Cyber Deception: Games, Defense, and Learning

Quanyan Zhu

December 1, 2022

NYU TANDON SCHOOL OF ENGINEERING
C2SMART CONNECTED CITIES WITH SMART TRANSPORTATION
CUSP CENTER FOR URBAN SCIENCE+PROGRESS
CENTER FOR CYBER SECURITY

Hello Quanyan Inbox ×

David [redacted] <katherinebrian31@gmail.com> Tue, Jul 21, 12:18 PM
to Quanyan

I would take little of your time today are you free? Send me a number to reach you.

[redacted] Professor of Electrical and Computer Engineering | Director, [redacted]
2064793905

Quanyan Zhu <quanyan.zhu@gmail.com> Tue, Jul 21, 12:20 PM
to David

I am free any time today before 4PM at 646 997 3371.

Best regards,

•••


· 37m
Brains behind new **5G** data communications networks described below! New Bill Gates sponsored **corona virus** vaccine, w/nano tech, will run everything and control everyone who are still necessary, like bots to serve the elite? Get your vaccine now?
Get the facts about COVID-19
berg
The Rise of AI
There's an AI revolution sweeping across the world.
Yet few people know the rea
youtube.com


1:02 / 1:12
You Won't Believe What Obama Says In This Video! 😉

Mission: Impossible - Ghost Protocol (2011) - Hallway Projection Scene
https://www.youtube.com/watch?v=7DkV8WE7DFA

Deception

To deceive $\stackrel{\text{def}}{=}$ to **intentionally** cause another agent to acquire or continue to have a false **belief**, or to be prevented from acquiring or cease to have a true **belief**.

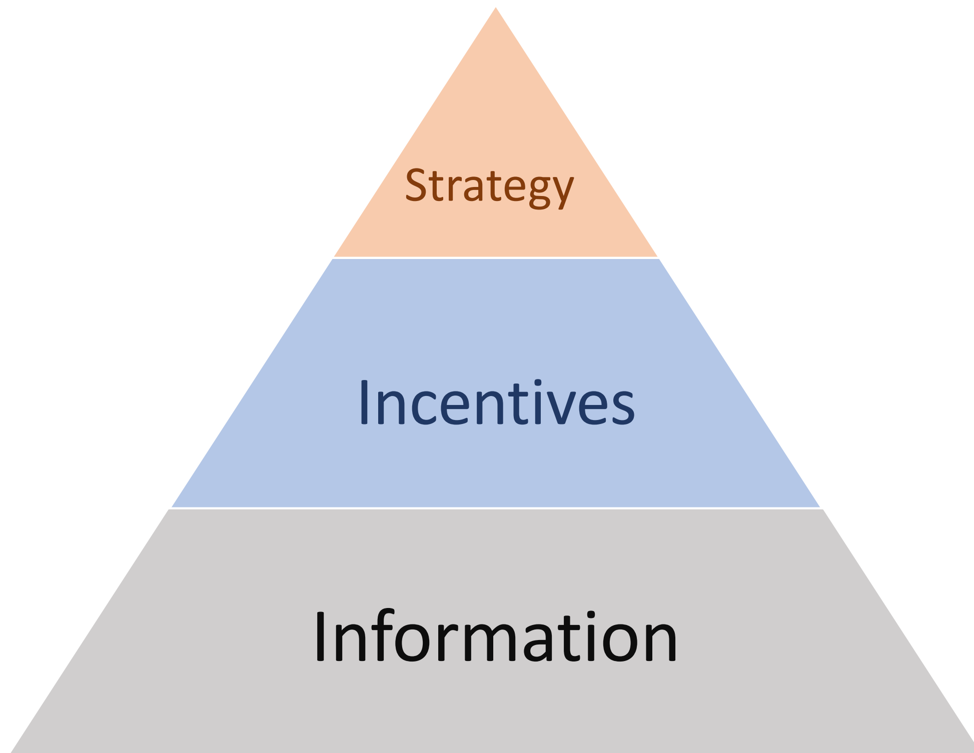
Strategy
Incentives
Information

Incentives: What is the Purpose of the Deception?

Mimetic Deception

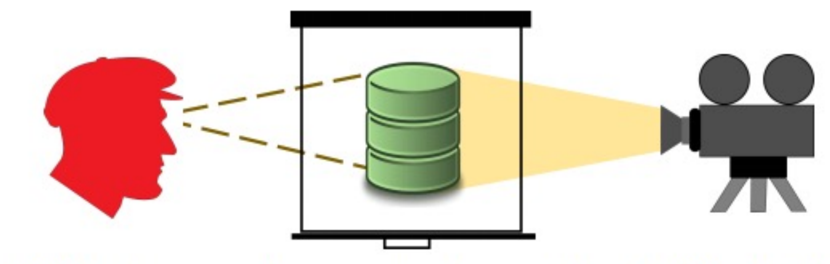

Honey-X

Cryptic Deception

Obfuscation

Strategies: Single Actor or Multiple Actors?

Intensive Deception

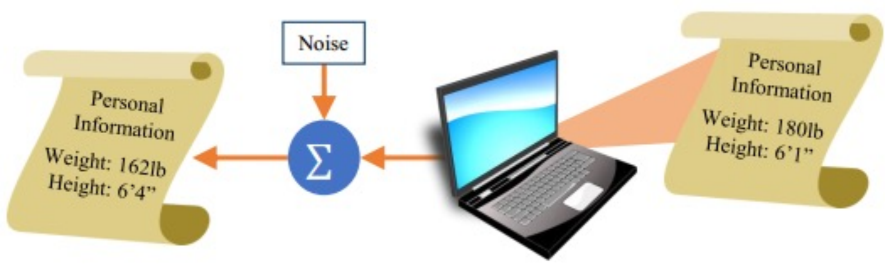

- Single target / actor

Perturbation

Extensive Deception

- Multiple targets / actors

Mixing

Defensive Deception: Taxonomy

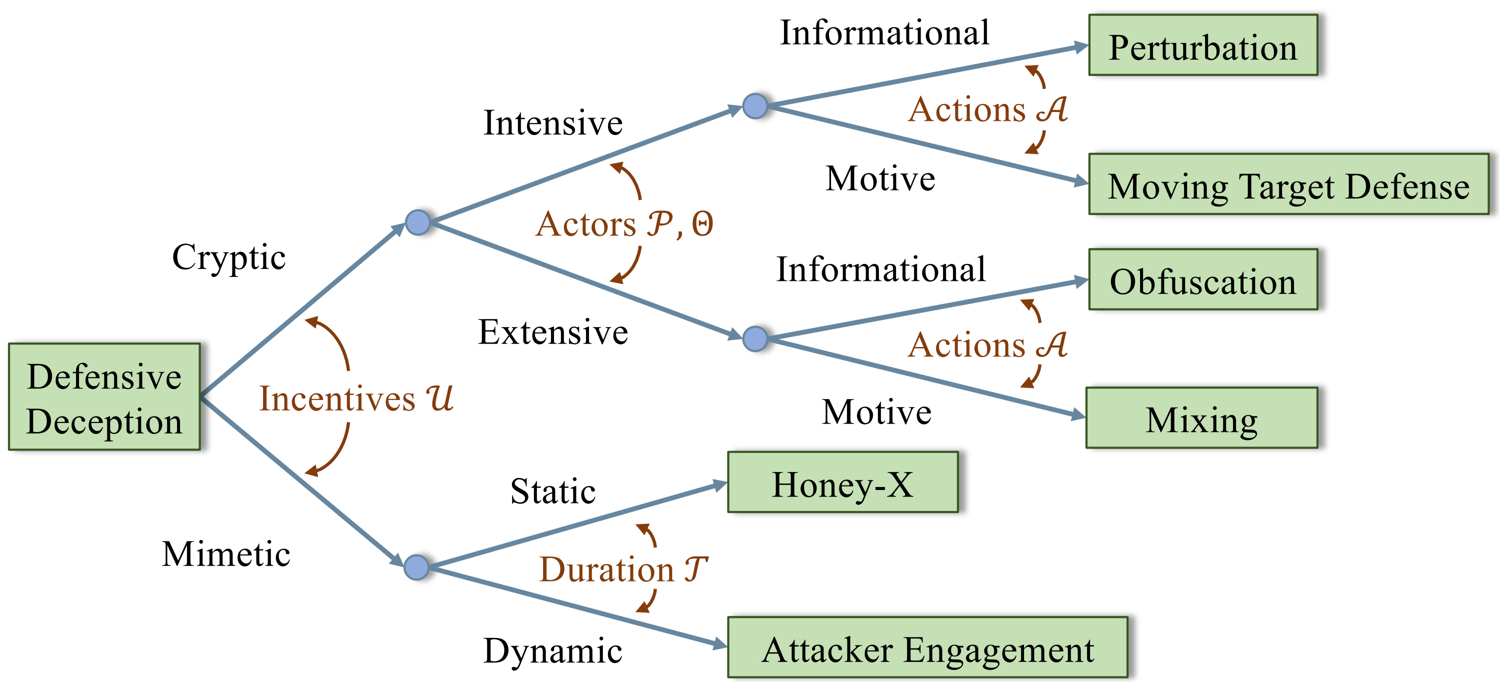

Pawlick J, Colbert E, Zhu Q. A game-theoretic taxonomy and survey of defensive deception for cybersecurity and privacy. ACM Computing Surveys (CSUR). 2019 Aug 30;52(4):1-28.

Talk Outline

1) Introduction
2) Taxonomy of defensive deception
3) **Signaling games for mimetic cyber deception**
   - Honey-X
   - Attack Engagement
4) Dynamic games for cyber-physical deception
   - Robotic Deception
   - Conjectural Meta-Learning
5) Future challenges

Taxonomy Based on Game Theoretic Principles

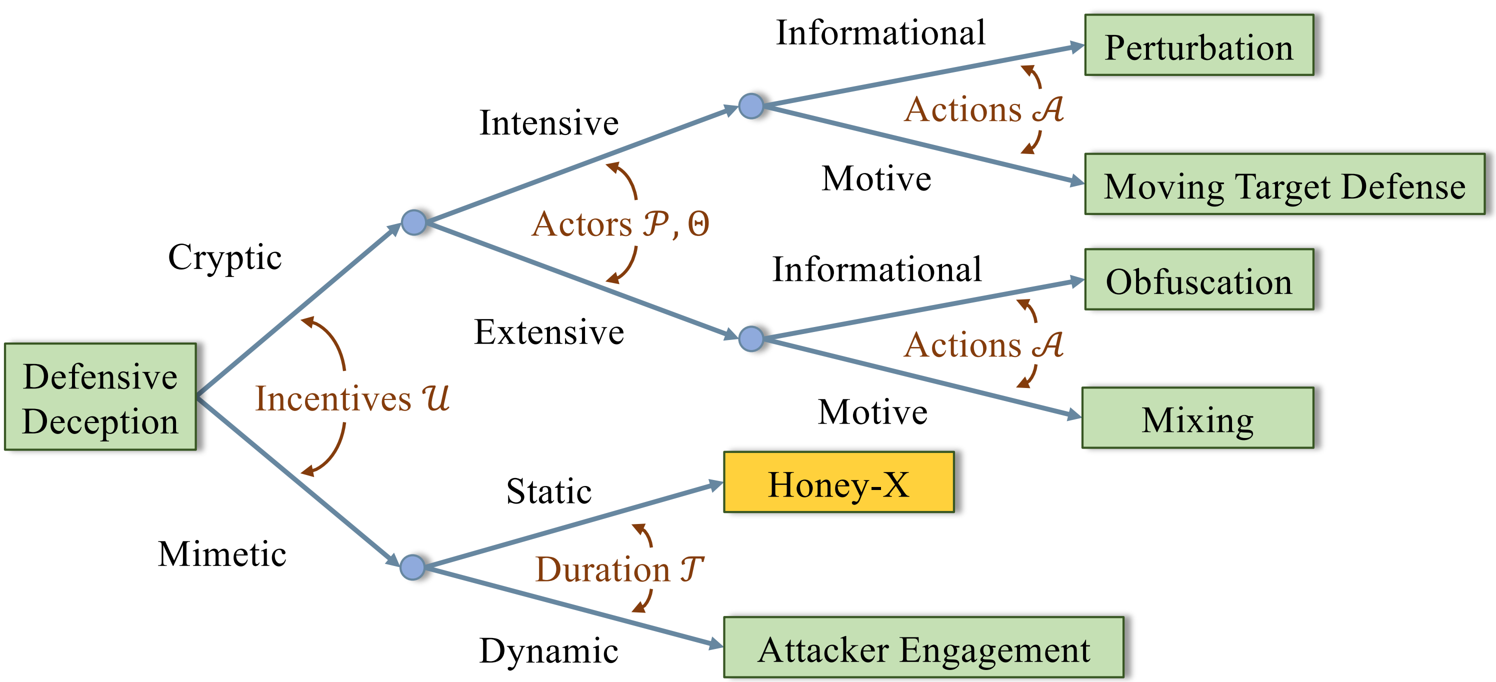

Pawlick J, Colbert E, Zhu Q. Modeling and analysis of leaky deception using signaling games with evidence. IEEE Transactions on Information Forensics and Security. 2018 Dec 12;14(7):1871-86.

Mimesis and Modeling Belief

- Signaling games model belief [Lewis 1969, Crawford & Sobel 1982].

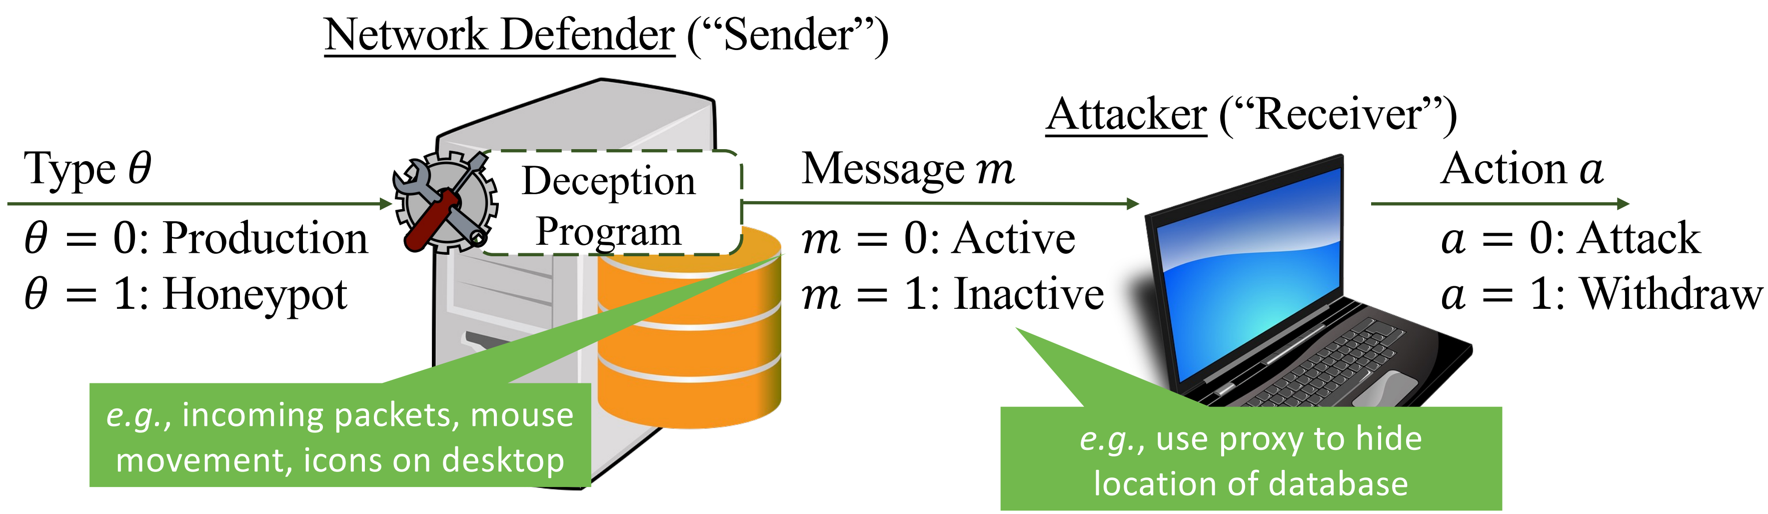

Mimesis and Modeling Belief

- But “deception program” may leak evidence.

