*Proceedings of the $12^{th}$ International Defense and Homeland Security Simulation Workshop (DHSS), 002*
$18^{th}$ International Multidisciplinary Modeling & Simulation Multiconference

2724-0363 © 2022 The Authors.
doi: 10.46354/i3m.2022.dhss.002

# Augmented probability simulation for adversarial risk analysis in general security games

Roi Naveiro[1,*], David Rios Insua[1] and Jose Manuel Camacho[1]

[1]Institute of Mathematical Sciences (ICMAT-CSIC), C/ Nicolás Cabrera 13-15, Madrid, 28049, Spain

[*]Corresponding author. Email address: roi.naveiro@icmat.es

## Abstract

Although widely used, standard game-theoretic approaches to security games face severe shortcomings, being the common knowledge assumption a critical one. Adversarial Risk Analysis (ARA) is an alternative modeling framework that mitigates such limitations. However, from a computational perspective, ARA is much more involved than its game theoretical counterparts. We propose an approach for finding ARA solutions to security games represented as bi-agent influence diagrams that is based on augmented probability simulation. We motivate this approach using two simple cases: sequential and simultaneous defend-attack models. We next provide the general framework and illustrate it in handling risks in a cybersecurity setting.

**Keywords**: Security games; adversarial risk analysis; augmented probability simulation

## 1. Introduction

Security games provide a flexible and strong modeling framework for strategic and operational defense and homeland security problems, as argued in Brown et al. (2006), Zhuang and Bier (2007), Brown et al. (2008) or Hausken (2011). These authors illustrate the analysis of security games from a standard game theoretic perspective based on approximating Nash equilibria and related refinements.

The common knowledge assumptions underlying such game theoretic approaches, critically reviewed in e.g. Raiffa et al. (2002) or Hargreaves-Heap and Varoufakis (2004), constitute a shortcoming in the security domain. In most games, it is assumed that agents know not only their own payoffs, preferences, beliefs and possible actions, but also those of their opponents. In games of incomplete information (Harsanyi, 1967), it is typically assumed that each agent has a joint probability distribution over their opponents' types which is known by all the players. Such assumptions allow for a symmetric joint normative analysis in which players maximise their expected utilities, and expect other players to do the same. However, in defense and homeland security, agents will not generally have so much knowledge about their opponents as players try to hide information. Adversarial Risk Analysis (ARA), Rios Insua et al. (2009), is an alternative modeling framework which mitigates these common knowledge assumptions. Rather than addressing the problem simultaneously for all agents, ARA provides prescriptive support to one of the decision makers, which we shall refer to as the defender (she), seeking for actions that maximize her expected utility: Nash equilibria notions are abandoned and the defender's problem is viewed as a decision analytic one instead. However, procedures which employ the game theoretic structure are used to estimate the probabilities of the opponent's (referred to as the attacker, he) actions. ARA thus make operational the Bayesian approach to games, as sketched in Kadane and Larkey (1982) or Raiffa (1982).

A main motivation for ARA comes from security and counter-terrorism analysis. Since its introduction, it has been used to model a variety of problems such as network routing for insurgency (Wang and Banks, 2011), international piracy (Sevillano et al., 2012), urban security resource allocation (Gil et al., 2016), adversarial classification (Naveiro et al., 2019), counter-terrorist online surveillance (Gil and Parra-Arnau, 2019), cyber-security (Rios Insua et al., 2019), insider threat modelling (Joshi et al., 2020) and combat modeling enhancement (Roponen and Salo, 2015), to name but a few.

However, a major problem with the implementation of ARA refers to its computational issues since it essentially entails a two-stage approach in which we first simulate the attacker's problem to compute a probabilistic forecast of the adversary's actions and then use that forecast to optimize the defender's problem. In Ekin et al. (2022), augmented probability simulation (APS, Bielza et al. (1999)) was explored as a solution method for simple sequential defend-attack games from an ARA perspective. In particular, it was shown that APS could be more efficient than standard Monte Carlo based techniques in situations in which the cardinality of the decision sets of the agents intervening in the security game is very big.

We present here how APS may be used to solve general security games described through bi-agent influence diagrams (BAIDs) (Banks et al., 2015) from an ARA perspective. For this, we first provide a brief introduction to APS (Section 2), and explain how this solution approach can be utilized for solving sequential defend-attack (Section 3) and simultaneous defend-attack (Section 4) games. Section 5 sketches the analysis for general BAIDs. We end up with an experiment illustrating the proposed approach (Section 6) and a brief discussion (Section 7).

Our aim is thus to support the defender in her decision making. For this, we need to forecast the attacker's intentions. Many different attacker rationalities may be considered in the ARA framework, see Rios Insua et al. (2016). However, we circumscribe here the defender to be a level-2 thinker, in the Stahl and Wilson (1995) sense: she will ponder how the attacker's strategy would adapt to her own strategy but presume that he will not conduct likewise. Thus, we assume that the attacker is an expected utility maximiser. We could predict his actions by finding his maximum expected utility policy; however, the uncertainty in our assessments about the attacker's probabilities and utilities propagates to his optimal decision that becomes random, providing us with the required probabilistic forecast over the attacks.

## 2. Augmented Probability Simulation

In decision analysis settings, the goal is usually to find an action or set of actions that maximize the expected utility of the decision maker. In general, expected utility cannot be computed analytically and needs to be approximated. Most solution methods proceed by first estimating the expected utility for every possible action (usually through Monte Carlo integration) and then optimizing such estimate. Instead, APS solves for expected utility maximization by converting the tasks of estimation and optimization into simulation from an augmented distribution over the joint space of decisions and outcomes. Suppose that we aim to maximize the expected utility $\psi(d) = \int u(d, \theta) \cdot p(\theta | d) \, d\theta$, where $\theta$ is a random outcome whose distribution is $p(\theta|d)$ and $u(d, \theta)$ is the utility perceived when choosing $d$ and obtaining outcome $\theta$. APS stems from the observation that if the utility is non-negative and integrable, we can define the augmented distribution $\pi(d, \theta) \propto u(d, \theta) \cdot p(\theta | d)$ on the augmented space of decisions and outcomes. Then, under mild conditions, the mode of the marginal augmented distribution in $d$ coincides with the maximum expected utility solution. This suggests a strategy to compute the optimal decision based on simulating from the augmented probability distribution (for which Markov chain Monte Carlo simulation methods French and Insua (2000) are instrumental) and, then, assessing the marginal mode using mode estimation methods Chacon (2020).

