

Handling Risk-aware Attackers in Security Games

Yundi Qian¹, William B. Haskell², Milind Tambe¹

¹ {yundi.qian, tambe}@usc.edu

University of Southern California

² isehwb@nus.edu.sg

National University of Singapore

Abstract. Stackelberg security games (SSGs) are now established as a powerful tool in security domains. In this paper^{3,4}, we consider a new dimension of security games: the risk preferences of the attacker. Previous work assumes a risk-neutral attacker that maximizes his expected reward. However, extensive studies show that the attackers in some domains are in fact risk-averse, e.g., terrorist groups in counter-terrorism domains. The failure to incorporate the risk aversion in SSG models may lead the defender to suffer significant losses. Additionally, defenders are uncertain about the degree of attacker’s risk aversion. Motivated by this challenge this paper provides the following five contributions: (i) we propose a novel model for security games against risk-averse attackers with uncertainty in the degree of their risk aversion; (ii) we develop an intuitive MIBLP formulation based on previous security games research, but find that it finds locally optimal solutions and is unable to scale up; (iii) based on insights from our MIBLP formulation, we develop our scalable BeRRA algorithm that finds globally ϵ -optimal solutions; (iv) we extend our BeRRA algorithm to handle other risk-aware attackers, e.g., risk-seeking criminals; (v) we show that we do not need to consider attacker’s risk attitude in zero-sum games.

Keywords: Risk-aware, Robust Stackelberg equilibrium, Security games.

1 introduction

Stackelberg security games (SSGs) are now established as a successful tool in the security domain [7, 1]. In this paper, we focus on a critical dimension of SSGs that has not yet been studied — the risk preferences of the attacker. Previous work on game theory for SSGs emphasizes a risk neutral attacker that is trying to maximize his expected reward. However, if the attacker is not risk-neutral but is actually risk-averse, then the failure to model this risk attitude may lead the defender to suffer significant losses in solution quality.

³ This paper extends our AAMAS main track paper [14]. We extend our BeRRA algorithm to other risk-aware attacker types and add the corresponding experimental results.

⁴ This research is supported by the United States Department of Homeland Security through the National Center for Risk and Economic Analysis of Terrorism Events (CREATE) under award number 2010-ST-061-RE0001 and MURI grant W911NF-11-1-0332.

A major motivating example of our work is the application of security games to the counter-terrorism domain, and there is a thread of work that studies terrorist risk attitudes [13, 12]. In [13], portfolio theory is applied to study a terrorist group’s decision making process, and this research argues that terrorist strategies are risk-averse and are highly sensitive to the group’s level of risk aversion. While this finding of risk aversion may appear to be counter-intuitive, notice that it is the terrorist groups (and the planners in these groups) that are found to be risk-averse due to resource limitation; not the individuals in the organization who finally launch an attack. [12] studies the risk preferences of Al Qaeda specifically and concludes that the group is risk-averse and consistently displays the same degree of risk aversion in their activities.

Risk aversion encompasses a wide range of behavior — so to say that attackers are risk-averse is not enough for the defender. To address this issue, we compute a robust defender strategy against risk-averse attackers with uncertainty in the degree of risk aversion. In this process, we provide the following contributions in this paper. First, we build a robust SSG framework against an attacker with uncertainty in level of risk aversion. Second, building on previous work on SSGs in mixed-integer programs, we provide a novel mixed-integer bilinear programming problem (MIBLP), and find that it only finds locally optimal solutions. While the MIBLP formulation is also unable to scale up, it provides key intuition for our new algorithm. This new algorithm, BeRRA (**B**inary search based **R**obust algorithm against **R**isk-**A**verse attackers) is our third contribution, and it finds globally ϵ -optimal solutions by solving $\mathcal{O}(n \log(\frac{1}{\epsilon}) \log(\frac{1}{\delta}))$ linear feasibility problems. The key idea of our BeRRA algorithm is to reduce the problem from maximizing the reward with a given number of resources to minimizing the number of resources needed to achieve a given reward. This transformation allows BeRRA to scale up via the removal of the bilinear terms and integer variables as well as the utilization of key theoretical properties that prove correspondence of its potential “attack sets” [7] with that of the maximin strategy. Fourth, although the BeRRA algorithm is designed for risk-averse attackers, it also applies to other risk-aware attackers, e.g., risk-seeking attackers. Finally, we also show that we do not need to consider attacker’s risk attitude in zero-sum games. Our experimental results show the solution quality and runtime advantages of our robust model and BeRRA algorithm.

2 Model

2.1 Stackelberg Security Games

An SSG [7, 1] is a two-player game between a defender and an attacker. We consider the problem with n targets where $\mathbb{T} = \{1, 2, \dots, n\}$ is the set of targets. The defender has a total number of m resources to allocate among these n targets to protect them from attack. The defender commits to a mixed strategy \mathbf{c} to protect these targets, where $c_i \in [0, 1]$ is the probability that target i is protected. We have the resource constraint $\sum_{i \in \mathbb{T}} c_i \leq m$. The attacker observes the defender’s strategy \mathbf{c} and then chooses one target to attack. If the attacker attacks a protected target i , this attack is unsuccessful and the attacker receives utility⁵ $U_a^c(i)$ while the defender receives utility $U_d^c(i)$. If the

⁵ Literature on risk denotes $U_a^c(i)/U_d^c(i)$ as values. However, we use the terminology utilities and mapped utilities for consistency with previous work on SSGs.

attacker attacks an unprotected target i , this attack is successful and the attacker receives utility $U_a^u(i)$ while the defender receives utility $U_d^u(i)$. Necessarily, $U_d^u(i) < U_d^c(i)$ and $U_a^c(i) < U_a^u(i), \forall i \in \mathbb{T}$. If $U_d^u(i) + U_a^u(i) = 0$ and $U_d^c(i) + U_a^c(i) = 0, \forall i \in \mathbb{T}$, this SSG is a zero-sum game.

