A Dynamic Incentive Mechanism for Security in Networks of Interdependent Agents

Farzaneh Farhadi[1,2](✉), Hamidreza Tavafoghi[1], Demosthenis Teneketzis[1], and Jamal Golestani[2]

[1] University of Michigan, Ann Arbor, USA
{ffarhadi,tava,teneket}@umich.edu
[2] Sharif University of Technology, Tehran, Iran
golestani@ieee.org

Abstract. We study a dynamic mechanism design problem for a network of interdependent strategic agents with coupled dynamics. In contrast to the existing results for static settings, we present a dynamic mechanism that is incentive compatible, individually rational, budget balanced, and social welfare maximizing. We utilize the correlation among agents' states over time, and determine a set of *inference signals* for all agents that enable us to design a set of incentive payments that internalize the effect of each agent on the overall network dynamic status, and thus, align each agent's objective with the social objective.

Keywords: Security games · Dynamic mechanism design · Epidemics over networks

1 Introduction

Recently there has been a growing body of literature studying the dynamic behavior of networked strategic agents, where each agent's state and utility is affected by his interactions with his neighbors in the network. This literature is motivated by various applications that include opinion dynamics in social networks, epidemics spreading in networks, dynamic adoption of new products and technologies over networks, and network security. In this paper, we study a model of dynamic networked agents motivated by a network security application.

We consider a dynamic network with strategic agents who privately observe their own security state and are only interested in maximizing their own utility. We formulate a mechanism design problem for a network manager whose objective is to dynamically allocate his limited security resources in the network so as to maximize the overall security of the whole network over time.

We assume that an agent's utility depends on his own private security state as well as the externality he receives from his neighbors in the network. Moreover, an agent's security state dynamically evolves over time; its evolution depends

This work was supported in part by the NSF grants CNS-1238962 and CCF-1111061.

L. Duan et al. (Eds.): GameNets 2017, LNICST 212, pp. 86–96, 2017.
DOI: 10.1007/978-3-319-67540-4_8

on the security resources the agent receives from the network manager, as well as direct external attacks launched from outside of the network and indirect internal attacks launched from his unsafe local neighbors in the network. Therefore, the network manager needs to design a dynamic incentive mechanism for agents with *correlated types* and *interdependent valuations* so as to align their selfish objectives with his own objective, which is the maximization of the overall security of the whole network.

We propose a dynamic incentive mechanism that is individually rational and budget balanced [3], and enables the network manager to achieve the socially efficient outcome. Our result is in contrast with the existing impossibility results for incentive mechanisms that are socially efficient, individually rational, incentive compatible, and budget balanced in the static settings [22]. We exploit the dynamic correlation among the agents' security states and determine a set of *inference signal*s for all agents over time. Utilizing the proposed set of inference signals, we characterize a dynamic incentive mechanism that ensures the agents' incentive compatibility and individual rationality, achieves a socially efficient outcome, and is ex-ante budget balanced.

There is a growing body of literature on network security games (see [18] and references therein). One set of papers assume that network agents are cooperative, and study the interactions between the network as a whole and an outside attacker as a two-player attacker-defender game [4,13,17]. Another set of papers assume an exogenously-fixed attack behavior from outside the network, and study the interactions between strategic agents within the network as a network game problem (see [9,15] and references therein). For instance, the work of [10] studies a network security game with strategic agents, and shows that the equilibrium outcome of the game can be very poor compared to the social optimum, and this gap tends to increase with the increase in network size and the agents' interdependence. In our work, we study the dynamic interactions among agents within the network. However, we take the mechanism design approach rather than analyzing the resulting security game for a given environment.

