Imbalanced Collusive Security Games

Han-Ching Ou, Milind Tambe, Bistra Dilkina, and Phebe Vayanos

Computer Science Department,
University of Southern California
{hanchino,tambe,dilkina,phebe.vayanos
}@usc.edu

Abstract. Colluding adversaries is a crucial challenge for defenders in many real-world applications. Previous literature has provided Collusive Security Games (COSG) to model colluding adversaries, and provided models and algorithms to generate defender strategies to counter colluding adversaries, often by devising strategies that inhibit collusion [6]. Unfortunately, this previous work focused exclusively on situations with perfectly matched adversaries, i.e., where their rewards were symmetrically distributed. In the real world, however, defenders often face adversaries where their rewards are asymmetrically distributed. Such inherent asymmetry raises a question as to whether human adversaries would attempt to collude in such situations, and whether defender strategies to counter such collusion should focus on inhibiting collusion. To address these open questions, this paper: (i) explores and theoretically analyzes Imbalanced Collusive Security Games (ICOSG) where defenders face adversaries with asymmetrically distributed rewards; (ii) conducts extensive experiments of three different adversary models involving 1800 real human subjects and (iii) derives novel analysis of the reason behind why bounded rational attackers models outperform perfectly rational attackers models. The key principle discovered as the result of our experiments is that: careful modeling of human bounded rationality reveals a key difference (when compared to a model using perfect rationality) in defender strategies for handling colluding adversaries which face symmetric vs asymmetric rewards. Whereas a model based on perfect rationality always attempts to break collusion among adversaries, a bounded rationality model acknowledges the inherent difficulty of breaking such collusion in symmetric situations and focuses only on breaking collusion in asymmetric situation, and only on damage control from collusion in the symmetric situation.

Keywords: Stackelberg Security Game, Collusion, Human Behavior Model, Amazon Mechanical Turk

1 Introduction

Motivated by threats on human and cyber-physical systems, game theoretic approaches have been devised to help solve real-life security problems in many domains. [1, 13, 19] In these applications, Stackelberg security game based models have proved to be both practical and effective in many scenarios, e.g., to protect airports [17, 19], train stations [22], and even wildlife [4, 5, 21].

In order to advance these applications, collusion among adversaries in security games is a vital issue that needs to be addressed. Instead of working alone, adversaries colluding with one another often attack targets more effectively. For example, smugglers in Colombia have developed from small criminal groups to huge drug cartels by working closely since the 1950s [18]. Terrorists received material support from criminal groups and carried out more severe violent actions in the US [8]. These examples show that if we do not break the collusion in an early stage, the attackers might collude to become a stronger threat to defenders.

Previous literature has provided Collusive Security Games (COSG) to model colluding adversaries [6]. It has provided algorithms to generate defender strategies to counter the colluding adversaries in such a model. However, this previous work focused exclusively on situations with perfectly matched adversaries, i.e., where their rewards were symmetrically distributed. In the real world, however, defenders often face adversaries where rewards are asymmetrically distributed [3].

Such inherent asymmetry raises a question as to whether human adversaries would attempt to collude in such situations, and whether defender strategies to counter such collusion should focus on inhibiting such collusion. To address these open questions, this paper: (i) explores and theoretically analyzes Imbalanced Collusive Security Games (ICOSG) expand from COSG where defenders face adversaries with asymmetrically distributed rewards; (ii) conducts extensive experiments involving 1800 real human subjects of three different adversary models and (iii) derives novel analysis of the reason behind why bounded rational attackers models outperform perfectly rational attackers models. The key principle discovered as the result of our experiments is that: Careful modeling of human bounded rationality reveals a key difference (when compared to a model using perfect rationality) in defender strategies for handling colluding adversaries which face symmetric vs asymmetric rewards. Whereas a model based on perfect rationality always attempts to break collusion among adversaries, a bounded rationality model acknowledges the inherent difficulty of breaking such collusion in symmetric situations and focuses only on breaking collusion in asymmetric situation, and only on damage control from collusion in the symmetric situation.

2 Imbalanced Collusive Security Games

The Stackelberg Security Game model is widely used in both literature and security applications. [2, 11, 15, 19] Models based on it usually consist of two stages, the defender decides her strategy in the first stage. After observing the defender strategy, the attacker chooses a strategy as the best response to it in the second stage. The objective of the defender is to use a limited number m of resources to protect several targets in a set T, with each target having a different value. The attacker on the other hand seeks to attack the target that gives him the highest utility while avoiding being caught by the defender. Knowing the attacker will observe her strategy[1], the best choice for the defender is to deploy a mixed strategy that defends each target stochastically (rather than using pure strategies to always defend the same set of targets). This mixed strategy can be

[1] By convention in the security game literature, we refer to the defender as a "she" and the adversary as a "he".

viewed as a probability distribution over pure strategies. We can equivalently represent such strategy with a vector c with elements $c_t \in [0,1]$ ($\sum_t c_t = 1$) [11] denoting the probability of target t being covered. As for the attacker, after observing the defender strategy, his strategy consists of choosing a particular target to attack. We express the attacker's choice using a vector α with element $\alpha_t \in \{0,1\}$, $t \in T$, equal to 1 iff target t is chosen.

For each target, we denote the utility that the defender receives when she successfully defends the target by $U_\Theta^S(t)$. The defender commits to a strategy and the adversary observes this strategy and each select a target to attack; accordingly, if she fails to protect the target, she receives the utility $U_\Theta^F(t)$. We also denote the utility of the attacker when he successfully attacks a target (without being caught) as $U_\Psi^S(t)$; accordingly, if he gets caught attacking a target, he receives the utility $U_\Psi^F(t)$.

The expected utilities of the defender and the attacker for a given defense strategy c and attack vector α can be respectively expressed as

$$U_\Theta(c,\alpha) = \sum_{t\in T} \alpha_t \left(c_t U_\Theta^S(t) + (1-c_t)U_\Theta^F(t)\right) \tag{1}$$

$$U_\Psi(c,\alpha) = \sum_{t\in T} \alpha_t \left(c_t U_\Psi^F(t) + (1-c_t)U_\Psi^S(t)\right). \tag{2}$$

The solution concept, known as Strong Stackelberg Equilibrium (SSE), assumes that the attackers maximize their own expected utility and break ties in favor of the defender [10, 19]. In such equilibrium, the defender plays the strategy that is the best response to the attacker's strategy $\overline{c}(\overline{\alpha})$ and the attacker plays the strategy that is best response the the defender strategy $\overline{\alpha}(\overline{c})$ such that $U_\Theta(\overline{c},\overline{\alpha}) \geq U_\Theta(c,\overline{\alpha})\ \forall\ c$ and $U_\Psi(\overline{c},\overline{\alpha}) \geq U_\Psi(\overline{c},\alpha)$ $\forall\ \alpha$.

Numerous security problems from the real world have been cast in the Stackelberg Security Game framework. For example, [23, 12] provide a model where the attacker has the ability to attack multiple targets, [7] models the cooperation methods between multiple attackers and finally, [6] models the possibility of collusion between identical individual attackers. In this paper, we exploit and expand the work from [6].

