Towards a Science of Security Games:

Key Algorithmic Principles, Deployed Systems, Research Challenges

Milind Tambe

Helen N. and Emmett H. Jones Professor in Engineering

University of Southern California

with:

Current/former PhD students/postdocs:

Bo An, Matthew Brown, Francesco Delle Fave, Fei Fang, Benjamin Ford, William Haskell, Manish Jain, Albert Jiang, Debarun Kar, Chris Kiekintveld, Rajiv Maheswaran, Janusz Marecki, Sara McCarthy, Thanh Nguyen, Praveen Paruchuri, Jonathan Pearce,James Pita, Yundi Qian, Aaron Schlenker, Eric Shieh, Jason Tsai, Pradeep Varakantham, Haifeng Xu, Amulya Yadav, Rong Yang, Zhengyu Yin, Chao Zhang

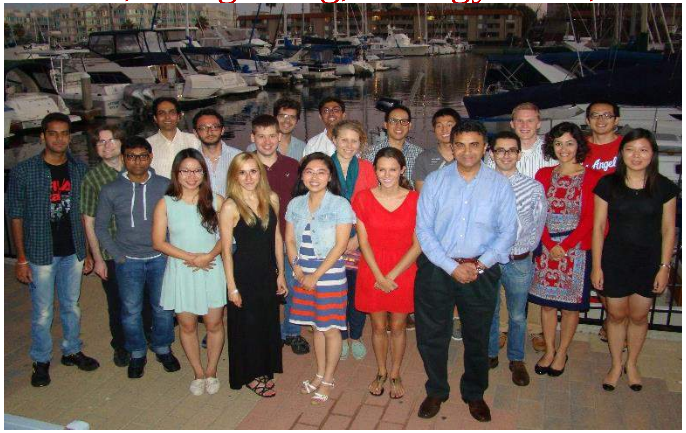

TEAMCORE
USC

Other collaborators:

Shaddin Dughmi (USC), Richard John (USC), David Kempe (USC), Nicholas Weller (USC) &

Craig Boutilier (Toronto), Jeff Brantingahm (UCLA), Vince Conitzer (Duke), Sarit Kraus (BIU, Israel), E Lam (Panthera), Andrew Lemieux (NCSR) Barney Long(WWF), Arnault Lyet (WWF), Kevin Leyton-Brown (UBC), F. Ordonez (U Chile), M. Pechoucek (CTU, Czech R), R. Pickles (Panthera), Andy Plumptre (WCS), Ariel Procaccia (CMU), Martin Short (GATech), P. Stone (Univ of Teaxas), Y. Vorobeychik (Vanderbilt), ….

Global Challenges for Security:
Game Theory for Security Resource Optimization

Example Model: Stackelberg Security Games

Security allocation:

- Targets have weights
- Adversary surveillance

Adversary

Defender

	Target #1	Target #2
Target #1	4, -3	-1, 1
Target #2	-5, 5	2, -1

Example Model: Stackelberg Security Games

Security allocation:

- Targets have weights
- Adversary surveillance

Adversary

Defender

	Target #1	Target #2
Target #1	4, -3	-1, 1
Target #2	-5, 5	2, -1

Example Model:
Stackelberg Security Games

Security allocation:

- Targets have weights
- Adversary surveillance

Adversary

Defender

	Target #1	Target #2
Target #1	4, -3	-1, 1
Target #2	-5, 5	2, -1

Stackelberg Security Games

Security Resource Optimization: *Not 100% Security*

- Random strategy:
  - *Increase cost/uncertainty to attackers*
- Stackelberg game:
  - *Defender commits to mixed strategy*
  - *Adversary conducts surveillance; responds*
- Stackelberg Equilibrium: Optimal random?

Adversary

Defender

	Target #1	Target #2
Target #1	4, -3	-1, 1
Target #2	-5, 5	2, -1

Security Games: Research & Applications

Massive scale games, Plan/reason under uncertainty

Human adversary behavior modeling, Repeated games, Learn from field data

Infrastructure Security Games


Coast Guard

Coast Guard


LAX

TSA


LA Sheriff

Green Security Games

Coast Guard


Panthera


WWF

Opportunistic Crime Security Games

ARMORWAY

Global Presence of Security Games Efforts

Key Lessons: Security Games

Decision aids based on computational game theory in daily use

- Optimize limited security resources against adversaries

Applications yield research challenges: Science of security games

- Scale-up: Incremental strategy generation & Marginals
- Uncertainty: Integrate MDPs, Robustness
- Human behavior models: Quantal response

Current applications: Global, interdisciplinary challenges

- Green security games: criminology, computation, conservation

Outline: Security Games Research (2007-)

Infrastructure Security Games

Green Security Games

Opportunistic Crime Security Games

Air travel

Ports

Trains

Fisheries

Wildlife

Urban Crime

2007 2011 2013 2014 2014 2015

Evaluation I: *AAAI, IJCAI, AAMAS papers…*

Evaluation II: *Real-world deployments (Patience)*

ARMOR Airport Security: LAX [2007]
First Stackelberg Security Game Application

GLASGOW 6/30/07

Template for Security Game Operations [2007]

ARMOR Example

Pita

Paruchuri

	Target #1	Target #2	Target #3
Defender #1	2, -1	-3, 4	-3, 4
Defender #2	-3, 3	3, -2	….
Defender #3	….	….	….

Mixed Integer Program

Pr(Canine patrol, 8 AM @Terminals 2,5,6) = 0.17

Pr(Canine patrol, 8 AM @ Terminals 3,5,7) = 0.33

Canine Team Schedule, July 28

	Term 1	Term 2	Term 3	Term 4	Term 5	Term 6	Term 7	Term 8
8 AM		Team1			Team3	Team5		
9 AM			Team1	Team2				Team4
10 AM		Team3		Team5		Team2		

Security Game MIP [2007]

Generate Mixed Strategy for Defender in ARMOR


Pita


Paruchuri

	Target #1	Target #2	Target #3
Defender #1	2, -1	-3, 4	-3, 4
Defender #2	-3, 3	3, -2	….
Defender #3	….	….	….

$$\max \sum_{i \in X} \sum_{j \in Q} R_{ij} \times x_i \times q_j$$

Maximize defender expected utility

$$s.t. \quad \sum_i x_i = 1$$

Defender mixed strategy

$$\sum_{j \in Q} q_j = 1$$

Adversary response

$$0 \le \left(a - \sum_{i \in X} C_{ij} x_i\right) \le (1 - q_j) M$$

Adversary best response

Security Game Payoffs [2007]

Previous Research Provides Payoffs in Security Game Domains

	Target #1	Target #2	Target #3
Defender #1	2, -1	-3, 4	-3, 4
Defender #2	-3, 3	3, -2	….
Defender #3	….	….	….

$$\max \sum_{i \in X} \sum_{j \in Q} R_{ij} \times x_i \times q_j$$

Maximize defender expected utility

+ Handling Uncertainty

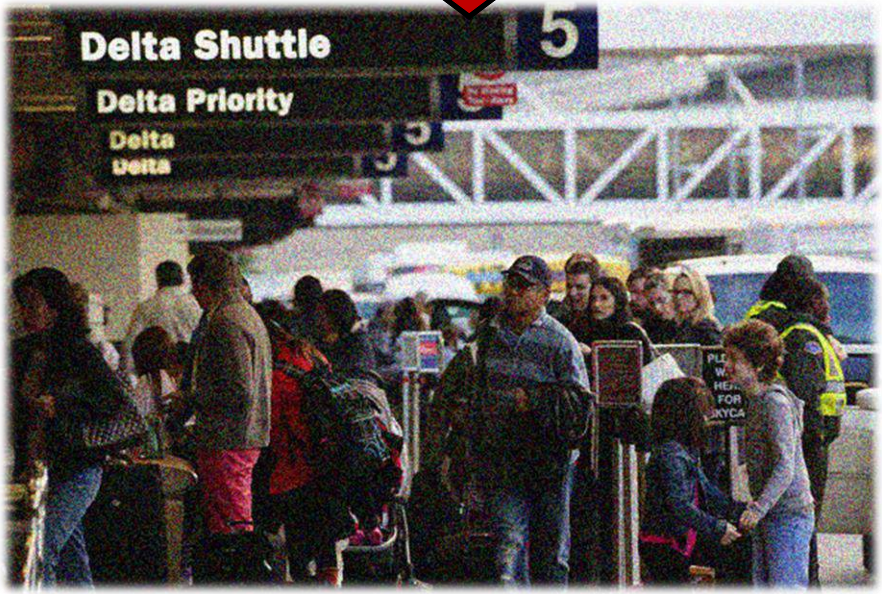


Newsweek
ARMOR…throws a digital cloak of invisibility….

IRIS: Federal Air Marshals Service [2009]
Scale Up Number of Defender Strategies

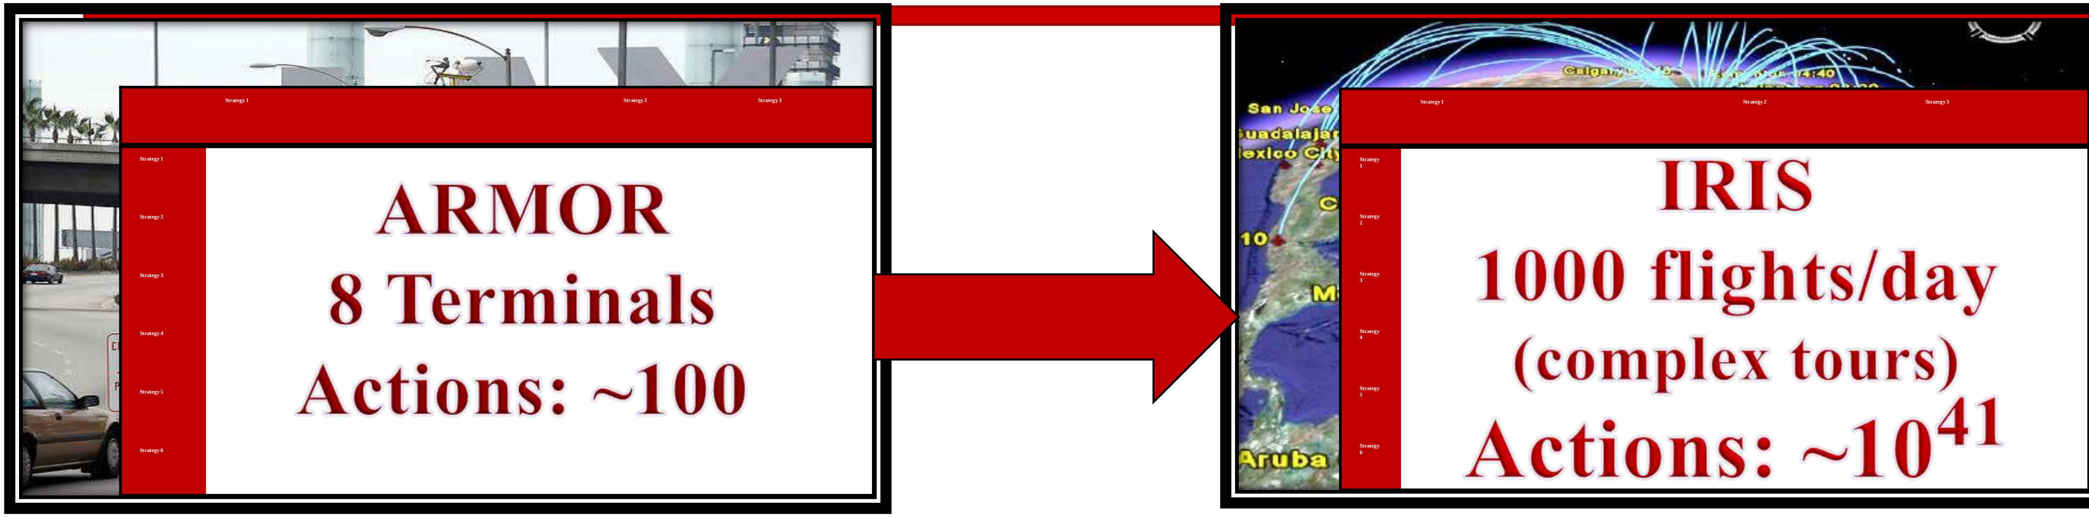

- *1000 Flights, 20 air marshals: 10^{41} combinations*
  - *ARMOR out of memory*
- *Incremental strategy generation:*
  - *Avoid enumerating all possible 10^{41} combinations*

IRIS: Scale Up [2009]
Small Support Set for Mixed Strategies


Jain Kiekintveld

Small support set size:

- Most xi variables zero

1000 flights, 20 air marshals:

10^{41} *combinations*

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

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

~~x123=0.0~~

x124=0.239

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

~~x135=0.0~~

x378=0.123

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

	Attack 1	Attack 2	Attack …	Attack 1000
~~1,2,3~~	~~5,-10~~	~~4,-8~~	~~…~~	~~-20,9~~
1,2,4..	5,-10	4,-8	…	-20,9
~~1,3,5..~~	~~5,-10~~	~~-9,5~~	~~…~~	~~-20,9~~
…				
…				

10^{41} rows

IRIS: Incremental Strategy Generation
Exploit Small Support


Jain Kiekintveld

Master

	Attack 1	Attack 2	Attack…	Attack 6
1,2,4	5,-10	4,-8	…	-20,9

	Attack 1	Attack 2	Attack…	Attack 6
1,2,4	5,-10	4,-8	…	-20,9
3,7,8	-8, 10	-8,10	…	-8,10

Slave (LP Duality Theory)

Best new pure strategy:

Minimum cost network flow

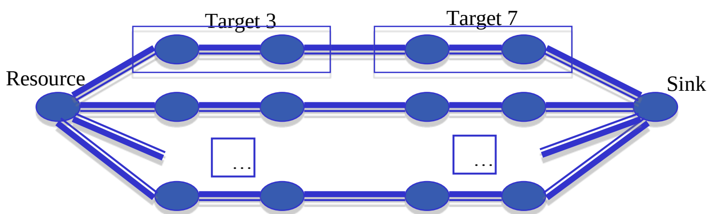

Converge:
GLOBAL
OPTIMAL

	Attack 1	Attack 2	Attack	Attack 6
1,2,4	5,-10			,9
3,7,8	-8, 10			0
…				

500 rows

NOT 10^{41}

IRIS: Deployed FAMS (2009-)

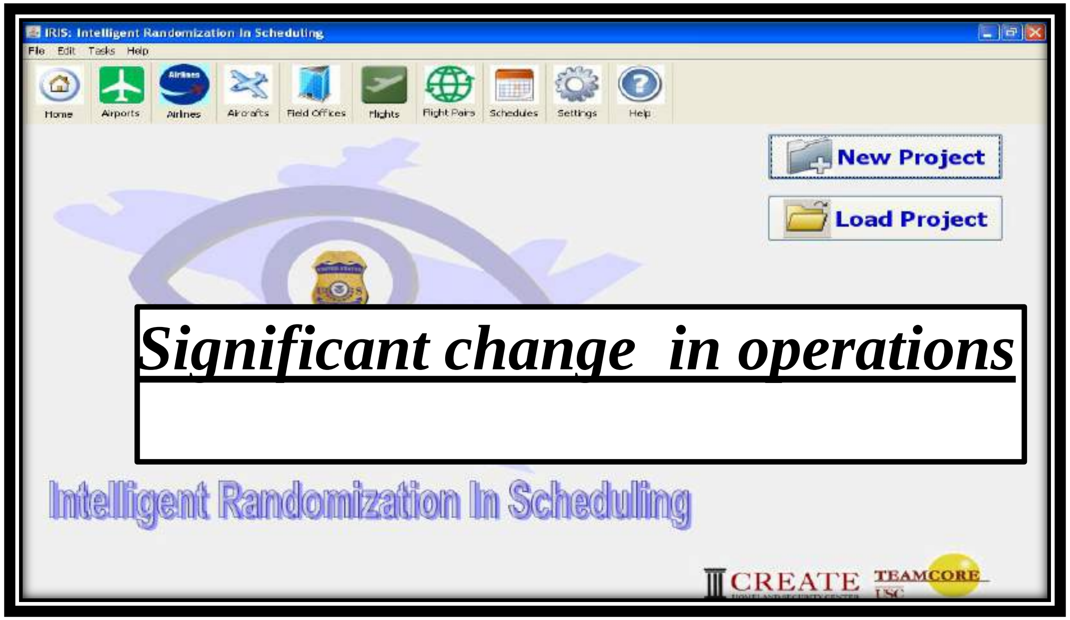

"…in 2011, the Military Operations Research Society selected a University of Southern California project with **FAMS on randomizing flight schedules for the prestigious Rist Award**…"

-R. S. Bray (TSA)

Transportation Security Subcommittee

US House of Representatives 2012

Road, Transportation, Social Networks[2013]
Results of Scale-up

Jain

Road networks:

20,000 roads, 15 checkpoints
(20 min)

