CMU 15-781

Lecture 23: Game Theory II

Teacher:
Gianni A. Di Caro

Game of chicken

http://youtu.be/u7hZ9jKrwvo

Each player, in attempting to secure his best outcome, risks the worst

Game of chicken

- Social welfare is the sum of utilities
- Pure NE: (C,D) and (D,C), social welfare = 5
- Mixed NE: both (1/2,1/2), social welfare = 4
- Optimal social welfare = 6
- Can we do better? Players are independent so far ...

	Dare	Chicken
Dare	0,0	4,1
Chicken	1,4	3,3

Correlated equilibrium

- A "trusted" authority / mediator chooses a pair of strategies (s_1, s_2) according to a distribution p over S^2 (it can be generalized to n players)

The mediator flips a coin and based on the outcome tells the players which pure strategy to use based on some distribution $p(\boldsymbol{s})$

Correlated equilibrium

- The trusted party only tells each player what to do, but it does not reveal what the other party is supposed to do
- The distribution p is known to the players: each player knows the probability of observing a strategy profile and assumes the other player will follow mediator's instructions
- It is a **Correlated Equilibrium (CE)** if no player wants to deviate from the trusted party's instructions, such that choices are *correlated*
- Find distribution p that guarantees a CE

Correlated equilibrium

- Distribution p (is CE)
  - (D,D): 0
  - (D,C): $\frac{1}{3}$
  - (C,D): $\frac{1}{3}$
  - (C,C): $\frac{1}{3}$

	Dare	Chicken
Dare	0,0	7,2
Chicken	2,7	6,6

- If Player 2 is told to play D, then 2 knows that the outcome must be (C,D) and that Player 1 will obey the instructions. Therefore, P1 plays C, and Player 2 has no incentive to change from playing D

Correlated equilibrium

- Distribution p (is CE)
  - (D,D): 0
  - (D,C): $\frac{1}{3}$
  - (C,D): $\frac{1}{3}$
  - (C,C): $\frac{1}{3}$

	Dare	Chicken
Dare	0,0	7,2
Chicken	2,7	6,6

- If Player 2 is told to play C, then 2 knows that the outcome must be (D,C) or (C,C) with equal probability. Player's 2 expected utility on playing C conditioned on the fact that he is told to play C (and Player 1 will obey instructions) is:

$$(1/2)^* u_2(\mathrm{D,C}) + (1/2)^* u_2(\mathrm{C,C}) = (1/2)^*2 + (1/2)^*6 = 4$$

- If Player 2 deviates from instructions and plays D: u_2=3.5 < 4
  - It's better to follow the instructions!

Correlated equilibrium

- Distribution p (is CE)
  - (D,D): 0
  - (D,C): $\frac{1}{3}$
  - (C,D): $\frac{1}{3}$
  - (C,C): $\frac{1}{3}$

	Dare	Chicken
Dare	0,0	7,2
Chicken	2,7	6,6

- Player 2 does not have incentive to deviate
- Since the game is symmetric, also Player 1 does not have incentive to deviate
- $\rightarrow$ Correlated equilibrium
- Expected reward per player: $(1/3)*7 + (1/3)*2 + (1/3)*6 = 5$
- Mixed strategy NE: $4*(2/3)$, which is < 5
- Social welfare: $30/3$

Correlated equilibrium

- Let $N = \{1,2\}$ for simplicity
- A mediator chooses a pair of strategies (s_1, s_2) according to a distribution p over S^2
- Reveals s_1 to player 1 and s_2 to player 2
- When player 1 gets $s_1 \in S$, he knows that the distribution over strategies of 2 is

$$\Pr[s_2 | s_1] = \frac{\Pr[s_1 \wedge s_2]}{\Pr[s_1]} = \frac{p(s_1, s_2)}{\sum_{s_2' \in S} p(s_1, s_2')}$$

Correlated equilibrium

- Player 1 is best responding if for all $s_1' \in S$

$$\sum_{s_2 \in S} \Pr[s_2 | s_1]\, u_1(s_1, s_2) \geq \sum_{s_2 \in S} \Pr[s_2 | s_1]\, u_1(s_1', s_2)$$

- Equivalently, replacing using Bayes' rule

$$\sum_{s_2 \in S} p(s_1, s_2) u_1(s_1, s_2) \geq \sum_{s_2 \in S} p(s_1, s_2) u_1(s_1', s_2)$$