## 3. Sequential Defend-Attack games

### 3.1. Basic template

Drawing on Ekin et al. (2022), this section sketches how APS-based methods can be utilized for finding optimal defenses in sequential defend-attack games under incomplete information. For illustration purposes, we briefly sketch also the complete information case.

Assume a Defender ($D$, she) who chooses her defense $d \in \mathcal{D}$, where $\mathcal{D}$ is her set of feasible alternatives. After having observed it, an Attacker ($A$, he) chooses his attack $a \in \mathcal{A}$, with $\mathcal{A}$ being his set of available alternatives. The consequences of the interaction for both agents depend on a random outcome $\theta \in \Theta$, with $\Theta$ the space of outcomes. Figure 1 displays a template for the corresponding BAID (Banks et al., 2015). Arc $D$-$A$ reflects that the Attacker observes the Defender's decision. As a motivating example, suppose that a defending organisation decides its protection level against, say, a cyber-attack. An attacking organisation observes such protection level and decides its attack intensity. The attack may be successful or not depending on its intensity and the protection level chosen by the defender. Increasing protection level (attack intensity) entail higher costs for the defender (attacker).

The agents have their own assessment of the probability of the random outcome, respectively modeled through $p_D(\theta | d, a)$ and $p_A(\theta | d, a)$. The Defender's utility $u_D(d, \theta)$ is a function of her chosen defense and the outcome. Similarly, the Attacker's utility function is $u_A(a, \theta)$.

### 3.2. Complete information. Game theoretic solution

In the standard complete information setup, the basic game-theoretic solution does not require $A$ to know $D$'s probabilities and utilities, as he observes her actions. However, the Defender must know $(u_A, p_A)$, the common knowledge assumption in this case. Then, both agents'

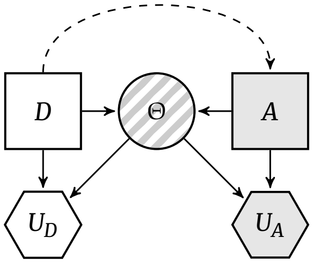

**Figure 1.** Template for basic two player sequential defend-attack game BAID. White nodes affect solely the Defender; grey nodes affect only the Attacker; striped nodes affect both agents.

expected utilities $\psi_A(d,a) = \int u_A(a,\theta)\, p_A(\theta|d,a)\, d\theta$ and $\psi_D(d,a) = \int u_D(d,\theta)\, p_D(\theta|d,a)\, d\theta$ can be computed. The Attacker's best response to $D$'s action $d$, is $a^*(d) = \arg\max_{a\in\mathcal{A}} \psi_A(d,a)$. Such response is used to find the Defender's optimal game-theoretic action $d^*_{\mathrm{GT}} = \arg\max_{d\in\mathcal{D}} \psi_D(d, a^*(d))$. The pair $(d^*_{\mathrm{GT}}, a^*(d^*_{\mathrm{GT}}))$ is a Nash equilibrium and, indeed, a sub-game perfect equilibrium (Hargreaves-Heap and Varoufakis, 2004).

In case of multiple best responses for the attacker given some defender action $d$, $a^*(d)$ becomes a set. Ties are generally broken either choosing the most favorable attack for the defender, leading to a *strong Stackelberg equilibrium (SSE)* or choosing the worst attack for the defender, leading to a *weak Stackelberg equilibrium*. In what follows, whenever necessary, we shall assume that ties are broken in favor of the defender, the standard solution in security games (Korzhyk et al., 2011).

## 3.3. Incomplete information. ARA solution

As discussed in the introduction, the complete information assumption will not hold in many security scenarios. When this is the case, the problem may be handled as an incomplete information game. The most common approach in such context approximates Bayes-Nash equilibria, see Hargreaves-Heap and Varoufakis (2004) for details. Alternatively, we use a decision analytic approach based on ARA, that facilitates the defender to acknowledge the uncertainty she might have about the attacker's judgements $(u_A, p_A)$. The Defender's problem is depicted in Figure 2a as an influence diagram, where $A$'s action appears as an uncertainty. With this, $D$'s expected utility is $\psi_D(d) = \int \psi_D(d,a)\, p_D(a|d)\, da$. Computing this requires estimating $p_D(a|d)$, $D$'s assessment of the probability that the Attacker will choose $a$ after having observed $d$. Then, her optimal decision is $d^*_{\mathrm{ARA}} = \arg\max_{d\in\mathcal{D}} \psi_D(d)$.

Eliciting $p_D(a|d)$ is complex, as it possesses a strategic element. It is facilitated by analyzing $A$'s problem from $D$'s perspective (Figure 2b). In order to accomplish this, the defender would use all information and judgment available about $A$'s utilities and probabilities. However, instead of using point estimates for $u_A$ and $p_A$ to find $A$'s best response to $d$, as in Secion 3.2, her uncertainty about the attacks would derive from modeling $(u_A, p_A)$ through a distribution $F = (U_A, P_A)$ on the space of utilities and probabilities. With no loss of generality, assume that $U_A$ and $P_A$ are defined on a common probability space $(\Omega, \mathcal{F}, \mathcal{P})$ with atomic elements $\omega \in \Omega$ (Chung, 2001). This induces a distribution over the Attacker's expected utility $\psi_A(d,a)$, where the random expected utility would be $\Psi^\omega_A(d,a) = \int U^\omega_A(a,\theta) P^\omega_A(\theta|d,a)\, d\theta$. In turn, this leads to a random optimal alternative defined through $A^*(d)^\omega = \arg\max_{x\in\mathcal{A}} \Psi^\omega_A(d,x)$. Then, the Defender would find $p_D(a|d) = \mathbb{P}_F\left[A^*(d) = a\right] = \mathcal{P}\left\{\omega : A^*(d)^\omega = a\right\}$ in the discrete case (and, similarly in the continuous one).

