Optimal Security Investments in a Prevention and Detection Game

Carlos Barreto, Carlos.BarretoSuarez@utdallas.edu
Alvaro A. Cárdenas, Alvaro.Cardenas@utdallas.edu
Alain Bensoussan, Alain.Bensoussan@utdallas.edu

University of Texas at Dallas

Hot Topics in the Science of Security Symposium 2017

Problem: How to invest in security?


WIRED SUBSCRIBE

BRUCE SCHNEIER BUSINESS 04.19.07 12:00 PM

HOW SECURITY COMPANIES SUCKER US WITH LEMONS

The White House

Office of the Press Secretary

For Immediate Release February 09, 2016

FACT SHEET: Cybersecurity National Action Plan

Taking bold actions to protect Americans in today's digital world.

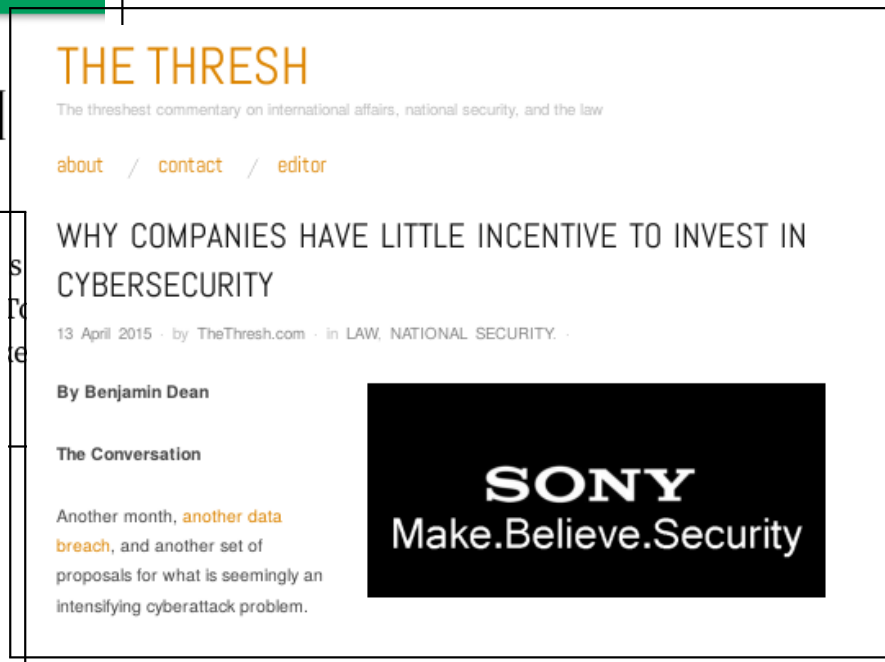
THE THRESH

The threshest commentary on international affairs, national security, and the law

about / contact / editor

WHY COMPANIES HAVE LITTLE INCENTIVE TO INVEST IN CYBERSECURITY

13 April 2015 · by TheThresh.com · in LAW, NATIONAL SECURITY.

By Benjamin Dean

The Conversation

Another month, another data breach, and another set of proposals for what is seemingly an intensifying cyberattack problem.

SONY
Make.Believe.Security

Although security is important, firms fail to protect systems because they

- underestimate their exposure
- lack incentives
- ignore the cost/benefit of security
- firms do not know the best way to protect a system

Related works

Previous work on increasing security investments:

Interdependences: Deal with the negative effects of networked systems, which create cooperation problems.

Cyber-Insurance: Tool that might give incentives to invest in protection.

How can we protect systems?[1]

[1] New York State Department of Financial Services: Report on Cyber Security in the Insurance Sector, Feb. 2015, URL: http://www.dfs.ny.gov/reportpub/dfs_cyber_insurance_report_022015.pdf.

Objective: Investigate the best investment strategy to protect a system

We propose a model of the interactions between a defender and an attacker where

Defender invest in two technologies

- Prevention
- Detection

Attacker invest its resources in

- Finding vulnerabilities
- Attacking the system

Questions:

How does the attacker's strategy change as a function of the defense strategy?

How does the defense strategy change with limited resources? With limited information?

Outline

Players

Attacker

Objective Maximize its profit attacking firms (e.g., stealing information)

Actions

- Find bugs (hack the system) $v_h \in [0, 1]$
- Exploit bugs $v_e \in [0, 1]$

Defender

Objective Minimize operation costs of a system. Balance between costs of attacks and cost of protection

Actions

- Prevent bugs in the system $v_p \in [0, 1]$ (e.g., secure code development)
- Detect attacks and correct failures $v_d \in [0, 1]$ (e.g., IDS)

The cost of each player is affected by the decisions of the adversary.

Security Model

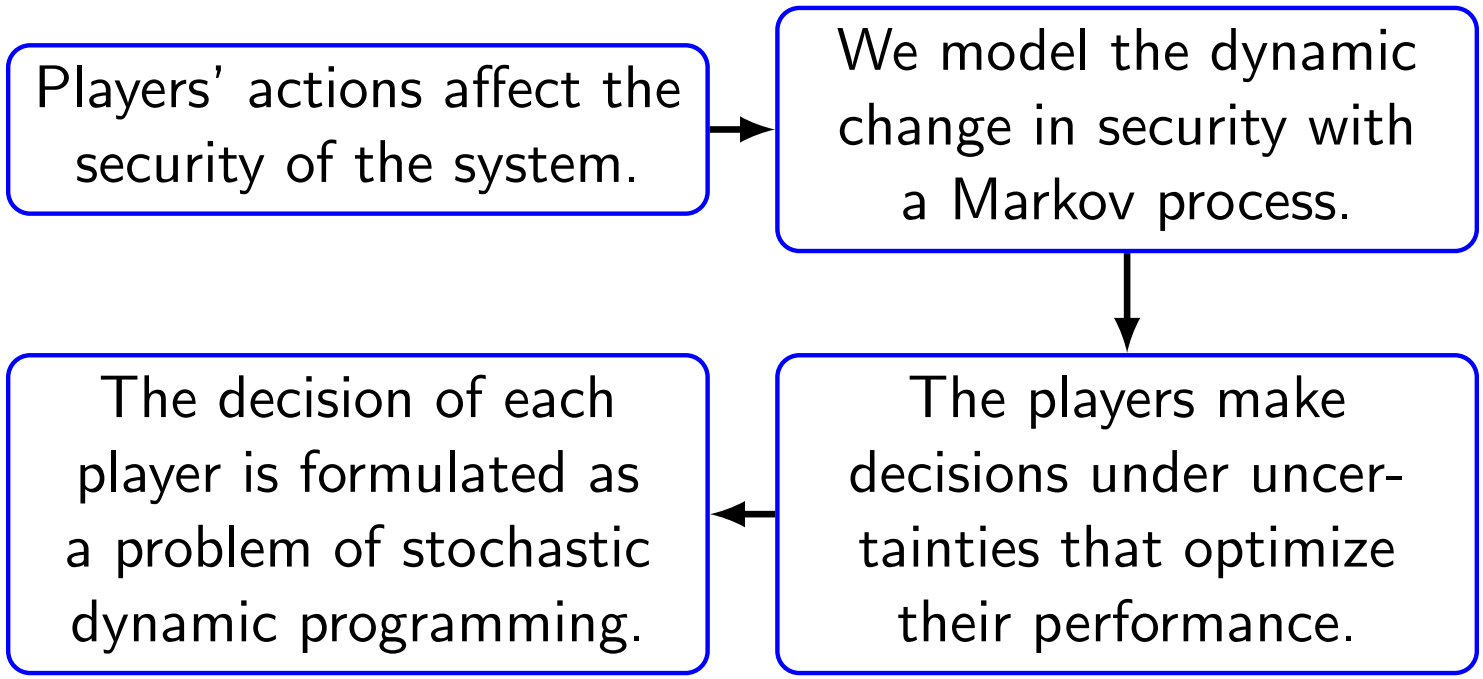

Problems of stochastic dynamic programming[2] involve solving iteratively a Bellman equation that describes the conditions of optimal decisions.

