Game Theory: introduction and applications to computer networks

Zero-Sum Games (follow-up)

Giovanni Neglia
INRIA – EPI Maestro
20 January 2014

Part of the slides are based on a previous course
with D. Figueiredo (UFRJ) and H. Zhang (Suffolk University)

Saddle Points main theorem

- The game has a saddle point iff
  $\max_v \min_w u(v,w) = \min_w \max_v u(v,w)$

Colin

Rose

	A	B	D	$\min_w$
A	12	-1	0	-1
B	5	1	-20	-20
C	3	2	3	2
D	-16	0	16	-16
$\max_v$	12	2	16	

- Rose C ε argmax $\min_w u(v,w)$ most cautious strategy for Rose: it secures the maximum worst case gain independently from Colin's action (the game ***maximin value***)

- Colin B ε argmin $\max_v u(v,w)$ most cautious strategy for Colin: it secures the minimum worst case loss (the game ***minimax value***)

Saddle Points main theorem

- Another formulation:
  - The game has a saddle point iff

    maximin = minimax,
- This value is called the **value of the game**

Saddle Points main theorem

- The game has a saddle point iff

  $\max_v \min_w u(v,w) = \min_w \max_v u(v,w)$

N.C.

Two preliminary remarks

1. It holds (always)

   $\max_v \min_w u(v,w) <= \min_w \max_v u(v,w)$

   because $\min_w u(v,w) <= u(v,w) <= \max_v u(v,w)$ for all v and w
2. By definition, (x,y) is a saddle point iff
   - $u(x,y) <= u(x,w)$ for all w in S_{Colin}
     - i.e. $u(x,y) = \min_w u(x,w)$
   - $u(x,y) >= u(v,y)$ for all v in S_{Rose}
     - i.e. $u(x,y) = \max_v u(v,y)$

Saddle Points main theorem

- The game has a saddle point iff
  $\max_v \min_w u(v,w) = \min_w \max_v u(v,w)$

1. $\max_v \min_w u(v,w) <= \min_w \max_v u(v,w)$
2. if (x,y) is a saddle point
   - $u(x,y)=\min_w u(x,w)$, $u(x,y)=\max_v u(v,y)$

N.C.

$u(x,y)=\min_w u(x,w)<=\max_v \min_w u(v,w)<=\min_w \max_v u(v,w)<=\max_v u(v,y)=u(x,y)$

Saddle Points main theorem

- The game has a saddle point iff
  $\max_v \min_w u(v,w) = \min_w \max_v u(v,w)$

S.C.

x in argmax $\min_w u(v,w)$

y in argmin $\max_v u(v,w)$

We prove that (x,y) is a saddle-point

w_0 in $\text{argmin}_w\, u(x,w)$ ($\max_v \min_w u(v,w)=u(x,w_0)$)

v_0 in $\text{argmax}_v\, u(v,y)$ ($\min_w \max_v u(v,w)=u(v_0,y)$)

$u(x,w_0)=\min_w u(x,w) <= u(x,y) <= \max_v u(v,y)=u(v_0,y)$

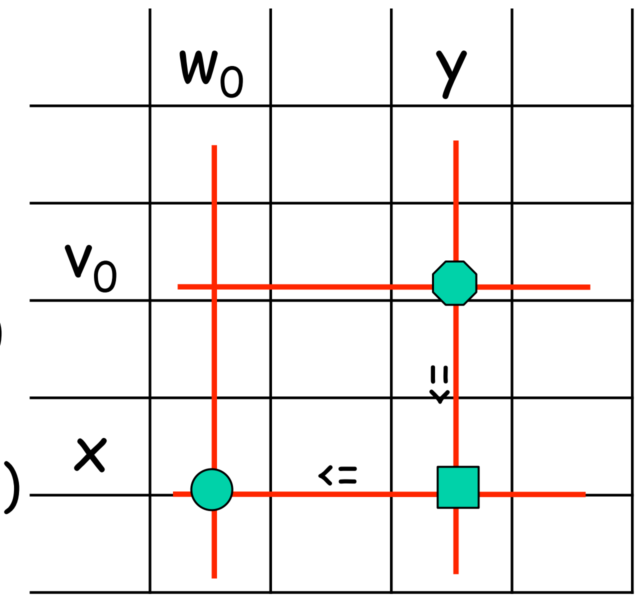

But $u(x,w_0)=u(v_0,y)$ by hypothesis, then

$u(x,y) = \min_w u(x,w) = \max_v (v,y)$

Saddle Points main theorem

- The game has a saddle point iff
  $\max_v \min_w u(v,w) = \min_w \max_v u(v,w)$