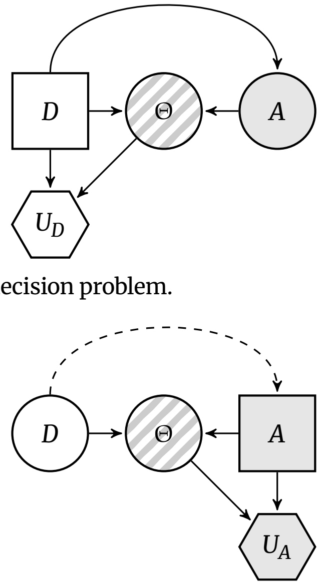

**(a)** Defender's decision problem.

**(b)** Defender analysis of Attacker problem.

**Figure 2.** Influence diagrams for Defender and Attacker problems.

## 3.4. Computation through APS

Computationally, ARA models entail integration and optimization procedures that can be challenging in many cases. Therefore, we explore the APS-based methods sketched in the introduction.

An augmented distribution for the Defender's problem is introduced as $\pi_D(d,a,\theta) \propto u_D(d,\theta)\, p_D(\theta|d,a)\, p_D(a|d)$. Its marginal $\pi_D(d) = \int\int \pi_D(d,a,\theta)\, da\, d\theta$ is proportional to the expected utility $\psi_D(d)$ and, consequently, $d^*_{\mathrm{ARA}} = \mathrm{mode}\,(\pi_D(d))$. Thus, one just needs to sample $(d,a,\theta) \sim \pi_D(d,a,\theta)$ and estimate its mode in $d$. In ? a Metropolis Hastings sampling procedure is introduced to provide a Markov chain with states $(d,\ a,\ \theta)$ whose stationary distribution is $\pi_D(d,a,\theta)$. Under mild conditions, a convenient mode estimate of the $d$ samples generated converges to the ARA solution of the sequential game. Constructing this Markov chain, requires the ability to sample from $p_D(a|d)$. To perform this sampling, an APS in the space of the Attacker's random utilities and probabilities is constructed. For a given $d$, the random augmented distribution built is $\Pi_A^\omega(a,\theta|d) \propto U_A^\omega(a,\theta) P_A^\omega(\theta|d,a)$, its marginal $\Pi_A^\omega(a|d) = \int \Pi_A^\omega(a,\theta|d)\,d\theta$ being proportional to $A$'s random expected utility $\Psi_A^\omega(d,a)$. Then, the random optimal attack $A^*(d)^\omega$ almost surely coincides with the mode of the marginal $\Pi_A^\omega(a|d)$. Consequently, by sampling $u_A(a,\theta) \sim U_A(a,\theta)$ and $p_A(\theta|d,a) \sim P_A(\theta|d,a)$, one can build $\pi_A(a,\theta|d) \propto u_A(a,\theta) p_A(\theta|d,a)$ which is a sample from $\Pi_A(a,\theta|d)$. Then, mode$(\pi_A(a|d))$ is a sample of $A^*(d)$, whose distribution is $\mathbb{P}_F\left[A^*(d) \leqslant a\right] = p_D(A \leqslant a|d)$, thus providing a mechanism to sample from it.

The type of computations underlying APS-based solution methods for ARA are illustrated in an example in Section 6.

## 4. Simultaneous Defend-Attack Games

### 4.1. Basic template

The simultaneous defend-attack game template is depicted in Figure 3, with the same semantics as that of Figure 1. Observe that arc $D \to A$ is lacking, reflecting that both agents act simultaneously in this case. As an example, consider a national aviation security organisation that randomly assigns undercover security guards on planes to deter terror acts. In turn, a terrorist organisation might board a member on certain planes to implement their malicious actions. The security consequences on a given plane would depend on both agents' decisions (presence of security guard and terrorist on the plane).

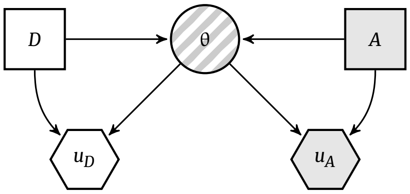

**Figure 3.** Template for basic two player simultaneous defend-attack game.

In this case, the ARA approach would support $D$ by solving $d^*_{\text{ARA}} = \arg\max_{d \in \mathcal{D}} \int \psi_D(d,a) \cdot p_D(a)\,da$ where $p_D(a)$ represents $D$'s beliefs about what action $a$ will the attacker implement. In order to forecast it, consider the attacker's problem $a^* = \arg\max_{a \in \mathcal{A}} \int \psi_A(d,a) \cdot p_A(d)\,dd$. Since we lack knowledge about the attacker's utilities and probabilities, we adopt a Bayesian approach and model our uncertainty through prior distributions (over possible utilities and probabilities) $U_A(a,\theta)$, $P_A(\theta|d,a)$ and $P_A(d)$ defined over a common probability space $(\Omega, \mathcal{F}, \mathcal{P})$ with atomic elements $\omega \in \Omega$. Then, the random optimal attack is

$$A^* = \arg\max_{x \in \mathcal{A}} \int \Psi_A^\omega(d,x) \cdot P_A^\omega(d)\,dd \qquad (1)$$

and make $p_D(A \leqslant a) = \mathcal{P}(A^* \leqslant a)$. This defines $p_D(a)$, the missing ingredient in the Defender problem.

Clearly, the above requires the specification of $U_A(a,\theta)$, $P_A(\theta|d,a)$, and $P_A(d)$]. Of these three ingredients, eliciting $P_A(d)$ has a recursive component as we need to think about how the attacker thinks about the defender, which in turn begs for thinking about how the defender thinks about how the attacker thinks about the defender, and so on. This leads to a recursion similar to the level-$k$ scheme described in Rios and Insua (2012). As mentioned in the Introduction, we shall only focus on level-2 thinking, although the scheme described extend to higher levels in the thinking hierarchy.

