

Games and Cyber Security

Seminar at University of Bristol

Manos Panaousis

June 3, 2015



University of Brighton

Outline

- ① Why and how game theory can support decision-making in cyber security
- ② Discuss fundamental concepts of game theory
- ③ Analyse how cyber security investments can be made by modeling and solving a Nash game

A simple example to start with...

Table: **The Malware Game**

Defender, Attacker	Install mw.	Don't Install mw.
Anti-mw. defences	$-c, -a$	$-c, 0$
No Mw. defences	$-d, r$	$0, 0$

- c : antimalware software cost
- d : damage when compromised
- a : cost for attacker installing malware and maybe getting arrested too through forensics analysis
- r : reward of the attacker

Note that without loss of generality and to be realistic $d > c$: The damage when compromised is higher than the money spent for the purchasing anti-malware software

Motivation

- **Objective:** Develop a quantitative framework for network security including
 - Intrusion detection
 - Security and privacy models and analysis
 - Optimised response
- **Tools**
 - Game theory
 - Artificial intelligence and machine learning
 - Optimisation
- We utilise **game theory** in order to analyse the interaction between **attackers** and **systems** (intrusion detection systems, security administrators) in the context of cyber security



Application domains

GT has been examined in the **context of**

- physical and MAC layer security
- application layer security in mobile networks
- intrusion detection systems
- anonymity and location privacy
- economics of cyber security
- cryptography

Advantages of a game theoretic approach

- its quantitative approach and comprehensive mathematical modeling capabilities
- paving the way to **automatic decision making** on
 - reconfiguration of security policies given the severity of attacks
 - allocation of limited resources in real time for detecting significant threats to vital subsystems in a large networked system



What is a game?





- Game theory involves **multi-person decision making**
- Autonomous parts of the networked systems (such as software agents) as well as malicious attackers and intrusion detection systems are modeled as players
- **Security games:**
 - ① **Players:** Defender and Attacker
 - ② **Action Space:** Set of defensive measures or attacks
 - ③ **Outcome:** Cost and benefit to players for each action-reaction or game branch
 - ④ **Information Structures:** Players fully or partially observe each other's actions

More on game theory...

- Discipline aiming at **modelling situations** in which actors have to make decisions which have **mutual, possibly conflicting, consequences** (i.e. economics, but also politics and biology)
- Most widespread kind of game: **non-cooperative** (meaning that the players do not attempt to find an agreement about their possible moves)
- Players interact and compete with each other on the same system (for limited and shared resources)
- Players are associated with **reward** (resp., **cost functions**), which they maximise (resp. minimise) by choosing a strategy from well defined strategy sets
- **Nash equilibrium** (NE) provides an appropriate solution concept, which is (approximately) optimal (wrt a global objective function)

My research

Motivation - limited cyber security resources

- ▶ A report published by Deloitte and NASCIO*, points out that:
 - ➔ Only 24% of CISOs are very confident in protecting their organisation's assets against external threats. 
 - ➔ The biggest concern CISOs face in addressing cybersecurity is a "Lack of sufficient funding". 
- ▶ Most organisations will have a fixed budget for the protection of their systems.  $\leq B$
- ▶ As such an organisation is interested in how to use the limited financial budget available to best protect them from various vulnerabilities given that the implementation of a cybersecurity control is associated with a direct cost 



Ross Anderson's view...



Ross Anderson first proposed the study of security from an economics perspective putting forward the idea that cyber security is bounded by other non-technical incentives

Spending less £ == less protected?

Anderson highlighted with an example that although some organizations spend less money on security they spend it more effectively therefore having put in place better cyber defences

In our work we share Andersons view!



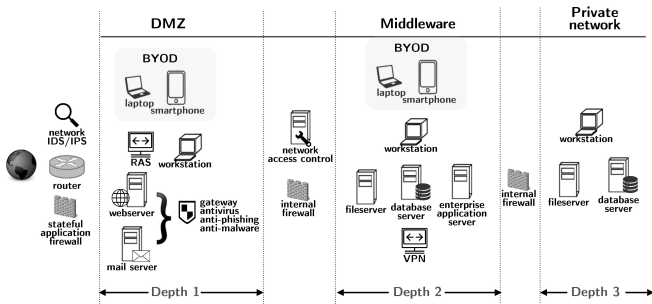
- Emmanouil Panaousis, Andrew Fielder, Pasquale Malacaria, Chris Hankin, Fabrizio Smeraldi. cyber security games and investments: A decision support approach. In Proceedings of the 5th Conference on Decision and Game Theory for Security. (GameSec 2014), Los Angeles, CA, USA, November 6-7, 2014.
- Andrew Fielder, Emmanouil Panaousis, Pasquale Malacaria, Chris Hankin, Fabrizio Smeraldi. Comparing Decision Support Approaches for Cyber Security Investment. arXiv:1502.05532 Subject: Computer Science and Game Theory (cs.GT); Cryptography and Security (cs.CR) (under journal review)