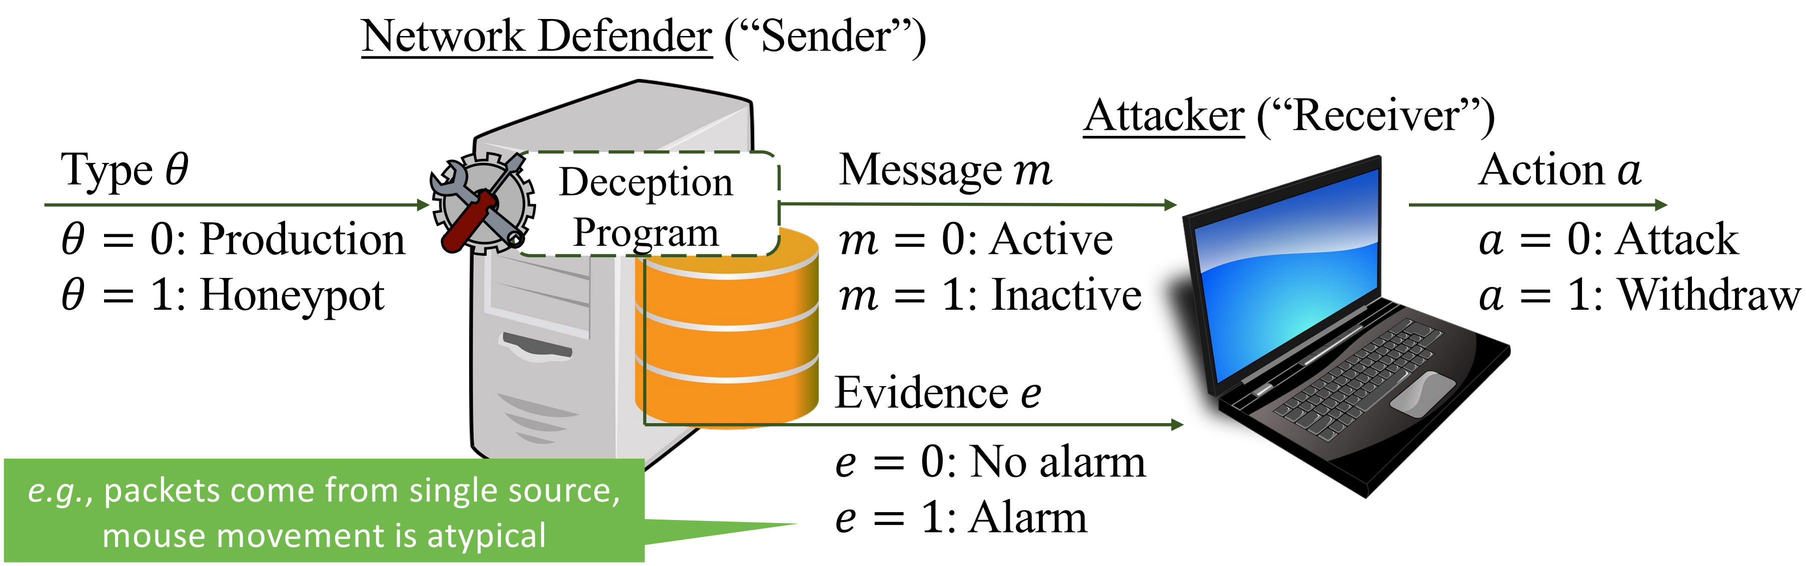

Mixed Strategies, Belief, and Expected Utility

- Attacker has (common) prior belief of system type θ with probability (wp) $p(\theta)$.
- Defender chooses message m wp $\sigma^S(m \mid \theta)$.
- Defender leaks evidence e wp $\lambda(e \mid \theta, m)$.
- Attacker forms belief $\mu^R(\theta \mid m, e)$ and chooses action a wp $\sigma^R(a \mid m, e)$.

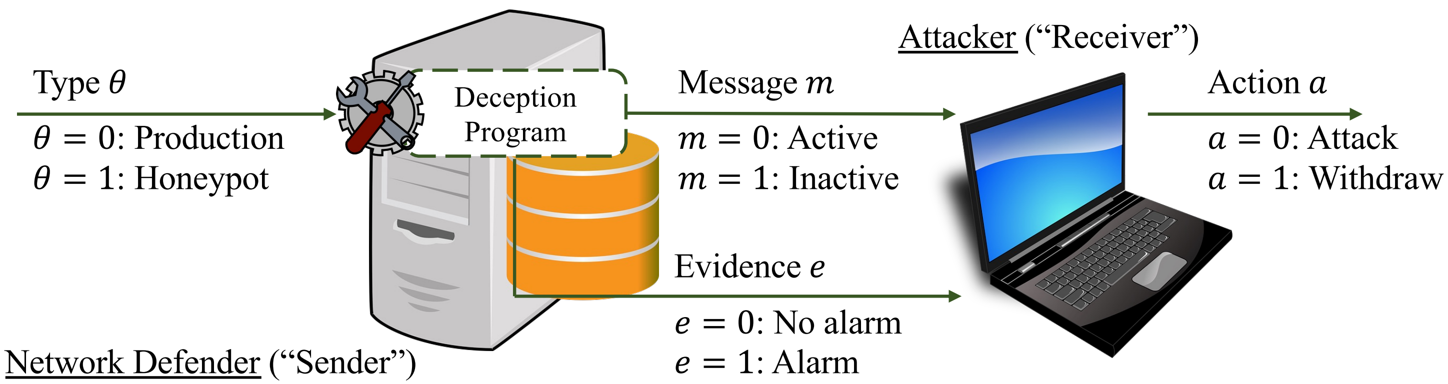

Mixed Strategies, Belief, and Expected Utility

- System of type θ has an expected utility of $U^S(\sigma^S, \sigma^R \mid \theta)$.
- Attacker that observes activity level m and evidence e has an expected utility of $\sum_{\theta\in\Theta} \mu^R(\theta \mid m, e) U^R(\sigma^R \mid \theta, m, e)$.

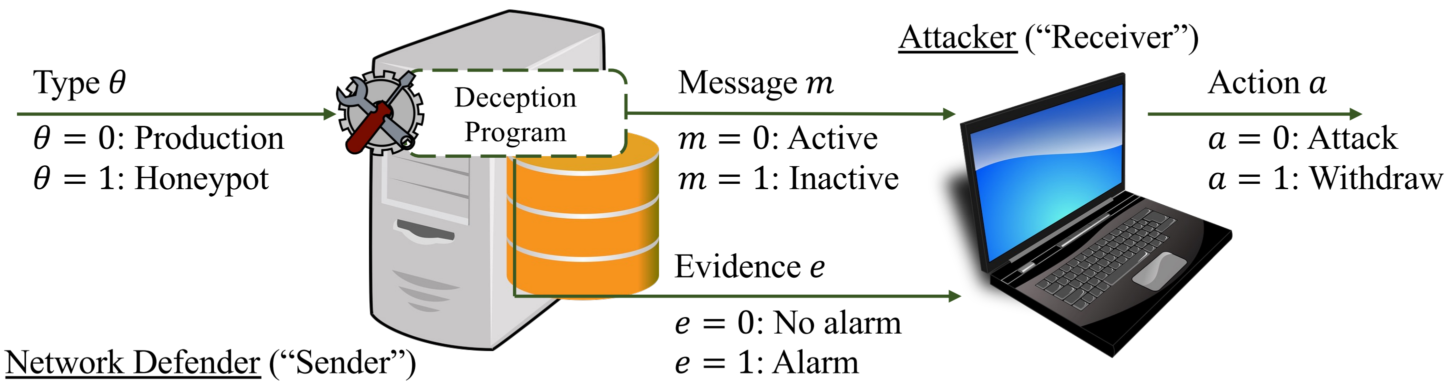

Perfect Bayesian Nash Equilibrium

A PBNE is a strategy profile $(\sigma^{S*}, \sigma^{R*})$ and posterior beliefs $\mu^R(\theta \mid m, e)$ such that:

$\forall \theta \in \Theta$,

$$\sigma^{S*} \in \text{argmax}_{\,\sigma^S \in \Gamma^S} U^S(\sigma^S, \sigma^{R*} \mid \theta),$$

$\forall m \in M, e \in \mathbb{EV}$,

$$\sigma^{R*} \in \text{argmax}_{\,\sigma^R \in \Gamma^R} \sum_{\theta \in \Theta} \mu^R(\theta \mid m, e) U^R(\sigma^R \mid \theta, m, e),$$

and

$$\mu^R(\theta \mid m, e) = \frac{\lambda(e \mid \theta, m)\sigma^S(m \mid \theta)p(\theta)}{\sum_{\tilde{\theta} \in \Theta} \lambda(e \mid \tilde{\theta}, m)\sigma^S(m \mid \tilde{\theta})p(\tilde{\theta})},$$

when that fraction is defined.

Detector and Utility Meta-Parameters

"Aggressive Detector"
A lot of false alarms (detect too quickly/aggressively)

"Conservative Detector"
A lot of miss detection (detect too slowly/conservatively)

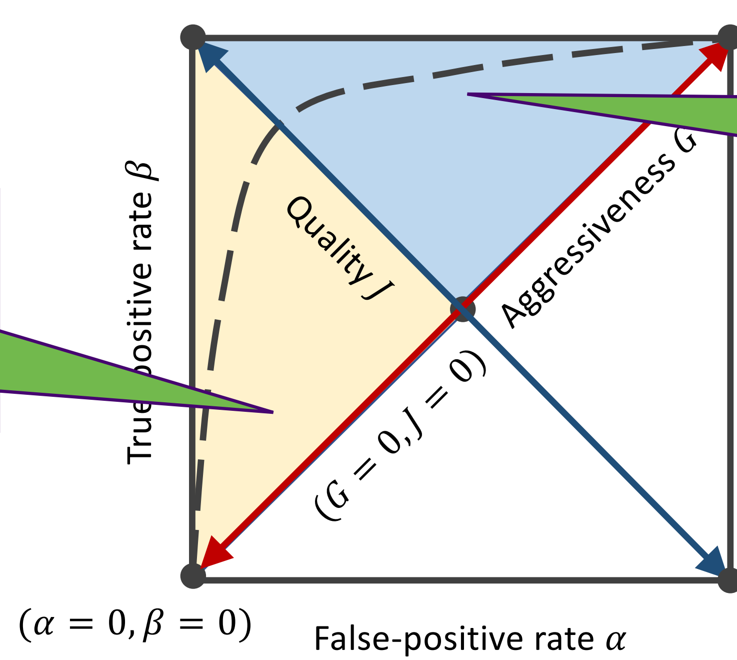

Utility Parameters:

$$\Delta_0^R = u^R(0, m, 0) - u^R(0, m, 1)$$

$$\Delta_1^R = u^R(1, m, 1) - u^R(1, m, 0)$$

Equilibrium Regions

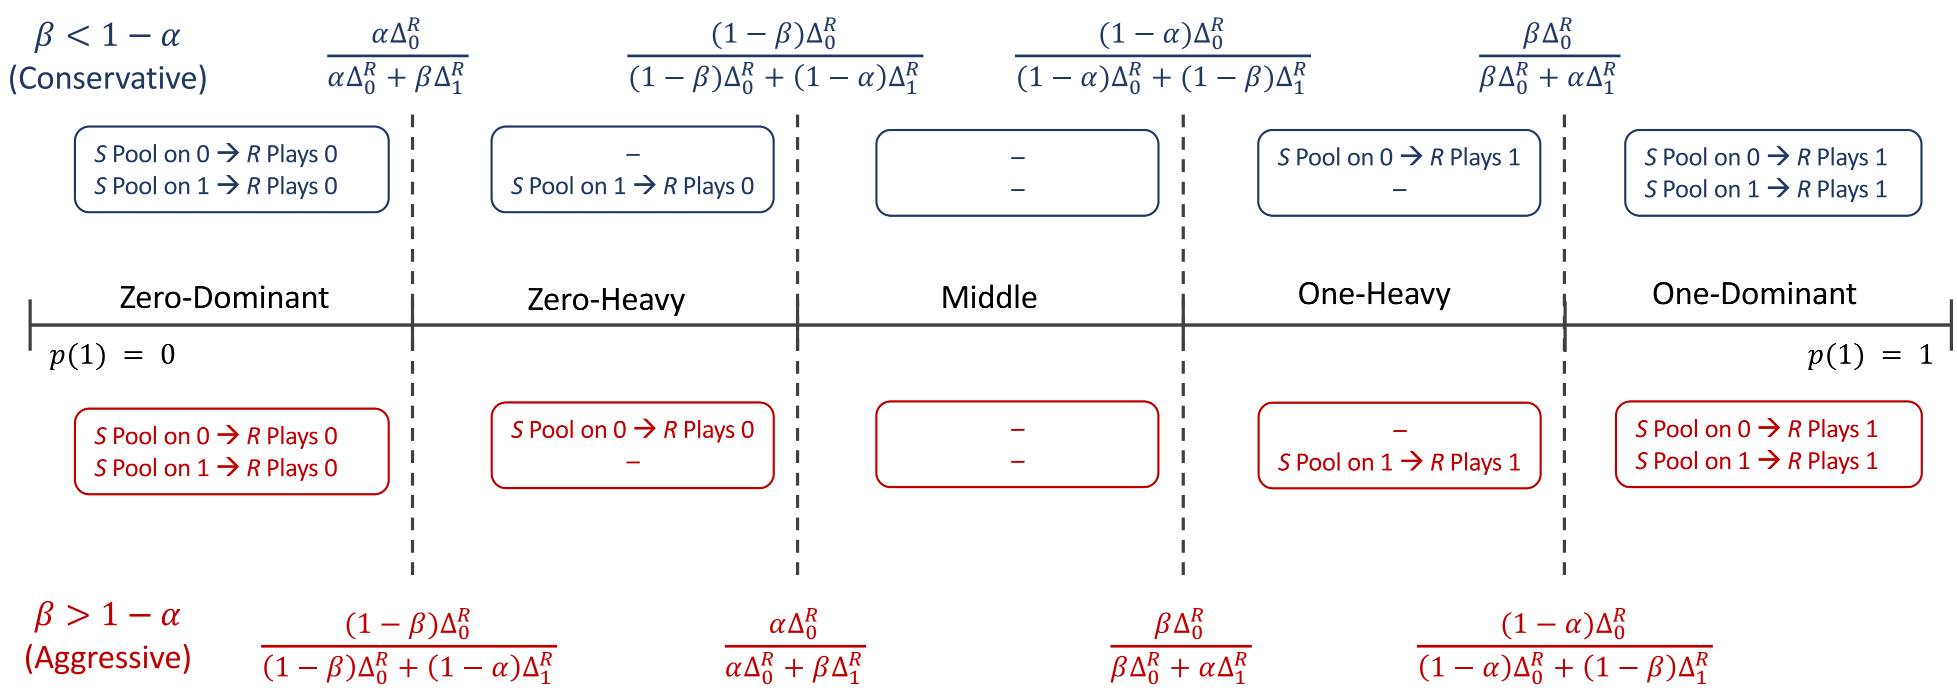

Partially-Separating Equilibria in the Middle Regime

Theorem (Aggressive Detectors). For $\beta > 1 - \alpha$, within the Middle regime, there exists a PBNE in which

$$\sigma^{S*}(m = 1|\theta = 0) = \frac{\overline{\alpha}\overline{\beta}\Delta_1^R}{(\overline{\alpha}^2 - \overline{\beta}^2)\Delta_0^{\mathrm{R}}}\left(\frac{p(1)}{1 - p(1)}\right) - \frac{\overline{\beta}^2}{\overline{\alpha}^2 - \overline{\beta}^2},$$

$$\sigma^{S*}(m = 1|\theta = 1) = \frac{\overline{\alpha}^2}{\overline{\alpha}^2 - \overline{\beta}^2} - \frac{\overline{\alpha}\overline{\beta}\Delta_0^R}{(\overline{\alpha}^2 - \overline{\beta}^2)\Delta_1^{\mathrm{R}}}\left(\frac{1 - p(1)}{p(1)}\right),$$

and

$$\sigma^{R*}(a = 1|m = 0, e = 0) = 0, \quad \sigma^{R*}(a = 1|m = 0, e = 1) = \frac{1}{\alpha + \beta},$$

$$\sigma^{R*}(a = 1|m = 1, e = 0) = 1, \quad \sigma^{R*}(a = 1|m = 1, e = 1) = \frac{\alpha + \beta - 1}{\alpha + \beta},$$

and the beliefs are computed by Bayes' Law in all cases. Here $\overline{\boldsymbol{x}} = \mathbf{1} - \boldsymbol{x}$.