[2] Alain Bensoussan: Dynamic programming and inventory control, vol. 3 (Studies in Probability, Optimization and Statistics), 2011;
Onésimo Hernández-Lerma/Jean B Lasserre: Discrete-time Markov control processes: basic optimality criteria, vol. 30, 2012.

System's Security as a Markov Decision Process

Vulnerable state S_0

An adversary can exploit a vulnerability.

Secure state S_1

The adversary must search a vulnerability to attack.

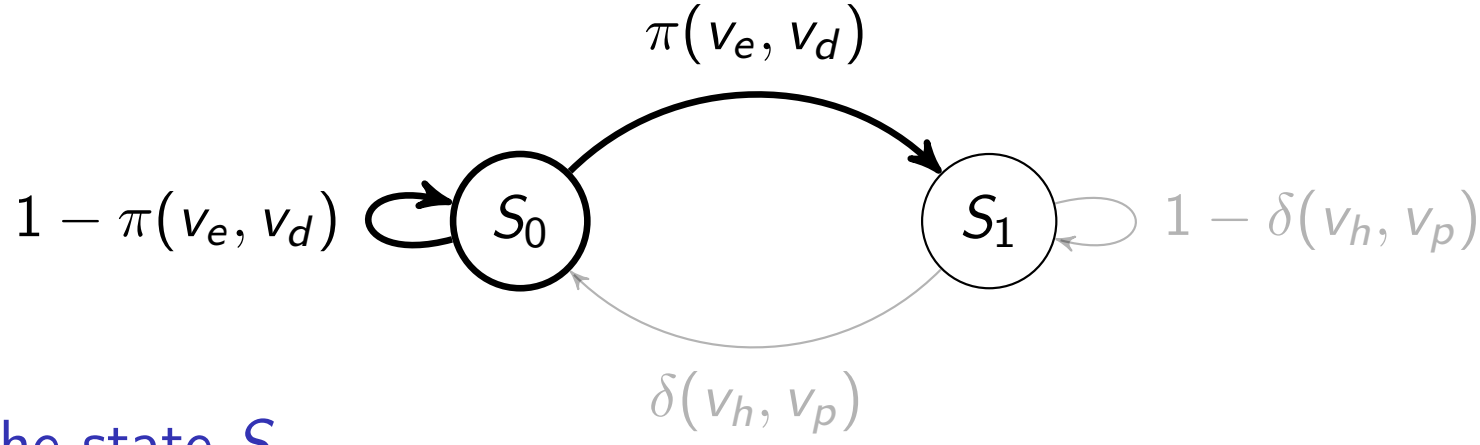

In the state S_0

Attacker

Gains: $g_a(v_e)$

Cost: C_0

$l_A = -ga(v_e) + C_0$

Defender

Loses: $g_d(v_e)$

Cost: $C_d(v_d) + C_p(v_p)$

$l_D = g_d(v_e) + C_d(v_d) + C_p(v_p)$

The defender detects the attack with probability $\pi(v_e, v_d)$, which increases with v_e and v_d

System's Security as a Markov Decision Process

Vulnerable state S_0

An adversary can exploit a vulnerability.

Secure state S_1

The adversary must search a vulnerability to attack.

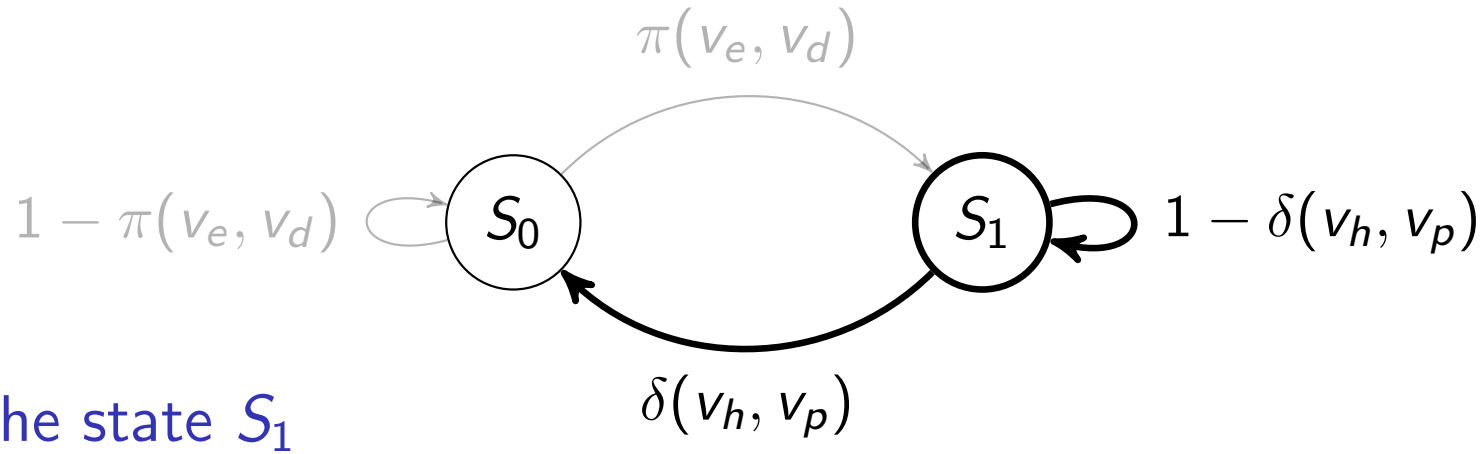

In the state S_1

Attacker

Gains: 0

Cost: C_v

$l_A = C_v$

Defender

Loses: 0

Cost $= C_d(v_d) + C_p(v_p)$

$l_D = C_d(v_d) + C_p(v_p)$

The attacker finds a vulnerability with probability $\delta(v_h, v_p)$.

- increases with the effort of the attacker v_h.
- decreases with the effort of the defender v_p.