Colin

Rose	A	B	D	$\min_w$
A	12	-1	0	-1
B	5	1	-20	-20
C	3	2	3	2
D	-16	0	16	-16
$\max_v$	12	2	16	

This result provides also another way to find saddle points

Properties

- Given two saddle points (x_1,y_1) and (x_2,y_2),
  - they have the same payoff (*equivalence property*):
    - it follows from previous proof:
      $u(x_1,y_1) = \max_v \min_w u(v,w) = u(x_2,y_2)$
  - (x_1,y_2) and (x_2,y_1) are also saddle points(*interchangeability property*):
    - as in previous proof

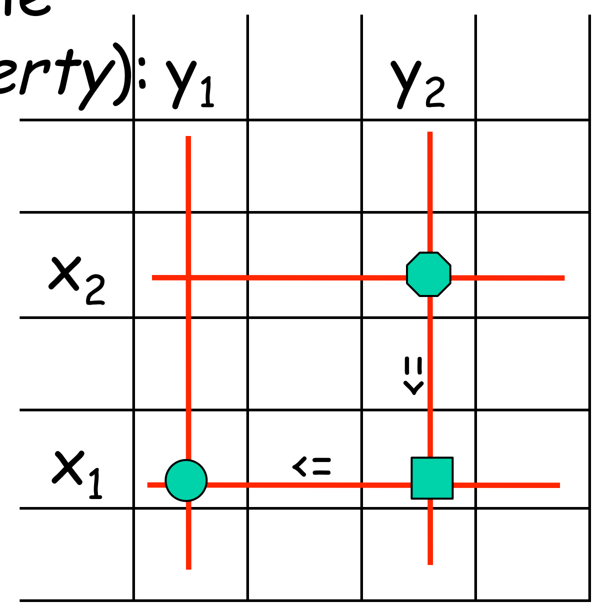

They make saddle point a very nice solution!

What is left?

- There are games with no saddle-point!
- An example?

	R	P	S	min
R	0	-1	1	-1
P	1	0	-1	-1
S	-1	1	0	-1
max	1	1	1	

maximin

minimax

maximin <> minimax

What is left?

- There are games with no saddle-point!
- An example? An even simpler one

	A	B	min
A	2	0	0
B	-5	3	-5
max	2	3	

maximin

minimax

Some practice: find all the saddle points

	A	B	C	D
A	3	2	4	2
B	2	1	3	0
C	2	2	2	2

	A	B	C
A	-2	0	4
B	2	1	3
C	3	-1	-2

	A	B	C
A	4	3	8
B	9	5	1
C	2	7	6

Games with no saddle points

		Colin	
		A	B
Rose	A	2	0
	B	-5	3

- What should players do?
  - resort to randomness to select strategies

Mixed Strategies

- Each player associates a probability distribution over its set of strategies
- Expected value principle: maximize the expected payoff

Colin	1/3	2/3
	A	B
Rose A	2	0
Rose B	-5	3

Rose's expected payoff when playing A = 1/3*2+2/3*0=2/3

Rose's expected payoff when playing B = 1/3*-5+2/3*3=1/3

- How should Colin choose its prob. distribution?