Partially-Separating Equilibria in the Middle Regime

Theorem (Conservative Detectors). For $\beta < 1 - \alpha$, within the Middle regime, there exists a PBNE in which

$$\sigma^{S*}(m = 1|\theta = 0) = \frac{\beta^2}{\beta^2-\alpha^2} - \frac{\alpha\beta\Delta_1^R}{(\beta^2-\alpha^2)\Delta_0^{\mathrm{R}}}\left(\frac{p(1)}{1-p(1)}\right),$$

$$\sigma^{S*}(m = 1|\theta = 1) = \frac{\alpha\beta\Delta_0^R}{(\beta^2-\alpha^2)\Delta_1^{\mathrm{R}}}\left(\frac{1-p(1)}{p(1)}\right) - \frac{\alpha^2}{\beta^2-\alpha^2},$$

and

$$\sigma^{R*}(a = 1|m = 0, e = 0) = \frac{1-\alpha-\beta}{2-\alpha-\beta}, \quad \sigma^{R*}(a = 1|m = 0, e = 1) = 1,$$

$$\sigma^{R*}(a = 1|m = 1, e = 0) = \frac{1}{2-\alpha-\beta}, \quad \sigma^{R*}(a = 1|m = 1, e = 1) = 0,$$

and the beliefs are computed by Bayes' Law in all cases.

Partially-Separating Strategies for S

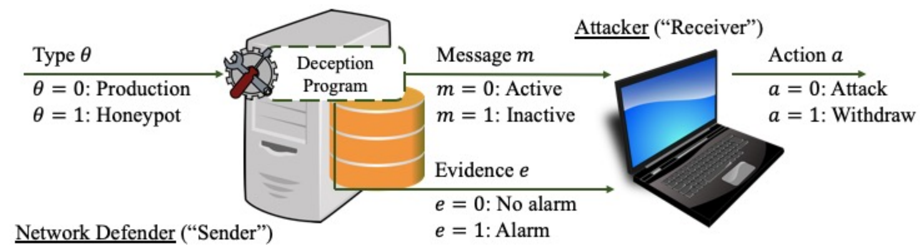

$\theta = 1$: "Reveal" honeypot $(m = 1)$ as inactive wp 0.47.

$\theta = 0$: "Reveal" production $(m = 0)$ as active wp 0.89.

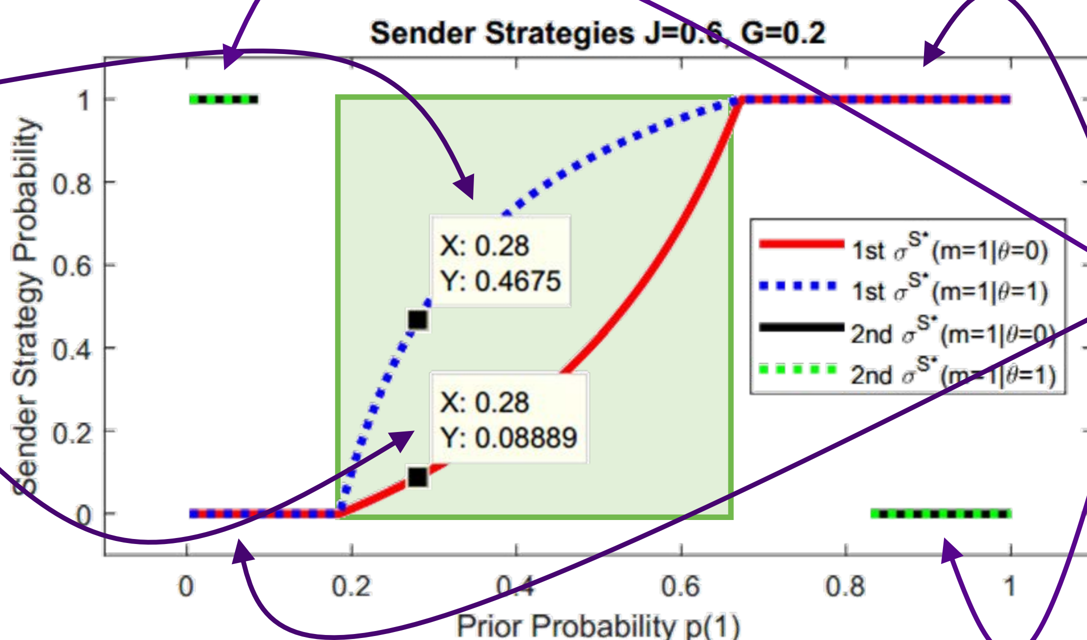

Coincident lines: Pooling Equilibrium #2

Middle Regime: Partially Separating

Comparative Statics: Detector Quality $J = \beta - \alpha$

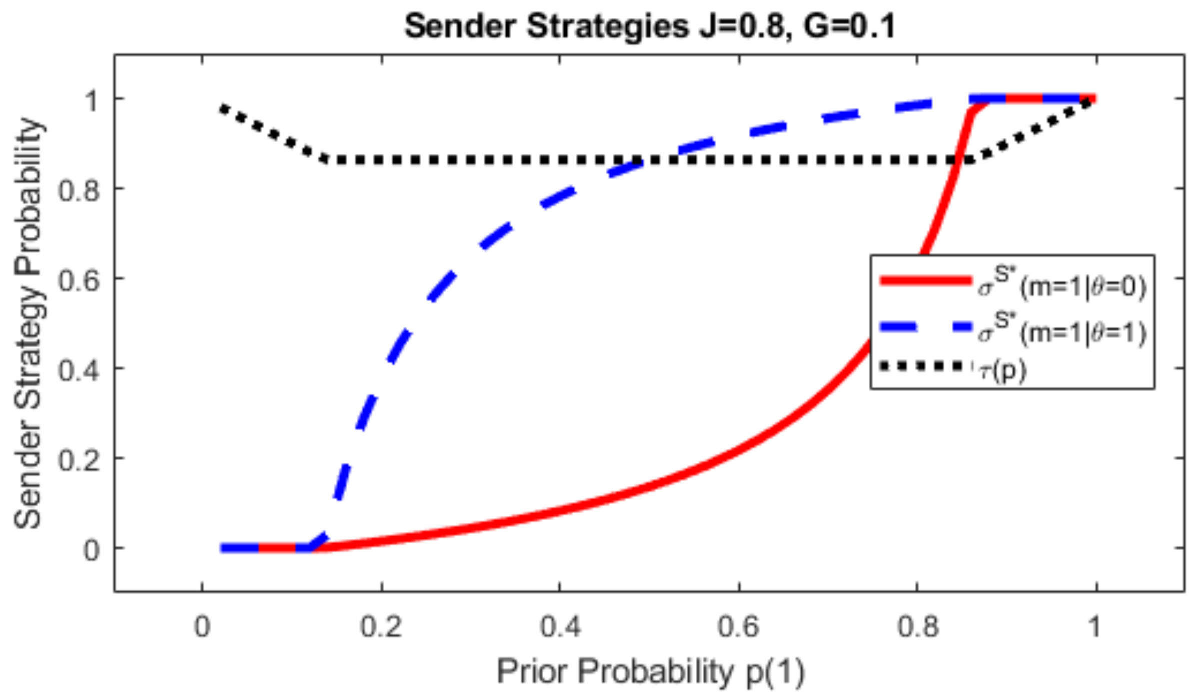

Comparative Statics: Aggressiveness $\mathrm{G} = \beta - (1-\alpha)$

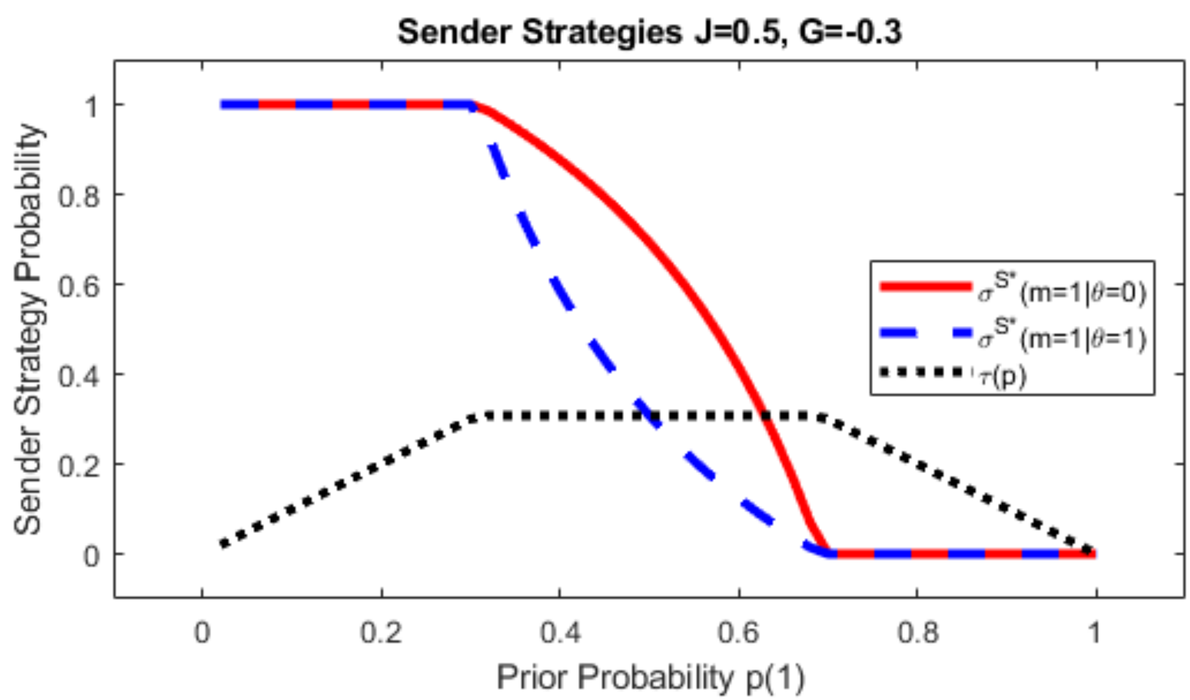

Truth Induction

Theorem (Truth Induction). Set $\Delta_0^R = \Delta_1^R$. Within regimes that feature unique PBNE, for all $J \in [0,1]$ and for any prior probability $p(\theta)$:

$$\tau(J, G, p) \geq \frac{1}{2} \text{ for } G \in [0,1),$$

$$\tau(J, G, p) \leq \frac{1}{2} \text{ for } G \in (-1,0],$$

Fraction of messages $m = \theta$

where

$$\tau(J, G, p) \triangleq \sum_{\theta \in \{0,1\}} p(\theta) \sigma^{S*}(m = \theta \mid \theta; p).$$

Aggressive detectors induce a *truth-telling convention*, while conservative detectors induce a *falsification convention*.

