Effective Premium Discrimination for Designing Cyber Insurance Policies with Rare Losses

Mohammad Mahdi Khalili, Xueru Zhang, and Mingyan Liu

University of Michigan, Ann Arbor, MI , USA
{khalili,xueru,mingyan}@umich.edu

Abstract. Cyber insurance like other types of insurance is a method of risk transfer, where the insured pays a premium in exchange for coverage in the event of a loss. As a result of the reduced risk for the insured and the lack of information on the insurer's side, the insured is generally inclined to lower its effort, leading to a worse state of security, a common phenomenon known as moral hazard. To mitigate moral hazard, a widely employed concept is premium discrimination, i.e., an agent/insured who exerts higher effort pays less premium. This, however, relies on the insurer's ability to assess the effort exerted by the insured. In this paper, we study two methods of premium discrimination that rely on two different types of assessment: pre-screening and post-screening. Pre-screening occurs before the insured enters into a contract and can be done at the beginning of each contract period; the result of this process gives the insurer an estimated risk on the insured, which then determines the contract terms. The post-screening mechanism involves at least two contract periods whereby the second-period premium is increased if a loss event occurs during the first period.
Prior work shows that both pre-screening and post-screening are generally effective in mitigating moral hazard and increasing the insured's effort. The analysis in this study shows, however, that the conclusion becomes more nuanced when loss events are rare. Specifically, we show that post-screening is not effective at all with rare losses, while pre-screening can be an effective method when the agent perceives them as rarer than the insurer does; in this case pre-screening improves both the agent's effort level and the insurer's profit.

Keywords: Cyber Insurance · Premium Discrimination · Risk Assessment · Rare Losses

1 Introduction

Facing increasingly common cyber attacks and data breaches, organizations and businesses big and small have to invest in cyber self-protection against a myriad of losses, such as business interruption induced by such incidents. Organizations

This work is supported by the NSF under grants CNS-1616575, CNS-1739517, and ARO W911NF1810208.

are also increasingly turning to cyber insurance as a form of protection for mitigating cyber risks by transferring all or part of their risks to the insurer through the purchase of a policy [1, 2]. Specifically, a cyber insurance contract is between a risk averse agent and an insurer; the agent pays a premium in exchange for the insurer to provide certain coverage in the event of a loss. Risk aversion on the agent's part makes him willing to buy insurance from the insurer to undertake the risk, resulting in reduced uncertainty for the insured agent.

One of the challenges in offering an insurance contract is the lack of information on the insured, which results in the well-known moral hazard issue. In other words, the insurer is unaware of the agent's effort in self-protection and therefore the latter's true risk. This, together with the fact that the agent is now (partially) covered in his loss, typically leads the agent to exert less effort toward securing himself. This results in a worse state of cyber risk as compared to a scenario with uninsured agents [3–6].

The problem of designing cyber insurance policies in the presence of moral hazard has been studied in the literature and it has been shown that the impact of the insurance contract on the state of the network security depends on the insurance market [7–10]. The key idea for mitigating moral hazard is premium discrimination, i.e., the agent who invests more in self-protection gets charged less premium. Yang *et al.* [9] consider a competitive insurance market and show that insurance cannot improve the state of network security in the presence of moral hazard; on the other hand, in the absence of moral hazard and with observable agents' actions, the insurer can premium discriminate and the state of security can improve as a result. Hofman [8] studies a monopoly insurance market in the presence of a welfare-maximizing insurer. In this case, the insurer premium discriminates among high and low risk agents using imperfect information that the insurer has, and the insurance contract can incentivize the agents to exert higher effort as compared to the no-insurance case.

In practice premium discrimination can be achieved in a number of ways. Traditional insurance products (e.g., auto, life, home, property) rely on actuarial models that estimate risks based on a variety of inputs obtainable through questionnaires or surveys. For instance, getting an auto insurance policy requires the submission of information on the model/year of the car being insured, the primary driver(s) of the vehicle, their age, gender, marital status, place of residence and so on, a process we refer to as *pre-screening* throughout the paper. The estimated risk based on this type of input directly determines the premium on the policy or a set of policies (with different choices of premium-deductible combinations) offered by the insurer. Furthermore, when a driver continues to purchase insurance over multiple years, then his/her previous driving and claims record also factor into future-year premium calculation. This latter element of premium discrimination, referred to as *post-screening* throughout the paper, has been shown to be effective in general: since an agent faces (potentially significantly) higher payments in the future, there is incentive for the agent to act responsibly (exert high effort) in the present time to avoid a loss event, see e.g., Rubinstein *et. al* [11].

