Mitigating Covert Compromises: A Game-Theoretic Model of Targeted and Non-Targeted Covert Attacks

Aron Laszka[1,2] Benjamin Johnson[3] Jens Grossklags[1]

[1]Pennsylvania State University

[2]Budapest University of Technology and Economics

[3]University of California, Berkeley

WINE 2013

Motivation

- Continuous covert attacks against resources
  - attackers often want to keep successful security compromises covert
  - examples
    - cyber-espionage: targets should not be aware that they are being spied on
    - botnets: targets should not be aware that their computers are infected

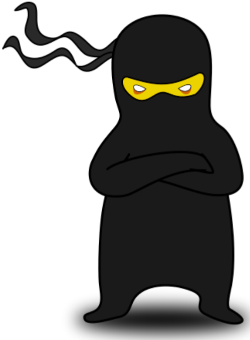

Motivation

- Continuous covert attacks against resources
  - mitigation of covert attacks
    - minimizing possible losses by resetting the resource to a secure state
    - e.g., resetting passwords, changing private keys, reinstalling servers
  - since the attacks are covert, the question arises: when to reset the resource?
    - what is the economically optimal frequency?
    - what is the optimal scheduling?

> traditionally, security is more concerned with what to do and how to do it

in practice: usually periodic password and key renewal policies

Motivation (contd.)

- Continuous covert attacks against resources
- Targeted and non-targeted attacks
  - extent to which the attack is customized for a particular target

	Targeted	Non-Targeted
Example	cyber-espionage	botnets
Number of targets	low	high
Number of attackers	low	high
Effort required for each attack	high	low
Success probability of each attack	high	low

Related Work

- Timing games:
  - since the cold-war era, games of timing have been studied with the tools of non-cooperative game theory
- `FlipIt` [1]:
  - in response to recent-high profile stealthy attacks, researchers at RSA proposed the `FlipIt` model
  - mitigation of targeted attacks
  - lesson: defender should play upredictably

[1] K. D. Bowers, M. van Dijk, R. Griffin, A. Juels, A. Oprea, R. L. Rivest, and N. Triandopoulos. Defending against the unknown enemy: Applying FlipIt to system security. In GameSec, pages 248–263, 2012

Model

- Strategic players:
  - defender (denoted by D)
  - targeting attacker (denoted by A)

+ non-strategic actors: non-targeting attackers (denoted by N)

Model

- Strategic players
- Resource:
  - some computing resource, e.g., user account, machine
  - having it compromised generates B_i benefit per unit of time for attacker i

Model

- Strategic players
- Resource:
  - some computing resource, e.g., user account, machine
  - having it compromised generates B_i benefit per unit of time for attacker i
- Time:
  - continuous
  - game starts at time $t = 0$ with the resource being uncompromised
  - and played indefinitely as $t \to \infty$

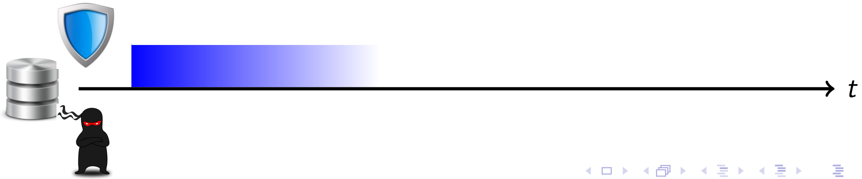

Model

- Strategic players
- Resource
- Time
- Moves:
  - at any time instance, player i may make a move, which costs her C_i
  - when the defender makes a move, the resource becomes uncompromised immediately, but the attackers will know of it
  - when the targeting attacker makes a move, she starts her attack, which takes some random amount of time
    - ⋆ distribution of the attack time is given by the cumulative function F_A,

  but the attackers' moves are stealthy (i.e., the defender does not know when the resource became compromised or if it is compromised at all)