The Eye of Providence

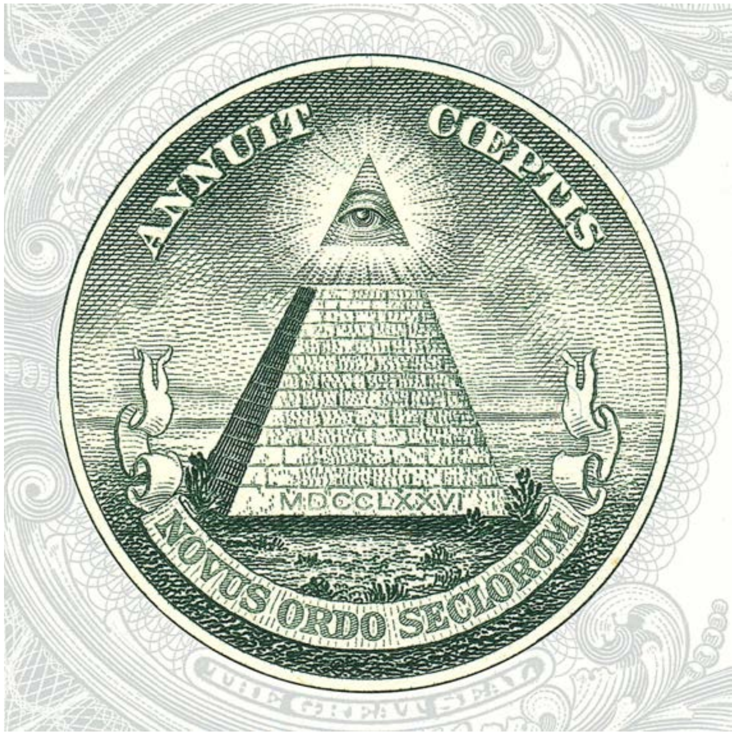

Offensive Deception

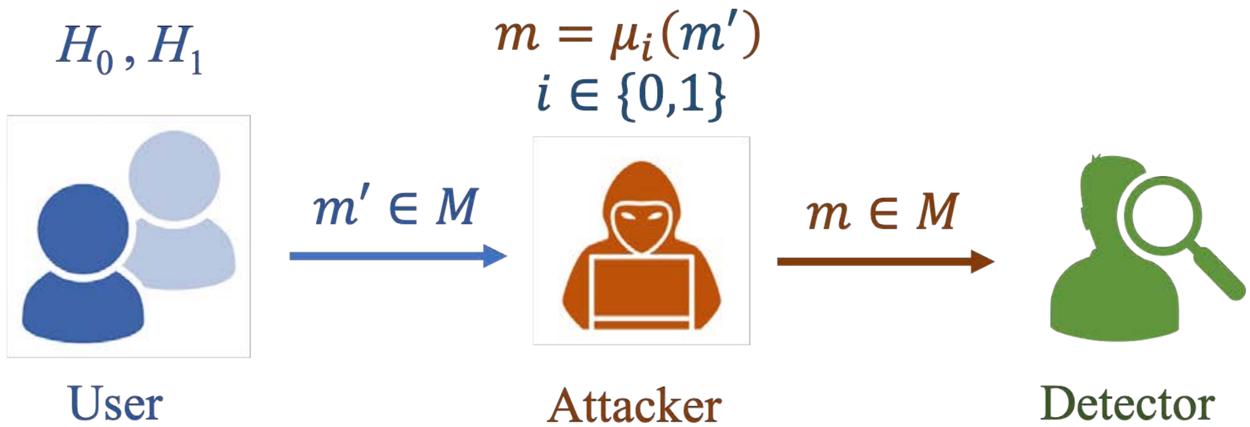

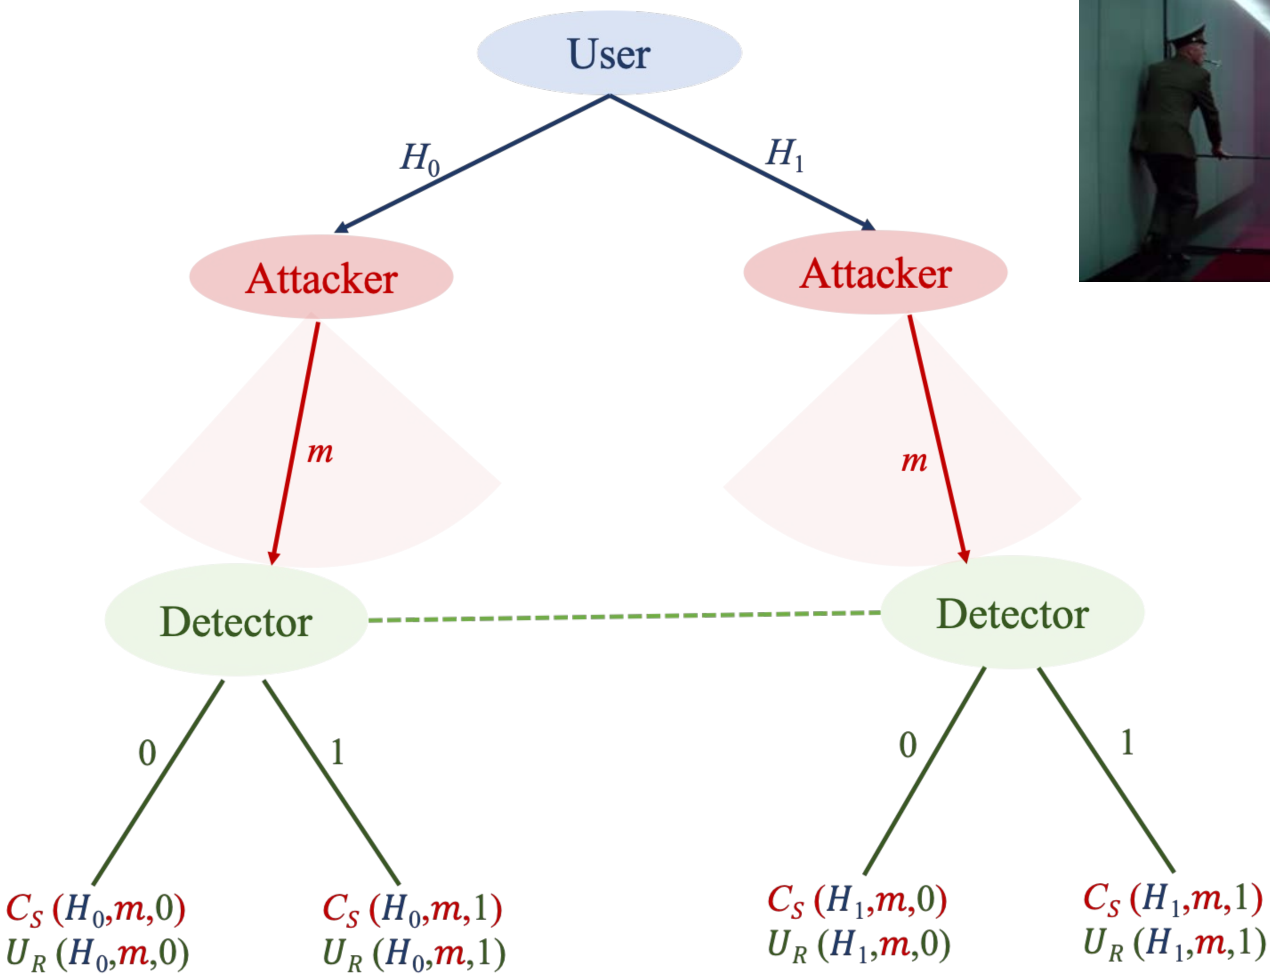

Hu Y, Zhu Q. Game-theoretic Neyman-Pearson detection to combat strategic evasion, in Proceedings of CDC 2022, arXiv preprint arXiv:2206.05276.

Defensive Deception:
Taxonomy Based on Game Theoretic Principles

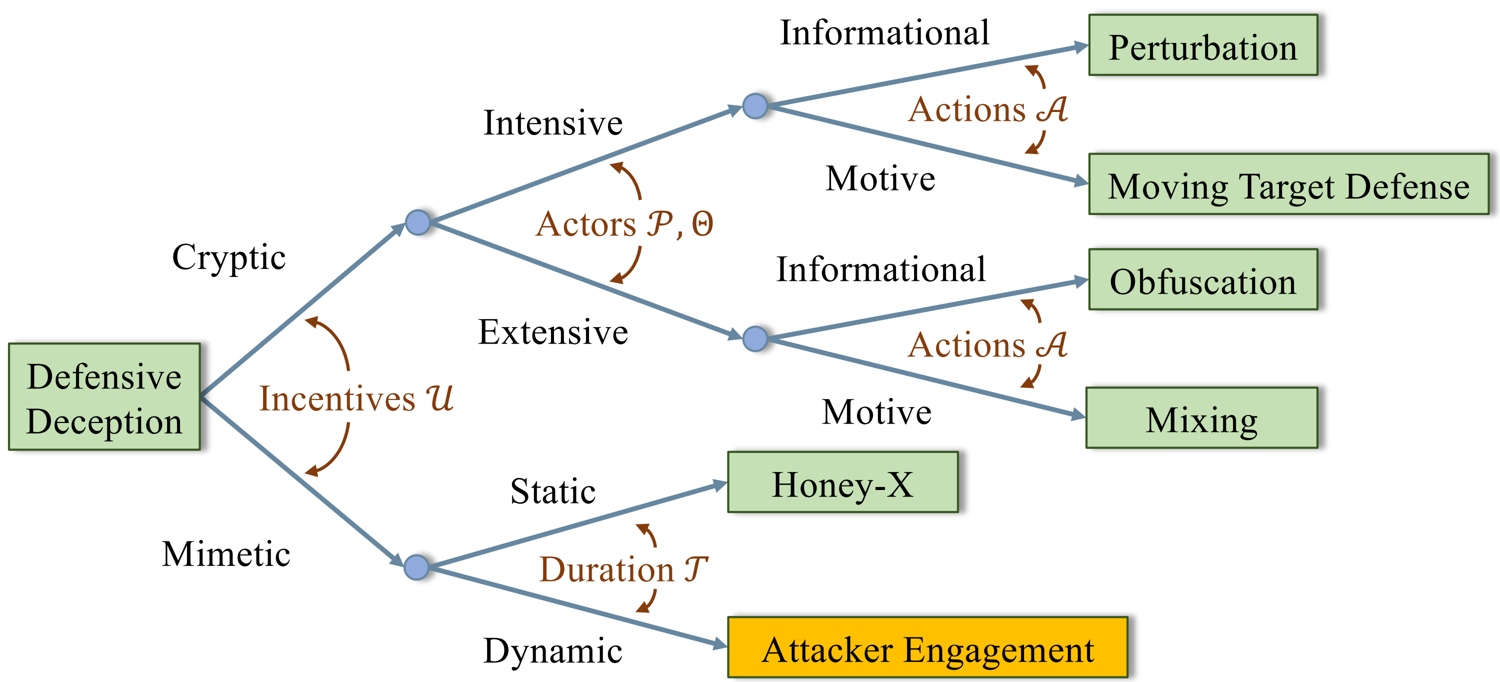

Dynamic Deception Model: One-Sided Partially Observable Markov Stochastic Games

- Two-player zero-sum
- Discounted infinite horizon

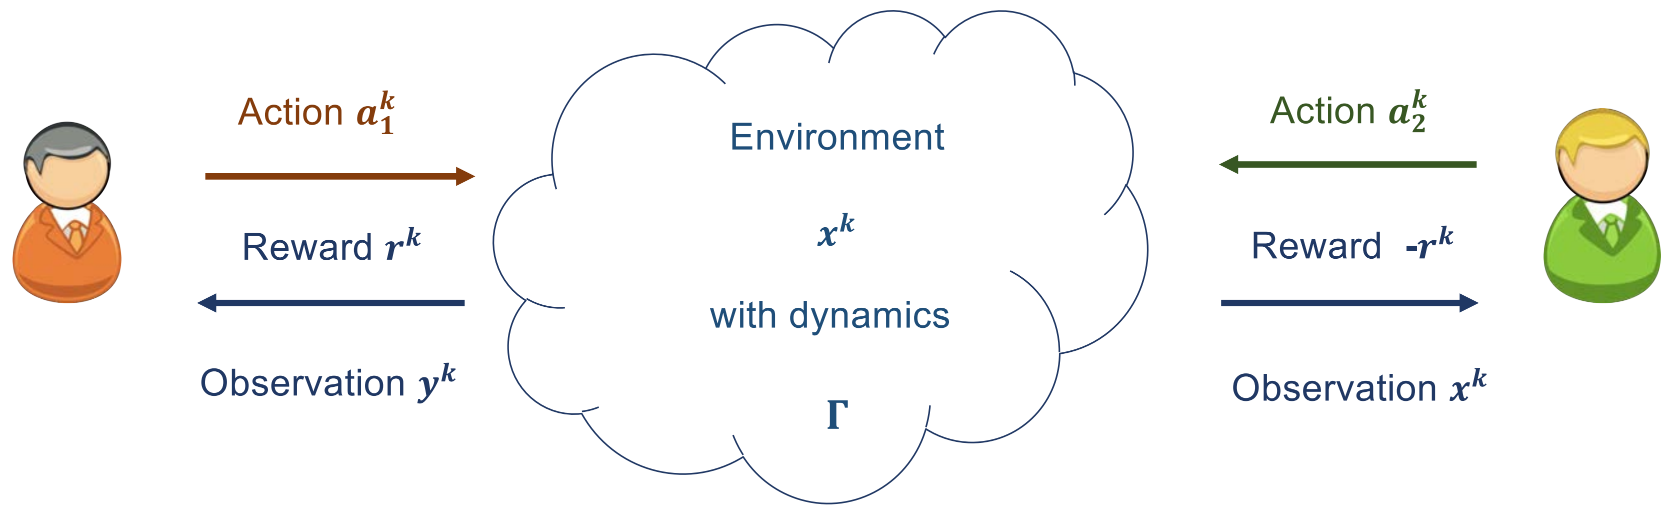

Horák K, Zhu Q, Bošanský B. Manipulating adversary's belief: A dynamic game approach to deception by design for proactive network security. International Conference on Decision and Game Theory for Security 2017 Oct 23 (pp. 273-294).

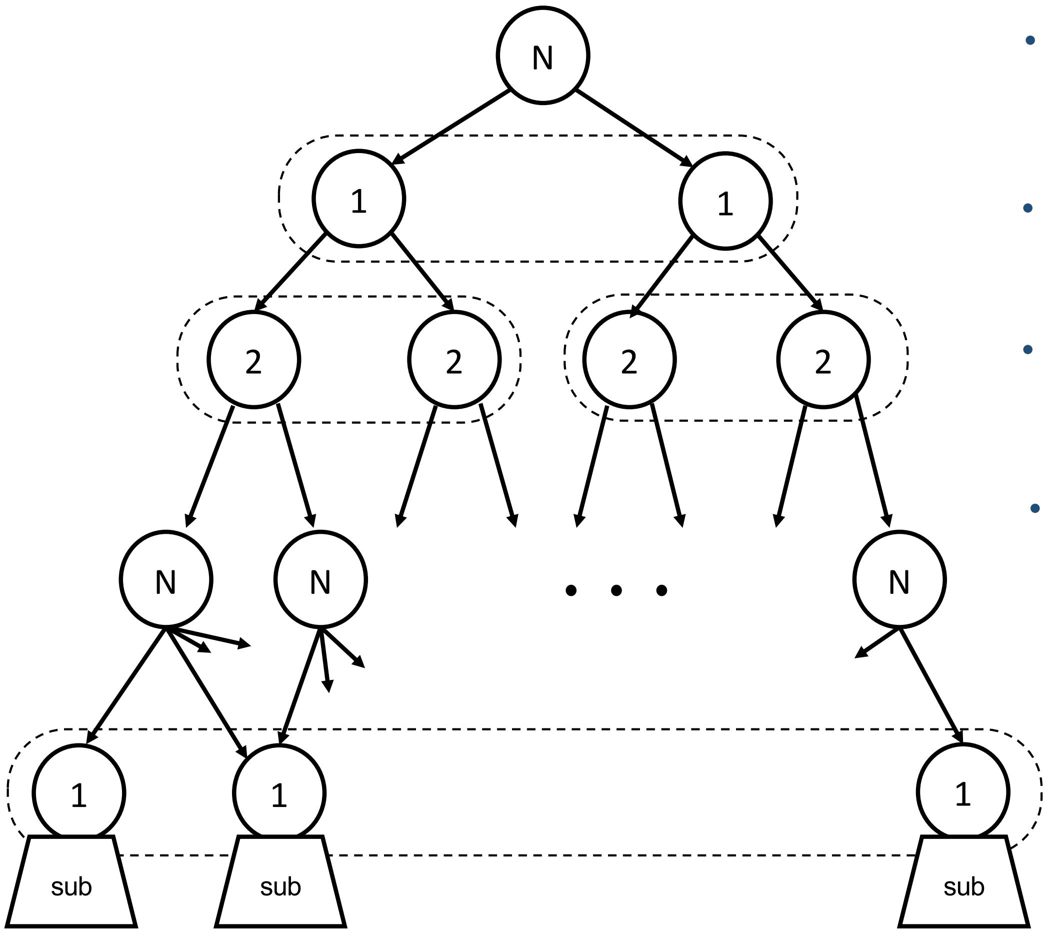

- An initial state is drawn from the initial belief $b^k \in \Delta(X)$.
- P2 observes x^k, P1 observes y^k.
- Players take simultaneous actions (a_1^k, a_2^k).
- Nature decides the state x^{k+1} and observation y^{k+1} at $k+1$ according to transition kernel $\Gamma_{x^k, a_1^k, a_2^k}(x^{k+1}, y^{k+1})$.

