SAND2021-15127C

Sandia National Laboratories

Stochastic Bayesian Games for the Cybersecurity of Critical Infrastructures

Lee T. Maccarone, PhD – Sandia National Laboratories

Daniel G. Cole, PhD – University of Pittsburgh

MORS Emerging Techniques Forum

December 6-9, 2021

U.S. DEPARTMENT OF ENERGY | NNSA

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

This research uses game theory to defend against cyber-attacks while considering uncertainty

Who is our adversary?

Disgruntled employee:

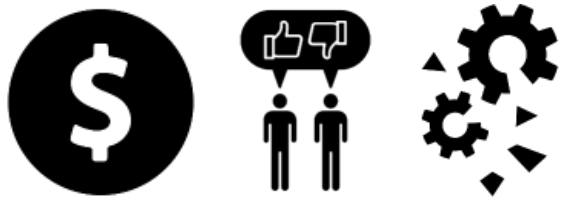

Terrorist:

How effective are our defenses?

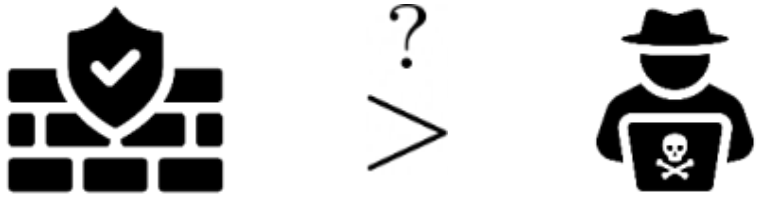

The goal of this research is to reduce the likelihood of successful cyber-attacks on nuclear power plants

1. Predict how an adversary might target a plant
2. Quantify nuclear power plant cybersecurity
3. Allocate cybersecurity resources to defend a plant

Stochastic game theory is used to analyze interactions where the outcome is uncertain

Bayesian games address uncertainty about the adversary

Type 1: Disgruntled Employee

	Attack 1	Attack 2
Defense 1	0, 2	-10, 1
Defense 2	-3, 10	0, 2

Probability = 0.2

Type 2: Terrorist

	Attack 1	Attack 2
Defense 1	0, -1	-10, 10
Defense 2	-3, 3	0, 4

Probability = 0.8

Stochastic Bayesian games consider uncertainty regarding the players and their interactions

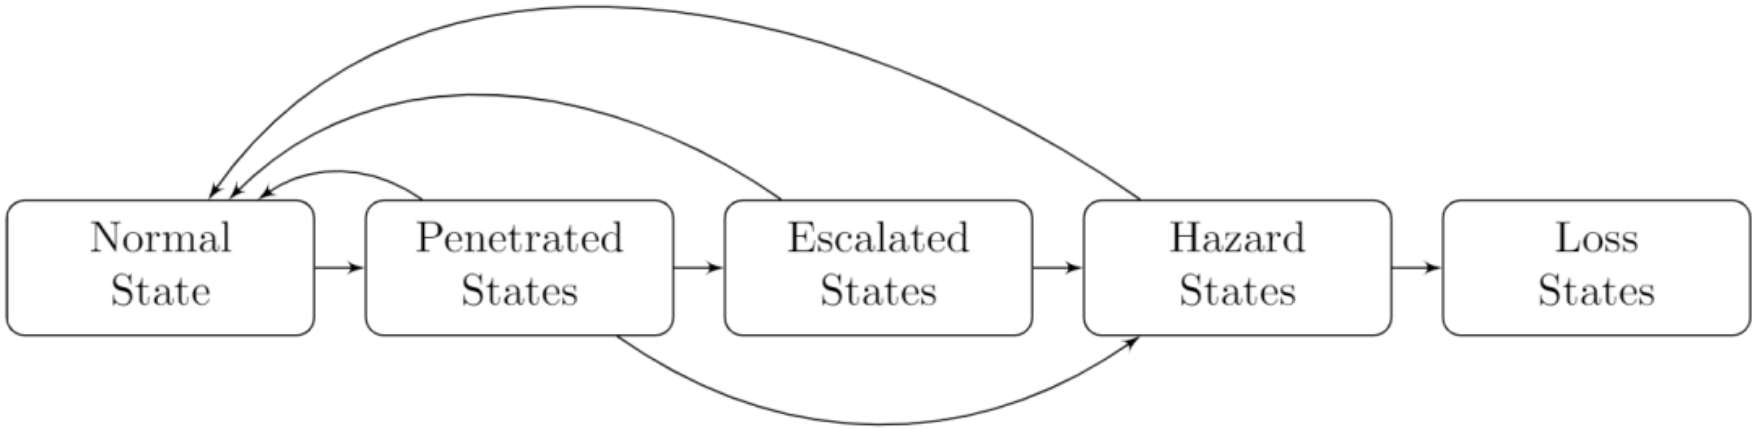

Adversary Types

- Disgruntled employee
- Government cyberwarrior
- Radical activist
- Terrorist

This research consists of two major tasks:

Game construction

Game solution

Approach is demonstrated for a notional system & notional adversaries

The residual heat removal system maintains reactor water level during a loss of coolant accident

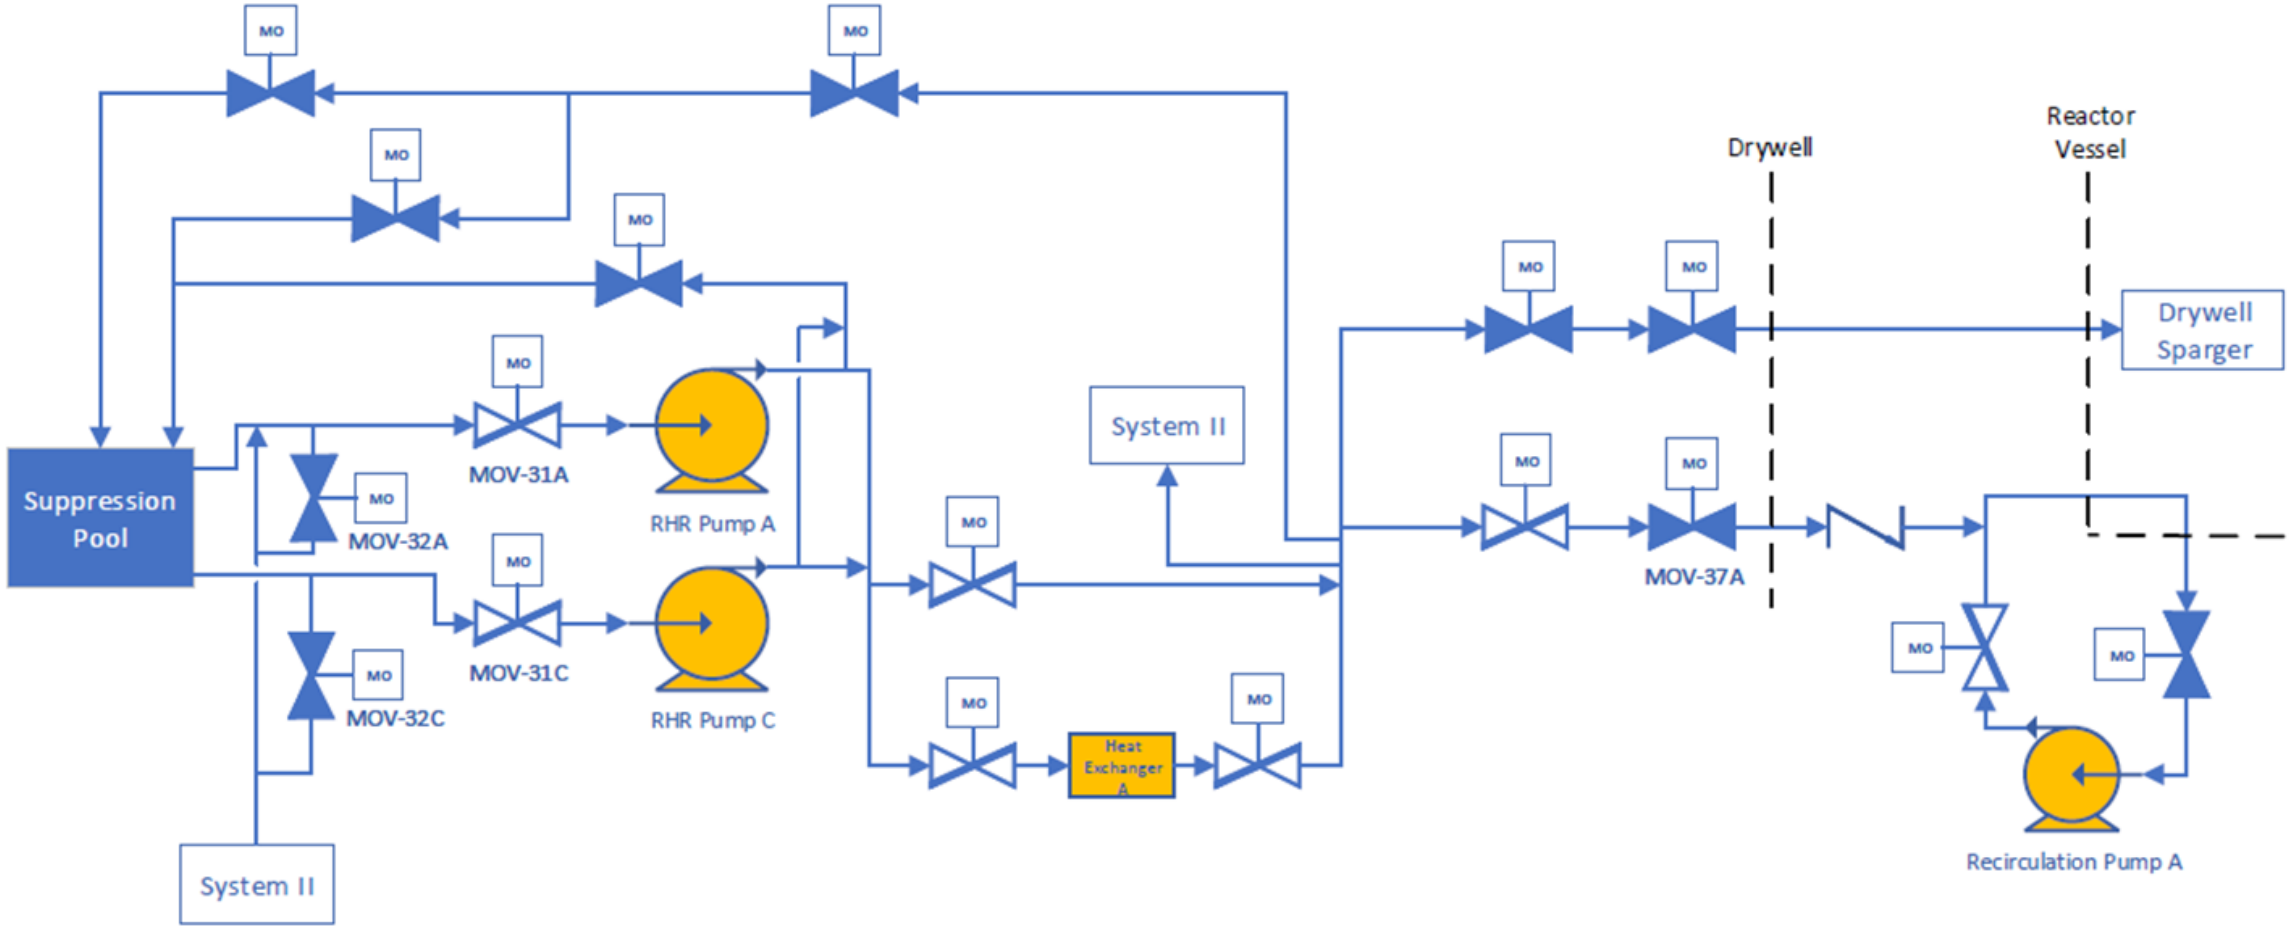

Adversary types were selected using Intel Corporation's Threat Agent Risk Assessment methodology

