Adversarial Risk Analysis: The Somali pirates case

Abstract

Some of the current world's biggest problems revolve around security issues. This has raised recent interest in resource allocation models to manage security threats, from terrorism to organized crime through money laundering. One of those approaches is adversarial risk analysis, which aims at dealing with decision making problems with intelligent opponents and uncertain outcomes. We show here how such framework may cope with a current important security issue in relation with piracy in the Somali coasts. Specifically, we describe how to support the owner of a ship in managing risks from piracy in that area. We illustrate how a sequential defend-attack-defend model can be used to formulate this decision problem and solve it for the ship owner. Emphasis will be put on explaining how we can model the Pirates thinking in order to anticipate their behavior and how it would lead to a predictive probability distribution, from the defender's perspective, over what the Pirates may do.

1 Introduction

As described in Lomborg (2008), several of the world's biggest problems revolve around security issues. These include global terrorism, conflicts, arms proliferation, money laundering and organized crime. Indeed, multi-billion euro investments are being made to increase safety and security, which has stirred public debate about the convenience of some of such measures, specially in a shrinking economy. This has motivated a great deal of interest in modeling issues in relation with security resource allocation, with varied techniques and tools which may be seen in Gutfraind (2009), Wein (2009), Parnell et al. (2008), Enders and Sandler (2009) or Devlin and Lorden (2007). The key feature of these problems is the presence of two or more intelligent opponents who make interdependent decisions whose outcomes are uncertain. Many models and frameworks combining risk analysis, decision analysis and game theory have been developed. Merrick and Parnell (2011) provide a recent review of several of these approaches, commenting favorably on the Adversarial Risk Analysis (ARA) framework, introduced in Rios Insua et al (2009). ARA has a Bayesian game theoretic flavor. In supporting one of the participants, the problem is viewed as a decision analytic one, but principled procedures which employ the game theoretical structure, and other information available, are used to assess probabilities on the opponents' actions.

In Rios and Rios Insua (2011), the ARA framework is used to analyze three important stylized counterterrorism resource allocation models and illustrated in simple numerical examples. In this paper, we shall consider how the ARA framework may cope with a realistic example, referring to defending a boat which needs to travel through the Aden Gulf in the face of piracy risks. Starting from the early 90's, piracy has been a threat to international marine transportation and fishing boats around the coasts of Somalia[1]. Since 2005, several international organizations have expressed their worries about the increase in piracy acts. Somali pirates, originally dedicated to fishing, have traditionally claimed that the actual pirates are the foreign fishermen who loot their fish. Nowadays, no boat is safe within several hundred miles from the Somali coast. Piracy in Somalia may be explained in purely business terms, see Carney (2009), as there is actually a whole system supporting such acts: the elderly act, de facto, as a government; local businessmen and islamist groups provide funding. There is a clear organization as attacks are undertaken by small groups of around ten individuals in fast offshore boats which depart from a mother ship. Once successful, around fifty pirates remain in the boarded boat, with around fifty more pirates providing support from the coast. Pirates have learnt that ransom is more profitable than theft and they reinvest part of their earnings in equipment and training. We shall assume that we support a (large tuna fish) ship owner in deciding what defensive resources to implement and, if attacked, how to respond. Somali pirates demand a ransom in exchange of the kidnapped boat and crew. Our intention is twofold. On one hand, we illustrate the ARA framework on a realistic problem; on the other hand, the proposed case study may serve as a template for other future ARA applications.

The structure of the paper is as follows. In Section 2 we structure this problem using a sequential defend-attack-defend model and represent it through an asymmetric game tree. In Section 3 we model the Defender's own beliefs and preferences. In Section 4, we explain how the ARA framework can be used to solve the Defender's problem, with emphasis on how may we model the Defender's beliefs over whether the Pirates will attack the boat given the initial defenses adopted. We then calculate the optimal strategy for the Defender. We end up with some discussion, drawing lessons that may be used to develop templates for other ARA applications.

