Stackelberg Games
with Applications to Security

Chris Kiekintveld **Bo An** **Albert Xin Jiang**



UTEP

TEAMCORE
USC

Outline

- ***Motivating real-world applications***
- Background and basic security games
- Scaling to complex action spaces
- Modeling payoff uncertainty: Bayesian Security Games
- Human behavior and observation uncertainty
- Evaluation and discussion

Motivation: Game Theory for Security

- Limited security resources: Selective checking
- Adversary monitors defenses, exploits patterns

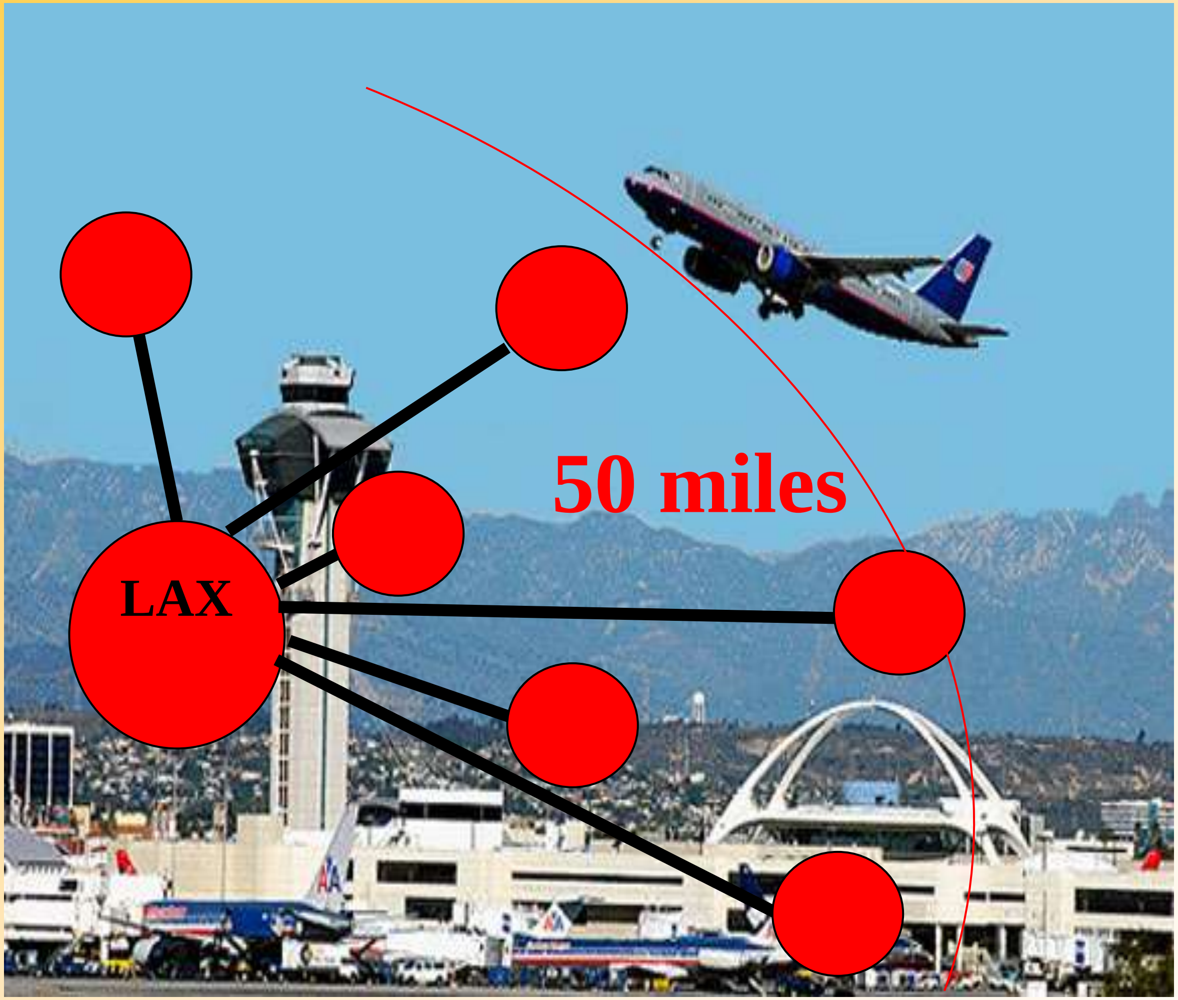

Many Targets Few Resources

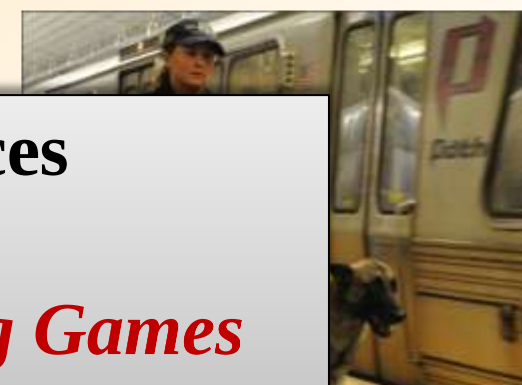

How to assign limited resources
to defend the targets?

Game Theory: Bayesian Stackelberg Games

Game Theory: Bayesian Stackelberg Games

- Security allocation: (i) Target weights; (ii) Opponent reaction
- *Stackelberg*: Security forces commit first
- *Bayesian*: Uncertain adversary types
- *Optimal security allocation:* Weighted random
- Strong Stackelberg Equilibrium (Bayesian)
  - *NP-hard (Conitzer/Sandholm '06)*

Adversary

Police

	Terminal #1	Terminal #2
Terminal #1	5, -3	-1, 1
Terminal #2	-5, 5	2, -1

ARMOR: Deployed at LAX 2007

- "Assistant for Randomized Monitoring Over Routes"
  - *Problem 1: Schedule vehicle checkpoints*
  - *Problem 2: Schedule canine patrols*
- Randomized schedule: (i) target weights; (ii) surveillance

ARMOR-Checkpoints

ARMOR-K9

ARMOR Canine: Interface

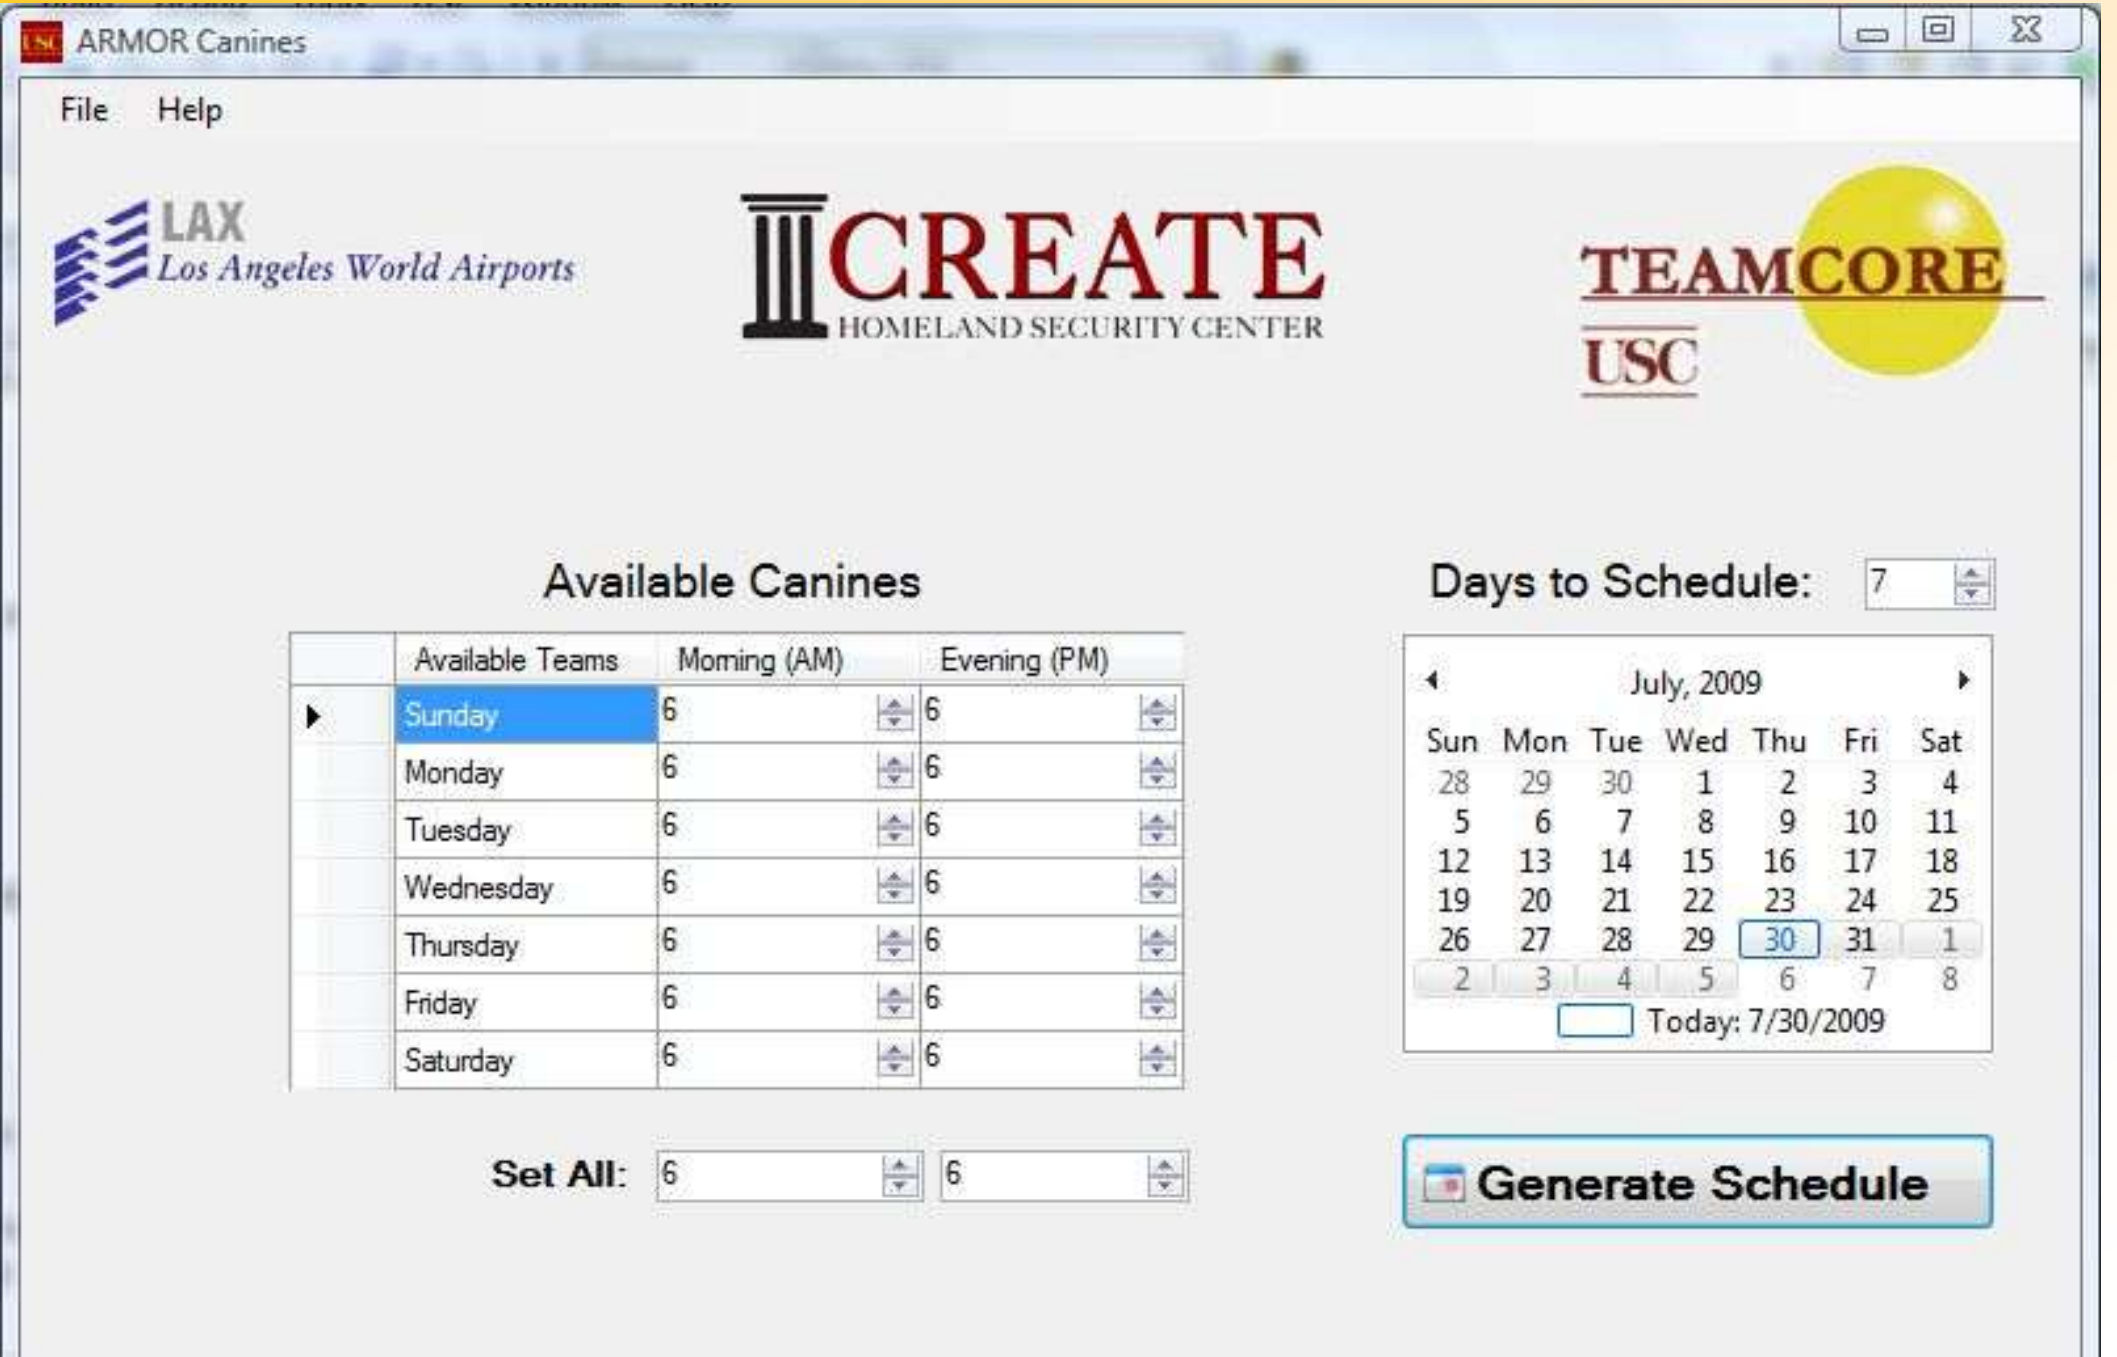

Federal Air Marshals Service (FAMS)

Undercover, in-flight
law enforcement

<u>Flights (each day)</u>
~27,000 domestic flights
~2,000 international flights

Not enough air marshals:
Allocate air marshals to flights?

International Flights from Chicago O'Hare

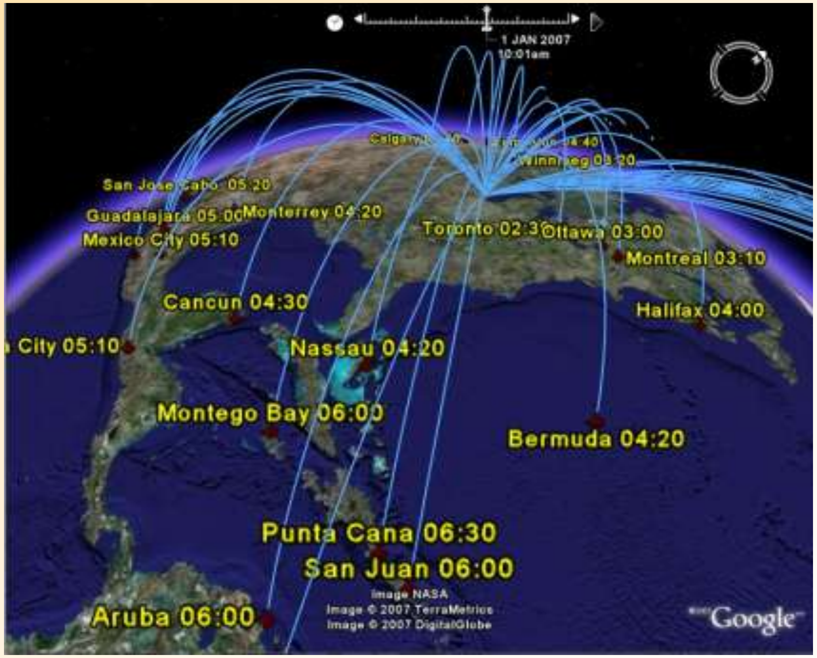

Federal Air Marshals Service (FAMS)

- Massive scheduling problem
- Adversary may exploit predictable schedules
- Complex constraints: tours, duty hours, off-hours

100 flights, 10 officers:

1.7×10^{13} combinations

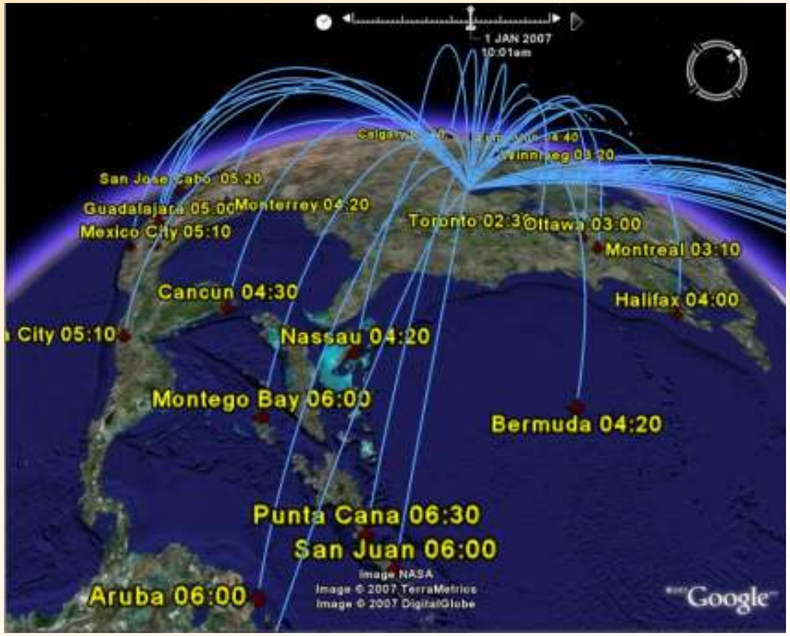

Overall problem: 30000 flights, 3000 officers

Our focus: international sector

IRIS: "Intelligent Randomization in International Scheduling" (Deployed 2009)

PROTECT (Boston and Beyond)

- US Coast Guard: *Port Resilience Operational / Tactical Enforcement to Combat Terrorism*
- Randomized patrols; deployed in Boston, with more to follow
- More realistic models of human behaviors

Application in Transition: GUARDS

- ***GUARDS: under evaluation for national deployment***
- Transportation Security Administration
  - *Protect over 400 airports*
    - Limited security resources
    - Numerous security measures
    - Diverse potential threats
  - *Adaptive adversary*

International Interest: Mumbai

- *Protect networks*

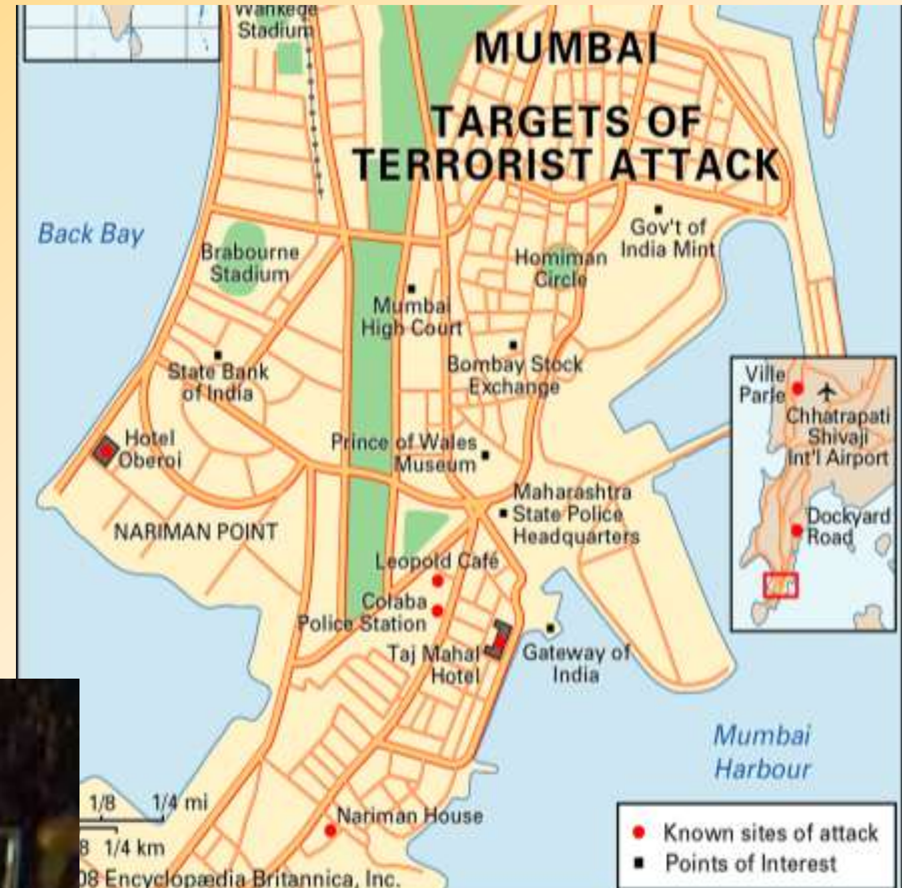

Urban Road Network Security

Southern Mumbai

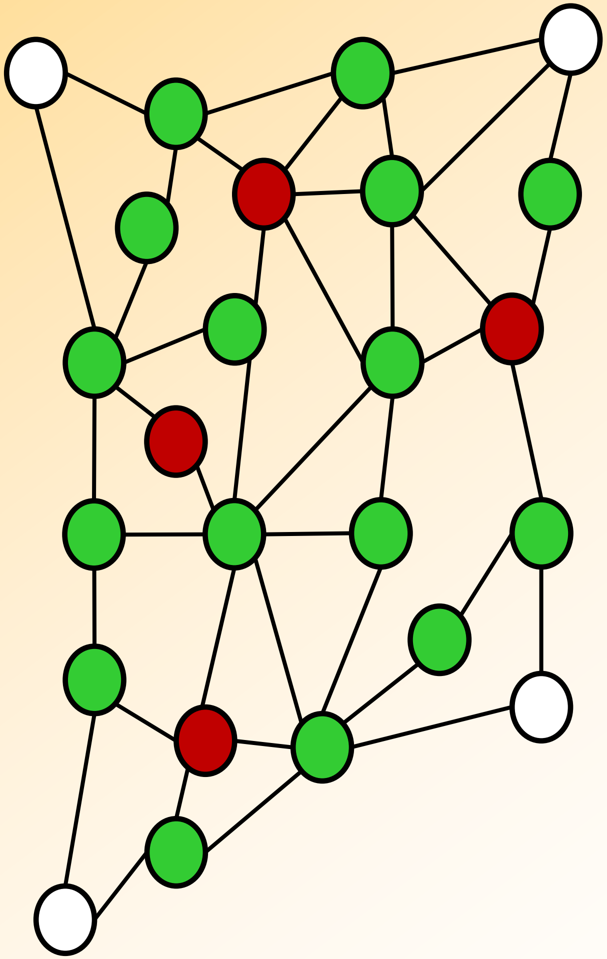

Beyond Counterterrorism: Other Domains

- LA Sheriff's dept *(Crime suppression & ticketless travelers):*

- Customs and Border Protection
- Cybersecurity
- Forest/environmental protection
- Economic leader/follower models

Research Challenges

- Scalable algorithms
- Rich representations; networks
- Payoff uncertainty, robustness
- Imperfect surveillance
- Evaluation of deployed systems
- Human behavior, bounded rationality
- Explaining game theory solutions
- …

Publications

Publications ~40 rigorously reviewed papers:

- AAMAS' [06-12: (15)]
- AAAI[08,10-12: (10)]
- IJCAI'11: (2)
- ECAI'12: (1)
- IAAI'12: (1)
- JAIR'11
- JAAMAS'12
- AI Journal'10, 12
- Interfaces'10
- AI Magazine'09,12…
- Journal ITM'09

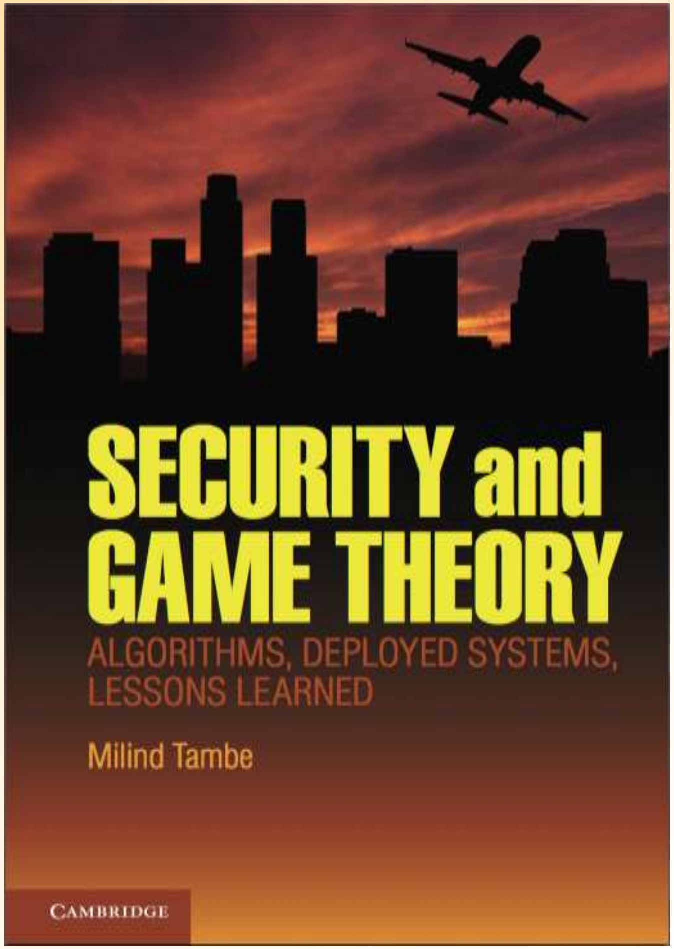

Outline

- Motivating real-world applications
- ***Background and basic security games***
- Scaling to complex action spaces
- Modeling payoff uncertainty: Bayesian Security Games
- Human behavior and observation uncertainty
- Evaluation and discussion

Games

- Players:
  - $1, \ldots, n$
  - *focus on 2 players*
- Strategies
  - $a_i \in A_i$
  - $a = (a_1,\ldots,a_n) \in A$
- Utility function
  - $u_i : A \rightarrow R$

Security Games

- Two players
  - *Defender: Θ*
  - *Attacker: ψ*
- Set of targets: T
- Set of resources: R
  - *Defender assigns resources to protect targets*
  - *Attacker chooses one target to attack*
- Payoffs define the reward/penalty for each player for a successful or unsuccessful attack on each target

Zero-Sum Payoffs?

- Are security games always zero-sum?
  - *NO!*
- In real domains attackers and defenders often have different preferences and criteria
  - *Weighting casualties, economic consequences, symbolic value, etc.*
  - *Player may not care about the other's cost (e.g., cost of security, cost of carrying out an attack)*
- We often make a weaker assumption:
  - *An attack on a defended target is better than an attack on the same target if it is undefended (for the defender)*
  - *The opposite holds for attackers (attackers prefer to attack undefended targets)*

Security Game

2 players
2 targets
1 defender resource

	Target1	Target 2
Target 1	1, -1	-2, 2
Target 2	-1, 1	2, -1

Game Solutions

Best Response

	Target1	Target 2
Target 1	1, -1	-2, 2
Target 2	-1, 1	2, -1

Game Solutions

Best Response

	Target1	Target 2
Target 1	1, -1	-2, 2
Target 2	-1, 1	2, -1

Game Solutions

Best Response

	Target1	Target 2
Target 1	1, -1	-2, 2
Target 2	-1, 1	2, -1

Game Solutions

Mixed Strategy

		Target1	Target 2
50%	Target 1	1, -1	-2, 2
50%	Target 2	-1, 1	2, -1

Game Solutions

Nash Equilibrium

A mixed strategy for each player such that no player benefits from a unilateral deviation

	Target1	Target 2
Target 1	1, -1	-2, 2
Target 2	-1, 1	2, -1

Game Solutions

Nash Equilibrium

A mixed strategy for each player such that no player benefits from a unilateral deviation

		67%	33%
		Target1	Target 2
40%	Target 1	1, -1	-2, 2
60%	Target 2	-1, 1	2, -1

Stackelberg Equilibrium

Attackers use surveillance in planning attacks

Defender commits to a mixed strategy

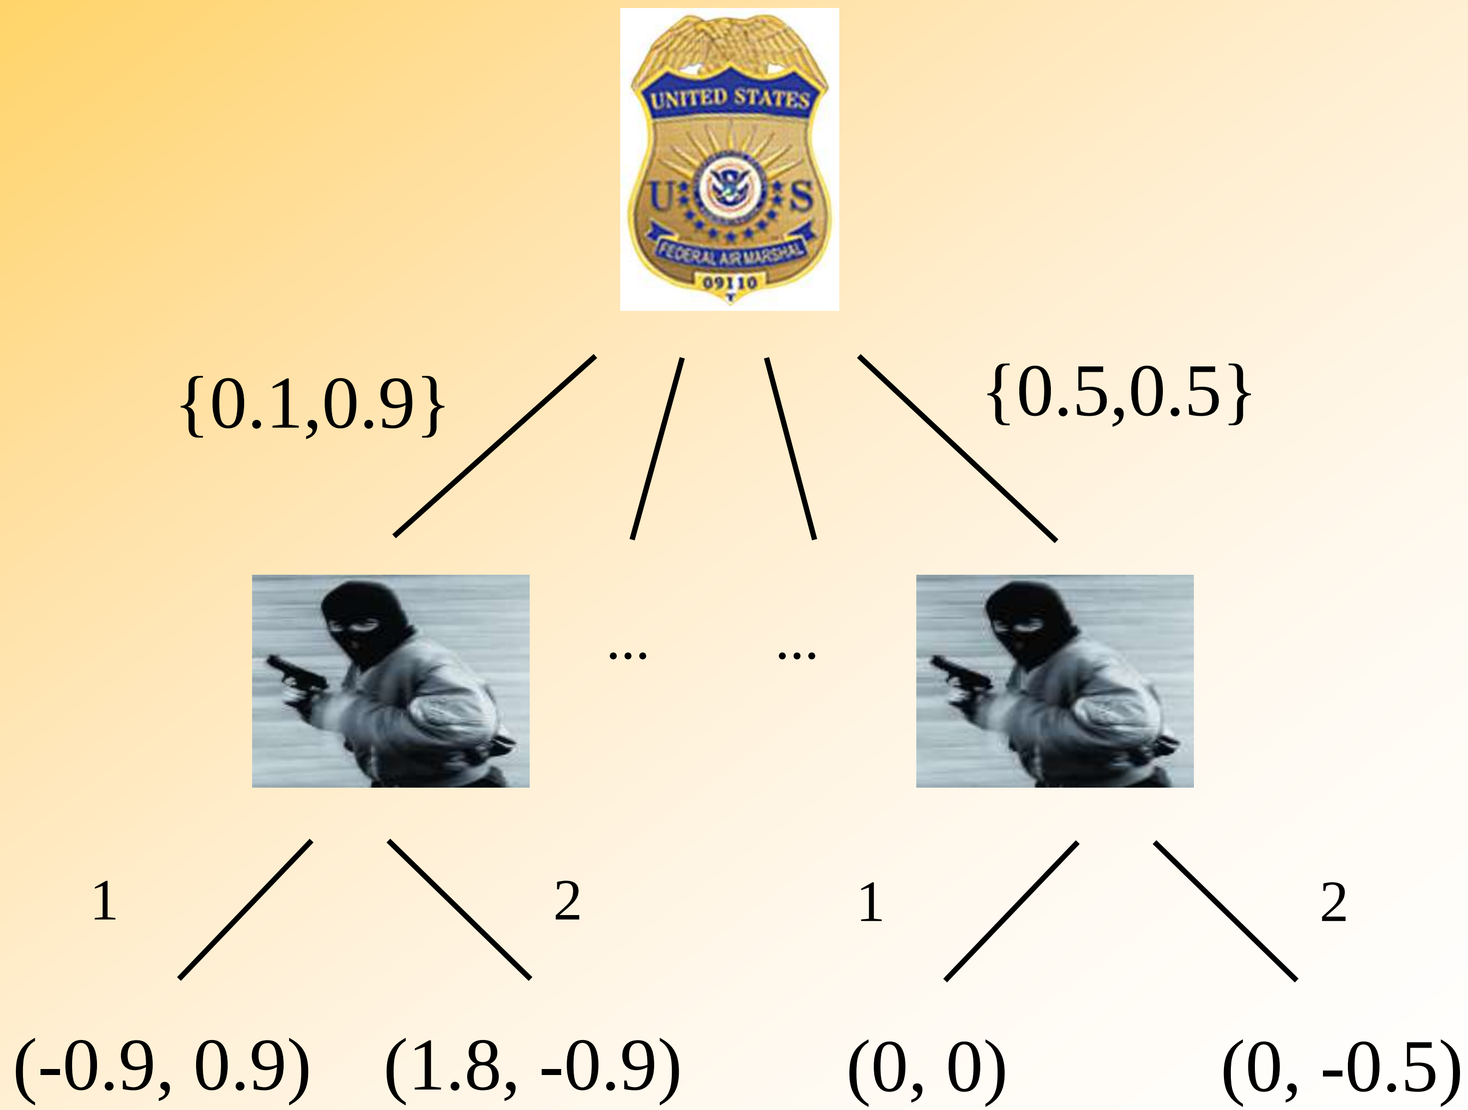

Strong Stackelberg Equilibrium

- Strong Stackelberg Equilibrium (SSE)
  - *Break ties in favor of the defender*
  - *Can often induce SSE by perturbing defender strategy*
- More robust concepts
  - *Weak Stackelberg Equilibrium not guaranenteed to exist*
  - *Payoff uncertainty*
  - *Quantal response*
  - *Equilibrium refinement*

Finding Stackelberg Equilibria

Multi-linear programming formulation
Conitzer and Sandholm, 2006

$$\max \sum_{s_1} p_{s_1} u_1(s_1, s_2)$$

$$\forall s_2', \quad \sum_{s_1} p_{s_1} u_2(s_1, s_2') \leq \sum_{s_1} p_{s_1} u_2(s_1, s_2)$$

$$\sum_{s_1} p_{s_1} = 1$$

$$p_{s_1} \geq 0$$

The formulation above gives the maximum utility of the leader when the follower chooses action *a*

The Stackelberg equilibrium is obtained by maximizing over all the possible pure strategies for player two

Single LP formulation
(Korzhyk & Conitzer 2011)

$$\max \sum_{s_1,s_2} x_{s_1,s_2} u_1(s_1,s_2)$$

$$\forall s_2, s_2', \quad \sum_{s_1} x_{s_1,s_2} u_2(s_1,s_2') \leq \sum_{s_1} x_{s_1,s_2} u_2(s_1,s_2)$$

$$\sum_{s_1,s_2} x_{s_1,s_2} = 1$$

$$x_{s_1,s_2} \geq 0$$

- Relaxation of the LP for correlated equilibrium
  - *removed player 1's incentive constraints*
- Corollary: SSE leader expected utility at least that of best CE

Research Challenges

- Scalability
  - *Large, complex strategy spaces*
- Robustness
  - *Payoff & observation uncertainty*
  - *Human decision-makers*
- Not in this talk:
  - *Stackelberg equilibria for dynamic games (Letchford & Conitzer 2010, Letchford et al. 2012)*
  - *Multiple objectives (Brown et al. 2012)*

Outline

- Motivating real-world applications
- Background and basic security games
- ***Scaling to complex action spaces***
- Modeling payoff uncertainty: Bayesian Security Games
- Human behavior and observation uncertainty
- Evaluation and discussion

Large Numbers of Defender Strategies

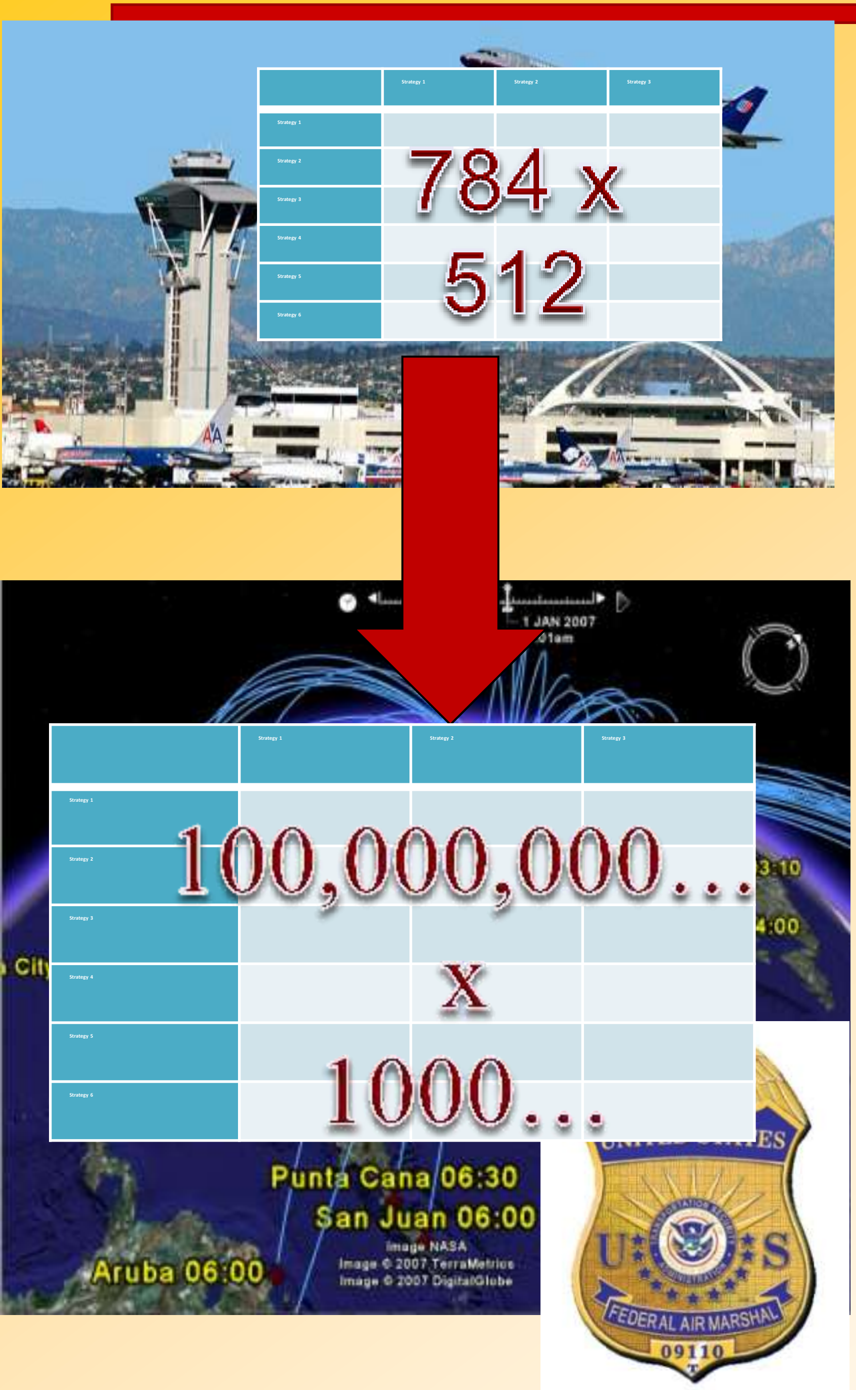

*FAMS: **Joint** Strategies or Combinations*

100 Flight tours
10 Air Marshals

1.73×10^{13} Schedules: ARMOR out of memory

Don't enumerate ALL joint strategies

- *Marginals* (IRIS I & II)
- *Branch and price* (IRIS III)