In [6], the authors defined a model called Collusive Security Games (COSG) in which each attacker needs to make an extra decision of whether or not to collude with each other besides choosing which target to attack. They have introduced a new solution for COSG, termed Collusive Security Equilibrium (CSE), which generalizes the SSE. This solution concept eliminates weak equilibria while preserving properties of SSE. In CSE, the defender and attackers form a Nash equilibrium that the attackers will both choose to collude if and only if they both receive strictly higher expected utility. In addition, the attacker breaks ties between equilibria of the colluding decision in favor of the defender. Next, we expand this concept to the Imbalanced Collusive Security Games model (ICOSG), which captures the imbalance in wealth of the attackers in COSG.

2.1 Problem Formulation

In this paper, we consider a problem similar to the classic Stackelberg Security Games [19], where the defender plays as the leader of the two-stage game and needs to deploy a defense strategy before knowing the action of the attacker.

In contrast to most works in this setting, our game consist of 1 defender Θ and N attackers $\Psi_1 \dots \Psi_N$.Given a set of target $T = \{t_1, t_2, ...t_n\}$ that consist of N disjoint sets T_1,...,T_N, each attacker Ψ_i is restricted to choosing one target within his own set T_i. By choosing and successfully attacking target t without being captured, the attacker Ψ_i will receive utility $U^S_{\Psi_i}(t)$, and penalty of being capture $U^F_{\Psi_i}(t)$ otherwise. Similarly for the defender Θ, if attacker Ψ_i choose to attack target t, she receives certain utility when successfully capturing the attacker of $U^S_{\Theta}(t)$ and otherwise $U^F_{\Theta}(t)$ as the penalty of failure to protect it. Both the defender and the attacker receive zero utilities from non-attacked targets.

We introduce the notion of a *wealth index*, which captures the relative wealth of each attacker mentioned above by evaluating the portion of total utility he could earn. The wealth index of attacker i is defined as

$$\lambda_i := \frac{\sum_{t \in T_i} U^S_{\Psi_i}(t)}{\sum_{k=1}^{N} \sum_{t \in T_k} U^S_{\Psi_k}(t)} \tag{3}$$

For the following paper, to streamline the presentation, we henceforth focus on ICOGS with a single defender and two attackers with zero-sum game structures. Without loss of generality we also assume $\lambda_1 \in [0.5, 1]$ and $\lambda_2 = 1 - \lambda_1 \in [0, 0.5]$ so that the first attacker is more "powerful" (has higher relative wealth). We define the utility that the defender will receive in each case as $U^F_{\Theta}(t) = -U^S_{\Psi_i}(t) := -R(t) \leq 0$ and $U^S_{\Theta}(t) = -U^F_{\Psi_i}(t) := -P \geq 0$ of constant P for all targets t.

After deciding which target to attack, each attacker can choose if he wants to offer collusion or not. The collusion (which can be thought of as an alliance) will only be established if both attackers agree to collude. If the collusion is established, they will receive some collusion bonus δ for each successful attack, which captures the extra benefit they gain through collusion. However, they will split their total reward based on their relative wealth.

Given that the collusion between attackers may result in a higher loss, the defender's strategy is no longer simply to defend the high-risk target with more resources. For example, the defender may benefit by allocating resources in a way that breaks the willingness of the adversaries to collude by allowing them to have a higher utility from working alone while still maintaining some level of defense against these attackers.

For defender strategy, due to the same reasons as SSG, we consider only mixed defender strategies and define them as coverage vectors c. As for attacker strategies, there is always an equilibrium in which all of them play only pure strategies after observing the defender's mixed strategy [10]. We represent the attack decision over the N target sets $T_1 \dots T_N$ of the N attackers as N vectors $\alpha_1, \dots, \alpha_N$ of different length $|T_i|$, where $\alpha_i \in \{0, 1\}^{|T_i|}$ with its tth element equal to 1 iff attacker i chooses to attack target t. We also encapsulate their decision about offering collusion as β, of which $\beta_i = 1$ if Ψ_i offers collude and 0 otherwise. Finally, we can represent

the strategy of each attacker Ψ_i as the Cartesian product of the two decision vector $g_i = \alpha_i \times \beta_i \in \{0,1\}^{|T_i|} \times \{0,1\}$ or $G = \alpha \times \beta \in \{0,1\}^{|T|} \times \{0,1\}^N$ that encapsulate strategies of all the attackers as $\alpha = \left[\alpha_1^{\mathrm{T}}, \alpha_2^{\mathrm{T}} \dots \alpha_N^{\mathrm{T}}\right]^{\mathrm{T}}$ and $\beta = [\beta_1, \beta_2 \dots \beta_N]^{\mathrm{T}}$. For the sake of easier expression, we represent U_Θ^S and U_Θ^F as vectors of length $|T|$ with their tth element as $U_\Theta^S(t)$ and $U_\Theta^F(t)$ respectively. Given the strategies of the defender and attackers, the expected utility of defender can be expressed as:

$$U_\Theta(c, G) = \alpha^{\mathrm{T}}(c \circ U_\Theta^S + (\mathbb{1}_{|T|} - c) \circ U_\Theta^F) - (N - \alpha^{\mathrm{T}} c)\delta \prod_{i=1}^{N} \beta_i \tag{4}$$

where $\circ$ denotes element wise product (Hadamard product) of same length vectors and $\mathbb{1}_{|T|}$ is the element vector of length $|T|$.

As for attacker reward calculation, if any of the players refuse to collude with others, the expected utilities they received, defined as $U_{\Psi_1}, U_{\Psi_2} \dots U_{\Psi_N}$, are expressed as

$$U_{\Psi_i}(c, g_i) = \sum_{t \in T_i} \alpha_i(t) \left((1 - c_t) U_{\Psi_i}^S(t) + c_t U_{\Psi_i}^F(t)\right) \tag{5}$$

If all the players choose to collude ($\beta_i = 1 \; \forall i$), the total reward they receive is calculated as

$$U_\Psi^*(c, G) = \sum_{i=1}^{N} \sum_{t \in T_i} \alpha_i(t) \left((1 - c_t) U_{\Psi_t}^{*S}(t) + c_t U_{\Psi_i}^{*F}(t)\right) \tag{6}$$

where $U_{\Psi_i}^{*S}(t) = U_{\Psi_i}^S(t) + \delta$ and $U_{\Psi_i}^{*F}(t) = U_{\Psi_i}^F(t)$. The final reward each attacker receives depends on their wealth index.

$$U_{\Psi_i}^* = \lambda_i U_\Psi^* \tag{7}$$

Also noted that in zero-sum reward structure settings, the above equation yields to $U_\Theta = -U_\Psi^*$ for $\prod \beta_i = 1$ and $U_\Theta = -U_\Psi$ for $\prod \beta_i = 0$ in the respective colluding and non-colluding cases.

The goal is to find the optimal strategy c to maximize the expected defender utility U_Θ by breaking the collusion of the attackers while maintaining good defense. However, such strategy diverges for different attacker behavior assumptions, which will be elaborated in the following section.

2.2 Defender Strategies

Perfectly rational model (PRM) By assuming each attacker to be perfectly rational, we assume each of them selects the strategy to maximize their expected utility U_{Ψ_i}. We applied the solution concept of Collusive Security Equilibrium (CSE) used in [6]. CSE requires that (i) the defender's strategy is a best response to each attacker's strategy, (ii) the attacker strategies form a Nash Equilibrium in their game, (iii) both attackers play collude if they obtain strictly greater utility in a (collude, collude) equilibrium than (not collude, not collude) equilibrium, and (iv) the attackers break ties between equilibria which satisfy (i)-(iii) in favor of the defender.

