Designing Cyber Insurance Policies: Mitigating Moral Hazard Through Security Pre-Screening

Mohammad Mahdi Khalili(✉), Parinaz Naghizadeh, and Mingyan Liu

Electrical and Computer Science Department,
University of Michigan, Ann Arbor, USA
{khalili,naghizad,mingyan}@umich.edu

Abstract. Cyber-insurance has been studied as both a method for risk-transfer, as well as a potential incentive mechanism for improving the state of cyber-security. However, in the absence of regulated insurance markets or compulsory insurance, the introduction of insurance deteriorates network security. This is because by transferring part of their risk to the insurer, the insured agents can decrease their levels of effort. In this paper, we consider the design of insurance contracts by an (unregulated) profit-maximizing insurer, and allow for voluntary participation. We propose the use of pre-screening to offer premium discounts to higher effort agents. We show that such premium discrimination not only helps the insurer attain higher profits, but also leads the agents to improve their efforts. We show that with interdependent agents, the incentivized improvement in efforts can compensate for the effort reduction resulting from risk transfer, thus improving the state of network security over the no-insurance scenario. In other words, the availability of pre-screening signals benefits both the insurer, as well as the state of network security, without the need to regulate the market or compulsory participation.

1 Introduction

Organizations and businesses big and small are facing increasingly more complex, costly and frequent cyber threats. Many technology based protection methods such as novel cryptography schemes and protection softwares have been developed to reduce the risk of cyber threats. In addition to a myriad of technology based protection methods, cyber-insurance has emerged as an accepted risk mitigation mechanism, that allows purchasers of insurance policies/contracts to transfer their residual risks to the insurer.

The impact of cyber insurance on firms' security investment has been quite extensively studied in the past few years. These studies include cyber-insurance as a method for risk transfer, as well as a possible incentive mechanism for risk reduction, see e.g., [1–8]. Many papers on cyber insurance markets have studied the impact of cyber-insurance on the state of network security. Existing literature has arrived at two seemingly contradictory conclusions about the potential

This work is partially supported by the NSF under grant CNS-1616575.

L. Duan et al. (Eds.): GameNets 2017, LNICST 212, pp. 63–73, 2017.
DOI: 10.1007/978-3-319-67540-4_6

of cyber-insurance as an incentive mechanism for risk reduction. The difference is mainly due to the underlying model of the insurer/insurance market. In particular, when the cyber-insurance market is modeled as a competitive market, e.g., [7,8], the insurance contracts are designed with the intention of attracting clients, and are hence not optimized to induce better security behavior. As a result, [7,8] show that the introduction of cyber-insurance deteriorates network security. Furthermore, as a consequence of the assumption of competitive markets, the insurers make no profit.

On the other hand, by considering a monopolist (profit-neutral) cyber-insurer, whose goal is to increase social welfare, [3–7] show that it is possible to design cyber-insurance contracts that lead users to improve their efforts toward securing their systems, and consequently, improve the state of security. The works in [5–7] propose *premium discrimination*; the idea is to assign less favorable contracts (i.e., higher premiums) to agents with worse types or lower efforts. These contracts can lead to an increase in social welfare and network security, as well as non-negative profit for the insurer. However, the underlying models assume that the insurer acts to increase social welfare (due to e.g., government regulation), and is therefore not profit-maximizing. In addition, participation by agents is assumed compulsory.

In this paper, we are similarly interested in the possibility of using cyber-insurance as an incentive mechanism for improved network security. We modify two of the key existing assumptions, in order to better capture the current state of cyber-insurance markets, by (1) considering a profit-maximizing cyber-insurer, and (2) ensuring that participation is voluntary, i.e., agents may opt out of purchasing a contract.

We propose the use of *pre-screening* (initial audit) by the insurer; pre-screening allows the insurer to evaluate the potential client's security posture, prior to offering the contract. This essentially allows the insurer to premium-discriminate the agents, based on their perceived/measured state of security. We provide sufficient conditions under which the introduction of pre-screening can lead to higher profits for the insurer, and that it also positively impacts the state of security. In other words, this type of pre-screening is a potential option for making cyber-insurance contracts better drivers for improved cyber-security.