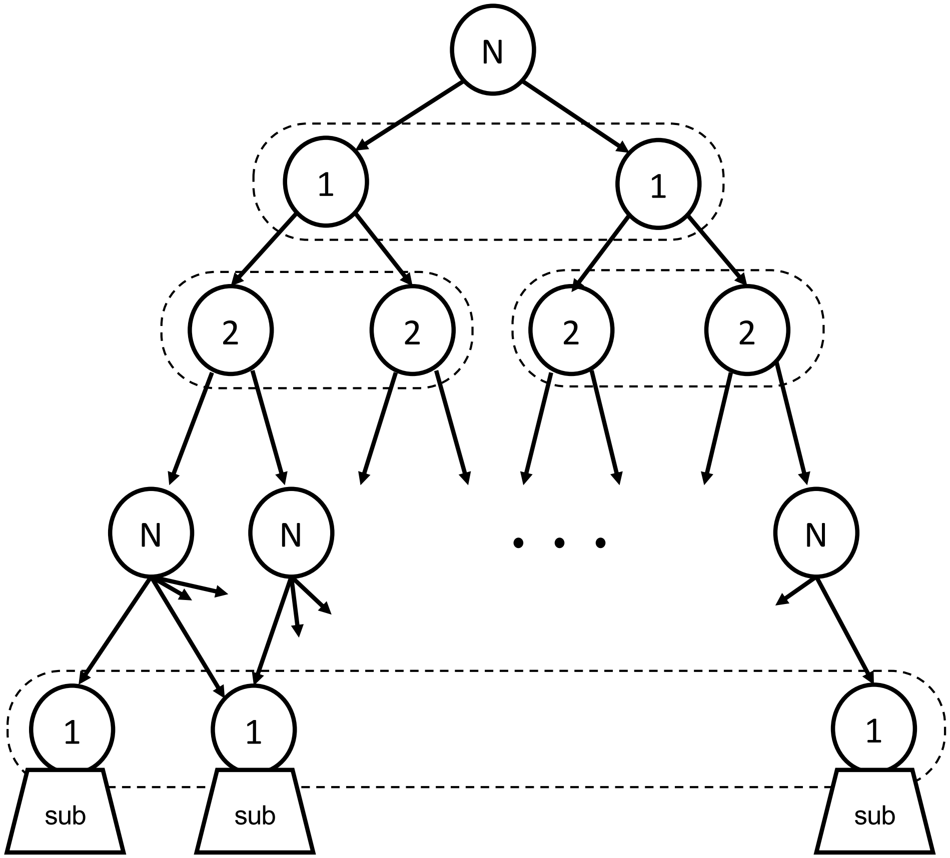

- P1's history $H_1^k := (A_1 \times Y)^k$
- P2's history $H_2^k := X \times (A_1 \times A_2 \times Y \times X)^k$
- Policy $\phi_i^k : H_i^k \mapsto \Delta(A_i)$
- Only need to keep track of belief for stationary policies
  - $\phi_1^{(b)} \in \Delta(A_1)$,
  - $\phi_2^{(b)} : X \mapsto \Delta(A_2)$
- P1's belief update under P2's policy $\phi_2^{(b)}$:

$$b_{\phi_2}^{a_1^k, y^k}(x^{k+1}) = \frac{\sum_{x^k \in X} \sum_{a_2^k \in A_2} \Gamma_{x^k, a_1^k, a_2^k}(x^{k+1}, y^k) b(x^k) \phi_2(x^k, a_2^k)}{\Pr(y^k | a_1^k, \phi_2)}$$

- Discounted-sum objective: $L = \sum_k \beta^k r^k$

- For zero-sum game:

$$\inf_{\phi_2} \sup_{\phi_1} L(\phi_1, \phi_2) = \sup_{\phi_1} \inf_{\phi_2} L(\phi_1, \phi_2)$$

- Convex value function v^* maps beliefs over the system state to the expected value.

$$v^*(b^k) = \min_{\phi_2} \max_{\phi_1} \left[\sum_{x^k, a_1^k, a_2^k} b^k(x^k) \phi_1(a_1^k) \phi_2(x^k, a_2^k) r^k(x^k, a_1^k, a_2^k) + \beta \sum_{a_1^k, y^k} \Pr(a_1^k, y^k | b^k, \phi_1, \phi_2) v^*(b^{k+1}) \right]$$

- Algorithms: LP and HSVI.

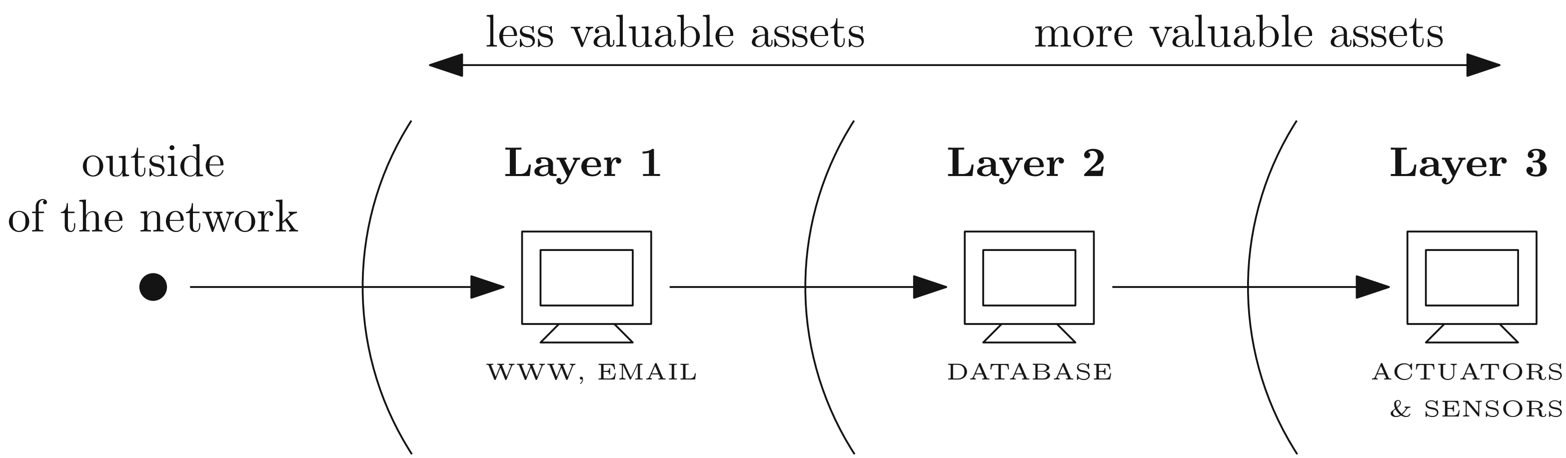

[Horak, Zhu, Bosansky, GameSec 2017]

Application in Network Security

- Defender has perfect information.
- Attacker has partial observation.
- Defender manipulates the attacker's belief to prevent him from succeeding.

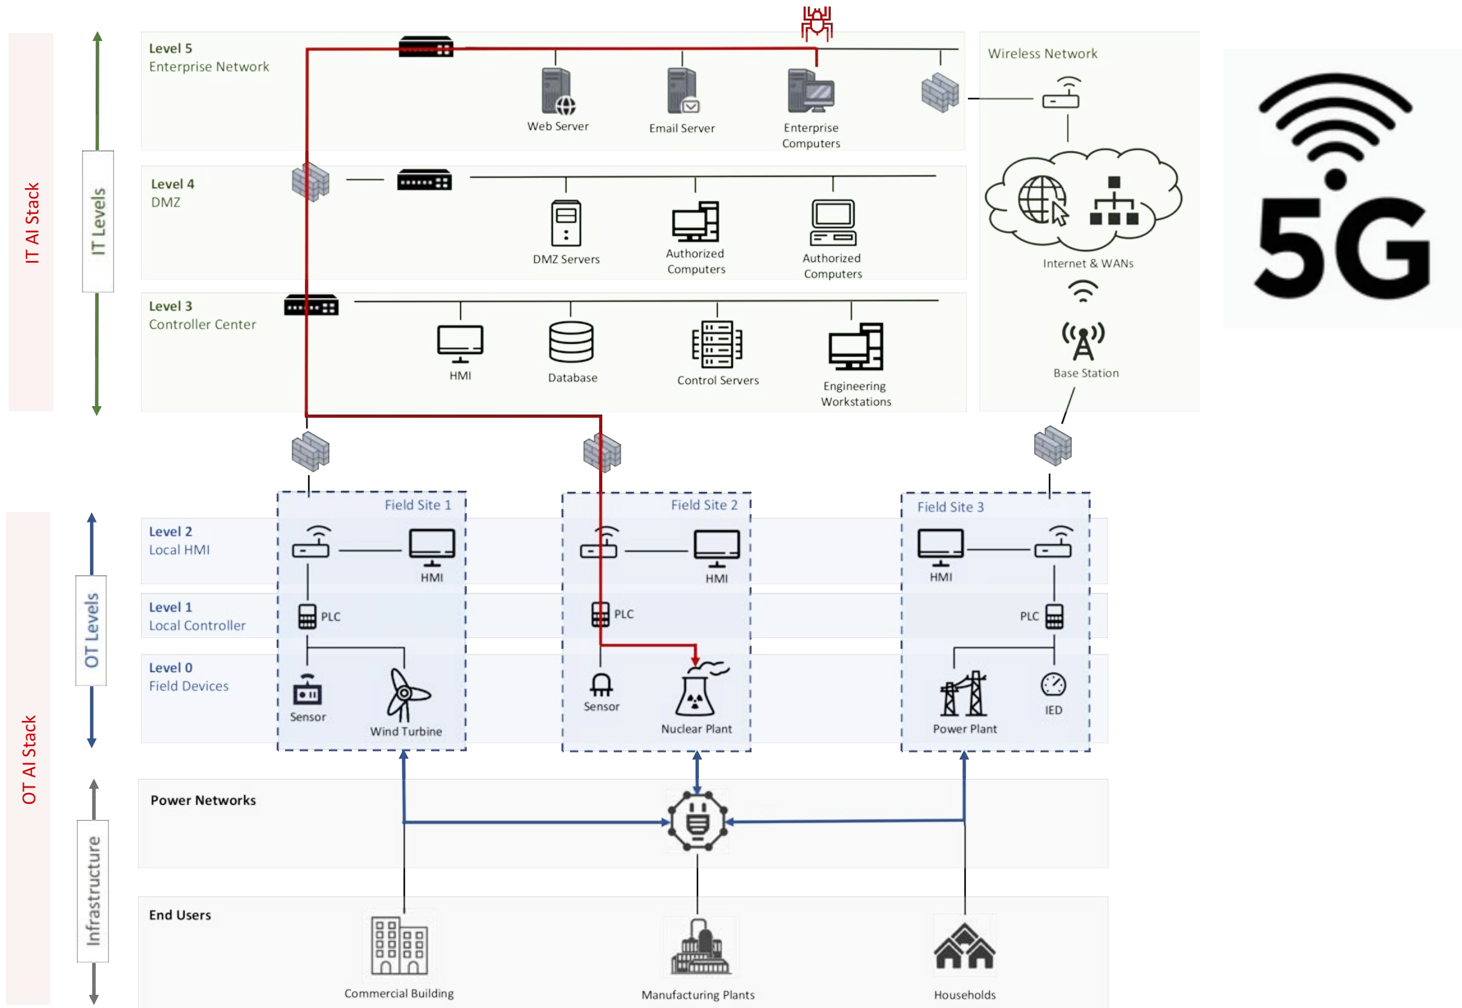

IT AI Stack
IT Levels
Level 5
Enterprise Network
Web Server
Email Server
Enterprise Computers
Level 4
DMZ
DMZ Servers
Authorized Computers
Authorized Computers
Level 3
Controller Center
HMI
Database
Control Servers
Engineering Workstations
Wireless Network
Internet & WANs
Base Station
5G
OT AI Stack
OT Levels
Field Site 1
Field Site 2
Field Site 3
Level 2
Local HMI
Level 1
Local Controller
Level 0
Field Devices
HMI
PLC
Sensor
Wind Turbine
Nuclear Plant
Power Plant
IED
Infrastructure
Power Networks
End Users
Commercial Building
Manufacturing Plants
Households

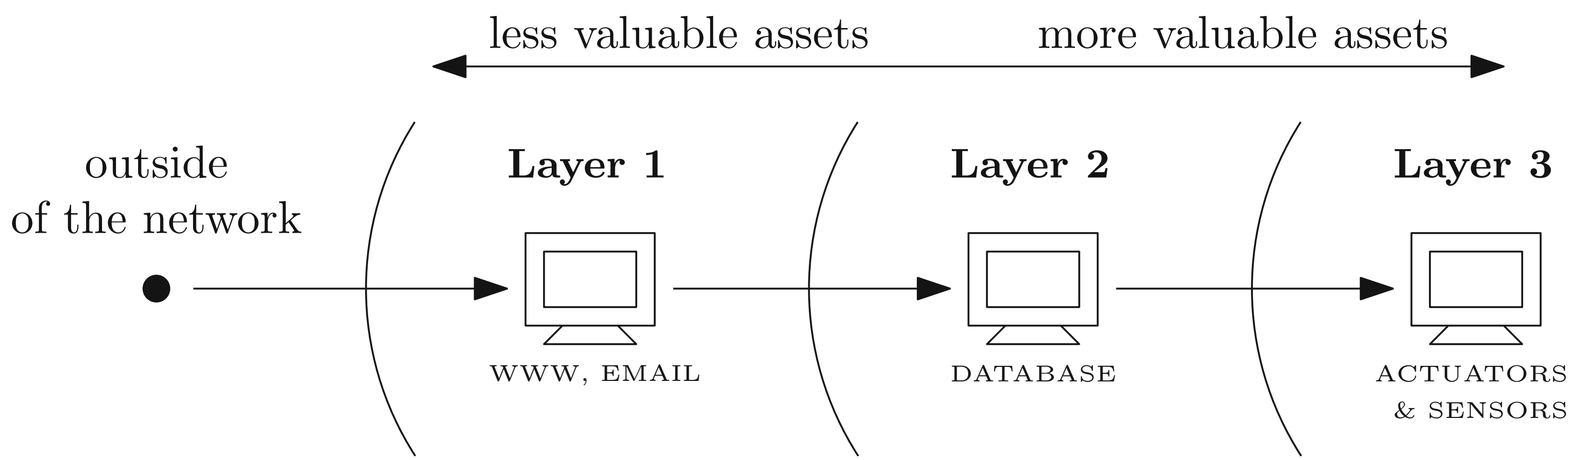

- Possible network topologies
- Attack vectors: $X_A = \{\emptyset, Layer_1, Layer_2, Layer_3\}$
- Defense vectors: $X_D = \{\emptyset\}$, i.e., deploy no dynamic resources
- Detection states: detected or not.