We define $U_a(i, \mathbf{c}) \triangleq c_i U_a^c(i) + (1 - c_i) U_a^u(i)$ to be the expected utility for the attacker when the defender's strategy is \mathbf{c} and the attacker chooses to attack target i ; similarly, $U_d(i, \mathbf{c}) \triangleq c_i U_d^c(i) + (1 - c_i) U_d^u(i)$ is the expected utility for the defender. Given the defender strategy \mathbf{c} , the attacker would attack the target that maximizes his expected utility. When there are ties, the attacker is assumed to break ties in favor of the defender. Thus, a mixed integer linear program (MILP) can be formulated to compute the defender's optimal strategy, as is shown in Problem (1). Here, $\{q_i\}_{i \in \mathbb{T}}$ are auxiliary variables to represent if target i is chosen by the attacker, and M is a constant orders of magnitude larger than all target utilities. The solution \mathbf{c} is called the Strong Stackelberg Equilibrium (SSE) strategy [3, 15] of the game.

$$\begin{aligned}
& \max_{\mathbf{c}, \{q_i\}_{i \in \mathbb{T}}, v, d} && v \\
& \text{s.t.} && 0 \leq c_i \leq 1, \forall i \in \mathbb{T} && \sum_{i \in \mathbb{T}} c_i \leq m \\
& && q_i \in \{0, 1\}, \forall i \in \mathbb{T} && \sum_{i \in \mathbb{T}} q_i = 1 \\
& && v \leq U_d(i, \mathbf{c}) + (1 - q_i)M, \forall i \in \mathbb{T} \\
& && 0 \leq d - U_a(i, \mathbf{c}) \leq (1 - q_i)M, \forall i \in \mathbb{T}
\end{aligned} \tag{1}$$

2.2 Stackelberg Security Games with Unknown Risk-averse Attackers

The SSE strategy provides the optimal defender strategy when the attacker is risk-neutral. However, as previously discussed, attackers are risk-averse rather than risk-neutral in several key domains. If the defender executes the SSE strategy against a risk-averse attacker, then the defender may suffer significant losses in solution quality. We provide an example in the online appendix⁶ showing that the losses can be arbitrarily large. This example strongly motivates the need to consider risk-averse attackers. However, real world defenders are uncertain about the attacker's degree of risk aversion, and the defender may suffer significant losses if she incorrectly estimates it. Therefore we focus on a robust strategy in this paper, i.e., our aim is to compute a defender strategy that is robust against all possible risk-averse attackers.

In literature on risk, the utility function f , which maps values to utilities, is used to specify the risk preference. f is concave for the risk-averse case and is convex for the risk-seeking case, while the risk-neutral case corresponds to the function $y = Cx, C > 0$. The agent makes decisions based on the mapped utilities.

In our problem, we define the mapping function \widehat{U} that maps the utilities $U_a^c(i)$ and $U_a^u(i)$ to the attacker's mapped utilities. We denote $\widehat{U}_a(i, \mathbf{c}) \triangleq c_i \widehat{U}(U_a^c(i)) + (1 - c_i) \widehat{U}(U_a^u(i))$ as the attacker's expected utility under the mapping \widehat{U} . We restrict \widehat{U} to be strictly increasing, concave and satisfying the equality $\widehat{U}(0) = 0$ — strictly increasing

⁶ <http://teamcore.usc.edu/people/yundiqia/web%20page/papers/OptMAS2015Appendix.pdf>

reflects the preference for more to less; concavity corresponds to risk aversion; and $\widehat{U}(0) = 0$ distinguishes between gains and losses. According to this definition, the risk-averse case includes the risk-neutral case.

We define \mathcal{U} to be the set of all valid mapping functions \widehat{U} . Problem (2) describes the robust defender strategy through a bilevel optimization problem. In the upper level, the defender chooses \mathbf{c} to maximize her expected utility $U_d(k, \mathbf{c})$. The constraint $k \in \arg \max_{i \in \mathbb{T}} \widehat{U}_a(i, \mathbf{c})$ requires target k to have the highest expected utility for the attacker under the utility mapping \widehat{U} when the defender’s strategy is \mathbf{c} . The lower level demonstrates that the defender maximizes her worst-case reward over all possible attacker responses with utility mapping functions $\widehat{U} \in \mathcal{U}$. The lower level also suggests that the attacker breaks ties against the defender due to the concept of robustness. We define the solution \mathbf{c} to be the Robust Stackelberg Equilibrium (RSE) strategy of the game.

$$\begin{aligned} \max_{\mathbf{c}} \min_{\widehat{U} \in \mathcal{U}, k} & \left\{ U_d(k, \mathbf{c}) : k \in \arg \max_{i \in \mathbb{T}} \widehat{U}_a(i, \mathbf{c}) \right\} \\ \text{s.t.} \quad & 0 \leq c_i \leq 1, \forall i \in \mathbb{T} \quad \sum_{i \in \mathbb{T}} c_i \leq m \end{aligned} \quad (2)$$

3 Related Work

Our work is related to handling uncertainty in SSGs. Previous approaches can be divided into two categories: 1) model uncertainty in terms of different attacker types and solve a resulting Bayesian Stackelberg game [11, 18]; 2) apply robust optimization techniques to optimize the worst case for the defender over the range of model uncertainty [17, 6, 9].

Bayesian Stackelberg Game Bayesian Stackelberg game models uncertainty by allowing different attacker types, where there is some prior probability corresponding to each attacker type. Although this method is used to model payoff uncertainty in previous work [11, 18], it can also be used to model different degrees of attacker risk aversion in SSGs. However, this approach requires a prior distribution of attacker types, which is usually inapplicable for many real-world security domains [9]. In addition, it is difficult to apply this approach to infinitely many attacker types. Therefore, we focus on the robust approach in this work.

Robust Approach The robust approach for SSGs, in line with robust optimization, computes a defender strategy that optimizes the worst case over the model uncertainty. Yin et al. [17] computes a defender strategy that is robust against defender execution uncertainty as well as uncertainty in the attacker’s observations of the defender’s strategy. Kiekintveld et al. [6] focus on interval uncertainty in the attacker’s payoffs. Nguyen et al. [9] develop a robust strategy that takes the attacker’s bounded rationality into account as well as the uncertainties [17, 6] discuss.