IRIS I & II: Marginals Instead of Joint Strategies

ARMOR: 10 tours, 3 air marshals

ARMOR Actions	Tour combos	Prob
1	1,2,3	x1
2	1,2,4	x2
3	1,2,5	x3
…	…	…
120	8,9,10	x120

Compact Action	Tour	Prob
1	1	y1
2	2	y2
3	3	y3
…	…	…
10	10	y10

Payoff duplicates: Depends on target covered

	Attack 1	Attack 2	Attack …
1,2,3		4,-8	…
1,2,4		4,-8	…
1,3,5		-9,5	…
…		…	…

$$\max_{x,q} \sum_{i \in X} \sum_{l \in L} \sum_{j \in Q} p^l R_{ij}^l x_i q_j^l$$

$$s.t. \sum_i x_i = 1, \sum_{j \in Q} q_j^l = 1$$

$$0 \leq (a^l - \sum_{i \in X} C_{ij}^l x_i) \leq (1 - q_j^l) M$$

$$x_i \in [0 \ldots 1], q_j^l \in \{0,1\}$$

MILP similar to ARMOR, y instead of x:

- 10 instead of 120 variables
- y1+y2+y3…+y10 = 3
- Sample from “y”, not enumerate “x”
- Only works for SIMPLE tours (Korzhyk et al. 2010)

IRIS II

Max Defender Payoff

$$\max \quad d \tag{5}$$

Attacker Strategy Definition

$$a_t \in \{0,1\} \quad \forall t \in T \tag{6}$$

$$\sum_{t \in T} a_t = 1 \tag{7}$$

Defender Strategy Definition

$$c_t \in [0,1] \quad \forall t \in T \tag{8}$$

$$\sum_{t \in T} c_t \leq m \tag{9}$$

Best Responses

$$d - U_{\Theta}(t,C) \leq (1-a_t) \cdot Z \quad \forall t \in T \tag{10}$$

$$0 \leq k - U_{\Psi}(t,C) \leq (1-a_t) \cdot Z \quad \forall t \in T \tag{11}$$

IRIS I

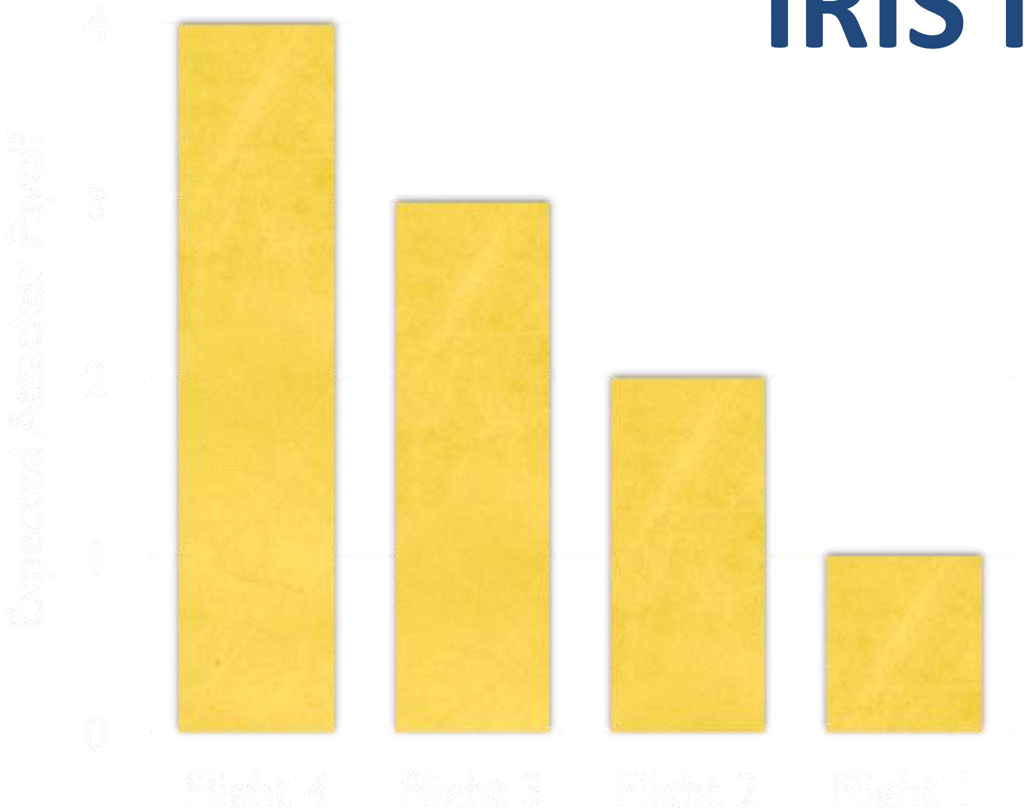

Four flights
One marshal

Zero Sum
Attacker payoffs

Uncovered	Covered
4	0
3	0
2	0
1	0

0	0	0	0

Coverage Probability

IRIS I

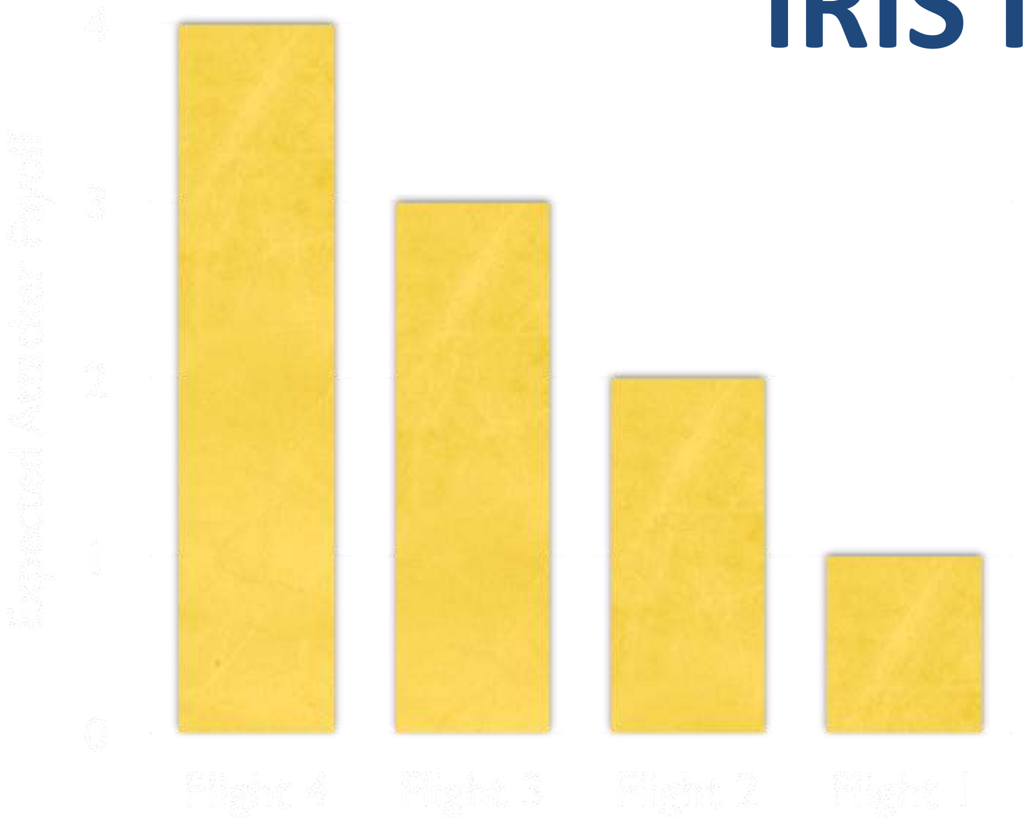

<u>Attack Set:</u>
Set of targets with maximal expected payoff for the attacker

0	0	0	0

Coverage Probability

IRIS I

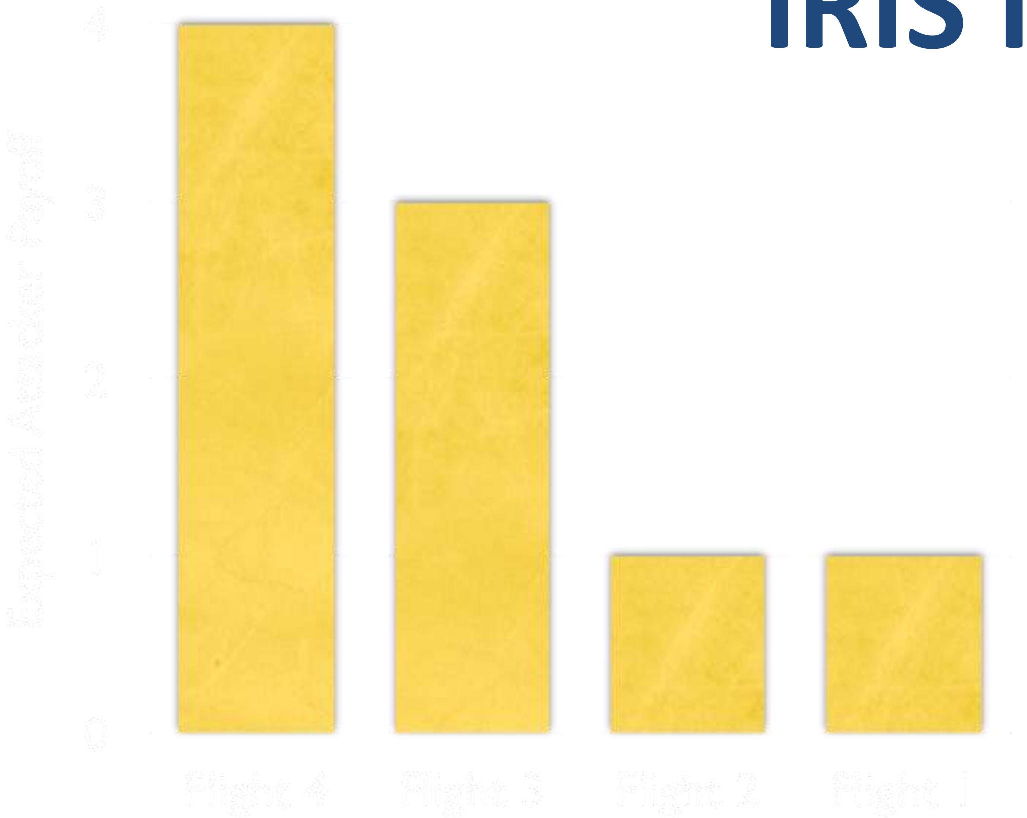

Observation 1
It never benefits
the defender to
add coverage outside the attack set.

0	0	0.5	0

Coverage Probability

IRIS I

Compute coverage necessary to make attacker indifferent between 3 and 4

0.25	0	0	0

Coverage Probability

IRIS I

<u>Observation 2</u>
It never benefits the defender to add coverage to a subset of the attack set.

0.5	0	0	0

Coverage Probability

IRIS I

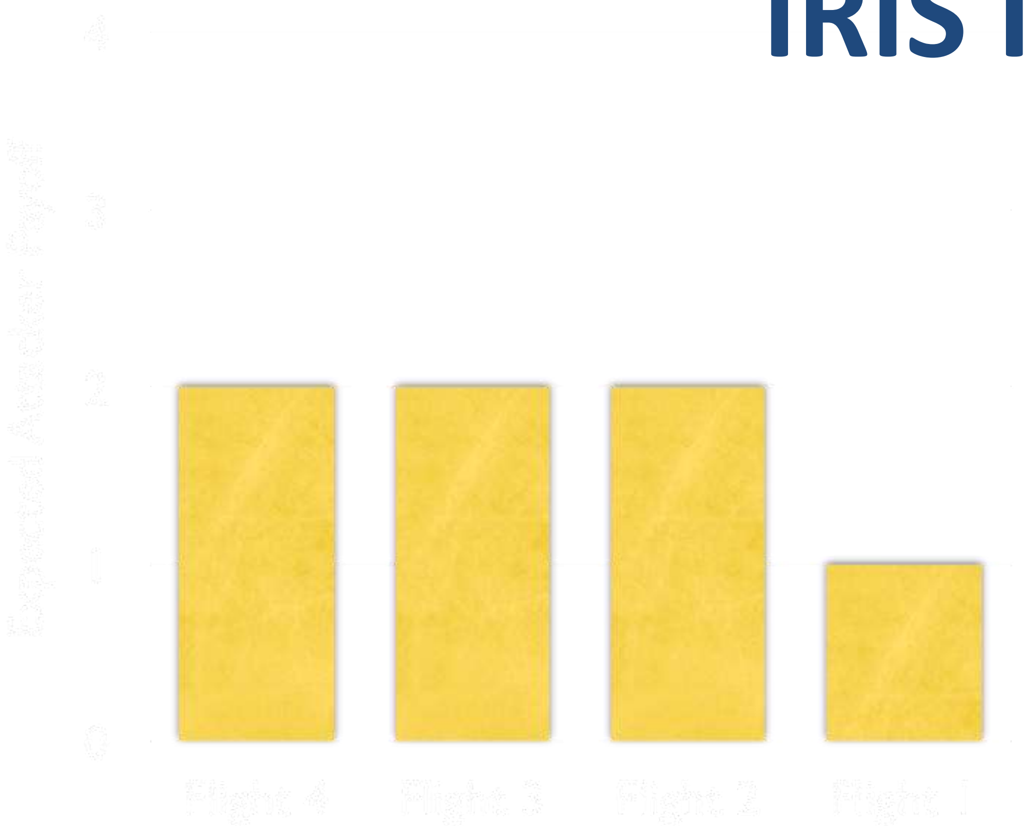

0.5	0.33	0	0

Coverage Probability

IRIS I

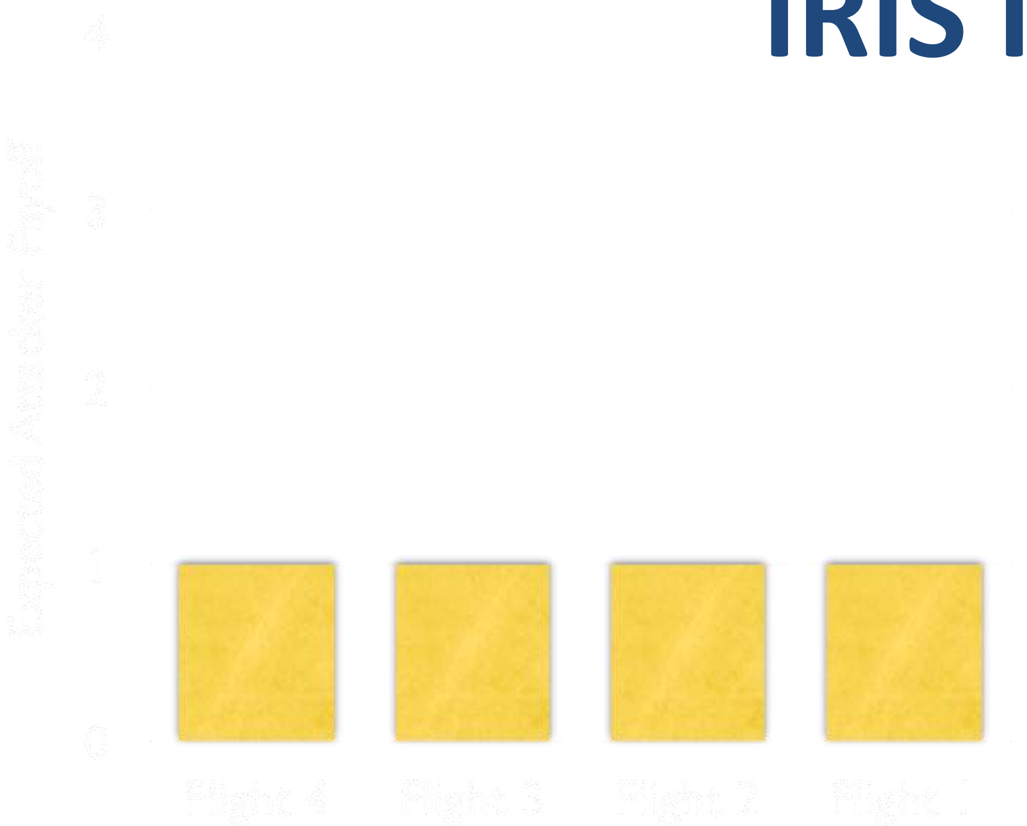

Need more than one air marshal!

0.75	0.66	0.5	0

Coverage Probability

IRIS I

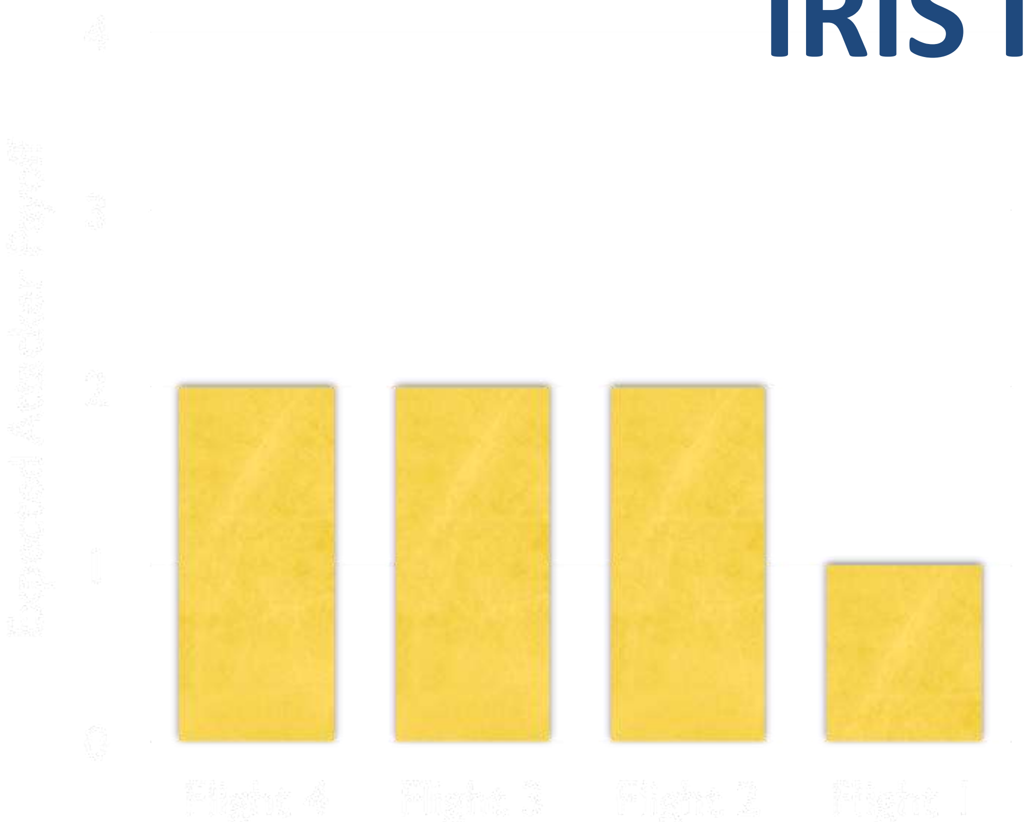

Can still assign 0.17

0.5	0.33	0	0

Coverage Probability

IRIS I

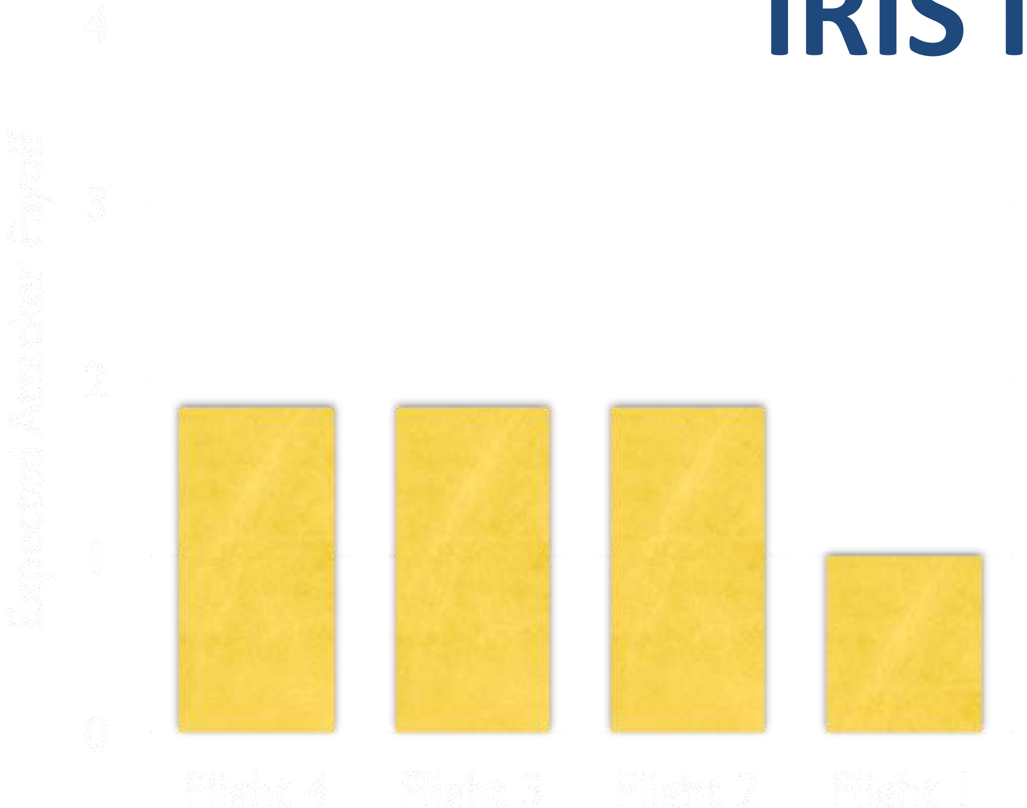

Allocate all remaining coverage to flights in the attack set

Fixed ratio necessary for indifference

0.54	0.38	0.08	0

Coverage Probability

IRIS Speedups

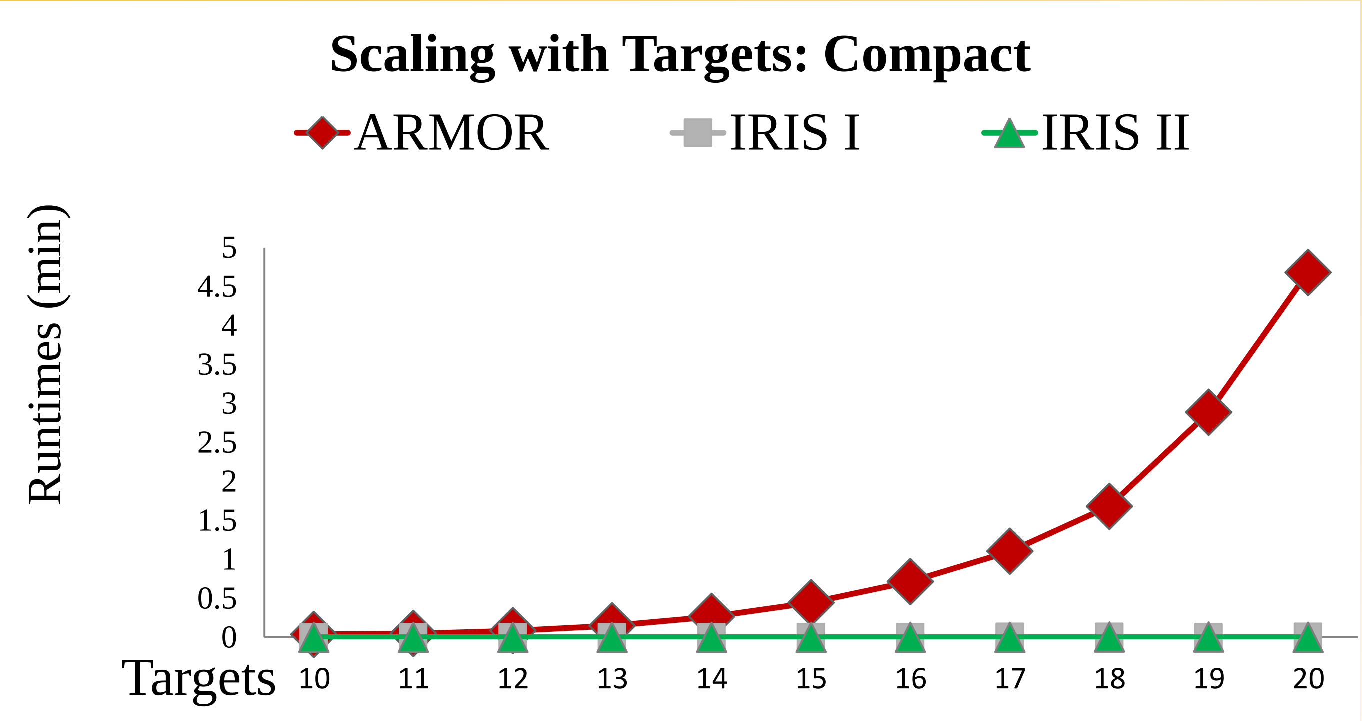

	ARMOR Actions	ARMOR Runtime	IRIS Runtime
FAMS Ireland	6,048	4.74s	0.09s
FAMS London	85,275	----	1.57s

IRIS III: Branch and Price: Tours of Arbitrary Size

Branch & Price: Branch & Bound + Column Generation

- Not out of the box

- Upper bounds: IRIS I
- Column generation:
  Network flow

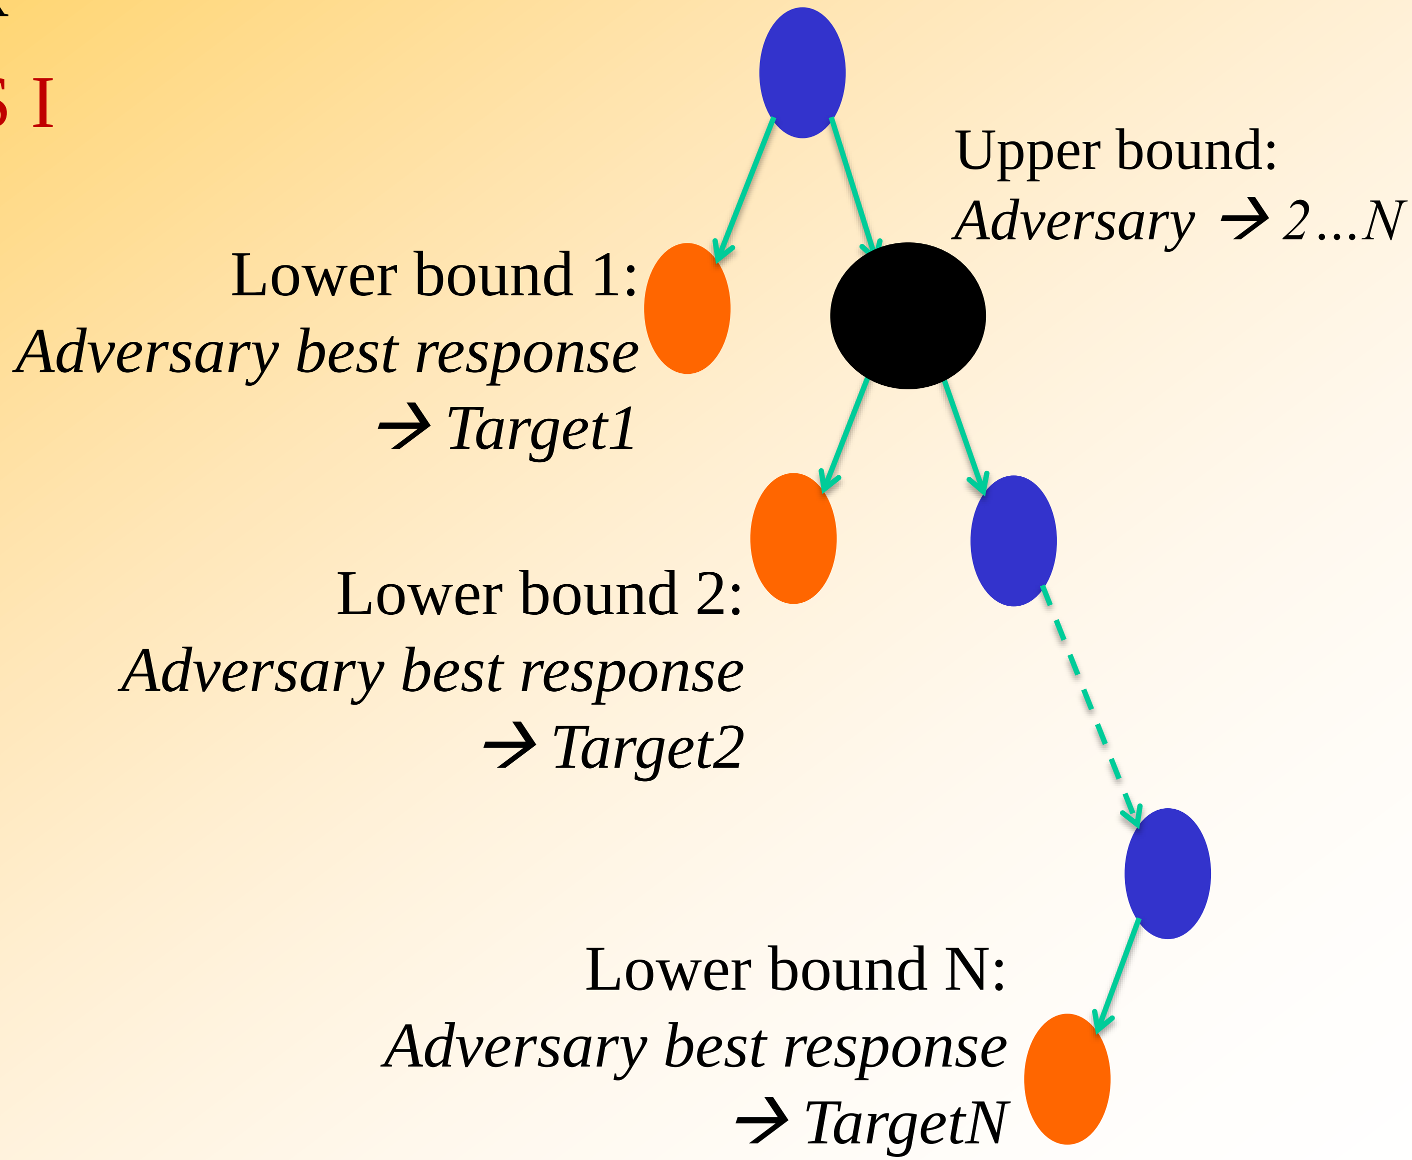

IRIS III: Branch & Price
Column Generation Quick Overview

LEAF NODE:
Incrementally build support for mixed strategy

“Master” Problem
(mixed integer program)

Solution supported by N pure strategies

$(N+1)^{th}$ pure Strategy

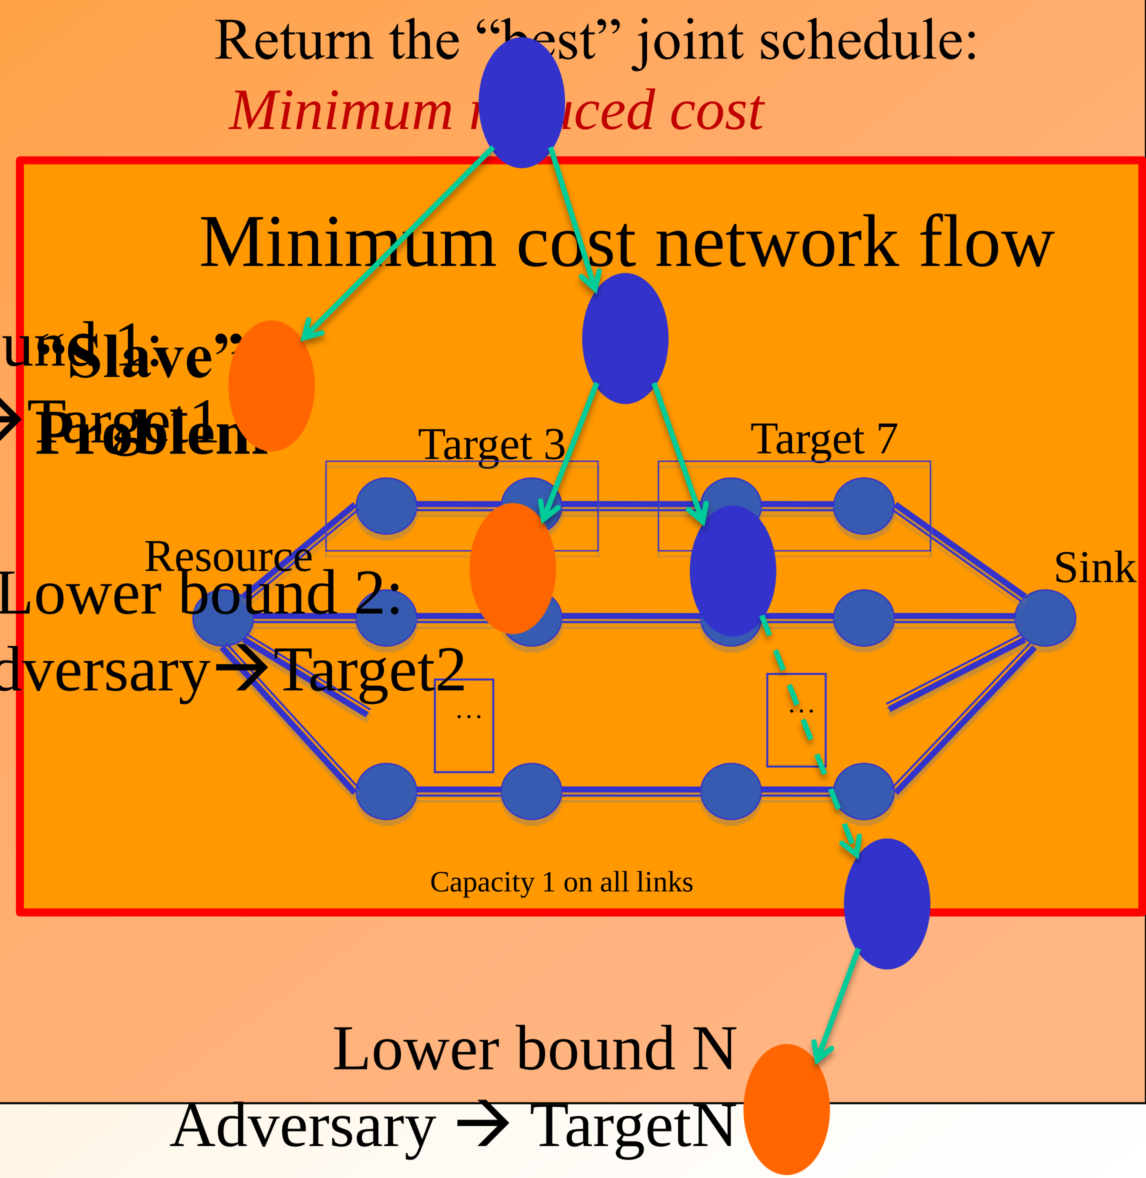

IRIS Results

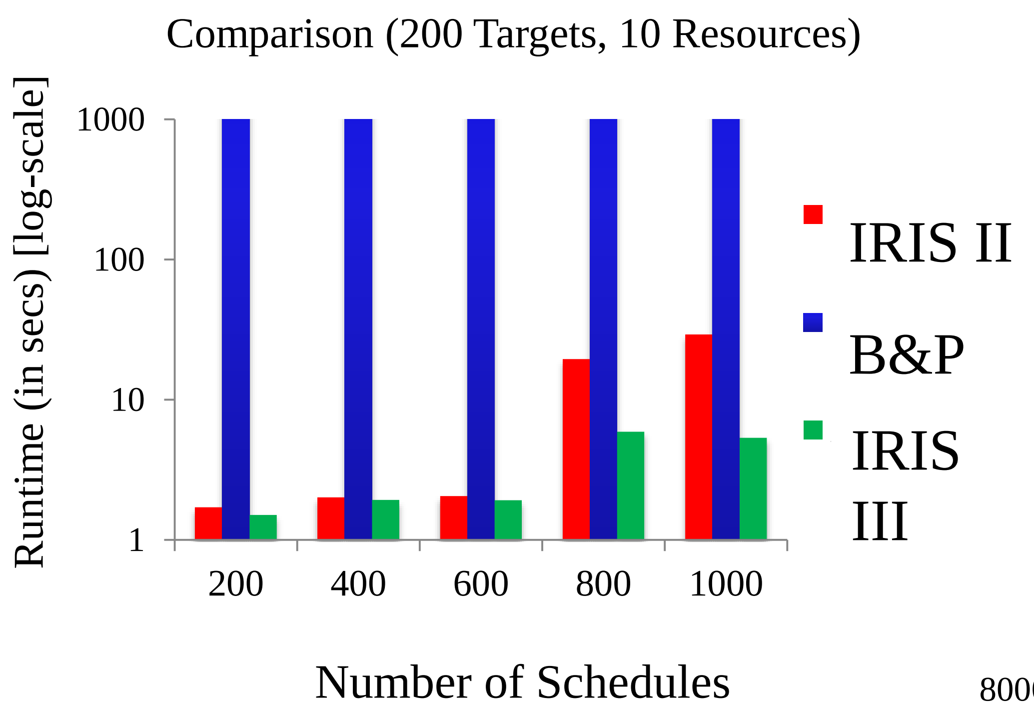

ARMOR

Runs out of memory

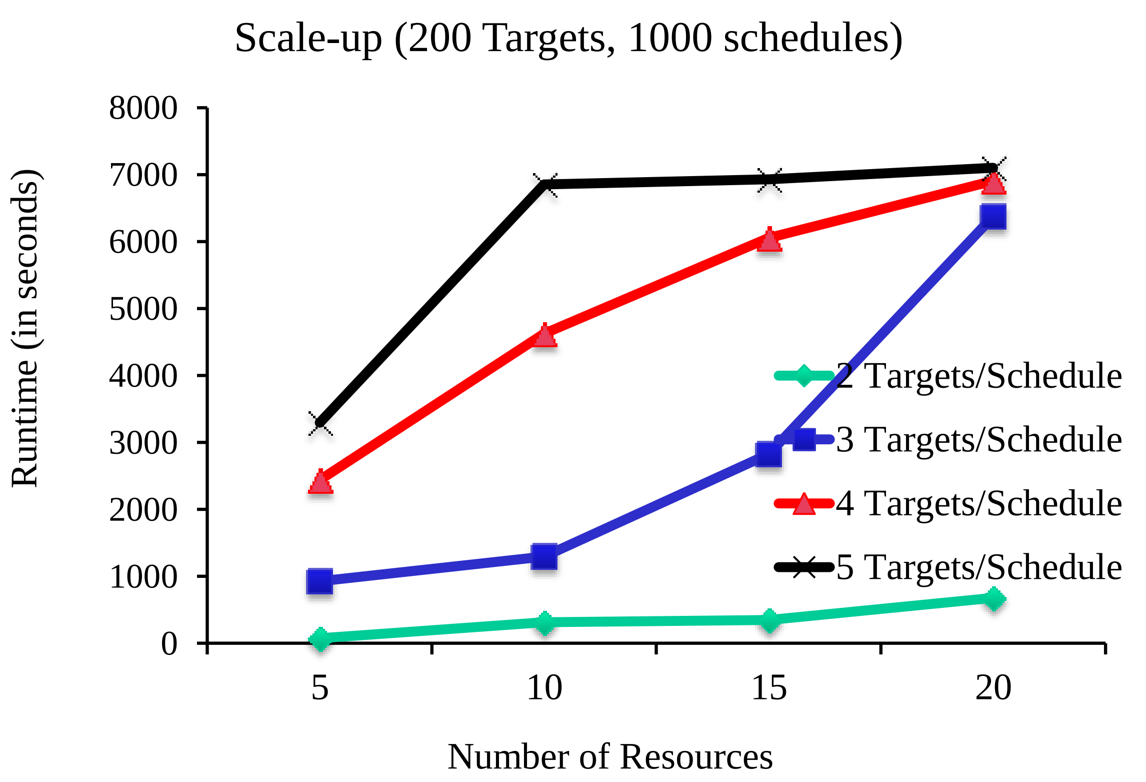

Fare Checking in LA Metro
(Yin et al. 2012)

- Los Angeles Metro Rail System
  - *Barrier-free system with random inspections*
  - *Approximately 300,000 daily riders, ≈6% fare evaders*
  - *Fare evasion costs ≈ $5.6 million annually (Booz Allen Hamilton 2007)*

TEAMCORE
USC

How to Model?


