Dynamic monitoring of adaptive criminals

Alae Baha

August 24, 2022

A thought experiment

A decision maker doubles the monitoring capacity

A thought experiment

A decision maker doubles the monitoring capacity

$\implies$ The policy intervention induced less detected fraud

A thought experiment

A decision maker doubles the monitoring capacity

$\implies$ The policy intervention induced less detected fraud

Is it a success in terms of fraud deterrence?

A thought experiment

A decision maker doubles the monitoring capacity

$\implies$ The policy intervention induced less detected fraud

Is it a success in terms of fraud deterrence?

Less detection implies either less fraud

A thought experiment

A decision maker doubles the monitoring capacity

$\Longrightarrow$ The policy intervention induced less detected fraud

Is it a success in terms of fraud deterrence?

Less detection implies either less fraud or fraudsters are frauding differently

A thought experiment

A decision maker doubles the monitoring capacity

$\implies$ The policy intervention induced less detected fraud

Is it a success in terms of fraud deterrence?

Less detection implies either less fraud or fraudsters are frauding differently (See Riley (2005))

A thought experiment

A decision maker doubles the monitoring capacity

$\implies$ The policy intervention induced less detected fraud

Is it a success in terms of fraud deterrence?

Less detection implies either less fraud or fraudsters are frauding differently (See Riley (2005))

Examples: Cyber security, border control, doping, tax evasion, money laundering, etc.

A thought experiment

A decision maker doubles the monitoring capacity

$\implies$ The policy intervention induced less detected fraud

Is it a success in terms of fraud deterrence?

Less detection implies either less fraud or fraudsters are frauding differently (See Riley (2005))

Examples: Cyber security, border control, doping, tax evasion, money laundering, etc.

To which extent can we reduce misbehavior in these environments?

Contribution:

This paper contributes to understanding the effect of monitoring policies on:

- Short term incentives to fraud:

Contribution:

This paper contributes to understanding the effect of monitoring policies on:

- Short term incentives to fraud: By studying the impact on fraud decisions
- Long term incentives to invest:

Contribution:

This paper contributes to understanding the effect of monitoring policies on:

- Short term incentives to fraud: By studying the impact on fraud decisions
- Long term incentives to invest: By studying their effect on technology adoption

Contribution:

This paper contributes to understanding the effect of monitoring policies on:

- Short term incentives to fraud: By studying the impact on fraud decisions
- Long term incentives to invest: By studying their effect on technology adoption

Contribution to the reputation literature: The state can be manipulated by both players

Outline

The model

Preview of the model

One attacker (player A)

Preview of the model

One attacker (player A)

One defender (player D)

Preview of the model

One attacker (player A)

One defender (player D)

Discrete time and infinite horizon

Preview of the model

One attacker (player A)

One defender (player D)

Discrete time and infinite horizon

The defender's has an endogenous ability to detect attacks θ_t

Preview of the model

One attacker (player A)

One defender (player D)

Discrete time and infinite horizon

The defender's has an endogenous ability to detect attacks θ_t

The attacker chooses:

Preview of the model

One attacker (player A)

One defender (player D)

Discrete time and infinite horizon

The defender's has an endogenous ability to detect attacks θ_t

The attacker chooses:

- An attack intensity a_t

Preview of the model

One attacker (player A)

One defender (player D)

Discrete time and infinite horizon

The defender's has an endogenous ability to detect attacks θ_t

The attacker chooses:

- An attack intensity a_t
- Investments in hiding technologies α_t

Preview of the model

One attacker (player A)

One defender (player D)

Discrete time and infinite horizon

The defender's has an endogenous ability to detect attacks θ_t

The attacker chooses:

- An attack intensity a_t
- Investments in hiding technologies α_t

The defender can invest in a detection technology δ_t

Preview of the model

One attacker (player A)

One defender (player D)

Discrete time and infinite horizon

The defender's has an endogenous ability to detect attacks θ_t

The attacker chooses:

- An attack intensity a_t
- Investments in hiding technologies α_t

The defender can invest in a detection technology δ_t

The monitoring ability θ_t

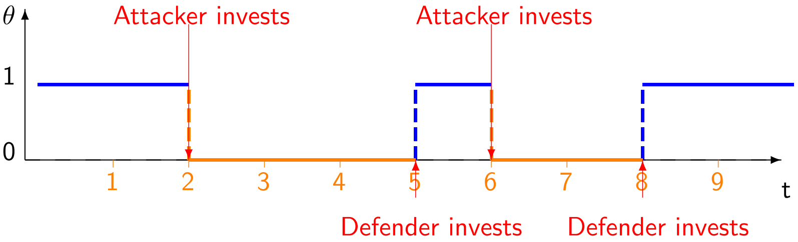

The monitoring ability as a function of time and investments

The monitoring ability θ_t

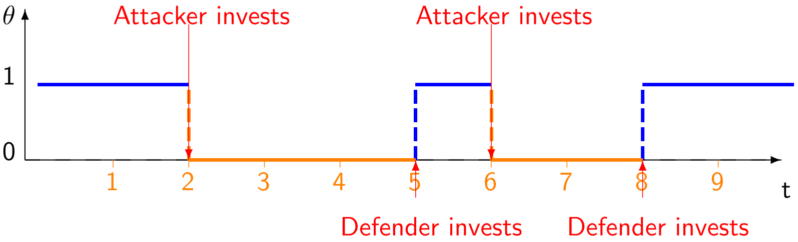

The monitoring ability as a function of time and investments

At each time $t \geq 0$:

- **Attacker** chooses investment $\boldsymbol{\alpha_t} \in \{0, 1\}$. Investment costs F^A
- **Defender** chooses investment $\boldsymbol{\delta_t} \in \{0, 1\}$, Investment costs F^D

The monitoring ability θ_t

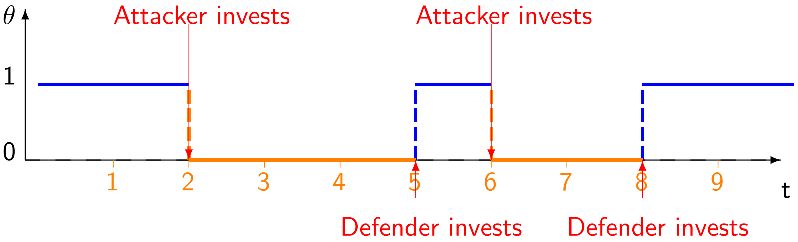

The monitoring ability as a function of time and investments

At each time $t \geq 0$:

- **Attacker** chooses investment $\boldsymbol{\alpha_t} \in \{0, 1\}$. Investment costs F^A
- **Defender** chooses investment $\boldsymbol{\delta_t} \in \{0, 1\}$, Investment costs F^D

The defender can only detect ongoing attacker ($\theta_t = 1$) if she invested last State

The monitoring ability θ_t

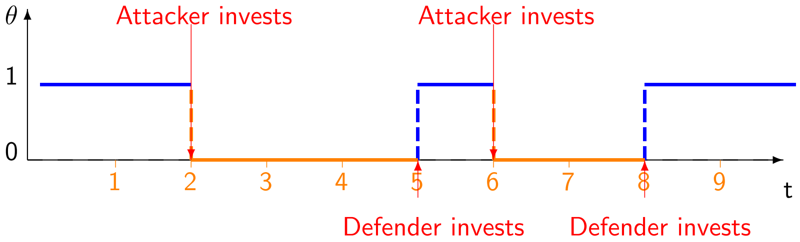

The monitoring ability as a function of time and investments

At each time $t \geq 0$:

- **Attacker** chooses investment $\boldsymbol{\alpha_t} \in \{0, 1\}$. Investment costs F^A
- **Defender** chooses investment $\boldsymbol{\delta_t} \in \{0, 1\}$, Investment costs F^D

The defender can only detect ongoing attacker $(\theta_t = 1)$ if she invested last State

The intensity of attacks

At each time t, the attacker chooses an attack intensity $a_t \in \mathbb{R}^+$

The intensity of attacks

At each time t, the attacker chooses an attack intensity $a_t \in \mathbb{R}^+$

A policy π is $(\lambda_\pi(a), u_\pi^D(a,\theta), u_\pi^A(a,\theta))$ such that:

The intensity of attacks

At each time t, the attacker chooses an attack intensity $a_t \in \mathbb{R}^+$

A policy π is $(\lambda_\pi(a), u_\pi^D(a, \theta), u_\pi^A(a, \theta))$ such that:

- Detection arrives at a rate $\theta_t \lambda_\pi(a_t)$

The intensity of attacks

At each time t, the attacker chooses an attack intensity $a_t \in \mathbb{R}^+$

A policy π is $(\lambda_\pi(a), u_\pi^D(a, \theta), u_\pi^A(a, \theta))$ such that:

- Detection arrives at a rate $\theta_t \lambda_\pi(a_t)$
- The defender earns expected flow payoffs: $u_\pi^D(a_t, \theta_t)$

The intensity of attacks

At each time t, the attacker chooses an attack intensity $a_t \in \mathbb{R}^+$

A policy π is $(\lambda_\pi(a), u_\pi^D(a, \theta), u_\pi^A(a, \theta))$ such that:

- Detection arrives at a rate $\theta_t \lambda_\pi(a_t)$
- The defender earns expected flow payoffs: $u_\pi^D(a_t, \theta_t)$
- The attacker earns expected flow payoffs: $u_\pi^A(a_t, \theta_t)$