In addition, the CSE of our problem can also be calculated by modifying the mixed integer linear program (MILP) in [6]. The MILP set is based on the ERASER formulation introduced by Kiekintveld et al. [10] that solves the equilibrium of traditional SSGs. More details can be found in [6].

This algorithm can return the CSE of any reward structure and gives us the equilibrium strategies of the defender $\overline{c}(\overline{g}_1, \overline{g}_2)$ and attackers $\overline{g}_1(\overline{c}, \overline{g}_2)$ and $\overline{g}_2(\overline{c}, \overline{g}_1)$. If the attackers select their strategies in a perfectly rational way, this method generates the optimal strategy for the defender.

Bounded rational model (BRM) In contrast to perfectly rational model, BRM assumes players perceive the utility in a bounded rational way. Instead of strictly maximizing their expected utility, it is often more effective to assume human adversaries choose strategies (i) *which grid to attack* (ii) *collude with another player or not* stochastically based on their perceived utility[14]. The features we applied to model the bounded rationality of human subjects, which were used and proven to be effective in [6] are:

1. SUQR model [16]
2. Prospect Theory [9, 20].

For the first feature, SUQR is an extension of Quantal Response (QR). Instead of expected utility, SUQR assumes humans make decisions stochastically based on their perceived utility, which is a weighted function of different factors. In addition, the bounded rationality of how people perceive probabilities is also considered using Prospect Theory (PT). PT proposes that individuals perceive the probability of success and failure in a non-linear way. Such nonlinearity can be captured by various functional forms [9, 20].

What follows are the details of how we construct and learn our BRM model. Given the defender strategy (c), reward ($U^S_{\Psi_i}$) and penalty ($U^F_{\Psi_i}$) for each target $t \in T_i$, the perceived utility of attacking it for attacker Ψ_i is defined as

$$\hat{U}^{\alpha}_{\Psi_i}(t, c) = \omega^{\alpha}_{c} \cdot \hat{c}_t(c) + \omega^{\alpha}_{R} \cdot U^S_{\Psi_i}(t) + \omega^{\alpha}_{P} \cdot U^F_{\Psi_i}(t) \tag{8}$$

Where $\hat{c}_{t_j}$ is the Prospect Theory modified perceived probability of the original probability c_{t_j}, defined as

$$\hat{c}_t = \frac{\eta c_t^{\gamma}}{\eta c_t^{\gamma} + (1 - c_t)^{\gamma}} \tag{9}$$

From the SUQR model, the probability of that the adversary Ψ_i will attack target $t \in T_i$ is given by:

$$\hat{\alpha}_i(t, c) = \frac{e^{\hat{U}^{\alpha}_{\Psi_i}(t, \hat{c})}}{\sum_{t \in T_i} e^{\hat{U}^{\alpha}_{\Psi_i}(t, \hat{c})}} \tag{10}$$

Another decision of the bounded rational attacker we need to model is the probability of collusion. Similar to attack probability, we define the perceived utility of colluding and not colluding as

$$\hat{U}^{*\beta}_{\Psi_i}(c) = \lambda_i \sum_{j=1}^{N} \frac{\omega^{\beta}_c \cdot \sum_{t\in T_j} c_t + \omega^{\beta}_R \cdot \sum_{t\in T_j} U^{*S}_{\Psi_j}(t) + \omega^{\beta}_P \cdot \sum_{t\in T_j} U^{*F}_{\Psi_j}(t)}{|T_j|} \quad (11)$$

$$\hat{U}^{\beta}_{\Psi_i}(c) = \frac{\omega^{\beta}_c \cdot \sum_{t\in T_i} c_t + \omega^{\beta}_R \cdot \sum_{t\in T_i} U^{S}_{\Psi_i}(t) + \omega^{\beta}_P \cdot \sum_{t\in T_i} U^{F}_{\Psi_i}(t)}{|T_i|} \quad (12)$$

Again, from the SUQR model, the probability adversary Ψ_i will offer collusion is given by

$$\hat{\beta}_i(c) = \frac{e^{\hat{U}^{*\beta}_{\Psi_i}(t_j,c)}}{e^{\hat{U}^{*\beta}_{\Psi_i}(t_j,c)} + e^{\hat{U}^{\beta}_{\Psi_i}(t_j,c)}} \quad (13)$$

There are a total of 5 parameters for attack probability and 3 parameters for collusion probability to be determined, which are $(\omega^{\alpha}_c, \omega^{\alpha}_R, \omega^{\alpha}_P, \eta, \gamma)$ and $(\omega^{\beta}_c, \omega^{\beta}_R, \omega^{\beta}_P)$ respectively. These parameters are estimated via Maximum Likelihood Estimation (MLE) using data collected from the human subject experiments of PRM strategy.

Note that in the bounded rational model applied for identical powerful adversaries in the previous work [6], it is assumed the grid attacking probabilities of an attacker are conditional probabilities of given his decision to collude or not and given which attacker he is. Thus it has a total number of 4×5 (4 condition of 5 parameters to model attack probabilities)+3 (parameters to model collusion probabilities)$= 23$ parameters to learn for each game. By assuming the $\hat{\alpha}$ and $\hat{\beta}$ to be independent, the modified model is able to reduce the number of parameters and still be applicable when either type of decision making data is missing for certain data points.

Given a learned parameter set and a defender strategy as input, BRM can generate the response "strategy" of a bounded rational attacker for each attacker Ψ_i, which can be expressed as the Cartesian product of the two decision probability vector of length $|T_i|$ and 1 as $\hat{g}_i = \hat{\alpha}_i \times \beta_i$ or $\hat{G} = \hat{\alpha} \times \hat{\beta}$ that encapsulate probabilities of all the attackers decisions as $\hat{\alpha} = \left[\hat{\alpha}^{\mathrm{T}}_1, \hat{\alpha}^{\mathrm{T}}_2 \ldots \hat{\alpha}^{\mathrm{T}}_N\right]^{\mathrm{T}}$ and $\hat{\beta} = \left[\hat{\beta}_1, \hat{\beta}_2 \ldots \hat{\beta_N}\right]^{\mathrm{T}}$. Given C as the feasible solution space of defender's coverage vector c, by replacing G with $\hat{G}$, we want to find $c = \arg\max_{c\in C} \left(U_{\Theta}(c, \hat{G}(c))\right)$ in equation 4, which we approximate by multiple runs of fmincon optimizer.

Simulations Figure 1 shows the simulation of the probabilities that collusion between two attackers is actually established. The number of defender resources is set to be $m = 3$. Along the rows are different wealth index combinations; along the columns are different values of delta, which gives the collusion bonus.