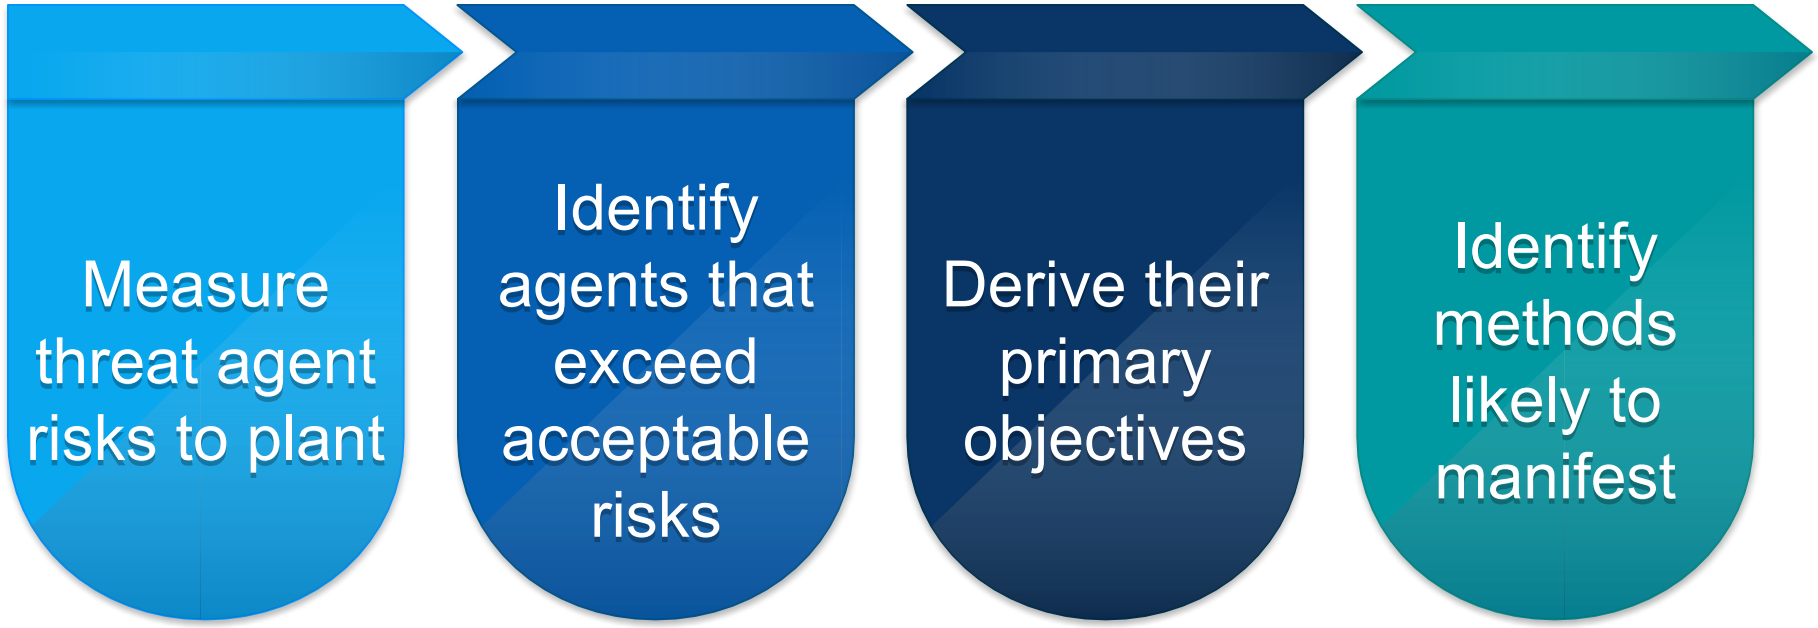

The radical activist, disgruntled employee, government cyberwarrior, and terrorist are the greatest risks