The existing literature on mechanism design for network security considers mainly static incentive design problems. For instance, the work of [16] investigates the role of cyber-insurance as an incentive instrument for agents to increase their security investment in self-protection. The work of [22] studies the mechanism design problem for general networks with strategic agents in static settings, and shows that there exists no incentive mechanism that can implement the socially efficient outcome, while ensuring individual rationality, incentive compatibility and (weak) budget balance. Our paper contributes to this set of literature by showing that this impossibility result does not hold for dynamic settings. The fact that the agents' incentive problem improves in dynamic settings has been previously shown by works that look at security games in repeated settings (see [10,21]). Our work is different from those that consider repeated game settings. First, we take a mechanism design approach rather than analyzing a repeated game setting. Second, in repeated game settings there is no system dynamics, and the existing results are based on the reputation that agents mutually form over time. Our work provides another insight for such improvements in dynamic settings by capitalizing on the coupling among the agents' security dynamics over time.

The model we consider in this paper is also related the literature on Susceptible-Infected-Susceptible (SIS) epidemic models over networks with strategic agents (see [23] and references therein). For instance, the works of [8,24] study different variations of SIS epidemic models over networks with strategic agents from a game-theoretical approach. The authors in [24], investigate a game setting where agents make one-time investment decisions in their security which then affect the epidemic process. The work of [8] studies a marketing problem on networks using a SIS epidemic model, and investigates a game problem between two firms which compete for market shares over the network.

The mechanism design problem we consider in this paper can also be viewed as a dynamic resource allocation mechanism with strategic agents. The work of [12] studies the resource allocation problem in networks with non-strategic agents. The authors in [6,11] consider the resource allocation problem in static networks with strategic agents, take an *implementation theory* approach, and propose resource allocation mechanisms that are social welfare maximizing, individually rational and budget balanced. In this paper, we consider a class of dynamic resource allocation problems with strategic agents, and we present a dynamic resource allocation mechanism that is social welfare maximizing, incentive compatible, individually rational and ex-ante budget balanced.

The rest of the paper is organized as follows. We present our model in Sect. 2. We formulate the dynamic incentive design problem and characterize its solution in Sect. 3. We show that the dynamic incentive mechanism proposed in this paper can implement the solution of the corresponding dynamic centralized optimal resource allocation problem. In Sect. 4, we formulate such a centralized resource allocation problem as a centralized stochastic control problem and provide its solutions for a set of specific network topologies. The proofs of all the results that appear in this paper can be found in [7].

2 Model

There are n strategic agents each one residing in a distinct node of an interconnected network interacting over time $t \in \mathcal{T} := \{0, 1, 2, \ldots\}$. At each time $t \in \mathcal{T}$, the security state of agent i is given by $\theta_t^i \in \Theta := \{0, 1\}$; the realization of θ_t^i is agent i's private information. Agent i's state is safe if $\theta_t^i = 1$ and is unsafe if $\theta_t^i = 0$. We refer to θ_t^i as agent i's *type* at time t. There is a network manager who takes security measures dynamically over time so as to defend the network against external attacks and/or propagation of internal attacks. The security state θ_t^i of agent i dynamically evolves over time; θ_t^i's evolution depends on the security state of his neighbors in the network, the network manager's actions, and the probability of external attacks.

System Dynamics. We represent the agents' network by a directed graph $G = (N, L)$ where $N = \{1, ..., n\}$ and $L \in \mathbb{R}_+^{n \times n}$ denote the set of agents and the set of directed links between them, respectively. The state θ_t^i of agent i

is affected by agent j if $l_{ji} > 0$. We define the set of agent i's neighbors as $N^i\colon = \{j : l_{ji} > 0\}$. During each time $t \in \mathcal{T}$, if agent i is in the safe state, *i.e.* $\theta_t^i = 1$, it may be attacked directly from outside with probability d_i, or indirectly from any of his unsafe neighbors $j \in N^i$ with probability l_{ji}. The topology of the network G and the probability of outside attacks d_i remains the same over time.

The goal of the network manager is to maximize the overall security of the network over time, *i.e.* maximize the social welfare. At each time t, the manager can choose one agent $a_t \in N$ and apply a security measure to him. As a result of applying the security measure to agent i, *i.e.* $a_t = i$, if agent i is in the unsafe state he will switch to the safe state with probability h. The security measure also protects the chosen agent against direct attacks from outside during time t with the same probability h, but it does not affect the indirect spread of attacks within the network.

