Security Games:

Key Algorithmic Principles, Deployed Systems, Research Challenges

Milind Tambe

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

Other collaborators:

Shaddin Dughmi (USC), Richard John (USC), David Kempe (USC), Nicole Sintov (USC), Nicholas Weller (USC) ….&

Craig Boutilier (Google), Vince Conitzer (Duke), Sarit Kraus (BIU, Israel), E Lam (Panthera), Andrew Lemieux (NCSR), Arnault Lyet (WWF), Kevin Leyton-Brown (UBC), Fernando Ordonez (U Chile), M. Pechoucek (CTU, Czech R), Rob Pickles (Panthera), Andy Plumptre (WCS), Ariel Procaccia (CMU), Tuomas Sandholm (CMU), Peter Stone (UT Austin), Y. Vorobeychik (Vanderbilt), ….&

Collaborators from the US Coast Guard, Transportation Security Administration, LA Sheriff's Dept, Uganda Wildlife Authority, …&

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

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
Game theory+Optimization+Uncertainty+Learning+…

- *Massive-scale games*
- *Reason with uncertainty*

- *Learn adversary behavior from data*
- *Repeated games*
- *+ Conservation biology, criminology*

Infrastructure Security Games


Coast Guard

Coast Guard


LAX

TSA

Argentina airport


USC


LA Sheriff

Green Security Games

Coast Guard


Panthera/WWF

IBM

Opportunistic Crime Security Games

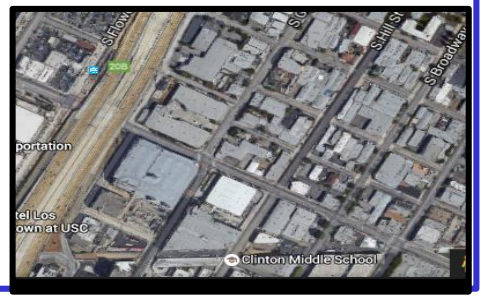

Cyber Security Games

Global Presence of Security Games Efforts

MILIND TAMBE'S ARMOR AND ITS MANY ITERATIONS ARE USED AROUND THE WORLD TO PROTECT AGAINST TERRORISM, POACHERS, ILLEGAL FISHING AND OTHER THREATS.

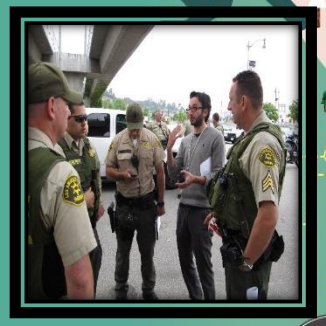

DEPLOYED

Ports — PROTECT
PROTECT intelligently randomizes U.S. Coast Guard patrols to optimize scarce resources to secure crowded piers, bridges and ferry terminals.
PROTECT is employed at:
- *Port of New York and New Jersey*
- *Port of Boston*
- *Port of Houston*
- *Port of Los Angeles-Long Beach*

Staten Island Ferry — PROTECT
PROTECT provides protection to the Staten Island Ferry, which carries up to 4,000 passengers at peak times.

Los Angeles International Airport — ARMOR
ARMOR intelligently randomizes schedules of checkpoints along the five roads that lead into the airport.

U.S. Air Traffic — IRIS
As part of its multipronged strategy to prevent attacks, the Transportation Security Administration (TSA) has since 2009 deployed Milind Tambe's IRIS system, which intelligently randomizes federal air marshals' flight schedules to make their air patrols unpredictable to would-be malefactors.

SUCCESSFULLY TESTED

Gulf of Mexico (Near Corpus Christi, Texas) — ARMOR-FISH
ARMOR-FISH intelligently randomized schedules for U.S. Coast Guard aerial patrols to thwart the illegal fishing of decimated shark and red snapper populations. (2014)

Los Angeles Metro — TRUSTS
The Los Angeles Sheriff's Department, which LA Metro subcontracts for security, employed TRUSTS to intelligently randomize patrol schedules to stop fare evasion. The Sheriff's Department later ran preliminary experiments to ascertain effectiveness in deploying scarce police personnel to deter crime and terrorism on LA Metro. (2011–2013)

Uganda — PAWS
Ugandan rangers tested PAWS at Queen Elizabeth National Park to intelligently randomize patrols to prevent the slaughter of animals, including Cape buffalo, waterbuck and giant forest hogs, which are served up locally and exported as "bush meat." (2014)

Malaysia — PAWS
Panthera, an NGO that is committed to ensuring the survival of tigers and other wild cats, in conjunction with the nonprofit group Rimba, began testing PAWS in forests in northeastern Malaysia, to evaluate its ability to generate effective patrols in the challenging, hilly terrain. (2014)

POSSIBLE FUTURE TEST SITES

Vietnam, Cambodia, Bangladesh, Indonesia — PAWS
authorities could employ ntly randomize police patrols vers, a big problem in this island nation.

Madagascar — PAWS
Milind Tambe, working with Meredith Gore, an associate professor of conservation social sciences at Michigan State University, and a Malagasy civil society group called Alliance Voahary Gasy (AVG), hopes to eventually employ PAWS in Madagascar to randomize patrol schedules for rangers, police and national park officials to reduce environmental crimes, especially illegal logging.

Startup: ARMORWAY

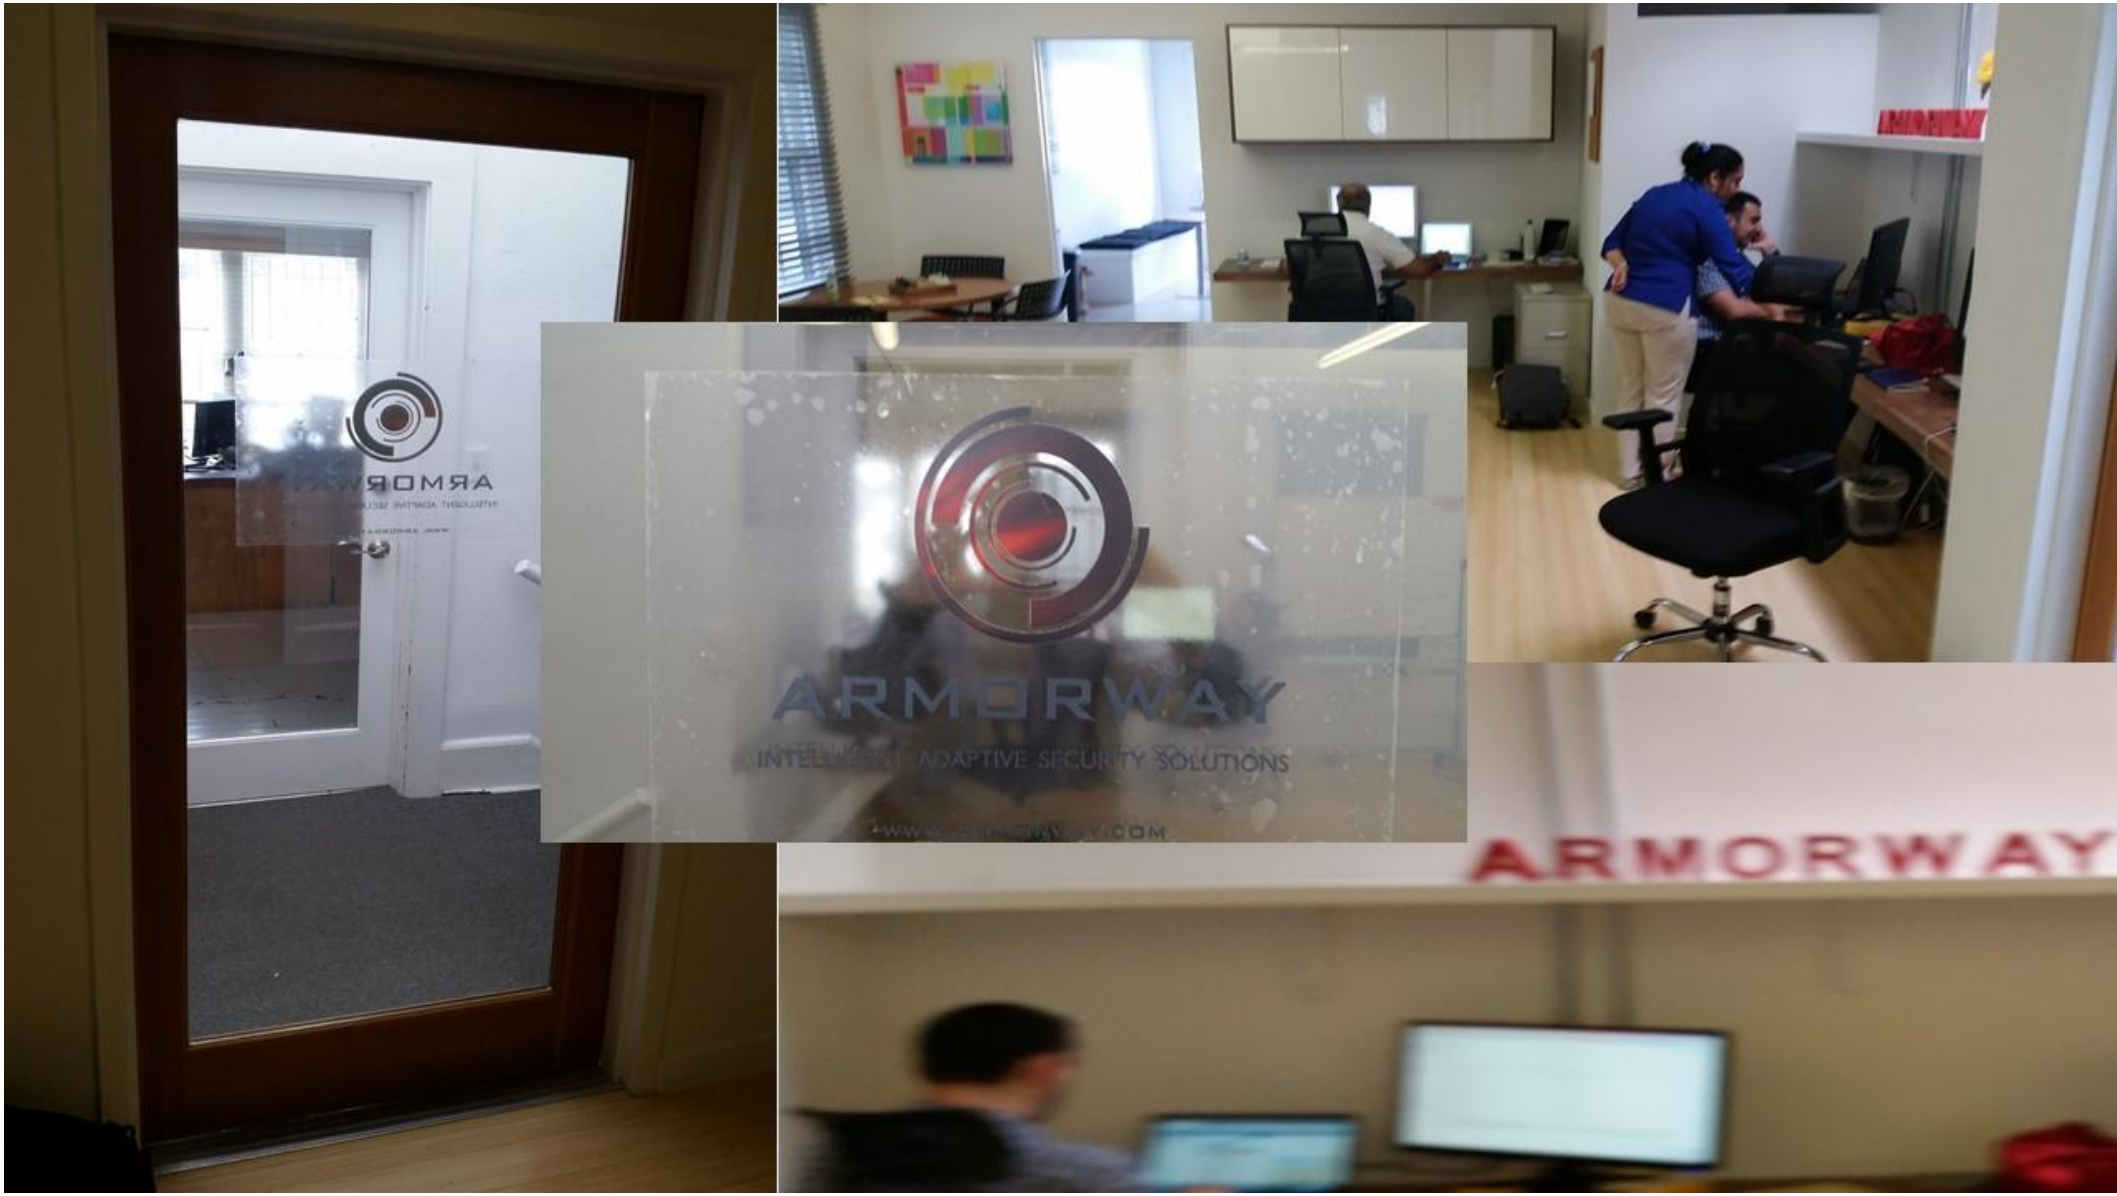

Outline: Security Games Research (2007-)

Infrastructure Security Games

Green Security Games

Opportunistic Crime Security Games

Air travel

Ports

Trains

Fisheries

Wildlife

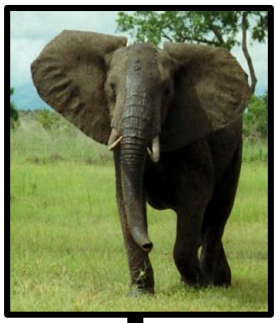

Urban Crime

2007 2011 2013 2014 2014 2015

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

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

ARMOR Airport Security: LAX [2007]
Basic "Stackelberg Security Game" Model

Basic Security Game Operation [2007]

Using ARMOR as an 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

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

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

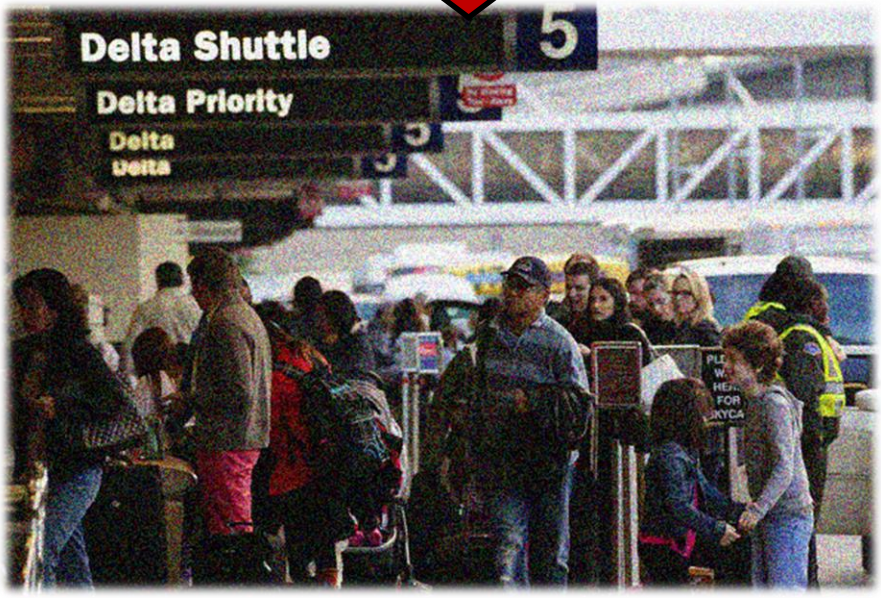


Newsweek
ARMOR…throws a digital cloak of invisibility….

