the number

Perfect information: an example

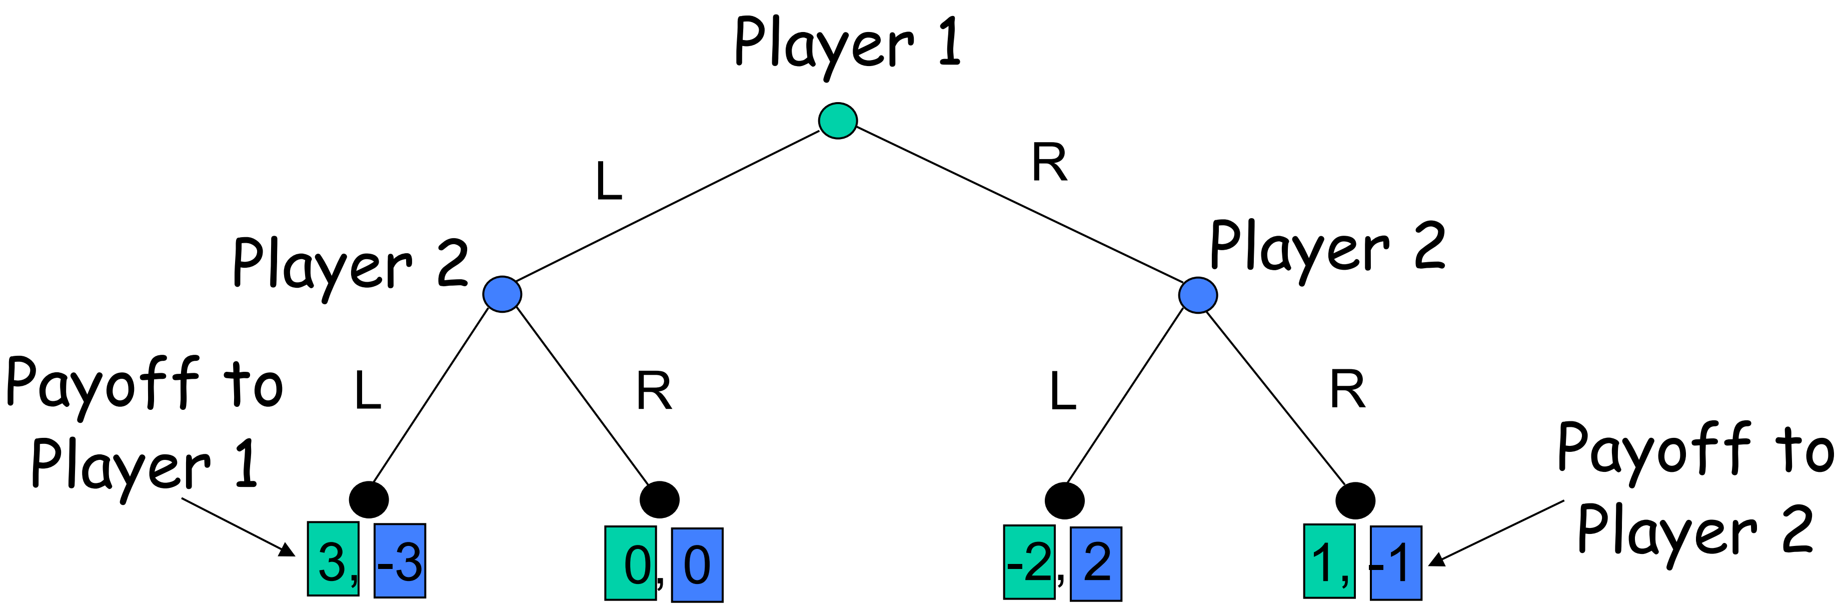

- Strategy sets
  - for Player 1: {L, R}
  - for Player 2: {LL,LR,RL,RR}
- Convert it to a matrix game and solve it

Converting to Matrix Game

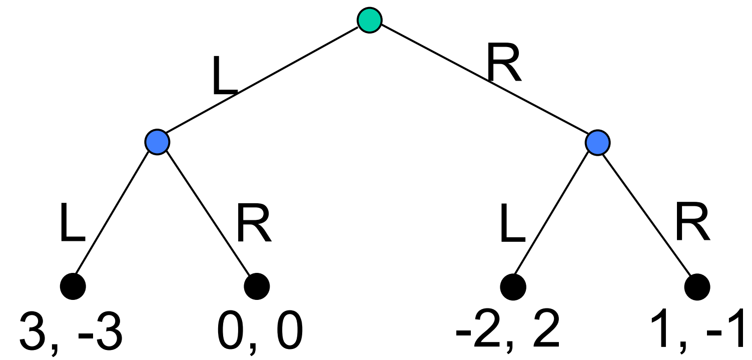

Player 2

Player 1		LL	LR	RL	RR
	L	3	3	0	0
	R	-2	1	-2	1

Solving the game by backward induction

- Starting from terminal nodes
  - move up game tree making best choice

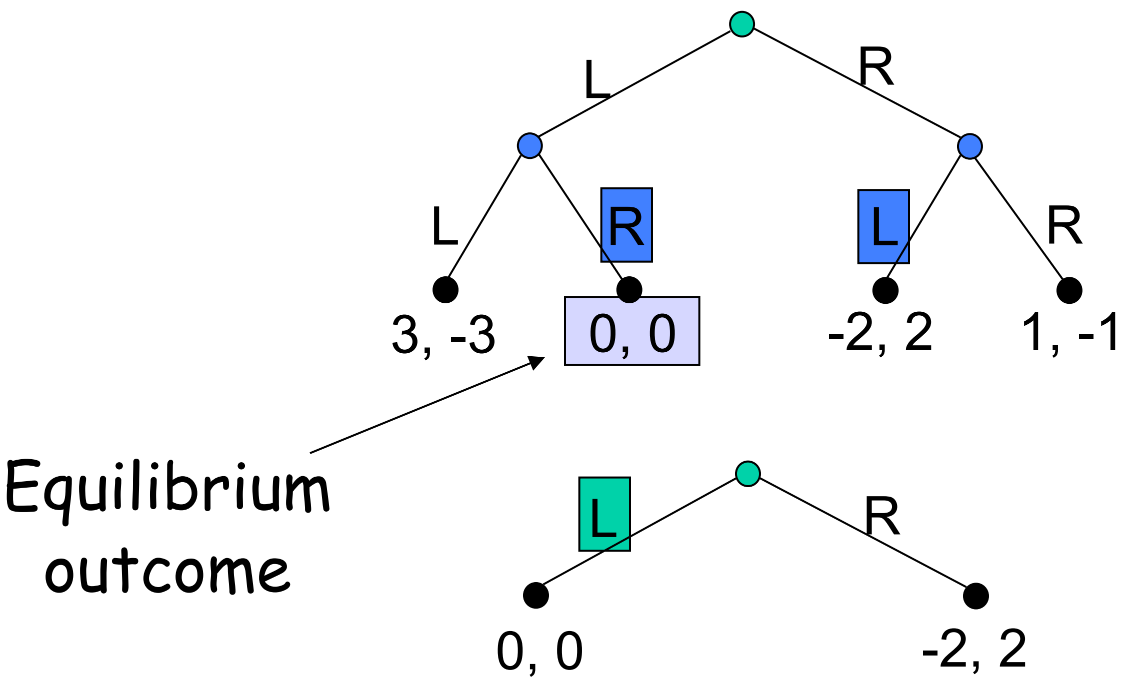

Best strategy for P2: RL

Best strategy for P1: L

- Saddle point:
  P1 chooses L, P2 chooses RL

Kuhn's Theorem

- Backward induction always leads to saddle point (on games with perfect information)
  - game value at equilibrium is unique (for zero-sum)
- Consequences for chess?
  - at the saddle point
    - or White wins, value = 1 -> White has winning strategy no matter what Black does
    - or Black wins, value = -1 -> Black has winning strategy, no matter what White does
    - or they draw, value = 0 -> Both White and Black have a strategy guaranteeing at least drawing

Chess is a *simple* game! (Zermelo 1913)

More on Game Trees

- We will talk more on about
  - games with imperfect information
  - and mixed strategies
- when presenting repeated games (a special case of game trees).