As in the sequential case, ARA models may be solved through APS constructing augmented distributions. The distribution corresponding to the Defender's problem is

$$\pi_D(d,\theta,a) \propto u_D(d,\theta) \cdot p_D(\theta|d,a) \cdot p_D(a) \qquad (2)$$

its marginal being $\pi_D(d) = \iint u_D(d,\theta) \cdot p_D(\theta|d,a) \cdot p_D(a)\,d\theta\,da$ which is proportional to the expected utility $\psi_D(d)$ and, consequently, $d^*_{ARA} = \text{mode}(\pi_D(d))$. As in the sequential case, we can solve the problem sampling $d,\theta$ and $a$ from $\pi_D(d,\theta,a)$, and computing the mode of the $d$-samples. This sampling is performed using MCMC methods. In particular, a Metropolis-Hastings is illustrated in Algorithm 1. Notice that this approach requires the ability to sample $a \sim p_D(a)$. To produce those samples, consider now the attacker random APS model, defined for each $\omega$ through

$$\Pi_A^\omega(d,\theta,a) \propto U_A^\omega(a,\theta) \cdot P_A^\omega(\theta|d,a) \cdot P_A^\omega(d) \qquad (3)$$

Then, the random mode of the marginal of $\Pi^\omega(d,\theta,a)$ in $a$ coincides with $A^*$, whose distribution is $p_D(a)$. Consequently, by sampling $u_A(a,\theta) \sim U_A(a,\theta)$, $p_A(\theta|d,a) \sim P_A(\theta|d,a)$ and $p_A(d) \sim P_A(d)$ one can build $\pi_A(a,\theta|d) \propto u_A(a,\theta) p_A(\theta|d,a) p_A(d)$ which is a sample from $\Pi_A(a,\theta,d)$. Then, mode$(\pi_A(a))$ is a sample of $A^*$, whose distribution is $\mathbb{P}_F\left[A^*(d) \leqslant a\right] = p_D(A \leqslant a|d)$, thus providing a mechanism to sample from such distribution. This procedure is illustrated in the `sample_attack` function of Algorithm 1. Under mild conditions, we can prove the convergence of the Algorithm to the ARA solution when using a consistent mode estimator, see Chacon (2020).

## 5. General Games

The previous sections dealt with relatively simple ARA templates with basic sequences of defense and attack movements. However, such stylized settings may not be sufficient to cope with the complexities of many actual defense and homeland security problems.

As an illustration, consider the BAID in Figure 4. In it, the incumbent authorities decide how to allocate a given number of patrols to prevent land access to a place of interest (e.g. an airport) by a certain terrorist organization. This decision is referred to as $D_1$. The terrorists observe the patrols' display over a certain period and thus have

```
function sample_attack($(M, K, g_A, U_A, P_A)$):
    initialize : $a^{(0)}$
    Draw $u_A(a, \theta) \sim U_A(a, \theta)$
    Draw $p_A(\theta | d, a) \sim P_A(\theta | d, a)$
    Draw $p_A(d) \sim P_A(d)$
    Draw $d^{(0)} \sim p_A(d)$
    Draw $\theta^{(0)} \sim p_A(\theta | d^{(0)}, a^{(0)})$
    for $i = 1$ to $M$ do                                    ▷ Inner APS
        Propose new attack $\tilde{a} \sim g_A(\tilde{a} | a^{(i-1)})$
        Draw $\tilde{d} \sim p_A(d)$
        Draw $\tilde{\theta} \sim p_A(\theta | \tilde{d}, \tilde{a})$
        Evaluate acceptance probability
          $\alpha = \min \left\{ 1, \frac{u_A(\tilde{a},\tilde{\theta}) \cdot g_A\left(a^{(i-1)} | \tilde{a}\right)}{u_A\left(a^{(i-1)},\theta^{(i-1)}\right) \cdot g_A\left(\tilde{a} | a^{(i-1)}\right)} \right\}$
        With probability $\alpha$ set
          $(d^{(i)}, a^{(i)}, \theta^{(i)}) = (\tilde{d}, \tilde{a}, \tilde{\theta})$. Otherwise, set
          $(d^{(i)}, a^{(i)}, \theta^{(i)}) = (d^{(i-1)}, a^{(i-1)}, \theta^{(i-1)})$.
    Discard first $K$ samples and record the mode of
      interest of the rest of draws $\{a^{(K+1)}, ..., a^{(M)}\}$ as
      $a^*$.
    return $a^*$
input: $U_A, P_A, M, K, N, R, g_D$ and $g_A$ proposal
        distributions
initialize : $d^{(0)}$
Draw $a^{(0)} \sim p_D(a)$ using sample_attack( $M, K, g_A,$
  $U_A, P_A$)
Draw $\theta^{(0)} \sim p_D(\theta | d^{(0)}, a^{(0)})$
for $i = 1$ to $N$ do                                        ▷ Outer APS
    Propose new defense $\tilde{d} \sim g_D(\tilde{d} | d^{(i-1)})$
    Draw $\tilde{a} \sim p_D(a)$ using sample_attack($M, K, g_A,$
      $U_A, P_A$)
    Draw $\tilde{\theta} \sim p_D(\theta | \tilde{d}, \tilde{a})$
    Evaluate acceptance probability
      $\alpha = \min \left\{ 1, \frac{u_D\left(\tilde{d},\tilde{\theta}\right) \cdot g_D\left(d^{(i-1)} | \tilde{d}\right)}{u_D\left(d^{(i-1)},\theta^{(i-1)}\right) \cdot g_D\left(\tilde{d} | d^{(i-1)}\right)} \right\}$
    With probability $\alpha$ set $(d^{(i)}, a^{(i)}, \theta^{(i)}) = (\tilde{d}, \tilde{a}, \tilde{\theta})$.
      Otherwise, set
      $(d^{(i)}, a^{(i)}, \theta^{(i)}) = (d^{(i-1)}, a^{(i-1)}, \theta^{(i-1)})$.
Discard first $R$ samples and estimate mode of the
  rest of draws $\{d^{(R+1)}, ..., d^{(N)}\}$ as $\widehat{d}^*_{\mathrm{ARA}}$
Record $\widehat{d}^*_{\mathrm{ARA}}$.
```

