Designing Cyber Insurance Policies: The Role of Pre-Screening and Security Interdependence

Mohammad Mahdi Khalili, Parinaz Naghizadeh, Mingyan Liu

***Abstract*—Cyber insurance is a viable method for cyber risk transfer. However, it has been shown that depending on the features of the underlying environment, it may or may not improve the state of network security. In this paper, we consider a single profit-maximizing insurer (principal) with voluntarily participating insureds/clients (agents). We are particularly interested in two distinct 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 not only on its own investment and effort, but also the efforts of others' in the same eco-system (i.e. externalities). The second is the fact that recent advances in Internet measurement combined with machine learning techniques now allow us to perform accurate quantitative assessments of security posture at a firm level. This can be used 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. We show that security interdependency leads to a "profit opportunity" for the insurer, created by the inefficient effort levels exerted by interdependent agents who do not account for the risk externalities when insurance is not available; this is in addition to risk transfer that an insurer typically profits from. Security pre-screening then allows the insurer to take advantage of this additional profit opportunity by designing the appropriate contracts which incentivize agents to increase their effort levels, allowing the insurer to "sell commitment" to interdependent agents, in addition to insuring their risks. We identify conditions under which this type of contracts lead to not only increased profit for the principal, but also an improved state of network security.**

***Index Terms*—Cybersecurity, Cyber Insurance, Pre-screening, Security Interdependence.**

I. Introduction

The market for cyber-insurance products has been growing steadily in recent years [3], [4], with over 70 carriers around the world and total premiums estimated over $3B and projected to reach $10B by 2020. These products enable organizations and businesses to manage their cyber-risks by transferring (part of) their risks to an insurer in return for paying premiums. This growing market has motivated an extensive literature (see e.g. [5]–[15]), which aims to understand the unique characteristics of these emerging contracts, their effect on the insureds' security expenditure, and the possibility of leveraging these contracts to shape users' behavior and improve the state of cybersecurity; see Section II for an overview of the related literature. The conclusions of these studies depend on the assumptions on the insurance market model (profit maker vs. welfare maximizing insurers), the agents' (insured's) participation decisions (compulsory vs. voluntary insurance), and the assumed model of interdependency among the insured.

In this paper, we are interested in analyzing the possibility of using cyber-insurance as an incentive for improving network security. We adopt two model assumptions which we believe better capture the current state of cyber insurance markets but differ from the majority of the existing literature; we shall assume a profit-maximizing cyber insurer, and voluntary participation, i.e., agents may opt out of purchasing a contract. Under this model, we focus on two features of cyber-insurance: (i) availability of risk assessmentfor mitigating moral hazard, and (ii) the interdependent nature of security.

The first feature is due to the fact that recent advances in Internet measurements combined with machine learning techniques now allow us to perform accurate, quantitative security posture assessments at a firm level [16]. This can be used as a tool to perform an initial security audit, or *pre-screening*, of a prospective client to mitigate moral hazard by premium discrimination and the design of customized policies. The second distinct feature, the interdependent nature of security, refers to the observation that the security standing of an entity often depends not only on its own effort towards implementing security metrics, but also on the efforts of other entities interacting with it within the eco-system; see e.g., [17]–[20]. Such interdependency is crucial for the insurer's contract design problem, as the insurer will need to offer coverage to each insured for both its losses due to direct breaches, as well as indirect losses caused by breaches of other entities.

To distinguish the effect of each feature on the cyber-insurance contract design problem, we begin by con-

Department of Electrical Engineering & Computer Science, University of Michigan, Ann Arbor. {khalili,naghizad,mingyan}@umich.edu.
Preliminary versions of this work appeared at GameNets17 [1] and NetEcon17 [2]. This work is supported by the NSF under grant CNS-1422211, CNS-1616575, CNS-1739517, and by the DHS via contract number HSHQPM17X00233.

sidering a single-agent; this allows us to remove the effects of risk interdependence and focus on the role of pre-screening. We consider both risk-neutral and risk-averse agents. We first show that when the agent is risk-neutral, a market for cyber-insurance does not exist, this is consistent with previous results, see e.g. [21], [22]. For the risk-averse agent on the other hand, a cyber-insurance market exists. We show that the agent's effort inside the contract increases as the quality of pre-screening increases, that is, the insurer can use pre-screening to mitigate moral hazard. Nevertheless, we show that even with perfect pre-screening, the agent's effort inside the contract remains below his effort before the introduction of insurance. In other words, for a single-agent, and even in the absence of moral hazard, the introduction of cyber-insurance deteriorates the state of network security.

We will next analyze the effect of risk interdependence by considering the design of cyber-insurance contracts for a network of two interdependent agents. We again consider both risk-neutral and risk-averse agents. Here, in contrast to the single agent case, we obtain a rather surprising result: an insurance market *exists* even for two risk-neutral agents. As there is no risk-transfer between the agents and the insurer in this scenario, we conclude that the emergence of a market is due to the agents' interdependence. We intuitively interpret this finding as follows. The interdependency among agents leads them to under-invest in security at the no-insurance equilibrium; this is commonly referred to as *free-riding*, see e.g., [20]. This inefficiency gap between the no-insurance equilibrium and the agents' utilities at more efficient investment levels creates a "profit opportunity" for the insurer. In particular, the insurer can use pre-screening to offer a pair of contracts that incentivize the agents to improve their levels of effort. In return for improving his effort level as prescribed by the contract, an insured is not only offered coverage in case of a loss, but further the "commitment" of the other agent to also improve its security, which will lead to further reduction in the insured's risks. Consequently, network security under these contracts is higher than the no-insurance equilibrium, which further benefits the insurer by lowering the risks of the insureds in its portfolio.

We will then consider the combined effect of risk transfer, interdependence, and security pre-screening, by considering a network of two interdependent risk-averse agents. Similar to the risk-neutral case, the interdependence leads to free-riding by agents in the absence of insurance. Consequently, the insurer can extract profit from both fronts: risk transfer, and taking advantage of the efficiency gap by incentivizing agents to exert higher effort. We identify a sufficient condition under which insurance leads to the improvement of network security compared to the no-insurance scenario. We illustrate these results in both a two-heterogeneous-agents model and an N-homogeneous-agents model. Lastly, we will discuss the effects of correlation in agents' losses, as well as a risk-averse insurer, on the cyber-insurance contracts, and illustrate our findings through numerical simulations.

Our main finding is that security interdependence among agents seeking cyber-insurance leads to a profit opportunity for the insurer. A cyber-insurer profits not only from risk-transfer, but also from *selling commitment* to interdependent agents: each agent will be required to improve its levels of investment in security, in return for the guarantee that other agents will do so as well. Security pre-screening allows the insurer to take advantage of this additional profit opportunity, by designing the appropriate contracts which incentivize agents to increase their effort levels. Together, these contracts can lead to an improvement in the state of network security.

Our analysis is primarily based on a two-agent model. While technically limited in scope, this simple model offers substantial conceptual insights, some of which are more generally applicable. We also use numerical examples to highlight where conclusions are expected to hold under more relaxed assumptions.

Preliminary versions of this work appeared in [1] and [2]. In [1], we studied the role of pre-screening in designing cyber insurance contracts between a risk-neutral insurer and a risk-averse agent, as well as two homogeneous interdependent risk-averse agents. In [2], we examined the problem of designing cyber insurance contracts using pre-screening between a risk-neutral insurer and agents in the following scenarios: (i) a single risk-neutral agent, (ii) two heterogeneous interdependent risk-neutral agents, (iii) two heterogeneous interdependent risk-averse agents, and (iv) N homogeneous interdependent risk-averse agents. In addition to a better exposition of our work by including numerical results and technical analysis in the appendix, we extend our previous work by considering a risk-averse insurer and N interdependent agents whose losses are correlated. In this scenario, we study the effect of correlated losses and insurer's risk-aversion on network security.

The remainder of the paper is organized as follows. We review related work in Section II. We present the single agent model in Section III, followed by the analysis in Section IV. We present the two-agent model and analysis in Section V. We discuss an N-homogeneous-agent case in Section VI, present numerical results in Section VII, and conclude in Section VIII.

II. Related work

We provide an overview of existing literature that is most closely related to this paper. These studies have considered either competitive or monopolistic insurers, as well as either mandatory or voluntary adoption by the insured. The works in [5]–[10] consider competitive insurance markets under compulsory insurance, and analyze the effect of insurance on agents' security expenditures. The authors of [5], [6] 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. [7], [8] study a network of heterogeneous agents and show that the introduction of insurance cannot improve the state of network security. Ogut *et al.* [9] study the impact of the degree of agents' interdependence, and show that agents' investments decreases as the degree of interdependence increases. Yang *et al.* [10] study a competitive market under the assumption of voluntary participation by agents, with and without moral hazard. In the absence of moral hazard, the insurer can observe agents' investments in security, and hence premium discriminates based on the observed investments. They show that such a market can provide incentives for agents to increase their investments in self protection. However, they show that under moral hazard, the market will not provide an incentive for improving agents' investments.