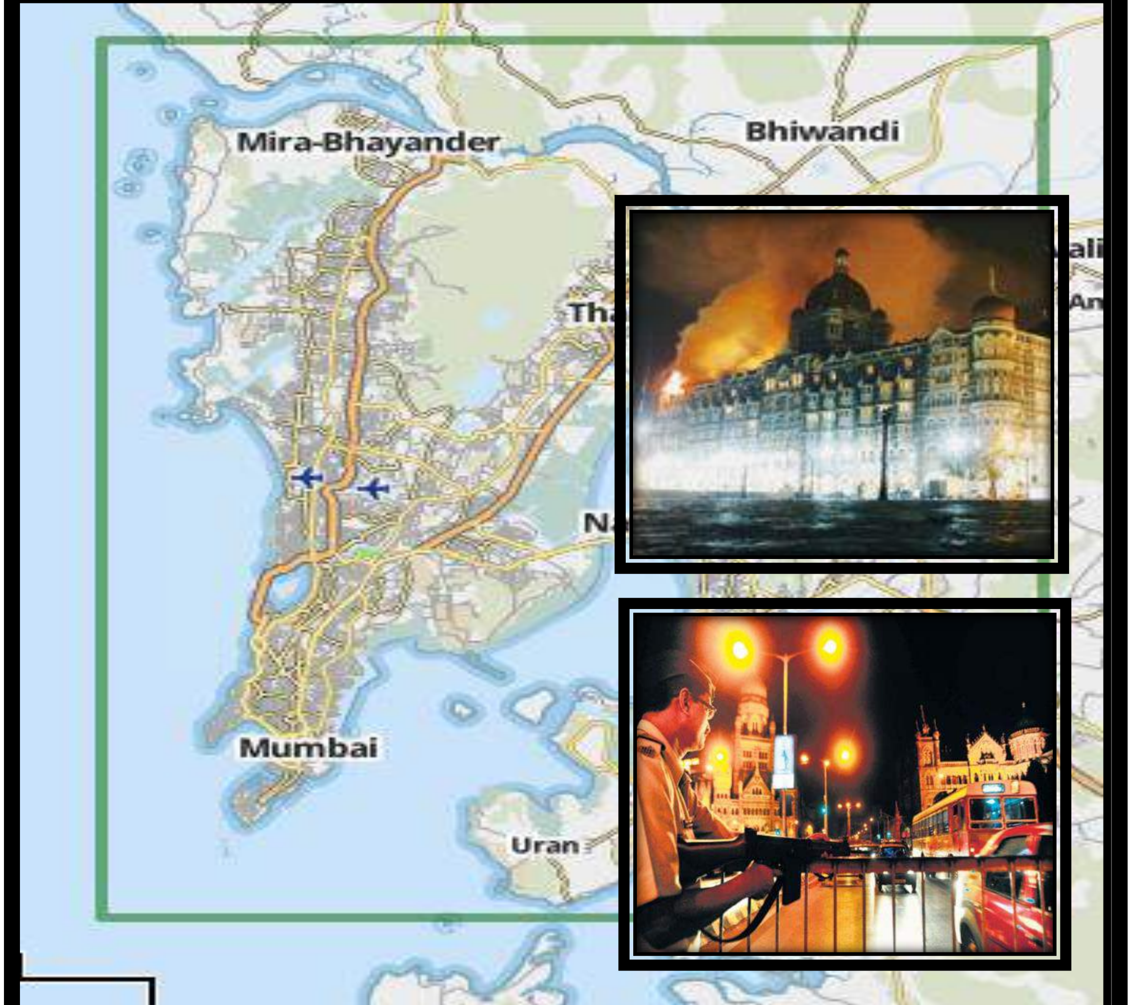

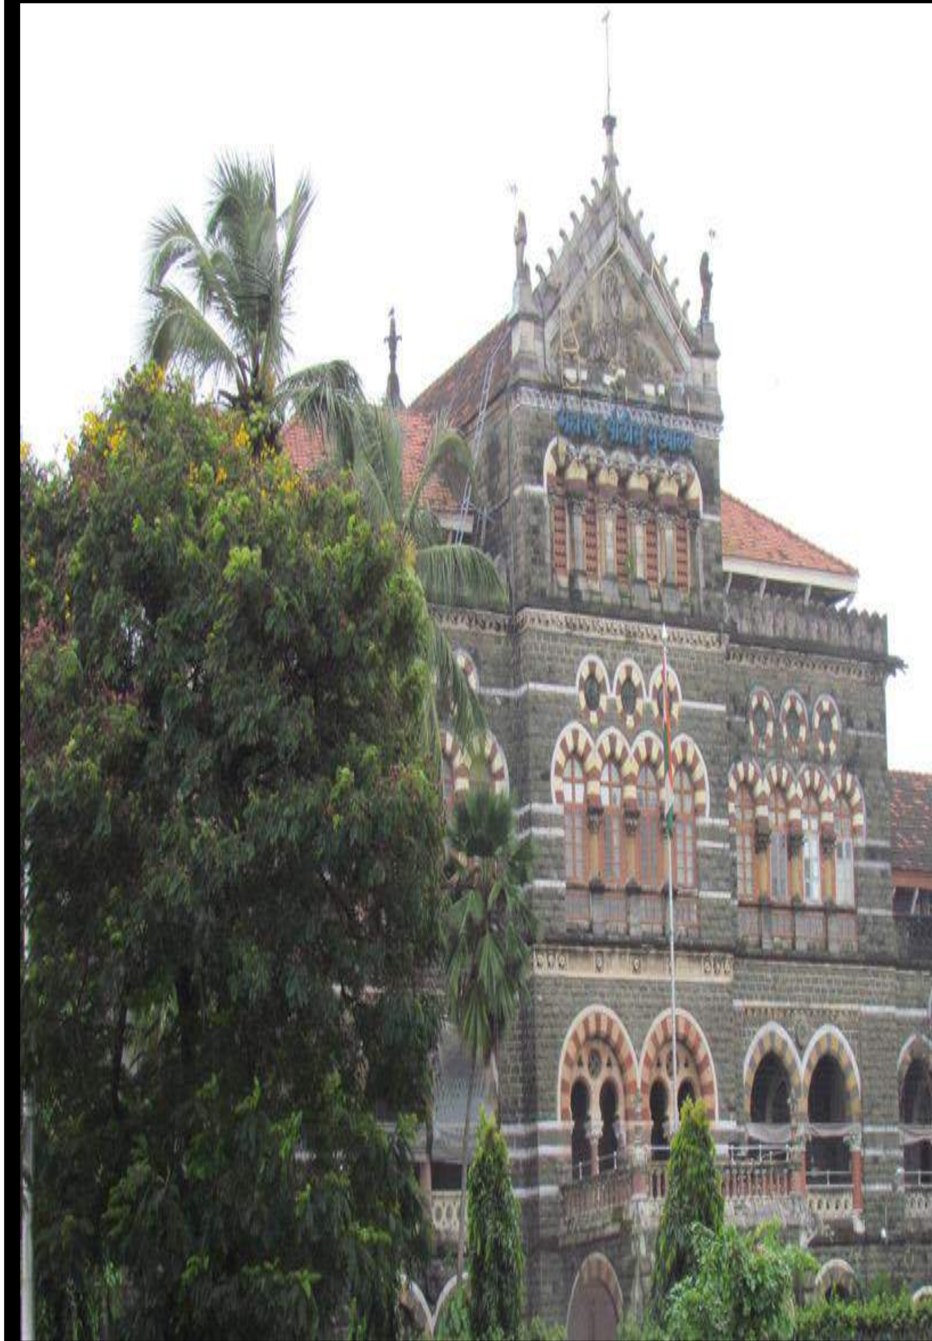

Outline: Security Games Research (2007-)

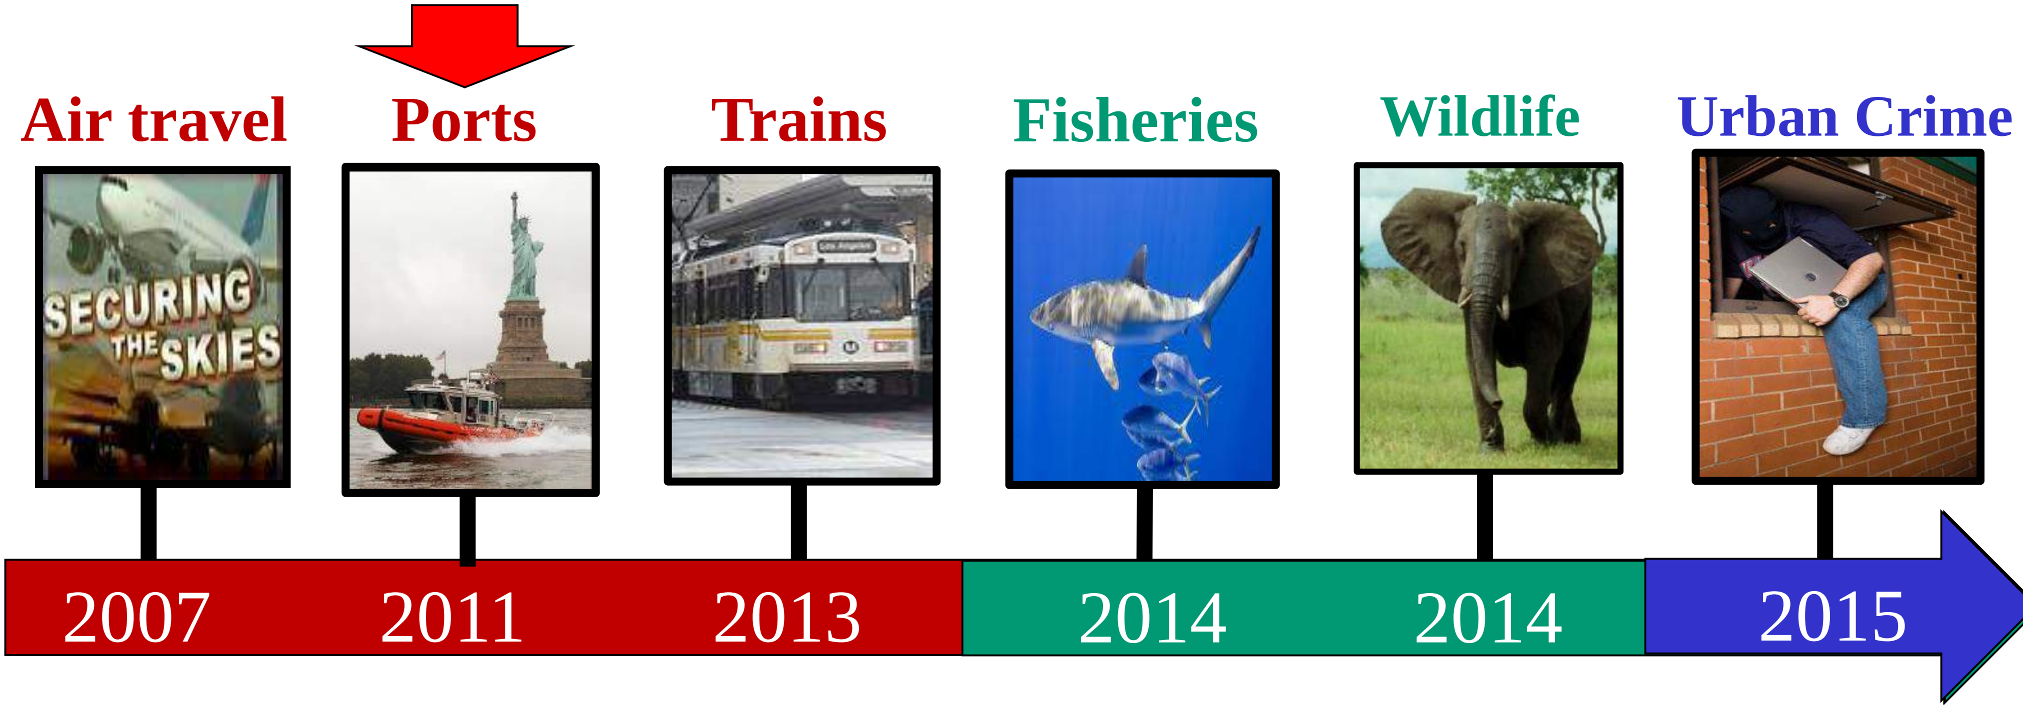

Port Security Threat Scenarios

USS *Cole* after suicide attack

French oil tanker hit by small boat

PROTECT: Randomized Patrol Scheduling [2013]

Port Protection (Scale-up) and Ferries (Continuous Space/time)

PROTECT: Randomized Patrol Scheduling [2013]

Port Protection (Scale-up) and Ferries (Continuous Space/time)

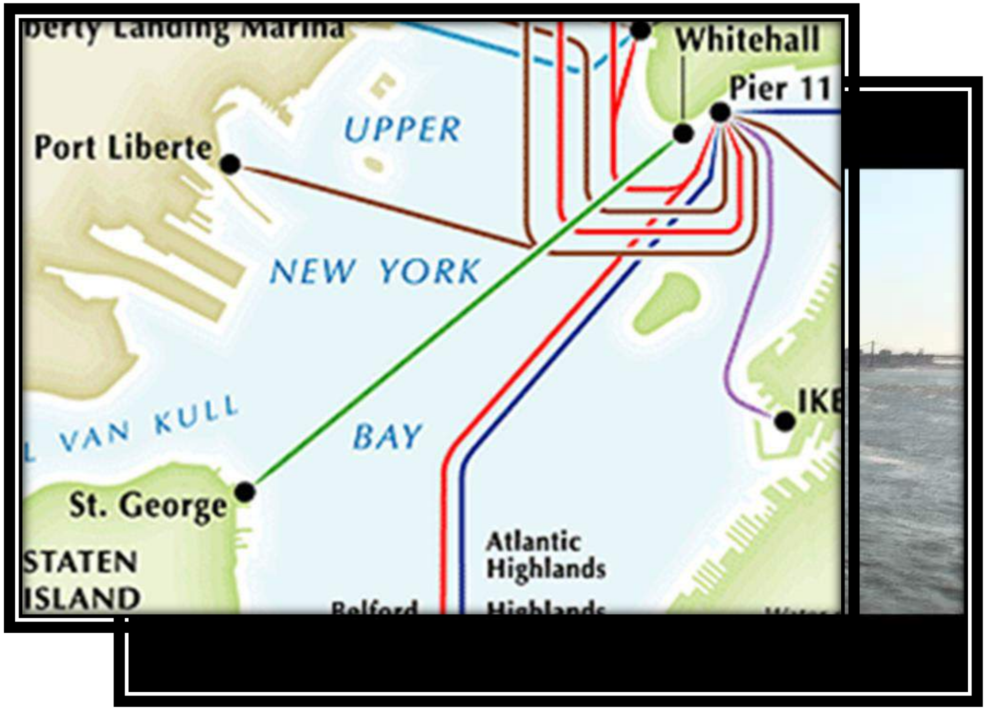

Ferries: Scale-up with Mobile Resources & Moving Targets
Transition Graph Representation

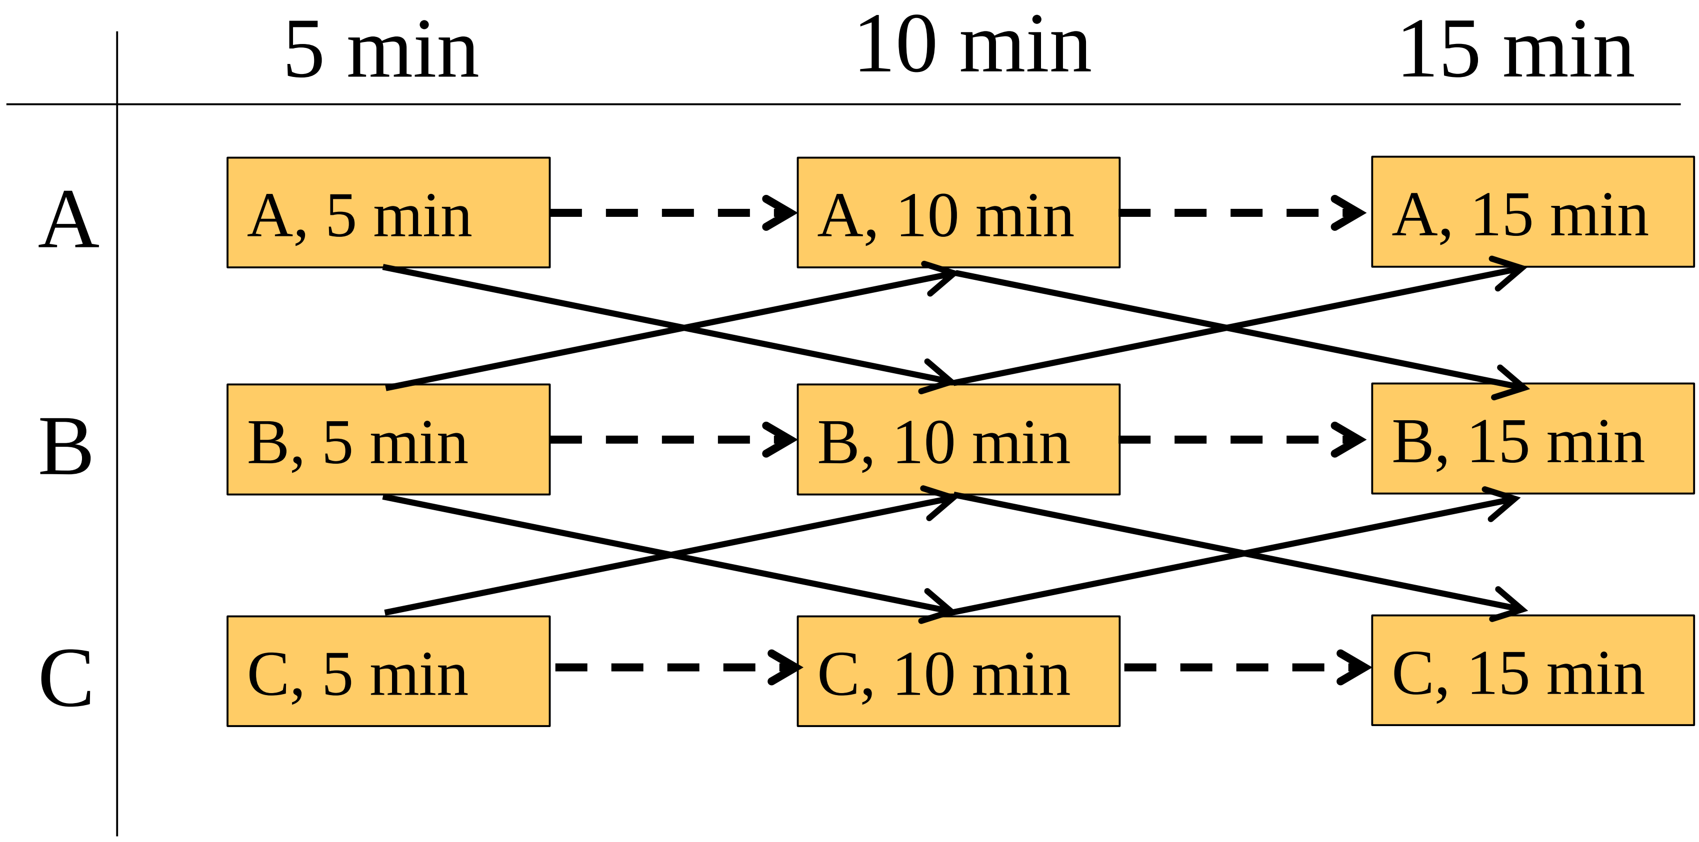

Ferries: Scale-up with Mobile Resources & Moving Targets
Transition Graph Representation