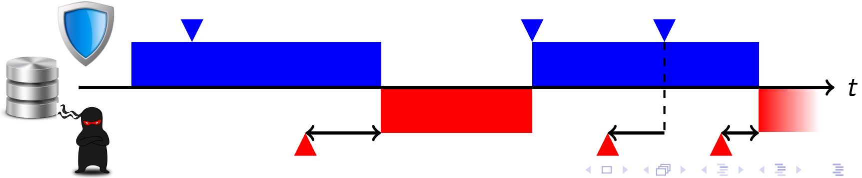

Model

- Strategic players
- Resource
- Time
- Moves
- Strategies:
  - set of rules, algorithm, etc. for making moves
  - in practice: defender's key or password update policy, targeting attacker's plan of attack, etc.

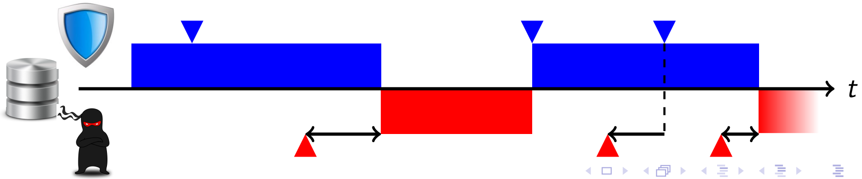

Model

- Strategic players
- Resource
- Time
- Moves
- Strategies
- Payoffs:
  - targeting attacker: $b_A - c_A$
  - defender: $-(b_A + b_N) - c_D$
  - benefit (loss) rate b_i: average fraction of time i has the resource compromised $\times$ unit benefit B_i
  - cost rate c_i: average number of moves per unit of time $\times$ move cost C_i

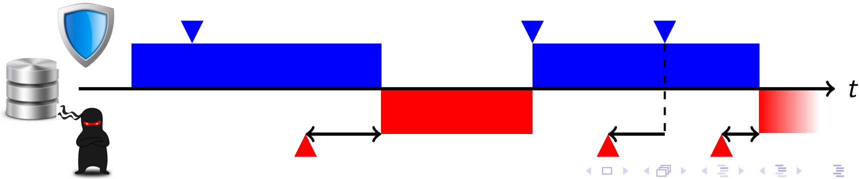

Strategies

- Adaptive strategies (for attackers):
  - an attacker uses an adaptive strategy if, after each move of the defender, she computes the time of her next move based on the defender's all previous moves using some non-deterministic function
  - this class is a simple representation of all the rational strategies available to an attacker

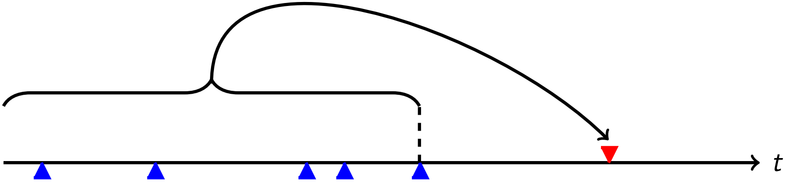

Strategies

- Adaptive strategies (for attackers)
- Renewal strategies:
  - player i uses a renewal strategy if the time intervals between her consecutive moves are identically distributed independent random variables
  - renewal strategies are well-motivated for the defender by the fact that the defender is playing blindly; thus, she has the same information available after each move

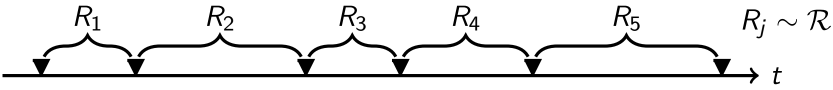

Strategies

- Adaptive strategies (for attackers)
- Renewal strategies
- Periodic strategies:
  - player i uses a periodic strategy if the time intervals between her consecutive moves are identical (this period is denoted by δ_i)

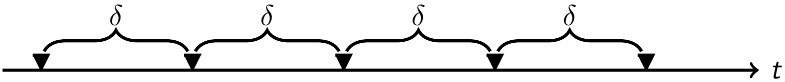

Strategies

- Adaptive strategies (for attackers)
- Renewal strategies
- Periodic strategies
- Not moving:
  - a player can choose to never move
  - while this might seem counter-intuitive, it is actually a best-response if the expected benefit from making a move is always less than the cost of moving

