Who Should Pay the Cost: A Game-theoretic Model for Government Subsidized Investments to Improve National Cybersecurity

Xinrun Wang[1] , **Bo An**[1] and **Hau Chan**[2]

[1]School of Computer Science and Engineering, Nanyang Technological University, Singapore
[2]University of Nebraska-Lincoln, Lincoln, NE, USA
{xwang033, boan}@ntu.edu.sg, hchan3@unl.edu

Abstract

Due to recent cyber attacks, cybersecurity is becoming more critical. A single attack (e.g., WannaCry ransomware attack) can cause as much as $4 billion in damage. However, the cybersecurity investment by companies is far from satisfactory. Therefore, governments (e.g., in the UK) launch grants and subsidies to help companies to boost their cybersecurity to create a safer national cyber environment. Computing the optimal allocation is challenging due to limited subsidies, the interdependence between companies and the presence of strategic cyber attackers. To tackle the government's allocation problem, we introduce a Stackelberg game model where the government first commits to an allocation and the companies/users and attacker simultaneously determine their protection and attack (pure or mixed) strategies, respectively. For the pure-strategy case, while there may not be a feasible allocation in general, we prove that computing an optimal allocation is NP-hard and propose a linear reverse convex program when the attacker can attack all users. For the mixed-strategy case, we show that there is a polynomial time algorithm to find an optimal allocation when the attacker has a single-attack capability. We then provide a heuristic algorithm, based on best-response-gradient dynamics, to find an effective allocation in general settings. Experimentally, we show that our heuristic algorithm is effective and significantly outperforms baselines on synthetic and real data.

1 Introduction

Cybersecurity has become one of the most important issues of the modern society, and cyber attacks can cause extremely high loss to many companies, organizations and governments. For instance, the 2017 WannaCry ransomware attack affected more than 300,000 computers across 150 countries with a total damage $4 billion [CBS, 2017]. It is obviously clear that, from these recent cyber attacks, businesses are not doing enough to protect themselves and/or do not have sufficient awareness of cyber threats. In fact, 80% of over 400 global companies do not know where their sensitive data is located and how to secure it, and 60% of them do not protect their privileged accounts adequately [Thycotic, 2017]. Because companies and organizations are interdependent in cyber space, cyber attacks could propagate from one company to other organizations. Thus, the lack of protections against cyber threats not only causes damages to companies and organizations but also decreases the national cybersecurity level[1].

To combat these issues, governments around the world have launched various initiatives to improve the national cybersecurity level. For example, the UK government, one of the pioneering countries on national cybersecurity practice, provided millions of subsidies, i.e., grants and vouchers, for businesses to boost their cybersecurity in 2015 and 2016 [UK, 2015]. The UK government's goal is [UK, 2016]

> *to intervene more actively and use increased investment, while continuing to support market forces to raise cyber security standards across the UK.*

The government's goal is optimally assigning limited subsidies to companies to protect against the attacker. However, this assignment task is highly challenging because cyber companies are self-interested and interdependent (induced by the spreadability of cyber attacks) when making cyber protection decisions. In this work, we adopt game-theoretic methodologies to analyze cyber interactions and help the government to improve the national cybersecurity through subsidies.

Our contributions. First, we formulate our cybersecurity setting as a Stackelberg game played between the government, interdependent companies/users and an attacker where the government moves first to allocate subsidies and the users and attacker move simultaneously to determine their protection and attack (pure or mixed) strategies, respectively. We comprehensively investigate three settings where the attacker has different capabilities (i.e., attack all users, a single user and multiple users). For the pure-strategy case, we show that computing an optimal allocation is NP-hard when the attacker is able to attack all users. We propose a linear reverse convex program to compute an optimal allocation in this setting. However, we show that there may not be a feasible allocation in the two other settings. For the mixed-strategy case, we show that there is a polynomial time algorithm to find an optimal allocation in the single-attack setting. We then provide a

[1] The national cybersecurity level in this paper is defined as the social welfare for companies across a country against cyber crimes.

heuristic algorithm, based on best-response-gradient dynamics, to find an effective allocation in the general setting. We extensively evaluate our model and the heuristic algorithm on synthetic and real data. The results show that our model captures the interdependent behaviors of users and our heuristic algorithm is effective and outperforms other baselines.

2 Related Works

There are several lines of the research related to this work. The first line is *interdependent security game* [Kunreuther and Heal, 2003; Kearns and Ortiz, 2004] where users are interdependent and can experience direct risk from internal contamination and indirect risk transferred from their neighbors and is widely applied to security settings such as airline baggage security, fire safety and computer virus. An extension of these works is *interdependent defense game* [Chan *et al.*, 2012; Chan *et al.*, 2017] where a strategic attacker is added into the game. However, they focus on the computation of Nash equilibria between users and the attacker and there is no *defender* who tries to optimize the global security level.

The second line of the related research is *interdependent information security game* [Laszka *et al.*, 2015]. The famous Gordon-Loeb model [Gordon and Loeb, 2002] is a single-agent decision model, whose parameters are similar to our model, and its extension [Gordon *et al.*, 2003] studies the information sharing between two agents, which differs from our model. Various schemes are proposed to improve the security level, which are based on either game-theoretic equilibrium improvements [Jiang *et al.*, 2011] or mechanisms such as mandatory or optional insurances [Böhme *et al.*, 2010], subsidies and fines [Grossklags *et al.*, 2010] and regulations [Omic *et al.*, 2009]. However, the mechanisms in these works are modeled as a part of agent's utility explicitly, which is not determined by a strategic agent as in our setting. Besides, there is no strategic attacker in all these works.

The third line of the related research is *Stackelberg security game* [Tambe, 2011; Fang *et al.*, 2016] for security problems where the defender allocates limited resources to protect valuable targets against the attacker. Some recent works extended the method to cybersecurity problems, such as allocating cyber alerts [Schlenker *et al.*, 2017] and deceiving cyber adversaries [Schlenker *et al.*, 2018]. Some works consider the games where multiple types of independent followers following a known distribution [Paruchuri *et al.*, 2008], the interdependence between targets [Vorobeychik and Letchford, 2015], the externalities of the protection [Gan *et al.*, 2015] and network security games [Guo *et al.*, 2016]. However, the interdependence between self-interested followers is not tackled in all previous works. Some works consider the multiple defenders against an attacker [Gan *et al.*, 2018], which also cannot be applied to our problem due to the lack of the global optimizer. Recent works [Basilico *et al.*, 2016; Coniglio *et al.*, 2017] proposed the game between a leader and multiple followers. Their algorithms for the optimistic case (corresponding to our model) needs to enumerate all pure strategies of followers which is impossible in our case because we adopt a more compact representation of the game.