Let $\boldsymbol{\theta}_t = (\theta_t^1, \ldots, \theta_t^n) \in \Theta^n$ denote the security state of the network at time t. As a result of the network manger's action a_t, new direct attacks from outside, and the spread of indirect attacks within the network during time t, the network state θ_{t+1} has the following Markovian dynamics:

$$\mathbb{P}\{\boldsymbol{\theta}_{t+1} = \boldsymbol{b} | \boldsymbol{\theta}_t, a_t\} = \prod_{i=1}^{n} \mathbb{P}\{\theta_{t+1}^i = b_i | \boldsymbol{\theta}_t, a_t\},\ \forall \boldsymbol{b} \in \Theta^n, \tag{1}$$

where,

$$\mathbb{P}\{\theta_{t+1}^i{=}1|\boldsymbol{\theta}_t, a_t\}{=}\begin{cases} 0, & \theta_t^i = 0, i \neq a_t \\ h(1 - d_i(1-h)) \prod_{j \in N^i : \theta_t^j = 0} (1 - l_{ji}), & \theta_t^i = 0, i = a_t \\ (1 - d_i) \prod_{j \in N^i : \theta_t^j = 0} (1 - l_{ji}), & \theta_t^i = 1, i \neq a_t \\ (1 - d_i(1-h)) \prod_{j \in N^i : \theta_t^j = 0} (1 - l_{ji}), & \theta_t^i = 1, i = a_t \end{cases}, \tag{2}$$

and $\mathbb{P}\{\theta_{t+1}^i = 0 | \boldsymbol{\theta}_t, a_t\} = 1 - \mathbb{P}\{\theta_{t+1}^i = 1 | \boldsymbol{\theta}_t, a_t\}$. We note that by (1) and (2) we assume that the outside attacks and attack spreads within network are independent across different agents, and thus, conditioned on previous state $\boldsymbol{\theta}_t$ and the network manager's action a_t, the agents' security states evolve independently as in (1). Equation (2) describes this evolution: (i) if agent i is in the unsafe state and is not receiving any security measure from the network manager at t, he remains in the unsafe state; (ii) if agent i is in the unsafe state and receives the security measure from the network manager, he will restore his security if the security measure is successful (prob. h), he is not the subject of new direct attacks (prob. $(1 - d_i(1-h))$), and he is not attacked by his unsafe neighbors (prob. $\prod_{j \in N^i : \theta_t^j = 0} (1 - l_{ji})$); (iii) similarly, if agent i is in the safe state and is not receiving a security measure, he will remain in the safe state if he is not attacked from outside (prob. $1 - d_i$) and he is not attacked by his unsafe neighbors (prob. $\prod_{j \in N^i : \theta_t^j = 0} (1 - l_{ji})$); (iv) if agent i is in the safe state and is receiving a security measure from the network manager, he will remain in the safe state if he is not attacked from outside (prob. $1 - d_i(1-h)$) and he is not attacked by his neighbors that are in an unsafe state (prob. $\prod_{j \in N^i : \theta_t^j = 0} (1 - l_{ji})$).

Agents' Utilities. Each agent $i \in N$ has a valuation for his security state θ_t^i as well as the security state of his neighbors θ_t^j, $j \in N^i$, and the security measures he receives from the network manager; this valuation is given by,

$$v^i(\boldsymbol{\theta}_t, a_t) = \theta_t^i + \frac{\alpha}{|N^i|}\mathbf{1}_{\{\theta_t^i=1 \text{ or } a_t=i\}} \sum_{j \in N^i} \theta_t^j, \tag{3}$$

where $0 < \alpha < 1$ captures the value of a safe neighborhood to an agent i. As a result of (3), agent i has a positive valuation for safe neighbors only if he is in the safe state or he is receiving a security measure at t, *i.e.* $\{\theta_t^i = 1 \text{ or } a_t = i\}$. Let p_t^i denote the monetary payment made by agent i to the network manager at t ($p_t^i \in \mathbb{R}$). Then the total utility of agent i at t is given by,

$$u_t^i(\boldsymbol{\theta}_t, a_t, p_t^i) = v^i(\boldsymbol{\theta}_t, a_t) - p_t^i, \tag{4}$$