Observable to attacker

Not observable to attacker

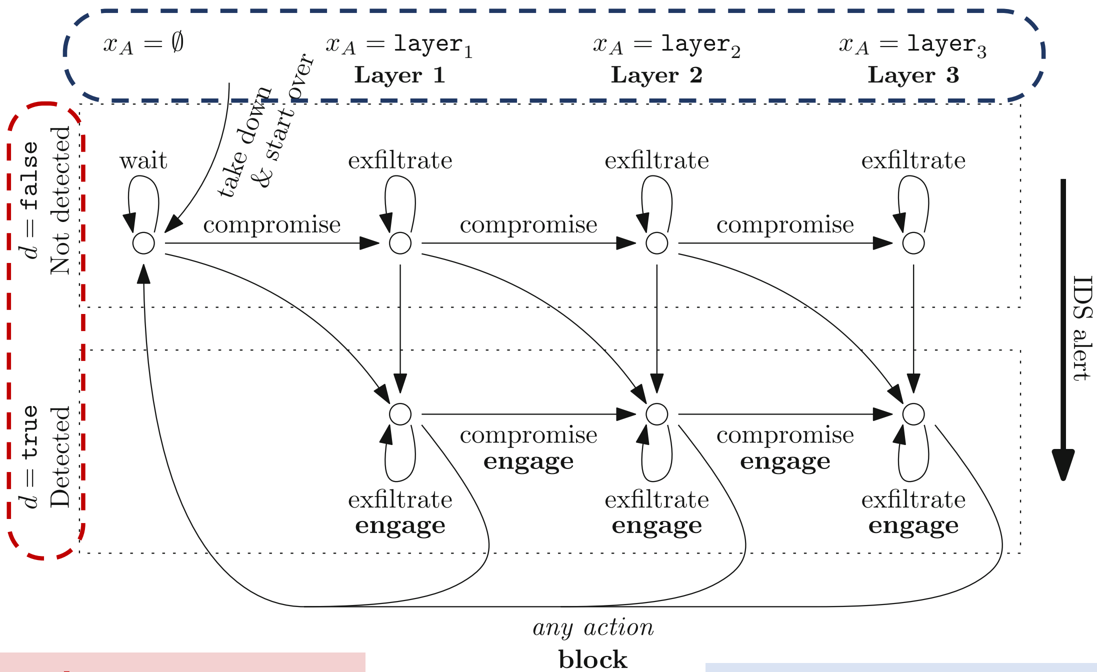

Attacker's actions

- Compromise: Go deeper.
- Exfiltrate: Stay and gain access to confidential info.
- Takedown: Incur immediate damage and get detected

Defender's Action

- If not detected, defender does nothing.
- If detected, defender's action
  - Block: remove the attacker
  - Engage: present falsified data to the attacker

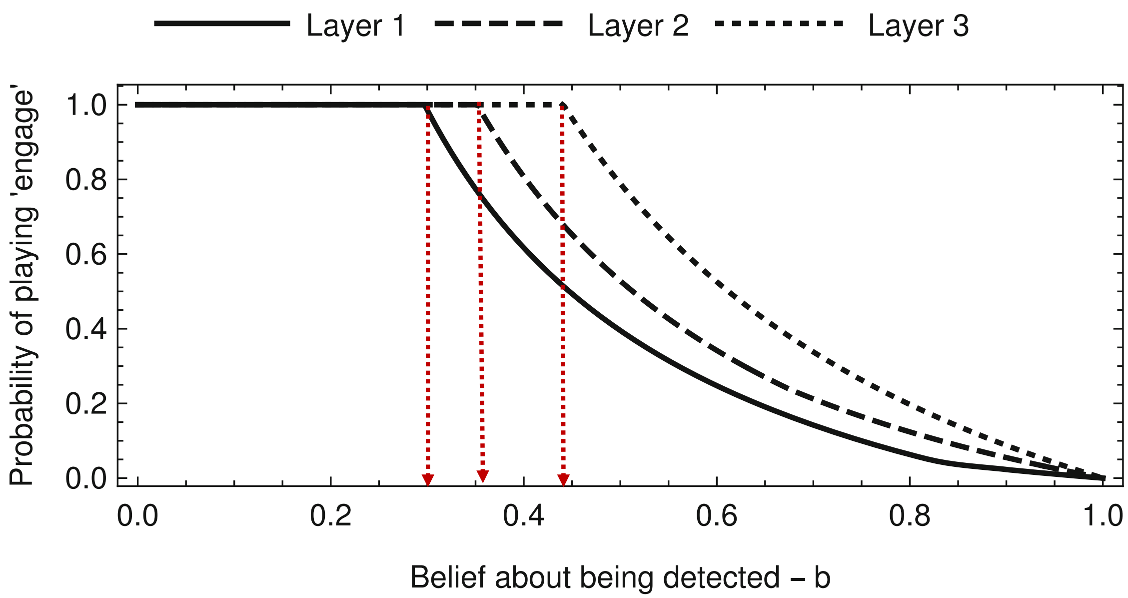

Blocking threshold

When attacker's confidence is below the threshold, the defender engages with prob. 1.

- Optimal defense strategy:
  - Engage the attacker who believes that he has not been detected
  - Block others
- Demise of the greedy:
  - The blocking threshold increases when the attacker is closer to the goal of deeper layer penetration.
  - Attacker cares less about being detected when getting closer to the asset.
  - Less stringent on the belief for engagement when closer to the asset.

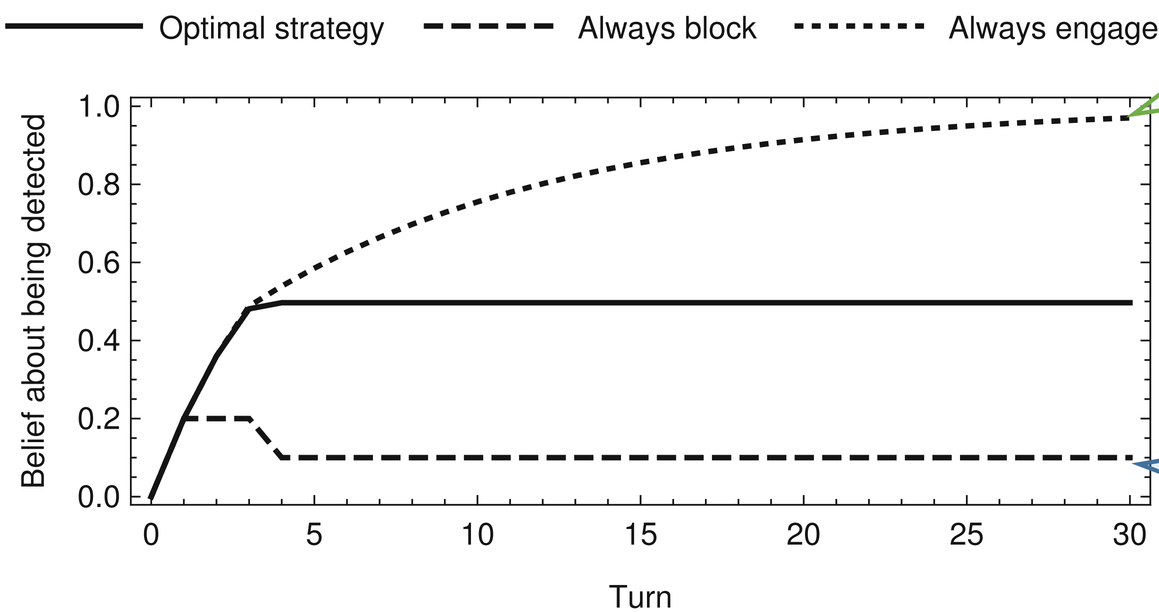

Always engaging strategy

- Too good to be true: the attacker realizes the engagement.
- A bird in the hand is worth two in the bush: the attacker chooses immediate damage.

Always blocking strategy

- Temporary solution: attackers can reenter later.
- Curse of exclusion: no blocking = not detected yet.

- Optimal strategy
  - Stabilize the attacker's belief at around 0.5.
  - Attacker's tradeoff of data exfiltration or being manipulated.

Talk Outline

1) Introduction
2) Taxonomy of defensive deception
3) Signaling games for mimetic cyber deception
   - Honey-X
   - Attack Engagement
4) Dynamic games for cyber-physical deception
   - Robotic Deception
   - Conjectural Meta-Learning
5) Future challenges

Robotic Deception

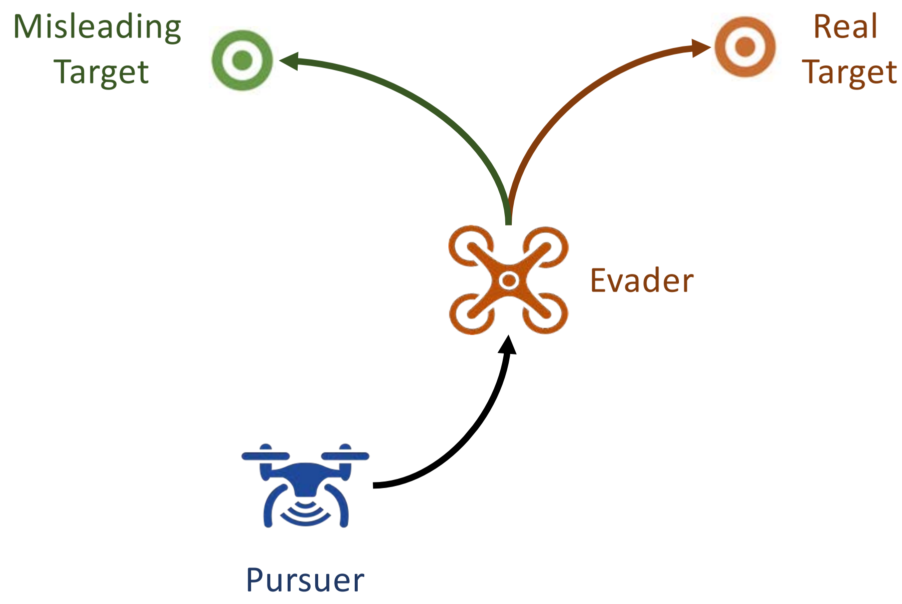

- The evader aims to reach his real target and keep a distance from the pursuer.
- The evader does not want to reveal his real target.
- The pursuer goes after the evader.

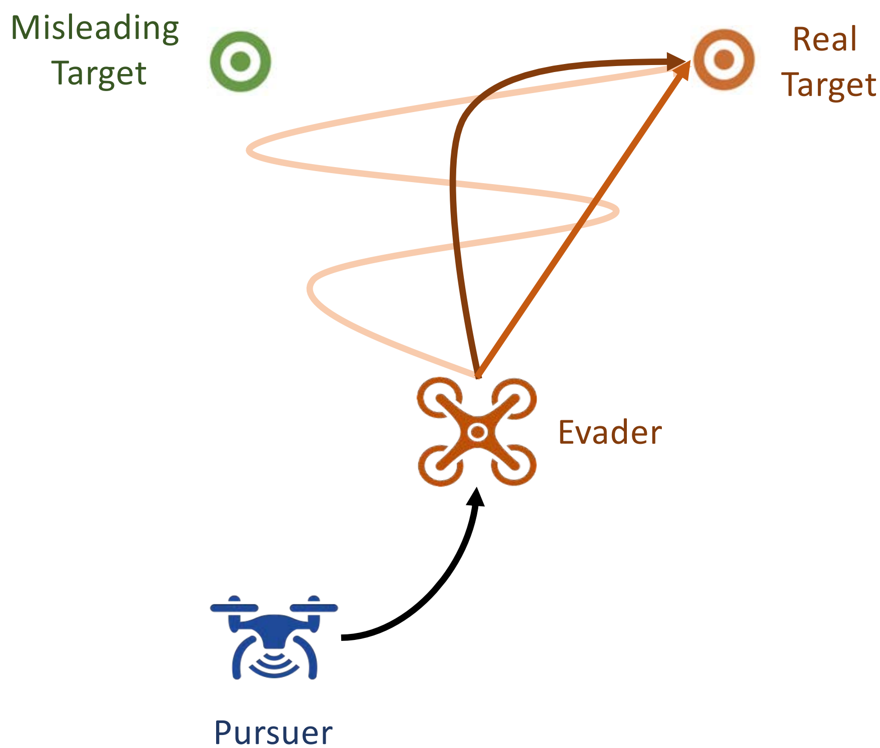
Misleading Target
Real Target
Evader
Pursuer

$$x^{k+1} = f^k(x^k, a_1^k, a_2^k, \theta_1, \theta_2, w^k)$$

$$\mathbb{E}_{\theta_{-i}, \mathbf{w}} J_i(\mathbf{x}, \mathbf{a}_1, \mathbf{a}_2, \theta_1, \theta_2)$$

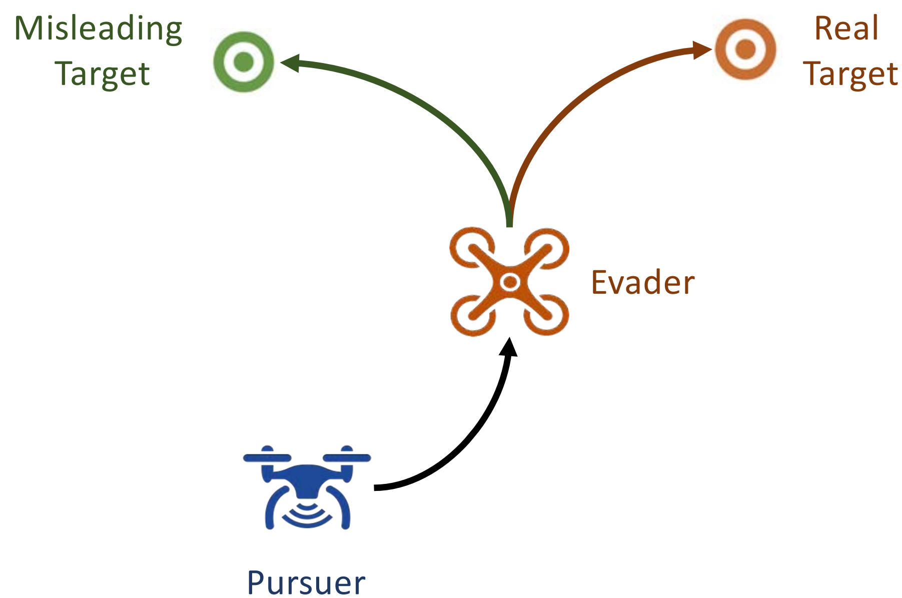

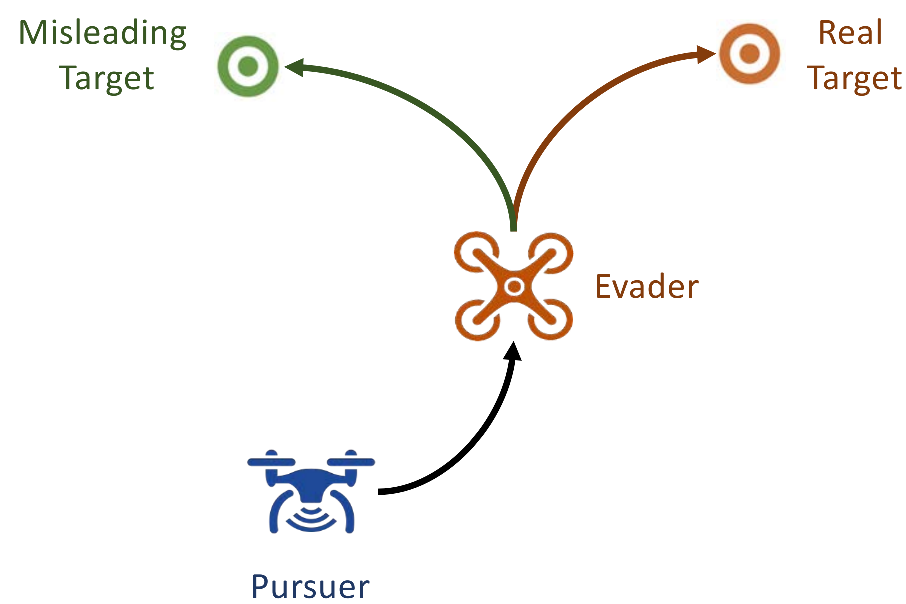

$$x^{k+1} = A^k(\theta)x^k + \sum_{i=1}^{N} B_i^k(\theta_i)u_i^k + w^k.$$

$$g_i^k(x^k, u^k, \theta_i) = (x^k - x_{d_i}^k)' D_i^k(\theta_i)(x^k - x_{d_i}^k) + \sum_{j=1}^{N} (u_j^k)' F_{ij}^k(\theta_i) u_j^k$$

Solution Concept (Informal, (Huang and Zhu, 2021))

- Sequential Rationality: Control $u^{*,0:K-1}$ is sequential rational for each player i under his belief sequence $b^{*,0:K-1}$.

