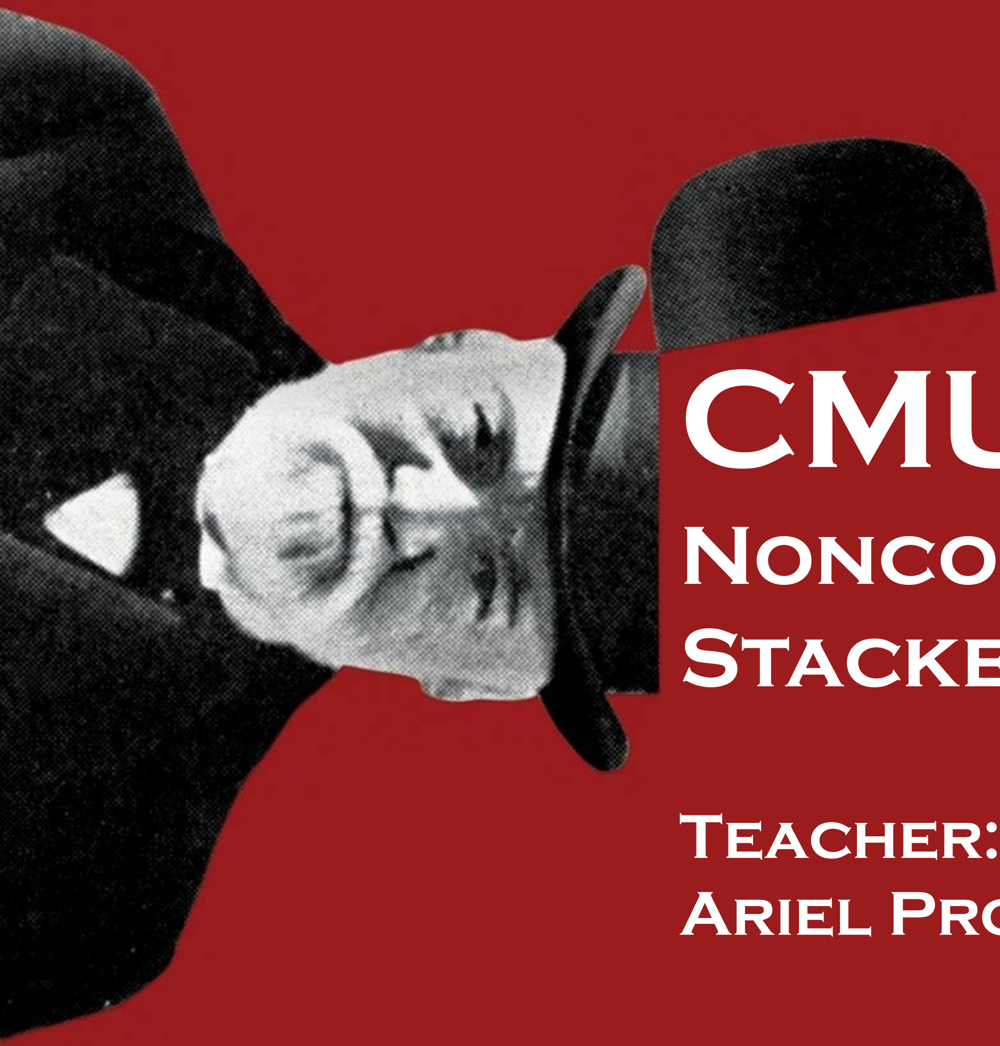

CMU 15-896

Noncooperative games 4: Stackelberg games

Teacher:
Ariel Procaccia

A curious game

- Playing up is a dominant strategy for row player
- So column player would play left
- Therefore, (1,1) is the only Nash equilibrium outcome

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

Commitment is good

- Suppose the game is played as follows:
  - Row player commits to playing a row
  - Column player observes the commitment and chooses column
- Row player can commit to playing down!