Game Theory: introduction and applications to computer networks

Two-person non zero-sum games

Giovanni Neglia
INRIA – EPI Maestro

Slides are based on a previous course
with D. Figueiredo (UFRJ) and H. Zhang (Suffolk University)

Outline

- Two-person zero-sum games
  - Matrix games
    - Pure strategy equilibria (dominance and saddle points), ch 2
    - Mixed strategy equilibria, ch 3
  - Game trees, ch 7
- Two-person non-zero-sum games
  - Nash equilibria…
    - …And its limits (equivalence, interchangeability, Prisoner's dilemma), ch. 11 and 12
  - Strategic games, ch. 14
  - Subgame Perfect Nash Equilibria (not in the book)
  - Repeated Games, partially in ch. 12
  - Evolutionary games, ch. 15
- N-persons games

Two-person Non-zero Sum Games

- Players are not strictly opposed
  - payoff sum is non-zero

Player 2

Player 1		A	B
	A	3, 4	2, 0
	B	5, 1	-1, 2

- Situations where interest is not directly opposed
  - players could cooperate
  - communication may play an important role
    - for the moment assume no communication is possible

What do we keep from zero-sum games?

- Dominance
- Movement diagram
  - pay attention to which payoffs have to be considered to decide movements

		Player 2	
		A	B
Player 1	A	5, 4	2, 0
	B	3, 1	-1, 2

- Enough to determine pure strategies equilibria
  - but still there are some differences (see after)

What can we keep from zero-sum games?

- As in zero-sum games, pure strategies equilibria do not always exist…

		Player 2	
		A	B
Player 1	A	5, 0	-1, 4
	B	3, 2	2, 1

- …but we can find mixed strategies equilibria

Mixed strategies equilibria

- Same idea of equilibrium
  - each player plays a mixed strategy (*equalizing strategy*), that equalizes the opponent payoffs
  - how to calculate it?

		Colin A	Colin B
Rose	A	5, 0	-1, 4
	B	3, 2	2, 1

Mixed strategies equilibria

- Same idea of equilibrium
  - each player plays a mixed strategy, that equalizes the opponent payoffs
  - how to calculate it?

		Colin	
		A	B
Rose	A	-0	-4
	B	-2	-1

Rose considers *Colin's game*

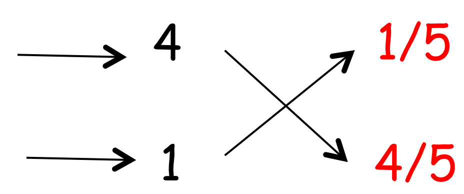

Mixed strategies equilibria

- Same idea of equilibrium
  - each player plays a mixed strategy, that equalizes the opponent payoffs
  - how to calculate it?

		Colin A	Colin B
Rose	A	5	-1
	B	3	2
		3/5	2/5

Colin considers *Rose's game*

Mixed strategies equilibria

- Same idea of equilibrium
  - each player plays a mixed strategy, that equalizes the opponent payoffs
  - how to calculate it?

		Colin A	Colin B
Rose	A	5, 0	-1, 4
	B	3, 2	2, 1

Rose playing (1/5,4/5)
Colin playing (3/5,2/5)
is an equilibrium

Rose gains 13/5
Colin gains 8/5

Good news: Nash's theorem [1950]

- Every two-person games has at least one equilibrium either in pure strategies or in mixed strategies
  - Proved using fixed point theorem
  - generalized to N person game
- This equilibrium concept called Nash equilibrium in his honor
  - A vector of strategies (a profile) is a Nash Equilibrium (NE) if no player can unilaterally change its strategy and increase its payoff

A useful property

- Given a finite game, a profile is a mixed NE of the game if and only if for every player i, every pure strategy used by i with non-null probability is a best response to other players mixed strategies in the profile
  - see Osborne and Rubinstein, *A course in game theory*, Lemma 33.2

Bad news: what do we lose?

- equivalence
- interchangeability
- identity of equalizing strategies with prudential strategies
- main cause
  - at equilibrium every player is considering the opponent's payoffs ignoring its payoffs.
- New problematic aspect
  - group rationality versus individual rationality (cooperation versus competition)
  - absent in zero-sum games
- we lose the idea of **the** solution

Game of Chicken

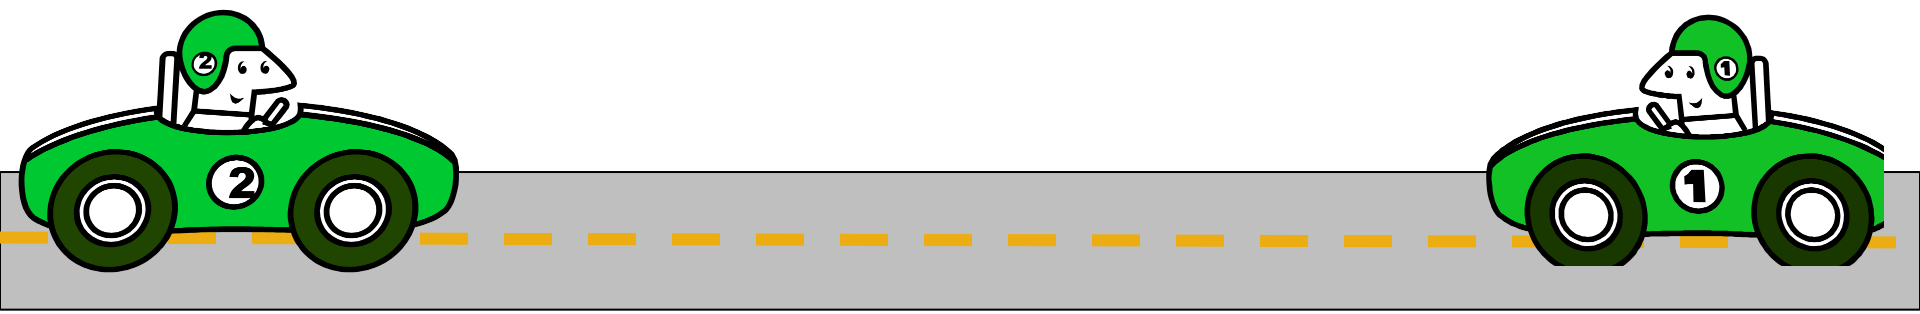

- Game of Chicken (aka. Hawk-Dove Game)
  - driver who swerves looses

		Driver 2	
		swerve	stay
Driver 1	swerve	0, 0	-1, 5
	stay	5, -1	-10, -10

Drivers want to do opposite of one another

Two equilibria:
not equivalent
not interchangeable!

