Risk Management and Game Theory: a case study on cybersecurity

Jesús Palomo

jesus.palomo@urjc.es
Rey Juan Carlos University

May 16^{th} – 20^{th} 2016
SAMSI

Table of contents

The problem

Characteristics of Organized Crime (OC)

- Cybercriminals are highly organized and coordinated in terms of their speciality.
  - In some sense, they are endowed. They would rather change target than change technique.
- OC fights among them to obtain illegally money from the Financial System.

Three negotiations

- The Preparation: OC looks for opportunities to seize and gather information units (Log/Pass, Credit cards, PINs...)
- The Dark market: OC exchanges pieces of information
- The Execution: OC actually uses the information and gathers money

Utility function components u_A: *anonymity*, *accessibility* and *value*.

Financial sector

Goals

- Define robust defense strategy against technological risk: it has an impact on customers.
- Powers:
  - Economic pain threshold: There are products that, if they become too risky, it can be closed immediately.
  - Monitors partially the activity. There are identity thefts, card theft, duplicated cards, etc. that they do not control. This provides information, but only partially.
  - It has the possibility to affect the three negotiations that OC has.

Example: Turkish Banks early 2003

Banks (on-line) in Istanbul noticed increase stolen accounts

- OC: installed keystroke logging spy sw on public computers.
- Banks: changed the web logging/pass. textbox by a virtual keyboard.
- OC: within days developed sw to capture clicks on screen.
- Banks: one-time password randomly generated by hw devices.

Outcome: OC activity shifted to other countries, and negative impact on customers because of the changes.

Alternative: US Banks

US Banks to increase customer use of on-line banking (saves money)

- Consider identity thefts as cost of doing business (compensated by cost saved by online banking)
  - Reasonable from risk management, but has unintended consequences.
- Banks attracted new international OC and the enriched cybercriminals developed more sophisticated attacks.

Apparently, in this setting (and in several other cases), players look for higher payoffs, and those with higher payoff will increase their population (or probability) and those with less payoff will decrease their population (or probability).

Defend-Attack models: Formulation

Defender has a discrete set $\mathcal{D}\{d_1, \ldots, d_m\}$ and the Attacker has a discrete strategy set $\mathcal{A} = \{a_1, \ldots, a_n\}$ (may include do-nothing or mix strategies)
For each (d_i, a_j), the probability that an attack is successful $\pi(\omega|d_i, a_j)$, the utility functions $u_D(a_i, d_j, \omega)$ and $u_A(a_i, d_j, \omega)$, and the expected payoffs

$$\psi_D(d, a) = \int u_D(a_i, d_j, \omega)\pi(\omega)\ d\omega,$$

$$\psi_A(d, a) = \int u_A(a_i, d_j, \omega)\pi(\omega)\ d\omega.$$

Defend-Attack models: Common knowledge

Each player knows $\psi(d, a)$ for each pair $(d, a) \in \mathcal{D} \times \mathcal{A}$ and this fact is known by all of the players. Furthermore, any other uncertainty in the game it is assumed that players have common probabilities over the uncertain variables.
A Nash Equilibrium (NE) $(d^*, a^*(d^*))$

$$\begin{aligned} a^*(d) &= \arg\max_{a \in \mathcal{A}} \psi_A(d, a), \ \forall d \in \mathcal{D}, \\ d^* &= \arg\max_{d \in \mathcal{D}} \psi_D(d, a^*(d)), \end{aligned}$$

$$\begin{aligned} \psi_D(d^*, a^*) &\geq \psi_D(d, a^*), \ \forall d \in \mathcal{D} \text{ and} \\ \psi_A(d^*, a^*) &\geq \psi_A(d^*, a), \ \forall a \in \mathcal{A}. \end{aligned}$$

Defend-Attack models: ARA

- When relaxing the common knowledge

$$\psi_D(d_i, a_j) = \int u_D(d_i, a_j, \omega) p_D(\omega | d_i, a_j) \, d\omega$$

and the same problem for the attacker $\psi_A(d_i, a_j)$, with $p_A(\omega | d_i, a_j)$

- $p_D(\omega | d_i, a_j)$, $p_A(\omega | d_i, a_j)$, $\pi_D(a|d)$, $\pi_A(d)$, and $u_A(d_i, a_j, \omega)$ are unknown.