The previous work has addressed neither attacker risk aversion nor ambiguity about the attacker risk profile. Although Kiekintveld et al. [6] try to capture uncertainty in attacker’s utilities, they are unable to fully capture the attacker’s risk aversion. The mapped utilities are coupled in our problem since they are mapping with the same utility

function \widehat{U} , and interval uncertainty is unable to model that. For example, Suppose target t_1 is of reward 1 and penalty -2 ; target t_2 is of reward 2 and penalty -1 . The coverage probability $c_1 = c_2 = 0.5$. A risk-averse attacker will always attack t_2 since \widehat{U} must be strictly increasing. However, the model with interval uncertainty 1 would consider both t_1, t_2 to be potential targets for attack. The weakness of the interval uncertainty model is also shown numerically with experiments later.

A third thread of related work is the research in game theory that explores human’s bounded rationality in decision making — humans do not necessarily choose the strategy that provides them the highest expected utility. Quantal response [8] argues that human are more likely to choose the strategy with a higher expected utility. Yang et al. [16] apply the concept of quantal response to security games. Nguyen et al. [10] propose the SUQR model by extending the quantal response concept with subjective utilities in security games. However, these approaches do not model risk aversion, and nor do they model uncertainty in risk aversion that we model in this paper. In fact, models such as SUQR essentially address concerns that are orthogonal to the issue of risk aversion; future research may thus consider integrating bounded rationality models with risk aversion.

4 Preliminaries

In its current form, the optimization problem (2) is not tractable because it is a bilevel programming problem that requires the solution of uncountably many inner optimization problems indexed by \mathcal{U} . To take steps towards tractability, in Section 4.1 and 4.2, we provide key concepts that are used in our MIBLP formulation (Section 5) and our BeRRA algorithm (Section 6).

4.1 Risk Aversion Modeling

In this section, we write the condition $\widehat{U} \in \mathcal{U}$ in a computationally tractable way via linear constraints. For any utility function $\widehat{U} \in \mathcal{U}$, we are actually only interested in its values at 0 and at the points of the attacker’s utility set $U_a^c(i)$ and $U_a^u(i)$, which we denote as Θ :

$$\Theta = \{U_a^u(i), U_a^c(i), \forall i \in \mathbb{T}\} \cup \{0\} = \{\theta_1, \dots, \theta_I\},$$

where $\theta_1 < \theta_2 < \dots < \theta_I$.

Lemma 1. Choose $\epsilon_u > 0$.⁷ $\widehat{U} \in \mathcal{U}$ is equivalent to satisfying the linear constraints (3) on the values $\{\widehat{U}(\theta)\}_{\theta \in \Theta}$, i.e., $\forall \widehat{U} \in \mathcal{U}$, \widehat{U} satisfies the constraints (3); $\forall \{\widehat{U}'(\theta)\}_{\theta \in \Theta}$

⁷ Since Problem (2) is invariant under scaling of \widehat{U} , i.e., the attacker makes the same decision under either \widehat{U} or $\alpha\widehat{U}$, $\forall \alpha > 0$. Thus, the value of ϵ_u does not affect the result.

that satisfies constraints (3), $\exists \hat{U} \in \mathcal{U}$ such that $\{\hat{U}(\theta) = \hat{U}'(\theta)\}_{\theta \in \Theta}$.

$$\begin{aligned} \frac{\hat{U}(\theta_2) - \hat{U}(\theta_1)}{\theta_2 - \theta_1} &\geq \frac{\hat{U}(\theta_3) - \hat{U}(\theta_2)}{\theta_3 - \theta_2} \\ &\geq \dots \geq \frac{\hat{U}(\theta_I) - \hat{U}(\theta_{I-1})}{\theta_I - \theta_{I-1}} \geq \epsilon_u \\ \hat{U}(0) &= 0 \end{aligned} \quad (3)$$

The proof of Lemma 1 can be found in the online appendix. Based on Lemma 1, the condition $\hat{U} \in \mathcal{U}$ is completely captured by constraints (3). From now on we denote the constraints (3) compactly as $\hat{U} \in \mathcal{U}$.

4.2 Possible Attack Set

In this section, to better understand Problem (2) we study the ‘‘possible attack set’’ $S_p(\mathbf{c})$ and its complement $S_i(\mathbf{c}) = \mathbb{T} - S_p(\mathbf{c})$.

Definition 1. Given the coverage probability \mathbf{c} , Possible Attack Set $S_p(\mathbf{c})$ is defined to be the set of targets that may be attacked by a risk-averse attacker, i.e., it is the set of targets that have the highest expected utility for the attacker for some $\hat{U} \in \mathcal{U}$.

$S_i(\mathbf{c}) = \mathbb{T} - S_p(\mathbf{c})$ is defined to be the set of targets that the attacker will never attack, i.e., the set of targets that for any $\hat{U} \in \mathcal{U}$, there always exists another target $i \in S_p(\mathbf{c})$ with a higher expected utility for the attacker.

Given the coverage probability \mathbf{c} , we can compute $S_p(\mathbf{c})$ and $S_i(\mathbf{c})$ by testing the feasibility of the following constraints for every target.

$$\begin{aligned} \hat{U}_a(i, \mathbf{c}) &\geq \hat{U}_a(j, \mathbf{c}), \forall j \in T, j \neq i \\ \hat{U} &\in \mathcal{U} \end{aligned} \quad (4)$$

If these constraints are feasible for a target i , there exists a mapping $\hat{U} \in \mathcal{U}$ under which target i has the highest expected utility for the attacker, and thus $i \in S_p(\mathbf{c})$; otherwise, $i \in S_i(\mathbf{c})$.

With this definition, Problem (2) can be written as

$$\begin{aligned} \max_{\mathbf{c}} \min_{i \in S_p(\mathbf{c})} \{U_d(i, \mathbf{c})\} \\ \text{s.t. } 0 \leq c_i \leq 1, \forall i \in \mathbb{T} \quad \sum_{i \in \mathbb{T}} c_i \leq m \end{aligned} \quad (5)$$

5 MIBLP Formulation

In this section, we formulate Problem (5) as an MIBLP problem to find the RSE strategy for the defender. While this approach does not scale up to large-scale games, it provides several insights for our BeRRA algorithm. The derivation can be found in the online appendix.

