Designing Cyber Insurance Policies in the Presence of Security Interdependence*

Mohammad Mahdi Khalili
EECS, University of Michigan
khalili@umich.edu

Parinaz Naghizadeh
EECS, University of Michigan
naghizad@umich.edu

Mingyan Liu
EECS, University of Michigan
mingyan@umich.edu

ABSTRACT

Cyber insurance is a method for risk transfer but may or may not improve the state of network security. In this work, we consider a profit-maximizing insurer with voluntarily participating insureds. We are particularly interested in two features of cybersecurity and their impact on the contract design problem. The first is the interdependent nature of cybersecurity, whereby one entity's state of security depends on its own effort and others' effort. The second is our ability to perform accurate quantitative assessment of security posture at a firm level by combining recent advances in Internet measurement and machine learning techniques. We observe that security interdependency leads to a "profit opportunity" for the insurer, created by the inefficient effort levels exerted by agents who do not account for risk externalities when insurance is not available; this is in addition to risk transfer that an insurer profits from. Security pre-screening allows the insurer to take advantage of this opportunity by designing appropriate contracts which incentivize agents to increase their effort levels, allowing the insurer to effectively "sell commitment" to interdependent agents, in addition to risk transfer. We identify conditions under which this type of contracts lead to an improved state of network security.

CCS CONCEPTS

•Networks → Network economics; •Security and privacy → Network security;

1 INTRODUCTION

Faced with a myriad of increasingly costly and frequent cyber threats, organizations and businesses not only invest in software security mechanisms such as firewalls and intrusion detection systems, but increasingly also turn to ways of better managing the risk. Within this context, cyber insurance has emerged as an accepted risk mitigation mechanism, that allows purchasers of insurance policies to transfer their residual risks to the insurer.

The design of cyber insurance contracts, and their effects on firms' security behavior, has been extensively studied in the literature [2, 5, 8, 11, 15, 17–20]. These studies show that the impact of cyber insurance on firms' investments, and the resulting state of network security, depend on the assumptions on the insurance market and the assumed model of interdependency among firms.

In particular, for the market model, existing literature has considered either competitive or monopolistic insurance markets. The works in [15, 17–19] study competitive insurance markets with interdependent agents under voluntary insurance. The authors of [18, 19] consider a competitive market with homogeneous agents, and show that insurance often deteriorates the state of network security as compared to the no-insurance scenario. [17] studies a network of heterogeneous agents and shows that the introduction of insurance cannot improve the state of network security.

The impact of insurance on the state of network security in the presence of a monopolistic profit neutral (welfare maximizing) insurer has been studied in [4, 5, 8, 16]. In these models, as the insurer's goal is to maximize social welfare, under the assumption that agents' participation in the market can be made compulsory, insurance contracts can lead to improvement of network security. Agents' incentives for improving their investments is due to premium discrimination, i.e., agents with higher investments in security pay lower premiums. An insurance market with a monopolistic profit maximizing insurer, under the assumption of voluntary participation, has been studied in [11], which shows that in the presence of moral hazard, insurance cannot improve network security as compared to the no-insurance scenario.

In this paper, we are similarly interested in understanding the role of cyber insurance and its efficacy as an incentive mechanism for improved network security. We adopt two key assumptions as in [11], which we believe better capture the current state of cyber insurance markets but differ from the majority of existing literature; we assume (1) a profit-maximizing cyber insurer, and (2) voluntary participation, i.e., agents may opt out of purchasing a contract.

Furthermore, we focus on the effects of two distinct features of cybersecurity in the context of cyber insurance. The first is the interdependent nature of cybersecurity, whereby one entity's state of security depends on not only its own investment and effort, but also on the investments and efforts of others in the same eco-system (i.e., externalities), see e.g., [6, 7, 10, 14]. In other words, the risk that an insured transfers to the insurer is not only a function of its own actions, but also of other entities' actions who may or may not be seeking to transfer risks. The second distinct feature is the fact that recent advances in Internet measurement combined with machine learning techniques now allow us to perform accurate, quantitative security posture assessments at a firm level [12]. Such assessments can be used to mitigate information asymmetry about agents' security posture, and also as audit tools, allowing for coordination on higher effort levels by interdependent agents, see e.g., [3]. Here, we are interested in the former, i.e., the use of security posture assessments as a tool to perform an initial security audit, or *pre-screening*, of a prospective client to better enable premium discrimination and the design of customized policies.

Towards this end, we present models that take into account both interdependence and security assessments, in addition to the profit-maximizing insurer and voluntary participation assumptions mentioned earlier. We are interested in understanding the impact

*The work is partially supported by the NSF under grants CNS-1422211 and CNS-1616575. A preliminary version of this work appeared in the International Conference on Game Theory for Networks (Gamenets) 2017 [9]. Numerical simulations and proofs are available in the online appendix [1].

of each of these elements, as well as their combined effect, on the profit of the insurer, the participation incentives, and last but not least, the state of security with and without insurance. In doing so we also try to separate out the impact of these elements from that induced by risk aversion and risk transfer.