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

- *ARMOR runs out of memory*

- *Incremental strategy generation:*
  - *Column generation: Not enumerate all 10^{41} actions*

IRIS: Incremental Strategy Generation
Column Generation


Jain Kiekintveld

Master

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

Slave

Best new pure strategy

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

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

GLOBAL OPTIMAL

500 rows

NOT 10^{41}

IRIS: Deployed FAMS (2009-)

Significant change in FAMS operations

"…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, Social Networks[2013]
Scale-up: Double Oracle

Jain

Road networks:
e.g., checkpoints

20416 Roads,15 checkpoints:
20 min

PROTECT: Port Protection Patrols Deployed 2011-
Using “Marginals” for Scale-up

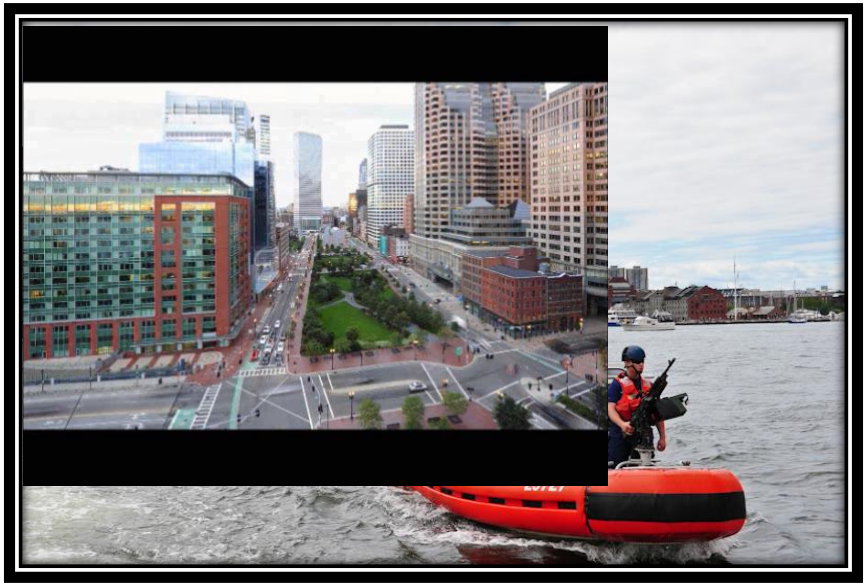

USS *Cole* after attack

French oil tanker hit by small boat

PROTECT: Ferry Protection Deployed 2013-
Using “Marginals” for Scale-up

CNN iReport

SIGN UP | LOG IN

Main Explore Assignments Profile Upload

g+1 Tweet Share Favorite

99 VIEWS | 0 COMMENTS | 6 SHARES

About this iReport

- Not vetted for CNN

Posted September 8, 2013 by shortysmom | Follow

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

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

Share on Facebook 0

I ride the Staten Island Ferry on a daily basis to and from work. We ferry riders

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

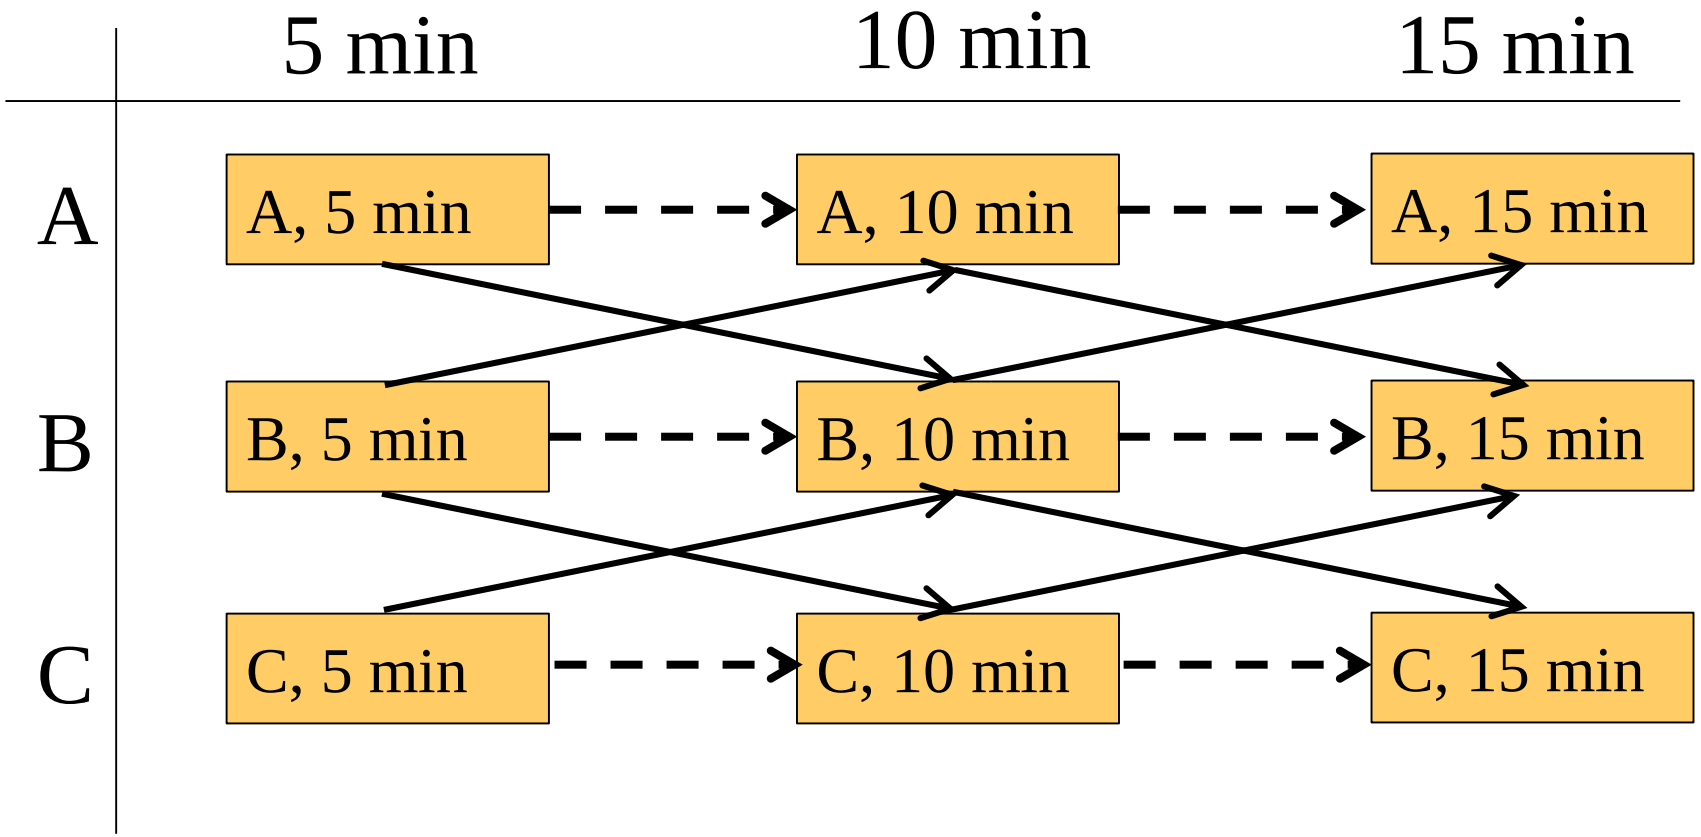

Ferries: Scale-up with Mobile Resources & Moving Targets
Transition Graph Representation