Let $\delta \in (0, 1)$ denote the common discount factor. Then the total discounted utility of agent $i \in N$, is

$$U^i = (1-\delta)\sum_{t=0}^{\infty} \delta^t u_t^i(\boldsymbol{\theta}_t, a_t, p_t^i) = (1-\delta)\sum_{t=0}^{\infty} \delta^t (v^i(\boldsymbol{\theta}_t, a_t) - p_t^i)\ . \tag{5}$$

The network manager's objective is to maximize the social welfare W given by,

$$W = \mathbb{E}\{(1-\delta)\sum_{t=0}^{\infty} \delta^t \sum_{i=1}^{n} v^i(\boldsymbol{\theta}_t, a_t)\}. \tag{6}$$

The network manager's problem would be a standard control problem (Markov decision problem) if the manager knew $\boldsymbol{\theta}_t$ for all t. However, $\boldsymbol{\theta}_t$ is not known to the manager; θ_t^i, $i \in N$, is agent i's private information. Thus, in order to take a security measure at any time t, the manager has to elicit information about each agent's security status. Since all agents are selfish (strategic) and want to maximize their own utility given by (5), they do not voluntarily reveal their information to the manager. Therefore, the manager needs to design an incentive mechanism so as to align the agents' objectives with his own objective. In this paper, we investigate such an incentive design problem, and formulate it as a mechanism design problem in Sect. 3.

3 Dynamic Incentive Design Problem

We invoke the revelation principle for dynamic games [20], and, without loss of generality, restrict attention to *direct revelation mechanisms* that are *incentive compatible*. In a direct revelation mechanism, at every $t \in \mathcal{T}$, the network manager asks agents to report their current security state. Let r_t^i denote agent i's report for time t, which is not necessarily the same as θ_t^i. Let $h_t := \{r_s^i, i \in N, s \le t\}$ denote the history of reports and $\mathcal{H}_t$ denote the set of all possible histories at t. A direct mechanism is captured by a set of functions $(\pi(.), p(.)) = \{\pi_t(\cdot), p_t^i(\cdot), i \in N, t \in \mathcal{T}\}$ that the network manager designs and commits to them, where $\pi_t : \mathcal{H}_t \to N$ determines which agent receives the security measure at t, and $p_t^i : \mathcal{H}_t \to \mathbb{R}$, $i \in N$, determines the monetary payment (or the negative of the monetary incentive) that agent i makes (receives) at time t based on the history up to t. A direct mechanism is incentive compatible (IC) if at every $t \in \mathcal{T}$ every agent is willing to report truthfully his security state given that the other agents report truthfully. That is, for every agent $i \in N$ and for all reporting strategies $\{\sigma_\tau^i : \Theta \times \mathcal{H}_\tau \to \Delta(\Theta), \tau \geq t\}$, truth telling results in higher expected utility at every $t \in \mathcal{T}$ and $h_t \in \mathcal{H}_t$, *i.e.*

$$\begin{aligned} \mathbb{E}\{(1-\delta)\textstyle\sum_{\tau=t}^{\infty} \delta^{\tau-t} \Big[v^i(\boldsymbol{\theta}_\tau, \pi_\tau(\boldsymbol{\theta}_\tau^{-i}, \theta_\tau^i)) - p_\tau^i(\boldsymbol{\theta}_\tau^{-i}, \theta_\tau^i)\Big]\} \geq \\ \mathbb{E}\{(1-\delta)\textstyle\sum_{\tau=t}^{\infty} \delta^{\tau-t} \Big[v^i(\boldsymbol{\theta}_\tau, \pi_\tau(\boldsymbol{\theta}_\tau^{-i}, \sigma_\tau(\theta_\tau^i, h_\tau^i))) - p_\tau^i(\boldsymbol{\theta}_\tau^{-i}, \sigma_\tau(\theta_\tau^i, h_\tau^i))\Big]\}, \end{aligned} \tag{7}$$

where $\Delta(\Theta)$ denotes the set of all probability distributions on Θ.

