CPS 570: Artificial Intelligence

Game Theory

Instructor: Vincent Conitzer

Penalty kick example

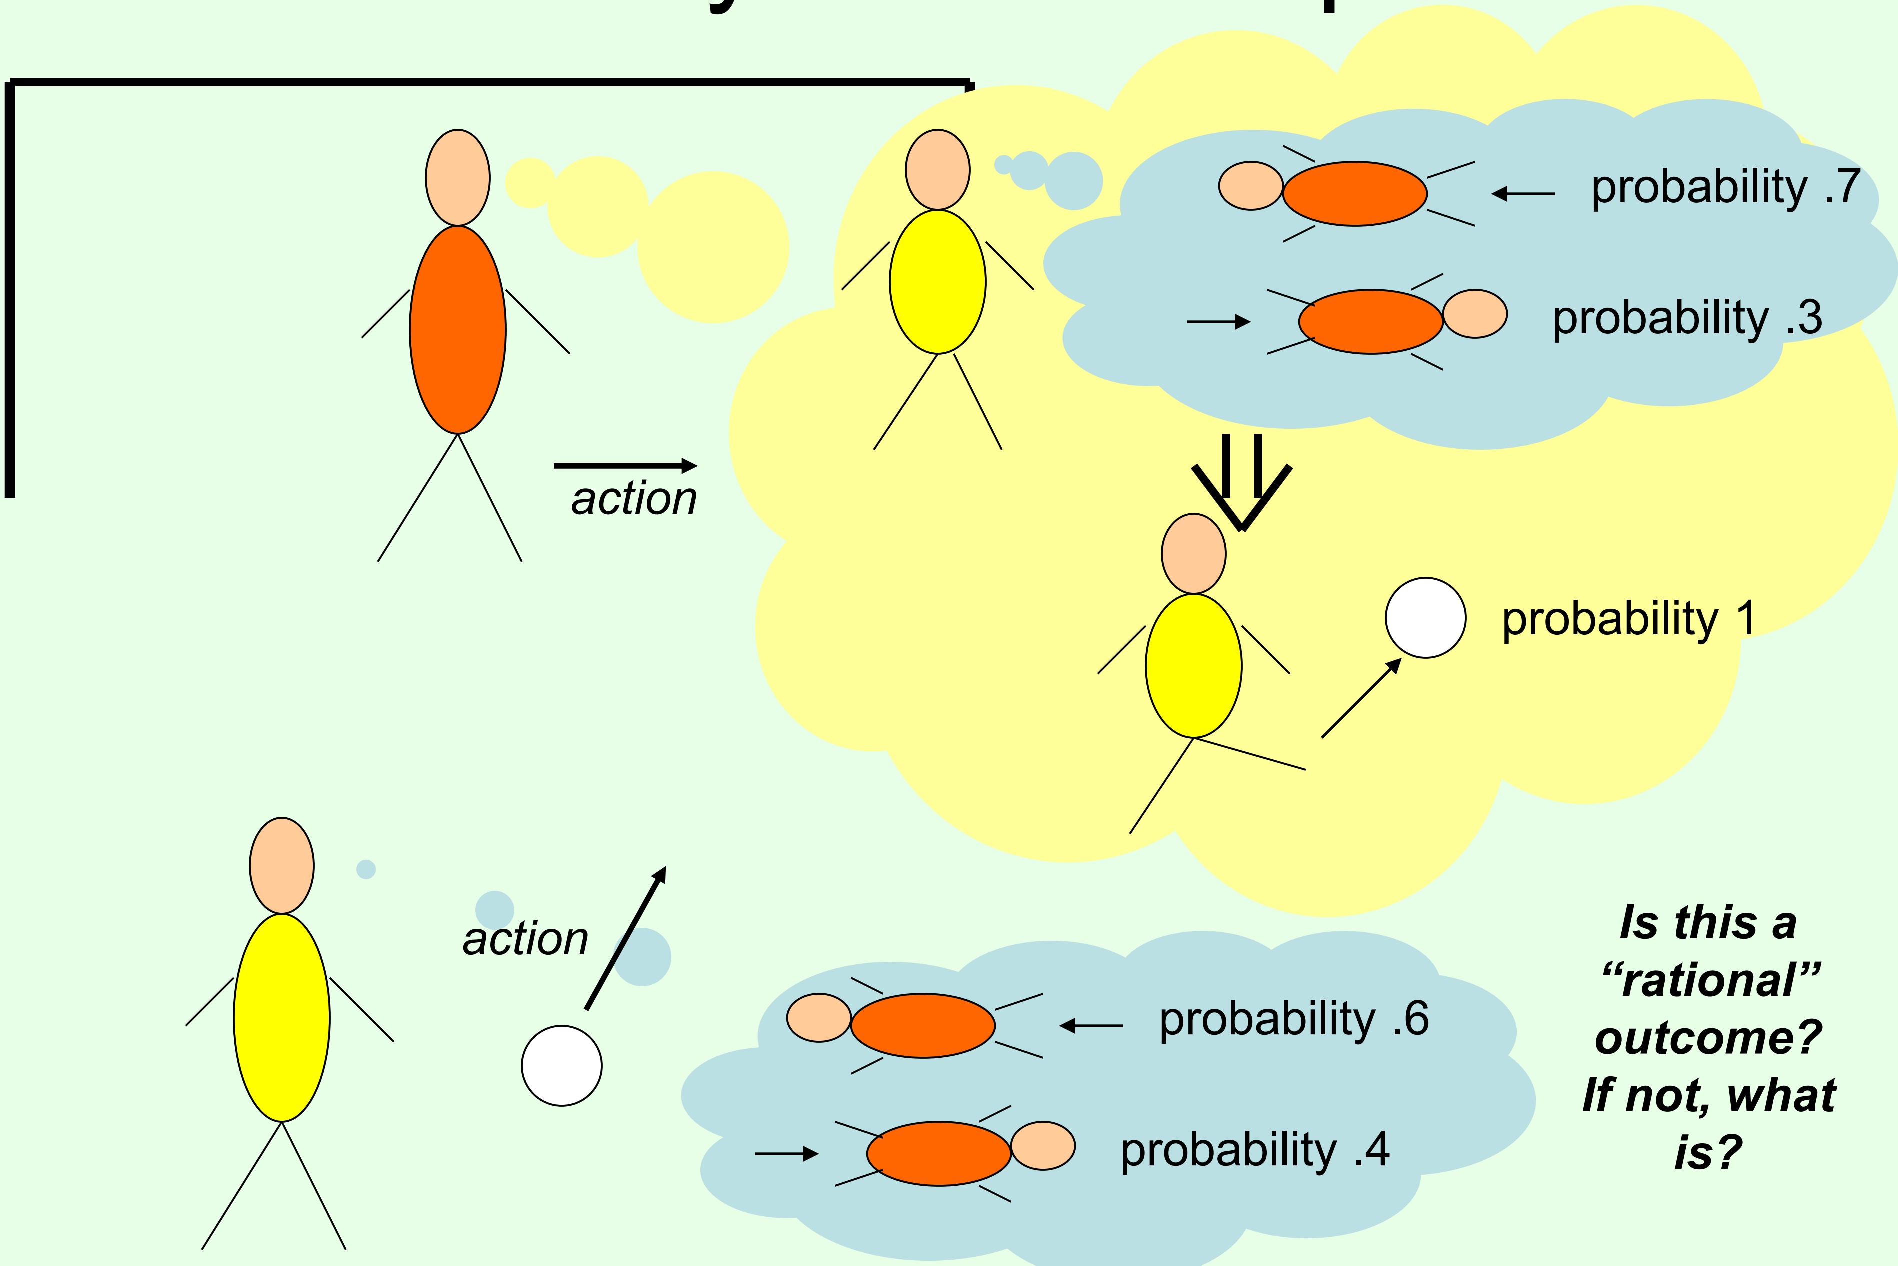

Is this a "rational" outcome? If not, what is?