- p is a correlated equilibrium (CE) if both players are best responding

Implementation of CE

- Instead of a mediator, use a hat!
- Balls in hat are labeled with "chicken" or "dare", each blindfolded player takes a ball
- Poll 1: Which balls implement the distribution of slide 6?
  1. 1 chicken, 1 dare
  2. 2 chicken, 1 dare
  3. 2 chicken, 2 dare
  4. 3 chicken, 2 dare

CE vs. NE

- Poll 2: What is the relation between CE and NE?
  1. CE $\Rightarrow$ NE
  2. NE $\Rightarrow$ CE
  3. NE $\Leftrightarrow$ CE
  4. NE $\parallel$ CE

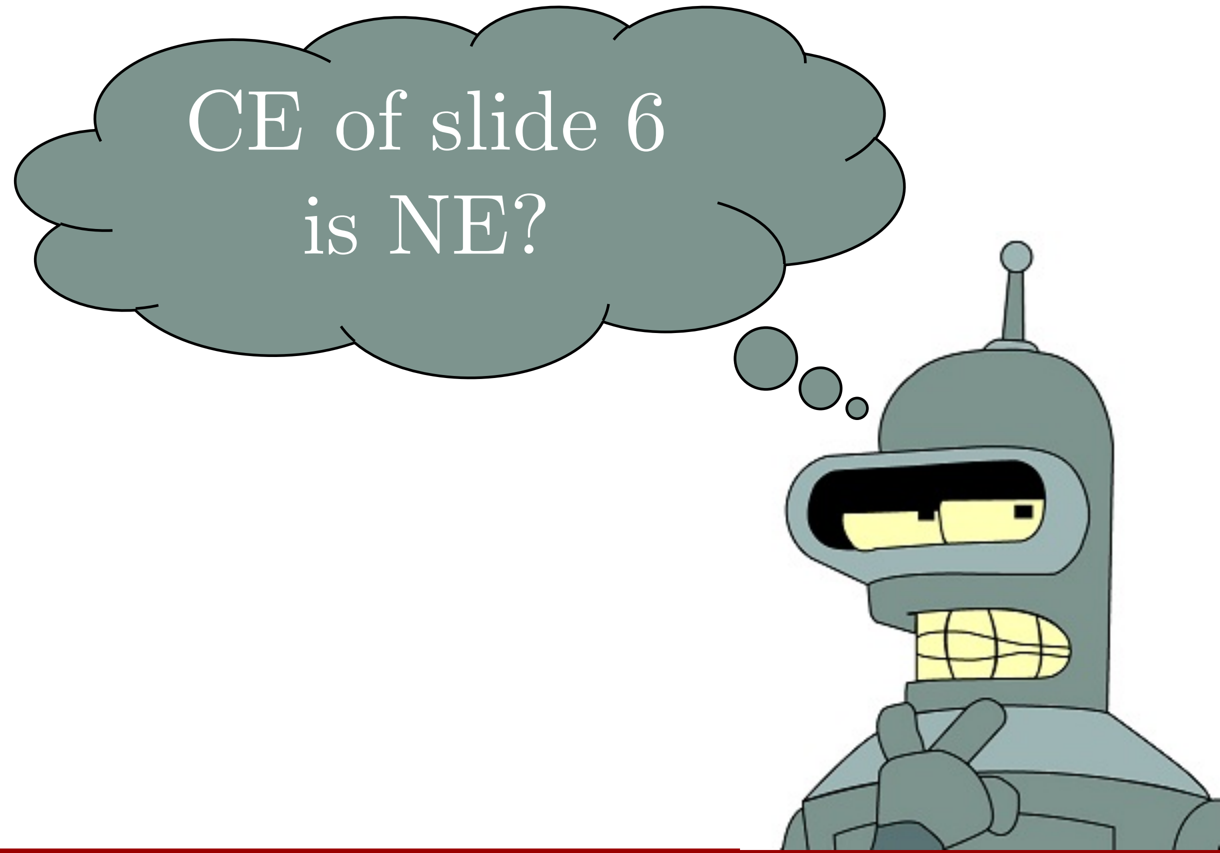

CE vs. NE

- For any pure strategy NE, there is a corresponding correlated equilibrium yielding the same outcome.
- For any mixed strategy NE, there is a corresponding correlated equilibrium yielding the same distribution of outcomes.
- From Nash theorem, "all" games have a mixed strategies NE. Since a NE implies a CE, a CE always exist