The intensity of attacks

At each time t, the attacker chooses an attack intensity $a_t \in \mathbb{R}^+$

A policy π is $(\lambda_\pi(a), u_\pi^D(a,\theta), u_\pi^A(a,\theta))$ such that:

- Detection arrives at a rate $\theta_t \lambda_\pi(a_t)$
- The defender earns expected flow payoffs: $u_\pi^D(a_t, \theta_t)$
- The attacker earns expected flow payoffs: $u_\pi^A(a_t, \theta_t)$

Example of a class of policies: Choice of monitoring rates m and punishment P: $\pi = (ma, -a, u(a) - \theta maP)$

The intensity of attacks

At each time t, the attacker chooses an attack intensity $a_t \in \mathbb{R}^+$

A policy π is $(\lambda_\pi(a), u_\pi^D(a,\theta), u_\pi^A(a,\theta))$ such that:

- Detection arrives at a rate $\theta_t \lambda_\pi(a_t)$
- The defender earns expected flow payoffs: $u_\pi^D(a_t, \theta_t)$
- The attacker earns expected flow payoffs: $u_\pi^A(a_t, \theta_t)$

Example of a class of policies: Choice of monitoring rates m and punishment P: $\pi = (ma, -a, u(a) - \theta maP)$

The timing for each $t \geq 0$

- Stage 0 (Belief updating): The state is inherited from the past, and the defender updates her belief about it,

The timing for each $t \geq 0$

- Stage 0 (Belief updating): The state is inherited from the past, and the defender updates her belief about it,
- Stage 1 (Investments): Both players simultaneously make investment decisions,

The timing for each $t \geq 0$

- Stage 0 (Belief updating): The state is inherited from the past, and the defender updates her belief about it,
- Stage 1 (Investments): Both players simultaneously make investment decisions,
- Stage 2 (Technology outcome): θ_t is determined and observed by the attacker,

The timing for each $t \geq 0$

- Stage 0 (Belief updating): The state is inherited from the past, and the defender updates her belief about it,
- Stage 1 (Investments): Both players simultaneously make investment decisions,
- Stage 2 (Technology outcome): θ_t is determined and observed by the attacker,
- Stage 3 (Attack): The attacker chooses an intensity of attack a_t,

The timing for each $t \geq 0$

- Stage 0 (Belief updating): The state is inherited from the past, and the defender updates her belief about it,
- Stage 1 (Investments): Both players simultaneously make investment decisions,
- Stage 2 (Technology outcome): θ_t is determined and observed by the attacker,
- Stage 3 (Attack): The attacker chooses an intensity of attack a_t,
- Stage 4 (Outcome): The outcome of detection is publicly observed, and stage payoffs are realized.

The timing for each $t \geq 0$

- Stage 0 (Belief updating): The state is inherited from the past, and the defender updates her belief about it,
- Stage 1 (Investments): Both players simultaneously make investment decisions,
- Stage 2 (Technology outcome): θ_t is determined and observed by the attacker,
- Stage 3 (Attack): The attacker chooses an intensity of attack a_t,
- Stage 4 (Outcome): The outcome of detection is publicly observed, and stage payoffs are realized.

Focus on **Markov perfect equilibria** that depend on the defender's beliefs ρ and the attacker's private information about the state θ_t Equilibrium

Results

Types of equilibria

Any equilibrium is:

- An entente equilibrium if the cost of developing hiding technologies is high relative to short-term gains from being undetectable
- Otherwise, the equilibrium is either an arms race or a complete hiding equilibrium

The equilibrium

Lemma 1: I equilibrium intensity of attacks is myopic
$a^*(\theta) = argmax_a u_\pi^A(a, \theta)$

The equilibrium

If $u_\pi^A(a^*(0),0) - u_\pi^A(a^*(1),1) > (1-e^{-r\Delta})F^A$ and $F^D < F^{D*}$, an arms race equilibrium exists.

The equilibrium

If $u_\pi^A(a^*(0),0) - u_\pi^A(a^*(1),1) > (1 - e^{-r\Delta})F^A$ and $F^D < F^{D^*}$, an arms race equilibrium exists.

Any such an equilibrium is characterized by an initial belief $\rho_0 \in (0,1)$ and a stopping belief ρ^* such that:

The equilibrium

If $u_\pi^A(a^*(0),0) - u_\pi^A(a^*(1),1) > (1-e^{-r\Delta})F^A$ and $F^D < F^{D^*}$, an arms race equilibrium exists.
Any such an equilibrium is characterized by an initial belief $\rho_0 \in (0,1)$ and a stopping belief ρ^* such that:
(i) The investment by the attacker $\alpha_0 \in (0,1)$ is :

The equilibrium

If $u^A_\pi(a^*(0),0) - u^A_\pi(a^*(1),1) > (1-e^{-r\Delta})F^A$ and $F^D < F^{D*}$, an arms race equilibrium exists.