Our main findings are as follows: (1) Security interdependency leads to a "profit opportunity" for the insurer, created by the inefficient effort levels exerted by free-riding agents when insurance is not available; the profit resulting from this inefficiency gap is in addition to the risk transfer that an insurer typically profits from given risk averse agents. We shall use risk neutral agent models to demonstrate this phenomenon, as in this case risk transfer does not exist and were it not for interdependency the market for insurance would not even exist; yet, we show that because of interdependence, there exists a profit opportunity for the insurer. In other words, because of the spill-over effect of security investments, the improved security posture of one agent not only benefits this agent itself but all other agents dependent on it, resulting in a multiplication effect of benefit to the insurer in the form of reduced total loss. (2) Security pre-screening allows the insurer to take advantage of this additional profit opportunity, by designing appropriate contracts that incentivize agents to increase their effort levels. We conclude that in the presence of security interdependency, the insurer can essentially *sell commitment* to agents with dependent risks by engaging them in contracts designed to incentive high levels of security effort: one is guaranteed of another's higher effort after purchasing a contract.

The remainder of the paper is organized as follows. We present the model and preliminaries in Section 2. We analyze the optimal contracts for two risk-neutral and two risk-averse agents in Section 3. We discuss an N-homogeneous-agent case and possible extensions in Section 4, and conclude in Section 5.

2 MODEL AND PRELIMINARIES

We consider a single-period contract design problem between a risk-neutral insurer and either one or two agents[1]; we refer the interested reader to [13] for an overview of contract theory. We are primarily interested in the contract design problem for risk-averse agents; however, we will also examine risk-neutral agents because it allows us to isolate risk transfer from other factors.

In a single agent model, 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 implies that increased effort 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 is observed by both the insurer and the agent through, e.g., claims filed by the agent and post-incident audit by the insurer. We further assume that $\lambda(e)$ is sufficiently small compared to $\mu(e)$, so that $Pr(L_e < 0)$ is negligible.[2]

[1] We will use she/her and he/his to refer to the insurer and agent(s), respectively.
[2] Extension to other types of risk distributions remains a direction of future work.

In the two interdependent agent model, we will assume agent i's loss is given by,

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

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 $0 \leq x < 1$.

When designing cyber insurance contracts for such agent(s), although an agent's exerted effort is not observable by the insurer (i.e., there is moral hazard), we assume that the insurer can conduct a *pre-screening* of the agent's security standing. Through pre-screening, the insurer obtains a pre-screening assessment or outcome $S_{e_i} = e_i + W_i,\ i = 1, 2$, on each agent, where W_i is a zero mean Gaussian noise with variance σ_i^2. We assume W_i's are independent, and that the agent(s) and the insurer know the distribution of S_{e_i}. We also assume S_{e_i} is conditionally independent of L_{e_i}, given e_i. The pre-screening outcome S_{e_i} will be used by the insurer in determining the terms of the contract.

We next present preliminary results on the contract design problem in the single-agent case.

Linear Contract and the Insurer's Payoff. We consider the design of a set of *linear* contracts. Specifically, the contract offered by the insurer to an agent 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 in the event of a loss. 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 $\overline{V}(p, \alpha, \beta, e) = p - \alpha e - \beta\mu(e)$.

Risk-Neutral Agent. The utility (and expected utility) of a risk-neutral agent is given by,

$$U^{\text{out}}(e) = -L_e - ce \ \Rightarrow\ \overline{U}^{\text{out}}(e) = \mathbb{E}[U^{\text{out}}(e)] = -\mu(e) - ce \quad (1)$$

If the agent chooses not to enter a contract, he bears the full cost of his effort as well as any realized loss. Therefore, the optimal effort of the agent outside the contract, m, is

$$m = \arg\min_{e \geq 0}\ \mu(e) + ce\ ,$$

and his expected utility outside the contract is $u^o := \overline{U}^{\text{out}}(m)$.

On the other hand, if the agent purchases a contract (p, α, β) from the insurer, then his utility, and expected utility, are given by:

$$U^{\text{in}}(p, \alpha, \beta, e) = -p + \alpha S_e - L_e + \beta L_e - ce \Rightarrow$$
$$\overline{U}^{\text{in}}(p, \alpha, \beta, e) = \mathbb{E}[U^{\text{in}}(p, \alpha, \beta, e)] = -p + (\alpha - c)e + (\beta - 1) \cdot \mu(e)$$

Risk-Averse Agent. For simplicity, we use the same notation for risk-averse agents as for risk-neutral agents. The utility of a risk-averse agent is given by:

$$U^{\text{out}}(e) = -\exp\{-\gamma \cdot (-L_e - ce)\}\ , \quad (2)$$

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:

$$\overline{U}^{\text{out}}(e) = \mathbb{E}[U^{\text{out}}(e)] = -\exp\{\gamma\mu(e) + \frac{1}{2}\gamma^2\lambda(e) + \gamma ce\} \ . \quad (3)$$

Using (3), the optimal effort for an agent outside the contract is given by $m := \arg\min_{e\geq 0}\left\{\mu(e) + \frac{1}{2}\gamma\lambda(e) + ce\right\}$. Again, let $u^o := \overline{U}^{\text{out}}(m)$ denote the expected payoff of the agent without a contract.

