CSC304 Lecture 7

Game Theory : Security games, Applications to security

Until now…

- Simultaneous-move Games

- All players act simultaneously

- Nash equilibria = stable outcomes

- Each player is best responding to the strategies of all other players

Sequential Move Games

- Focus on two players: "leader" and "follower"

1. Leader commits to a (possibly mixed) strategy x_1
   - Cannot change later
2. Follower learns about x_1
   - Follower must believe that leader's commitment is credible
3. Follower chooses the best response x_2
   - Can assume to be a pure strategy without loss of generality
   - If multiple actions are best response, break ties in favor of the leader

Sequential Move Games

- Wait. Does this give us anything new?
  - Can't I, as player 1, commit to playing x_1 in a simultaneous-move game too?
  - Player 2 wouldn't believe you.

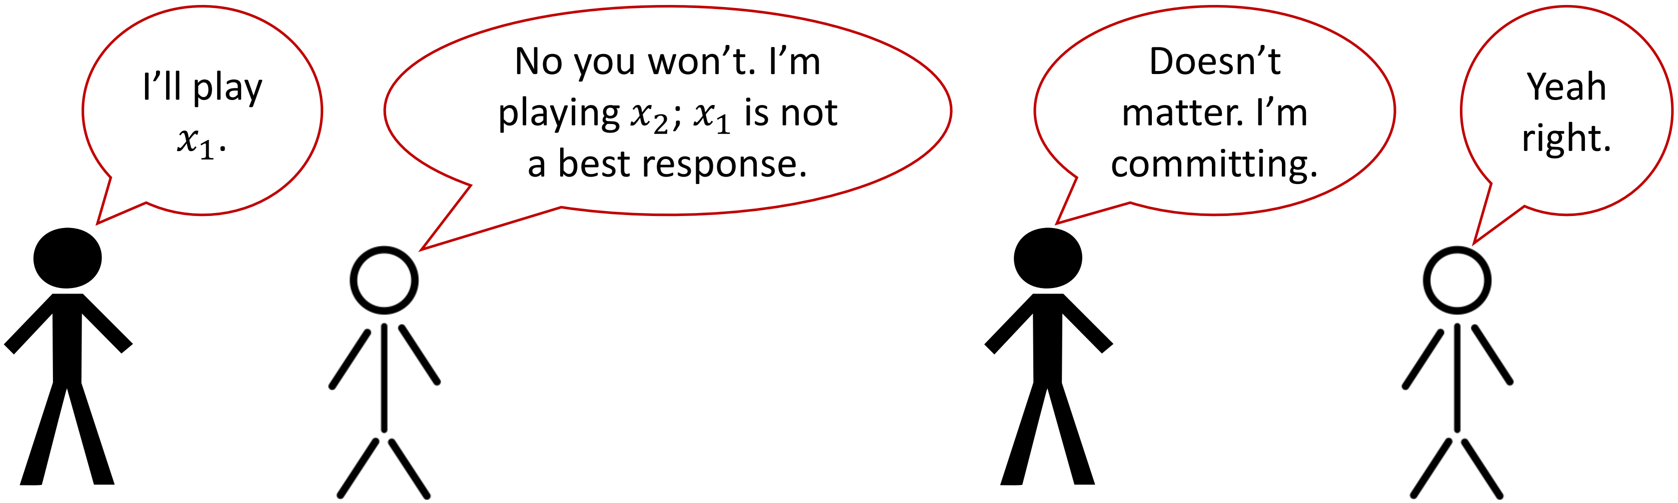

That’s unless…

- You’re as convincing as this guy.

How to represent the game?

- Extensive form representation
  - Can also represent "information sets", multiple moves, …