- playing an equilibrium strategy does not lead to equilibrium

The Prisoner's Dilemma

- One of the most studied and used games
  - proposed in 1950
- Two suspects arrested for joint crime
  - each suspect when interrogated separately, has option to confess

		Suspect 2	
		NC	C
Suspect 1	NC	2, 2	10, 1
	C	1, 10	5, 5

payoff is years in jail
(**smaller is better**)

better outcome

single NE

Pareto Optimal

		Suspect 2 NC	C
Suspect 1	NC	2, 2	10, 1
	C	1, 10	5, 5

Pareto Optimal (highlighted cells: 2, 2; 10, 1; 1, 10)

- Def: outcome o* is Pareto Optimal if no other outcome would give to all the players a payoff not smaller and a payoff higher to at least one of them
- Pareto Principle: to be acceptable as a solution of a game, an outcome should be Pareto Optimal
  - the NE of the Prisoner's dilemma is not!
- Conflict between group rationality (Pareto principle) and individual rationality (dominance principle)

Payoff polygon

		Colin	
		A	B
Rose	A	5, 0	-1, 4
	B	3, 2	2, 1

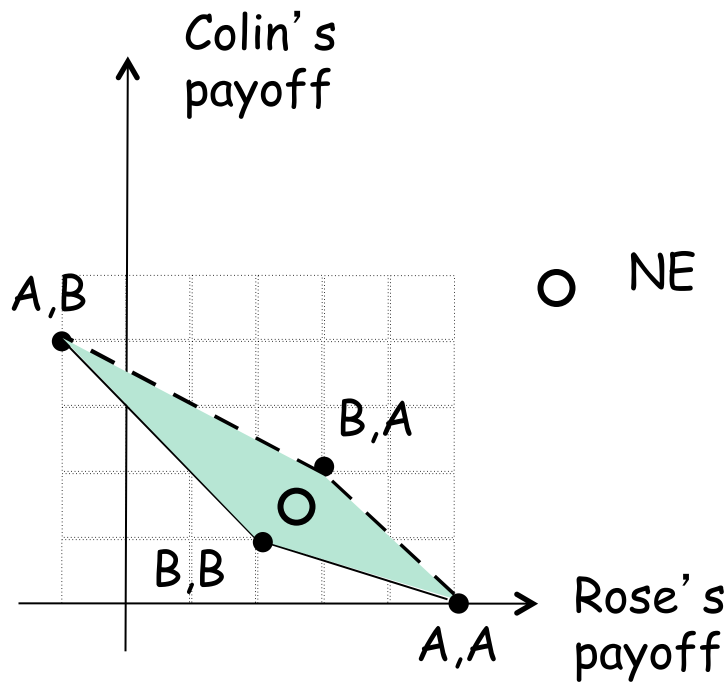

- All the points in the convex hull of the pure strategy payoffs correspond to payoffs obtainable by mixed strategies
- The north-east boundary contains the Pareto optimal points

Another possible approach to equilibria

- NE ⇔ equalizing strategies
- What about prudential strategies?

Prudential strategies

- Each player tries to minimize its maximum loss (then it plays in its own game)

		Colin	
		A	B
Rose	A	5, 0	-1, 4
	B	3, 2	2, 1

Prudential strategies

- Rose assumes that Colin would like to minimize her gain
- Rose plays in Rose's game
- Saddle point in BB
- B is Rose's prudential strategy and guarantees to Rose at least 2 (Rose's *security level*)

		Colin	
		A	B
Rose	A	5	-1
	B	3	2

Prudential strategies

- Colin assumes that Rose would like to minimize his gain (maximize his loss)
- Colin plays in Colin's game
- mixed strategy equilibrium,
- (3/5,2/5) is Colin's prudential strategy and guarantees Colin a gain not smaller than 8/5

		Colin A	Colin B
Rose	A	0	-4
	B	-2	-1

Prudential strategies

- Prudential strategies
  - Rose plays B, Colin plays A w. prob. 3/5, B w. 2/5
  - Rose gains 13/5 (>2), Colin gains 8/5
- Is it stable?
  - No, if Colin thinks that Rose plays B, he would be better off by playing A (Colin's *counter-prudential strategy*)

		Colin A	Colin B
Rose	A	5, 0	-1, 4
	B	3, 2	2, 1

Prudential strategies

- are not the solution neither:
  - do not lead to equilibria
  - do not solve the group rationality versus individual rationality conflict
- dual basic problem:
  - look at your payoff, ignoring the payoffs of the opponents

Exercises

- Find NE and Pareto optimal outcomes:

	NC	C
NC	2, 2	10, 1
C	1, 10	5, 5

	A	B
A	2, 3	3, 2
B	1, 0	0, 1

	swerve	stay
swerve	0, 0	-1, 5
stay	5, -1	-10, -10

	A	B
A	2, 4	1, 0
B	3, 1	0, 4

Game Trees Revisited

- Microsoft and Mozilla are deciding on adopting new browser technology (.net or java)
  - Microsoft moves first, then Mozilla makes its move

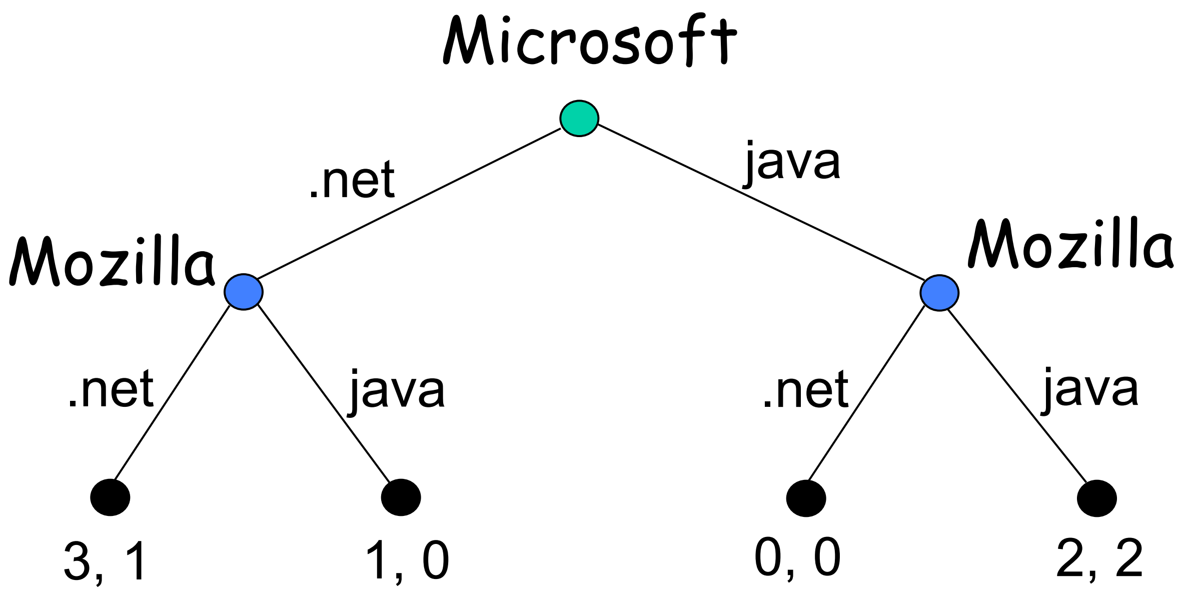

- Non-zero sum game
  - what are the NEs?
  - remember: a (pure) strategy has to specify the action at each information set

NE and Threats

- Convert the game to normal form

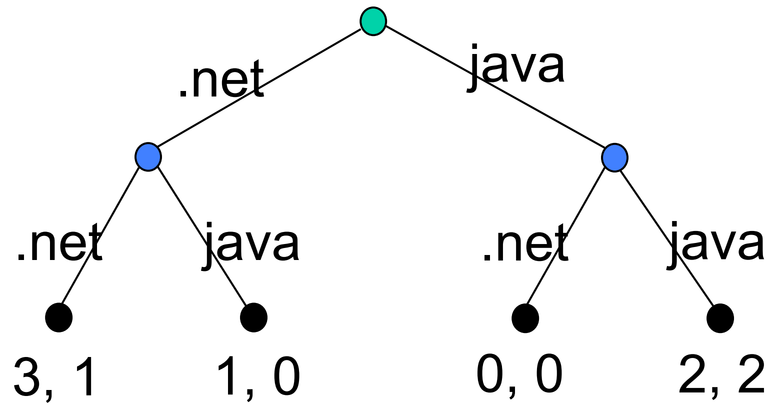

		Mozilla			
		NN	NJ	JN	JJ
Microsoft	.net	3,1	3,1	1,0	1,0
	java	0,0	2,2	0,0	2,2

- A strategy specifies the action in each information set
  - "NN" = Mozilla chooses .net in both the information sets, i.e. both if Microsoft chooses .net and if it chooses java