If a risk-averse agent accepts a contract (p, α, β), his utility is given by:

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

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

$$\overline{U}^{\text{in}}(p, \alpha, \beta, e) = \mathbb{E}[U^{\text{in}}(p, \alpha, \beta, e)] =$$
$$-\exp\{\gamma(p + (c-\alpha)e + \tfrac{1}{2}\alpha^2\gamma\sigma^2 + (1-\beta)\mu(e) + \tfrac{1}{2}\gamma(1-\beta)^2\lambda(e))\} \ .$$

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 expected 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,

$$\begin{aligned} \max_{p, \alpha\geq 0, 0\leq\beta\leq 1, e\geq 0} \quad & \overline{V}(p, \alpha, \beta, e) = p - \alpha e - \beta\mu(e) \\ \text{s.t. (IR)} \quad & \overline{U}^{\text{in}}(p, \alpha, \beta, e) \geq u^o \qquad (4) \\ \text{(IC)} \quad & e \in \arg\max_{e'\geq 0} \overline{U}^{\text{in}}(p, \alpha, \beta, e') \end{aligned}$$

The above optimization problem can be simplified, for risk-neutral and risk-averse agents, respectively. First, note that as the base premium is a constant in the contract, the (IC) constraint for a risk-neutral agent can be rewritten as follows:

$$e \in \arg\min_{e'\geq 0} \ (c-\alpha)e' + (1-\beta)\mu(e') \ .$$

Similarly, the (IC) constraint for a risk-averse agent is:

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

Next, we can simplify the (IR) constraint using the following lemma; the proof can be found in the online appendix [1].

Lemma 2.1. *The (IR) constraint is binding in the optimal contract.*

By lemma 2.1, the (IR) constraint of a risk-neutral agent can be written as:

$$p + (c-\alpha)e + (1-\beta)\mu(e) = -u^o \ ,$$

and, for a risk-averse agent,

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

where $w^o := \frac{\ln(-u^o)}{\gamma} = \min_{e\geq 0}\{\mu(e) + \frac{1}{2}\gamma\lambda(e) + ce\}$.

Using the above expressions to substitute for the base premium p in the objective function in (4), and using the simplified expressions for the (IC) constraints, we rewrite the insurer's contract design problem as follows.

Insurer's problem with a risk-neutral agent:

$$\begin{aligned} \max_{\alpha\geq 0, 0\leq\beta\leq 1, e\geq 0} \quad & -u^o - \mu(e) - ce \\ \text{s.t.,} \quad & e \in \arg\min_{e'\geq 0} \ (c-\alpha)e' + (1-\beta)\mu(e') \qquad (5) \end{aligned}$$

Insurer's problem with a risk-averse agent:

$$\max_{\alpha\geq 0, 0\leq\beta\leq 1, e\geq 0} w^o - \mu(e) - \frac{1}{2}\gamma(1-\beta)^2\lambda(e) - ce - \frac{1}{2}\alpha^2\gamma\sigma^2$$
$$\text{s.t.,} \ \ e \in \arg\min_{e'\geq 0} \ (c-\alpha)e' + (1-\beta)\mu(e') + \frac{1}{2}\gamma(1-\beta)^2\lambda(e') \quad (6)$$

We now solve the contract design problems posed in (5) and (6).

Risk-Neutral Agent (Problem (5)*).* In this case, the objective function of the insurer is given by $-u^o - \mu(e) - ce$. However, note that $u^o = \max_{e\geq 0}\{-\mu(e) - ce\}$, and therefore the insurer's profit is at most zero. A contract with $(p = 0, \alpha = 0, \beta = 0)$ will yield a payoff of zero, making it an optimal contract. We thus conclude that it is optimal for the insurer to not offer a contract to a risk-neutral agent. Also note that in this case the quality of pre-screening, or indeed the availability of pre-screening regardless of the quality, plays no role in either the insurer's or agent's decisions.

Risk-Averse Agent (Problem (6)*).* We have the following theorem on the state of network security, defined as the effort exerted by the agent, before and after the purchase of an insurance contract.

Theorem 2.2. *Assume that $(\hat{\alpha}, \hat{\beta}, \hat{e})$ solves optimization problem* (6)*. Then $\hat{e} \leq m$, where m is the level of effort outside the contract.*

Theorem 2.2 illustrates the inefficiency of cyber insurance as a tool for improving the state of security. Existing work in [16, 18] have also arrived at a similar conclusion when studying competitive/unregulated cyber insurance markets. Note also that Theorem 2.2 holds regardless of the pre-screening quality. We next examine the effect of pre-screening on the insurer's profit.