Non-Targeted Attacks

- in practice, the number of non-targeting attackers is very large

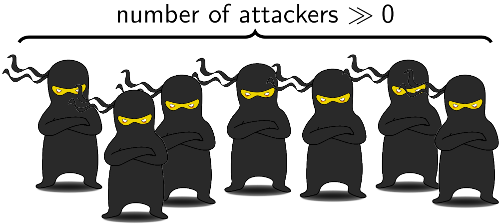

Non-Targeted Attacks

- in practice, the number of non-targeting attackers is very large, but the expected number of attacks in any time interval is finite

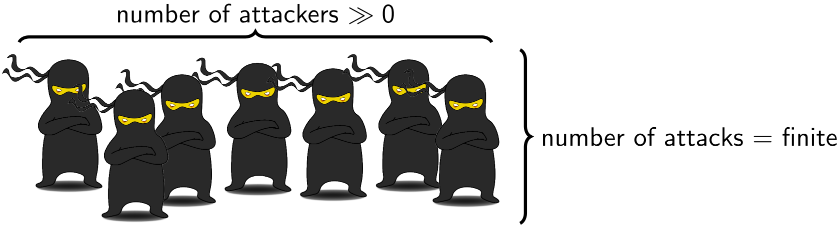

Non-Targeted Attacks

- in practice, the number of non-targeting attackers is very large, but the expected number of attacks in any time interval is finite
  $\longrightarrow$ the probability that a given non-targeting attacker targets the defender approaches zero

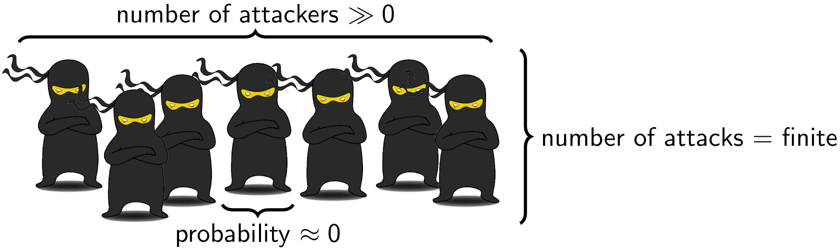

Non-Targeted Attacks

- in practice, the number of non-targeting attackers is very large, but the expected number of attacks in any time interval is finite
  $\longrightarrow$ the probability that a given non-targeting attacker targets the defender approaches zero
- since non-targeting attackers operate independently, the number of successful attacks in any time interval depends solely on the length of the interval
  $\longrightarrow$ arrival of <u>non-targeted attacks</u> follows a <u>Poisson process</u>

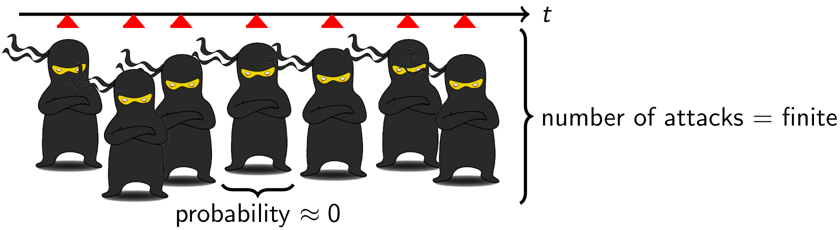

Non-Targeted Attacks (contd.)