NE and Threats

- Convert the game to normal form

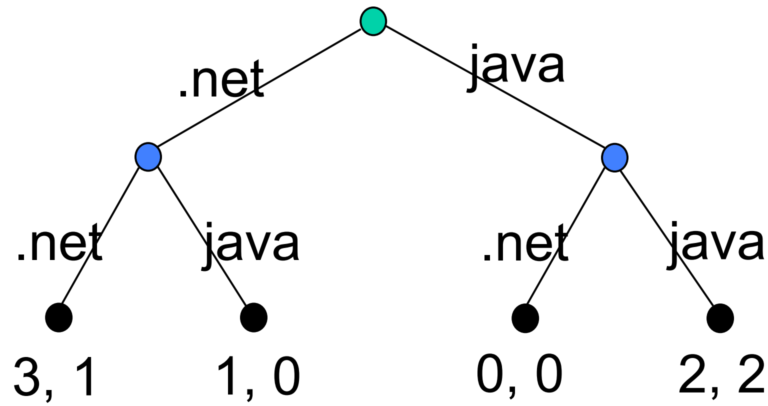

Mozilla

Microsoft		NN	NJ	JN	JJ
	.net	3,1	3,1	1,0	1,0
	java	0,0	2,2	0,0	2,2

NE

- Mozilla's JJ is a threat to Microsoft
  - I will play Java, no matter what you do
  - harmful to Microsoft, but also to Mozilla if Microsoft plays .net

NE and Threats

- Convert the game to normal form

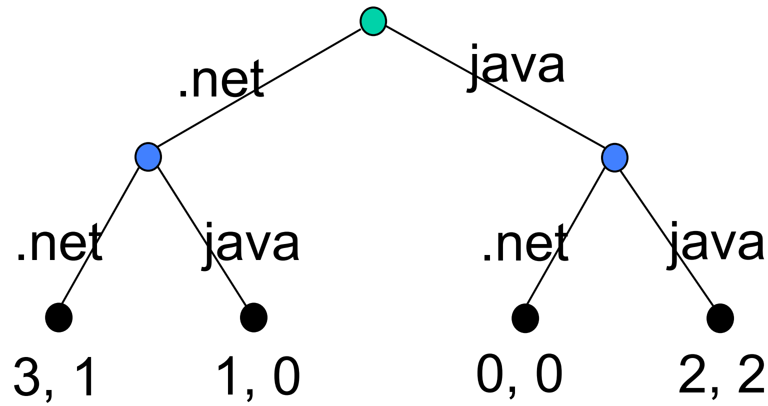

Mozilla

Microsoft	NN	NJ	JN	JJ
.net	3,-1	3,-1	1,0	1,0
java	0,-2	2,2	0,-2	2,2

NE (shaded cells: .net/NN, .net/NJ, java/JJ)

- Mozilla's JJ is a threat to Microsoft
- Mozilla may declare that it will never adopt .net (loss of image when adopting .net equal to -2)

NE and Incredible Threats

- Convert the game to normal form

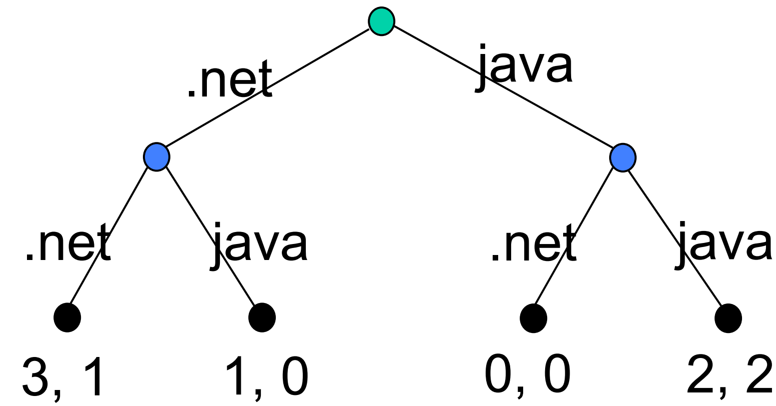

Mozilla

Microsoft	NN	NJ	JN	JJ
.net	3,1	3,1	1,0	1,0
java	0,0	2,2	0,0	2,2

NE

- Mozilla's JJ is a threat to Microsoft
- If loss of image is negligible, the threat is **incredible**
- Even if the threat is incredible, (java,JJ) is still a NE
  - How to get rid of this unconvincing NE?

Removing Incredible Threats and other poor NE

- Apply backward induction to game tree

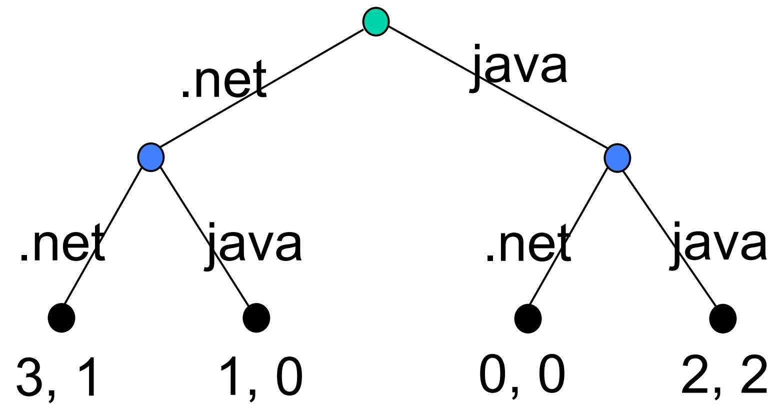

- Single NE remains
  .net for Microsoft,
  .net, java for Mozilla

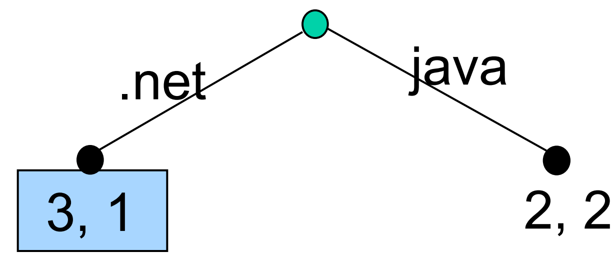

- In general, multiple NEs are possible after backward induction
  - cases with no strict preference over payoffs
- Corollary: be careful with reduction to normal form, when the game is not zero-sum!

Subgame Perfect Nash Equilibrium

- Def: a subgame is any subtree of the original game that also defines a proper game
  - only it makes sense in games with perfect information
- Def: a NE is *subgame perfect* if its restriction to *every* subgame is also a NE of the subgame
- The one deviation property: *s** is a Subgame Perfect Nash Equilibrium (SPNE) if and only if no player can gain by deviating from *s** in a single stage.
- Kuhn's Thr: every finite extensive form game with complete information has one SPNE
  - based on backward induction

NE and Incredible Threats

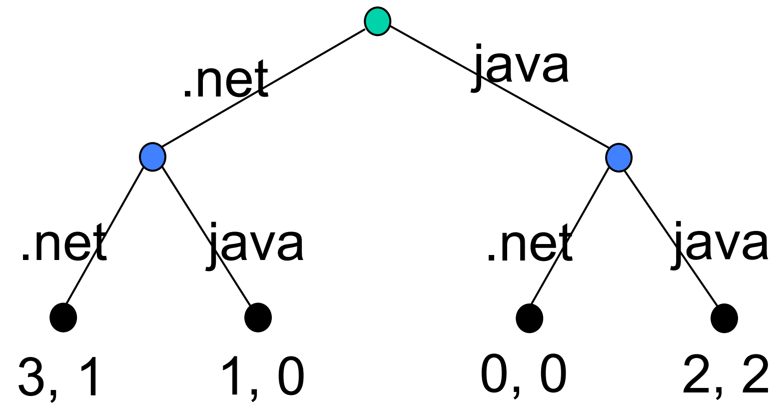

Mozilla

Microsoft		NN	NJ	JN	JJ
	.net	3,1	3,1	1,0	1,0
	java	0,0	2,2	0,0	2,2

NE: 3,1 under NN; 2,2 under JJ. SPNE: 3,1 under NJ.