Theorem 2.3. *Let $v(\alpha, \beta, e, \sigma^2) = w^o - \mu(e) - \frac{1}{2}\gamma(1-\beta)^2\lambda(e) - c \cdot e - \frac{1}{2}\alpha^2\gamma\sigma^2$ denote the payoff of the principal, at a contract (α, β) when the agent exerts effort e, and the noise of pre-screening is σ^2. Let $z(\sigma^2) := \{\max_{\alpha\geq 0, 0\leq\beta\leq 1, e\geq 0} v(\alpha, \beta, e, \sigma^2)$, s.t.(IC)$\}$ be the principal's payoff under the optimal contract as a function of the pre-screening noise. We then have $z(\sigma_1^2) \leq z(\sigma_2^2)$, $\forall\sigma_1^2 \geq \sigma_2^2$. In other words, better pre-screening improves the insurer's profit.*

The following theorem presents a sufficient condition under which the availability of a pre-screening assessment 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 the appendix, by setting $\sigma = \infty$, the insurer's optimal choice will be to set $\alpha = 0$, which effectively removes the effects of pre-screening.

Theorem 2.4. *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 $0 \leq \alpha_1 \leq \alpha_2 \leq c$, then $e_1 \geq e_2$ if $\sigma_1 \leq \sigma_2$. In other words, better pre-screening improves network security. In addition, if $k(e, 0)^2\lambda(e) - k(e, \alpha)^2\lambda(e)$ is non-decreasing in e for all $0 \leq \alpha \leq c$, then $e_1 \geq e_\infty$. In other words,* *availability of a pre-screening assessment improves network security over the no pre-screening scenario.*

Several instances of $\mu(e)$ and $\lambda(e)$, e.g., ($\mu(e) = \frac{1}{e}, \lambda(e) = \frac{1}{e^2}$), and ($\mu(e) = \exp\{-e\}, \lambda(e) = \exp\{-2e\}$), satisfy the condition of Theorem 2.4.

Our analysis of the single-agent case leads to the following conclusions. First, as expected, a market for insurance for a single agent (or *independent* agents) exists if and only if agents are risk-averse. We further observe that when the market exists, the introduction of pre-screening benefits the insurer (Theorem 2.3) as well the state of network security (Theorem 2.4). In the remainder of the paper, we are interested in investigating similar properties of the contract design problem in the presence of *interdependent* agents.

3 ROLE OF INTERDEPENDENCE IN A NETWORK OF TWO AGENTS

In order to design optimal contracts, and especially evaluating agents' participation constraints, an insurer needs to evaluate the agents' utilities and contracts in the following three scenarios:

(i) neither agent enters a contract;
(ii) one of the agents enters a contract, while the other one opts out; and
(iii) both agents purchase contracts.

Here, 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 (Case (iii)), we first need to find the optimal contracts and agents' payoffs in Cases (i) and (ii). In the remainder of this section, we shall first consider risk-neutral agents. The intention of this case is to exclude the effect of risk transfer and solely focus on the effect of interdependence. We will then examine the combined effect of risk transfer and risk interdependency with risk-averse agents.

3.1 Case (i): Neither enters a contract

Let G^{oo} denote the game between two risk-neutral agents neither of which have purchased cyber insurance contracts. The expected payoffs of a risk-neutral agent, with unit cost of effort $c_i > 0$ at the effort profile (e_1, e_2), is given by:

$$\overline{U}_i^{\text{out}}(e_1, e_2) = -\mu(e_i + xe_{-i}) - c_i e_i \ .$$

The best response of agent i is therefore given by,

$$B_i^{\text{out}}(e_{-i}) = \arg\max_{e_i \geq 0} \ -\mu(e_i + xe_{-i}) - c_i e_i \ .$$

The above optimization problem is convex, and has the following solution:

$$B_i^{\text{out}}(e_{-i}) = (m_i - xe_{-i})^+, \ m_i = \arg\min_{e \geq 0} \ \mu(e) + c_i e, \ i = 1, 2.$$

where $(a)^+ := \max\{a, 0\}$. The Nash equilibrium is given by the fixed point of the best-response mappings $B_1^{\text{out}}(e_2)$ and $B_2^{\text{out}}(e_1)$, i.e., the following set of equations:

$$e_1 = (m_1 - xe_2)^+, \ e_2 = (m_2 - xe_1)^+ \ . \tag{7}$$

It is straightforward to show that given $0 \leq x < 1$, the above equations have a unique fixed point, leading to the effort of agent i, $e_i^*(m_i, m_{-i})$, at the unique Nash equilibrium:

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

Therefore, $u_i^{oo} = \overline{U}_i^{\text{out}}(e_1^*(m_1, m_2), e_2^*(m_2, m_1))$ 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 u_i^{oo} to evaluate agents' outside options when proposing optimal contracts.