The above in principle applies to the domain of cyber insurance, but with two challenges. (1) pre-screening is much harder to do for lack of actuarial models, and (2) while cyber attacks are increasingly common collectively, for a single organization it remains a relatively rare occurrence with high losses and damages, which means post-screening may seldom come into effect. In this study we shall examine these two mechanisms separately and attempt to understand under what conditions are these effective in incentivizing the agent to exert higher effort, thereby improving the state of security. Note that rare cyber incidents are different from natural disasters that have been studied in the literature [12, 13]; the latter are also rare incidents with high losses but differ from cyber incident in the following sense. The agents/insureds cannot prevent natural disasters by exerting effort. For instance, the authors in [12, 13] do not consider the agent's effort in their models as it does not affect the probability of natural disaster occurrence. On the other hand, an agent can actively and proactively work toward decreasing his chance of being attacked or an attack being successful by investing in security and addressing vulnerability.

In this paper, we shall assume that data breach and loss incidents are rare for each agent but the amount of loss from a breach is extremely large. This model is reasonably borne out by recent events such as the Equifax data breach, which affected 143 million American consumers and incurred $68.6 billion in loss for the company [14]; most of these events have been unprecedented in the respective victim's company history.

Our main finding in this paper is that post-screening (which involves at least two contract periods) is not effective at all with rare loss incidents. On the other hand, pre-screening can be an effective method if the agent perceives loss incidents as rarer than the insurer does; in this case *sufficiently accurate* pre-screening can be effective and improves the state of security as well as the insurer's profit as compared to not using premium discrimination.

The organization of this paper is as follows. In Section 2 we introduce the model and contract design problem. Section 3 summarizes prior results (but recast under our model) on designing cyber insurance policies when incidents are not rare. In Section 4 we examine the effect of pre-screening and post-screening on both the state of security and the insurer's profit with rare losses. Section 5 discusses how pre-screening may be used to enable an active policy, as well as dependent cyber risks. Section 6 presents numerical results and Section 7 concludes the paper.

2 Model

We consider the cyber insurance design, a principal-agent problem, between a profit-maximizing, risk-neutral insurer/principal and a risk-averse insured/agent. The agent exerts effort e toward securing himself, incurring linear cost $c \cdot e$.

Let $p(e)$ denote the probability of a loss incident, assumed to be strictly decreasing and strictly convex. Decreasing and convexity imply that the initial effort toward security leads to a considerable reduction in probability of a loss

incident, and strict convexity implies that the probability of the loss incident cannot be zero even if the agent exerts high effort [15]. Specifically, we assume that probability of a loss incident has the following form,

$$p(e) = t \cdot \exp\{-\alpha \cdot e\}, \tag{1}$$

where t is the nominal probability of a successful attack to the agent if he exerts zero effort ($e = 0$) and α is a constant. Larger α implies that investment in security is more efficient and $p(.)$ converges to zero faster. Note that t and α both are constants and cannot be modified by the agent or the insurer.

When a loss occurs, the agent suffers the amount of loss l, also a constant. This is obviously a simplification; however, our qualitative conclusions remain the same for a random loss given by a known distribution. The expected utility of the agent without any insurance contract is given by:

$$U(e) = p(e)f(-l - ce) + (1 - p(e))f(-ce), \tag{2}$$

where $f(.)$ is a concave function that captures the agent's risk aversion. To make the analysis concrete, we will further assume $f(.)$ is an exponential function with constant absolute risk aversion γ:

$$f(y) = 1 - \exp\{-\gamma \cdot y\}\ , \tag{3}$$

where γ is referred to as the agent's risk attitude; the higher the risk attitude the more risk averse the agent.

2.1 Agent's Effort & Utility Without Insurance

Without insurance, the agent exerts an effort level e^o to maximize his utility:

$$e^o = \arg\max_{e \geq 0} U(e)\ . \tag{4}$$

It is easy to see that if $\gamma c \geq \alpha$, then $e^o = 0$. Intuitively, $\gamma c \geq \alpha$ implies that the cost of effort is higher than its benefit, and the agent is not able to improve his utility by exerting effort. If $\alpha > \gamma c$, then e^o is given by the first order condition. Together, we have

$$e^o = \begin{cases} 0 & \text{if } \gamma c \geq \alpha \\ (\frac{1}{\alpha} \ln t \cdot \frac{(\alpha - \gamma c)(\exp\{\gamma l\} - 1)}{\gamma c})^+ & \text{if } \gamma c < \alpha \end{cases} \tag{5}$$

where $(a)^+ = \max\{0, a\}$. As a result, the maximum utility of the agent outside the contract is given by,

$$u^o = U(e^o) = \begin{cases} 1 - \frac{\alpha}{\alpha - \gamma c}(t \cdot \frac{\alpha - \gamma c}{\gamma c}(\exp\{\gamma l\} - 1))^{\frac{\gamma c}{\alpha}} & \text{if } e^o > 0 \\ t \cdot (1 - \exp\{\gamma l\}) & \text{if } e^o = 0\ . \end{cases} \tag{6}$$

$p(e)$ can be written as $t \cdot (\exp\{-\alpha\})^e$ which is a function consistent with the exponential probability function introduced in [16].