$$a(\theta) = \arg\max_{a \in \mathcal{A}} \sum_{d \in \mathcal{D}} \psi_A^\theta(d, a) \pi_A(d)$$

and selects the strategy that maximizes her expected utility

$$d(\theta) = \arg\max_{d \in \mathcal{D}} \sum_{a \in \mathcal{A}} \psi_D^\theta(d, a) \pi_D(a|d)$$

Defend-Attack models: ARA

When the attacker is rational

$$\begin{aligned} a(\theta, d)^* &= \arg\max_{a \in \mathcal{A}} \psi_A^\theta(d, a), \ \forall d \in \mathcal{D} \\ d(\theta)^* &= \arg\max_{d \in \mathcal{D}} \psi_D^\theta(d, a(\theta, d)^*) \end{aligned}$$

corresponding Nash Equilibria $(d(\theta)^*, a(\theta)^*)$. By taking the expectation with respect to θ it can be approximated the distribution $\pi_D(a|d) = E_\theta(a(\theta)^*)$. Finally, the defender maximizes her expected utility by playing d^* as follows

$$d^* = \arg\max_{d \in \mathcal{D}} \sum_{a \in \mathcal{A}} \psi_D(d, a) \pi_D(a|d)$$

Alternative rational thinking

Let's review the following example under common knowledge

		Attacker	
		a_1	$\mathbf{a_2}$
Defender	d_1	-10,2	-2,**3**
	d_2	-3,8	-9,**9**

NE is on $p = 1$, $q = 0$, but for an attacker, playing a mix strategy $q \neq 0$ has huge impact on the Defender for a small expected payoff reduction.

Alternative rational thinking

Attacker's mix strategy is not optimal from the NE viewpoint. In fact, there is a non stable equilibrium. The only situations that clearly will be avoided, and this is common knowledge, are:

a) Defender has no incentives to play d_2:
- if $q = 0$ we find the NE (d_1, a_2)
- if $0 < q < 0.5$ then $\psi_D(d_1, a) > \psi_D(d_2, a)$, and
- if $0.5 \leq q < 1$ then $\psi_D(d_1, a) \leq \psi_D(d_2, a)$,$\psi_D(d_2, a) = -9 + 6q \geq \psi_D(d_1, a) = -2 - 8q$., but defender, playing d_2, is highly increasing the attacker payoff $\psi_A(d_1, a) = 3 - q$ and $\psi_A(d_2, a) = 9 - q$.
- If $q = 1$, then $\psi_D(d_2, a) > \psi_D(d_1, a)$, but in this case the attacker has incentives to switch to a_2 since $\psi_A(d_2, a_2) > \psi_A(d_2, a_1)$.

b) Attacker has no incentives to play a_1 in pure strategies since a_2 strictly dominates a_1 no matter what the defender does.

ARA for $\pi_D(a|d)$: Observational viewpoint

Assuming attacker chooses at random (although strategic) historical data can be collected. assuming, for example, a $\pi_D(a|d) \sim \mathfrak{D}(\alpha_1, \ldots, \alpha_n)$ where the Dirichlet parameters are updated with the observations from an stage game. Repeated sequential of plays of same game would not be appropriate (learning process)

Based on the observations, the posterior distribution could be approximated by $\pi_D(a|\cdot) \sim Be(\alpha, \beta)$ For the example, $\alpha = 2$ and $\beta = 5$.

ARA for $\pi_D(a|d)$: Strategic opponent

There are authors that assume certain rationality in the strategies of the opponents that goes beyond the common models, e.g. Nash equilibria, level-k thinking, mirroring equilibria, etc.

Paté-Cornell and Guikema (2002) recognize that the model for $\pi_D(a|d)$ must capture the dynamics and game-analytics of the strategic decisions. In particular, the probability of a particular attack is proportional to the ease of the execution and the damage or profit obtained for the attacker.

$$\pi_D(a_j|d) = \frac{\psi_A(d, a_j)}{\sum_j \psi_A(d, a_j)}.$$

ARA for $\pi_D(a|d)$: Strategic opponent

From a survival point of view, these probabilities justify the strategy of a population (the attackers) that seek survival interests trough attacks.