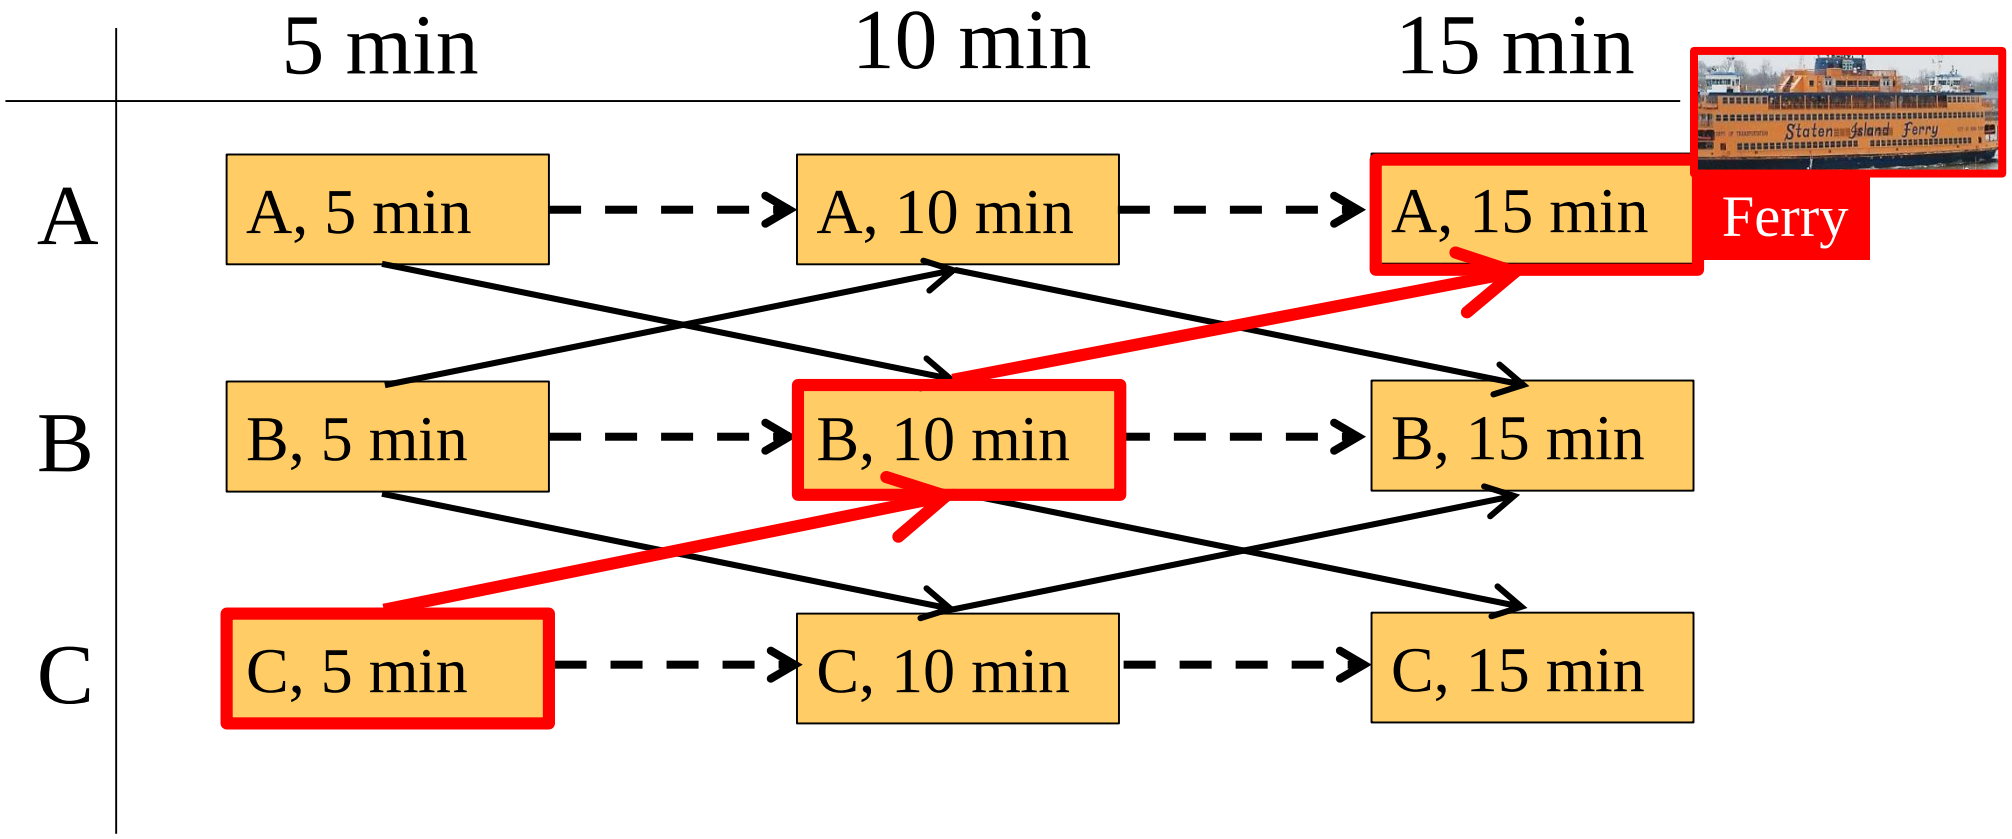

Ferries: Patrol Routes as Variables
Exponential Numbers of Patrol Routes

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

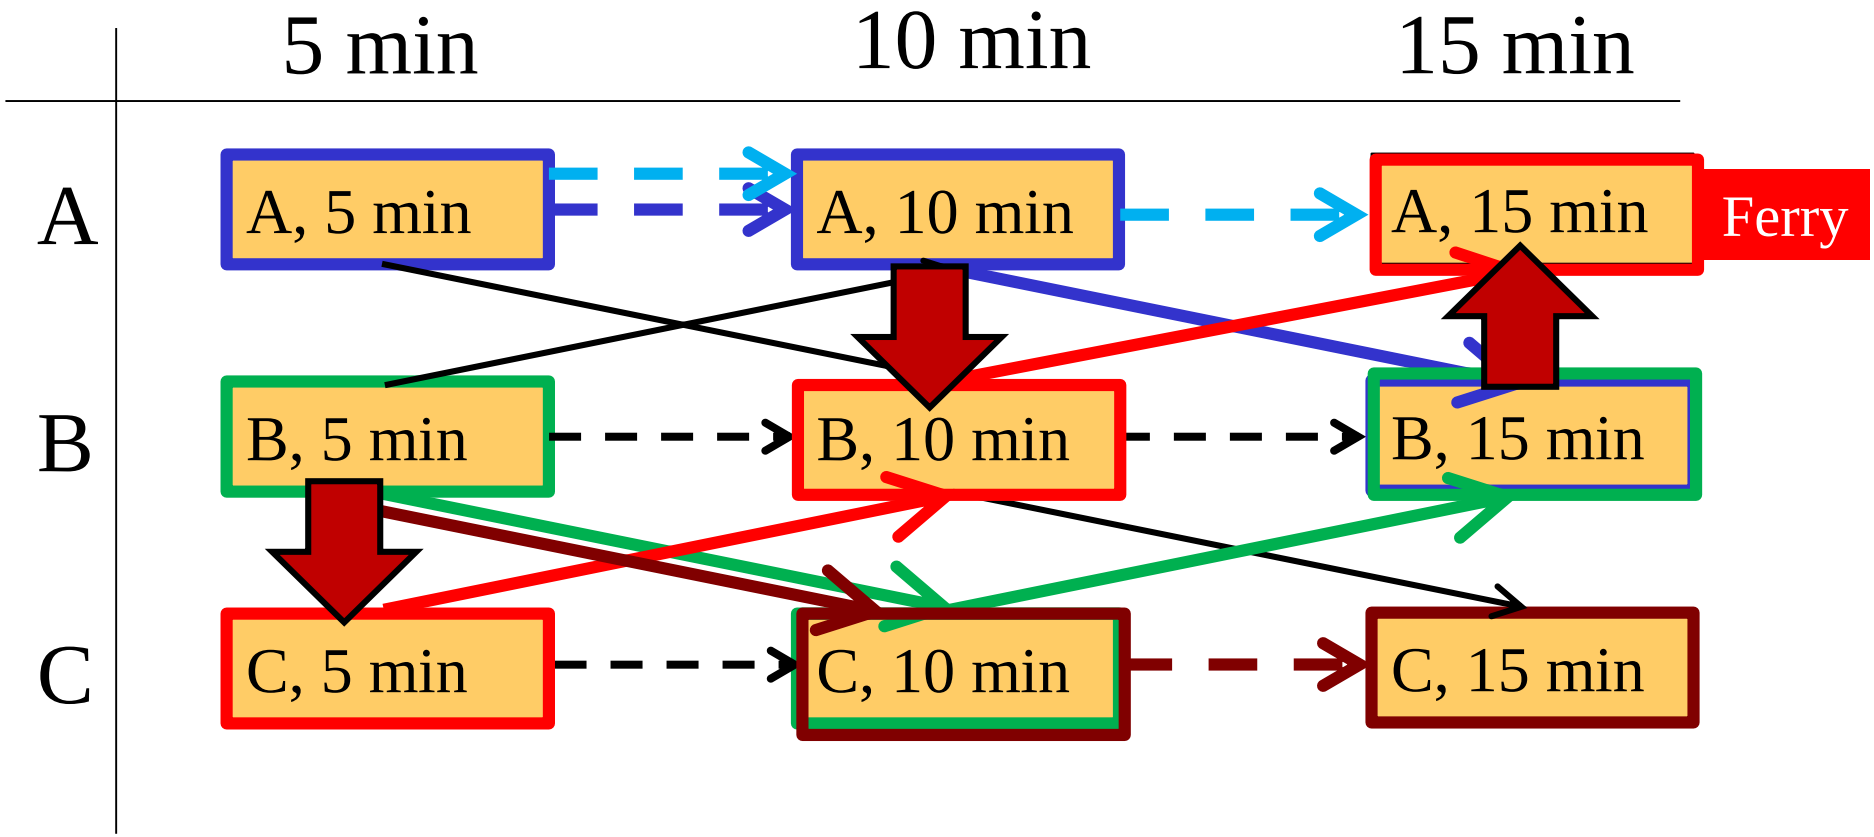

Ferries: Patrol Routes as Variables
Exponential Numbers of Patrol Routes



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*

N^T variables

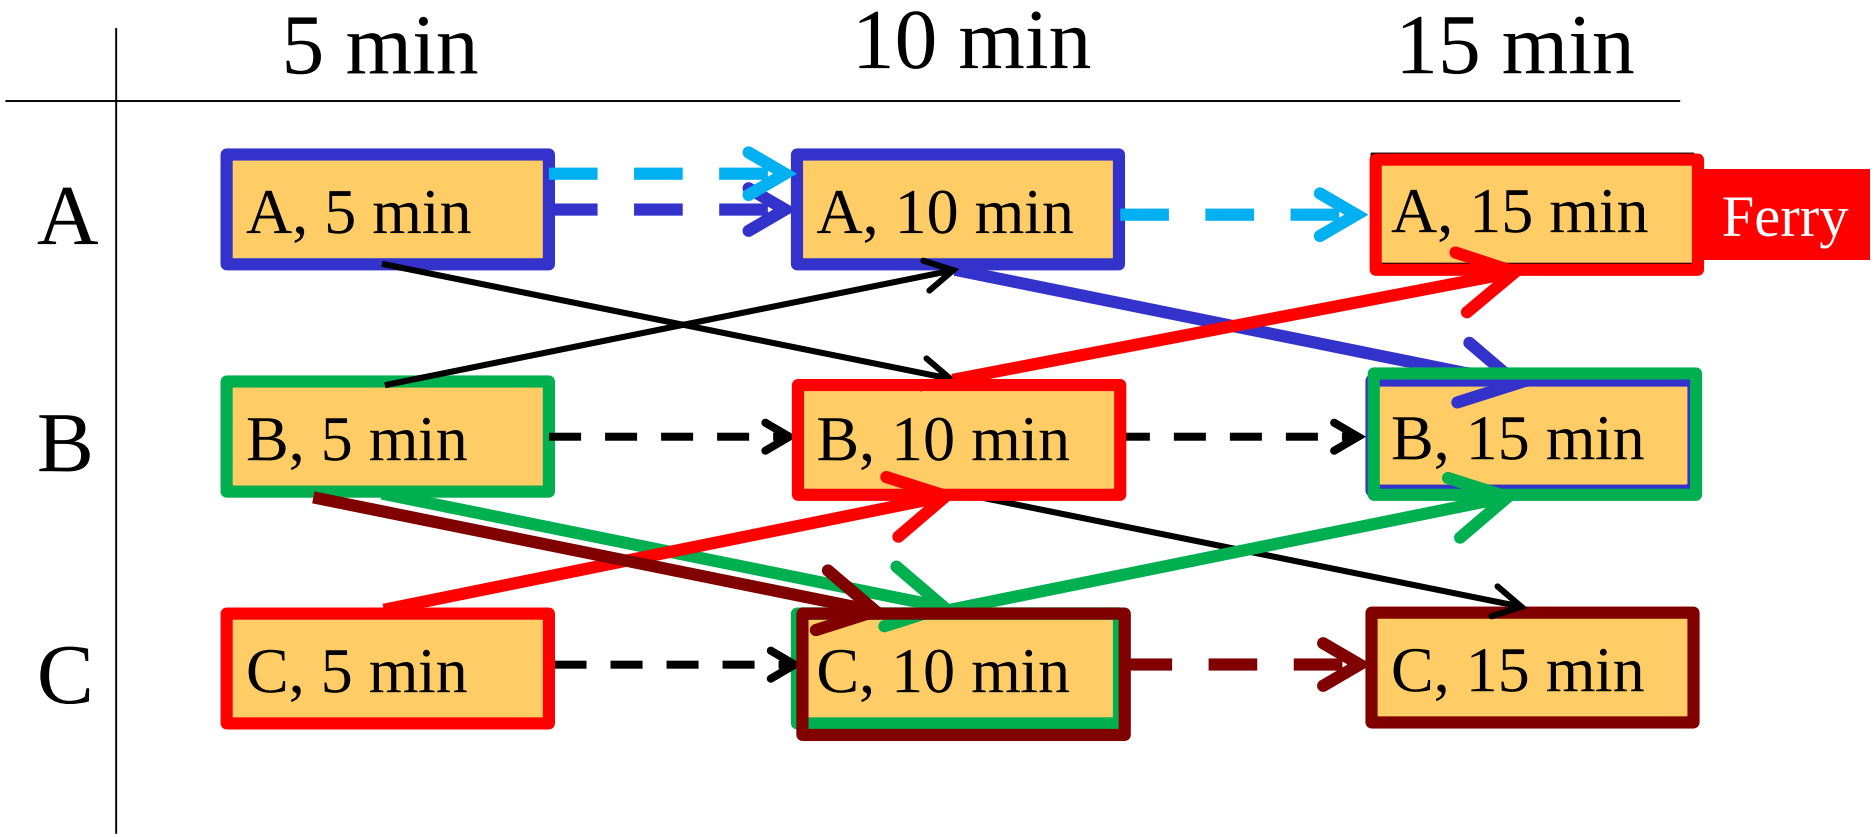

Ferries: Scale-up
Marginal Probabilities Over Segments


Fang

Jiang

Kiekintveld

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

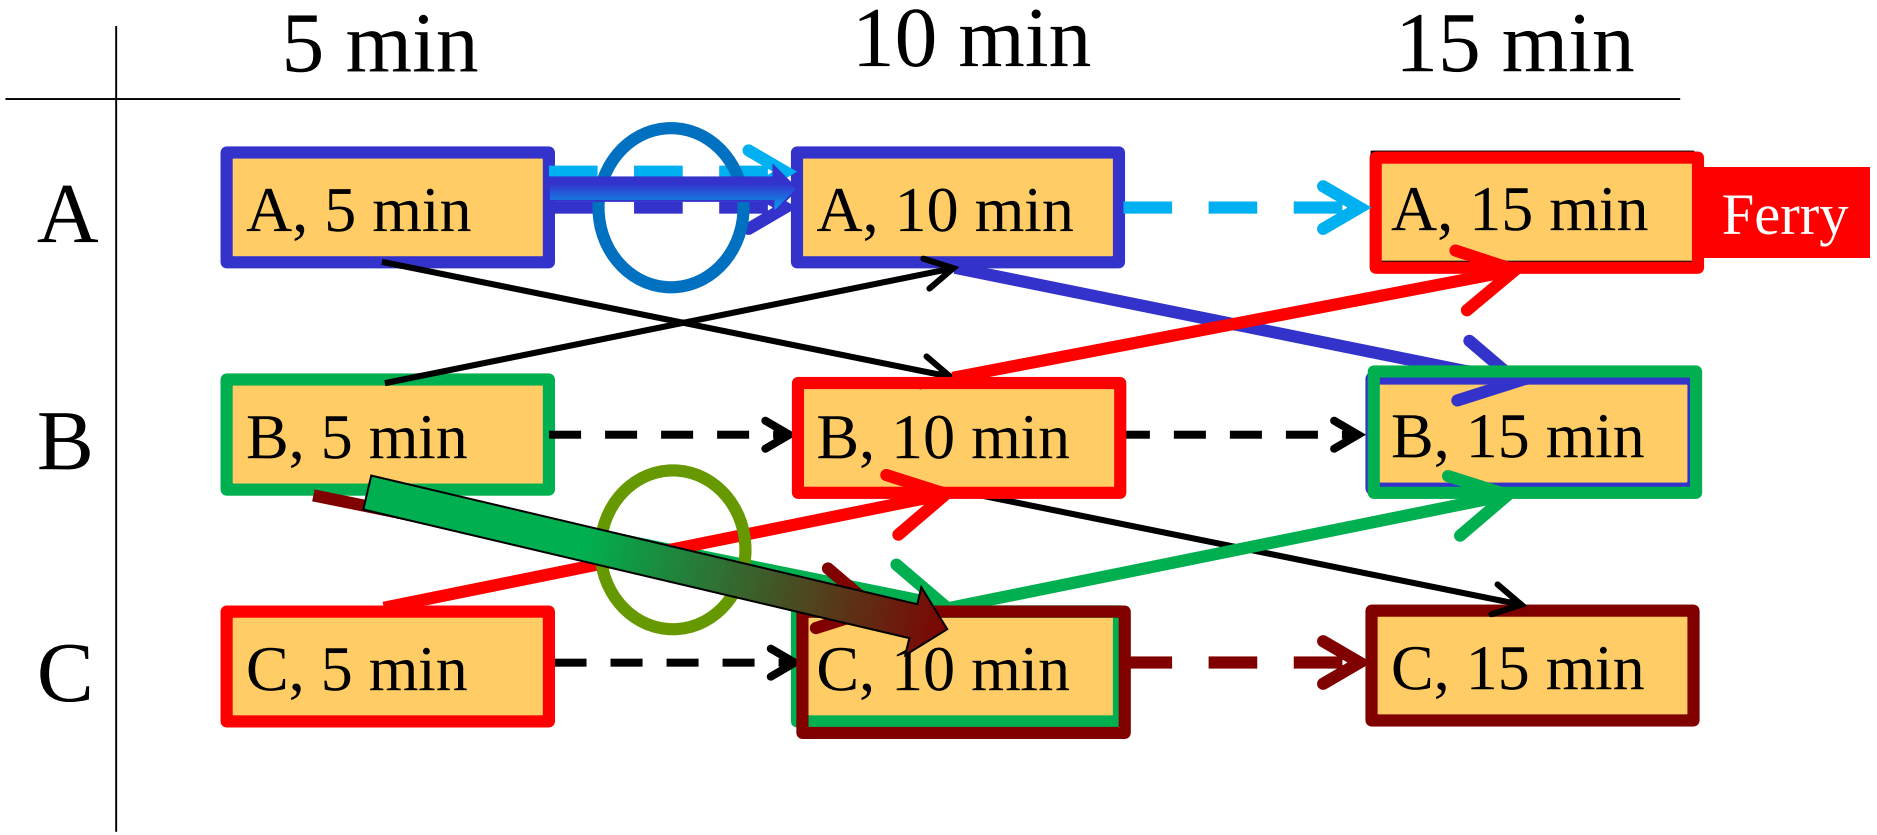

Ferries: Scale-up with Marginals Over Separable Segments Significant Speedup

- *Obtain marginal probabilities over segments*

$N^2.T$ *variables*

N^T *variables*

Extract: [(A, 5), (C, 10), (C,15)]) = 0.47

[(B, 5), (C,10), (B,15)]) =0.23

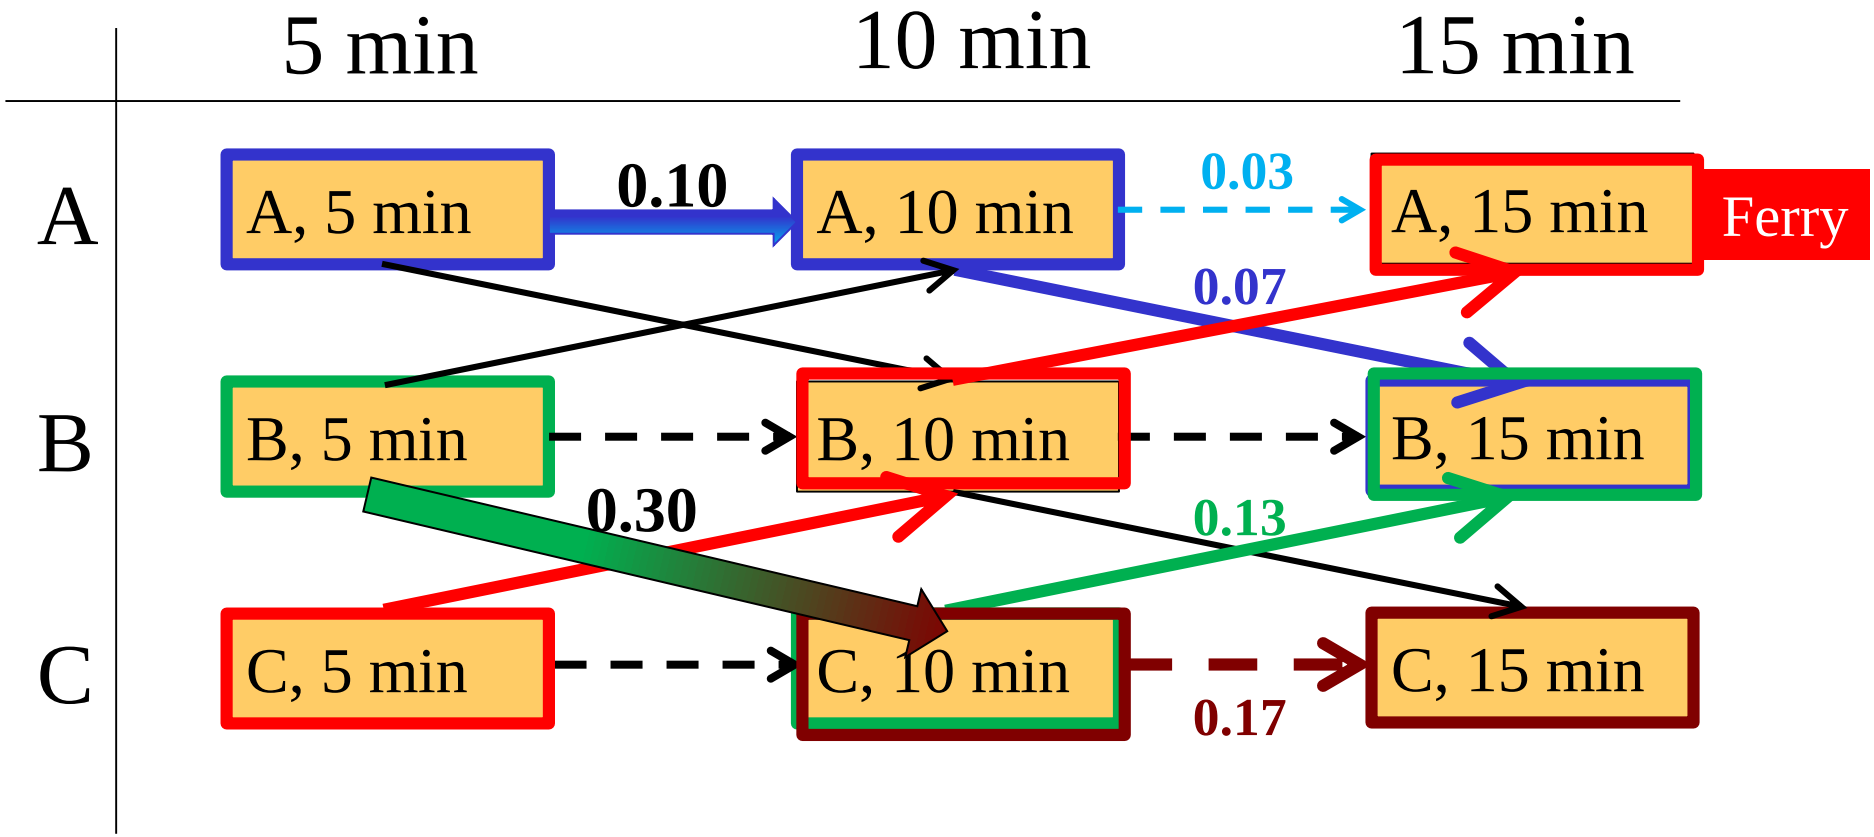

Outline: “Security Games” Research (2007-)

