Warm-up

- Design an AI to play Rock-Paper-Scissors for T rounds

```
function playRPS(roundID,T)
Input: $T$, roundID ∈ $\{1..T\}$
Output: action a ∈ {Rock, Paper, Scissors}

Return a
```

Announcement

- Assignments
  - P5: Ghostbusters and Bayes' Nets; Due 11/25 Mon, 10 pm
  - HW11 (online); Due 11/27 Wed, 10 pm

- Piazza post for in-class questions

AI: Representation and Problem Solving

Game Theory: Equilibrium

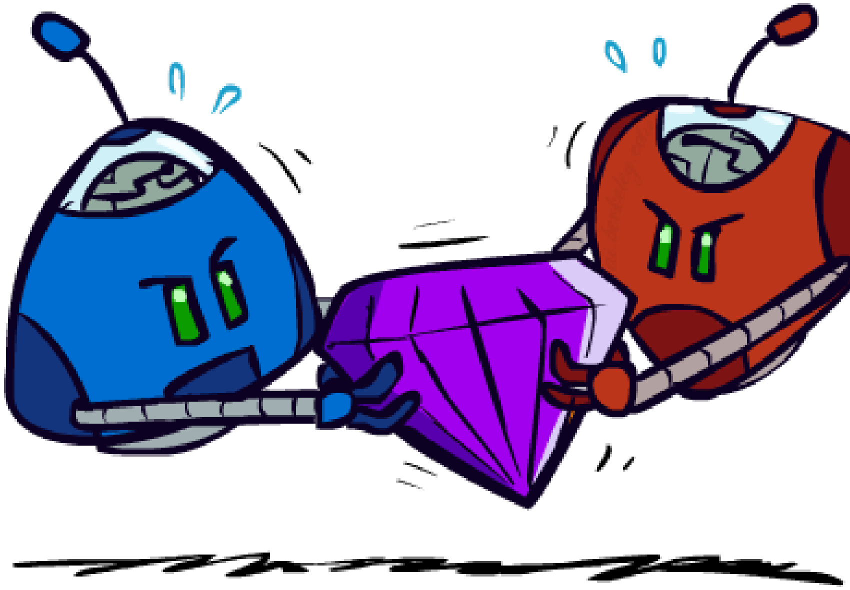

Instructors: Fei Fang & Pat Virtue

Slide credits: CMU AI and http://ai.berkeley.edu

Learning Objectives

- Formulate a problem as a game
- Describe and compare the basic concepts in game theory
  - Normal-form game, extensive-form game
  - Zero-sum game, general-sum game
  - Pure strategy, mixed strategy, support, best response, dominance
  - Dominant strategy equilibrium, Nash equilibrium, Minimax strategy, maximin strategy, Stackelberg equilibrium
- Describe iterative removal algorithm
- Describe minimax theorem
- Compute equilibria for bimatrix games
  - Pure strategy Nash equilibrium
  - Mixed strategy Nash equilibrium (including using LP for zero-sum games)
  - Stackelberg equilibrium (only pure strategy equilibrium is required)

From Games to Game Theory

- The study of mathematical models of conflict and cooperation between intelligent decision makers
- Used in economics, political science etc