The network manager also needs to ensure that agents voluntarily participate in the direct mechanism $(\pi(.), p(.))$. Let $U_0^i \geq 0$ denote agent i's expected utility by opting out of the mechanism. Then, agents' voluntary participation is ensured by the following individual rationality (IR) constraints as follows,

$$\mathbb{E}\{(1-\delta)\sum_{\tau=0}^{\infty} \delta^\tau \Big[v_\tau^i(\boldsymbol{\theta}_\tau, \pi_\tau(\boldsymbol{\theta}_\tau^{-i}, \theta_\tau^i)) - p_\tau^i(\boldsymbol{\theta}_\tau^{-i}, \theta_\tau^i)\Big]\} \geq U_0^i, \forall i \in N. \tag{8}$$

Therefore, we can formulate the dynamic incentive design problem for the network manager as follows:

$$\max_{\pi(\cdot), p(\cdot)} \mathbb{E}\{(1-\delta)\sum_{t=0}^{\infty} \delta^t \sum_{i=1}^{n} v^i(\boldsymbol{\theta}_t, a_t)\} \tag{9}$$
$$\textit{subject to IC constraints}\,(7)\,\textit{and}\,\textit{IR}\,\textit{constraints}(8)$$

The incentive design problem formulated above is a dynamic mechanism design problem with correlated types and interdependent valuations. It is a dynamic mechanism design (in the strategic sense) since agents' incentive constraints at any time t depend on their strategic decisions at other times. Moreover, since the evolution of security states, given by (2), are coupled among agents, the agents' types are correlated with each other and over time. Furthermore, each agent's utility, given by (3), depends on his neighbor's security states in addition to his own security state, thus, agents have interdependent valuations. As a result of the correlation among agents' types and agents' interdependent valuations, the dynamic generalizations of the Vickrey–Clarke–Groves (VCG) mechanism [2] and that of d'Aspremont and Gerard-Varet (AGV) mechanism [1] cannot be used to solve the network manager's problem (9).

In this paper, we present an alternative approach to the dynamic incentive design problem by the network manager. We utilize the correlation among agents' security states over time to form a set of *cross inference signals* that enable us to internalize the effect of each agent's security state on the overall network security through incentive payments. The idea of utilizing the correlation among agents' types to extract their private information was first exploited by Cremer and McLean in a static setting [5]. They formed a cross inference signal for each agent by utilizing the correlation among the realization of agents' types, determined appropriate incentive payments that depend on the cross inference signals, and extracted the agents' private information. Liu [19] considered a dynamic setting with coupled dynamics, and utilized the inter-temporal correlation among agents' types to form cross inference signals for each agent that lead to truthful reporting at each time instant.

We provide a similar approach as the one in [19]. We utilize the inter-temporal correlation between agent i's security state θ_t^i at t and other agents' security state θ_{t+1}^j, $j \neq i$, at $t+1$ and form a cross inference signal that determines agent i's payment over time. We show that such cross inference signals enable the network manager to align the agents' self-interests with the overall social interest, and maximize the social welfare W.

3.1 Specification of the Mechanism

In this section we present a 'Dynamic Cross Inference' (DCI) mechanism that maximizes the social welfare subject to the IC and IR constraints (9). The description of our mechanism is divided into two parts: the allocation policy $\{\pi_t(\cdot), t \in \mathcal{T}\}$, and the monetary transfers $\{p_t^i(\cdot), i \in N, t \in \mathcal{T}\}$.

Allocation Policy. The specification of the allocation policy is based on the premise that the mechanism is incentive compatible. In an incentive compatible mechanism the agents report their security states truthfully. Therefore, the network manager is faced with a stochastic control problem with complete information. We design the allocation policy of our mechanism to be an optimal solution to this problem which we denote by π^*, i.e., $\pi_t = \pi^*(r_t)$, $\forall t \in \mathcal{T}$. In Sect. 4, we discuss how the network manager can find such an optimal policy.