Decision-support tools vs. Market efficiencies and inefficiencies

However our approach is quite different as we focus on developing cyber security decision support tools to assist security managers on how to spend a cyber security budget in terms of different controls acquisition and implementation

Organisational structure

- We follow the network architecture as proposed in the *SANS Critical Security Control 19-1* entitled “Secure Network Engineering”



Definition (Commodity Attacks)

Commodity attacks are attack methods where the attack tools can be purchased by a user, where the adversaries do not develop the attacks themselves, and only configure the tools for their own use.

Network architecture

DMZ

An organization's assets that can be accessed from the Internet are placed in the *DMZ*, and they should not contain any highly sensitive data

Middleware & Private Network

Any asset with highly sensitive data must be located at the *Private Network*, and communicate with the outside world only through a proxy which resides on the *Middleware*

Definition (Depth)

The *depth* of an asset, denoted by d , is the location of this asset within an organization's network architecture

- ① Depths are separated from each other by a set of network security software, e.g., firewalls, IDS
- ② A depth determines
 - the level of security that needs to be breached or bypassed in order for an attack to successfully exploit a vulnerability at this depth, and,
 - the importance of the data asset compromised if an attack is successful

Targets

Definition (Target)

A *cyber security target* is defined as a (*vulnerability*, *depth*) pair, i.e. $t_i = (v_z, d)$

- A target abstracts any *data asset*, located at depth $d \in \{1, \dots, n\}$, that an attack threatens to compromise by exploiting $v_z \in V$.
- We define the set of all targets as

$$T = \{(v_z, d) | v_z \in V, d \in \{1, \dots, n\}\}$$

- Each network architecture has its own set of targets
- We specify that *data assets* located at the *same depth* and having the same *vulnerabilities* are abstracted by the *same target*, and they are worth the *same value* to the organization

Cyber security controls

- A *cyber security control* is the defensive mechanism that can be put in place to alleviate the risk from one or more attacks by reducing the probability of these attacks successfully exploiting vulnerabilities
- The Defender can choose to implement a control c_j at a certain level l

Definition (Cyber Security Process)

A cyber security process is the implementation of a control at a certain level, and we denote by p_{jl} the cyber security process that implements the control c_j at level l

- A cyber security process p_{jl} has a *degree of mitigation* for each target t_i which equals the effectiveness of the cyber security process on this target, denoted by $e(t_i, p_{jl}) \in [0, 1)$

What are interested in...

- We are interested in how cyber security processes are combined in a proportional manner to give an implementation plan for this control
- We call this **a cyber security plan** which allows us to examine advanced ways of mitigating vulnerabilities

Definition (Cyber Security Plan)

A cyber security plan is a probability distribution over different cyber security processes

Map to reality...

- Consider a security control entitled **Vulnerability Scanning and Automated Patching**
- Assume 5 different implementation levels i.e. $\{0, 1, 2, 3, 4\}$ where level 4 corresponds to *real-time scanning* while level 2 to *regular scanning*
- A mixed strategy $[0, 0, \frac{7}{10}, 0, \frac{3}{10}]$ determines a cyber security plan that dictates the following:
 - ① $\frac{3}{10} \mapsto$ real-time scanning for the 30% of the most important devices
 - ② $\frac{7}{10} \mapsto$ regular scanning for the rest 70% of devices
- This mixed strategy can be realized more as an advice to a security manager on how to undertake different control implementations rather than a rigorous set of instructions related only to a time factor

Risks

- The perceived impact of a successful attack against targets at a given depth
- The **target risks** express the damage incurred to the Defender when the Attacker succeeds in compromising one or more targets
- The different risks we consider are:
 - ① Data Loss (DL)
 - ② Business Disruption (BD), and
 - ③ Reputation (RE)
- Each risk factor depends on the depth d that the attack targets
- We denote by DL_d , RE_d , and BD_d the risk values associated with a depth d