CE As LP

- Can compute CE via linear programming in polynomial time!

$$\textbf{find } \forall s_1, s_2 \in S, p(s_1, s_2)$$

$$\textbf{s.t. } \forall s_1, s_1', s_2 \in S, \sum_{s_2 \in A} p(s_1, s_2) u_1(s_1, s_2) \geq \sum_{s_2 \in A} p(s_1, s_2) u_1(s_1', s_2)$$

$$\forall s_1, s_2, s_2' \in S, \sum_{s_1 \in A} p(s_1, s_2) u_2(s_1, s_2) \geq \sum_{s_1 \in A} p(s_1, s_2) u_2(s_1, s_2')$$

$$\sum_{s_1, s_2 \in S} p(s_1, s_2) = 1$$

$$\forall s_1, s_2 \in S, p(s_1, s_2) \in [0,1]$$

Best Welfare CE

- Adding an objective (linear) function f, the best correlated equilibrium (e.g., max welfare) can be found

$$\textbf{max}\ \forall s_1, s_2 \in S, f(p(s_1, s_2); u_1, u_2)$$

$$\textbf{s.t.}\ \forall s_1, s_1', s_2 \in S, \sum_{s_2 \in A} p(s_1, s_2) u_1(s_1, s_2) \geq \sum_{s_2 \in A} p(s_1, s_2) u_1(s_1', s_2)$$

$$\forall s_1, s_2, s_2' \in S, \sum_{s_1 \in A} p(s_1, s_2) u_2(s_1, s_2) \geq \sum_{s_1 \in A} p(s_1, s_2) u_2(s_1, s_2')$$

$$\sum_{s_1, s_2 \in S} p(s_1, s_2) = 1$$

$$\forall s_1, s_2 \in S, p(s_1, s_2) \in [0,1]$$

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

	L	R
U	1,1	3,0
D	0,0	2,1

Commitment is good

- Suppose the game is played *sequentially* as follows:
  - Row player commits to playing a row
  - Column player observes the commitment and chooses column
- Row player can commit to playing down: Column player will play R and the Row player gets now a better reward!

	L	R
U	1,1	3,0
D	0,0	2,1

Commitment to mixed strategy

- By committing to a mixed strategy, row player can get even better and guarantee a reward of almost 2.5
- Called a Stackelberg strategy (1934)
- Rooted in duopoly scenarios

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

- Player 1 (*Leader*) moves at the start of the game. Then use backward induction to find the subgame perfect equilibrium.
- First, for any output of leader, find the strategy of *Follower* that maximizes its payoff (its expected best reply).
- Next, find the strategy of leader that maximizes player 1 utility, given the strategy of follower

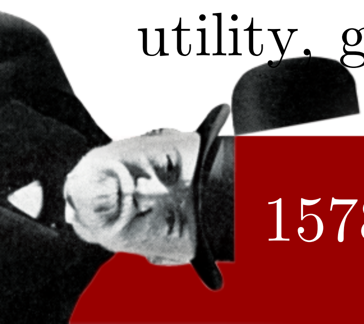

COMPUTING STACKELBERG

- Theorem [Conitzer and Sandholm, EC 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

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

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

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

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

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 $T = \{1, \ldots, n\}$
- 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 (follower) chooses one target to attack