- JJ is an incredible threat and java-JJ is not an SPNE
- NN is not really a threat (it motivates more Microsoft to play net), but net-NN is not an SPNE

Weakness of SPNE
(or when GT does not predict people's behaviour)

- Centipede Game
  - two players alternate decision to continue or stop for k rounds
  - stopping gives better payoff than next player stopping in next round (but not if next player continues)

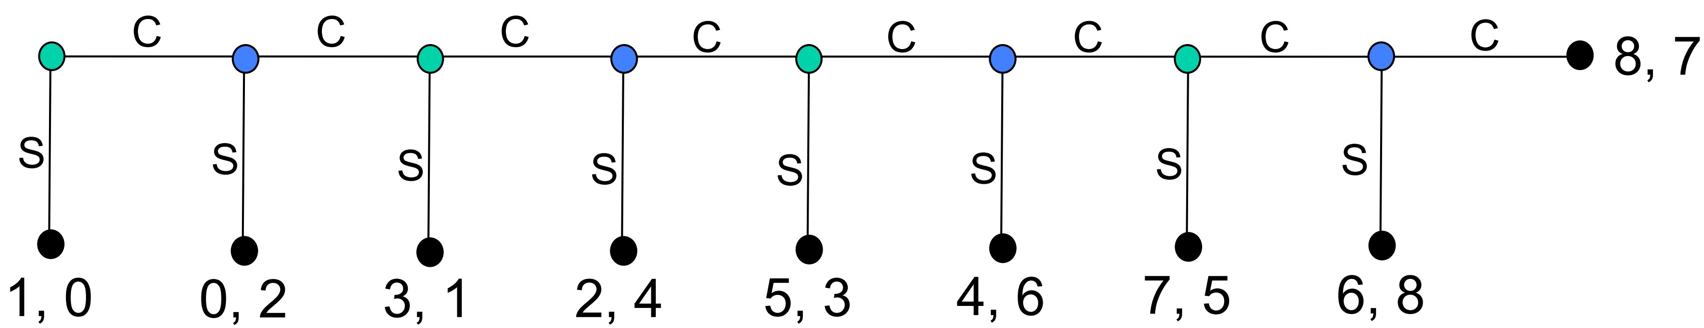

- Backward induction leads to unique SPNE
  - both players choose S in every turn
- How would you play this game with a stranger?
  - empirical evidence suggests people continue for many rounds

Stackelberg Game

- A particular game tree
- Two moves, leader then follower(s)
  - can be modeled by a game tree

- Stackelberg equilibrium
  - Leader chooses strategy knowing that follower(s) will apply best response
  - It is a SPNE for this particular game tree

Stackelberg Game and Computer Networking

- "Achieving Network Optima Using Stackelberg Routing Strategies."

  Yannis A. Korilis, Aurel A. Lazar, Ariel Orda. IEEE/ACM Transactions on Networking, 1997.

- "Stackelberg scheduling strategies".

  Tim Roughgarden. STOC 2001.

Promises

- Example: in a sequential prisoner's dilemma "I will not confess, if you not confess".

Suspect 1 \ Suspect 2	NC	C
NC	2, 2	10, 1
C	1, 10	5, 5

- Similar issues about credibility as for threats

Outline

- Two-person zero-sum games
  - Matrix games
    - Pure strategy equilibria (dominance and saddle points), ch 2
    - Mixed strategy equilibria, ch 3
  - Game trees, ch 7
  - About utility, ch 9
- Two-person non-zero-sum games
  - Nash equilibria…
    - …And its limits (equivalence, interchangeability, Prisoner's dilemma), ch. 11 and 12
  - Strategic games, ch. 14
  - Subgame Perfect Nash Equilibria (not in the book)
  - Repeated Games, partially in ch. 12
  - Evolutionary games, ch. 15
- N-persons games

Repeated games

- players face the same "stage game" in every period, and the player's payoff is a weighted average of the payoffs in each stage.
- moves are simultaneous in each stage game.
- finitely repeated (finite-horizon) and infinitely repeated (infinite-horizon) games
- in this talk, we assume:
  - players perfectly observed the actions that had been played.