[1]See the wikipedia page on piracy in Somalia: http://en.wikipedia.org/wiki/Piracy_in_Somalia

2 The Somali Pirates case: Structure

We shall use the sequential defend-attack-defend model, see Brown et al. (2006), Parnell et al. (2010) or Rios and Rios Insua (2011), to formulate the Somali pirates case. We describe here how to support a ship owner in managing risks from piracy in the coast of Somalia. The ship owner will pro-actively decide on a defensive strategy to reduce piracy risks, ranging from different levels of deployed armed security to sailing through an alternative much longer route avoiding such area. The Pirates will respond to the Defender's move by launching (or not) an attack with the intention of taking over the ship and asking for a ransom. If the Pirates' operation is successful, the ship owner will have to decide on paying or not the ramson, or even sending armed forces to release the ship. The boat is assumed to be of Spanish ownership for the economical computations

Specifically, we shall assume that the ship owner (the Defender, she) initially decides on one of the following four alternative defense actions (elements of $\mathcal{D}_1$):

d_1^1 : Do nothing, i.e. no defensive action is taken,

d_1^2 : Use private protection with an armed person,

d_1^3 : Use private protection with a team of two armed persons,

d_1^4 : Do the trip through the Cape of Good Hope, rather than the Suez Channel, thus avoiding the Somali coast and an eventual kidnapping.

Once the Defender has made her initial decision, the pirates (the Attacker, he) observe this and decide whether attacking ($a^1 \in \mathcal{A}$) or not ($a^0 \in \mathcal{A}$) the Defender's boat. The attack results in either the boat kidnapped ($S = 1$) or not ($S = 0$) by the pirates, with likelihoods depending on the initial defense action chosen by the boat's owner. If $S = 1$, the Defender has the option of responding by either (elements of $\mathcal{D}_2$):

d_2^1 : Doing nothing, i.e. not responding to the pirates' demands, assuming all entailed costs.

d_2^2 : Pay the amount demanded by the pirates, thus recovering the boat and the crew.

d_2^3 : Ask the Navy for support to release the boat and crew.

The asymmetric game tree shown in Figure 1 represents the sequence of decisions and events faced by Defender and Pirates in this case, where nodes D_1 and D_2 correspond to the Defender's first and second decisions, respectively, node A represents the Attacker's decision, and chance node S represents the outcome of the attack.

Figure 1: Game tree for the Somali pirates case

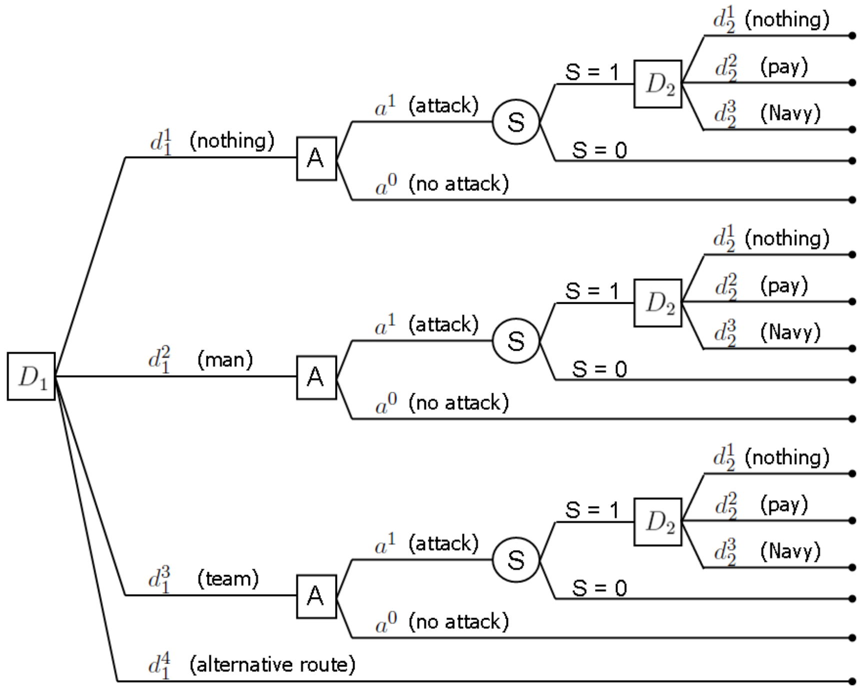

3 The Defender's own preferences and beliefs

In supporting the Defender, her decision problem is seen as the decision tree in Figure 2, in which the Attacker's decision node [A] has been replaced by an uncertain node (A). This reflects that the Attacker's decision is seen as an uncertainty from the Defender's viewpoint, with the bulk of the modeling work consisting of assessing her probabilities over such node. Thus, to solve her decision problem, she needs to assess $p_D(A \mid d_1)$, her predictive probability of an attack against each $d_1 \in \mathcal{D}_1$, in addition to the (more standard) assessments $p_D(S \mid d_1, a^1)$ and $u_D(c_D)$, with c_D representing her monetary cost equivalent of the multi-attribute consequences associated with each leave of the tree. We specify these assessments, starting with her own preferences and beliefs.