The impact of insurance on the state of network security in the presence of a monopolistic welfare maximizing insurer has been studied in [11]–[13], [23], [24]. In these models, as the insurer's goal is to maximize social welfare, assuming compulsory insurance, agents are incentivized through premium discrimination, i.e., agents with higher investments in security pay lower premiums. As a result, these studies show that insurance can lead to improvement of network security. An insurance market with a monopolistic profit maximizing insurer, under the assumption of voluntary participation, has been studied in [14], which shows that in the presence of moral hazard, insurance cannot improve network security as compared to the no-insurance scenario.

Our assumptions on the model, namely a profit-maximizing insurer and voluntary participation, are similar to [14]. Our work differs from [14], as well as other existing work, in that we illustrate (i) the role of pre-screening in mitigating moral hazard, and (ii) the possibility of designing contracts that leverage sufficiently accurate pre-screening and agents' interdependence to improve the state of network security.

III. Model and Preliminaries: Single Agent

We begin by considering the single-period contract design problem between a single risk-neutral insurer and a single agent[1]; we refer the interested reader to [22] for an overview of contract theory. The analysis of the single-agent case allows us to study solely the role of pre-screening by excluding the interdependency, and later, in conjunction with the analysis of Section V-B and V-C, to uncover the role of interdependency.

An 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[2], with mean $\mu(e) \geq 0$ and variance $\lambda(e) \geq 0$.[3] We assume both $\mu(e)$ 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. The convexity assumption suggests that while initial investment in security leads to considerable reduction in loss, the marginal benefit decreases as effort increases. In other words, it is not possible to reduce risk from cyber attacks to zero even if the agent exerts very large effort [25], [26]. We further preclude the possibility of misclaims by assuming that the realized loss is observed perfectly by both the insurer and the agent.

In general, the effort exerted by an agent is not observable by the insurer; this information asymmetry is formally referred to as moral hazard. We assume that in order to reduce this asymmetry and attain better information about the agent, 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* $S_e = e + W$, where W is a zero mean Gaussian noise with variance σ^2. We assume both agent and insurer know the distribution of S_e; such assessment can be obtained through a range of possible methods and (Internet) measurement techniques, information from initial surveys filled out by the agent, external audits, or internal audits conducted by a third party firm. We assume S_e is conditionally independent of L_e, given e. The pre-screening outcome S_e will be used by the insurer in determining the terms of the contract.

A. Linear Contract and the Insurer's Payoff

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 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:

$$V(p, \alpha, \beta, e) = p - \alpha \cdot S_e - \beta \cdot L_e \ . \tag{1}$$

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

[2]The normal assumption on L_e is to some extent justified by the fact that L_e is meant to capture the sum total of losses from a variety of sources, such as hacking, malware, insider threats, etc.

[3]For ease of exposition, we assume that $\lambda(e)$ is sufficiently small compared to $\mu(e)$, so that $Pr(L_e < 0)$ is negligible.

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

B. Risk-Neutral Agent

The utility of a risk-neutral agent is given by,

$$U(e) = -L_e - ce \Rightarrow \overline{U}(e) = E(U(e)) = -\mu(e) - ce \tag{2}$$

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 (m) of the agent outside the contract is $m = \arg\min_{e\geq 0} \quad \mu(e) + ce$ and his expected utility outside the contract is $u^o := \overline{U}(m)$.

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:

$$\begin{aligned} U^{in}(p, \alpha, \beta, e) &= -p + \alpha S_e - L_e + \beta L_e - ce \\ \overline{U}^{in}(p, \alpha, \beta, e) &= E(U^{in}(p, \alpha, \beta, e)) \\ &= -p + (\alpha - c)e + (\beta - 1)\mu(e) \end{aligned} \tag{3}$$

C. Risk-Averse Agent

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

$$U(e) = -\exp\{-\gamma \cdot (-L_e - ce)\} \ , \tag{4}$$

where γ denotes the *risk attitude* of the agent; a higher γ implies more risk aversion. We assume 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:

$$\begin{aligned} \overline{U}(e) &= E(-\exp\{-\gamma \cdot (-L_e - ce)\}) \\ &= -\exp\{\gamma \cdot \mu(e) + \frac{1}{2}\gamma^2\lambda(e) + \gamma ce\} \ . \end{aligned} \tag{5}$$

Using (5), 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 $u^o = \overline{U}(m)$ denote the maximum expected payoff of the agent without a contract.

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

$$\begin{aligned} &U^{in}(p, \alpha, \beta, e) = \\ &-\exp\{-\gamma \cdot (-p + \alpha \cdot S_e - L_e + \beta \cdot L_e - ce)\} \ . \end{aligned} \tag{6}$$

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

$$\begin{aligned} \overline{U}^{in}(p, \alpha, \beta, e) = -\exp\{\gamma(p + (c - \alpha)e + \frac{1}{2}\alpha^2\gamma\sigma^2 + \\ (1 - \beta)\mu(e) + \frac{1}{2}\gamma(1 - \beta)^2\lambda(e))\} \ . \end{aligned} \tag{7}$$

D. 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,

$$\begin{aligned} \max_{p, \alpha\geq 0, 0\leq\beta\leq 1, e\geq 0} \quad & \overline{V}(p, \alpha, \beta, e) = p - \alpha \cdot e - \beta \cdot \mu(e) \\ \text{s.t.} \quad (\text{IR}) \quad & \overline{U}^{in}(p, \alpha, \beta, e) \geq u^o \\ (\text{IC}) \quad & e \in \arg\max\nolimits_{e'\geq 0} \quad \overline{U}^{in}(p, \alpha, \beta, e') \end{aligned} \tag{8}$$

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

$$e \in \arg\min_{e'\geq 0} \quad (c - \alpha)e' + (1 - \beta)\mu(e') \ . \tag{9}$$

Similarly, the (IC) constraint for a risk-averse agent can be rewritten as:

$$e \in \arg\min_{e'\geq 0}(c - \alpha)e' + (1 - \beta)\mu(e') + \frac{\gamma}{2}(1 - \beta)^2\lambda(e') \tag{10}$$

Next, we can simplify the (IR) constraint using the following lemma; proofs can be found in the online appendix [27].

Lemma 3.1: The (IR) constraint is binding in the optimal contract.

By lemma 3.1, the (IR) constraint of a risk-neutral agent can be written as $-p-(c-\alpha)\cdot e-(1-\beta)\mu(e) = u^o$ and, for a risk-averse agent,

$$p+(c-\alpha)e+\frac{\gamma}{2}\alpha^2\sigma^2+(1-\beta)\mu(e)+\frac{\gamma}{2}(1-\beta)^2\lambda(e) = w^o, \tag{11}$$

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

Using the above expressions to substitute for the base premium p in the objective function in (8), and using the simplified expressions for the (IC) constraints, we re-write the insurer's contract design problem as follows.

Insurer's problem with a risk-neutral agent:

$$\begin{aligned} &\max\nolimits_{\alpha\geq 0, 0\leq\beta\leq 1, e\geq 0} \quad -u^o - \mu(e) - c \cdot e \\ &\text{s.t.,} \ e = \arg\min\nolimits_{e'\geq 0} \ (c - \alpha)e' + (1 - \beta)\mu(e') \end{aligned} \tag{12}$$

Insurer's problem with a risk-averse agent:

$$\begin{aligned} &\max\nolimits_{\alpha\geq 0, 0\leq\beta\leq 1, e\geq 0} \\ &\quad w^o - \mu(e) - \tfrac{\gamma}{2}(1 - \beta)^2\lambda(e) - ce - \tfrac{\gamma}{2}\alpha^2\sigma^2 \\ &\text{s.t.,} \quad e = \arg\min\nolimits_{e'\geq 0} \\ &\quad (c - \alpha)e' + (1 - \beta)\mu(e') + \tfrac{\gamma}{2}(1 - \beta)^2\lambda(e') \end{aligned} \tag{13}$$

IV. Role of Pre-Screening for a Single Agent

We now solve the optimal contract problem posed in (12) and (13), respectively.

A. Risk-Neutral Agent (Problem (12))

In this case, the objective function of the insurer is given by $-u^o - \mu(e) - c \cdot e$. 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.

B. Risk-Averse Agent (Problem (13))

We start with the following theorem on the state of network security, defined as the effort exerted by the agent, before and after the purchase of a contract.

Theorem 4.1: Assume that $(\hat{\alpha}, \hat{\beta}, \hat{e})$ solves optimization problem (13). Then $\hat{e} \leq m$, where m is the level of effort outside the contract; in other words, insurance decreases network security.