John von Neumann

John Nash

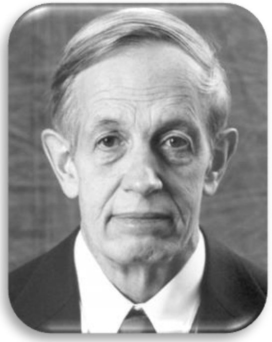

Heinrich Freiherr von Stackelberg

Winners of Nobel Memorial Prize in Economic Sciences

Recall: Adversarial Search

- Zero-sum, perfect information, two player games with turn-taking moves

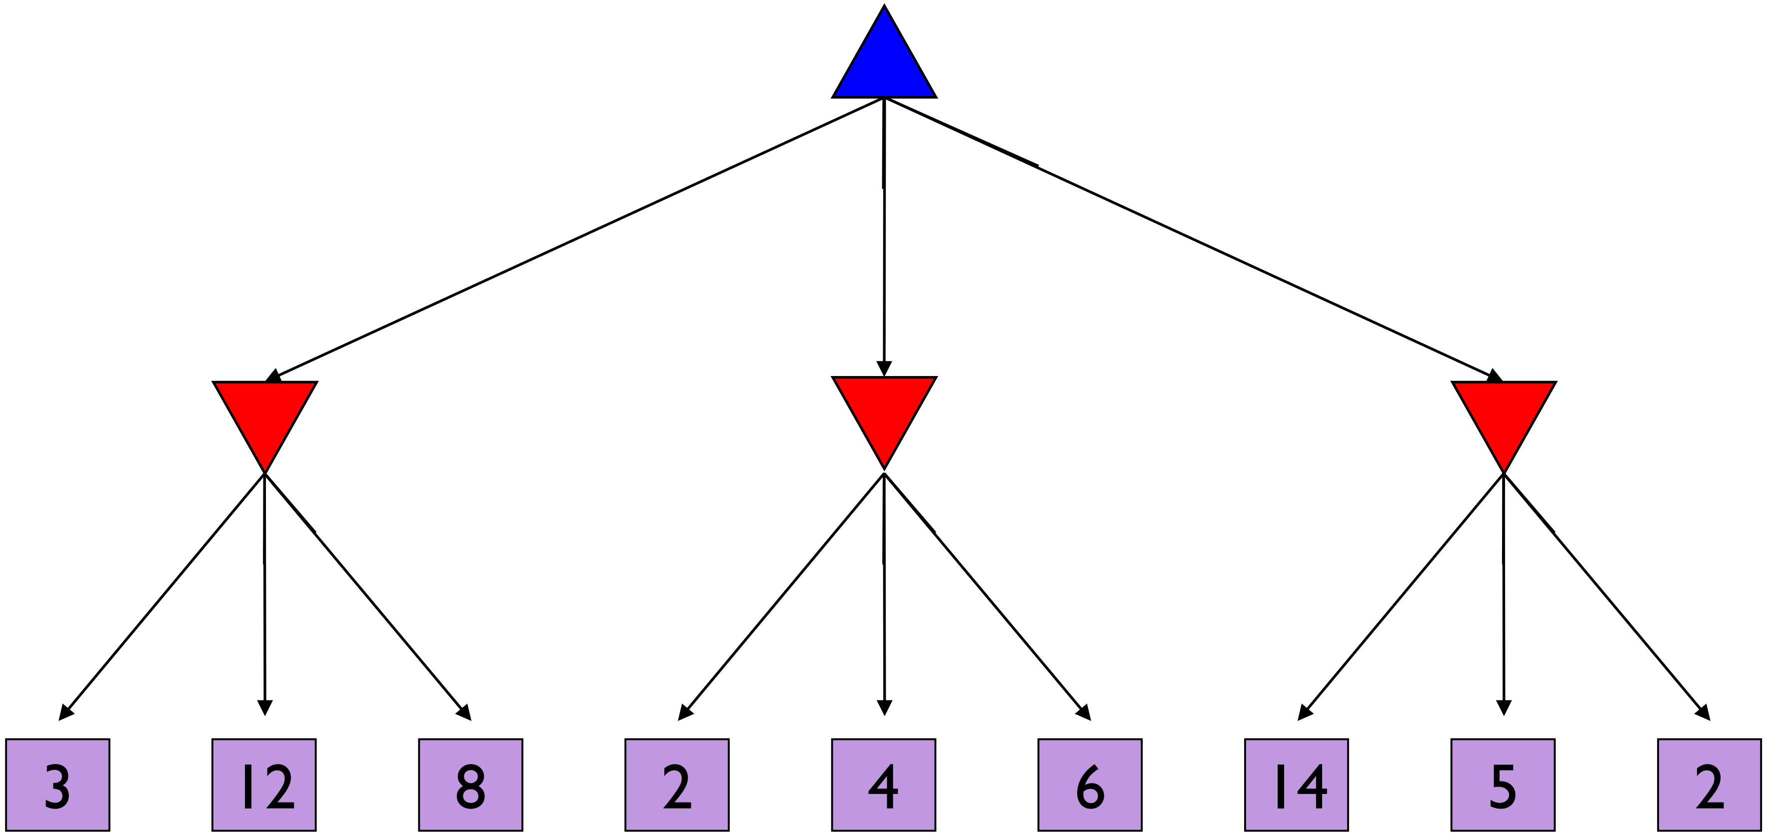

Classical Games and Payoff Matrices

- Rock-Paper-Scissors (RPS)
  - Rock beats Scissors
  - Scissors beats Paper
  - Paper beats Rock

Player 2

Player 1

	Rock	Paper	Scissors
Rock	0, 0	-1, 1	1, -1
Paper	1, -1	0, 0	-1, 1
Scissors	-1, 1	1, -1	0, 0

Classical Games and Payoff Matrices

- Prisoner's Dilemma (PD)
  - If both Cooperate: 1 year in jail each
  - If one Defect, one Cooperate: 0 year for (D), 3 years for (C)
  - If both Defect: 2 years in jail each
  - Let's play!

Player 2

Player 1

	Cooperate	Defect
Cooperate	-1,-1	-3,0
Defect	0,-3	-2,-2

Variation: Split or Steal

https://youtu.be/p3Uos2fzIJ0

Classical Games and Payoff Matrices

- Football vs Concert (FvsC)
  - Historically known as Battle of Sexes
  - If football together: Alex ☺☺, Berry ☺
  - If concert together: Alex ☺, Berry ☺☺
  - If not together: Alex ☹, Berry ☹

Fill in the payoff matrix

Berry

Alex

	Football	Concert
Football		
Concert		

Classical Games and Payoff Matrices

- Football vs Concert (FvsC)
  - Historically known as Battle of Sexes
  - If football together: Alex ☺☺, Berry ☺
  - If concert together: Alex ☺, Berry ☺☺
  - If not together: Alex ☹, Berry ☹

Berry

		Football	Concert
Alex	Football	2,1	0,0
	Concert	0,0	1,2

Normal-Form Games

- A game in normal form consists of the following elements
  - Set of players
  - Set of actions for each player
  - Payoffs / Utility functions
    - Determines the utility for each player given the actions chosen by all players (referred to as action profile)
  - Bimatrix game is special case: two players, finite action sets
- Players move simultaneously and the game ends immediately afterwards

What are the players, set of actions and utility functions of Football vs Concert (FvsC) game?

Warm-up

- Design an AI to play Rock-Paper-Scissors for T rounds

If T=1, roundID=1, what action does your AI choose?

```
function playRPS(roundID,T)
Input: $T$, roundID $\in \{1..T\}$
Output: action a $\in$ {Rock, Paper, Scissors}

Return a
```

Strategy

- Pure strategy: choose an action deterministically
- Mixed strategy: choose actions according to a probability distribution
  - Notation: $s = (0.3, 0.7, 0)$
  - Support: set of actions chosen with non-zero probability

Notation Alert! We use s to represent strategy here (not states)

Does your AI play a deterministic strategy or a mixed strategy?

What is the support size of your AI's strategy?

Player 2

		Rock	Paper	Scissors
Player 1	Rock	0, 0	-1, 1	1, -1
	Paper	1, -1	0, 0	-1, 1
	Scissors	-1, 1	1, -1	0, 0

Zero-sum vs General-sum

- Zero-sum Game
  - No matter what actions are chosen by the players, the utilities for all the players sum up to zero or a constant
- General-sum Game
  - The sum of utilities of all the players is not a constant

Which ones are general-sum games?

Player 2

Player 1

	Rock	Paper	Scissors
Rock	0,0	-1,1	1,-1
Paper	1,-1	0,0	-1,1
Scissor	-1,1	1,-1	0,0

Player 2

Player 1

	Cooperate	Defect
Cooperate	-1,-1	-3,0
Defect	0,-3	-2,-2

Berry

Alex

	Football	Concert
Football	2,1	0,0
Concert	0,0	1,2

Zero-sum vs General-sum

- Zero-sum Game
  - No matter what actions are chosen by the players, the utilities for all the players sum up to zero or a constant
- General-sum Game
  - The sum of utilities of all the players is not a constant

Which ones are general-sum games?

Player 2

Player 1

	Rock	Paper	Scissors
Rock	0,0	-1,1	1,-1
Paper	1,-1	0,0	-1,1
Scissor	-1,1	1,-1	0,0

Player 2

Player 1

	Cooperate	Defect
Cooperate	-1,-1	-3,0
Defect	0,-3	-2,-2

Berry

Alex

	Football	Concert
Football	2,1	0,0
Concert	0,0	1,2

Expected Utility

Notation Alert!
Use a, s, u to represent action, strategy, utility of a player
Use $\mathbf{a}, \mathbf{s}, \mathbf{u}$ to represent action, strategy, utility profile

- Given the strategies of all players,

Expected Utility for player i u_i=

$$\sum_{\mathbf{a}} \text{Prob(action profile } \mathbf{a}) \times \text{Utility for player } i \text{ in } \mathbf{a}$$

Can skip action profiles with probability 0 or utility 0

If Alex's strategy s_A=$(\frac{1}{2},\frac{1}{2})$, Berry's strategy s_B=$(1,0)$

What is the probability of action profile $\mathbf{a}$ =(Concert,Football)?