3.2 Case (ii): Only one enters a contract

Assume without loss of generality 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 the uninsured agent 2. The expected payoff of the risk-neutral agents in this case is as follows:

$$\overline{U}_1^{\text{in}}(e_1, e_2, p_1, \alpha_1, \beta_1) = -p_1 - (c_1 - \alpha_1)e_1 - (1 - \beta_1)\mu(e_1 + xe_2)$$
$$\overline{U}_2^{\text{out}}(e_1, e_2) = -\mu(e_2 + xe_1) - c_2 e_2$$

Let $B_1^{\text{in}}(e_2)$ denote the best response of agent 1. The following optimization problem finds the best response of the first agent,

$$B_1^{\text{in}}(e_2) = \arg\max_{e_1 \geq 0} -p_1 - (c_1 - \alpha_1)e_1 - (1 - \beta_1)\mu(e_1 + xe_2) \ .$$

The above optimization problem is convex, and has a solution given by,

$$B_1^{\text{in}}(e_2) = (m_1(\alpha_1, \beta_1) - xe_2)^+,$$
$$\text{where } m_1(\alpha_1, \beta_1) = \arg\min_{e \geq 0} \{(c_1 - \alpha_1)e + (1 - \beta_1)\mu(e)\} \ .$$

For the uninsured agent 2, it is easy to see that the best-response function is given by $B_2^{\text{out}}(e_1)$, the same best response function in game G^{oo}.

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 $e_2^*(m_2, m_1(\alpha_1, \beta_1))$, as defined in (8). For notational convenience, we denote these efforts by e_1^*, e_2^*.

Similar to Lemma 2.1, we can show that the (IR) constraint is binding under the optimal contract. Therefore, similar to the single agent case, we can write the insurer's problem by replacing the base premium p_1, leading to,

$$\begin{aligned} &\max_{\alpha_1 \geq 0, 0 \leq \beta_1 \leq 1, e_1^* \geq 0, e_2^* \geq 0} \ -u_1^{oo} - \mu(e_1^* + xe_2^*) - c_1 e_1^* \\ &\text{s.t., (IC) } e_1^*, e_2^* \text{ are the agents' efforts in the NE of } G^{io} \end{aligned} \tag{9}$$

Let u_2^{io} be the second agent's utility when the insurer offers the *optimal contract* (determined by solving problem (9)) to the first agent and the second agent opts out. Similarly, u_1^{oi} denotes the first agent's utility when he opts out and the second agent purchases the *optimal contract*. The insurer uses her knowledge of u_2^{io} and u_1^{oi} in designing a pair of contracts to attract both agents.

3.3 Case (iii): Both purchase contracts

Let G^{ii} denote the game between the two agents when they are both in a contract. Assume the insurer offers each agent i a contract (p_i, α_i, β_i). The expected utility of the agents when both purchase contracts is given by

$$\overline{U}_i^{\text{in}}(e_1, e_2, p_i, \alpha_i, \beta_i) = -p_i - (c_i - \alpha_i)e_i - (1 - \beta_i)\mu(e_i + xe_{-i}) \ .$$

Following steps similar to those in Section 3.2, the best-response function of agent i, denoted B_i^{in}, is given by

$$B_i^{\text{in}}(e_{-i}) = (m_i(\alpha_i, \beta_i) - xe_{-i})^+ ,$$

where $m_i(\alpha_i, \beta_i)$ is the solution to the following equation,

$$m_i(\alpha_i, \beta_i) = \arg\min_{e\geq 0} \{(c_i - \alpha_i)e + (1 - \beta_i)\mu(e)\} .$$

The agents' efforts at the Nash equilibrium are again the fixed point of the best-response mappings, and will be given by $e_i^*(m_i(\alpha_i, \beta_i), m_{-i}(\alpha_{-i}, \beta_{-i}))$, with $e_i^*(.,.)$ defined in (8). For notational convenience, we will denote these as e_i^*.

To write the insurer's problem, note that the outside option of agent 1 (resp. 2) from this game is his utility in the game G^{oi} (resp. G^{io}). Similar to Lemma 2.1, we can show that the (IR) constraints are binding. Therefore, we can write the insurer's contract design problem as follows,

$$\begin{aligned} v^{ii} := \max_{\alpha_1\geq 0, 0\leq\beta_1\leq 1, \alpha_2\geq 0, 0\leq\beta_2\leq 1, e_1^*\geq 0, e_2^*\geq 0} & -u_1^{oi} - u_2^{io} \\ & -\mu(e_1^* + xe_2^*) - c_1e_1^* - \mu(e_2^* + xe_1^*) - c_2e_2^* \qquad (10) \\ \text{s.t., (IC) } & e_1^*, e_2^* \text{ are the agents' efforts in the NE of } G^{ii} \end{aligned}$$

3.4 Optimal Contracts in a Network of Two Risk-Neutral Agents

We now analyze the properties of the contracts designed based on the optimization problem (10), and their impact on agents' efforts.