Outline

Attacker's Discounted Payoff

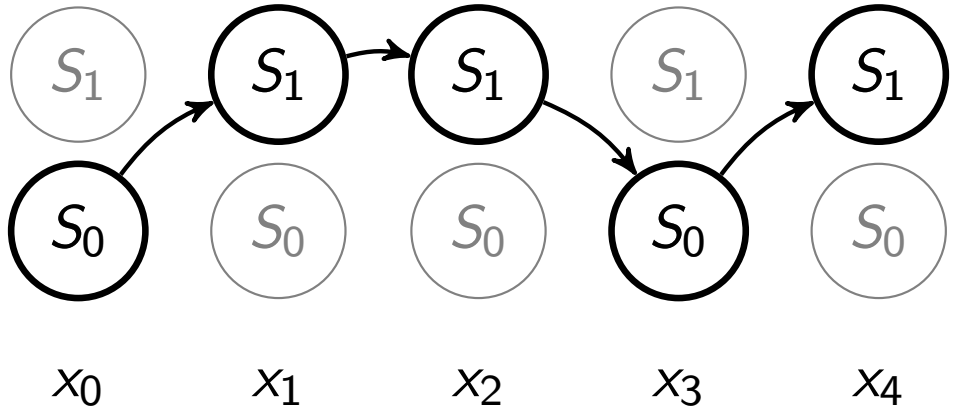

The discounted payoff of the attacker with the attack and defense strategies $v_A = (v_e, v_h)$ and $v_D = (v_d, v_p)$ is

$$
\begin{aligned}
J^A(x_0, v_A, v_D) = l_A(x_0, v_A) + \\
\beta \mathbb{E}_{x_0}^{v_A, v_D} \{ l_A(x_1, v_A) + \\
\beta \mathbb{E}_{x_1}^{v_A, v_D} \{ l_A(x_2, v_A) + \\
\beta \mathbb{E}_{x_2}^{v_A, v_D} \{ l_A(x_3, v_A) + \\
\vdots \\
+ \beta \mathbb{E}_{x_{n-1}}^{v_A, v_D} \{ l_A(x_n, v_A) + \ldots \} \} \} \}
\end{aligned}
$$

The discount factor β relates future costs with the present.

Attacker's Discounted Payoff

We consider an infinite horizon problem in which the attacker wants to find the best attack strategy v_A. The cost functional can be written as

$$J^A(x_0, v_A, v_D) = \overbrace{l_A(x_0, v_A)}^{\text{Present Cost}} + \beta \overbrace{\mathbb{E}_{x_0}^{v_A, v_D}\{J^A(x_1, v_A, v_D)\}}^{\text{Future Cost}},$$

where x_0 is the initial state.
The minimum cost is given by the Bellman equation

$$u^A(x_0, v_D) = \min_{v_A} J^A(x_0, v_A, v_D) =$$
$$\min_{v_A} \left\{ l_A(x_0, v_A) + \beta \mathbb{E}_{x_0}^{v_A, v_D} \left\{ u^A(x_0, v_D) \right\} \right\}$$

The optimal attack strategy v_A^* satisfies

$$u^A(x_0, v_D) = J^A(x_0, v_A^*, v_D)$$

Optimal Attack strategy: Procedure

1. Show that the cost functional is a contraction mapping
2. From the Banach Fixed point theorem we can approximate the cost functional as

$$u_{n+1}(x, v_d) = \inf_{v_n \in [0,1]} \{ l_A(x, v_n) + \beta \mathbb{E}_x^{v_n, v_D} \{ u_n(x, v_d) \} \},$$

   where $u_n(x, v_d) \to u(x, v_d)$ as $n \to \infty$.
3. We can analyze the optimal actions of the attacker with the approximated function.

Optimal Attack strategy

Theorem: Optimal strategy of the attacker

1. $v_a = 0$ and $v_h = 0$ if $K > 0$,
2. $v_a = 1$ and $v_h = 0$ if $K < 0$ and $B > 0$,
3. $v_a = 1$ and $v_h = 1$ if $K < 0$ and $B < 0$,

where

$$K = \underbrace{C_0 - g_a(1)}_{\text{Independent of } v_D}, \qquad B = \underbrace{C_v + \beta \frac{K}{1 + \beta\pi(1, v_d) - \beta} \delta(1, v_p)}_{\text{Increases with } v_d, v_p}.$$

Notes

- The decision to attack the system in S_0 ($v_a = 1$) depends on the profitability of the attack, not on the defense strategy.
- The defender affects the decision to hack the system through its defense strategy. B increases with both v_d and v_p.