3 Motivation

We use the UK as a motivating example to illustrate the applicability of our model. From 2015 to 2020, the UK government will invest £1.9 billion to improve its cybersecurity through various schemes, such as grants, vouchers and subsidies [UK, 2015]. The UK government also builds the national cyber security centre (NCSC)[2] to manage cyber incidents, analyze cyber threats and provide tailored expertise support to businesses [UK, 2016]. With the help of NCSC, companies can obtain accurate information of cyber space and learn optimal cyber investment strategies. The information includes the vulnerabilities of businesses and the connections between businesses. The information is provided with the existence of the strategic cyber attacker (e.g., hacktivist) [UK, 2016].

We assume that companies are self-interested and there is no cooperation or coordination among them. It is worth noting that it is difficult for the attacker to observe companies' strategies before the attacks due to the lack of transparency of the cyber space. For example, the attacker cannot observe anti-spam filtering systems of companies before sending spam emails in phishing attacks. Therefore, we assume that companies and the attacker move simultaneously. We note that Nash equilibrium is a canonical solution concept for non-cooperative simultaneous-move settings.

4 Cybersecurity Investment Game

A **S**tackelberg **C**ybersecurity **I**nvestment **G**ame (SCIG) is played between a *government* (e.g., cybersecurity agency), a set of interdependent *users* (e.g., companies and organizations) and an *attacker*. The whole procedure of the game can be divided into two stages: i) the government allocates budgeted subsidies to users, ii) each user, after obtaining the government's subsidy amount, and the attacker decide their own strategies simultaneously. Both users and the attacker are termed as *followers*. The interactions between followers exactly follow the assumptions and models in *interdependent security (defense) game*, which is widely investigated in [Kunreuther and Heal, 2003; Kearns and Ortiz, 2004; Chan *et al.*, 2012; Chan *et al.*, 2017].

Formally, we consider N interdependent users, of which each user $i \in [N] = \{1, 2, ..., N\}$ is characterized by a tuple with parameters $\langle p_i, c_i, l_i \rangle$ where p_i is the probability that user i will be compromised if being attacked by the attacker, i.e., *direct attack*, c_i is the cost of user i to get a cybersecurity system such as installing firewalls, building intrusion detection systems and investigation programs, to prevent himself from the contamination, and l_i is the loss user i may suffer if he is compromised. To avoid the trivial case, we assume that $c_i < p_i l_i, \forall i \in [N]$. We use $\boldsymbol{q} = \langle q_{ji} \rangle$ to model the spread of the attack between users. For each pair $j, i \in [N], j \neq i$, let q_{ji} denote the probability that user i is compromised as a result of a transfer of the attack from j, i.e., *indirect attack*.

Strategies. The government's strategy is denoted by the vector $\boldsymbol{x} = \langle x_i \rangle$ where x_i is the subsidy assigned to the user i, constrained by the budget B (i.e., $\sum_{i=1}^{N} x_i \leq B$). The subsidies can only be used to invest in cybersecurity. The

[2] https://www.ncsc.gov.uk/

users' pure strategies are denoted by $\boldsymbol{a} = \langle a_i \rangle$ where $a_i = 1$ if user i invests, otherwise $a_i = 0$. The users' mixed strategies are denoted by $\boldsymbol{y}$ where $y_i \in [0, 1]$ is the probability that user i will invest in cybersecurity. The attacker's pure strategy is denoted by the vector $\boldsymbol{b} = \langle b_i \rangle$ where $b_i = 1$ implies the attacker attacks user i, otherwise $b_i = 0$. Note that $\sum_{i=1}^{N} b_i \leq K$ where K is the number of users the attacker can attack. We use $\mathcal{B}^K$ to denote the set of pure strategies of the attacker. The mixed strategy of the attacker is denoted by $\boldsymbol{z}$ which is a distribution over all pure strategies in $\mathcal{B}^K$.

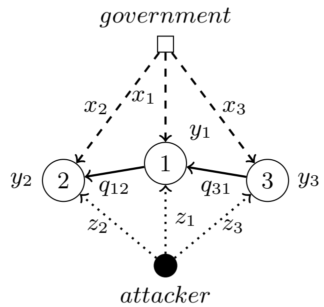

Figure 1: An illustrative example of SCIG where the square node is the government who assigns the subsidy to users (circle nodes) first, and users and the attacker (filled circle node) make their decisions simultaneously. Solid arrows between users indicate indirect risks.

Utilities. Given a strategy profile $\langle \boldsymbol{x}, \boldsymbol{a}, \boldsymbol{b} \rangle$ where the followers take pure strategies, the users' utilities are defined as

$$U_i^u(\boldsymbol{x}, \boldsymbol{a}, \boldsymbol{b}) = a_i(x_i - c_i) - \left[1 - (1 - b_i p_i)^{(1-a_i)}\right] l_i$$
$$-[(1-b_i p_i)^{(1-a_i)}] \left[1 - \prod\nolimits_{j \neq i} (1 - b_i q_{ji})^{(1-a_j)}\right] l_i. \quad (1)$$

The first term of Eq.(1) is the investment cost of user i with the subsidy x_i assigned by the government where the user will obtain the subsidy if he invests, i.e., $a_i = 1$. The second term is the expected loss to i due to the direct attack, i.e., $b_i p_i l_i$ if $a_i = 0$ and 0 if $a_i = 1$. The third term is the expected loss due to the indirect attack transferred from others, i.e., $(1 - b_i p_i) \left[1 - \prod_{j \neq i} (1 - b_i q_{ji})^{(1-a_j)}\right] l_i$ if $a_i = 0$ and $\left[1 - \prod_{j \neq i} (1 - b_i q_{ji})^{(1-a_j)}\right] l_i$ if $a_i = 1$ where $1 - \prod_{j \neq i} (1 - b_i q_{ji})^{(1-a_j)}$ is the probability that there is at least one of user i's neighbors will transfer the contamination to him. The utilities (omitting the x_i's) and terms are defined the same way as in [Kearns and Ortiz, 2004; Chan *et al.*, 2017]. The government's utility is defined to be the negative of the summation of the expected loss of users

$$U^d(\boldsymbol{x}, \boldsymbol{a}, \boldsymbol{b}) = \sum\nolimits_{i=1}^{N} [U_i^u(\boldsymbol{x}, \boldsymbol{a}, \boldsymbol{b}) - a_i(x_i - c_i)] \qquad (2)$$

We use the user's utilities to define the government's utility for simplicity and note that the government's utility is independent of the terms $a_i(x_i - c_i)$, which are canceled out when substituting Eq.(1) into Eq.(2). The attacker's utility is defined as $U^a(\boldsymbol{x}, \boldsymbol{a}, \boldsymbol{b}) = -U^d(\boldsymbol{x}, \boldsymbol{a}, \boldsymbol{b})$. Analogously, given a strategy profile $\langle \boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z} \rangle$ where the followers take mixed strategies, the users' utility is

$$U_i^u(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z}) = y_i[(x_i - c_i) - \Psi_i^1 l_i] - (1 - y_i)[\Psi_i^2 p_i + \Psi_i^1] l_i$$