Rock-paper-scissors

Column player aka. player 2 (simultaneously) chooses a column

	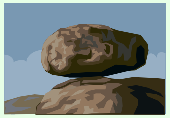	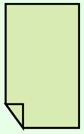	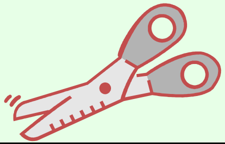
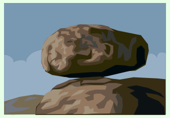	0, 0	-1, 1	1, -1
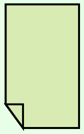	1, -1	0, 0	-1, 1
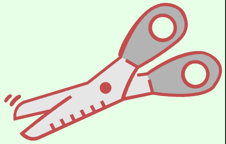	-1, 1	1, -1	0, 0

Row player aka. player 1 chooses a row

A row or column is called an action or (pure) strategy

Row player's utility is always listed first, column player's second

Zero-sum game: the utilities in each entry sum to 0 (or a constant)
Three-player game would be a 3D table with 3 utilities per entry, etc.

A poker-like game

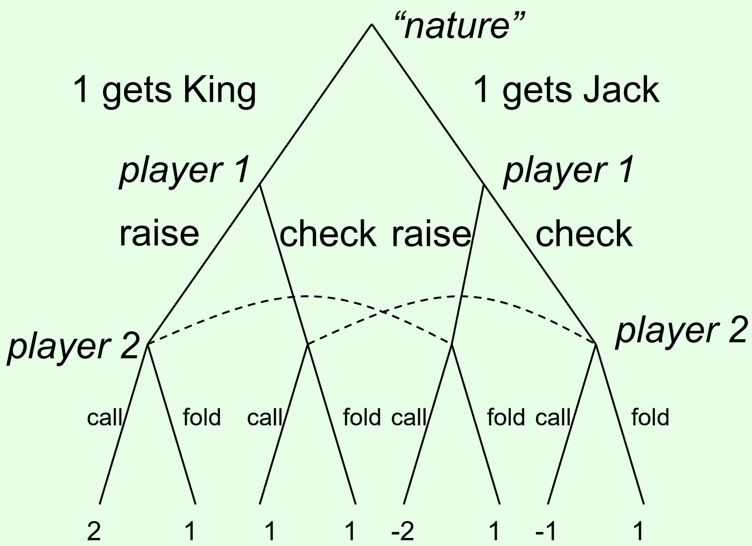

	cc	cf	fc	ff
rr	0, 0	0, 0	1, -1	1, -1
rc	.5, -.5	1.5, -1.5	0, 0	1, -1
cr	-.5, .5	-.5, .5	1, -1	1, -1
cc	0, 0	1, -1	0, 0	1, -1

“Chicken”

- Two players drive cars towards each other
- If one player goes straight, that player wins
- If both go straight, they both die