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

Commitment to mixed strategy

- By committing to a mixed strategy, row player can guarantee a reward of 2.5
- Called a Stackelberg (mixed) strategy

	0	1
.49	1,1	3,0
.51	0,0	2,1

Computing Stackelberg

- Theorem [Conitzer and Sandholm 2006]: In 2-player normal form games, an optimal Stackelberg strategy can be found in poly time
- Theorem [ditto]: the problem is NP-hard when the number of players is ≥ 3

Tractability: 2 players

- For each pure follower strategy s_2, we compute via the LP below a strategy x_1 for the leader such that
  - Playing s_2 is a best response for the follower
  - Under this constraint, x_1 is optimal
- Choose x_1^* that maximizes leader value

$$\mathbf{max} \ \sum_{s_1 \in S} x_1(s_1) u_1(s_1, s_2)$$

$$\mathbf{s.t.} \quad \forall s_2' \in S, \ \sum_{s_1 \in S} x_1(s_1) u_2(s_1, s_2) \geq \sum_{s_1 \in S} x_1(s_1) u_2(s_1, s_2')$$

$$\sum_{s_1 \in S} x_1(s_1) = 1$$

$$\forall s_1 \in S, x_1(s_1) \in [0,1]$$

Application: security

- Airport security: deployed at LAX
- Federal Air Marshals
- Coast Guard
- Idea:
  - Defender commits to mixed strategy
  - Attacker observes and best responds

SECURITY GAMES

- Set of targets $T = \{1, \dots, n\}$
- Set of m security resources Ω available to the defender (leader)
- Set of schedules $\Sigma \subseteq 2^T$
- Resource ω can be assigned to one of the schedules in $A(\omega) \subseteq \Sigma$
- Attacker chooses one target to attack

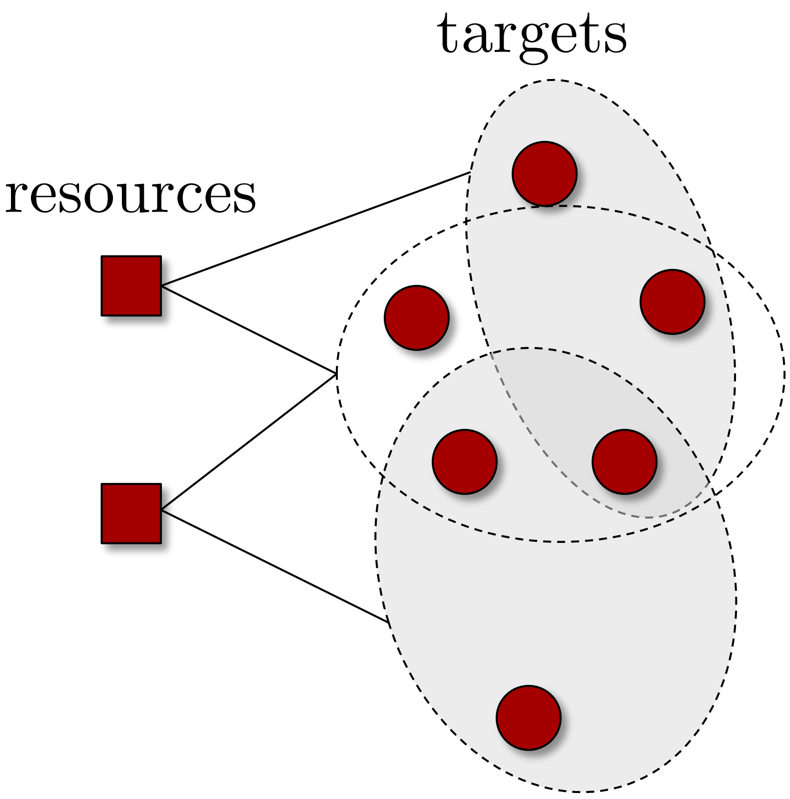

SECURITY GAMES

- For each target t, there are four numbers: $u_d^c(t) \geq u_d^u(t)$, and $u_a^c(t) \leq u_a^u(t)$
- Let $\boldsymbol{c} = (c_1, \ldots, c_n)$ be the vector of coverage probabilities
- The utilities to the defender/attacker under $\mathbf{c}$ if target t is attacked are $u_d(t, \boldsymbol{c}) = u_d^c(t) \cdot c_t + u_d^u(t)(1 - c_t)$
- $u_a(t, \boldsymbol{c}) = u_a^c(t) \cdot c_t + u_a^u(t)(1 - c_t)$