2x2 game

Rose \ Colin	A (p)	B (1-p)
A	2	0
B	-5	3

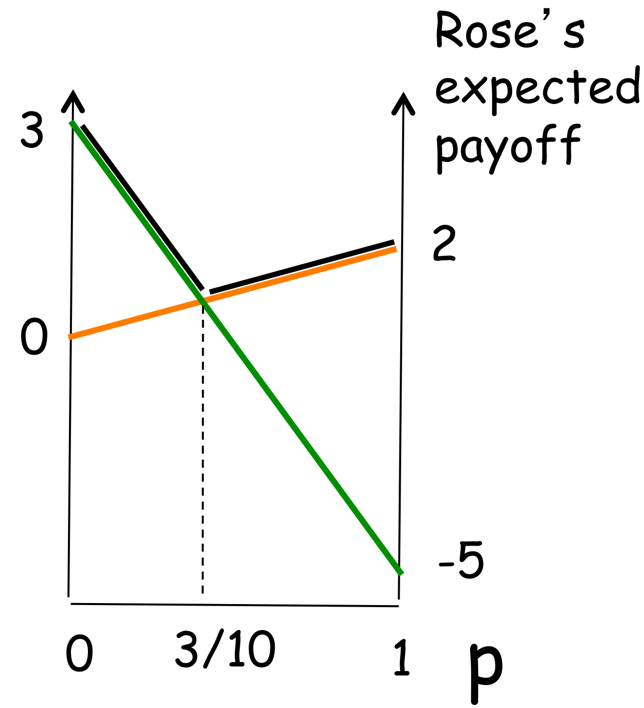

Rose's exp. gain when playing A = 2p + (1-p)*0 = 2p

Rose's exp. gain when playing B = -5*p + (1-p)*3 = 3-8p

- How should Colin choose its prob. distribution?
  - Rose cannot take advantage of p=3/10
  - for p=3/10 Colin guarantees a loss of 3/5, what about Rose's?

2x2 game

		Colin	
		A	B
Rose q	A	2	0
Rose $1-q$	B	-5	3

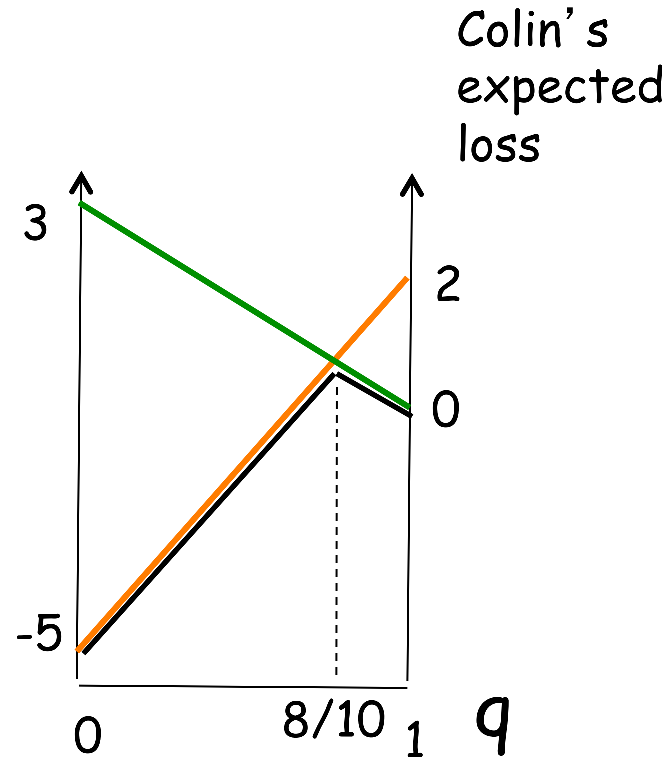

Colin's exp. loss when playing A = $2q - 5*(1-q) = 7q-5$

Colin's exp. loss when playing B = $0*q+3*(1-q) = 3-3q$

- How should Rose choose its prob. distribution?
  - Colin cannot take advantage of $q=8/10$
  - for $q=8/10$ Rose guarantees a gain of?