Proof. Assume that $(\hat{\alpha}, \hat{\beta}, \hat{e})$ solves optimization problem (13), and that, by contradiction, $\hat{e} > m \geq 0$.

First, recall that the agent's optimal effort m outside the contract is given by $m := \arg\min_{e\geq 0}\left\{\mu(e) + \frac{1}{2}\gamma\lambda(e) + ce\right\}$. For m to be the minimizer, we should have $c+\mu'(m)+\frac{1}{2}\gamma\lambda'(m) \geq 0$. Next, consider the following two cases:

(i) $\hat{\alpha} = 0$. Starting from the first order condition (FOC) on the (IC) constraint, we have,

$$\begin{aligned} &(1-\hat{\beta})\mu'(\hat{e}) + \tfrac{1}{2}\gamma(1-\hat{\beta})^2\lambda'(\hat{e}) + c = 0 \\ &\Rightarrow \mu'(\hat{e}) + \tfrac{1}{2}\gamma\lambda'(\hat{e}) + c < 0 \\ &\Rightarrow \mu'(m) + \tfrac{1}{2}\gamma\lambda'(m) + c < 0 \end{aligned} \tag{14}$$

Here, the second line follows from the decreasing nature of $\mu(\cdot)$ and $\lambda(\cdot)$, and the third line follows from their convexity. The last inequality is impossible given the optimality of the effort m outside the contract. This contradiction shows that we cannot have $\hat{e} > m$.

(ii) $\hat{\alpha} > 0$. Given the assumption that $\hat{e} > m$, and $\mu(\cdot)$ and $\lambda(\cdot)$ are strictly convex, we have,

$$\begin{aligned} 0 \quad &\leq c + \mu'(m) + \tfrac{1}{2}\gamma\lambda'(m) \\ &\leq c + \mu'(m) + \tfrac{1}{2}\gamma(1-\hat{\beta})^2\lambda'(m) \\ &< c + \mu'(\hat{e}) + \tfrac{1}{2}\gamma(1-\hat{\beta})^2\lambda'(\hat{e}) \end{aligned} \tag{15}$$

Therefore, if the insurer decreases $\hat{\alpha}$, the agent decreases his effort (this can be seen from the IC constraint), and consequently the insurer's utility increases, as the objective function of the insurer, $w^o - \mu(e) - \frac{1}{2}(1-\hat{\beta})^2\lambda(e) - ce - \frac{1}{2}\gamma\alpha^2\sigma^2$, is decreasing in e and α at $e = \hat{e}, \alpha = \hat{\alpha}$. Therefore, $(\hat{\alpha}, \hat{\beta}, \hat{e})$ is not the optimal contract. Again by contradiction, we conclude that the agent's effort in the optimal contract should be less than or equal to m. ■

Theorem 4.1 illustrates the inefficiency of cyber insurance as a tool for improving the state of security. Existing work in [6], [24] have also arrived at a similar conclusion when studying competitive/unregulated cyber insurance markets. Note also that Theorem 4.1 holds regardless of the pre-screening quality. We next examine the role of pre-screening in this model. We first analyze its impact on the insurer's profit.

Theorem 4.2: Let $v(\alpha, \beta, e, \sigma^2)$ 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 $z(\sigma^2) := \{\max_{\alpha\geq 0, 0\leq\beta\leq 1, e\geq 0} \; v(\alpha, \beta, e, \sigma^2), \text{ s.t. (IC)}\}$ be the principal's payoff under the optimal contract as a function of the pre-screening noise. We then have $z(\sigma_1^2) \leq z(\sigma_2^2), \forall \sigma_1^2 \geq \sigma_2^2$. That is, $z(\sigma^2)$ is a decreasing function of the pre-screening noise.

Proof. Let $v(\alpha, \beta, e, \sigma^2)$ be 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, and let $z(\sigma^2)$ be the insurer's profit at the optimal contract as a function of the pre-screening noise. We have,

$$\begin{aligned} z(\sigma_1^2+\sigma_2^2) &= \max_{\alpha, 0\leq\beta\leq 1, e\geq 0, \text{IC}} v(\alpha, \beta, e, \sigma_1^2+\sigma_2^2) \\ &\leq \max_{\alpha, 0\leq\beta\leq 1, e\geq 0, \text{IC}} v(\alpha, \beta, e, \sigma_1^2) + \\ &\quad \max_{\alpha, 0\leq\beta\leq 1, e\geq 0, \text{IC}}\{-\tfrac{1}{2}\alpha^2\gamma\sigma_2^2\} \leq \\ &\quad \max_{\alpha, 0\leq\beta\leq 1, e\geq 0, \text{IC}} v(\alpha, \beta, e, , \sigma_1^2) = z(\sigma_1^2) \end{aligned} \tag{16}$$

Therefore, $z(\sigma_1^2+\sigma_2^2) \leq z(\sigma_1^2), \;\; \forall \sigma_2^2$. That is, $z(\sigma^2)$ is a decreasing function of the pre-screening noise. ■

The above result is intuitively to be expected, as a strategic insurer can leverage improved pre-screening to better mitigate moral hazard and attain a higher 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 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 online appendix [27], by setting $\sigma = \infty$, the insurer's optimal choice will be $\alpha = 0$, which removes the effects of pre-screening.

Theorem 4.3: 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 $0 \leq \alpha_1 \leq \alpha_2 \leq c$, 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 $e_1 \geq e_\infty$. That is, the availability of a pre-screening improves network security over the no pre-screening scenario.

Sketch of Proof. The proof proceeds in the following steps:

- We first show that $0 \leq \alpha_i \leq c$ using the KKT conditions for the (IC) constraint of (13), given by

$$\begin{aligned} &(1-\beta_i)\mu'(e_i) + \tfrac{1}{2}\gamma(1-\beta_i)^2\lambda'(e_i) \\ &\qquad\qquad + c - \alpha_i - v_i = 0 \qquad (17) \\ &v_i \cdot e_i = 0, \; e_i \geq 0 \end{aligned}$$

- We next show that $\alpha_1 \geq \alpha_2$; this follows from the inequalities determining the optimality of the contracts at their respective pre-screening noises. In other words, as pre-screening noise decreases, the insurer offers higher discount factor.
- We then proceed by contradiction, assuming $0 \leq e_1 < e_2$. As $e_2 > 0$, by (17) we have,

$$\begin{aligned} &(1-\beta_2)\mu'(e_2) + \tfrac{\gamma(1-\beta_2)^2\lambda'(e_2)}{2} + c - \alpha_2 = 0 \\ &1-\beta_2 = \tfrac{\mu'(e_2)+\sqrt{\mu'(e_2)^2 - 2\gamma(c-\alpha_2)\lambda'(e_2)}}{-\gamma\lambda'(e_2)} := k(e_2,\alpha_2) \qquad (18) \end{aligned}$$

In addition, as $e_1 < e_2$ and $\alpha_1 \geq \alpha_2$, we can show that $\alpha_1 > 0$ and $e_1 > 0$. With $e_1 > 0$, by (17) we have,

$$\begin{aligned} &(1-\beta_1)\mu'(e_1) + \tfrac{\gamma(1-\beta_1)^2\lambda'(e_1)}{2} + c - \alpha_1 = 0 \\ &1-\beta_1 = \tfrac{\mu'(e_1)+\sqrt{\mu'(e_1)^2 - 2\gamma(c-\alpha_1)\lambda'(e_1)}}{-\gamma\lambda'(e_1)} := k(e_1,\alpha_1) \qquad (19) \end{aligned}$$

- Lastly, we show that if $(k(e,\alpha_2)^2 - k(e,\alpha_1)^2)\lambda(e)$ is non-decreasing, then α_1 and e_1 are not the maximizer of the insurer's profit when $\sigma^2 = \sigma_1^2$. This is a contradiction. Therefore, we conclude that $e_1 \geq e_2$. ■

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 4.3.

C. Comparison

By comparing the contracts in the risk-neutral and risk-averse agent cases, we observe that a market exists and the insurer makes profit only when offering a contract to a risk-averse agent. This is indeed to be expected, as insurance is primarily a method for risk transfer; risk-averse agents are willing to pay premiums that are higher than their expected loss, in order to reduce the uncertainty in their loss, consequently allowing the risk-neutral insurer to make a profit. We further observe that when the market exists, the introduction of pre-screening benefits the insurer (Theorem 4.2) as well the state of network security (Theorem 4.3).

V. Model and Analysis for Two Agents

We next study the contract design problem between the insurer and two agents. In particular, we analyze the impact of interdependency and pre-screening on the optimal contract and agents' effort, in the case of two risk neutral and two risk averse agents, respectively, with the former allowing us to exclude the effect of risk aversion and focus on the effect of interdependence.