2 A Single Risk-Averse Agent

We first consider a single-period contract design problem between a risk-neutral cyber-insurer and a risk-averse agent.[1] The agent exerts *effort* $e \in [0, +\infty)$ towards securing his system, incurring a cost of c per unit of effort. Let L_e denote the loss, a random variable, that the agent experiences given his effort e. We assume L_e has a normal distribution, with mean $\mu(e) \geq 0$ and variance $\lambda(e) \geq 0$. We assume $\mu(e)$ and $\lambda(e)$ are strictly convex, strictly decreasing, and twice differentiable. The decreasing assumption entails that increased effort

[1] Throughout the paper, we use she/her and he/his to refer to the insurer and agent(s), respectively.

reduces the expected loss, as well as its unpredictability, for the agent. The convexity assumption suggests that while initial investment in security leads to considerable reduction in loss, the marginal benefit decreases as effort increases. We assume once a loss L_e is realized, it will be observed by both the cyber-insurer and agent through e.g., reporting and auditing. We further assume $\lambda(e)$ is small compared to $\mu(e)$, so that $Pr(L_e < 0)$ is negligible. Finally, when the agent exerts an effort e, the insurer observes a *pre-screening* signal $S_e = e + W$, where W is a zero mean Gaussian noise with variance σ^2. This signal can be attained through, e.g., external audits or initial surveys filled out by the agent. We assume S_e is conditionally independent of L_e, given e.

Linear Contract and Insurer's Payoff: In this paper, we consider the design of a set of *linear* contracts. Specifically, the contract offered by the insurer consists of a base premium p, a discount factor α, and a coverage factor β. The agent pays a premium $p - \alpha \cdot S_e$, and receives $\beta \cdot L_e$ as coverage. We let $0 \leq \beta \leq 1$, i.e., coverage never exceeds the actual loss. Thus the insurer's utility (profit) is given by:

$$V(p, \alpha, \beta, e) = p - \alpha \cdot S_e - \beta \cdot L_e.$$

The insurer's expected profit is then given by $\overline{V}(p, \alpha, \beta, e) = p - \alpha e - \beta \mu(e)$.

Agent's Payoff without a Contract: If the agent chooses not to enter a contract, he bears the full cost of his effort as well as any loss. We assume

$$U(e) = -\exp\{-\gamma \cdot (-L_e - ce)\}, \tag{1}$$

where γ denotes the *risk attitude* of the agent; a higher γ implies more risk aversion. We shall assume that γ is known to the insurer, thereby eliminating adverse selection and solely focusing on the moral hazard aspect of the problem.

Using basic properties of the normal distribution, we have the following expected utility for the agent:

$$\overline{U}(e) = E(-\exp\{-\gamma \cdot (-L_e - ce)\}) = -\exp\{\gamma \cdot \mu(e) + \frac{1}{2}\gamma^2 \lambda(e) + \gamma ce\}. \tag{2}$$

Using (2), the optimal effort for an agent outside the contract is given by $m := \arg\min_{e \geq 0} \left\{\mu(e) + \frac{1}{2}\gamma\lambda(e) + ce\right\}$. Let $U^o := \bar{U}(m)$ denote the maximum expected payoff of the agent without a contract.

Agent's Payoff with a Contract: If the agent accepts a contract, his utility is given by:

$$U^c(p, \alpha, \beta, e) = -\exp\{-\gamma \cdot (-p + \alpha \cdot S_e - L_e + \beta \cdot L_e - ce)\}.$$

Noting that S_e and L_e are conditionally independent, his expected utility is

$$\begin{aligned}\overline{U}^c(p, \alpha, \beta, e) &= E(-\exp\{-\gamma \cdot (-p + \alpha \cdot S_e - L_e + \beta \cdot L_e - ce)\}) \\ &= -\exp\left\{\gamma(p + (c - \alpha)e + \tfrac{1}{2}\alpha^2\gamma\sigma^2 + (1 - \beta)\mu(e) + \tfrac{1}{2}\gamma(1 - \beta)^2\lambda(e))\right\}\end{aligned}$$

The Insurer's Problem: The insurer designs the contract (p, α, β) to maximize her expected payoff. In doing so, the insurer also has to satisfy two constraints:

Individual Rationality (IR), and Incentive Compatibility (IC). The first stipulates that a rational agent will not enter a contract with payoff less than his outside option U^o, and the second that the effort desired by the insurer should maximize the agent's expected utility under that contract. Formally,

$$\begin{aligned} \max_{p,\alpha,0\leq\beta\leq 1} \quad & \bar{V}(p,\alpha,\beta,e) = p - \alpha\cdot e - \beta\cdot\mu(e) \\ \text{s.t.} \quad (\text{IR}) \quad & \bar{U}^c(p,\alpha,\beta,e) \geq U^o \\ (\text{IC}) \quad & e \in \arg\max_{e'\geq 0} \bar{U}^c(p,\alpha,\beta,e') \end{aligned} \tag{3}$$