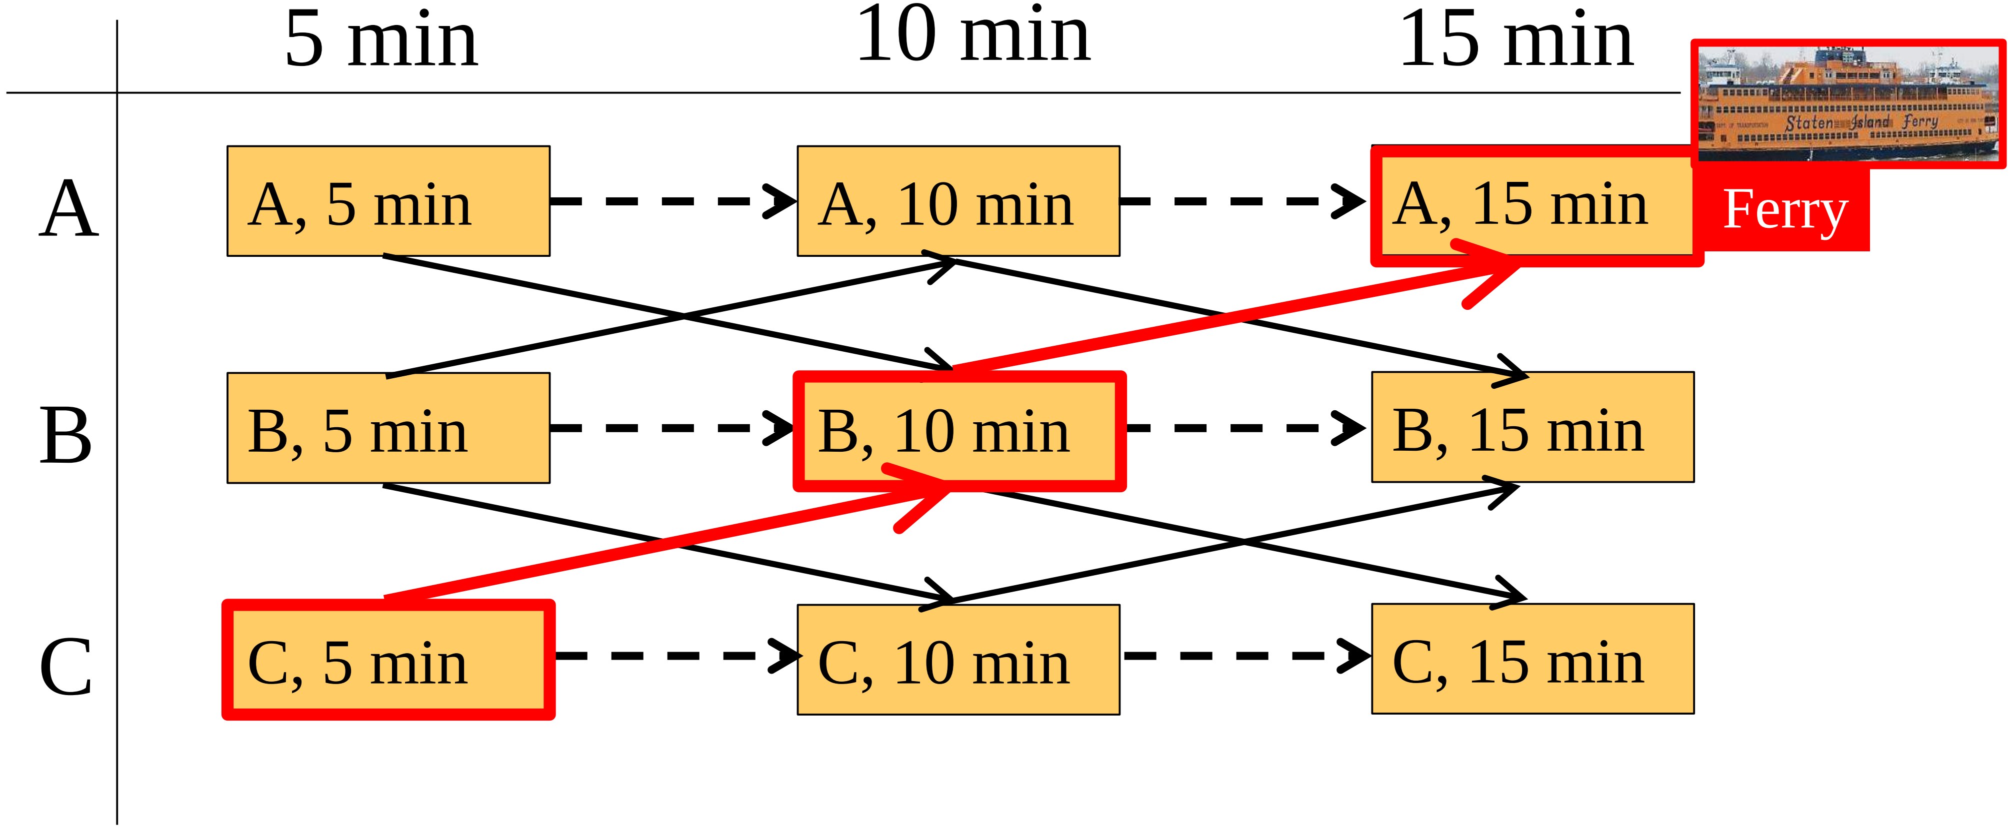

Ferries: Patrol Routes as Variables
Exponential Numbers of Patrol Routes

- Patrols protect nearby ferry location; Solve as done in ARMOR

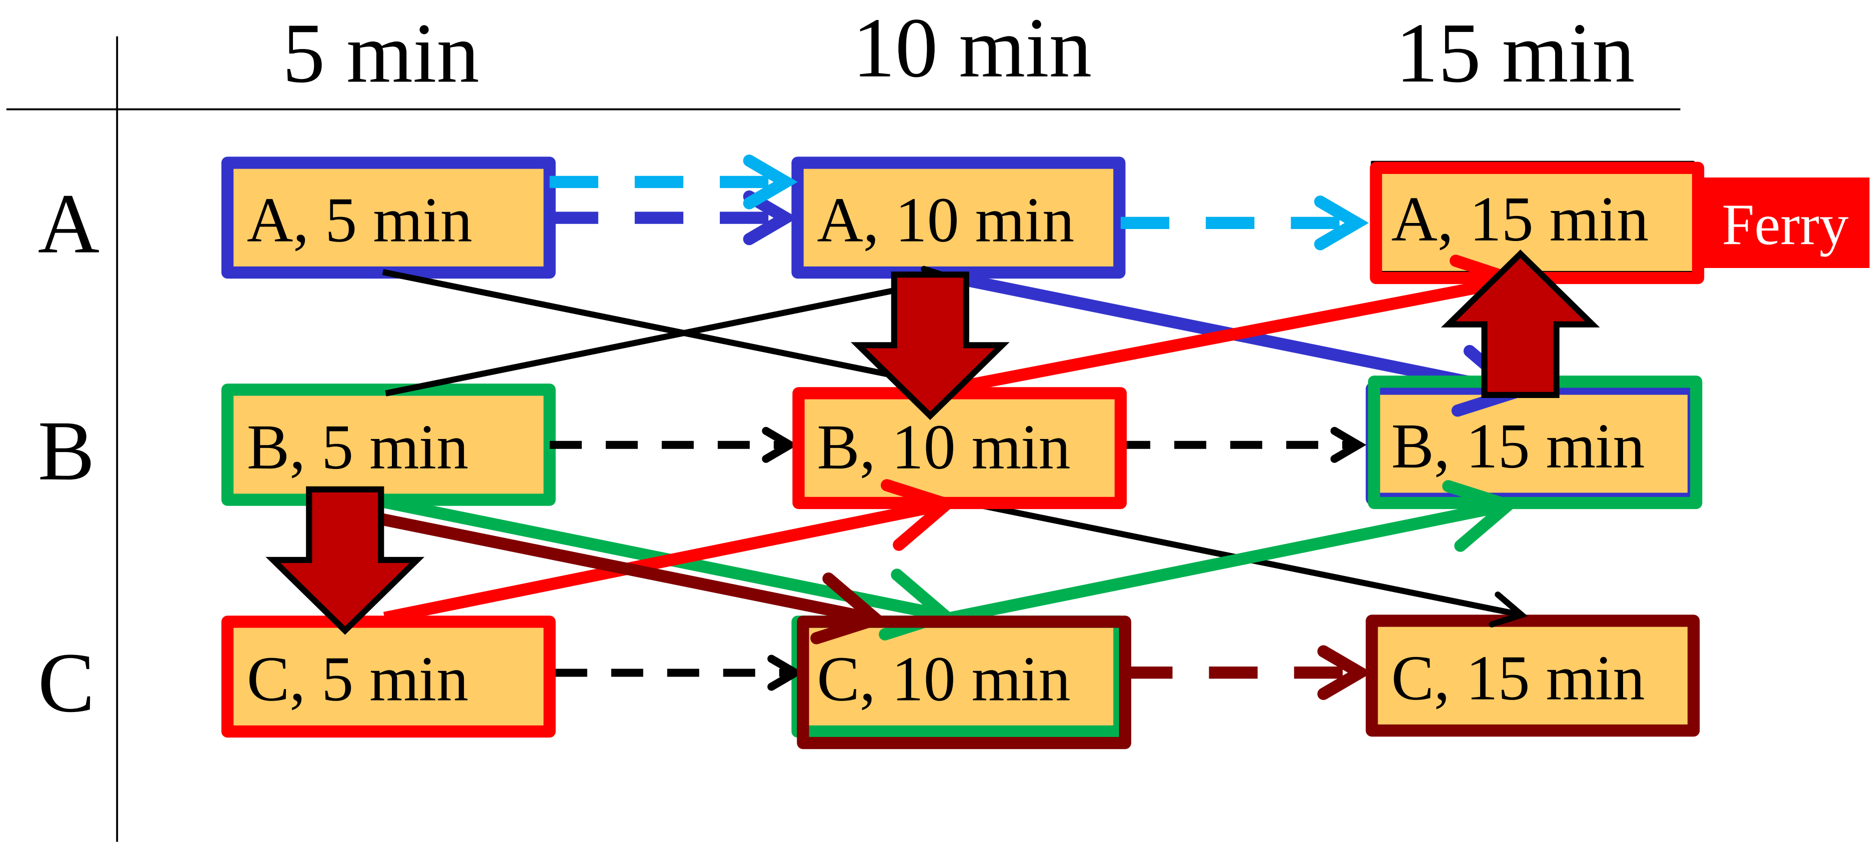

Ferries: Patrol Routes as Variables
Exponential Numbers of Patrol Routes


Fang


Jiang


Kiekintveld

- Patrols protect nearby ferry location; Solve as done in ARMOR
  - *Pr([(B,5), (C, 10), (C,15)]) = 0.47*
  - *Pr([(B,5), (C,10), (B,15)]) =0.23*
  - *Pr([(A,5), (A,10), (B,15)]) = 0.17*
  - *Pr([(A,5), (A,10), (A,15)]) = 0.13*