A. A model of two agents

The two agents are interdependent, in that the effort exerted by one agent affects not only himself, but also the loss that the other agent experiences. We model the interdependence between these two agents as follows:

$$L^{(i)}_{e_1,e_2} \sim \mathcal{N}(\mu(e_i + x \cdot e_{-i}), \lambda(e_i + x \cdot e_{-i})) \; . \qquad (20)$$

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 $x \in [0,1)$. Note that this is not a unique modeling choice and is indeed a simplification; a more general way of expressing correlated risks would be to model the losses as jointly distributed; more on extensions is discussed in Section VIII.

We assume the agents' utilities are again given by (2) and (4) for risk-neutral and risk-averse agents, respectively, with the loss distributions replaced by the above expression. We allow the two agents to have different effort cost c_1, c_2, as well as different risk attitudes γ_1, γ_2.

The insurer can again conduct a pre-screening assessment, $S_{e_i} = e_i + W_i$, 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 [4], 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.

Similar to the single agent case, we need to evaluate the agents' outside options from purchasing a contract. These will then be used to impose the individual rationality constraints in determining the terms of the contracts. However, compared to the single agent case, the outside option of one agent is now influenced by the participation choice of the other agent as well. Specifically, we need to evaluate the agents' utilities as well as potential contracts in the following three scenarios:

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

Here, Case (ii) is the outside option for agents in Case (iii), and Case (i) is the outside option for agents in

[4] An example and discussion on correlated pre-screening noises can be found in the online appendix [27].

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). We therefore evaluate the agents' utilities for each case, and subsequently solve the insurer's contract design problem, in Sections V-B and V-C for risk-neutral and risk-averse agents, respectively.

B. Two Risk-Neutral Agents

Our first two-agent model is for risk-neutral agents to solely focus on the effect of interdependence. As mentioned above, in order to evaluate the agents' opt-out options and finding the optimal contract, the insurer's problem and the agents' utilities need to be studied under three different cases. We begin by analyzing these three cases, and then proceed to discussing the role of pre-screening and the contracts' effect on network security.

1) Case (i): Neither agent enters a contract: Let G^{oo} denote the game between two risk-neutral agents which have purchased cyber insurance contracts. In this game, Agents' efforts e_1, e_2 are their actions, and the expected payoffs of risk-neutral agents, with unit cost of effort $c_1, c_2 > 0$, are given by:

$$\overline{U}_i(e_1, e_2) = -\mu(e_i + xe_{-i}) - c_i e_i \ . \qquad (21)$$

The best response of each agent is therefore given by

$$B_i^{out}(e_{-i}) = \arg\max_{e_i \geq 0} \ -\mu(e_i + xe_{-i}) - c_i e_i \ . \qquad (22)$$

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

$$\begin{aligned} &m_i = \arg\min\nolimits_{e\geq 0} \ \mu(e) + c_i e, \qquad i = 1, 2 \ , \\ &B_i^{out}(e_{-i}) = (m_i - xe_{-i})^+ \ , \end{aligned} \qquad (23)$$

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

$$e_1 = (m_1 - xe_2)^+, \text{and } e_2 = (m_2 - xe_1)^+ \qquad (24)$$

To find a fixed point, we consider three cases,

• $e_1 = 0, e_2 \geq 0$: In this case, $e_2 = m_2$. Also, this case is valid if $m_1 - xm_2 \leq 0$ otherwise e_1 should be nonzero.

• $e_2 = 0, e_1 \geq 0$: similar to previous case, $e_1 = m_1$. This case is valid if $m_2 - xm_1 \leq 0$ otherwise e_2 should be nonzero.

• $e_1 > 0, e_2 > 0$: In this case, we solve the following system of equations:

$$e_1 = m_1 - xe_2, \text{ and } e_2 = m_2 - xe_1 \qquad (25)$$

The solutions of above equations is given by,

$$\begin{aligned} e_1 &= \tfrac{m_1 - x\cdot m_2}{1-x^2} \\ e_2 &= \tfrac{m_2 - x\cdot m_1}{1-x^2} \end{aligned} \qquad (26)$$

Notice that this case is valid if $\frac{m_1 - x\cdot m_2}{1-x^2} > 0$ and $\frac{m_2 - x\cdot m_1}{1-x^2} > 0$. Therefore, given $0 \leq x < 1$, system of equations (24) has a unique fixed point, and agent i's effort, $e_i^*(m_i, m_{-i})$, at the unique Nash equilibrium:

$$e_i^*(m_i, m_{-i}) = \begin{cases} \frac{m_i - x\cdot m_{-i}}{1-x^2} & \text{if } m_i \geq x\cdot m_{-i} \text{ and} \\ & \quad m_{-i} \geq x \cdot m_i \\ 0 & \text{if } m_i \leq x \cdot m_{-i} \\ m_i & \text{if } m_{-i} \leq x \cdot m_i \end{cases} \qquad (27)$$

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

2) Case (ii): One and 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 uninsured agent 2. The agents' expected payoff in this case is:

$$\begin{aligned} &\overline{U}_1^{in}(e_1, e_2, p_1, \alpha_1, \beta_1) = \\ &-p_1 - (c_1 - \alpha_1)e_1 - (1-\beta_1)\mu(e_1 + xe_2) \\ &\overline{U}_2(e_1, e_2) = -\mu(e_2 + xe_1) - c_2 e_2 \end{aligned} \qquad (28)$$

Let $B_1^{in}(e_2)$ denote the best response of agent 1. The following optimization problem finds its best response:

$$\begin{aligned} &B_1^{in}(e_2) = \arg\max\nolimits_{e_1\geq 0} \overline{U}_1^{in}(e_1, e_2, p_1, \alpha_1, \beta_1) = \\ &\arg\max\nolimits_{e_1 \geq 0} \\ &\qquad -p_1 - (c_1 - \alpha_1)e_1 - (1-\beta_1)\mu(e_1 + xe_2) \ . \end{aligned} \qquad (29)$$

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

$$\begin{aligned} &m_1(\alpha_1, \beta_1) = \arg\min\nolimits_{e\geq 0}\{(c_1 - \alpha_1)e + (1-\beta_1)\mu(e)\} \\ &B_1^{in}(e_2) = (m_1(\alpha_1, \beta_1) - xe_2)^+ \end{aligned} \qquad (30)$$

For the uninsured agent 2, it is easy to see that the best-response function is given by $B_2^{out}(e_1)$, 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 (27). For notational convenience, we denote these efforts by e_1^*, e_2^*.

Let $\overline{V}^{io}(p_1, \alpha_1, \beta_1, e_1, e_2)$ denote the insurer's utility, when agent 2 opts out and the insurer 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 to the participating agent is the solution to,

$$\begin{aligned}&\max_{p_1,\alpha_1,0\le\beta_1\le 1,e_1 e_2^*} \overline{V}^{io}(p_1,\alpha_1,\beta_1,e_1^*,e_2^*) = \\ &\quad p_1-\alpha_1 e_1^*-\beta_1\cdot\mu(e_1^*+x\cdot e_2^*)\\ &s.t.,\quad \text{(IR) } \overline{U}_1^{in}(e_1^*,e_2^*,p_1,\alpha_1,\beta_1)\ge u_1^{oo},\\ &\text{(IC) } e_1^*,e_2^* \text{ are the agents' efforts in NE of } G^{io}\end{aligned} \tag{31}$$

Similar to Lemma 3.1, we can show that the (IR) constraint is binding under the optimal contract. Therefore, we can re-write the insurer's problem by replacing the base premium p_1, leading to,

$$\begin{aligned}&\max_{\alpha_1,0\le\beta_1\le 1,e_1^*,e_2^*} -u_1^{oo}-\mu(e_1^*+xe_2^*)-c_1e_1^*\\ &s.t., \text{(IC) } e_1^*,e_2^* \text{ are the agents' efforts in NE of } G^{io}\end{aligned} \tag{32}$$

Let u_2^{io} be the second agent's utility when the insurer offers the *optimal contract* to the first agent and the second agent opts out. The insurer can calculate u_2^{io} by finding the optimal contract in problem (32) and the resulting Nash equilibrium of game G^{io}. 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) Case (iii): Both agents 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

$$\begin{aligned}&\overline{U}_i^{in}(e_1,e_2,p_i,\alpha_i,\beta_i) = \\ &\quad -p_i-(c_i-\alpha_i)e_i-(1-\beta_i)\mu(e_i+x\cdot e_{-i}).\end{aligned} \tag{33}$$

Following steps similar to those in Section V-B2, B_i^{in}, the best-response function of agent i, is given by

$$B_i^{in}(e_{-i}) = (m_i(\alpha_i,\beta_i)-xe_{-i})^+ \ , \tag{34}$$

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

$$m_i(\alpha_i,\beta_i) = \arg\min_{e\ge 0}\{(c_i-\alpha_i)e+(1-\beta_i)\mu(e)\} \ . \tag{35}$$