One interesting observation is that it is easier to break the collusion for higher wealth imbalance in both PRM and BRM simulations. In fact in PRM, there is a transition λ

λ_1-λ_2 \ δ	1	2	3	4	5	6
0.8-0.2	0	0	0	0	0	0
0.6-0.4	0	0	0	0	0	1
0.5-0.5	0	0	0	0	1	1

(a) Perfetly rational attackers

λ_1-λ_2 \ δ	1	2	3	4	5	6
0.8-0.2	0.645674	0.833459	0.907055	0.946411	0.9693	0.9826
0.6-0.4	0.64303	0.876889	0.962063	0.98841	0.9994	0.9989
0.5-0.5	0.636445	0.889326	0.972065	0.993309	0.99381	0.9986

(b) Bounded rational attackers

Fig. 1. Simulation of actual collusion probability (attacker 1 collusion offering probability $\times$ attacker 2 collusion offering probability) with different collusion bonuses (δ) and wealth indexes (λ_1 and λ_2) of the two attackers. Larger bonuses and closer wealth indexes yield higher probabilities of collusion.

that determines if the collusion is breakable or not in the CSE for some given structure, which we will elaborate on the next section.

Another interesting observation is that in the simulations, PRM always breaks the collusion between perfectly rational attackers when the collusion bonus is low, whereas BRM predicts that even with low collusion bonuses, the bounded rational attackers are still going to collude with high probability.

Our experiment focused on the first column of the simulation, which is collusion bonus $\delta = 1$ for different wealth imbalance. The bonus value is far from the value that PRM starts to give up on breaking the collusion ($\delta \geq 5$). If the human subjects are perfectly rational, PRM strategy should be able to break the collusion completely.

3 Effect of Imbalance

In this section, we provide an analysis of the effect of imbalance, and use the perfectly rational model for simplicity. For N player ICOSGs, the expected utility of the defender and each attacker can be expressed as equations 4$\sim$ 7. By using a MILP, we can solve the equilibrium strategy of the defender. However, it is complicated to analyze the effect of the parameters due to the complexity of reward structure, as the general structure does not have a closed form of utility gain.

To analyze the effect of imbalance degree λ, we start with a more straightforward case. We denote the total value of targets as $\sum_t U^F_\Theta(t) = -R_\Theta$ and $\sum_i \sum_t U^S_{\Psi_i}(t) = R_\Psi$, the reward/penalty of catching/being caught as $U^S_\Theta(t) = P_\Theta$ and $U^F_{\Psi_i}(t) = -P_\Psi$ $\forall t \in T$ for the defender and attackers respectively. All of the parameters above (R_Θ, R_Ψ, P_Θ and P_Ψ) are non-negative number to avoid confusion. For zero-sum game, $R_\Theta = R_\Psi$ and $P_\Theta = P_\Psi$. The total number of defender resources is set to be m.

3.1 Uniform Distribution Reward Structure

Assume we have a uniform distribution of value allocated on each target set T_i with density $U^S_{\Psi_1}(t) = \lambda R_\Psi$ $\forall t \in T_1$ and $U^S_{\Psi_2}(t) = (1-\lambda) R_\Psi$ $\forall t \in T_2$. We simplify each field to a single target as they all have the same utility.

Since the structure is simplified, the only decision that the attackers have to make is to collude or not. The only decision the defender has to make is how many resources to allocate to each attacker, denote as m_i. We separate the colluding (U^*_Θ and $U^*_{\Psi_i}$) and non-colluding (U_Θ and U_{Ψ_i}) case and rewrite the expected utility in equation 4$\sim$ 7 as:

$$U^*_{\Theta}(c) = -\sum_{i=1}^{N}(1-m_i)\lambda_i R_{\Theta} + mP_{\Theta} - (N-m)\delta \tag{14}$$

$$U_{\Theta}(c) = -\sum_{i=1}^{N}(1-m_i)\lambda_i R_{\Theta} + mP_{\Theta} \tag{15}$$

$$U^*_{\Psi_i}(c) = \lambda_i \left(\sum_{j=1}^{N}(1-m_j)\lambda_j R_{\Psi} - mP_{\Psi} + (N-m)\delta \right) \tag{16}$$

$$U_{\Psi_i}(c) = (1-m_i)\lambda_i R_{\Psi} - m_i P_{\Psi} \tag{17}$$

Proposition 1 *In two attackers imbalanced COSGs with uniform reward distribution and negligibly small penalties of failing the attack, the best defender strategy is to allocate all its resources to the attacker with the largest wealth index λ, regardless of other parameters.*

Proof. Suppose the optimal defender strategy is c^* in the CSE and the defender resources it deployed on Ψ_1 and Ψ_2 are m_1^* and m_2^* respectively. Without loss of generality, assume $\lambda_1 > 0.5 > \lambda_2$. We prove the proposition by showing $m_1^* \geq m_2^*$ first. Then we show that if $m_1^* \geq m_2^*$, allocating more resources to m_1^* always results in higher defender utility until Ψ_1 is fully covered.

First, we prove that $m_1^* \geq m_2^*$. Suppose $m_1^* < m_2^*$, consider another defender strategy $\overline{c}$ such that $\overline{m}_1 = m_2^*$ and $\overline{m}_2 = m_1^*$. We prove that this alternate defender strategy returns higher defender utility thus $m_1^* < m_2^*$ can not be optimal. To be clear, attacker strategy *collude* means both attackers choose to collude ($\beta_1\beta_2 = 1$) and attacker strategy *not collude* means at least one of the attackers refuse to collude($\beta_1\beta_2 = 0$). If the attacker strategies against the two defender strategies(c^* and $\overline{c}$) are both *collude*, both *not collude* or *collude* against c^* and *not collude* against $\overline{c}$, the new strategy $\overline{c}$ returns higher defender utility in all three cases thus c^* can not be the optimal strategy. Since $U^*_{\Theta}(c^*) < U^*_{\Theta}(\overline{c})$, $U^*_{\Theta}(c^*) < U_{\Theta}(\overline{c})$ and $U_{\Theta}(c^*) < U_{\Theta}(\overline{c})$ for $\lambda_1 > \lambda_2$ and $-(N-m)\delta < 0$ in equation (14) and (15).

As for the last case, the attackers playing *not collude* against c^* and *collude* against $\overline{c}$, we prove that such a scenario is not possible. The condition of breaking the collusion is $U_{\Psi_i} \geq U^*_{\Psi_i}$ for any i. In the two attackers case, since P_{Ψ} is negligibly small, it can be derived from equation 16 and 17 that the condition of breaking the collusion is to satisfy one of the following two inequalities:

$$m_2 \leq \frac{m\lambda_1\lambda_2 R_{\Psi} + m\lambda_2 P_{\Psi} - (2-m)\lambda_2\delta}{2\lambda_1\lambda_2 R_{\Psi} + P_{\Psi}} \sim \frac{m}{2} - \frac{(2-m)\delta}{2\lambda_1 R_{\Psi}} \tag{18}$$

$$m_1 \leq \frac{m\lambda_1\lambda_2 R_{\Psi} + m\lambda_1 P_{\Psi} - (2-m)\lambda_1\delta}{2\lambda_1\lambda_2 R_{\Psi} + P_{\Psi}} \sim \frac{m}{2} - \frac{(2-m)\delta}{2\lambda_2 R_{\Psi}} \tag{19}$$

Note that these two equations cannot be satisfied simultaneously since $m_1 + m_2 = m$ and only one of them can be less than $m/2$. This suggests that at least one of the attackers is willing to offer the collusion when the penalties are negligibly small. If

$m_1^* < m_2^*$ in c^* and the two attackers are not colluding, c^* must satisfy equation 19. However, since $\overline{m}_2 = m_1^*$ and $\lambda_1 > \lambda_2$, we have

$$\overline{m}_2 = m_1^* \leq \frac{m}{2} - \frac{(2-m)\delta}{\lambda_2 R_\Psi} < \frac{m}{2} - \frac{(2-m)\delta}{\lambda_1 R_\Psi}$$

thus for defender strategy $\overline{c}$ satisfy equation 18 and the collusion will be break as well.