Air travel

2007

Urban Crime

2015

TRUSTS: Frequent adversary interaction games
Patrols Against Fare Evaders

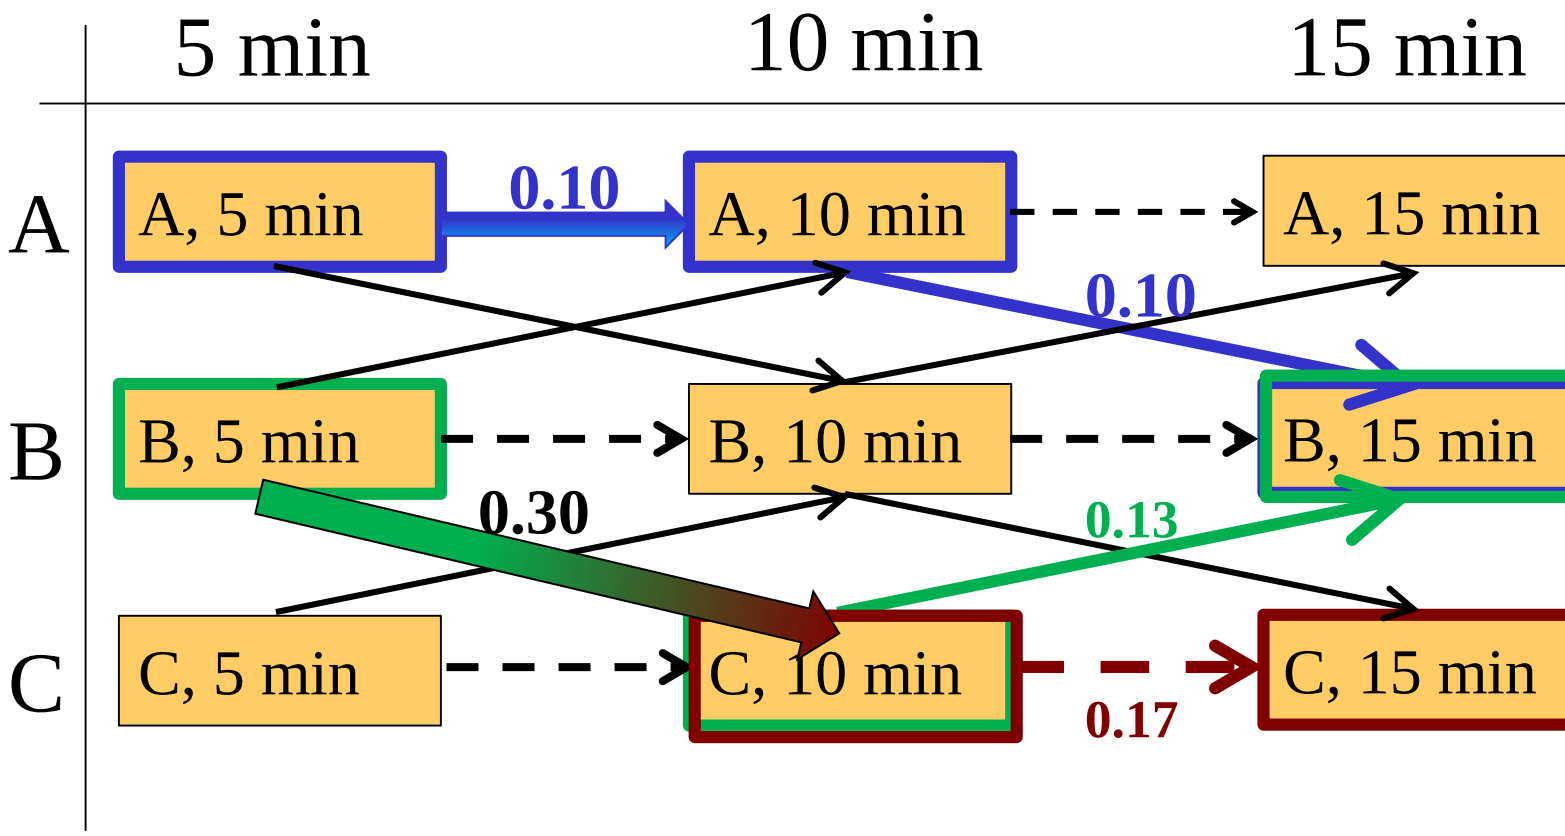

TRUSTS: Patrols Against Fare Evaders Uncertainty in Defender Action Execution


Jiang *Delle Fave*

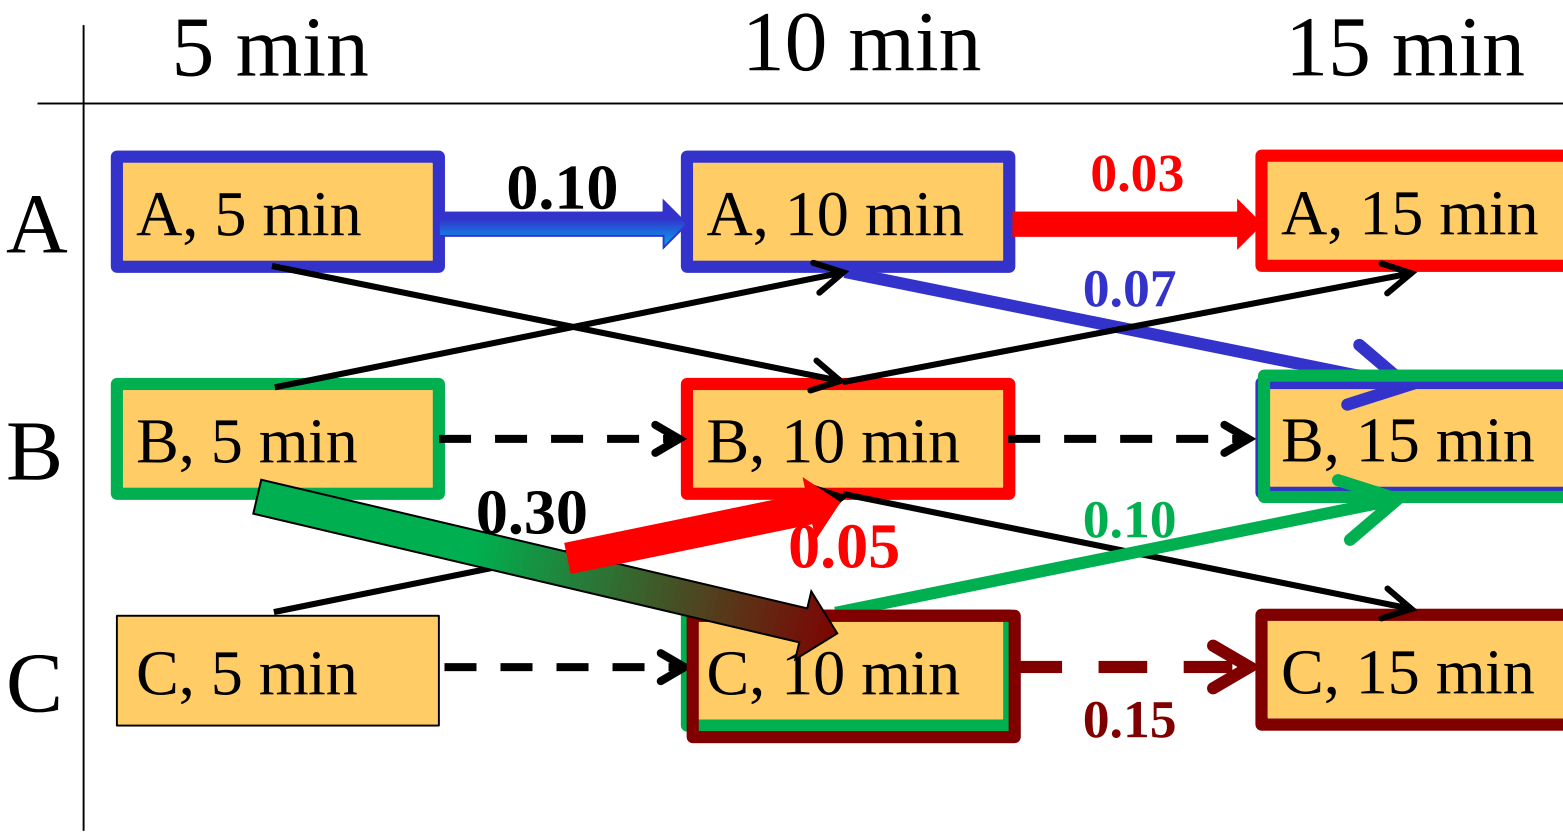

TRUSTS: Patrols Against Fare Evaders Uncertainty in Defender Action Execution


Jiang *Delle Fave*

- Markov Decision Problems *in Security games*

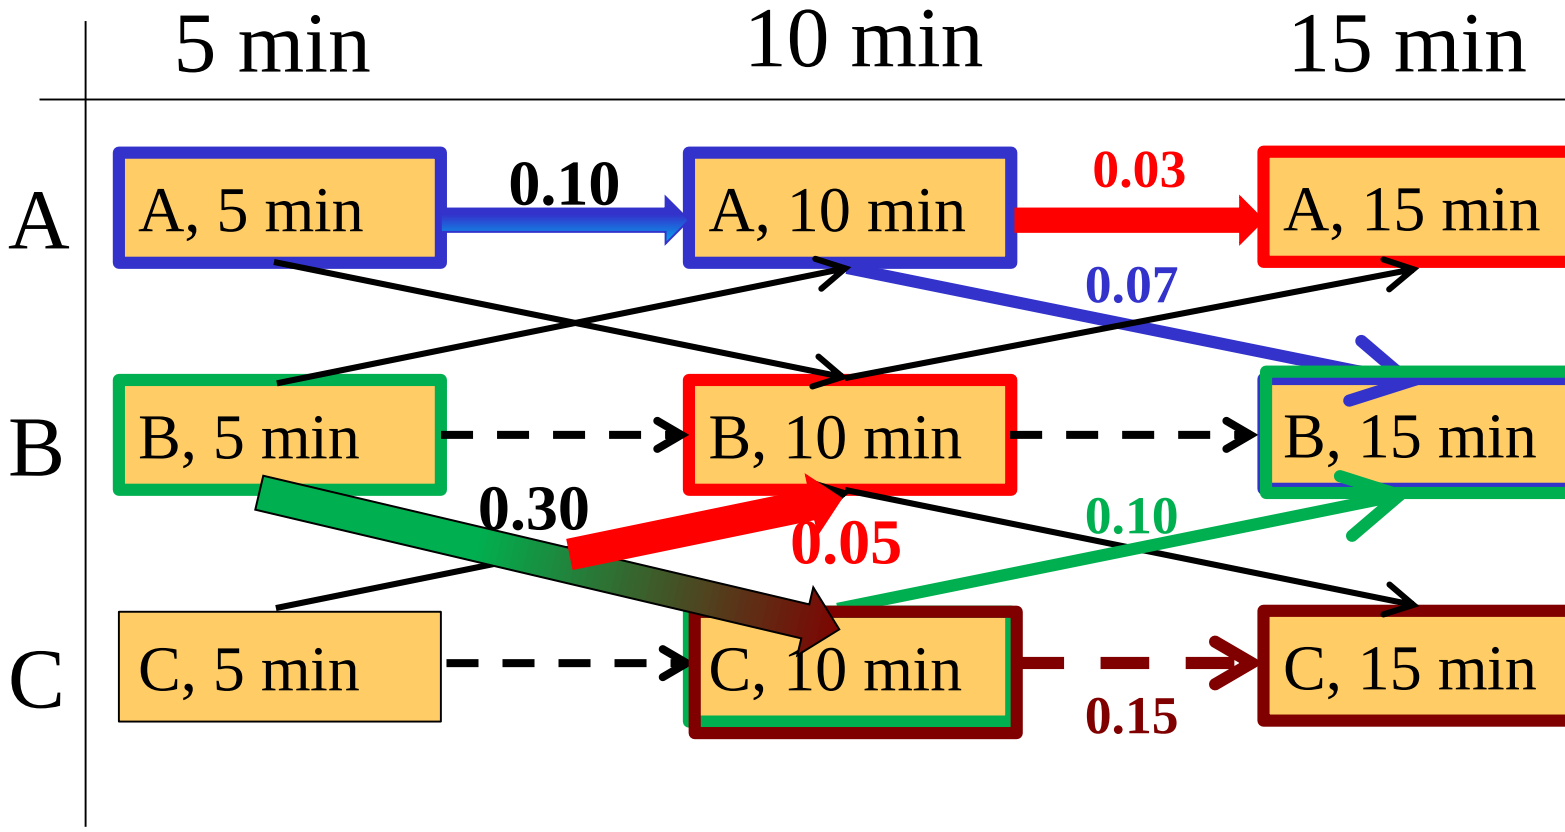

Uncertainty Space Algorithms: Bayesian and Robust Approaches

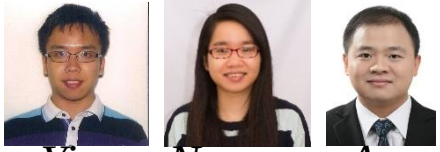

Adversary payoff uncertainty

Payoff interval; Not point estimate

GMC

BRASS

HUNTER

RECON

URAC

Monotonic Maximin (Monotonic adversary)

Adversary observation & defender execution uncertainty

Adversary rationality uncertainty

Bayesian

Robust

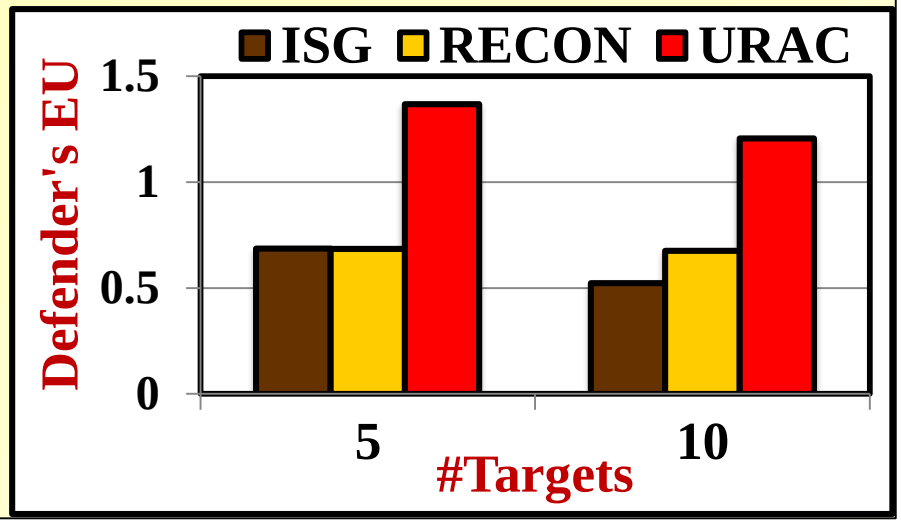

Outline: Security Games Research (2007-)