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 (27). 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}). Then, the optimal contracts offered by the insurer to the agents is the solution to the following optimization problem:

$$\begin{aligned}&\max_{p_1,\alpha_1,0\le\beta_1\le 1,p_2,\alpha_2,0\le\beta_2\le 1,e_1^*,e_2^*}\\ &p_1-\alpha_1e_1^*-\beta_1\cdot\mu(e_1^*+x\cdot e_2^*)+\\ &\qquad p_2-\alpha_2e_2^*-\beta_2\cdot\mu(e_2^*+x\cdot e_1^*)\\ &\text{s.t., (IR) } \overline{U}_j^{in}(e_1^*,e_2^*,p_j,\alpha_j,\beta_j)\ge u_j^{oi},\ j=1,2\\ &\text{(IC) } e_1^*,e_2^* \text{ are the agents' efforts in NE of } G^{ii}\end{aligned} \tag{36}$$

The (IR) constraints can again be shown to be binding. Therefore, the insurer's contract design problem for two risk-neutral agents is given by,

$$\begin{aligned}&v^{ii} := \max_{\alpha_1,0\le\beta_1\le 1,\alpha_2,0\le\beta_2\le 1,e_1^*,e_2^*} -u_1^{oi}-u_2^{io}\\ &-\mu(e_1^*+x\cdot e_2^*)-c_1\cdot e_1^*-\mu(e_2^*+x\cdot e_1^*)-c_2\cdot e_2^*\\ &\text{s.t., } e_1^*,e_2^* \text{ are the agents' efforts in NE of } G^{ii}\end{aligned} \tag{37}$$

4) Optimal Contracts for Two Risk-Neutral Agents: We now analyze the properties of the contracts designed based on the optimization problem (37), and their impact on agents' efforts.

Theorem 5.1: 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 (37), 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 $i=1,2$. That is, the agents exert socially optimal effort levels in the optimal contract.

(ii) $e_1^{in}+e_2^{in}\ge e_1^o+e_2^o$. That is, when both agents purchase optimal insurance contracts, the overall effort exerted toward security increases compared to the no-insurance scenario.

(iii) $v^{ii}\ge\overline{U}_1(\tilde{e}_1,\tilde{e}_2)+\overline{U}_2(\tilde{e}_1,\tilde{e}_2)-\overline{U}_1(e_1^o,e_2^o)-\overline{U}_2(e_1^o,e_2^o)$. That is, the principal's profit is higher than the gap between agents' welfare at the socially optimal solution and the no-insurance equilibrium.

Theorem 5.1, implies the following. Firstly, recall that, as discussed in Section IV-C, the insurer cannot make profit from offering contracts to a single risk-neutral agent, as there is no risk transfer from risk-neutral agents to an insurer. However, we observe that the insurer can make profit when offering contracts to interdependent risk-neutral agents. We conclude that this improvement is due to the agents' interdependency, and can be interpreted as follows. Due to interdependency, agents under invest in security at the no-insurance equilibrium. This leads to a profit opportunity for the insurer, in which she uses her (accurate) 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 its contracts. In addition, this effect can be viewed as the insurer "selling commitment" to agents. That is, the insurer is also providing each agent with the commitment of the other agent to exert higher effort, if he also commits to exerting high effort.

Secondly, 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 from the contract 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 also 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.

Finally, note that the statements of this theorem do not depend on the pre-screening noises $\sigma_i < \infty$. This is because the expected utilities and consequent effort choices of risk-neutral agents are only sensitive to the mean, but not the variances of uncertainties in the problem parameters. As such, under the assumption of zero mean noise in the pre-screening assessments, agents' behavior will be independent of σ.

C. Two Risk-Averse Agents

We next analyze the case of two risk-averse agents. Again, as discussed in Section V, in order to evaluate the agents' individual rationality constraints and finding the optimal contracts, we need to account for three possible cases based on the agents' participation alternatives.

The ensuing analysis is similar to that presented in Section V-B, by replacing the agent's utility functions with their risk-averse versions and solving the resulting optimization problems. We thus present the details in the online appendix [27]. Following the analysis, the simplified insurer's optimization problem is given by

$$\begin{aligned} v^{ii} = \max_{\alpha_1, 0\leq\beta_1\leq 1, \alpha_2, 0\leq\beta_2\leq 1, e_1^*\geq 0, e_2^*\geq 0} & \; w_1^{oi} + w_2^{io} \\ -\mu(e_1^* + x\cdot e_2^*) - \tfrac{1}{2}\gamma_1(1-\beta_1)^2\lambda(e_1^* + x\cdot e_2^*) & \\ -c_1\cdot e_1^* - \tfrac{1}{2}\alpha_1^2\gamma_1\sigma_1^2 & \\ -\mu(e_2^* + x\cdot e_1^*) - \tfrac{1}{2}\gamma_2(1-\beta_2)^2\lambda(e_2^* + x\cdot e_1^*) & \\ -c_2\cdot e_2^* - \tfrac{1}{2}\alpha_2^2\gamma_2\sigma_2^2 & \\ s.t., e_1^*, e_2^* \text{ are the agents' efforts in NE of game } G^{ii} & \end{aligned} \tag{38}$$

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

We now discuss how different problem parameters, particularly the availability of pre-screening, affect the insurer's profit in the optimal contracts, as well as the system's state of security. 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 that in the optimal contract without pre-screening; i.e., the availability of pre-screening is in the insurer's interest.

We now turn to the effect of pre-screening on the state of network security, which we shall measure by the total effort toward security, $e_1 + e_2$.

Theorem 5.2:

Let $m_i = \arg\min_{e\geq 0} \mu(e) + \frac{1}{2}\gamma_i\lambda(e) + c_i e$. Let e_i 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, $e_1 + e_2 \geq e_1^o + e_2^o$, if,

$$\begin{aligned} &1.\ \mu'(m_i) < \tfrac{-c_i + xc_{-i}}{1-x^2}, \quad i = 1,2 \\ &2.\ (\mu')^{-1}\left(\tfrac{-c_i + xc_{-i}}{1-x^2}\right) \geq x(\mu')^{-1}\left(\tfrac{-c_{-i} + xc_i}{1-x^2}\right), \quad i = 1,2 \end{aligned} \tag{39}$$

That is, under these conditions, insurance improves network security compared to the no-insurance scenario.

(ii) Assume both pre-screening assessments are uninformative. i.e., $\sigma_1 = \sigma_2 = \infty$. Then $e_1 + e_2 \leq e_1^o + e_2^o$. That is, the insurance contract without pre-screening worsens network security as compared to the no-insurance scenario.

The results of Theorem 5.2 can be intuitively interpreted as follows. By Theorem 4.1, 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 5.2 shows that with interdependent agents network security can improve, while the insurer continues to make profit. Therefore, it is agents' interdependency that plays a role in the improvement of security. To see why, note that the insurer uses pre-screening and offers premium discounts accordingly in order to incentivize the interdependent agents to increase their effort levels. Providing such incentives is in the insurer's interest, as higher effort exerted by the agent decreases both agents' risk, and consequently, the coverage required by the insurer once losses are realized. Note also that it is the availability of (accurate) pre-screening that provides the required tools for the insurer in designing such incentives; otherwise, as shown in part (ii) of the theorem, improving network security is no longer possible.

The conditions of part of (i) of the theorem can also be interpreted as follows. The first condition imposes a restriction on the derivative of μ, so that the decrease in loss as a function of effort is faster than the normalized cost of effort; as a result, the insurer will have the option to make more profit through loss reduction (by encouraging agents to exert higher effort). The second condition imposes a restriction on the agents' cost of effort and guarantees that both agents exert positive effort (see proof of Theorem 5.2). Specifically, when the two agents' effort costs are sufficiently similar, this condition is satisfied, and both agents exert non-zero effort.

VI. N Homogeneous Agents, Correlated Losses, and Risk-Averse Insurer

In this section we show a number of extensions of our results. First, in SectionVI-A we study the optimal contracts in a network of N homogeneous risk-averse agents. In Section VI-B, we examine the case where the losses of these agents are not only distributionally dependent but also correlated in their realizations; we will also consider the impact of risk aversion on the part of the insurer on the resulting contract.