Mission
Fillmore
Del Mar
Memorial Park
Lake
Allen
Sierra Madre Villa
PASADENA
Highland Park
Southwest Museum
Heritage Sq
Lincoln/Cypress
Chinatown
Union Station
LAC+USC Medical Ctr
Cal State LA
El Monte
San B
EL MONTE
Civic Ctr
Pershing Sq
7th St/Metro Ctr
Little Tokyo/
Arts Dist
EAST LOS ANGELES
Pico/Aliso
Mariachi Plaza
Soto
Indiana
Maravilla
East LA Civic Ctr
Atlantic

How to Model?

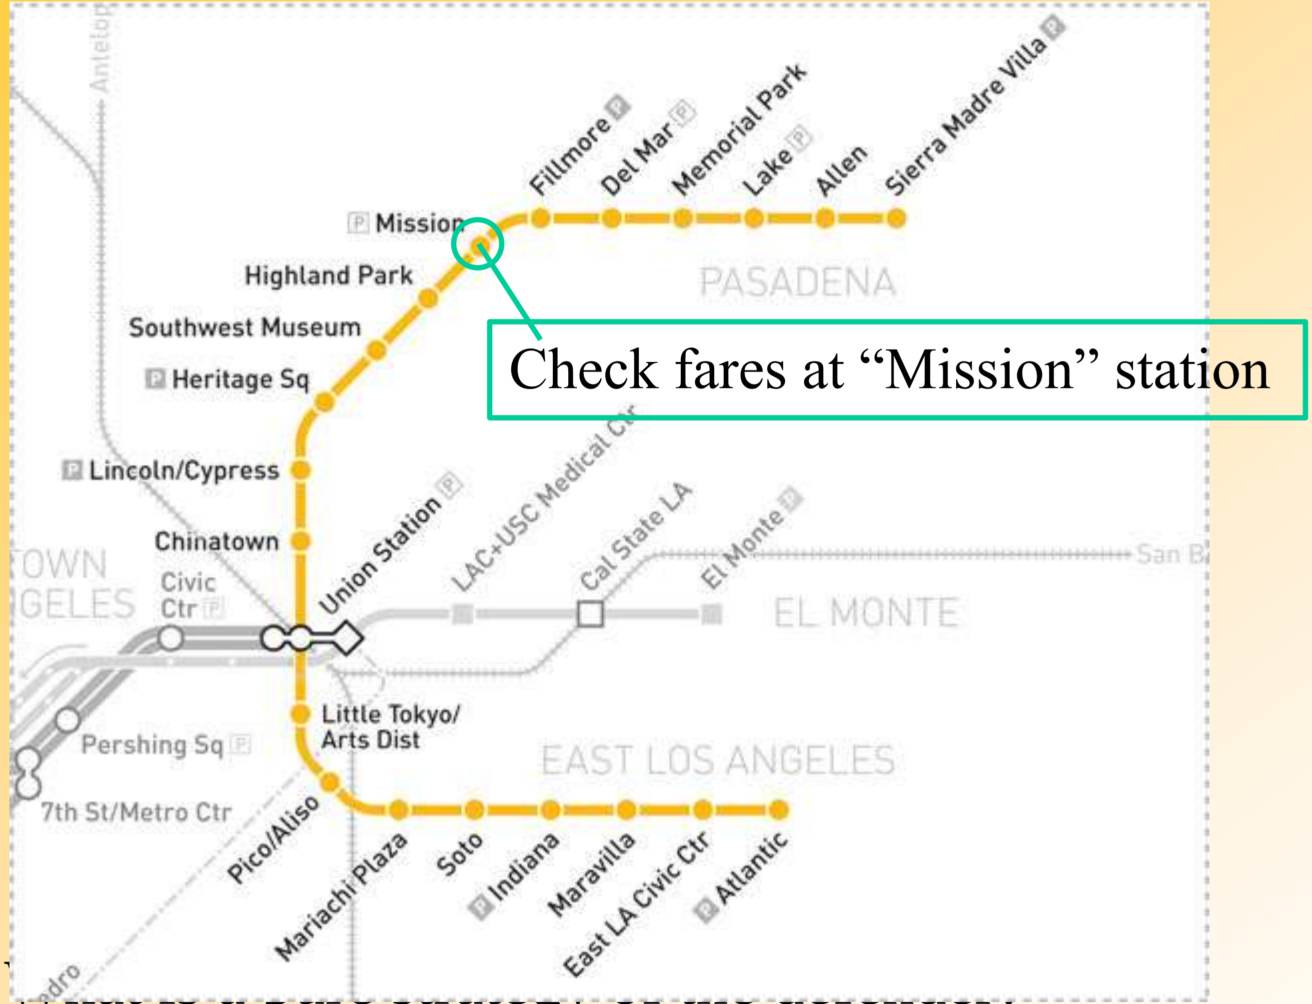

- What is a pure strategy of the defender?

How to Model?

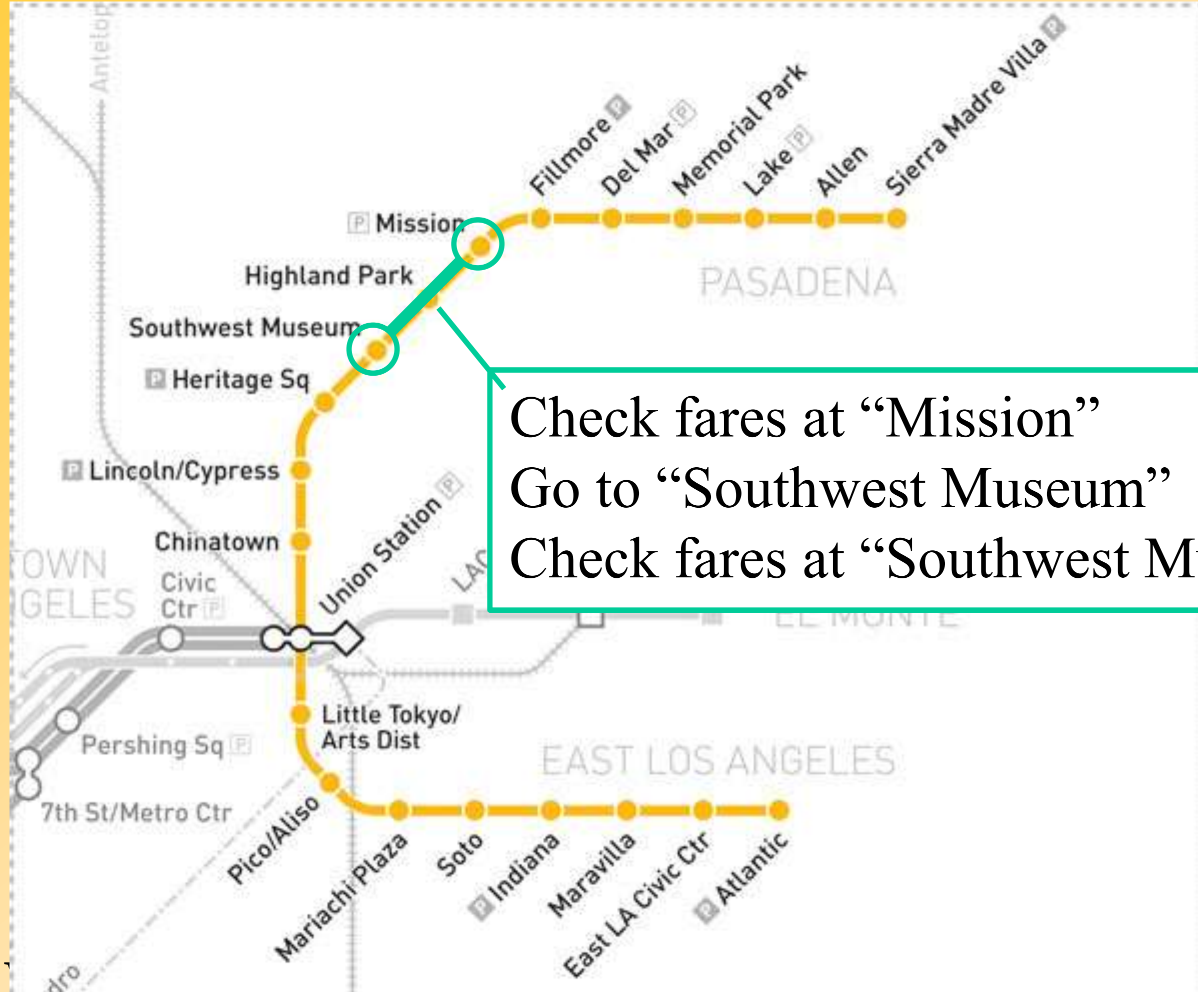

Check fares at "Mission"
Go to "Southwest Museum"
Check fares at "Southwest Museum"

TEAMCORE
USC

How to Model?

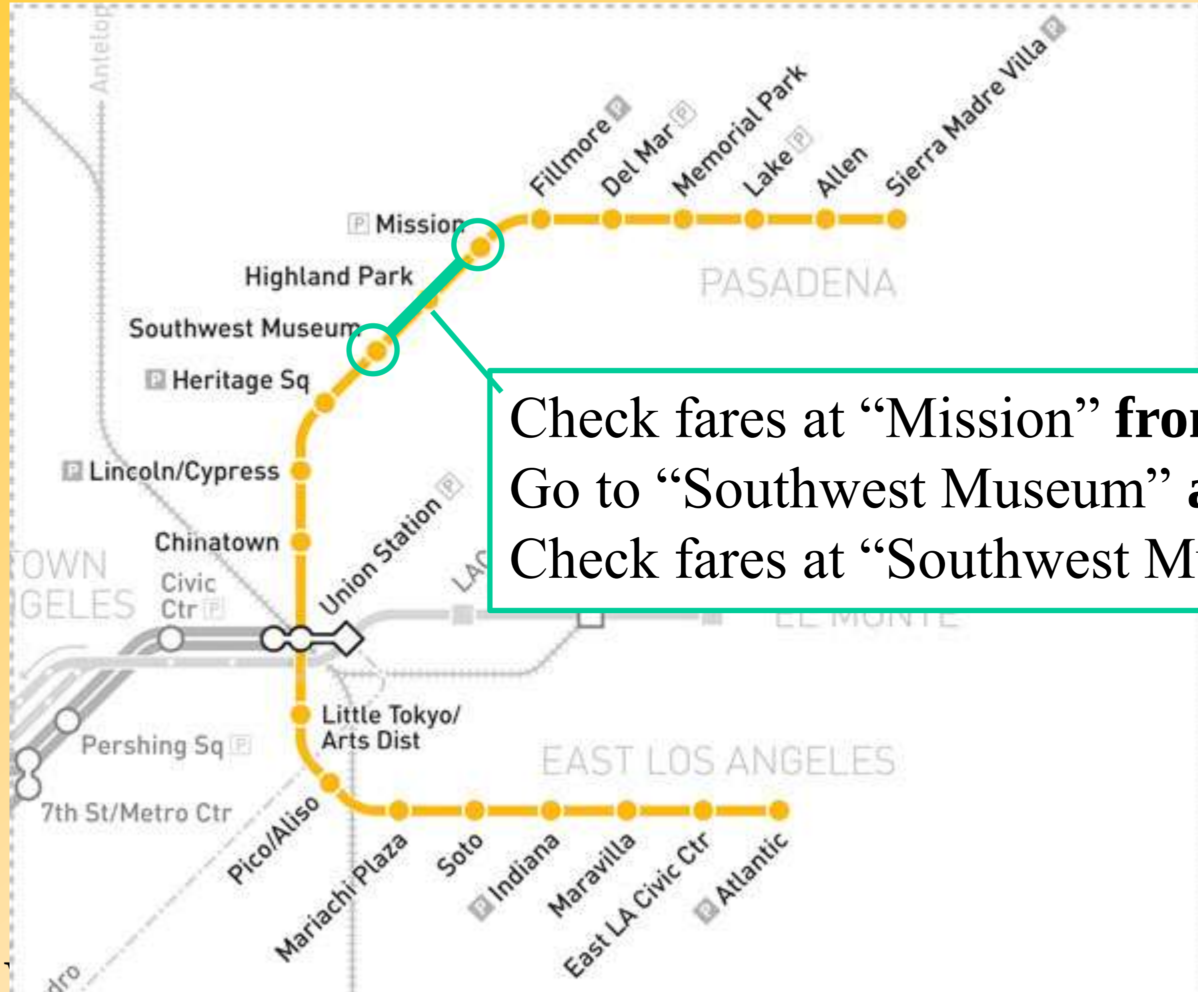

Check fares at “Mission” **from 7am to 7:50am**
Go to “Southwest Museum” **at 7:50am**
Check fares at “Southwest Museum” **from 8am to 9am**

TEAMCORE USC

How to Model?

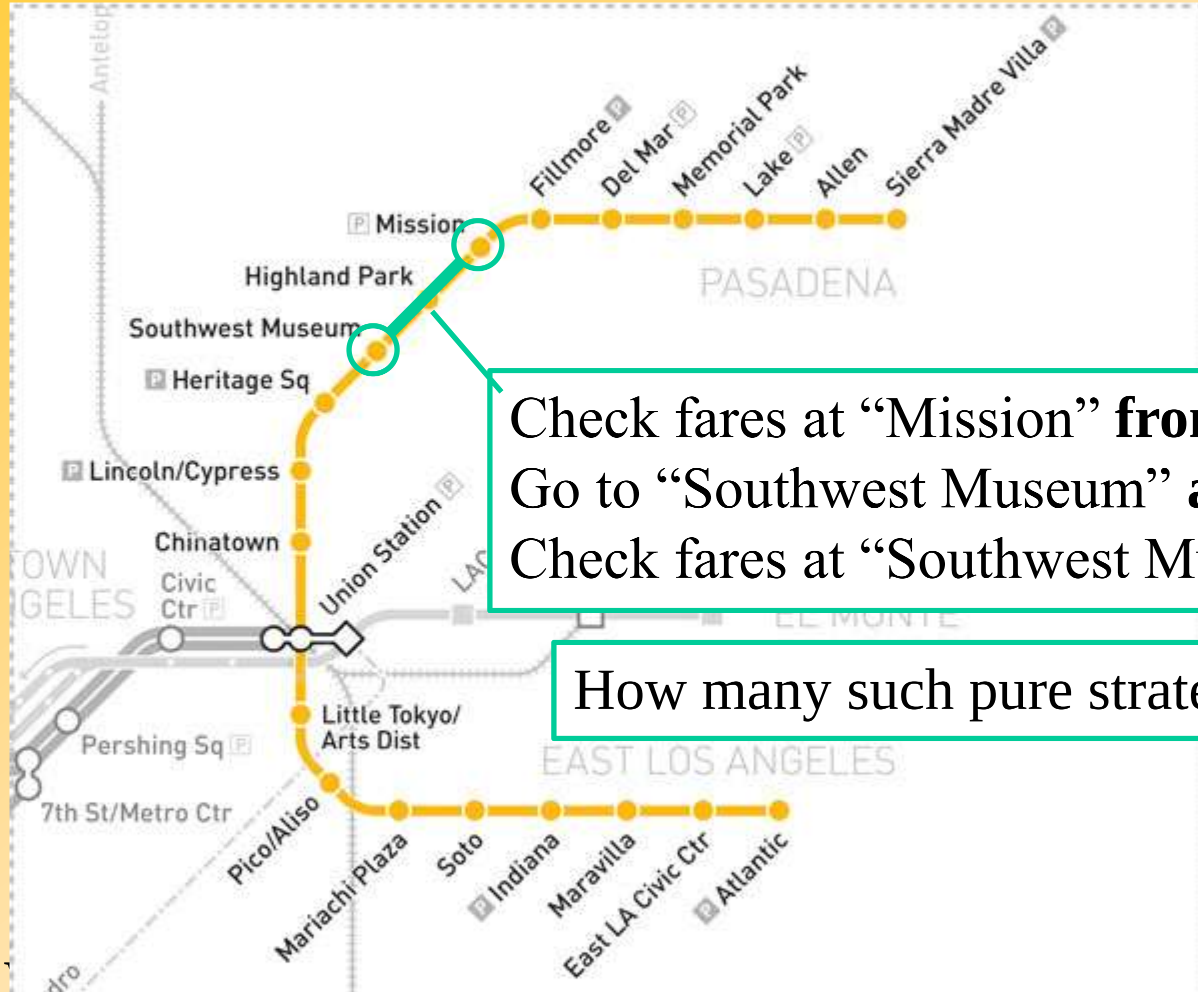

Check fares at “Mission” **from 7am to 7:50am**
Go to “Southwest Museum” **at 7:50am**
Check fares at “Southwest Museum” **from 8am to 9am**

How many such pure strategies?

TEAMCORE USC

Problem Setting

- *Transition graph*

Vertex: *station and time pair*

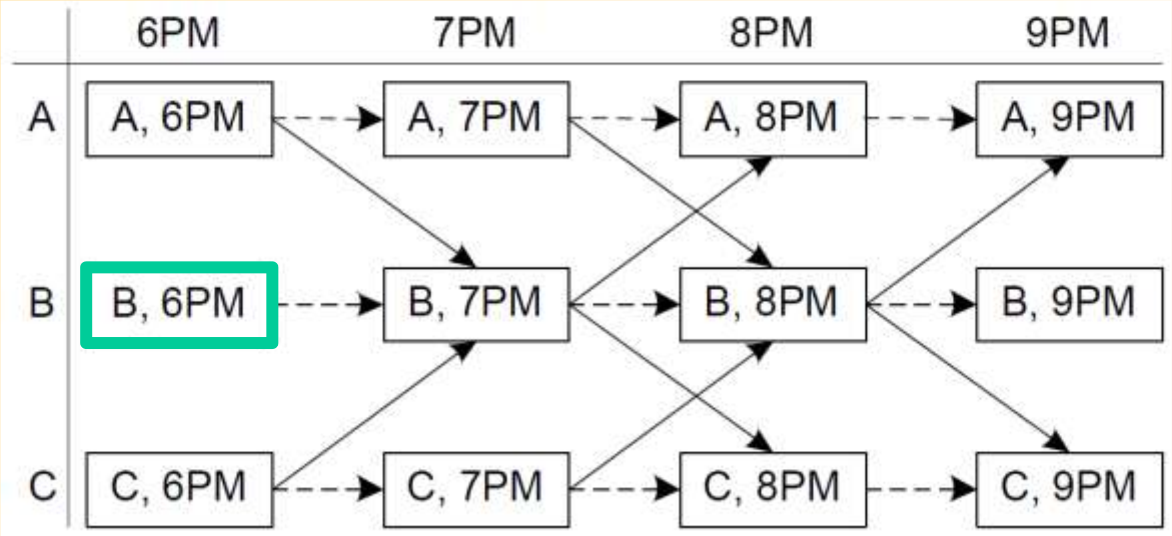

Problem Setting

- *Transition graph*

Edge: *inspection action*

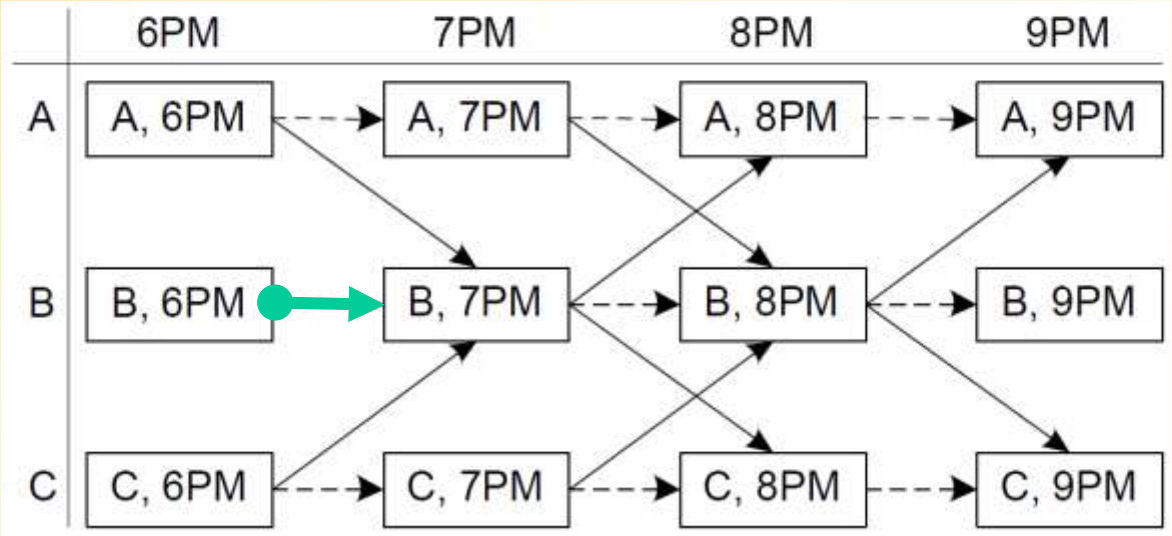

Problem Setting

- *Transition graph*

Edge: *inspection action*

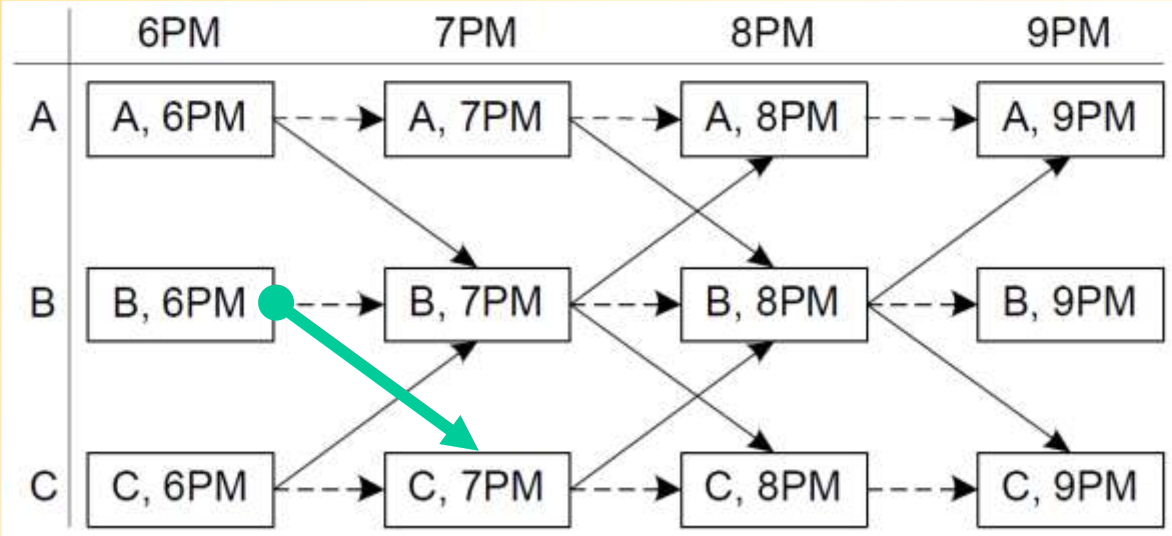

Problem Setting

- *Transition graph*

Edge: *inspection action*
l_e - action duration
f_e - fare-check effectiveness

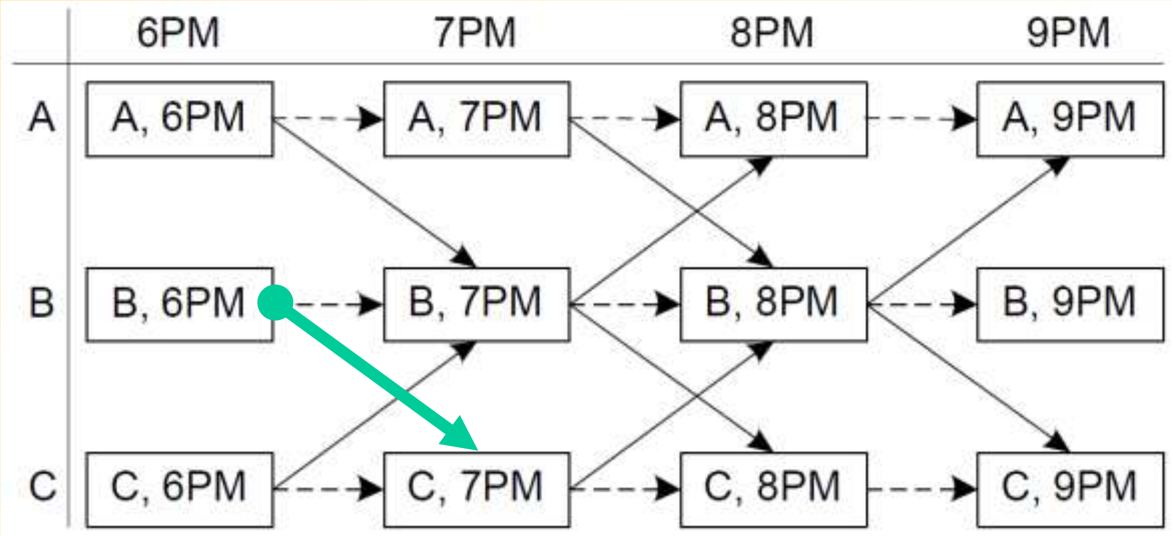

TEAMCORE
USC

Problem Setting

- *Transition graph*

Patrols: *bounded-length paths*

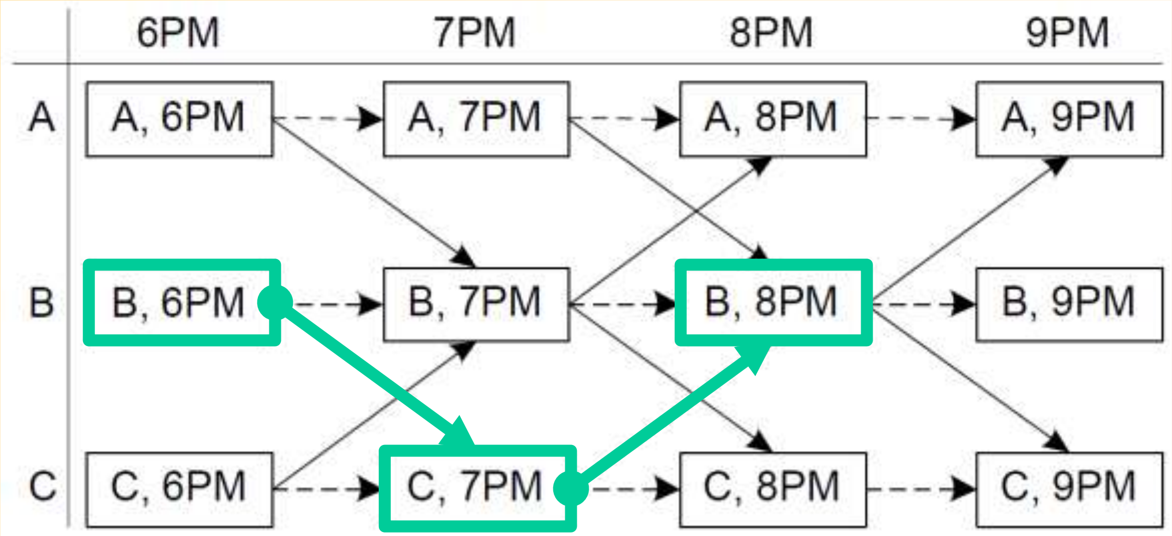

TEAMCORE USC

Problem Setting

- *Transition graph*

Patrols: *bounded-length paths*

γ – patrol units

κ – patrol hours per unit

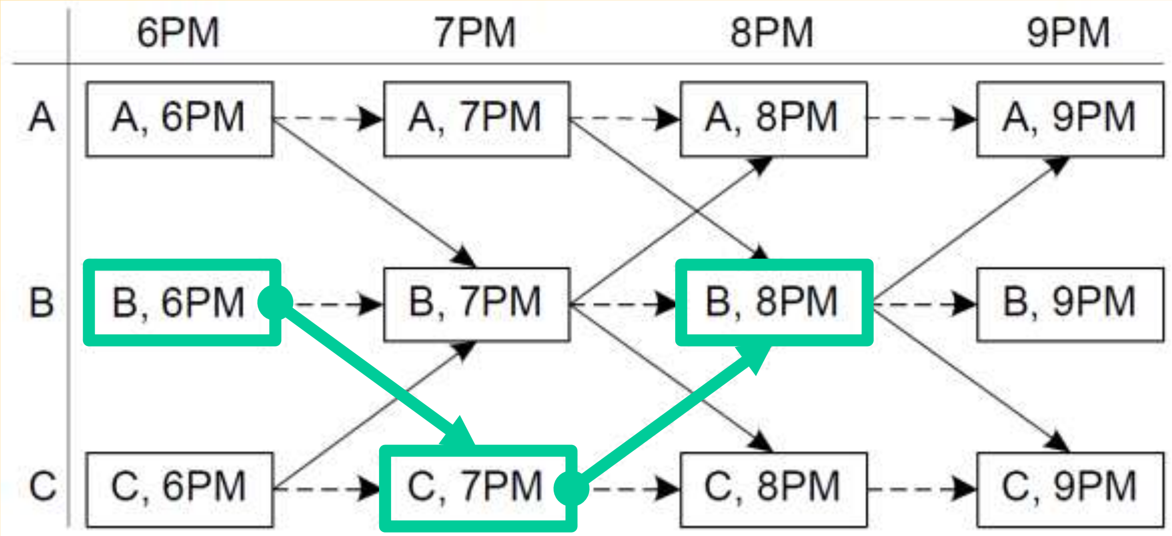

Problem Setting cont.

- Riders: *multiple types*
  - *Each type takes fixed route*
  - *Fully observes the probability of being inspected*
  - *Binary decision: buy or not buy the ticket*
  - *Perfectly rational and risk-neutral*

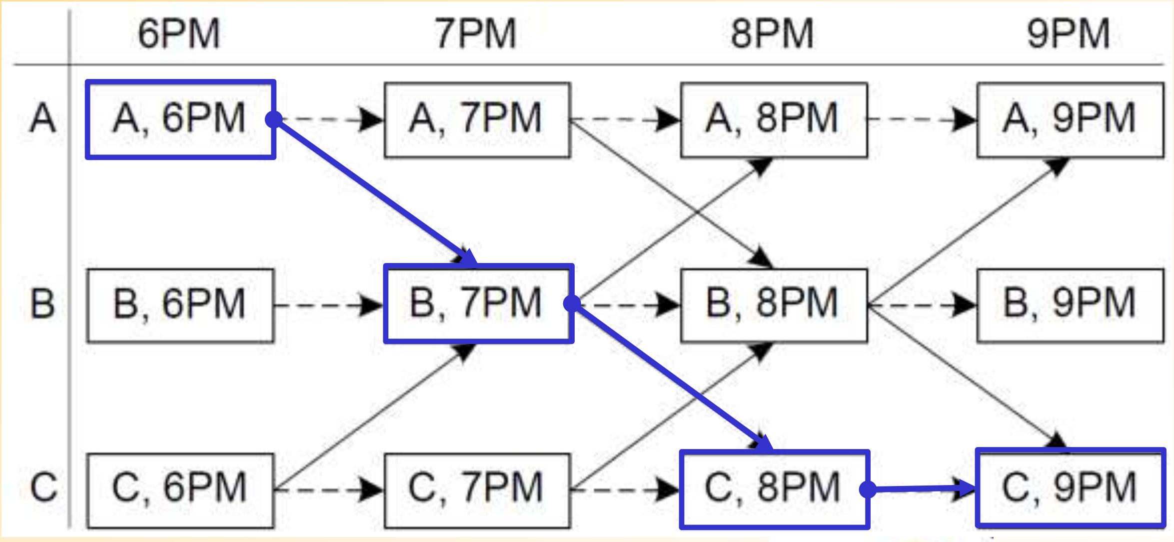

Problem Setting cont.

- Riders: *multiple types*
  - *Each type takes fixed route*
  - *Fully observes the probability of being inspected*
  - *Binary decision: buy or not buy the ticket*
  - *Perfectly rational and risk-neutral*

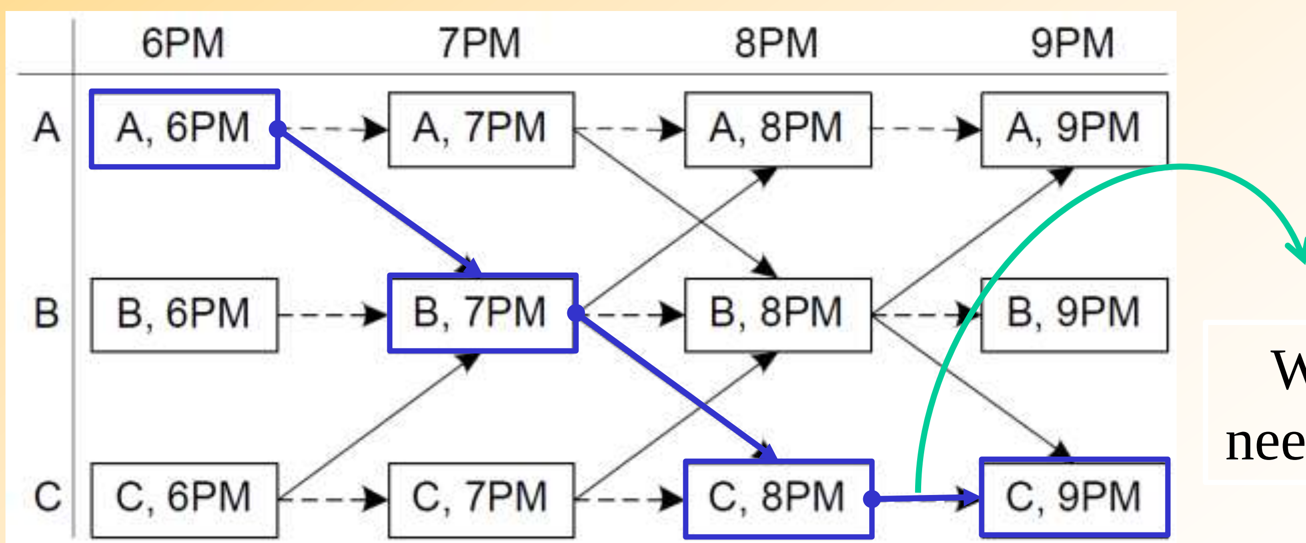

Why do we need this edge?