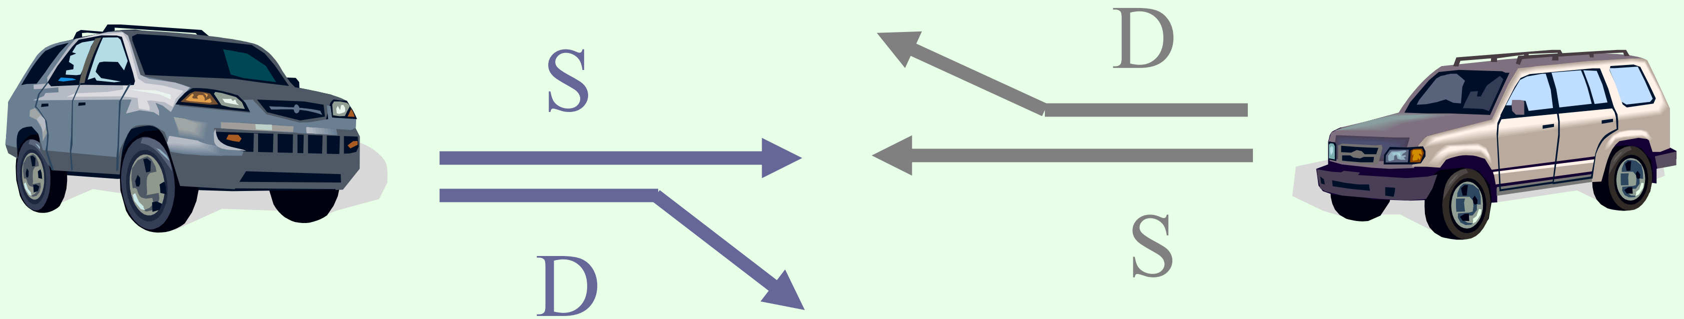

	D	S
D	0, 0	-1, 1
S	1, -1	-5, -5

not zero-sum

“2/3 of the average” game

- Everyone writes down a number between 0 and 100
- Person closest to 2/3 of the average wins
- Example:
  - A says 50
  - B says 10
  - C says 90
  - Average(50, 10, 90) = 50
  - 2/3 of average = 33.33
  - A is closest (|50-33.33| = 16.67), so A wins

Rock-paper-scissors – Seinfeld variant

MICKEY: All right, rock beats paper!
(Mickey smacks Kramer's hand for losing)
KRAMER: I thought paper covered rock.
MICKEY: Nah, rock flies right through paper.
KRAMER: What beats rock?
MICKEY: (looks at hand) Nothing beats rock.

		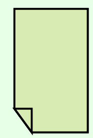	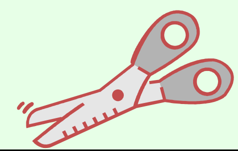
	0, 0	1, -1	1, -1
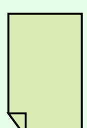	-1, 1	0, 0	-1, 1
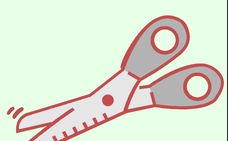	-1, 1	1, -1	0, 0

Dominance

- Player i's strategy s_i **strictly dominates** s_i' if
  - for any s_{-i}, $u_i(s_i, s_{-i}) > u_i(s_i', s_{-i})$
- s_i **weakly dominates** s_i' if
  - for any s_{-i}, $u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$; and
  - for some s_{-i}, $u_i(s_i, s_{-i}) > u_i(s_i', s_{-i})$

-i = "the player(s) other than i"

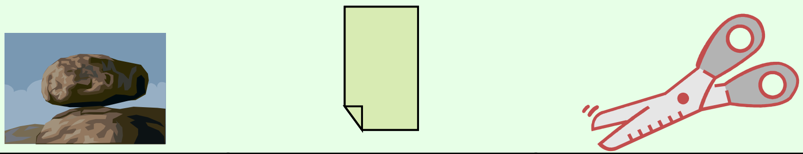

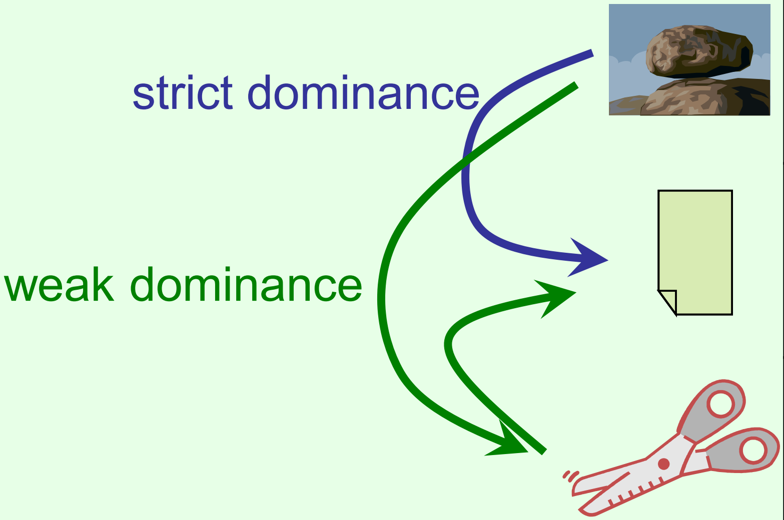

0, 0	1, -1	1, -1
-1, 1	0, 0	-1, 1
-1, 1	1, -1	0, 0

Prisoner's Dilemma

- Pair of criminals has been caught
- District attorney has evidence to convict them of a minor crime (1 year in jail); knows that they committed a major crime together (3 years in jail) but cannot prove it
- Offers them a deal:
  - If both confess to the major crime, they each get a 1 year reduction
  - If only one confesses, that one gets 3 years reduction

	confess	don't confess
confess	-2, -2	0, -3
don't confess	-3, 0	-1, -1

“Should I buy an SUV?”