The first term $y_i[(x_i - c_i) - \Psi_i^1 l_i]$ is the user's utility when he chooses to invest and $\Psi_i^1 = \sum_{\boldsymbol{b} \in \mathcal{B}^K} z(\boldsymbol{b}) \left[1 - \prod_{j \neq i} (1 - (1 - y_j) b_j q_{ji})\right]$ is the indirect risk of the user where $1 - \prod_{j \neq i} (1 - (1 - y_j) b_j q_{ji})$ is the probability that the user will be compromised when other users take $\boldsymbol{y}$ and the attacker takes $\boldsymbol{b}$, i.e., the probability that there is at least one of the neighbors of user i will transfer the contamination to him. The second term $-(1 - y_i)[\Psi_i^2 p_i + \Psi_i^1] l_i$ is the user's utility when he chooses not to invest where Ψ_l^1 is the same as introduced before and $\Psi_i^2 = \sum_{\boldsymbol{b} \in \mathcal{B}^K} z(\boldsymbol{b}) b_i \left[\prod_{j \neq i} (1 - (1 - y_j) b_j q_{ji})\right]$ denotes the probability that user i is attacked and there is no neighbor of user i will transfer the contamination to him because for $\boldsymbol{b} \in \mathcal{B}^K$, if $b_i = 1$, which means that user i is attacked and we add $\prod_{j \neq i} (1 - (1 - y_j) b_j q_{ji})$ into Ψ_i^2, which is the probability that no user transfers the contamination to user i. We can check all the three cases of the user i will suffer when he does not invest to verify the correctness that $\Psi_i^2 p_i + \Psi_i^1$, which are i) user i is successfully attacked, $\sum_{\boldsymbol{b} \in \mathcal{B}^K} z(\boldsymbol{b}) b_i p_i$, no matter whether the neighbors transfer the contamination to user i or not because the user can only be compromised once, ii) user i is attacked but not compromised and there is at least one of user i's neighbors transfers the contamination to him, $\sum_{\boldsymbol{b} \in \mathcal{B}^K} z(\boldsymbol{b}) b_i (1 - p_i) \left[1 - \prod_{j \neq i} (1 - (1 - y_j) b_j q_{ji})\right]$ and iii) user i is not attacked and there is at least one of user i's neighbors transfers the contamination to him, $\sum_{\boldsymbol{b} \in \mathcal{B}^K} z(\boldsymbol{b}) (1 - b_i) \left[1 - \prod_{j \neq i} (1 - (1 - y_j) b_j q_{ji})\right]$. Readers can obtain the term $\Psi_i^2 p_i + \Psi_i^1$ by summing all three cases. The government's utility for mixed strategies $U^d(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z})$ is defined analogously, also the attacker's utility $U^a(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z})$.

Equilibrium. Our objective is to find the optimal strategy for the government to assign the subsidy to the users, as well as the users' and the attacker's strategies. In particular, we are interested in the Stackelberg equilibrium between the government and followers (the users and the attacker), and Nash equilibrium between followers. We denote this notion as **S**tackelberg-**P**ure-**N**ash **E**quilibrium (SPNE) when the followers take pure strategies and **S**tackelberg-**M**ixed-**N**ash **E**quilibrium (SMNE) when the followers take mixed strategies, respectively. We assume that users break ties in favor of the government, i.e., if there are multiple equilibria, followers would select the one maximizing the government's utility. This is the standard assumption in Stackelberg security game [Tambe, 2011]. As the attacker's utility is the negation of the government's utility, the attacker does not need to break ties.

5 Solution Approach

In this section, we consider computing SPNE and SMNE in three settings where the attacker has different capabilities. We first show that SPNE always exists in all-user attacks and does not exist in the other two settings. We show that computing an SPNE is NP-hard and propose a linear reverse convex formulation to compute an SPNE. In the single-user attack setting, we provide an exact polynomial algorithm to compute an SMNE. Then, we provide an effective heuristic algorithm to compute an SMNE in all-user and multiple-user attacks.

5.1 All-user Attack: $K \geq N$

In this section, we consider the case where the attacker is able to attack all users. We note that the attacker will always attack all users due to the non-negative utilities an attack can obtain. We focus on the computation of SPNE in this section. SMNE will be discussed in the multiple-user attack case.

Given the strategy $\boldsymbol{x}$, there is at least one PNE, which can be computed in time $O(N^2)$, while computing all PNE of users is NP-complete [Kearns and Ortiz, 2004]. Therefore, there always exists at least a PNE for all-user attacks. However, Theorem 1 proves that computing an SPNE is NP-hard.

Theorem 1. *Computing an SPNE is NP-hard.*

Proof. We reduce from the knapsack problem which is known to be NP-hard. In a knapsack problem, we are given a set of items $[N]$ and a budget $W > 0$. Each item $i \in [N]$ has a value $f_i > 0$ and weight $w_i > 0$. The goal is to find $S \subseteq [N]$ such that $\sum_{i \in S} w_i \leq W$ and $\sum_{i \in S} f_i$ is maximum.

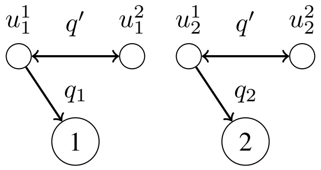

Figure 2: Reduction of the Knapsack problem to SCIG

We reduce the knapsack problem to our problem. For each item $i \in [N]$, we introduce a user i with $\langle c_i, p_i, l_i \rangle$ and two additional connected users, i.e., u_i^1 and u_i^2 and connects u_i^1 to user i, as displayed in Figure 2. For pairs of users u_i^1 and u_i^2, we set $c_i^1 = c_i^2 = c'$ and $p_i^1 = p_i^2 = p'$, $l_i^1 = l_i^2 = l'$ and $q_i^{12} = q_i^{21} = q'$. We set that $p'(1-q')l' < c' < q'l'$ and $c' - p'(1-q')l' > B$, i.e., any of the additional users will not invest even with the government's subsidy. For each user $i \in [N]$, the connection between u_i^1 and i is denoted as q_i and we set $p_i(1-q_i)l_i < c_i < p_i l_i$, $c_i - p_i(1-q_i)l_i = w_i$ and $p_i(1-q_i)l_i = f_i$, i.e., by assigning w_i to user i, the user would invest and the government gains a positive utility $p_i(1-q_i)l_i$. The government's budget B is set to be equal to W. Given a solution of the knapsack problem $S \subseteq [N]$, the solution is also ensured to be optimal to our problem, and vice versa. Note that the government will never assign the subsidy to the additional users. □