Any such an equilibrium is characterized by an initial belief $\rho_0 \in (0,1)$ and a stopping belief ρ^* such that:

(i) The investment by the attacker $\alpha_0 \in (0,1)$ is :

- $$\alpha(\rho) = \begin{cases} 0 \ \forall \rho \in (\rho^*, \rho_0) \\ 1 - \rho_0 \text{ if } \rho \leq \rho^* \\ 1 - \frac{\rho_0}{\rho} \text{ if } \rho \geq \rho_0 \end{cases}$$

The equilibrium

If $u_\pi^A(a^*(0),0) - u_\pi^A(a^*(1),1) > (1-e^{-r\Delta})F^A$ and $F^D < F^{D*}$, an arms race equilibrium exists.
Any such an equilibrium is characterized by an initial belief $\rho_0 \in (0,1)$ and a stopping belief ρ^* such that:
(i) The investment by the attacker $\alpha_0 \in (0,1)$ is :

- $\alpha(\rho) = \begin{cases} 0 \ \forall \rho \in (\rho^*, \rho_0) \\ 1-\rho_0 \text{ if } \rho \leq \rho^* \\ 1-\frac{\rho_0}{\rho} \text{ if } \rho \geq \rho_0 \end{cases}$

(ii) The investment strategy by the defender:

The equilibrium

If $u_\pi^A(a^*(0),0) - u_\pi^A(a^*(1),1) > (1-e^{-r\Delta})F^A$ and $F^D < F^{D^*}$, an arms race equilibrium exists.

Any such an equilibrium is characterized by an initial belief $\rho_0 \in (0,1)$ and a stopping belief ρ^* such that:

(i) The investment by the attacker $\alpha_0 \in (0,1)$ is :

- $$\alpha(\rho) = \begin{cases} 0 \ \forall \rho \in (\rho^*, \rho_0) \\ 1 - \rho_0 \text{ if } \rho \leq \rho^* \\ 1 - \frac{\rho_0}{\rho} \text{ if } \rho \geq \rho_0 \end{cases}$$

(ii) The investment strategy by the defender:

- $$\delta(\rho) = \begin{cases} 1 \text{ if } \rho \leq \rho^* \\ 0 \text{ otherwise} \end{cases}$$

The equilibrium

If $u_\pi^A(a^*(0),0) - u_\pi^A(a^*(1),1) > (1-e^{-r\Delta})F^A$ and $F^D < F^{D^*}$, an arms race equilibrium exists.

Any such an equilibrium is characterized by an initial belief $\rho_0 \in (0,1)$ and a stopping belief ρ^* such that:

(i) The investment by the attacker $\alpha_0 \in (0,1)$ is :

- $$\alpha(\rho) = \begin{cases} 0 \ \forall \rho \in (\rho^*, \rho_0) \\ 1 - \rho_0 \text{ if } \rho \leq \rho^* \\ 1 - \frac{\rho_0}{\rho} \text{ if } \rho \geq \rho_0 \end{cases}$$

(ii) The investment strategy by the defender:

- $$\delta(\rho) = \begin{cases} 1 \text{ if } \rho \leq \rho^* \\ 0 \text{ otherwise} \end{cases}$$

The equilibrium

If $u_\pi^A(a^*(0),0) - u_\pi^A(a^*(1),1) > (1-e^{-r\Delta})F^A$ and $F^D < F^{D*}$, an arms race equilibrium exists.

The equilibrium

If $u_\pi^A(a^*(0), 0) - u_\pi^A(a^*(1), 1) > (1 - e^{-r\Delta})F^A$ and $F^D < F^{D*}$, an arms race equilibrium exists.
Any such an equilibrium is characterized by an initial belief $\rho_0 \in (0, 1)$ and a stopping belief ρ^* such that:

The equilibrium

If $u_\pi^A(a^*(0),0) - u_\pi^A(a^*(1),1) > (1-e^{-r\Delta})F^A$ and $F^D < F^{D*}$, an arms race equilibrium exists.

Any such an equilibrium is characterized by an initial belief $\rho_0 \in (0,1)$ and a stopping belief ρ^* such that:

(iii) An equilibrium length of the cycle:

$t^A = \frac{1}{r} ln(1 + \frac{rF^A}{u_\pi^A(a^*(0),0) - u_\pi^A(a^*(1),1) - rF^A})$

The equilibrium

If $u_\pi^A(a^*(0),0) - u_\pi^A(a^*(1),1) > (1-e^{-r\Delta})F^A$ and $F^D < F^{D^*}$, an arms race equilibrium exists.