2.2 Contract Design

We will assume that in the event of a loss, a contract covers the full amount l. This is again a simplification but it allows us to get to the essence of our analysis more straightforwardly without affecting the main qualitative conclusions. Because a loss is covered in full, the agent will exert zero effort after entering an insurance contract. Thus the insurer will have to use premium discrimination to incentivize the insured to exert a higher effort in exchange for lower premium. We next describe in detail the resulting contract design problem under two different methods of premium discrimination: post-screening and pre-screening.

Post-screening In this case the contract design problem is framed in a two-period setting where the insurer is able to assess premium in the second period based on what happens in the first period. Such a contract is given by three parameters (π_1, π_2, π_3): π_1 is the first-period premium; in the second period, the agent pays premium π_2 if a loss happened (and was covered in full) during the first period and pays π_3 otherwise. Obviously $\pi_3 \leq \pi_2$.

In this case, the agent may exert non-zero effort in the first period to decrease the chance of a loss in order to reduce the likelihood of paying a higher premium in the second period. In the second period, on the other hand, the agent will always exert zero effort as the loss is fully covered and he faces no more future punishment.

We assume that when an agent enters such a contract he commits to both periods. The agent's utility inside a contract (π_1, π_2, π_3) with post-screening is thus the summation of his utility in each period:

$$U^{in}(e, \pi_1, \pi_2, \pi_3) = f(-\pi_1 - ce) + p(e)f(-\pi_2) + (1 - p(e))f(-\pi_3), \tag{7}$$

where e is the effort in the first period.

The insurer's problem is to maximize her profit subject to the Individual Rationality (IR) constraint and Incentive Compatibility (IC) constraint:

$$\begin{aligned} V = & \max_{\{\pi_1, \pi_2, \pi_3, e\}} \pi_1 - p(e)l + p(e)(\pi_2 - p(0)l) + (1 - p(e))(\pi_3 - p(0)l) \\ \text{s.t.} \quad & (IR) \quad U^{in}(e, \pi_1, \pi_2, \pi_3) \geq 2 \cdot u^o, \ i = 1, 2 \\ & (IC) \quad e \in \arg\max U^{in}(e, \pi_1, \pi_2, \pi_3). \end{aligned} \tag{8}$$

The (IR) constraint ensures that the agent enters the contract only if he gets no lower utility than his outside option. Note that since the contract covers two periods, the comparison here is between his utility inside the contract over two periods and outside the contracts over two periods. The (IC) constraint suggests that the agent acts in self-interest: he exerts an effort level maximizing his utility given the policy parameters.

Our analysis can be extended to a multi-period setting where the premium of each period depends on the agent's history of losses, i.e., the agent's third-period premium depends on his loss events in the first and second periods and so on.

Under the contract (π_1, π_2, π_3), by the first order condition, the agent's optimal effort e^{in} is given by:

$$e^{in}(\pi_1,\pi_2,\pi_3) = \begin{cases} \left(\frac{1}{\alpha+\gamma c}\ln(t\cdot\frac{\alpha}{\gamma c}\frac{\exp\{\gamma\pi_2\}-\exp\{\gamma\pi_3\}}{\exp\{\gamma\pi_1\}})\right)^+ & \text{if } \pi_2 > \pi_3 \\ 0 & \text{if } \pi_2 \le \pi_3 \end{cases} . \quad (9)$$

For notational convenience, we use e^{in} instead of $e^{in}(\pi_1, \pi_2, \pi_3)$, while noting the dependency. We have the following lemma on the (IR) constraint.

Lemma 1. *The (IR) constraint in the optimization problem* (8) *is binding.*

The above lemma implies that at the optimal solution, the agent is indifferent between entering vs. not entering the contract, as expected.

Pre-screening We now turn to the case of pre-screening. We assume the insurer can conduct a risk assessment prior to determining the contract terms; the determination mechanism is known to the agent so this is again a game of perfect information. We assume the outcome of the pre-screening is given by an assessment $S = e + N$, where N is a zero-mean Gaussian noise with variance σ^2. There are various ways to achieve pre-screening in practice, using surveys, penetration tests, or advanced Internet measurement techniques, see e.g., [17].