A. N-homogeneous risk-averse agents

Consider a network of N homogeneous risk-averse agents given by $\gamma_i = \gamma$, $c_i = c$, and $\sigma_i = \sigma$, $\forall i$. The assumption of homogeneity simplifies the insurer's problem, allowing us to obtain additional insights about the contracts and their impact on network security. Let $\boldsymbol{e} = (e_1, e_2, \cdots, e_N)$ denote the vector of efforts of all agents. The loss of agent i is given by,

$$L_{\boldsymbol{e}}^{(i)} \sim \mathcal{N}(\mu(e_i + x\sum_{j\neq i} e_j), \lambda(e_i + x\sum_{j\neq i} e_j)) \ . \quad (40)$$

The agents' expected utility outside the contract is,

$$\overline{U}_i(\boldsymbol{e}) = E(-\exp\{-\gamma(-L_{\boldsymbol{e}}^{(i)} - ce_i)\}) = -\exp\{\gamma(\mu(e_i + x\textstyle\sum_{j\neq i} e_j) + \frac{\gamma\lambda(e_i + x\sum_{j\neq i} e_j)}{2} + ce_i)\} \quad (41)$$

Let $m = \arg\min_{e\geq 0} \mu(e) + \frac{1}{2}\gamma\lambda(e) + ce$. Then, the best response mapping of agent i is given by,

$$B_i^{out}(\boldsymbol{e}_{-i}) = (m - x\sum_{j\neq i} e_j)^+ \ , \quad (42)$$

where $(x)^+ = \max\{0, x\}$. The Nash equilibrium is the fixed point of the above best response functions, leading to efforts $e = \frac{m}{1+(N-1)x}$ by each agent at the symmetric Nash equilibrium.

When agent i purchases a contract (p, α, β), his expected utility will be given by,

$$\begin{aligned}&\overline{U}_i^{in}(\boldsymbol{e}, p, \alpha, \beta) = \\ &E(-\exp\{-\gamma(-p + \alpha\cdot S_{e_i} - L_{\boldsymbol{e}}^{(i)} + \beta L_{\boldsymbol{e}}^{(i)} - c\cdot e_i)\}) \\ &= -\exp\{\gamma(p + (c-\alpha)e_i + \tfrac{1}{2}\alpha^2\gamma\sigma^2 + \\ &(1-\beta)\mu(e_i + x\textstyle\sum_{j\neq i} e_j) + \frac{\gamma(1-\beta)^2\lambda(e_i + x\sum_{j\neq i} e_j)}{2})\}\end{aligned} \quad (43)$$

Therefore, the best response of agent i, when he enters the contract, is as follows,

$$\begin{aligned}&B_i^{in}(\boldsymbol{e}_{-i}) = (m(\alpha,\beta) - x\textstyle\sum_{j\neq i} e_j)^+ \\ &m(\alpha,\beta) = \arg\min_{e\geq 0} \\ &\qquad (1-\beta)\mu(e) + \tfrac{1}{2}(1-\beta)^2\gamma\lambda(e) + (c-\alpha)e \ .\end{aligned} \quad (44)$$

Similar to the two-agent case, we can write the insurer's contract design problem as follows,

$$\begin{aligned}&\max_{\alpha,\beta,e} N\cdot\{p - \alpha e - \beta\mu(e + x(N-1)e)\} \\ &s.t., \ \text{(IR)}\ \overline{U}_i^{in}(\boldsymbol{e}, p, \alpha, \beta) \geq u^{out} \\ &\quad \text{(IC)}\ \boldsymbol{e} = (e, \cdots, e) \text{ is the effort of the agents} \\ &\quad \text{at the NE where all agents purchase contracts}\end{aligned} \quad (45)$$

Here, u^{out} denotes the utility of an agent when he is opts out of purchasing a contract, while all other agents purchase contracts. We can again show that the individual rationality constraints in the above problem are binding at the optimal contract. Consequently, the insurer's optimization problem simplifies to:

$$\begin{aligned}&\max_{\alpha,\beta,m'} N\cdot \\ &\{w^{out} - \mu(m') - \tfrac{(1-\beta)^2\gamma\lambda(m')}{2} - \tfrac{c\cdot m'}{1+(N-1)x} - \tfrac{\gamma\alpha^2\sigma^2}{2}\} \\ &s.t., \quad (IC) \quad m' = \arg\min_{e\geq 0} \\ &\qquad (1-\beta)\mu(e) + \tfrac{(1-\beta)^2\gamma\lambda(e)}{2} + (c-\alpha)e\end{aligned} \quad (46)$$

where $w^{out} = \frac{\ln(-u^{out})}{\gamma}$. Note also that problem (46) prescribes identical contracts for all agents.

We now analyze the effect of the pre-screening noise, σ, on the state of network security, defined as the sum of all agents' efforts; with homogeneous agents, this is equivalent to each agent's effort.

Theorem 6.1: 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 no 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, is similar to the statements of Theorem 5.2 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 5.2 and $N = 2$ in Theorem 6.1.

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

Theorem 6.2: 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.

B. The case of risk averse insurer and correlated losses

We next study the problem of designing cyber-insurance policies in a network of N homogeneous risk-averse agents with perfect pre-screening (i.e., $\gamma_i = \gamma$ and $c_i = c$ and $\sigma_i = \sigma = 0$) with correlated losses defined as follows.

Let θ be the covariance between any two losses, that is,

$$Cov(L^i_{\boldsymbol{e}}, L^j_{\boldsymbol{e}}) = \theta, \quad \forall i \neq j \tag{47}$$

We further assume that the insurer is risk-averse, with risk attitude $\delta \geq 0$ and the vector $(L^1_{\boldsymbol{e}}, \cdots, L^N_{\boldsymbol{e}})$ has the multivariate Gaussian distribution. The insurer can conduct a pre-screening of each agent's security posture and receives the pre-screening outcome $S_i = e_i$ as the pre-screening is perfect. Similar to (45), we can write the insurer's problem as follows,

$$\begin{aligned}
&\max_{p,\alpha,\beta,e} \\
&\quad E(-\exp\{-\delta(\textstyle\sum_{i=1,\cdots,N} p - \alpha S_i - \beta L^i_{\boldsymbol{e}})\}) \\
&= -\exp\{N\delta(-p + \alpha e + \beta\mu(e + x(N-1)e) + \\
&\qquad \tfrac{\delta\beta^2\lambda(e+x(N-1)e)}{2} + \tfrac{(N-1)}{2}\delta\beta^2\theta)\} \\
&s.t., \text{ (IR) } \overline{U}^{in}_i(\boldsymbol{e}, p, \alpha, \beta) \geq u^{out} \\
&\qquad \text{(IC) } \boldsymbol{e} = (e, e, \cdots, e) \text{ is the effort of the agents} \\
&\qquad\quad \text{at the NE where all the agents purchase contrcts}
\end{aligned} \tag{48}$$

As the (IR) constraint is binding, similar to (46), we have

$$\begin{aligned}
&\max_{\alpha,\beta,e} w^{out} - \mu(m') - \tfrac{\beta^2\delta+(1-\beta)^2\gamma}{2}\lambda(m') \\
&\qquad -\tfrac{c}{1+(N-1)x}m' - \tfrac{(N-1)}{2}\delta\beta^2\theta \\
&s.t., \; m' \in \arg\min_{e\geq 0} \\
&\qquad (1-\beta)\mu(e) + \tfrac{\gamma(1-\beta)^2\lambda(e)}{2} + (c-\alpha)e
\end{aligned} \tag{49}$$

The following theorem characterizes the effect of pre-screening in the presence of a risk averse insurer.

Theorem 6.3: Let $m = \arg\min_{e\geq 0} \mu(e) + \frac{\gamma}{2}\lambda(e) + c$ and assume $\theta = 0$ and $m > 0$. Then the agents exerts higher effort than their effort outside the contract if and only if $\mu'(m) + \frac{1}{2}\frac{\delta\gamma}{\gamma+\delta}\lambda'(m) + \frac{c}{1+(N-1)x} < 0$.

Notice that the condition of Theorem 6.3 reduces to the condition of Theorem 6.1 if we set $\delta = 0$. Also, notice that the condition of Theorem 6.3 is more likely to be satisfied for larger values of δ. For instance, if $\delta = \infty$, the condition is always satisfied, and the agents exert higher effort inside the contract. In other words, if the insurer is more risk averse, it is more likely that she encourages agents to exert higher effort as compared to their efforts outside of the contract.

We close this section by characterizing the effect of correlation on agents' efforts given perfect pre-screening.

Theorem 6.4: Assume $\theta \geq 0$, i.e., positive correlation between losses. Then, agents' efforts inside the contract increase as θ increases.

Theorem 6.4 implies that if agents' losses are more correlated, a risk averse insurer encourages the agents to exert more effort. This is because with correlated losses, it is more likely for losses to happen simultaneously as compared to a scenario with independent losses. Note that when $\delta = 0$ in (49), i.e., when the insurer is risk neutral, the problem becomes independent of θ, meaning that the covariance between any two losses does not affect the optimal contract or the agents' efforts if the insurer is risk neutral.