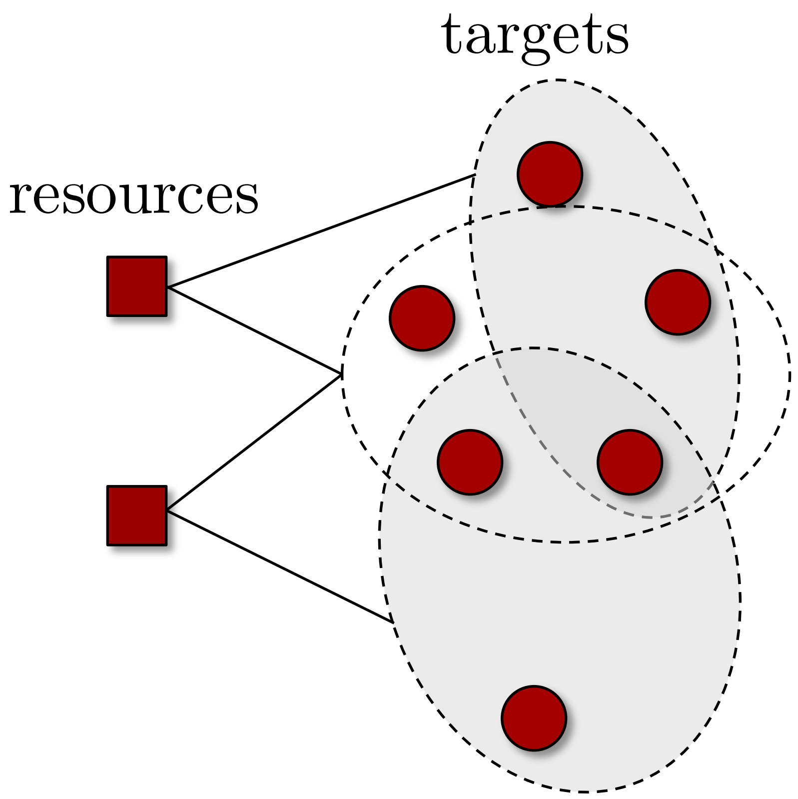

Security games

- For each target t, there are four numbers: $u_d^+(t) \geq u_d^-(t)$, and $u_a^+(t) \leq u_a^-(t)$
- 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

$$u_d(t, \boldsymbol{c}) = u_d^+(t) \cdot c_t + u_d^-(t)(1 - c_t)$$
$$u_a(t, \boldsymbol{c}) = u_a^+(t) \cdot c_t + u_a^-(t)(1 - c_t)$$

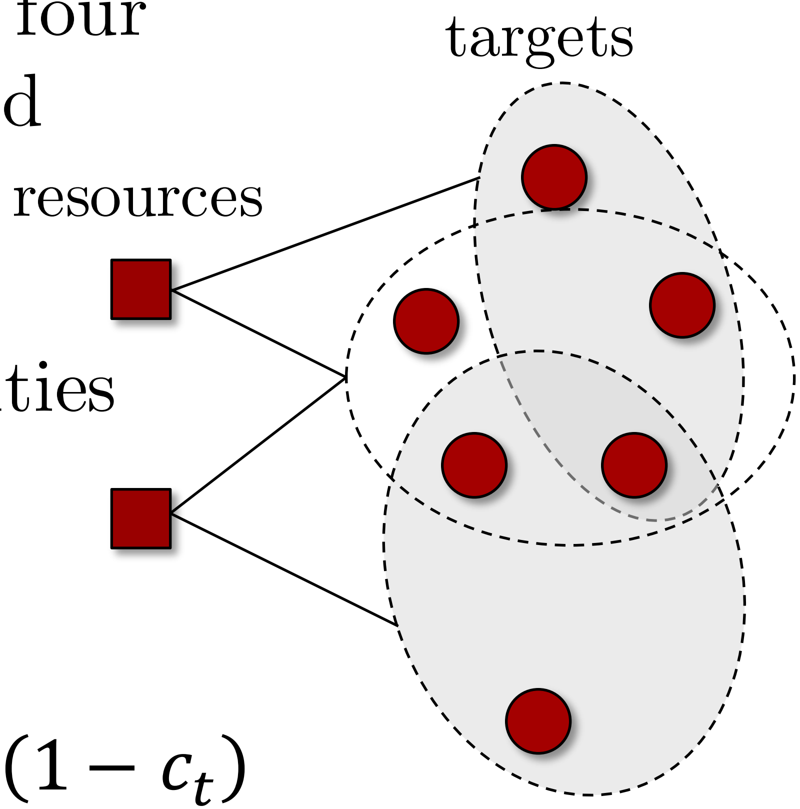

This is a 2-player Stackelberg game, so we can compute an optimal strategy for the defender in polynomial time...?

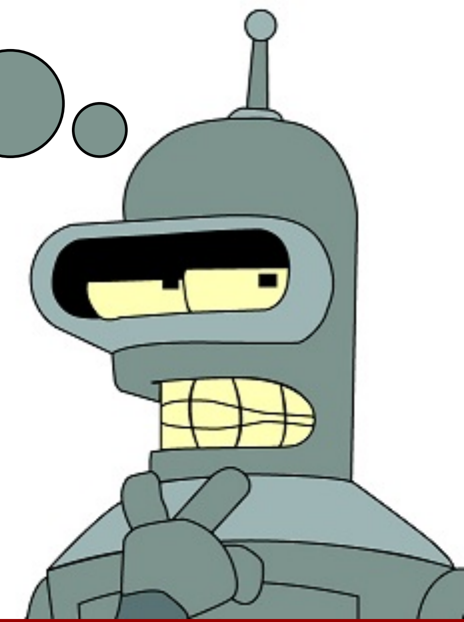

SOLVING SECURITY GAMES

- 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?