The relevant consequences for the Defender in this problem are

- The loss of the boat
- The costs (in millions of euros) associated with defending and responding to an eventual

Figure 2: Decision tree representing the Defender's decision problem for the Somali pirates case

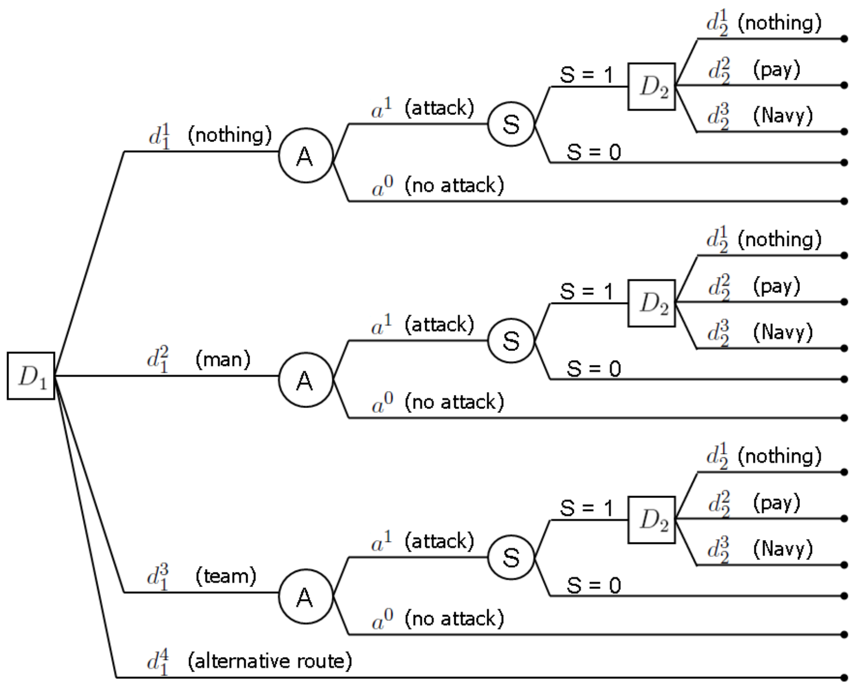

attack

- The number of deaths on her side

As far as the costs associated with implementing her protective and response actions against an eventual kidnapping are concerned, we have that for the defensive actions in $\mathcal{D}_1$ these costs are:

- 0 euros, if she chooses to do nothing, d_1^1,
- 0.05M euros, if she chooses to use one armed person, d_1^2. This cost corresponds to the salary of the armed person for 6 months, plus equipment.
- 0.15M euros, if she chooses an armed team, d_1^3. This cost corresponds to the salary of two armed persons with better equipment for 6 months.
- 0.5M euros, if she chooses d_1^4, going through the Good Hope Cape. This cost is consequence of the longer distance of the trip, notwithstanding eventualities due to bad weather.

The costs associated with the defense actions in $\mathcal{D}_2$ are:

- 0 euros, for option d_2^1, doing nothing.
- 2.3M euros, for option d_2^2, paying the ransom. We have estimated it through the average on the last ransoms paid, which have been growing higher over time.

- 0.2M euros, for option d_2^3, calling for the Navy. This estimation is based on a military intervention using the Spanish Navy boats already deployed there.

As far as human lives are concerned, we consider that if the boat is attacked and the attack is aborted ($S = 0$), there are no lives lost. If the attack is successful ($S = 1$), we assume that the armed Defenders have died and that, depending of the chosen response at D_2, there might be additional lives lost, specifically:

- If the response to the kidnapping is doing nothing, d_2^1, the kidnappers might kill part of the crew as a warning for future kidnappings. We estimate this to be 4 crew members.
- If the ship owner decides to pay the ransom, d_2^2, there will be no human losses.
- If the kidnapped boat is rescued by the Navy, d_2^3, we estimate that there might be 2 casualties due to both collateral damage during the intervention and/or because the pirates feel threatened and kill some of the crew during the operation.

We quantify now economically the value of the ship and a human life in monetary terms. The type of boats operating in that area focus on tuna fishing, with a length between 80 and 110m and lots of technology built in. A new boat of this characteristics costs between 9M and 12M euros. Assuming some depreciation because of time, we shall assume that the incumbent boat is valued in 7M euro. As far as quantifying the value of a human life, we shall use the concept of statistical value of a life, see Martinez et al (2009), which, according to Riera et al (2007), was estimated in 2.04 M euros for a Spanish person.