What is Alex's utility in this action profile?

		Berry	
		Football	Concert
Alex	Football	2,1	0,0
	Concert	0,0	1,2

Expected Utility

Notation Alert!
Use a, s, u to represent action, strategy, utility of a player
Use $\mathbf{a}, \mathbf{s}, \mathbf{u}$ to represent action, strategy, utility profile

- Given the strategies of all players,

Expected Utility for player i u_i=

$$\sum_{\mathbf{a}} \text{Prob(action profile } \mathbf{a}) \times \text{Utility for player } i \text{ in } \mathbf{a}$$

Can skip action profiles with probability 0 or utility 0

If Alex's strategy s_A=$(\frac{1}{2}, \frac{1}{2})$, Berry's strategy s_B=$(1,0)$

What is the probability of action profile $\mathbf{a}$ =(Concert,Football)?

$$\frac{1}{2} \times 1 = \frac{1}{2}$$

What is Alex's utility in this action profile?

$$0$$

		Berry	
		Football	Concert
Alex	Football	2,1	0,0
	Concert	0,0	1,2

Piazza Poll 1

- In Rock-Paper-Scissors, if $s_1 = (\frac{1}{3}, \frac{2}{3}, 0), s_2 = (0, \frac{1}{2}, \frac{1}{2})$, how many non-zero terms need to be added up when computing the expected utility for player 1?
  - A: 9
  - B: 6
  - C: 4
  - D: 3
  - E: I don't know

Player 2

Player 1

	Rock	Paper	Scissors
Rock	0, 0	-1, 1	1, -1
Paper	1, -1	0, 0	-1, 1
Scissors	-1, 1	1, -1	0, 0

Piazza Poll 1

- In Rock-Paper-Scissors, if $s_1 = (\frac{1}{3}, \frac{2}{3}, 0), s_2 = (0, \frac{1}{2}, \frac{1}{2})$, how many non-zero terms need to be added up when computing the expected utility for player 1?
  - A: 9
  - B: 6
  - C: 4
  - D: 3
  - E: I don't know

$$\begin{aligned} u_1 = {} & 0 \times \frac{1}{3} \times 0 + (-1) \times \frac{1}{3} \times \frac{1}{2} + 1 \times \frac{1}{3} \times \frac{1}{2} \\ & +1 \times \frac{2}{3} \times 0 + 0 \times \frac{2}{3} \times \frac{1}{2} + (-1) \times \frac{2}{3} \times \frac{1}{2} \\ & +(-1) \times 0 \times 0 + 1 \times 0 \times \frac{1}{2} + 0 \times 0 \times \frac{1}{2} = -\frac{1}{3} \end{aligned}$$

Player 2

Player 1

	Rock	Paper	Scissors
Rock	0, 0	-1, 1	1, -1
Paper	1, -1	0, 0	-1, 1
Scissors	-1, 1	1, -1	0, 0

Best Response

- Best Response (BR): Given the strategies or actions of all players but player i (denoted as $\mathbf{s}_{-i}$ or $\mathbf{a}_{-i}$), Player i's best response to $\mathbf{s}_{-i}$ or $\mathbf{a}_{-i}$ is the set of actions or strategies of player i that can lead to the highest expected utility for player i

In RPS, what is Player 1's best response to Rock (i.e., assuming Player 2 plays Rock)?

In Prisoner's Dilemma, what is Player 1's best response to Cooperate? What is Player 1's best response to Defect?

Player 2

Player 1	Cooperate	Defect
Cooperate	-1,-1	-3,0
Defect	0,-3	-2,-2

Best Response

- Best Response (BR): Given the strategies or actions of all players but player i (denoted as $\mathbf{s}_{-i}$ or $\mathbf{a}_{-i}$), Player i's best response to $\mathbf{s}_{-i}$ or $\mathbf{a}_{-i}$ is the set of actions or strategies of player i that can lead to the highest expected utility for player i

In RPS, what is Player 1's best response to Rock (i.e., assuming Player 2 plays Rock)? Paper

In Prisoner's Dilemma, what is Player 1's best response to Cooperate? What is Player 1's best response to Defect?

Defect, Defect

Player 2

Player 1

	Cooperate	Defect
Cooperate	-1,-1	-3,0
Defect	0,-3	-2,-2

Best Response

- Best Response (BR): Given the strategies or actions of all players but player i (denoted as $\mathbf{s}_{-i}$ or $\mathbf{a}_{-i}$), Player i's best response to $\mathbf{s}_{-i}$ or $\mathbf{a}_{-i}$ is the set of actions or strategies of player i that can lead to the highest expected utility for player i

What is Alex's best response to Berry's mixed strategy $s_B = (\frac{1}{2}, \frac{1}{2})$?

Berry

Alex	Football	Concert
Football	2,1	0,0
Concert	0,0	1,2

Best Response

- Best Response (BR): Given the strategies or actions of all players but player i (denoted as $\mathbf{s}_{-i}$ or $\mathbf{a}_{-i}$), Player i's best response to $\mathbf{s}_{-i}$ or $\mathbf{a}_{-i}$ is the set of actions or strategies of player i that can lead to the highest expected utility for player i

What is Alex's best response to Berry's mixed strategy $s_B = (\frac{1}{2}, \frac{1}{2})$?