- Belief consistency: Each player *i's* belief sequence $b^{*,0:K-1}$ is consistent with rationality under control $u^{*,0:K-1}$.

Huang L, Zhu Q. A dynamic game framework for rational and persistent robot deception with an application to deceptive pursuit-evasion. IEEE Transactions on Automation Science and Engineering. 2021.

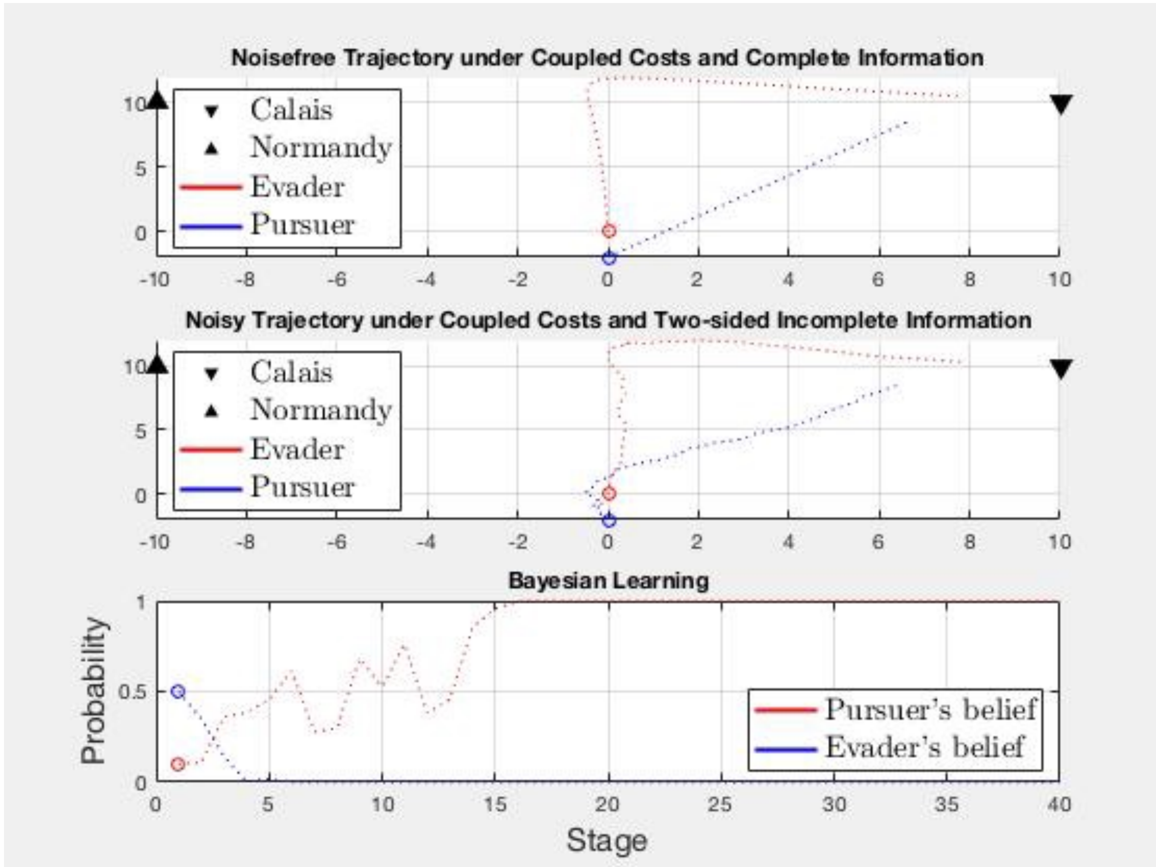
Noisefree Trajectory under Coupled Costs and Complete Information
Calais
Normandy
Evader
Pursuer
Noisy Trajectory under Coupled Costs and Two-sided Incomplete Information
Calais
Normandy
Evader
Pursuer
Bayesian Learning
Probability
Pursuer's belief
Evader's belief
Stage

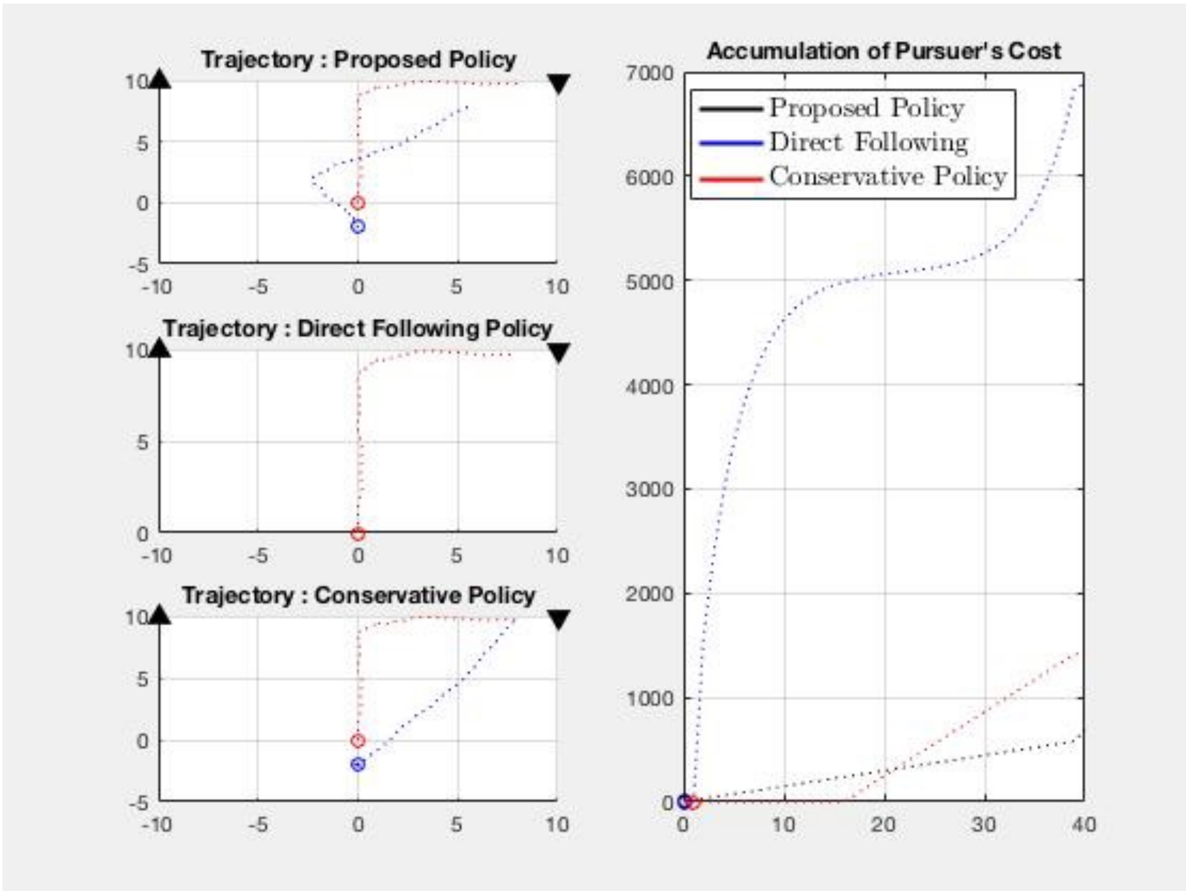

Trajectory : Proposed Policy
Trajectory : Direct Following Policy
Trajectory : Conservative Policy
Accumulation of Pursuer's Cost
Proposed Policy
Direct Following
Conservative Policy

Future Challenges: Learning-Based Solutions

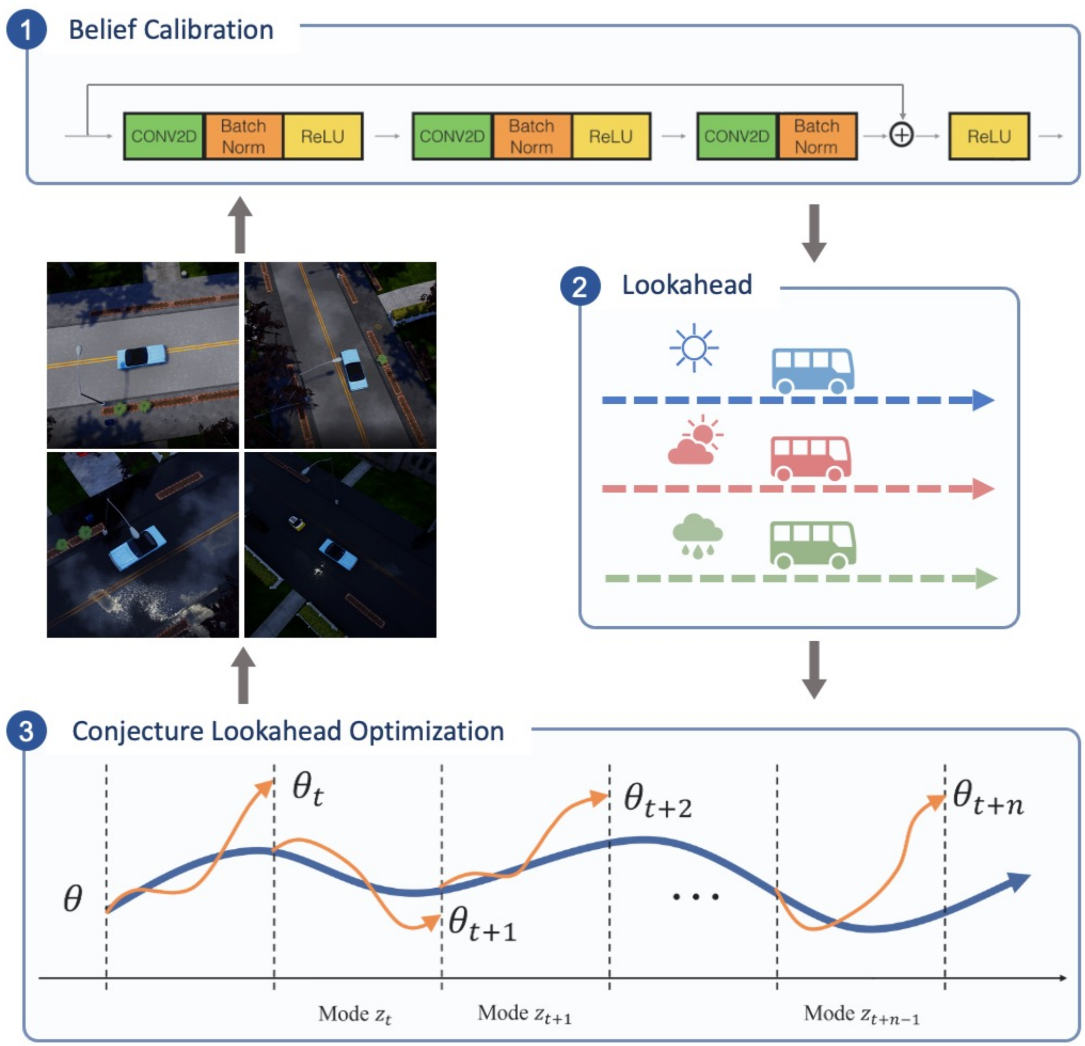

Li T, Lei H, Zhu Q. Self-Adaptive Driving in Nonstationary Environments through Conjectural Online Lookahead Adaptation. arXiv preprint arXiv:2210.03209. 2022 Oct 6.

Future Challenges: Human

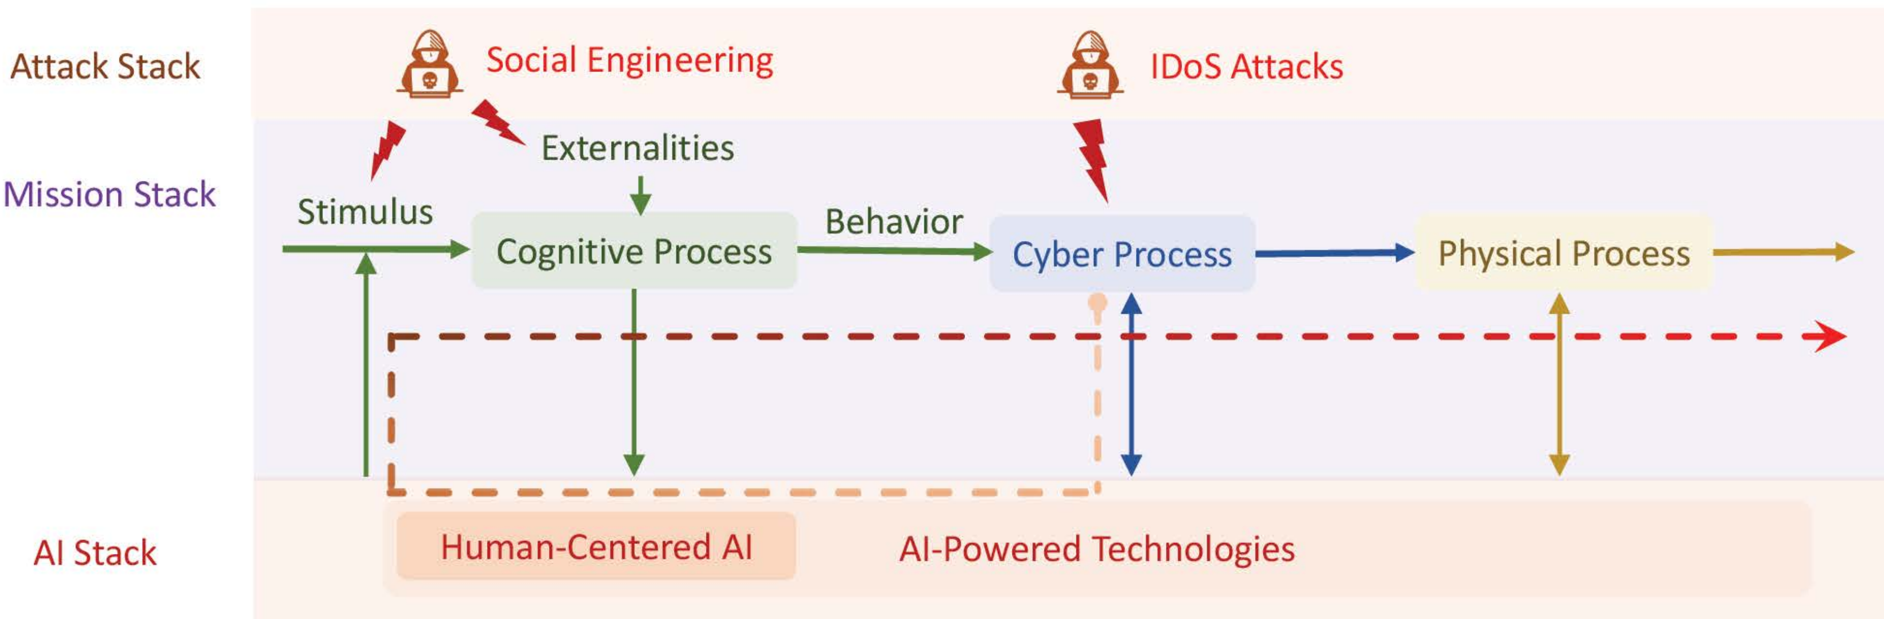
Attack Stack
Social Engineering
IDoS Attacks
Mission Stack
Externalities
Stimulus
Cognitive Process
Behavior
Cyber Process
Physical Process
AI Stack
Human-Centered AI
AI-Powered Technologies