Note that the (IR) constraint can be re-written as follows,

$$p + (c-\alpha)\cdot e + \tfrac{1}{2}\alpha^2\cdot\gamma\sigma^2 + (1-\beta)\mu(e) + \tfrac{1}{2}\gamma(1-\beta)^2\lambda(e) \leq u^o\ .$$

where, $u^o := \frac{\ln(-U^o)}{\gamma} = \min_{e\geq 0}\{\mu(e) + \frac{1}{2}\gamma\lambda(e) + c\cdot e\}$. Similarly, the (IC) constraint can be rearranged as follows,

$$e \in \arg\min_{e'\geq 0} \quad (c-\alpha)\cdot e' + (1-\beta)\mu(e') + \tfrac{1}{2}\gamma(1-\beta)^2\lambda(e').$$

3 The Role of Pre-screening in a Single Agent System

In this section, we first solve the optimization problem in (3). We then study the impact of several problem parameters, particularly the accuracy of pre-screening, on the optimal contract.

Lemma 1. *The (IR) constraint is binding in the optimal contract.*

By Lemma 1, an optimal contract satisfies the following equation:

$$p + (c-\alpha)\cdot e + \frac{1}{2}\alpha^2\cdot\gamma\sigma^2 + (1-\beta)\mu(e) + \frac{1}{2}\gamma(1-\beta)^2\lambda(e) = u^o\ .$$

We use the above expression to substitute for the base premium p in the objective function of (3), and re-writing the insurer's problem as follows,

$$\begin{aligned} &\max_{\alpha,0\leq\beta\leq 1,e\geq 0} f(\beta,e,\alpha) = u^o - \mu(e) - \tfrac{1}{2}\gamma(1-\beta)^2\lambda(e) - c\cdot e - \tfrac{1}{2}\alpha^2\gamma\sigma^2 \\ &\text{s.t.,} \qquad e = \arg\min_{e'\geq 0}(c-\alpha)\cdot e' + (1-\beta)\mu(e') + \tfrac{1}{2}\gamma(1-\beta)^2\lambda(e') \end{aligned} \tag{4}$$

We now turn to the issue of network security. We consider the effort level of the agent as the metric for evaluating the change in network security. We start with the following theorem on the state of network security, before and after the purchase of an insurance contract.

Theorem 1. *The effort exerted by the agent in the optimal contract is less than or equal to the level of effort outside the contract. In other words, insurance decreases network security as compared to the no-insurance scenario.*

Theorem 1 illustrates the inefficiency of cyber-insurance as a tool for improving the state of security. Existing work in [8,9] have also arrived at a similar conclusion when studying competitive/unregulated cyber-insurance markets. Nevertheless, as cyber-insurance is a profitable market, especially given risk-averse users, a market for cyber-insurance exists, and its growth is conceivable. We therefore ask whether the introduction of a pre-screening signal can lead to higher profits for the insurer, while also positively impacting the state of security, over the case of *no pre-screening*. We first analyze the impact of a pre-screening signal on the insurer's profit.

Theorem 2. *The insurer's payoff in the optimal contract increases as σ decreases. That is, the insurer's profit is increasing in the quality of the pre-screening signal.*

The above result is intuitively to be expected, as we predict that a strategic insurer can leverage the improved pre-screening information to her benefit, and attain better payoff. The more interesting observation is on the effect of pre-screening on the state of network security. The following theorem presents a sufficient condition under which the availability of a pre-screening signal improves network security, compared to the no pre-screening scenario. Note that we use $\sigma = \infty$ for evaluating the no pre-screening scenario. The equivalence follows from the fact that, as shown in [11], by setting $\sigma = \infty$, the insurer's optimal choice will be to set $\alpha = 0$, which effectively removes the effects of pre-screening.