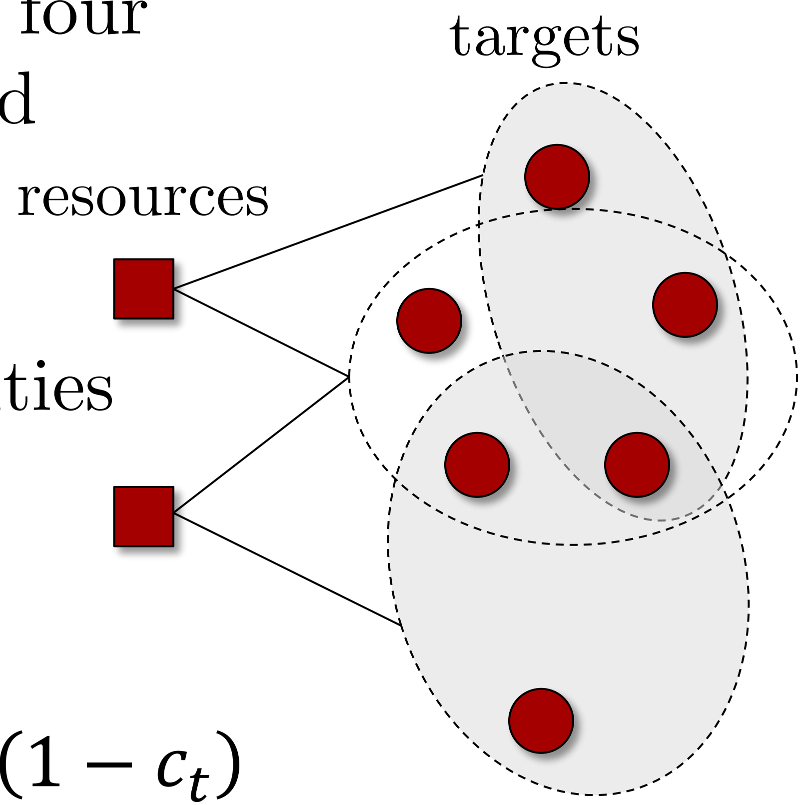

This is a 2-player Stackelberg game. Can we compute an optimal strategy for the defender in polynomial time?

Solving security games

- Consider the case of $\Sigma = T$, i.e., resources are assigned to individual targets, i.e., schedules have size 1
- Nevertheless, number of leader strategies is exponential
- Theorem [Korzhyk et al. 2010]: Optimal leader strategy can be computed in poly time

A compact LP

- LP formulation similar to previous one
- Advantage: logarithmic in #leader strategies
- Problem: do probabilities correspond to strategy?

$$
\begin{aligned}
\textbf{max} \quad & u_d(t^*, c) \\
\textbf{s.t.} \quad & \forall \omega \in \Omega, \forall t \in A(\omega), 0 \leq c_{\omega,t} \leq 1 \\
& \forall t \in T, c_t = \sum_{\omega \in \Omega : t \in A(\omega)} c_{\omega,t} \leq 1 \\
& \forall \omega \in \Omega, \sum_{t \in A(\omega)} c_{\omega,t} \leq 1 \\
& \forall t \in T, u_a(t, \boldsymbol{c}) \leq u_a(t^*, \boldsymbol{c})
\end{aligned}
$$