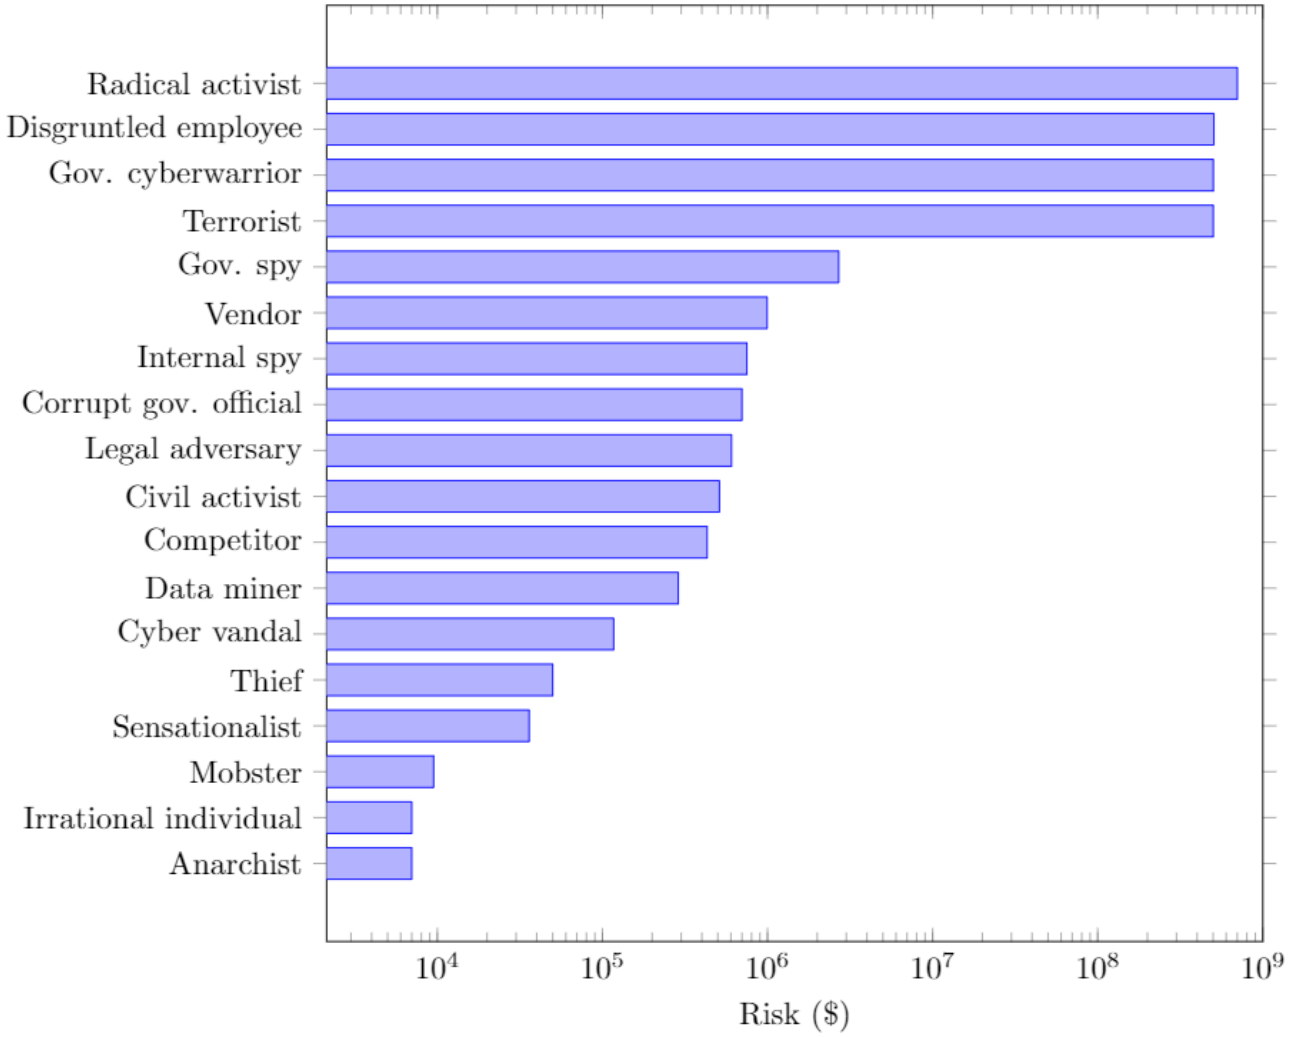

The states of the game were defined using System-Theoretic Process Analysis

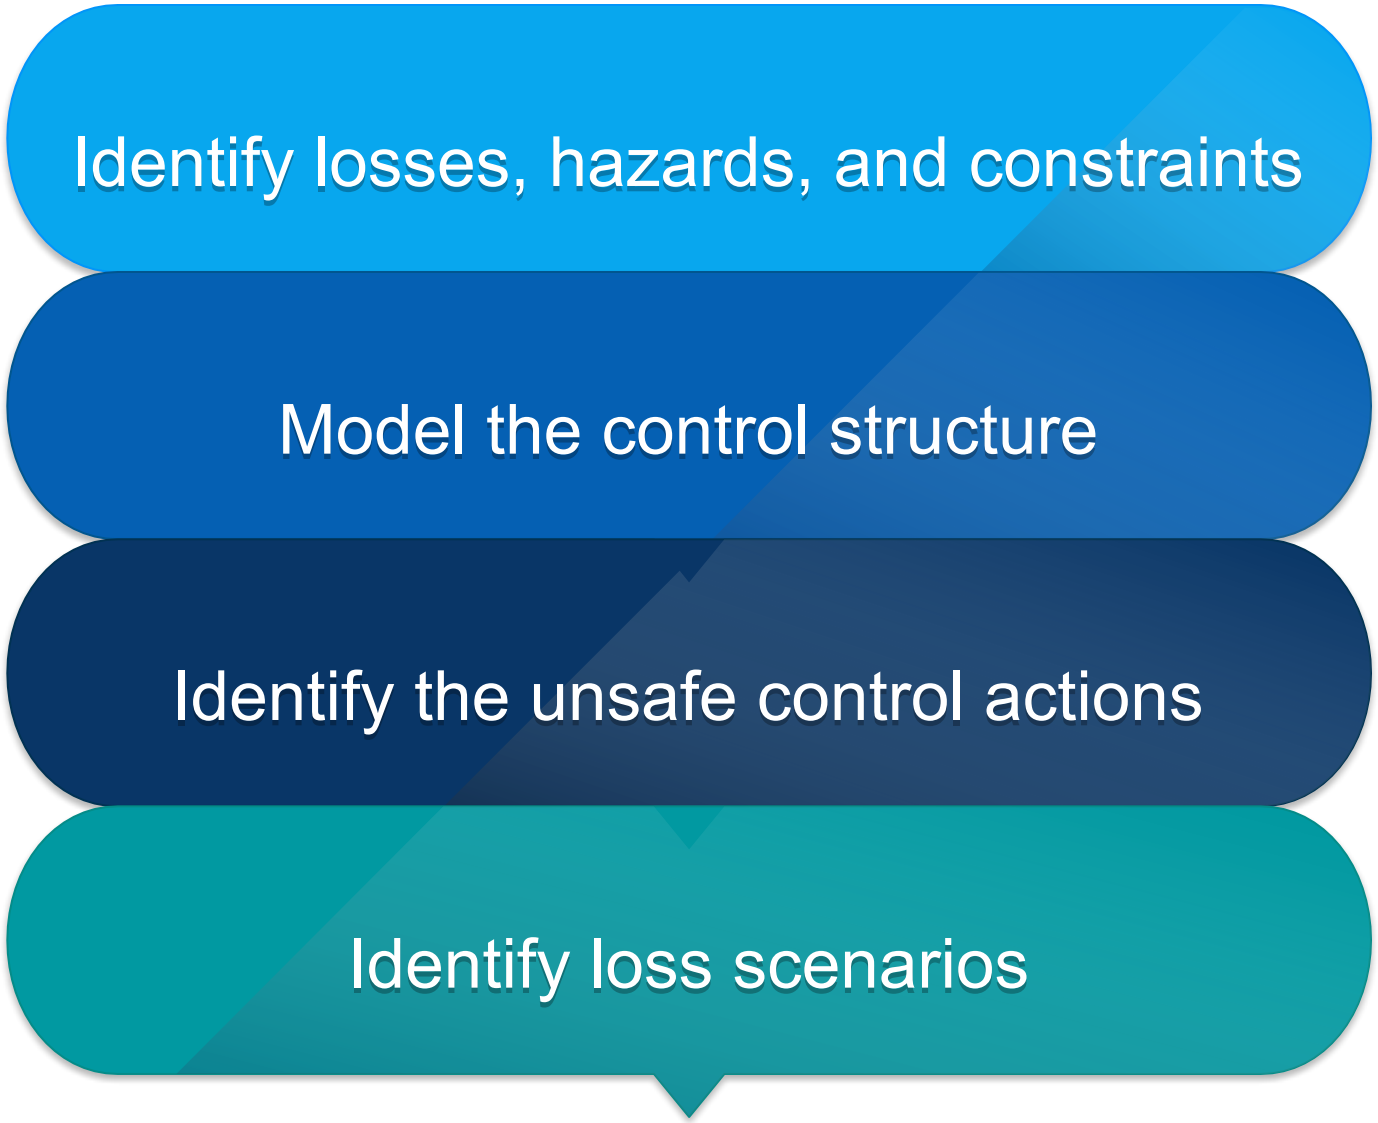

The state space consists of 45 states

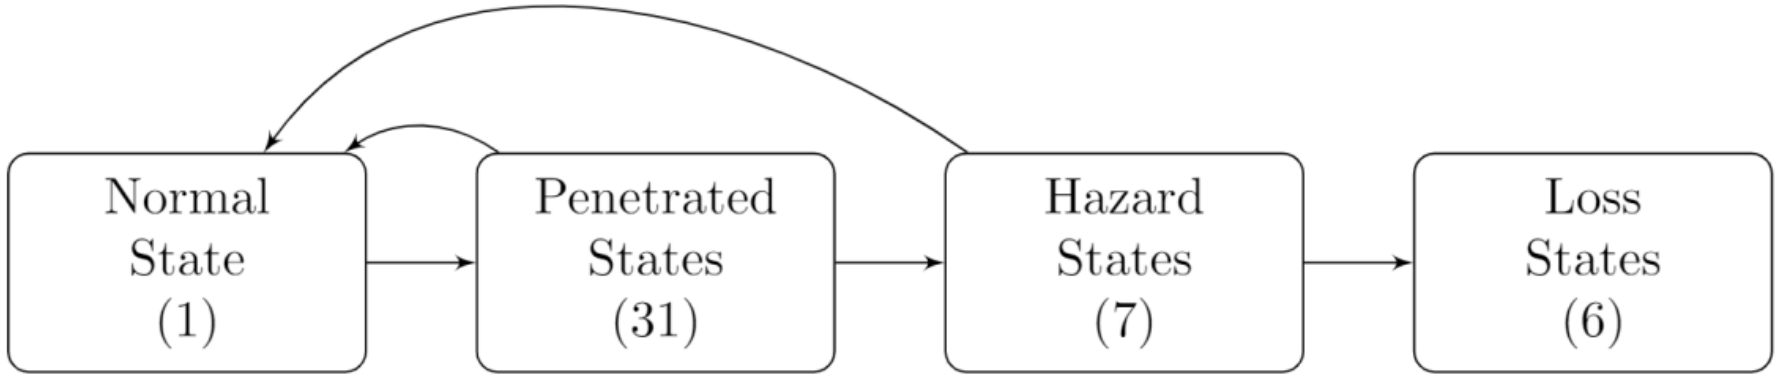

Hazards:

1. Loss of flow path
2. Damage to RHR pumps
3. Removal of suppression pool inventory
4. Reactor trip
5. RHR does not initiate
6. Inadequate flow
7. Cooling provided when not needed

Losses:

1. Loss of power generation
2. Environmental damage
3. Personnel injury
4. Damaged public opinion
5. Major equipment damage
6. Core damage

State transition probabilities were estimated using the Common Vulnerability Scoring System

CVSS Exploitability Metrics

1. **Attack vector**
2. **Attack complexity**
3. **Privileges required**
4. **User interaction**

$$p(s_j|s_i, a_D, a_A)$$

The attacker chooses between the best accessible hazard or pursuing a better hazard

H^* = most desired hazard
H = best accessible hazard

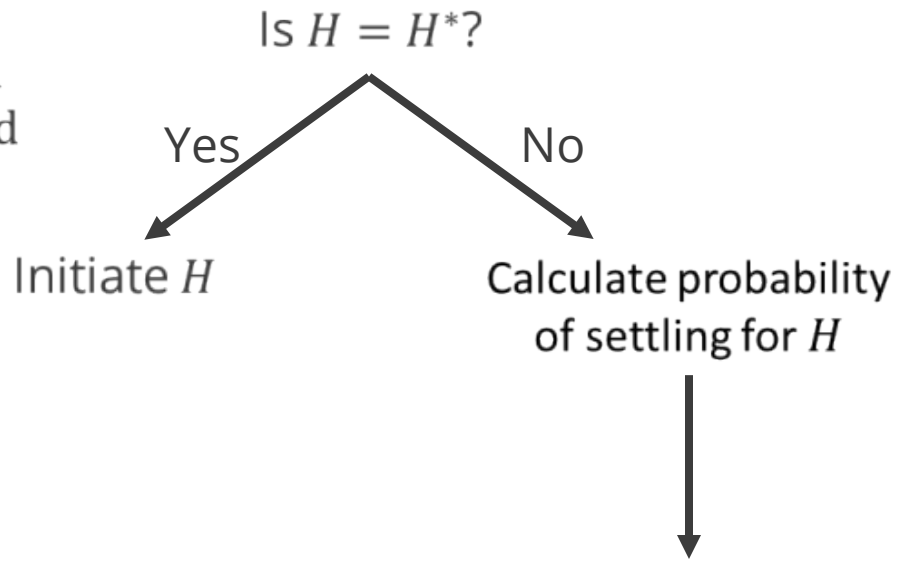

$$p(\text{action}) = p(\text{not settling}) \sum_{\text{Hazards}} p(\text{action}|\text{hazard})p(\text{hazard}|\text{not settling})$$