$\boxed{N^T \textit{ variables}}$

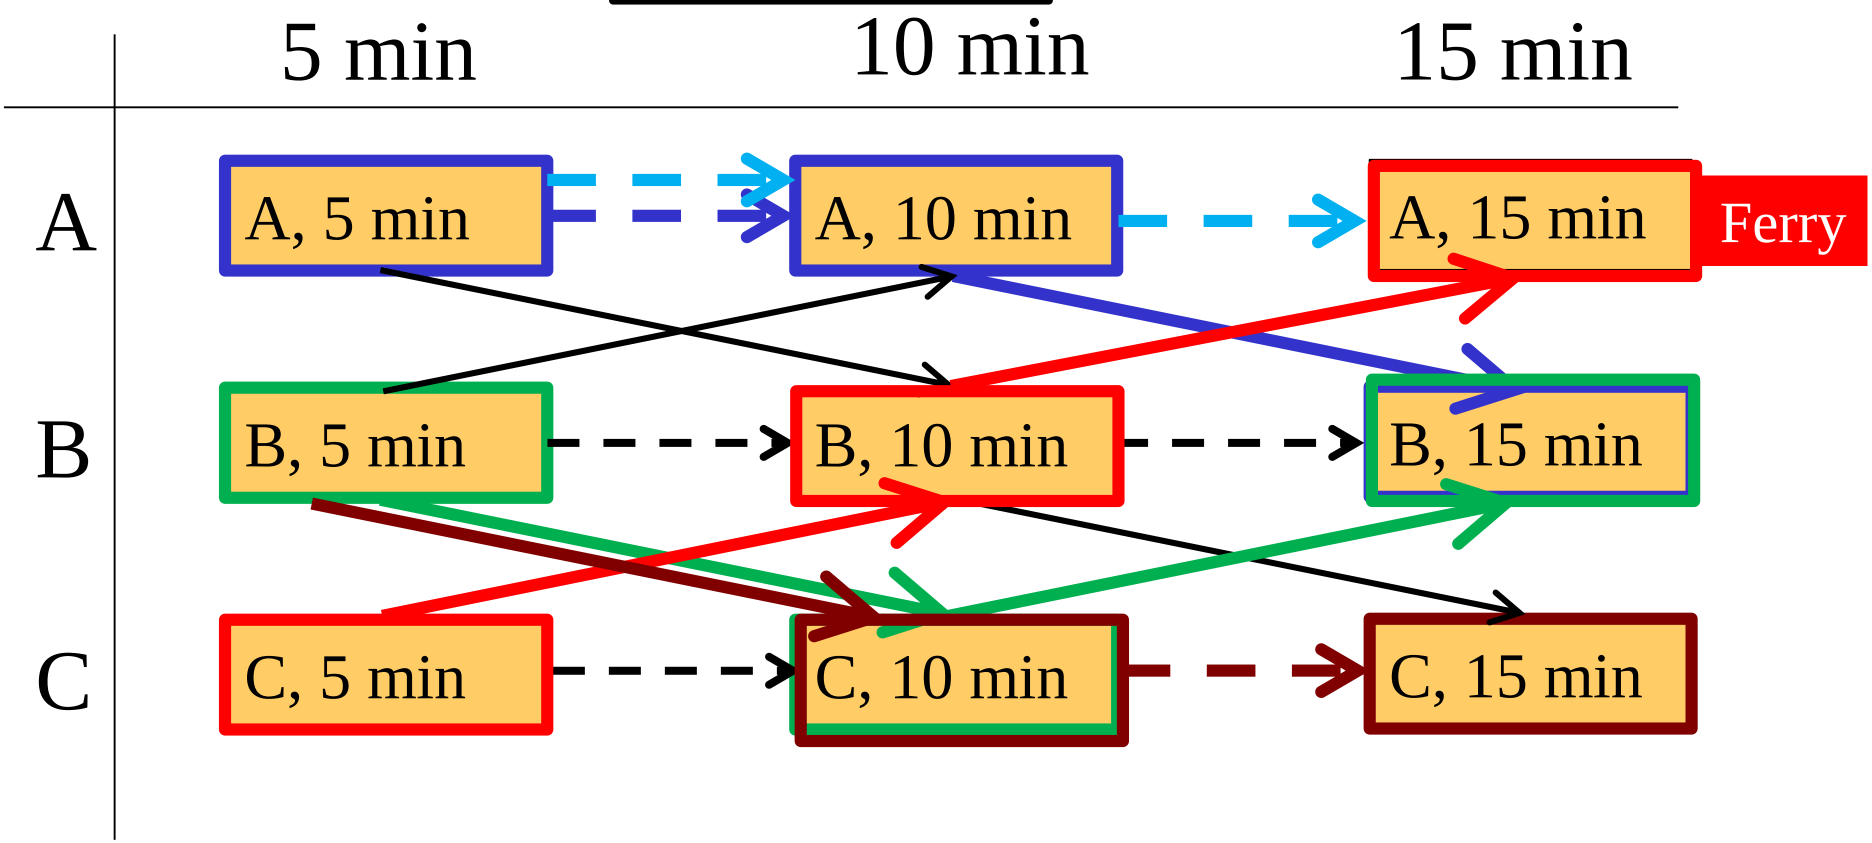

Ferries: Scale-up
Marginal Probabilities Over Segments

Fang *Jiang* *Kiekintveld*

- Variables: NOT routes, but probability flow over each segment

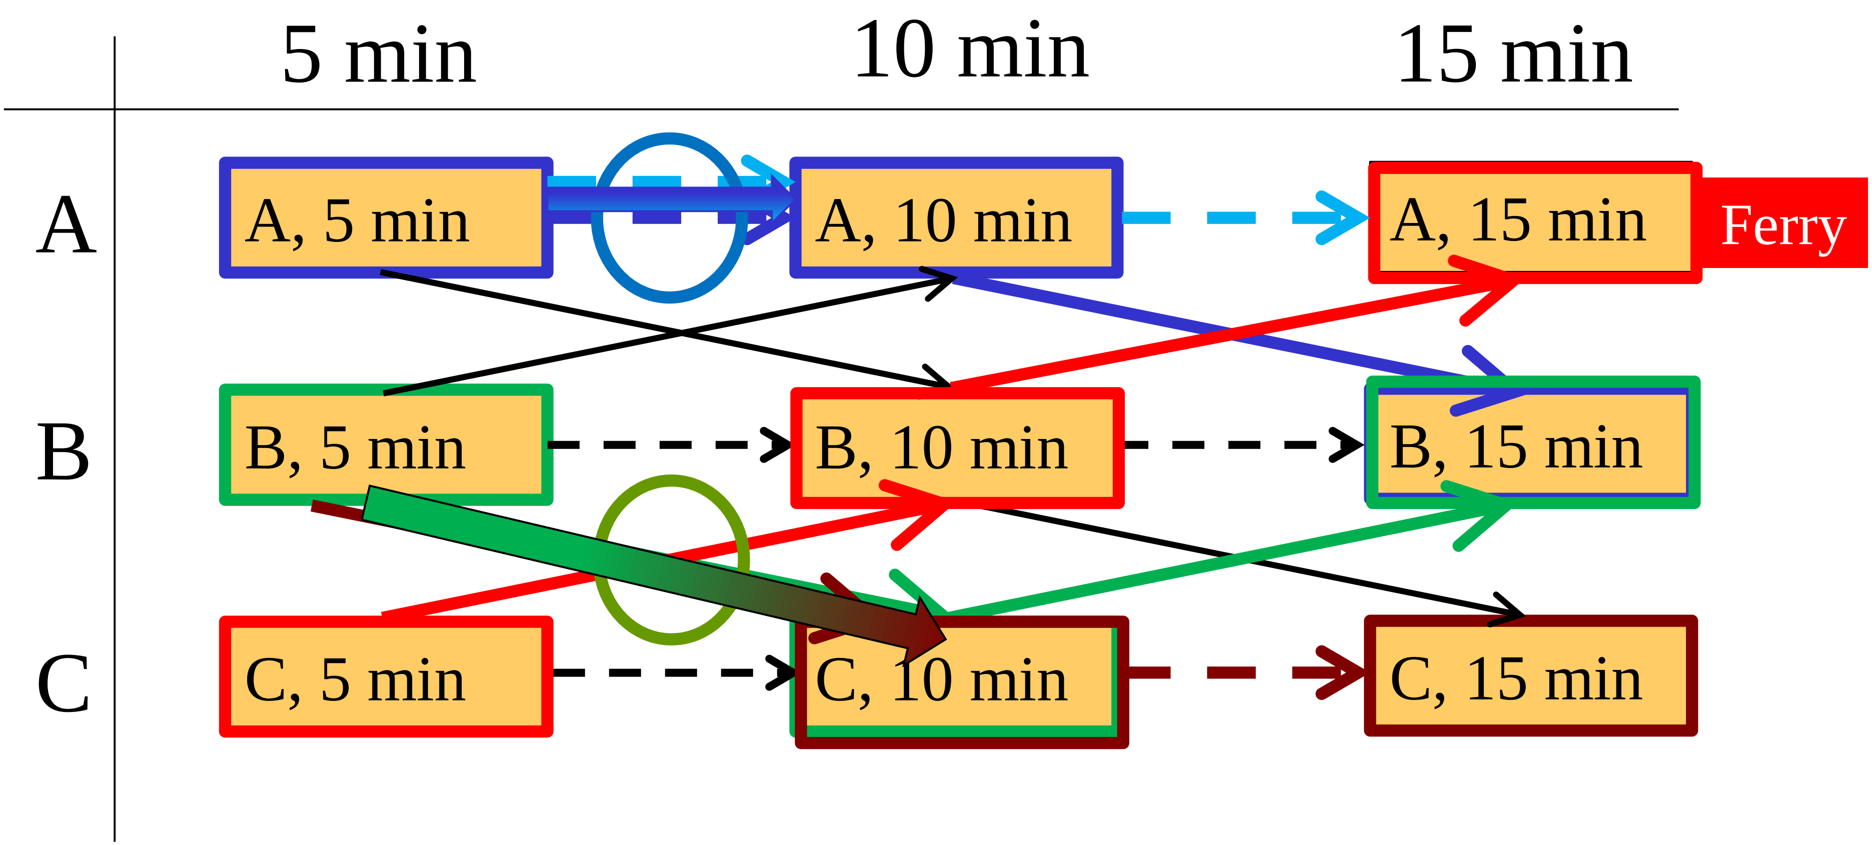

Ferries: Scale-up with Marginals Over Separable Segments
Significant Speedup

- *Obtain marginal probabilities over segments*

N^2.T variables

~~N^T variables~~

Extract: ...(B,5), (C, 10), (C,15)]) = 0.47
...), (C,10), (B,15)]) =0.23

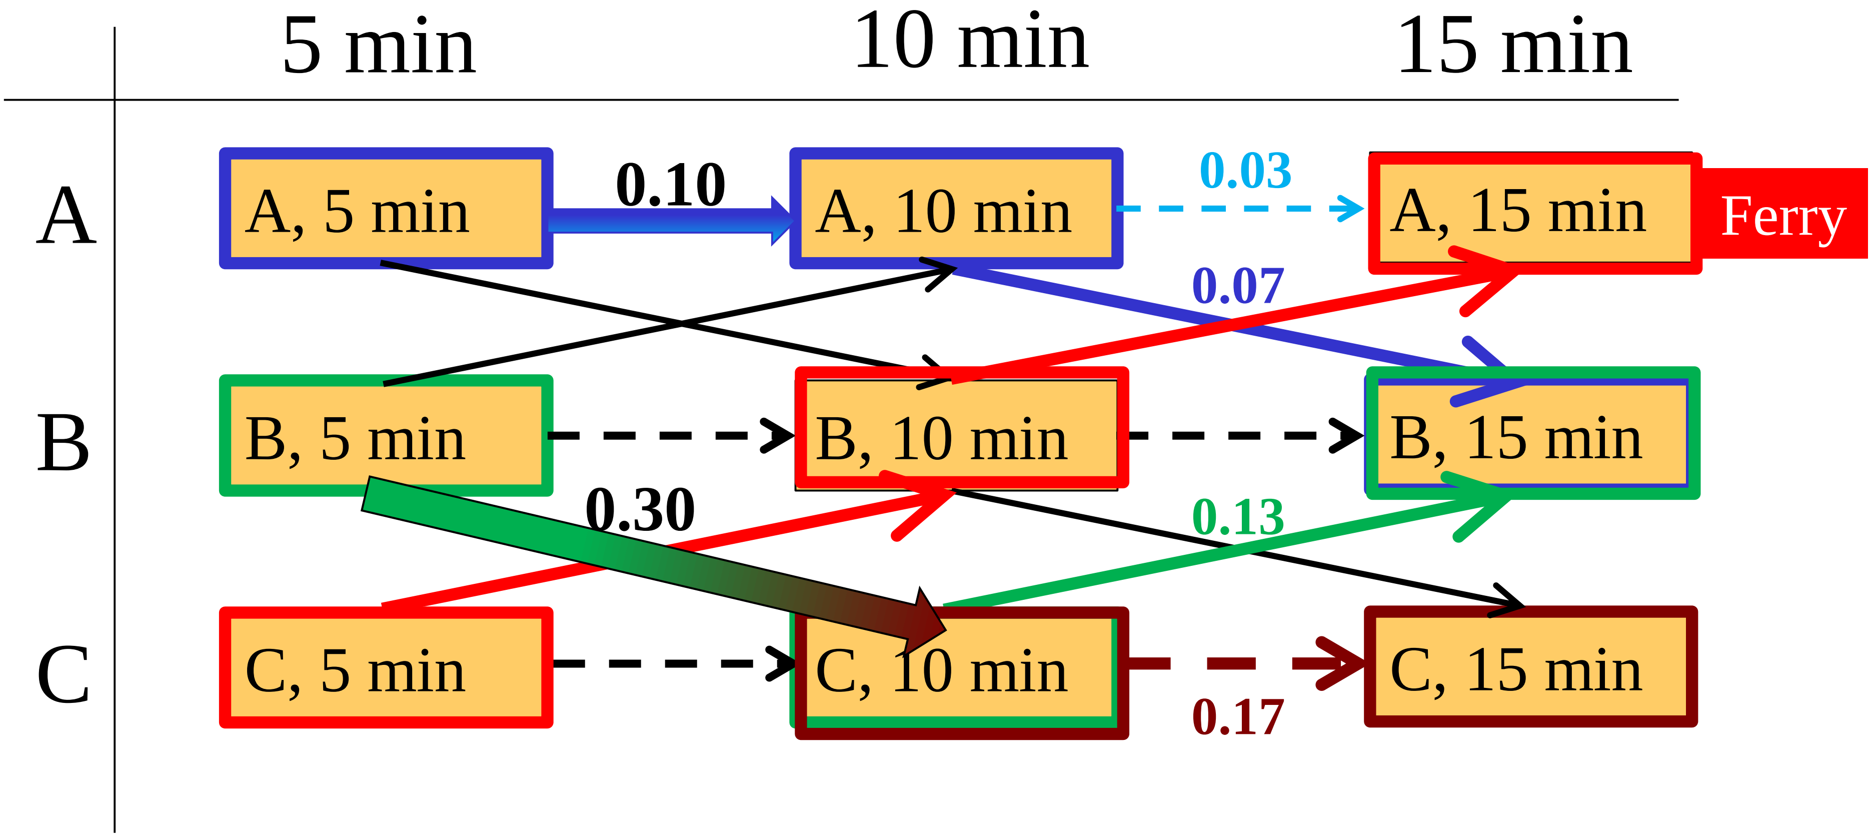

3 of 5

U.S. Coast Guard protects the Staten Island Ferry: I feel safe!

By shortysmom | Posted September 8, 2013 | Staten Island, New York

About this iReport

- Not vetted for CNN

Posted September 8, 2013 by
shortysmom | Follow
Location
Staten Island, New York

I ride the Staten Island Ferry on a daily basis to and from work. We ferry riders have our own personal protectors in the form of the U.S. Coast Guard. The

Outline: “Security Games” Research (2007-)