In the previous example, attacker's mix strategy is $q = \frac{8-6p}{17-12p}$ as opposite to the NE ($q = 0$).
Since the defender plays d_1 ($p = 1$) as best option, so the attacker uses a_1 with probability $q = \frac{2}{5}$.

Although this approach does not formally consider the attacker as an expected utility maximizer, it does capture an evolutionary dynamic of a player. Attackers are players that seek profit/reduce losses causing harm while they are involved in a survival game with both their congeners and opponents.

Evolutionary Game Theory: motivation

It allows to deal with strategic interactions between entities (groups, species,etc.). It is a formal model of strategic interaction over time in which:

- Survival of the fittest maxim: Higher payoff strategies tend over time to displace lower payoff strategies,
- there is some inertia, i.e. no revolutionary change so aggregated behavior does not change abruptly, and
- players do not systematically attempt to influence other players' future actions, i.e. evolution is not repeated games.

The aim is to study the evolution of the strategies in the population, where individuals are endowed to play certain strategies. More successful strategies survive with higher probability.

Evolutionary Game Theory: motivation

- The objective is to characterize a stable population mix of strategies. Those individuals that use a strategy may come and go, but the mix of strategies in the population can persist.
- EGs offer a different perspective on information conditions than *rational* static complete information (NE) and incomplete information (Bayes-NE) models.
- EGs offer a built-in selection criteria among NE and predictions of relatively rapid or slow convergence (or non-convergence) to NE.

Evolutionary Game Theory: motivation

- An individual member cannot modify his behavior, the proportion of members who use a strategy can evolve.
- The expected utilities (payoffs) are now in terms of fitness: Darwinian reproductive success of the involved individuals.
- Mixed strategies are reinterpreted as population frequencies.
- Rationality is substituted by genetic transmission on behavioral strategies to offspring.
- A population is stable when it is resilient to a *mutation*.
- For simplicity, we focus on homogeneous populations, so that all members use same strategy. For simplicity, again, two-person games with an added wrinkle (the opponent is randomly selected from a given population).

Evolutionary Game Theory: The conditions

- If entry and exit, or resource redistribution, or learning tends to increase the prevalence of high-payoff actions over time in each population, then an evolutionary game model may be appropriate.
- Time paths in such models quite generally converge to evolutionary equilibria EE, a subset of the Nash equilibria NE for the underlying state game.
- The learning (or other dynamic) process, together with the historical given initial state, determines which EE is attained when there are more than one NE.

Evolutionary Game Theory: Some examples

- (Arce and Sandler 2003) use evolutionary game models to show the conditions under which moderates within a society adopt extremist preferences in order to fit within the extremist group.
- (d'Artigues and Vignolo 2003) mimetic rivalry on Terrorism.
- (Sadler 2009) Games and Terrorism and the choice between proactive and defensive countermeasures.
- (Harrington 2009) Revision on Games, strategies and decision making.
- (Frey and Luechinger 2003) Alternatives to deterrence in fighting terrorism.
- (Eid, El-adaway, Coatney 2015) Strategy for post-disaster insurance.
- (Kasthurirathna and Piraveenan 2015) strategies in network games.

Evolutionary Game Theory: the EE

A evolutionary stable strategy ESS must satisfy the condition that it be resistant to the arrival of a small mutation deploying a different strategy. For example, $C > V > 0$

		Criminal 2	
		Aggressive	Cautious
Criminal 1	Aggressive	$\frac{V-C}{2}, \frac{V-C}{2}$	$V, 0$
	Cautious	$0, V$	$\frac{V}{2}, \frac{V}{2}$

This game has 3 NE, (*aggressive,cautious*), (*cautious,aggressive*) and $p = \frac{1}{2}$.

Evolutionary Game Theory: the EE

The expected fitness of an individual endowed *aggressive* when meeting a small proportion ϵ of *mutants* that play *cautious* is

$$\phi(aggressive) = (1-\epsilon)\left(\frac{V-C}{2}\right) + \epsilon V$$

for those mutants endowed *cautious*

$$\phi(cautious) = (1-\epsilon)\cdot 0 + \epsilon\left(\frac{V}{2}\right)$$

aggressive population will drive out the *mutants* when $\phi(aggressive) > \phi(cautious)$. But if $\epsilon < \frac{C-V}{C}$ *mutants* can invade the population of *aggressive* individuals. So

- Population of all *aggressive* is not stable ($\epsilon < \frac{C-V}{C}$)
- Population of all *cautious* is not stable ($\epsilon < \frac{V}{C}$)
- NE (*aggressive*,*cautious*), (*cautious*,*aggressive*) are not stable.

