Intro to Game Theory

CPS 170
Ron Parr
(with many slides courtesy of Vince Conitzer)

What is game theory?

- Game theory studies settings where multiple parties (agents) each have
  - different preferences (utility functions),
  - different actions that they can take
- Each agent's utility (potentially) depends on all agents' actions
  - What is optimal for one agent depends on what other agents do
    - Very circular!
- Game theory studies how agents can rationally form beliefs over what other agents will do, and (hence) how agents should act
  - Useful for acting as well as (potentially) predicting behavior of others
- Game theory does not directly aim to be a descriptive theory

Real World Game Theory Examples

- War
- Auctions
- Animal behavior
- Networking protocols, peer to peer networking behavior
- Road traffic

- Mechanism design: Suppose we want people to do X? How do we engineer the situation so that they will act that way?

Penalty kick example

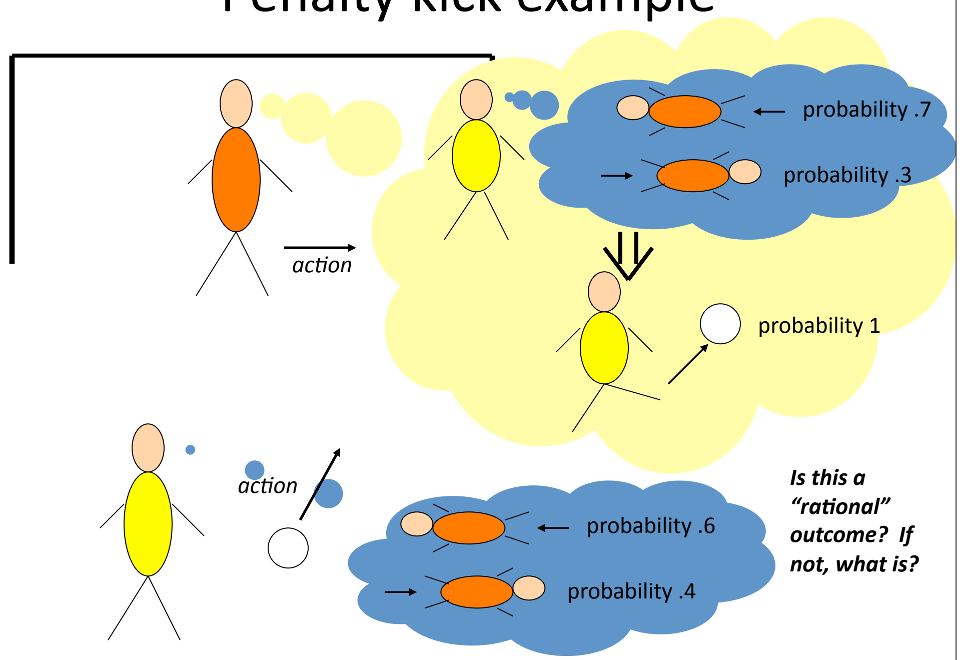

Rock-paper-scissors

Column player AKA player 2 (simultaneously) chooses a column

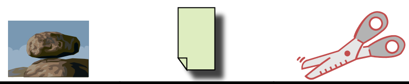

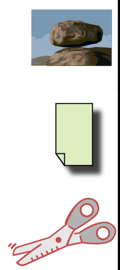

0, 0	-1, 1	1, -1
1, -1	0, 0	-1, 1
-1, 1	1, -1	0, 0

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.

"Chicken"

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

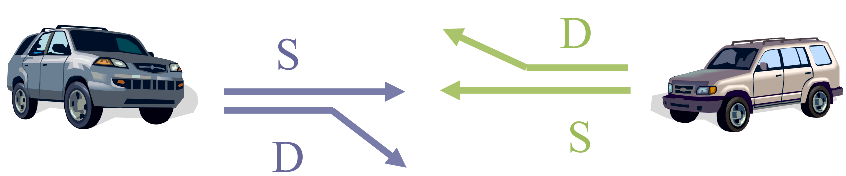

	D	S
D	0, 0	-1, 1
S	1, -1	-5, -5

not zero-sum

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.

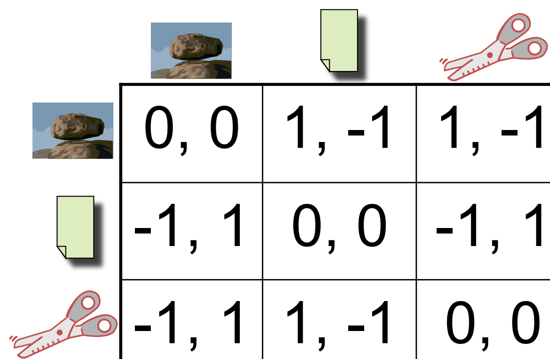

	0, 0	1, -1	1, -1
	-1, 1	0, 0	-1, 1
	-1, 1	1, -1	0, 0

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"