The insurer then offers the agent a contract given by two parameters (π, β), where π is the base premium and β is the assessment-dependent discount factor: the agent pays $\pi - \beta S$ in exchange for full coverage in the event of a loss. The agent's total cost inside the contract (π, β) while exerting effort e is:

$$X^{in} = \pi - \beta \cdot S + c \cdot e \ . \quad (10)$$

As X^{in} follows a Gaussian distribution, using moment-generating function the agent's *expected* utility under the contract is given by:

$$U^{in}(\pi,\beta,e) = E(f(-X^{in})) = 1 - \exp\{\gamma\pi + \gamma(c-\beta)e + \tfrac{\gamma^2\beta^2\sigma^2}{2}\}. \quad (11)$$

Therefore, the insurer's design problem using pre-screening is as follows:

$$\begin{aligned} \max_{\pi,\beta,e} \ & E\{\pi - \beta S\} - p(e)\cdot l \\ \text{s.t.} \ \ (IR) \quad & U^{in}(\pi,\beta,e) \ge u^o, \\ (IC) \quad & e \in \arg\max_{e'\ge 0} U^{in}(\pi,\beta,e') \end{aligned} \quad (12)$$

Similar as in Lemma 1, we can show that the (IR) constraint is binding in this case. Thus we have the following relation between optimal contract parameters $(w^o = \frac{1}{\gamma}\ln(1-u^o))$:

$$\pi = w^o + \beta e - ce - \frac{\gamma\beta^2\sigma^2}{2} \ . \quad (13)$$

The analysis can be extended to other noise distributions.

Using (13), the insurer's problem can be simplified as follows:

$$V(\sigma) = \max_{\beta,e} \; w^o - ce - \frac{\gamma\beta^2\sigma^2}{2} - p(e)l$$
$$\text{s.t. } (IC) \quad e \in \arg\min_{e'\geq 0}(c-\beta)e' + \frac{\gamma\beta^2\sigma^2}{2}, \tag{14}$$

We next summarize (known) results on these two types of premium discrimination in terms of their effectiveness in incentivizing efforts.

3 State of Security and Optimal Contract when Losses Are Not Rare

Post-screening: Post-screening has been studied in the literature. Rubinstein *et.al.* in [11] showed that post-screening can improve the agent's effort inside the contract compared to the one-period contract without post-screening.

This can be similarly observed in our model. In particular, in Theorem 1 below we introduce a sufficient condition under which the agent exerts non-zero effort in the first period of a contract with post-screening. In Section 6, we also provide an example where the agent inside a contract with post-screening exerts higher effort as compared to the no-insurance scenario.

Theorem 1. *Let* $(\hat{\pi}_1, \hat{\pi}_2, \hat{\pi}_3, \hat{e})$ *be the solution of the optimization problem* (8). *Suppose that* $t = 1$ *and* $\left[\frac{(\alpha-\gamma c)(\exp\{\gamma l\}-1)}{\gamma c}\right] > 1$, *then* $\hat{e} > 0$.

Theorem 1 suggests that post-screening can be an effective mechanism to incentivize non-zero effort. Note that the condition $\left[\frac{(\alpha-\gamma c)(\exp\{\gamma l\}-1)}{\gamma c}\right] > 1$ in theorem 1 can be satisfied if loss l is sufficiently large.

Pre-screening: Our previous work [4] shows that pre-screening can simultaneously incentivize the agent to exert non-zero effort and improve the insurer's utility. This is characterized for the present model in the following theorem.

Theorem 2. *Pre-screening incentivizes non-zero effort if and only if*

$$\frac{c}{\alpha \cdot t \cdot l} < 1 \tag{15}$$
$$\sigma^2 \leq \frac{-\frac{c}{\alpha} - \frac{c}{\alpha}\ln\frac{c}{\alpha\cdot t\cdot l} + t\cdot l}{0.5\cdot c^2\gamma}\ . \tag{16}$$

Theorem 2 suggests pre-screening is effective if and only if it is sufficiently accurate (Eqn (16)) and the expected loss $t \cdot l$ is sufficiently large (Eqn (15)). Further, the next theorem identifies the relation between insurer's profit and pre-screening accuracy.

Theorem 3. *Let* $V(\sigma)$ *be the insurer's maximum utility. That is,*

$$V(\sigma) = \max_{\beta,e} w^o - ce - \frac{\gamma\beta^2\sigma^2}{2} - p(e)l \quad s.t. \quad IC\ constraint \tag{17}$$

Then, $V(\sigma)$ *is decreasing in* σ.