Evolutionary Game Theory: the EE

Definition: An ESS, when individuals are randomly matched to play a symmetric two-player game, is defined as one for which either

$$\phi(s,s) > \phi(r,s),\ \forall r,\ \textit{or} \quad (1)$$

$$\phi(s,s) = \phi(r,s),\ \text{for some } r, \text{ then } \phi(s,r) > \phi(r,r). \quad (2)$$

s resists all possible *mutations* r either because (1) they are less fit or (2) they are fit at the current state (where they are rare) but less fit when they are prevalent.

$$\phi(s,r_\epsilon) > \phi(r,r_\epsilon),\ \text{for } \epsilon > 0 \text{ sufficiently small,}$$
$$\text{and } r_\epsilon = (1-\epsilon)s + \epsilon r$$

Evolutionary Game Theory: the EE

Conditions:

- $\phi(s,s) \geq \phi(r,s),\ \forall r$ is the NE. It is clear from a comparison of the definitions that every ESS is a NE.
- every ESS is a NE
- if an ESS is not a pure strategy, then it's a mild-ESS.
- The existence of a ESS is not guaranteed.
  - Without an ESS, population will never settle on a single common strategy. Mutants always have a chance to survive. Ex: Rock-Paper-Scissors. If all play rock.

Evolutionary Game Theory: example none ESS

		Criminal 2		
		a_1	a_2	a_3
Criminal 1	a_1	0,0	-1,1	1,-1
	a_2	1,-1	0,0	-1,1
	a_3	-1,1	1,-1	0,0

In this game, the only NE is $\{\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\}$, but it is not an ESS, since, e.g. $\{\frac{1}{2}, \frac{1}{2}, 0\}$ can do as good.

Evolutionary Game Theory: dynamics

The direction of change of the current state $s = (p, 1-p)$ is entirely determined by the sign of the fitness difference between the two pure strategies:

$$\begin{aligned} d(s(p)) &= f\big((1,0),s\big) - f\big((0,1),s\big) = f\big((1,-1),s\big). \\ d(s(p)) = D(p) &= f((1,-1),s) = (1-p)(a_{12}-a_{22}) - p(a_{21}-a_{11}). \end{aligned}$$

- if $D(p) > 0$, then p increases
- if $D(p) < 0$, then p decreases
- $D(p)$ is a straight line with intercept ($p = 0$) at $a_{12} - a_{22}$ and maximum $p = 1$ at $a_{21} - a_{11}$.

Evolutionary Game Theory: dynamics

Example when $a_{12} - a_{22} > 0$ and $a_{21} - a_{11} > 0$:

		Criminal 2	
		Aggressive	Cautious
Criminal 1	Aggressive	$\frac{V-C}{2}, \frac{V-C}{2}$	$V, 0$
	Cautious	$0, V$	$\frac{V}{2}, \frac{V}{2}$

for the state strategy $s = (1, 0)$, $r = (0, 1)$ and $r_\epsilon = (1 - \epsilon)s + \epsilon r$

$$\phi\big((1,0),(1-\epsilon,\epsilon)\big) < \phi\big((0,1),(1-\epsilon,\epsilon)\big)$$

we find that strategy s is not an ESS. Since $a_{12} - a_{22} > 0$ and $a_{21} - a_{11} > 0$, then $p^* = \frac{V}{C}$ ESS.

Evolutionary Game Theory: dynamics

When $a_{12} - a_{22} < 0$ and $a_{21} - a_{11} < 0$:
$p^* = \frac{a_{12}-a_{22}}{a_{12}-a_{22}+a_{21}-a_{11}}$ is NE but not ESS. The pure strategies ($p = 0$ and $p = 1$) are ESS+NE. When $D(p) < 0$ (when $p < p^*$) then p decreases. When $D(p) > 0$ (when $p > p^*$) then p increases. So p^* is unstable (not an attractor) since p moves away from p^*.
Example, when $a_{12} - a_{22} < 0$ and $a_{21} - a_{11} < 0$, often called *symmetric coordination* games.