- the arrival of non-targeted attacks follows a Poisson process
- furthermore, since the economic decisions of the non-targeting attackers depend on a very large pool of possible targets, the effect of the defender's strategy choice on the non-targeting attackers' strategies is negligible
  $\longrightarrow$ non-targeting attackers' strategies can be considered exogenously given
- that is, the expected number of arrivals that occur per unit of time, denoted by λ_N, is exogenously given

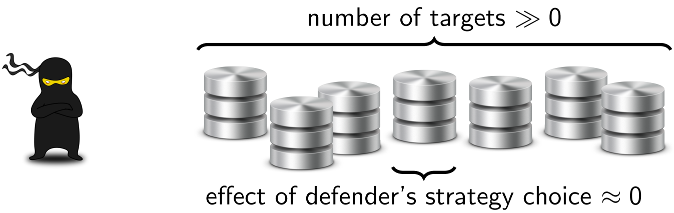

Game-Theoretic Analysis

- Defender has to play "blindly"

  $\longrightarrow$ after each one of her moves, she has the same information (and can be assumed to make her decision the same way)

  $\longrightarrow$ defender plays a renewal strategy

Game-Theoretic Analysis

- Defender has to play "blindly"

  $\longrightarrow$ after each one of her moves, she has the same information (and can be assumed to make her decision the same way)

  $\longrightarrow$ defender plays a renewal strategy
- Since the defender plays a renewal strategy (which is memoryless), the attacker also has the same information after each of the defender's moves (and uses the same non-deterministic function to choose the wait time until her next move)

  $\longrightarrow$ the attacker uses a fixed wait time distribution
  - in the analysis, we use the sum of the wait and attack times, whose cumulative distribution function is denoted by F_S

Defender's Best Response

Lemma

Suppose that the non-targeted attacks arrive according to a Poisson process with rate λ_N, and the targeting attacker uses an adaptive strategy with a fixed wait time distribution given by the cumulative function F_W. Then,

- *not moving is the only best response if $C_D = \mathcal{D}(I)$ has no solution for $I > 0$, where*

$$\mathcal{D}(I) = B_A \left(I F_S(I) - \int_{s=0}^{I} F_S(s)\, ds \right) + B_N \left(-I e^{-\lambda_N I} + \frac{1 - e^{-\lambda_N I}}{\lambda_N} \right);$$

- *the periodic strategy whose period is the unique solution to $C_D = \mathcal{D}(I)$ is the only best response otherwise.*

Attacker's Best Response

Lemma

Against a defender who uses a periodic strategy with period δ_D,

- *never attacking is the only best response if $C_A > \mathcal{A}(\delta_D)$, where*

$$\mathcal{A}(\delta) = B_A \int_{a=0}^{\delta} F_A(a)da \; ;$$

- *attacking immediately after the defender has moved is the only best response if $C_A < \mathcal{A}(\delta_D)$;*
- *both not attacking and attacking immediately are best responses otherwise.*

Equilibrium

Theorem

Suppose that the defender uses a renewal strategy and the targeting attacker uses an adaptive strategy. Then, the equilibria of the game can be described as follows.

1. *If* $C_D = \mathcal{D}^A(I)$ *does not have a solution for* I*, then the* **attacker has an advantage***: there is a unique equilibrium in which the defender does not move and the targeting attacker moves once at the beginning.*
2. *If* $C_D = \mathcal{D}^A(I)$ *does have a solution* δ_D *for* I*:*
   (a) *If* $C_A \leq \mathcal{A}(\delta_D)$*, then* **no player has an advantage***: there is a unique equilibrium in which the defender plays a periodic strategy with period* δ_D*, and the targeting attacker moves immediately after each of the defender's moves.*
   (b) *If* $C_A > \mathcal{A}(\delta_D)$*, then the* **defender has an advantage***:*
      i. *if* $C_D = \mathcal{D}^N(I)$ *has a solution* δ'_D *for* I*, and* $C_A \geq \mathcal{A}(\delta'_D)$*, then there is a unique equilibrium in which the defender plays a periodic strategy with period* δ_D*, and the targeting attacker never moves;*
      ii. *otherwise, there is no equilibrium.*