Monetary Transfers. To obtain an incentive compatible mechanism, we design monetary transfers so that they exactly align the incentives of each agent with the social welfare. Since agents' valuations are interdependent, we cannot use the idea of Groves' mechanism by simply paying each agent i the total valuations of other agents, because the valuations of agents except i depend directly on the report of agent i, and this creates incentive for misreporting. To fix this, we utilize the correlation between agent i's security state θ_t^i at t and other agents' security states θ_{t+1}^j, $j \neq i$, at $t+1$ and form a cross inference signal about the security state of agent i which is independent of his own reports. We use this cross inference signal to align the objective of agent i with the social welfare.

Specifically, let $\boldsymbol{r}_t^{-i}$ denote the report profile of all agents except agent i at time t. We define the cross inference signal for agent i at time t as follows:

$$m_t^i = \begin{cases} 0, & \text{if } r_{t+1}^j = 0, \forall j \in O^i, \\ 1, & \text{otherwise}, \end{cases} \tag{10}$$

where $O^i := \{j \in N : i \in N^j\}$ is the set of *output neighbors* of agent i. If at time $t+1$, all output neighbors of agent i report to be unsafe, the manager interprets this as a signal that agent i was unsafe at time t. Otherwise, he assesses agent i as a safe agent.

By using the cross inference signal m_t^i, we construct payments p_{t+1}^i such that, in expectation, at time $t+1$ agent i receives the sum of time-t flow valuations of all other agents. So agent i's continuation payoff at time t is equal to the social surplus from time t onward. With this in mind, we define the tax $p_{t+1}^i(m_t^i, \boldsymbol{r}_t^{-i}, a_t)$ to be paid by each agent i at time $t+1$, as the solution to the following system of linear equations:

$$\mathcal{P}(m_t^i = 0|\theta_t^i, \boldsymbol{r}_t^{-i}, a_t)p_{t+1}^i(0, \boldsymbol{r}_t^{-i}, a_t) + \mathcal{P}(m_t^i = 1|\theta_t^i, \boldsymbol{r}_t^{-i}, a_t)p_{t+1}^i(1, \boldsymbol{r}_t^{-i}, a_t) = \\ -\frac{1}{\delta}\sum_{j \neq i} v^j(\boldsymbol{\theta}_t, a_t), \forall \theta_t^i \in \mathit{\Theta}, \quad (11)$$

where $\mathcal{P}(m_t^i|\theta_t^i, \boldsymbol{r}_t^{-i}, a_t)$ is the probability of m_t^i given θ_t^i, $\boldsymbol{r}_t^{-i}$ and a_t, assuming truthful reports of agents except i, i.e. $\boldsymbol{r}_\tau^{-i} = \boldsymbol{\theta}_\tau^{-i}$, $\tau = t, t+1$.

Lemma 1. For any a_t and $\boldsymbol{r}_t^{-i}$, the system of equations (11) has a solution.

Therefore, payments p_{t+1}^i are always well-defined. Using these payments the network manager is able to align the objective of each agent with the social welfare since,

$$v^i(\boldsymbol{\theta}_t, a_t) - \delta\ \mathbb{E}\{p_{t+1}^i(m_t^i, \boldsymbol{\theta}_t^{-i}, a_t)\} = \sum_{j \in N} v^j(\boldsymbol{\theta}_t, a_t). \quad (12)$$

This feature is the key to proving the main result of this paper stated below.

Theorem 1. The DCI mechanism maximizes the social welfare and satisfies the IC and IR constraints, therefore, it is an optimal solution to the dynamic incentive design problem (9) for the network manager.

3.2 Budget Balance

The DCI mechanism proposed in Sect. 3.1 efficiently solves the problem network manager faces (9), however, the transfers are not budget balanced. When the agents adopt truthful strategies, the total amount of monetary transfers the network manager receives from the agents is negative. This means that the mechanism runs large deficits subsidizing agents. In this section we show that this budget deficit can be alleviated by introducing a set of participation fees.