Theorem 3. *Let e_1, e_2, e_∞ denote the optimal effort of the agent in the optimal contract when $\sigma = \sigma_1$, $\sigma = \sigma_2$ and $\sigma = \infty$, respectively. Let $k(e,\alpha) = \frac{\mu'(e)+\sqrt{\mu'(e)^2-2\gamma(c-\alpha)\lambda'(e)}}{-\gamma\lambda'(e)}$. If $k(e,\alpha_1)^2\lambda(e) - k(e,\alpha_2)^2\lambda(e)$ is non-decreasing in e for all $0 \le \alpha_1 \le \alpha_2 \le c$, then $e_1 \ge e_2$ if $\sigma_1 \le \sigma_2$. In other words, better pre-screening signals improve network security.*

In addition, if $k(e,0)^2\lambda(e) - k(e,\alpha)^2\lambda(e)$ is non-decreasing in e for all $0 \le \alpha \le c$, then $e_1 \ge e_\infty$. In other words, availability of a pre-screening signal improves network security over the no pre-screening scenario.

In the above theorem, $k(e,\alpha)$ is in fact equivalent to $1-\beta$. Consequently, $k(e,\alpha)^2\lambda(e)$ is the variance of the uncovered loss in a contract as a function (e,α). Therefore, Theorem 3 introduces a sufficient condition for improvement of network security based on the change in the variance of the uncovered loss.

Several instances of $\mu(e)$ and $\lambda(e)$ satisfy the condition of Theorem 3; for instance, $(\mu(e) = \frac{1}{e}, \lambda(e) = \frac{1}{e^2})$ or $(\mu(e) = \exp\{-e\}, \lambda(e) = \exp\{-2e\})$. Theorems 2 and 3 together imply that the introduction of a pre-screening signal benefits the insurer, as well the state of network security.

4 A Network of Two Risk Averse Agents

We next consider the one period contract problem between one risk-neutral insurer and two risk-averse agents. We assume the agents' utilities are again given by (1), and let γ_1, γ_2 denote the risk attitudes of the agents. We assume that the two agents are interdependent; the effort exerted by an agent affects not only himself, but further affects the loss that the other agent experiences. This assumption captures the fact that viruses, worms, etc., can spread from an infected agent to others. We model the interdependence between these two agents as follows,

$$L^{(i)}_{e_1,e_2} \sim \mathcal{N}(\mu(e_i + x \cdot e_{-i}), \lambda(e_i + x \cdot e_{-i}))$$

Here, $\{-i\} = \{1, 2\} - \{i\}$, and $L^{(i)}_{e_1,e_2}$ is a random variable denoting the loss that agent i experiences, given both agents' efforts. The *interdependence factor* is denoted by x, and we let $0 \leq x < 1$.

The insurer can observe the result of pre-screening audit $S_{e_i} = e_i + W_i$ on each agent i, where W_i is a zero mean Gaussian noise with variance σ_i^2. We assume that W_1 and W_2 are independent, and that $S_{e_1}, S_{e_2}, L^{(1)}_{e_1,e_2}, L^{(2)}_{e_1,e_2}$ are conditionally independent given e_1, e_2.

We next separately analyze the following three cases, based on whether agents purchase cyber-insurance contracts.

(i) Neither agent enters a contract

(ii) One of the agents enters a contract, while the other one opts out

(iii) Both agents purchase contracts

Note that Case (ii) is the outside option for agents in Case (iii), and Case (i) is the outside option for agents in Case (ii). Therefore, in order to evaluate the participation constraints of agents when both purchase insurance contracts, we first need to find the optimal contract and agents' payoffs in Cases (i) and (ii).

4.1 Case (i): Neither Agent Enters a Contract

We start by considering the game G_{oo} between two agents, neither of which have purchased cyber-insurance contracts. The expected payoffs of these agents, with unit costs of effort $c_1, c_2 > 0$, are given by,

$$\bar{U}_i(e_1, e_2) = -\exp\{\gamma_i \mu(e_i + x \cdot e_{-i}) + \tfrac{1}{2}\gamma_i^2 \lambda(e_i + x \cdot e_{-i}) + \gamma_i \cdot c_i \cdot e_i\}$$

The best-response of each agent, when both opt out, can be found by solving the following optimization problem,

$$\begin{aligned} B_i^{oo}(e_{-i}) &= \arg\max\nolimits_{e_i \geq 0} -\exp\{\gamma_i \mu(e_i + x \cdot e_{-i}) + \tfrac{1}{2}\gamma_i^2 \lambda(e_i + x \cdot e_{-i}) + \gamma_i \cdot c_i \cdot e_i\} \\ &= \arg\min\nolimits_{e_i \geq 0} \mu(e_i + x \cdot e_{-i}) + \tfrac{1}{2}\gamma_i \lambda(e_i + x \cdot e_{-i}) + c_i \cdot e_i \ . \end{aligned} \tag{5}$$