		Criminal 2	
		Cooperative	Individual
Criminal 1	Cooperative	5,5	-1,4
	Individual	4,-1	1,1

Evolutionary Game Theory: dynamics

When $|a_{12} - a_{22}| + |a_{21} - a_{11}| > 0$ or $(a_{12} - a_{22}) \cdot (a_{21} - a_{11}) \leq 0$:

- When $D(p) > 0 \ \forall p \in [0, 1]$, then p increases, so the NE and ESS is $p = 1$
- When $D(p) < 0 \ \forall p \in [0, 1]$, then p decreases, so the NE and ESS is $p = 0$

For example,

		Criminal 2	
		Cooperate	Defect
Criminal 1	Cooperative	1,1	-1,2
	Defect	2,-1	0,0

$D(p) < 0 \ \forall p \in [0, 1]$, then p decreases (until $p = 0$) so the NE and ESS is the strategy $p = 0$

Evolutionary Game Theory: Evolution of a Spite

Example: 4 Banks investing on two security policies

		Bank 2	
		Best	Standard
Bank 1	Best	-9,-9	-2,-7
	Standard	-7,-2	-1,-1

Standard policy ($p = 0$), to be ESS, must verify that a mutant bank playing *Best* ($q = 1$) will not succeed

$$\frac{2}{3}\phi(p,p) + \frac{1}{3}\phi(p,q) \not> \phi(q,p) \Rightarrow \frac{2}{3} - 1 + \frac{1}{3} - 7 \not> -2$$

However, the NE is not ESS. In this small population example , (*Best, Best*) is ESS.

Evolutionary Game Theory: Replicator Dynamics

The proportion of a population using a strategy increases (decreases) when it produces more (less) fitness over the average fitness of the population.

$$h_p^{t+1} = \frac{h_p^t \phi(p)}{h_p^t \phi(p) + (1 - h_P^t)\phi(q)}$$

Evolutionary Game Theory: Replicator Dynamics

		Criminal 2	
		Aggressive	Cautious
Criminal 1	Aggressive	$\frac{V-C}{2}, \frac{V-C}{2}$	$V, 0$
	Cautious	$0, V$	$\frac{V}{2}, \frac{V}{2}$

$$
\begin{aligned}
\phi(aggressive)^t &= h_p^t \cdot \frac{V-C}{2} + (1-h_p^t) \cdot V, \\
\phi(cautious)^t &= h_p^t \cdot 0 + (1-h_p^t) \cdot \frac{V}{2} \\
\bar{\phi}^t &= h_p^t \cdot \phi^t(aggressive) + (1-h_p^t) \cdot \phi^t(cautious) \\
h_p^{t+1} &= h_p^t \left(\frac{\phi^t(aggressive)}{\bar{\phi}^t} \right)
\end{aligned}
$$