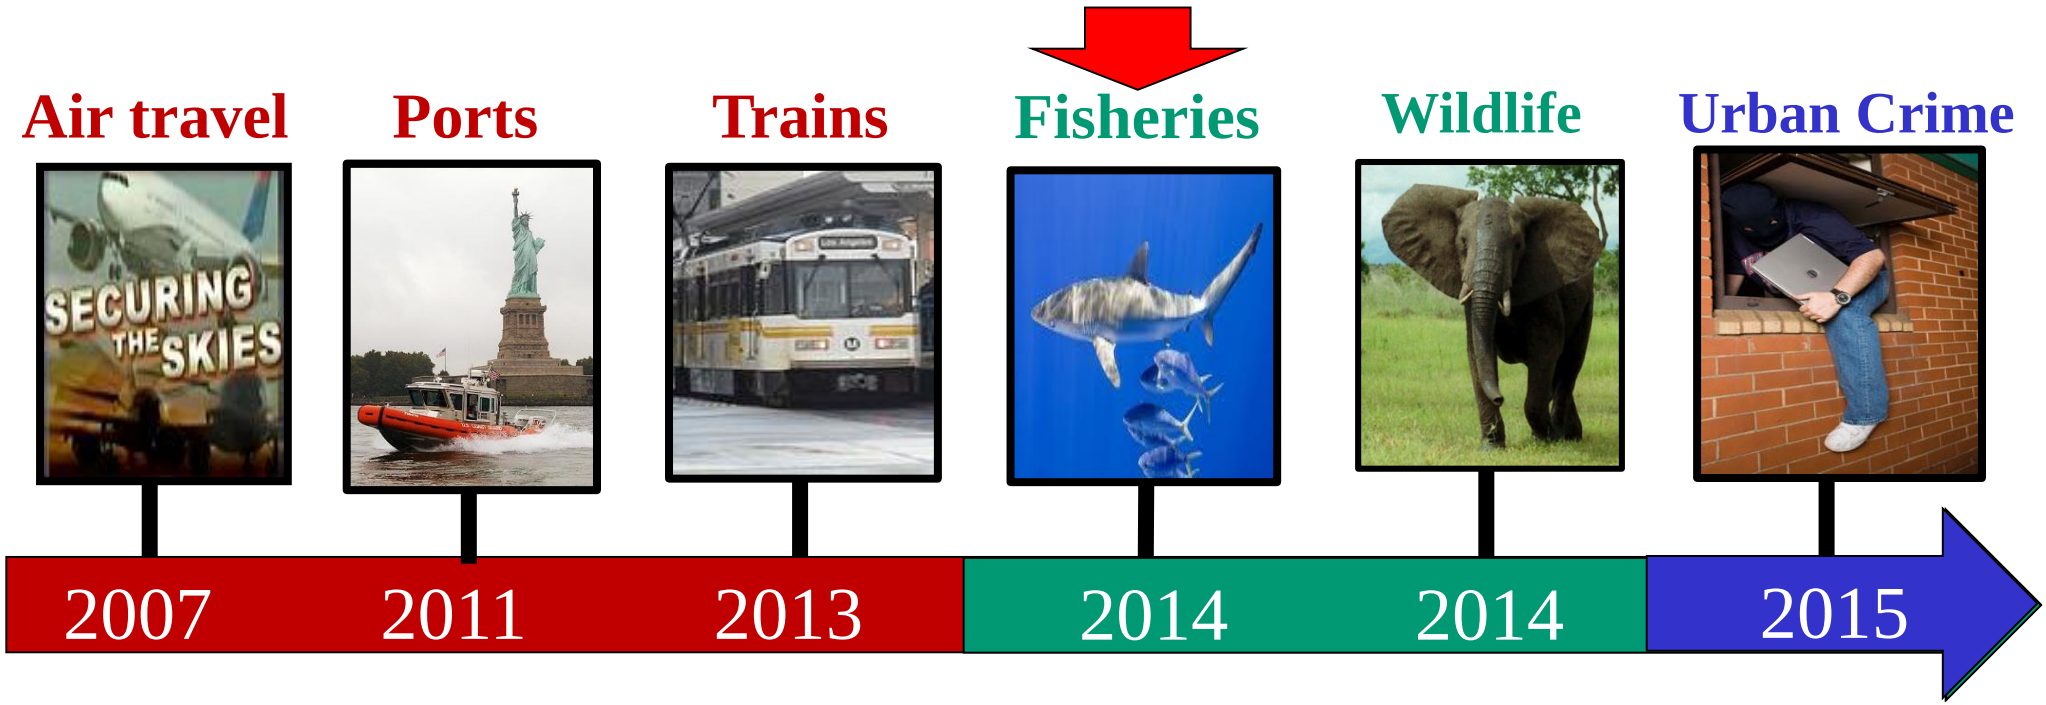

Protecting Forests, Fish, Rivers & Wildlife: Green Security Games



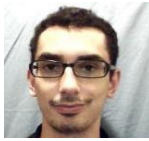

McCarthy *Ford* *Brown*

River Pollution Prevention



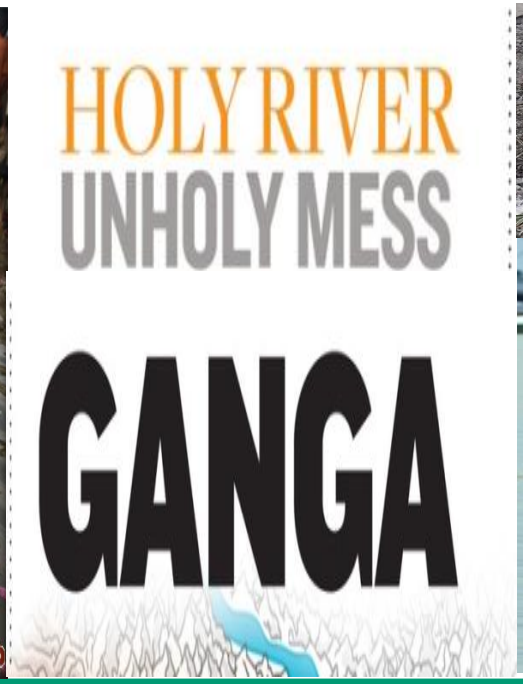


Fishery Protection

Forest Protection

Wildlife Protection:
Murchison Falls National Park, Uganda

Green Security Games: Repeated Stackelberg Game

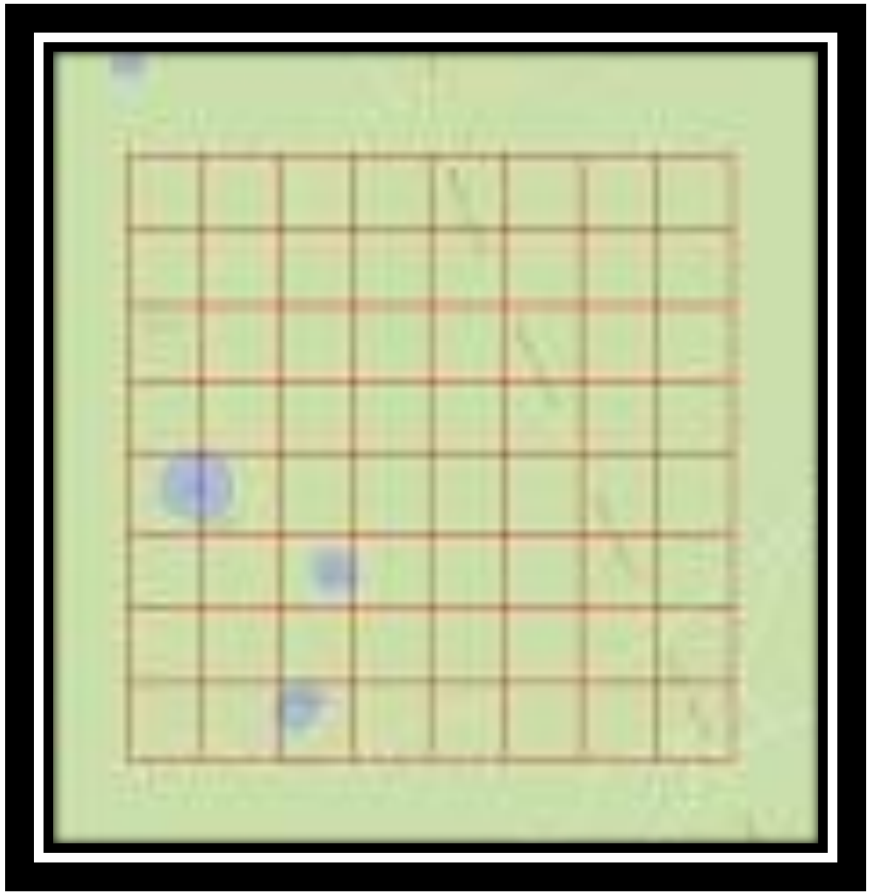

Bounded rationality model of poachers

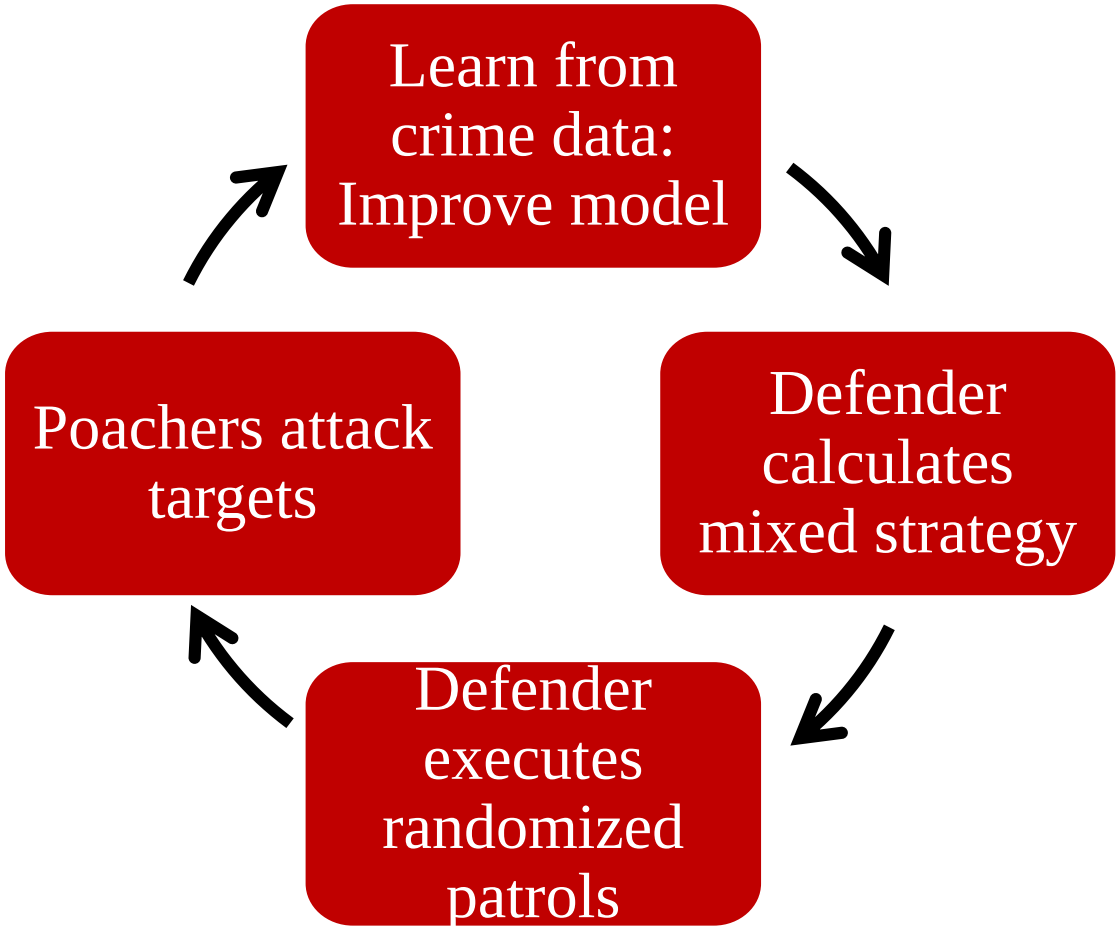

Uncertainty in Adversary Decision: Bounded Rationality
Human Subjects as Poachers