Second, we prove that for $m_1 \geq m_2$, allocating more resources on attacker 1 before it is fully covered will result in higher defender utility. Similar to the first part of proof, suppose the optimal defender strategy is c^* and $1 > m_1^* \geq m_2^* > 0$, consider another defender strategy $\overline{c}$ such that $\overline{m}_1 = m_1^* + \epsilon$ and $\overline{m}_2 = m_2^* - \epsilon$. If c^* breaks the collusion, c^* must satisfy equation 18 as $m_1^* > m_2^*$ in the optimal strategy. Since $\overline{m}_2 < m_2^*$, the new strategy $\overline{c}$ must break the collusion as well. Again, from equation 14 and 15, $\overline{c}$ always returns higher defender utility in the other three possible cases. Thus we have proven the proposition.

Proposition 2 *Define the transition threshold λ^* as the two attackers will not collude in the equilibrium if and only if $\lambda \geq max(\lambda^*, 0.5)$ for fixed total R_Ψ. In two attackers ICOSGs with uniform distribution, assuming none of the attackers is fully covered, the transition threshold is*

$$\lambda^* = \frac{(2-m)\delta}{mR_\Psi} - \frac{P_\Psi}{R_\Psi} \tag{20}$$

Proof. By proposition 1, we can replace m_2 with 0 and m_1 with m for $m \leq 1$ in equation 18 and derive the above transition threshold.

This equation indicates whether the collusion is breakable or not in the uniform distribution game for a specific parameter set. The collusion becomes harder to break when collusion bonus δ is higher and easier to break when defender resource m, penalty P_Ψ and the wealth index of stronger attacker λ is higher. In other words, when other conditions are the same, higher wealth imbalance makes the collusion easier to break.

Unfortunately, the transition threshold λ^* does not have a closed form in general structure game. However, it is still obtainable using numerical approach.

3.2 Uniform Scale Affine Transformation Reward Structure

It is difficult to derive closed-form analysis for the reward structure of the general distribution. However, one class of distribution; which is what we used in the latter experiments as the example figure 3 shows, have some nice properties to be explored.

Definition 1. *Uniform Scale Affine Transformation Reward Structure*

Given a base reward structure with $|T|/N$ targets $t_1 ... t_{|T|/N}$ of a general distribution define as U_Ψ^S such that $\sum_t U_\Psi^S(t) = R_\Psi$. Each attacker has the same number of targets to choose from($|T_i| = |T|/N| \ \forall i$). The reward structure of each attacker is the uniform scale affine transformation of the base reward structure, in which the scale is given by $U_{\Psi_i}^S(t_j) = \lambda_i U_\Psi^S(t_j)$ for $j = 1 \cdots |T_i|$ and the penalty $U_{\Psi_i}^F(t_j) = P$ is an ignorable small constant for all targets.

Define $\Lambda = (\lambda_1, \lambda_2, ...\lambda_N)$ as a set of wealth indexes. For the same base reward structure, we define $U_{(\Theta,\Lambda)}(c, G)$ as the defender utility in game with wealth index set Λ. We decompose the first term of the right hand side in equation 4 and rewrite the equation as $U_{(\Theta,\Lambda)}(\overline{c}_\Lambda, G) = \sum_{i=1}^{N} E_\Theta(\Psi_i, \lambda_i, \overline{c}_\Lambda) - (N - \alpha^{\mathrm{T}}\overline{c}_\Lambda)\delta \prod_{i=1}^{N} \beta_i$. The term $E_\Theta(\Psi_i, \lambda_i, \overline{c}_\Lambda)$ represents the expected utility of defender gain from defending attacker Ψ_i and $\overline{c}_\Lambda$ represent the best response of the defender for game with wealth index set Λ. This term has the following three properties

1. $E_\Theta(\Psi_i, \frac{1}{N}, \overline{c}_\Lambda) \geq E_\Theta(\Psi_j, \frac{1}{N}, \overline{c}_\Lambda)$ for $m_i \geq m_j$
2. $E_\Theta(\Psi_i, \hat{\lambda}, \overline{c}_\Lambda) = \frac{\hat{\lambda}}{\lambda} E_\Theta(\Psi_i, \lambda, \overline{c}_\Lambda)$ for any $\hat{\lambda}, \lambda > 0$
3. $m_i \geq m_j$ in $\overline{c}_\Lambda$ for $\lambda_i \geq \lambda_j$

The proof of the first property is straightforward. In game with identical wealth index attackers, the more resources the defender allocate to the attacker, the more expected utility she will gain from him. As for the second term, since we have the same reward distribution and same strategy, the reward on every target is proportional to λ and yields to the expected reward proportional to λ. Finally, term 3 can be proved using the same method in the first part of proposition 1. Based on the above properties, we are now able to prove proposition 3.

Proposition 3 *In two attackers ICOSGs with uniform or zero-sum uniform scale affine transformation reward structure, a larger wealth imbalance results in a smaller defender loss.*

Proof. The uniform distribution part is straightforward, as it can be inferred that the defender in the game with larger wealth imbalance could always break the collusion if the defender in the game with lower wealth could break the collusion from proposition 2 thus its utility is higher from equation 14, 15 and proposition 1.

As for zero-sum uniform scale affine transformation reward structure, assume $\Lambda = (\lambda, (1-\lambda))$ and $\lambda \geq 0.5$, we want to prove that $U_{(\Theta,\hat{\Lambda})} > U_{(\Theta,\Lambda)}$ for any $\hat{\Lambda} = (\hat{\lambda}, (1 - \hat{\lambda})), \hat{\lambda} > \lambda$. From above properties and the fact that the defender gain higher utilities when playing the best response, assume $\beta_1\beta_2 = 0$ in both games, we have:

$$
\begin{aligned}
U_{(\Theta,\Lambda)} =& E_\Theta(\Psi_1, \lambda, \overline{c}_\Lambda) + E_\Theta(\Psi_2, (1-\lambda), \overline{c}_\Lambda) \\
=& 2\lambda E_\Theta(\Psi_1, \frac{1}{2}, \overline{c}_\Lambda) + 2(1-\lambda) E_\Theta(\Psi_2, \frac{1}{2}, \overline{c}_\Lambda) \\
<& 2\hat{\lambda} E_\Theta(\Psi_1, \frac{1}{2}, \overline{c}_\Lambda) + 2(1-\hat{\lambda}) E_\Theta(\Psi_2, \frac{1}{2}, \overline{c}_\Lambda) \\
=& E_\Theta(\Psi_1, \hat{\lambda}, \overline{c}_\Lambda) + E_\Theta(\Psi_2, (1-\hat{\lambda}), \overline{c}_\Lambda) \\
\leq& E_\Theta(\Psi_1, \hat{\lambda}, \overline{c}_{\hat{\Lambda}}) + E_\Theta(\Psi_2, (1-\hat{\lambda}), \overline{c}_{\hat{\Lambda}}) \\
=& U_{(\Theta,\hat{\Lambda})}
\end{aligned}
$$

Since $\lambda < \hat{\lambda}$, for some transition threshold λ^*, the only possible relations of the three parameters are either $\lambda \leq \hat{\lambda} \leq \lambda^*$ ($\beta_1\beta_2 = 1$ for both game), ($\lambda \leq \lambda^* \leq \hat{\lambda}$)($\beta_1\beta_2 = 1$

for Λ and $\beta_1\beta_2 = 0$ for $\hat{\Lambda}$) or $(\lambda^* \leq \lambda \leq \hat{\lambda})(\beta_1\beta_2 = 0$ for both game). Similar inequalities can be derived for remaining two cases by adding the collusion bonus term. Thus we have proved proposition 3.