Table 1 summarizes the estimated consequences and aggregated monetary costs c_D for the Defender associated with each scenario consisting of a path in the tree shown in Figure 1. Note that if there is no attack ($a = a^0$), then necessarily we have that $S = 0$.

D_1	S	D_2	Boat loss	Action costs	Lives lost	c_D
d_1^1 (nothing)	$S=1$	d_2^1 (nothing)	1	0 + 0	0 + 4	15.16
d_1^1 (nothing)	$S=1$	d_2^2 (pay)	0	0 + 2.3M	0 + 0	2.3
d_1^1 (nothing)	$S=1$	d_2^3 (Navy)	0	0 + 0.2M	0 + 2	4.28
d_1^1 (nothing)	$S=0$		0	0	0	0
d_1^2 (man)	$S=1$	d_2^1 (nothing)	1	0.05M + 0	1 + 4	17.25
d_1^2 (man)	$S=1$	d_2^2 (pay)	0	0.05M + 2.3M	1 + 0	4.39
d_1^2 (man)	$S=1$	d_2^3 (Navy)	0	0.05M + 0.2M	1 + 2	6.37
d_1^2 (man)	$S=0$		0	0.05M	0	0.05
d_1^3 (team)	$S=1$	d_2^1 (nothing)	1	0.15M + 0	2 + 4	19.39
d_1^3 (team)	$S=1$	d_2^2 (pay)	0	0.15M + 2.3M	2 + 0	6.53
d_1^3 (team)	$S=1$	d_2^3 (Navy)	0	0.15M + 0.2M	2 + 2	8.51
d_1^3 (team)	$S=0$		0	0.15M	0	0.15
d_1^4 (alternative route)			0	0.5M	0	0.5

Table 1: Consequences of various tree paths for the Defender

We shall assume that the Defender is constant risk averse with respect to monetary costs. Thus, her utility function is (strategically equivalent to) $u_D(c_D) = -\exp(c \times c_D)$, with $c > 0$. We shall study what happens when $c \in \{0.1, 0.4, 2\}$ as a way to perform sensitivity analysis.

Based on information from Carney (2009), we shall assume that the Defender's beliefs about an attack being successful conditional on her initial defense maneuver are

- $p_D(S = 1|a^1, d_1^1) = 0.40$, for the case in which no initial defensive action is taken,
- $p_D(S = 1|a^1, d_1^2) = 0.10$, for the case in which she uses private protection with an armed person, and
- $p_D(S = 1|a^1, d_1^3) = 0.05$, for the case in which she uses private protection with two armed persons.

4 The Defender's beliefs over the Attacker's beliefs and preferences

We describe now how the Defender assesses the probability of being attacked, given her implemented initial defense. Carney (2009) estimates that the probability of being attacked is 0.005

based on historical data on Piracy in the cost of Somalia. However, this estimate does not take into account that some boats may be more desirable than others for the Pirates, which typically use observers and informers to decide upon their targets. Our objective is to assess a predictive probability of attack $p_D(A = a^1|d_1)$ for the type of boat owned by the Defender conditional on each possible initial protective defense $d_1 \in \mathcal{D}_1 \setminus \{d_1^4\}$ taken. To do so, alternatively to Carney (2009), we assume the Pirates behave as an expected utility maximizer and derive the Defender's uncertainty about the Pirates' decision from her uncertainty about the Pirates' probabilities and utilities. Thus, the Defender must analyze the decision problem faced by the Pirates from her perspective, as shown in Figure 3. Note that the set of alternatives for the Attacker has been expanded to include alternatives $a^i \in \mathcal{A}$, for $i = 2, \ldots, n$, representing the Pirates' option to attack other boats that are not owned by our Defender. We have also added new chance nodes D_2 at the end of the tree paths starting at a^i, representing the response of boat $i = 2, \ldots, n$ to an eventual kidnapping, which is considered an uncertainty from the perspective of the Attacker. This analysis of the Pirates decision must lead to the probabilistic assessment of the perceived Pirates' preferences (the uncertainty of the Defender over the Attacker's preferences, modeled through a random variable $U_A(a, s, d_2)$, where now $a \in \mathcal{A} = \{a^0, a^1, \ldots, a^n\}$) and beliefs (the uncertainty of the Defender over the Attacker's beliefs, modeled through the random variables $P_A(S = 1 \mid a^1, d_1)$ and $P_A(D_2 \mid d_1, a^1, S = 1)$ as well as $P_A(S = 1 \mid a^i)$ and $P_A(D_2 \mid a^i, S = 1)$ for $i = 2, \ldots, n$).