Indirect costs

- ① **System Performance Cost** (SPC): anything related to system performance being affected by a cyber security process (e.g. processing speed affected by anti-malware scanning)
- ② **Morale Cost** (MOC) accounts for morale issues that higher levels of security can cause to users' happiness and job satisfaction
 - The stricter the security measures that an organization implements, the more likely an individual will circumvent them
 - e.g. Having a control about different passwords for everything, might annoy users \Rightarrow MOC \uparrow
 - The user picking weak, memorable passwords which can often be cracked by dictionary or brute force attacks
- ③ **Re-Training Cost** (RTC): cost for re-training users, including system administration, so they can either perform the cyber security process in the right way or be able to continue using all systems after a security update

We express the different indirect costs of a cyber security process p_{jl} by SPC_{jl} , RTC_{jl} , and MOC_{jl}

Direct costs

Each cyber security process has a direct cost which refers to the **budget the organization must spend** to implement the control c_j at a level l

- ① **Capital Cost** (CAC) is related to hardware or software that must be purchased for the implementation of a control at some level
- ② **Labour Cost** (LAC) is the direct cost for having system administrators implementing the control such as

$$(\text{hours spent}) \times (\text{cost/hour})$$

Vulnerability factors

- The Council on cyber security has published a set of 25 **software weaknesses** (i.e. vulnerabilities) and their factors: *Prevalence* (PR), *Attack Frequency* (AF), *Ease of Detection* (ED), and *Attacker Awareness* (AA)
- The level of a factor determines its contribution towards an overall **vulnerability assessment score**

We assume Commodity Attacks

Commodity attacks are attack methods where the attack tools can be purchased by a user, where the adversaries do not develop the attacks themselves, and only configure the tools for their own use \Rightarrow To facilitate our analysis when solving a complete information game

Vulnerability factors

For a commodity attack, one can argue that

- AA measures whether the average adversary would know that a malicious script is for sale,
- ED is a measure of the computational cost of the attack discovery process,
- PR indicates the number of times the weakness is found in the system (e.g. only 30% of windows systems ever downloaded a given patch), and
- AF dictates the number of times someone actually tries to exploit it (e.g. how many random SQL injection probes a day)
- PR and AF accounting for threats that are currently widespread (**current threats**), and
- ED and AA for threats that have the most potential for future attack vectors (**future potential threats**)

Organisational profile

- Characteristics unique to the Company or organisation, that dictate how they perceive aspects of their concerns outside of technical knowledge
- Different importance weights for risks RE, DL, BD
- Different importance weights for current and future threats

Game theoretic formulation

- We define a two-player non-cooperative game where there is a negative functional correlation between the Attacker and the defender payoffs; the idea is that **the more an Attacker gains the more the Defender loses**
- The Defender \mathcal{D} abstracts any cyber security decision-maker (e.g., security manager) which defends an organisation's data assets by minimising cyber security risks with respect to the indirect costs of the cyber security processes
- The Attacker \mathcal{A} abstracts all adversaries that aim to benefit from compromising the Defender's data assets by using commodity attacks

Control game

- We consider $|L|$ implementation levels for the j th control (c_j)
- For each $\lambda \in \{1, 2, \dots, |L|\}$ we define the λ -control subgame, denoted by $\mathcal{G}_{j\lambda}$

Pure strategies in $\mathcal{G}_{j\lambda}$

A control-subgame $\mathcal{G}_{j\lambda}$ is a game where

- 1 \mathcal{D} 's pure strategies correspond to **consecutive implementation levels** of the control c_j starting always from 0 (i.e. fictitious control-game), and including all levels up to λ , and
- 2 \mathcal{A} 's pure strategies are the different **targets** akin to pairs of vulnerabilities and depths

Why do we consider control subgames and not just control games?

Nash cyber security plan

- Each control-subgame has a **Nash cyber security plan**, which dictates how to implement a control in the form of a mixed strategy
- For the finite nonempty set of implementation levels, let Δ_λ represent the set of all probability distributions over it, i.e., $\Delta_\lambda := \{\delta \in \mathbb{R}^{+R} \mid \sum_{l \in \{1, \dots, \lambda\}} \delta(l) = 1\}$, likewise we define Δ_T for the targets
- The mixed strategy $\delta \in \Delta_\lambda$ of the Defender is a probability distribution over the different implementation levels
- The mixed strategy $\alpha \in \Delta_T$ of the Attacker is a probability distribution over the different targets