Theorem 1. *Problem (2) is (approximately)⁸ equivalent to*

$$\begin{aligned}
& \max \quad v \\
& \text{s.t.} \quad 0 \leq c_i \leq 1, \forall i \in \mathbb{T} \quad \sum_{i \in \mathbb{T}} c_i \leq m \\
& \quad \quad q_i \in \{0, 1\}, \forall i \in \mathbb{T} \\
& \quad \quad v \leq U_d(i, \mathbf{c}) + (1 - q_i)M, \forall i \in \mathbb{T} \\
& \quad \quad \widehat{U}_a^i(j, \mathbf{c}) \leq \widehat{U}_a^i(i, \mathbf{c}) + (1 - q_i)M, \forall i \in \mathbb{T}, \forall j \in \mathbb{T}, j \neq i \\
& \quad \quad \widehat{U}^i \in \mathcal{U}, \forall i \in \mathbb{T} \\
& \quad \quad \beta^i \epsilon_u \geq \epsilon_c - q_i M, \forall i \in \mathbb{T} \\
& \quad \quad (\boldsymbol{\alpha}^i, \beta^i, \boldsymbol{\gamma}^i, \kappa^i) \in \mathcal{D}, \forall i \in \mathbb{T}
\end{aligned} \tag{9}$$

We have converted Problem (2) into Problem (9), which is an MIBLP: $\{q_i\}_{i \in \mathbb{T}}$ are integer variables; $\widehat{U}_a^i(j, \mathbf{c}) = c_j \widehat{U}^i(U_a^c(j)) + (1 - c_j) \widehat{U}^i(U_a^u(j))$ contains bilinear terms since both c_j and $\widehat{U}^i(U_a^c(j))/\widehat{U}^i(U_a^u(j))$ are variables. Problem (9) is a non-convex optimization problem and lacks efficient solvers. We used a powerful nonlinear solver — KNITRO to search for local optimal solutions to Problem (9). However, this approach does not scale up — the two-target scenario takes about 1 minute and the three-target scenario takes about 15 minutes to solve. Faced with this scalability issue, we develop the BeRRA algorithm that finds the ϵ -optimal solution and provides significant scalability.

6 BeRRA Algorithm

Problem (9) has two main hindrances to scaling up: the presence of $\Theta(n^2)$ bilinear terms and the presence of n integer variables. Thus, eliminating these bilinear terms and integer variables should allow us to scale the problem up. The bilinear terms in Problem (9) have two components: the coverage probability c_i and the mapped attacker utilities $\widehat{U}(U_a^c(i))/\widehat{U}(U_a^u(i))$. Intuitively, we can avoid the bilinearity by fixing one of these two terms. In addition, if the coverage probability \mathbf{c} is fixed, then $S_p(\mathbf{c})$ is also fixed and we no longer need the integer variables $\{q_i\}_{i \in \mathbb{T}}$ to represent if $i \in S_p(\mathbf{c})$. Based on the idea of fixing the coverage probability \mathbf{c} , we develop the BeRRA algorithm. This algorithm computes an ϵ -optimal RSE strategy where ϵ can be made arbitrarily small.

The main idea of the BeRRA algorithm is to reduce the problem to computing the minimum amount of resources needed to achieve a given reward, which can be solved efficiently by using special properties of the problem. With this reduction, we use binary search to find the highest reward that the defender can achieve with the given number of resources. The high-level intuition of this reduction is that a fixed defender's reward leads to fixed defender maximin strategy, which eliminates the bilinear terms and integer variables. Additionally, optimal strategy can be derived efficiently from the maximin strategy via the property $S_p(\mathbf{c}^{\max}) = S_p(\mathbf{c}^{\text{opt}})$.

⁸ The approximation is due to the introduction of ϵ_c .

6.1 Binary Search Reduction

Algorithm 1 lists the steps of our BeRRA algorithm. The input to Algorithm 1 is the number of defender resources m and the defender's and the attacker's utilities \mathbf{U} . The output is the defender's RSE strategy \mathbf{c} and her reward lb . The lower bound lb and upper bound ub are first set to be the lowest and the highest possible rewards, respectively, that the defender may achieve (Line 2). The function $\text{MinimumResources}(r, \mathbf{U})$ returns the strategy \mathbf{p} that uses the minimum number of resources for the defender to achieve reward r . This function will be discussed in detail in Section 7. During the binary search phase (Lines 3 ~ 11), the lower bound is set to be the defender's achievable reward (the strategy \mathbf{p} returned by the MinimumResources function is the solution) and the upper bound is set to be an unachievable reward. Therefore, the BeRRA algorithm achieves the ϵ -optimal solution and we can set ϵ arbitrarily small to get arbitrarily near-optimal solutions.

Algorithm 1 BeRRA Algorithm

```

1: function BERRA ( $m, \mathbf{U}$ )
2:    $lb \leftarrow \min_{i \in \mathbb{T}} U_d^u(i), ub \leftarrow \max_{i \in \mathbb{T}} U_d^c(i)$ 
3:   while  $ub - lb \geq \epsilon$  do
4:      $\mathbf{p} \leftarrow \text{MINIMUMRESOURCES}(\frac{lb+ub}{2}, \mathbf{U})$ 
5:     if  $\sum_{i \in \mathbb{T}} p_i \leq m$  then
6:        $lb \leftarrow \frac{lb+ub}{2}$ 
7:        $\mathbf{c} \leftarrow \mathbf{p}$ 
8:     else
9:        $ub \leftarrow \frac{lb+ub}{2}$ 
10:    end if
11:  end while
12:  return ( $\mathbf{c}, lb$ )
13: end function

```

7 Minimum Resources

We present Algorithm 2 in this section. This algorithm computes the defender strategy that requires as few resources as possible to achieve a given reward r , i.e., the MinimumResources function in Algorithm 1. We call this resource-minimizing strategy the optimal strategy and denote it as \mathbf{c}^{opt} for succinctness.