2x2 game

		Colin: p	Colin: $1-p$
		A	B
Rose: q	A	2	0
Rose: $1-q$	B	-5	3

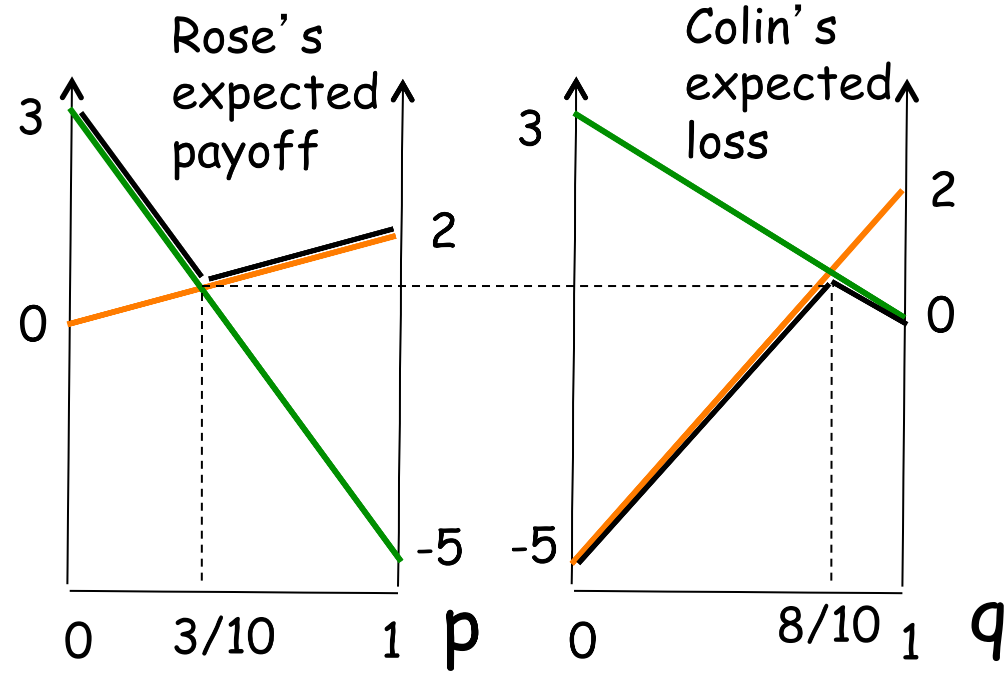

- Rose playing the mixed strategy (8/10,2/10) and Colin playing the mixed strategy (3/10,7/10) is the equilibrium of the game
  - No player has any incentives to change, because any other choice would allow the opponent to gain more
  - Rose gain 3/5 and Colin loses 3/5

mx2 game

			Colin	
			p	1-p
			A	B
Rose	x	A	2	0
	y	B	-5	3
	1-x-y	C	3	-5

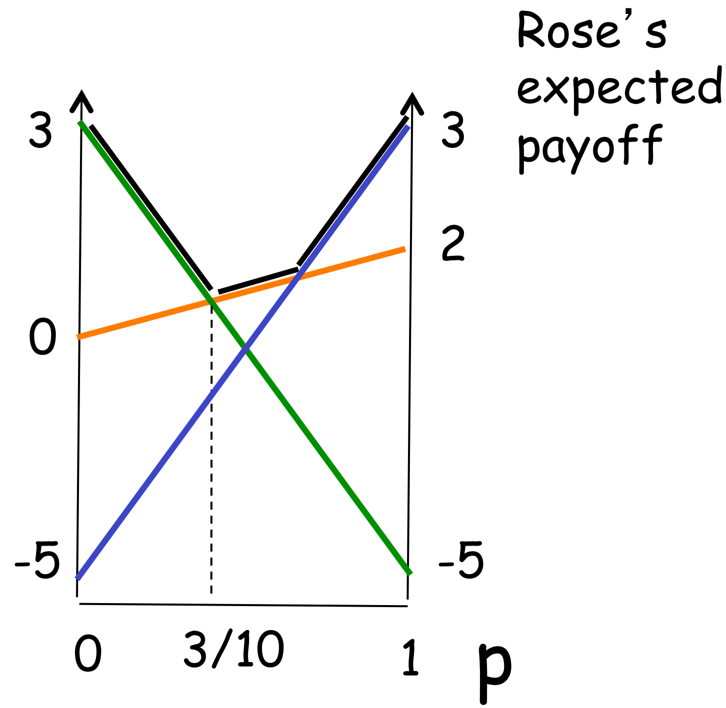

- By playing p=3/10, Colin guarantees max exp. loss = 3/5
  - it loses 3/5 if Rose plays A or B, it wins 13/5 if Rose plays C
- Rose should not play strategy C