**Algorithm 1:** MH based APS to approximate ARA solution in the simultaneous defend-attack game.

knowledge of $D_1$. Based on this, the terrorists decide when to access the place of interest with the intention of committing an attack. This decision is referred to as $A_1$. Whether the terrorists manage to access the airport without being detected by the patrols is modeled through variable $S_1$. Provided that the terrorist have a successful access, their vehicles can still be detected before the attack is committed (for instance, via helicopters supervising the area). Whether or not the terrorists' vehicles are detected after land access is modeled with the variable $S_2$. If the terrorists are indeed detected, they choose to change their attacking strategy (for instance, modify their target). The new attacking decision is denoted as $A_2$. Simultaneously, the authorities decide to allocate reinforcement patrols to further protect the airport. Based on these two decisions, the outcome $S_3$ models the (random) consequences of an eventual attack. For each agent, there is a utility node which assesses their consequences, respectively represented by $u_D$ and $u_A$ in Figure 4. In particular, the defender's utility $u_D$ (respectively, the attacker's utility $u_A$) will depend on her decisions $D_1$ and $D_2$ (respectively, his decisions $A_1$ and $A_2$) and the results $S_1$, $S_2$ and $S_3$. Moreover, this, or similar sequences of defense-attack movements, could be repeated across time, spanning over several planning periods.

Note that within the general layout of the BAID in Figure 4, we identify patterns of defence and attack moves as those earlier described as basic templates. In particular, nodes $D_1$–$A_1$ correspond to the Seq D-A template. Similarly, nodes $D_1$–$A_2$ reproduce the backbone structure of a Sim D-A template.

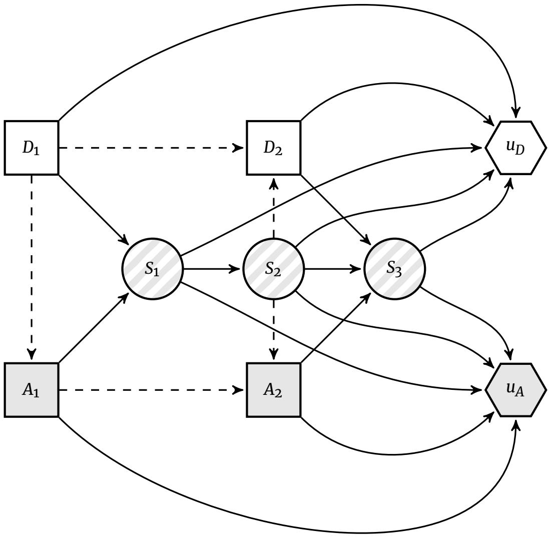

**Figure 4.** ARA modeling of general BAID example.

In González-Ortega et al. (2019), a scheme for dealing with these structures was introduced adapting strategic relevance concepts from computational game theory (Koller, 2003), ARA methods (Banks et al., 2015), and classic influence diagram reductions from Shachter (1986). However, the basic reduction operations, were essentially analytic. Inspired by the reasoning in Appendix B, we replace here such operations with APS schemes.

### 5.1. Simulation based BAID reduction operations

The methods in González-Ortega et al. (2019) assess analytically a BAID from an ARA perspective. First, they are used to decide whether we deal with a sequential or a simultaneous block. Once such decision is made, we decide which ID reduction to implement, according to Shachter (1986)'s conditions. At a third level, we then implement the required ID reductions until the incumbent block is reduced. In such case, we start again finding the next block to be treated, until all the Defender's decision nodes are reduced.

ID reductions in the Defender problem (*D-reductions*) are as in Shachter (1986). On the contrary, ID reductions in the Attacker problem (*A-reductions*) need to be modified to take into account the uncertainty about the attacker's probabilities and utilities. As before, without loss of generality, assume that all involved random utilities $U_A$ and probabilities $P_A$ are defined over a common probability space $(\Omega, \mathcal{F}, \mathcal{P})$ with atomic elements $\omega$. Thus, when invoking them in reference to $\omega$, we designate them $U_A^{\omega}$ and $P_A^{\omega}$, respectively. The remaining notation is adapted from Shachter (1986) where: superscripts *old* and *new* refer to an element, respectively, before or after a BAID transformation; and, for any node $i$, $C(i)$ designates its conditional predecessors (chance and value nodes), $I(i)$ its informational predecessors (decision nodes) and $S(i)$ its direct successors. Initially, we denote $\Psi_A^{\omega}(x_{C(v)}) = U_A^{\omega}(x_{C(v)})$, where $v$ refers to the value node. González-Ortega et al. (2019) provided their theoretical analytical descriptions which we briefly recall:

- *A-Arc inversion*. An arc $(i, j)$ between chance nodes $i$ and $j$ (for the attacker), satisfying Shachter (1986) conditions may be reduced with node inheritance as in conventional IDs. (Random) probability assessments are changed by applying Bayes' formula parametrically so that

$$P_A^{\omega\, new}(x_j|x_{C\,new(j)}) = \int P_A^{\omega\, old}(x_j|x_{C\,old(j)})\, P_A^{\omega\, old}(x_i|x_{C\,old(i)})\; dx_i, \tag{4}$$

$$P_A^{\omega\, new}(x_i|x_{C\,new(i)}) = \frac{P_A^{\omega\, old}(x_j|x_{C\,old(j)})\, P_A^{\omega\, old}(x_i|x_{C\,old(i)})}{P_A^{\omega\, new}(x_j|x_{C\,new(j)})}. \tag{5}$$

- *A-Chance node removal*. A chance node $i$ (for the attacker) satisfying Shachter (1986) conditions may be removed with node inheritance as in conventional IDs. Note, however, that expectations have to be taken parametrically so that we obtain (random) expected utilities. Prior to node $i$ reduction, we have a (random) expected utility $\Psi_A^{\omega\, old}(x_{C\,old(v)})$ associated with each combination $x_{C\,old(v)}$ of values for predecessors of $v$. After the reduction, we associate with $v$ a new (random) expected utility defined by

$$\Psi_A^{\omega\, new}(x_{C\,new(v)}) = \int \Psi_A^{\omega\, old}(x_{C\,old(v)})\, P_A^{\omega}(x_i|x_{C(i)})\; dx_i. \tag{6}$$