Algorithm 2 consists of two functions: Maximin and Reduce. The Maximin function computes the maximin strategy \mathbf{c}^{max} for which the defender achieves reward r , as well as the corresponding sets $S_p(\mathbf{c}^{\text{max}})$ and $S_i(\mathbf{c}^{\text{max}})$. The variable *flag* is set to *false* when the input reward is not achievable for any amount of defender resources. In this case, Algorithm 2 returns $(\infty, \infty, \dots, \infty)^\top$ (Lines 3 ~ 5) so that Algorithm 1 knows r is not achievable. We will prove in Theorem 2 that if the reward r is achievable, then $S_p(\mathbf{c}^{\text{max}}) = S_p(\mathbf{c}^{\text{opt}})$ and $S_i(\mathbf{c}^{\text{max}}) = S_i(\mathbf{c}^{\text{opt}})$. Based on this property, the Reduce function derives the optimal strategy \mathbf{c}^{opt} from the maximin strategy \mathbf{c}^{max} . Section 7.1 and Section 7.2 discuss these two functions in detail.

Algorithm 2 Minimum Resources

```
1: function MINIMUMRESOURCES( $r, \mathbf{U}$ )
2:   ( $flag, \mathbf{c}^{\max}, S_p(\mathbf{c}^{\max}), S_i(\mathbf{c}^{\max})$ )  $\leftarrow$  MAXIMIN( $r, \mathbf{U}$ )
3:   if  $flag = false$  then
4:     return  $(\infty, \infty, \dots, \infty)^\top$ 
5:   end if
6:    $\mathbf{c}^{\text{opt}} \leftarrow$  REDUCE( $\mathbf{U}, \mathbf{c}^{\max}, S_p(\mathbf{c}^{\max}), S_i(\mathbf{c}^{\max})$ )
7:   return  $\mathbf{c}^{\text{opt}}$ 
8: end function
```

7.1 Maximin Function

The Maximin function is summarized in Algorithm 3. It first computes the maximin strategy \mathbf{c}^{\max} for which the defender achieves reward r (Lines 2 ~ 4) and then it assigns each target to either $S_p(\mathbf{c}^{\max})$ or $S_i(\mathbf{c}^{\max})$ (Lines 5 ~ 15). If the reward r is not achievable for any amount of resources, then it returns $flag = false$ (Line 10).

Algorithm 3 Maximin

```
1: function MAXIMIN( $r, \mathbf{U}$ )
2:   for  $i = 1 \rightarrow n$  do
3:      $c_i^{\max} \leftarrow \min\{1, \max\{\frac{r - U_d^u(i)}{U_d^c(i) - U_d^u(i)}, 0\}\}$ 
4:   end for
5:    $S_p(\mathbf{c}^{\max}), S_i(\mathbf{c}^{\max}) \leftarrow \emptyset$ 
6:   for  $i = 1 \rightarrow n$  do
7:     if Problem (4) is feasible for target  $i$  given  $\mathbf{c}^{\max}$  then
8:        $S_p(\mathbf{c}^{\max}) \leftarrow S_p(\mathbf{c}^{\max}) \cup \{i\}$ 
9:       if  $r > U_d^c(i)$  then
10:        return ( $false, \mathbf{c}^{\max}, S_p(\mathbf{c}^{\max}), S_i(\mathbf{c}^{\max})$ )
11:       end if
12:     else
13:        $S_i(\mathbf{c}^{\max}) \leftarrow S_i(\mathbf{c}^{\max}) \cup \{i\}$ 
14:     end if
15:   end for
16:   return ( $true, \mathbf{c}^{\max}, S_p(\mathbf{c}^{\max}), S_i(\mathbf{c}^{\max})$ )
17: end function
```

Lines 2 ~ 4 compute the maximin strategy for a given reward r . Given a coverage probability \mathbf{c} , the maximin setting assumes that the attacker attacks target $i = \arg \min_{j \in \mathbb{T}} U_d(j, \mathbf{c})$, and thus the defender's reward will be $\min_{i \in \mathbb{T}} U_d(i, \mathbf{c})$. For the defender to achieve reward r , we should have $U_d(i, \mathbf{c}) \geq r, \forall i \in \mathbb{T}$ so that $c_i^{\max} = \frac{r - U_d^u(i)}{U_d^c(i) - U_d^u(i)}$ (which is bounded by $[0, 1]$).

Given the maximin strategy \mathbf{c}^{\max} , Lines 5 ~ 15 iterate through all targets and assign them to either $S_p(\mathbf{c}^{\max})$ or $S_i(\mathbf{c}^{\max})$ by testing the feasibility of constraints (4). If these constraints are feasible, then $i \in S_p(\mathbf{c})$; otherwise, $i \in S_i(\mathbf{c})$. Next in

Lemma 3 we prove that $\exists i \in S_p(\mathbf{c}^{\max})$ that satisfies $r > U_d^c(i)$ if and only if reward r is not achievable. In that case, Algorithm 3 returns $flag = false$ (Lines 9 ~ 11).

Lemma 2. *Given coverage probability \mathbf{c} , the defender's reward is $\min_{i \in S_p(\mathbf{c})} U_d(i, \mathbf{c})$.*

Proof. Follows from the form of problem 5.

The proof of the following Lemma 3 can be found in the online appendix.

Lemma 3. *Reward r is infeasible if and only if Algorithm 3 returns $flag = false$.*

Theorem 2 demonstrates why we compute \mathbf{c}^{\max} , $S_p(\mathbf{c}^{\max})$ and $S_i(\mathbf{c}^{\max})$. We see that $S_p(\mathbf{c}^{\max}) = S_p(\mathbf{c}^{\text{opt}})$ and $S_i(\mathbf{c}^{\max}) = S_i(\mathbf{c}^{\text{opt}})$. Therefore, we get $S_p(\mathbf{c}^{\text{opt}})$ and $S_i(\mathbf{c}^{\text{opt}})$ by computing $S_p(\mathbf{c}^{\max})$ and $S_i(\mathbf{c}^{\max})$. We introduce supporting lemmas before proving Theorem 2.