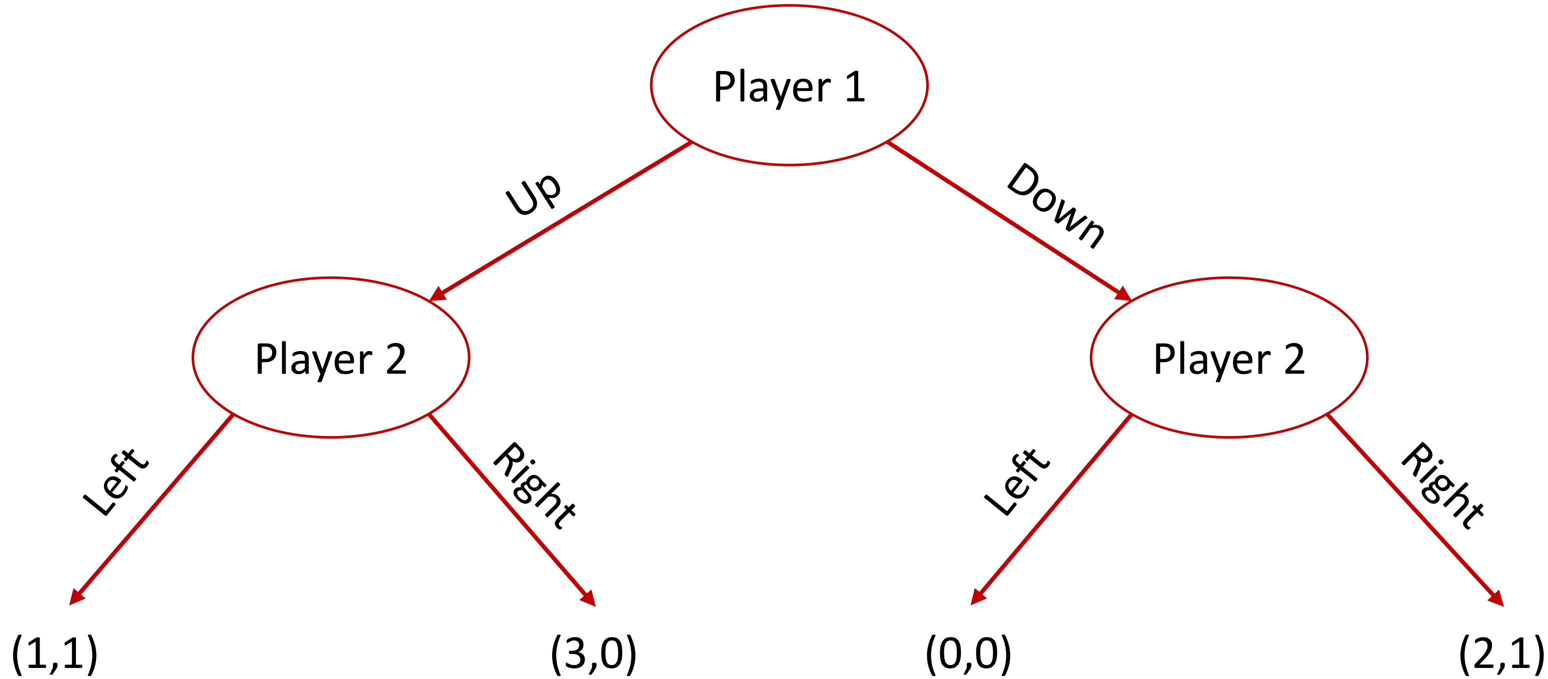

A Curious Case

P2 / P1	Left	Right
Up	(1 , 1)	(3 , 0)
Down	(0 , 0)	(2 , 1)

- Q: What are the Nash equilibria of this game?
- Q: You are P1. What is your reward in Nash equilibrium?

A Curious Case

P1 \ P2	Left	Right
Up	(1 , 1)	(3 , 0)
Down	(0 , 0)	(2 , 1)

- Q: As P1, you want to commit to a pure strategy. Which strategy would you commit to?

- Q: What would your reward be now?

Commitment Advantage

P2 / P1	Left	Right
Up	(1 , 1)	(3 , 0)
Down	(0 , 0)	(2 , 1)

- Reward in the unique Nash equilibrium = 1
- Reward when committing to Down = 2

Commitment Advantage

P2 / P1	Left	Right
Up	(1 , 1)	(3 , 0)
Down	(0 , 0)	(2 , 1)

- Higher reward in committing to a mixed strategy
  - P1 commits to: Up w.p. $0.5 - \epsilon$, Down w.p. $0.5 + \epsilon$
  - P2 is still better off playing Right
  - $\mathbb{E}[\text{Reward}]$ to P1 ≈ 2.5
  - Note: If P1 plays both actions with probability exactly 0.5, we assume P2 plays Right (break ties in favor of leader)

Stackelberg vs Nash

- Committing first is always better than playing a simultaneous-move game?
- Yes!
  - If (x_1^*, x_2^*) is a NE, P1 can always commit to x_1^*, ensure that P2 will play x_2^*, and achieve the reward in the NE
  - P1 may be able to commit to a better strategy than x_1^*

- Applications to security
  - Law enforcement is better off committing to a mixed patrolling strategy, and announcing the strategy publicly!