mx2 game

			Colin	
			p	1-p
			A	B
Rose	x	A	2	0
	y	B	-5	3
	1-x-y	C	3	-5

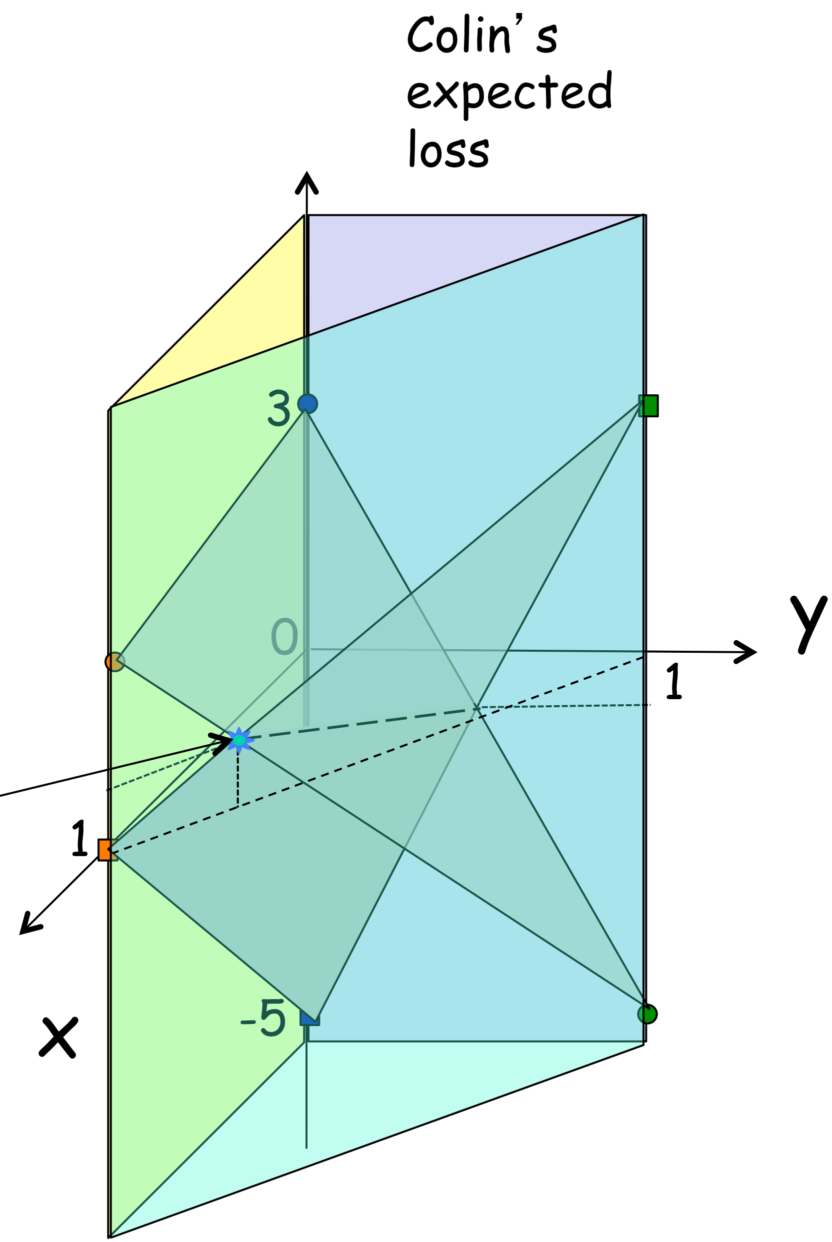

(8/10,2/10,3/5)

- Then Rose should play mixed strategy(8/10,2/10,0)
- guaranteeing a gain not less than 3/5

Minimax Theorem

- Every two-person zero-sum game has a solution, i.e, there is a unique value v (*value of the game*) and there are optimal (pure or mixed) strategies such that
  - Rose's optimal strategy guarantees to her a payoff >= v (no matter what Colin does)
  - Colin's optimal strategies guarantees to him a payoff <= v (no matter what Rose does)
- This solution can always be found as the solution of a kxk subgame
- Proved by John von Neumann in 1928!
  - birth of game theory…

How to solve mxm games

- if all the strategies are used at the equilibrium, the probability vector is such to make equivalent for the opponent all its strategies
  - a linear system with m-1 equations and m-1 variables
  - if it has no solution, then we need to look for smaller subgames

Colin

Rose		A	B	C
x	A	2	0	1
y	B	-5	3	-2
1-x-y	C	3	-5	3

Example:

- $2x-5y+3(1-x-y)=0x+3y-5(1-x-y)$
- $2x-5y+3(1-x-y)=1x-2y+3(1-x-y)$

How to solve 2x2 games

- If the game has no saddle point
  - calculate the absolute difference of the payoffs achievable with a strategy
  - invert them
  - normalize the values so that they become probabilities

		Colin	
		p	$1-p$
Rose		A	B
q	A	2	0
$1-q$	B	-5	3

$\longrightarrow |2-0|=2$ $\quad$ 8 $\longrightarrow$ 8/10

$\longrightarrow |-5-3|=8$ $\quad$ 2 $\longrightarrow$ 2/10

How to solve mxn matrix games

1. Eliminate dominated strategies
2. Look for saddle points (solution of 1x1 games), if found stop
3. Look for a solution of all the hxh games, with h=min{m,n}, if found stop
4. Look for a solution of all the (h-1)x(h-1) games, if found stop
5. …

h+1. Look for a solution of all the 2x2 games, if found stop

Remark: when a potential solution for a specific kxk game is found, it should be checked that Rose's m-k strategies not considered do not provide her a better outcome given Colin's mixed strategy, and that Colin's n-k strategies not considered do not provide him a better outcome given Rose's mixed strategy.

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