Basic Compact Formulation

- Based on *transition graph*
- Strategy representation: *marginal coverage on edges*

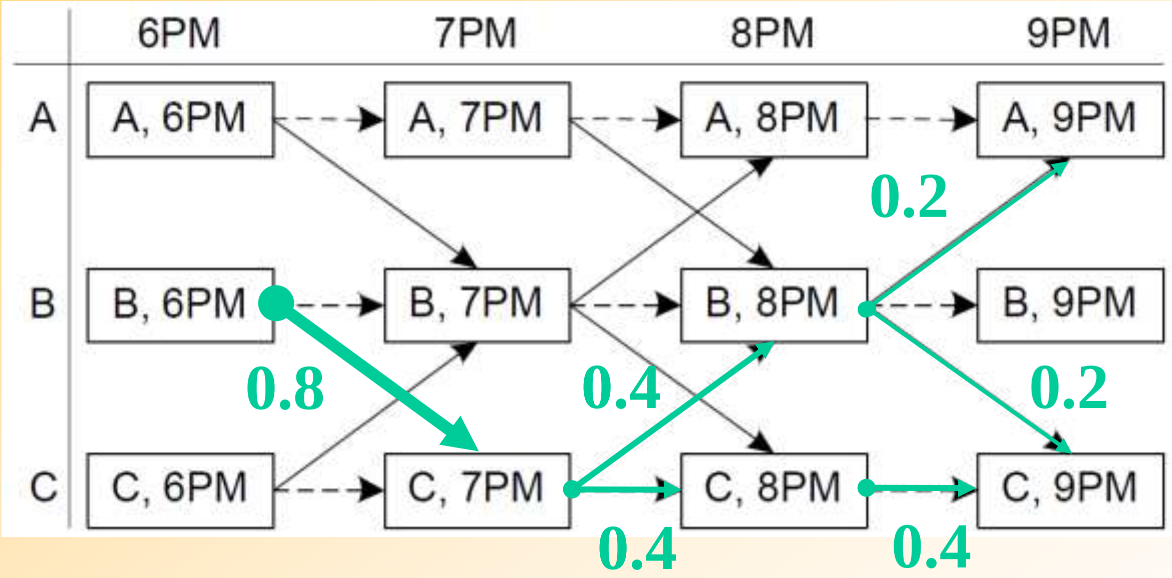

Basic Compact Formulation

- Based on *transition graph*
- Strategy representation: *marginal coverage on edges*

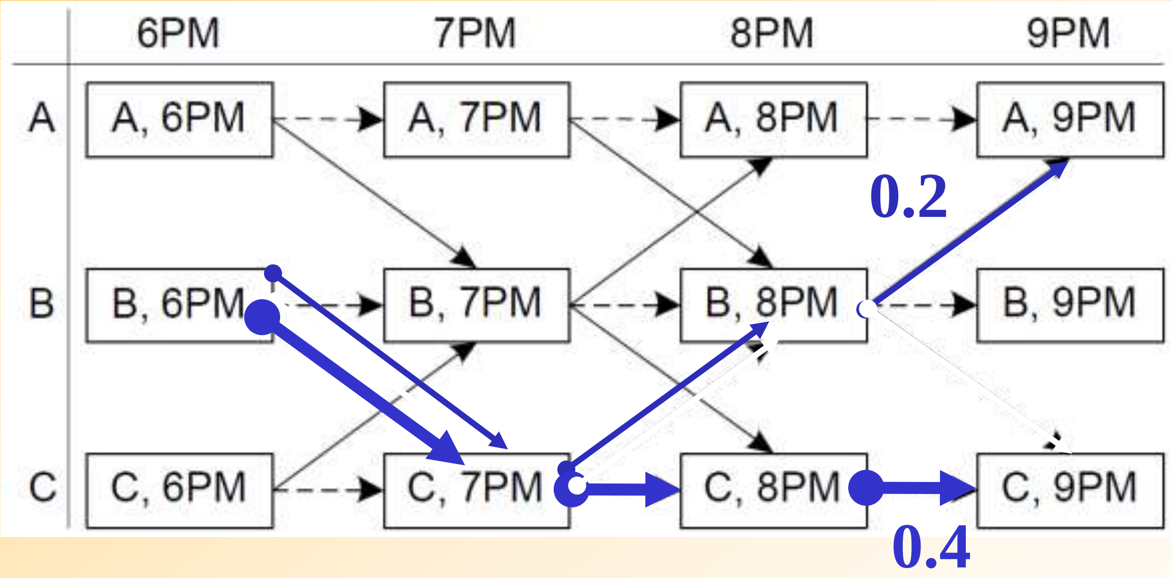

TEAMCORE USC

Basic Compact Formulation

- *Transition graph*: G = <V, E>
  - *Dummy source v$^+$, possible starting vertices V$^+$*
  - *Dummy sink v$^-$, possible ending vertices V$^-$*

$$\max_{\mathbf{x},\mathbf{u}} \sum_{\lambda\in\Lambda} p_\lambda u_\lambda \quad (2)$$

$$\text{s.t. } u_\lambda \le \min\{\rho,\ \tau \sum_{e\in\lambda} x_e f_e\},\ \text{for all } \lambda \in \Lambda \quad (3)$$

$$\sum_{v\in V+} x_{(v^+,v)} = \sum_{v\in V-} x_{(v,v^-)} \le \gamma \quad (4)$$

$$\sum_{(v',v)\in E} x_{(v',v)} = \sum_{(v,v^\dagger)\in E} x_{(v,v^\dagger)},\ \text{for all } v \in V \quad (5)$$

$$\sum_{e\in E} l_e \cdot x_e \le \gamma \cdot \kappa, 0 \le x_e \le \alpha, \forall e \in E \quad (6)$$

TEAMCORE
USC

Issues with Basic Compact Formulation

- Patrol length may not be bounded by κ
  - *E.g., $\gamma = 1$, $\kappa = 1$*

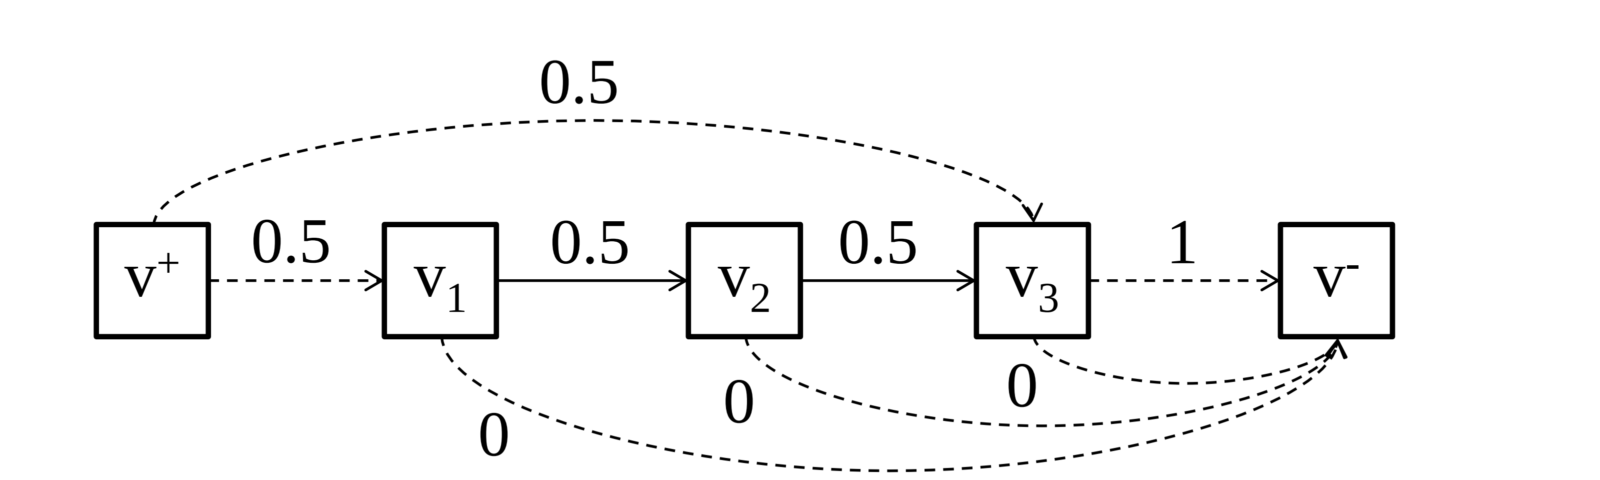

$$\sum_{v \in V+} x_{(v^+,v)} = \sum_{v \in V-} x_{(v,v^-)} \leq \gamma \quad (4)$$

$$\sum_{(v',v) \in E} x_{(v',v)} = \sum_{(v,v^\dagger) \in E} x_{(v,v^\dagger)}, \text{ for all } v \in V \quad (5)$$

$$\sum_{e \in E} l_e \cdot x_e \leq \gamma \cdot \kappa, 0 \leq x_e \leq \alpha, \forall e \in E \quad (6)$$

Issues with Basic Compact Formulation

- Patrol length may not be bounded by κ
  - *E.g.,* $\gamma = 1, \kappa = 1$

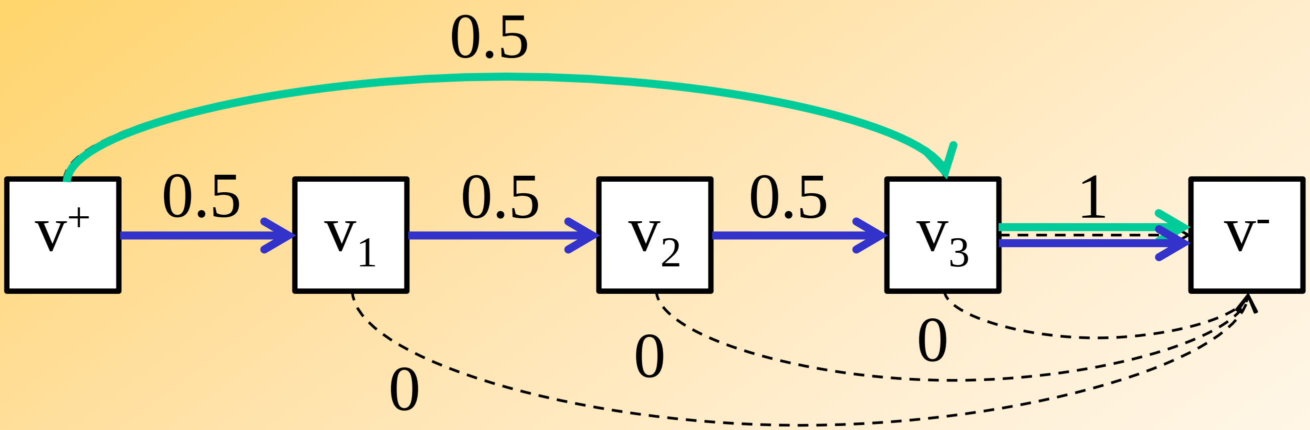

  - **0.5,** $v^+ \rightarrow v_3 \rightarrow v^-$
  - **0.5,** $v^+ \rightarrow v_1 \rightarrow v_2 \rightarrow v_3 \rightarrow v^-$

TEAMCORE USC

Extended Compact Formulation

- *History-duplicate transition graph*
  - *Store history information in vertices*
  - *Access necessary patrol information without exponential blowup*

Extended Compact Formulation cont.

- *History-duplicate transition graph*
  - *Store history information in vertices*
  - *Access necessary patrol information without exponential blowup*
- E.g., to forbid patrols longer than *2* hours
  - *What information should be duplicated?*

TEAMCORE
USC

Extended Compact Formulation cont.

- *History-duplicate transition graph*
  - *Store history information in vertices*
  - *Access necessary patrol information without exponential blowup*
- E.g., to forbid patrols longer than *2* hours
  - *2 subgraphs corresponding to 2 starting time:* ***6pm*** *and* ***7pm***

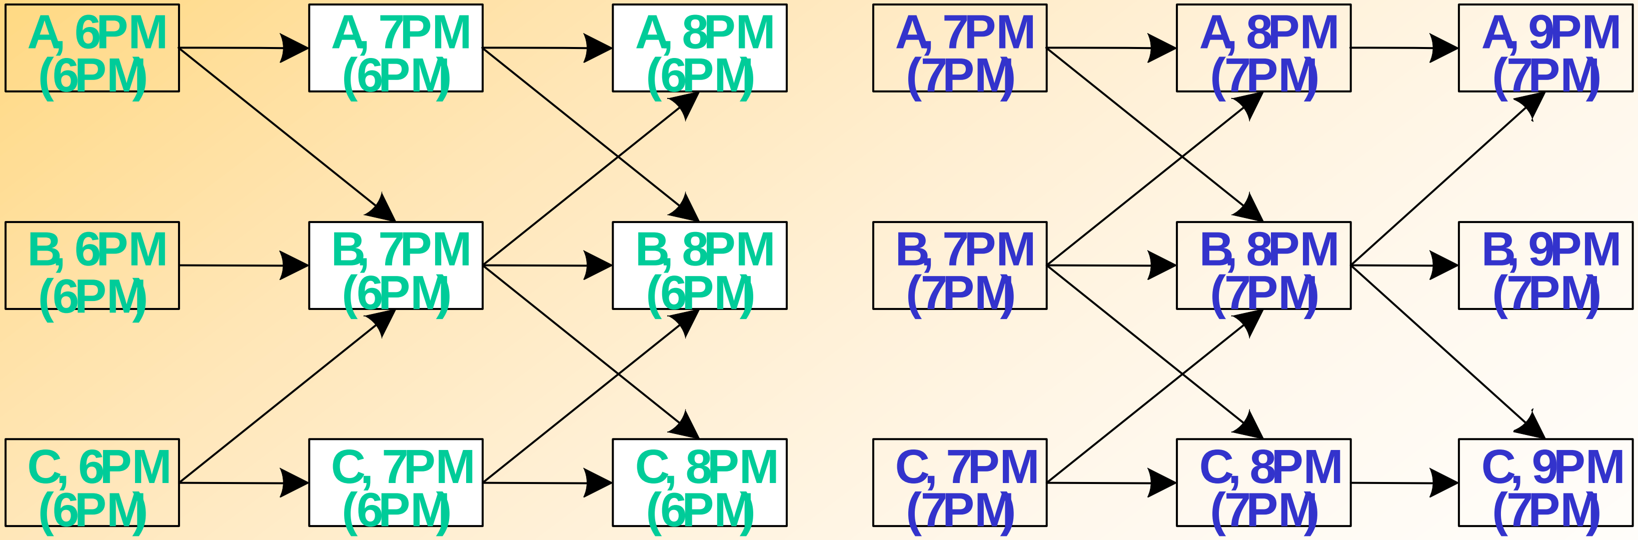

TEAMCORE
USC

Outline

- Motivating real-world applications
- Background and basic security games
- Scaling to complex action spaces
- ***Modeling payoff uncertainty: Bayesian Security Games***
- Human behavior and observation uncertainty
- Evaluation and discussion

Robustness

	Target 1	Target 2	Target 3	Target 4
Defender Reward	1	0	-1	3
Defender Penalty	-1	-4	-6	-10
Attacker Penalty	-2	-3	-3	-5
Attacker Reward	1	3	5	9

How do we know the model is correct?

If it is not exactly correct, how robust is the solution?

Estimating Target Values

What is the **attacker's value** for a successful attack on a particular target?

- *What is the likely number of casualties?*
- *What is the economic cost?*
- *What is the value of the media exposure?*
- *What is the symbolic value of the attack?*
- *How should these factors be weighted?*

Answers can only be *estimated*

Modeling Choices

Players

- How many?
- Model organizations as individuals?
- Specific people or generic types of people?
- Are players rational?
- If not, how do they behave?

Actions

- What is the set of feasible actions?
- Do players know all of the actions?
- If the set is infinite, how do we represent it?
- Are some actions similar to others?
- Are actions sequential?

Payoffs

- How do we determine payoffs?
- Are payoffs known to all players?
- What is the uncertainty about the payoffs?
- Are payoffs deterministic or stochastic?
- Do players care about risk?

Solution concepts

- What to do if there are multiple equilibria?
- Do we care about the worst case?
- Bounded rationality
- Limited observability
- Can the solution be computed?

Robustness Perspectives

- Game theorist's perspective
  - *The model is given, and known to everyone*
  - *We can model uncertainty explicitly by making the model more complex*
- Engineer's perspective:
  - *Do the math*
  - *Add a "fudge factor" to for safety*
  - *The cost is worth the risk reduction*
  - *"Unknown unknowns"*
  - *Confidence is critical*

Real problems force us to deal with robustness

Research on Robustness

- Payoff uncertainty
  - *Conitzer et al 2006, Paruchuri et al 2008, Kiekintveld et al 2011, Jain et al 2011, Yin et al 2012, Kiekintveld at al 2012, Brown et al 2012, …*
- Human behavior
  - *Jain et al 2008, Pita et al 2009, Pita et al 2010, Yang et al 2011, Pita et al 2012, Yang et al 2012, …*
- Observation/Execution uncertainty
  - *Yin et al 2010, Pita et al 2011, Yin et al 2011, An et al 2012, …*

Diverse Techniques

Bayesian Models

Finite Models
Infinite Models

Interval Models

Modified Strategy Models

Finite Bayesian Games

P=0.3

	Term #1	Term #2
Term#1	5, -3	-1, 1
Term#2	-5, 5	2, -1

P=0.5

	Term #1	Term #2
Term#1	2, -1	-3, 4
Term#2	-2, 1	3, -3

P=0.2

	Term #1	Term #2
Term#1	4, -2	-1,0.5
Term#2	-4, 3	1.5, -0.5

	111	121	112	211	…	…	…	222
Terminal #1	3.3,-2.2	2.3,…						
Terminal #2	-3.8,2.6	…,…						

Harsanyi Transformation

NP-Hard

Multiple LPs Method

[Conitzer and Sandholm 2006]

- First optimization formulation for FBSG
- Basic idea:
  - *Enumerate attacker pure strategies*
  - *Solve an LP to maximize leader's payoff*

$$
\begin{aligned}
\max_{a \in A_2} \max_{\sigma_1} \quad & \textstyle\sum_{a' \in A_1} p_1(a') u_1(a', a) \\
\text{s.t.} \quad & \textstyle\sum_{a' \in A_1} p_1(a') u_2(a', a) \geq \displaystyle\sum_{a' \in A_1} p_1(a') u_2(a', a'') \quad \forall a'' \in A_1 \\
& \textstyle\sum_{a \in A_1} p_1(a') = 1 \\
& p_1(a) \geq 0 \quad \forall a \in A_1
\end{aligned}
$$

Finite Bayesian Stackelberg Games

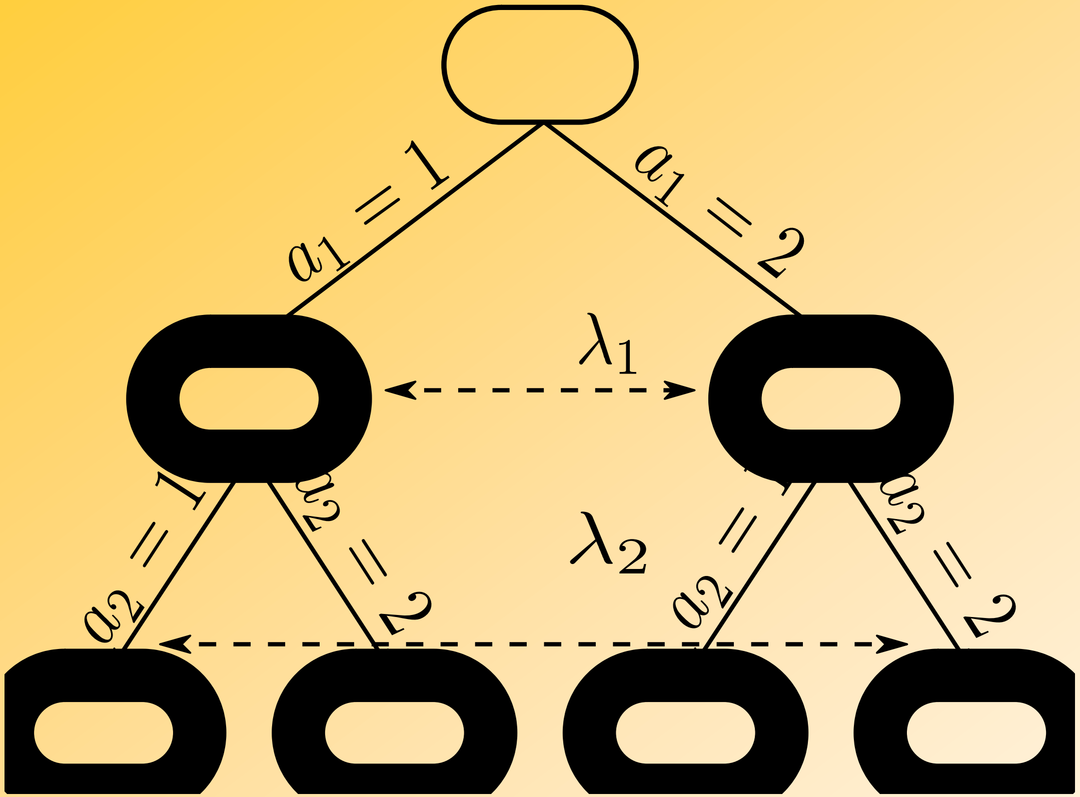

Attacker

Type λ_1

	a_1	a_2
d_1	5, -3	-1, 1
d_2	-5, 5	2, -1

p_1

Type λ_2

	a_1	a_2
d_1	1, -2	-2, 3
d_2	-3, 5	3, -1

p_2

Challenge: Exponential number of type combinations

Handling Multiple Adversary Types: ARMOR

P=0.3 

	Term #1	Term #2
Term#1	5, -3	-1, 1
Term#2	-5, 5	2, -1

P=0.5 

	Term #1	Term #2
Term#1	2, -1	-3, 4
Term#2	-[illegible], 1	3, -3

P=0.2 

	Term #1	Term #2
Term#1	4, -2	-1,0.5
Term#2	-4, 3	1.5, -0.5

$$\max_{x,q} \sum_{i\in X}\sum_{l\in L}\sum_{j\in Q} p^l R^l_{ij} x_i q^l_j$$

$$s.t. \sum_i x_i = 1, \sum_{j\in Q} q^l_j = 1$$

$$0 \le (a^l - \sum_{i\in X} C^l_{ij} x_i) \le (1-q^l_j)M$$

$$x_i \in [0\ldots1], q^l_j \in \{0,1\}$$

ARMOR: Run-time Results

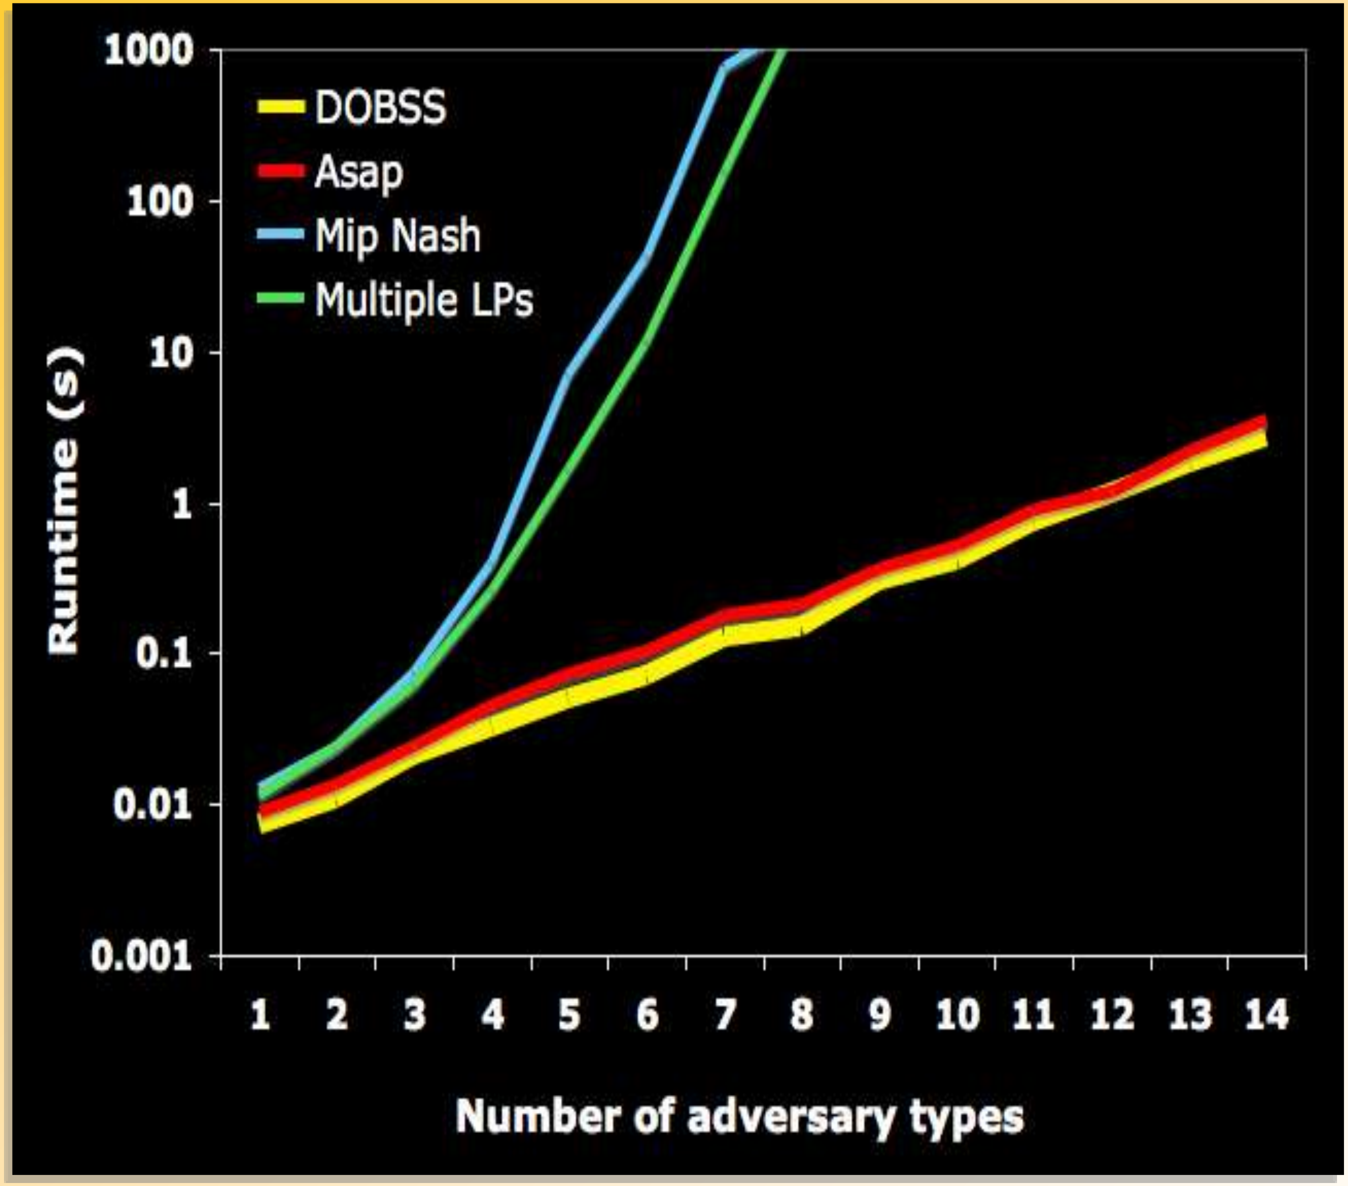

- *Multiple LPs (Conitzer & Sandholm'06)*
- *MIP-Nash (Sandholm et al'05)*
- *Sufficient for LAX*

Scaling Up: Hierarchical Solver (HBGS)

[Jain et al. 2011]

- Efficient tree search
  - Bounds and pruning
  - Branching heuristics
- Evaluate fewer LPs
- Column generation
  - Consider restricted games
  - Solve much smaller LPs

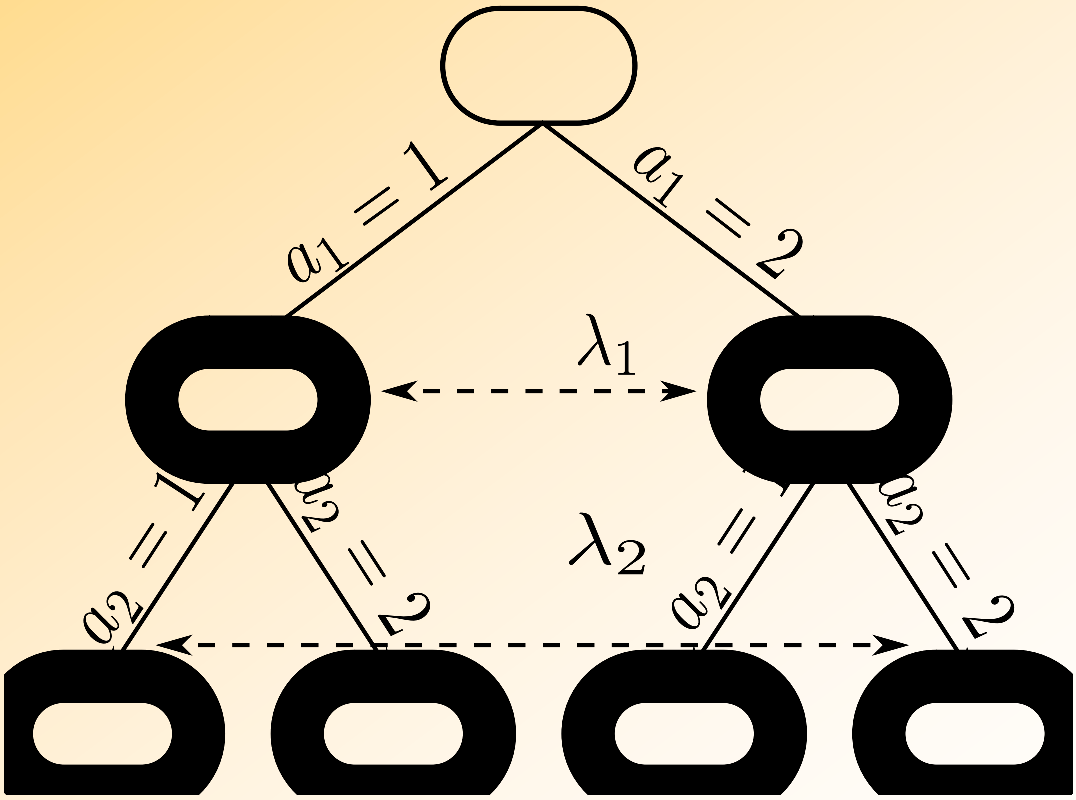

Scaling Up: Hierarchical Solver (HBGS)

- Key Idea: solve restricted games (few types)
- Use solutions to generate bounds/heuristics

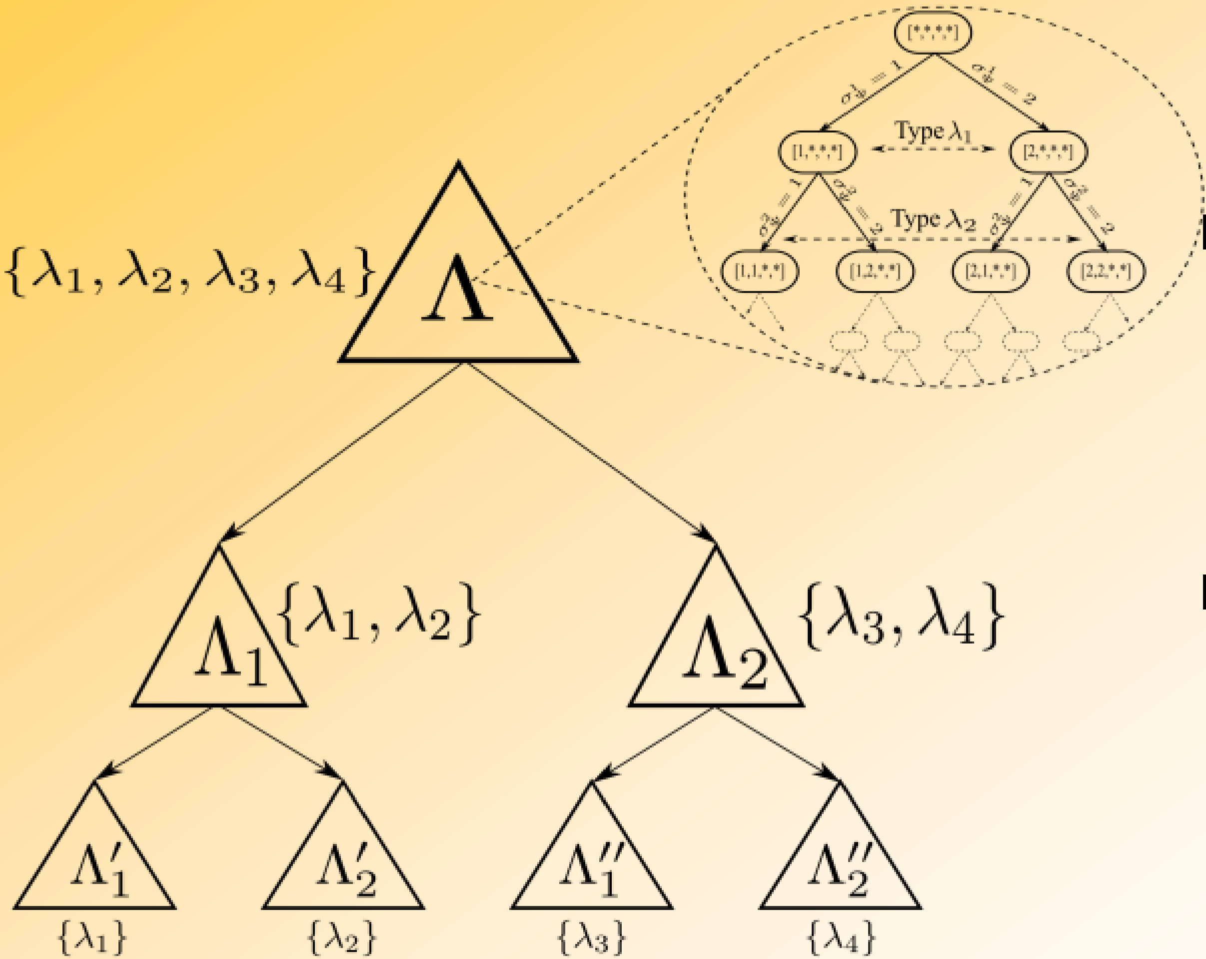

- Each node in this tree represents a full Bayesian Stackelberg game
- Can use column generation to solve these nodes

Pruning

- *Theorem 1:* If a pure strategy is infeasible in a "restricted" game, all its combinations are infeasible in the Bayesian game.

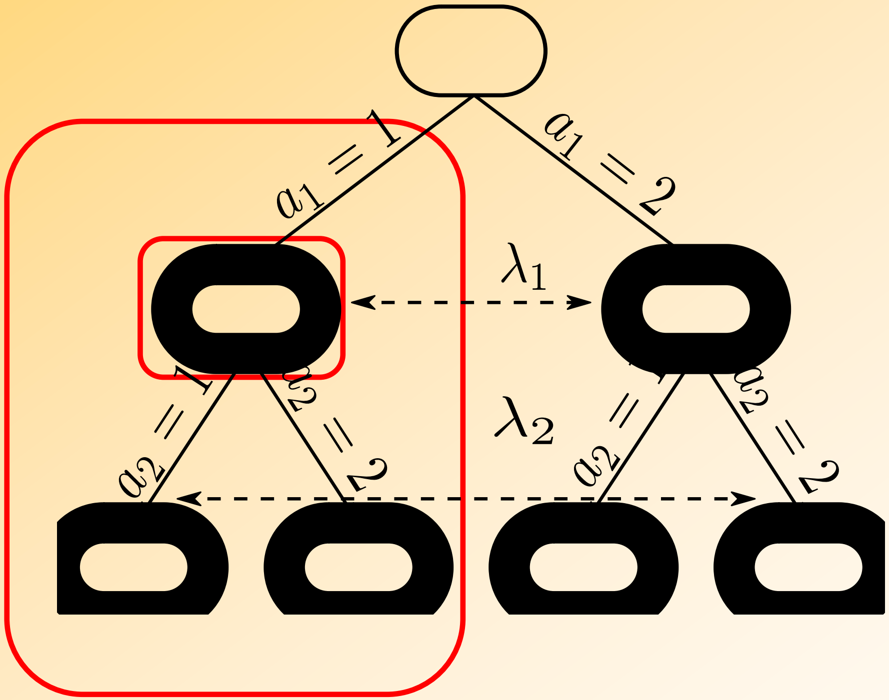

Bounds and Branching Rules

- *Theorem 2:* Leader payoff in the Bayesian game is upper bounded by the sum of leader payoffs in the corresponding restricted games.