4 State of Security and Optimal Contract When Losses Are Rare

We next consider the case when loss events are rare, by assuming its likelihood diminishes (i.e., $t \to 0$) but that the loss amount is high in such an event (i.e., $l \to \infty$). This model is motivated by recent data breaches that result in extremely high losses and damages but remain relatively rare for a single organization as mentioned earlier.

Furthermore, we would like to explicitly capture a common asymmetry in perception between the insurer and the agent, i.e., the latter tends to think of loss as rarer than the former does. Specifically, let t_a and t_p denote the nominal attack probability from the agent and the insurer's perspective, respectively. By our assumption, both t_a and t_p go to zero and l goes to infinity. For tractability, we adopt the following assumptions on t_a and t_p and l,

$$\begin{array}{l} \lim_{\{t_a \to 0, l \to \infty\}} t_a \cdot \exp\{\gamma l\} = \exp\{\gamma l_a\} \\ \lim_{\{t_p \to 0, l \to \infty\}} t_p \cdot l = l_p, \end{array} \tag{18}$$

where l_a and l_p are the perceived expected loss from the agent and the insurer's perspective when the agent exerts zero effort, respectively. It is worth noting that Eqn (18) implies that the expected loss is always limited. Otherwise the cyber insurance market may not exist. Moreover, (18) implies that $t_a = \frac{\exp\{\gamma l_a\}}{\exp\{\gamma l\}}$ goes to zero exponentially while $t_p = \frac{l_p}{l}$ goes to zero slower than t_a as l goes to infinity, i.e, $t_a > t_p$ as $l \to \infty$. Therefore, the agent thinks the loss is rarer than the insurer does.

4.1 Post-screening

With the above rare loss assumptions, we have the following theorem on post-screening.

Theorem 4. *Using post-screening and given $t \to 0$,*

1. *the agent always exerts zero effort inside the contract, and*
2. *at the optimal contract we have,*

$$\pi_1 = \pi_3 = \frac{1}{\gamma} \ln\left[1 - u^o\right], \quad \pi_2 \in \mathcal{R}^+$$

.

Theorem 4 implies that premium discrimination in the second period based on the first period is not at all effective and the insurer is not able to improve the agent's effort or her utility by post-screening as compared to a contract without premium discrimination.

By assuming that t goes to zero, the entire probability of a loss incident (i.e., $p(e) = t\exp(\alpha(e))$ goes to zero.

If the agent exert effort e, then $l_a \exp\{-\alpha \cdot e\}$ and $l_p \exp\{-\alpha \cdot e\}$ are the perceived expected loss from the agent and the insurer's perspective.

4.2 Pre-screening

For pre-screening, it turns out perception asymmetry makes a difference. The following theorem characterizes the optimal contract and introduces a sufficient condition under which pre-screening can incentivize the agents to exert non-zero effort inside the optimal contract.

Theorem 5. *Pre-screening can incentivize non-zero effort under the rare loss model, if and only if*

$$\frac{c}{\alpha l_p} < 1 \tag{19}$$

$$\sigma^2 \leq \frac{-\frac{c}{\alpha} - \frac{c}{\alpha} \ln\left[\frac{c}{\alpha \cdot l_p}\right] + l_p}{0.5 \cdot c^2 \gamma} \ . \tag{20}$$

Note that conditions in Theorem 2 reduce to those in Theorem 5 if we substitute tl with l_p in (15) and (16). Theorem 5 implies that pre-screening incentivizes effort if and only if the pre-screening is sufficiently accurate and insurer's perceived loss l_p is sufficiently large.

5 Discussion

5.1 Contingencies on Periodic Pre-screening: Active Policy

So far we have assumed that the agent exerts a one-shot effort level, which applies to the entire policy period. Under this assumption, pre-screening helps incentivize non-zero effort. In reality, keeping risk at a certain level typically requires sustained effort throughout the period, and it is conceivable that the insured may choose to lower his effort after the initial risk assessment (yet another form of moral hazard). If so then our results on pre-screening suggests that it has to be performed more often, whereby premium adjustment is made following each screening. This effectively means that the initial contract is an *active policy* with built-in contingencies, and the actual premium payable is realized over time dependent on the screening results. We illustrate this idea using the following example with one additional, mid-term, screening.