4 Empirical investigation using human subjects

4.1 Imbalanced Wildlife Poaching Game

To investigate imbalance in COSG, we developed the imbalanced wildlife poaching game and asked human subjects to play the role of poachers in a national park of Africa. We recruited 1800 unique participants from Amazon Mechanical Turk (AMT) and offered them bonus rewards as an incentive for them to perform well. Figure. 2 shows the interface of the game used in our human subject experiments. We will elaborate the detail of experiment design in the following section.

Fig. 2. Imbalanced wildlife poaching game: the human subject is assigned to the left side of the park. His partner attacker is assigned to the right side. The probabilities of being caught can be observed by human subject through interacting with the game.

4.2 Human Subject Experiments

Our wildlife poaching game is a three-player security game with $|T_1| = |T_2| = 9$ targets available to each adversary. There are a total of $|T| = 18$ grids that contains some fixed number of animals. Each attacker is able to attack 9 targets in a 3×3 reward distribution. A total of two reward structures of three different wealth imbalance has been deployed on AMT as figure 3 shows. The penalty of the attacker getting caught is set to be $P = -1$. The total number of rangers is set to be $m = 3$, and the collusion bonus is set to be $\delta = 1$.

For each wave of the experiment, we deployed different defender mixed strategies against human adversaries played as each side of Game 1, Game 2 and Game 3 of both

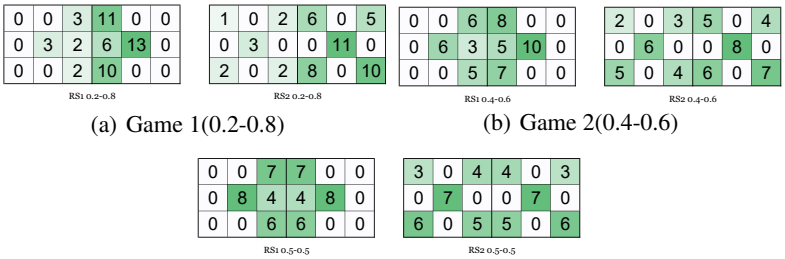

(a) Game 1(0.2-0.8) (b) Game 2(0.4-0.6)

(c) Game 3(0.5-0.5)

Fig. 3. Reward (animal density) structures of different wealth imbalance (λ_1-λ_2) deployed on AMT.

reward structures. Note that although we have a symmetric reward structure in Game 3, the defender will still deploy the defender resources asymmetrically to break the collusion.

Each participant was asked to play three rounds of carefully designed games, which are the trial game, test game, and the main game. The score that participants gain from the test game and main game (displayed as round 1 and round 2 for the participants) were accumulated as the bonus payment to the participants to incentivize the players to perform well. The bonus of each participant was calculated as $0.5 + 0.05 \times$ (*points earned in the test game and main game*), the points earned could be negative if the participant got caught in both games.

Before playing the trial game, participants were provided with a background story and detailed instructions about the game. After reading the instructions, the participants next played the trial game that has an obvious choice of the grid to attack and collusion decision to make sure they comprehended the game. When they finished the trial game, the participants could either choose to reread the instructions or begin to play the round 1 game (test game) and earn points. The test game, acts as a validation game, having an apparent yet opposite choice for the collusion decision to the test game to avoid any bias.

The test game serves two purposes. The first purpose is for us to validate if the participants understand the game or not. The data of the participant was excluded if it does not meet certain criteria in the test game. The second purpose is to balance the total reward payment of different settings in the main game. For example, the participant played as 0.2 side of game 3 has a limit potential to earn points in round 2 (the main game). Thus he/she will be assigned to a higher potential reward in round 1 (test game) to be fair and avoid bias as much as possible.

Finally, the second round game that participants played was the main game that we used to collect data of their decision making. After the game, the participants were asked to take a survey about their experience of the game and their personalities.

In each individual game, the human player is given a set amount of time to make decisions about: (i) whether to collude with the other player or not and (ii) which region of the park to place their snare. To make the first decision, a question appears on the screen which asks whether the human player is inclined to collude or not. After answer-

ing this question, a message appears on the screen that indicates whether collusion was preferred by both players or not. Collusion occurs only if it is preferred by both players. It is worth noting that the human participant has no opportunity to communicate with or learn about the other player. Next, players are asked to choose a target in their own region to attack. As before, players cannot communicate about which target to attack.

Note that whereas the human player plays as one of the adversaries, we designed a computer agent with rational behavior to play as the second adversary; thus there is an algorithm generating defender strategies, and two adversaries (one a human and one a computer agent). Choosing a computer agent as a second player let us to avoid requiring coordination between two human players in the experiments. While the other player is a computer, it is suggested to the human player that they are actually another human. The computer agent rationally chooses its decision to collude. To simplify the analysis, we assume that the second stage of decision making (where each adversary chooses a target to attack) depends on his own inclination for collusion and does not depend on the attitude of the other adversary.

There are total of 50 participants for each game set. For each strategy, 12 sets of games ($\{$Game Structure$|$ RS1, RS2$\} \times \{\lambda_1 | 0.2, 0.4, 0.5A, 0.5DA, 0.6, 0.8\}$) were deployed, in which RS1 represent reward structure 1, RS2 represent reward structure 2, $0.5A$ represents human player playing the side with less coverage in the symmetric game ($\lambda_1 = \lambda_2 = 0.5$) and $0.5DA$ represents otherwise for the sake of distinguishing. Thus, a total of 600 human players participated in the experiments for each strategy, and we deployed three strategies in total. We show and analyze our results in the following section.

4.3 Numerical Results

Three waves of experiments have been conducted. In the first wave, we deployed the optimal strategy acquired from PRM and asked human subjects to play the security poaching game described above. We collected the human decision data of the attacking decision α and the collusion decision β on these 12 sets of games.

Two models that assume the attackers are bounded rational have been learned using the data collected in the first wave. For the first BRM strategy, we used maximum likelihood estimation on all the data collected and learned the 8 parameters required to generate the strategy, which are $(\omega_C^\alpha, \omega_R^\alpha, \omega_P^\alpha, \eta, \gamma)$ and $(\omega_C^\beta, \omega_R^\beta, \omega_P^\beta)$. The model is named "BRM" in the figure below.

As for the second BRM strategy (called BRM05), we only used data collected from symmetric case in the first wave ($\{$Game Structure$|$ RS1, RS2$\} \times \{\lambda_1 | 0.5A, 0.5DA\}$) and learned the 8 parameters that generate the third strategy.