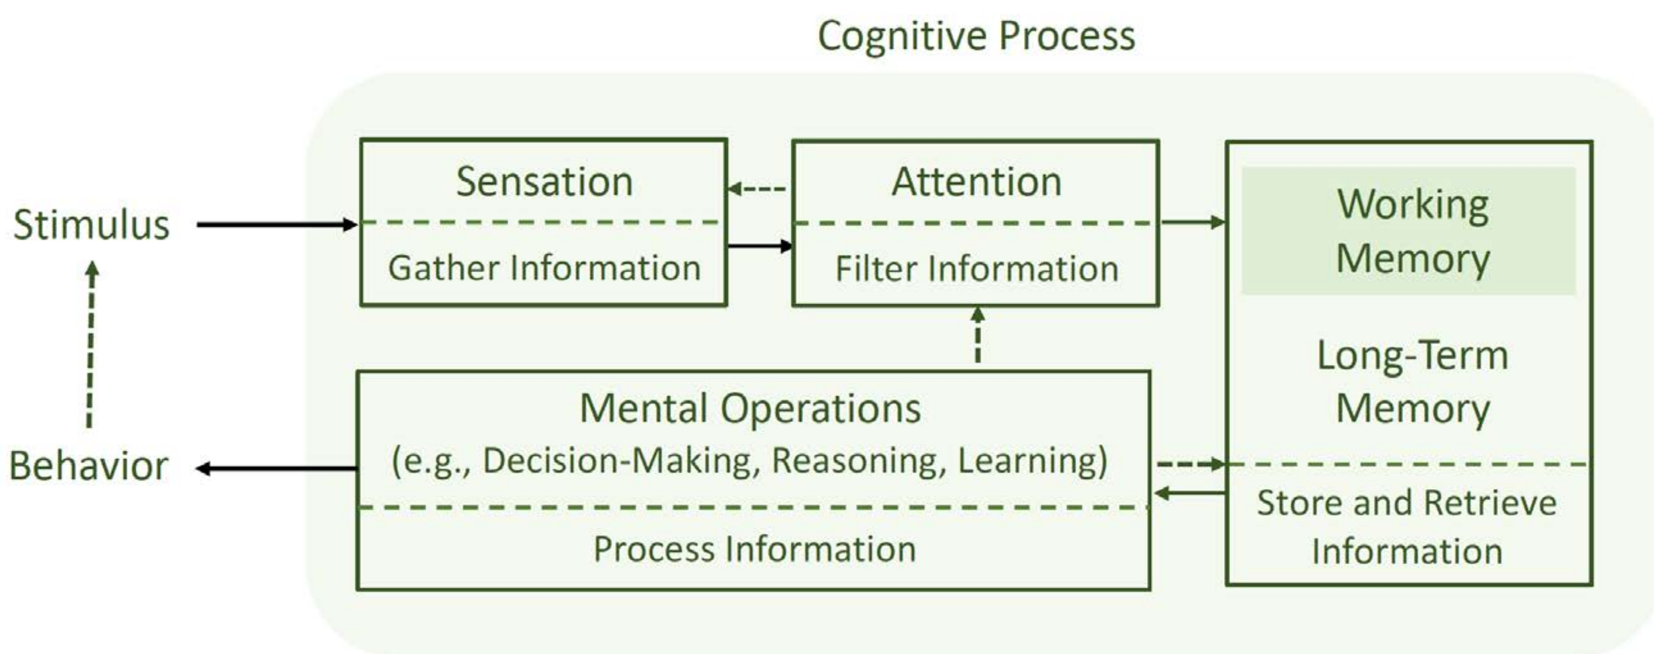

[Huang and Zhu, 2023]

Future Challenges: Mechanism Design

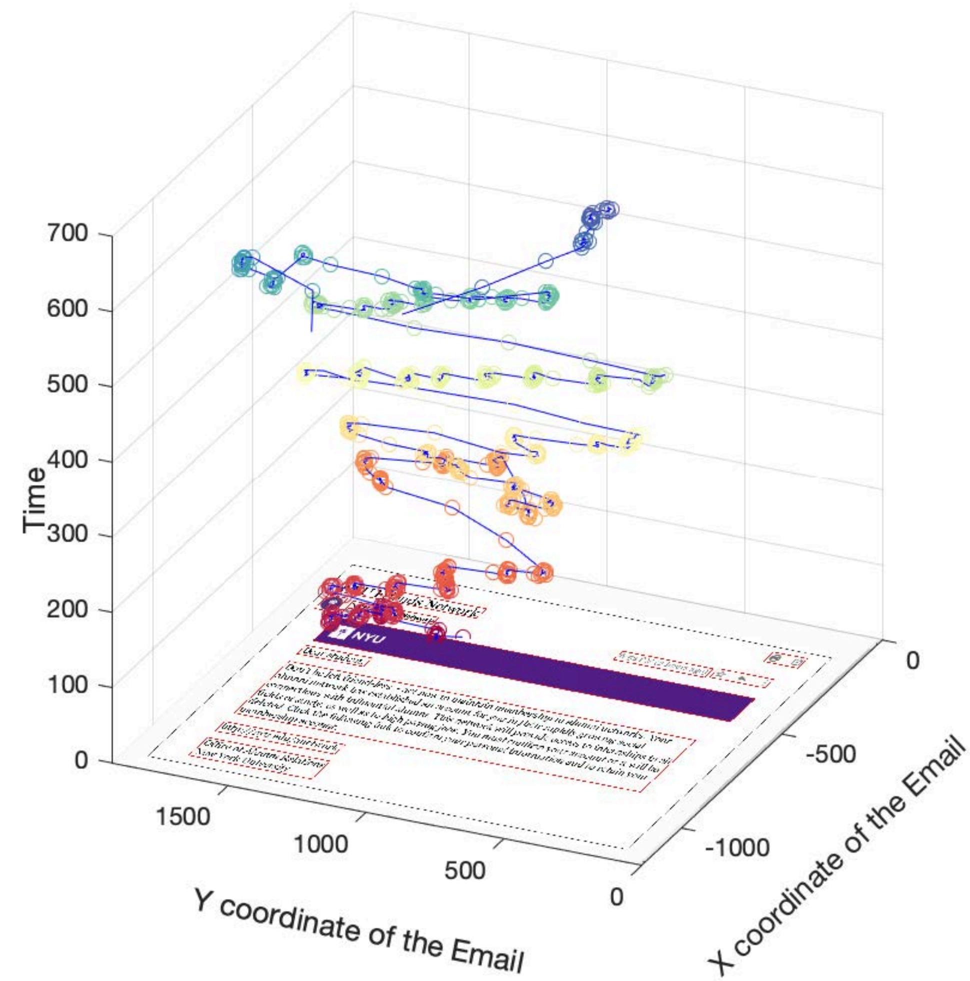

Time
700
600
500
400
300
200
100
0
1500
1000
500
0
Y coordinate of the Email
0
-500
-1000
X coordinate of the Email
NYU

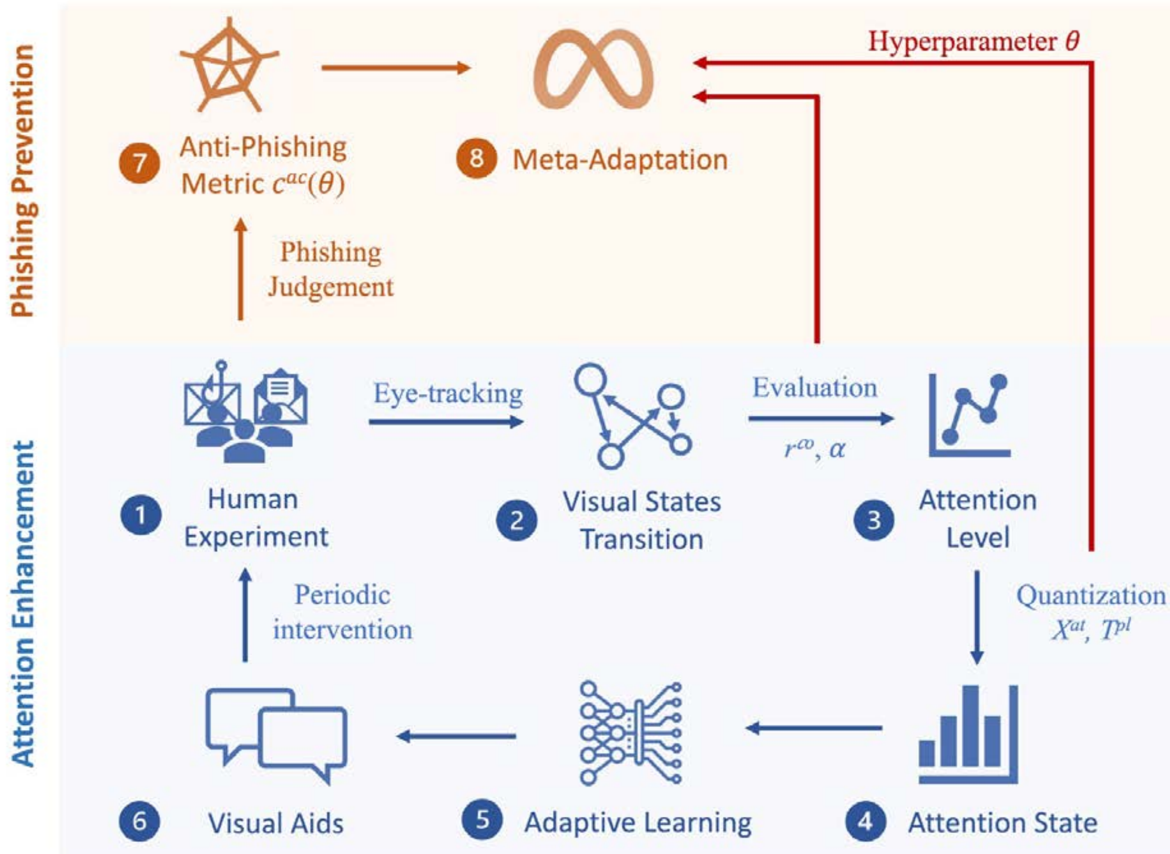

Huang L, Zhu Q. RADAMS: Resilient and adaptive alert and attention management strategy against informational denial-of-service (IDoS) attacks. Computers & Security. 2022 Oct 1;121:102844.

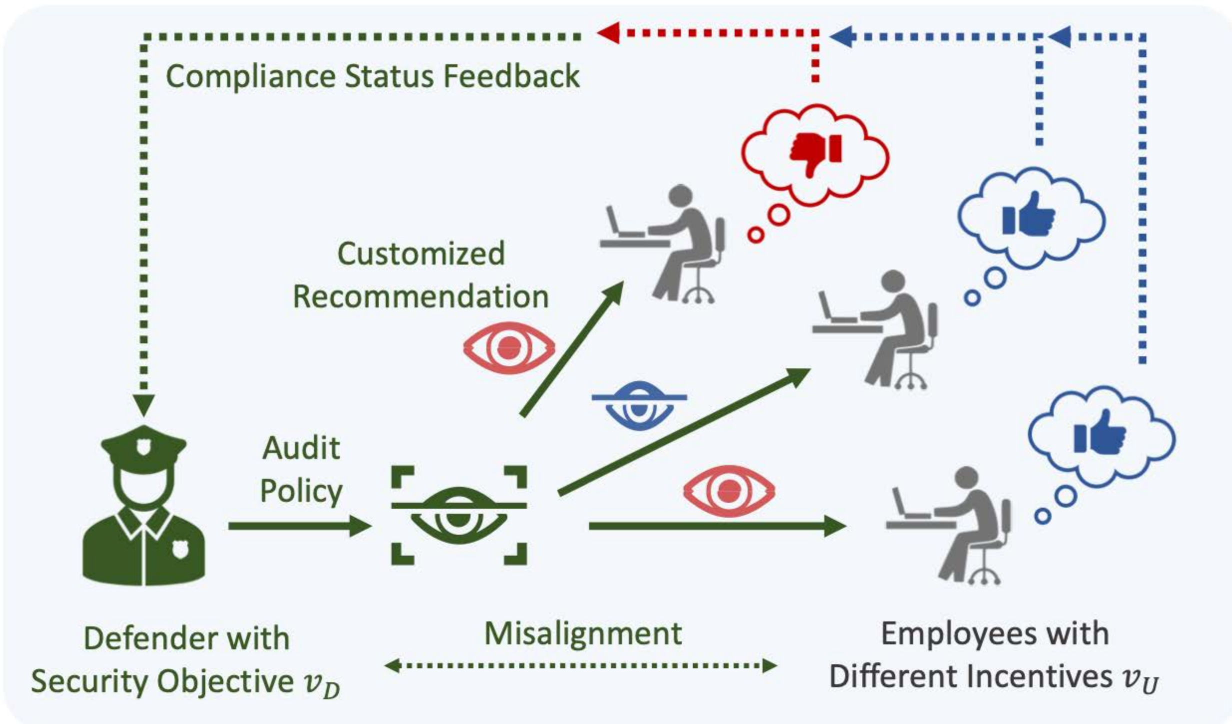

Huang L, Zhu Q. Duplicity games for deception design with an application to insider threat mitigation. IEEE Transactions on Information Forensics and Security. 2021 Oct 8;16:4843-56.

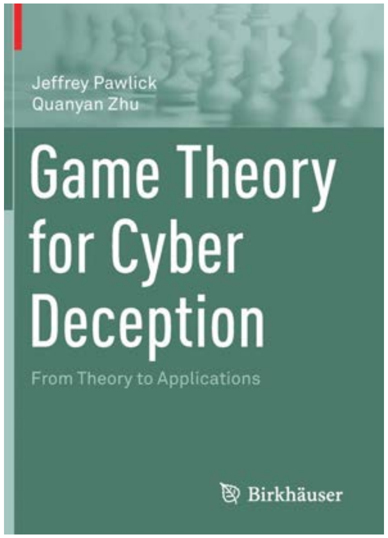

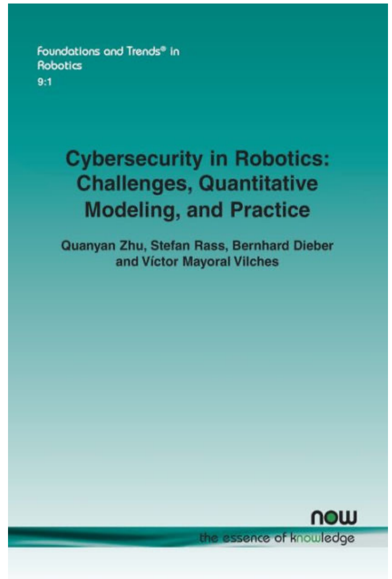

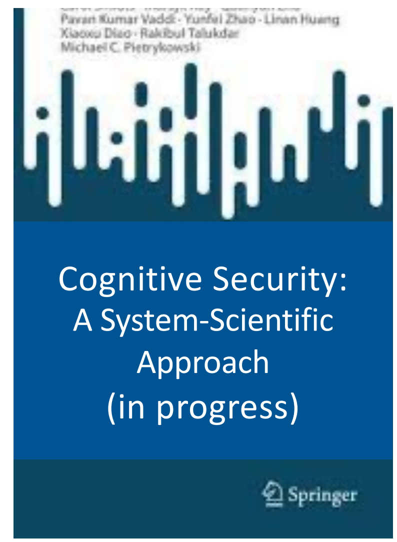

Contact: quanyan.zhu@nyu.edu