Equilibrium - Illustration

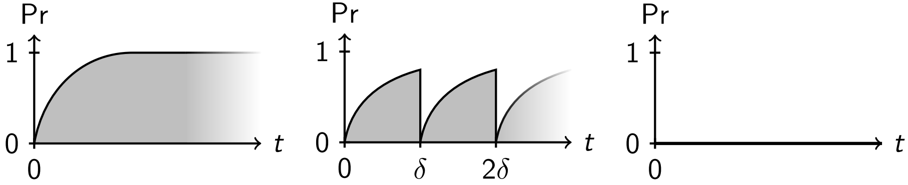

Case 1. | Case 2. (a) | Case 2. (b) i.
attacker has advantage | no player has advantage | defender has advantage

The probability that the targeting attacker has compromised the resource (vertical axis) as a function of time (horizontal axis) in various equilibria.

Sequential Game: Deterrence by Committing to a Strategy

- in practice, the defender can publicly commit to a strategy $\longrightarrow$ sequential game, in which the defender chooses her strategy first and the attacker chooses second
- in this model, we restrict the defender to periodic strategies

Theorem

Let δ_1 be the solution of $C_D = \mathcal{D}^A(\delta)$ (if any), δ_2 be the maximal period δ for which $C_A = \mathcal{A}(\delta)$, and δ_3 be the solution of $C_D = \mathcal{D}^N(\delta)$ (if any). In a subgame perfect equilibrium, the defender's strategy is one of the following:

- *not moving,*
- *periodic strategies with periods $\{\delta_1, \delta_2, \delta_3\}$.*