Repeated games are game trees

		P2 left	P2 right
P1	L	3,1	1,0
	R	0,0	2,2

- normal form simultaneous game
  - transform it in a game tree

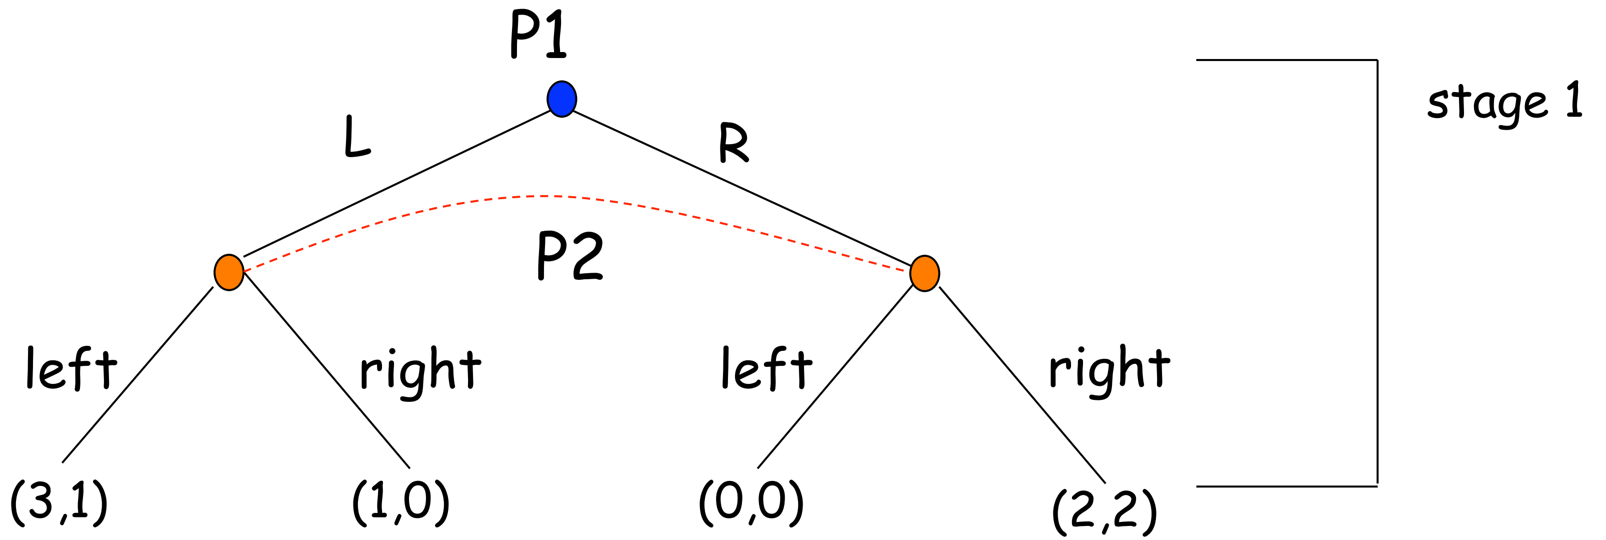

Repeated games are game trees

		P2 left	P2 right
P1	L	3,1	1,0
	R	0,0	2,2

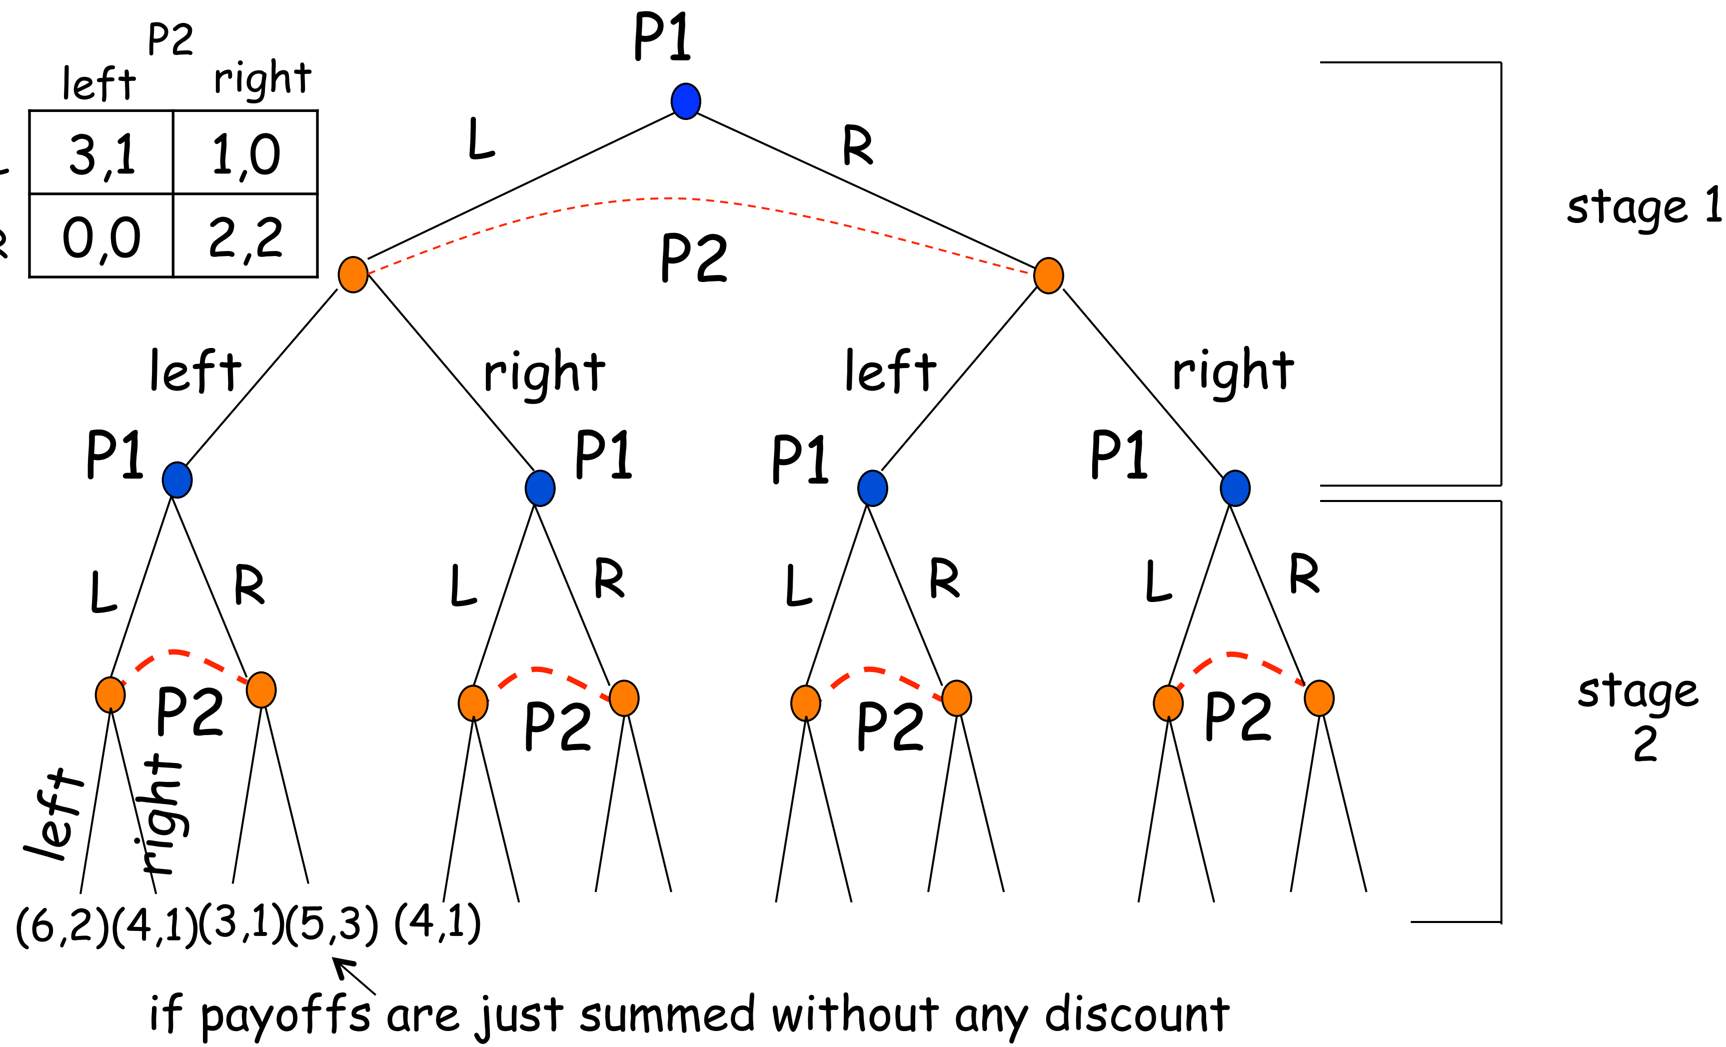

Repeated games

- $A_i=(a_{i1}, a_{i2}, \ldots, a_{i|Ai|})$: action space for player i at each stage.
- $a^t=(a_1^t,\ldots, a_n^t)$: the actions that are played in stage t.
- $h^t=(a^0, a^1\ldots, a^{t-1})$: the history of stage t, the realized choices of actions at all stages before t.
- As common in game trees a pure strategy s_i for player i maps all its information sets to actions a_i in A_i
  - in this case it means mapping possible stage-t histories h^t to actions a_i in A_i
  - player strategy needs to specify his actions also after histories that are **impossible** if he carries out his plan (see Osborne and Rubinstein section 6.4)

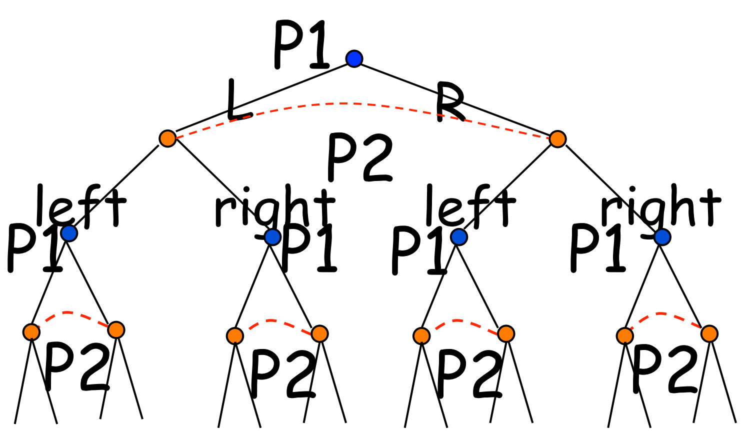

5 possible information sets and two actions available for each player.