The above optimization problem is a convex optimization problem and has a unique solution. In order to find $B_i^{oo}(e_{-i})$, we first define m_i as follows,

$$m_i := \arg\min_{e \geq 0} \{\mu(e) + \frac{1}{2}\gamma_i \lambda(e) + c_i \cdot e\} \tag{6}$$

Using (6), the solution to (5) is given by,

$$B_i^{oo}(e_{-i}) = \begin{cases} m_i - x \cdot e_{-i} & \text{if } m_i \geq x \cdot e_{-i} \\ 0 & \text{if } m_i \leq x \cdot e_{-i} \end{cases} \tag{7}$$

The Nash equilibrium is given by the fixed point of the best-response mappings $B_1(e_2)$ and $B_2(e_1)$. Let $e_i^*(m_i, m_{-i})$ denote the effort of agent i at the *unique* Nash equilibrium. We have,

$$e_i^*(m_i, m_{-i}) = \begin{cases} \frac{m_i - x \cdot m_{-i}}{1-x^2} & \text{if } m_i \geq x \cdot m_{-i} \text{ and } m_{-i} \geq x \cdot m_i \\ 0 & \text{if } m_i \leq x \cdot m_{-i} \\ m_i & \text{if } m_{-i} \leq x \cdot m_i \end{cases} \tag{8}$$

Therefore, $\bar{U}_i^{*oo} = \bar{U}_i(e_1^*(m_1, m_2), e_2^*(m_2, m_1))$ is the utility of agent i in the equilibrium when agents do not choose to enter the contract. As we will see shortly, an insurer uses her knowledge of $\bar{U}_i^{*oo}$ to evaluate agents' outside options when proposing optimal contracts.

4.2 Case (ii): One of the Agents Enters a Contract

Assume that agent 1 enters a contract, while agent 2 opts out. We use G_{io} to denote the game between the insured agent 1 and uninsured agent 2. The expected payoffs of agents in this game are as follows,

$$\begin{aligned} U_1^{io}(e_1, e_2, p_1, \alpha_1, \beta_1) = & \\ & E(-\exp\{-\gamma_1 \cdot (-p_1 + \alpha_1 \cdot S_{e_1} - L_{e_1,e_2}^{(1)} + \beta_1 \cdot L_{e_1,e_2}^{(1)} - c_1 \cdot e_1)\}) \\ & = -\exp\{\gamma_1 \cdot (p_1 + (c_1 - \alpha_1) \cdot e_1 + \tfrac{1}{2}\alpha_1^2 \cdot \gamma_1 \sigma_1^2 \\ & +(1-\beta_1)\mu(e_1 + x \cdot e_2) + \tfrac{1}{2}\gamma_1(1-\beta_1)^2\lambda(e_1 + x \cdot e_2))\} \\ U_2^{io}(e_1, e_2) & = E(-\exp\{-\gamma_2(-L_{e_1,e_2}^{(2)} - c_2 \cdot e_2)\}) \\ & = -\exp\{\gamma_2\mu(e_2 + x \cdot e_1) + \tfrac{1}{2}\gamma_2^2\lambda(e_2 + x \cdot e_1) + \gamma_2 \cdot c_2 \cdot e_2\} \end{aligned}$$

In order to find the Nash Equilibrium of G_{io}, we first calculate the best response of each agent. Let $B_i^{io}(e_{-i})$ denote the best response of agent i. We have,

$$\begin{aligned} & B_1^{io}(e_2) = \arg\max_{e_1 \geq 0} - \exp\{\gamma_1 \cdot (p_1 + (c_1 - \alpha_1) \cdot e_1 \\ & + \tfrac{1}{2}\alpha_1^2 \cdot \gamma_1\sigma_1^2 + (1-\beta_1)\mu(e_1 + x \cdot e_2) + \tfrac{1}{2}\gamma_1(1-\beta_1)^2\lambda(e_1 + x \cdot e_2))\} \\ & = \arg\min_{e_1 \geq 0}(c_1 - \alpha_1) \cdot e_1 + (1-\beta_1)\mu(e_1 + x \cdot e_2) + \tfrac{1}{2}\gamma_1(1-\beta_1)^2\lambda(e_1 + x \cdot e_2) \end{aligned} \tag{9}$$