Air travel

Urban Crime

2007

2015

TRUSTS: Frequent adversary interaction games
Patrols Against Fare Evaders

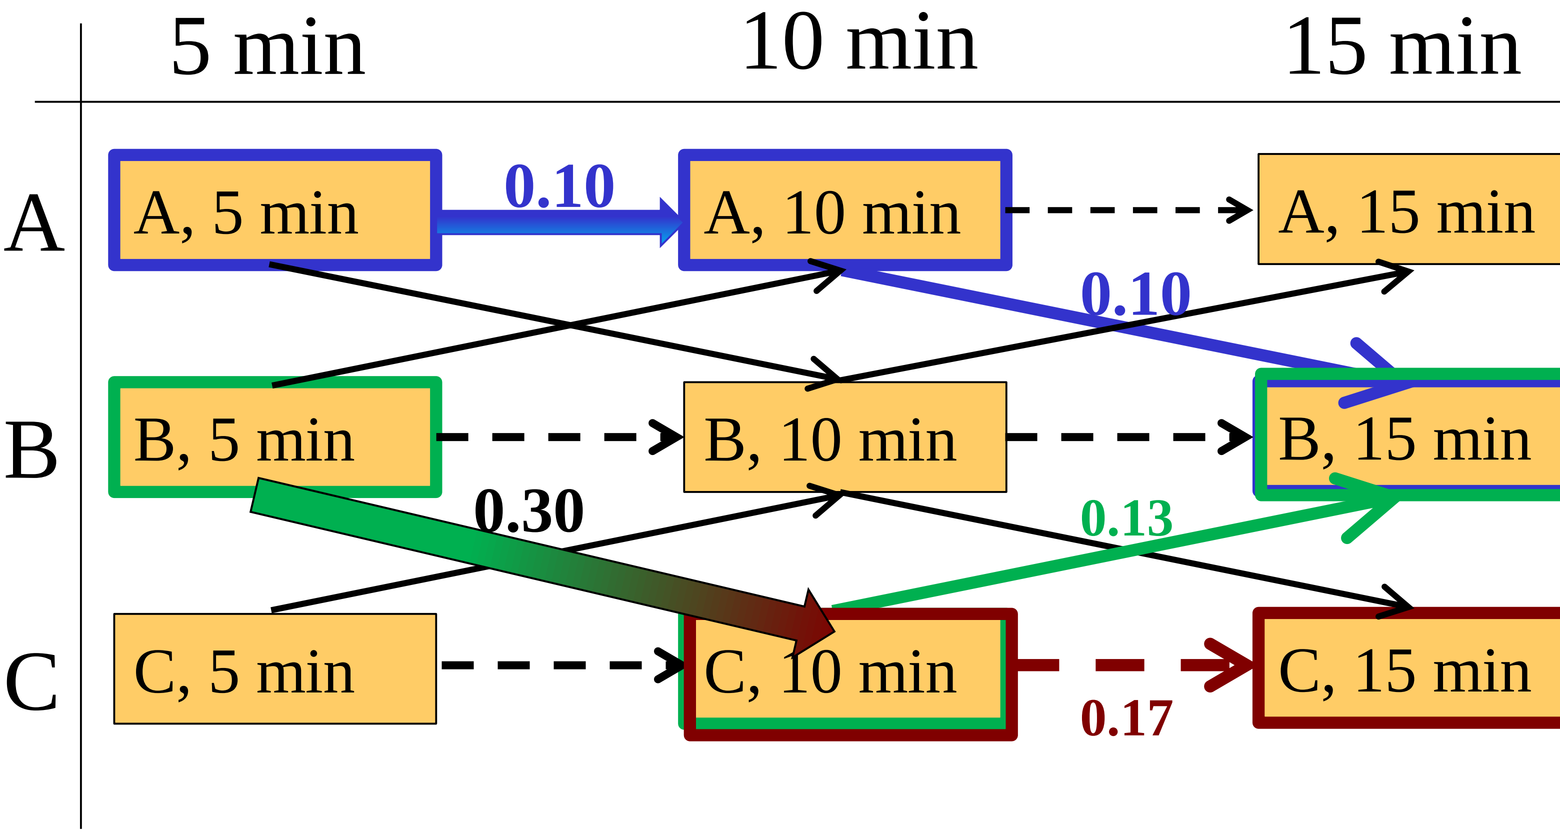

TRUSTS: Patrols Against Fare Evaders
Uncertainty in Defender Action Execution

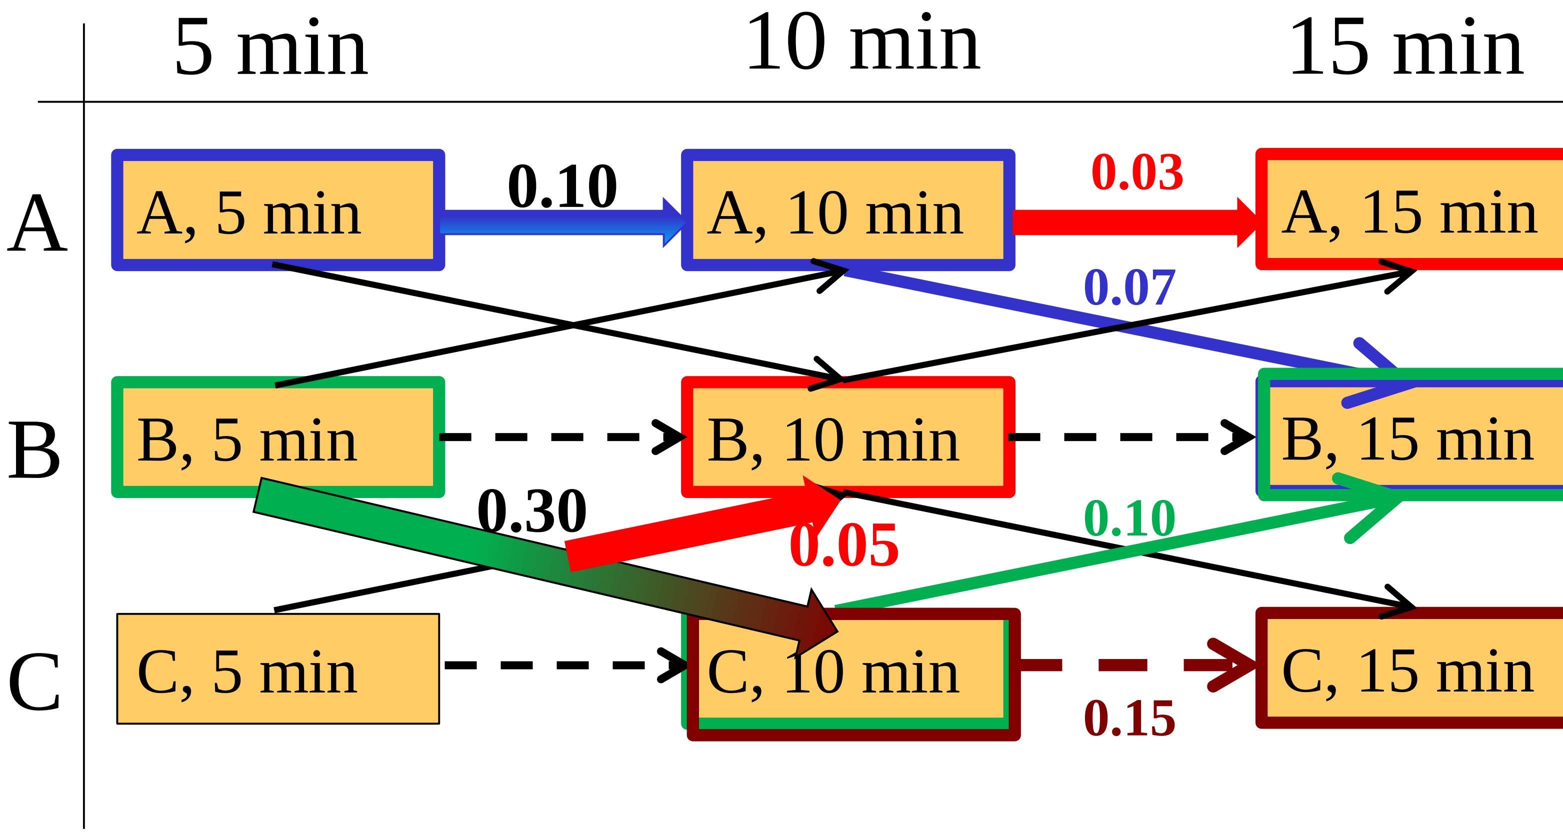

TRUSTS: Patrols Against Fare Evaders
Uncertainty in Defender Action Execution

- Markov Decision Problems *in Security games*

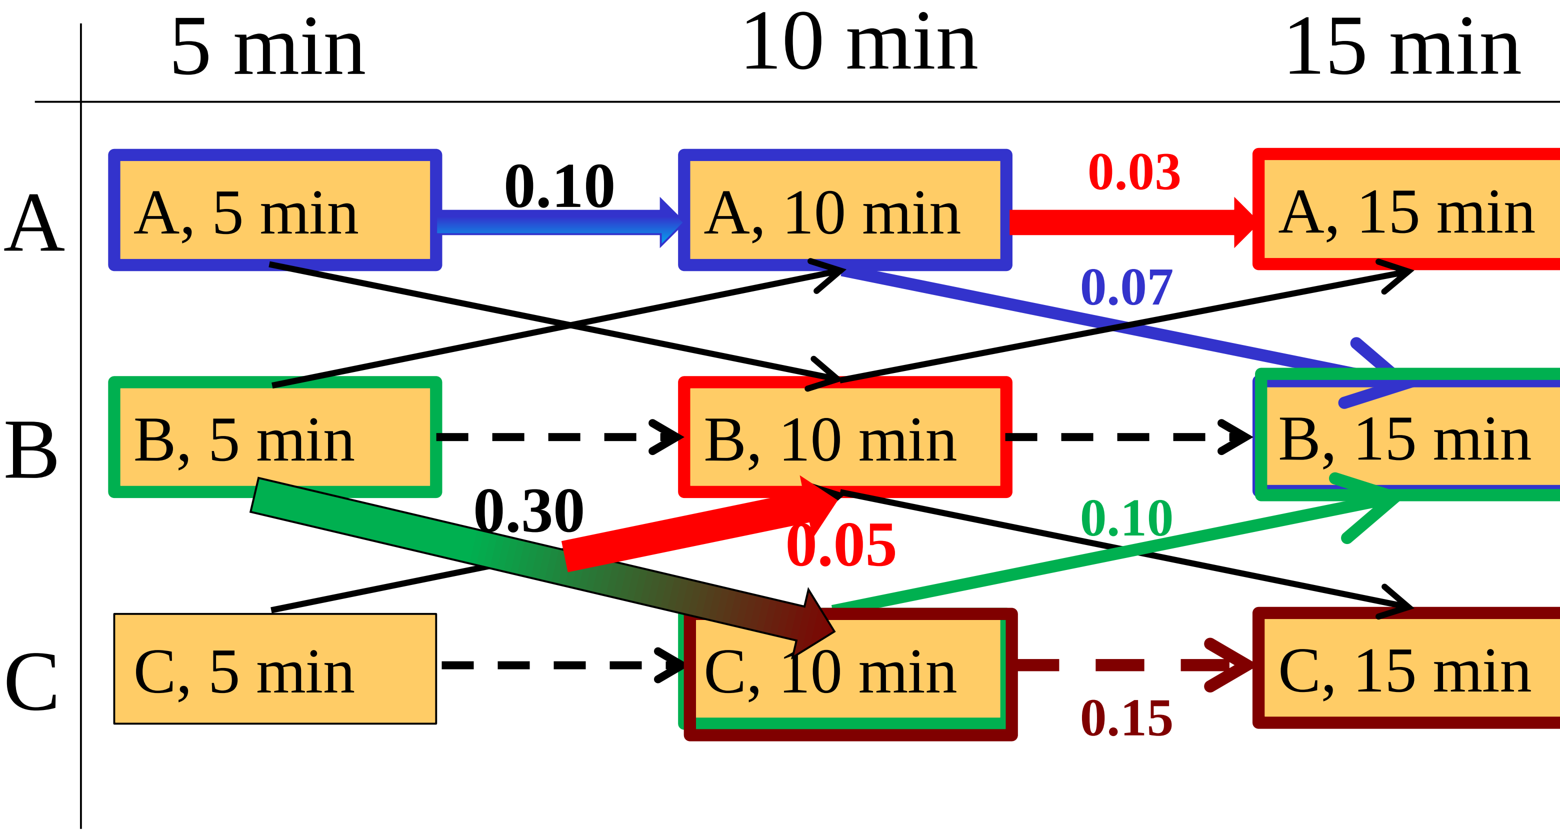

Uncertainty Space Algorithms: Bayesian and Robust Approaches

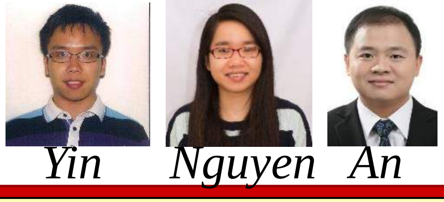

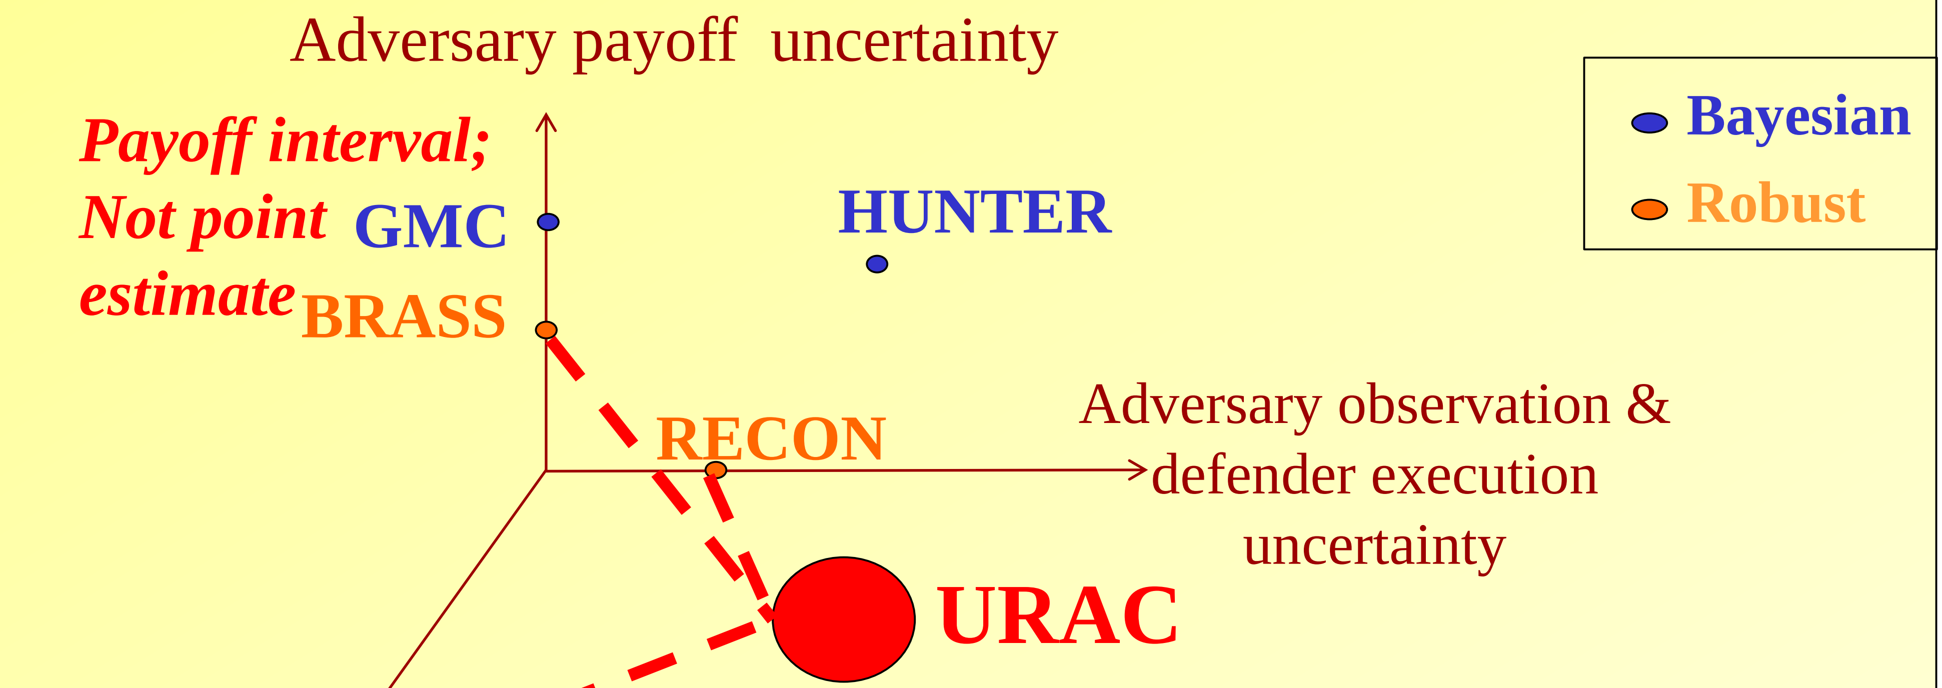

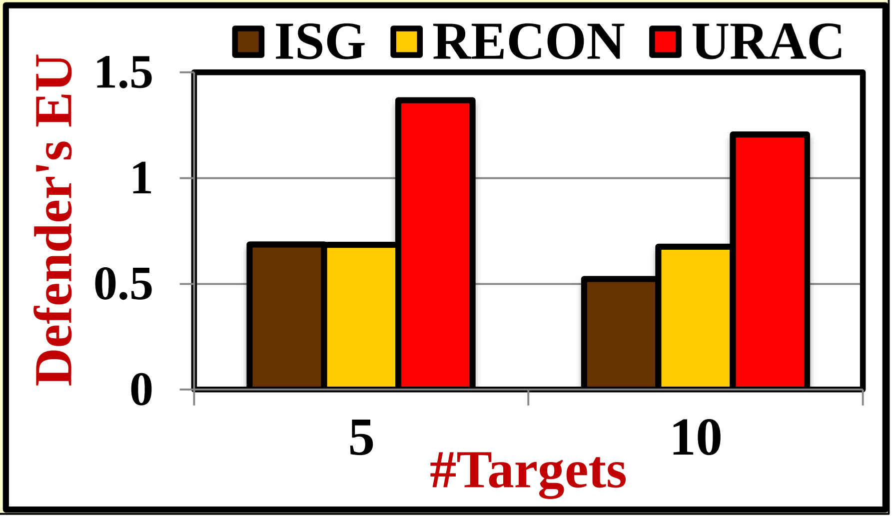

Outline: Security Games Research (2007-)

Air travel

Ports

Trains

Fisheries

Wildlife

Urban Crime

2007 2011 2013 2014 2014 2015

Protecting Forests, Fish and Wildlife: Green Security Games

Yang

Ford

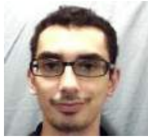

Brown

Wildlife Protection

Forest Protection

Fishery Protection

Murchison Falls National Park, Uganda




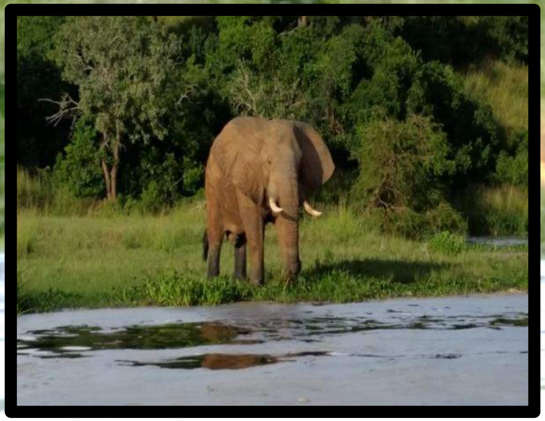

Green Security Games: Repeated Stackelberg Game

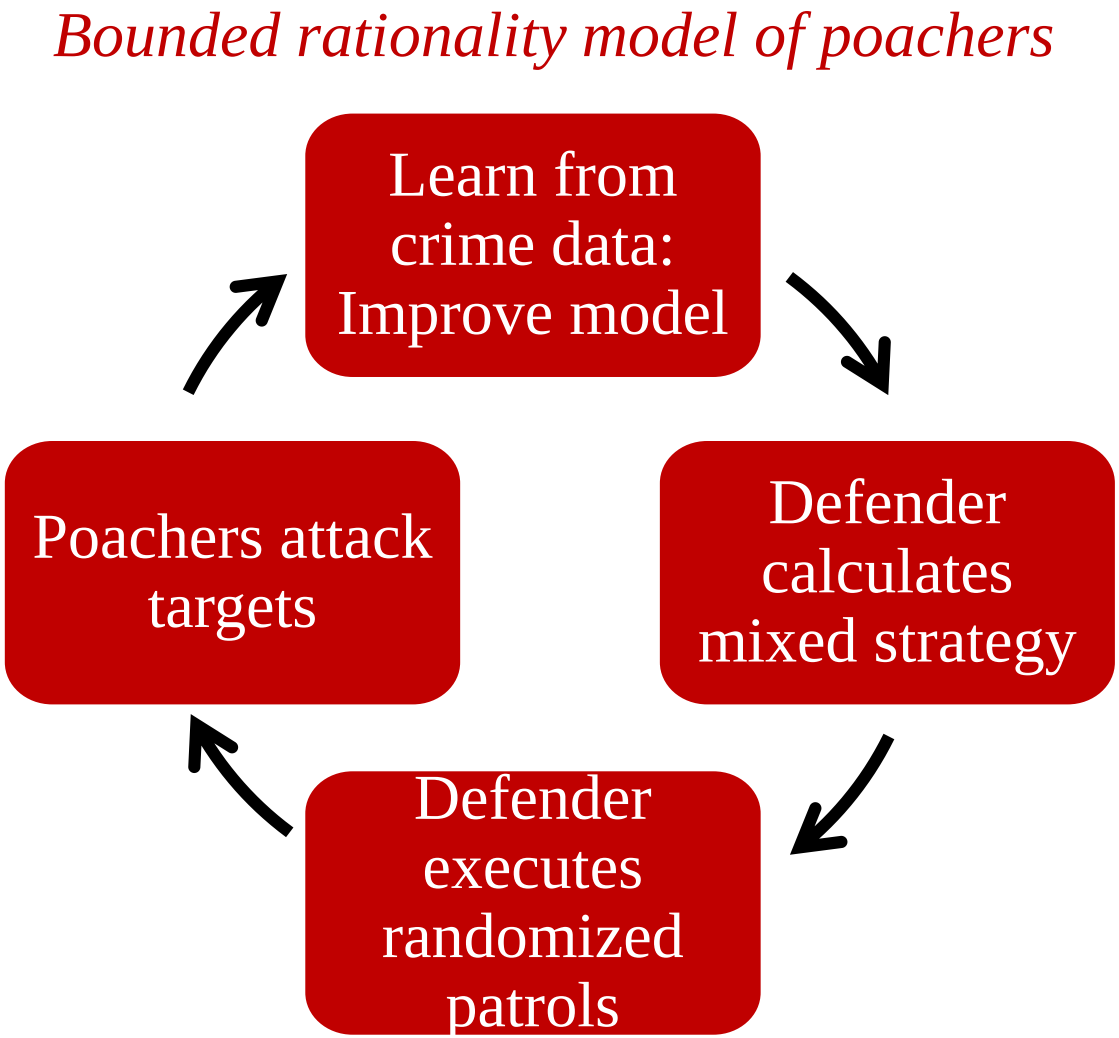

Uncertainty in Adversary Decision: Bounded Rationality
Human Subjects as Poachers