Stackelberg in Zero-Sum

- Recall the minimax theorem:

$$\max_{x_1} \min_{x_2} x_1^T A \, x_2 = \min_{x_2} \max_{x_1} x_1^T A \, x_2$$

- P1 goes first → P1 chooses her minimax strategy
- P2 goes first → P2 chooses her minimax strategy
- Minimax Theorem: It doesn’t make a difference!
  - Simultaneous-move, P1 going first, and P2 going first are essentially identical scenarios.

Stackelberg in General-Sum

- 2-player non-zero-sum game with reward matrices A and $B \neq -A$ for the two players

$$\max_{x_1} \; x_1^T \, A \;\; f(x_1)$$

$$\text{where } f(x_1) = \operatorname*{argmax}_{x_2} \; x_1^T \, B \;\; x_2$$

- How do we compute this?

Example

P2 / P1	Left	Right
Up	(1 , 1)	(3 , 0)
Down	(0 , 0)	(2 , 1)

- Let us separately maximize the reward of P1 in 2 cases:
  - Strategies that cause P2 to play Left
  - Strategies that cause P2 to play Right

- Suppose P1 commits to Up w.p. p, Down w.p. $1 - p$

Example

P2 / P1	Left	Right
Up	(1 , 1)	(3 , 0)
Down	(0 , 0)	(2 , 1)

- Strategies that cause P2 to play Left

$$\text{Max} \quad p \cdot 1 + (1-p) \cdot 0$$
$$s.t.$$
$$p \cdot 1 + (1-p) \cdot 0 \geq p \cdot 0 + (1-p) \cdot 1$$
$$p \in [0,1]$$

Reward of P1 assuming P2 plays Left

Condition that causes P2 to play Left

Example

P2 / P1	Left	Right
Up	(1 , 1)	(3 , 0)
Down	(0 , 0)	(2 , 1)

- Strategies that cause P2 to play Left

$$\text{Max} \quad p$$
$$s.t.$$
$$p \geq 1 - p$$
$$p \in [0,1]$$

Answer=1

Example

P2 / P1	Left	Right
Up	(1 , 1)	(3 , 0)
Down	(0 , 0)	(2 , 1)

- Strategies that cause P2 to play Right

$$\text{Max} \quad p \cdot 3 + (1-p) \cdot 2$$

Answer=2.5

$$s.t.$$

$$p \cdot 1 + (1-p) \cdot 0 \leq p \cdot 0 + (1-p) \cdot 1$$

$$p \in [0,1]$$

Stackelberg via LPs

- High-level Idea:
  - For each action s_2^* of P2…
  - Write a *linear program* with the mixed strategy x_1 of P1 as the unknown, which…
  - Maximizes the reward of P1 when P1 plays x_1, P2 responds with s_2^*…
  - Subject to the constraint that x_1 in fact incentivizes P2 to play s_2^*

Stackelberg via LPs

- S_1, S_2 = sets of actions of leader and follower
- $|S_1| = m_1, |S_2| = m_2$
- $x_1(s_1)$ = probability of leader playing s_1
- π_1, π_2 = reward functions for leader and follower

$\max \Sigma_{s_1 \in S_1} x_1(s_1) \cdot \pi_1(s_1, s_2^*)$

subject to

$\forall s_2 \in S_2, \;\; \Sigma_{s_1 \in S_1} x_1(s_1) \cdot \pi_2(s_1, s_2^*) \;\geq$

$\Sigma_{s_1 \in S_1} x_1(s_1) \cdot \pi_2(s_1, s_2)$

$\Sigma_{s_1 \in S_1} x_1(s_1) = 1$

$\forall s_1 \in S_1, x_1(s_1) \geq 0$

- One LP for each s_2^*, take the maximum over all m_2 LPs
- The LP corresponding to s_2^* optimizes over all x_1 for which s_2^* is the best response

Real-World Applications

- Security Games
  - Defender (leader) has k identical patrol units
  - Defender wants to defend a set of n targets T
  - In a pure strategy, each resource can protect a subset of targets $S \subseteq T$ from a given collection $\mathcal{S}$
  - A target is covered if it is protected by at least one resource
  - Attacker wants to select a target to attack