Let's assume that the agent exerts effort e before the first screening, resulting in assessment outcome $S = e + N$ as before, and then he lowers the effort to e'. Accordingly, let $S' = e' + N'$ be the outcome of the second, mid-term screening, where N' is a zero-mean Gaussian noise with variance σ^2. We assume that N, N' are independent random variables. Below we show that the insurer is able to incentivize the agent not to decrease the effort level through the second screening, i.e., to ensure $e' = e$. Consider an active contract with three parameters (π, β, β') offered to the agent, where β' is a penalty factor, β a discount factor, and π the base premium. The total cost of the agent is given by,

$$X^{in} = \pi - \beta \cdot S + ce + \beta'(S - S') - b(e - e'), \tag{21}$$

where $0 \leq b \leq c$ and b is the benefit of lowering the effort, and $\beta'(S - S')$ is the penalty that the insured would pay after the second risk assessment.

Similar to (11), the agent's expected utility under contract (π, β, β') is:

$$U^{in}(\pi, \beta, e, \beta', e') = E(f(-X^{in})) =$$
$$1 - \exp\{\gamma\pi + \gamma(c - b + \beta' - \beta)e + \gamma(-\beta' + b)e' + \gamma^2\sigma^2 \frac{(\beta - \beta')^2 + (\beta')^2}{2}\}. \tag{22}$$

The insurer's problem can be written as follows:

$$\begin{aligned} R(\sigma) = \max_{\{\pi,\beta,e,\beta',e'\}} \quad & [E\{\pi - \beta S + \beta'(S - S')\} - p(e')l] \\ \text{s.t.} \quad (IR) \quad & U^{in}(\pi, \beta, e, \beta', e') \geq u^o \\ (IC) \quad & (e, e') \in \arg\max_{\tilde{e},\tilde{e}'} U^{in}(\pi, \beta, \tilde{e}, \beta', \tilde{e}'),\ e' \leq e \end{aligned} \tag{23}$$

The following theorem shows that the second risk assessment is effective in preventing the agent from lowering his effort.

Theorem 6. *Let $\hat{e}$ and $\hat{e}'$ be the agent's effort level at the solution to* (23)*, and $\overline{e}$ be the optimal effort level in optimization problem* (14)*. Then, we have $\hat{e} = \hat{e}'$, and the optimal contract parameters are $\beta = c$ and $\beta' = b$ if $\hat{e} > 0$ otherwise they are $\beta = \beta' = 0$. Moreover, if $\overline{e} > 0$, then $\hat{e} = \overline{e}$.*
Lastly, we have $V(\sigma) \leq R(\sigma)$, where $V(\sigma)$ is obtained from (14) *by assuming there is only one pre-screening and the agent does not lower his effort afterward, with equality achieved if $b = c$.*

The last part of the theorem above suggests that performing the second screening helps the insurer to improve profit even when the agent may be assumed not to lower his effort. This is because second pre-screening decreases the variance and uncertainty in agent's utility. Therefore, a risk averse agent is willing to pay more premium when the uncertainty and variance on his side decreases.

5.2 Insuring Interdependent Agents

So far we have assumed that the probability of a loss incident is solely determined by the effort of the agent. On the other hand, risk dependency is a unique feature of cyber risks: the incident probability for an agent may depend on the effort levels of other agents (the former's vendors or service providers, etc.). In our previous work [4], we considered a cyber insurance market in the presence of risk dependency, and showed that the insurer can achieve higher profit as compared to a network of independent agents; moreover, pre-screening in such a case increases the agents' efforts as compared to the no insurance scenario. If we introduce security dependency into our rare loss model, it can be shown that post-screening is not able to incentivize non-zero effort while pre-screening can. Table 1 summarizes the role of dependency and rare loss on the agents' effort, where $(*)$ indicates the associated result holds under certain conditions.

	Pre-screening	Post-screening
Rare Loss, dependent agents	$e^{in} > e^o$ $(*)$	$e^{in} = 0$
Rare Loss, independent agents	$e^{in} > e^o$ $(*)$	$e^{in} = 0$
Frequent loss, dependent agents	$e^{in} > e^o$ $(*)$	$e^{in} > e^o$ $(*)$
Frequent loss, independent agents	$e^o \geq e^{in} \geq 0$	$e^{in} > e^o$ $(*)$

Table 1: Comparing agent's effort inside (e^{in}) and outside (e^o) a contract

6 Numerical Result

We show a number of numerical examples with the following parameters $\gamma = c = 1$, $\alpha = 1.5$.

6.1 Frequent Losses: Post-screening

Our first example shows when post-screening may be effective in incentivizing the agent to exert higher effort as compared to the no-insurance scenario.