Football

		Berry	
		Football	Concert
Alex	Football	2,1	0,0
	Concert	0,0	1,2

Piazza Poll 2

- In Rock-Paper-Scissors, if $s_1 = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, which actions or strategies are player 2’s best responses to s_1?
  - A: Rock
  - B: Paper
  - C: Scissors
  - D: $s_2 = \left(\frac{1}{2}, \frac{1}{2}, 0\right)$
  - E: $s_2 = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$
  - F: I don’t know

Player 2

Player 1

	Rock	Paper	Scissors
Rock	0, 0	-1, 1	1, -1
Paper	1, -1	0, 0	-1, 1
Scissors	-1, 1	1, -1	0, 0

Best Response

- Theorem 1 (Nash 1951): A mixed strategy is BR iff all actions in the support are BR

Player 2

Player 1

	Rock	Paper	Scissors
Rock	0, 0	-1, 1	1, -1
Paper	1, -1	0, 0	-1, 1
Scissors	-1, 1	1, -1	0, 0

Dominance

- s_i and $s_i{}'$ are two strategies for player i
- s_i strictly dominates $s_i{}'$if s_i is always better than $s_i{}'$, no matter what strategies are chosen by other players

s_i strictly dominates s_i' if

$$u_i\,(s_i, \mathbf{s}_{-i}) > u_i(s_i', \mathbf{s}_{-i}), \forall \mathbf{s}_{-i}$$

always better

s_i very weakly dominates s_i' if

$$u_i\,(s_i, \mathbf{s}_{-i}) \geq u_i(s_i', \mathbf{s}_{-i}), \forall \mathbf{s}_{-i}$$

never worse

s_i weakly dominates s_i' if

$$u_i\,(s_i, \mathbf{s}_{-i}) \geq u_i(s_i', \mathbf{s}_{-i}), \forall \mathbf{s}_{-i}$$

and $\exists \mathbf{s}_{-i}, u_i\,(s_i, \mathbf{s}_{-i}) > u_i(s_i', \mathbf{s}_{-i})$

never worse and sometimes better

Dominance

Can you find any dominance relationships between the pure strategies in these games?

Player 2

Player 1

	Rock	Paper	Scissors
Rock	0,0	-1,1	1,-1
Paper	1,-1	0,0	-1,1
Scissor	-1,1	1,-1	0,0

Player 2

Player 1

	Cooperate	Defect
Cooperate	-1,-1	-3,0
Defect	0,-3	-2,-2

Berry

Alex

	Football	Concert
Football	2,1	0,0
Concert	0,0	1,2

Dominance

- If s_i strictly dominates $s_i', \forall s_i' \in S_i \backslash \{s_i\}$, is s_i a best response to $\mathbf{s}_{-i}, \forall \mathbf{s}_{-i}$?

Yes, because for any $\mathbf{s}_{-i}$ and for any other strategy s_i'

$$u_i\ (s_i, \mathbf{s}_{-i}) > u_i(s_i', \mathbf{s}_{-i})$$

That is, s_i leads to the highest utility, and is a best response

Solution Concepts in Games

- How should one player play and what should we expect all the players to play?
  - Dominant strategy and dominant strategy equilibrium
  - Nash Equilibrium
  - Minimax strategy
  - Maximin strategy
  - (Stackelberg Equilibrium)

Dominant Strategy

- A strategy that is always better / never worse / never worse and sometimes better than any other strategy
- s_i is a (strictly/very weakly/weakly) dominant strategy if it (strictly/very weakly/weakly) dominates $s_i', \forall s_i' \in S_i \backslash \{s_i\}$
- Focus on single player's strategy
- Not always exist

Is there a strictly dominant strategy for player 1 in PD?

Player 2

		Cooperate	Defect
Player 1	Cooperate	-1,-1	-3,0
	Defect	0,-3	-2,-2

Dominant Strategy Equilibrium

- Sometimes called dominant strategy solution
- Every player plays a dominant strategy
- Focus on strategy profile for all players
- Not always exist

What is the dominant strategy equilibrium for PD?

Player 2

Player 1

	Cooperate	Defect
Cooperate	-1,-1	-3,0
Defect	0,-3	-2,-2

Solution Concepts in Games

- How should one player play and what should we expect all the players to play?
  - Dominant strategy and dominant strategy equilibrium
  - Nash Equilibrium
  - Minimax strategy
  - Maximin strategy
  - (Stackelberg Equilibrium)

Nash Equilibrium

- Nash Equilibrium (NE)
  - Every player’s strategy is a best response to others’ strategy profile
  - In other words, one cannot gain by unilateral deviation
  - Pure Strategy Nash Equilibrium (PSNE)
    - $a_i \in BR(\mathbf{a}_{-i}), \forall i$
  - Mixed Strategy Nash Equilibrium
    - At least one player use a randomized strategy
    - $s_i \in BR(\mathbf{s}_{-i}), \forall i$

Nash Equilibrium

What are the PSNEs in these games?

What is the mixed strategy NE in RPS?

Player 2

Player 1

	Rock	Paper	Scissors
Rock	0,0	-1,1	1,-1
Paper	1,-1	0,0	-1,1
Scissor	-1,1	1,-1	0,0

Player 2

Player 1

	Cooperate	Defect
Cooperate	-1,-1	-3,0
Defect	0,-3	-2,-2

Berry

Alex

	Football	Concert
Football	2,1	0,0
Concert	0,0	1,2

Nash Equilibrium

What are the PSNEs in these games?

RPS: None. Prisoner's Dilemma: (D,D). Football vs Concert: (F,F),(C,C)

What is the mixed strategy NE in RPS?

(1/3,1/3,1/3) for both players

Player 2

Player 1

	Rock	Paper	Scissors
Rock	0,0	-1,1	1,-1
Paper	1,-1	0,0	-1,1
Scissor	-1,1	1,-1	0,0

Player 2

Player 1

	Cooperate	Defect
Cooperate	-1,-1	-3,0
Defect	0,-3	-2,-2

Berry

Alex

	Football	Concert
Football	2,1	0,0
Concert	0,0	1,2

Nash Equilibrium

- **Theorem 2** (Nash 1951): NE always exists in finite games
  - Finite number of players, finite number of actions
  - NE: can be pure or mixed
  - Proof: Through Brouwer's fixed point theorem

Find PSNE

- Find pure strategy Nash Equilibrium (PSNE)
  - Enumerate all action profile
  - For each action profile, check if it is NE
    - For each player, check other available actions to see if he should deviate
  - Other approaches?

Player 2

Player 1

	L	C	R
U	10,3	1,5	5,4
M	3,1	2,4	5,2
D	0,10	1,8	7,0

Find PSNE

- A strictly dominated strategy is one that is always worse than some other strategy
- Strictly dominated strategies cannot be part of an NE Why?

Which are the strictly dominated strategies for player 1?
How about player 2?

Player 2

Player 1

	L	C	R
U	10,3	1,5	5,4
M	3,1	2,4	5,2
D	0,10	1,8	7,0

Find PSNE

- A strictly dominated strategy is one that is always worse than some other strategy
- Strictly dominated strategies cannot be part of an NE   Why?

Which are the strictly dominated strategies for player 1?
How about player 2?

Player 2

Player 1

	L	C	R
U	10,3	1,5	5,4
M	3,1	2,4	5,2
D	0,10	1,8	7,0

≥

Find PSNE through Iterative Removal

- Remove strictly dominated actions (pure strategies) and then find PSNE in the remaining game

Can have new strictly dominated actions in the remaining game!