We can formulate the problem to compute an SPNE as Program 3, which is a bilevel optimization problem:

$$\max\nolimits_{\boldsymbol{x},\boldsymbol{a}} U^d(\boldsymbol{x},\boldsymbol{a},\mathbf{1}) \tag{3a}$$

$$\text{s.t. } a_i \in \arg\max\nolimits_{a_i \in \{0,1\}} \{U_i^u(\boldsymbol{x},\boldsymbol{a},\mathbf{1})\} \tag{3b}$$

$$\sum\nolimits_{i=1}^{N} x_i \leq B \tag{3c}$$

To make the problem computable, we reduce the bilevel formulation to single level optimization by rewriting each user i's utility as

$$U_i^u(\boldsymbol{x},\boldsymbol{a},\mathbf{1}) = a_i(x_i - c_i) - l_i$$

$$+ (1-p_i)^{(1-a_i)} \prod\nolimits_{j \neq i} (1-q_{ji})^{(1-a_j)} l_i \tag{4}$$

$$= a_i \left[x_i - c_i + p_i l_i \prod\nolimits_{j \neq i} (1-q_{ji})^{(1-a_j)} \right] + \Omega_i \tag{5}$$

Eq.(4) can be easily obtained from Eq.(1). The reduction from Eq.(4) to Eq.(5) is due to the fact that $(1-p_i)^{(1-a_i)} = 1-(1-a_i)p_i$ when $a_i \in \{0,1\}$ and $\Omega_i = (1-p_i)\prod_{j \neq i}(1-q_{ji})^{(1-a_j)} - l_i$, which is independent of a_i and can be ignored when users maximize their own utilities. Therefore, we introduce Program 6 which is a single level optimization problem.

$$\max\nolimits_{\boldsymbol{x},\boldsymbol{a}} U^d(\boldsymbol{x},\boldsymbol{a},\mathbf{1}) \tag{6a}$$

$$\text{s.t. } \left(x_i - c_i + p_i l_i \prod\nolimits_{j \neq i} (1-q_{ji})^{(1-a_j)} \right) a_i \geq 0 \tag{6b}$$

$$\sum\nolimits_{i=1}^{N} x_i \leq B \tag{6c}$$

$$a_i \in \{0,1\} \tag{6d}$$

Eq.(6b) ensures that if $x_i - c_i + p_i l_i \prod_{j \neq i}(1-q_{ji})^{(1-a_j)} < 0$, the user will not invest, i.e., $a_i = 0$, otherwise $a_i = 1$, which is straightforward from Eq.(5). Theorem 2 proves that Program 6 can be reformulated as a linear reverse convex program with a linear objective and constraints as $g(x) \geq 0$ where $g(x)$ is a convex function [Horst and Tuy, 2013].

Theorem 2. *Program 6 can be reformulated as a linear reverse convex program.*

Proof. The utility function of user i can be reformulated as

$$U_i^u(\boldsymbol{x},\boldsymbol{a},\mathbf{1}) = a_i(x_i - c_i) + \prod\nolimits_{j \in [N]} (1-q_{ji})^{(1-a_j)} l_i - l_i \tag{7}$$

For simplicity, we denote p_i as q_{ii}. Thus, the government's utility, i.e., the objective of Program 6, is reformulated as

$$U^d(\boldsymbol{x},\boldsymbol{a},\mathbf{1}) = \sum\nolimits_{i \in [N]} \left[\prod\nolimits_{j \in [N]} (1-q_{ji})^{1-a_j} \right] - \sum\nolimits_{i \in [N]} l_i.$$

It can be easily verified by computing the Hessian matrix of $U^d(\boldsymbol{x},\boldsymbol{a},\mathbf{1})$ that $U^d(\boldsymbol{x},\boldsymbol{a},\mathbf{1})$ is convex and $\sum_{i \in [N]} l_i$ is constant. Then, Eq.(6b) can be equivalently rewritten as

$$v_i - a_i c_i + z_i \geq 0 \tag{8a}$$

$$0 \leq v_i \leq x_i \tag{8b}$$

$$x_i - (1-a_i)M \leq v_i \leq a_i M \tag{8c}$$

$$0 \leq z_i \leq p_i l_i \cdot w_i \tag{8d}$$

$$p_i l_i \cdot w_i - (1-a_i)M \leq z_i \leq a_i M \tag{8e}$$

$$w_i \leq \prod\nolimits_{j \neq i} (1-q_{ji})^{1-a_j} \tag{8f}$$

Eqs.(8b)-(8c) ensure that if $a_i = 0$, $v_i = x_i$ and $v_i = 0$ otherwise. Eq.(8d)-(8e) ensure that if $a_i = 0$, $z_i = p_i l_i \cdot w_i$ and $v_i = 0$ otherwise, where M is a big constant. To maximize the government utility, w_i will always be equal to $\prod_{j \neq i}(1-q_{ji})^{1-a_j}$ in Eq.(8f), which is a reverse convex constraint. Besides, as the variables a_i are binary, we can also add a reverse convex constraint into the program for each a_i

$$a_i^2 - a_i \geq 0, 0 \leq a_i \leq 1 \tag{9}$$

Furthermore, we introduce an auxiliary variable U as the linear objective with an additional reverse convex constraint

$$U \leq \sum\nolimits_{i \in [N]} \left[\prod\nolimits_{j \in [N]} (1-q_{ji})^{1-a_j} \right] l_i - \sum\nolimits_{i \in [N]} l_i \tag{10}$$

Thus, we obtain the linear reverse convex program. □

The reformulation in Theorem 2 ensures that all terms in the objective and constraints are convex, which makes it easier to solve. In this work, we use global optimization solvers, e.g., BARON, to solve the problem.

5.2 Single-user Attack: $K = 1$

We consider the game where the attacker can only attack one user. The setting without the government is extensively investigated in [Chan *et al.*, 2012; Chan *et al.*, 2017], which provides some useful results for us to develop our algorithm. The following proposition proves that, except the case where all users invest, there is no other pure NE between followers.

Proposition 3. *Given the strategy $\boldsymbol{x}$, if $\langle \boldsymbol{a}, \boldsymbol{b}\rangle = \langle \boldsymbol{1}, *\rangle$ is not an NE, then there is no pure NE between followers.*

Proof. As there is at most a user being attacked, we can write the attacker's utility as

$$U^a(\boldsymbol{x}, \boldsymbol{a}^*, \boldsymbol{b}^*) = \max_{i\in[N]} \left\{(1 - a^*_i)\left(p_i l_i + \sum\nolimits_{j\neq i} q_{ij} l_j\right)\right\}.$$