Attacker's Hack Decision Boundary

Attacker's gain $g_a(1) = 2.5$

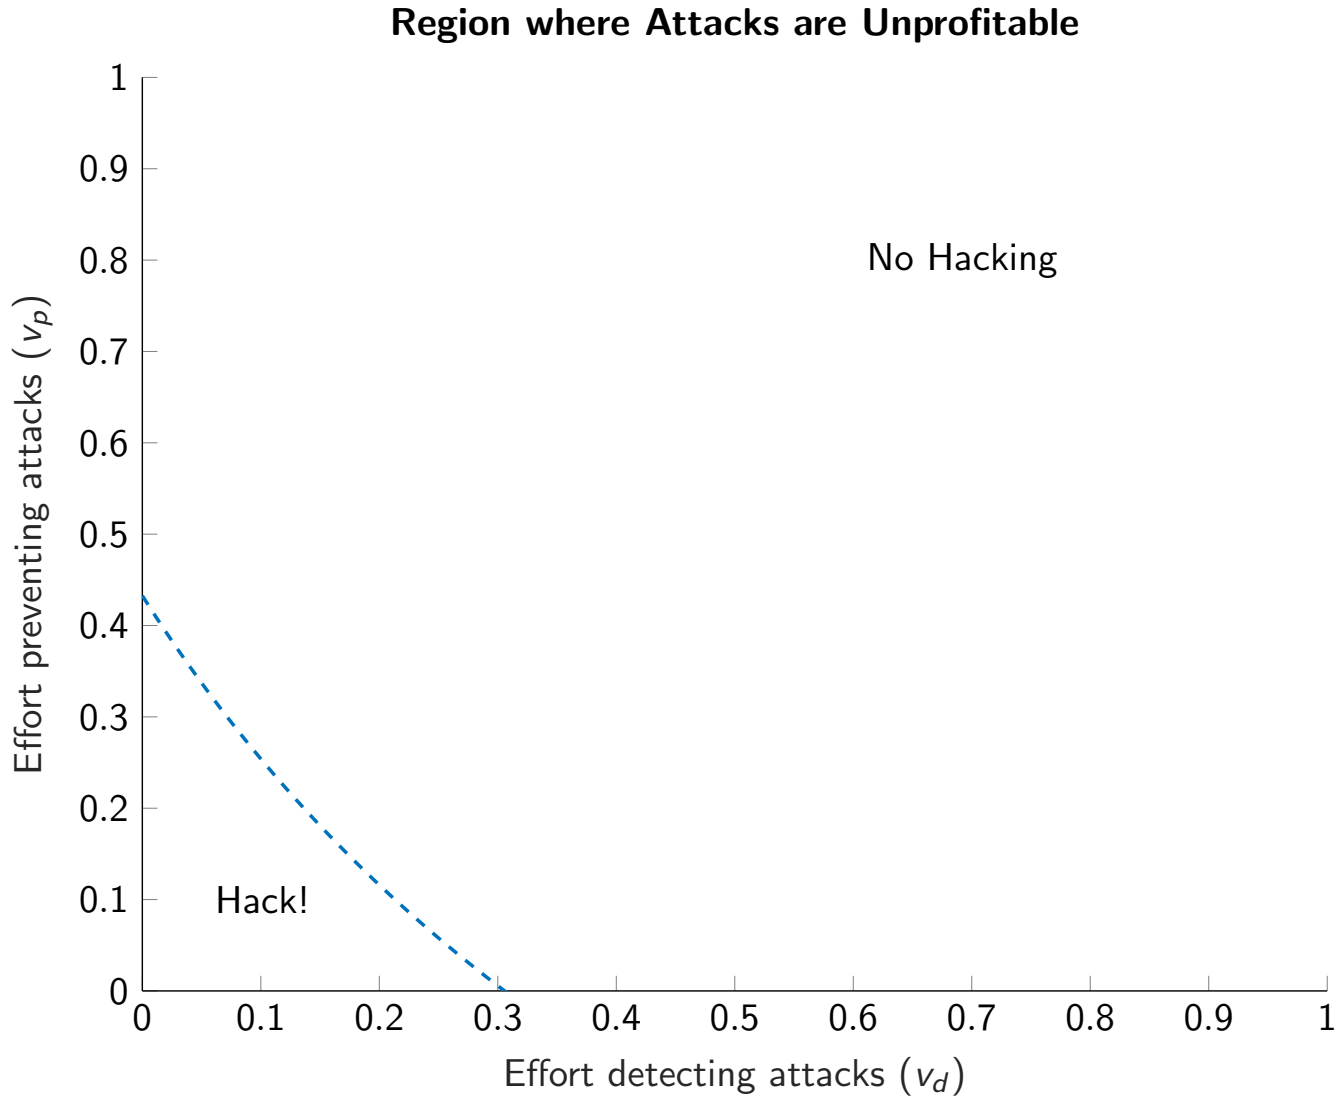

Attacker's Hack Decision Boundary

Attacker's gain $g_a(1) = 4$

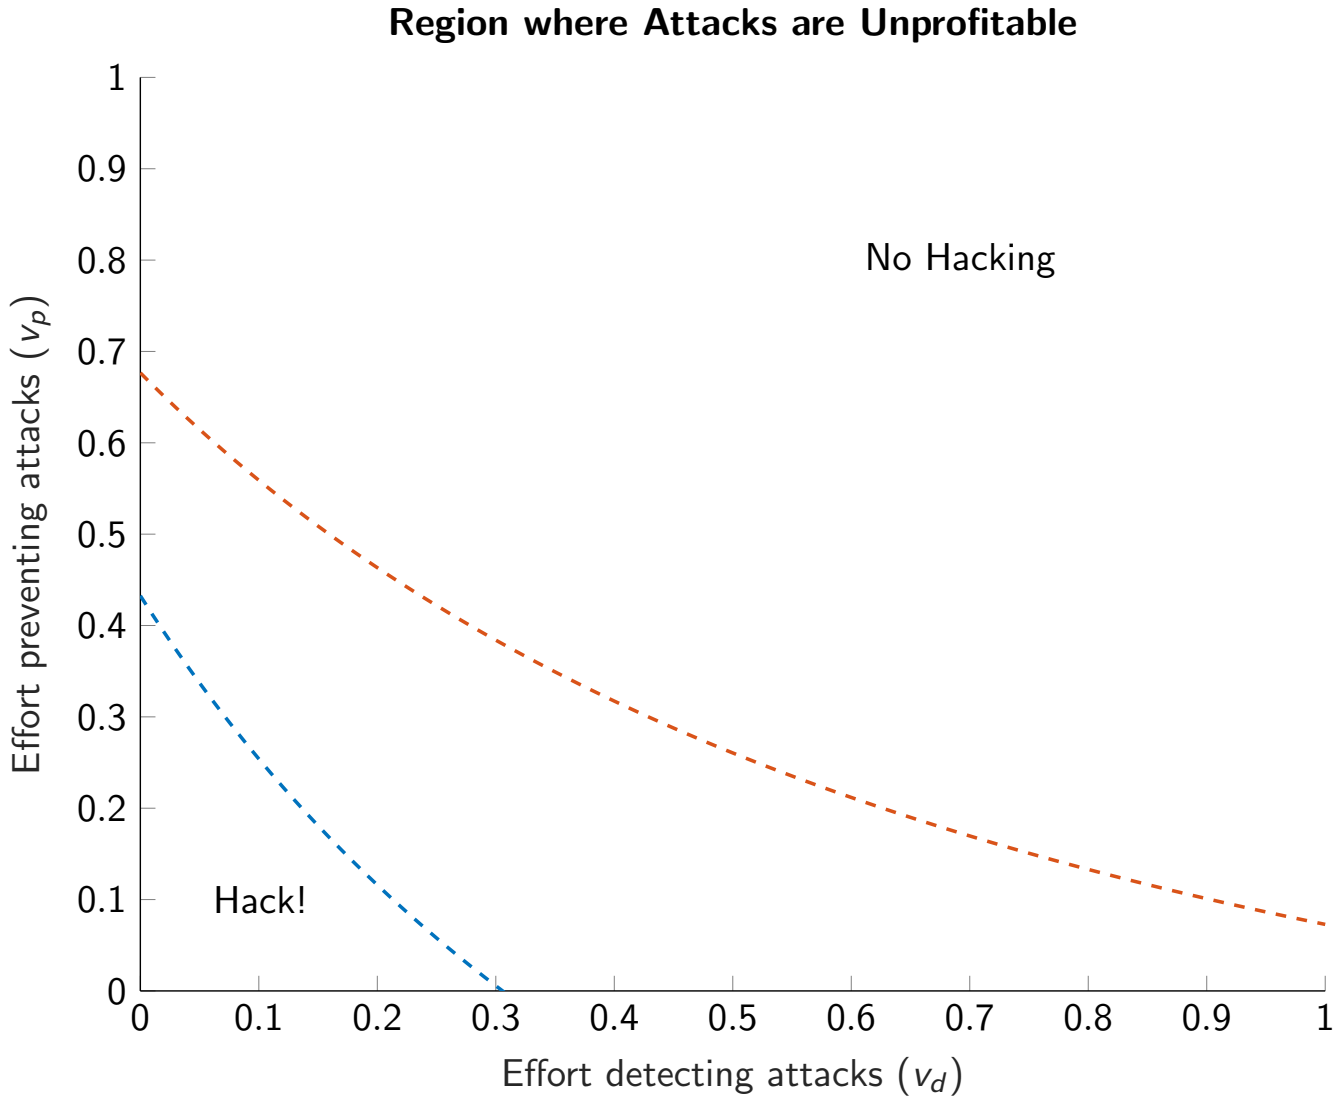

Attacker's Hack Decision Boundary

Attacker's gain $g_a(1) = 5$

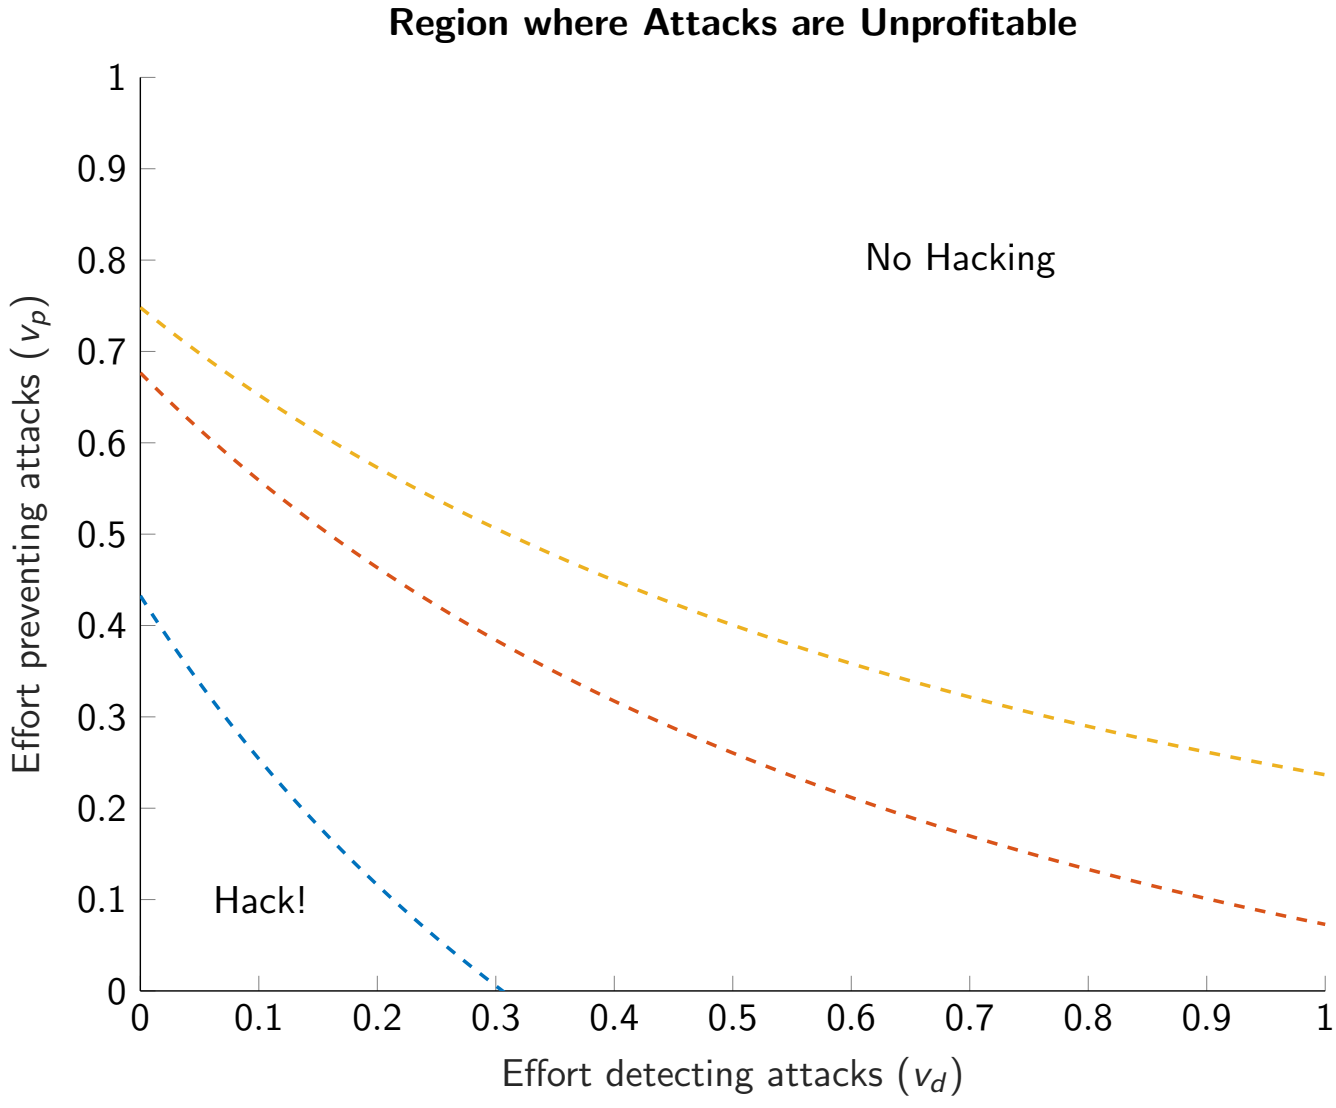

Attacker's Hack Decision Boundary

Attacker's gain $g_a(1) = 6$