Stable GT Decision on ARA

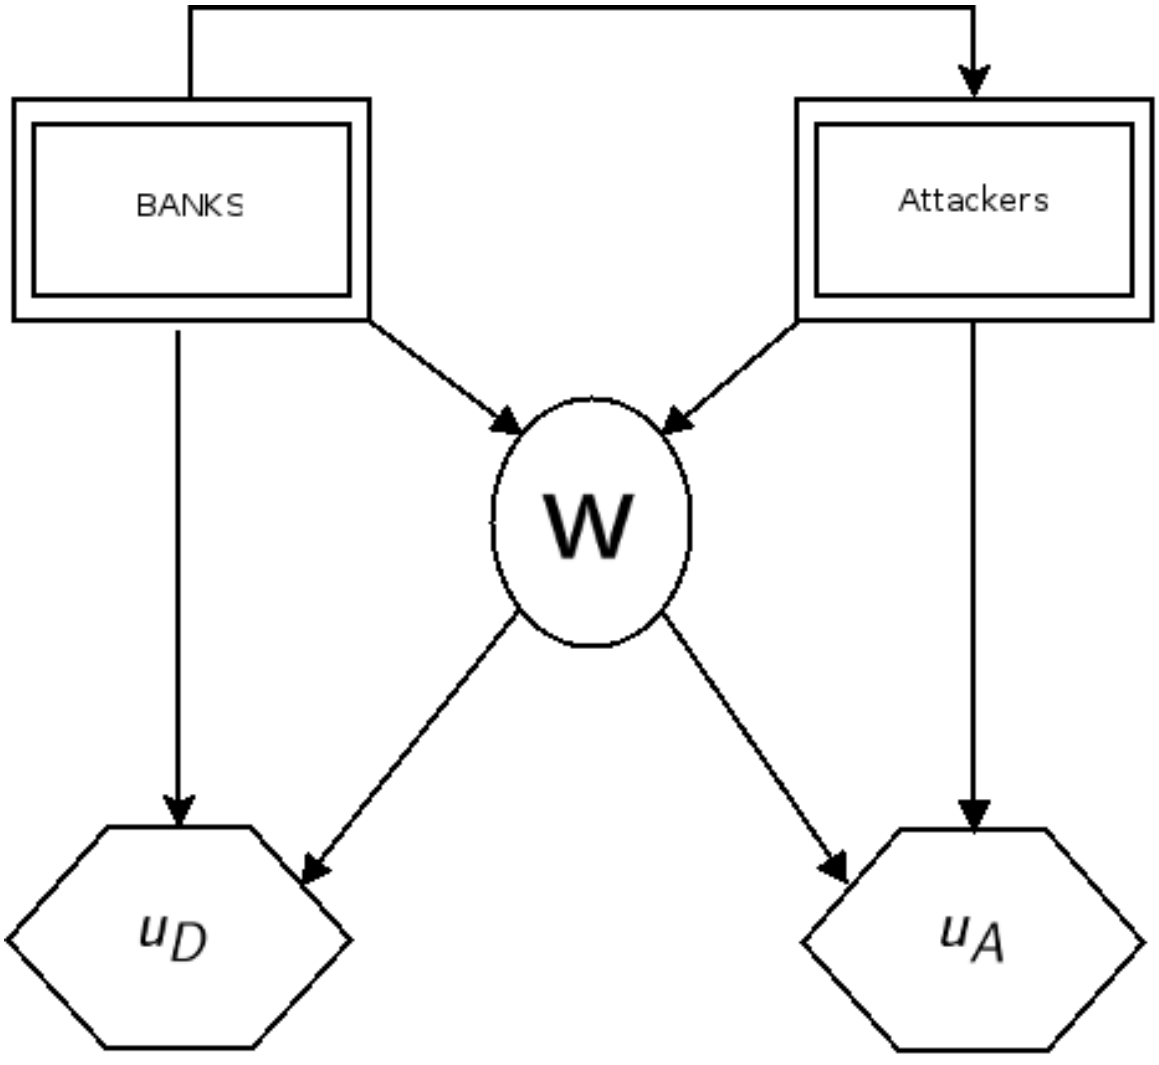

Stable GT Decision on ARA

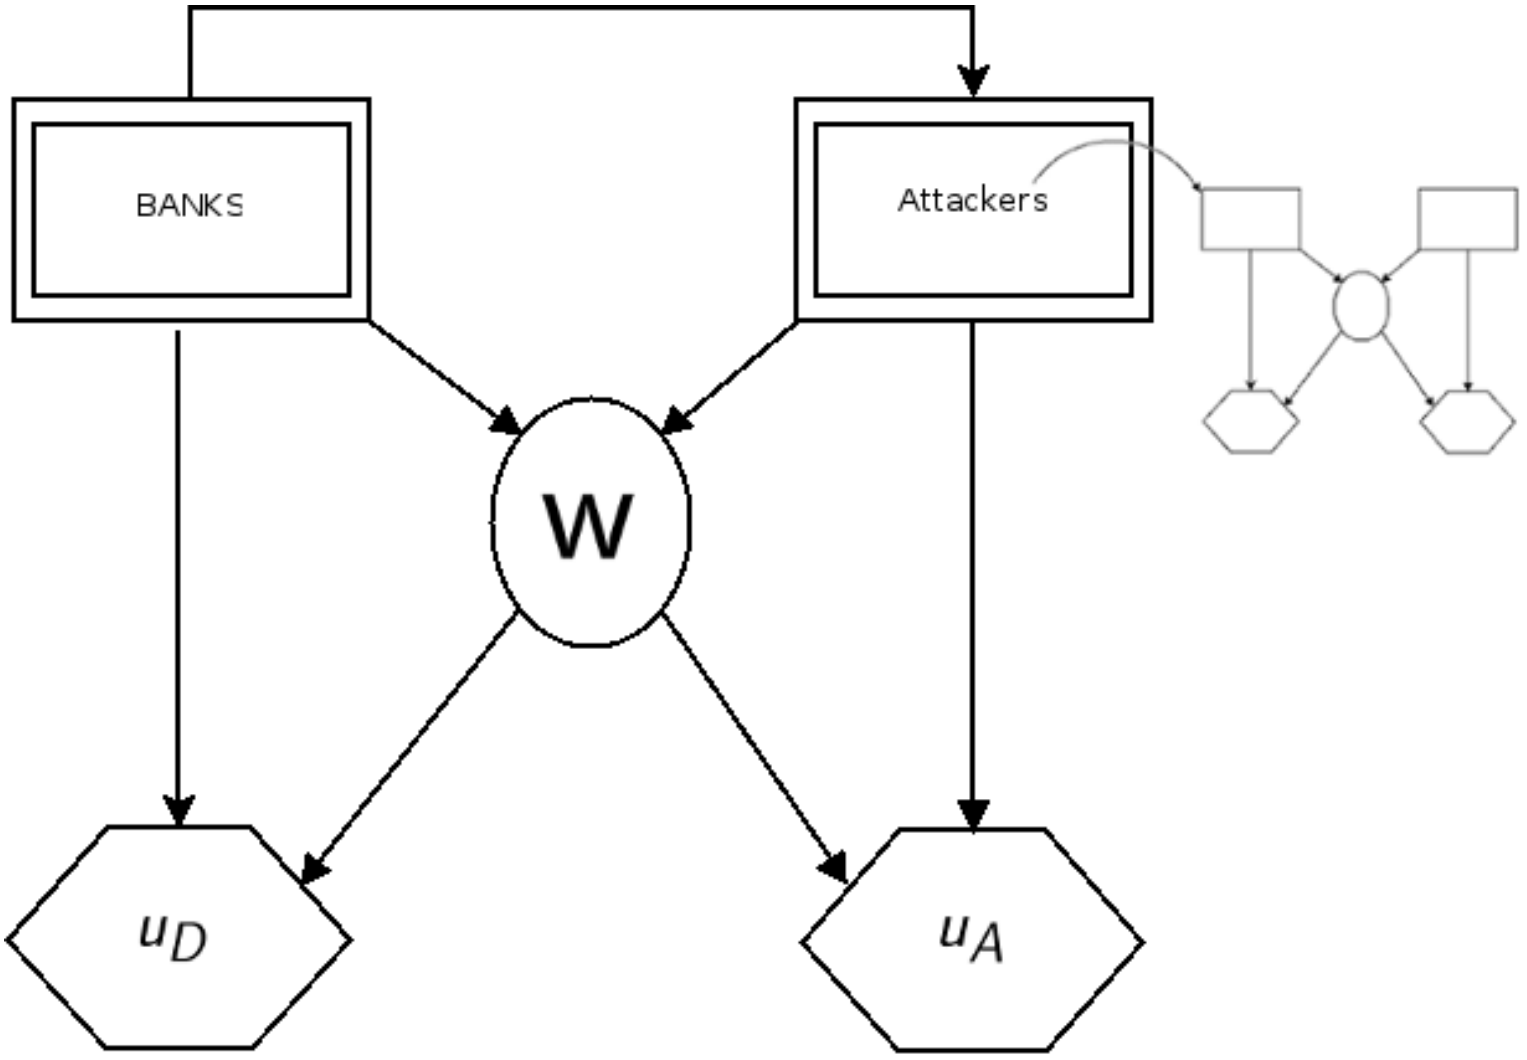

Stable GT Decision on ARA

- It provides from the defender's viewpoint an alternative more robust/stable decision making when assuming attackers taking stable (or evolutionary) strategies.
- Attackers adapt very fast to the conditions but maintaining the type of attack.
- Banks can 'affect' in different ways the evolution of populations of attackers.
- Trying to predict the probability of a mutant could be interesting, but here the focus is on:
  - is it profitable? then, will it be sustainable? how fast will that strategy grow?, ...

Thanks for your attention!