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

* Just for fun

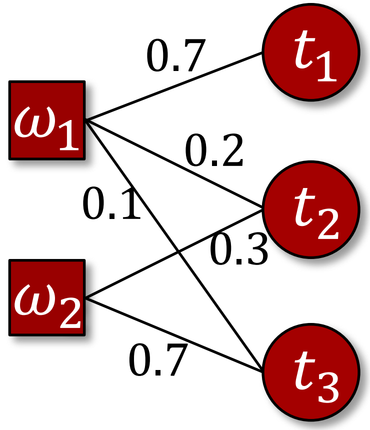

	t_1	t_2	t_3
ω_1	0.7	0.2	0.1
ω_2	0	0.3	0.7

<u>0.1</u>

	t_1	t_2	t_3
ω_1	0	0	1
ω_2	0	1	0

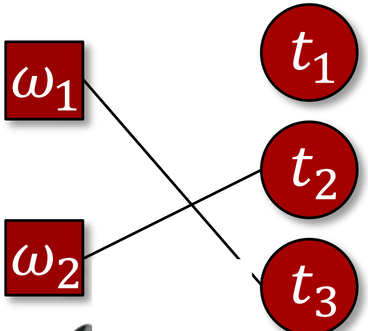

<u>0.2</u>

	t_1	t_2	t_3
ω_1	0	1	0
ω_2	0	0	1

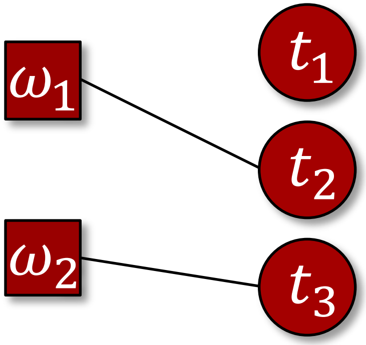

<u>0.2</u>

	t_1	t_2	t_3
ω_1	1	0	0
ω_2	0	1	0

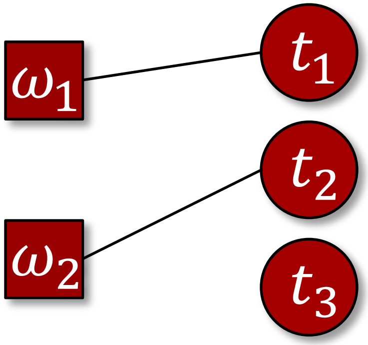

<u>0.5</u>

	t_1	t_2	t_3
ω_1	1	0	0
ω_2	0	0	1

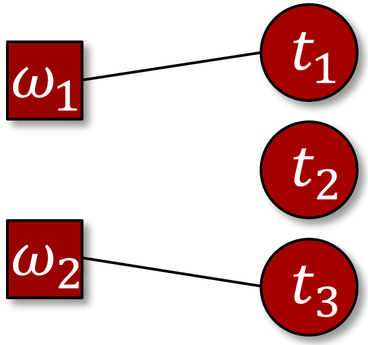

* Just for fun

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

* Just for fun

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

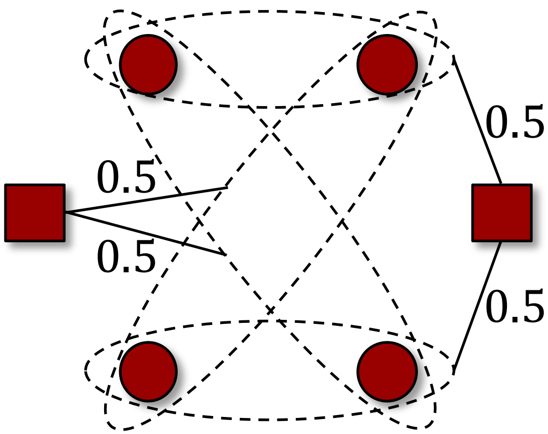

* Just for fun

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.