Numerical Illustrations - Varying the Unit Benefit B_A

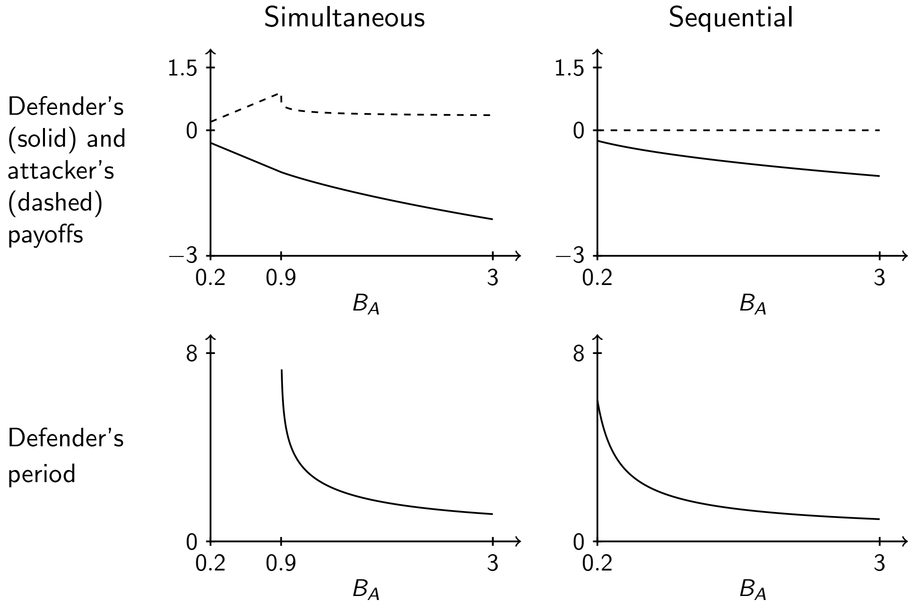

Numerical Illustrations - Varying the Defender's Cost C_D

Defender's (solid) and attacker's (dashed) payoffs

Defender's period

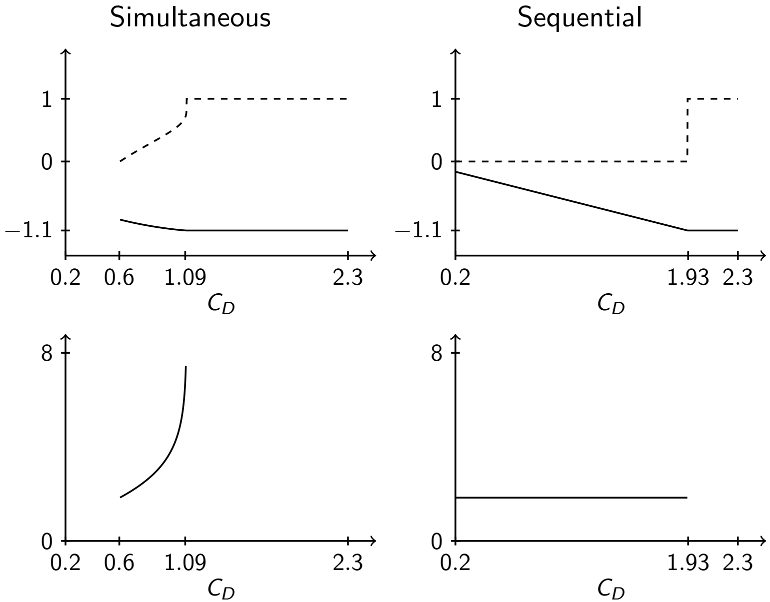

Conclusions and Lessons Learned

- most effective against both types of attacks is the periodic strategy
  - contradicts the lesson learned from the `FlipIt` model [1], which suggests that the defender should use an unpredictable strategy against an adaptive strategy

    Pay attention to what assumptions you make!
  - but justifies the practice of periodic password and key renewal policies

Conclusions and Lessons Learned

- most effective against both types of attacks is the periodic strategy
  - contradicts the lesson learned from the `FlipIt` model [1], which suggests that the defender should use an unpredictable strategy against an adaptive strategy

    Pay attention to what assumptions you make!
  - but justifies the practice of periodic password and key renewal policies
- substantial difference between simultaneous and sequential equilibria
  - defender should not try to keep her strategy secret, but rather publicly commit to it

Conclusions and Lessons Learned

- most effective against both types of attacks is the periodic strategy
  - contradicts the lesson learned from the `FlipIt` model [1], which suggests that the defender should use an unpredictable strategy against an adaptive strategy

    Pay attention to what assumptions you make!
  - but justifies the practice of periodic password and key renewal policies
- substantial difference between simultaneous and sequential equilibria
  - defender should not try to keep her strategy secret, but rather publicly commit to it
- defender is more likely to stay in play and bear the cost of periodic risk mitigation if she is threatened by both types of attacks
  - however, a very high level of either threat type can force the defender to abandon all hope and stop moving

Thank you for your attention!

Questions?

laszka@crysys.hu, johnsonb@ischool.berkeley.edu, jensg@ist.psu.edu

Acknowledgements

We gratefully acknowledge the support of the Penn State Institute for Cyber-Science.

References I

[1] K. D. Bowers, M. van Dijk, R. Griffin, A. Juels, A. Oprea, R. L. Rivest, and N. Triandopoulos.
Defending against the unknown enemy: Applying FlipIt to system security.
In GameSec, pages 248–263, 2012.

Comparison with FlipIt

Contrary to the FlipIt model [1], we assume the following.

- Defender's moves are not stealthy:
  - for most covert attacks with continuous benefits, the attacker knows whether she is in control of the resource
- Targeting attacker's moves are not instantaneous:
  - in practice, an attack requires some (non-deterministic) amount of time and effort to be carried out
- Defender faces multiple attackers:
  - a large range of targets must optimize their defense strategies for both types of attacks

Defender's Best Response Revisited

- recall that, to any attacker strategy, the defender's best response is determined by

$$\mathcal{D}(I) = B_A \left(I F_S(I) - \int_{s=0}^{I} F_S(s)\, ds \right) + B_N \left(-I e^{-\lambda_N I} + \frac{1 - e^{-\lambda_N I}}{\lambda_N} \right)$$

- for particular attacker strategies, we can simplify this formula
  - to not moving, the defender's best response is determined by

$$\mathcal{D}^N(I) = B_N \left(-I e^{-\lambda_N I} + \frac{1 - e^{-\lambda_N I}}{\lambda_N} \right)$$