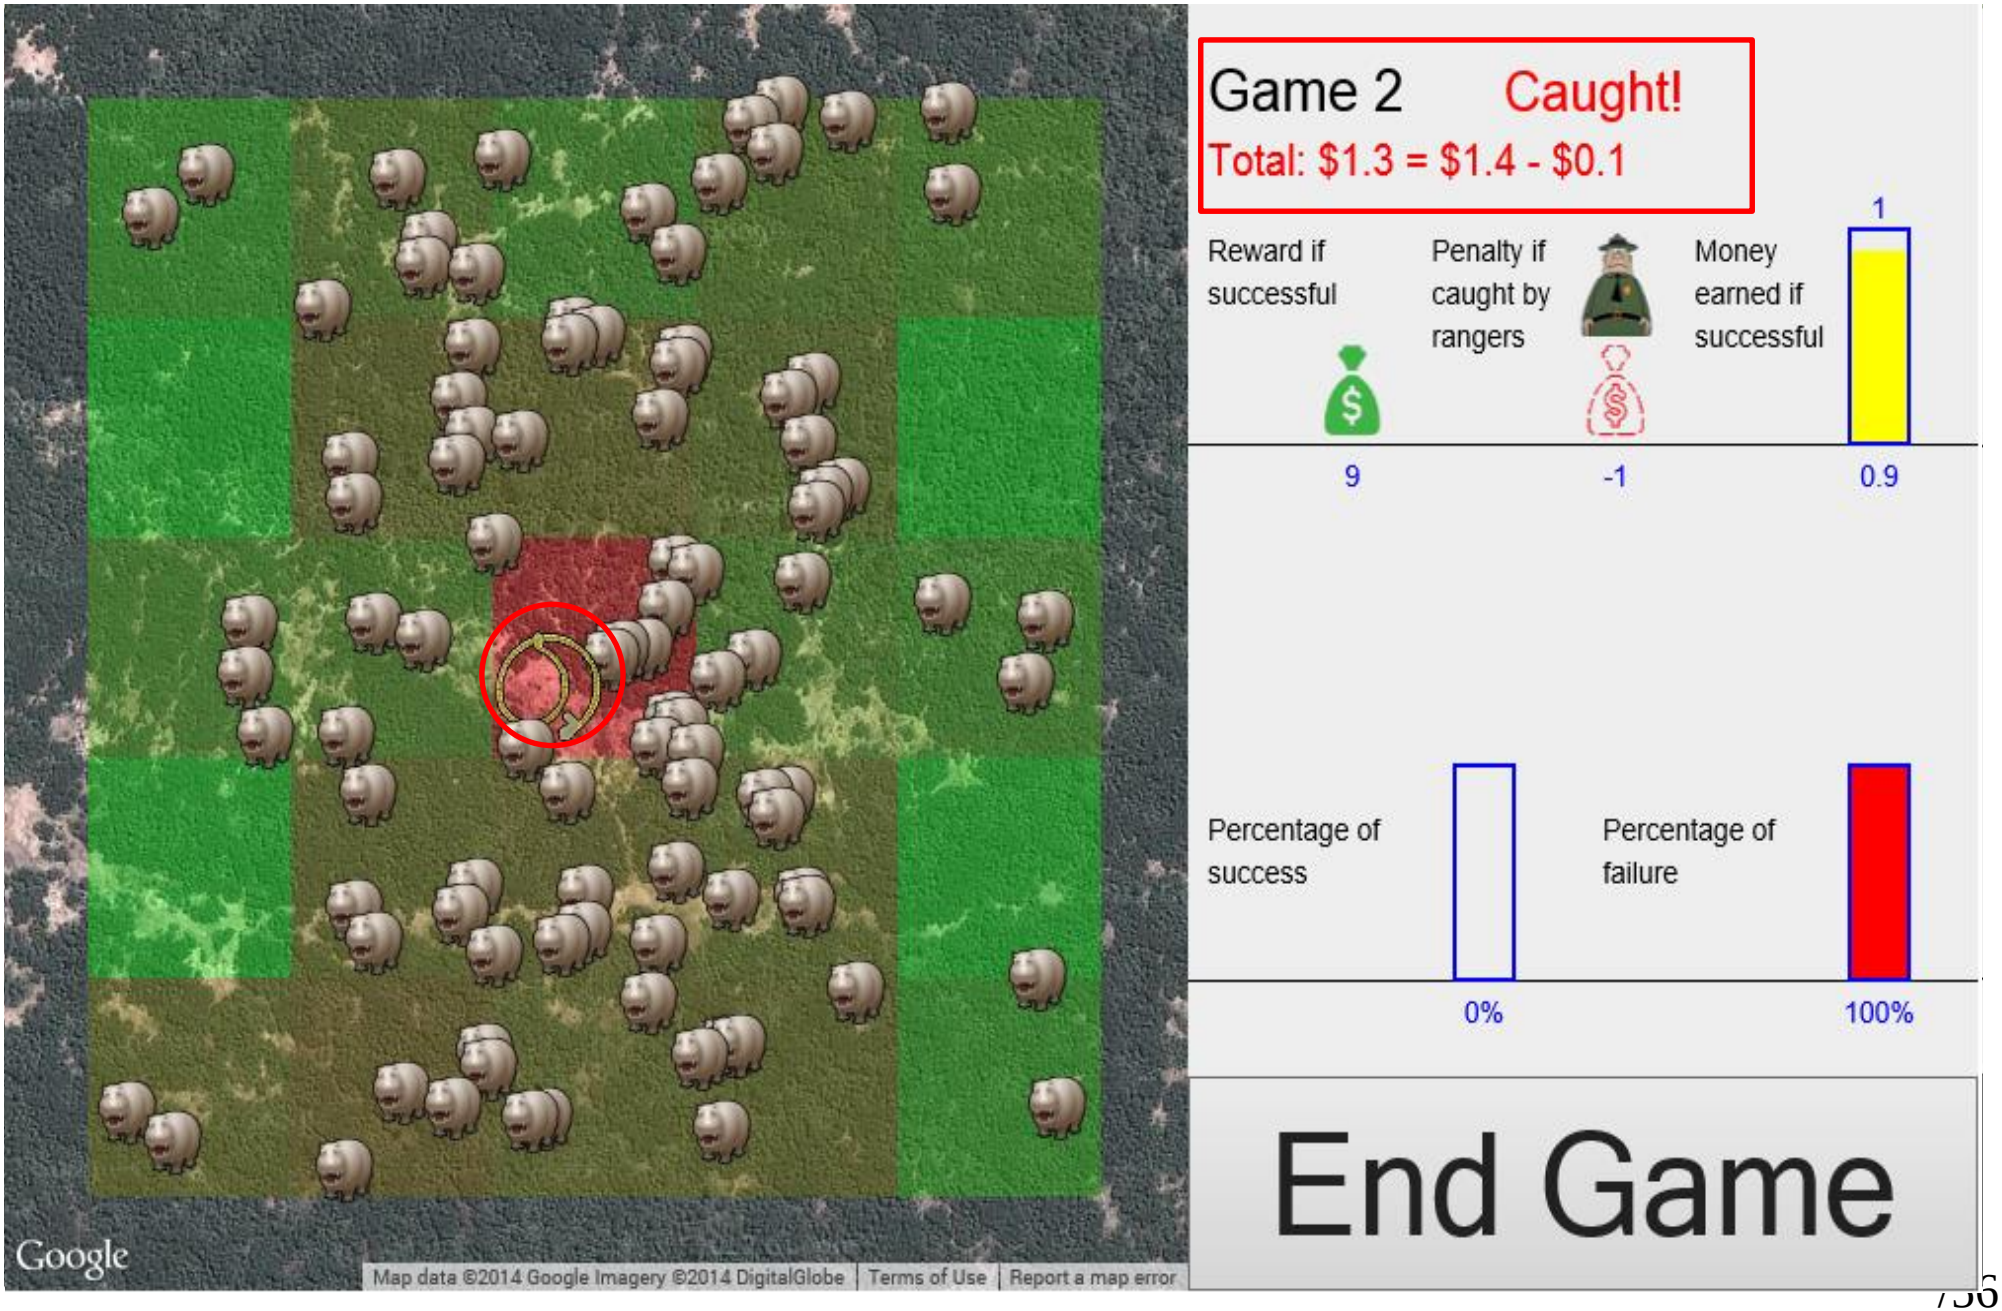

Lesson 1: Quantal Response [2011]: Models of Bounded Rationality


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

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

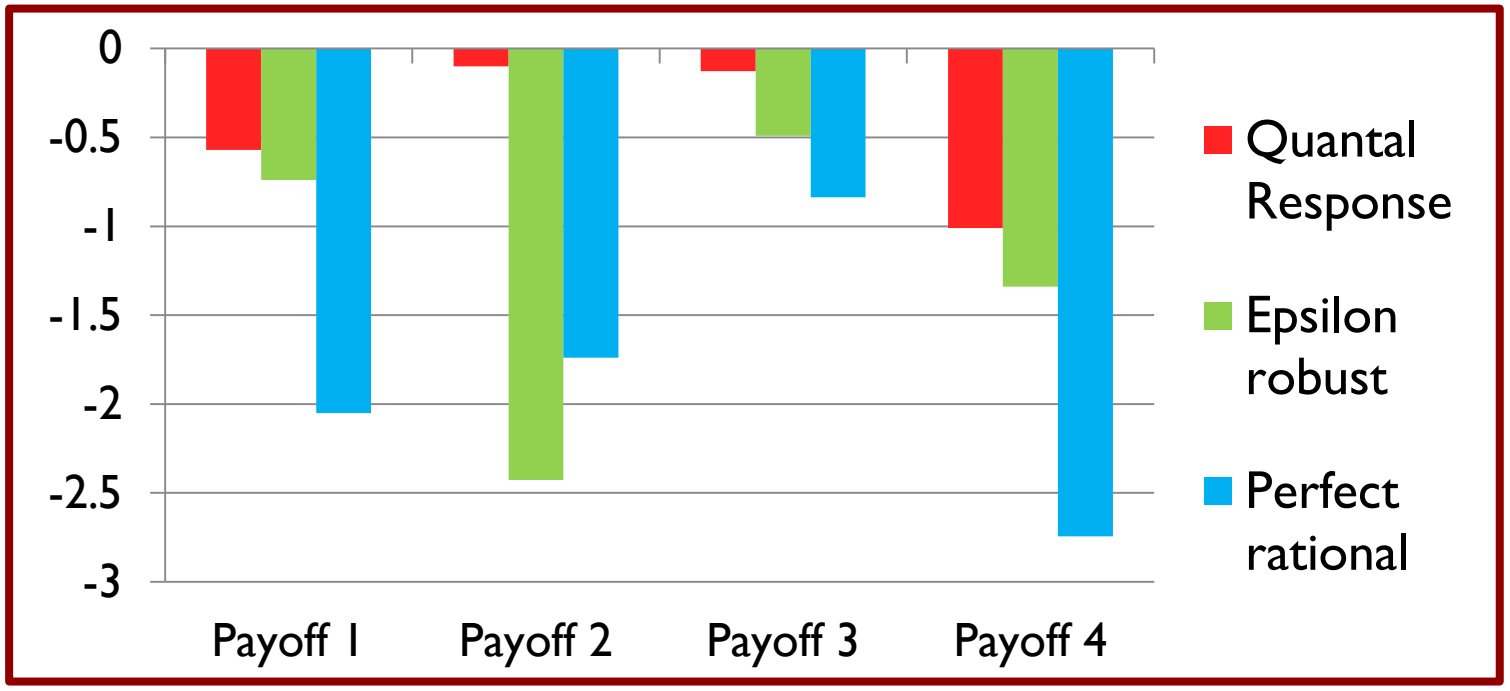

Lesson 2: Subjective Utility Quantal Response Models of Bounded Rationality

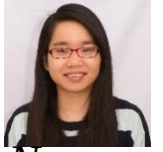

Subjective Utility Quantal Response(SUQR) ***[Nguyen 13]:***

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

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

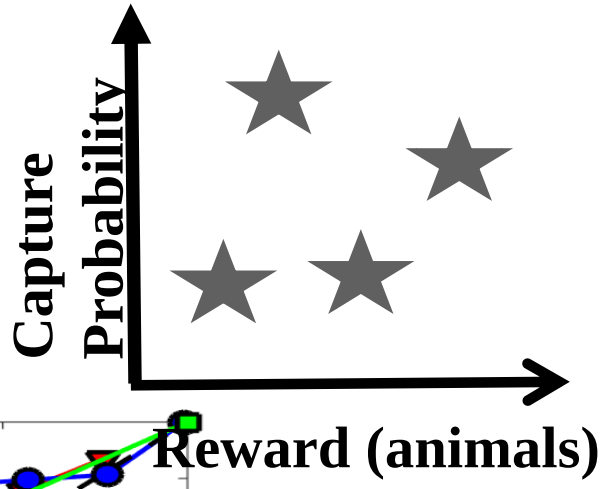

UWA data from Uganda: Year 2012 Predicting poaching

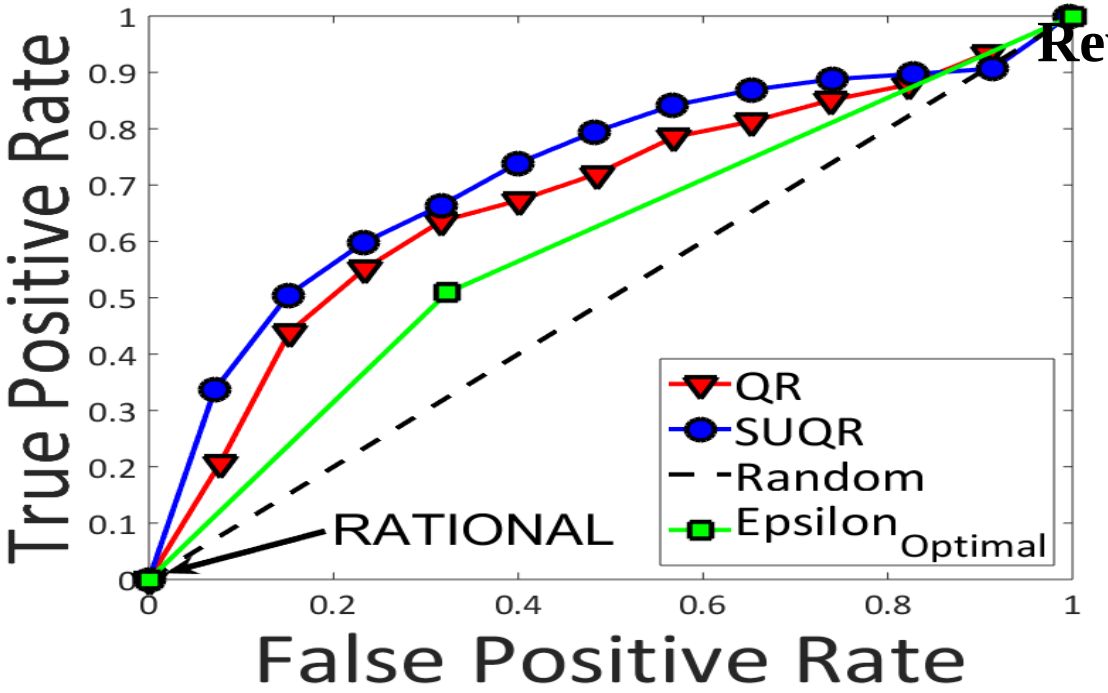

Green Security Games[2015]
Testing SUQR: From One-Shot to Repeated Games

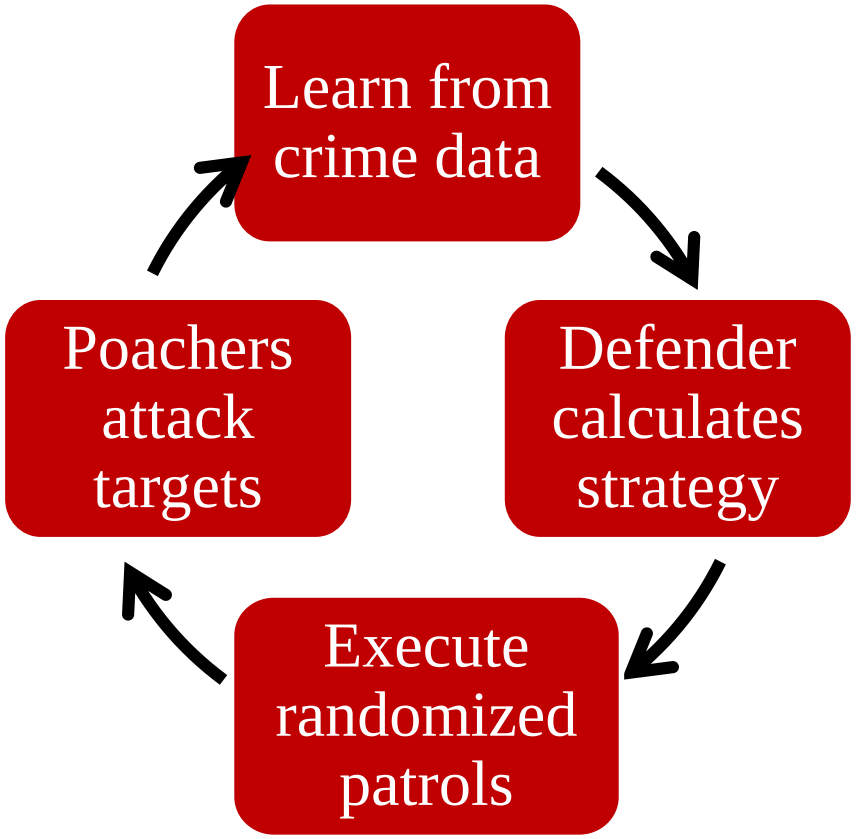

Repeated games on AMT:
35 weeks, 40 human subjects
10,000 emails!

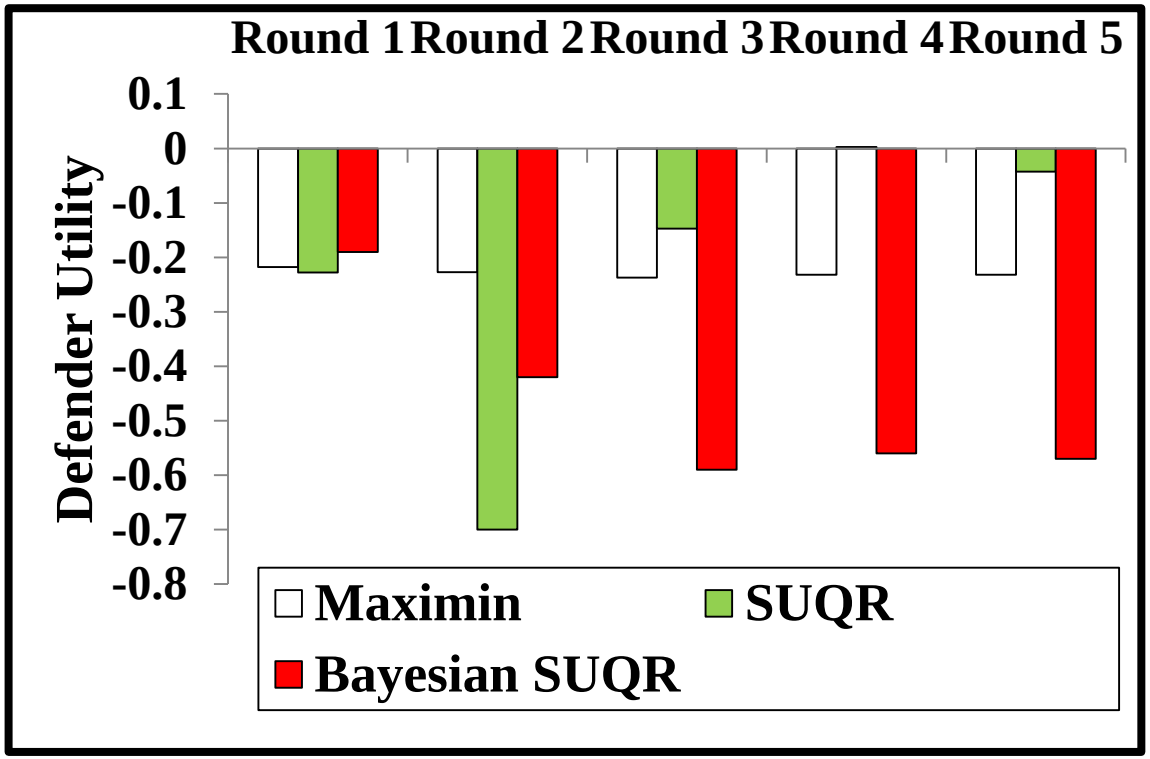

Lesson 3: SHARP and Repeated Stackelberg Games
Incorporate Past Success/Failure in SUQR