- Repeat the process until no actions can be removed
- This is the Iterative Removal algorithm (also known as Iterative Elimination of Strictly Dominated Strategies)

Find PSNE in this game using iterative removal

Player 2

Player 1

	L	C	R
U	10,3	1,5	5,4
M	3,1	2,4	5,2
D	0,10	1,8	7,0

Find PSNE through Iterative Removal

- Remove strictly dominated actions (pure strategies) and then find PSNE in the remaining game

Can have new strictly dominated actions in the remaining game!

- Repeat the process until no actions can be removed
- This is the Iterative Removal algorithm (also known as Iterative Elimination of Strictly Dominated Strategies)

Find PSNE in this game using iterative removal

	L	C	R
U	10,3	1,5	5,4
M	3,1	2,4	5,2
D	0,10	1,8	7,0

	L	C
U	10,3	1,5
M	3,1	2,4
D	0,10	1,8

	L	C
U	10,3	1,5
M	3,1	2,4

	C
U	1,5
M	2,4

	C
M	2,4

Find PSNE through Iterative Removal

- When the algorithm terminates, if the remaining game has only one action for each player, then that is the unique NE of the game and the game is called dominance solvable
  - It may not be a dominant strategy equilibrium
- When the remaining game has more than one action for some players, find PSNE in the remaining game
- Order of removal does not matter

Player 2

Player 1

	L	C	R
U	10,3	1,5	5,4
M	3,1	2,4	5,2
D	0,10	1,8	7,0

Find Mixed Strategy Nash Equilibrium

- Can we still apply iterative removal?
  - Yes! The removed strategies cannot be part of any NE
  - You can always apply iterative removal first

Find Mixed Strategy Nash Equilibrium

- How to find mixed strategy NE (after iterative removal)?

Berry

Alex

	Football	Concert
Football	2,1	0,0
Concert	0,0	1,2

If $s_A = (p, 1-p)$ and $s_B = (q, 1-q)$ with $0 < p, q < 1$ is a NE, what are the necessary conditions for p and q?

Find Mixed Strategy Nash Equilibrium

- How to find mixed strategy NE (after iterative removal)?

Berry

		Football	Concert
Alex	Football	2,1	0,0
	Concert	0,0	1,2

Is $s_A = (\frac{2}{3}, \frac{1}{3})$ and $s_B = (\frac{1}{3}, \frac{2}{3})$ an NE?

Find Mixed Strategy Nash Equilibrium

- How to find mixed strategy NE (after iterative removal)?

Berry

Alex

	Football	Concert
Football	2,1	0,0
Concert	0,0	1,2

Is $s_A = (\frac{2}{3}, \frac{1}{3})$ and $s_B = (\frac{1}{3}, \frac{2}{3})$ an NE?

$$u_A(s_A, s_B) = \frac{2}{3} * \frac{1}{3} * 2 + \frac{1}{3} * \frac{2}{3} * 1 = 2/3$$

$$u_A(F, s_B) = 2 * \frac{1}{3} = \frac{2}{3}, u_A(C, s_B) = 1 * \frac{2}{3} = \frac{2}{3},$$

So for any strategy $s_A' = (\epsilon, 1 - \epsilon), u_A(s_A', s_B) = \epsilon u_A(F, s_B) + (1 - \epsilon) u_A(C, s_B) = 2/3$

So Alex has no incentive to deviate (u_A cannot increase). Similar reasoning goes for u_B

Find Mixed Strategy Nash Equilibrium

- How to find mixed strategy NE (after iterative removal)?

Berry

Alex

	Football	Concert
Football	2,1	0,0
Concert	0,0	1,2

If $s_A = (p, 1 - p)$ and $s_B = (q, 1 - q)$ with $0 < p, q < 1$ is a NE, what are the necessary conditions for p and q?

Find Mixed Strategy Nash Equilibrium

- How to find mixed strategy NE (after iterative removal)?

Berry

Alex

	Football	Concert
Football	2,1	0,0
Concert	0,0	1,2

If $s_A = (p, 1-p)$ and $s_B = (q, 1-q)$ with $0 < p, q < 1$ is a NE, what are the necessary conditions for p and q?

$$u_A(F, s_B) = u_A(C, s_B)$$
$$u_B(s_A, C) = u_B(s_A, C)$$

Why?

For Alex: $s_A \in BR(s_B)$, so according to Theorem 1, F$\in BR(s_B)$, C$\in BR(s_B)$, so $u_A(F, s_B) = u_A(C, s_B)$

Find Mixed Strategy Nash Equilibrium

- How to find mixed strategy NE (after iterative removal)?

Berry

Alex

	Football	Concert
Football	2,1	0,0
Concert	0,0	1,2

If $s_A = (p, 1-p)$ and $s_B = (q, 1-q)$ with $0 < p, q < 1$ is a NE, what are the necessary conditions for p and q?

$u_A(F, s_B) = u_A(C, s_B)$ $u_B(s_A, C) = u_B(s_A, C)$

Find Mixed Strategy Nash Equilibrium

- How to find mixed strategy NE (after iterative removal)?

Berry

Alex

	Football	Concert
Football	2,1	0,0
Concert	0,0	1,2

If $s_A = (p, 1-p)$ and $s_B = (q, 1-q)$ with $0 < p, q < 1$ is a NE, what are the necessary conditions for p and q?

$u_A(F, s_B) = u_A(C, s_B)$ $\quad$ $u_B(s_A, C) = u_B(s_A, C)$

$$u_A(F, s_B) = 2 * q + 0 * (1-q) = u_A(C, s_B) = 0 * q + 1 * (1-q)$$

So $2q = 1 - q$, we get $q = \frac{1}{3}$

Similarly, $u_B(s_A, F) = 1 * p + 0 * (1-p) = u_B(s_A, C) = 0 * p + 2 * (1-p)$

So $p = 2(1-p)$, we get $p = \frac{2}{3}$

Piazza Poll 3

- If $s_A = (p, 1-p)$ and $s_B = (q, 1-q)$ with $0 < p, q < 1$ is a NE of the game, which equations should p and q satisfy?
- A: $2q = 3(1-q)$
- B: $2p = 3(1-p)$
- C: $q = 2(1-q)$
- D: $p = 2(1-p)$
- E: I don't know