Any such an equilibrium is characterized by an initial belief $\rho_0 \in (0,1)$ and a stopping belief ρ^* such that:

(iii) An equilibrium length of the cycle:

$t^A = \frac{1}{r} ln(1 + \frac{rF^A}{u_\pi^A(a^*(0),0) - u_\pi^A(a^*(1),1) - rF^A})$

(iv) The stopping belief $\rho^*(\rho_0)$ is reached at time t^D such that:

The equilibrium

If $u^A_\pi(a^*(0),0) - u^A_\pi(a^*(1),1) > (1-e^{-r\Delta})F^A$ and $F^D < F^{D^*}$, an arms race equilibrium exists.

Any such an equilibrium is characterized by an initial belief $\rho_0 \in (0,1)$ and a stopping belief ρ^* such that:

(iii) An equilibrium length of the cycle:

$t^A = \frac{1}{r} ln(1 + \frac{rF^A}{u^A_\pi(a^*(0),0) - u^A_\pi(a^*(1),1) - rF^A})$

(iv) The stopping belief $\rho^*(\rho_0)$ is reached at time t^D such that:

$$r + \frac{X}{\rho_0} = \frac{\lambda_\pi(a^*(1))(1-e^{-rt^D}) + \frac{X}{\rho_0(1-\rho_0)}}{e^{\lambda_\pi(a^*(1))t^D} - 1}$$

The equilibrium

If $u_\pi^A(a^*(0),0) - u_\pi^A(a^*(1),1) > (1-e^{-r\Delta})F^A$ and $F^D < F^{D*}$, an arms race equilibrium exists.
Any such an equilibrium is characterized by an initial belief $\rho_0 \in (0,1)$ and a stopping belief ρ^* such that:

(iii) An equilibrium length of the cycle:
$t^A = \frac{1}{r} ln(1 + \frac{rF^A}{u_\pi^A(a^*(0),0) - u_\pi^A(a^*(1),1) - rF^A})$

(iv) The stopping belief $\rho^*(\rho_0)$ is reached at time t^D such that:

$$r + \frac{X}{\rho_0} = \frac{\lambda_\pi(a^*(1))(1-e^{-rt^D}) + \frac{X}{\rho_0(1-\rho_0)}}{e^{\lambda_\pi(a^*(1))t^D} - 1}$$