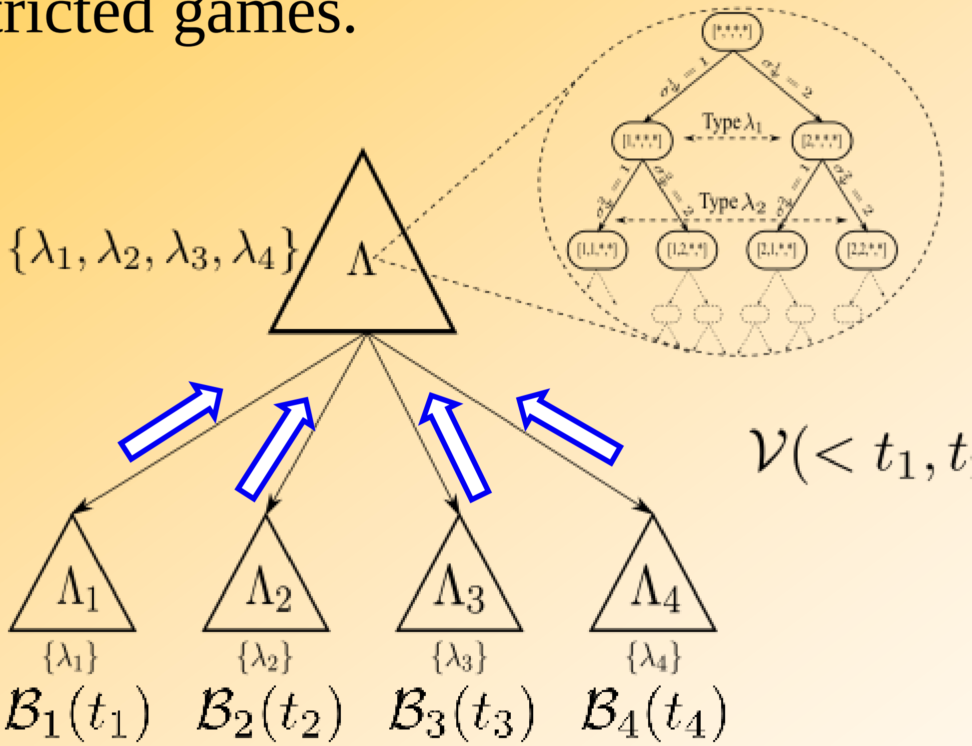

$$\mathcal{V}(< t_1, t_2, t_3, t_4 >) \leq \sum_{\lambda \in \Lambda} p_\lambda \mathcal{B}_\lambda(t_\lambda)$$

Column Generation

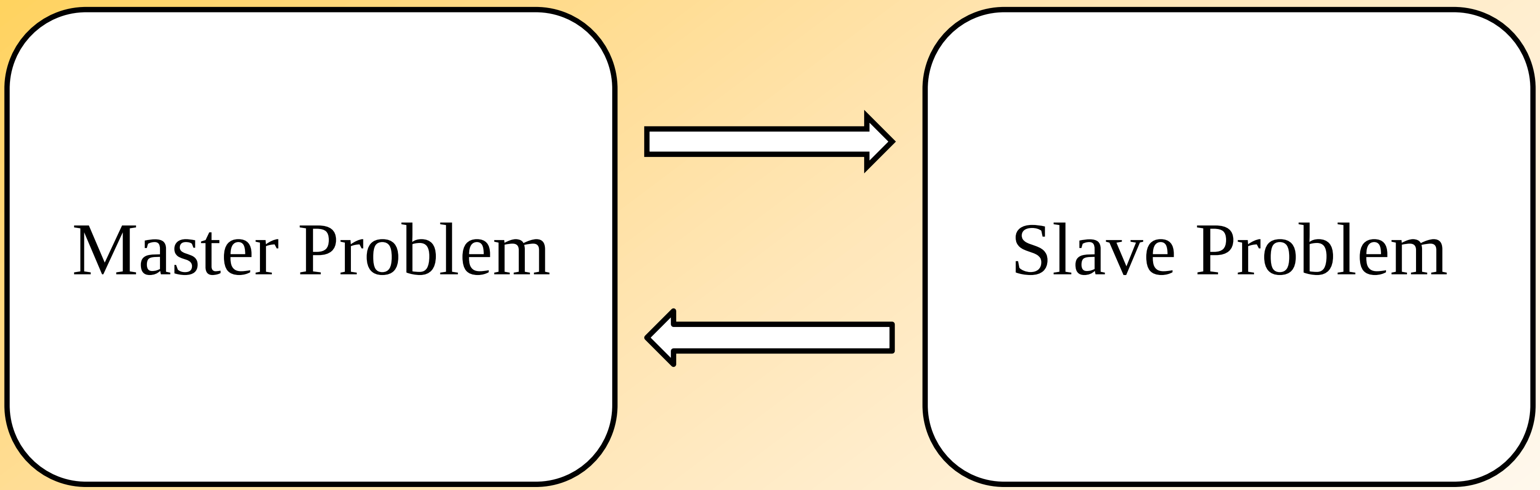

Defender and Attacker Optimization Constraints

Scheduling Constraints

HBGS Results

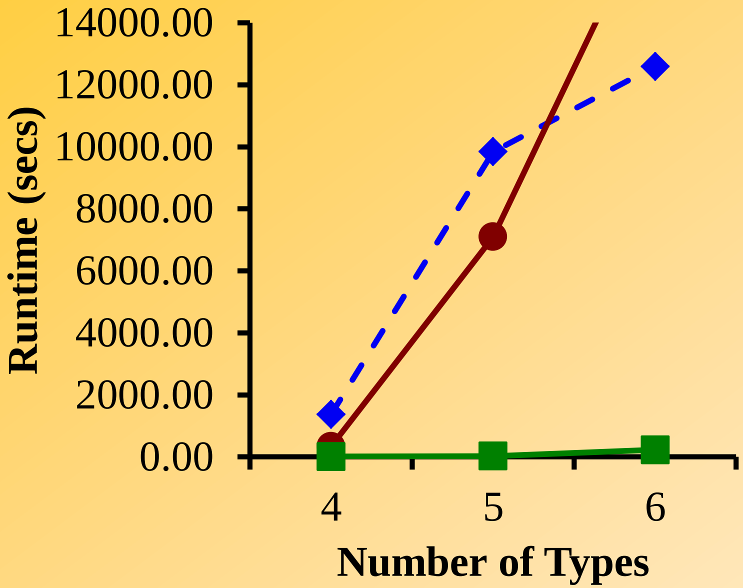

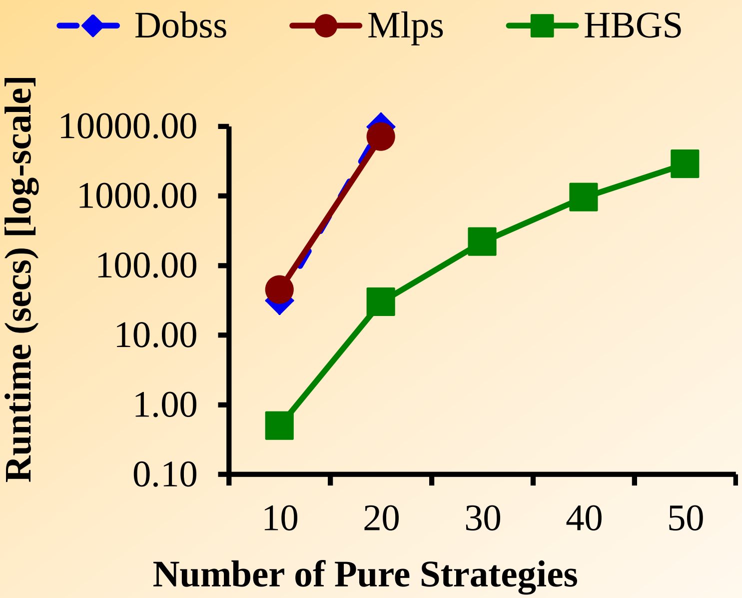

Types	Follower Pure Strategy Combinations	Runtime (secs)
10	9.7e7	0.41
20	9.5e13	16.33
30	9.3e20	239.97
40	9.1e27	577.49
50	8.9e34	3321.68

Approximation

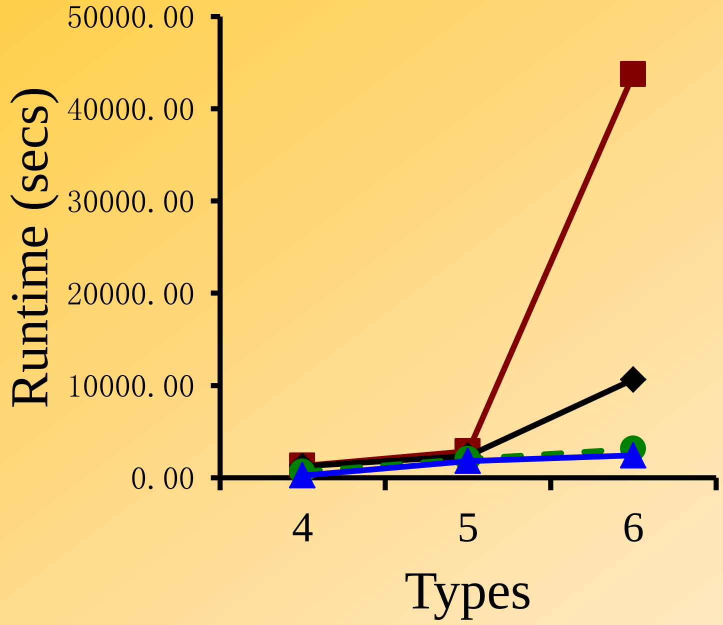

Approximation

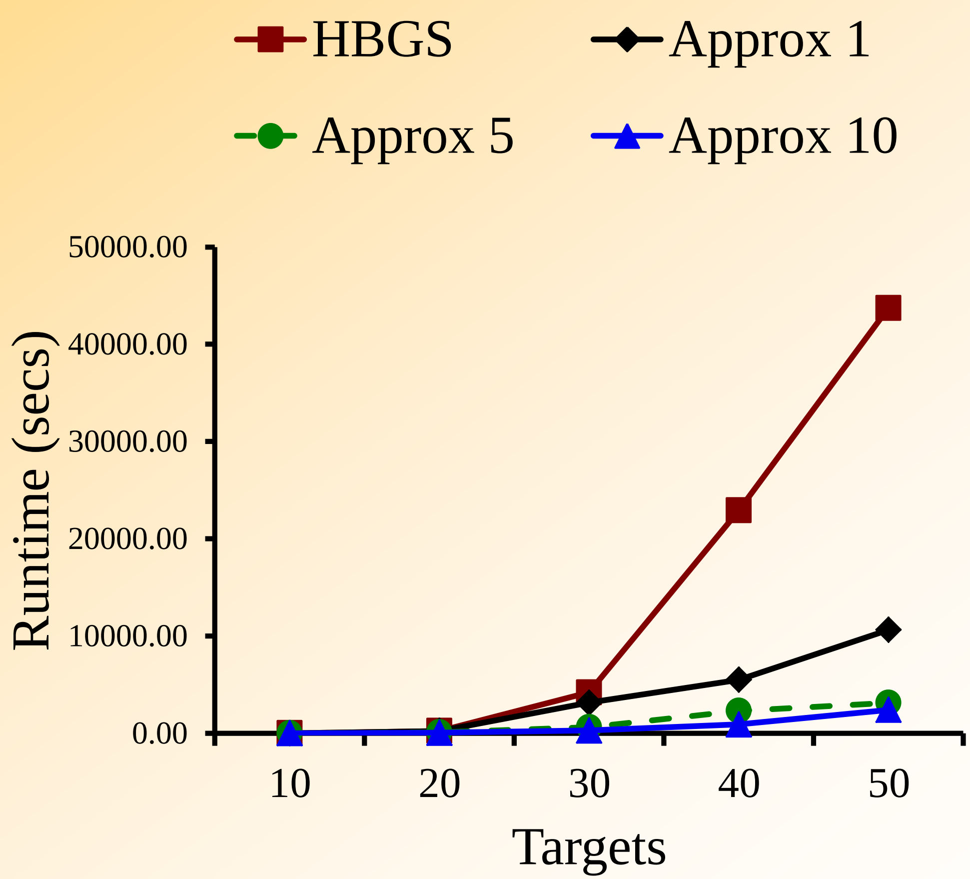

HUNTER

[Yin et al. 2012]

- Improves on tree search from HBGS
- Improved bounds (convex hulls on types)
- Bender's decomposition on LPs

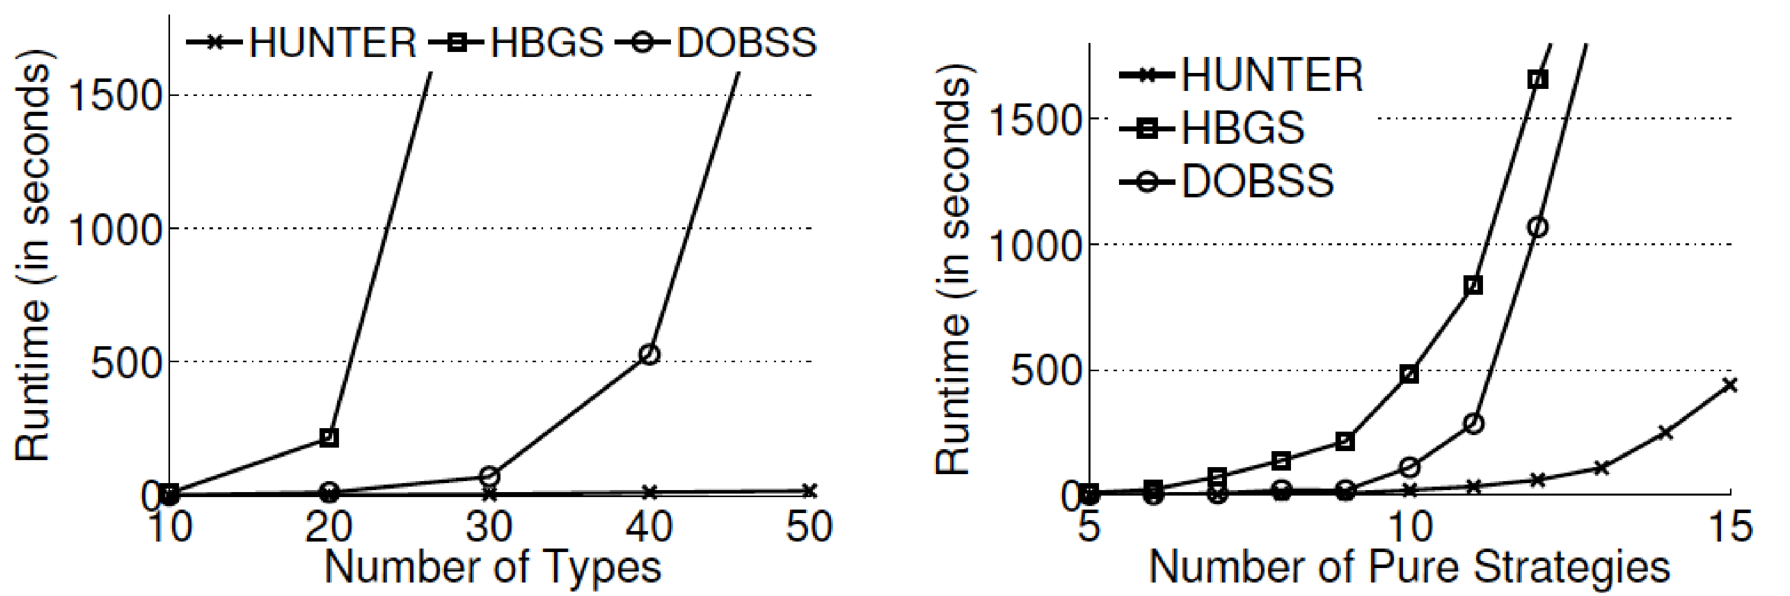

(a) Scaling up types. (b) Scaling up pure strategies.

Finite vs Infinite BSG

- Finite games capture distinct attacker types
  - *Terrorists vs. local criminal activity*
  - *Attackers with different motivations*

- Infinite games capture distributional uncertainty
  - *E.g., Gaussian, Uniform distributions*
  - *Natural for expressing beliefs over possible values*
  - *Useful in knowledge acquisition from experts*

Distributional Payoff Representation

[Kiekintveld et al. 2011]

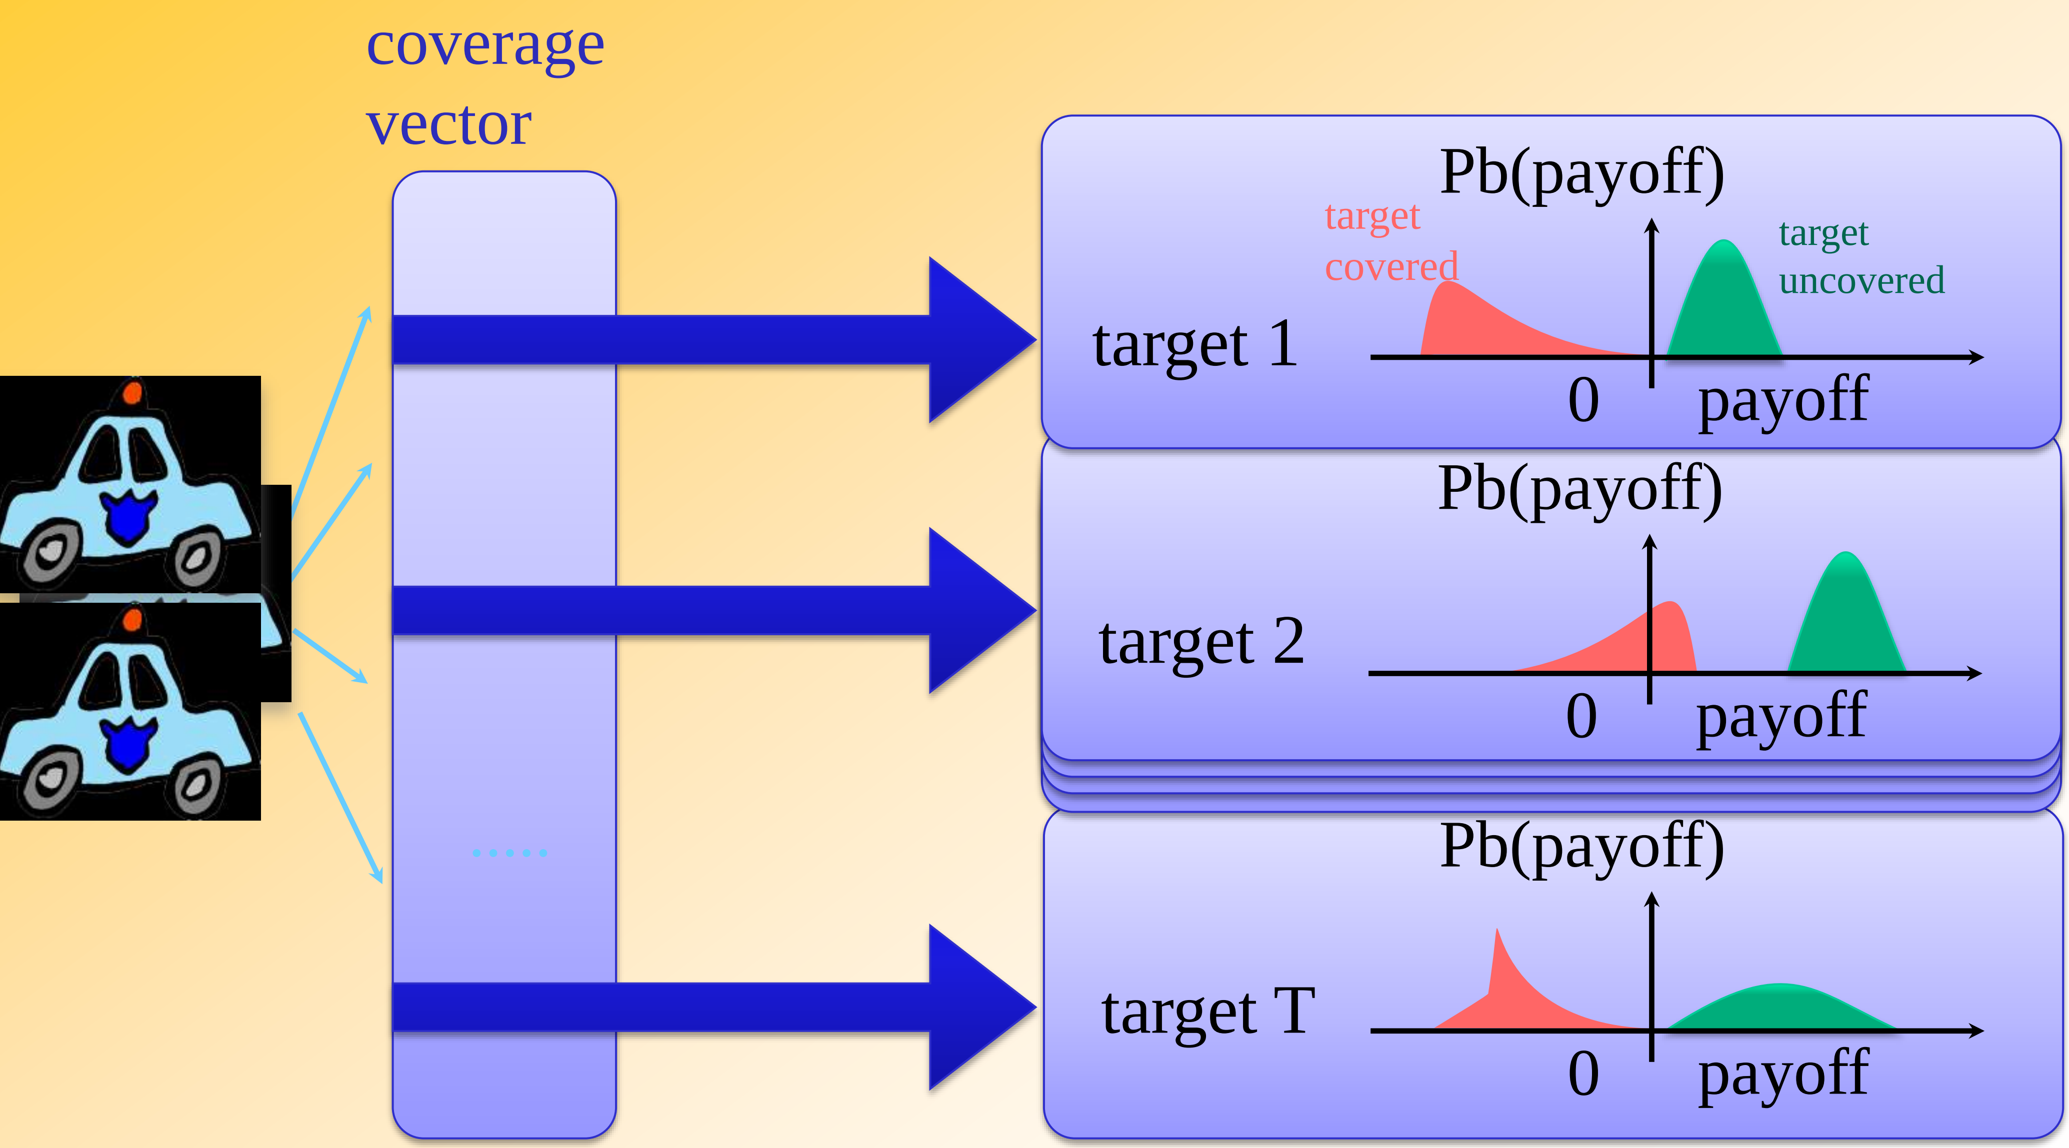

Problem 1 of 2

given a coverage vector C...

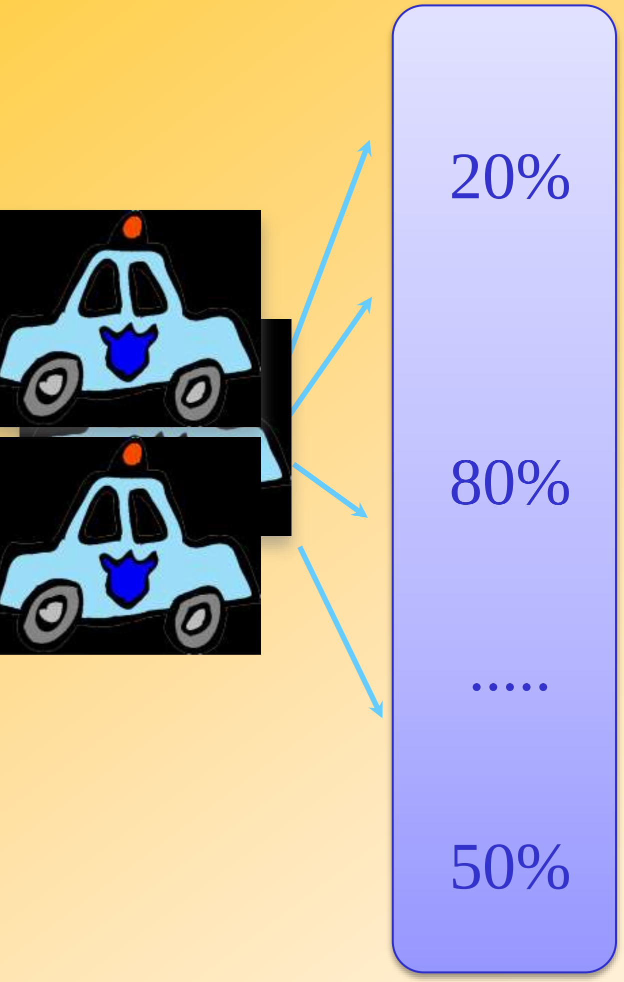

...and payoff distributions

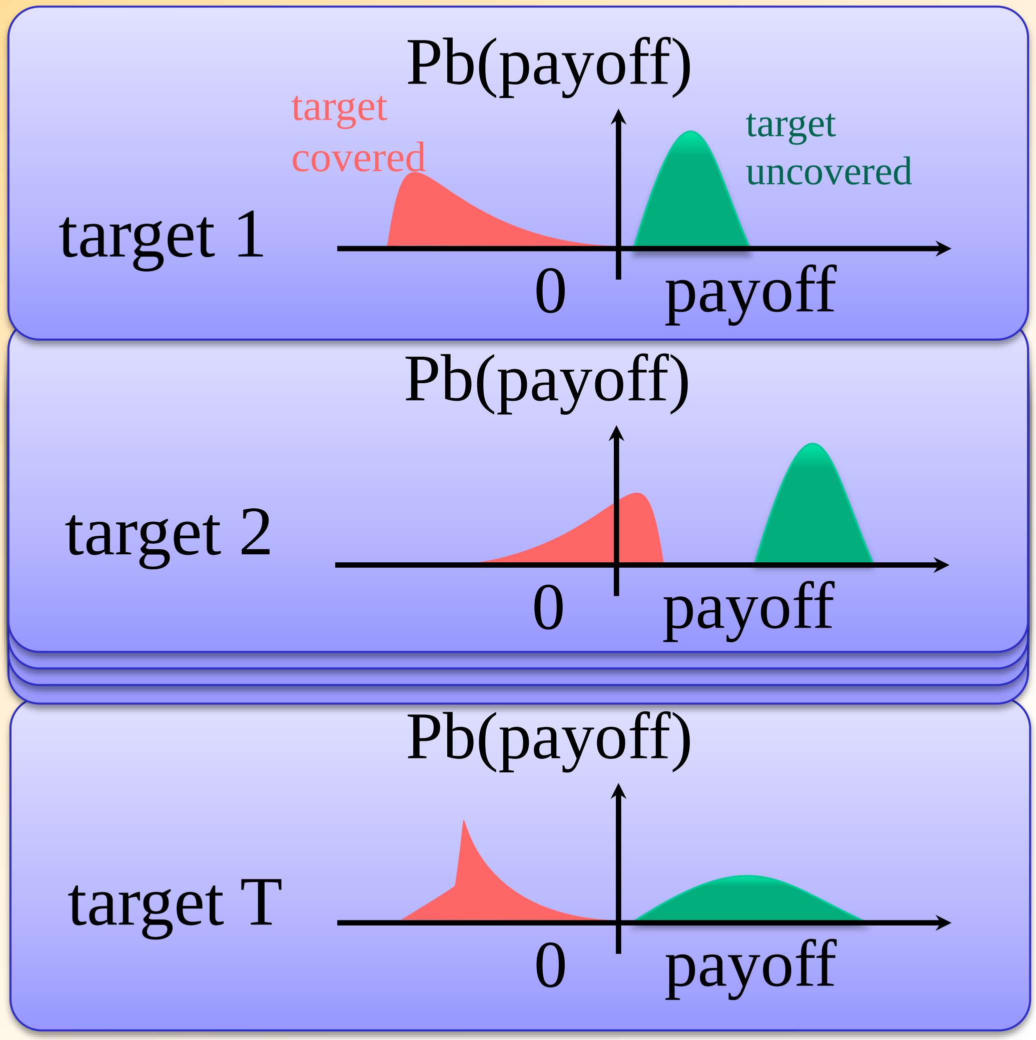

Problem 1 of 2

given a coverage vector C...

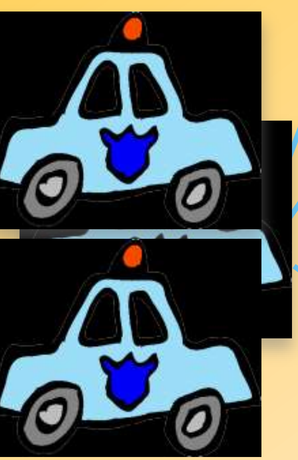

20%

80%

.....

50%

attack vector A(C)

?

...and payoff distributions

target 1

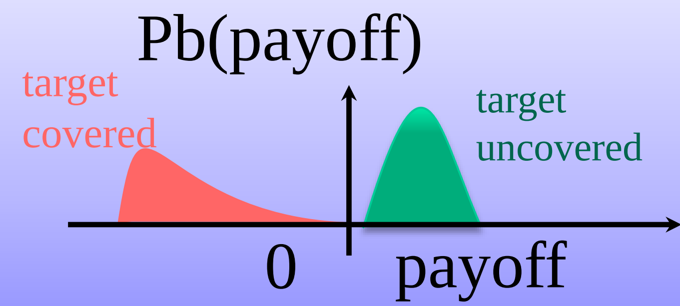

target 2

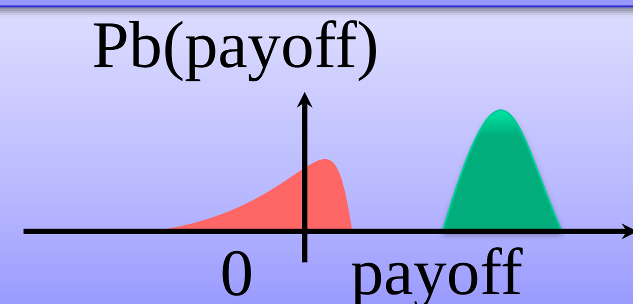

target T

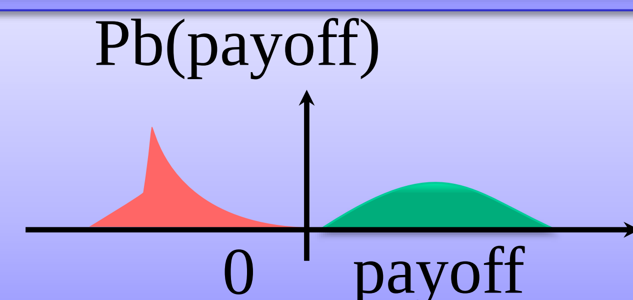

Problem 2 of 2

find the optimal coverage vector C*.

... given A(C) for every C

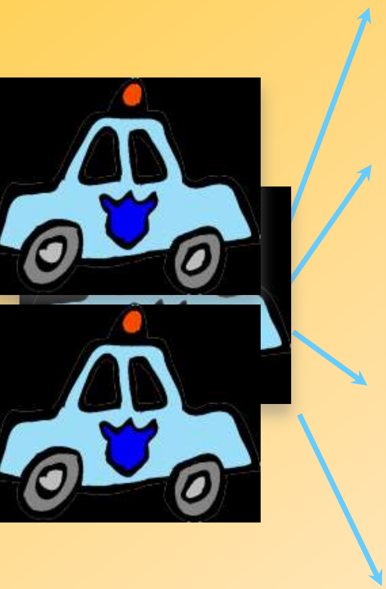

?

$a_1(C)$

$a_2(C)$

$a_3(C)$

...

$a_T(C)$

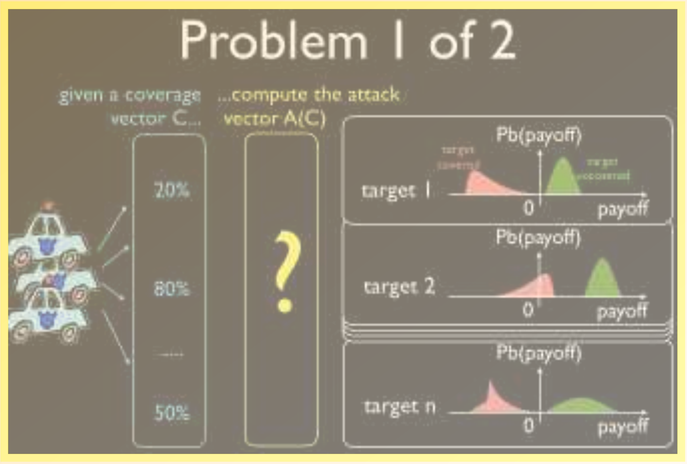

Approach

Coverage Vector

Attack Vector

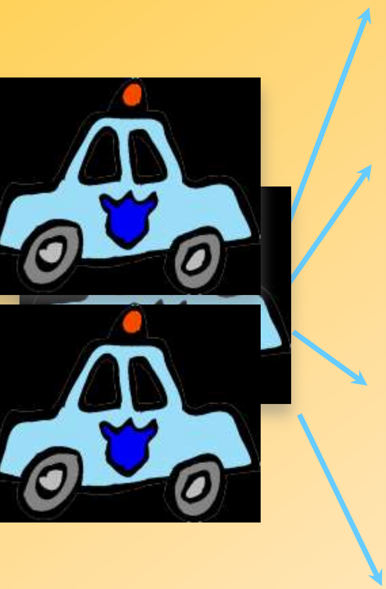

?

?

(1) Monte-Carlo estimation
(2) Numerical methods

(1) Optimal Finite Algorithms
(2) Sampled Replicator Dynamics
(3) Greedy Monte-Carlo
(4) Decoupled Target Sets

Attacker Response Estimation

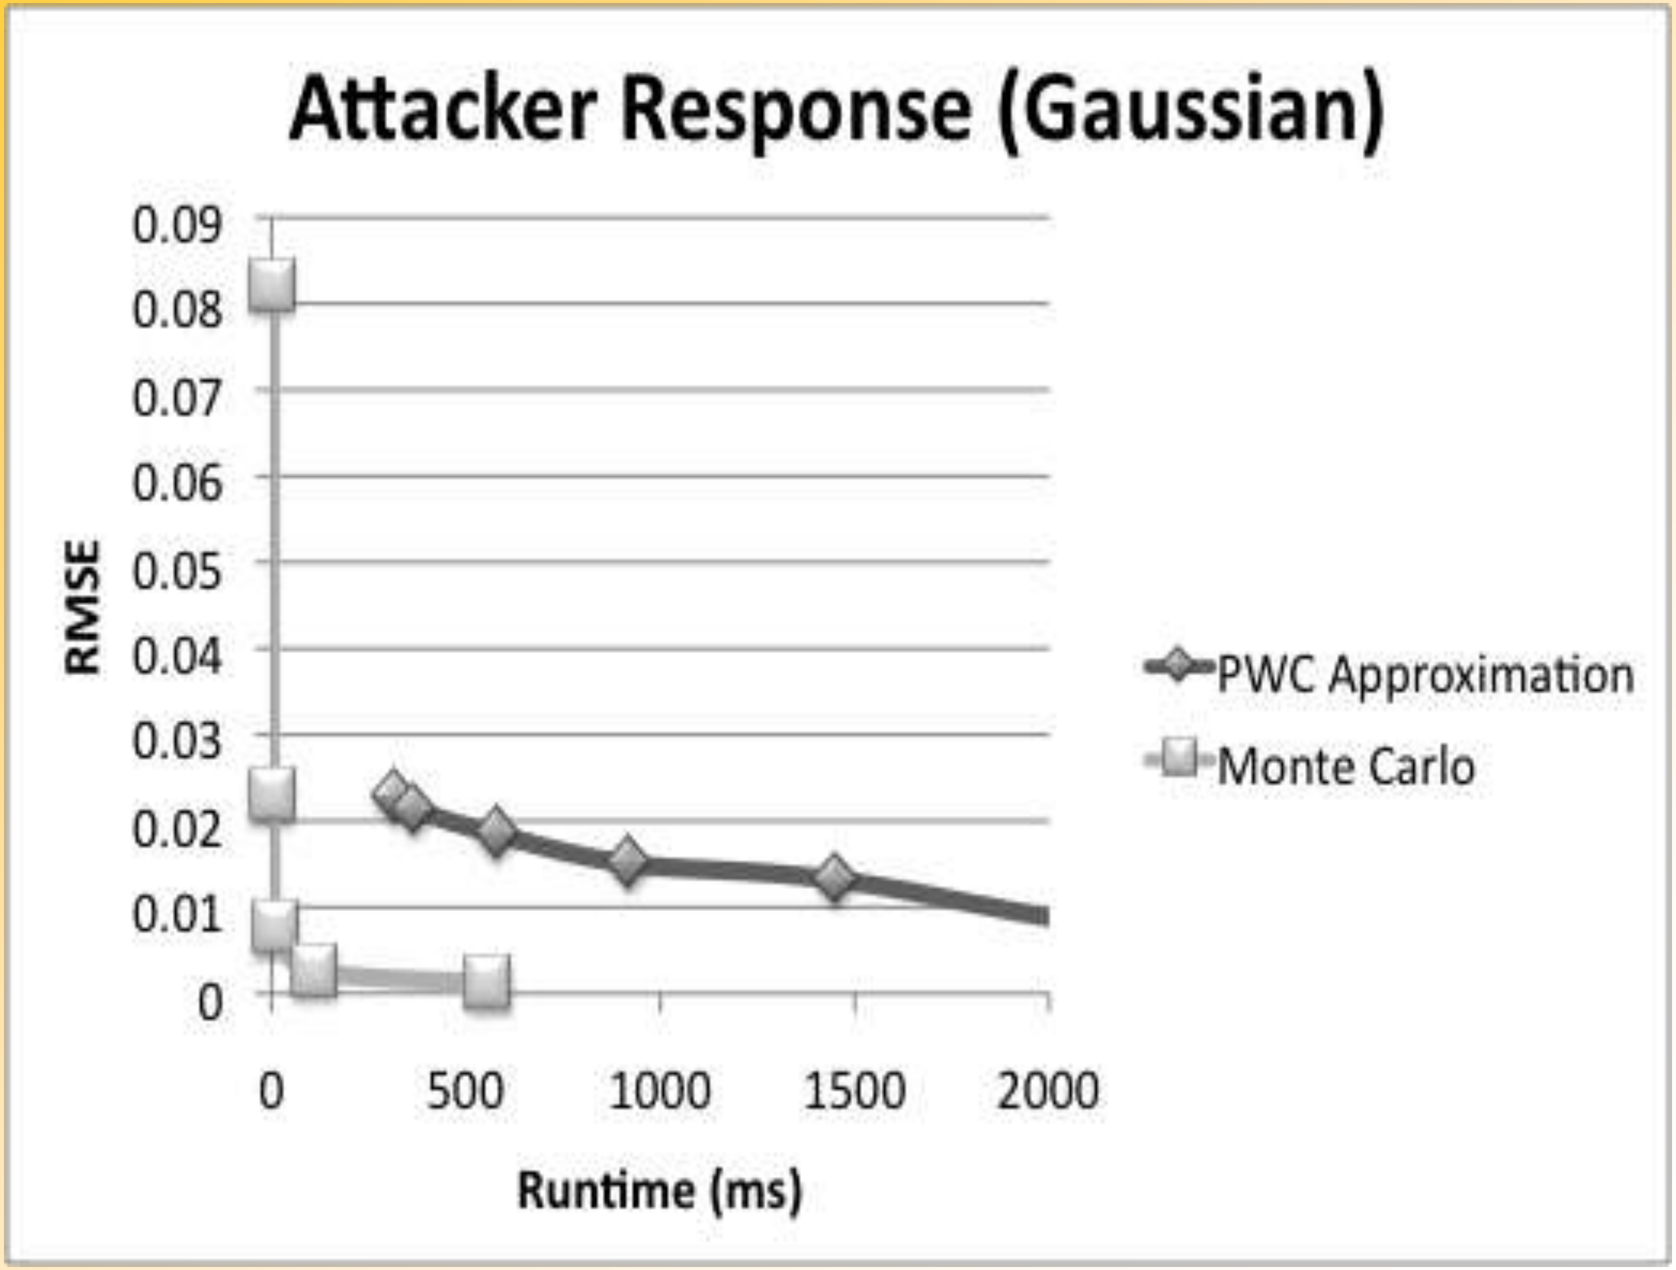

Computing Coverage Vectors

- Baselines
  - *Mean (ignore uncertainty)*
  - *Uniform Random*
- Exact optimization given sampled types
  - *SBE (ARMOR variation)*
- Worst-case optimization
  - *BRASS*
- Approximate optimization
  - *Replicator Dynamics (SRD)*
  - *Greedy Monte Carlo (GMC)*
  - *Decoupled Target Sets (DTS)*

Results for Distributional Games

Sample Types
Approx Optimization

Sample Types
Exact Optimization

Assume
Perfect Information

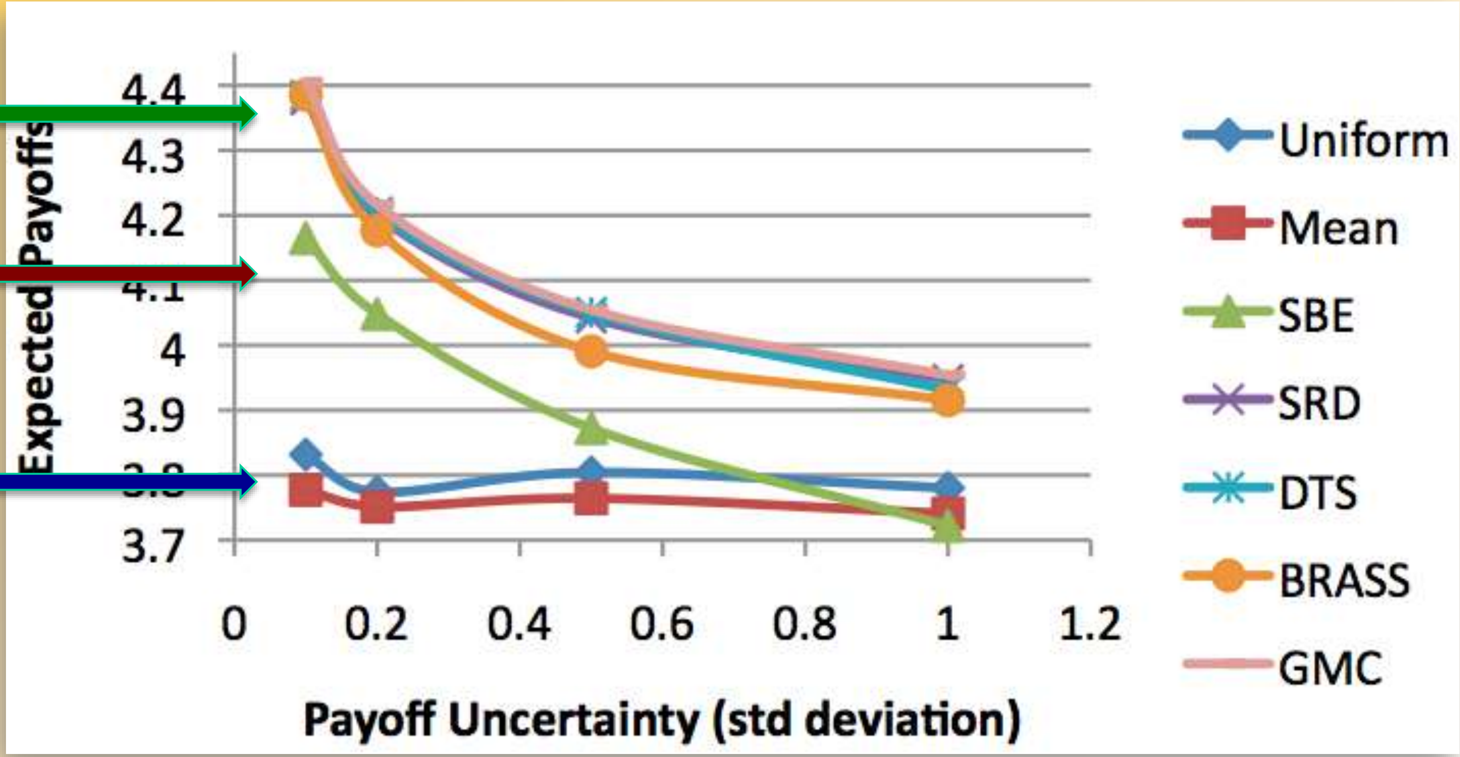

Assuming perfect information is very brittle

Approximate both type distribution and optimization

Beyond Bayesian Games

- Bayesian games are powerful
  - *General framework for model uncertainty*
  - *Exact behavior predictions based on uncertainty*

- Some limitations
  - *Require distributional information*
    - Even MORE parameters to specify!
    - What if these are wrong?
  - *Computational challenges (NP-hard)*
  - *Uncertainty about human decision making is hard to capture in Bayesian models*

Interval Security Games

[Kiekintveld et al. 2012]

	Target 1	Target 2	Target 3	Target 4
Defender Reward	0	0	0	0
Defender Penalty	-1	-4	-6	-10
Attacker Penalty	0	0	0	0
Attacker Reward	[1,3]	[2,5]	[4,7]	[6,10]

- Attacker payoffs represented by intervals
- Maximize worst case for defender
- Distribution-free

Polynomial Interval Solver

[Kiekintveld et al. 2012]

- Fast feasibility checks
  - Given resource constraint, can the defender guarantee a given payoff?
  - Exploits structure of security games
- Binary search on defender payoffs
- Polynomial time: $O(n^2 * \log(1/\varepsilon))$

Attacker Payoffs

0	0	0	0	0

Defender Coverage

5 Targets

Bars represent range of possible attacker payoffs

Attacker Payoffs

0	0	0	0.5	0.5

Defender Coverage

When targets are covered, payoffs decrease and range shrinks

Potential Attack Set

0	0	0	0.5	0.5

Defender Coverage

Given a coverage strategy, which set of targets *could* be attacked?