Figure 3: Decision tree representing the perceived decision problem of the Somali pirates

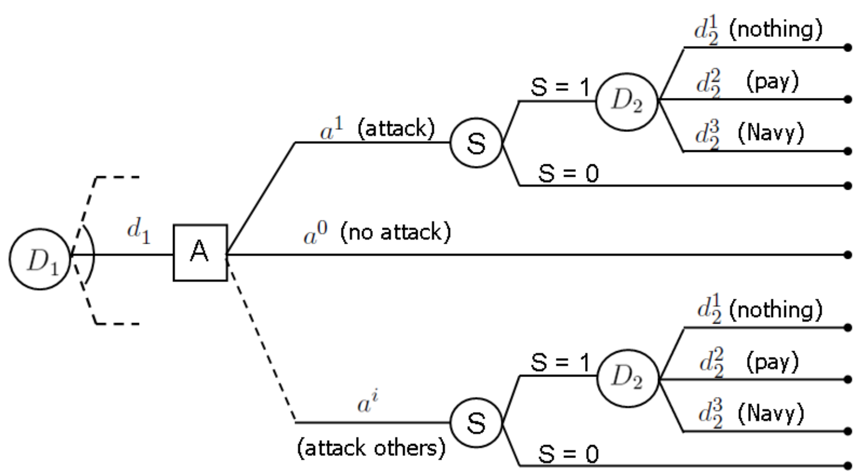

The Defender considers that the relevant consequences for the Pirates are:

- Whether they keep the boat or not.

- The amount of money earned.
- The number of pirates' lives lost.

As described in Carney (2009), the estimated average cost of an attack operation is around 30000 euros (0.03 M euros). The eventual benefits if the Defender pays the ransom are the above mentioned 2.3 M euros. As human lives are concerned, we shall assume that two pirates are dead if the attack is repelled and, if successful, no pirates' lives are lost in the attack. However, if the Defender responds by sending the Navy, we shall assume five pirates will be killed.

If the pirates keep the boat, it will not have the same monetary value as for the Defender. They may use the machinery or the technological instruments, or they could use it as a mother boat. We shall assess its economic value for the pirates as 1M euros. We shall assume that they put a value equivalent of 0.25 M euros to a pirate's life. Table 2 summarizes the consequences for the Pirates for various attack scenarios, including the aggregate monetary equivalent c_A in the last column. Note that we have assumed that for the Pirates there are no differences between the consequences of attacking our Defender's boat (a^1) and those from attacking others boats ($a^i, i = 2, \ldots, n$). Thus, for $i = 1, \ldots, n$

A	S	D_2	Boat kept	Profit	Lives lost	c_A
a^0 (no attack)			0	0	0	0
a^i (attack)	$S = 1$	d_2^1 (nothing)	1	-0.03M	0	0.97
a^i (attack)	$S = 1$	d_2^2 (pay rescue)	0	2.27M	0	2.27
a^i (attack)	$S = 1$	d_2^3 (Navy sent)	0	-0.03M	5	-1.28
a^i (attack)	$S = 0$		0	-0.03M	2	-0.53

Table 2: Consequences for the Pirates of various tree paths of their decision problem

At a qualitative level, the Defender thinks the Pirates are risk prone over money, specifically she assumes the Pirates have increasing constant risk prone preferences for money and, therefore, she uses a utility function (strategically equivalent to) $u_A(c_A) = \exp(c \times c_A)$, with $c > 0$ to model the pirates' preferences and risk attitude. However, she is not sure which c determines the Pirates's utility function, but she thinks that $c \sim \mathcal{U}(0, 20)$. This uncertainty over c induces the uncertainty over u_A to provide U_A.

As the Pirates most likely have access to the same information than the Defender, she assesses the following probabilities for the Pirates' beliefs over an attack on her boat being successful, conditional on her initial defense move.