- *A-Decision node removal*. Under Shachter (1986) conditions, we may remove a decision node $i$ (for the attacker) with node inheritance as in conventional IDs, by computing the expected utility of the (random) optimal alternatives, conditional on the values of its predecessors. (Random) optimal alternatives are stored through

$$A_i^{\omega *}(x_{C\,new(v)}) = \arg\max_{x_i} \Psi_A^{\omega\, old}(x_i, x_{C\,new(v)}), \tag{7}$$

whereas their (random) expected utilities are

$$\Psi_A^{\omega\, new}(x_{C\,new(v)}) = \max_{x_i} \Psi_A^{\omega\, old}(x_i, x_{C\,new(v)}). \tag{8}$$

Now the crucial observation is that for expected utility optimization purposes, the denominator in Bayes' formula operates as a constant and we can just use the potential on the right of the proportionality condition

$$P_A^{\omega\, new}(x_i|x_{C\,new(i)}) \propto P_A^{\omega\, old}(x_j|x_{C\,old(j)})\, P_A^{\omega\, old}(x_i|x_{C\,old(i)}),. \tag{9}$$

thus avoiding the costly arc inversions. Then, we can group all arc inversions and chance node removals prior to a decision node removal conceptually and, operationally, by multiplying the corresponding parameterized potentials and distributions. Next, we just multiply them by the parameterized utility function to obtain the parameterized augmented probability distribution, from which we may sample to obtain a forecast of the attacker's action as required.

In a similar fashion, we can determine all arc inversions and decision nodes prior to a defender decision node removal, multiply her utility function by the product of potentials and distributions to obtain an augmented probability distribution from which to sample the defender's optimal decision at such stage.

These APS modified operations would be integrated in the operation and block scheduling scheme in González-Ortega et al. (2019) to provide an APS approach to general BAIDs.

## 6. Example

An example illustrates the type of computations undertaken in APS for ARA in a simple sequential Defend-Attack game. Consider an organization ($D$) that has to choose among ten security protocols: $d = 0$ (no extra defensive action); $d = i$ (level $i$ protocol with increasing protection), $i = 1, \ldots, 8$; $d = 9$ (safe but cumbersome protocol). $A$ has two alternatives: attack ($a = 1$) or not ($a = 0$). Successful (unsuccessful) attacks are denoted as $\theta = 1$ ($\theta = 0$). When there is no attack, $\theta = 0$.

|  | $\theta$ |  |
|---|---|---|
| $d$ | 0 | 1 |
| 0 | 0.05 | 7.05 |
| 1 | 0.10 | 7.10 |
| 2 | 0.15 | 7.15 |
| 3 | 0.20 | 7.20 |
| 4 | 0.25 | 7.25 |
| 5 | 0.30 | 7.30 |
| 6 | 0.35 | 7.35 |
| 7 | 0.40 | 7.40 |
| 8 | 0.45 | 7.45 |
| 9 | 0.50 | 7.50 |

**(a)**

|  | $a$ |  |
|---|---|---|
| $d$ | 0 | 1 |
| 0 | 0.0 | 0.50 |
| 1 | 0.0 | 0.40 |
| 2 | 0.0 | 0.35 |
| 3 | 0.0 | 0.30 |
| 4 | 0.0 | 0.25 |
| 5 | 0.0 | 0.20 |
| 6 | 0.0 | 0.15 |
| 7 | 0.0 | 0.10 |
| 8 | 0.0 | 0.05 |
| 9 | 0.0 | 0.01 |

**(b)**

|  | $\theta$ |  |
|---|---|---|
| $a$ | 0 | 1 |
| 0 | 0.00 | 0.00 |
| 1 | −0.53 | 1.97 |

**(c)**

| $d$ | $\alpha_d$ | $\beta_d$ |
|---|---|---|
| 0 | 50.0 | 50.0 |
| 1 | 40.0 | 60.0 |
| 2 | 35.0 | 65.0 |
| 3 | 30.0 | 70.0 |
| 4 | 25.0 | 75.0 |
| 5 | 20.0 | 80.0 |
| 6 | 15.0 | 85.0 |
| 7 | 10.0 | 90.0 |
| 8 | 5.0 | 95.0 |
| 9 | 1.0 | 99.0 |

**(d)**

**Table 1.** a Def. net costs; b Successful attack probs.; c Att. net benefits; d Beta dist. parameters

*Defender non strategic judgments.* Table 1a presents costs $c_D$ associated with each decision and outcome, based on a 7M€ business valuation; and 0.05M€ base security cost plus 0.05M€ per each security level increase. Upon successful attack, $D$ loses the entire business value. The probability $p_D(\theta = 1|d, a)$ of successful attack is in Table 1b (complementary values for unsuccessful attacks). $D$ is constant risk averse in costs, with utility strategically equivalent to $u_D(c_D) = -\exp(0.4 \times c_D)$.

*Attacker judgments.* The average attack cost is 0.03M€. The average attack benefit is 2M€. An unsuccessful attack has an extra cost of 0.5M €. Table 1c presents the Attacker's net benefit $c_A(a, \theta)$. $D$ thinks that $A$ is constant risk prone over benefits, with utility strategically equivalent to $u_A(c_A) = \exp(e \times c_A)$, with $e > 0$.

*Defender strategic judgments.* $D$'s beliefs over $A$'s judgments are described through $P_A$ and $U_A$. Assume $A$'s random probability of success is modeled as $P_A(\theta = 1|d, a = 1) \sim Beta(\alpha_d, \beta_d)$ with $\alpha_d$ and $\beta_d$ in Table 1d (with expected values the $p_D(\theta = 1|d, a)$ in Table 1b). In addition, $A$'s risk coefficient $e$ is uncertain, with $e \sim \mathcal{U}(0, 2)$, inducing the random utility $U_A(c_A)$.

In this case, for APS, as the cardinality of $\mathcal{D}$ is small, one can estimate the value of $p_D(a|d)$ for each $d$ sampling from this distribution and counting frequencies of different attacks. Figure 5a presents such estimates $\hat{p}_D(a|d)$. Next, the ARA solution for the Defender is computed. Figure 5b presents the frequency of samples from the marginal $\pi_D(d)$. Its mode coincides with the optimal defense, $d^*_{\text{ARA}} = 9$.