Berry

Alex

	Football	Concert
Football	2,1	0,0
Concert	0,0	3,2

Piazza Poll 3

- If $s_A = (p, 1-p)$ and $s_B = (q, 1-q)$ with $0 < p, q < 1$ is a NE of the game, which equations should p and q satisfy?
- A: $2q = 3(1-q)$
- B: $2p = 3(1-p)$
- C: $q = 2(1-q)$
- D: $p = 2(1-p)$
- E: I don't know

$$u_A(F, s_B) = u_A(C, s_B)$$

$$u_B(s_A, C) = u_B(s_A, C)$$

$$u_A(F, s_B) = 2 * q + 0 * (1-q)$$
$$u_A(C, s_B) = 0 * q + 3 * (1-q)$$

$$u_B(s_A, F) = 1 * p + 0 * (1-p)$$
$$u_B(s_A, C) = 0 * p + 2 * (1-p)$$

Berry

Alex

	Football	Concert
Football	2,1	0,0
Concert	0,0	3,2

Solution Concepts in Games

- How should one player play and what should we expect all the players to play?
  - Dominant strategy and dominant strategy equilibrium
  - Nash Equilibrium
  - Minimax strategy
  - Maximin strategy
  - (Stackelberg Equilibrium)

Maximin and Minimax Strategy

- Both focus on single player's strategy

- Maximin Strategy
  - Maximize worst case expected utility
  - Maximin value (also called safety level)

- Minimax Strategy
  - Minimize best case expected utility for the other player (just want to harm your opponent)
  - Minimax value

Minimax Theorem

- Theorem 3 (von Neumann 1928, Nash 1951):
  - Minimax=Maximin=NE in 2-player zero-sum games
  - All NEs leads to the same utility profile in a two-player zero-sum game

  - Formally, every two-player zero-sum game has a unique value v such that
    - Player 1 can guarantee value at least v
    - Player 2 can guarantee loss at most v
    - v is called value of the game or game value

Solution Concepts in Games

- How should one player play and what should we expect all the players to play?
  - Dominant strategy and dominant strategy equilibrium
  - Nash Equilibrium
  - Minimax strategy
  - Maximin strategy
  - (Stackelberg Equilibrium)

Power of Commitment

- What's the PSNEs in this game and the players' utilities?
- What action should player 2 choose if player 1 commits to playing b? What is player 1's utility?
- What action should player 2 choose if player 1 commits to playing a and b uniformly randomly? What is player 1's utility?

Player 2

Player 1

	c	d
a	2,1	4,0
b	1,0	3,2

Power of Commitment

- What's the PSNEs in this game and the players' utilities? (a, c), 2
- What action should player 2 choose if player 1 commits to playing b? What is player 1's utility? d, 3
- What action should player 2 choose if player 1 commits to playing a and b uniformly randomly? What is player 1's utility?

d, 3.5

Player 2

Player 1

	c	d
a	2,1	4,0
b	1,0	3,2

Stackelberg Equilibrium

- Stackelberg Game
  - Leader commits to a strategy first
  - Follower responds after observing the leader's strategy
- Stackelberg Equilibrium
  - Follower best responds to leader's strategy
  - Leader commits to a strategy that maximize her utility assuming follower best responds

Player 2

Player 1

	c	d
a	2,1	4,0
b	1,0	3,2

Stackelberg Equilibrium

- If the leader can only commit to a pure strategy, or you know that the leader's strategy in equilibrium is a pure strategy, the equilibrium can be found by enumerating leader's pure strategy
- In general, the leader can commit to a mixed strategy and $u^{SSE} \geq u^{NE}$ (first-mover advantage)!

		Berry	
		Football	Concert
Alex	Football	2,1	0,0
	Concert	0,0	1,2

		Player 2	
		c	d
Player 1	a	2,1	4,0
	b	1,0	3,2

Warm-up

- Design an AI to play Rock-Paper-Scissors for T rounds

If T=2, roundID=1, what action does your AI choose in each round?

```
function playRPS(roundID,T)
Input: $T$, roundID ∈ $\{1..T\}$
Output: action a ∈ {Rock, Paper, Scissors}

Return a
```

Extensive-Form Games

- Normal-form game: players, actions, utilities for each action profile, simultaneous movement
- Extensive-form game
  - Set of players
  - Sequencing of players' possible moves
  - Player’s actions at every decision point
  - (Possibly imperfect) information each player has about the other player's moves when they make a decision
  - Payoffs for all possible game outcomes