- $P_A(S=1|a^1, d_1^1) \sim \mathcal{B}e(40,60)$, for no defensive action taken

- $P_A(S=1|a^1, d_1^2) \sim \mathcal{B}e(10,90)$, for private protection with an armed person,

- $P_A(S=1|a^1, d_1^3) \sim \mathcal{B}e(50,950)$, for private protection with a team of two armed persons

Note how the expected value of each of these distributions corresponds with the estimated probabilistic beliefs of the Defender on the same uncertainty. Likewise, acknowledging lack of information, the Defender assesses that the probabilities representing the Pirates' beliefs of a successful attack on boats $i = 2, \ldots, n$ are:

- $P_A(S=1|a^i) \sim \mathcal{B}e(1,1)$, for boats $i = 2, \ldots, n$.

Now, the Defender holistically assesses how the Pirates think she will respond to a successful attack. Thus, we will obtain her beliefs on the Pirates' probabilistic beliefs $p_A(D_2 \mid d_1, A = a^1, S = 1)$. Specifically, the Defender thinks that the Pirates expects her to respond along the same lines of the defense chosen by her at the first stage. Thus, a tough deterring defense at her first move is expected to produce an eventual response of similar harshness. Specifically, the Defender assesses the following Dirichlet distributions over d_2^1 (doing nothing), d_2^2 (pay), and d_2^3 (Navy) $\in \mathcal{D}_2$:

- $P_A(D_2 \mid d_1^1, A = a^1, S = 1) \sim Dir(1,1,1)$: If $d_1 = d_1^1$ (doing nothing) and the boat is kidnapped, any response in this case is perceived equally likely by the Attacker.

- $P_A(D_2 \mid d_1^2, A = a^1, S = 1) \sim Dir(0.1,4,6)$: If $d_1 = d_1^2$ (protect with an armed man) and the boat is kidnapped, it is perceived that the Attacker expects the Defender to respond doing something, with sending the Navy more likely than paying the ransom.

- $P_A(D_2 \mid d_1^3, A = a^1, S = 1) \sim Dir(0.1,1,10)$: If $d_1 = d_1^3$ (protect with an armed team) and the boat is kidnapped, it is perceived that it would be even more likely for the Attacker to believe that the Defender will respond sending the Navy.

Finally, the Defender assesses that $P_A(D_2 \mid a^i, S = 1) \sim Dir(1,1,1)$ for $i = 2, \ldots, n$, suggesting lack of information.

Based on her above assessments, the Defender may solve the perceived Pirates' decision problem using backward induction over the decision tree in Figure 3, propagating the uncertainty of her assessed random preferences and beliefs of the Pirates as follows:

- Compute the random expected utilities associated with the Pirates choosing a^1 conditional on each of her initial protective defenses $d_1 \in \mathcal{D}_1 \setminus \{d_1^4\}$

$$\Psi_A(d_1, a^1) = P_A(S = 1 \mid d_1, a^1) \sum_{d_2 \in \mathcal{D}_2} U_A(a^1, S = 1, d_2)\ P_A(D_2 = d_2 \mid d_1, a^1, S = 1) + \\ P_A(S = 0 \mid d_1, a^1)\ U_A(a^1, S = 0)$$

- Compute the random expected utilities associated with the Pirates choosing a^i for $i = 2, \ldots, n$

$$\Psi_A(a^i) = P_A(S = 1 \mid a^i) \sum_{d_2 \in \mathcal{D}_2} U_A(a^i, S = 1, d_2)\ P_A(D_2 = d_2 \mid a^i, S = 1) + \\ P_A(S = 0 \mid a^i)\ \ U_A(a^i, S = 0)$$

- Compute the Defender's predictive probabilities of being attacked $(A = a^1)$ conditional on each of her initial defenses $d_1 \in \mathcal{D}_1 \setminus \{d_1^4\}$

$$p_D(A = a^1 \mid d_1) = \Pr(\ \Psi_A(d_1, a^1) > \max\{U_A(a^0),\ \Psi_A(a^2), \ldots, \Psi_A(a^n)\}\ )$$

This probabilities can be approximated by Monte Carlo simulation by drawing a sample from the Pirates' random utilities and probabilities assessed by the Defender and solving for each drawn the Pirates' decision problem as before. This generates a sample of when a^1 is the optimal decision for the Pirates, and then we approximate $p_D(A = a^1 \mid d_1)$ by

$$\frac{\#\{1 \leq k \leq N : \psi_A^k(d_1, a^1) > \max\{u_A^k(a^0),\ \psi_A^k(a^2), \ldots, \Psi_A^k(a^n)\}\}}{N},$$

with N the sample size.