Consider a scenario where the nominal probability of attack $t = 1$. Figure 1 illustrates the agent's effort in the first period as a function of loss l. We note that post-screening can be an effective mechanism to incentivize the agents to exert non-zero effort inside a contract with full coverage. In this example, the agent exerts higher effort as compared to the no insurance scenario when $l \leq 0.7$. This is because since the loss is relatively low, even without insurance the agent is not willing to exert substantial effort as the cost of effort is higher than the actual loss. Within a contract, the insurer is able to incentivize the agent to exert higher effort by imposing a large penalty (a much higher premium in the second period).

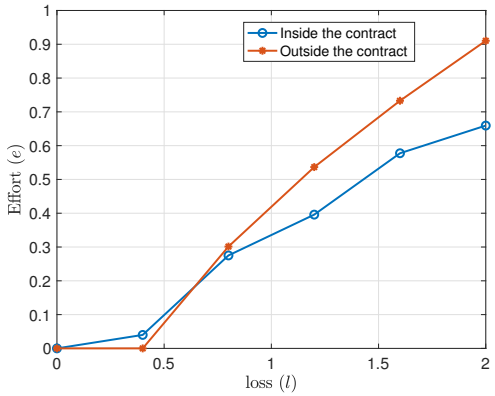

Fig. 1: Post-screening: agent's effort v.s. loss (l)

6.2 Rare Losses: Pre-screening

Our second example examines the effect of pre-screening on the agent's effort. Consider a scenario where t_a, t_p go to zero and l goes to infinity. Moreover, assume $l_a = 5$ and $\sigma = 0.1$. Figure 2 illustrates the agent's effort inside and outside the insurance contract with pre-screening. We see that the agent exerts non-zero effort inside the insurance contract and the effort increases as l_p increases. Note that outside a contract the agent's effort is a function of his perceived loss l_a and does not change with l_p. On the other hand, inside the contract, as the insurer's perceived loss l_p increases, the insurer incentivizes the agent to increase his efforts using premium discrimination (high premium for low pre-screening outcomes).

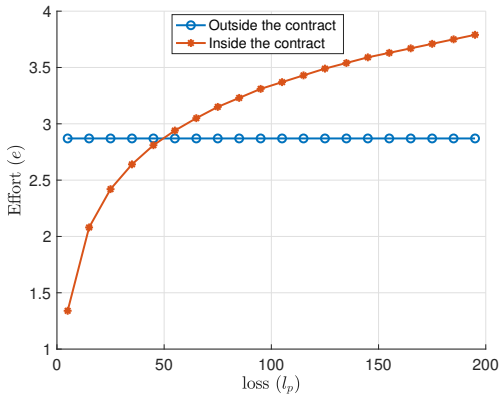

Fig. 2: Pre-screening: agent's effort v.s. loss (l_p)

Figure 3 illustrates the insurer's utility as a function of l_p. This figure implies that the insurer's utility is negative for $l_p \geq 85$. Therefore, she does not insure the agent if $l_p \geq 85$. Also, as expected, the insurer's utility decreases as the perceived expected loss l_p increases. The reason is that as the perceived expected loss increases, the insurer expects to pay more coverage and make less profit.

7 Conclusion

We studied the problem of designing cyber insurance contracts between a single profit-maximizing and risk-neutral insurer and a risk-averse agent. We showed that multi-period contract is an effective method of premium discrimination if loss incidents are frequent. We then considered rare but severe losses which is a common theme of cyber risks. In this case, we showed that multi-period contract is not effective in improving the agent's effort: the agent exerts zero effort inside a contract with full coverage. By contrast, pre-screening is shown to allow the

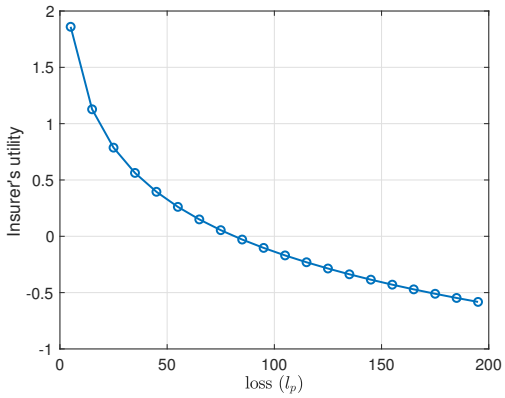

Fig. 3: Insurer's utility v.s. loss (l_p)

insurer to assess the agent's state of security and premium discriminates properly so as to incentivize effort by the agent within the contract.

We further discussed how the pre-screening result enables a type of active policy where periodic pre-screening within the same contract term can not only ensure the agent does not lower his effort after the initial assessment but also allows the insurer to improve her profit.