Minimum attacker payoff is R

Any target with a possible value greater than R is in the *potential attack set*

Polynomial Algorithm

- Main Idea:
  - *Design fast feasibility check to determine if a given defender payoff is possible*
  - *Use binary search on defender payoffs*
  - *Necessary resources increases monotonically with defender payoff*

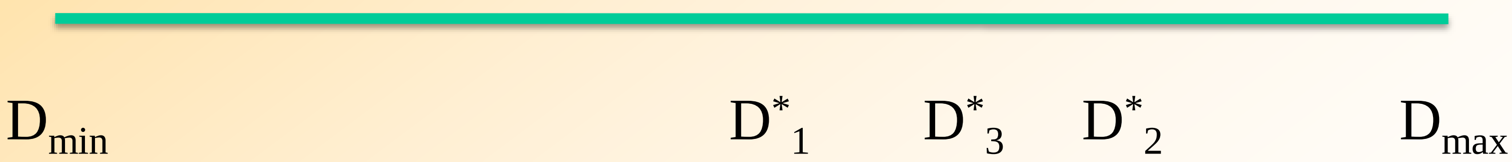

Feasibility Checks

Determine whether we can guarantee a defender payoff of D* using *m* or fewer resources

Challenge: potential attack set depends on coverage, and number of possible sets is combinatorial

Solution Idea

For any potential attack set, there is some target t' that determines the value of R

We will guess which target is t' and *construct* a minimal solution for this guess (n choices)

As soon as we find a choice of t' that works, we have a feasible solution

Constructing a Solution

0	0	0	0	0

Defender Coverage

Consider the selection
t' = t_2

Since t' is in the PAS,
it must give D* if attacked

Calculate minimal
coverage on t' using:

$$c_i^1 = \max(0, 1 - \frac{D^*}{U_\Theta^u(t_i)})$$

Constructing a Solution

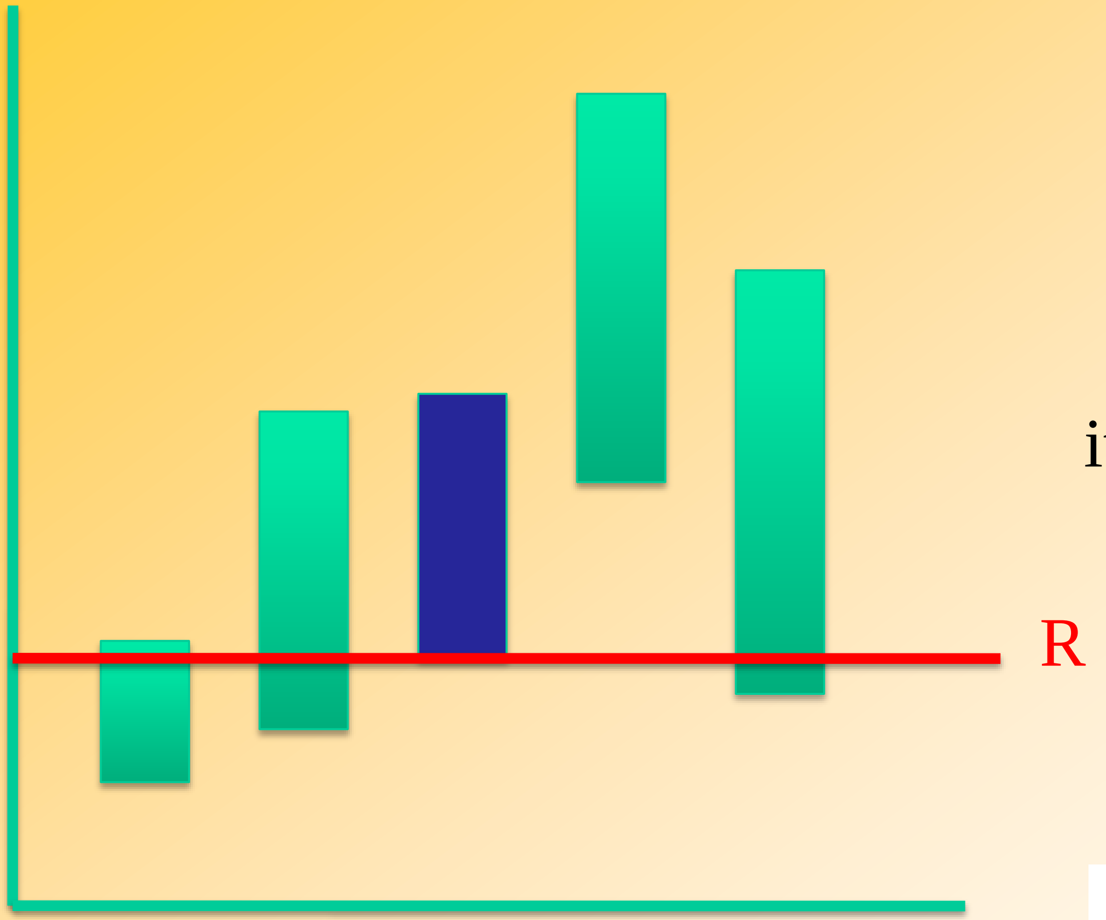

0	0	0.3	0	0

Defender Coverage

Consider the selection
t’ = t_2

Since t’ is in the PAS,
it must give D* if attacked

Calculate minimal coverage on t’ using:

$$c_i^1 = \max(0, 1 - \frac{D^*}{U_\Theta^u(t_i)})$$

Constructing a Solution

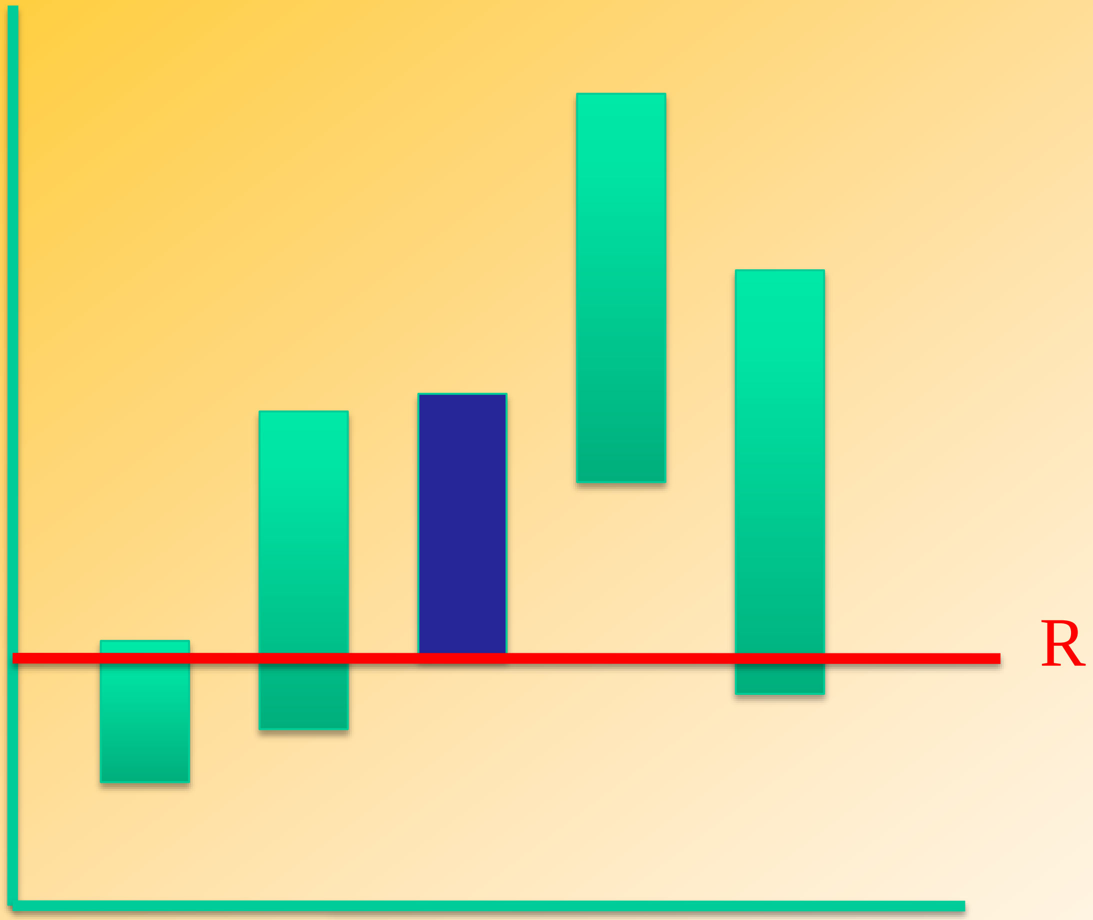

0	0	0.3	0	0

Defender Coverage

For every other target t", consider two cases:

1) Target is in the PAS
2) Target is not in the PAS

Constructing a Solution

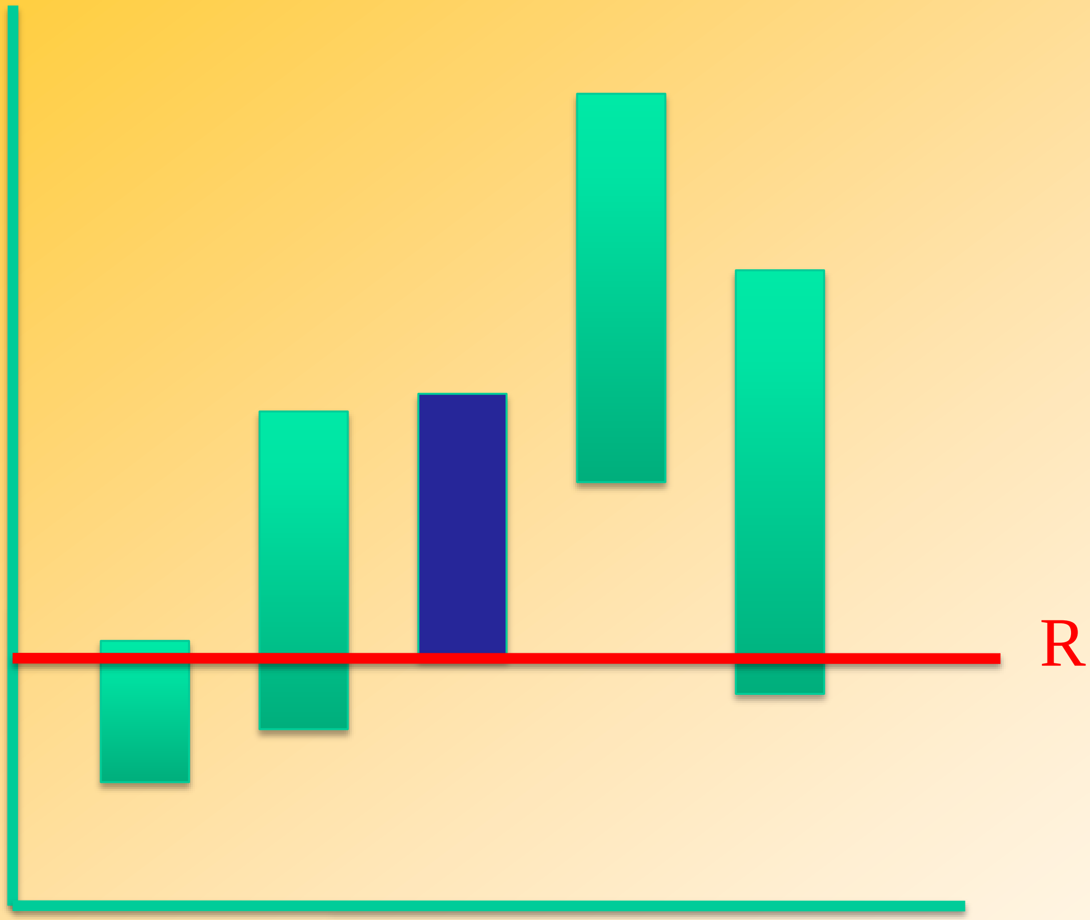

0	0	0.3	0	0

Defender Coverage

For every other target t'', consider two cases:

1) Target is in the PAS
2) Target is not in the PAS

<u>Case 1</u>

Payoff for t'' must be at least D*

$$c_i^1 = \max(0, 1 - \frac{D^*}{U_\Theta^u(t_i)})$$

Constructing a Solution

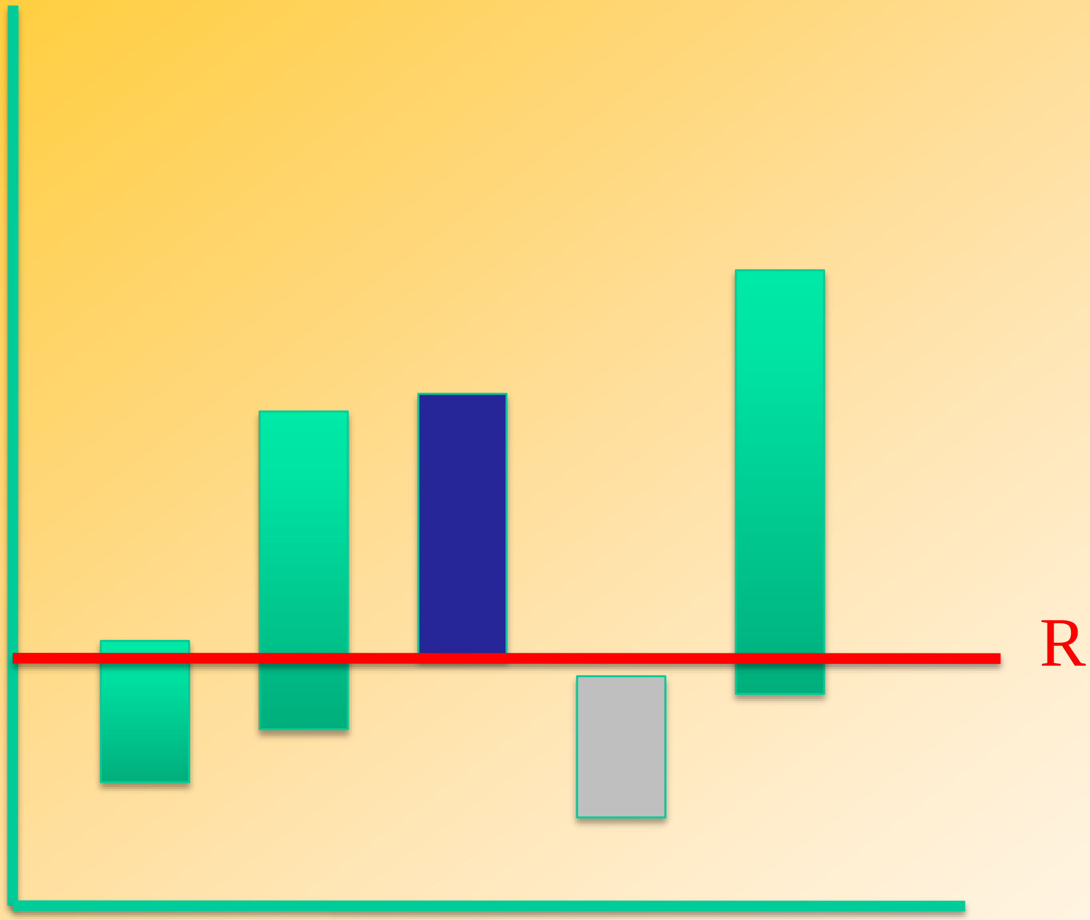

0	0	0.3	0.7	0

Defender Coverage

For every other target t'', consider two cases:

1) Target is in the PAS
2) Target is not in the PAS

<u>Case 2</u>

Max payoff to attacker for t'' must be < R

$$c_i^2 = \max(0, 1 - \frac{R}{U_\Psi^{u,max}(t_i)})$$

Constructing a Solution

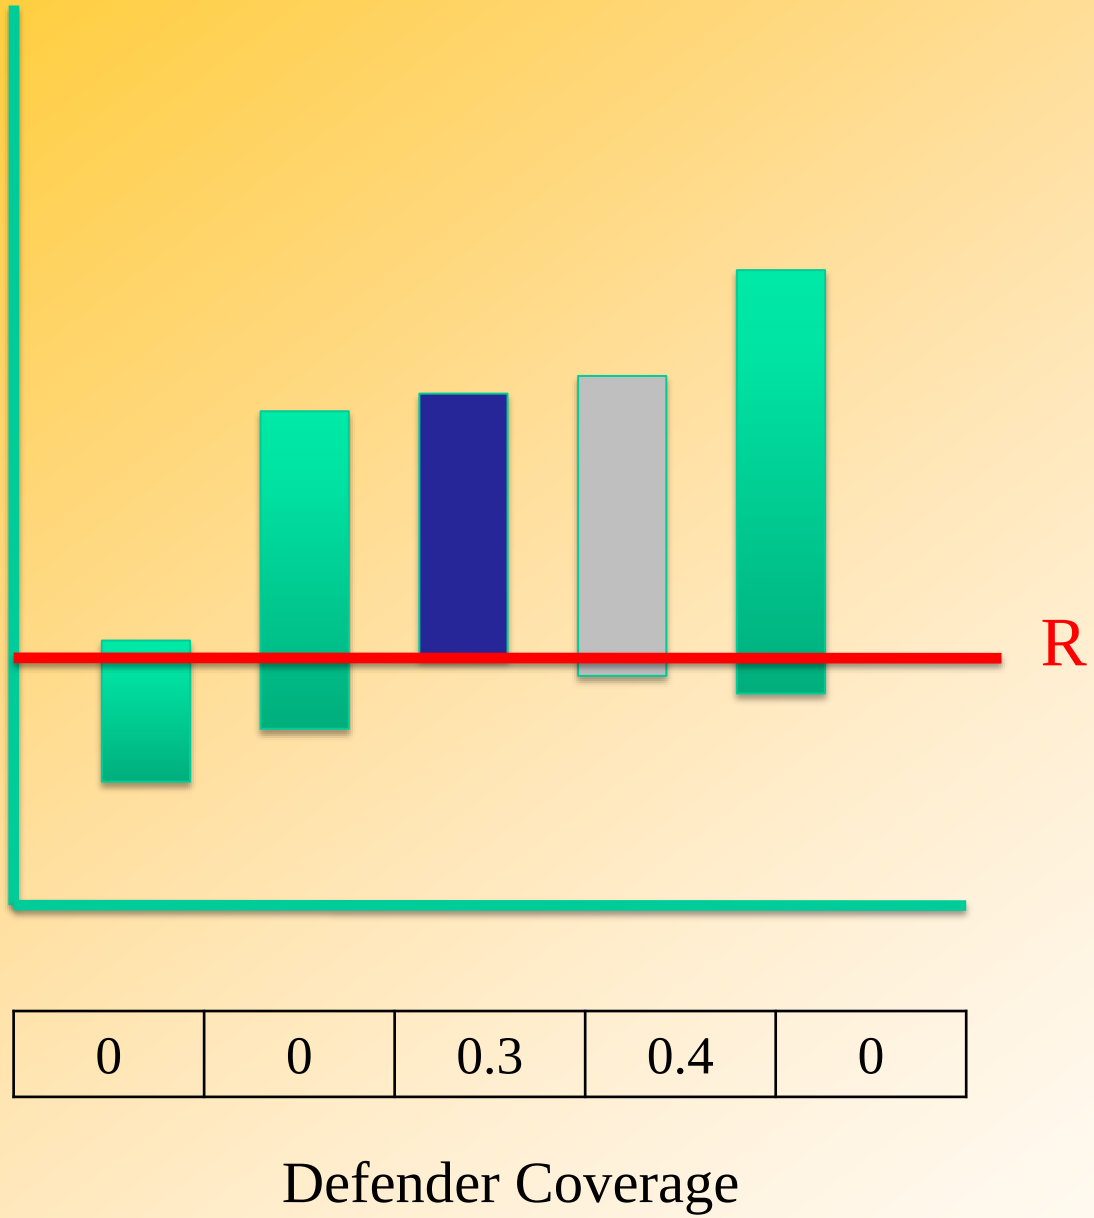

Final consistency check

No target other than t’ can have a higher minimum attacker payoff

Otherwise, t’ does not set R contradicting the initial assumption

$$c_i^3 = \max(0, 1 - \frac{R}{U_{\Psi}^{u,min}(t_i)})$$

Constructing a Solution

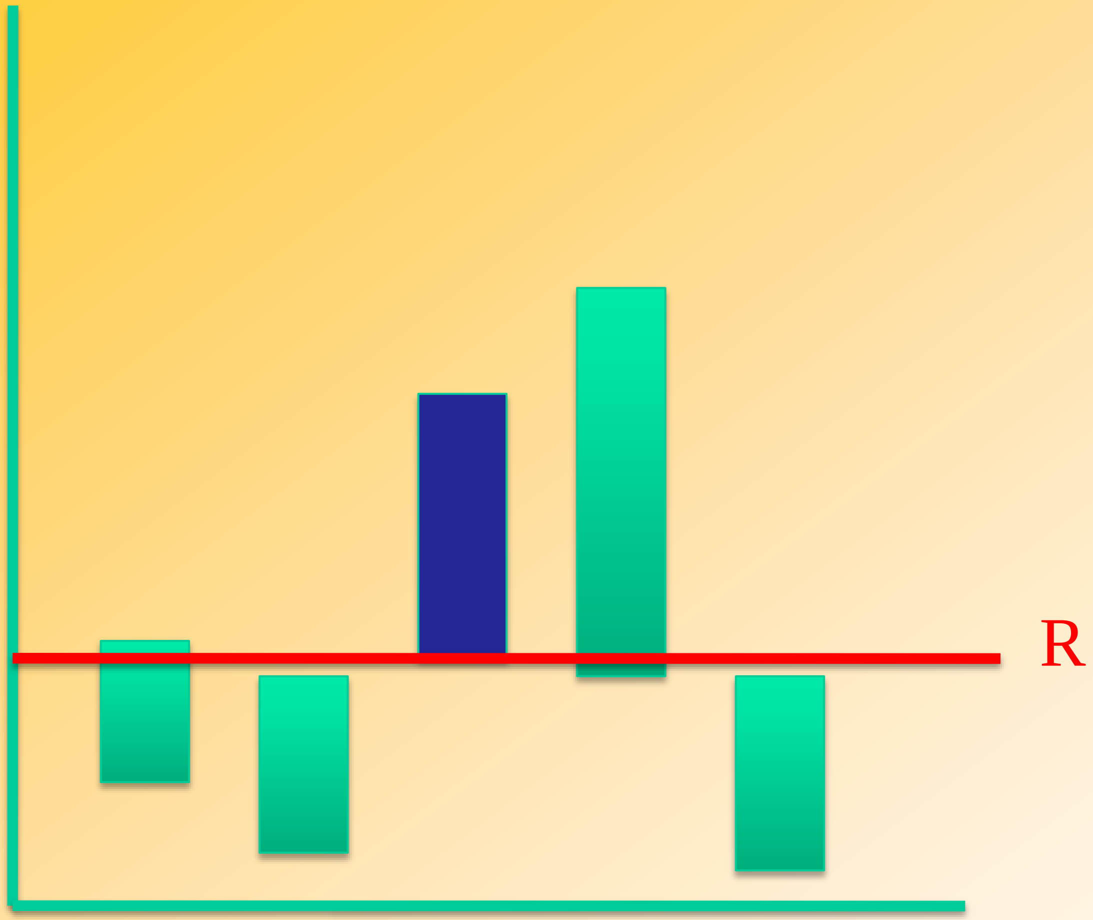

0.2	0.5	0.3	0.4	0.7

Defender Coverage

For each target, compute three coverage values

c^1: coverage for D*
c^2: coverage not in PAS
c^3: consistency with R

Best value given by:

$$\max(c_i^3, \min(c_i^1, c_i^2))$$

Analysis

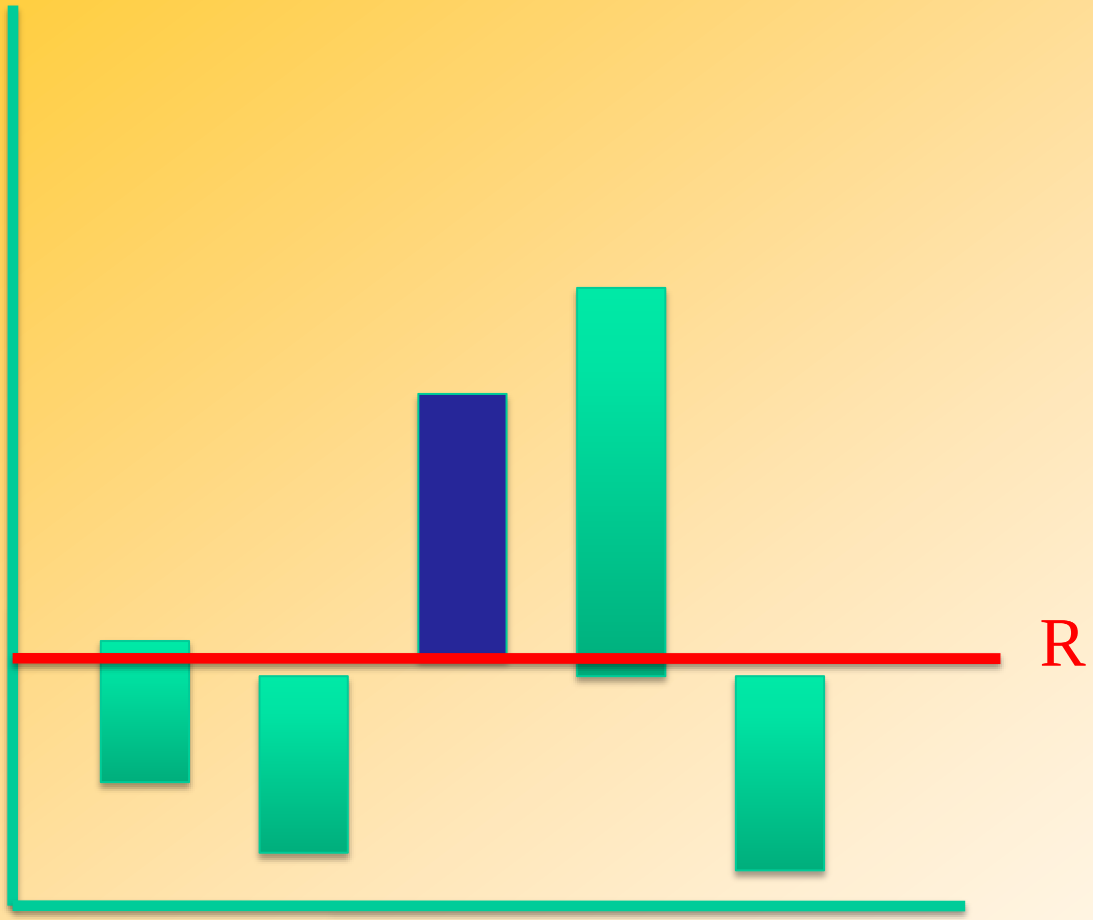

0.2	0.5	0.3	0.4	0.7

Defender Coverage

Need to check each target as t'

$O(n^2)$ worst case to test feasibility for D*

Binary search on D

$O(n^2 * \log(1/\varepsilon))$
where ε is error term

Interval Solver Scalability

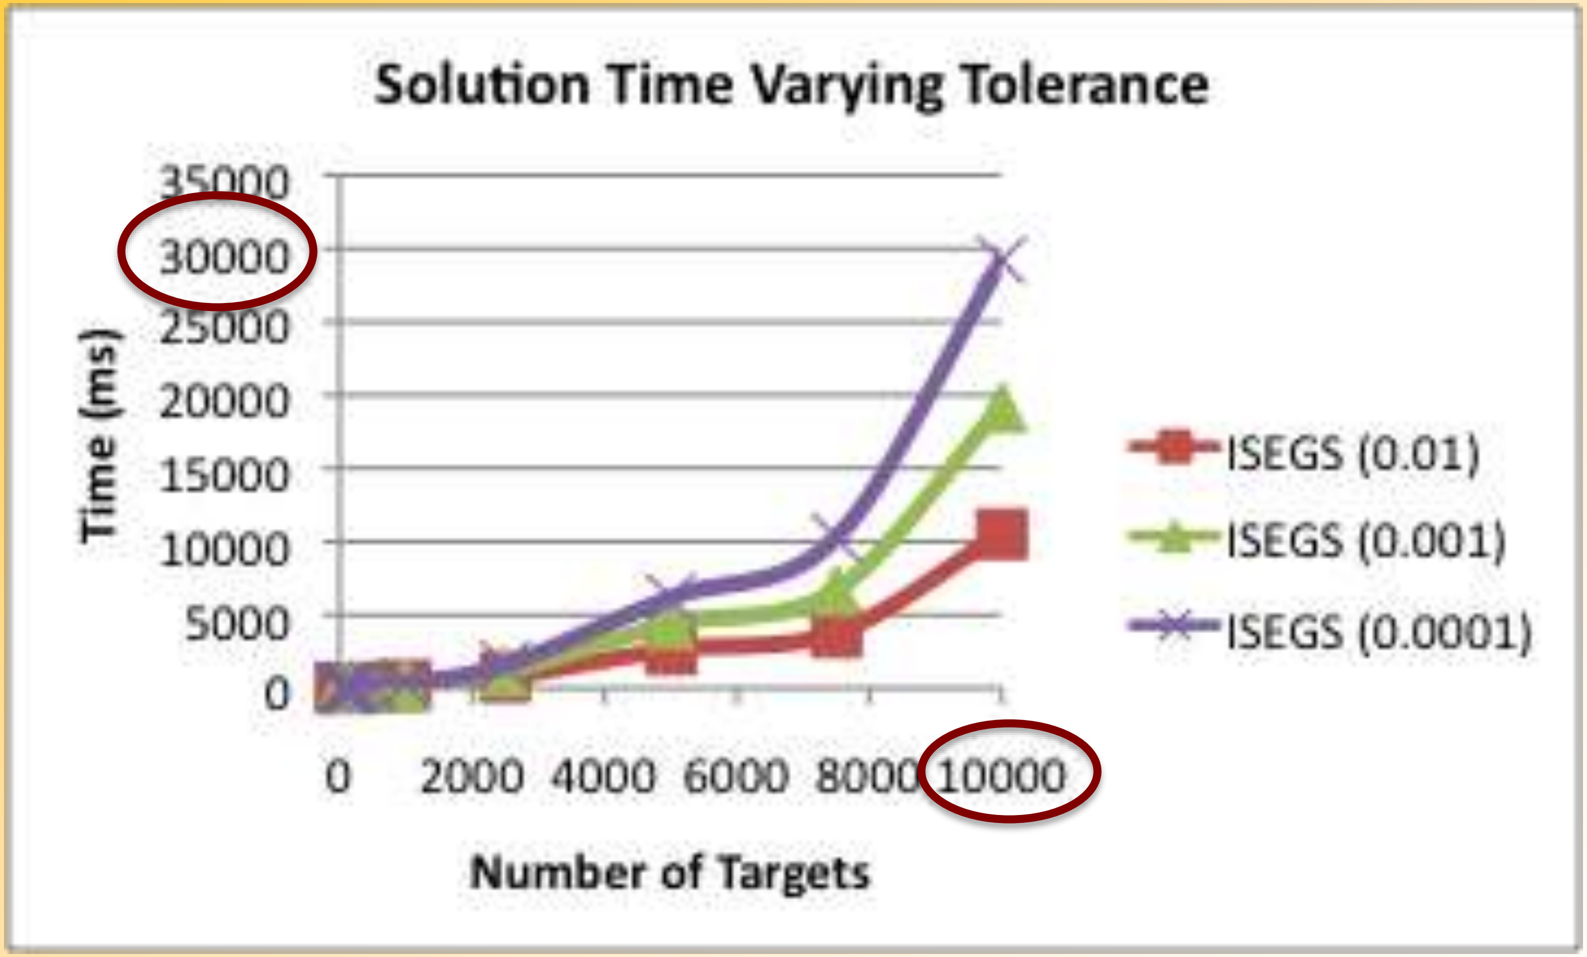

Fastest Bayesian solvers (HBGS, HUNTER)
scale only to 10s or 100s of targets

Outline

- Motivating real-world applications
- Background and basic security games
- Scaling to complex action spaces
- Modeling payoff uncertainty: Bayesian Security Games
- ***Human behavior and observation uncertainty***
- Evaluation and discussion

Key Topics

- PART I: Integrate models of human decision making as attacker's response
  - *Key model used:*
    - Anchoring bias and epsilon-bounded rationality
    - Prospect Theory [Kahneman and Tvesky, 1979]
    - Quantal Response [McKelvey and Palfrey, 1995]
  - *New efficient algorithms*
  - *Results from experiments with human subjects*
    - Quantal Response (QRE) outperforms other algorithms

- PART II: Impact of limited observations assuming rational attacker

Uncertainty: Attacker Decision Bounded Rationality & Observations: Experimental Setup

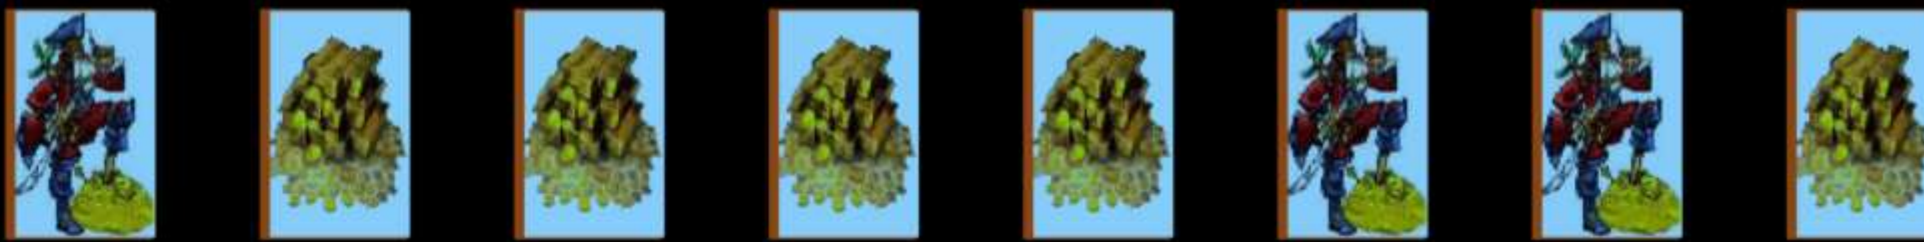

	1	2	3	4	5	6	7	8
Your Rewards:	8	5	3	10	1	3	9	4
Your Penalties:	−3	−2	−3	−2	−3	−3	−2	−3
Pirate's Rewards:	4	3	1	5	1	2	5	2
Pirate's Penalties:	−8	−10	−1	−8	−1	−3	−11	−5