The defender uses Bayesian learning to estimate the attacker's utility of a loss

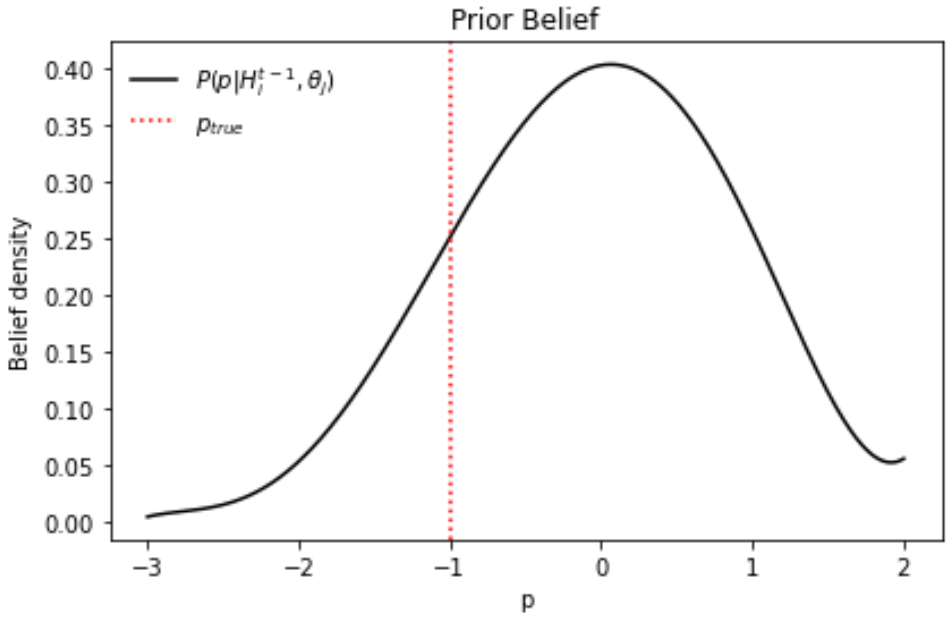

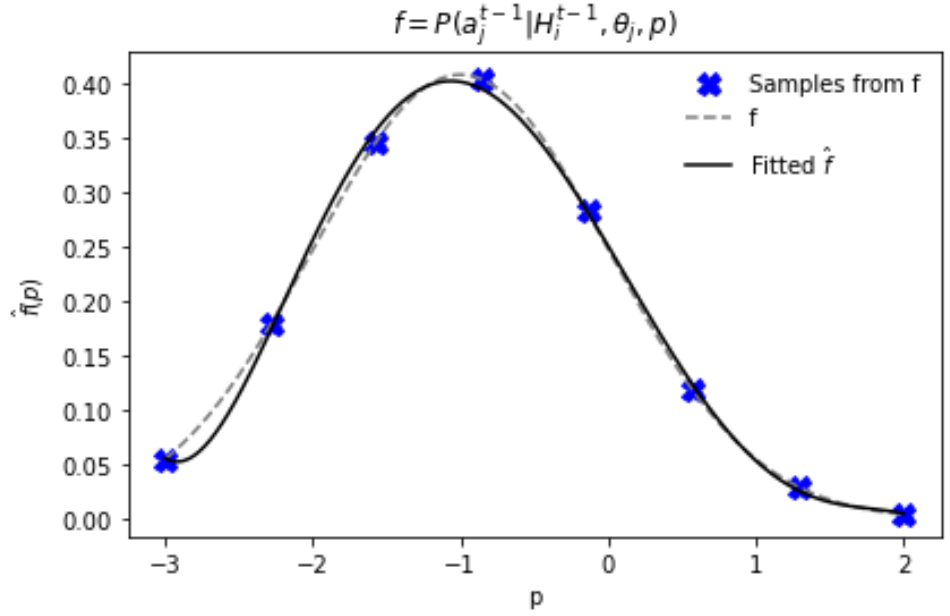

The defender uses Bayesian learning to estimate the attacker's utility of a loss

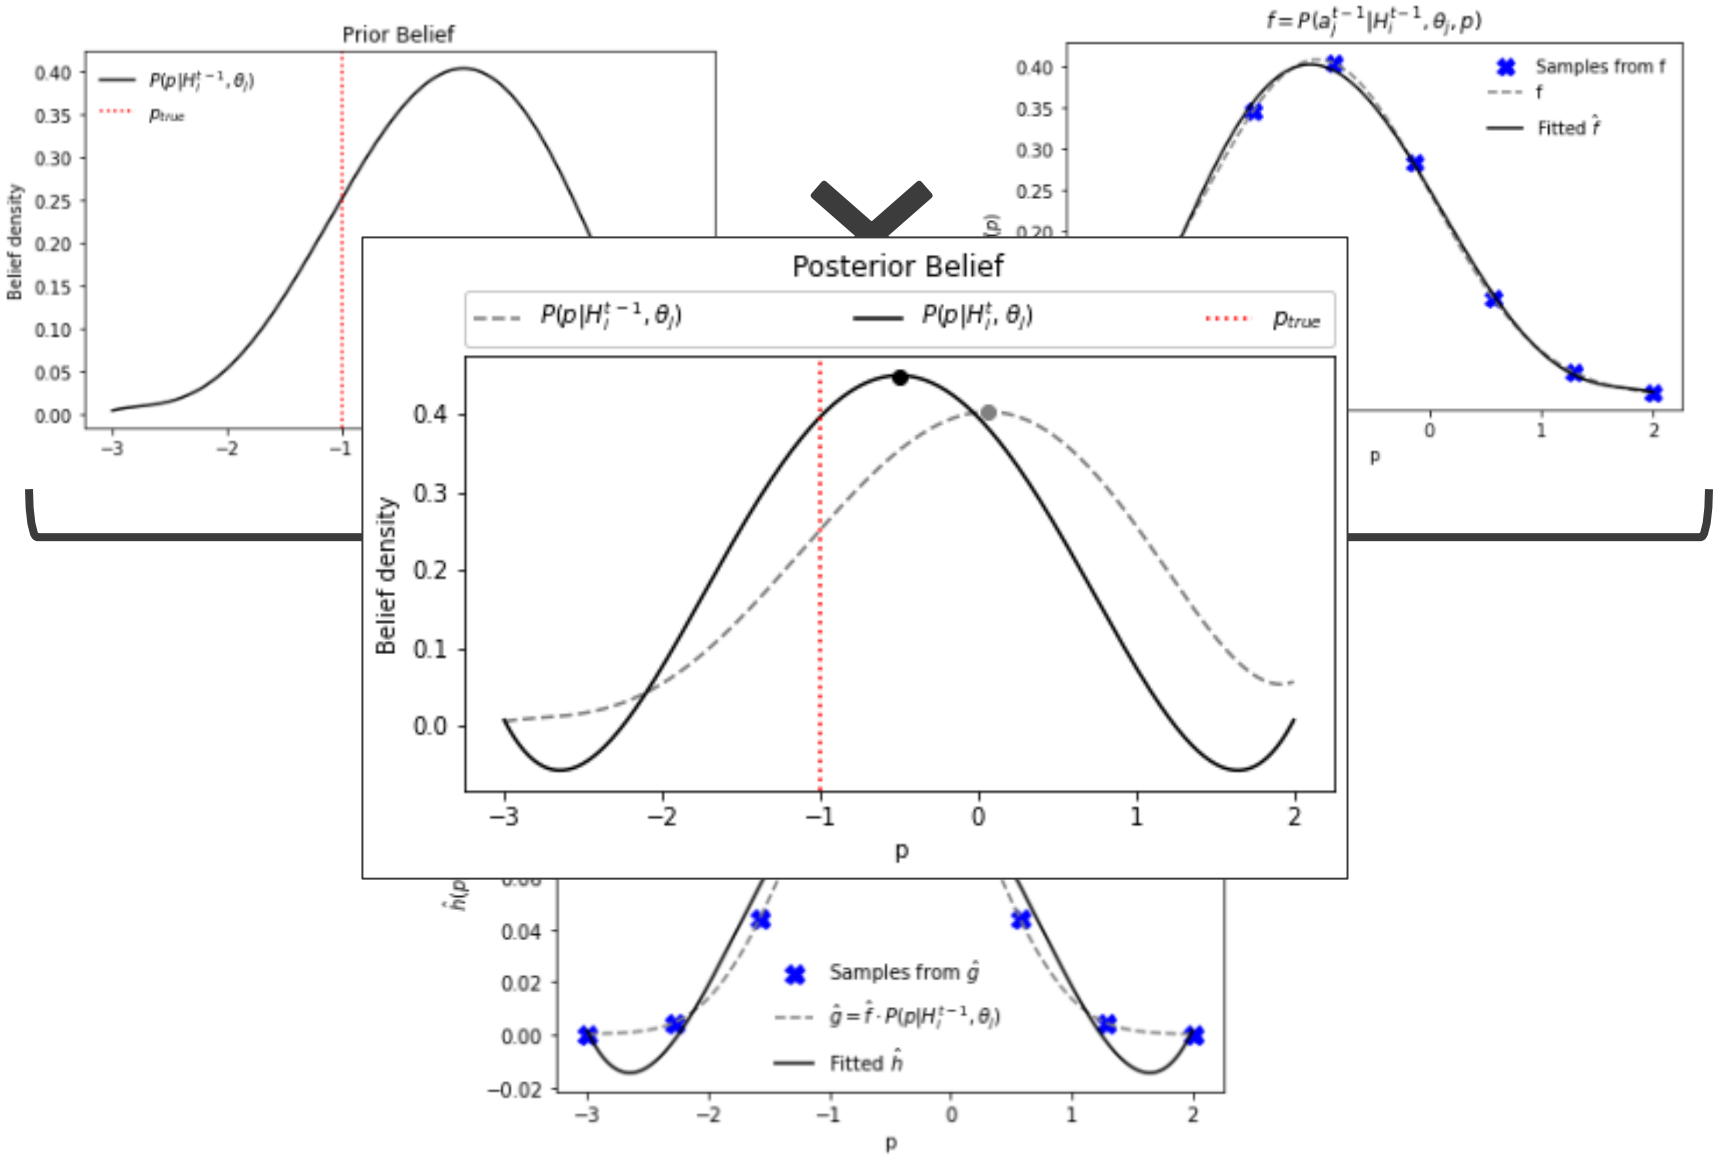

The defender uses Harsanyi-Bellman ad hoc coordination to select a defensive action

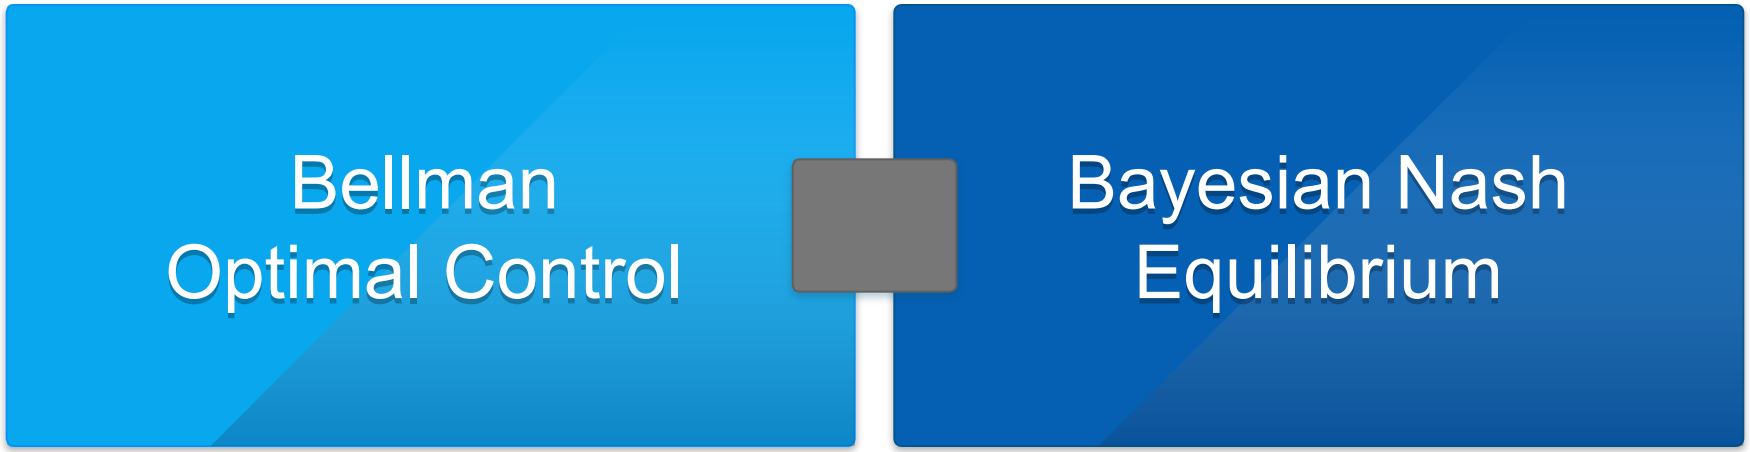

Advantage: Situational strategy selection

Challenge: Computationally infeasible for this game

HBA can be approximated with path sampling methods

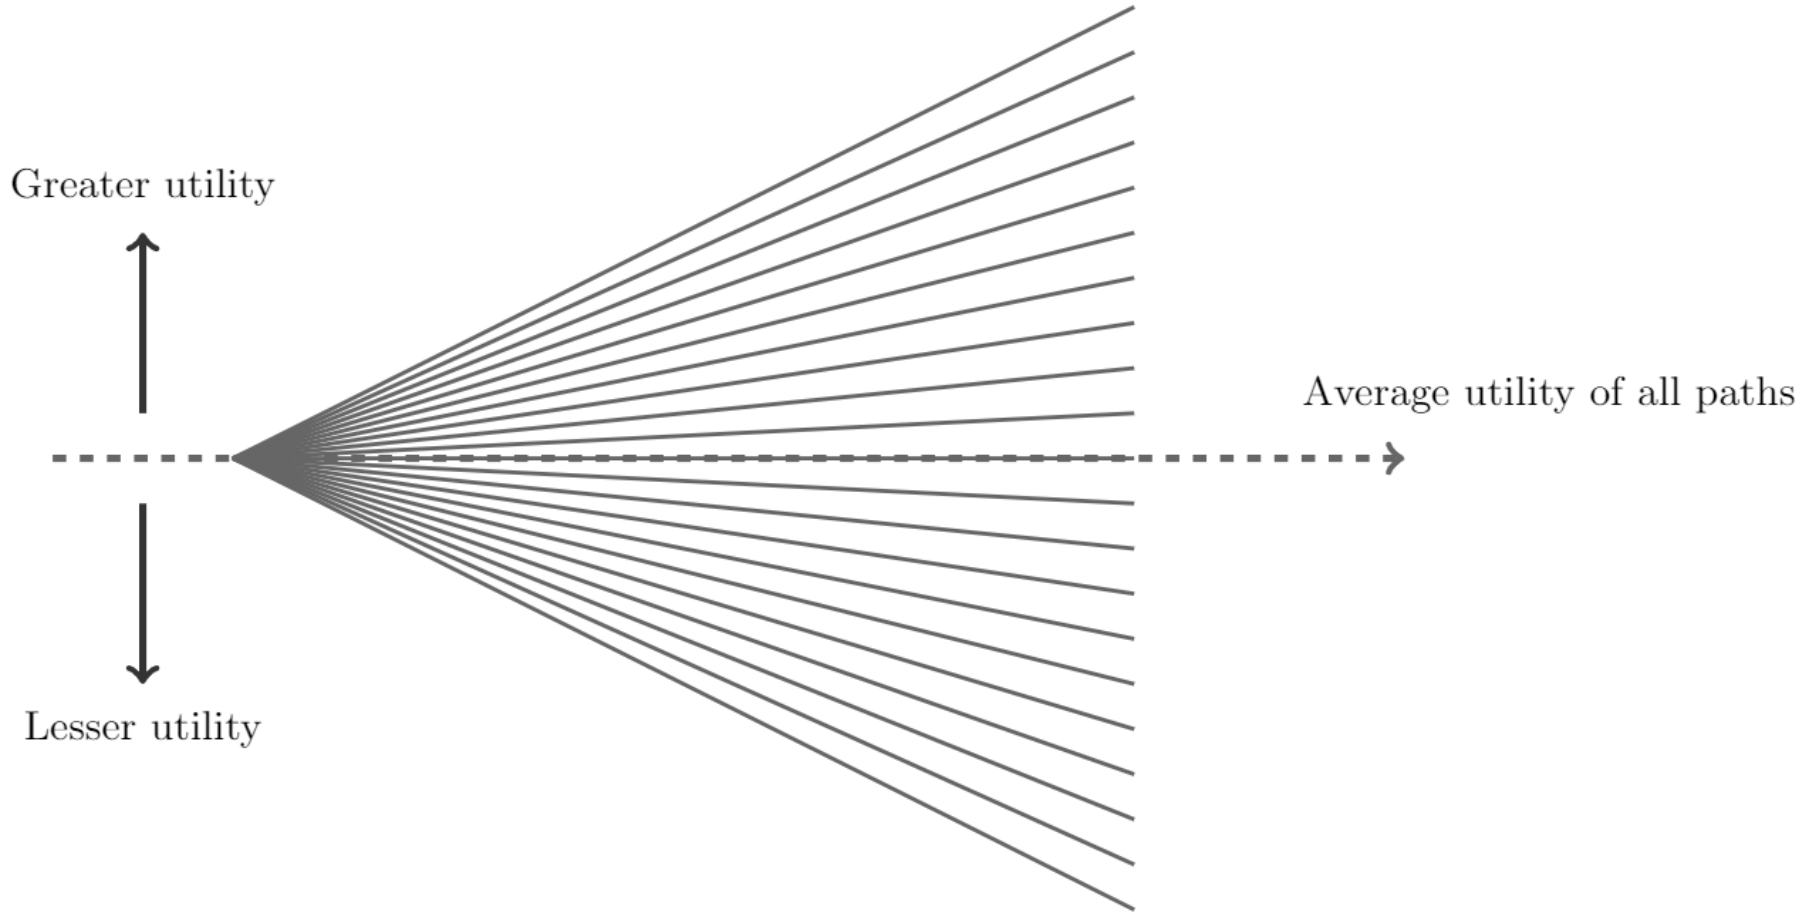

Rewards are aggregated throughout the game to quantify the outcome for each player