Fixing the probabilities

- Theorem [Birkhoff-von Neumann]: Consider an $m \times n$ matrix M with real numbers $a_{ij} \in [0,1]$, such that for each i, $\sum_j a_{ij} \leq 1$, and for each j, $\sum_i a_{ij} \leq 1$ (M is kinda doubly stochastic). Then there exist matrices $M^1, \ldots, M^q$ and weights $w^1, \ldots, w^q$ such that:
  1. $\sum_k w^k = 1$
  2. $\sum_k w^k M^k = M$
  3. For each k, M^k is kinda doubly stochastic and its elements are in $\{0,1\}$
- The probabilities $c_{\omega,t}$ satisfy theorem's conditions
- By 3, each M^k is a deterministic strategy
- By 1, we get a mixed strategy
- By 2, gives right probs

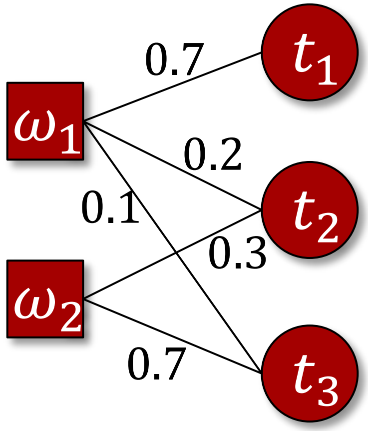

	t_1	t_2	t_3
ω_1	0.7	0.2	0.1
ω_2	0	0.3	0.7

<u>0.1</u>

	t_1	t_2	t_3
ω_1	0	0	1
ω_2	0	1	0

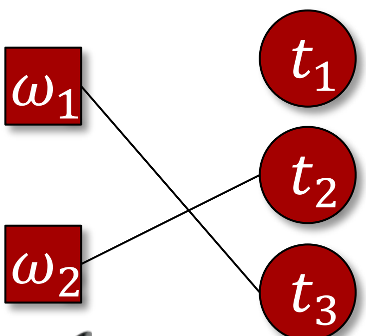

<u>0.2</u>

	t_1	t_2	t_3
ω_1	0	1	0
ω_2	0	0	1

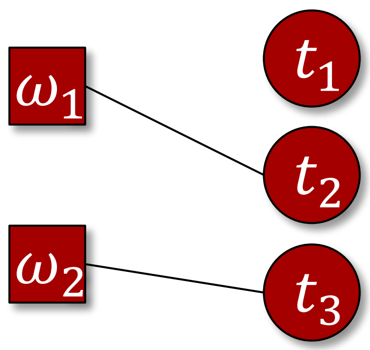

<u>0.2</u>

	t_1	t_2	t_3
ω_1	1	0	0
ω_2	0	1	0

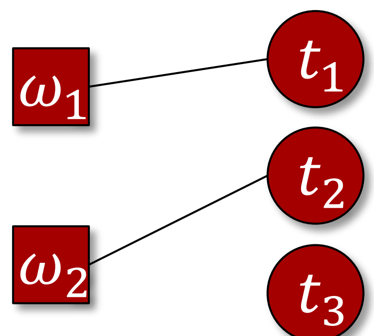

<u>0.5</u>

	t_1	t_2	t_3
ω_1	1	0	0
ω_2	0	0	1

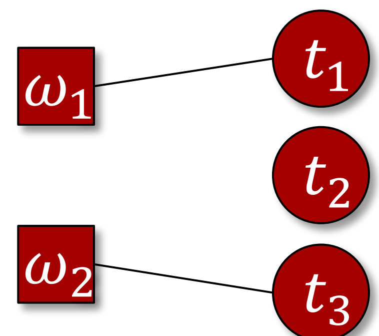

Generalizing?