If $U^a(\boldsymbol{x}, \boldsymbol{a}^*, \boldsymbol{b}^*) = 0$, it is the trivial case which means that all users invest under the given strategy $\boldsymbol{x}$, i.e., $\langle \boldsymbol{a}, \boldsymbol{b}\rangle = \langle \boldsymbol{1}, *\rangle$ is an NE where $*$ means the attacker's strategy can be any valid pure strategy because all pure strategies bring the same utility to the attacker, i.e., 0. We can easily check whether the strategy $\boldsymbol{x}$ can make $\langle \boldsymbol{a}, \boldsymbol{b}\rangle = \langle \boldsymbol{1}, *\rangle$ as an NE or not. For the case where $U^a(\boldsymbol{x}, \boldsymbol{a}^*, \boldsymbol{b}^*) > 0$, suppose that $b^*_k = 1$ for some $k \in [N]$, which implies that $a^*_k = 0$, i.e., user k does not invest. As $a^*_k = 0$ is user k's best-response to the attacker, we have $c_k \geq x_k + p_k l_k$. As we have $x_k \geq 0$, we obtain $c_k \geq p_k l_k$, which contradicts our assumption $c_i < p_i l_i$. □

The case that all users invest occurs when the budget is high, which is not the case in general. Therefore, we then focus on the mixed NE between followers. With a slight abuse of notation, we denote z_i as the probability of attacking i. Therefore, we can rewrite the user's utility as

$$\begin{aligned} U^u_i(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z}) &= y_i[(x_i - c_i) - \Psi^1_i l_i] - (1 - y_i)[\Psi^2_i p_i + \Psi^1_i] l_i \\ &= y_i(x_i - c_i) - \Psi^1_i l_i - (1 - y_i)\Psi^2_i p_i l_i \quad (11) \\ &= y_i(x_i - c_i + z_i p_i l_i) - \Psi^1_i l_i - \Psi^2_i p_i l_i \quad (12) \end{aligned}$$

where Ψ^1_i and Ψ^2_i are defined previously. The reduction from Eq.(11) to Eq.(12) is based on the key observation that $\Psi^2_i = z_i$, so we can ignore the last two terms of Eq.(12) which are independent of the user's strategy. Then, we define that $\Delta_i = \frac{c_i - x_i}{p_i l_i}$ and $L_i = p_i l_i + \sum_{j\neq i} q_{ij} l_j$. We denote that $V(\boldsymbol{x}) = \{i | \Delta_i > 0, i \in [N]\}$ for a given $\boldsymbol{x}$ because if $\Delta_i \leq 0$, the user will definitely invest, so we can remove it from our game. Proposition 4 characterizes the set of mixed Nash equilibria between followers in the three cases where $\sum_{i\in V(\boldsymbol{x})} \Delta_i$ is less, equal and larger than 1, which plays a central role to compute the optimal strategy of the government.

Proposition 4. *Given $\boldsymbol{x}$, the strategy profile $\langle \boldsymbol{y}, \boldsymbol{z}\rangle$ is a mixed Nash equilibrium (MNE) in the game where*

1. $\sum_{i\in V(\boldsymbol{x})} \Delta_i < 1$, *if and only if* $\forall i \in V(\boldsymbol{x}), y_i = 1, z_i \geq \Delta_i$ *and* $\sum_{i\in[N]} z_i = 1$;
2. $\sum_{i\in V(\boldsymbol{x})} \Delta_i = 1$, *if and only if* $\forall i \in V(\boldsymbol{x}), y_i = 1 - L/L_i, 0 \leq L \leq \min_{i\in V(\boldsymbol{x})} L_i, z_i = \Delta_i$;
3. $\sum_{i\in V(\boldsymbol{x})} \Delta_i > 1$, *if and only if there is a non-singleton, nonempty subset* $I \subset V(\boldsymbol{x})$ *such that* $\min_{i\in I} L_i \geq \max_{j\in V(\boldsymbol{x})\setminus I} L_j$ *and the followings hold: i)* $\forall i \in V(\boldsymbol{x}) \setminus I$, $y_i = z_i = 0$; *ii) let* $J = \arg\min_{i\in I} L_i$, $\forall i \in J$, $y_i = 0$, $0 \leq z_i \leq \Delta_i$; *iii)* $\forall i \in I \setminus J$, $y_i = 1 - \min_{j\in I} L_j/L_i$, $z_i = \Delta_i$.

Proof. We can rewrite the attacker's utility as

$$U^a(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z}) = \sum\nolimits_{i=1}^{N} z_i \left[(1 - y_i)\left(p_i l_i + \sum\nolimits_{j\neq i} q_{ij} l_j\right)\right] \quad (13)$$

The key observation is $\Psi^1_i = \sum_{j\neq i} z_j(1 - y_j) q_{ji}$. Then the proof can be adapted from the proof of Proposition 5 in [Chan *et al.*, 2017] by setting the attack cost to zero. Generally speaking, the proof of the proposition is based on three facts: i) when $z_i = \Delta_i$, the users will be indifferent to the investment which is observed from Eq.(12), ii) under any of the equilibria, the values $(1 - y_i)\left(p_i l_i + \sum_{j\neq i} q_{ij} l_j\right)$ are equal for all users with $z_i > 0$ and iii) for users with $z_i = 0, i \in V(\boldsymbol{x}), y_i = 0$. □

Proposition 5. *As the followers break ties in favor of the government, for the game where $\sum_{i\in V(\boldsymbol{x})} \Delta_i \leq 1$, the users will fully invest, i.e., $\boldsymbol{y} = \boldsymbol{1}$, and the attacker will choose $z_i \geq \Delta_i, \forall i \in V(\boldsymbol{x})$ and $z_i, i \in [N] \setminus V(\boldsymbol{x})$ are irrelevant.*

Proof. According to Proposition 4, when $\sum_{i\in V(\boldsymbol{x})} \Delta_i < 1$, the only NE strategy for users is $\boldsymbol{y} = \boldsymbol{1}$ and the attacker's utility is 0 for all NE strategies of the attacker. For the case where $\sum_{i\in V(\boldsymbol{x})} \Delta_i = 1$, the only NE strategy for the attacker is $z_i, i \in [N] \setminus V(\boldsymbol{x})$. As we suppose the users will break ties in favor of the government, i.e., minimize the attacker's utility. Therefore, the users will choose $\boldsymbol{y} = \boldsymbol{1}$ as their NE strategy, which leads to the case where the attacker's utility is 0. Thus, for the game where $\sum_{i\in V(\boldsymbol{x})} \Delta_i \leq 1$, the users will fully invest, i.e., $\boldsymbol{y} = \boldsymbol{1}$, and the attacker will choose $z_i \geq \Delta_i, \forall i \in V(\boldsymbol{x})$ and $z_i, i \in [N]\setminus V(\boldsymbol{x})$ are irrelevant. □

Given Proposition 5, we find that if $\sum_{i\in V(\boldsymbol{x})} \Delta_i \leq 1$, the users will always invest. And for the case where $\sum_{i\in V(\boldsymbol{x})} \Delta_i > 1$, Proposition 6 proves that all MNE can be computed in polynomial time in this case.