THEOREM 3.1. *Let e_i^o denote the effort of agent i when insurance is not available, and e_i^{in} denote the effort of agent i in the solution to* (10), *i.e., when purchasing the optimal contract. Also, let $\tilde{e}_i$ denote the effort level of agent i in the socially optimal outcome (i.e, the efforts maximizing the sum of agents' utilities). Then, the insurer offers contracts to both agents, with the following properties,*

(i) $e_i^{in} = \tilde{e}_i$, for $i = 1, 2$.
(ii) $e_1^{in} + e_2^{in} \geq e_1^o + e_2^o$.
(iii) $v^{ii} \geq \overline{U}_1(\tilde{e}_1, \tilde{e}_2) + \overline{U}_2(\tilde{e}_1, \tilde{e}_2) - \overline{U}_1(e_1^o, e_2^o) - \overline{U}_2(e_1^o, e_2^o)$.

Theorem 3.1, implies the following. First, in contrast to the single risk-neutral scenario, the insurer can make profit by offering a contract to interdependent risk-neutral agents. To see why, note that due to interdependency, agents under-invest in security at the no-insurance equilibrium. This leads to an opportunity for the insurer, in which she uses her pre-screening assessments to offer premium discounts and (full) coverage of losses, and in turn requires the agents to exert higher efforts (in this particular case, the socially optimal levels of effort). This increase in efforts is in the insurer's interest, as it lowers the risks of both of her contracts. This effect can be viewed as the insurer "selling commitment" to agents, by providing each agent with the commitment of the other agent to exert higher effort.

Part (iii) of the theorem shows that the profit opportunity for the insurer is even higher than the welfare gap between the socially optimal and Nash equilibrium outcomes. This is due to the fact that the outside option for agent i is an outcome in which the insurer offers a contract (only) to agent $-i$. The insurer will select this contract in a way that it requires agent $-i$ to exert low effort and get high coverage, effectively forcing agent i to bear the full cost of effort, leading to a utility lower than the no-insurance Nash equilibrium for agent i. Consequently, as agents' (IR) constraints are binding, it follows that the insurer's profit is in fact the gap between welfare attained under the optimal contract, and the welfare at these low payoff, unilateral opt-out outcomes.

3.5 Optimal Contracts in a Network of Two Risk-Averse Agents

We next analyze the case of two risk-averse agents. Again, in order to evaluate the agents' participation constraints and finding the optimal contracts, we need to account for three possible cases based on the agents' alternatives. The ensuing analysis is similar to that presented earlier in this section, by replacing the agent's utility functions with their risk-averse versions and solving the resulting optimization problems. The details are thus provided in the appendix. The simplified insurer's optimization problem is given by

$$\begin{aligned} v^{ii} := \max_{\alpha_1\geq 0, 0\leq\beta_1\leq 1, \alpha_2\geq 0, 0\leq\beta_2\leq 1, e_1^*\geq 0, e_2^*\geq 0} & \; w_1^{oi} + w_2^{io} \\ & -\mu(e_1^* + xe_2^*) - \tfrac{1}{2}\gamma_1(1-\beta_1)^2\lambda(e_1^* + xe_2^*) - c_1e_1^* - \tfrac{1}{2}\alpha_1^2\gamma_1\sigma_1^2 \\ & -\mu(e_2^* + xe_1^*) - \tfrac{1}{2}\gamma_2(1-\beta_2)^2\lambda(e_2^* + xe_1^*) - c_2e_2^* - \tfrac{1}{2}\alpha_2^2\gamma_2\sigma_2^2 \\ \text{s.t., (IC) } & e_1^*, e_2^* \text{ are the agents' efforts in the NE of } G^{ii} \end{aligned}$$

where $w_1^{oi} = \frac{\ln(-u_1^{oi})}{\gamma_1}$ and $w_2^{io} = \frac{\ln(-u_2^{io})}{\gamma_2}$.

We first consider the utility of the insurer. Note that the insurer always has the option to not use the outcome of pre-screening by setting $\alpha = 0$ in the contract. Therefore, the insurer's utility in the optimal contract with pre-screening is larger than her utility in the optimal contract without pre-screening.

We now turn to the effect of pre-screening on the state of network security. We again use the total effort towards security, $e_1 + e_2$, as the metric for evaluating network security.