purchasing + gas cost

cost: 5

cost: 3

accident cost

cost: 5  cost: 5

cost: 8 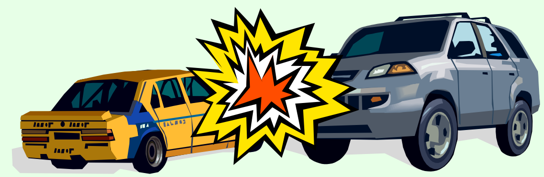 cost: 2

cost: 5  cost: 5

		
	-10, -10	-7, -11
	-11, -7	-8, -8

Back to the poker-like game

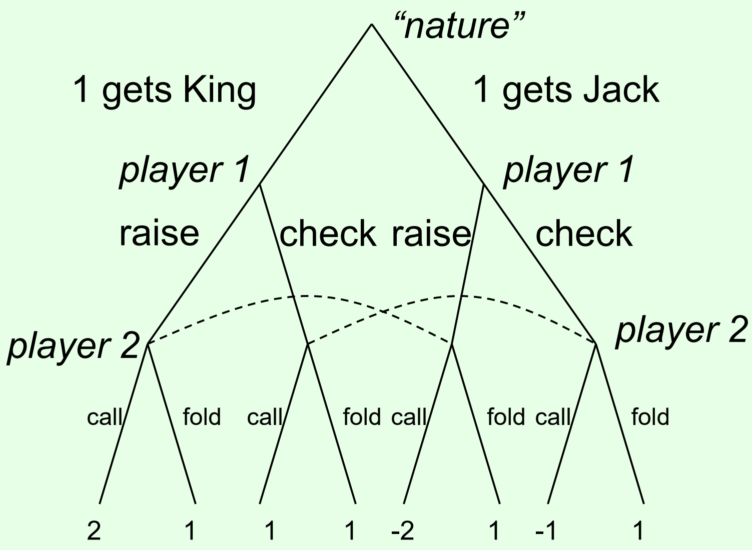

	cc	cf	fc	ff
rr	0, 0	0, 0	1, -1	1, -1
rc	.5, -.5	1.5, -1.5	0, 0	1, -1
cr	-.5, .5	-.5, .5	1, -1	1, -1
cc	0, 0	1, -1	0, 0	1, -1

Iterated dominance

- Iterated dominance: remove (strictly/weakly) dominated strategy, repeat
- Iterated strict dominance on Seinfeld's RPS:

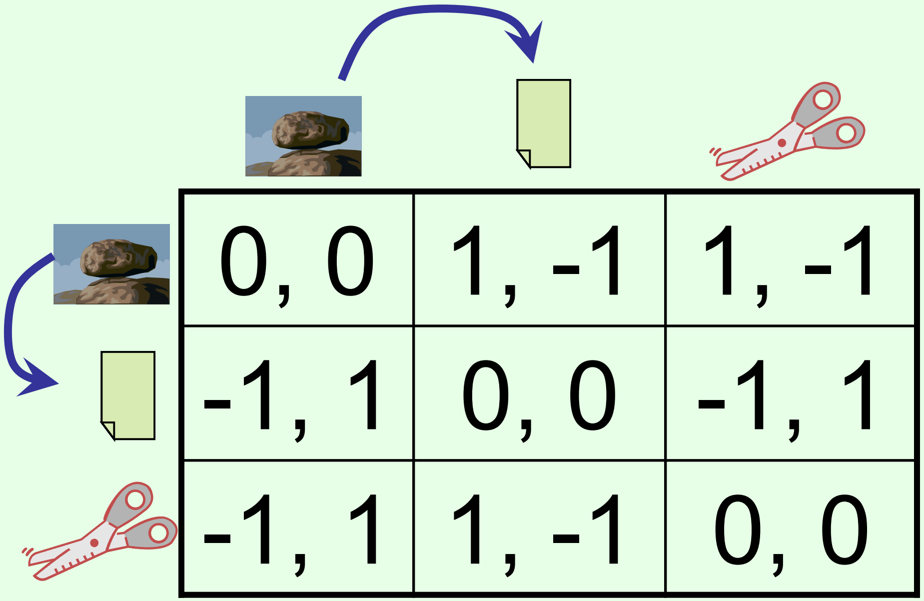

0, 0	1, -1	1, -1
-1, 1	0, 0	-1, 1
-1, 1	1, -1	0, 0

→

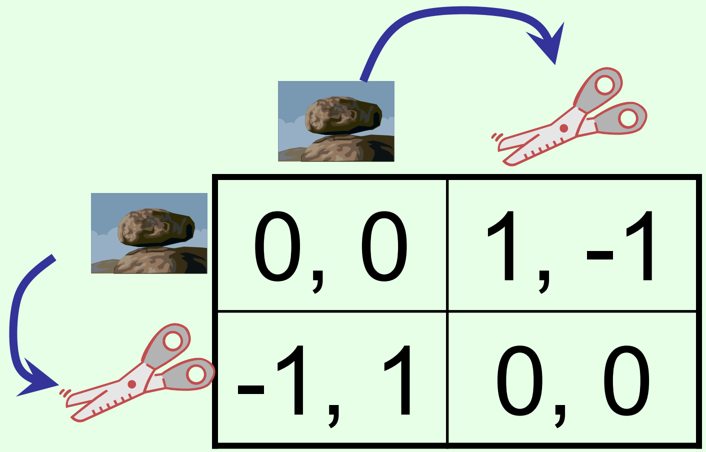

0, 0	1, -1
-1, 1	0, 0

“2/3 of the average” game revisited

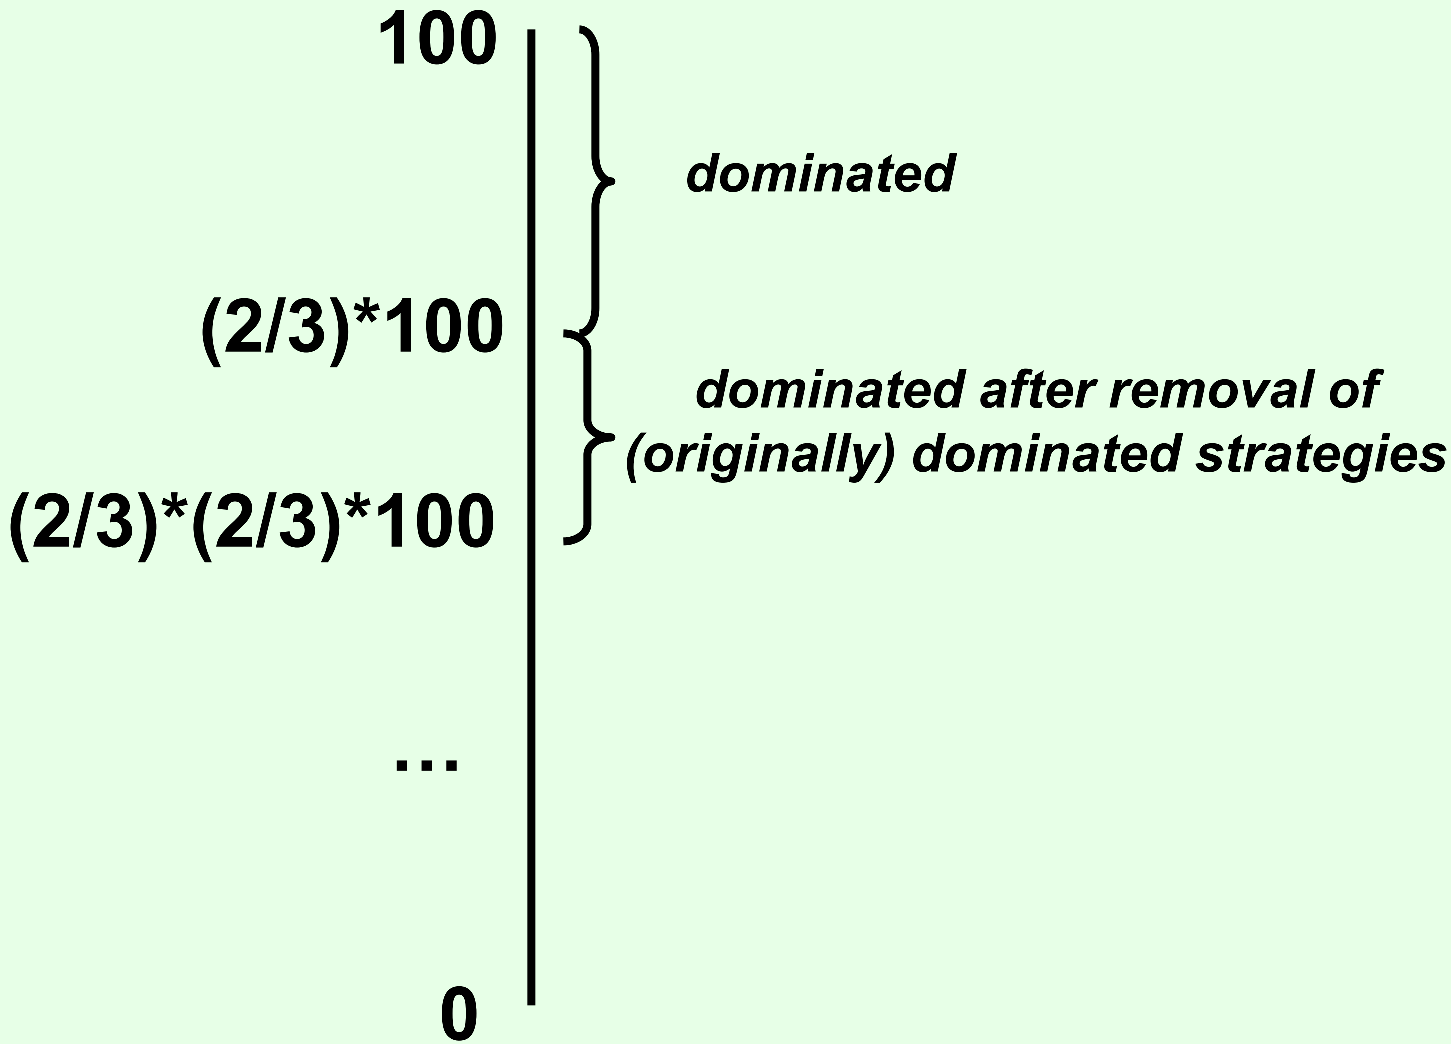

Mixed strategies

- Mixed strategy for player i = probability distribution over player i's (pure) strategies
- E.g. 1/3 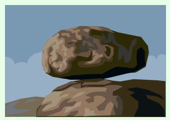 , 1/3 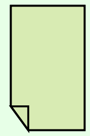 , 1/3 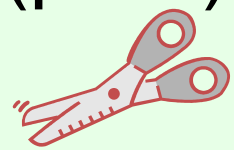
- Example of dominance by a mixed strategy:

1/2	3, 0	0, 0
1/2	0, 0	3, 0
	1, 0	1, 0

Nash equilibrium [Nash 1950]

- A profile (= strategy for each player) so that no player wants to deviate

	D	S
D	0, 0	-1, 1
S	1, -1	-5, -5

- This game has another Nash equilibrium in mixed strategies…

Rock-paper-scissors

	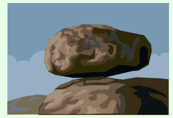	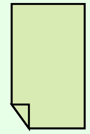	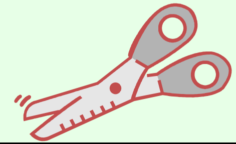
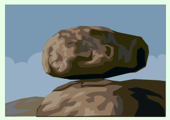	0, 0	-1, 1	1, -1
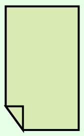	1, -1	0, 0	-1, 1
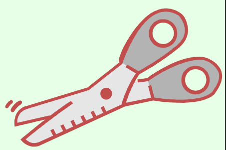	-1, 1	1, -1	0, 0