Proposition 6. *For the game where $\sum_{i\in V(\boldsymbol{x})} \Delta_i > 1$, given $\boldsymbol{x}$, all MNE can be computed in polynomial time.*

Proof. The algorithm can be adapted from Theorem 1 in [Chan *et al.*, 2017] by setting the attack cost to zero. Roughly speaking, the algorithm sorts users in descending order according to L_i and then finds t such that $1 - \Delta_{idx_1(t)} \leq \sum_{j=1}^{t-1} \Delta_{idx_1(j)} < 1$, where $idx_1(\cdot)$ and $val_1(\cdot)$ are the index and the value in the sorted list. We denote $I = \{idx_1(k) | k \leq t\}$ and specify followers' strategies by Proposition 4. □

Given Proposition 6, we now prove that the government's optimal strategy can be computed in polynomial time. Denote $\Delta = \sum_{i=1}^{N} \Delta_i$. We first check whether the government can make $\Delta \leq 1$ by greedily assigning c_i to users in the ascent order of $p_i l_i$. For the case where the government cannot make $\Delta \leq 1$, we can see from Proposition 6 that the government's utility is equal to $val_1(t)$ in the sorted list such that $1 - \Delta_{idx_1(t)} \leq \sum_{j=1}^{t-1} \Delta_{idx_1(j)} < 1$. Therefore, maximizing the government's utility is equivalent to minimizing $val_1(t)$, which is depicted in Algorithm 1.

Proposition 7. *Algorithm 1 computes an SMNE in polynomial time when the attacker has a single-attack capability.*

Algorithm 1: Compute government's optimal strategy

```
1  $\boldsymbol{x} = \mathbf{0}, t = 0, t' = N$;
2  Sort $L_i$ in descending order and let $val_1(i)$ and $idx_1(i)$
       be the $i$th value and index in the sorted list, respectively;
3  Sort $p_i l_i$ in descending order and let $val_2(i)$ and $idx_2(i)$
       be the $i$th value and index in the sorted list, respectively;
4  Find $t'$ such that $1 - \Delta_{idx_1(t')} \leq \sum_{j=1}^{t-1} \Delta_{idx_1(j)} < 1$;
5  while $t' > t$ do
6      $b = B, \boldsymbol{x} = 0, t = t'$;
7      while true do
8          $j \in \arg\min_{idx_1(i) \leq t, i \in V(\boldsymbol{x})} val_2(i)$;
9          if $j == $ null then break;
10         else if $b \geq c_j$ then $x_j = c_j$;$b = b - c_j$;
11         else $c_j = b$; break;
12     Find $t'$ where $1 - \Delta_{idx_1(t')} \leq \sum_{j=1}^{t'-1} \Delta_{idx_1(j)} < 1$;
13 return $\boldsymbol{x}$
```

Proof. We first prove that the loop in lines 7-11 can return the optimal solution if we only consider the users with $idx_1(i) \leq t$. As the attacker will attack the users with probability Δ_i, when the government assigns the subsidy with the minimal $p_i l_i$, which is denoted by j, i.e., $j \in \arg\min_{idx_1(i) \leq t} \{p_i l_i\}$. Then, user j will definitely invest and the attacker will assign the amount Δ_j to users with $idx_1(i) > t$. Thus, the greedy assignment of the subsidy is optimal due to the fact that assigning the subsidy to user j will reduce the largest amount of the probability that the attacker allocates to users with $idx_1(i) \leq t$. Then, we prove that when $t' = t$, the assignment is globally optimal. Observing that $t' \geq t$ at line 12, if $t' = t$, there is no user with $idx_1(i) > t$ being attacked by the attacker, given the optimal result of the loop in lines 7-11, the returned solution is globally optimal and we terminate the algorithm. The two sorting processes take $O(N \log N)$. The two while loops in line 5 and 7 run at most N times. The finding minimum operation in line 8 needs at most N comparisons. Thus, the runtime is $O(N^3 + N \log N) = O(N^3)$. □

5.3 Multiple-user Attack: $1 < K < N$

For multiple-user attacks, an argument similar to Proposition 3 can prove that there is no PNE between followers, so we focus on the mixed NE case. As claimed in [Chan *et al.*, 2017], the computation of MNE in general case is intractable. Therefore, we adopt *best-response-gradient dynamics* (BRGD) [Fudenberg and Levine, 1998] to compute an ϵ-MNE and greedily assign the subsidy to users. The algorithm is presented in Algorithm 2. BRGD initializes the users' strategies with 0.5 and the attacker's strategy uniformly with $\sum_{\boldsymbol{b} \in \mathcal{B}^K} z_{\boldsymbol{b}} = 1$. At each round, BRGD updates $y_i = y_i + \alpha(U_i^u(\boldsymbol{x}, \boldsymbol{y}_{-i}, 1, \boldsymbol{z}) - U_i^u(\boldsymbol{x}, \boldsymbol{y}_{-i}, 0, \boldsymbol{z}))$ and $z_b = z_b + \alpha(U^a(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{b}) - U^a(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z}))$ where $\boldsymbol{y}_{-i}$ is the strategies of other users except i, α is the step size of BRGD and $U_i^u(\boldsymbol{x}, \boldsymbol{y}_{-i}, \sigma, \boldsymbol{z}), \sigma \in \{0, 1\}$ is user i's utility when i takes the pure strategy σ, given that the other users take $\boldsymbol{y}_{-i}$ and the attacker takes $\boldsymbol{z}$. We define that for each user, $r_i = |\frac{U_i^u - U_i^u(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z})}{U_i^u(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z})}|$ where $U_i^u = \max\{U_i^u(\boldsymbol{x}, \boldsymbol{y}_{-i}, 1, \boldsymbol{z}), U_i^u(\boldsymbol{x}, \boldsymbol{y}_{-i}, 0, \boldsymbol{z})\}$ and $r_a = |\frac{U^a - U^a(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z})}{U^a}|$ where $U^a = \max_{\boldsymbol{b} \in \mathcal{B}^K}\{U^a(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{b})\}$ for the attacker. BRGD will terminate when $r < \epsilon$ where $r = \max\{r_1, ..., r_N, r_a\}$ and the solution is ensured to be an ϵ-MNE. To obtain ϵ-MNE using BRGD, for each iteration (Lines 3-11), we assign the $\delta = B/D$ to the user who can increase the government's utility the most, where D is the maximum number of iterations. We note that this algorithm can also be applied to approximate SMNE of all-user attacks.