Extensive-Form Games

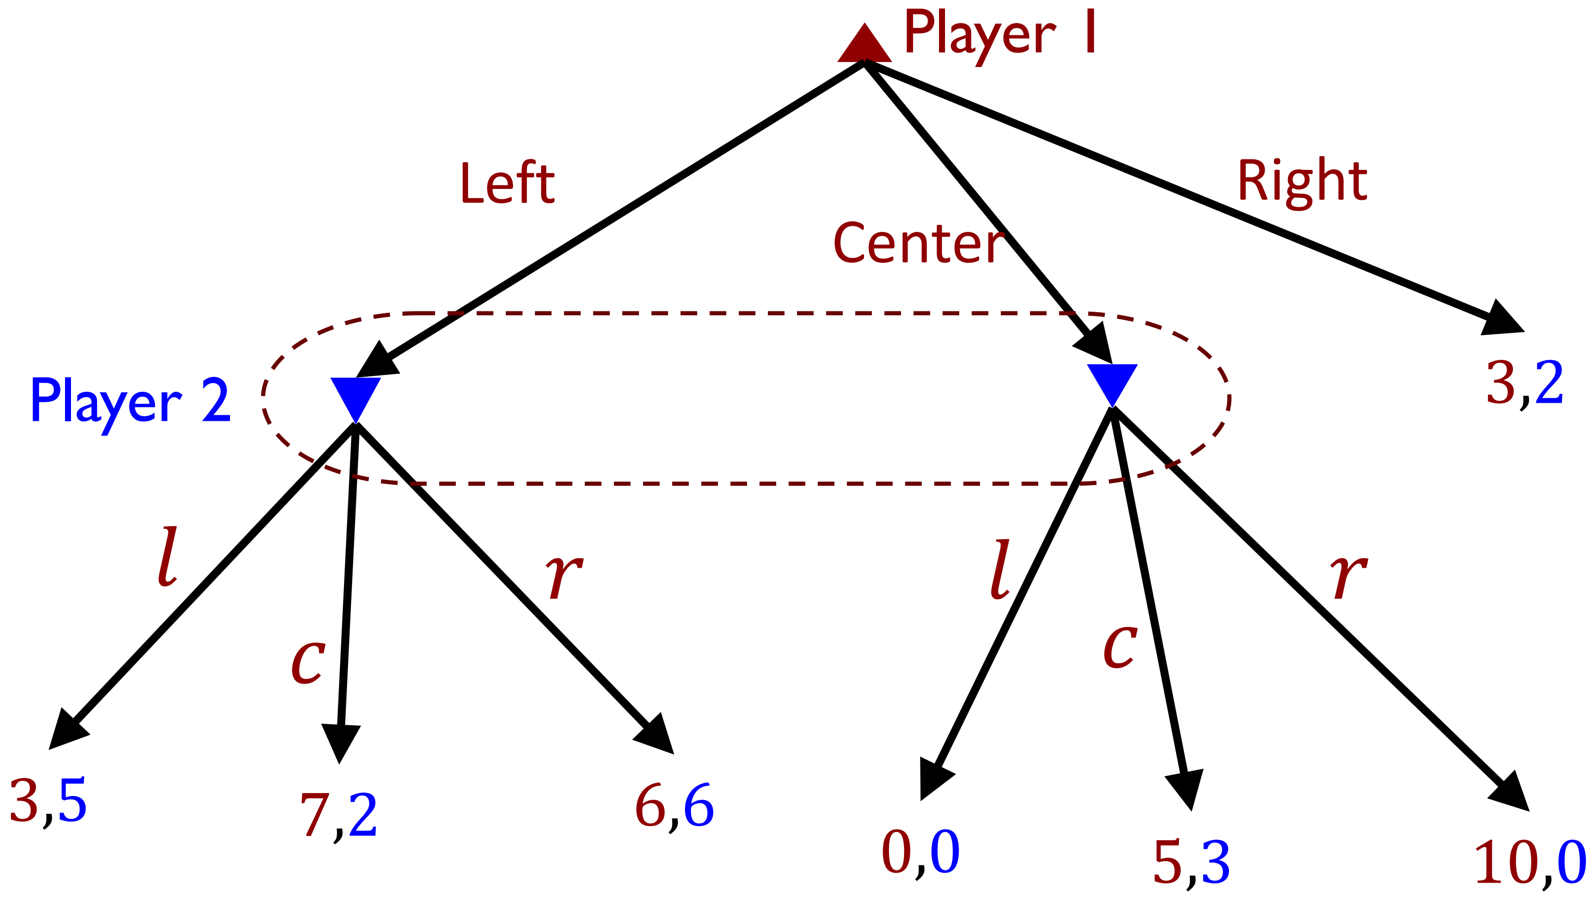

Can we represent this game in normal form?

Can we solve this game using algs for adversarial search?