Kar

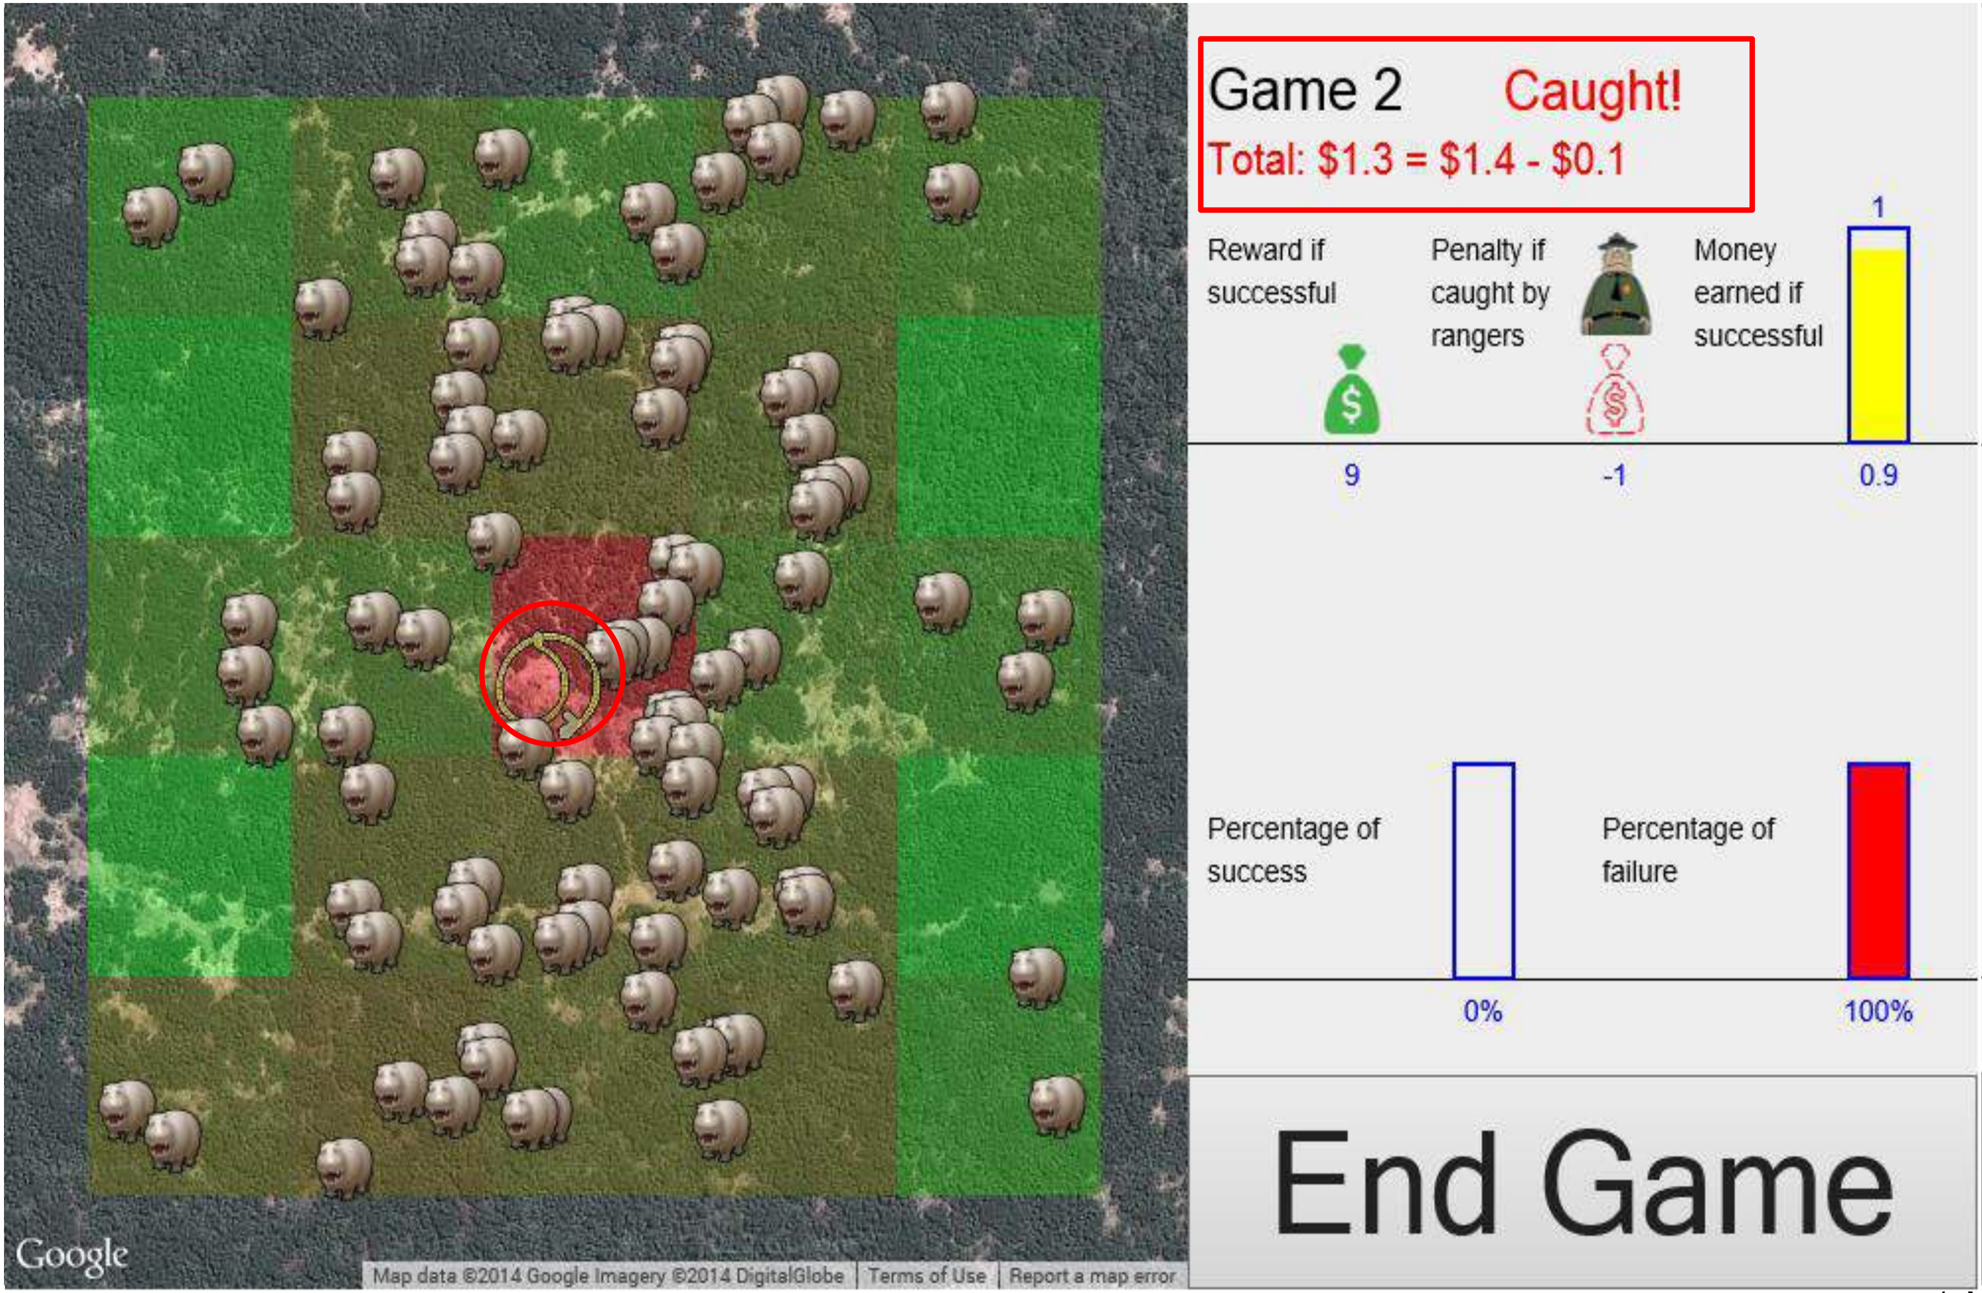

Lesson 1: Quantal Response [2011]: Models of Bounded Rationality

Pita

Yang

Perfect:

$$EU^{adversary}(j) = Capture\text{Prob} \times \text{Penalty} + (1 - Capture\,\text{Prob}) \times \text{Reward}$$

Quantal Response(QR) ***[McFadden 73]:*** Stochastic Choice, Better Choice More likely

Adversary's probability of **choosing target j** $= \dfrac{e^{\lambda \cdot (EU^{adversary}(x,j))}}{\sum_{j'=1}^{T} e^{\lambda \cdot (EU^{adversary}(x,j'))}}$

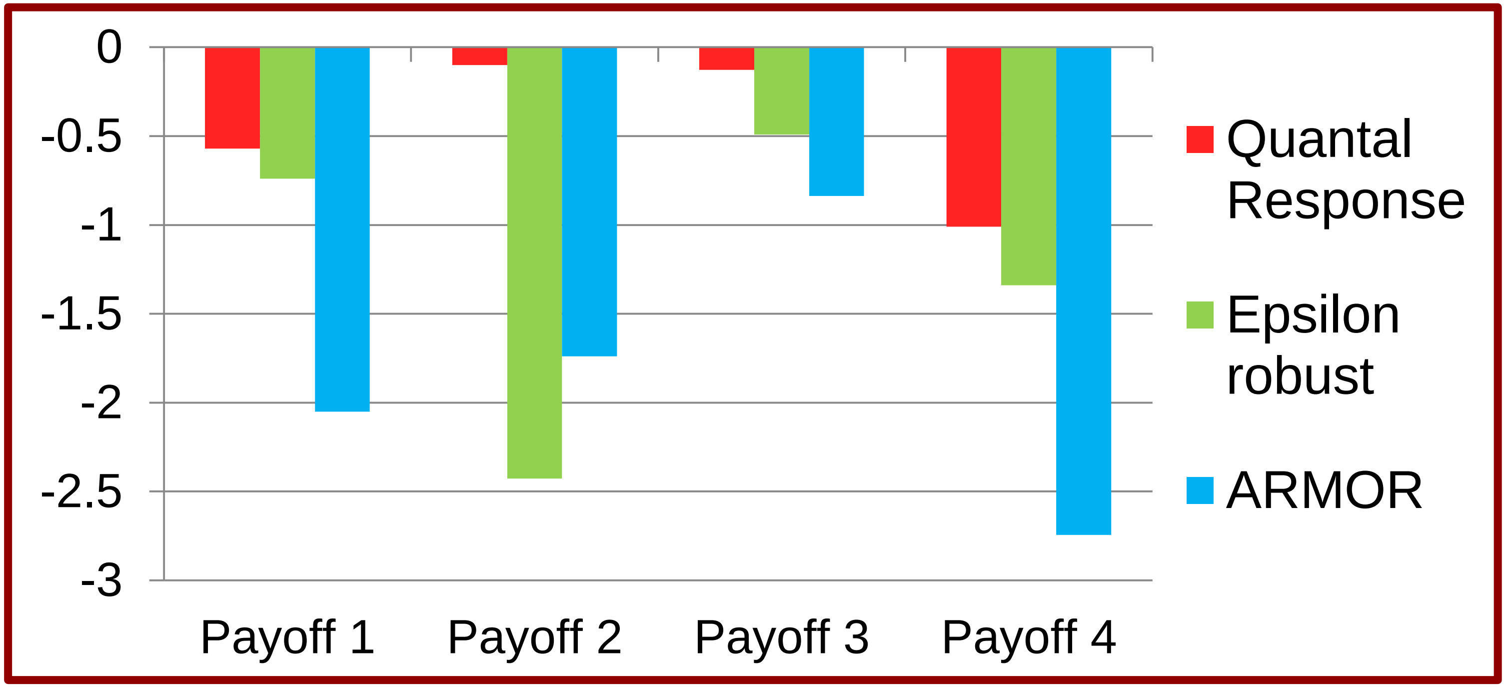

Lesson 2: Subjective Utility Quantal Response Models of Bounded Rationality

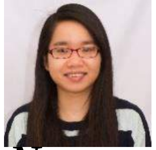
Nguyen

Subjective Utility Quantal Response(SUQR) [Nguyen 13]:

$$SEU^{adversary}(j) = w_1 \times Capture\,\mathrm{Prob} + w_2 \times \mathrm{Reward} + w_3 \times \mathrm{Penalty}$$

Adversary's probability of **choosing target j** $= \dfrac{e^{SEU^{adversary}(x,j)}}{\sum_{j'=1}^{M} e^{SEU^{adversary}(x,j')}}$

UWA data from Uganda: Year 2012 Predicting poaching

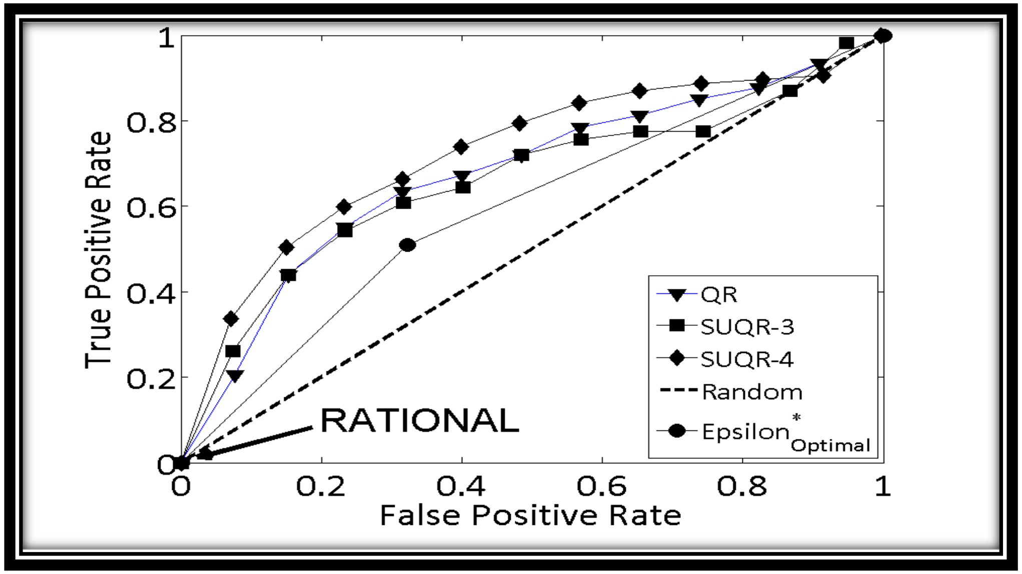

Green Security Games[2015]
Repeated Stackelberg Game

Bayesian SUQR:
Heterogeneous Poachers

$$\frac{e^{SEU_i(W1,W2,W3)}}{\sum_k e^{SEU_k(W1,W2,W3)}}$$

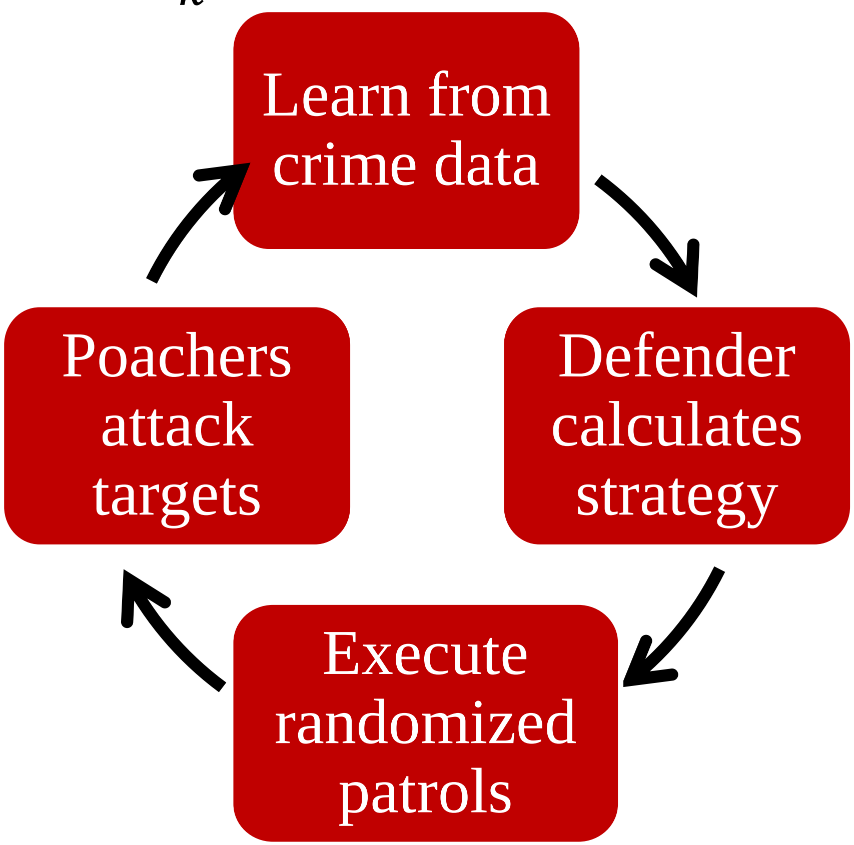

Logitudinal study on AMT:
35 weeks, 40 human subjects
10,000 emails!