Algorithm 2: Allocation for multiple attacks

```
1  $\boldsymbol{x} = 0, \delta = B/D$;
2  $\langle \boldsymbol{y}, \boldsymbol{z} \rangle \leftarrow$ BRGD$(\boldsymbol{x}, \epsilon)$;
3  for $d = 1 : D$ do
4      $\boldsymbol{icr} = \mathbf{0}$;
5      for $i = 1 : N$ do
6          $\boldsymbol{x}' = \boldsymbol{x}, x'_i = x_i + \delta$;
7          $\langle \boldsymbol{y}', \boldsymbol{z}' \rangle \leftarrow$ BRGD$(\boldsymbol{x}')$;
8          $icr_i = U^d(\boldsymbol{x}', \boldsymbol{y}', \boldsymbol{z}') - U^d(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z})$;
9      $j = \arg\max_{i \in [N], icr_i > 0} icr_i$;
10     if $j = $ null then break;
11     else $x_j = x_j + \delta$;
12 return $\boldsymbol{x}$
```

6 Experimental Results

We evaluate our algorithms through extensive experiments. All computations are performed on a 64-bit PC with 4.0 GB RAM and a 2.40 GHz CPU unless otherwise specified. All games are randomly generated with $c_i = 10^7 + [0, 10^8]$, $l_i = 10^8 + [0, 10^9]$ and $p_i, q_{ji} \in [0, 1]$ (Similar to [Chan *et al.*, 2012]). As there is an exact polynomial algorithm for single-user attacks, we focus on the other two cases in this section.

6.1 Results of All-user Attack for SPNE

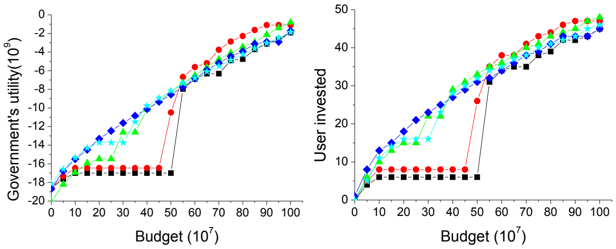

(a) Government's utility (b) Number of users investing

Figure 3: Results of all-user attack. Best viewed in color.

We use global nonlinear online solver BARON, provided by NEOS server [Czyzyk *et al.*, 1998], to find a global or local optimal solution of the problem. Figure 3 displays five generated instances of 50 users with different government budget. BARON can compute the government's optimal strategy in 5 minutes on average. Figure 3a and 3b show the government's utility and the number of users who invest against budget, respectively. There are two observations from the results: i) smooth transition where the number of users who invest and the government's utility increase smoothly (see blue, green and cyan lines) and ii) phase transition where increasing the budget leads to a significant increment in the number of invested users and government's utility (16 in red line and 25 in black line). The conclusion is supported by more instances.

6.2 Results of Multiple-user Attack for SMNE

We now present experimental results for multiple-user attacks. The convergence criteria is $\epsilon = 0.05$, the step size is $\alpha = 0.05/(K \cdot \max_{i\in[N]} l_i)$, the government's budget is $B = 5 \cdot 10^8$ and the number of iterations D is 10. To generate the synthetic data, we vary $K \in \{4, 6\}$ and $N \in \{12, 14, 16, 18, 20\}$. Each case is averaged over 30 instances.

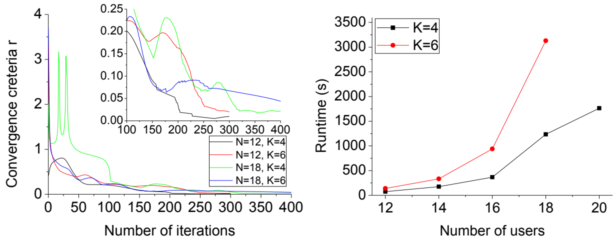

(a) Convergence of BRGD (b) Runtime

Figure 4: Convergence and runtime results of Algorithm 2.

Convergence and runtime. We first investigate the convergence of BRGD. Figure 4a shows the convergence results and, to make the figure readable, only four among all cases are depicted. All cases reach the termination in less than 350 iterations. We note that the number of iterations before the termination does not significantly depend on N or K. Furthermore, we display the runtime results of the algorithm with a cap of 3600 seconds in Figure 4b, where when K is large, the algorithm takes more time to terminate. Note that when $N = 20$ and $K = 6$, the runtime is beyond the cap.

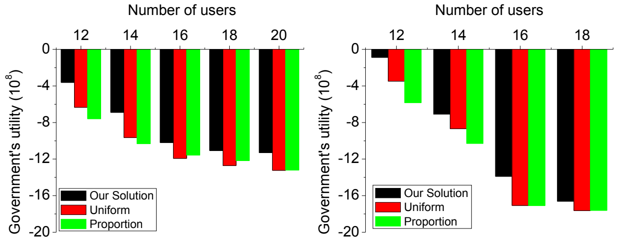

(a) $K = 4$ (b) $K = 6$

Figure 5: Solution quality of Algorithm 2.

Solution quality. We compare our solution with two other baselines: **uniform** where the government divides the subsidy to users uniformly and **proportion** where the government assigns the subsidy proportionally to the user's loss l_i. Figure 5 shows the solution quality with different values of K where our solution is better than the two baselines. The advantage of our solution is reduced when K is smaller, i.e., users are more reluctant to invest in safer environments.

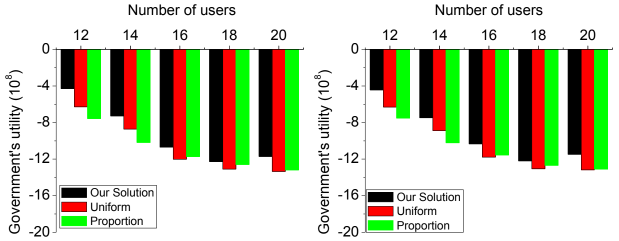

(a) Transfer probability (b) Number of K

Figure 6: Robustness of Algorithm 2 with $K = 4$.

Robustness. We consider two kinds of uncertainties: i) transfer probability where $\tilde{q}_{ji} = (1+\delta)q_{ji}, \delta \in [-10\%, 10\%]$ where $\tilde{q}_{ji}$ is the real transfer probability and q_{ji} is the transfer probability that the government uses to compute his strategy; and ii) the number of K where the real number of users the attacker can attack $\tilde{K}$ is either $K+1$ or $K-1$ with a probability of 10%, respectively, where K is the number used by the government to decide his strategy. Figure 6 shows that our solution outperforms the two baselines against the uncertainties for $K = 4$. Results for $K = 6$ are similar and omitted.

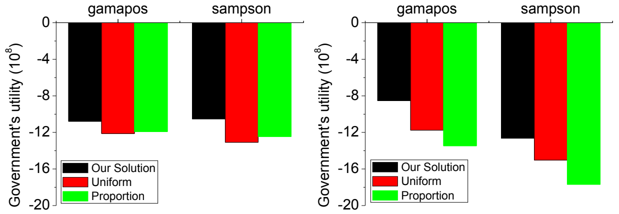

(a) $K = 4$ (b) $K = 6$

Figure 7: Experiment results on real network data.

Experiments on real networks. As business networks are often generated socially, we consider two real social networks. The networks are **gamapos**, a highland tribe network with 16 nodes and 58 edges, and **sampson**, a monastery monk network with 18 nodes and 55 edges[3]. Figure 7 shows that our solution outperforms the baselines. The results also demonstrate that our heuristic algorithm are better than baselines.