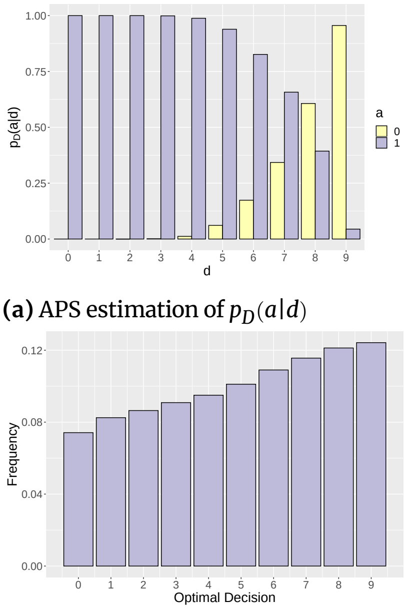

**(a)** APS estimation of $p_D(a|d)$

**(b)** APS solution

**Figure 5.** Estimation of $p_D(a|d)$ through ARA

For this same example, the game theoretic solution under the complete information assumption is $d^*_{\text{GT}} = 8$. Both solutions differ as the informational assumptions are different in both settings: the ARA decision appears to be more conservative, as it suggests a safer but more expensive defense.

## 7. Discussion

APS based methods have already been proven successful in sequential Defend-Attack games Ekin et al. (2022). In this paper, we have extended APS to deal with general games represented as BAIDs from an ARA perspective. Sequential and simultaneous defend-attack models paved the way to deal with general problems modeled as bi-agent influence diagrams. In Appendix A, we sketch how this methodology could be extended to work with n-stages sequential games. However, in this setting, APS is likely to suffer from the typical computational shortcomings of dynamic programming. Thus, studying scalable methods for solving n-stage sequential games under the ARA perspective is an interesting avenue for future research. In addition, exploring the use of Hamiltonian Monte Carlo methods for APS in situations in which gradients of the utilities are available is also a promising future research line.

We would like to finish acknowledging that the proposed approach will become specially relevant in a recent class of security game theoretic models, especially under incomplete information, appearing in the emergent field of adversarial machine learning (Ríos Insua et al., 2020). Efficient scalable algorithmic approaches to solve typical problems in this new domain are essential. Moreover, ARA solutions of such games turns out to be important, as common knowledge assumptions rarely hold in AML.

## 8. Funding

Research supported by the AXA-ICMAT Chair in Adversarial Risk Analysis(PI David Rios Insua), the EOARD-AFOSR project RC2APD, GRANT 13324227 (PIs David Rios Insua and Tahir Ekin) and the STARLIGHT H2020 project 101021797 (PI Nizar Touleimat). None of the funding bodies influenced on the design of this research and the writing of the manuscript.

## References

Banks, D. L., Aliaga, J. R., and Insua, D. R. (2015). *Adversarial Risk Analysis*. CRC Press.

Bielza, C., Müller, P., and Insua, D. R. (1999). Decision analysis by augmented probability simulation. *Management Science*, 45(7):995–1007.

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

Brown, G. G., Carlyle, W. M., and Wood, R. K. (2008). *Optimizing Department of Homeland Security defense investments: Applying Defender-Attacker(-Defender) optimization to terror risk assessment and mitigation*. National Academies Press, Washington, DC:. Appendix E.

Chacon, J. (2020). The modal age of statistics. *International Statistical Review*, 88(1):122–141.

Chung, K. L. (2001). *A Course in Probability Theory*. Academic Press.

Ekin, T., Naveiro, R., Insua, D. R., and Torres-Barrán, A. (2022). Augmented probability simulation methods for sequential games. *European Journal of Operational Research*.

French, S. and Insua, D. R. (2000). *Statistical decision theory*. Wiley.

Gil, C. and Parra-Arnau, J. (2019). An Adversarial-Risk-Analysis Approach to Counterterrorist Online Surveillance. *Sensors*, 19(3).

Gil, C., Rios Insua, D., and Rios, J. (2016). Adversarial Risk Analysis for Urban Security Resource Allocation. *Risk Analysis*, 36(4):727–741.

González-Ortega, J., Insua, D. R., and Cano, J. (2019). Adversarial risk analysis for bi-agent influence diagrams: An algorithmic approach. *European Journal of Operational Research*, 273(3):1085–1096.

Hargreaves-Heap, S. and Varoufakis, Y. (2004). *Game theory: a critical introduction*. New York, Routledge.

Harsanyi, J. C. (1967). Games with incomplete information played by "Bayesian" players, I–III Part I. the basic model. *Management science*, 14(3):159–182.

Hausken, K. (2011). Strategic defense and attack of series systems when agents move sequentially. *IIE Transactions*, 43(7):483–504.

Joshi, C., Aliaga, J. R., and Insua, D. R. (2020). Insider threat modeling: An adversarial risk analysis approach. *IEEE Transactions on Information Forensics and Security*, 16:1131–1142.

Kadane, J. B. and Larkey, P. D. (1982). Subjective probability and the theory of games. *Management Science*, 28(2):113–120.

Koller, D. Milch, B. (2003). Multi-agent influence diagrams for representing and solving games. *Games and Economic Behavior*, 45(1):181–221.

Korzhyk, D., Yin, Z., Kiekintveld, C., Conitzer, V., and Tambe, M. (2011). Stackelberg vs. Nash in security games: An extended investigation of interchangeability, equivalence, and uniqueness. *Jour. Art. Intell. Res.*, 41:297–327.

Naveiro, R., Redondo, A., Insua, D. R., and Ruggeri, F. (2019). Adversarial classification: An adversarial risk analysis approach. *International Journal of Approximate Reasoning*, 113:133 – 148.

Raiffa, H. (1982). *The art and science of negotiation (2003 ed.)*. Cambridge, MA: Harvard University Press.

Raiffa, H., Richardson, J., and Metcalfe, D. (2002). *Negotiation analysis: The science and art of collaborative decision making (2002 ed.)*. Cambridge, MA: Harvard University Press.

Rios, J. and Insua, D. R. (2012). Adversarial risk analysis for counterterrorism modeling. *Risk Analysis*, 32(5):894–915.