Dynamics of beliefs under an arms race policy

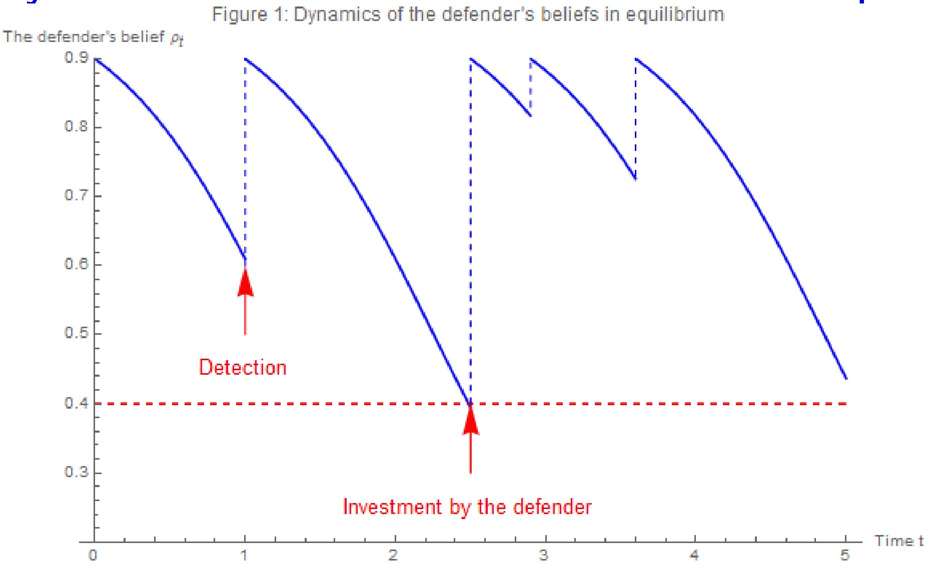

Dynamics of beliefs

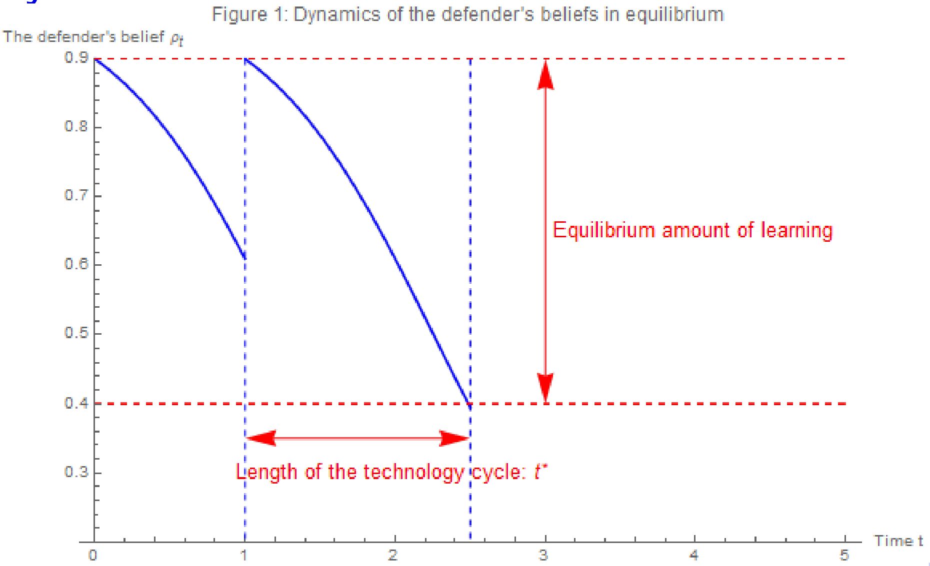

Effect of policy intervention

Effect of raising penalties

Consider the illustrative example:

- Detection arrives at a rate $\lambda_\pi = am$

Effect of raising penalties

Consider the illustrative example:

- Detection arrives at a rate $\lambda_\pi = am$
- The attacker's expected flow payoffs: $U_\pi^A(a, \theta) = u(a) - \theta a m P$

Effect of raising penalties

Consider the illustrative example:

- Detection arrives at a rate $\lambda_\pi = am$
- The attacker's expected flow payoffs: $U^A_\pi(a, \theta) = u(a) - \theta am P$
- The defender earns expected flow payoffs $U^D_\pi(a, \theta) = -a$

Effect of raising penalties

Higher penalties lead to:

- A deterrence effect: Less intense detectable attacks
- An increase in per-period gains from being undetectable
  $\implies$ Shorter technology cycles
- More investments by the attacker in equilibrium

Trade-off: More deterrence of detectable attacks versus less investments

Effect of more informative policies

Consider the illustrative example:

- Detection arrives at a rate am

Effect of more informative policies

Consider the illustrative example:

- Detection arrives at a rate am
- The attacker's expected flow payoffs: $(u(a) - \theta amP)$

Effect of more informative policies

Consider the illustrative example:

- Detection arrives at a rate am
- The attacker's expected flow payoffs: $(u(a) - \theta amP)$
- The defender earns expected flow payoffs: $-a$

Effect of more informative policies

Consider two policies π and π' such that $mP = m'P'$ with $m > m'$. These policies:

Effect of more informative policies

Consider two policies π and π' such that $mP = m'P'$ with $m > m'$. These policies:

- Lead to the same per-period gains from being undetectable

Effect of more informative policies

Consider two policies π and π' such that $mP = m'P'$ with $m > m'$. These policies:

- Lead to the same per-period gains from being undetectable
- Same short-term payoff functions

Effect of more informative policies

Consider two policies π and π' such that $mP = m'P'$ with $m{>}m'$. These policies:

- Lead to the same per-period gains from being undetectable
- Same short-term payoff functions
- The policy π leads to:

Effect of more informative policies

Consider two policies π and π' such that $mP = m'P'$ with $m > m'$. These policies:

- Lead to the same per-period gains from being undetectable
- Same short-term payoff functions
- The policy π leads to:
  1. A more aggressive investment strategy by the defender

Effect of more informative policies

Consider two policies π and π' such that $mP = m'P'$ with $m > m'$. These policies:

- Lead to the same per-period gains from being undetectable
- Same short-term payoff functions
- The policy π leads to:
  1. A more aggressive investment strategy by the defender
  2. Less investments in hiding technologies in equilibrium

Effect of more informative policies

Consider two policies π and π' such that $mP = m'P'$ with $m > m'$. These policies:

- Lead to the same per-period gains from being undetectable
- Same short-term payoff functions
- The policy π leads to:
  1. A more aggressive investment strategy by the defender
  2. Less investments in hiding technologies in equilibrium
  3. Less intense attacks on average

Effect of more informative policies

Consider two policies π and π' such that $mP = m'P'$ with $m > m'$. These policies:

- Lead to the same per-period gains from being undetectable
- Same short-term payoff functions
- The policy π leads to:
  1. A more aggressive investment strategy by the defender
  2. Less investments in hiding technologies in equilibrium
  3. Less intense attacks on average

Related literature

- Reputation/monitoring literature: Board & Meyer-Ter-Vehn (2013), Board and Meyer-Ter-Vehn (2020), Dilmé (2019), Dilmé & Garrett (2019) , Marinovic & Szidlowski (2019), Halac & Prat (2016), Varas, Marinovic, and Skrzypacz (2020)
- Experimentation with Poisson bandits: Bergemann, & Valimaki (2006), Keller, Rady & Cripps (2005)
- Optimal enforcement: Becker (1968), Polinsky, & Shavell (2000)
- Steganography: Cabaj, Caviglione, Mazurczyk, Wendzel, Woodward, and Zander (2018)
- Crime displacement: Gonzalez-Navarro (2013), Johnson, Guerette, and Bowers (2014), Ladegaard (2019), Yang (2008)