The next two lemmas explain how the set $S_p(\mathbf{c})$ changes when the coverage probability for a certain target decreases. The proofs can be found in the online appendix. Lemma 4 shows that if the coverage probability for a target $i \in S_p(\mathbf{c})$ decreases, then the set $S_p(\mathbf{c})$ “shrinks”. Lemma 5 shows that if the coverage probability for a target $i \in S_i(\mathbf{c})$ decreases, then the set $S_p(\mathbf{c})$ also “shrinks” but target i might be added to it.

Lemma 4. *Given coverage probability \mathbf{c} and another coverage probability \mathbf{c}' which satisfies $c'_i < c_i$ for a target $i \in S_p(\mathbf{c})$ and $c'_j = c_j, \forall j \in \mathbb{T}, j \neq i$, we have $S_p(\mathbf{c}') \subseteq S_p(\mathbf{c})$.*

Lemma 5. *Given coverage probability \mathbf{c} and another coverage probability \mathbf{c}' which satisfies $c'_i < c_i$ for a target $i \in S_i(\mathbf{c})$ and $c'_j = c_j, \forall j \in \mathbb{T}, j \neq i$, we have $S_p(\mathbf{c}') \subseteq S_p(\mathbf{c}) \cup \{i\}$.*

The next two lemmas discuss key properties of \mathbf{c}^{opt} , and the proofs are in the online appendix. Lemma 6 shows that the coverage probability for a target $i \in S_p(\mathbf{c}^{\text{opt}})$ must be $\max\{\frac{r - U_d^u(i)}{U_d^c(i) - U_d^u(i)}, 0\}$; Lemma 7 shows that the coverage probability for a target $i \in S_i(\mathbf{c}^{\text{opt}})$ is at most $\min\{1, \max\{\frac{r - U_d^u(i)}{U_d^c(i) - U_d^u(i)}, 0\}\}$. This property is used in the Reduce function that derives \mathbf{c}^{opt} from \mathbf{c}^{\max} , as well as in the proof of Theorem 2.

Lemma 6. *Given a feasible reward r , all $i \in S_p(\mathbf{c}^{\text{opt}})$ must satisfy $U_d^c(i) \geq r$ and have expected reward $\max\{U_d^u(i), r\}$ for the defender, i.e., $c_i^{\text{opt}} = \max\{\frac{r - U_d^u(i)}{U_d^c(i) - U_d^u(i)}, 0\}, \forall i \in S_p(\mathbf{c}^{\text{opt}})$.*

Lemma 7. *Given a feasible reward r , $\forall i \in S_i(\mathbf{c}^{\text{opt}})$, i has expected reward at most $\min\{U_d^c(i), \max\{U_d^u(i), r\}\}$ for the defender, i.e., $c_i^{\text{opt}} \leq \min\{1, \max\{\frac{r - U_d^u(i)}{U_d^c(i) - U_d^u(i)}, 0\}\}, \forall i \in S_i(\mathbf{c}^{\text{opt}})$.*

We are now ready to combine these preliminary lemmas to prove Theorem 2. The proof is in the online appendix.

Theorem 2. *Given a feasible reward r , $S_p(\mathbf{c}^{\max}) = S_p(\mathbf{c}^{\text{opt}})$ and $S_i(\mathbf{c}^{\max}) = S_i(\mathbf{c}^{\text{opt}})$.*

7.2 Reduce Function: Derive \mathbf{c}^{opt} from \mathbf{c}^{max}

Section 7.1 demonstrated that Algorithm 2 returns $(\infty, \infty, \dots, \infty)^\top$ if the reward r is infeasible; if the reward r is feasible, then $S_p(\mathbf{c}^{\text{max}})$ and $S_i(\mathbf{c}^{\text{max}})$ are the same as $S_p(\mathbf{c}^{\text{opt}})$ and $S_i(\mathbf{c}^{\text{opt}})$. It follows that Algorithm 4 correctly derives the optimal strategy \mathbf{c}^{opt} from \mathbf{c}^{max} .

Algorithm 4 Computing \mathbf{c}^{opt} by reducing \mathbf{c}^{max}

```

1: function REDUCE( $\mathbf{U}, \mathbf{c}^{\text{max}}, S_p(\mathbf{c}^{\text{max}}), S_i(\mathbf{c}^{\text{max}})$ )
2:    $\mathbf{c}^{\text{opt}} = \mathbf{c}^{\text{max}}$ 
3:   for every  $i \in S_i(\mathbf{c}^{\text{max}})$  do
4:      $lb \leftarrow 0, ub \leftarrow c_i^{\text{opt}}$ 
5:     while  $ub - lb \geq \delta$  do
6:        $c_i^{\text{opt}} \leftarrow \frac{lb+ub}{2}$ 
7:       if Problem (4) is feasible for target  $i$  given  $\mathbf{c}^{\text{opt}}$  then
8:          $lb \leftarrow \frac{lb+ub}{2}$ 
9:       else
10:         $ub \leftarrow \frac{lb+ub}{2}$ 
11:      end if
12:    end while
13:     $c_i^{\text{opt}} \leftarrow ub$ 
14:  end for
15:  return  $\mathbf{c}^{\text{opt}}$ 
16: end function

```

Given $\mathbf{c}^{\text{max}}, S_p(\mathbf{c}^{\text{max}})$ and $S_i(\mathbf{c}^{\text{max}})$, Algorithm 4 returns $c_i^{\text{opt}} = c_i^{\text{max}}$ for $i \in S_p(\mathbf{c}^{\text{max}})$. For $i \in S_i(\mathbf{c}^{\text{max}})$, Algorithm 4 uses binary search to find the minimum coverage probability c_i such that any further decrease⁹ in coverage probability would add target i to the set $S_p(\mathbf{c}^{\text{opt}})$. The next lemma shows that this mechanism leads to the optimal strategy \mathbf{c}^{opt} , and the proof can be found in the online appendix.