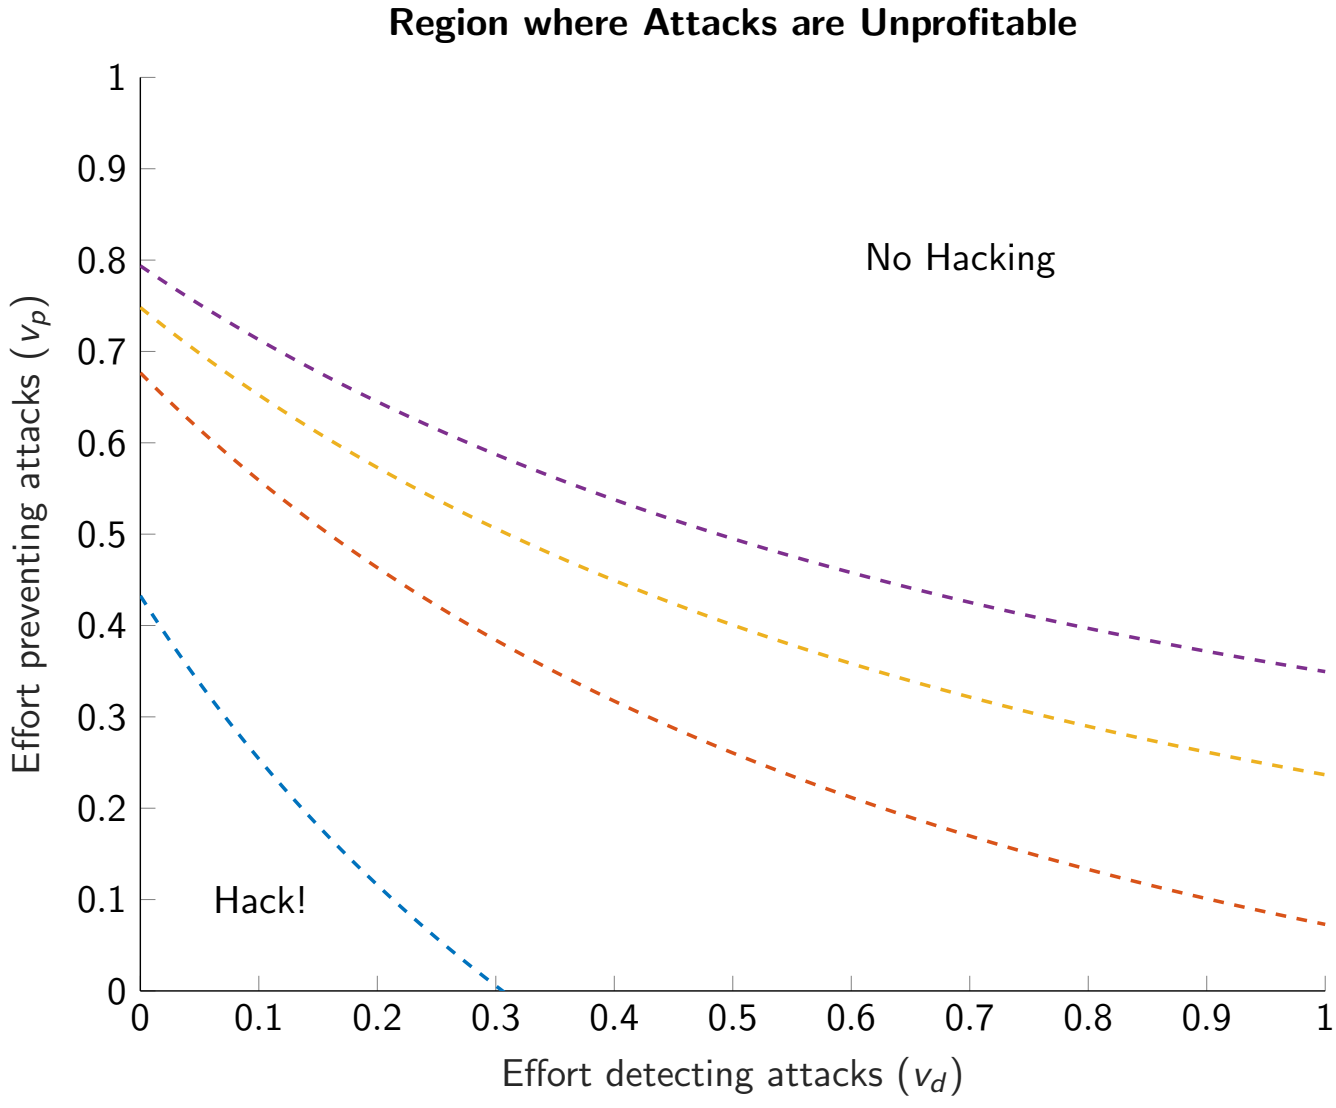

Attacker's Hack Decision Boundary

Attacker's gain $g_a(1) = 7$

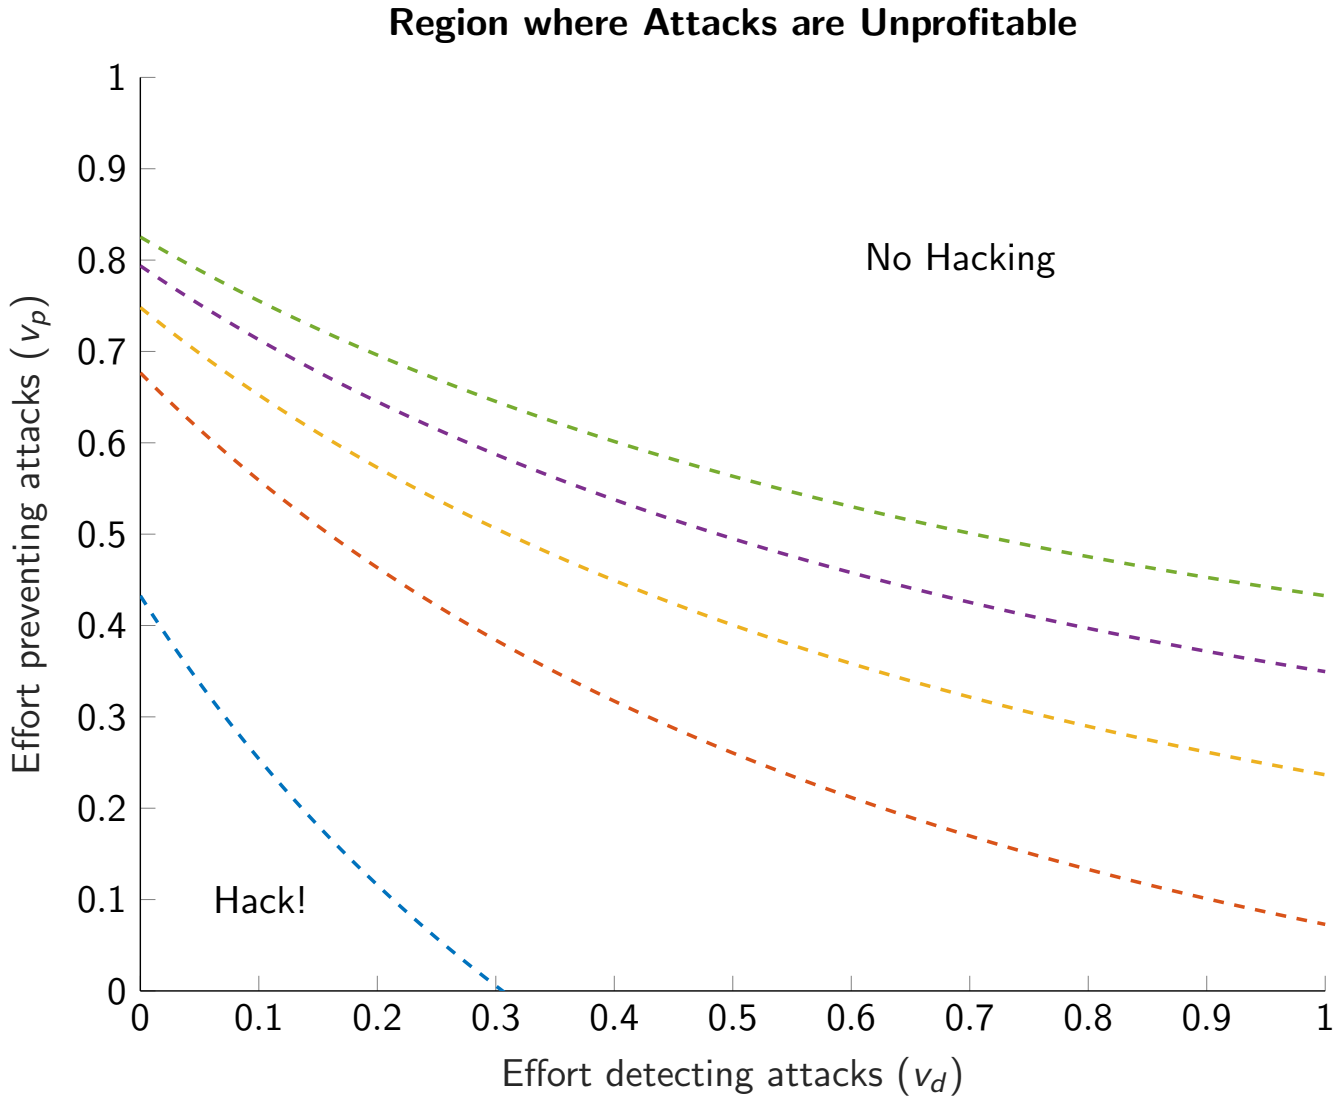

Outline

Defender Payoff

The cost of implementing the defense strategy $v_D = (v_d, v_p)$ in a time period is

$$l_D(x, v_A, v_D) = \begin{cases} \overbrace{g_d(v_e)}^{\text{Defender loss}} + C_p(v_p) + C_d(v_d) & \text{if } x = S_0, \\ \underbrace{C_p(v_p) + C_d(v_d)}_{\text{Protection cost}} & \text{if } x = S_1, \end{cases}$$

loss caused by an attack $g_d(v_e)$ is increasing with v_e.
The cost to prevent ($C_p(v_p)$) and detect ($C_d(v_d)$) attacks increase with v_p and v_d.

Defender's Objective: Full Information

The defender observes the state of the system (i.e., knows when the system is compromised, but does not know the precise cause).