As the above optimization problem is a convex problem, it has a unique solution. We next define $m_1(\alpha_1, \beta_1)$ as follows,

$$m_1(\alpha_1, \beta_1) = \arg\min_{e \geq 0}\{(c_1 - \alpha_1)e + (1-\beta_1)\mu(e) + \frac{1}{2}\gamma_1(1-\beta_1)^2\lambda(e)\}$$

Similar to (7), we use $m_1(\alpha_1, \beta_1)$ to find $B_1^{io}(e_2)$ as follows,

$$B_1^{io}(e_{-i}) = \begin{cases} m_1(\alpha_1, \beta_1) - x \cdot e_2 & \text{if } m_1(\alpha_1, \beta_1) \geq x \cdot e_2 \\ 0 & \text{if } m_1(\alpha_1, \beta_1) \leq x \cdot e_2 \end{cases} \tag{10}$$

For the uninsured agent 2, it is easy to see that the best-response function is given by $B_2^{io}(e_1) = B_2^{oo}(e_1)$.

We can now find the Nash equilibrium as the fixed point of the best-response mappings. Agents' efforts at the equilibrium are $e_1^*(m_1(\alpha_1,\beta_1), m_2)$ and $e_2^*(m_2, m_1(\alpha_1,\beta_1))$ which are defined in (8). For notational convenience, we denote these efforts by e_1^*, e_2^*. Let $\bar{U}_i^{*io}$ denote the utility of agent i at effort levels e_1^*, e_2^*, in an equilibrium where only agent 1 purchases a contract, that is,

$$\bar{U}_1^{*io}(p_1,\alpha_1,\beta_1) = U_1^{io}(e_1^*, e_2^*, p_1, \alpha_1, \beta_1), \quad \bar{U}_2^{*io}(\alpha_1,\beta_1) = U_2^{io}(e_1^*, e_2^*)$$

Let $\bar{V}^{io}(p_1,\alpha_1,\beta_1,e_1,e_2)$ denote the insurer's utility, when she offers contract (p_1,α_1,β_1) to agent 1, and agents exert efforts e_1, e_2. The optimal contract offered by the insurer is the solution to the following optimization problem:

$$\begin{aligned} V^{*io} &= \max\nolimits_{p_1,\alpha_1,\beta,e_1^*,e_2^*} \bar{V}^{io}(p_1,\alpha_1,\beta_1,e_1^*,e_2^*) = p_1 - \alpha_1 e_1^* - \beta_1 \cdot \mu(e_1^* + x \cdot e_2^*) \\ s.t., &\quad \text{(IR) } \bar{U}_1^{*io}(p_1,\alpha_1,\beta_1) \geq \bar{U}_1^{*oo}, \\ &\quad \text{(IC) } e_1^*, e_2^* \text{ are effort of the agents in Nash equilibrium of game } G_{io} \end{aligned}$$

We first re-write the (IR) constraint for agent 1 as follows,

$$p_1 + (c_1-\alpha_1)\cdot e_1^* + \frac{1}{2}\alpha_1^2\gamma_1\sigma_1^2 + (1-\beta_1)\mu(e_1^*+x\cdot e_2^*) + \frac{1}{2}\gamma_1(1-\beta_1)^2\lambda(e_1^*+x\cdot e_2^*) \leq u_1^{oo},$$

where $u_1^{oo} = \frac{\ln(-\bar{U}_1^{*oo})}{\gamma_1}$.

Similar to Lemma 1, we can conclude that (IR) constraint is binding in the optimal contract. Therefore, we can re-write the insurer's problem by replacing for the base premium p, similar to the single agent problem in Sect. 2.

4.3 Case (iii): Both Agents Purchase Contracts

Assume the insurer offers each agent i a contract (p_i, α_i, β_i). The expected utility of agents when both purchase contracts is given by,

$$\begin{aligned} U_j^{(ii)}&(e_1,e_2,p_j,\alpha_j,\beta_j) = \\ &E(-\exp\{-\gamma_j\cdot(-p_j + \alpha_j\cdot S_{e_j} - L_{e_1,e_2}^{(j)} + \beta_j\cdot L_{e_1,e_2}^{(j)} - c_j\cdot e_j)\}) \\ &= -\exp\{\gamma_j\cdot(p_j + (c_j-\alpha_j)\cdot e_j + \tfrac{1}{2}\alpha_j^2\cdot\gamma_j\sigma_j^2 + \\ &(1-\beta_j)\mu(e_j + x\cdot e_{-j}) + \tfrac{1}{2}\gamma_j(1-\beta_j)^2\lambda(e_j + x\cdot e_{-j}))\} \end{aligned}$$