At time $t = 0$ and before realizing the first period's security states θ_0^i, each agent i can decide whether or not to participate in the mechanism[1]. If he decides to participate, he should pay a participation fee $\tilde{p}_0^i$. We construct participation fees such that in expectation, their total amount is equal to the total amount of future subsidies. We define the participation fee of agent i by

$$\tilde{p}_0^i = \frac{-1}{N-1}\sum_{j \neq i}\mathbb{E}\{\sum_{t=0}^{\infty}\delta^t p_t^j\}, \quad (13)$$

[1] Equivalently, we can assume that all agents start from the safe state $\theta_0^i = 1$.

where the expectation is taken with respect to agents' strategies determined by the mechanism, the initial distribution of the security states which is assumed to be known to the network manager and the agents, and the dynamics of the security network. Adding these fees balances the budget as

$$\sum_i \tilde{p}_0^i + \mathbb{E}\{\sum_{t=0}^{\infty} \delta^t p_t^i\} = -\mathbb{E}\{\sum_{t=0}^{\infty} \delta^t p_t^i\} + \mathbb{E}\{\sum_{t=0}^{\infty} \delta^t p_t^i\} = 0. \tag{14}$$

Therefore, the DCI mechanism with participation fees is ex-ante budget balanced. With the introduction of the participation fees, an agent might rather stay out of the mechanism to avoid paying the participation fee while he still enjoys the positive externality that he receives from other agents' participation in the mechanism. Below, we show that for sufficiently patient agents, all agents voluntarily participate in the DCI mechanism with participation fees.

Theorem 2. For δ sufficiently close to 1, the DCI mechanism with participation fees is ex-ante budget balanced, satisfies the IC and IR constraints, and maximizes the social welfare W.

4 Dynamic Optimal Policy for the Network Manager

In this section, we study the control problem that the network manager must solve to find an optimal allocation policy π^*, when the agents reveal their security states $\{\boldsymbol{\theta}_t\}$ truthfully. In this case, the network manager is faced with a Markov decision process (MDP) with perfect observations, where the transition probabilities are given by (1) and (2) and the instantaneous reward is given by $r(\boldsymbol{\theta}_t, a_t) := \sum_{i=1}^{n} v^i(\boldsymbol{\theta}_t, a_t)$. Using dynamic programming [14], the network manager can solve this problem numerically, and find an optimal policy. However, there are some settings where qualitative properties of an optimal policy can be derived analytically. In the following, we discover qualitative properties of an optimal policy within the context of a specific network topology.

Example. Consider a circular network with $n = 4$ agents, where $h = 1$, $d_i = 0$, and $l_{ij} = l \leq 0.5$, for all i, j that are adjacent agents. The next proposition fully describes an optimal policy for this setting and the behavior of the corresponding value function.

Proposition. (i) An optimal policy π^* applies the security measure to one of the head ends of the shortest 'run of unsafe agents'. A run of unsafe agents of length k is a succession of k unsafe agents consecutively located between two safe agents.

(ii) The value function $V^*(.)$ induces a complete ordering on the set of states, such that a state with a greater number of safe agents is strictly preferred to a state with smaller number of safe agents. In the case of equality, the state with a longer run of unsafe agents is strictly preferred.

The above proposition provides two metrics in comparing security states: (1) the number of safe agents and (2) how close the unsafe agents are to one another. Numerical results show that these two metrics still work in symmetric circular networks with an arbitrary number of agents. This means that if l is below a certain threshold, an optimal policy tries to first maximize the number of safe agents, and then, bring the unsafe agents close to one another. To do so, the network manager applies the security measure to one of the head ends of the shortest run of unsafe agents.