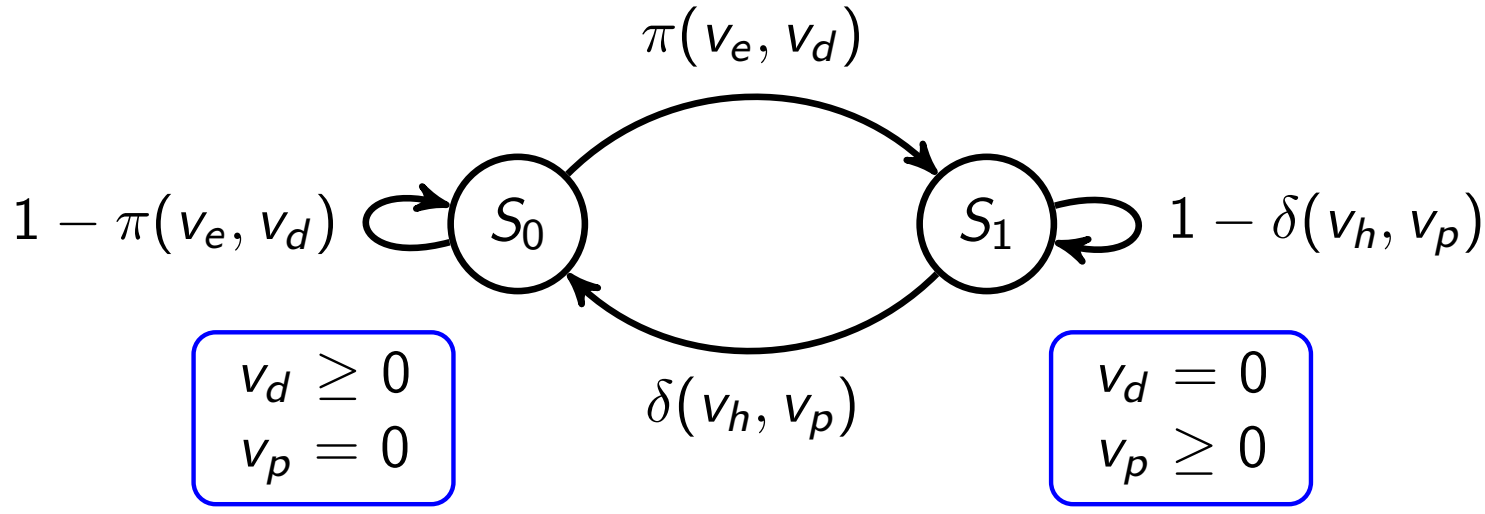

The cost functional is defined as

$$J^D(x_0, v_A, v_D) = l_D(x_0, v_A, v_D) + \beta \mathbb{E}_{x_0}^{v_A, v_D}\{J^D(x_1, v_A, v_D)\}.$$

Defender's Objective: Asymmetric Information

The defender cannot observe the state of the system, instead, has some belief about the initial state.

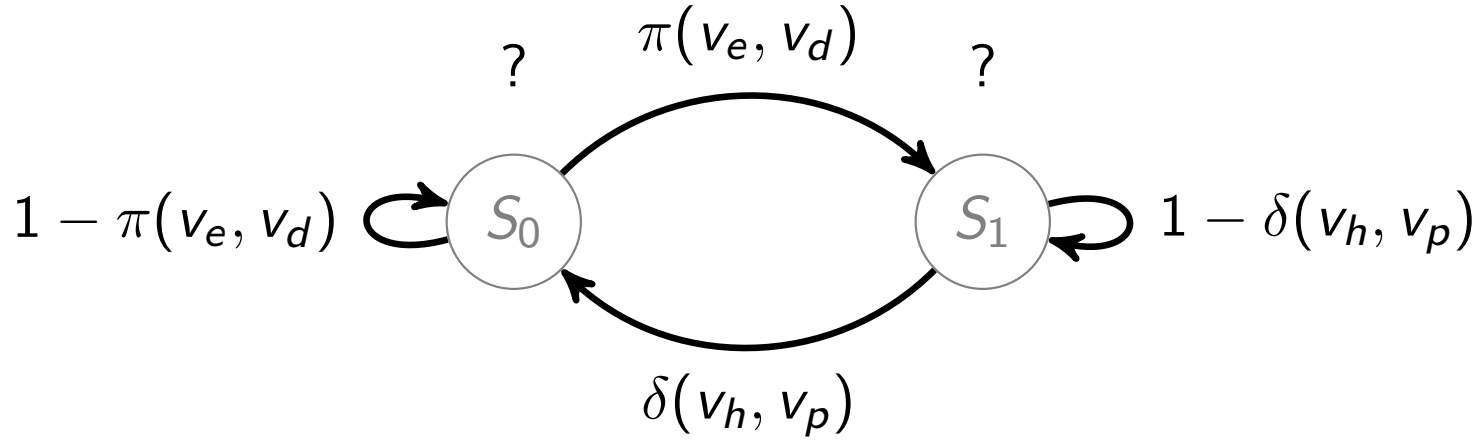

The cost function becomes

$$\hat{J}^D(v_A, v_D) = \mathbb{P}(x = S_0) l_D(S_0, v_A, v_D) + \mathbb{P}(x = S_1) l_D(S_1, v_A, v_D) + \beta \hat{J}^D(v_A, v_D)$$

Defender's cost function: Full information

Theorem: Defender's cost function with full information

The defender's discounted cost function is equal to

$$J^D(S_0, v_A, v_D(S_0)) = \frac{Q(v_d)}{1-\beta} + \frac{\beta}{1-\beta}\frac{\pi(v_a, v_d)(W(v_p) - Q(v_d))}{1+\beta(\pi(v_a, v_d) + \delta(v_h, v_p) - 1)}$$

and

$$J^D(S_1, v_A, v_D(S_1)) = \frac{W(v_p)}{1-\beta} + \frac{\beta}{1-\beta}\frac{\delta(v_h, v_p)(Q(v_d) - W(v_p))}{1+\beta(\pi(v_a, v_d) + \delta(v_h, v_p) - 1)},$$

where $v_D(S_0) = (0, v_p)$ and $v_D(S_1) = (v_d, 0)$,
$Q(v_d) = g_d(v_a) + C_d(v_d)$, and $W(v_d) = C_p(v_p)$.

Defender's cost function: Asymmetric information

Theorem: Defender's cost function with asymmetric information

$$\hat{J}^D(v_A, v_D) = \frac{g_d(v_a)}{1-\beta}\gamma(v_A, v_D) + \frac{C_d(v_d) + C_p(v_p)}{1-\beta}$$

where

$$\gamma(v_A, v_D) = \begin{cases} \frac{1}{1-\beta}\frac{\delta}{\pi+\delta} & \text{if } 0 < \pi + \delta < 2 \\ \frac{1}{2}\frac{1}{1-\beta} & \text{otherwise} \end{cases}$$

and $\delta = \delta(v_h, v_p)$ and $\pi = \pi(v_a, v_d)$.