Next, two waves of experiments were deployed, one using BRM strategy and one using BRM05 strategy. Each wave involved another 600 human subjects playing the 12 sets of games. To analyze which model is the more effective one, we looked at two perspective accuracy and performance.

Accuracy The human decision data from the wave 1 experiment acts as the training data for the BRM and BRM05 models. To be fair, we compare the prediction accuracy

of the three models on wave 2 human decision data, which has not been used for the training of any of the models. In figure 4, we vary the wealth imbalance($\lambda_1 - \lambda_2$) along the x-axis and shows the prediction of each model versus the actual defender loss along y-axis in the wave 2 games we conducted. The loss here refers to the total reward the defender faces due to collusion and target choices of the attackers. Note that in this game, the penalty is constant and low; and hence the defender usually faces a loss and must reduce it as much as possible. This is intentionally designed to be consistent with real world situation involving poaching in national parks.

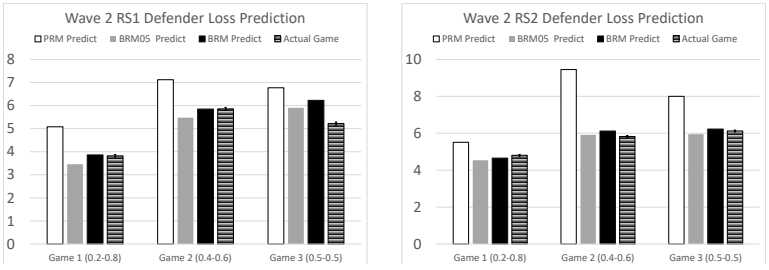

(a) RS1 (b) RS2

Fig. 4. Prediction of the defender loss by each model and the actual defender loss in wave 2 experiments. The error bars of the actual defender loss are small due to large sample size.

In terms of accuracy, both BRM models outperforms the PRM model. The inaccuracy of PRM model comes from two reasons, overestimating the defender loss from the attacker target choice α and underestimating the defender loss from collusion offering probability β. While the latter factor sometimes helps in reducing the error, overall its performance suffers compared to the BRM models.

The reason PRM overestimates the defender loss from the attacker target choice is because it assumed the attackers to be perfectly rational and always choose the grid with highest expected value to attack. However, in the experiments, we observe that human subjects avoids high-risk high-reward grid cells and choose some safer yet lower expected reward grid with high probability. The two BRM models are able to capture and exploit this along with other bounded rational behavior as explained in section 2.3 and thus leads to a relatively more accurate prediction on α.

The two BRM models also perform better than PRM in predicting the probability of collusion. In figure 5, we vary the wealth index of the attacker along the x-axis and show the prediction of collusion offering probabilities of players versus actual collusion offering probability in the wave 2 experiments along the y-axis .

In RS1, PRM predicts the strategy applied in wave 2 can always break the collusion (for situations shown in figure 5 (a)) by making the weaker attacker collude with probability 0. In the real experiments, however, human subjects still offer collusion with high probability, even if collusion results in a lower expected utility.

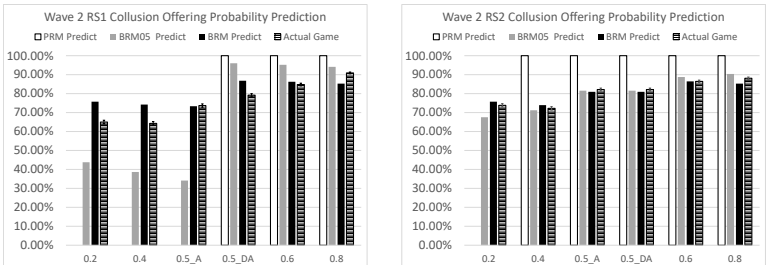

(a) RS1 (b) RS2

Fig. 5. Prediction of the collusion offering probabilities by each model and the actual collusion offering probabilities. PRM predicts 0 collusion offering probability of adversaries with wealth index 0.2, 0.4 and 0.5A in RS1 hence no bar is shown. Same for adversaries with wealth index 0.2 in RS2.

Given the prediction toward wave 2 experiments as an example, BRM are better than PRM at predicting attacker strategy. We now look into the performance of the strategy they generated in the actual experiments.

Performance In each wave, games with two reward structures of three wealth imbalance with multiple human subjects playing as both the weaker and stronger sides have been conducted with the three strategies. Figure 6 shows the performance of the three strategies in the experiments against real human attackers.

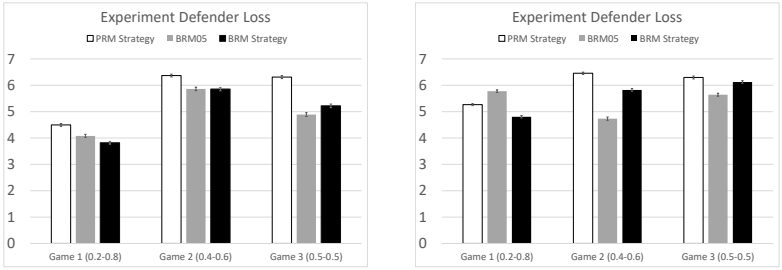

(a) RS1 (b) RS2

Fig. 6. Average defender loss from experiments with 100 participants (50 on each side) in each game per strategy (each bar). Error bars are small due to large sample size.

The BRM strategy outperforms PRM strategy in every game with different wealth imbalance of both reward structures. The error bars shown in the graph are small due to large sample size. The BRM05 however, is more unstable. It outperforms the other two

strategies in some games with lower wealth imbalance. However, it performs poorly for high wealth imbalance and even lost to PRM in RS2. This fact suggests that it could not capture some properties about the attacker behavior of the high imbalance game.

Another phenomenon worth noticing is that the defender loss did decrease as the wealth imbalance increases as proposition 3 suggested for both PRM and BRM in asymmetric games. Interestingly, BRM deployed some surprisingly different strategies when dealing with symmetrically powerful adversaries, which will be analyzed in the next section.

Strategy Difference Other than better prediction and the exploitation of bounded rational behavior when choosing grids to attack, there is another crucial reason for the BRM to perform well. Figure 7 shows the actual collusion probability of PRM and BRM strategy in real experiment.

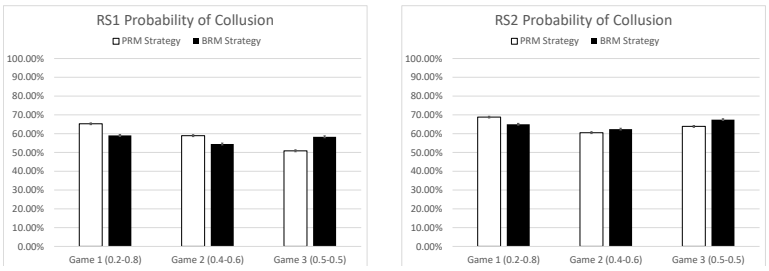

(a) RS1 (b) RS2

Fig. 7. Comparison of the actual collusion probabilities between the PRM and BRM defender strategies.

In asymmetric games, BRM is able to break the collusion with higher probabilities than PRM. Surprisingly, in the symmetric game(game 3), BRM did not break more collusion than PRM. To investigate this, we compare the strategy deployed in PRM and BRM for such game in figure 8. It can be observed that BRM did not try to break the collusion at all by defending both side symmetrically. By accepting the fact that human adversaries are still going to collude and the resources it has are too little to spare, BRM is able to keep the collusion probability within an acceptable amount without sacrificing one side too much.