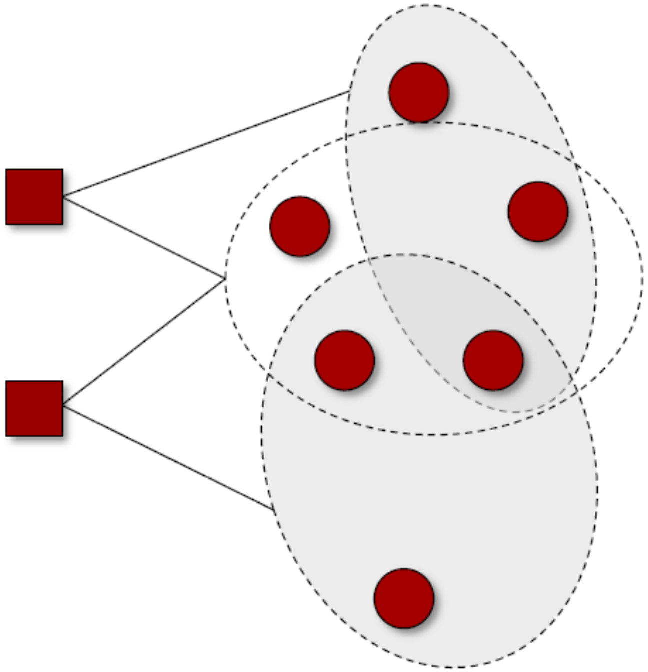

Real-World Applications

- Security Games
  - For each target, the defender and the attacker have two utilities: one if the target is covered, one if it is not.
  - Defender commits to a mixed strategy; attacker follows by choosing a target to attack.

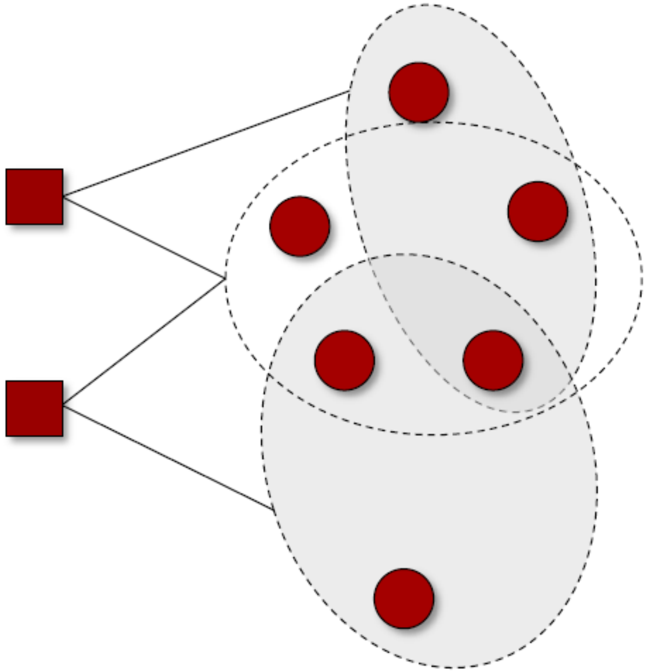

Ah!

- Q: Because this is a 2-player Stackelberg game, can we just compute the optimal strategy for the defender in polynomial time…?

- Time is polynomial in the number of pure strategies of the defender
  - In security games, this is $|\mathcal{S}|^k$
  - Exponential in k

- Intricate computational machinery required…

LAX

Newsweek National News

Subscribe Now | Make Newsweek Your Homepage | Newsletters | RSS

The Element of Surprise

To help combat the terrorism threat, officials at Los Angeles Inter Airport are introducing a bold new idea into their arsenal: randon of security checkpoints. Can game theory help keep us safe?

WEB EXCLUSIVE

By Andrew Murr
Newsweek
Updated: 1:00 p.m. PT Sept 28, 2007

Security forces work the sidewalk

Sept. 28, 2007 - Security officials at Los Angeles International Airport now have a new weapon in their fight against terrorism: complete, baffling randomness. Anxious to thwart future terror attacks in the early stages while plotters are casing the airport, LAX security patrols have begun using a new software program called ARMOR, NEWSWEEK has learned, to make the placement of security checkpoints completely unpredictable. Now all airport security officials have to do is press a button labeled "Randomize," and they can throw a sort of digital cloak of invisibility over where they place the cops' antiterror checkpoints on any given day.

Real-World Applications

- Protecting entry points to LAX
- Scheduling air marshals on flights
  - Must return home
- Protecting the Staten Island Ferry
  - Continuous-time strategies
- Fare evasion in LA metro
  - Bathroom breaks !!!
- Wildlife protection in Ugandan forests
  - Poachers are not fully rational
- Cyber security

…