7 Conclusion

We propose a novel Stackelberg cybersecurity investment game between the government, interdependent users and an attacker. We investigate three cases where the attacker can attack all, single and multiple users and propose a reverse convex formulation, an exact polynomial algorithm and a heuristic algorithm for the three cases, respectively. Experimentally, we show that our heuristic algorithm is effective and outperforms baselines on synthetic and real data.

Acknowledgements

This research is supported by NCR2016NCR-NCR001-0002, MOE, and NTU.

[3] Both networks are from http://konect.uni-koblenz.de/networks/

References

[Basilico *et al.*, 2016] Nicola Basilico, Stefano Coniglio, and Nicola Gatti. Methods for finding leader-follower equilibria with multiple followers. In *AAMAS*, pages 1363–1364, 2016.

[Böhme *et al.*, 2010] Rainer Böhme, Galina Schwartz, et al. Modeling cyber-insurance: towards a unifying framework. In *WEIS*, 2010.

[CBS, 2017] CBS. Wannacry ransomware attack losses could reach $4 billion. https://www.cbsnews.com/news/wannacry-ransomware-attacks-wannacry-virus-losses/, 2017. Last accessed on 25 Feb 2019.

[Chan *et al.*, 2012] Hau Chan, Michael Ceyko, and Luis E. Ortiz. Interdependent defense games: Modeling interdependent security under deliberate attacks. In *UAI*, pages 152–162, 2012.

[Chan *et al.*, 2017] Hau Chan, Michael Ceyko, and Luis Ortiz. Interdependent defense games with applications to internet security at the level of autonomous systems. *Games*, 8(1):13, 2017.

[Coniglio *et al.*, 2017] Stefano Coniglio, Nicola Gatti, and Alberto Marchesi. Pessimistic leader-follower equilibria with multiple followers. In *IJCAI*, pages 171–177, 2017.

[Czyzyk *et al.*, 1998] Joseph Czyzyk, Michael P Mesnier, and Jorge J Moré. The NEOS server. *IEEE Computational Science and Engineering*, 5(3):68–75, 1998.

[Fang *et al.*, 2016] Fei Fang, Thanh H Nguyen, Rob Pickles, Wai Y Lam, Gopalasamy R Clements, Bo An, Amandeep Singh, Milind Tambe, and Andrew Lemieux. Deploying PAWS: Field optimization of the protection assistant for wildlife security. In *IAAI*, pages 3966–3973, 2016.

[Fudenberg and Levine, 1998] Drew Fudenberg and David K Levine. *The Theory of Learning in Games*, volume 2. MIT press, 1998.

[Gan *et al.*, 2015] Jiarui Gan, Bo An, and Yevgeniy Vorobeychik. Security games with protection externalities. In *AAAI*, pages 914–920, 2015.

[Gan *et al.*, 2018] Jiarui Gan, Edith Elkind, and Michael Wooldridge. Stackelberg security games with multiple uncoordinated defenders. In *AAMAS*, pages 703–711, 2018.

[Gordon and Loeb, 2002] Lawrence A Gordon and Martin P Loeb. The economics of information security investment. *ACM Transactions on Information and System Security (TISSEC)*, 5(4):438–457, 2002.

[Gordon *et al.*, 2003] Lawrence A Gordon, Martin P Loeb, and William Lucyshyn. Sharing information on computer systems security: An economic analysis. *Journal of Accounting and Public Policy*, 22(6):461–485, 2003.

[Grossklags *et al.*, 2010] Jens Grossklags, Svetlana Radosavac, Alvaro A Cárdenas, and John Chuang. Nudge: Intermediaries' role in interdependent network security. In *TRUST*, pages 323–336, 2010.

[Guo *et al.*, 2016] Qingyu Guo, Bo An, Yair Zick, and Chunyan Miao. Optimal interdiction of illegal network flow. In *IJCAI*, pages 2507–2513, 2016.

[Horst and Tuy, 2013] Reiner Horst and Hoang Tuy. *Global Optimization: Deterministic Approaches*. Springer Science & Business Media, 2013.

[Jiang *et al.*, 2011] Libin Jiang, Venkat Anantharam, and Jean Walrand. How bad are selfish investments in network security? *IEEE/ACM Transactions on Networking (TON)*, 19(2):549–560, 2011.

[Kearns and Ortiz, 2004] Michael Kearns and Luis E Ortiz. Algorithms for interdependent security games. In *NIPS*, pages 561–568, 2004.

[Kunreuther and Heal, 2003] Howard Kunreuther and Geoffrey Heal. Interdependent security. *Journal of Risk and Uncertainty*, 26(2-3):231–249, 2003.

[Laszka *et al.*, 2015] Aron Laszka, Mark Felegyhazi, and Levente Buttyan. A survey of interdependent information security games. *ACM Computing Surveys*, 47(2):23, 2015.

[Omic *et al.*, 2009] Jasmina Omic, Ariel Orda, and Piet Van Mieghem. Protecting against network infections: A game theoretic perspective. In *INFOCOM*, pages 1485–1493, 2009.

[Paruchuri *et al.*, 2008] Praveen Paruchuri, Jonathan P Pearce, Janusz Marecki, Milind Tambe, Fernando Ordonez, and Sarit Kraus. Playing games for security: An efficient exact algorithm for solving bayesian Stackelberg games. In *AAMAS*, pages 895–902, 2008.

[Schlenker *et al.*, 2017] Aaron Schlenker, Haifeng Xu, Mina Guirguis, Chris Kiekintveld, Arunesh Sinha, Milind Tambe, Solomon Sonya, Darryl Balderas, and Noah Dunstatter. Don't bury your head in warnings: A game-theoretic approach for intelligent allocation of cyber-security alerts. In *IJCAI*, pages 381–387, 2017.

[Schlenker *et al.*, 2018] Aaron Schlenker, Omkar Thakoor, Haifeng Xu, Fei Fang, Milind Tambe, Long Tran-Thanh, Phebe Vayanos, and Yevgeniy Vorobeychik. Deceiving cyber adversaries: A game theoretic approach. In *AAMAS*, pages 892–900, 2018.

[Tambe, 2011] Milind Tambe. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press, 2011.

[Thycotic, 2017] Thycotic. 2017 state of cyber security metrics annual report. https://thycotic.com/wp-content/uploads/2013/03/2017-Cyber-Security-Strategy-Executive-Summary.pdf, 2017. Last accessed on 25 Feb 2019.

[UK, 2015] UK. *National Security Strategy and Strategic Defence and Security Review 2015*. Government of the United Kingdom, 2015.

[UK, 2016] UK. *National Cyber Security Strategy 2016-2021*. Government of the United Kingdom, 2016.

[Vorobeychik and Letchford, 2015] Yevgeniy Vorobeychik and Joshua Letchford. Securing interdependent assets. *Autonomous Agents and Multi-Agent Systems*, 29(2):305–333, 2015.