Outline

Impact of C_p: With Full information there is a NE in which the attacker does not hack the system

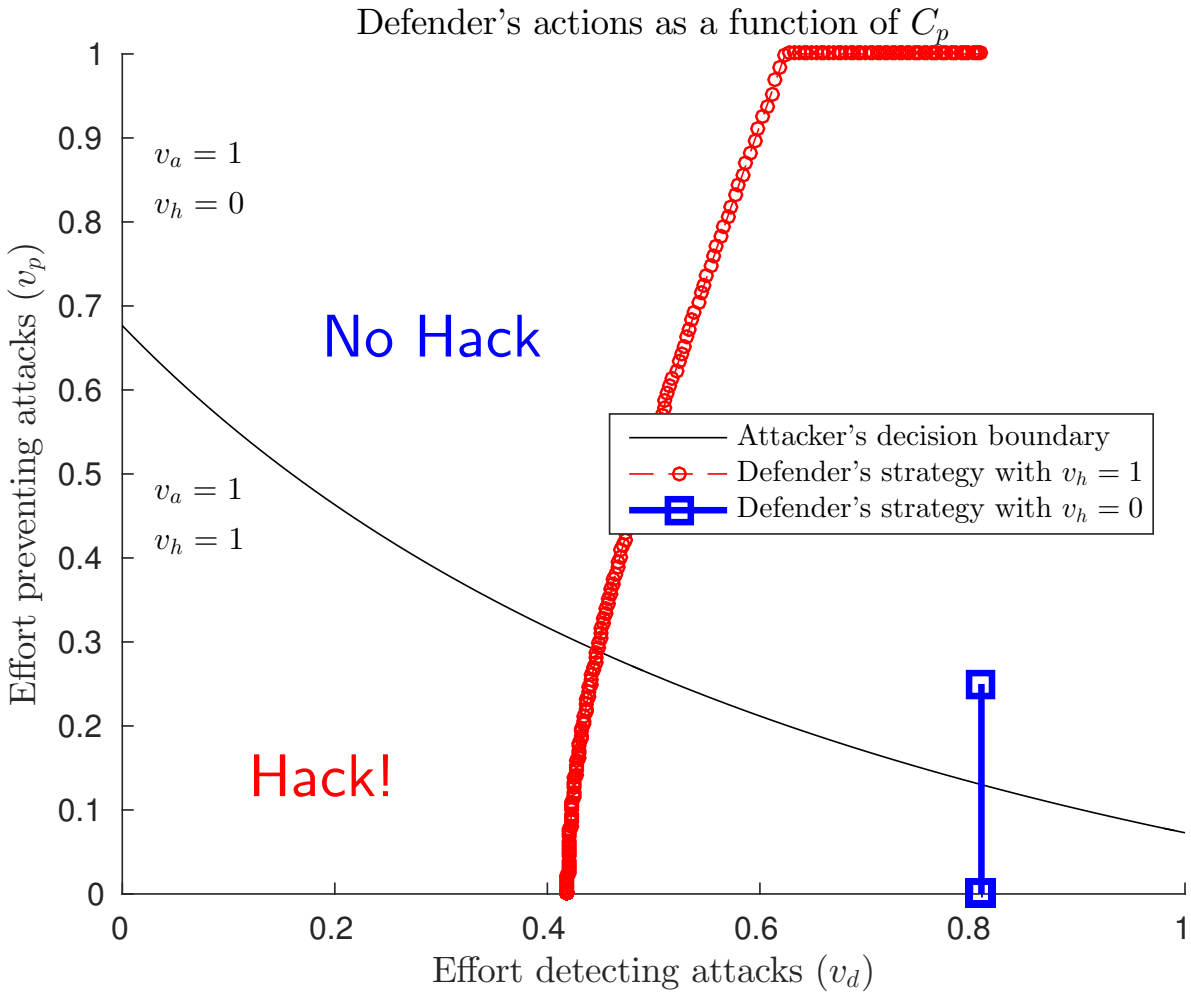

Defender's strategy with limited resources

$$\begin{aligned} \underset{v_D}{\text{Minimize}} \quad & \text{Defender's discounted cost} \\ \text{subject to} \quad & \\ & C_d(v_d) + C_p(v_p) \leq E, \\ & v_p, v_d \in [0, 1]. \end{aligned}$$

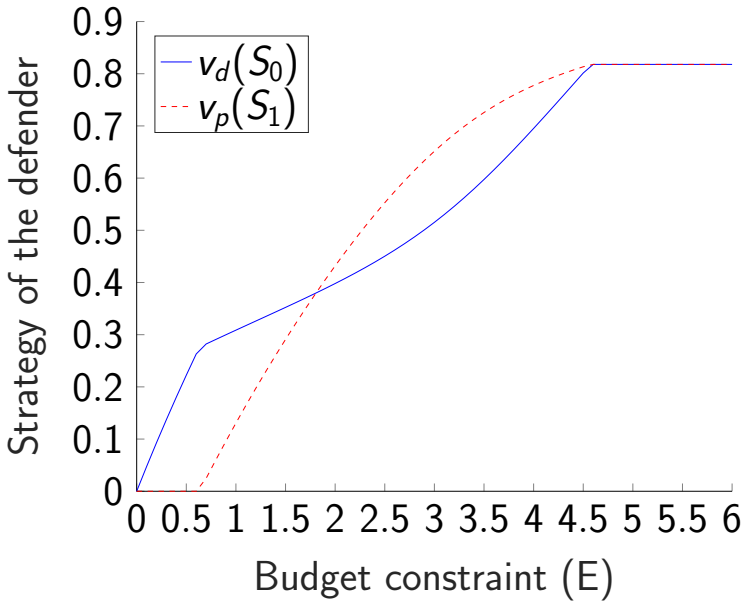

(a) Full information

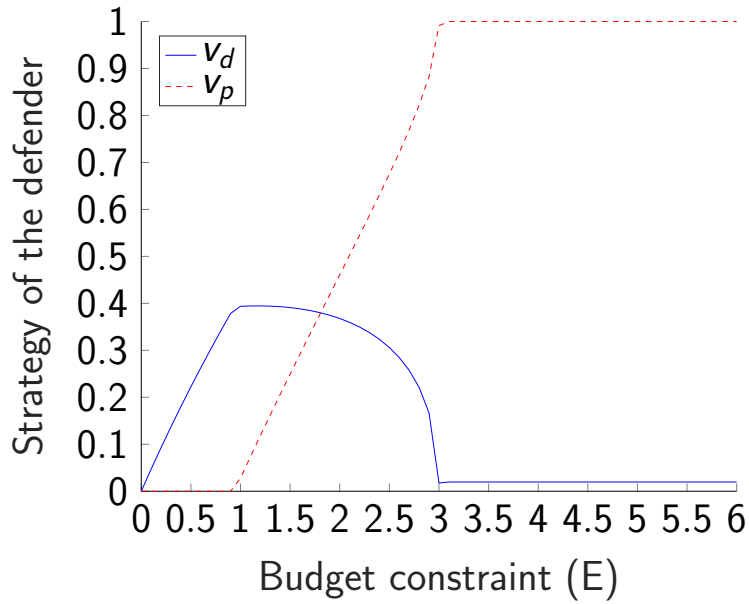

(b) Asymmetric information

Conclusions

- Detection alone can prevent attacks on systems that return low profit to the attacker.
- Prevention becomes more important for critical systems.
- With few resources the best strategy is to prioritize detection over prevention.
- With limited information the defender tends to invest only in detection (or maximum prevention when the cost of prevention is low or the losses are high).

Future work:

- We plan to adapt our models to allow investments in other risk mitigation strategies, such as cyber-insurance.

Thank You

Questions?

Contact:
Carlos Barreto, carlos.barretosuarez@utdallas.edu