- Any pure-strategy Nash equilibria?
- But it has a mixed-strategy Nash equilibrium:

  Both players put probability 1/3 on each action
- If the other player does this, every action will give you expected utility 0
  - Might as well randomize

Nash equilibria of "chicken"…

	D	S
D	0, 0	-1, 1
S	1, -1	-5, -5

- Is there a Nash equilibrium that uses mixed strategies? Say, where player 1 uses a mixed strategy?
- If a mixed strategy is a best response, then all of the pure strategies that it randomizes over must also be best responses
- So we need to make player 1 indifferent between D and S
- Player 1's utility for playing D = $-p^c_S$
- Player 1's utility for playing S = $p^c_D - 5p^c_S = 1 - 6p^c_S$
- So we need $-p^c_S = 1 - 6p^c_S$ which means $p^c_S = 1/5$
- Then, player 2 needs to be indifferent as well
- Mixed-strategy Nash equilibrium: ((4/5 D, 1/5 S), (4/5 D, 1/5 S))
  - People may die! Expected utility -1/5 for each player

The presentation game

	Pay attention (A)	*Do not pay attention (NA)*
Put effort into presentation (E)	2, 2	-1, 0
Do not put effort into presentation (NE)	-7, -8	0, 0

- Pure-strategy Nash equilibria: (E, A), (NE, NA)
- Mixed-strategy Nash equilibrium:

  ((4/5 E, 1/5 NE), (1/10 A, 9/10 NA))

  – Utility -7/10 for presenter, 0 for audience

Back to the poker-like game, again

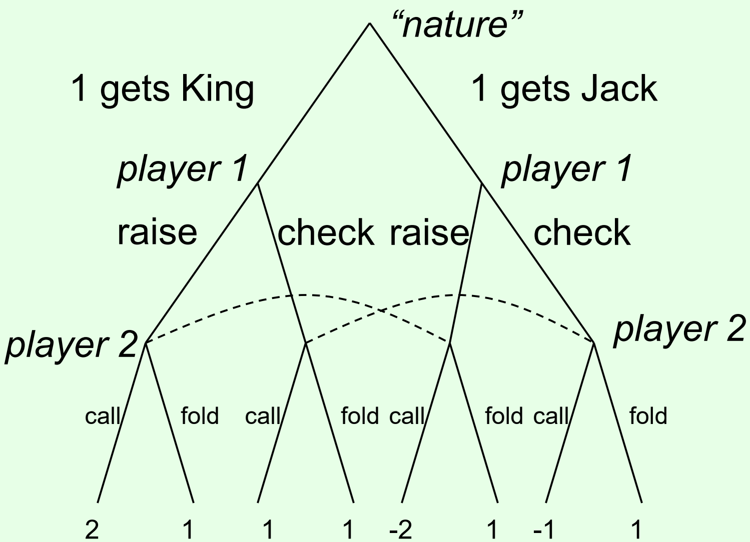

		2/3 cc	cf	1/3 fc	ff
1/3	rr	0, 0	0, 0	1, -1	1, -1
2/3	rc	.5, -.5	1.5, -1.5	0, 0	1, -1
	cr	-.5, .5	-.5, .5	1, -1	1, -1
	cc	0, 0	1, -1	0, 0	1, -1

- To make player 1 indifferent between rr and rc, we need:
  utility for rr = 0*P(cc)+1*(1-P(cc)) = .5*P(cc)+0*(1-P(cc)) = utility for rc
  That is, P(cc) = 2/3
- To make player 2 indifferent between cc and fc, we need:
  utility for cc = 0*P(rr)+(-.5)*(1-P(rr)) = -1*P(rr)+0*(1-P(rr)) = utility for fc
  That is, P(rr) = 1/3

Real-world security applications

Milind Tambe's TEAMCORE group (USC)

Airport security

Where should checkpoints, canine units, etc. be deployed?

Federal Air Marshals

Which flights get a FAM?

US Coast Guard

Which patrol routes should be followed?

Wildlife Protection

Where to patrol to catch poachers or find their snares?