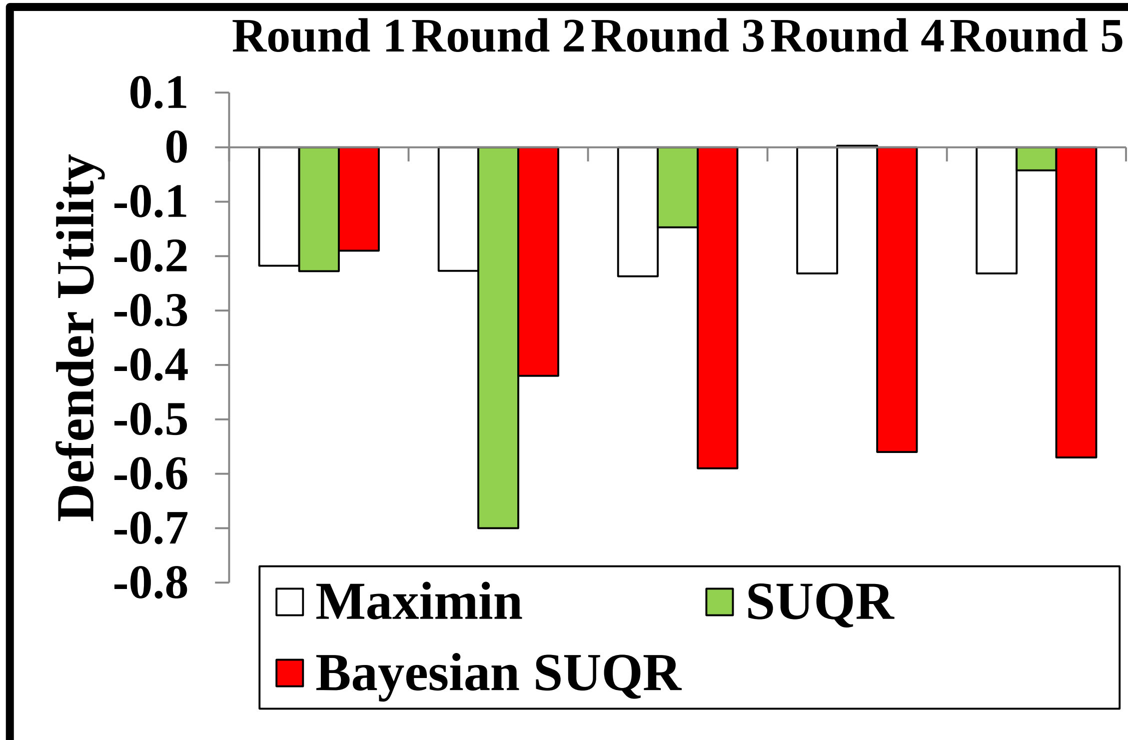

Lesson 3: SHARP and Repeated Stackelberg Games
Incorporate Past Success/Failure in SUQR


Kar

Bayesian SUQR:
Heterogeneous Poachers

$$\frac{e^{SEU_i(W1,W2,W3)}}{\sum_k e^{SEU_k(W1,W2,W3)}}$$

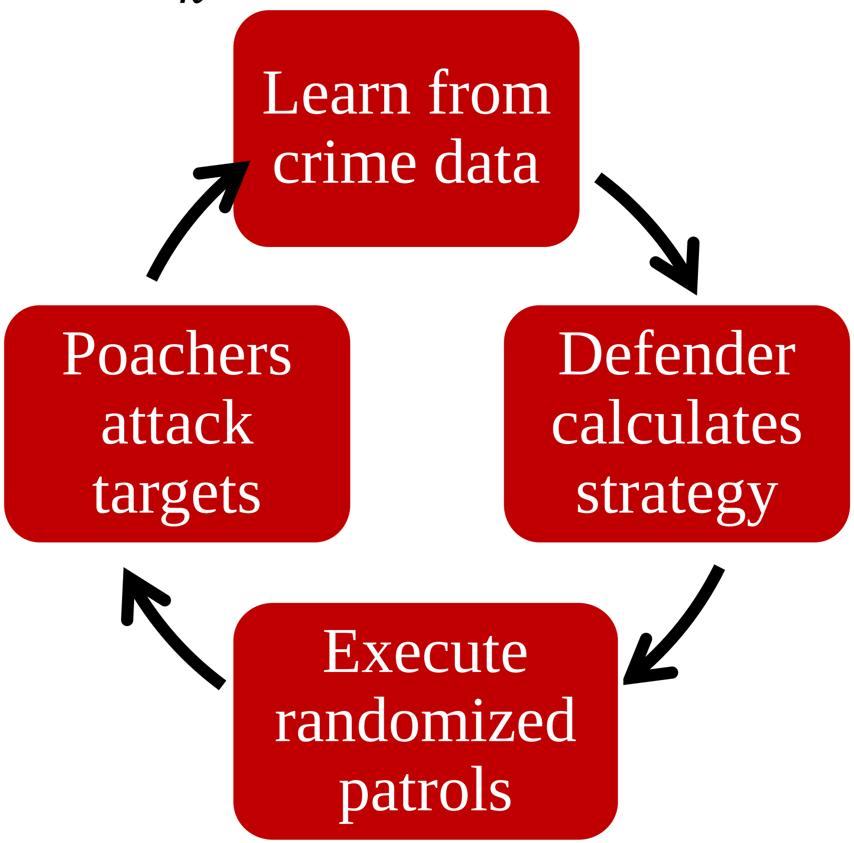

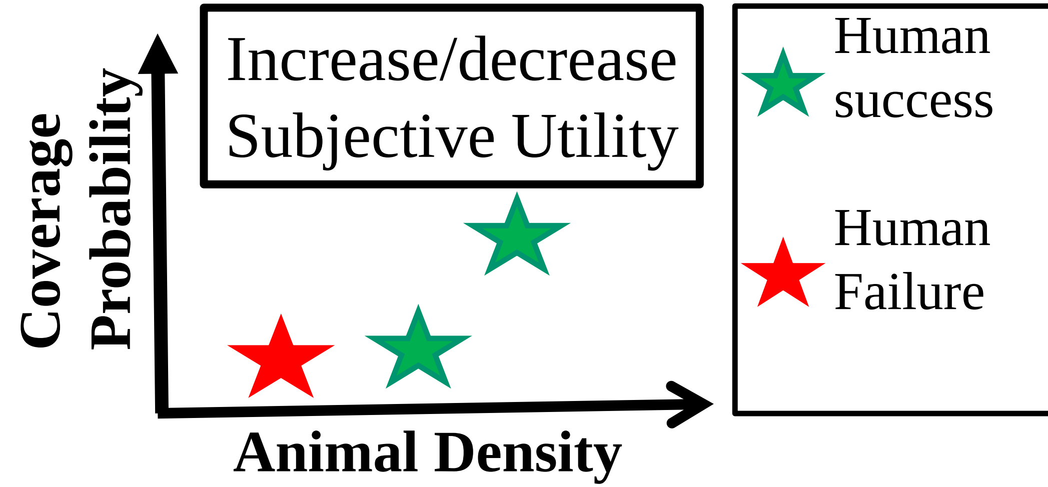

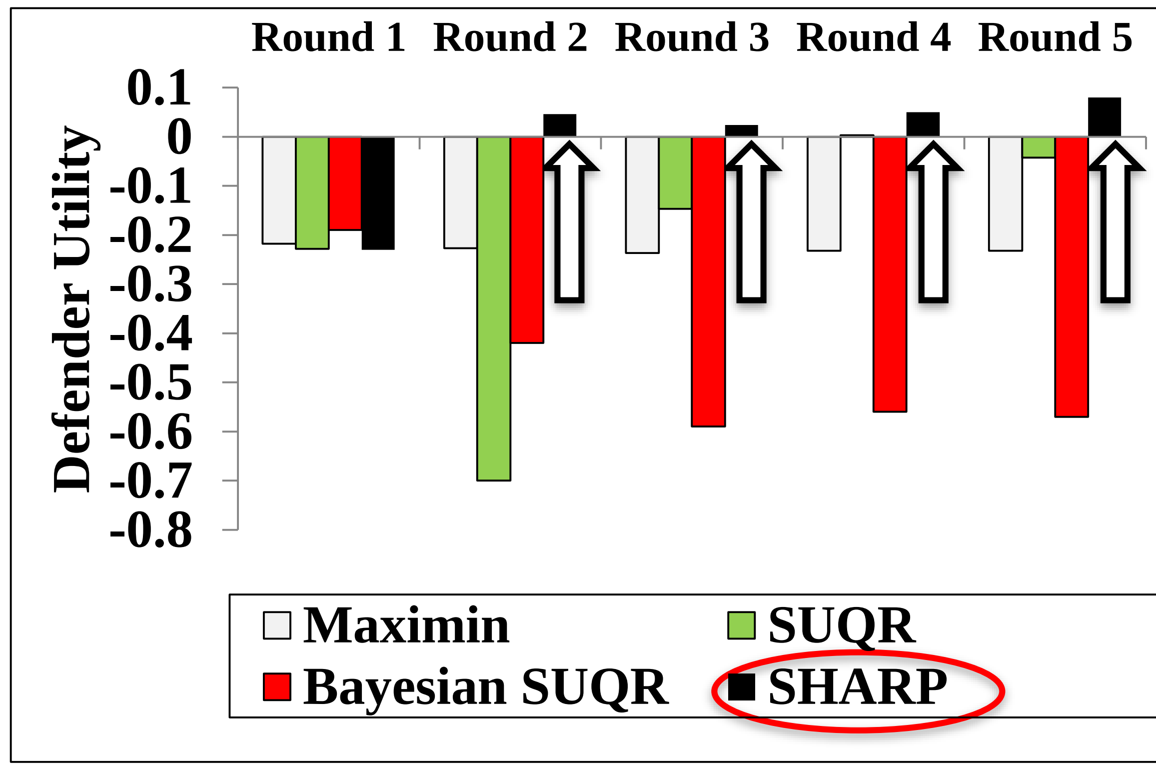

PAWS: Protection Assistant for Wildlife Security Trials in Uganda and Malaysia: [2014]

- **Important lesson: Geography!**


Uganda

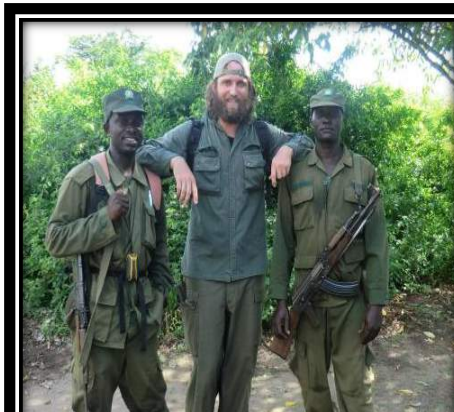
Andrew Lemieux


Malaysia


Panthera


Malaysia

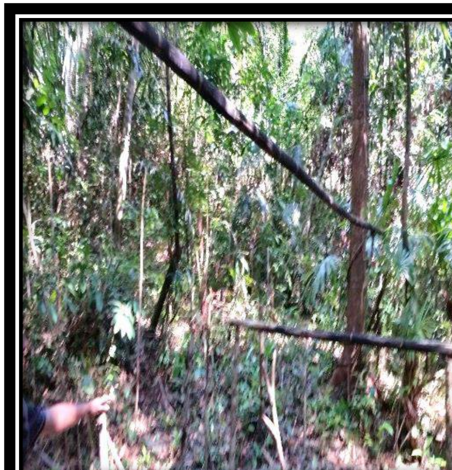

PAWS Next Steps 2015: Southeast Asia (with Panthera and WWF)

PAWS Version 2: Features

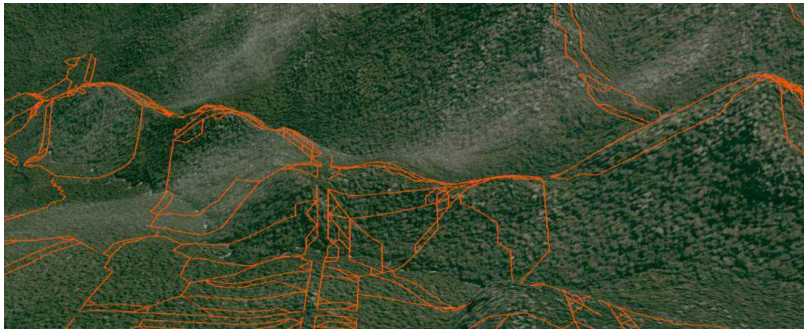

- Street map
  - *Ridgelines, rivers/streams*
- Species Distribution Models (SDMs)
  - *From data points to distribution map*

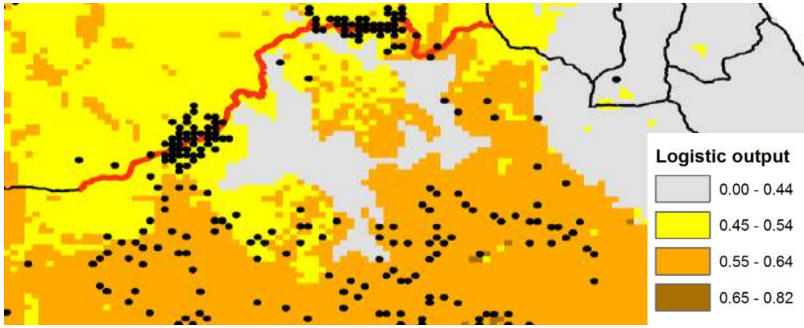

Indonesia

Malaysia

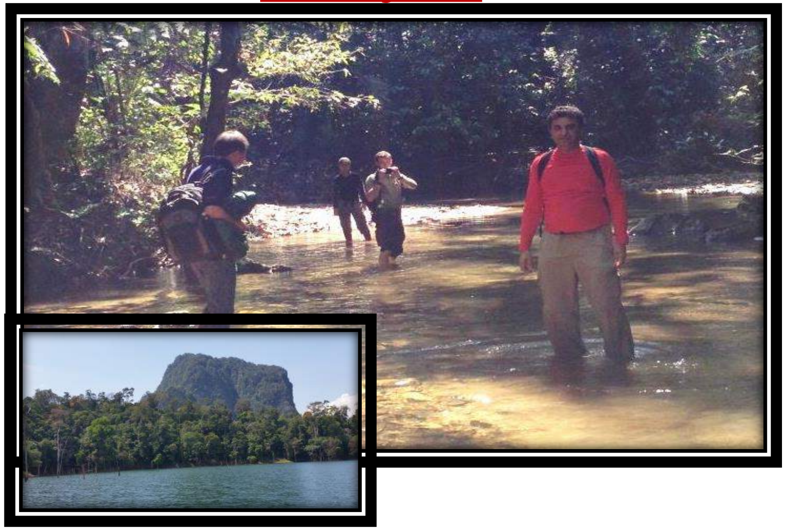

Opportunistic Crime Security Game[2015]
Urban Crime


Zhang


Sinha

- Crime prediction: uses past defender (police) allocation data

University of Southern California

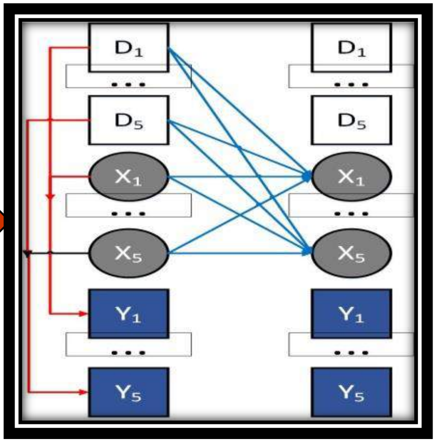

Results of learning

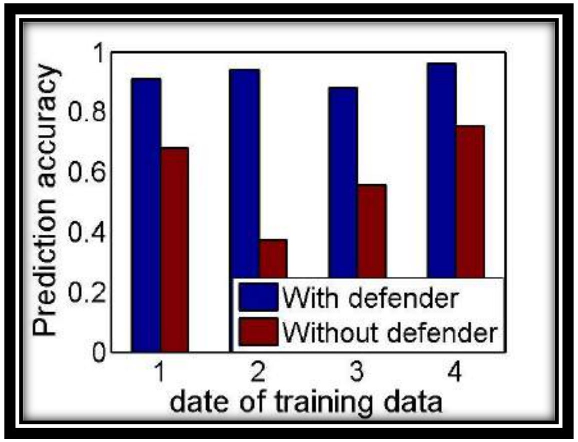

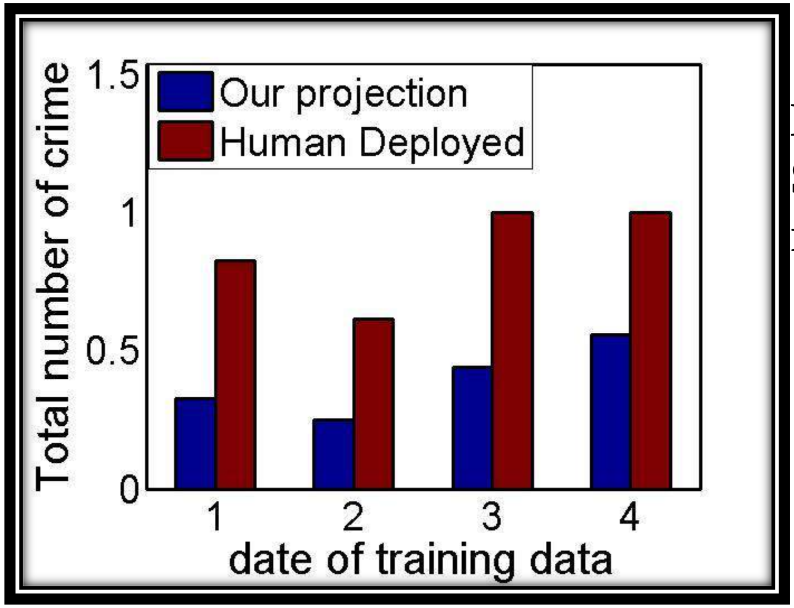

Best Simulation results

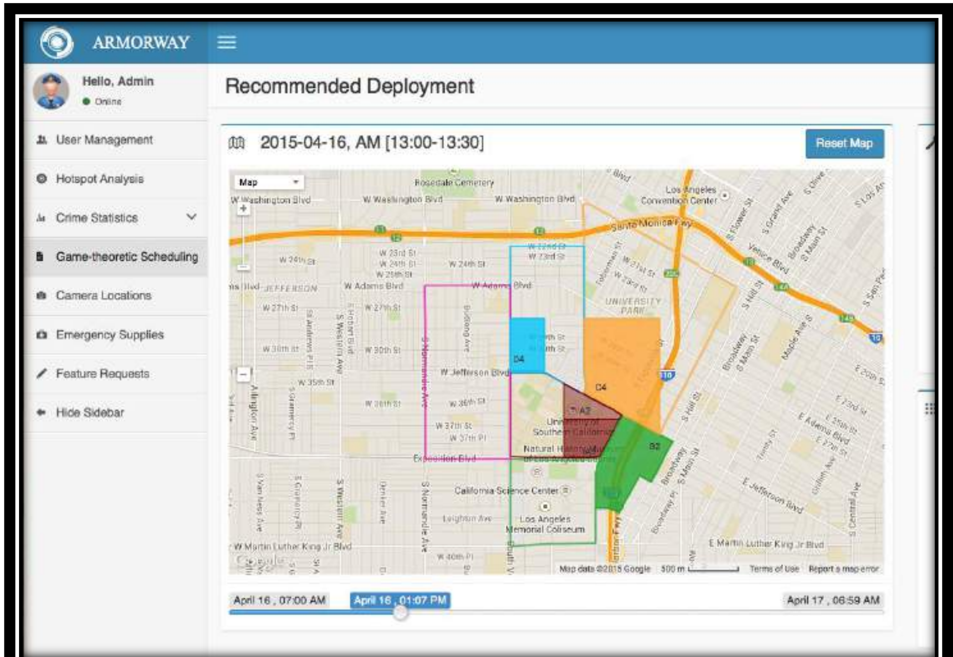

Security Resource Optimization:
Evaluating *Deployed Security Systems* Not Easy

- Game theory: Improvement over previous approaches
  - *Previous: Human schedulers or "simple random"*