LIMITATIONS

- The defender knows the utility function of the attacker
  - Solution: machine learning
- The attacker perfectly observes the defender's randomized strategy
  - MDPs, although this may not be a major concern
- The attacker is perfectly rational, i.e., best responds to the defender's strategy
  - Solution: bounded rationality models

Testing bounded rationality

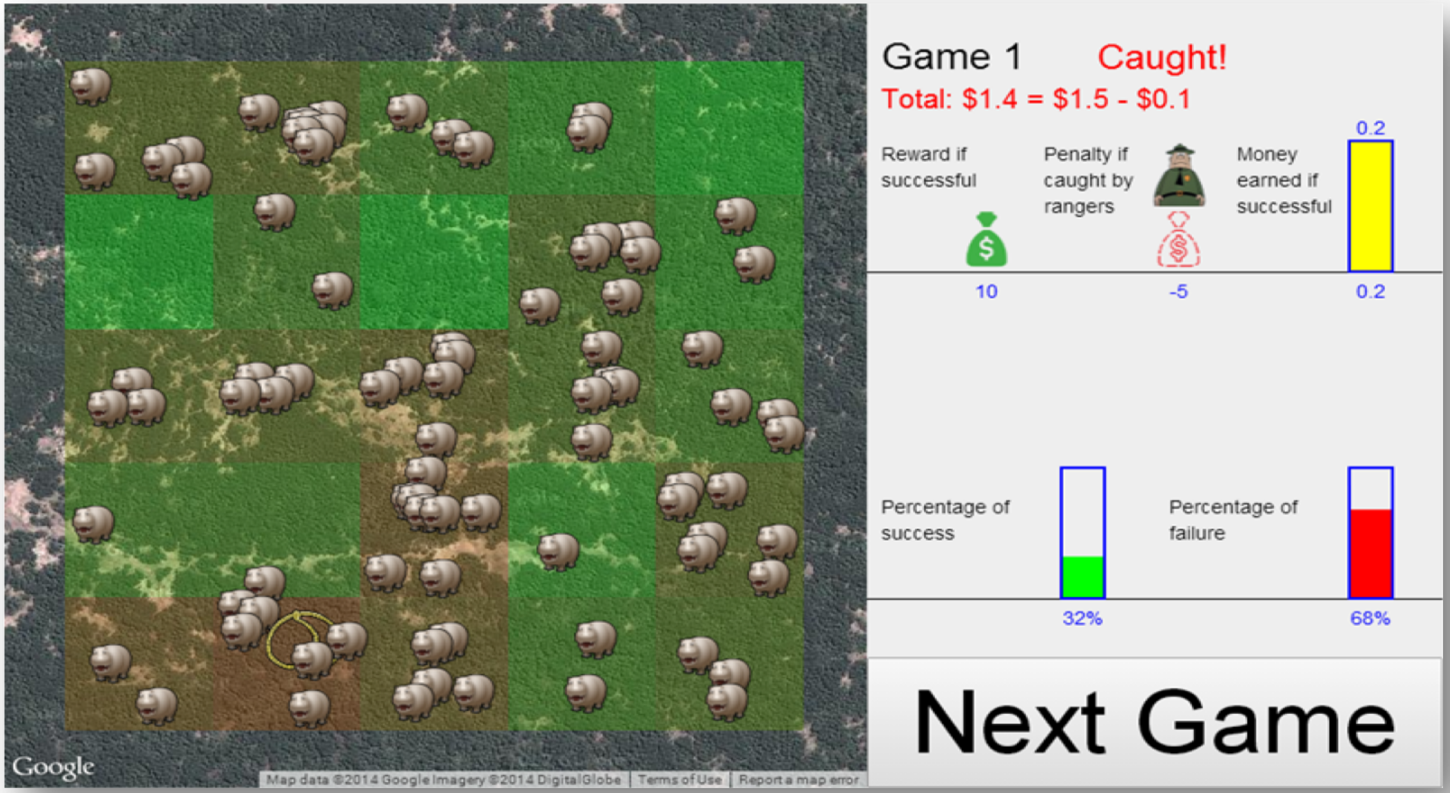

[Kar et al., 2015]

Summary

- Terminology:
  - Correlated equilibrium
  - Stackelberg game
  - Security game
- Nobel-prize-winning ideas:
  - Correlated equilibrium ☺
- Other big ideas:
  - Stackelberg games for security