- player 1 has 2^5 pure strategies
- player 2 has 2^5 pure strategies

Repeated games

- A mixed strategy x_i is a probability distribution over all possible pure strategies.
- A behavioral strategy b_i is a function which assigns to each information set a probability distribution over available actions, that is, randomizing over the actions available at each node.
  - see Osborne and Rubinstein, section 11.4

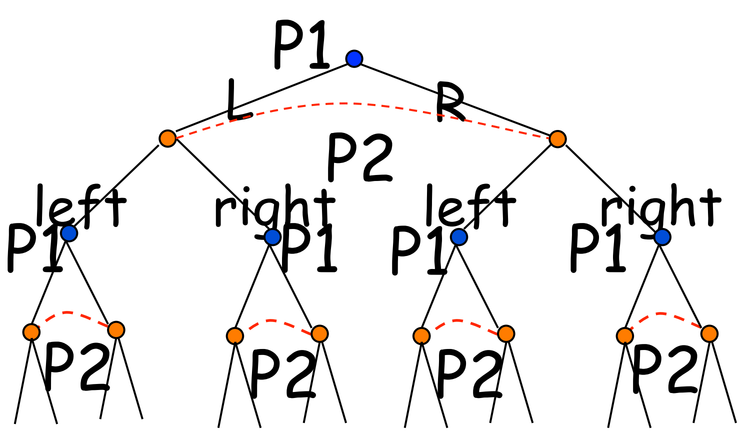

5 possible information sets and two actions available for each player.

- a mixed strategy for player 1 is specified by 2^5-1 values in [0,1]
- a behavioral strategy for player 1 is specified by 5 values in [0,1]

Repeated games

- behavioral strategies are outcome-equivalent to mixed strategies and vice versa in games with perfect recall,
  - perfect recall=a player remembers whatever he knew in the past
- two games with imperfect recall
  1. P1 forgets that he has already played
  2. P1 forgets what he played

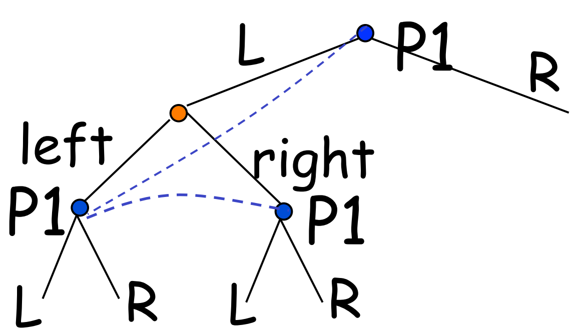

P1 behavioral strategy: play L with prob. p
• can give LL with prob. p^2, LR with prob. p(1-p)
P1 pure strategies: play L and play R
• no mixed strategy can be outcome equivalent to the behavioral strategy

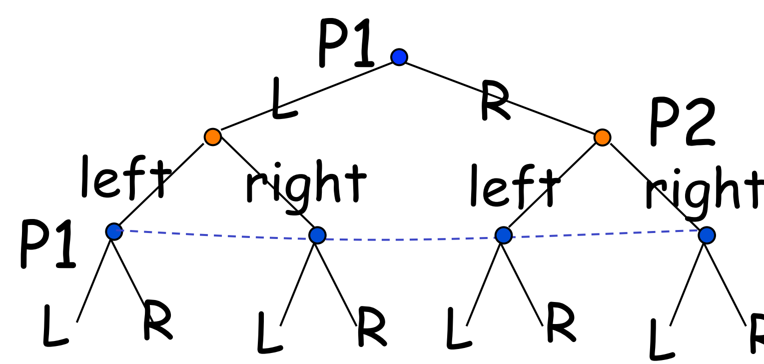

A possible P1 mixed strategy: play LL with prob. 1/2, RR with prob. 1/2
P1 behavioral strategy: 1st time play L with prob. p, 2nd time play L with prob. q
• can give LL with prob. pq,
RR with prob. (1-p)(1-q)
• not possible to obtain the mixed strategy

Infinite-horizon games

- stage games are played infinitely.
- payoff to each player is the sum of the payoffs over all periods, weighted by a discount factor δ, with $0< \delta <1$.
  - δ can be interpreted also as the probability to continue the game at each stage ($1-\delta$ is the prob. to stop playing)

- Central result: Folk Theorem.

Nash equilibrium in repeated game

- We may have new equilibrium outcomes that do not arise when the game is played only once.
  - Reason: players' actions are observed at the end of each period, players can condition their play on the past play of their opponents.
  - Example: cooperation can be a NE in Prisoner's Dilemma Game in infinitely repeated game.

Finite-horizon Prisoner's dilemma

Prisoner's Dilemma Game (Payoff Matrix)		P2	
		Cooperate	Defect
P1	Cooperate	5, 5	-3, 8
	Defect	8, -3	0, 0

- A Prisoner's Dilemma game is played 100 times.
- At the last play, $h=2^{99}\text{x}2^{99}\approx 4\text{x}10^{59}$ histories, so there are 2^h pure strategies !
- One unique subgame perfect NE: always "defect"
  - same criticism that for the centipede game (people play differently)

Infinite-horizon Prisoner's Dilemma

Prisoner's Dilemma Game (Payoff Matrix)		P2	
		Cooperate	Defect
P1	Cooperate	5, 5	-3, 8
	Defect	8, -3	0, 0

- How to find Nash equilibrium?
  - we cannot use Backward induction.
- Let's guess: trigger strategy can be subgame perfect NE if δ (discount factor) is close to one.

Trigger Strategy

- Def: follow one course of action until a certain condition is met and then follow a different strategy for the rest of the repeated game.
- Idea: each player will be deterred from abandoning the cooperative behavior by being punished. Punishments from other player are triggered by deviations
- examples:
  - trigger strategy 1: I cooperate as long as the other player cooperates, and I defect forever if the other player defects in one stage.
  - trigger strategy 2: I alternates C, D, C, … as long as the other player alternates D, C, D, … , if the other player deviates from this pattern, then I deviate forever.

Infinite-horizon Prisoner's Dilemma

- Trigger strategy 1: cooperate as long as the other player cooperates, and defect forever if the other player defects in one stage.
- Trigger strategy 1 can be subgame perfect NE if the discount factor δ is close to one.

  Proof:
  - if both players cooperate, then payoff is $5/(1-\delta)=5*(1+\delta+\delta^2+\ldots)$
  - suppose one player could defect at some round, in order to discourage this behavior, we need $5/(1-\delta) \geq 8$, or $\delta \geq 3/8$.
  - so, as long as $\delta \geq 3/8$, the pair of trigger strategies is subgame perfect NE

Cooperation can happen at Nash equilibrium !

Infinite-horizon Prisoner's Dilemma

- Trigger strategy 2: player 1 alternates C, D, C, … as long as player 2 alternates D, C, D, … , if player 2 deviates from this pattern, then player 1 deviates forever. This is also true for player 2.
- This pair of trigger strategies is also subgame perfect NE if δ is sufficiently close to one.

- In fact, there are lots of subgame perfect NEPs if δ is sufficiently close to one.
- What is happening here?