VII. Numerical results

We next present numerical examples of the findings of Sections IV-VI. Our main focus is on the impact of pre-screening noise in various scenarios. Throughout the first part of this section we use the following parameters:

$$\mu(e) = \frac{10}{e+1}, \quad \lambda(e) = \frac{10}{(e+1)^2}, \quad c = 2, \quad \gamma = 1 \;. \tag{50}$$

A. Impact of Agent's Risk Attitude γ

Figure 1 illustrates the optimal contract as a function of γ. As the agent becomes more risk-averse, the insurer can set a higher base premium p, offer a lower discount factor α, and offer a higher coverage β. In other words, pre-screening becomes less important as the agent's risk-aversion increases, as more risk-averse agents are most interested in transferring more of their risk to the insurer, making their own efforts less important.

Figure 2 illustrates network security (agent's effort), both inside and outside of a contract, vs. his risk attitude γ. First, we see that as suggested by Theorem 4.1, the agent's effort in the contract is less that his effort outside of the contract. In other words, insurance decreases network security. Intuitively, as the agent transfers his risk to the insurer, he does not have the incentive to exert high effort. We also observe that the agent's effort in the optimal contract is a decreasing function of γ. This is due to the fact that as shown in Fig. 1, as the agent becomes more risk-averse, he transfer more risk to the insurer, and further decreases his effort. Finally, when the agent is outside of the contract, he can only decrease his risks by exerting higher effort. Therefore, we observe that as an agent without insurance becomes more risk-averse, he exerts higher effort.

B. Impact of Pre-Screening Noise

A Single Risk-Averse Agent: Figure 3 illustrates the insurer's profit as a function of the pre-screening

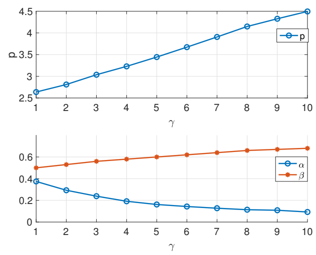

Figure 1: Parameters of the optimal contract v.s. risk aversion level γ

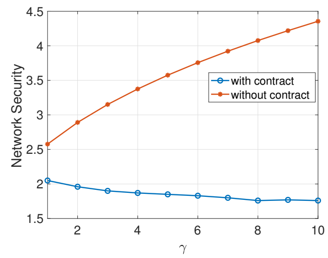

Figure 2: Effort of agent vs. risk aversion level γ

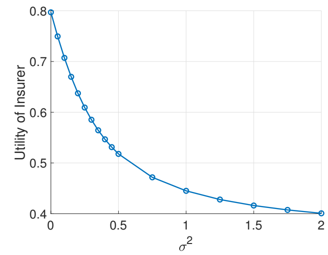

Figure 3: Insurer's profit vs. pre-screening noise σ^2 with a single risk-averse agent

noise σ^2. The observation is consistent with Theorem 4.2, which states that the insurer's profit is a decreasing function of σ^2. Figure 4 illustrates the effort of the agent inside and outside the contract as a function of σ^2. We see that the effort outside the contract is independent of the pre-screening noise, while it decreases inside the contract as σ^2 increases. This highlights that as the insurer becomes less accurate in her observation of the agent's effort, she starts to place less importance on the pre-screening outcome; as a result, it becomes less beneficial for the agent to exert high effort without receiving sufficient discount. In other words, low quality pre-screening dampens its effectiveness in mitigating moral hazard; consequently, network security worsens. A second observation here is that as the participation constraint is always binding, the constant effort outside the contract also means that the agent's utility remains constant regardless of the pre-screening noise. Thus, it is only the insurer who benefits from pre-screening.

Two Homogeneous Risk-Averse Agents: We next consider two homogeneous agents with interdependence factor $x = 0.5$. Figure 5 shows the insurer's utility as a function of the quality of pre-screening, which illustrates the insurer's profit decreases when the pre-screening accuracy decreases. Figure 6 shows the network security as a function of pre-screening noise. Here, the conditions of Theorem 6.1 is satisfied. As we can see, security under the contract is higher than that without insurance for small values of σ; but as σ increases, security worsens and drops below that without contract.

Two Heterogeneous Risk-Averse Agents: We next consider two heterogeneous agents with the following parameters:

$$\mu(e) = \frac{10}{e+1}, \quad \lambda(e) = \frac{10}{(e+1)^2}, \quad c_1 = 1, \quad c_2 = 1.1$$
$$\gamma_1 = 1.2 \;\; \gamma_2 = 1, \quad x = 0.5 \tag{51}$$

We assume that the pre-screening noise (σ^2) is the same for both agents. These parameters together satisfy the condition of Theorem 5.2. Figure 7 shows that the introduction of insurance can indeed improve the state of network security provided the pre-screening is sufficiently accurate. Figure 8 shows that the insurer's profit decreases as pre-screening becomes less accurate.

C. On the Sufficient Conditions of Theorem 5.2

Consider an example with parameters similar to those given in (51), except that $\gamma_1 = 1.5$ and $c_2 = 1.5$. In this case, it can be verified that the conditions of Theorem 5.2 do not hold. However, Figure 9 shows that network security improves after the introduction of insurance. This example shows that the sufficient conditions in Theorem 5.2 are not necessary.

Consider again the same parameters given in (51), except $x = 0.15$. In this case, it can again be verified that the conditions of Theorem 5.2 do not hold. Figure 10 shows that the network security worsens with the introduction of insurance and thus the sufficient conditions are meaningful.

D. Loss with Exponential Distribution and Pre-Screening with Uniform Distribution: An example

Single Risk-Averse Agent: Throughout our analysis, we assumed that losses and pre-screening outcomes are normally distributed. In this section, we provide a numerical example under the assumption of exponentially distributed losses and uniformly distributed pre-screening outcomes. We illustrate how our previous observations hold in this instance as well. Let,

$$\begin{aligned} &\gamma = 0.9, \;\; c = 0.25, \;\; E(L_e) = \mu(e) = \tfrac{1}{1+e}, \\ &L_e \sim \exp(\tfrac{1}{\mu(e)}), \\ &S_e = e + W, \;\; W \sim \;\; Unif(-b, b) \end{aligned} \tag{52}$$

Figure 11 illustrates the agent's effort when pre-screening noise W is uniformly distributed in interval $[-b, b]$. This figure shows that even though the loss and pre-screening outcome are not normally distributed, the agent's effort inside the contract is less than outside the contract; similarly, it remains a decreasing function of b.

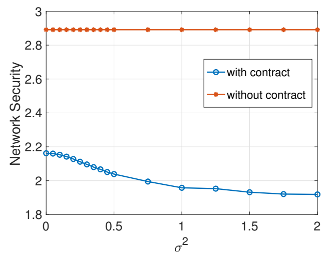

Figure 4: Agent's effort vs. pre-screening noise σ^2 with a single risk-averse agent

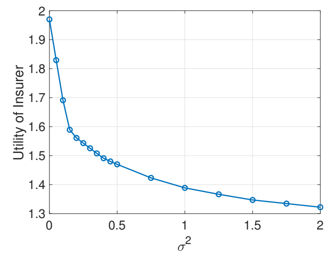

Figure 5: Principal's utility vs. σ^2 with two homogeneous risk-averse agents

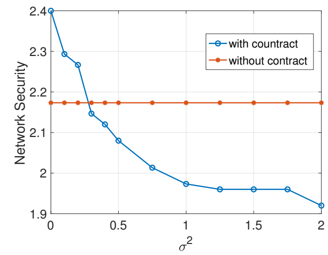

Figure 6: Network security $(e_1 + e_2)$ vs. σ^2 with two homogeneous risk-averse agents

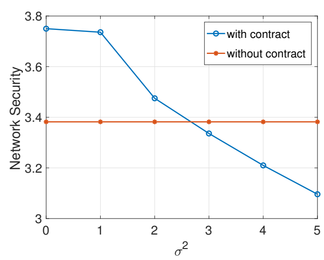

Figure 7: Network security $(e_1 + e_2)$ vs. σ^2 with two heterogeneous risk-averse agents

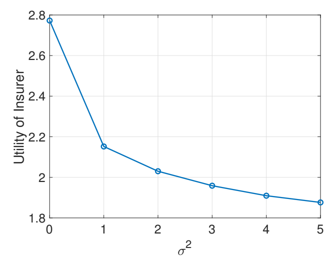

Figure 8: Principal's profit vs. σ^2 with two heterogeneous risk-averse agents

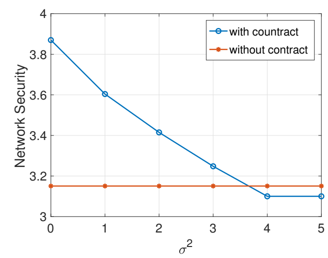

Figure 9: Network security $(e_1 + e_2)$ vs. σ^2 with two heterogeneous risk-averse agents. In this example, the conditions of Theorem 5.2 do not hold but network security improves after the introduction of insurance.