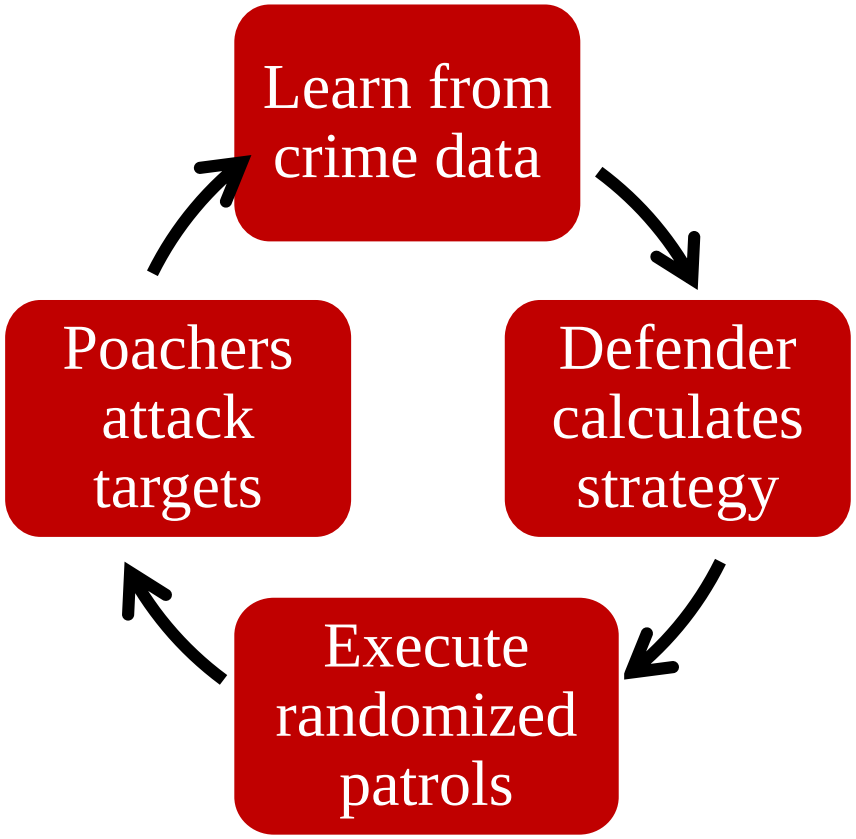

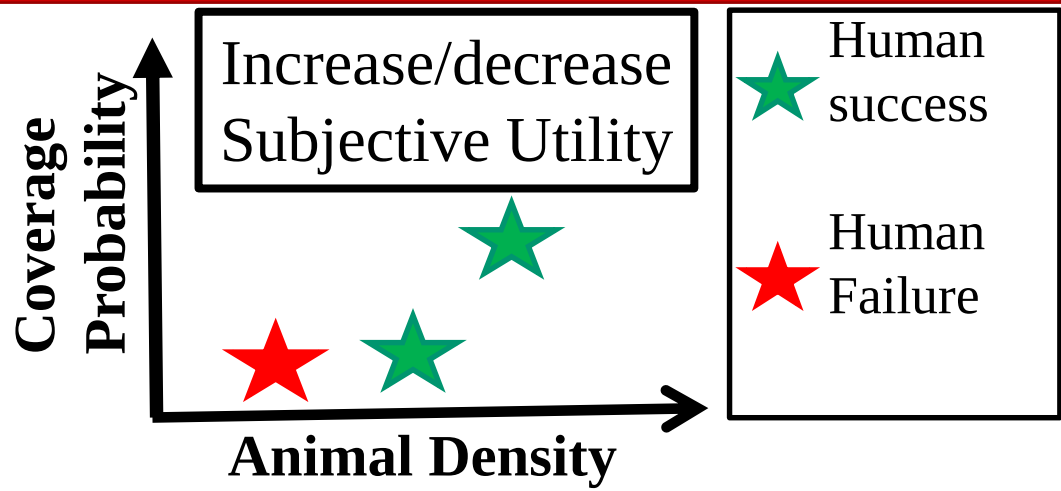

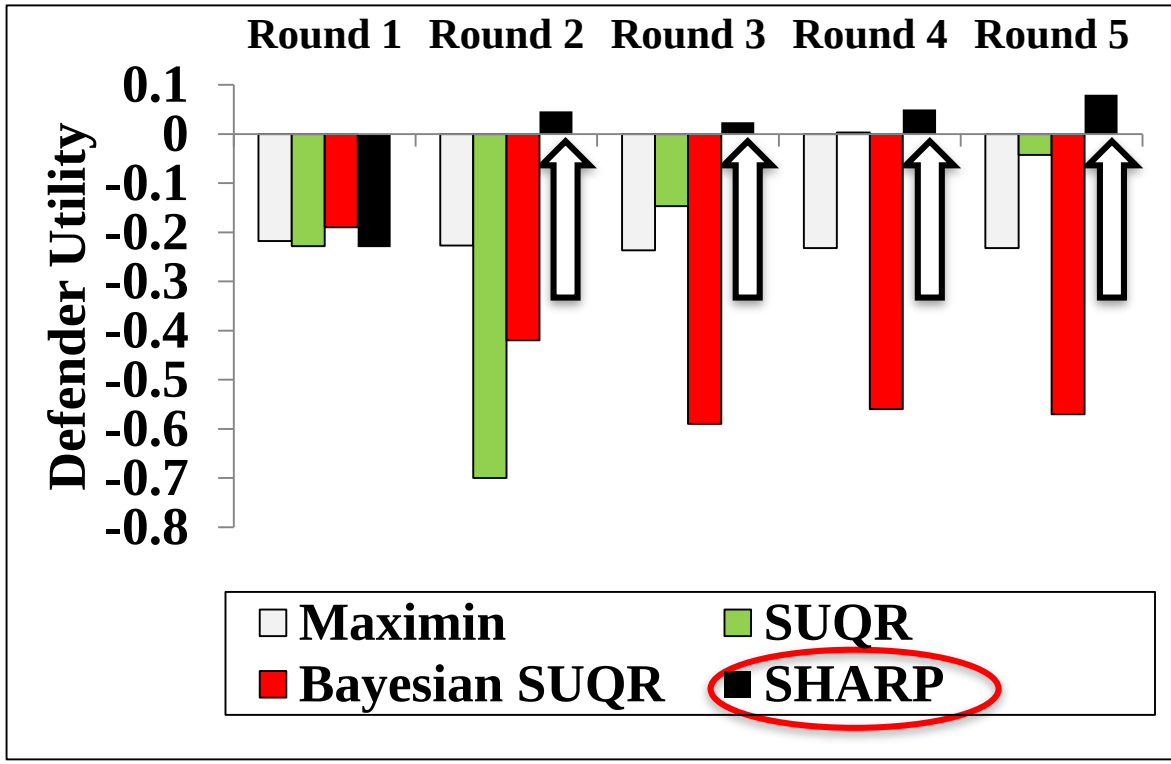

Learned Probability Weighting Function

- Adversary's probability weighting function is S-shaped.
  - *Contrary to Prospect Theory (Kahneman '79).*

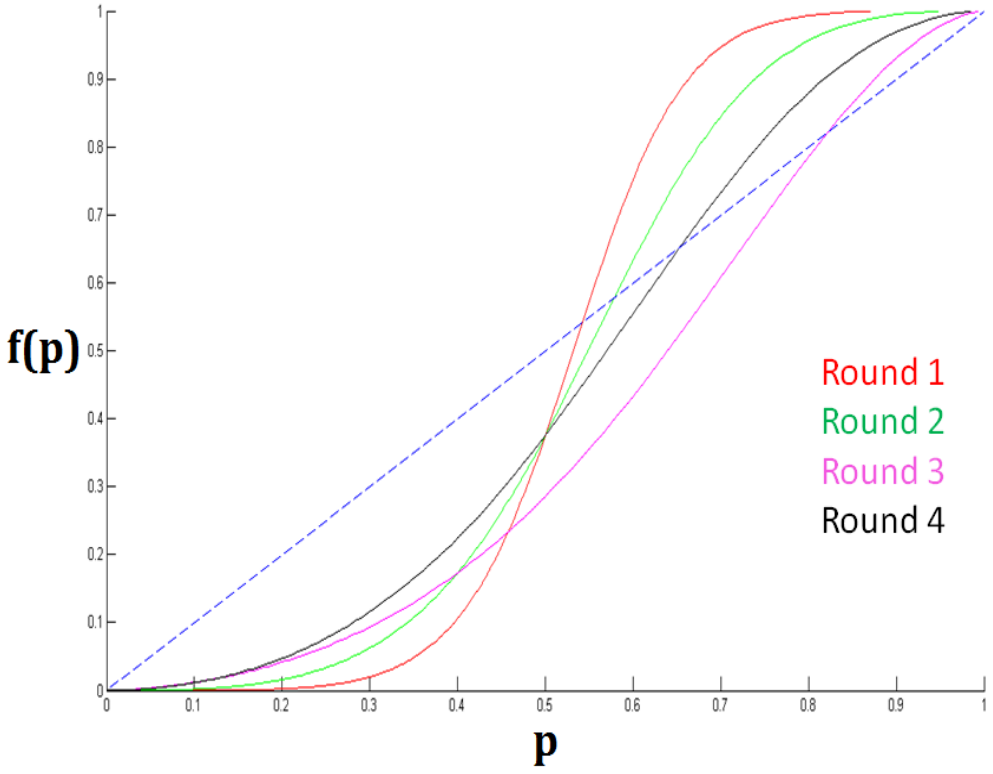

PAWS: Protection Assistant for Wildlife Security Trials in Uganda and Malaysia: [2014]

- **Important lesson: Geography!**

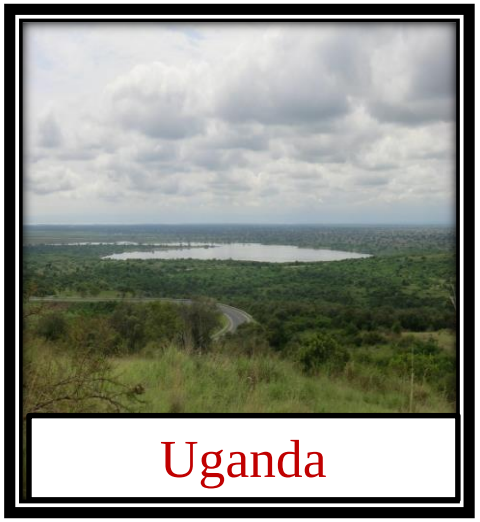
Uganda

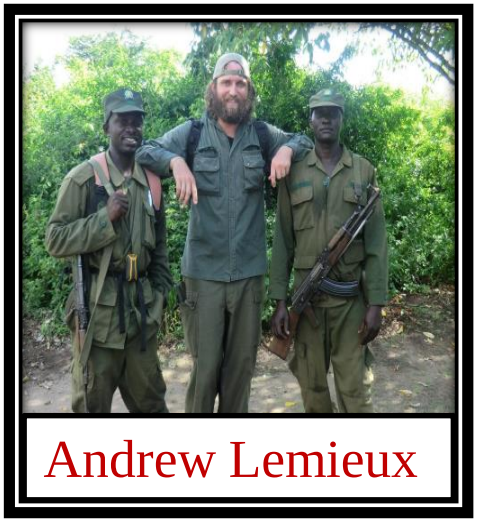
Andrew Lemieux

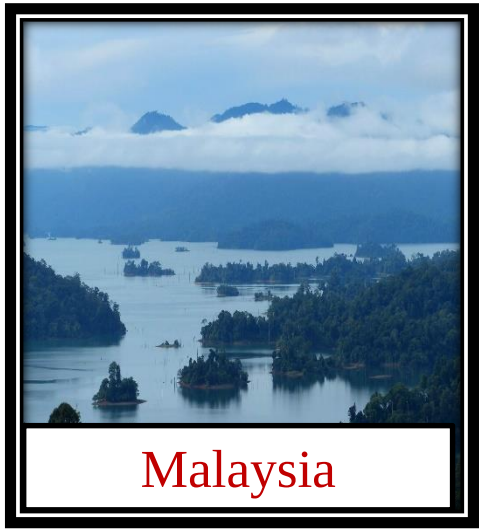
Malaysia

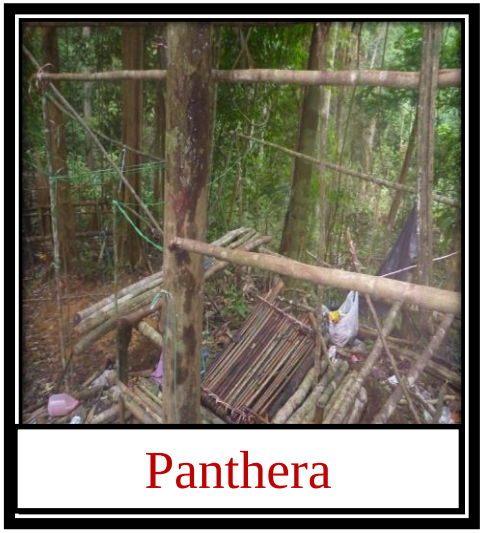
Panthera

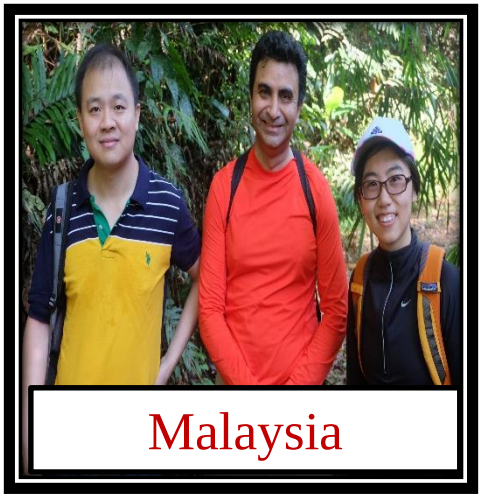
Malaysia

PAWS for Wildlife Security: Scale, Uncertainty in Green Security Games


Fang

- Scale: Hierarchical model
  - *Hierarchical: Grid +"Street map"*
- Species location uncertainty
- ***In regular use in Malaysia***

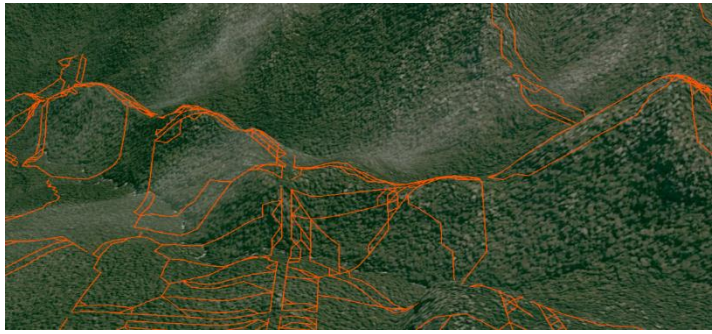

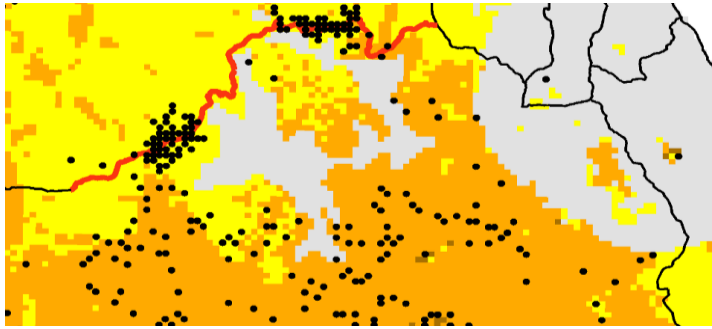

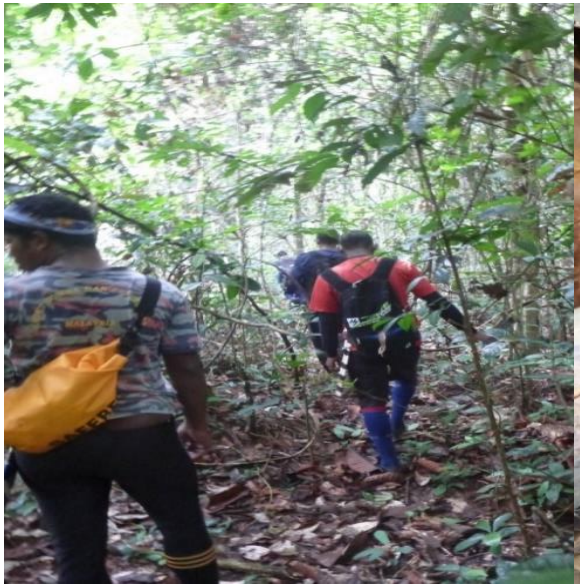

Opportunistic Crime Security Game[2015] Integrating Learning in Basic Security Game Model


Zhang


Sinha

- Crime prediction: use past crime & police allocation data

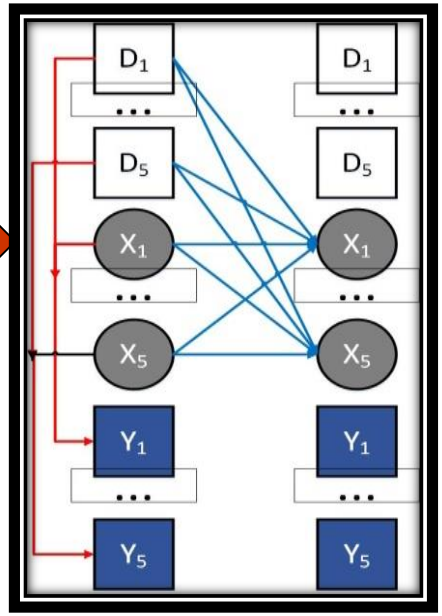

Best Simulation Results

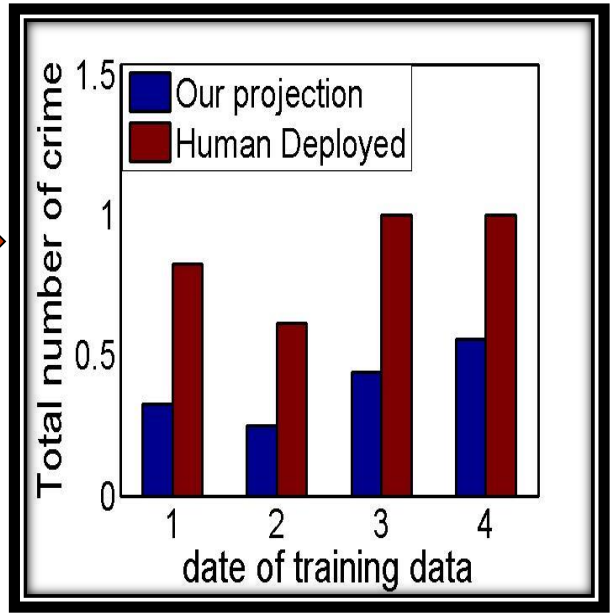

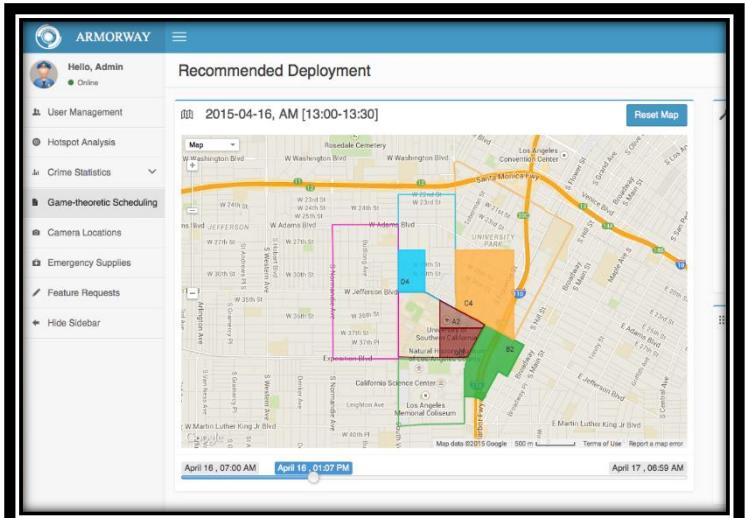

Evaluating *Deployed Security Systems* Not Easy:
Are Security Games Better at Optimizing Limited Resources

- Security games improve over humans (or simple) planners
  - *E.g., humans fall into predictable patterns; high cognitive load*

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	

Field Evaluation of Schedule Quality:
Improved Patrol Unpredictability & Coverage for Less Effort