References

1. Athey, S., Segal, I.: An efficient dynamic mechanism. Econometrica **81**(6), 2463–2485 (2013)
2. Bergemann, D., Välimäki, J.: Information acquisition and efficient mechanism design. Econometrica **70**(3), 1007–1033 (2002)
3. Börgers, T., Krahmer, D., Strausz, R.: An Introduction to the Theory of Mechanism Design. Oxford University Press, Oxford (2015)
4. Chen, L., Leneutre, J.: A game theoretical framework on intrusion detection in heterogeneous networks. IEEE Trans. Inf. Forensics Secur. **4**(2), 165–178 (2009)
5. Cremer, J., McLean, R.P.: Full extraction of the surplus in bayesian and dominant strategy auctions. Econometrica **56**, 1247–1258 (1988)
6. Farhadi, F., Golestani, S.J.: Mechanism design for network resource allocation: a surrogate optimization approach (2016). http://ee.sharif.edu/~farhadi/Surrogate-Optimization-Approach.pdf
7. Farhadi, F., Tavafoghi, H., Teneketzis, D., Golestani, S.J.: A dynamic incentive mechanism for security in networks of interdependent agents (2016). http://ee.sharif.edu/~farhadi/Dynamic-Incentive-Mechanism.pdf
8. Fazeli, A., Ajorlou, A., Jadbabaie, A.: Competitive diffusion in social networks: quality or seeding? IEEE Trans. Control Netw. Syst. **PP**(99), 1 (2016)
9. Jackson, M., Zenou, Y.: Games on networks. In: Handbook of Game Theory with Economic Applications, vol. 4. Elsevier (2015)
10. Jiang, L., Anantharam, V., Walrand, J.: How bad are selfish investments in network security? IEEE/ACM Trans. Netw. **19**(2), 549–560 (2011)
11. Kakhbod, A., Teneketzis, D.: An efficient game form for unicast service provisioning. IEEE Trans. Autom. Control **57**(2), 392–404 (2012)
12. Kelly, F., Maulloo, A., Tan, D.: Rate control for communication networks: shadow prices, proportional fairness and stability. J. Oper. Res. Soc. **49**, 237–252 (1998)
13. Khouzani, M.H.R., Sarkar, S., Altman, E.: A dynamic game solution to malware attack. In: IEEE INFOCOM, April 2011
14. Kumar, P.R., Varaiya, P.: Stochastic Systems: Estimation, Identification, and Adaptive Control, vol. 75. SIAM (2015)
15. Laszka, A., Felegyhazi, M., Buttyan, L.: A survey of interdependent information security games. ACM Comput. Surv. **47**(2), 1–38 (2014)
16. Lelarge, M., Bolot, J.: Economic incentives to increase security in the internet: the case for insurance. In: IEEE INFOCOM, pp. 1494–1502 (2009)
17. Li, M., Koutsopoulos, I., Poovendran, R.: Optimal jamming attacks and network defense policies in wireless sensor networks. In: IEEE INFOCOM, May 2007
18. Liang, X., Xiao, Y.: Game theory for network security. IEEE Commun. Surv. Tutor. **15**(1), 101–120 (2013)

19. Liu, H.: Efficient dynamic mechanisms in environments with interdependent valuations. SSRN 2504731 (2014)
20. Myerson, R.B.: Multistage games with communication. Econom. J. Econom. Soc. **54**(2), 323–358 (1986)
21. Naghizadeh, P., Liu, M.: On the role of public and private assessments in security information sharing agreements. arXiv preprint arXiv:1604.04871 (2016)
22. Naghizadeh, P., Liu, M.: Opting out of incentive mechanisms: a study of security as a non-excludable public good. IEEE Trans. Inf. Forensics Secur. **11**(12), 2790–2803 (2016)
23. Nowzari, C., Preciado, V.M., Pappas, G.J.: Analysis and control of epidemics: a survey of spreading processes on complex networks. arXiv:1505.00768 (2015)
24. Trajanovski, S., Hayel, Y., Altman, E., Wang, H., Mieghem, P.V.: Decentralized protection strategies against SIS epidemics in networks. IEEE Trans. Control Netw. Syst. **2**(4), 406–419 (2015)