THEOREM 3.2. *Let $m_i = \arg\min_{e\geq 0} \mu(e) + \frac{1}{2}\gamma_i\lambda(e) + c_ie$. Let e_i^{in} and e_i^o denote the effort of agent i in the optimal contract and in the no-insurance equilibrium, respectively.*

(i) Assume perfect pre-screening, i.e., $\sigma_1 = \sigma_2 = 0$. Then, $e_1^{in} + e_2^{in} \geq e_1^o + e_2^o$, if,

1. $\mu'(m_i) < \frac{-c_i + xc_{-i}}{1-x^2}, \quad i = 1, 2$
2. $(\mu')^{-1}(\frac{-c_i + xc_{-i}}{1-x^2}) \geq x(\mu')^{-1}(\frac{-c_{-i} + xc_i}{1-x^2}), \quad i = 1, 2$

(ii) Assume both pre-screening assessments are uninformative, i.e., $\sigma_1 = \sigma_2 = \infty$. Then $e_1^{in} + e_2^{in} \leq e_1^o + e_2^o$.

Recall that by Theorem 2.2, with a single risk-averse agent, the insurer profits from the agent's interest in risk transfer. However, the introduction of insurance always reduces network security. In contrast, Theorem 3.2 shows that in the case of interdependent agents, network security can improve, while the insurer continues to make profit. Therefore, it is the agents' interdependency that plays a role in the improvement of network security. Note that in this case, the insurer is making profit from the agents' risk aversion, as well as their interdependency.

4 DISCUSSION

The previous results can also be extended to a network of N homogeneous risk-averse agents, in which $\gamma_i = \gamma$, $c_i = c$, and $\sigma_i = \sigma$, $\forall i$, and agent i's loss is given by,

$$L_{\boldsymbol{e}}^{(i)} \sim \mathcal{N}(\mu(e_i + x\sum_{j\neq i} e_j), \lambda(e_i + x\sum_{j\neq i} e_j)) .$$

where $\boldsymbol{e} = (e_1, e_2, \cdots, e_N)$ denotes the vector of agents' efforts. The insurer again gathers assessments $S_i = e_i + W$, where W is a zero mean Gaussian noise with variance σ^2.

In the following two theorems, we show that under a sufficient condition, insurance improves network security as compared to no insurance scenario for sufficiently small pre-screening noise.

Theorem 4.1. *Assume N homogeneous agents purchase contracts from an insurer, and let $m = \arg\min_{e \geq 0} \mu(e) + \frac{1}{2}\gamma\lambda(e) + ce$. Let e^o be the effort of an agent in the no-insurance symmetric equilibrium, e' and $\hat{e}$ denote the effort in the optimal contract with perfect pre-screening, and without pre-screening, respectively. Then,*

(i) If pre-screening is accurate, i.e., $\sigma = 0$, and $m > 0$, then $e' \geq e^o$ if and only if $\mu'(m) < -\frac{c}{1+(N-1)x}$. *That is, network security improves after the introduction of insurance with prefect pre-screening.*

(ii) If pre-screening is uninformative, i.e., $\sigma = \infty$, then $e^o \geq \hat{e}$. That is, network security worsens after the introduction of insurance without pre-screening.

Note that this theorem, as well as its interpretation, are similar to the statements of Theorem 3.2 for two heterogeneous agents. In particular, it is straightforward to check that the conditions of part (i) of these theorems are equivalent when setting $c_i = c$ in Theorem 3.2 and $N = 2$ in Theorem 4.1.

The next theorem shows that with sufficiently accurate, yet imperfect prescreening, the use of pre-screening can lead to improvement of the state of network security compared to the no-insurance equilibrium.

Theorem 4.2. *Assume N homogeneous agents purchase contracts from an insurer. Let $m = \arg\min_{e \geq 0} \mu(e) + \frac{1}{2}\gamma\lambda(e) + ce$, and assume $\mu'(m) < -\frac{c}{1+(N-1)x}$. Let $\hat{e}$ and e^o be the effort level of agents in the optimal contract and at the no-insurance equilibrium, respectively. Let $\tilde{m}$ be the effort at which $\mu'(\tilde{m}) = -\frac{c}{1+(N-1)x}$. Then, if $\sigma \leq \frac{\mu(m) + \frac{c}{1+(N-1)x} m - \mu(\tilde{m}) - \frac{c}{1+(N-1)x} \tilde{m}}{0.5\gamma c^2}$, $\hat{e} \geq e^o$. That is, introducing pre-screening improves network security as compared to the no-insurance equilibrium.*

There are a number of other directions to pursue to extend our results. Firstly, all our results are derived under the assumption of perfect information. Studying the same contract design problem with pre-screening under partial information assumptions would be an important direction of future research; this would include imperfect knowledge of the agents' type by the principal as well as imperfect knowledge of the interdependence relationship by the agents and the principal. Other modeling choices such as alternative use of pre-screening assessment (as opposed to linear discounts on premiums), and capturing correlated risks (e.g., joint distribution of losses as opposed to average loss being a function of joint effort), would also be of great interest.

It should be noted that our key finding in the paper, that interdependent risk presents an opportunity to be embraced is counter to current practices, where underwriters typically avoid risk dependencies for fear of concurrent losses. This reflects a certain degree of risk aversion on the insurer's part, which our model of expected utility maximization does not capture. This presents another important direction of future work.