Infinite-horizon Prisoner's Dilemma

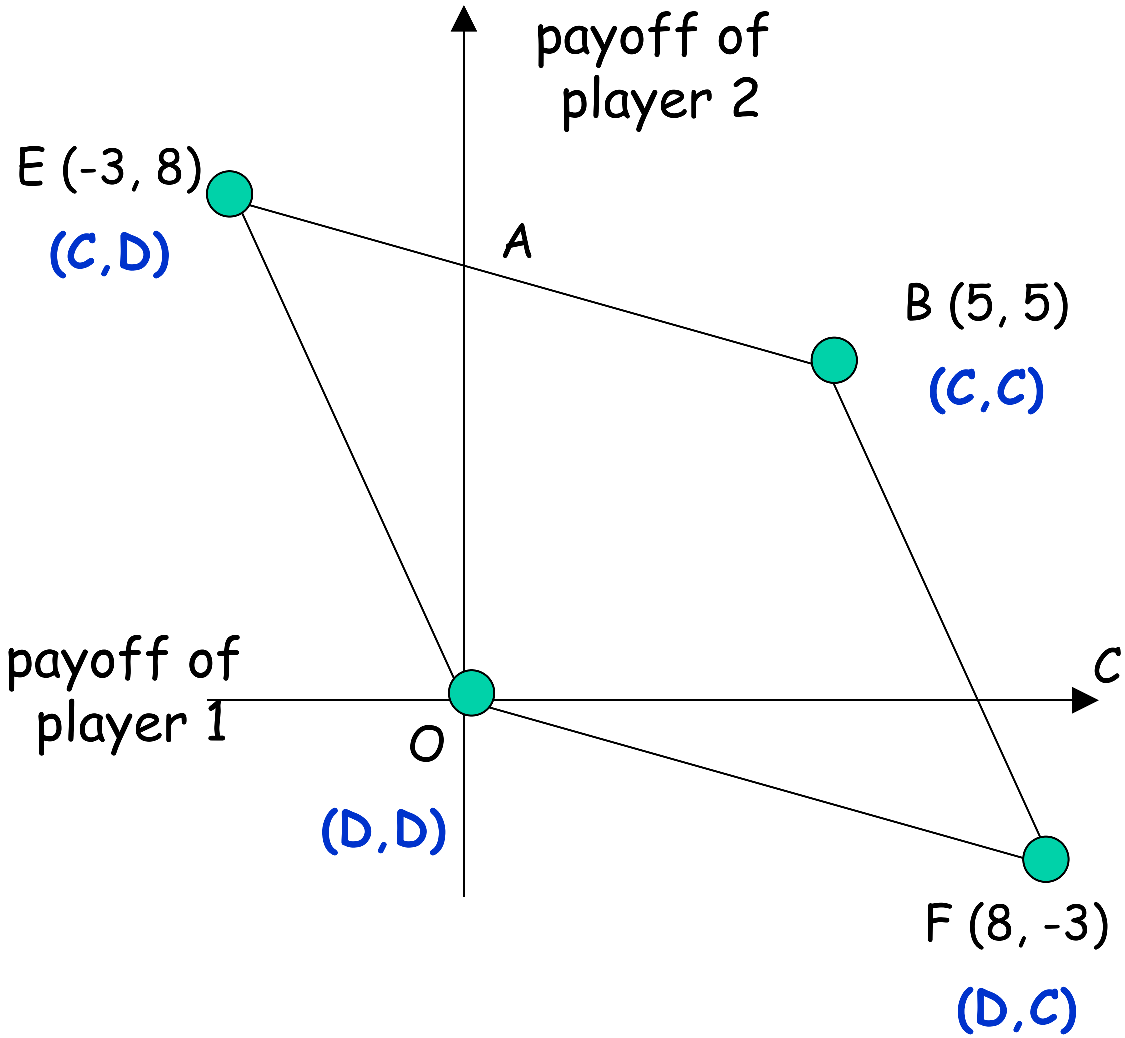

Region EOFBE contains the payoffs of all possible mixed strategy pairs.

Infinite-horizon Prisoner's Dilemma

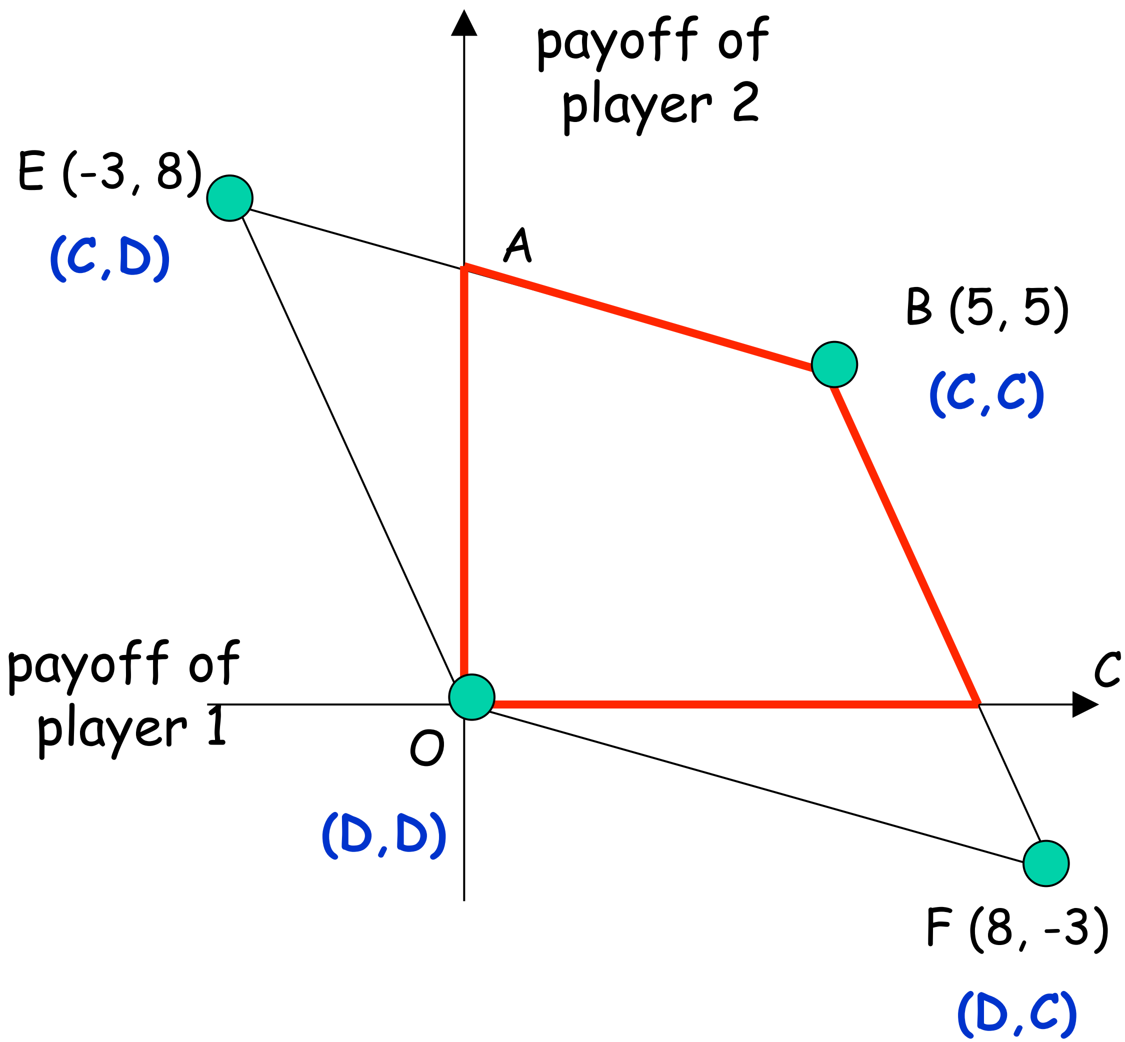

Any point in the region OABC can be sustained as a subgame perfect NE of the repeated game given the discount factor of the players is close to one (that is, players are patient enough) !

Folk Theorem

- For any two-player stage game with a Nash equilibrium with payoffs (a, b) to the players. Suppose there is a pair of strategies that give the players (c, d). Then, if c>=a and d>=b, and the discount factors of the players are sufficiently close to one, there is a subgame perfect NE with payoffs (c, d) in each period.