Example of a mixed strategy

- Consider a security control related to **password policy**, and its 5 different implementation levels i.e. $\{0, 1, 2, 3, 4\}$
- For instance, level 4 corresponds to **strong passwords that must change monthly** - you can define the other similarly!
- We assume an organisation with 1,000 employees among which 90 are senior managers (SM), 10 senior system administrators (SSA), and 900 other employees (OE) lower in hierarchy than SM and SSA
- The level of each class of users is determined by the **importance of data** their accounts have access to
- A **mixed strategy** akin to cyber security plan $[0, 0, \frac{7}{10}, 0, \frac{3}{10}]$ says to implement level 2 of the control for 70% of the employees and for the rest 30% the control must be implemented at level 4

Payoffs

- Typically Defender payoffs are determined by:
 - 1 $S(l, t)$: **expected security loss** at target t when a control has been implemented at level l (e.g. DL, BD, RE)
 - 2 $C(l)$: **indirect cost** of implementing the control at level l (e.g. SPC, RTC, MOC)
- A Toy Game:**
 - Defender (row player) can implement the control at two different levels l, l'
 - Attacker (column player) can choose two targets t, t'

Table: **Defender's Game Matrix**

	t	t'
l	$S(l, t) - C(l)$	$S(l, t') - C(l)$
l'	$S(l', t) - C(l')$	$S(l', t') - C(l')$

Payoffs

- For a Pure strategy profile: $U_{def}(l, t) = S(l, t) - C(l) = \text{impact} \times \text{threat} \times (\text{1-mitigation}) - \text{indirect cost}$
- **impact** = $r_1 DL + r_2 RE + r_3 BD$, where (r_1, r_2, r_3) expresses the organisational profile
- **threat** = $\tau_1 (PR + AF) + \tau_2 (ED + AA)$, where τ_1, τ_2 are the weights given by the company to *current threats* and *current threats*, respectively
- For a Mixed strategy profile:

$$U_{def}(\delta, \alpha) = \mathbb{E}_{\sim \delta, \alpha} U_{def}(l, t) = \sum_{l \in \{1, \dots, \lambda\}} \sum_{t \in T} U_{def}(l, t) \delta(l) \alpha(t)$$

Game structure

- We consider **zero-sum games**, i.e. the loss of the Defender is the Attacker gain.
- More realistic scenarios where the Attacker gain is a fraction of the Defender loss (stolen goods are worth a fraction of their original value) still **behave as zero-sum games**
- Zero-sum games allow us to consider worst-case scenarios (the most aggressive Attacker)
- In a zero sum game **Nash equilibrium**, **MiniMax** and **MaxMin** coincide

MaxMin Defence Strategy $\max_{\delta} \min_{\alpha} U_{def}(\delta, \alpha)$

In security terms MaxMin means that we are considering the best defence against all possible attackers, even irrational ones - any other strategy would score worst on at least one target

Cyber security investment optimisation

- For a given C set of cyber security controls we have, in total, $|L| \times |C|$ **Nash cyber security plans**
- Let us assume β is the budget that a company has available to spend on cyber security
- We examine how to optimally invest in the different plans by choosing at most one plan per control, i.e. $|L|$ plans
- We model this cyber security investment optimization problem as a **0-1 Multiple-Choice Multi-Objective Knapsack Problem**

Combinatorial optimisation: classic Knapsack

- One target (=objective), a budget and a set of resources (cyber security plans) each with a cost (financial cost, i.e. direct cost) and a benefit (how much it improves the Defender's payoff on each target) wrt the target
- **Problem:** find the “optimal” set of cyber security plans given a budget constraint

Multiple choice multi objective 0-1 Knapsack

- Given multiple targets and a partition of resources (cyber security plans), where each resource may benefit multiple targets
- **Problem:** find a subset of resources (let us call it **investment solution**) with exactly 1 element in each block of the partition (i.e. 1 cyber security plan per control) satisfying budget constraints so that any other choice will leave at least one target worse off
- In more detail, **each investment solution has a score** determined by the maximum expected damage across all targets (**weakest link in cyber security**) - **not the sum of damages!**
- In Security terms: find an investment within budget such that any other investment will leave at least one target more exposed

Optimisation problem

$$\max_{\vec{z}} \min_{t_i} \left\{ 1 - \sum_{j=1}^C \sum_{l=0}^m E(Q_{j,l}, t_i) z_{j,l} \right\} I(t_i) T(t_i)$$

$$\text{s.t. } \sum_{j=1}^C \sum_{l=0}^L \Gamma(Q_{j,l}) z_{j,l} \leq B, \forall j = 1, \dots, C$$

$$\sum_{l=0}^L z_{j,l} = 1, z_{j,l} \in \{0, 1\}, \forall j = 1, \dots, C$$

- $E(Q_{j,l}, t_i)$ is the effectiveness of plan $Q_{j,l}$ on protecting t_i
- Candidates are solutions from the control subgames (this is why indirect costs C are not present because they are part of the games solutions)
- $\Gamma(Q_{j,l})$ is the direct cost of plan $Q_{j,l}$
- $z_{j,l}$ is the decision variable, i.e. $Q_{j,l} \in \text{solution}$ then $z_{j,l} = 1$.