  - to moving immediately, the defender's best response is determined by

$$\mathcal{D}^A(I) = B_A \left(I F_A(I) - \int_{a=0}^{I} F_A(a)\, da \right) + B_N \left(-I e^{-\lambda_N I} + \frac{1 - e^{-\lambda_N I}}{\lambda_N} \right)$$

Model (extended description)

- Strategic players:
  - defender (denoted by D)
  - targeting attacker (denoted by A)

+ non-strategic actors: non-targeting attackers (denoted by N)

Model (extended description)

- Strategic players
- Resource:
  - some computing resource, e.g., user account, machine
  - having it compromised generates B_i benefit per unit of time for attacker i

Model (extended description)

- Strategic players
- Resource:
  - some computing resource, e.g., user account, machine
  - having it compromised generates B_i benefit per unit of time for attacker i
- Time:
  - continuous
  - game starts at time $t = 0$ with the resource being uncompromised
  - and played indefinitely as $t \to \infty$

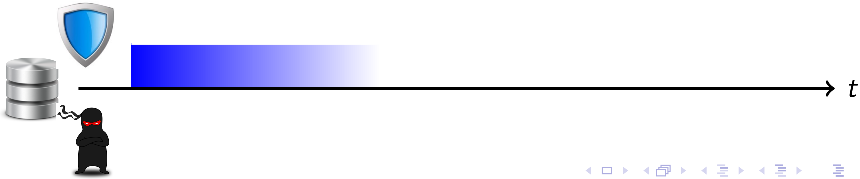

Model (extended description)

- Strategic players
- Resource
- Time
- Moves:
  - at any time instance, player i may make a move, which costs her C_i
  - when the defender makes a move, the resource becomes uncompromised immediately, but the attackers will know of it
  - when the targeting attacker makes a move, she starts her attack, which takes some random amount of time
    - ★ distribution of the attack time is given by the cumulative function F_A,

    but the attackers' moves are stealthy (i.e., the defender does not know when the resource became compromised or if it is compromised at all)

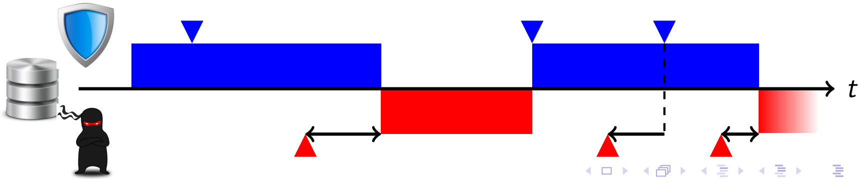

Model (extended description - contd.)

- Strategy:
  - set of rules, algorithm, etc. for making moves
  - in practice: defender's key or password update policy, targeting attacker's plan of attack, etc.

Model (extended description - contd.)

- Strategy:
  - set of rules, algorithm, etc. for making moves
  - in practice: defender's key or password update policy, targeting attacker's plan of attack, etc.
- Cost rate $c_i(t)$:
  - for player i up to time t, the cost rate $c_i(t)$ is the number of moves per unit of time made by player i up to time t, multiplied by the cost per move C_i

Model (extended description - contd.)

- Strategy:
  - set of rules, algorithm, etc. for making moves
  - in practice: defender's key or password update policy, targeting attacker's plan of attack, etc.
- Cost rate $c_i(t)$:
  - for player i up to time t, the cost rate $c_i(t)$ is the number of moves per unit of time made by player i up to time t, multiplied by the cost per move C_i
- Benefit rate $b_i(t)$:
  - for attacker i, the benefit rate $b_i(t)$ up to time t is the fraction of time up to t that the resource has been compromised by i, multiplied by the unit benefit B_i (note that if multiple attackers have compromised the resource, they all receive benefits until the defender's next move)
  - for the defender D, the benefit rate $b_D(t)$ up to time t is

$$-\sum_{i\in\{A,N\}} b_i(t)$$

- Payoff: player i's payoff is defined as

$$\liminf_{t\to\infty}\ b_i(t) - c_i(t)\ .$$