Following steps similar to those in Sect. 4.2, the best-response function of player j, denoted B_j^{ii}, is given by,

$$B_j^{ii}(e_{-j}) = \begin{cases} m_j(\alpha_j,\beta_j) - x\cdot e_{-j} & \text{if } m_j(\alpha_j,\beta_j) \geq x\cdot e_{-j} \\ 0 & \text{if } m_j(\alpha_j,\beta_j) \leq x\cdot e_{-j} \end{cases}$$

where $m_j(\alpha_j,\beta_j)$ is the solution of the following equation,

$$m_j(\alpha_j,\beta_j) = \arg\min\nolimits_{e\geq 0}(1-\beta_j)\mu(e) + \tfrac{1}{2}\gamma_j(1-\beta_j)^2\lambda(e) + (c_j-\alpha_j)\cdot e. \quad (11)$$

Agents' efforts at the *unique* Nash equilibrium are $e_i^*(m_i(\alpha_i, \beta_i), m_{-i}(\alpha_{-i}, \beta_{-i}))$, with $e_i^*(.,.)$ defined in (8). For notational convenience, we simply denote these by e_i^*.

To write the insurer's problem, note that the outside option of agent 1 (resp. 2) from this game is the game G_{oi} (resp. G_{io}). The IR constraints can again be shown to be binding, simplifying the insurer's problem to,

$$\begin{aligned}
&V^{*ii} = \max_{\alpha_1, 0\leq\beta_1\leq 1, \alpha_2, 0\leq\beta_2\leq 1, e_1^*\geq 0, e_2^*\geq 0} u_1^{oi} - \mu(e_1^* + x\cdot e_2^*) \\
&-\tfrac{1}{2}\gamma_1(1-\beta_1)^2\lambda(e_1^* + x\cdot e_2^*) - c_1\cdot e_1^* - \tfrac{1}{2}\alpha_1^2\gamma_1\sigma_1^2 \\
&+u_2^{io} - \mu(e_2^* + x\cdot e_1^*) - \tfrac{1}{2}\gamma_2(1-\beta_2)^2\lambda(e_2^* + x\cdot e_1^*) - c_2\cdot e_2^* - \tfrac{1}{2}\alpha_2^2\gamma_2\sigma_2^2 \\
&s.t., e_1^*, e_2^* \text{ are the agents' effort in the equilibrium of game } G_{ii}
\end{aligned}$$

where $u_1^{oi} = \frac{\ln(-\bar{U}_1^{*oi})}{\gamma_1}$ and $u_2^{io} = \frac{\ln(-\bar{U}_2^{*io})}{\gamma_2}$, defined in Sect. 4.2.

5 The Role of Pre-screening in a Two Agent Network

We next discuss how different problem parameters, particularly the accuracy of pre-screening, affect the insurer's profit, as well as the system's state of security.

We first consider the utility of the insurer. As the insurer always has the option to not use the outcome of pre-screening by setting $\alpha = 0$ in the contract, the insurer's profit in the optimal contract with pre-screening is larger than her profit in the optimal contract without pre-screening; i.e., the availability of pre-screening is in the insurer's interest and improves insurer's profit.

We now return to the effect of pre-screening on the state of network security. We choose the total effort towards security, $e_1 + e_2$, as the metric for evaluating network security. The following two theorems characterize the impact of pre-screening on network security when the two agents are homogeneous ($\gamma_1 = \gamma_2 = \gamma, c_1 = c_2 = c, \sigma_1 = \sigma_2 = \sigma$). Theorem 4 shows that fully accurate pre-screening can improve network security over the no insurance scenario. Theorem 5 shows that under certain additional conditions, the improvement is still possible for sufficiently, yet not fully, accurate pre-screening.

Theorem 4. *Assume two homogeneous agents purchase (identical) contracts from an insurer, and let* $m = \arg\min_{e\geq 0} \mu(e) + \frac{1}{2}\gamma\lambda(e) + ce$.