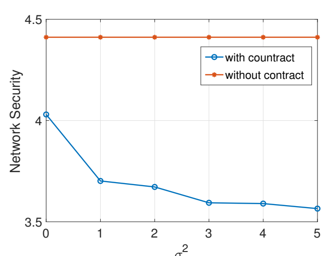

Figure 10: Network security $(e_1 + e_2)$ vs. σ^2 with two heterogeneous risk-averse agents. In this example, the conditions of Theorem 5.2 do not hold, and network security worsens after introduction of insurance.

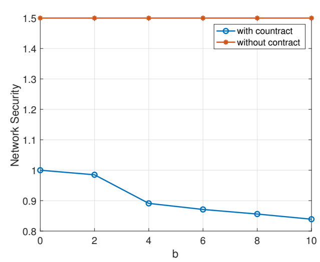

Figure 11: Agent's effort vs. σ^2 with a single risk-averse agent and exponentially distributed loss.

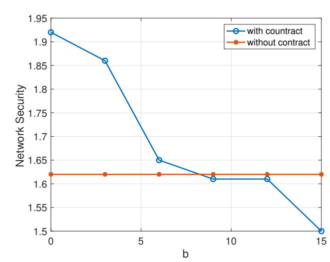

Figure 12: Network security $(e_1 + e_2)$ vs. σ^2 with two heterogeneous risk-averse agents with exponentially distributed interdependent losses.

Model with Two Risk-Averse Agents: We further consider a network of two risk-averse agents with the following parameters,

$$\begin{gathered}\gamma_1 = \gamma_2 = 0.9,\ c_1 = 0.25,\ c_2 = 0.5, x = 0.5 \\ E(L^i_{e_1,e_2}) = \mu(e_i + xe_{-i}) = \frac{1}{1+e_i+xe_{-i}} \\ L^i_{e_1,e_2} \sim \exp(\frac{1}{\mu(e_i+xe_{-i})}), \\ S_{e_i} = e_i + W_i,\ W_i \sim\ Unif(-b,b),\ i = 1,2\end{gathered} \tag{53}$$

Where, W_1 , W_2 are independent and uniformly distributed in interval $[-b, b]$.

Figure 12 illustrates network security in a network of two risk-averse agents with exponentially distributed interdependent losses and uniformly distributed pre-screening outcomes. In this example, when pre-screening is sufficiently accurate (b is sufficiently small), by exploiting agents' interdependence, the insurer can design contracts in a way that network security inside the contract is higher than prior to the introduction of insurance. In contrast, when pre-screening is not accurate enough (b is large), network security inside the contract falls bellow network security outside the contract. Again, these observations are consistent with our results under normally distributed losses and pre-screening.

VIII. Conclusions and Discussions

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 the result could be different in a network of interdependent agents. Specifically, 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 and essentially selling commitment to interdependent agents. We show under what conditions this type of contracts leads to not only increased profit for the principal and utility for the agents, but also improved state of network security.

There are a number of directions to pursue to extend the above results. As mentioned earlier, all our results are derived under the assumption of perfect information. Studying the 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 more general ways of 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. Finally, a competitive market setting and its effects on network security is also worth studying.

Appendix: Proofs are given in [27].

References

[1] M. M. Khalili, P. Naghizadeh, and M. Liu, "Designing cyber insurance policies: Mitigating moral hazard through security pre-screening," in *The 7th International EAI Conference on Game Theory fo Networks (Gamenets)*, 2017.

[2] M. M. Khalili, P. Naghizadeh, and M. Liu, "Designing cyber insurance policies in the presence of security interdependence," in *The 12th Workshop on the Economics of Networks, Systems and Computation (NetEcon)*, 2017.

[3] C. Hemenwa, "ABI Research: Cyber insurance market to reach $10B by 2020," *www.advisenltd.com/2015/07/30/abi-research-cyber-insurance-market-to-reach-10b-by-2020/*, 2015.

[4] Insurance Information Institute, "U.S. cyber insurance market demonstrates growth, innovation in wake of high profile data breaches," *www.iii.org/pressrelease/us-cyber-insurance-market-demonstrates-growthinnovation-in-wake-of-high-profile-data-breaches-102015*, 2015.

[5] N. Shetty, G. Schwartz, and J. Walrand, "Can competitive insurers improve network security?," in *The Third International Conference on Trust and Trustworthy Computing (TRUST)*, 2010.

[6] N. Shetty, G. Schwartz, M. Felegyhazi, and J. Walrand, "Competitive cyber-insurance and internet security," *Economics of Information Security and Privacy*, pp. 229–247, 2010.

[7] G. Schwartz, N. Shetty, and J. Walrand, "Cyber-insurance: Missing market driven by user heterogeneity," 2010.

[8] G. A. Schwartz and S. S. Sastry, "Cyber-insurance framework for large scale interdependent networks," in *The Third International Conference on High Confidence Networked Systems*, 2014.

[9] H. Ogut, N. Menon, and S. Raghunathan, "Cyber insurance and it security investment: Impact of interdependence risk," in *The Workshop on the Economics of Information Security*, 2005.

[10] Z. Yang and J. C. S. Lui, "Security adoption and influence of cyber-insurance markets in heterogeneous networks," *Performance Evaluation*, vol. 74, pp. 1–17, Apr. 2014.

[11] J. Kesan, R. Majuca, and W. Yurcik, "The economic case for cyberinsurance," University of Illinois Legal Working Paper Series uiuclwps-1001, University of Illinois College of Law, 2004.

[12] J. P. Kesan, R. P. Majuca, and W. Yurcik, "Cyber-insurance as a market-based solution to the problem of cybersecurity," in *The Workshop on the Economics of Information Security*, 2005.

[13] A. Hofmann, "Internalizing externalities of loss prevention through insurance monopoly: an analysis of interdependent risks," *The Geneva Risk and Insurance Review*, 2007.

[14] M. Lelarge and J. Bolot, "Economic incentives to increase security in the internet: The case for insurance," in *IEEE International Conference on Computer Communications (INFOCOM)*, 2009.

[15] S. Romanosky, L. Ablon, A. Kuehn, and T. Jones, "Content analysis of cyber insurance policies: How do carriers write policies and price cyber risk?," 2017.

[16] Y. Liu, A. Sarabi, J. Zhang, P. Naghizadeh, M. Karir, M. Bailey, and M. Liu, "Cloudy with a chance of breach: Forecasting cyber security incidents," in *USENIX Security Symposium*, 2015.

[17] R. A. Miura-Ko, B. Yolken, N. Bambos, and J. Mitchell, "Security investment games of interdependent organizations," in *The 46th Annual Allerton Conference on Communication, Control, and Computing*, 2008.

[18] B. Johnson, J. Grossklags, N. Christin, and J. Chuang, "Are security experts useful? bayesian nash equilibria for network security games with limited information," in *The 15th European Symposium on Research in Computer Security, (ESORICS)*, 2010.

[19] B. Johnson, J. Grossklags, N. Christin, and J. Chuang, "Uncertainty in interdependent security games," in *The First Conference on Decision and Game Theory for Security (GameSec)*, 2010.

[20] M. Lelarge, "Coordination in network security games: A monotone comparative statics approach.," *IEEE Journal on Selected Areas in Communications*, vol. 30, no. 11, pp. 2210–2219, 2012.

[21] R. Böhme and G. Schwartz, "Modeling cyber-insurance: Towards a unifying framework.," in *The Workshop on the Economics of Information Security*, 2010.

[22] A. Mas-Colell, M. D. Whinston, and J. R. Green, *Microeconomic theory*. New York: Oxford University press, 1995.

[23] J. Bolot and M. Lelarge, "Cyber insurance as an incentive for internet security," in *The Workshop on the Economics of Information Security*, 2008.

[24] R. Pal, L. Golubchik, K. Psounis, and P. Hui, "Will cyber-insurance improve network security? a market analysis," in *IEEE International Conference on Computer Communications (INFOCOM)*, 2014.

[25] A. Laszka, M. Felegyhazi, and L. Buttyan, "A survey of interdependent information security games," *ACM Comput. Surv.*, vol. 47, pp. 23:1–23:38, Aug. 2014.

[26] L. Jiang, V. Anantharam, and J. Walrand, "How bad are selfish investments in network security?," *IEEE/ACM Transactions on Networking*, vol. 19, pp. 549–560, April 2011.

[27] M. M. Khalili, P. Naghizadeh, and M. Liu, "Designing cyber insurance policies: The role of pre-screening and security interdependence." Accessed on 2017-10-13. [Online]. Available: http://bit.ly/2zkK2HJ.