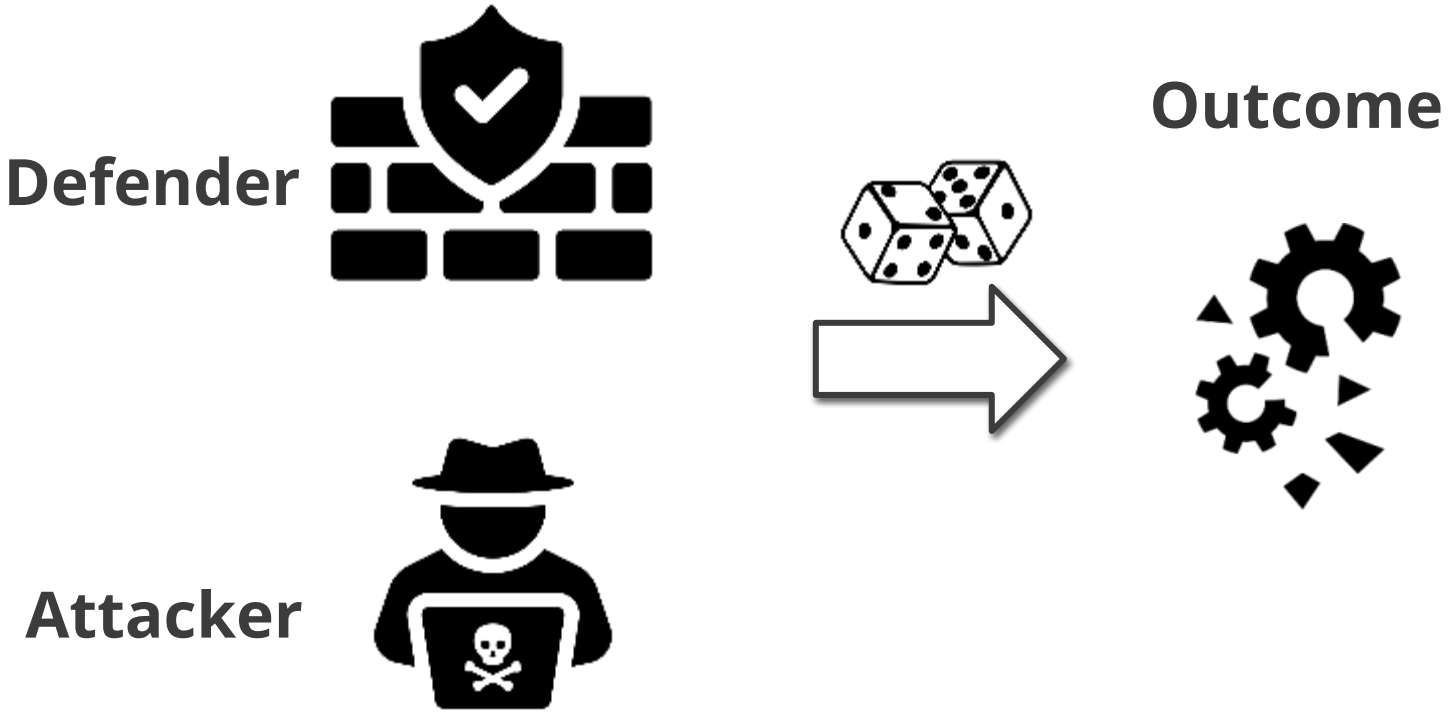

The players continue to select actions until either a loss occurs or the time limit is reached

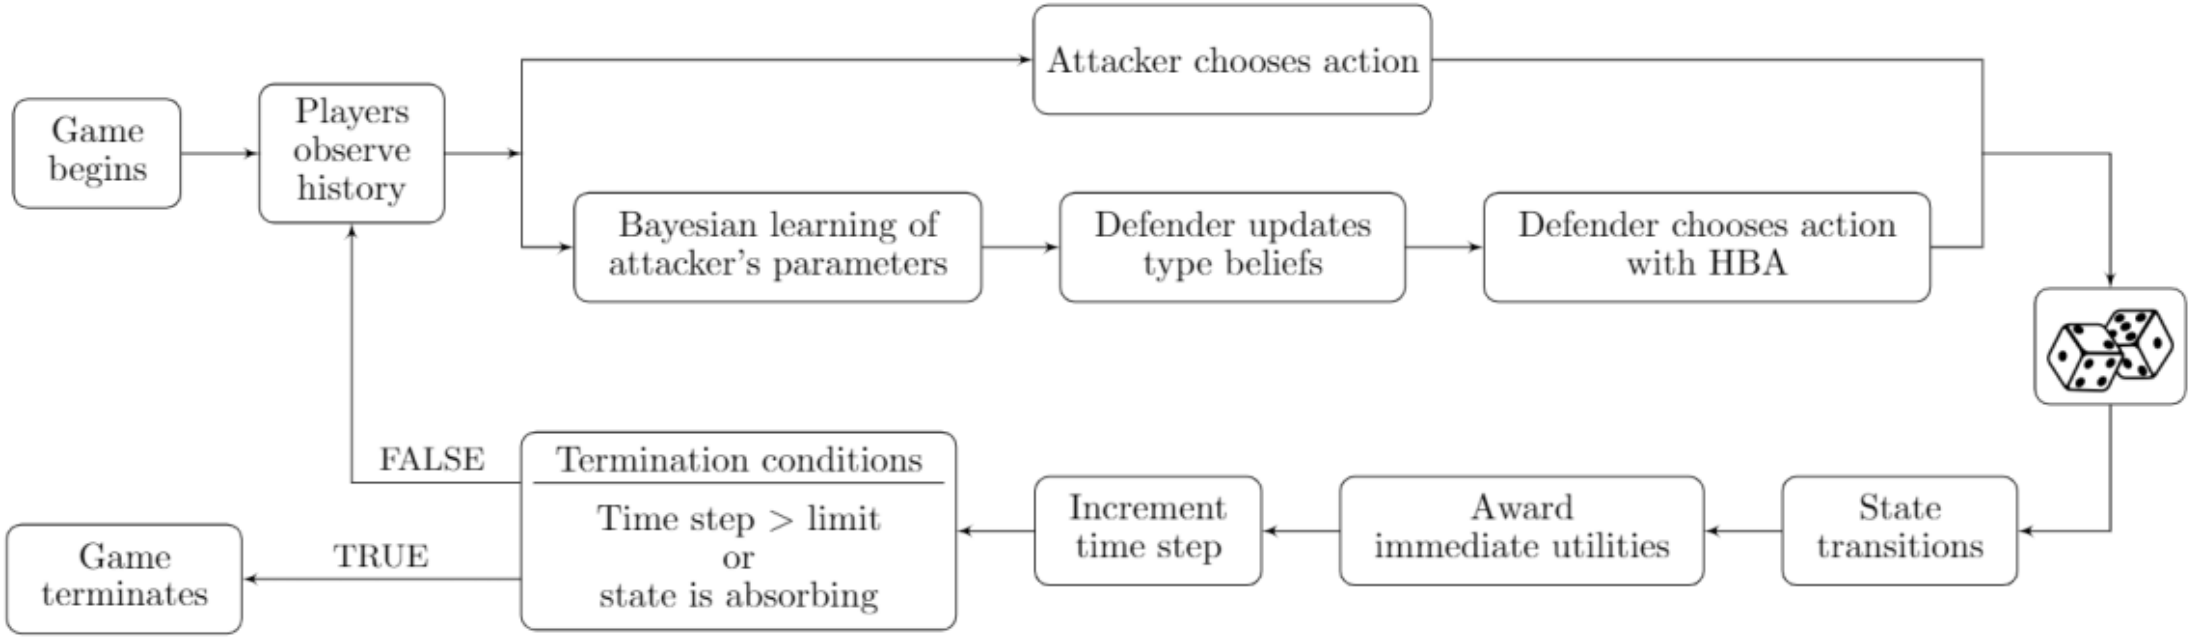

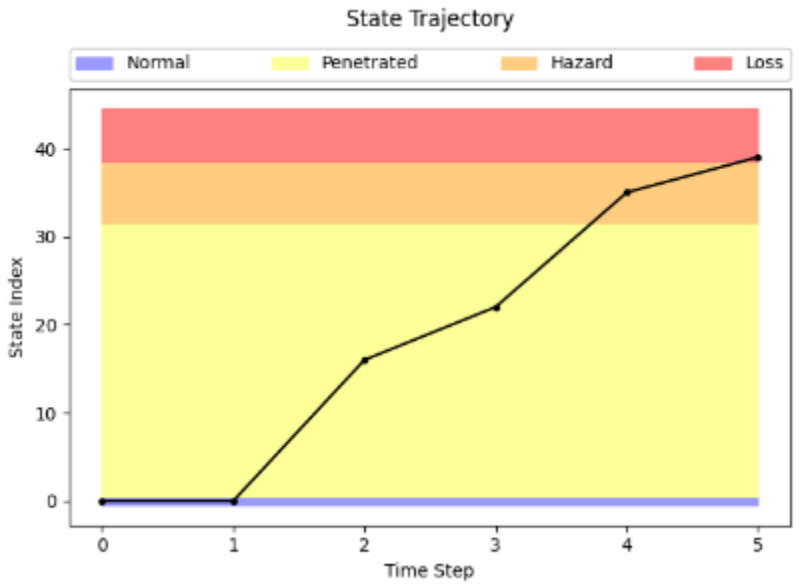
State Trajectory
Normal
Penetrated
Hazard
Loss
State Index
Time Step

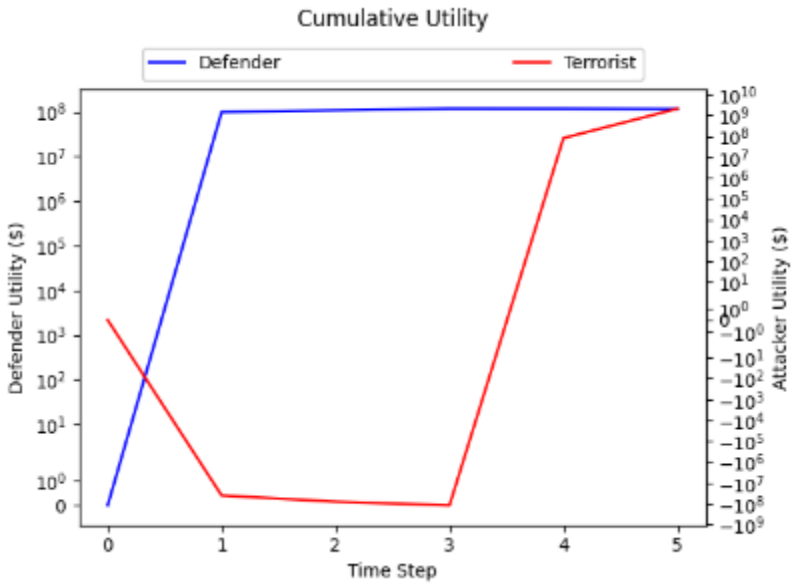
Cumulative Utility
Defender
Terrorist
Defender Utility ($)
Attacker Utility ($)
Time Step

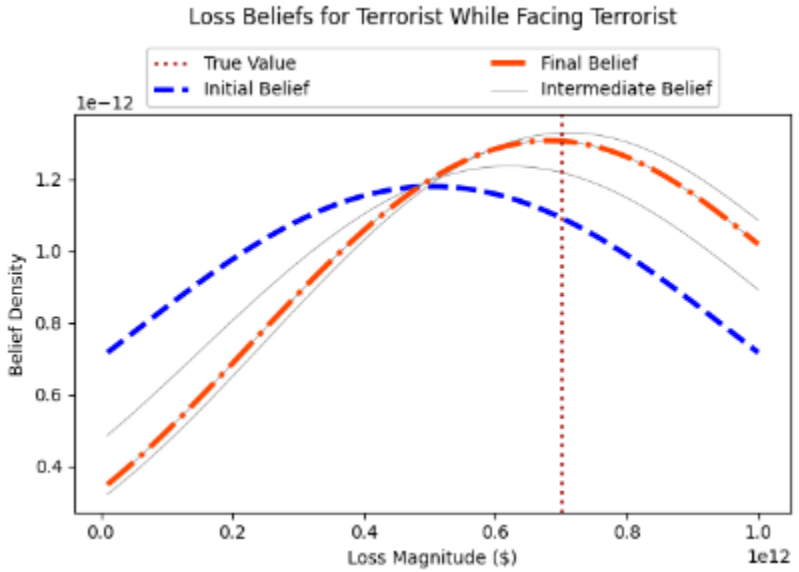
Loss Beliefs for Terrorist While Facing Terrorist
True Value
Initial Belief
Final Belief
Intermediate Belief
1e−12
Belief Density
Loss Magnitude ($)
1e12

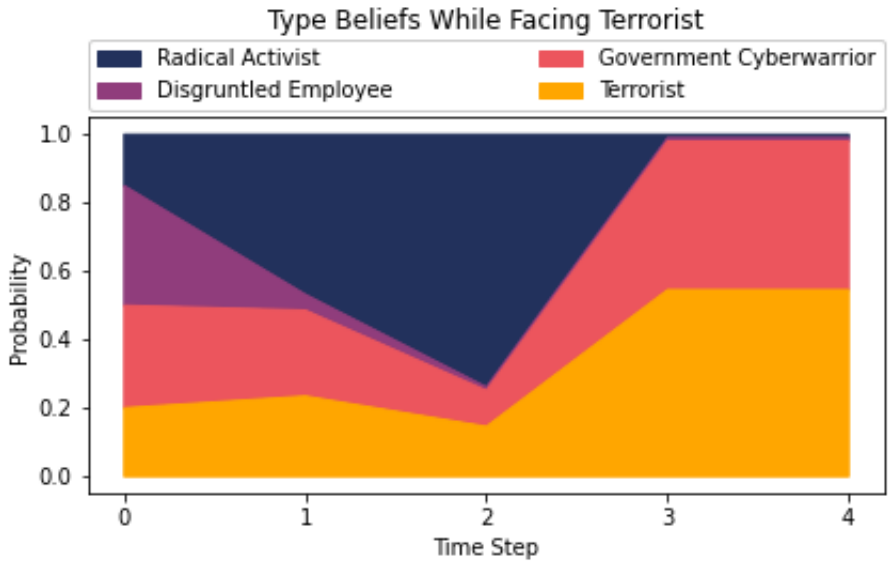
Type Beliefs While Facing Terrorist
Radical Activist
Disgruntled Employee
Government Cyberwarrior
Terrorist
Probability
Time Step