For illustrative purposes, let us assume that $n = 4$: there will be other 3 boats (of similar characteristics) exposed to the risk of being seized by the pirates at the time period in which the Defender's boat will sail through the Gulf of Aden. Based on 1000 Monte Carlo iterations, we get the following estimates for the probability of a Defender's boat being attacked given that she chooses as initial defense action $d_1 \in \mathcal{D}_1 \setminus \{d_1^4\}$

- $\hat{p}_D(A = a^1 \mid d_1^1\ ) = 0.1931$, for the case in which the Defender does not take any protective action initially,

- $\hat{p}_D(A = a^1 \mid d_1^2) = 0.0181$, for the case in which she uses private protection with an armed person,

- $\hat{p}_D(A = a^1 \mid d_1^3) = 0.0002$, for the case in which she uses private protection with a team of two armed persons.

Note that the probability of attacking the Defender's boat gets smaller if her boat is protected with an armed man and even smaller if she choose to protect it with an armed team. We can also check that the probability of attacking the Defender's boat would decrease if n gets bigger. We could also analyze the impact of differences in the vulnerabilities among different boats if we elicit different $P_A(S = 1|a^i)$ for each type of boat, $i = 2, \dots, n$. But we have assumed that they are all the same for this case.

5 Finding the optimal defense strategy

We now have all the inputs needed to solve the decision problem for the Defender. Given these, the Defender can solve her decision problem working backwards the tree in Figure 2. At node D_2, she can compute her maximum utility action conditional on each $d_1 \in \mathcal{D}_1 \setminus \{d_1^4\}$

$$d_2^*(d_1, a^1, S = 1) = \text{argmax}_{d_2 \in \mathcal{D}_2} u_D(c_D(d_1, S = 1, d_2)).$$

Afterwards, she will obtain at node S her expected utilities

$$\begin{aligned}\psi_D(d_1, a^1) = & p_D(S = 1 \mid d_1, a^1)\ u_D(c_D(d_1, S = 1, d_2^*(d_1, a^1, S = 1))) + \\ & p_D(S = 0 \mid d_1, a^1)\ u_D(c_D(d_1, S = 0)).\end{aligned}$$

At this point, she will use her probabilistic assessments of being attacked conditional on her initial defense moves, $\hat{p}_D(A = a^1 \mid d_1)$, to compute for each $d_1 \in \mathcal{D}_1 \setminus \{d_1^4\}$ her expected utility at node A

$$\psi_D(d_1) = \psi_D(d_1, a^1)\ \hat{p}_D(A = a^1 \mid d_1)\ +\ \psi_D(d_1, S = 0)\ (1 - \hat{p}_D(A = a^1 \mid d_1)).$$

Finally, she can find her maximum expected utility decision at node D_1

$$d_1^* = \text{argmax}_{d_1 \in \mathcal{D}_1} \psi_D(d_1),$$

where $\psi_D(d_1^4) = u_D(c_D(d_1^4))$, obtained from Table 1. The Defender's best strategy is then to first choose d_1^* at node D_1, and, if the the boat is successfully attacked, respond by choosing $d_2^*(d_1^*, a^1, S = 1)$ at node D_2 .

We obtained that the defense strategies of maximum expected utility, for each of the considered risk aversion coefficients determining her utility function, are:

- $c = 0.1$: protect with an armed man ($d_1^* = d_1^2$), and if kidnapped ($S = 1$), pay the ransom ($d_2^* = d_2^2$).
- $c = 0.4$: protect with an armed man ($d_1^* = d_1^2$), and if kidnapped ($S = 1$), pay the ransom ($d_2^* = d_2^2$).
- $c = 2$: avoid the Somali cost by going through Good Hope Cape ($d_1^* = d_1^4$).

We see that choosing to *go through the GH Cape* emerge as optimal decision when the risk aversion coefficient of the Defender is $c = 2$, that is when the Defender becomes more risk averse. This suggests that some security measures are required, but if the Defender is too risk averse it is better for him to change the route. The optimal Defender's response action to a kidnapping is always paying the ransom. This is possibly because it is the option that allows the Defender to keep the boat and minimize the lost of lives. Clearly, this neglects the political implications of this action, but recall that we are dealing with this problem from the boat owner perspective. One possibility to acknowledge such fact would be to add some extra cost if the ransom is actually paid, reflecting the negative political implication associated with it.

6 Conclusions

We have described an application of the Adversarial Risk Analysis framework to solve the Somali pirates case using the Defend-Attack-Defend model. The program produced to solve this case may serve as a template for decision makers facing this same problem, by just modifying the value of their boats, the value of their national lives, the estimated ransom, their risk aversion coefficient and other relevant inputs.