Protecting Staten Island Ferry

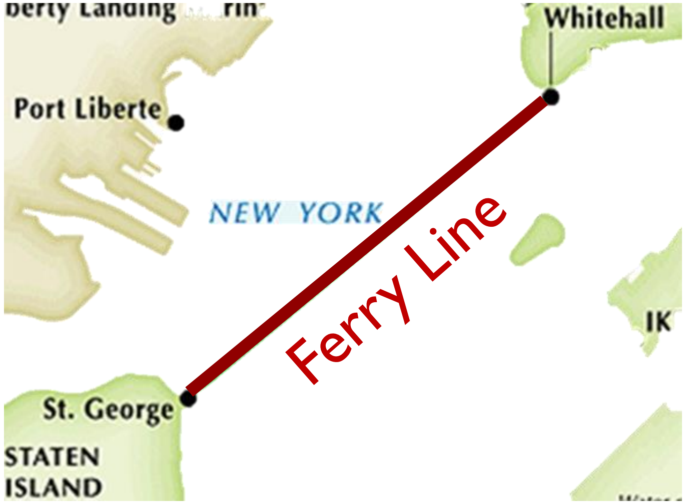

Protecting Staten Island Ferry

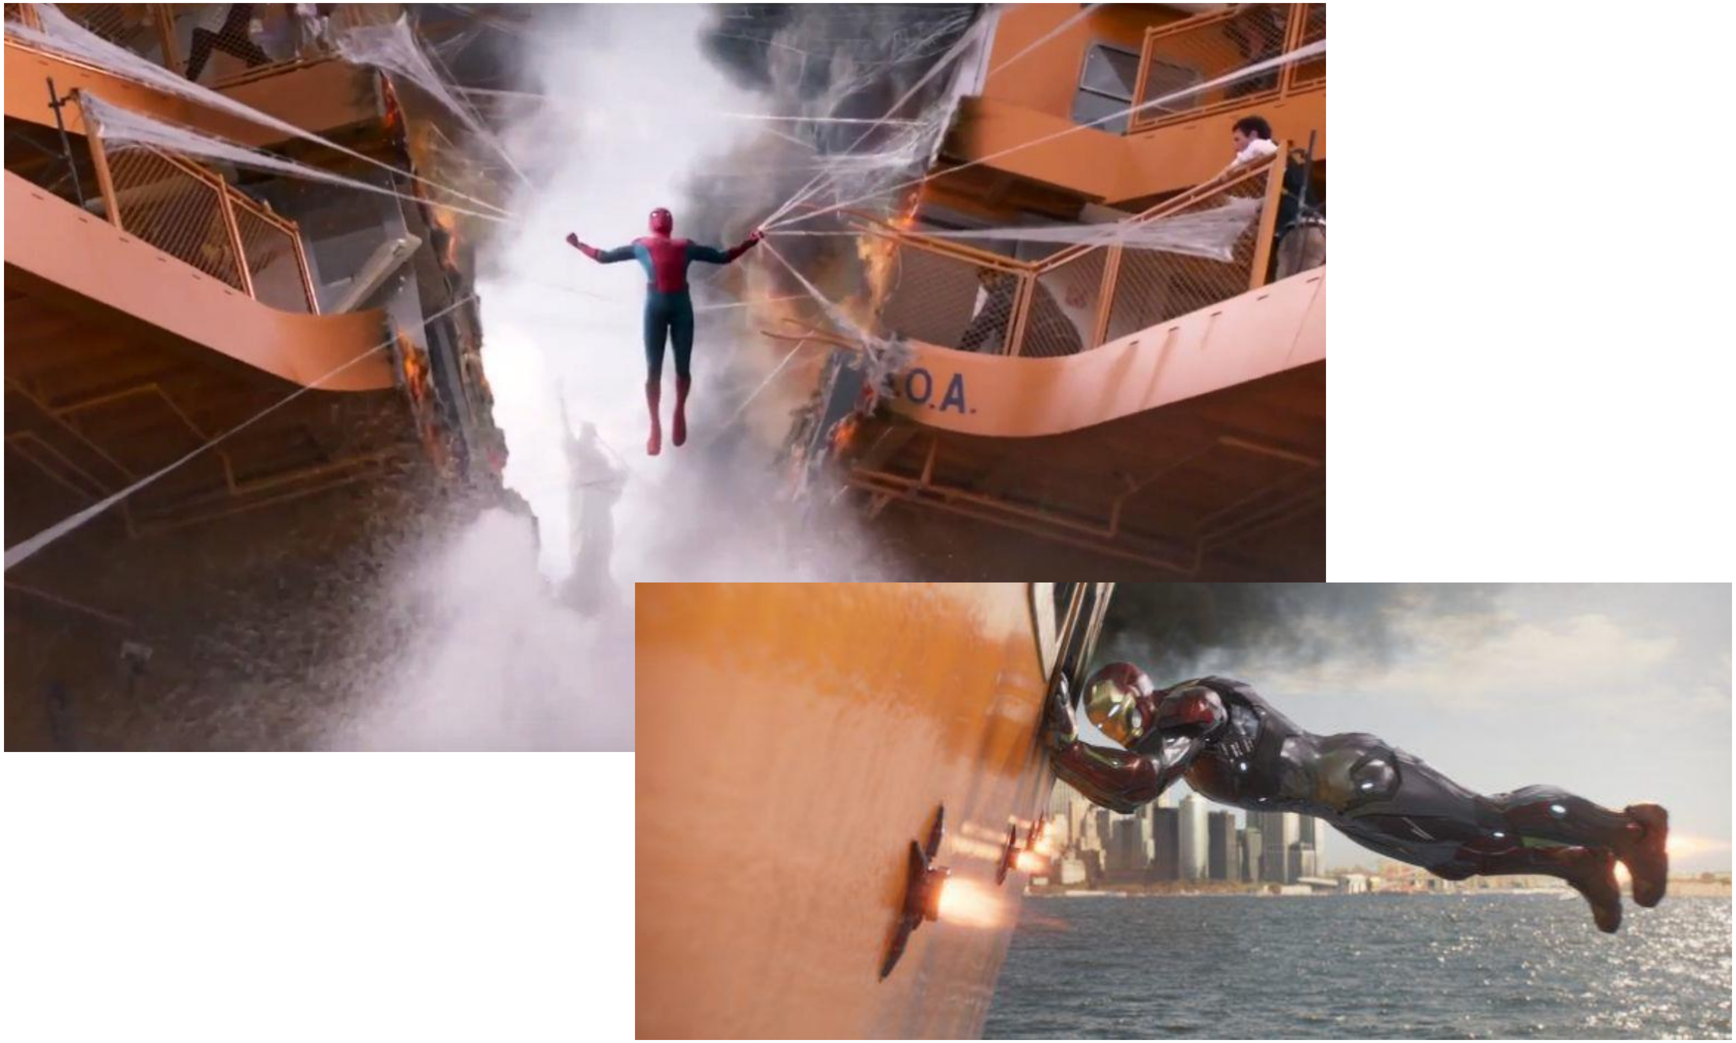

Protecting Staten Island Ferry

Previous USCG Approach

Game-Theoretic Patrols

In collaboration with US Coast Guard (Contact: Joe Direnzo, Craig Baldwin)

Problem

Optimal Patrol Strategy for Protecting Moving Targets with Multiple Mobile Resources
Fei Fang, Albert Xin Jiang, Milind Tambe
In AAMAS-13: The Twelfth International Conference on Autonomous Agents and Multiagent Systems, May 2013

Game Model and Linear Programming-based Solution

- Stackelberg game: Leader – Defender, Follower – Attacker
- Attacker's payoff: $u_i(t)$ if not protected, 0 otherwise
- Zero-sum → Strong Stackelberg Equilibrium=Nash Equilibrium =Minimax (Minimize Attacker's Maximum Expected Utility)

$$\min_{p_r, v} v$$

$$\text{s.t. } v \geq \mathbb{E}[U^{att}(i,t)] \;= u_i(t) \times \mathbb{P}[unprotected(i,t)], \forall i, t$$

Adversary

Defender

p_r

		10:00:00 AM Target 1	10:00:01 AM Target 1	…	10:30:00 AM Target 3	…
30%	Purple Route	-5, 5	-4, 4		0, 0	
40%	Orange Route					
20%	Blue Route					
	……					

$$\sum_r p_r \leq 1$$

Evaluation: Simulation & Real-World Feedback

Reduce potential risk by 50%

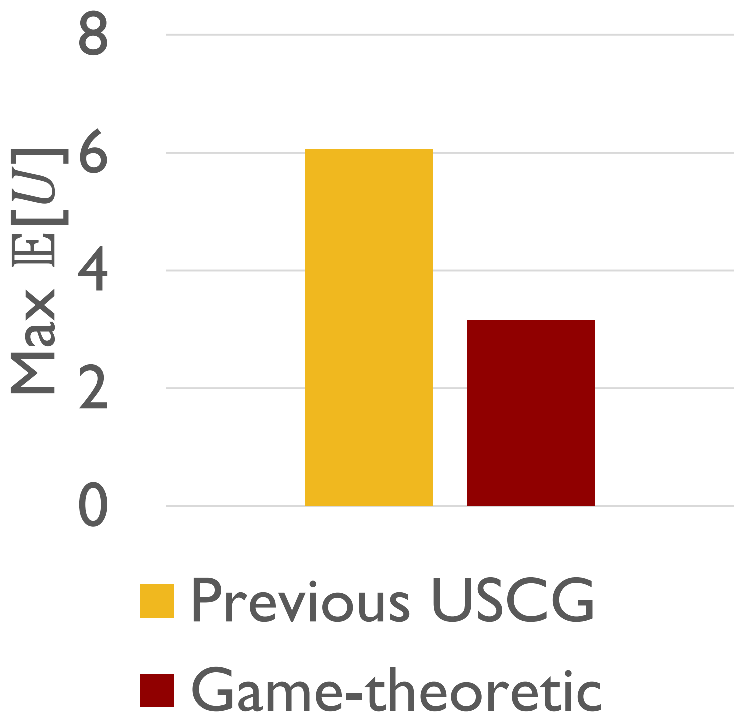

- Deployed by US Coast Guard
- USCG evaluation
  - Point defense to zone defense
  - Increased randomness
- Professional mariners:
  - Apparent increase in Coast Guard patrols