Uncertainty: Human Bounded Rationality and Observations

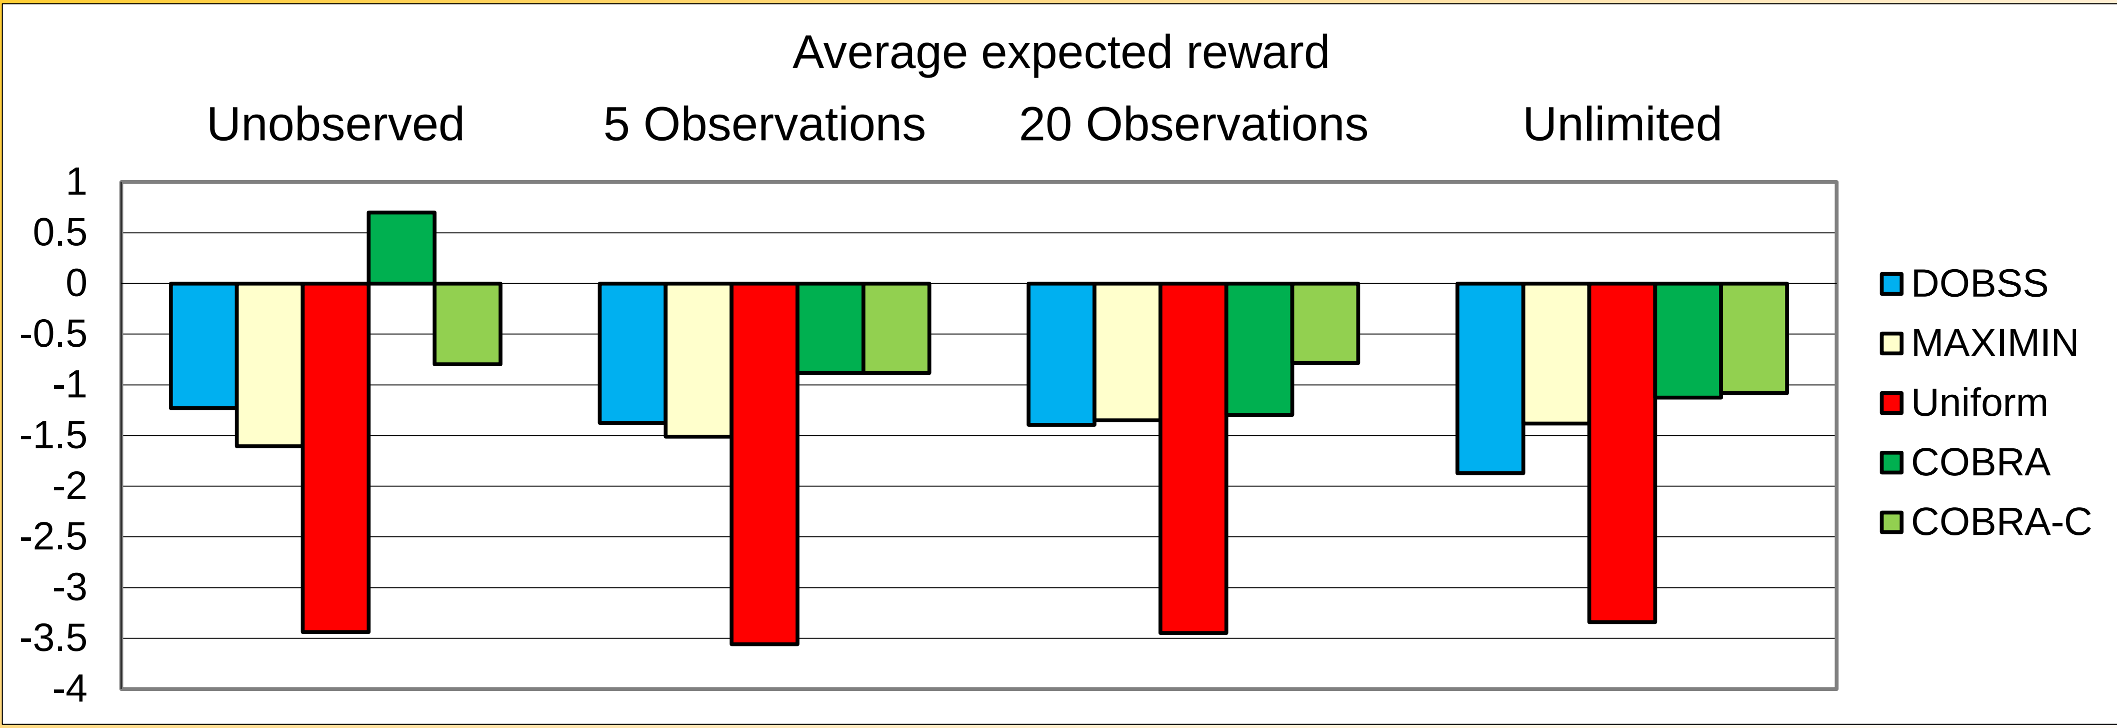

- *178 total subjects, 2480 trials, 40 subjects for each setting*
- *Four reward structures, four observation conditions*

- **DOBSS:** Outperforms uniform random, similar to Maximin

Uncertainty: Human Bounded Rationality and Observations

- **COBRA:**
  - “epsilon optimality”
  - Anchoring bias: Full observation vs no observation: **α**

Choosing **α** (even for unlimited observations?)

No observation: **α = 1**

Full observation: **α = 0**

$$\max_{x,q} \sum_{i\in X} \sum_{l\in L} \sum_{j\in Q} p^l R^l_{ij} x_i q^l_j$$

$$s.t.\quad x' = (1-\alpha)x + \alpha(1/|X|)$$

$$\varepsilon(1-q^l_j) \le (a^l - \sum_{i\in X} C^l_{ij} x'_i) \le \varepsilon + (1-q^l_j)M$$

Unlimited Observations: Choosing α

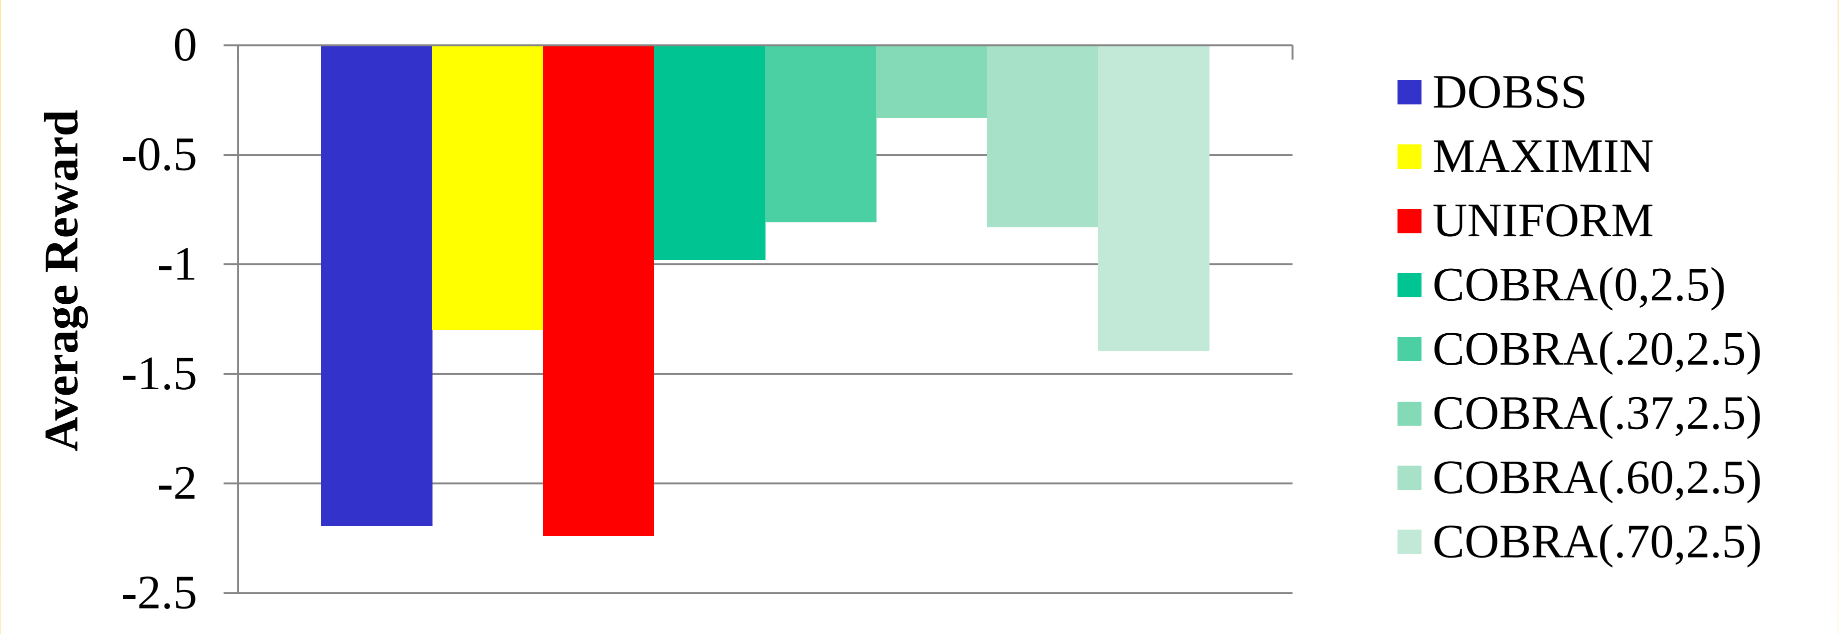

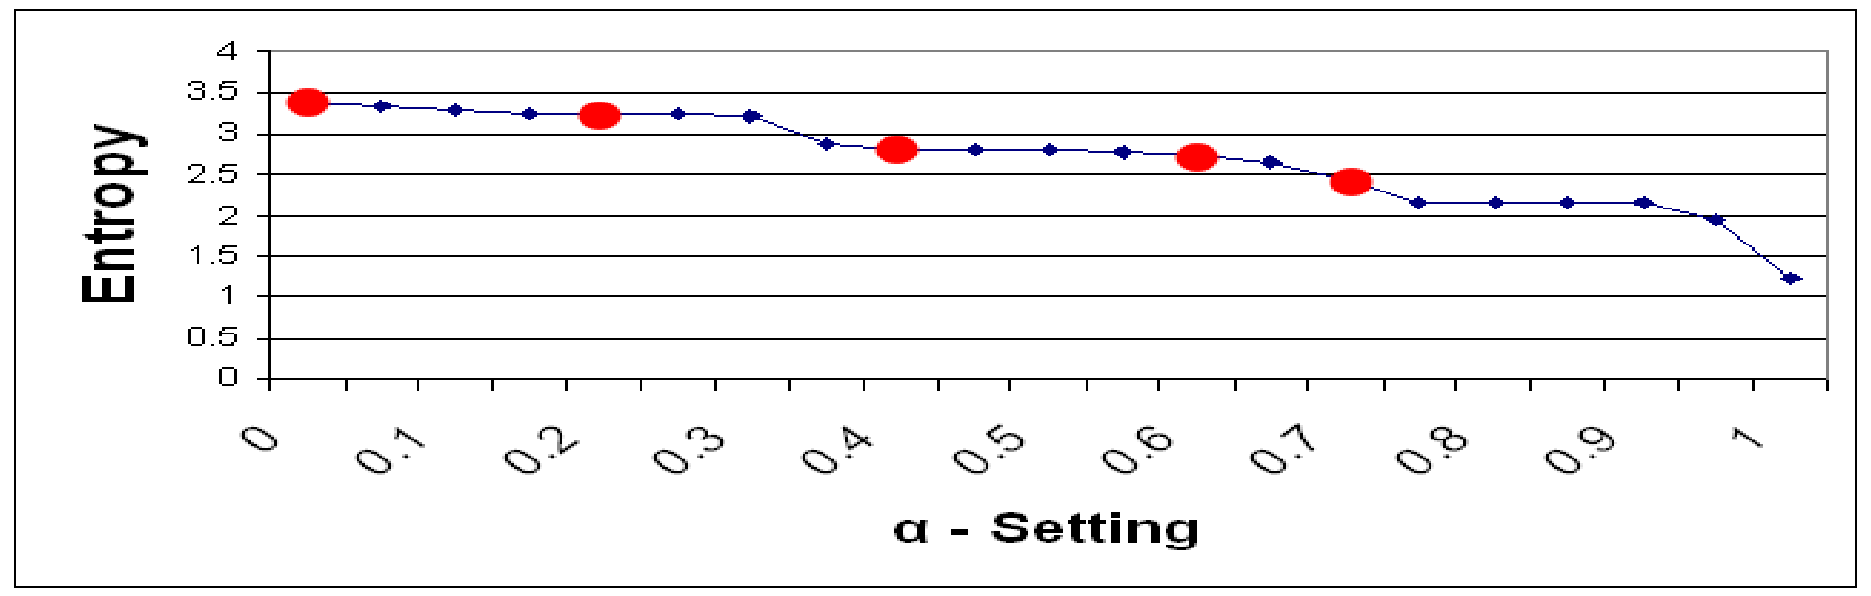

Prospect Theory

- Model human decision making under uncertainty
- Maximize the 'prospect' [Kahneman and Tvesky, 1979]

$$\text{prospect} = \sum_{i \in AllOutcomes} \pi(x_i) \cdot V(C_i)$$

  - $\pi(\cdot)$: *weighting function*
  - $V(\cdot)$: *value function*

Empirical weighting function

- Slope gets steeper as x gets closer to 0 and 1
- Not consistent with probability definition
  - $\pi(x)+\pi(1-x) < 1$
- Empirical value:

γ=0.64 (0<γ<1)

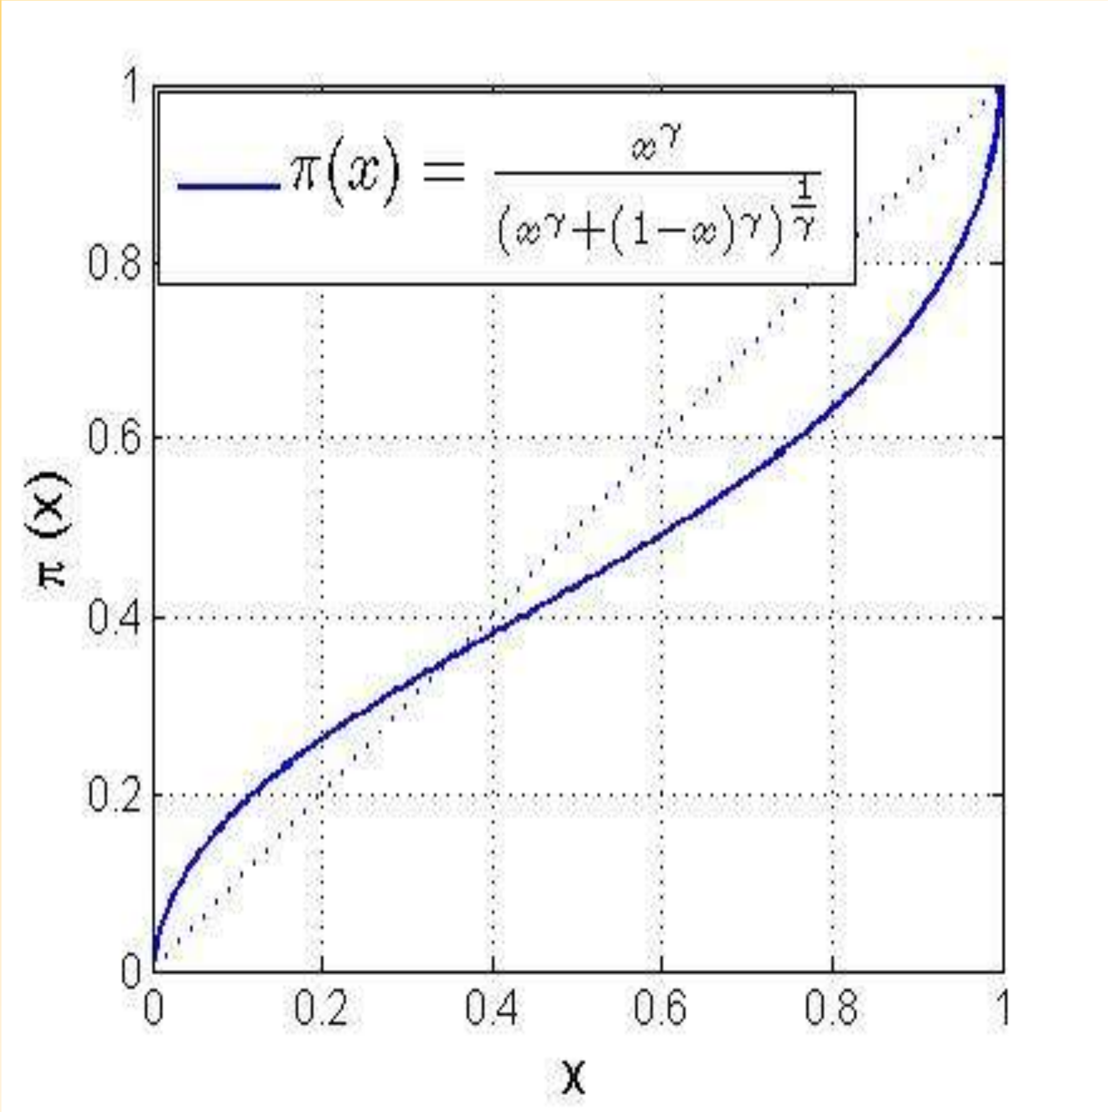

Compute Defender Strategy

- Piecewise Linear Approximation

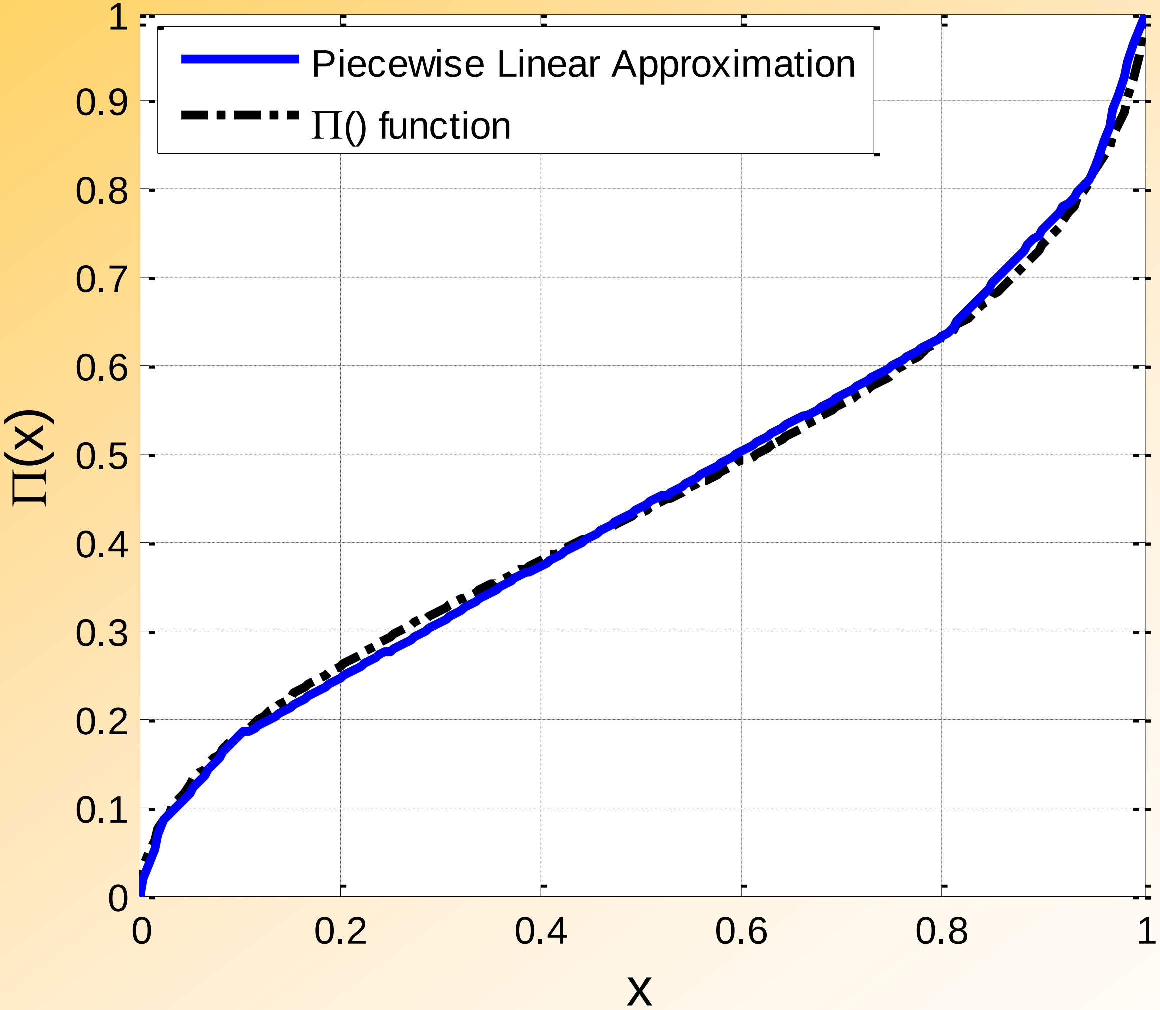

Empirical value function

- Risk averse regarding gain
- Risk seeking regarding loss
- Empirical value:

  α=β=0.88, λ=2.25

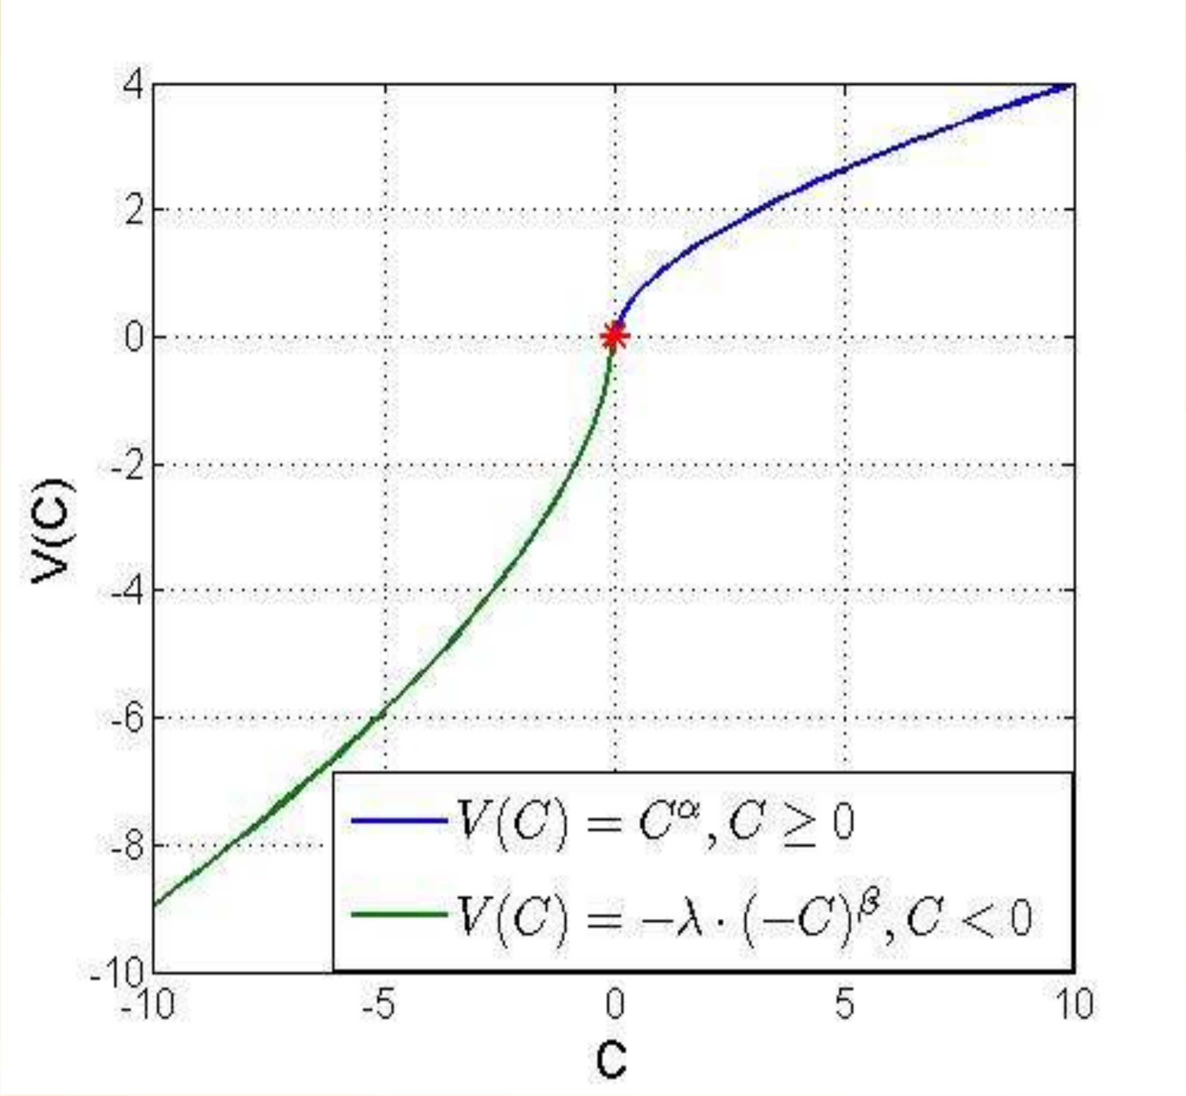

BRPT: Best Response to PT

- Mixed-Integer Linear Program
- Goal: maximize defender expected utility

$$\max_{x} \quad \text{DefenderUtility}$$

$$s.t \quad \sum_{i \in X} x_i \leq \text{Total\_Resources} \qquad (1)$$

Weighting Function

$$\pi(x_i) = \sum_{k=1..5} b_k \cdot x_{ik} \qquad (2)$$

$$\sum_{j \in Q} q_j = 1 \qquad (3)$$

Maximize prospect

$$0 \leq \text{AdversaryProspect} - \sum_{i \in X} \pi(x_i) \cdot V(C_{ij}) \leq M \cdot (1 - q_j), \forall j \in Q \qquad (4)$$

$$\text{DefenderUtility} - \sum_{i \in X} x_i \cdot R_{ij} \leq M \cdot (1 - q_j) \qquad (5)$$

Quantal Response Equilibrium

- Error in individual's response
  - *Still: more likely to select better choices than worse choices*
- Probability distribution of different responses
- Quantal best response:

$$q_j = \frac{e^{\lambda \cdot U(j,x)}}{\sum_{k=1}^{M} e^{\lambda \cdot U(k,x)}}$$

- λ: represents error level (=0 means uniform random)
  - *Maximal likelihood estimation (λ=0.76)*

Optimal Strategy against QR

- Solve the Nonlinear optimization problem

$$\max_{x} \quad \frac{\sum_{j\in Q}\sum_{i\in X} x_i R_{ij} \cdot \prod_{l\in X} e^{\lambda C_{lj} x_l}}{\sum_{k\in Q}\prod_{l\in X} e^{\lambda C_{lk} x_l}}$$

$$s.t. \qquad \sum_{\mathrm{i}\in \mathrm{X}} x_i \le \mathrm{Total\_Resource}$$

$$0 \le x_i \le 1, \qquad \forall i \in X$$

The Online Game

- Subjects are given $8 as the starting budget
- For each point they gain, $0.1 real money is paid

Gates	Gate 1	Gate 2	Gate 3	Gate 4	Gate 5	Gate 6	Gate 7	Gate 8
Your Rewards	1	9	5	6	7	1	10	3
Your Penalties	-2	-4	-3	-3	-3	-2	-4	-3
Probability of No Guard	1.00	0.41	0.54	0.49	0.44	1.00	0.39	0.73
Probability of Guard	0.00	0.59	0.46	0.51	0.56	0.00	0.61	0.27
Guards' Rewards	1	4	2	3	4	1	5	2
Guards' Penalties	-5	-8	-1	-6	-5	-1	-7	-7

Please choose a door to attack. Press the Submit Button to confirm your selection.

Submit

Experiment Setting

- 7 payoff structures
  - *4 new, 3 from previous tests with COBRA*
- 5 strategies for each payoff structure
  - *New methods: BRPT, RPT and BRQR*
  - *Leading contender: COBRA*
  - *Perfect rational baseline: DOBSS*
- Subjects play all games (randomized orders)
- No feedback until subject finishes all games

Average Defender Expected Utility

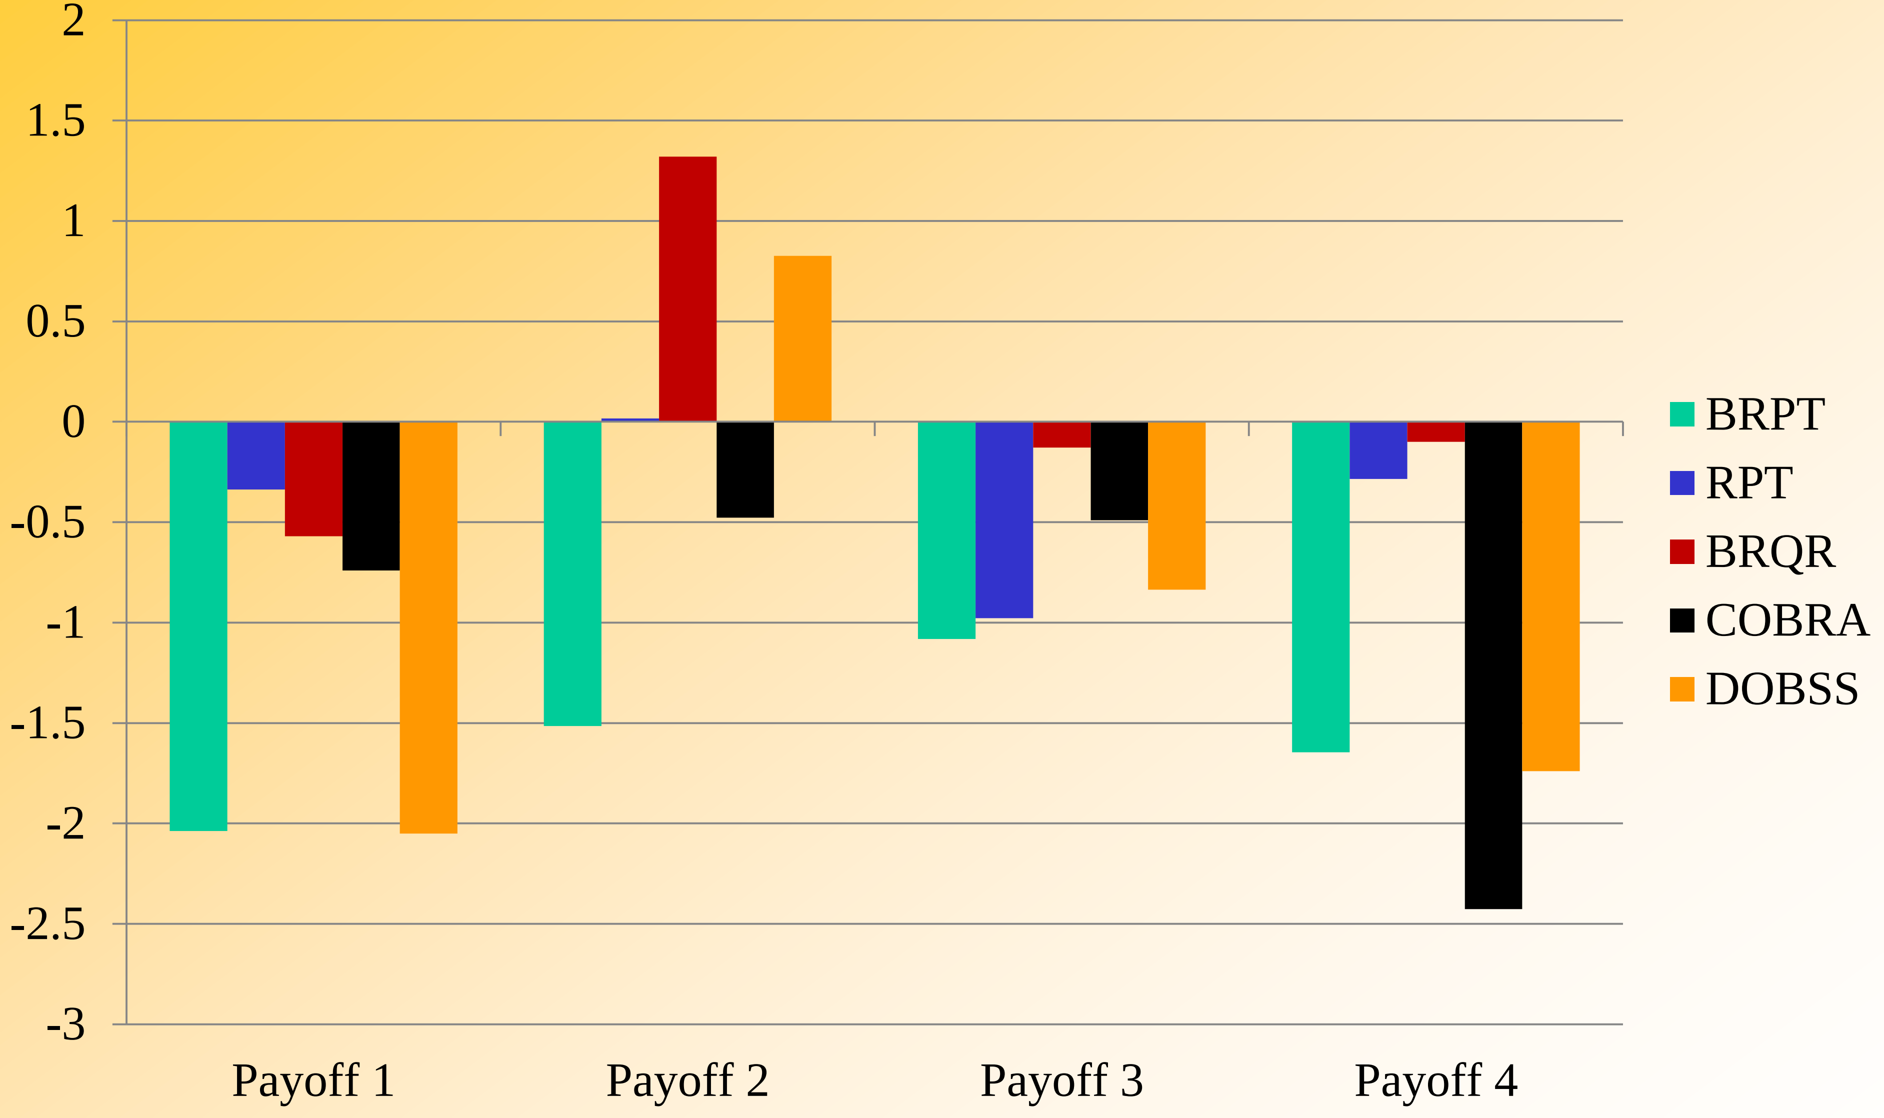

Average Defender Expected Utility

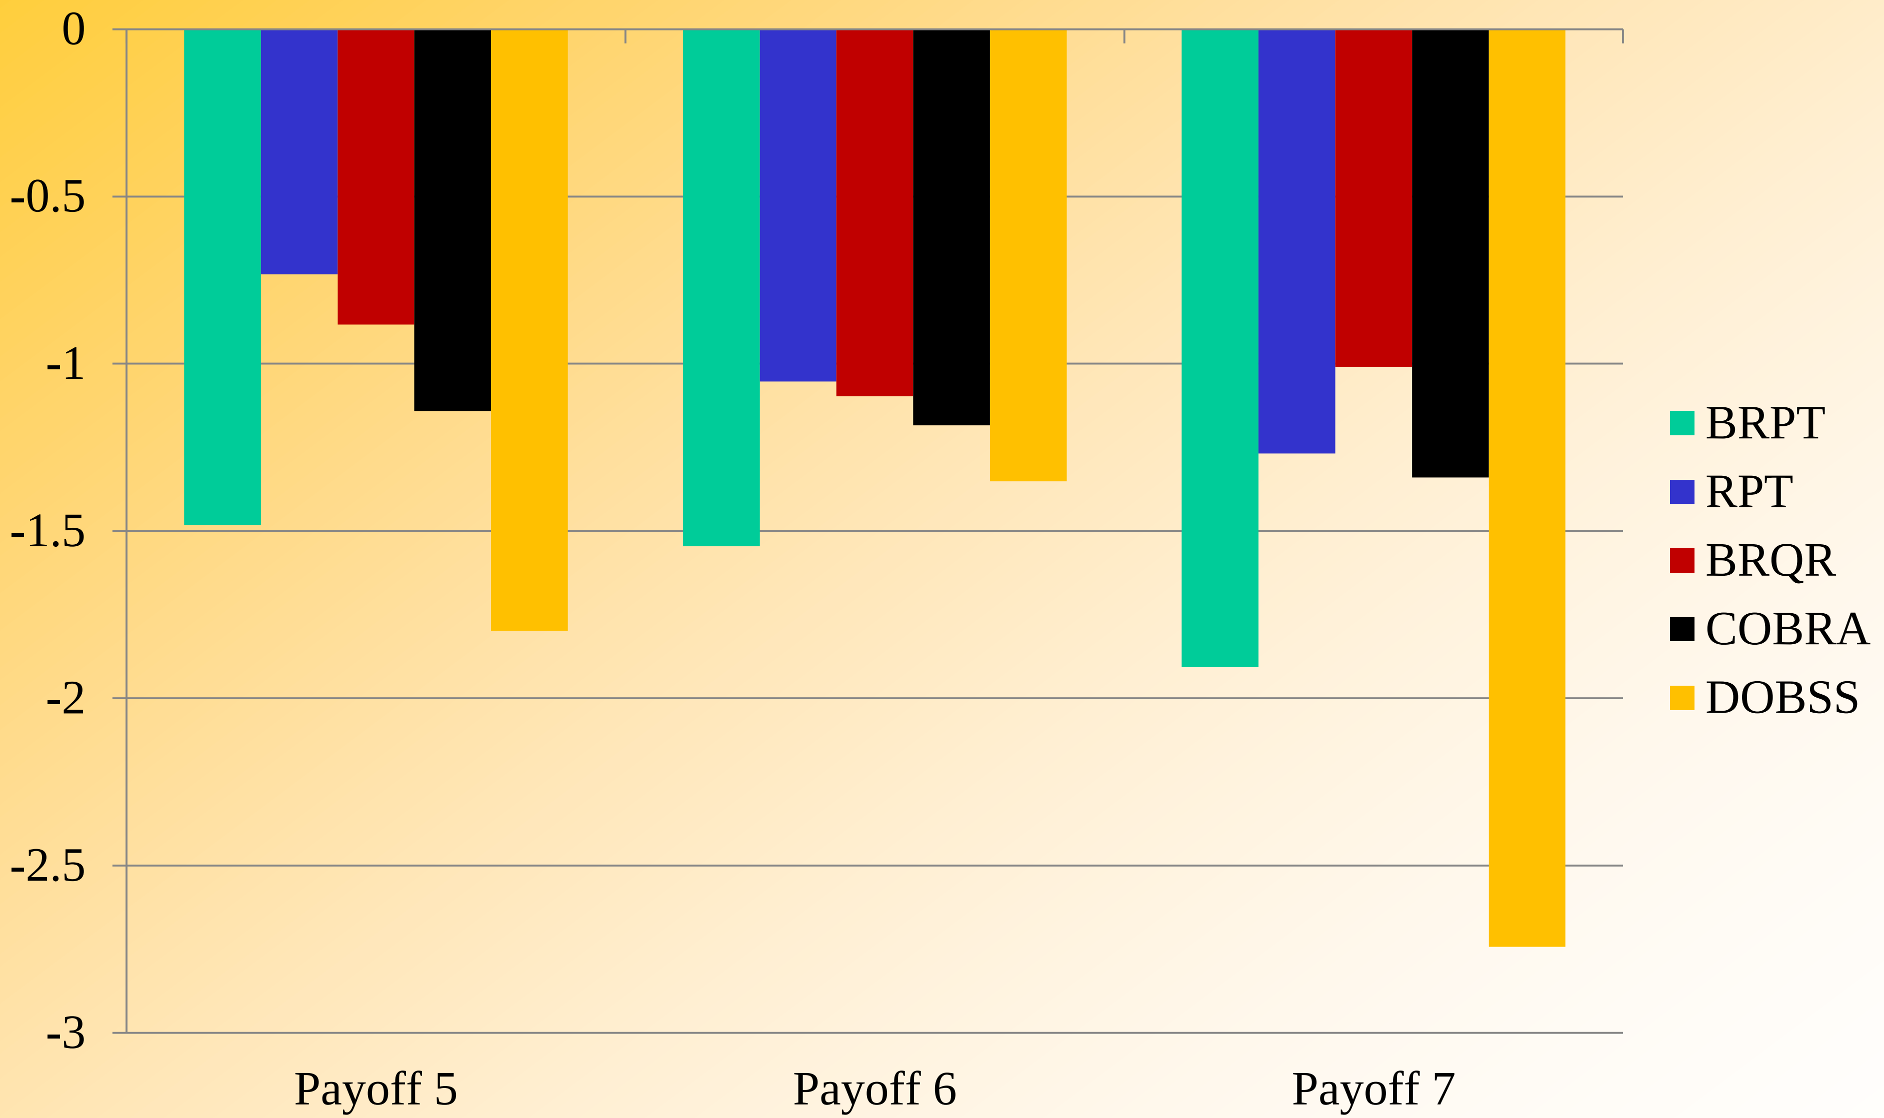

Result Summary

- **BRQR** outperforms **DOBSS** in all 7 payoffs
  - *In payoff 1,3 and 4, the result is statistically significant*
- **BRQR** outperforms **COBRA** in all 7 payoffs
  - *In payoff 2,3 and 4, the result is statistically significant*
- The poor performance BRPT is surprising!