The greatest mean time-to-loss occurred when facing the radical activist

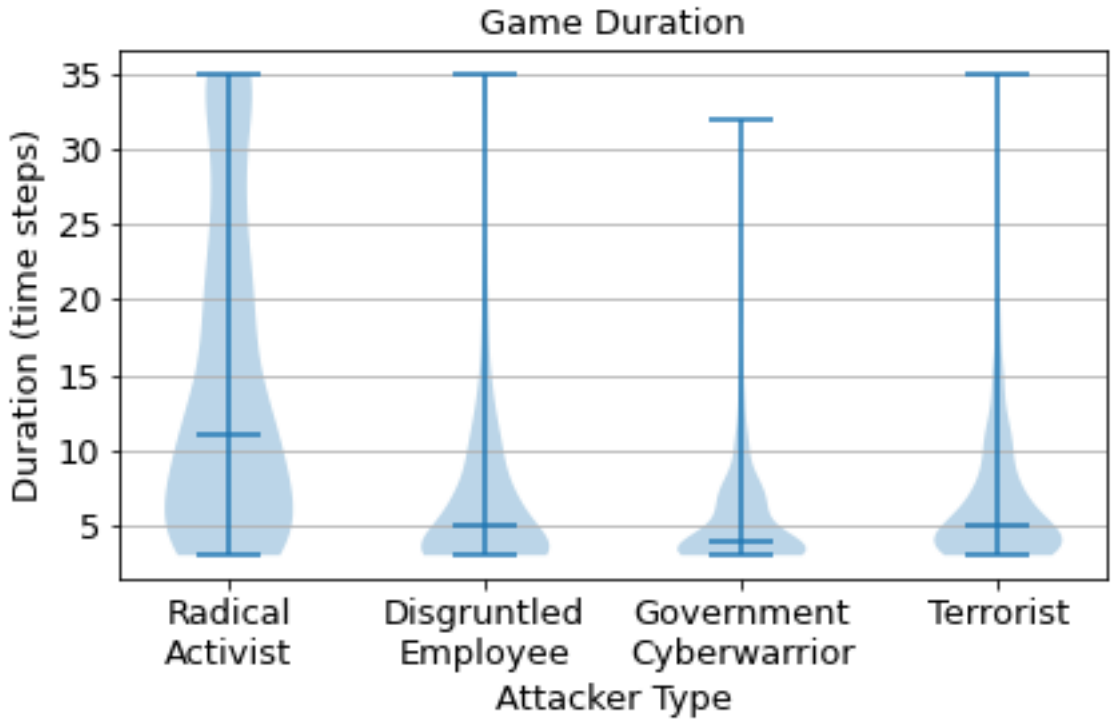

The greatest availability occurs when facing the terrorist

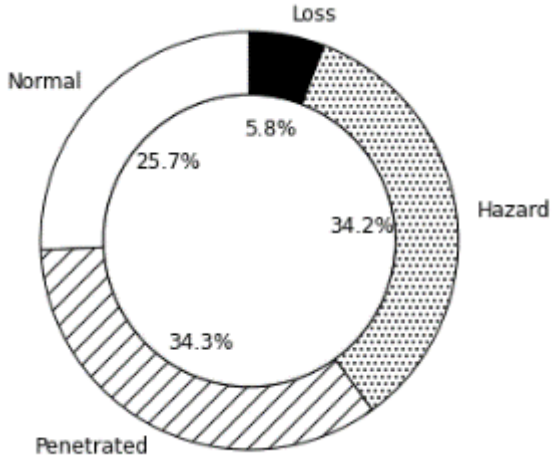

Radical Activist (60.0%)

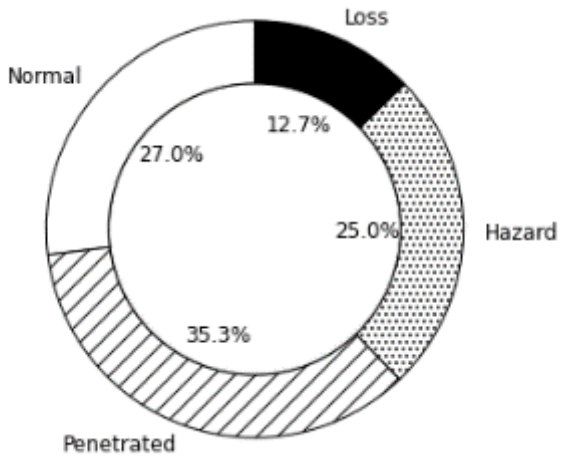

Disgruntled Employee (62.3%)

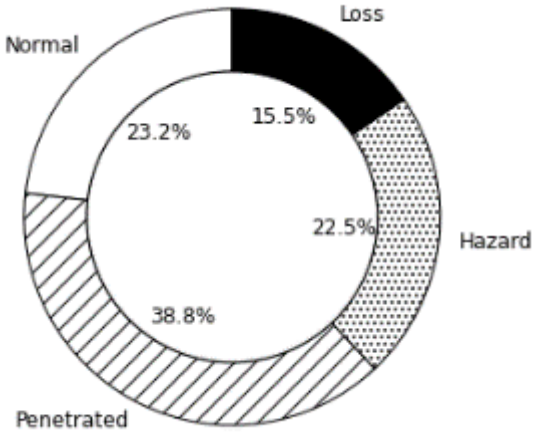

Government Cyberwarrior (62.0%)

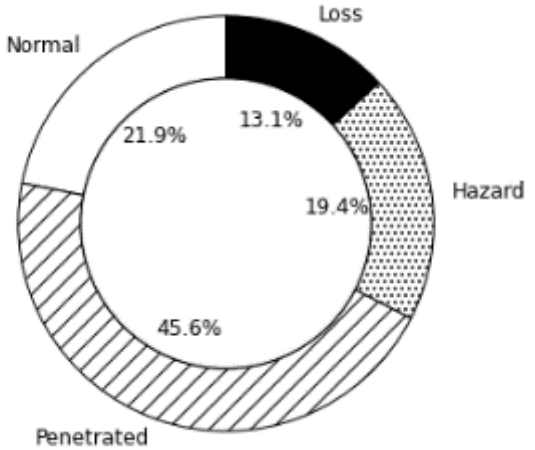

Terrorist (67.5%)

The defender's median utility was greatest when facing the radical activist

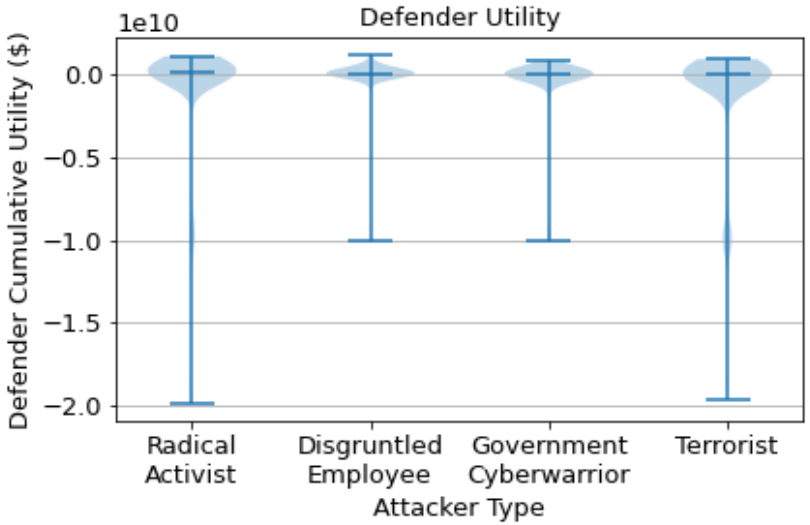

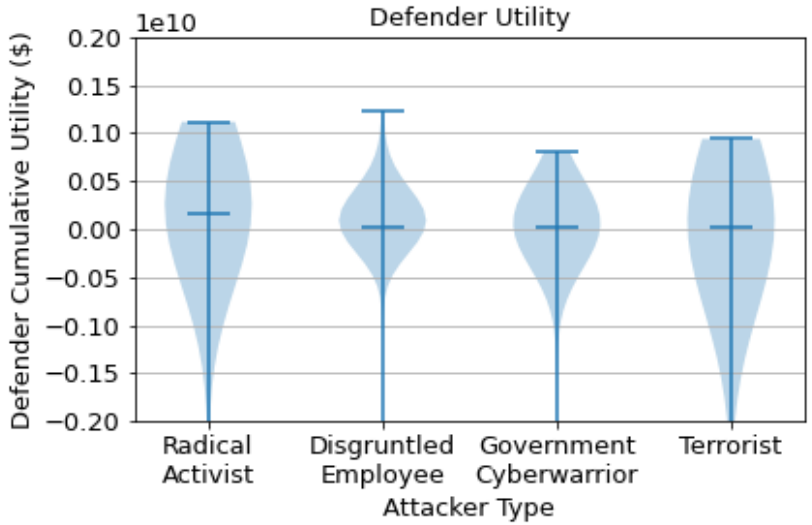

The spread of the attacker's utility varied significantly among the attacker types