Rios Insua, D., Banks, D., and Rios, J. (2016). Modeling opponents in adversarial risk analysis. *Risk Analysis*, 36(4):742–755.

Rios Insua, D., Couce-Vieira, A., Rubio, J. A., Pieters, W., Labunets, K., and G. Rasines, D. (2019). An adversarial risk analysis framework for cybersecurity. *Risk Analysis*, 41(1):16–36.

Ríos Insua, D., Naveiro, R., Gallego, V., and Poulos, J. (2020). Adversarial Machine Learning: Perspectives from adversarial risk analysis. *arXiv e-prints*, page arXiv:1908.06901.

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

Roponen, J. and Salo, A. (2015). Adversarial risk analysis for enhancing combat simulation models. *Journal of Military Studies*, 6(2):82–103.

Sevillano, J. C., Insua, D. R., and Rios, J. (2012). Adversarial Risk Analysis: The Somali Pirates Case. *Decision Analysis*, 9(2):86–95.

Shachter, R. D. (1986). Evaluating influence diagrams. *Operations research*, 34(6):871–882.

Stahl, D. O. and Wilson, P. W. (1995). On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 10(1):218–254.

Wang, S. and Banks, D. (2011). Network routing for insurgency: An adversarial risk analysis framework. *Naval Research Logistics (NRL)*, 58(6):595–607.

Zhuang, J. and Bier, V. M. (2007). Balancing terrorism and natural disasters—defensive strategy with endogenous attacker effort. *Operations Research*, 55(5):976–991.

## A. The n-stage Sequential Game

The proposed scheme for the two stage sequential Defend-Attack game extend to multiple stage games, informing our general approach in Section 4. Consider a general $n$-stage sequential game under incomplete information in which the defender moves first by choosing decision $d_1 \in \mathcal{D}_1$. After observing $d_1$, the attacker selects $a_2 \in \mathcal{A}_2$. Subsequently, the defender observes $a_2$, and decides $d_3 \in \mathcal{D}_3$. Interactions between players continue until stage $n$. After all actions have been chosen, the uncertainty $\theta$ is resolved. In general, the distribution of $\theta$ depends on all actions taken. For simplicity, denote by $S_i$ the sequence of actions taken from stage $i$. For instance, if the defender moves at stage $i$, $S_i = \{d_i, a_{i+1}, d_{i+2}, \dots\}$, and $S_{i-1} = S_i \cup a_{i-1}$. Then, the defender's beliefs about $\theta$, given the sequence $S_1$, are encoded in the distribution $p_D(\theta|S_1)$. In addition, denote the Defender's (Attacker's) utility at the $i$-th stage as $u_{D_i}(d_i, \theta)$ ($u_{A_i}(a_i, \theta)$).

With this, at the first stage, the defender needs to solve for $d_1^* = \arg\max_{d_1} \int u_{D_1}(d_1, \theta) \cdot p_D(\theta|S_1) \cdot p_D(S_2|d_1)\, \mathrm{d}\theta\, \mathrm{d}S_2$, where $p_D(S_2|d_1)$ encodes the Defender's beliefs about the sequence of actions starting at the second stage, given that she chooses $d_1$ at the first stage. We define an augmented distribution over the space of decisions and uncertainties $\pi_D(d_1, \theta, S_2) \propto u_{D_1}(d_1, \theta) \cdot p(\theta|S_1) \cdot p_D(S_2|d_1)$: its mode in $d_1$ coincides with $d_1^*$. Thus, we could solve the problem generating samples $d_1, \theta, S_2 \sim \pi_D(d_1, \theta, S_2)$ and computing the mode of the $d_1$ samples.

This requires generating sequences of actions given decision $d_1, S_2 \sim p_D(S_2|d_1$. Note that we can write $p_D(S_2|d_1) = p_D(a_2|d_1) \cdot p_D(d_3|d_1, a_2) \cdots$ Thus, $S_2$ can be sampled sequentially, by first sampling $a_2 \sim p_D(a_2|d_1)$, then $d_3 \sim p_D(d_3|d_1, a_2)$, and so on. Note also that:

1. For every $i$, $p_D(d_i|d_1, a_2, \dots, a_{i-1})$ is a point mass centered at the solution of the $i$-th stage problem, which can be computed using APS, as we did in the first stage problem.
2. For every $i$, sampling $a_i \sim p_D(a_i|d_1, a_2, \dots, d_{i-1})$, requires solving the Attacker's $i$-th stage problem. Let us illustrate how this could be done leveraging the ARA methodology with the Attacker's second stage problem. In such stage, after observing $d_1$, a expected utility maximizer attacker would choose

$$
\begin{aligned}
a_2^* &= \arg\max_{a_2} \int u_{A_2}(a_1, \theta) \cdot p_A(\theta|d_1, a_2, S_3) \\
&\cdot \; p_A(S_3|d_1, a_2)\, \mathrm{d}\theta\, \mathrm{d}S_3.
\end{aligned}
$$

However, the Defender is uncertain about the Attacker's utilities and probability judgments. As we did in the 2 stage game, for a given $d_1$, the random augmented distribution built is $\Pi^{\omega}_{A_2}(a_2, S_3\theta|d_1) \propto U^{\omega}_{A_2}(a_2, \theta) P^{\omega}_A(\theta|d_1, a_2, S_3) P^{\omega}_A(S_3|d_1, a_2)$, its marginal on $a_2$ is proportional to $A$'s random expected utility. Consequently, by sampling $u_{A_1}(a_2, \theta) \sim U_{A_2}(a_2, \theta)$, $p_A(\theta|d_1, a_2, S_3) \sim P_A(\theta|d_1, a_2, S_3)$ and $p_A(S_3|d_1, a_2) \sim P_A(S_3|d_1, a_2)$ one can build $\pi_{A_2}(a_2, S_3\theta|d_1)$ which is a sample from $\Pi_{A_2}(a_2, S_3\theta|d_1)$. Then, the mode of its marginal in $a_2$ is distributed as $p_D(a_2|d_1)$, thus providing a mechanism to sample from such distribution. Notice that modeling the Defender's uncertainty about $p_A(S_3|d_1, a_2)$, entails constructing a hierarchy of decision-making problems.