Lab Evaluation	**Field Evaluation: Patrol quality Unpredictable? Cover?**	**Field Evaluation: Tests against adversaries**
Simulated adversary	Compare real schedules	"Mock attackers"
Human subject adversaries	Scheduling competition	Capture rates of real adversaries
	Expert evaluation	

Why Does Game Theory Perform Better? Weaknesses of Previous Methods

- Human schedulers:
  - *Predictable patterns, e.g., US Coast Guard*
  - *Scheduling effort & cognitive burden*

- Simple random (e.g., dice roll):
  - *Wrong weights/coverage, e.g. officers to sparsely crowded terminals*
  - *No adversary reactions*

- Multiple deployments over multiple years: without us forcing them

Field Evaluation of Schedule Quality:

Improved Patrol Unpredictability & Coverage for Less Effort

PROTECT (Coast Guard): *350% increase defender expected utility*

Patrols Before PROTECT: Boston

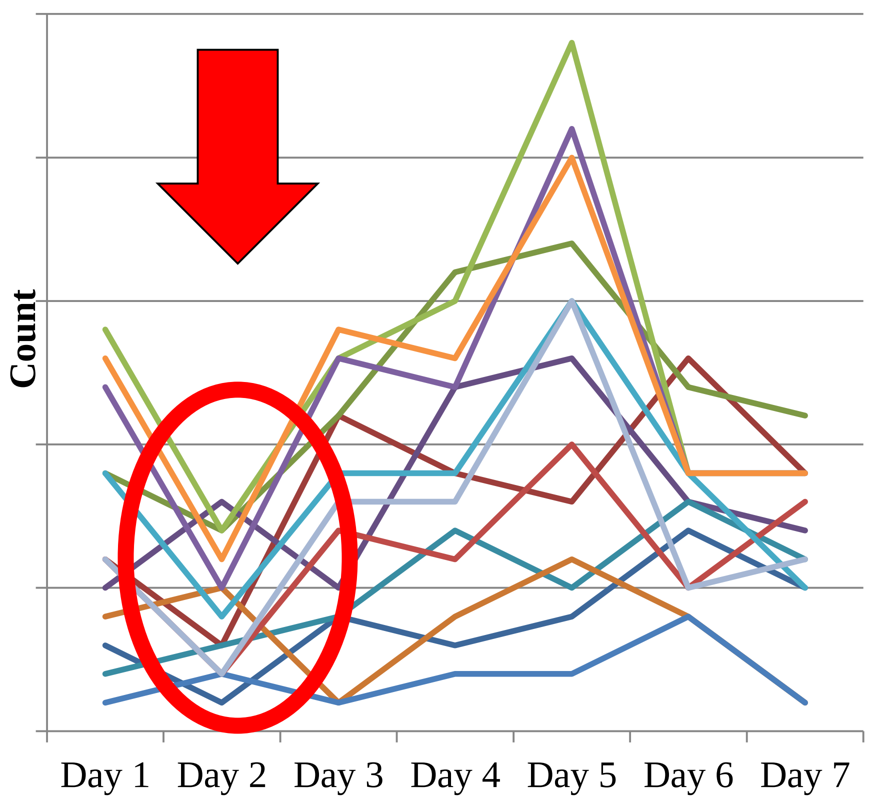

Patrols After PROTECT: Boston

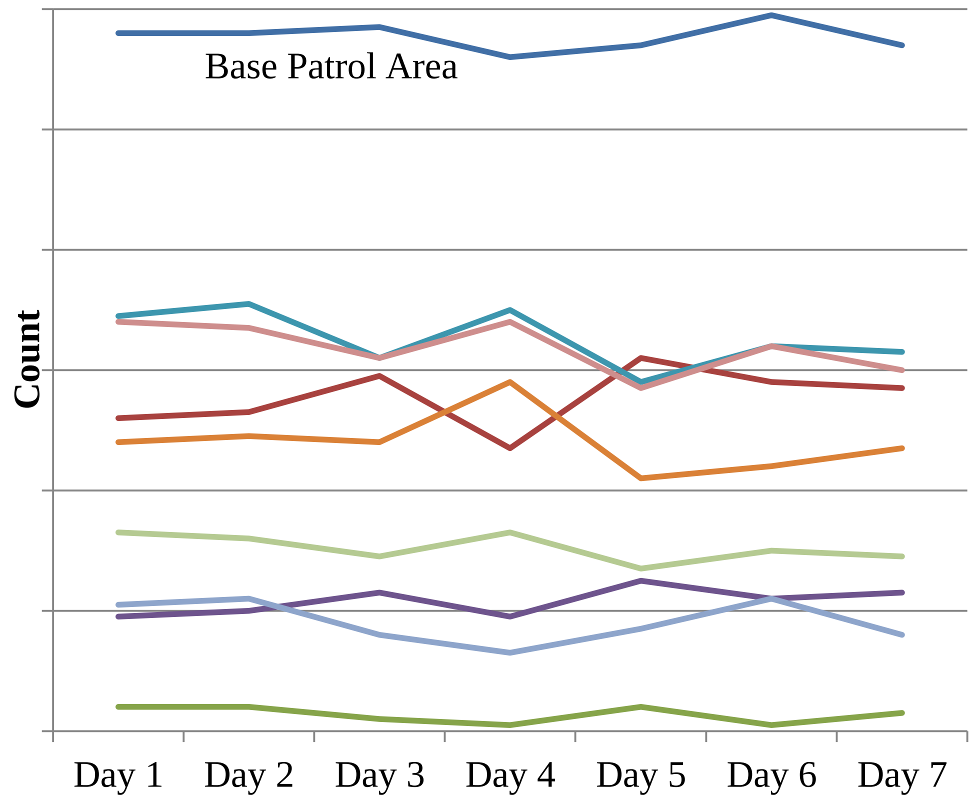

Field Evaluation of Schedule Quality:

Improved Patrol Unpredictability & Coverage for Less Effort

FAMS: IRIS Outperformed expert human over six months

Report:GAO-09-903T

Trains: TRUSTS outperformed expert humans schedule 90 officers on LA trains

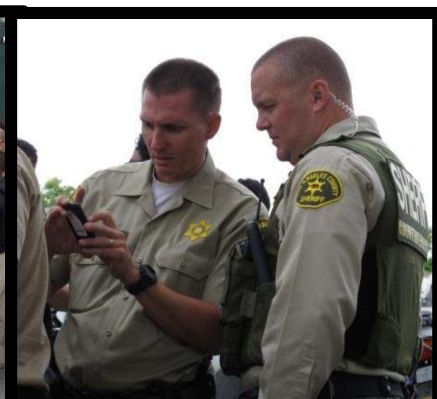

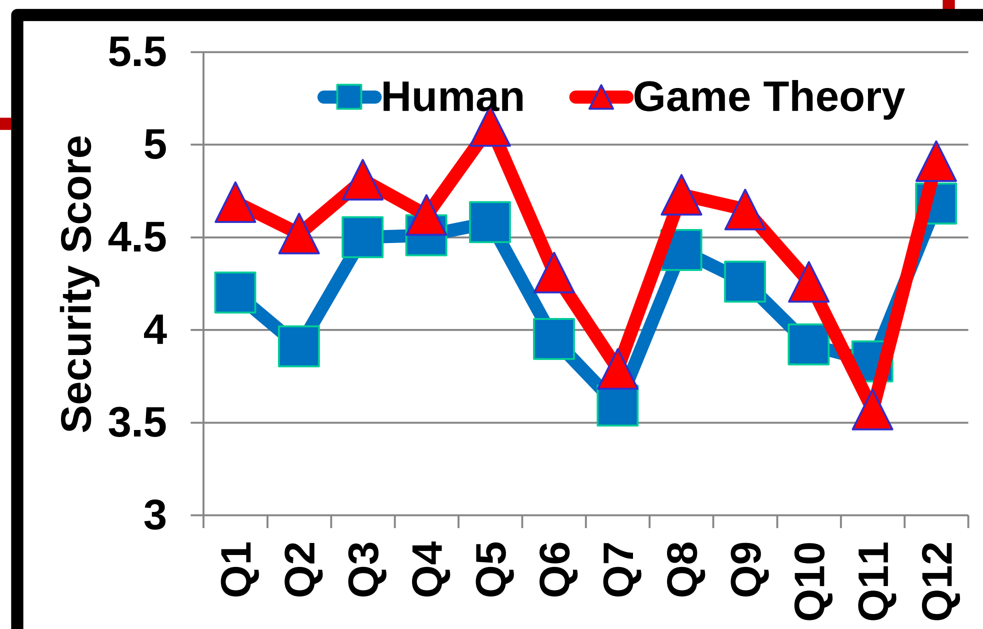

Field Test Against Adversaries: Mock Attackers Example from PROTECT

- "Mock attacker" team deployed in Boston
  - *Comparing PRE- to POST-PROTECT: "deterrence" improved*

- Additional real-world indicators from Boston:
  - *Boston boaters questions:*
    - "..has the Coast Guard recently acquired more boats"
  - *POST-PROTECT: Actual reports of illegal activity*

Field Tests Against Adversaries
Computational Game Theory in the Field

Controlled

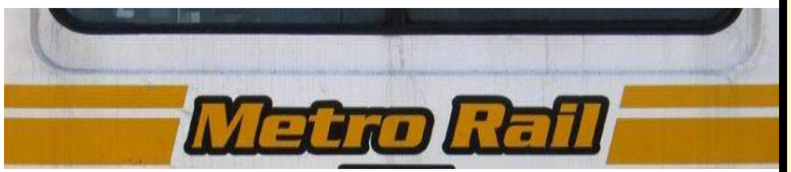

- *Game theory vs Random*
- *21 days of patrol*
- *Identical conditions*
- *Random + Human*

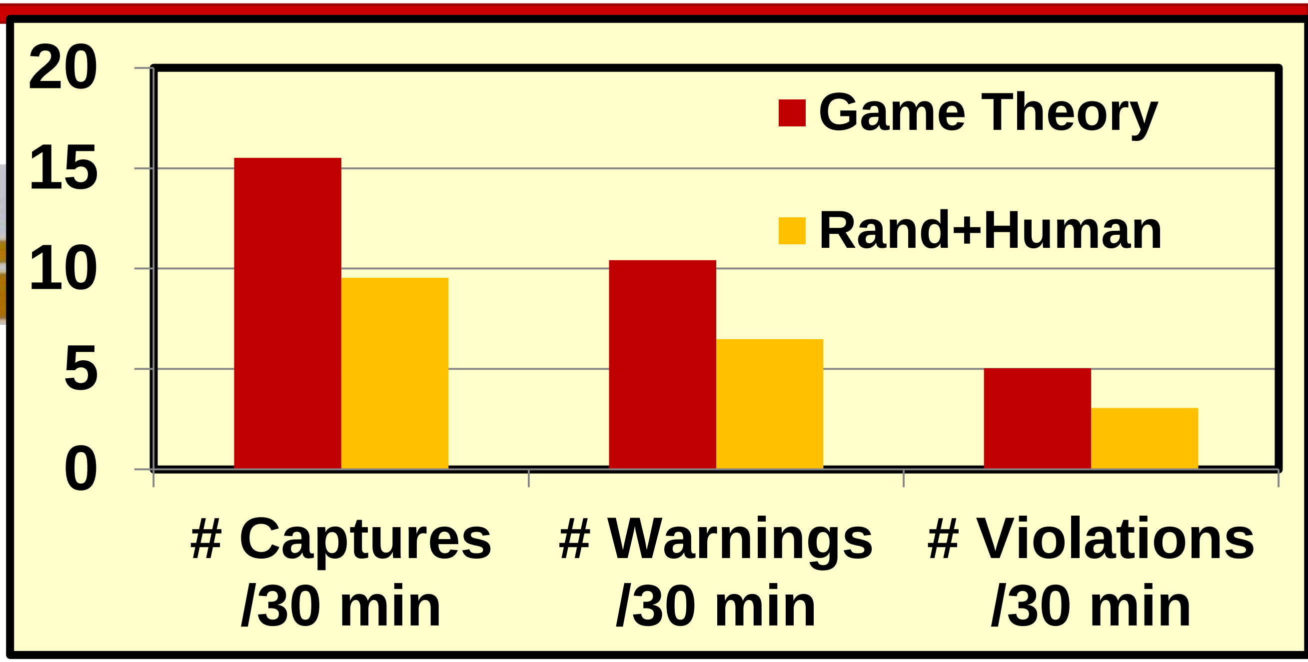

Not controlled

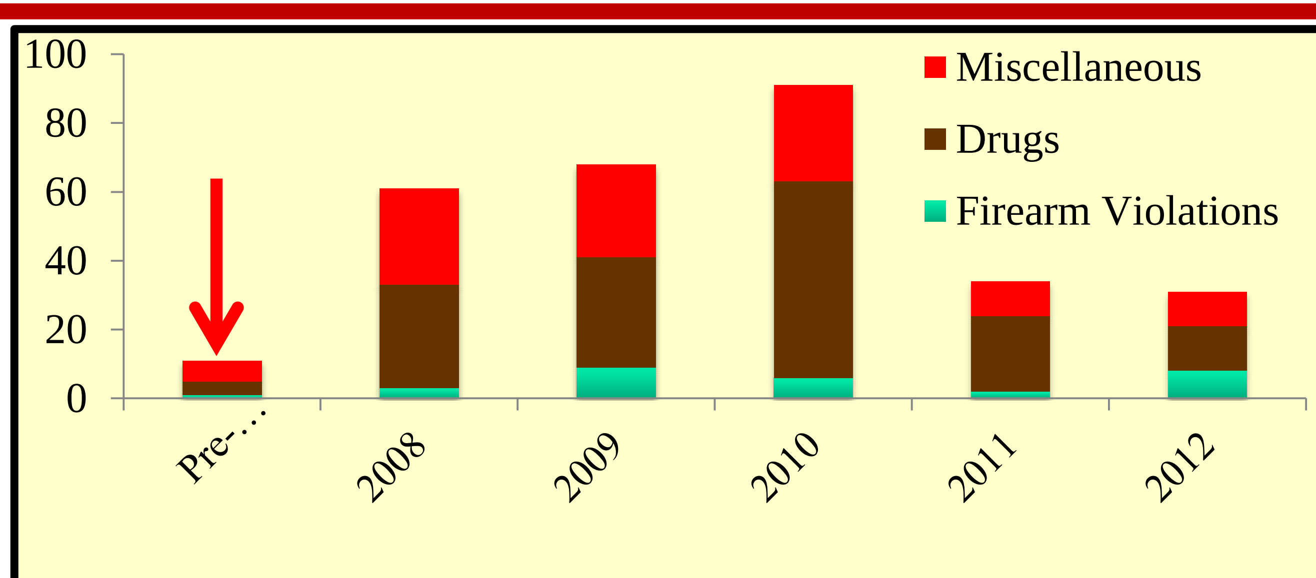

Expert Evaluation
Example from ARMOR, IRIS & PROTECT

June 2013: Meritorious Team Commendation from Commandant (US Coast Guard)

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

September 2011: Certificate of Appreciation (Federal Air Marshals)

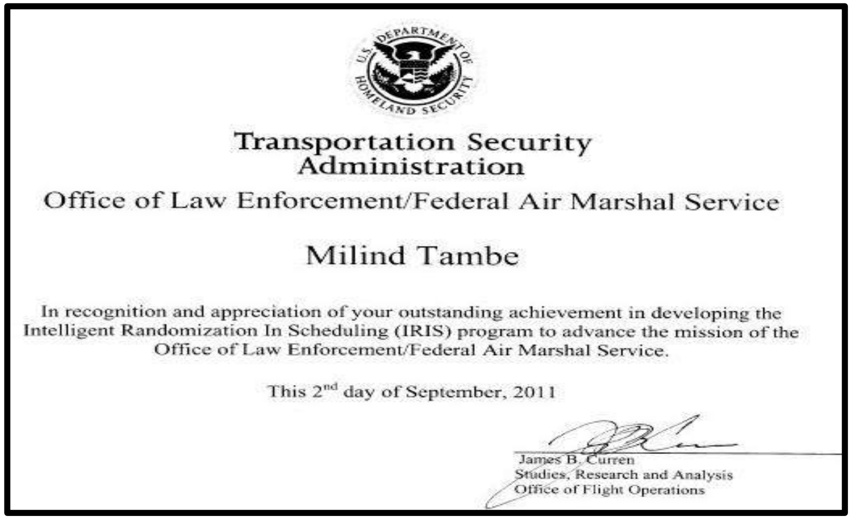

Transportation Security Administration

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

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

Key Lessons: Security Games

Decision aids based on computational game theory in daily use

- Optimize limited security resources against adversaries

Applications yield research challenges: Science of security games

- Scale-up: Incremental strategy generation & Marginals
- Uncertainty: Integrate MDPs, Robustness
- Human behavior models: Quantal response

Current applications: Global, interdisciplinary challenges

- Green security games: criminology, computation, conservation

Global Efforts on Security Games: Yet Just the Beginning…

Thank you to sponsors:

CREATE HOMELAND SECURITY CENTER

Transportation Security Administration

Army Research Office

THANK YOU

tambe@usc.edu

http://teamcore.usc.edu/security

Lab Evaluation via Simulations: Example from IRIS (FAMS)