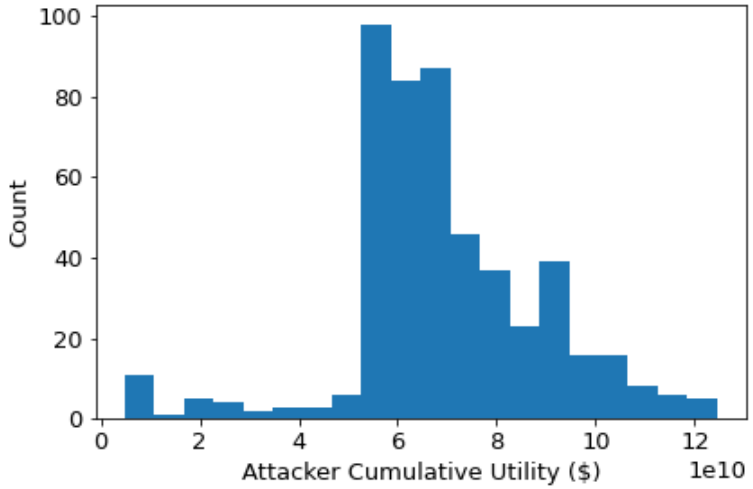

Radical Activist

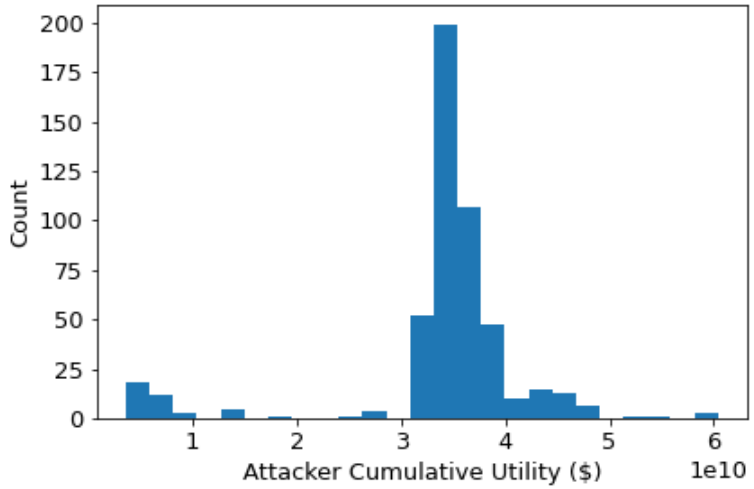

Disgruntled Employee

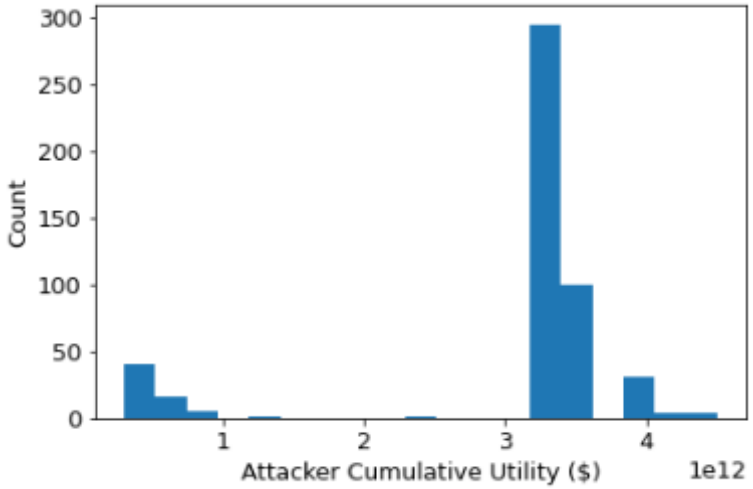

Government Cyberwarrior

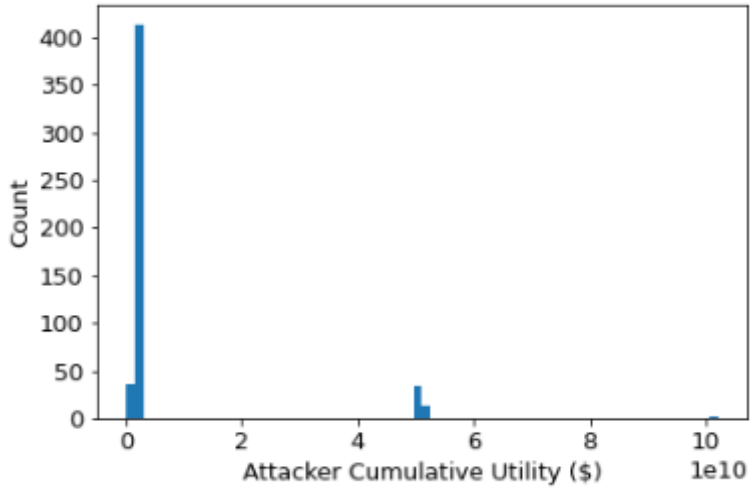

Terrorist

The estimation of the attacker's loss utility was best when facing the terrorist

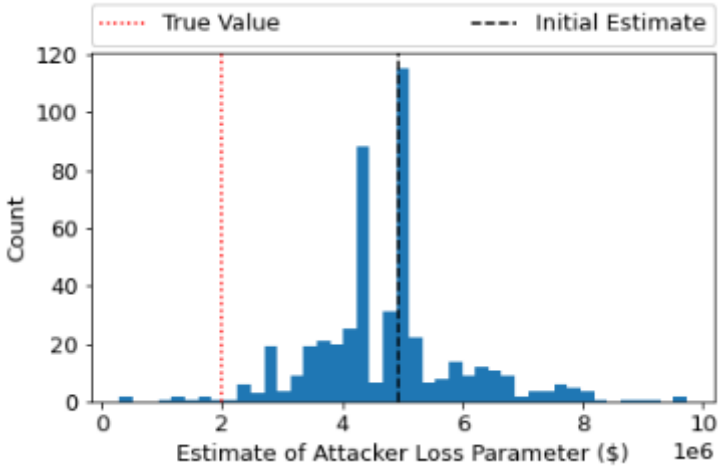

Radical Activist

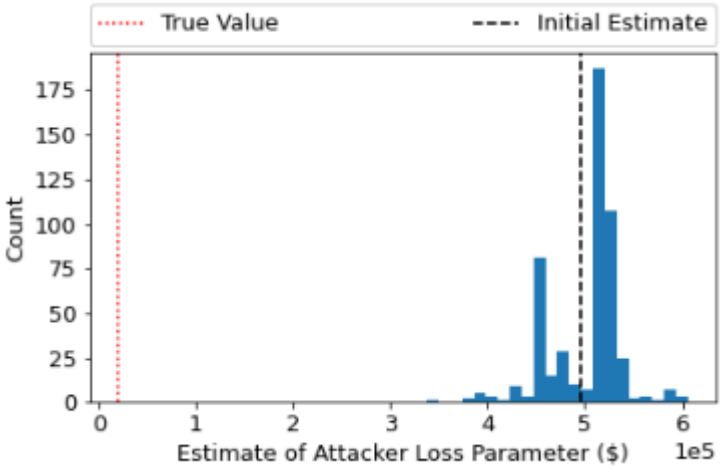

Disgruntled Employee

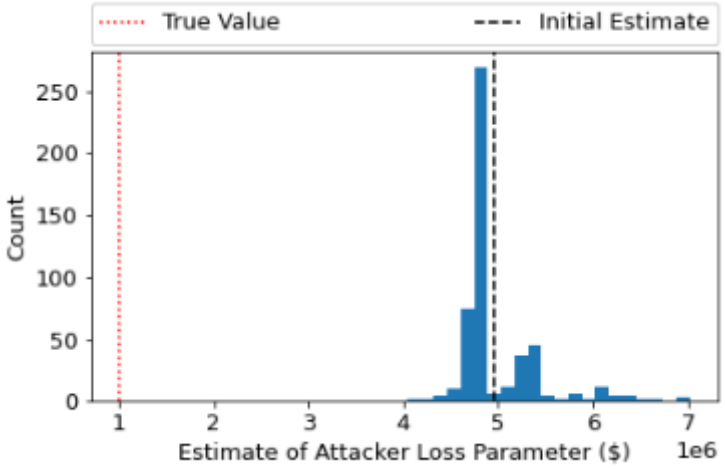

Government Cyberwarrior

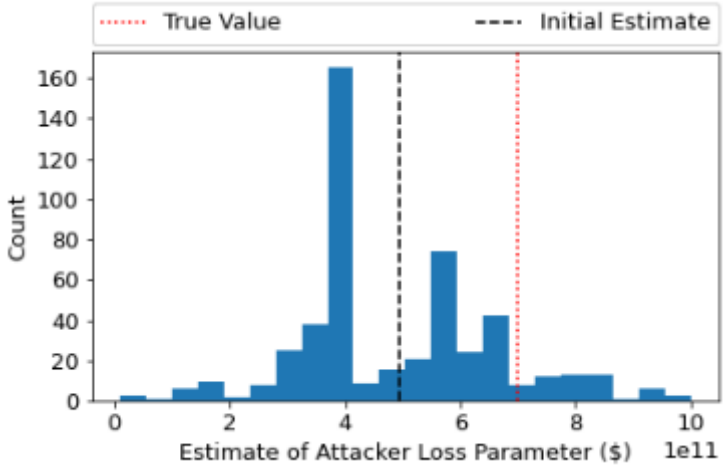

Terrorist

Estimation of the attacker's loss utility was inconsistent for several reasons

1. Short duration of games

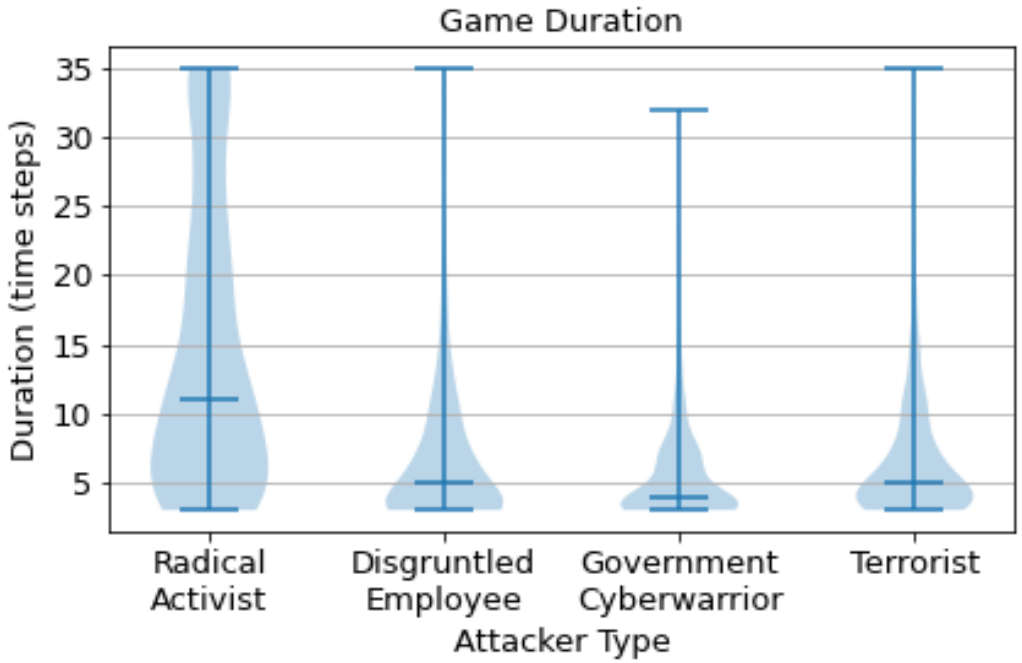

Estimation of the attacker's loss utility was inconsistent for several reasons

1. Short duration of games
2. Lack of new information as the game is played

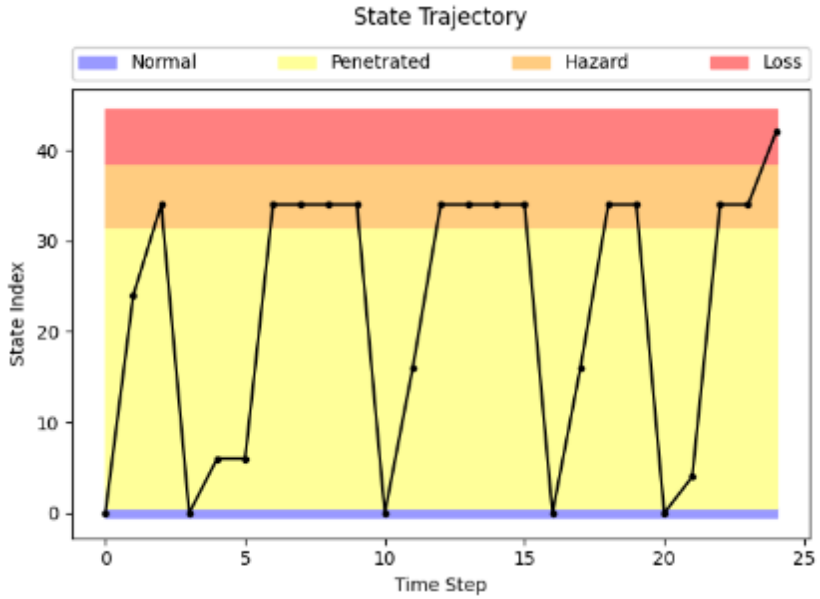

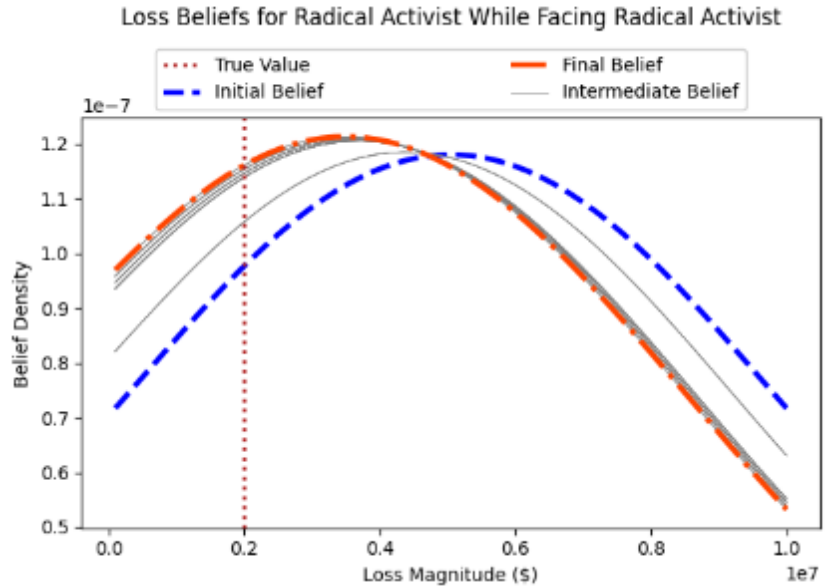

Estimation of the attacker's loss utility was inconsistent for several reasons

1. Short duration of games
2. Lack of new information as the game is played
3. Insensitivity of attacker's decision-making to the loss utility

	Radical Activist	Disgruntled Employee	Government Cyberwarrior	Terrorist
Loss Rank	5th (tie)	6th	4th (tie)	1st