5 CONCLUSION

We studied the problem of designing cyber insurance contracts by a single profit-maximizing insurer, for both risk-neutral and risk-averse agents. While the introduction of insurance worsens network security in a network of independent agents, we showed that security interdependency leads to a profit opportunity for the insurer, created by the inefficient effort levels exerted by free-riding agents when insurance is not available but interdependency is present; this is in addition to risk transfer that an insurer typically profits from. We showed that security pre-screening then allows the insurer to take advantage of this additional profit opportunity by designing the right contracts to incentivize the agents to increase their effort levels.

REFERENCES

[1] *Online Appendix.* Available at https://www.dropbox.com/sh/ek4p20ornmcio56/AADjDafFU1CbbHtMB4tX7qDea?dl=0.

[2] Rainer Böhme. 2005. Cyber-insurance revisited. In *Proceedings of the Workshop on the Economics of Information Security (WEIS).*

[3] Rainer Böhme. 2012. Security audits revisited. In *International Conference on Financial Cryptography and Data Security.* Springer, 129–147.

[4] Jean Bolot and Marc Lelarge. 2009. Cyber insurance as an incentive for Internet security. In *Managing information risk and the economics of security.* Springer.

[5] Annette Hofmann. 2007. Internalizing externalities of loss prevention through insurance monopoly: an analysis of interdependent risks. *The Geneva Risk and Insurance Review* 32, 1 (2007), 91–111.

[6] Benjamin Johnson, Jens Grossklags, Nicolas Christin, and John Chuang. 2010. Are security experts useful? Bayesian Nash equilibria for network security games with limited information. In *European Symposium on Research in Computer Security.* Springer, 588–606.

[7] Benjamin Johnson, Jens Grossklags, Nicolas Christin, and John Chuang. 2010. Uncertainty in interdependent security games. In *International Conference on Decision and Game Theory for Security.* Springer, 234–244.

[8] Jay P. Kesan, Ruperto P. Majuca, and William Yurcik. 2005. Cyber-insurance as a market-based solution to the problem of cybersecurity-a case study. In *Proceedings of the Workshop on the Economics of Information Security (WEIS).*

[9] Mohammad Mahdi Khalili, Parinaz Naghizadeh, and Mingyan Liu. 2017. Designing cyber insurance policies: Mitigating moral hazard through security pre-screening. In *the 5th International Conference on Game Theory for Networks (GameNets).* IEEE.

[10] Marc Lelarge. 2012. Coordination in network security games: a monotone comparative statics approach. *IEEE Journal on Selected Areas in Communications* 30, 11 (2012), 2210–2219.

[11] Marc Lelarge and Jean Bolot. 2009. Economic incentives to increase security in the Internet: The case for insurance. In *Proceedings of IEEE INFOCOM.* 1494–1502.

[12] Yang Liu, Armin Sarabi, Jing Zhang, Parinaz Naghizadeh, Manish Karir, Michael Bailey, and Mingyan Liu. 2015. Cloudy with a Chance of Breach: Forecasting Cyber Security Incidents. In *Proceedings of the 24th USENIX Security Symposium.*

[13] Andreu Mas-Colell, Michael Dennis Whinston, and Jerry R. Green. 1995. *Microeconomic theory.* Oxford University press, New York.

[14] R. Ann Miura-Ko, Benjamin Yolken, Nicholas Bambos, and John Mitchell. 2008. Security investment games of interdependent organizations. In *Proceedings of 46th Annual Allerton Conference on Communication, Control, and Computing.* 252–260.

[15] Hulisi Ogut, Nirup Menon, and Srinivasan Raghunathan. 2005. Cyber insurance and IT security investment: Impact of interdependence risk. In *Proceedings of the Workshop on the Economics of Information Security (WEIS).*

[16] Ranjan Pal, Leana Golubchik, Konstantinos Psounis, and Pan Hui. 2014. Will cyber-insurance improve network security? A market analysis. In *Proceedings of IEEE INFOCOM.* 235–243.

[17] Galina A Schwartz and S Shankar Sastry. 2014. Cyber-insurance framework for large scale interdependent networks. In *Proceedings of the 3rd international conference on high confidence networked systems.* ACM, 145–154.

[18] Nikhil Shetty, Galina Schwartz, Mark Felegyhazi, and Jean Walrand. 2010. Competitive cyber-insurance and internet security. In *Economics of Information Security and Privacy.* Springer, 229–247.

[19] Nikhil Shetty, Galina Schwartz, and Jean Walrand. 2010. Can competitive insurers improve network security?. In *International Conference on Trust and Trustworthy Computing.* Springer, 308–322.

[20] Zichao Yang and John CS Lui. 2014. Security adoption and influence of cyber-insurance markets in heterogeneous networks. *Performance Evaluation* 74 (2014), 1–17.