Lemma 8. *Given a feasible reward r , Algorithm 4 returns the optimal strategy \mathbf{c}^{opt} .*

Theorem 3. *Given reward r , Algorithm 2 either detects its infeasibility or provides the optimal strategy \mathbf{c}^{opt} .*

Proof. Follows from Lemmas 3 and 8.

8 Discussions

8.1 Computational Cost of BeRRA


The main computational cost of our BeRRA algorithm comes from evaluating the feasibility of the linear constraints (4), which is a linear feasibility problem and can be

⁹ δ can be arbitrarily small

solved in polynomial time. Algorithm 2 is called $\mathcal{O}(\log(\frac{1}{\epsilon}))$ times, and every call to Algorithm 2 involves solving Problem (4) $\mathcal{O}(n + |S_i(\mathbf{c}^{\max})| \log(\frac{1}{\delta}))$ times, which is bounded by $\mathcal{O}(n \log(\frac{1}{\delta}))$. Thus Problem (4) is solved $\mathcal{O}(n \log(\frac{1}{\epsilon}) \log(\frac{1}{\delta}))$ times in total for our BeRRA algorithm.

8.2 Extensions of BeRRA to General Risk Awareness

Notice that we only require \mathcal{U} to be increasing in the preceding proofs and algorithms. Thus, our BeRRA algorithm can also be used to compute the optimal robust strategy against other kinds of risk-aware attacker types, e.g., risk-seeking criminals [2, 4]. Besides, the attacker may be risk-averse for gains and risk-seeking for losses (S-shaped utility mapping function as is shown in the right figure) as is suggested in the nobel-prize-winning prospect theory (PT)¹⁰ model [5]. Here we use risk-seeking as an example to show how to apply BeRRA to other attacker types. If the attacker is risk-seeking, \hat{U} should be a strictly increasing, convex function and satisfies $\hat{U}(0) = 0$. Therefore, when adapting our BeRRA algorithm to deal with risk-seeking attackers, the only difference is in testing feasibility of constraints (4), where the condition $\hat{U} \in \mathcal{U}$ in constraints (4) should be written as:



$$\begin{aligned} \epsilon_u &\leq \frac{\hat{U}(\theta_2) - \hat{U}(\theta_1)}{\theta_2 - \theta_1} \leq \frac{\hat{U}(\theta_3) - \hat{U}(\theta_2)}{\theta_3 - \theta_2} \\ &\leq \dots \leq \frac{\hat{U}(\theta_I) - \hat{U}(\theta_{I-1})}{\theta_I - \theta_{I-1}} \\ &\hat{U}(0) = 0 \end{aligned} \quad (10)$$

8.3 Zero-sum Game

For zero-sum games, the utilities for the defender and the attacker are strongly correlated with correlation coefficient -1 , i.e., $U_a^c(i) = -U_d^c(i)$ and $U_a^u(i) = -U_d^u(i)$, $\forall i \in \mathbb{T}$. Based on this, we obtain the following theorem. The proof is in the online appendix.

Theorem 4. *For zero-sum games, the defender's Robust Stackelberg Equilibrium (RSE) strategy and Maximin strategy are the same.*

It is known that the solution concepts of Nash Equilibrium, minimax, maximin, and SSE all give the same answer for finite two-person zero-sum games. Therefore, Theorem 4 adds RSE to this equivalence list.

9 Experimental Evaluation

We will evaluate the performance of our algorithms in this section through extensive numerical experiments. This section is mainly for our new experiment results. For the

¹⁰ PT also involves a mapping of the probability. We don't consider this factor in this paper.

old results in the conference paper [14], we will provide a brief discussion here to make this paper self contained while more detailed results and discussions can be found in the online appendix. Unless otherwise stated, all of the experiment results are averaged over 20 instances. $U_d^c(i)$ and $U_a^u(i)$ are generated as random variables between 11 and 40; $U_d^u(i)$ and $U_a^c(i)$ are generated as random variables between -11 and -40 . $-\alpha$ is the correlation coefficient between $U_a^c(i)(U_a^u(u))$ and $U_d^c(i)(U_d^u(i))$. $\alpha = 1$ corresponds to zero-sum games. n is the number of targets in the game and m is the number of resources the defender has.

9.1 MIBLP vs BeRRA

The main conclusion is that BeRRA algorithm is much faster and has a higher average reward compared with MIBLP. We observe that the runtime of BeRRA increases almost linearly with the number of targets n , and the game with 50 targets only takes about 2 minutes to solve, which demonstrates BeRRA’s ability to scale up to larger problems. For MIBLP, it takes about 15 minutes for the very trivial case $n = 3$, which means it cannot scale up at all.

9.2 Performance Evaluation of RSE strategy against other attacker types

In this section, we will evaluate solution quality of the RSE strategy against risk-seeking attackers and prospect theory attackers (S-shaped utility mapping function) in detail. Since our BeRRA algorithm shows advantages in both solution quality and runtime compared with MIBLP, we use BeRRA to evaluate the performance of RSE strategy.

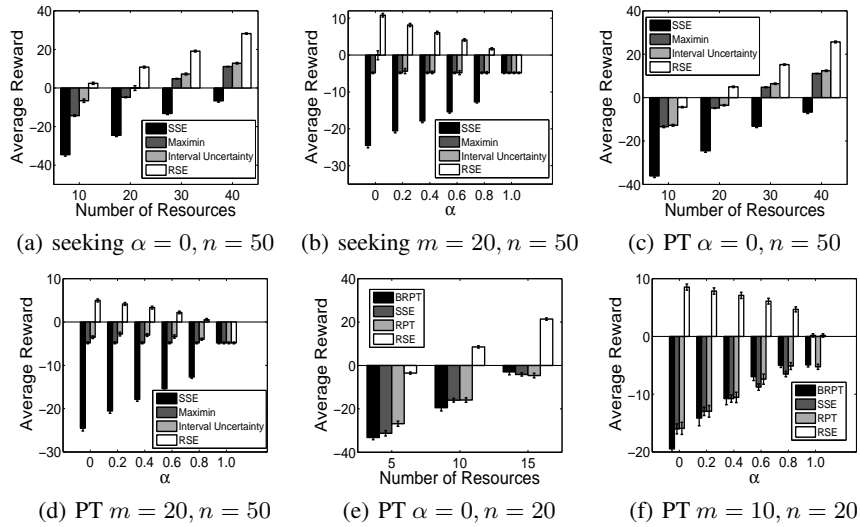


Fig. 3. Solution Quality of RSE

Figures 3(a) and 3(b) show the solution quality of RSE strategy against risk-seeking attackers. We compare its performance with the SSE strategy, Maximin strategy and

the robust strategy against interval uncertainty of $U_a^c(i)$ and $U_a^u(i)$ [6]. For values of the intervals, we tried different intervals ranging from 1 to 20 and pick the best one.