Conclusion

Conclusion

- I study monitoring game with endogenous ability to detect misbehavior

Conclusion

- I study monitoring game with endogenous ability to detect misbehavior
- I show that high deterrence increases incentives to invest in hiding technologies and leads to an arms race

Conclusion

- I study monitoring game with endogenous ability to detect misbehavior

- I show that high deterrence increases incentives to invest in hiding technologies and leads to an arms race

- Empirical predictions: 1 2
    1. A technological response to harsher policies(Bustos et. al. (2022))

Conclusion

- I study monitoring game with endogenous ability to detect misbehavior
- I show that high deterrence increases incentives to invest in hiding technologies and leads to an arms race
- Empirical predictions: 1 2
  1. A technological response to harsher policies(Bustos et. al. (2022))
  2. Investments increase as a function of penalties

Conclusion

- I study monitoring game with endogenous ability to detect misbehavior
- I show that high deterrence increases incentives to invest in hiding technologies and leads to an arms race
- Empirical predictions: 1 2
    1. A technological response to harsher policies(Bustos et. al. (2022))
    2. Investments increase as a function of penalties
    3. Fraud can increase after harsher policies

Conclusion

- I study monitoring game with endogenous ability to detect misbehavior
- I show that high deterrence increases incentives to invest in hiding technologies and leads to an arms race
- Empirical predictions: 1 2
  1. A technological response to harsher policies(Bustos et. al. (2022))
  2. Investments increase as a function of penalties
  3. Fraud can increase after harsher policies
  4. Investments are made by bigger attackers

Thank you

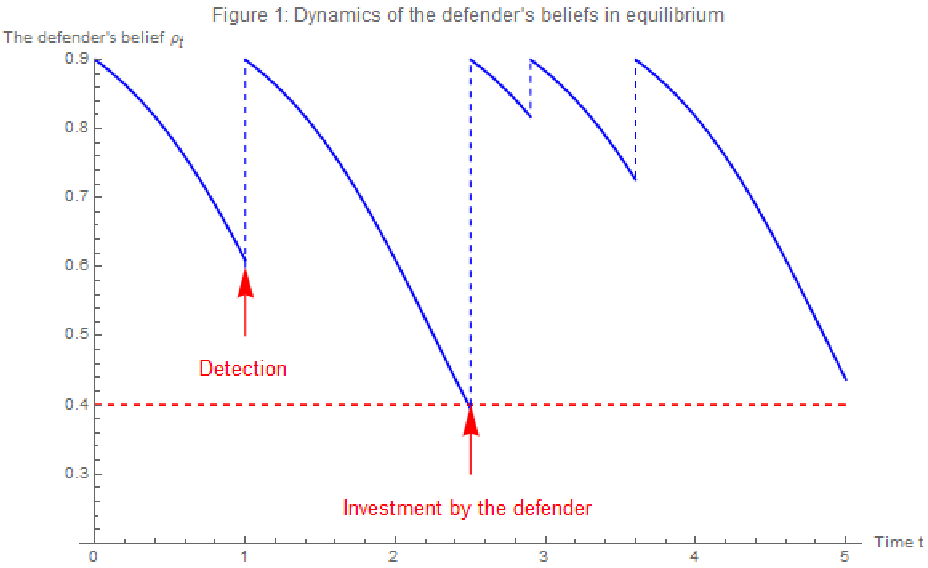
Figure 1: Dynamics of the defender's beliefs in equilibrium
The defender's belief ρ_t
0.9
0.8
0.7
0.6
0.5
0.4
0.3
Detection
Investment by the defender
0
1
2
3
4
5
Time t

The monitor's ability

Define t_i the date of last investment by player i

Then $\theta_t = \mathbf{1}_{t_D > t_A}$

Equilibrium notion

A deterministic Markov policy for the defender is:

$$\sigma^D : [0,1] \times \{0,1\}$$
$$\rho \to \delta^D$$

Equilibrium notion

A deterministic Markov policy for the defender is:

$$\sigma^D : [0,1] \times \{0,1\} \to \{0,1\}$$
$$\rho \to \delta^D$$

A deterministic Markov policy for the attacker is:

$$\sigma^A : [0,1] \times \{0,1\} \to \{0,1\} \times [0,\bar{a}] \times \{0,1\}$$
$$\rho \times \theta \to \alpha \times a$$