(i) If $\mu'(m) < -\frac{c}{1+x}$ *and both pre-screening signals are accurate, i.e.,* $\sigma_1 = \sigma_2 = 0$, *then network security improves after the introduction of insurance.*

(ii) If both of the pre-screening signals are uninformative, i.e., $\sigma_1 = \infty$ *and* $\sigma_2 = \infty$, *network security worsens after the introduction of insurance.*

Theorem 5. *Assume two homogeneous agents purchase (identical) contracts from an insurer. Let* $m = \arg\min_{e\geq 0} \mu(e)+\frac{1}{2}\gamma\lambda(e)+ce$, $u_{max} = \mu(m)+\frac{1}{2}\gamma\lambda(m)+cm$, *and* $h(m', \beta) = c\cdot m' + (1-\beta)\mu(m') + \frac{1}{2}\gamma(1-\beta)^2\lambda(m')$. *If* $\mu'(m) < -\frac{c}{1+x}$, *then there exists an upper bound* $\sigma_{max}^2 := \min\{\frac{-\mu(m) - \frac{c\cdot m}{1+x} + \mu(0)}{0.5c^2\gamma}, \frac{-\mu'(m) - \frac{c}{1+x}}{M\gamma}\}$, *where*

$$M := \max_{0\leq\beta\leq 1, 0\leq m'\leq \frac{(1+x)u_{max}}{c}} \{\frac{\partial h(m', \beta)}{\partial m'} \cdot \frac{\partial^2 h(m', \beta)}{\partial m'^2}\},$$

such that if $\sigma_1^2 = \sigma_2^2 \leq \sigma_{max}^2$*, the existence of pre-screening improves network security as compared to the no insurance scenario.*

6 Conclusion

We studied the problem of designing cyber-insurance contracts by a single profit-maximizing insurer, for both a single agent, as well as two interdependent agents. The introduction of insurance decreases network security in general, as agents reduce their effort after transferring part of their risks to an insurer. We propose the use of pre-screening signals on agents' efforts to prevent such reduction in effort after the introduction of insurance contracts, by offering premium discounts to agents with higher perceived efforts. We show that the availability of these pre-screening signals not only benefits the insurer by increasing her profit, but also improves network security, as compared to the no pre-screening scenario. Furthermore, when agents are interdependent and pre-screening is highly accurate, under a set of sufficient conditions, the incentivized improved efforts can increase network security not only over no pre-screening, but also compared to the no-insurance scenario. Therefore, introduction of pre-screening signals can be in the interest of the insurer, as well as the state of network security.

An important extension of this work is to consider arbitrary alternatives for including the pre-screening signals (as opposed to only linear discounts on premiums), and verify their role in improving network security. Considering multiple profit-maximizing insurers is another direction of future work.

Online Appendix. Numerical simulations and proofs are given in [11].

References

1. Bohme, R.: Cyber-insurance Revisited. Workshop on Economics Information Security (WEIS) (2005)
2. Bohme, R., Schwartz, G.: Modeling Cyber-Insurance: Towards a Unifying Framework. Workshop on Economics Information Security (WEIS) (2010)
3. Kesan, J., Majuca, R., Yurcik, W.: The economic case for cyberinsurance. In: Securing Privacy in the Internet Age. Stanford University Press (2005)
4. Kesan, J., Majuca, R., Yurcik, W.: Cyberinsurance as a Market-Based Solution to the Problem of Cybersecurity: a Case Study. Workshop on Economics Information Security (WEIS) (2005)
5. Hofmann, A.: Internalizing externalities of loss prevention through insurance monopoly. In: Proceeding Annual Meeting of American Risk and Insurance Association (2006)
6. Bolot, J., Lelarge, M.: Cyber-Insurance as an Incentive for Internet Security. Workshop on Economics Information Security (WEIS) (2008)
7. Pal, R., Golubchik, L., Psounis, K., Hui, P.: Will cyber-insurance improve network security? A market analysis. In: IEEE INFOCOM (2014)
8. Shetty, N., Schwartz, G., Felegyhazi, M., Walrand, J.: Competitive cyber-insurance and internet security. In: Economics of Information Security and Privacy (2010)

9. Shetty, N., Schwartz, G., Felegyhazi, M., Walrand, J.: Can competitive insurers improve network security?. In: International Conference on Trust and Trustworthy Computing (2010)
10. Osborne, M.J., Rubinstein, A.: A Course in Game Theory. MIT Press, Cambridge (1994)
11. www.dropbox.com/sh/euux09td56kqdnh/AAB6PGXFqa3BSbYyN4l1s5FZa?dl=0