Figure 3(a) shows how the performance comparison changes with different number of resources m . The RSE strategy significantly outperforms all of the other strategies. Since the robust strategy against interval uncertainty considers some type of “robustness”, it outperforms the SSE strategy and the Maximin strategy. However, since the interval uncertainty does not fully capture the risk awareness of the attacker, it is worse than the RSE strategy. The Maximin strategy is a more conservative strategy compared with the SSE strategy, leading to better performance when compared with SSE.

Figure 3(b) shows the performance comparison with different α . The main observation is that the difference between these four strategies becomes less as α increases, and when $\alpha = 1$, these four strategies perform the same, as is proved in Theorem 4¹¹.

Figures 3(c) and 3(d) show the solution quality of RSE strategy against unknown prospect-theory attackers (S-shaped utility mapping function). They show similar patterns with the risk-seeking case. The experiment results against risk-averse attackers are also similar and can be found in the online appendix.

Figure 3(e) and 3(f) show the performance comparison of the RSE strategy and the BRPT and RPT algorithm [16] against unknown prospect theory attackers. The RSE strategy computes the robust strategy against any S-shaped utility mapping attackers (risk-averse for gains and risk-seeking for losses). BRPT algorithm computes the defender’s best response against a specific utility mapping function. RPT algorithm adds some “robustness” against interval uncertainty on the basis of the BRPT algorithm. We can see from Figure 3(e) and Figure 3(f) that the performance of BRPT and RPT is similar to that of SSE, and is much worse than the RSE strategy. The reason is that both BRPT and RPT fail to capture the uncertainty in the degree of risk-awareness.

10 Conclusion

This paper presents a model and algorithm to compute robust defender strategy in security games against risk-aware attackers with uncertainty in the degree of risk awareness. We start with risk-averse attackers and then extend the algorithm to other risk-aware attackers. We find that the intuitive MIBLP formulation only finds locally optimal solutions and is unable to scale up. Inspired by the MIBLP formulation, we develop our BeRRA algorithm which finds globally ϵ -optimal solutions by solving $\mathcal{O}(n \log(\frac{1}{\epsilon}) \log(\frac{1}{\delta}))$ linear feasibility problems. The key idea of our BeRRA algorithm is to reduce the problem to minimizing the number of resources needed for a given reward, which can be solved efficiently using special properties of the problem. In addition, we also show that we do not need to consider the attacker’s risk attitude in zero-sum games. The experimental results show the advantage of BeRRA compared with MIBLP, as well as the promising solution quality of the RSE strategy.

¹¹ The proof of Interval Uncertainty = Maximin is similar to that of Theorem 4.

Bibliography

- [1] Basilio, N., Gatti, N., Amigoni, F.: Leader-follower strategies for robotic patrolling in environments with arbitrary topologies. In: AAMAS (2009)
- [2] Block, M.K., Gerety, V.E.: Some experimental evidence on differences between student and prisoner reactions to monetary penalties and risk. *The Journal of Legal Studies* (1995)
- [3] Conitzer, V., Sandholm, T.: Computing the optimal strategy to commit to. In: EC (2006)
- [4] Grogger, J.: Certainty vs. severity of punishment. *Economic Inquiry* (1991)
- [5] Kahneman, D., Tversky, A.: Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society* (1979)
- [6] Kiekintveld, C., Islam, T., Kreinovich, V.: Security games with interval uncertainty. In: AAMAS (2013)
- [7] Kiekintveld, C., Jain, M., Tsai, J., Pita, J., Ordóñez, F., Tambe, M.: Computing optimal randomized resource allocations for massive security games. In: AAMAS (2009)
- [8] McKelvey, R.D., Palfrey, T.R.: Quantal response equilibria for normal form games. *Games and economic behavior* (1995)
- [9] Nguyen, T., Jiang, A., Tambe, M.: Stop the compartmentalization: Unified robust algorithms for handling uncertainties in security games. In: AAMAS (2014)
- [10] Nguyen, T.H., Yang, R., Azaria, A., Kraus, S., Tambe, M.: Analyzing the effectiveness of adversary modeling in security games. In: AAAI (2013)
- [11] Paruchuri, P., Pearce, J.P., Marecki, J., Tambe, M., Ordonez, F., Kraus, S.: Playing games with security: An efficient exact algorithm for bayesian stackelberg games. In: AAMAS (2008)
- [12] Phillips, P.: The preferred risk habitat of al-qa'ida terrorists. *European Journal of Economics, Finance and Administrative Sciences* (2010)
- [13] Phillips, P.J.: Applying modern portfolio theory to the analysis of terrorism. computing the set of attack method combinations from which the rational terrorist group will choose in order to maximise injuries and fatalities. *Defence and Peace Economics* (2009)
- [14] Qian, Y., Haskell, W.B., Tambe, M.: Robust strategy against unknown risk-aware attackers in security games. In: AAMAS (2015)
- [15] Von Stengel, B., Zamir, S.: Leadership with commitment to mixed strategies (2004)
- [16] Yang, R., Kiekintveld, C., Ordonez, F., Tambe, M., John, R.: Improving resource allocation strategy against human adversaries in security games. In: IJCAI (2011)
- [17] Yin, Z., Jain, M., Tambe, M., Ordonez, F.: Risk-averse strategies for security games with execution and observational uncertainty. In: AAAI (2011)
- [18] Yin, Z., Tambe, M.: A unified method for handling discrete and continuous uncertainty in bayesian stackelberg games. In: AAMAS (2012)