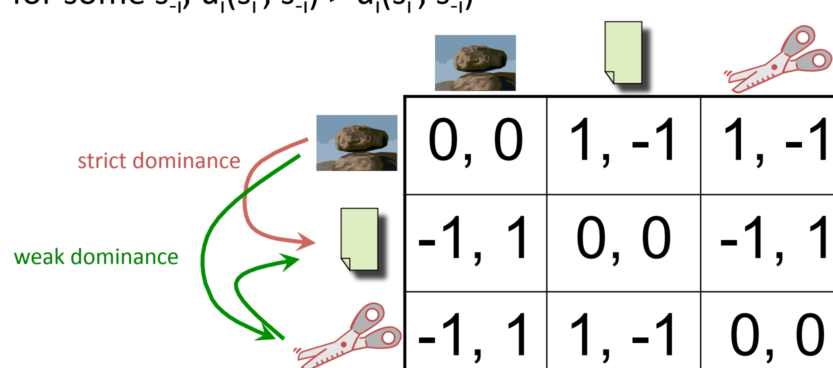

	0, 0	1, -1	1, -1
	-1, 1	0, 0	-1, 1
	-1, 1	1, -1	0, 0

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

accident cost

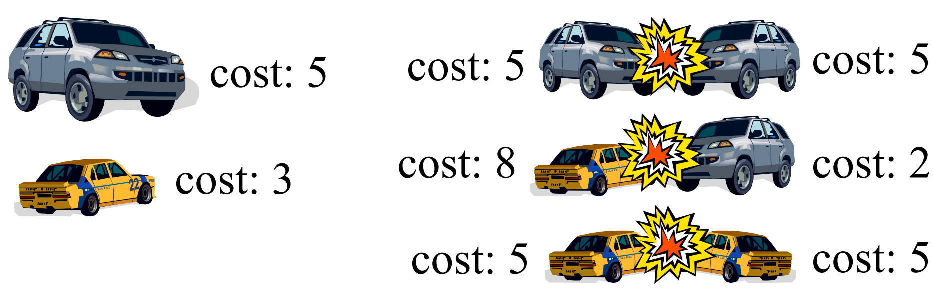

		
	-10, -10	-7, -11
	-11, -7	-8, -8

“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

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

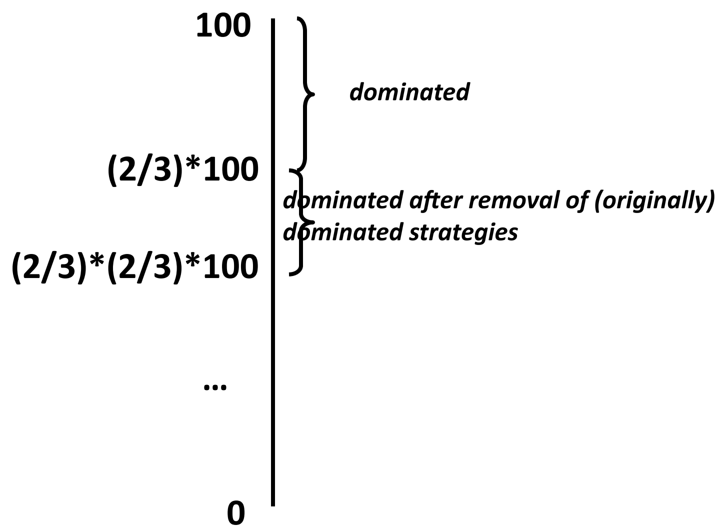

Mixed strategies

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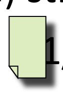 1/3 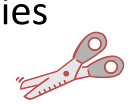
- Example of dominance by a mixed strategy:

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

Nash equilibrium [Nash 50]

- A vector of strategies (one for each player) is called a strategy profile
- A strategy profile (σ_1, σ_2 , …, σ_n) is a Nash equilibrium if each σ_i is a best response to σ_{-i}
  - That is, for any i, for any σ_i', $u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma_i', \sigma_{-i})$
- Note that this does not say anything about multiple agents changing their strategies at the same time
- In any (finite) game, at least one Nash equilibrium (possibly using mixed strategies) exists [Nash 50]

- (Note - singular: equilibrium, plural: equilibria)

Nash equilibria of “chicken”

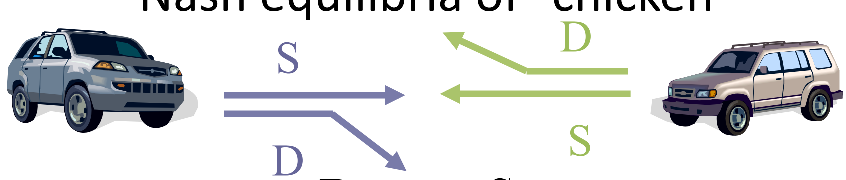

	D	S
D	0, 0	-1, 1
S	1, -1	-5, -5

- (D, S) and (S, D) are Nash equilibria
  - They are pure-strategy Nash equilibria: nobody randomizes
  - They are also strict Nash equilibria: changing your strategy will make you strictly worse off
- No other pure-strategy Nash equilibria

Rock-paper-scissors

		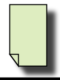	
	0, 0	-1, 1	1, -1
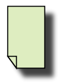	1, -1	0, 0	-1, 1
	-1, 1	1, -1	0, 0

- 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

Computational Issues

- Zero-sum games can be solved efficiently as linear programs (see slides from earlier in the semester)
- General sum games may require exponential time (in # of actions) to find a single equilibrium (non known efficient algorithm and good reasons to suspect that none exists)

- Some better news: Despite bad worst-case complexity, many games can be solved quickly

Game Theory Issues

- How descriptive is game theory?
  – Some evidence that people play equilibria
  – Some evidence that people act irrationally
  – If it is computationally intractable to solve for equilibria of large games, it would seem unlikely that people are doing this
- How reasonable is game theory?
  – Are payoffs known?
  – Are situations really simultaneous move with no information about how the other player will act?
  – Are situations really single-shot

Extensions

- Partial information (just as MDPs are extended to POMDPs)
- Uncertainty about the game parameters, e.g., payoffs (Bayesian games)
- Multistep games with distributions over next states (game theory + MDPs = stochastic games)
- Multistep + partial information (Partially observable stochastic games)

- Game theory is so general, that it can encompass essentially all aspects of strategic, multiagent behavior, e.g., negotiating, threats, bluffs, coalitions, bribes, etc.