- What about schedules of size 2?
- Air Marshals domain has such schedules: outgoing+incoming flight (bipartite graph)
- Previous apporoach fails
- Theorem [Korzhyk et al. 2010]: problem is NP-hard

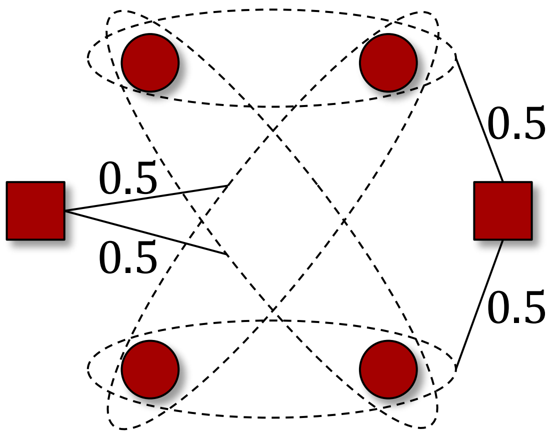

Newsweek National News

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: random
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.

Criticisms

- Problematic assumptions:
  1. The attacker exactly observes the defender's mixed strategy
  2. The defender knows the attacker's utility function
  3. The attacker behaves in a perfectly rational way
- We will focus on relaxing assumption #1

Limited surveillance

- Let us compare two worlds:
  1. Status quo: The defender optimizes against an attacker with unlimited observations (i.e., complete knowledge of the defender's strategy), but the attacker actually has only k observations
  2. Ideal: The defender optimizes against an attacker with k observations, and, miraculously, the attacker indeed has exactly k observations

Limited surveillance

- Theorem [Blum et al. 2014]: Assume that utilities are normalized to be in $[-1,1]$. For any m, d, k such that $2md \geq \binom{2k}{k}$, and any $\epsilon > 0$, there is a zero-sum security game such that the difference between worlds 2 and 1 is $1/2 - \epsilon$
- Lemma: If $|A| = \binom{2k}{k}$, there exists $\mathcal{D} = \{D_1, \dots, D_{2k}\} \subseteq 2^A$ such that:
  1. $\forall i, |D_i| = |A|/2$
  2. Each $a \in A$ is in exactly k members of $\mathcal{D}$
  3. If $\mathcal{D}' \subset \mathcal{D}$ and $|\mathcal{D}'| \leq k$ then $\cup \mathcal{D}' \neq A$

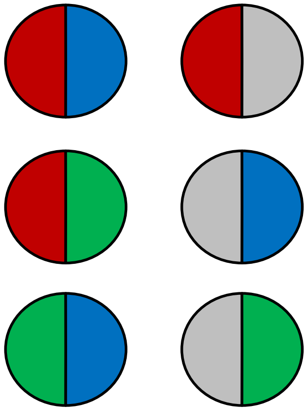

$k = 2$

Proof of theorem

- m resources, each can defend any d targets, $n = \left\lceil \frac{md}{\epsilon} \right\rceil$ targets
- For any target i, zero-sum utilities with $U_d^c(i) = 1$ and $U_d^u(i) = 0$
- Optimal strategy assuming unlimited surveillance: defend every target with probability $\frac{md}{n} \leq \epsilon$

Proof of theorem

- Next we define a much better strategy against an attacker with k observations
- $A =$ subset of targets $\left\{1, \ldots, \binom{2k}{k}\right\} \subseteq T$
- Define $\{D_1, \ldots D_{2k}\}$ as in the lemma
- Pure strategy S_i covers D_i; this is valid because $|D_i| = |A|/2 \leq md$ (by property 1)
- Let S^* be the uniform distribution over $S_1, \ldots, S_{2k}$
- By property 2, S^* covers each target in A with probability ½
- By property 3, k observations from S^* would show some target in A never being covered; that target is attacked ■

Limited surveillance

- Theorem [Blum et al. 2014]: For any zero-sum security game with n targets, m resources, and a set of schedules with max coverage d, and for any k observations, the difference between the two worlds is at most

$$O\left(\sqrt{\frac{\ln(mdk)}{k}}\right)$$