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

Patrols Before PROTECT: Boston

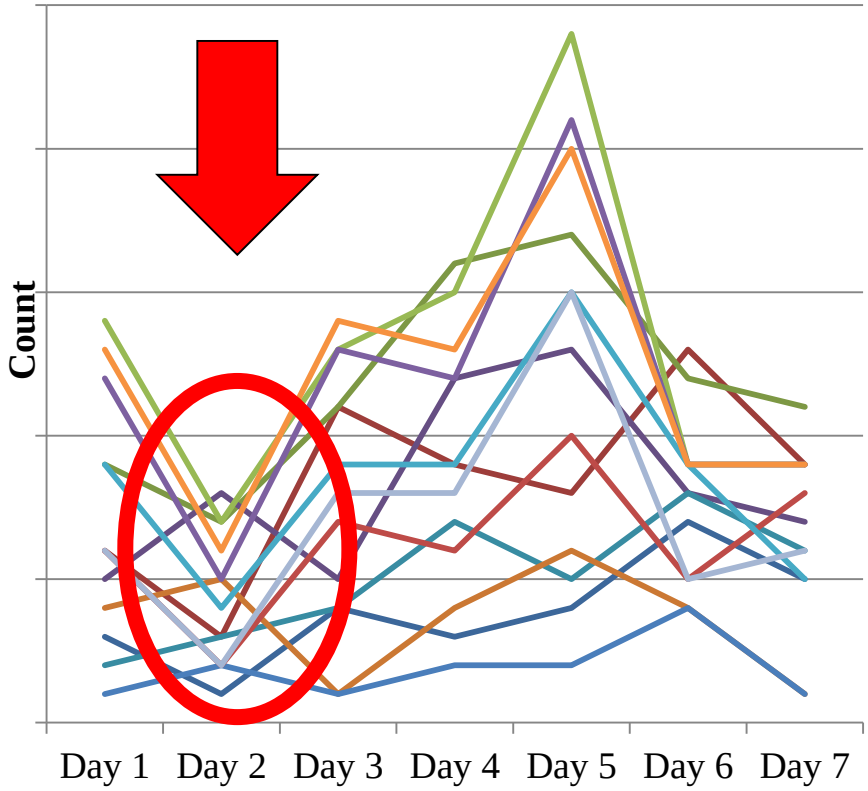

Patrols After PROTECT: Boston

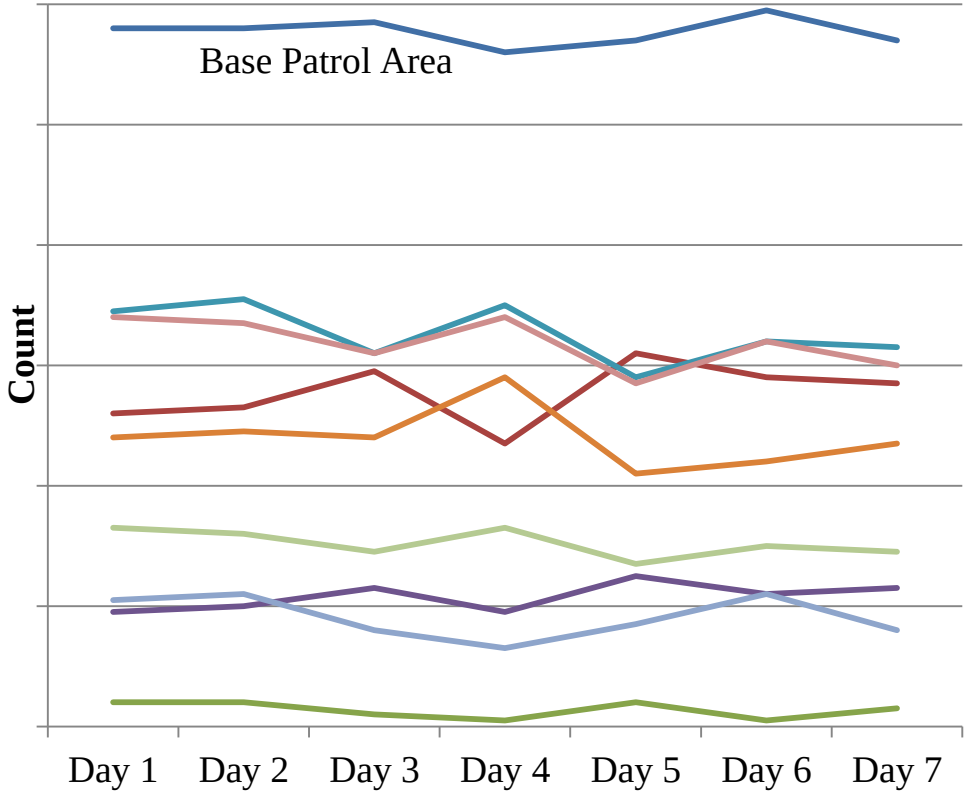

Field Evaluation of Schedule Quality:

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

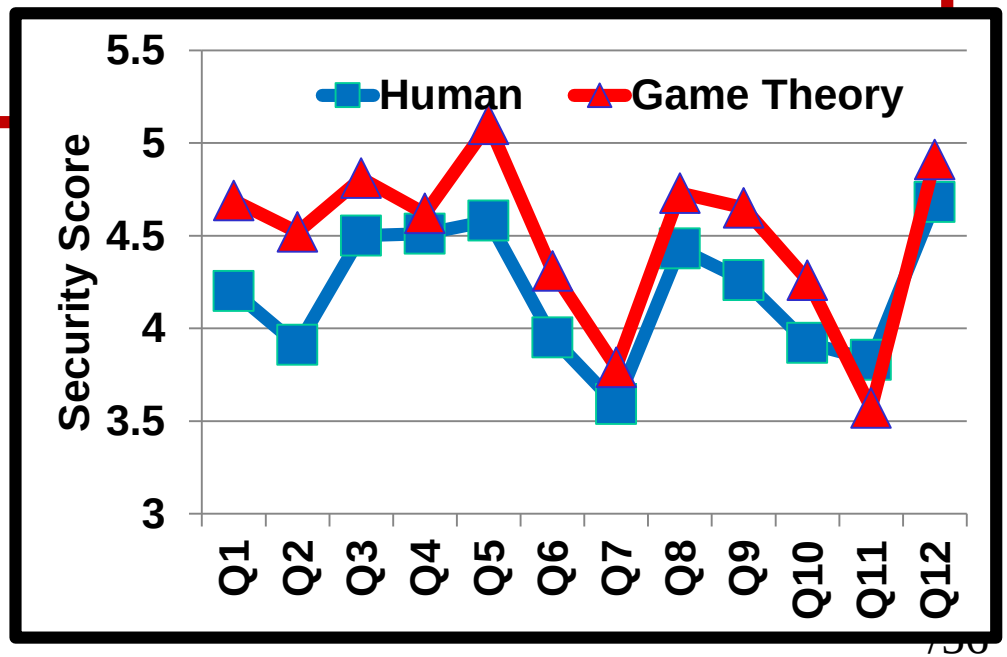

Field Test Against Adversaries: Mock Attackers Example from PROTECT

- "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"

Field Tests Against Adversaries
Computational Game Theory in the Field

Controlled

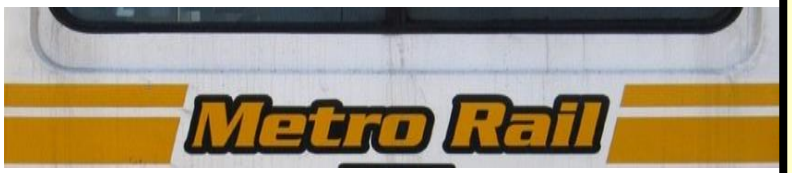

- *Game theory vs Random*
- *21 days of patrol*
- *Identical conditions*
- *Random + Human*

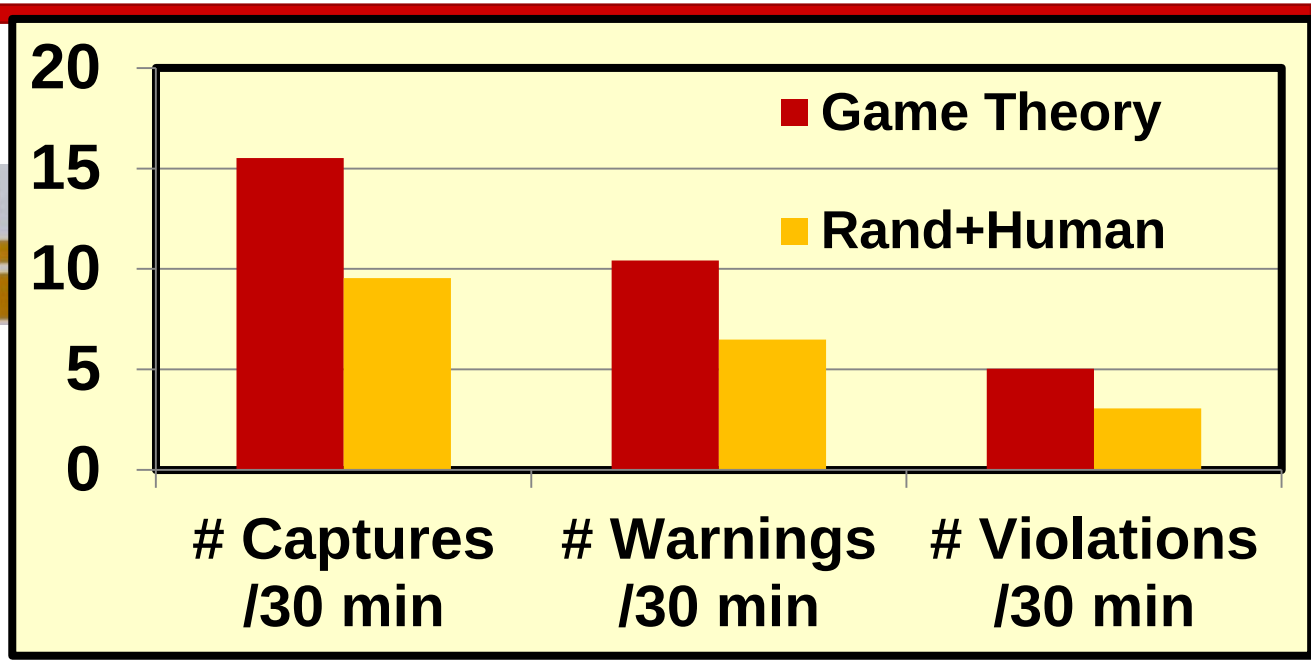

Not controlled

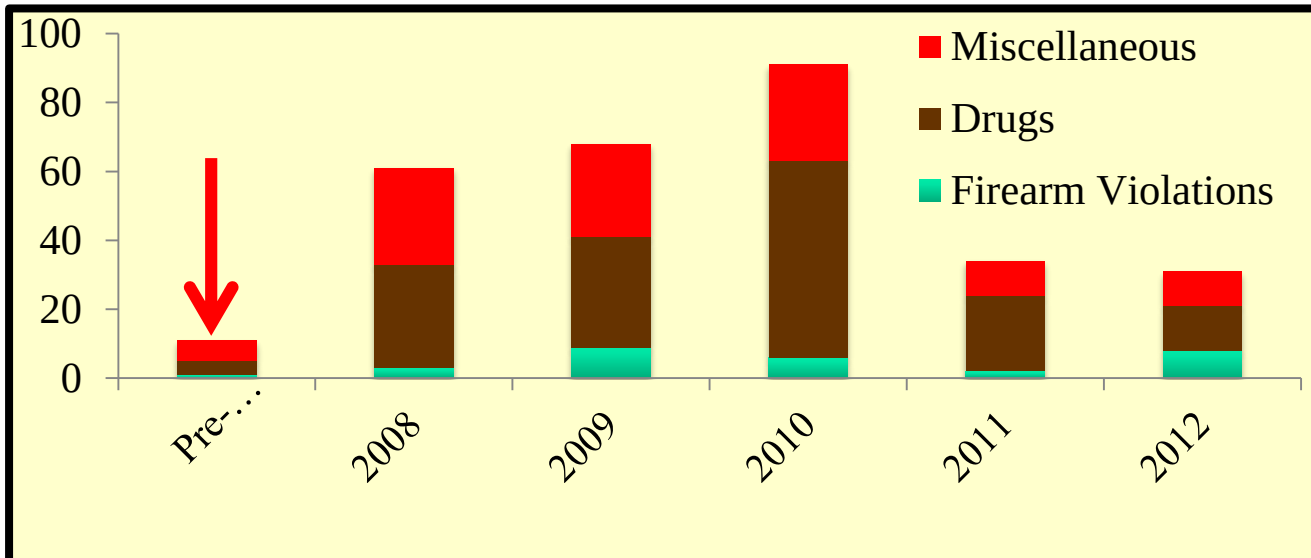

User Feedback
Example from ARMOR, IRIS & PROTECT

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)

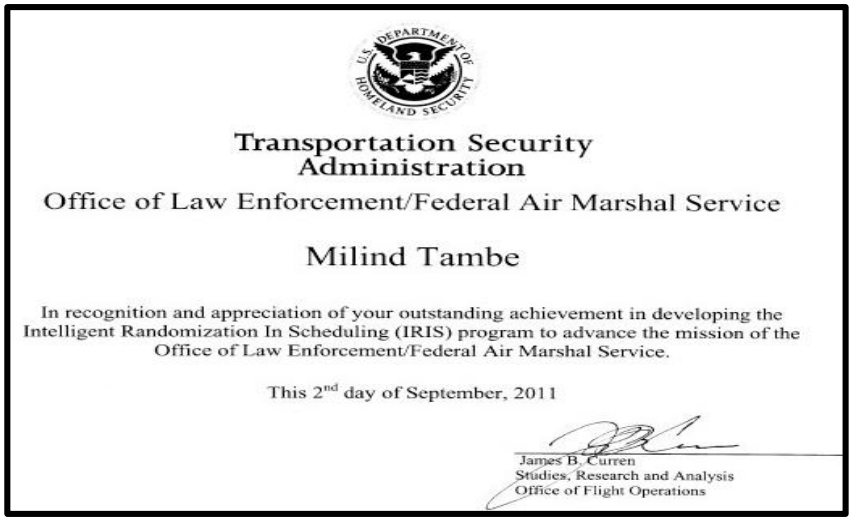

Transportation Security Administration

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)

Global Efforts on Security Games: Yet Just the Beginning…

MILIND TAMBE'S ARMOR AND ITS MANY ITERATIONS ARE USED AROUND THE WORLD TO PROTECT AGAINST TERRORISM, POACHERS, ILLEGAL FISHING AND OTHER THREATS.

SUCCESSFULLY TESTED

Gulf of Mexico (Near Corpus Christi, Texas) — ARMOR-FISH
ARMOR-FISH intelligently randomized schedules for U.S. Coast Guard aerial patrols to thwart the illegal fishing of decimated shark and red snapper populations. (2014)

Los Angeles Metro — TRUSTS
The Los Angeles Sheriff's Department, which LA Metro subcontracts for security, employed TRUSTS to intelligently randomize patrol schedules to stop fare evasion. The Sheriff's Department later ran preliminary experiments to ascertain effectiveness in deploying scarce police personnel to deter crime and terrorism on LA Metro. (2011–2013)

Uganda — PAWS
Ugandan rangers tested PAWS at Queen Elizabeth National Park to intelligently randomize patrols to prevent the slaughte animals, including Cape buffalo, waterbu and giant forest hogs, which are served u locally and exported as "bush meat." (20

Malaysia — PAWS
Panthera, an NGO that is committed to ensuring the survival of tigers and other wild cats, in conjunction with the nonprofit group Rimba, began testing PAWS in forests in northeastern Malaysia, to evaluate its ability to generate effective patrols in the challenging, hilly terrain. (2014)

Vietnam, Cambodia, Bangladesh, Indonesia — PAWS

aten Island Ferry — PROTECT
OTECT provides protection to the Staten and Ferry, which carries up to 4,000 ssengers at peak times.

s Angeles International Airport — ARMOR
RMOR intelligently randomizes schedules of eckpoints along the five roads that lead into e airport.

U.S. Air Traffic — IRIS
As part of its multipronged strategy to prevent attacks, the Transportation Security Administration (TSA) has since 2009 deployed Milind Tambe's IRIS system, which intelligently randomizes federal air marshals' flight schedules to make their air patrols unpredictable to would-be malefactors.

agascar — PAWS
d Tambe, working with Meredith Gore, an ciate professor of conservation social nces at Michigan State University, and a gasy civil society group called Alliance

Voahary Gasy (AVG), hopes to eventually employ PAWS in Madagascar to randomize patrol schedules for rangers, police and national park officials to reduce environmental crimes, especially illegal logging.

Thank you to sponsors:

CREATE HOMELAND SECURITY CENTER

Transportation Security Administration

Army Research Office

Research & Development Center · United States Coast Guard

THANK YOU

tambe@usc.edu
http://teamcore.usc.edu/security