As for asymmetric adversaries, figure 9 shows the difference between the strategies PRM and BRM deployed. The first thing to notice is that although the reward structure is not uniformly distributed on each side, both PRM and BRM agrees that more defender resources should be deployed on the attacker with more wealth as proposition 1. In these games, BRM tries to break the collusion harder than PRM by deploying more defender resources on stronger attacker than PRM.

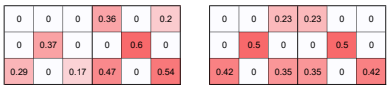

Fig. 8. Strategies of PRM and BRM for symmetric adversaries(Game 3(0.5-0.5)).

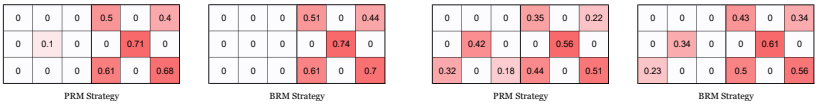

(a) Game 1(0.2-0.8) (b) Game 2(0.4-0.6)

Fig. 9. Strategies of PRM and BRM for asymmetric adversaries.

5 Conclusions

This paper modeled and addressed a security game problem that focused on breaking the collusion between asymmetric adversaries, which is often the case in the real world. Questions as to whether human adversaries would attempt to collude in such situations, and whether defender strategy to counter such collusion should focus on inhibiting such collusion were addressed in this paper by: (i) theoretically analyzing Imbalanced Collusive Security Games (ICOSG) where defenders face adversaries with asymmetrically distributed rewards; (ii) conducting extensive experiments of three different adversary models involving 1800 real human subjects and (iii) deriving novel analysis of the reason behind why bounded rational attacker models outperform perfectly rational attacker models. (iv) analyze the essential difference between balanced and imbalanced adversaries game. The key principle we found is that: Careful modeling of human bounded rationality reveals a key difference(when compared to a model using perfect rationality) in defender strategies for handling colluding adversaries which face symmetric vs asymmetric rewards. Whereas a model based on perfect rationality always attempts to break collusion among adversaries, a bounded rationality model acknowledges the inherent difficulty of breaking such collusion in symmetric situations and focuses only on breaking collusion in asymmetric situation, and only on damage control from collusion in the symmetric situation.

6 Acknowledgments

We gratefully acknowledge DARPA contract grant FA8650-15-D-6583, and subcontract from Lockheed Martin supported this research.

References

1. Basilico, N., Gatti, N., Amigoni, F.: Leader-follower strategies for robotic patrolling in environments with arbitrary topologies. In: Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 1. pp. 57–64. International Foundation for Autonomous Agents and Multiagent Systems (2009)
2. Bucarey, V., Casorrán, C., Figueroa, Ó., Rosas, K., Navarrete, H., Ordóñez, F.: Building real stackelberg security games for border patrols. In: International Conference on Decision and Game Theory for Security. pp. 193–212. Springer (2017)

3. Clark, R., Houde, J.F.: Collusion with asymmetric retailers: Evidence from a gasoline price-fixing case. American Economic Journal: Microeconomics 5(3), 97–123 (2013)
4. Fang, F., Nguyen, T.H., Pickles, R., Lam, W.Y., Clements, G.R., An, B., Singh, A., Tambe, M., Lemieux, A., et al.: Deploying paws: Field optimization of the protection assistant for wildlife security. In: AAAI. pp. 3966–3973 (2016)
5. Fang, F., Stone, P., Tambe, M.: When security games go green: Designing defender strategies to prevent poaching and illegal fishing. In: IJCAI (2015)
6. Gholami, S., Wilder, B., Brown, M., Thomas, D., Sintov, N., Tambe, M.: Divide to defend: Collusive security games. In: International Conference on Decision and Game Theory for Security. pp. 272–293. Springer (2016)
7. Guo, Q., An, B., Vorobeychik, Y., Tran-Thanh, L., Gan, J., Miao, C.: Coalitional security games. In: Proceedings of AAMAS. pp. 159–167 (2016)
8. Harms, J.: The war on terror and accomplices: An exploratory study of individuals who provide material support to terrorists. Security Journal 30(2), 417–436 (2017)
9. Kahneman, D., Tversky, A.: Prospect theory: An analysis of decision under risk. Econometrica: Journal of the Econometric Society pp. 263–291 (1979)
10. Kiekintveld, C., Jain, M., Tsai, J., Pita, J., Ordóñez, F., Tambe, M.: Computing optimal randomized resource allocations for massive security games. In: AAMAS (2009)
11. Korzhyk, D., Conitzer, V., Parr, R.: Complexity of computing optimal stackelberg strategies in security resource allocation games. In: AAAI (2010)
12. Korzhyk, D., Conitzer, V., Parr, R.: Security games with multiple attacker resources. In: IJCAI Proceedings. vol. 22, p. 273 (2011)
13. Letchford, J., Vorobeychik, Y.: Computing randomized security strategies in networked domains. Applied Adversarial Reasoning and Risk Modeling 11(06) (2011)
14. McFadden, D.L.: Quantal choice analaysis: A survey. In: Annals of Economic and Social Measurement, Volume 5, number 4, pp. 363–390. NBER (1976)
15. Nguyen, T.H., Kar, D., Brown, M., Sinha, A., Tambe, M., Jiang, A.X.: Towards a science of security games. New Frontiers of Multi-Disciplinary Research in STEAM-H (2015)
16. Nguyen, T.H., Yang, R., Azaria, A., Kraus, S., Tambe, M.: Analyzing the effectiveness of adversary modeling in security games. In: AAAI (2013)
17. Pita, J., Jain, M., Marecki, J., Ordóñez, F., Portway, C., Tambe, M., Western, C., Paruchuri, P., Kraus, S.: Deployed armor protection: the application of a game theoretic model for security at the los angeles international airport. In: Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems: industrial track. pp. 125–132. International Foundation for Autonomous Agents and Multiagent Systems (2008)
18. Restrepo, A.L., Guizado, Á.C.: From smugglers to warlords: twentieth century colombian drug traffickers. Can. J. Lat. Am. Caribb. Stud. 28(55-56), 249–275 (2003)
19. Tambe, M.: Security and game theory: algorithms, deployed systems, lessons learned. Cambridge University Press (2011)
20. Tversky, A., Kahneman, D.: Advances in prospect theory: Cumulative representation of uncertainty. Journal of Risk and uncertainty 5(4), 297–323 (1992)
21. Yang, R., Ford, B., Tambe, M., Lemieux, A.: Adaptive resource allocation for wildlife protection against illegal poachers. In: Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems. pp. 453–460. International Foundation for Autonomous Agents and Multiagent Systems (2014)
22. Yin, Z., Jiang, A.X., Johnson, M.P., Kiekintveld, C., Leyton-Brown, K., Sandholm, T., Tambe, M., Sullivan, J.P.: Trusts: Scheduling randomized patrols for fare inspection in transit systems. In: IAAI (2012)
23. Yin, Z., Korzhyk, D., Kiekintveld, C., Conitzer, V., , Tambe, M.: Stackelberg vs. nash in security games: Interchangeability, equivalence, and uniqueness. In: AAMAS (2010)