Comparisons

- We have compared the above **Hybrid model** (Game theory and Multi-objective knapsack) with two other investment strategies
- **Pure game**: a game where the Defender has available all possible controls implemented at all levels and attacker all possible weaknesses
 - This would be a very large game; also it would have no budget constraints on the solutions;
 - We use affordable schedules as pure strategies instead where a schedule is affordable if it is within budget
 - A schedule is a tuple of $|C| \times |L|$ bits where each bit represents the implementation of a control at a particular level, 1 stands for implemented and 0 for not implemented
- **Multiple-choice multi-objective knapsack**: given the partition of all cyber security plans, each plan with a direct cost and a benefit, find the optimal investment solution i.e. any other investment within budget will leave at least one target more exposed

Multiple-choice multi-objective knapsack

$$\begin{aligned} \max_{\vec{z}} \min_{t_i} & \left\{ 1 - \sum_{j=1}^C \sum_{l=0}^L E(Q_{j,l}, t_i) z_{j,l} \right\} I(t_i) T(t_i) - \left(\sum_{j=1}^C \sum_{l=0}^L C(Q_{j,l}) z_{j,l} \right) \\ \text{s.t.} & \sum_{j=1}^C \sum_{l=0}^L \Gamma(Q_{j,l}) z_{j,l} \leq B, \forall j = 1, \dots, C \\ & \sum_{l=0}^L z_{j,l} = 1, z_{j,l} \in \{0, 1\}, \forall j = 1, \dots, C \end{aligned}$$

What the solutions look like

- A solution would be a set of controls at particular levels: e.g. A solution would be a set of controls at particular levels: e.g. (patch every week, change password every 2 months, monitor and log analysis biweekly ...)
- Game theoretical solutions may also include mixed strategies, this could look like: [(70% of devices to be patched every week, 30% of devices to be patched every day), ..., 90% logs must be investigated weekly, 10% of logs must be investigated monthly), ...]

Case study - Initial results

TABLE V: Case Study Controls.


Control	Levels
Inventory of Authorised and Unauthorised Devices (1)	3
Inventory of Authorised and Unauthorised Software (2)	3
Secure Configuration for Hardware and Software on Devices (3)	5
Continuous Vulnerability Assessment and Remediation (4)	4
Malware Defences (5)	6
Application Software Security (6)	2
Controlled Use of Administrative Privileges (12)	6

TABLE VI: Case Study Vulnerabilities.

<i>v_z</i> : Vulnerability (CWE-code)	PR	AF	ED	AA	Vulnerability	PR	AF	ED	AA
<i>v₁</i> :SQLi (89)	2	3	3	3	<i>v₇</i> :Missing encryption (311)	2	2	3	2
<i>v₂</i> :OS command injection (78)	1	3	3	3	<i>v₈</i> :Unrestricted upload (434)	1	2	2	3
<i>v₃</i> :Buffer overflow (120)	2	3	3	3	<i>v₉</i> :Unnecessary privileges (250)	1	2	2	2
<i>v₄</i> :XSS (79)	2	3	3	3	<i>v₁₀</i> :CSRF (352)	2	3	2	3
<i>v₅</i> :Missing authentication (306)	1	2	2	3	<i>v₁₁</i> :Path traversal (22)	3	3	3	1
<i>v₆</i> :Missing authorization (862)	2	3	2	2	<i>v₁₂</i> :Unchecked code (494)	1	1	2	3

Case study which includes vulnerabilities (i.e. CWE) and cyber security controls published by the Council on cyber security - SANS 20 Critical Security Controls)

Case study



As budget increases the **Hybrid Method scores better** because it includes the impact of indirect costs in the decisions regarding the optimality of the deployment of the control at each level, the pure Knapsack includes the indirect cost, as a whole, in the outcome of the optimisation, and the Full Game applies the indirect

Thank you



I would particularly like to thank my collaborators: Dr Andrew Fielder (IC), Prof Chris Hankin (IC), Prof Pasquale Malacaria (QM), Dr Fabrizio Smeraldi (QM), Dr Tansu Alpcan (UoM), and Dr Arman Khuzani (QM) and the host Prof Nigel Smart

More info about me on: <https://www.panaousis.com>