The utility of the uncertain loss affected the terrorist's hazard ranking

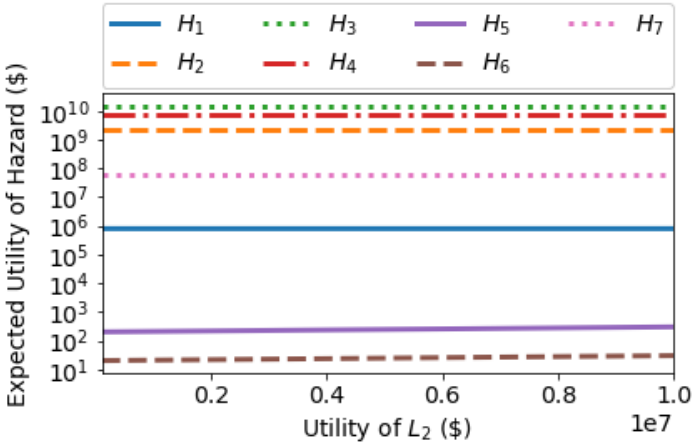

Radical Activist

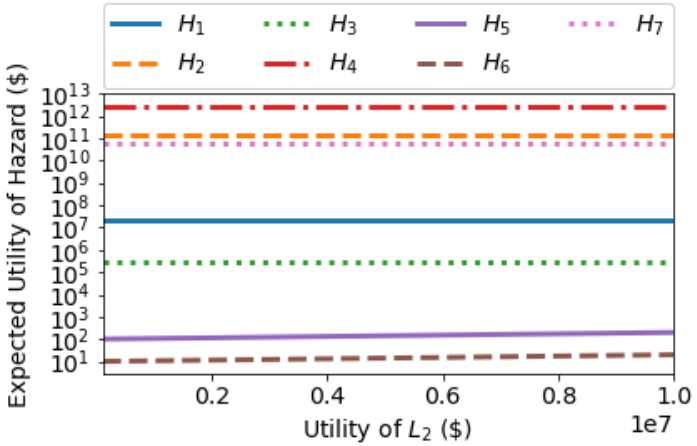

Disgruntled Employee

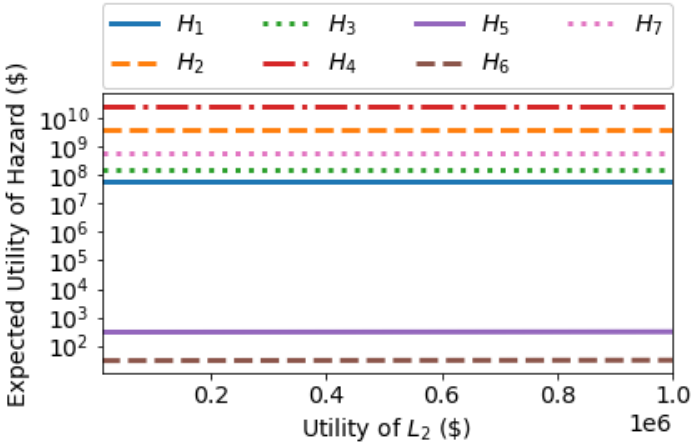

Government Cyberwarrior

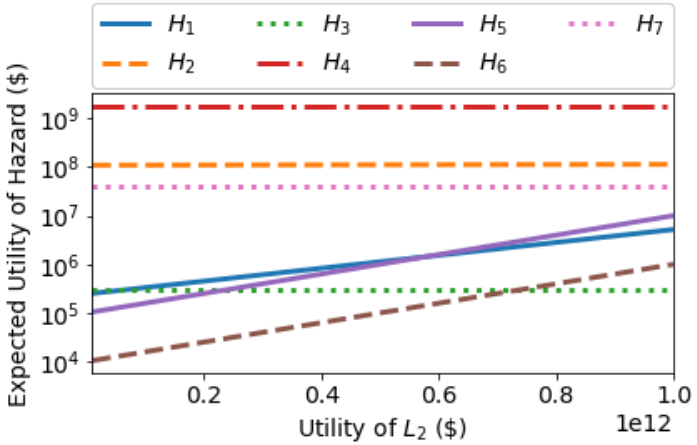

Terrorist

Detection of the attacker's true type was most effective against the radical activist and disgruntled employee

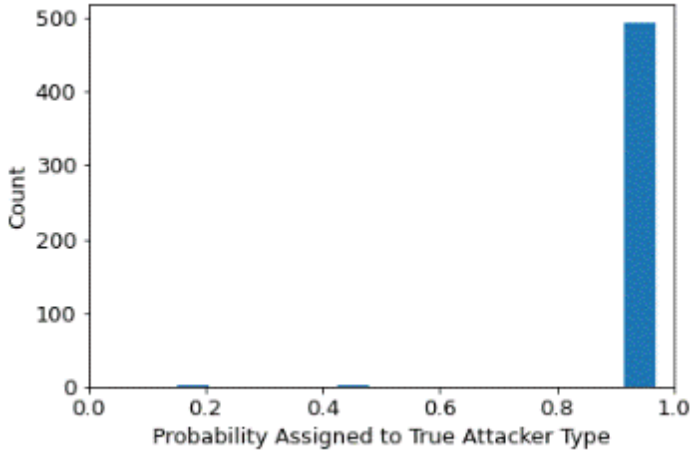

Radical Activist

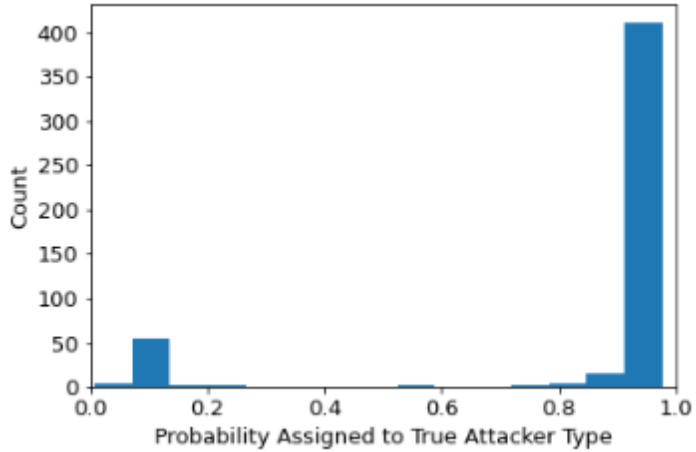

Disgruntled Employee

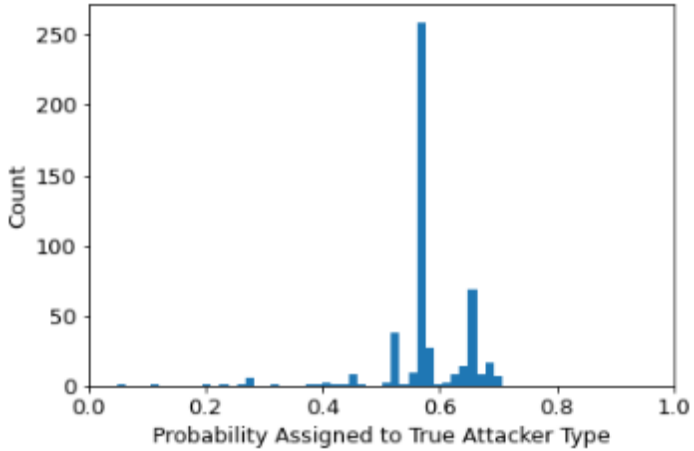

Government Cyberwarrior

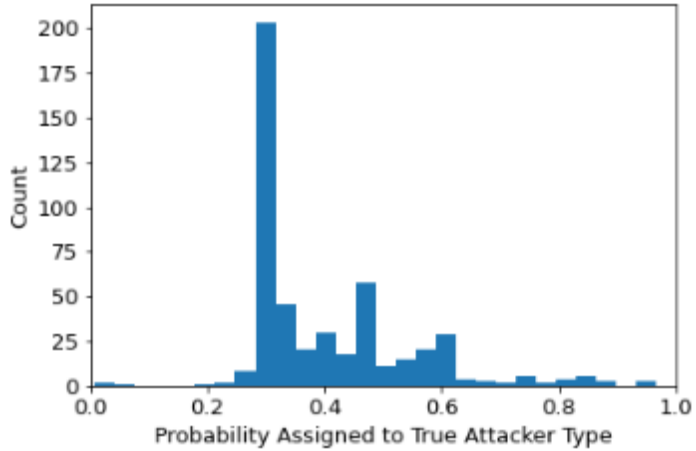

Terrorist

The goal of this research is to reduce the likelihood of successful cyber-attacks on nuclear power plants

1. Predict how an adversary might target a plant
   - SBG construction
   - SBG simulation

2. Quantify nuclear power plant cybersecurity
   - SBG simulation: time-to-loss, availability, utility

3. Allocate cybersecurity resources to defend a plant
   - HBA and Bayesian learning

This work can help infrastructure defenders to cost-effectively allocate security resources given uncertainty

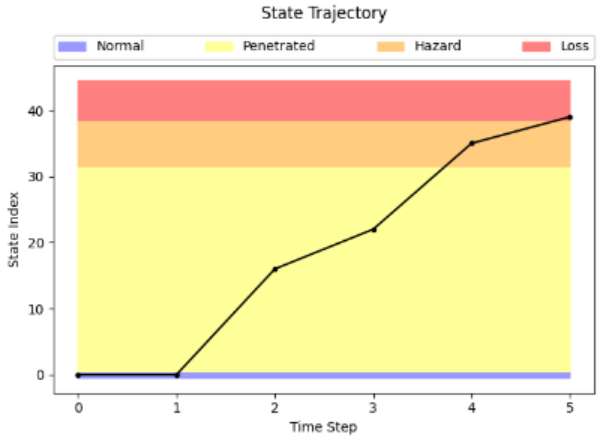

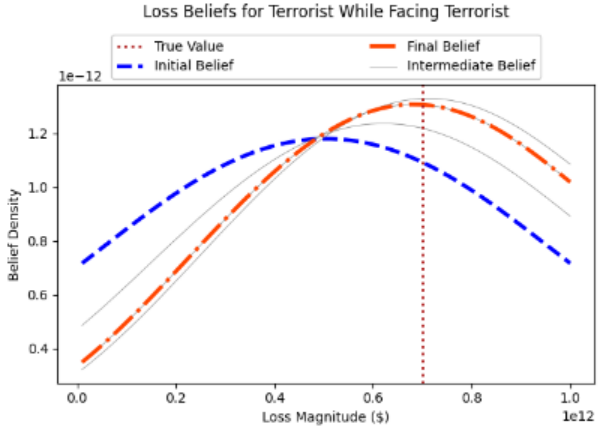

Questions?

Lee T. Maccarone
Postdoctoral Appointee
lmaccar@sandia.gov

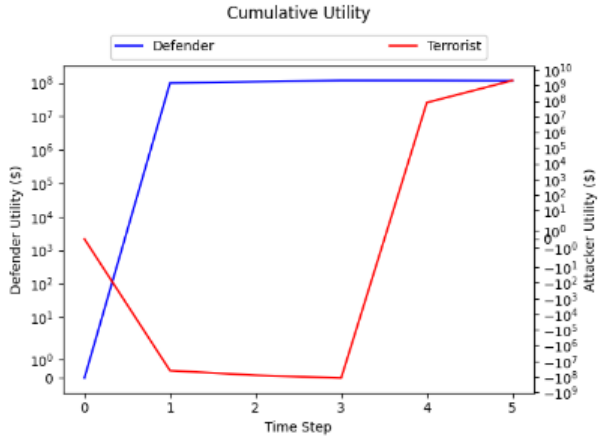

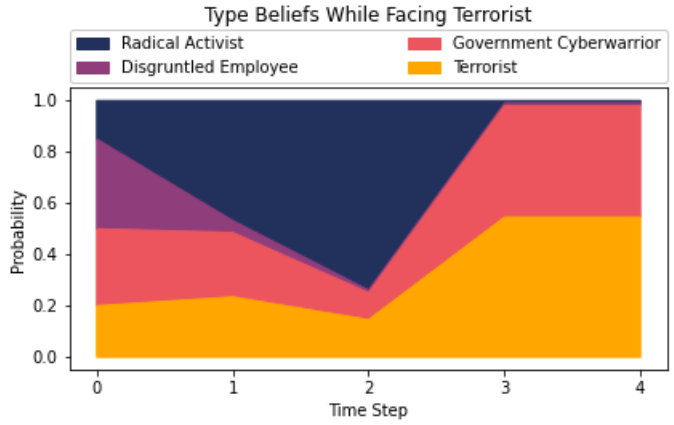

This material is based upon work supported under (i) a DOE IUP Graduate Fellowship, (ii) Sandia National Laboratories Contract 2084579: Civilian Nuclear Power Plant Systems Game and Risk Modeling, and (iii) the University of Pittsburgh Center for Research Computing