This analysis may also serve as template for other security resource allocation problems adaptable to the sequential Defend-Attack-Defend model. Essentially, the Defender would:

1. Build a game tree as in Figure 1.
2. Assess her consequences as in Table 1 and the Attacker's consequences as in Table 2.
3. Assess her utility function, her vulnerabilities and her beliefs over successful attacks as in Section 3.
4. Assess her beliefs over the preferences and beliefs of the Attacker as in Section 4.
5. (Simulate to) estimate her probabilities of various attacks as in Section 4.

6. Solve her decision tree to find her optimal defenses as in Section 5 and perform a sensitivity analysis.

In some instances the consequences in step 2 might be random and described through random variables [2]. This would require one further random node in the decision tree solved in Step 6 and one further loop in the MC simulation addressed in Step 5. Other problems might entail continuous defenses and/or attacks. That would typically entail solving nonlinear programming problems at decision nodes and discrete approximations at random nodes. But the methodology would remain essentially the same.

The assessment of the probability of being attacked $p_D(A = a^1 \mid d_1)$ is straightforward as far as the Defender is able to assess the random probabilities $P_A(D_2 \mid d_1, a^1, S = 1)$ and $P_A(D_2 \mid a^i, S = 1)$ for $i = 2, \ldots, n$. However, these assessments could be problematic, as the Defender may want to exploit information available to her about how the Pirates analyze how she and other boat owners would respond to an eventual successful attack. This potentially may lead to an infinite regress, as described in Rios and Rios Insua (2011). However, realistically, we argue that this recursive analysis will always stop at some point when the Defender has no more information that can be accommodated into the analysis. At this point, the Defender would use a noninformative distribution, as we have done for some of these assessments in our example.

Acknowledgements

We are grateful to SAMSI-NSF for providing the initial support for this research through its *Risk Analysis, Extreme Event and Decision Theory* program and the COST-ESF ALGODEC program for its later stages. DRI is grateful to the MICINN e-COLABORA and the RIESGOS-CM programs. Discussions with David Banks are gratefully acknowledged.

References

Brown, G., M. Carlyle, J. Salmeron, and K. Wood (2006) Defending critical infrastructure, *Interfaces*, 36, 6, 530-544.

Carney, S. (2009) An Economic Analysis of the Somali Pirate Business Model, *Wired.*

Devlin, K. and G. Lorden (2007) *Solving Crime with Mathematics*, Plume.

[2]Indeed, this would be the case in the example in this paper, but we have preferred to simplify this for clarity and because of space limits.

Enders, W. and T. Sandler (2009) *The Political Economy of Terrorism*, Cambridge U.P.

Gutfraind, A. (2009) Terrorism as a mathematical problem, *SIAM News*, 10, 12.

Lomborg, B. (2008) *Solutions for the World's Biggest Problems*, Cambridge U.P.

Martinez, J.E., and Mendez, I. (2009) Que podemos saber sobre el Valor Estadistico de la Vida en Espana utilizando datos laborales?, *Hacienda Publica Espanola*, 190, 73-93.

Merrick, J. and G. Parnell (2011) A comparative analysis of PRA and intelligent adversary methods for counterterrorism management, to appear in *Risk Analysis*.

Parnell, G., D. Banks, L. Borio, G. Brown, L. A. Cox, J. Gannon, E. Harvill, H. Kunreuther, S. Morse, M. Pappaioanou, S. Pollack, N. Singpurwalla, and A. Wilson (2008). *Report on Methodological Improvements to the Department of Homeland Security's Biological Agent Risk Analysis*, National Academies Press.

Parnell, G., Smith, C. and D. Moxley (2010) Intelligent adversary risk analysis: a bioterrorism risk management model, *Risk Analysis*, 30, 1, 32-48.

Riera, A., Ripoll, A.M. and J. Mateu (2007), Estimacion del valor estadistico de la vida en Espana: Una aplicacion del Metodo de Salarios Hedonicos, *Hacienda Publica Espanola*, 181, 29-48.

Rios, J. and D. Rios Insua (2011) Adversarial risk analysis for counterterrorism modeling, *submitted*.

Rios Insua, D., J. Rios and D. Banks (2009) Adversarial risk analysis, *Journal of the American Statistical Association*, 104, 486, 841-854.

Wein, L. (2009) Homeland security: from mathematical models to policy implementation, *Operations Research*, 57, 801-811.