Uncertainty in Adversary Decision: MATCH

Builds on QR, exploiting security game structure:

- Like QR: Adversary response error; better choice more likely
- Bound loss to defender on adversary deviation

Results on 100 games

	MATCH wins	Draw	QR wins
$\alpha = .05$	42	52	6

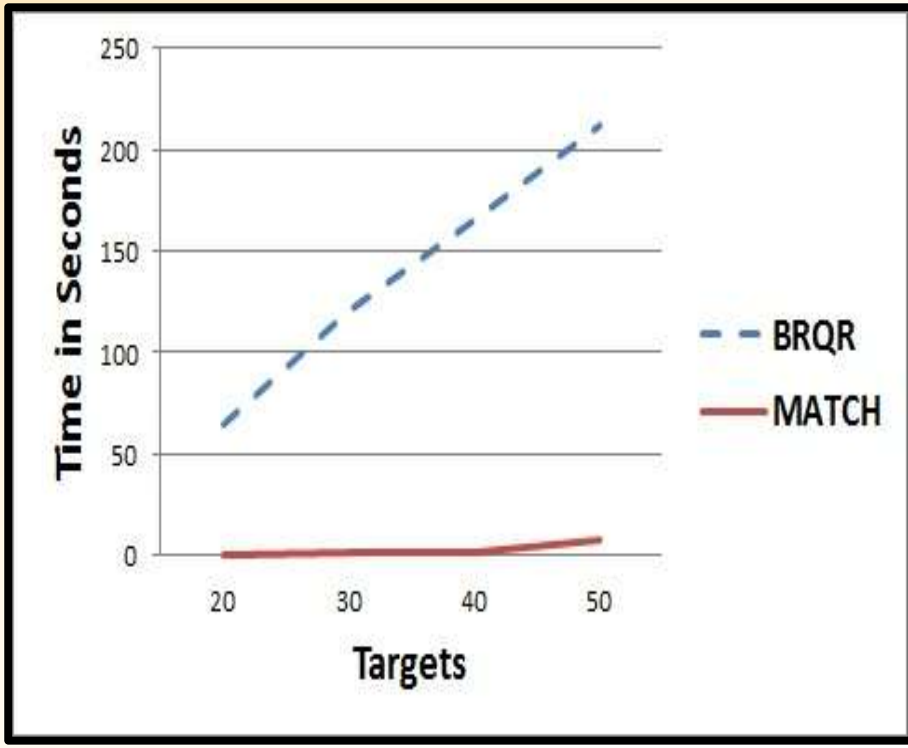

Uncertainty in Attacker Surveillance: Stackelberg vs Nash

- Defender commits first:
  - *Attacker conducts surveillance*
  - *Stackelberg (SSE)*
- Simultaneous move game:
  - *Attacker conducts no surveillance*
  - *Mixed strategy Nash (NE)*

How should a defender compute her strategy?

For security games (*):

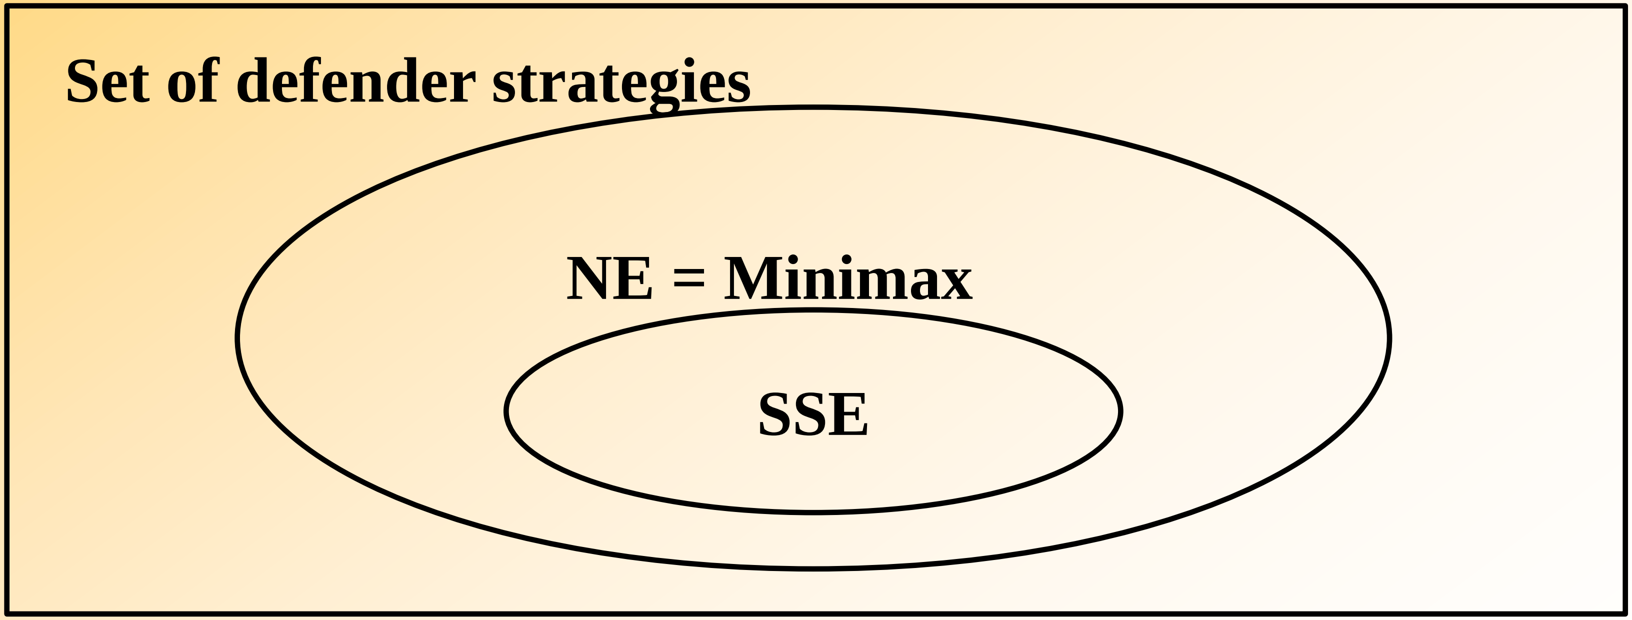

Action Execution & Observation Uncertainty

- RECON:
  - *Worst-case protection against action-execution & observation uncertainty*
  - *Efficient MILP and heuristics*

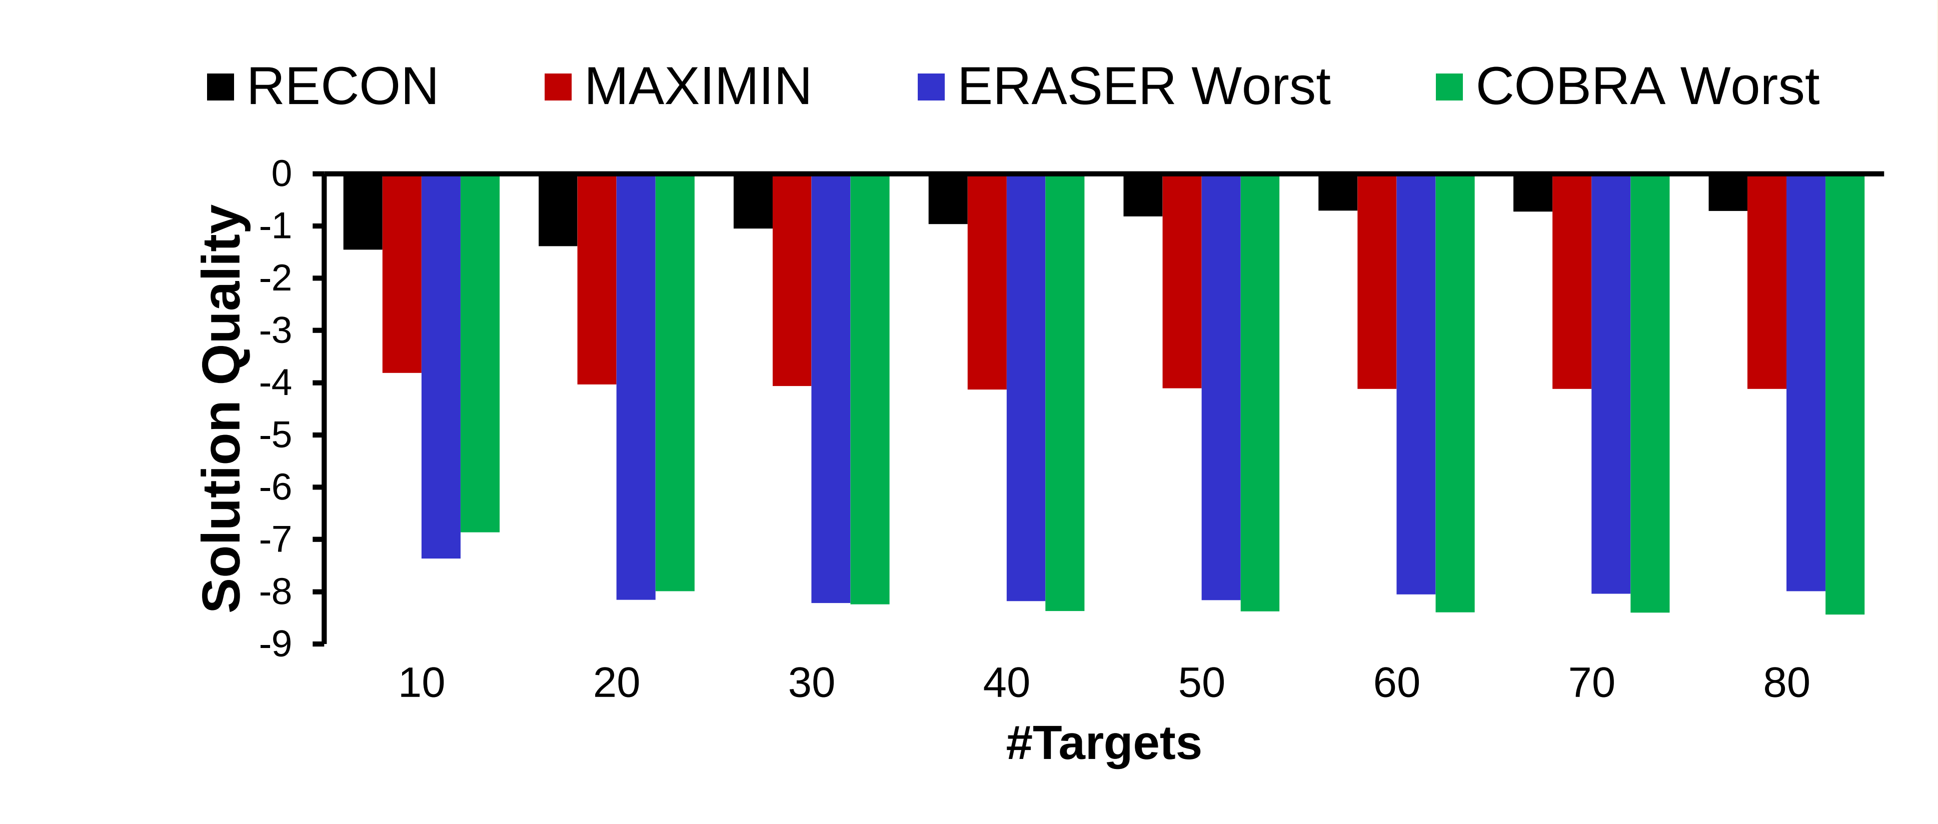

Outline

- Motivating real-world applications
- Background and basic security games
- Scaling to complex action spaces
- Modeling payoff uncertainty: Bayesian Security Games
- Human behavior and observation uncertainty
- ***Evaluation and discussion***

How Do We Evaluate Deployed Systems?

- "Main" vs "Application track": Evaluating deployed systems not easy
  - *Cannot switch security on/off for controlled experiments*
  - *Cannot show we are "safe" (no 100% security)*

- Are our systems useful: Are we better off than previous approaches?
  1. *Models and simulations*
  2. *Human adversaries in the lab*
  3. *Actual security schedules before vs after*
  4. *Expert evaluation*
  5. *"Adversary" teams simulate attack*
  6. *Supportive data from deployment*
  7. *Future deployments*

Key Conclusions

- Human schedulers:
  - *Predictable patterns, e.g. LAX, FAMS (GAO-09-903T)*
  - *Scheduling burden*
- Uniform random:
  - *Non-weighted, e.g. officers to sparsely crowded terminals*
- Simple weighted random:
  - *No adversary reactions, & enumerate large number of combinations?*

Systems in use for a number of years: without us "forcing" use

- *Internal evaluations, e.g. LAX evaluation by FBI, foreign experts*

1. Models and Simulations: Example from IRIS (FAMS)

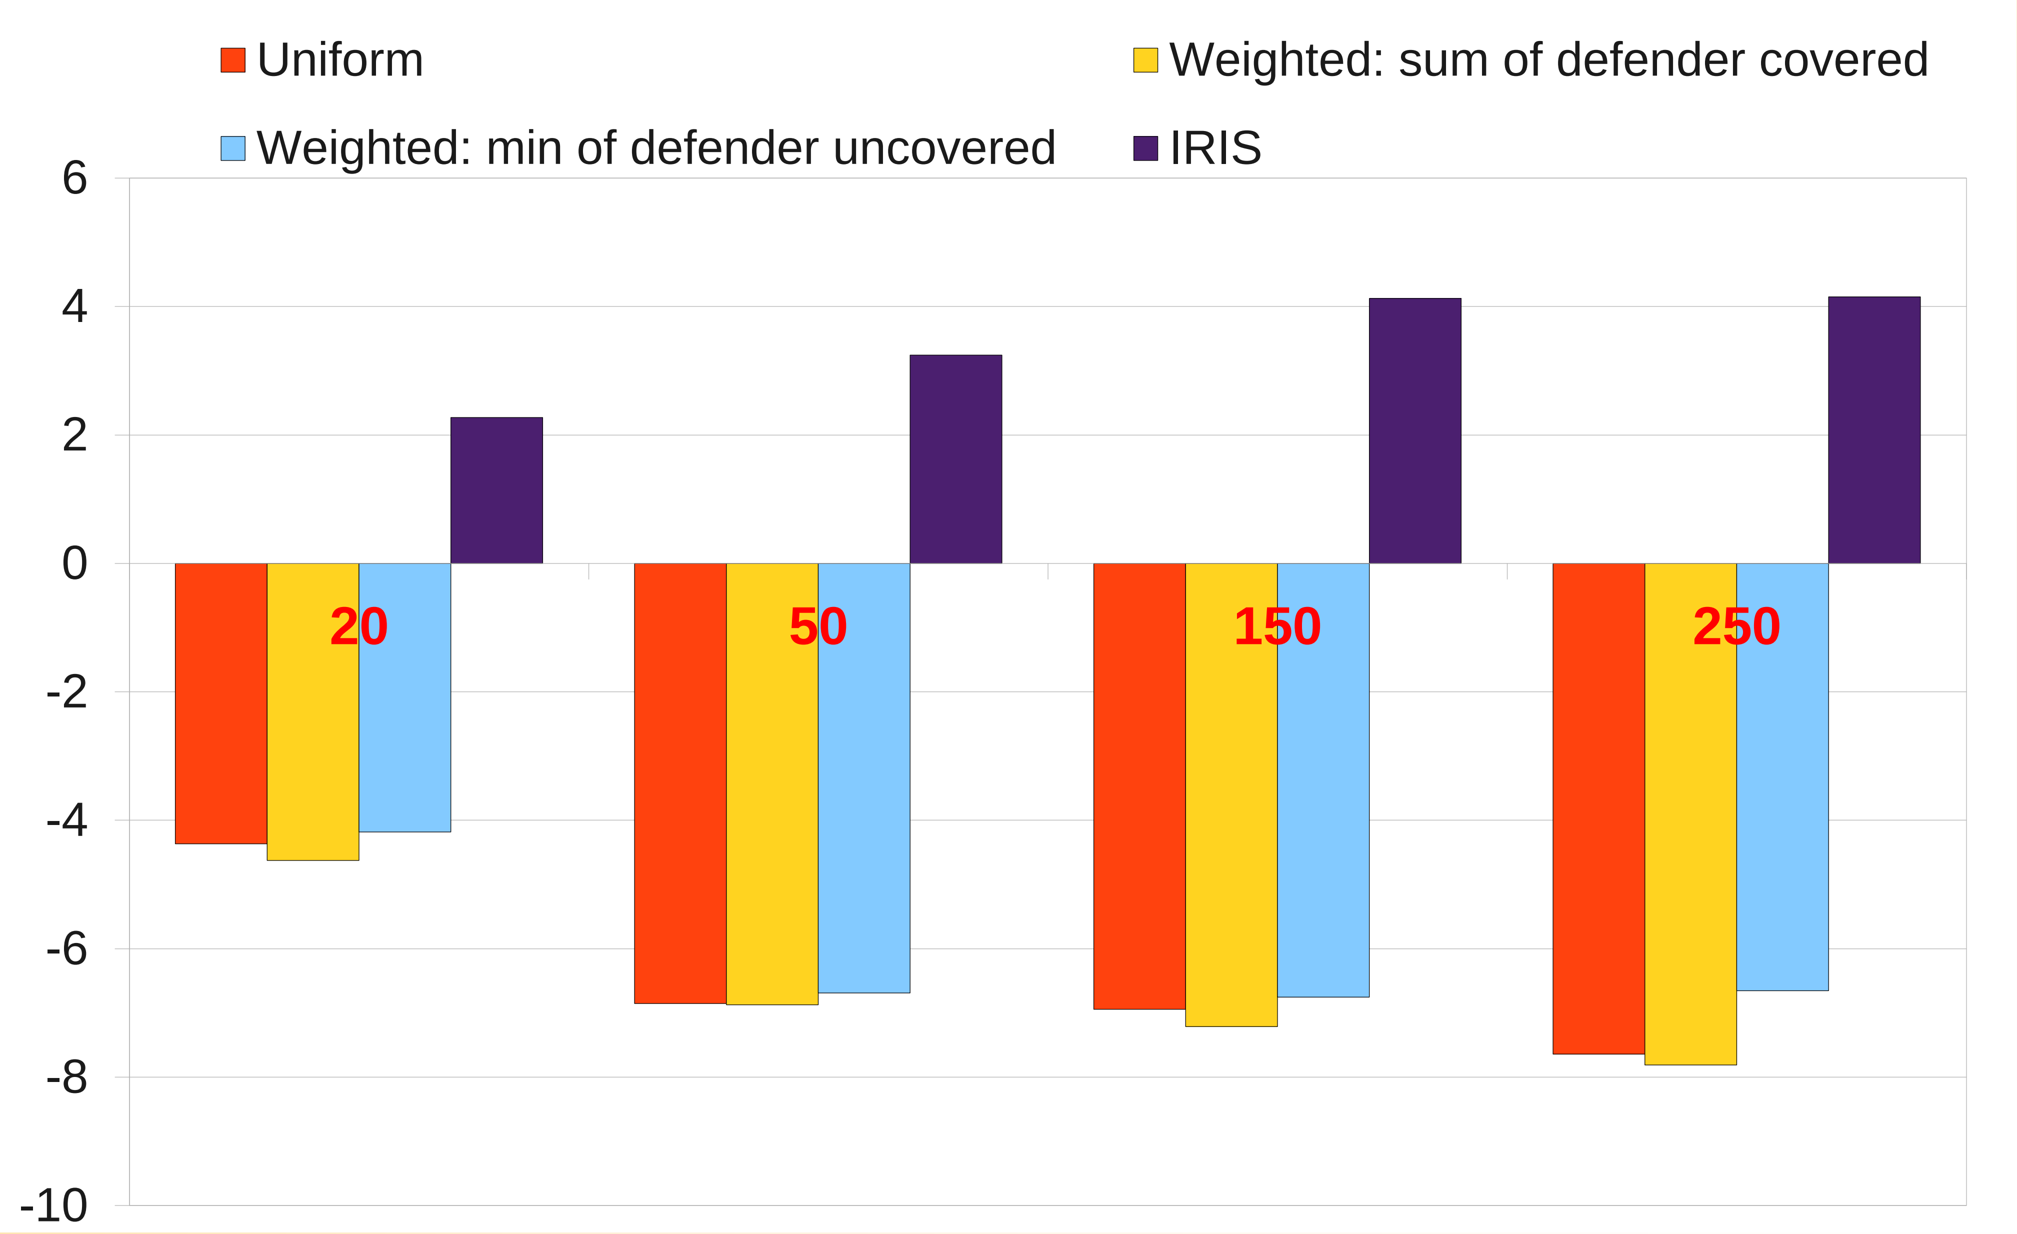

3. Actual Security Schedules Before vs After: Example from PROTECT (Coast Guard)

Patrols Before PROTECT: Boston

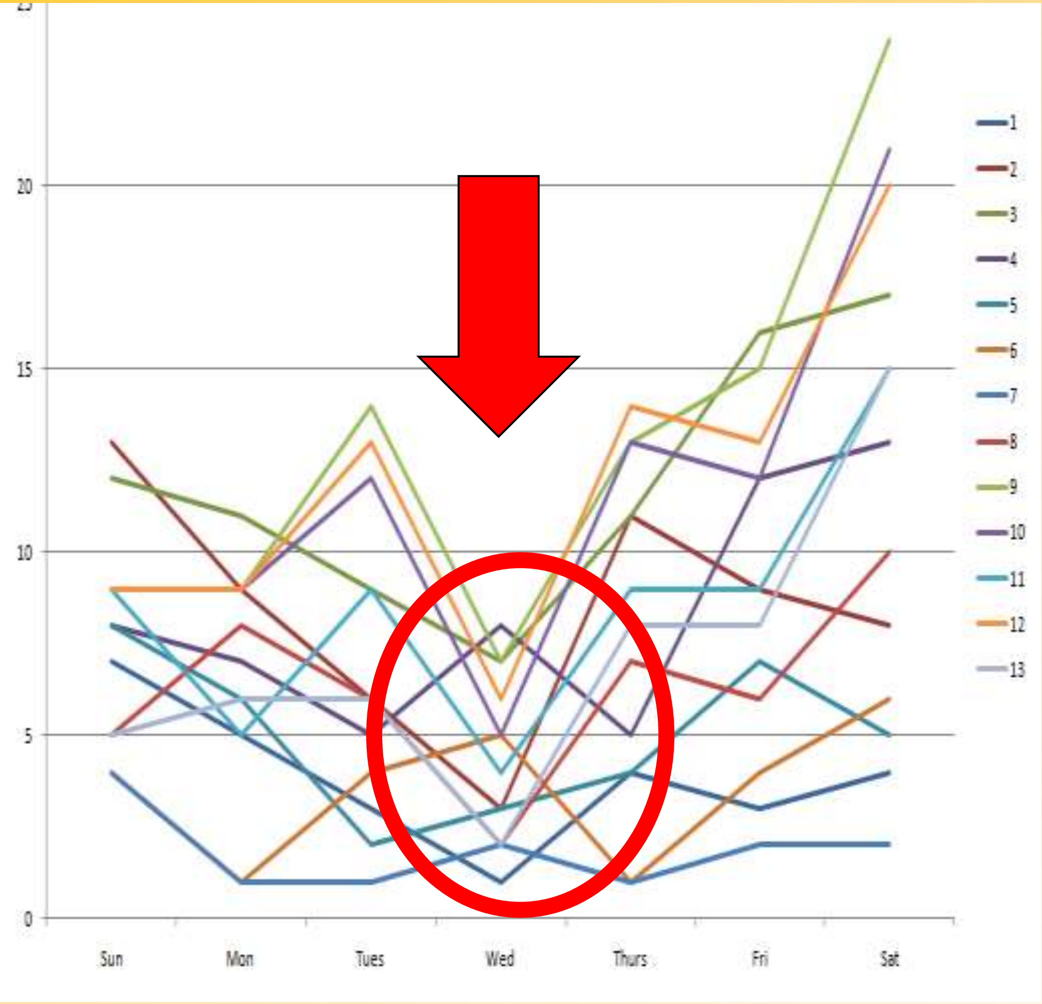

Patrols After PROTECT: Boston

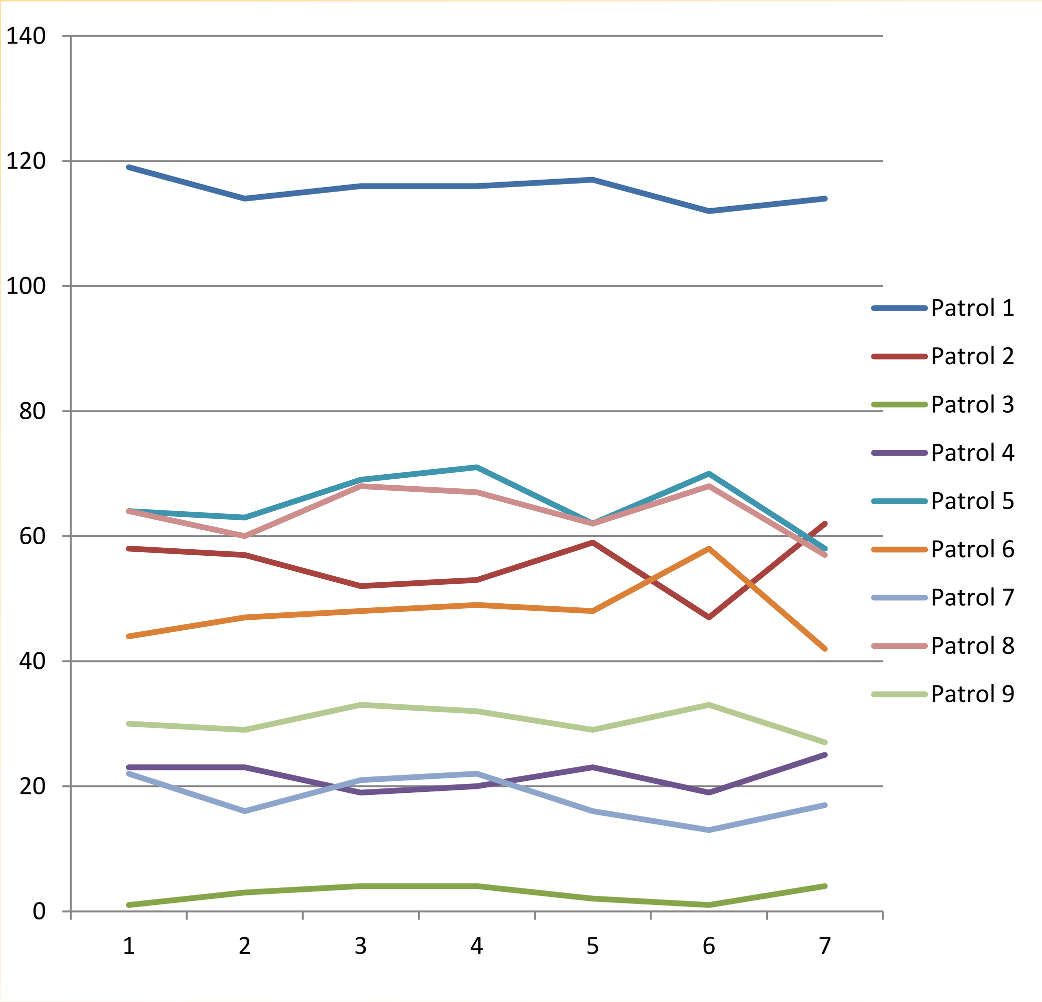

4. Expert Evaluation
Example from ARMOR, IRIS & PROTECT

February 2009: Commendations LAX Police (City of Los Angeles)

July 2011: Operational Excellence Award (US Coast Guard, Boston)

September 2011: Certificate of Appreciation (US Federal Air Marshals Service)

U.S. DEPARTMENT OF HOMELAND SECURITY

Transportation Security Administration

Office of Law Enforcement/Federal Air Marshal Service

Milind Tambe

In recognition and appreciation of your outstanding achievement in developing the Intelligent Randomization In Scheduling (IRIS) program to advance the mission of the Office of Law Enforcement/Federal Air Marshal Service.

This 2nd day of September, 2011

James B. Curren
Studies, Research and Analysis
Office of Flight Operations

5. "Red" Teaming, Supportive data Example from PROTECT

- "Mock attacker" team deployed in Boston
  - *Incorporated adversary's known intent, capability*
  - *Comparing PRE- to POST-PROTECT: "deterrence" improved*

- Additional real-world indicators from Boston:
  - *PRE- to POST-PROTECT: Actual reports of illicit activity*
  - *Industry port partners comments:*
    - **"The Coast Guard seems to be everywhere, all the time."**

(With no actual increase in the number of resources)

6. What Happened at Checkpoints before and after ARMOR -- Not a Controlled Experiment!

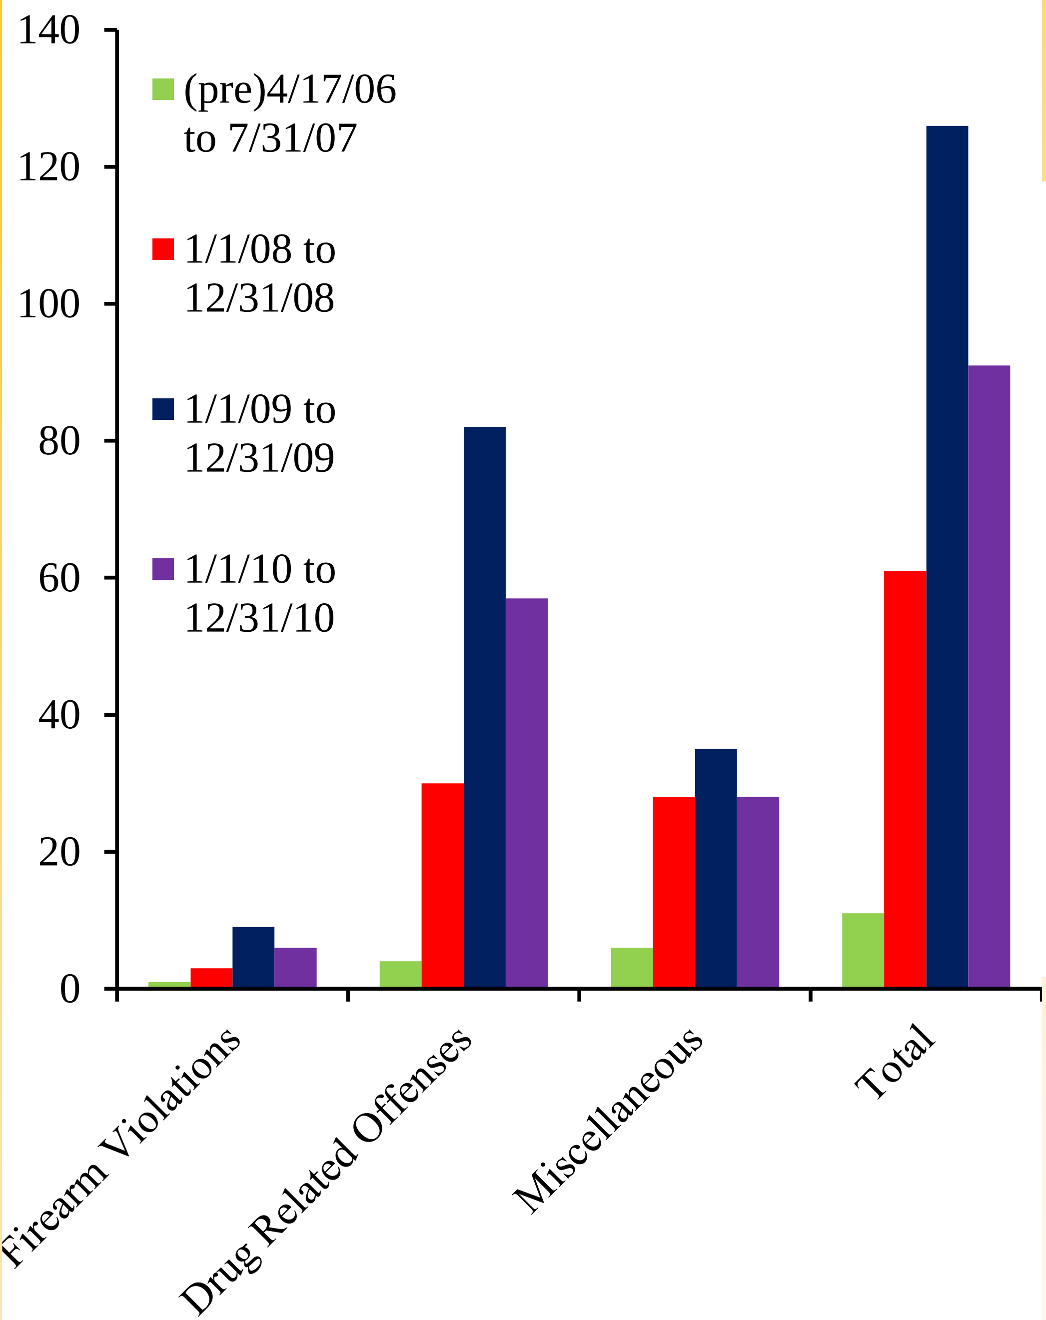

January 2009

- January 3rd — *Loaded 9/mm pistol*
- January 9th — *16-handguns,*
  4-rifles
  1-assault rifle;
  1000 rounds of ammo
- January 10th — *Two unloaded shotguns*
- January 12th — *Loaded 22/cal rifle*
- January 17th — *Loaded 9/mm pistol*
- January 22nd — *Unloaded 9/mm pistol*

Deployed Applications:
ARMOR, IRIS, PROTECT, GUARDS

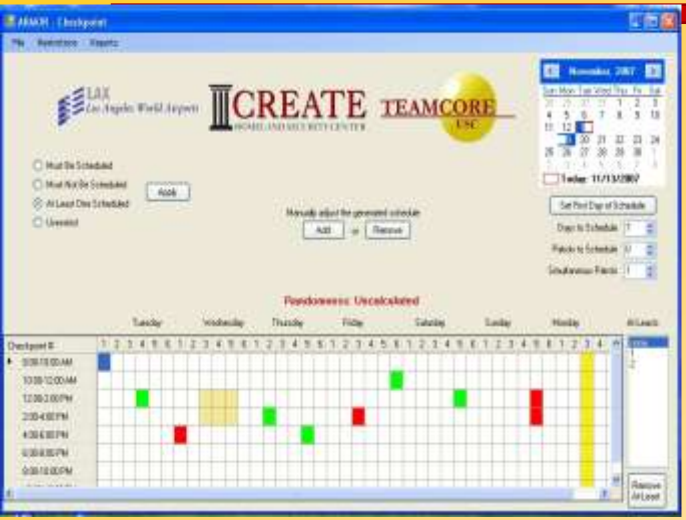

- <u>**Research challenges**</u>
  - *Efficient algorithms:* Scale-up to real-world problems
  - *Observability:* Adversary surveillance uncertainty
  - *Human adversary:* Bounded rationality, observation power
  - *Uncertainty…*

Thank you!

Chris Kiekintveld **Bo An** **Albert Xin Jiang**

cdkiekintveld@utep.edu
boa@usc.edu
jiangx@usc.edu
http://teamcore.usc.edu/security