Equilibrium notion

A deterministic Markov policy for the defender is:

$$\sigma^D : [0,1] \times \{0,1\}$$
$$\rho \to \delta^D$$

A deterministic Markov policy for the attacker is:

$$\sigma^A : [0,1] \times \{0,1\} \to \{0,1\} \times [0,\bar{a}] \times \{0,1\}$$
$$\rho \times \theta \to \alpha \times a$$

Example of a failing policy:

Border control (Riley, 2005): Higher monitoring intensities in the US-Mexican border lead to:

Example of a failing policy:

Border control (Riley, 2005): Higher monitoring intensities in the US-Mexican border lead to:

- Small impact on drug smuggling

Example of a failing policy:

Border control (Riley, 2005): Higher monitoring intensities in the US-Mexican border lead to:

- Small impact on drug smuggling
- Displacement to unguarded parts of the border

Example of a failing policy:

Border control (Riley, 2005): Higher monitoring intensities in the US-Mexican border lead to:

- Small impact on drug smuggling
- Displacement to unguarded parts of the border
- Adoption of better hiding technologies: Submarines, lightplanes, mules etc.

Example of a failing policy:

Border control (Riley, 2005): Higher monitoring intensities in the US-Mexican border lead to:

- Small impact on drug smuggling
- Displacement to unguarded parts of the border
- Adoption of better hiding technologies: Submarines, lightplanes, mules etc.

Monitoring policies impact technology adoption

Example of a failing policy:

Border control (Riley, 2005): Higher monitoring intensities in the US-Mexican border lead to:

- Small impact on drug smuggling
- Displacement to unguarded parts of the border
- Adoption of better hiding technologies: Submarines, lightplanes, mules etc.

Monitoring policies impact technology adoption

Evaluating policies based on detected fraud can be misleading in the short run

The arms race equilibrium:

If $u_\pi^A(a^*(0),0) - u_\pi^A(a^*(1),1) > (1-e^{-r\Delta})F^A$ and $F^D < F^{D^*}$, an arms race equilibrium exists.

Any such an equilibrium is characterized by an initial belief $\rho_0 \in (0,1)$ and a stopping belief ρ^* such that:

(i) The investment by the attacker $\alpha_0 \in (0,1)$ is :

- $\alpha(\rho) = \begin{cases} 0 \ \forall \rho \in (\rho^*, \rho_0) \\ 1-\rho_0 \text{ if } \rho \leq \rho^* \\ 1-\frac{\rho_0}{\rho} \text{ if } \rho \geq \rho_0 \end{cases}$

(ii) The investment strategy by the defender:

- $\delta(\rho) = \begin{cases} 1 \text{ if } \rho \leq \rho^* \\ 0 \text{ otherwise} \end{cases}$

The arms race equilibrium:

(iii) An equilibrium length of the cycle:
$t^A = \frac{1}{r} ln(1 + \frac{rF^A}{u_\pi^A(a^*(0),0) - u_\pi^A(a^*(1),1) - rF^A})$
(iv) The stopping belief $\rho^*(\rho_0)$ is reached at time t^D such that:

$$r + \frac{X}{\rho_0} = \frac{\lambda_\pi(a^*(1))(1 - e^{-rt^D}) + \frac{X}{\rho_0(1-\rho_0)}}{e^{\lambda_\pi(a^*(1))t^D} - 1}$$

(v) The initial belief ρ_0 is such that $t^* = t^A = t^D$

Empirical application

Steps of a cyber attack:

Empirical application

Steps of a cyber attack:

- Phase 1: The intrusion phase

Empirical application

Steps of a cyber attack:

- Phase 1: The intrusion phase (Affected by S)

Empirical application

Steps of a cyber attack:

- Phase 1: The intrusion phase (Affected by S)
- Phase 2: Exploitation phase

Empirical application

Steps of a cyber attack:

- Phase 1: The intrusion phase (Affected by S)
- Phase 2: Exploitation phase (Affected by m)

Empirical application

Steps of a cyber attack:

- Phase 1: The intrusion phase (Affected by S)
- Phase 2: Exploitation phase (Affected by m)

Policy: Increase in investments in cybersecurity under the Biden administration

Empirical application

Security programs detect patterns of code

Empirical application

Security programs detect patterns of code

New malwares are often a modification of old ones (A mutation)

Empirical application

Security programs detect patterns of code

New malwares are often a modification of old ones (A mutation)

Avtest institute registers and classifies (450 000 daily) new malwares

Empirical application

Security programs detect patterns of code

New malwares are often a modification of old ones (A mutation)

Avtest institute registers and classifies (450 000 daily) new malwares

The new policy should lead to:

- No change for some types of malwares

Empirical application

Security programs detect patterns of code

New malwares are often a modification of old ones (A mutation)

Avtest institute registers and classifies (450 000 daily) new malwares

The new policy should lead to:

- No change for some types of malwares
- An increase in the frequency at which other ones are created