References

1. D. K. Tosh, I. Vakilinia, S. Shetty, S. Sengupta, C. A. Kamhoua, L. Njilla, and K. Kwiat, "Three layer game theoretic decision framework for cyber-investment and cyber-insurance," in *Decision and Game Theory for Security*, S. Rass, B. An, C. Kiekintveld, F. Fang, and S. Schauer, Eds. Cham: Springer International Publishing, 2017, pp. 519–532.
2. I. Vakilinia and S. Sengupta, "A coalitional cyber-insurance framework for a common platform," *IEEE Transactions on Information Forensics and Security*, vol. 14, no. 6, pp. 1526–1538, 2018.
3. M. Lelarge and J. Bolot, "Economic incentives to increase security in the internet: The case for insurance," in *Proceedings of IEEE INFOCOM*, 2009, pp. 1494–1502.
4. M. M. Khalili, P. Naghizadeh, and M. Liu, "Designing cyber insurance policies: The role of pre-screening and security interdependence," *IEEE Transactions on Information Forensics and Security*, vol. PP, no. 99, pp. 1–1, 2018.
5. N. Shetty, G. Schwartz, and J. Walrand, "Can competitive insurers improve network security?" in *International Conference on Trust and Trustworthy Computing*. Springer, 2010, pp. 308–322.
6. G. Schwartz, N. Shetty, and J. Walrand, "Cyber-insurance: Missing market driven by user heterogeneity," *www. eecs. berkeley. edu/nikhils/SecTypes. pdf*, 2010.
7. R. Zhang, Q. Zhu, and Y. Hayel, "A bi-level game approach to attack-aware cyber insurance of computer networks," *IEEE Journal on Selected Areas in Communications*, vol. 35, no. 3, pp. 779–794, March 2017.

8. A. Hofmann, "Internalizing externalities of loss prevention through insurance monopoly: an analysis of interdependent risks," *The Geneva Risk and Insurance Review*, vol. 32, no. 1, pp. 91–111, 2007.
9. Z. Yang and J. C. Lui, "Security adoption and influence of cyber-insurance markets in heterogeneous networks," *Performance Evaluation*, vol. 74, pp. 1–17, 2014.
10. M. M. Khalili, M. Liu, and S. Romanosky, "Embracing and controlling risk dependency in cyber insurance policy underwriting," in *The Annual Workshop on the Economics of Information Security (WEIS)*, 2018.
11. A. Rubinstein and M. E. Yaari, "Repeated insurance contracts and moral hazard," *Journal of Economic Theory*, vol. 30, no. 1, pp. 74 – 97, 1983. [Online]. Available: http://www.sciencedirect.com/science/article/pii/0022053183900947
12. P. Slovic, B. Fischhoff, S. Lichtenstein, B. Corrigan, and B. Combs, "Preference for insuring against probable small losses: Insurance implications," *The Journal of Risk and Insurance*, vol. 44, no. 2, pp. 237–258, 1977. [Online]. Available: http://www.jstor.org/stable/252136
13. P. A. Raschky and H. Weck-Hannemann, "Charity hazard-a real hazard to natural disaster insurance?" *Environmental Hazards*, vol. 7, no. 4, pp. 321 – 329, 2007. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S174778910700049X
14. J. Cox, "Equifax stung with multibillion-dollar class-action lawsuit after massive data breach," *http://www.thedailybeast.com/equifax-stung-with-multi-billion-dollar-class-action-lawsuit-after-massive-data-breach*, 2017.
15. L. Jiang, V. Anantharam, and J. Walrand, "How bad are selfish investments in network security?" *IEEE/ACM Transactions on Networking*, vol. 19, no. 2, pp. 549–560, 2010.
16. L. A. Gordon and M. P. Loeb, "The economics of information security investment," *ACM Trans. Inf. Syst. Secur.*, vol. 5, no. 4, pp. 438–457, Nov. 2002. [Online]. Available: http://doi.acm.org/10.1145/581271.581274
17. 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 *Proceedings of the 24th USENIX Security Symposium*, 2015.

Appendix

Proof (Lemma 1). Proof by contradiction. Let $(\hat{\pi}_1, \hat{\pi}_2, \hat{\pi}_3)$ be the solution of optimization problem (8), and assume that the (IR) constraint is not binding at the optimal contract $(\hat{\pi}_1, \hat{\pi}_2, \hat{\pi}_3)$. Because the (IR) constraint is not binding, the insurer can increase her utility by increasing $\hat{\pi}_2, \hat{\pi}_3$ while she keeps $\exp\{\gamma\hat{\pi}_2\} - \exp\{\gamma\hat{\pi}_3\}$ fixed. Therefore, based on (9) the agent's effort inside the contract does not change, but the insurer's profit increases. As a result, $(\hat{\pi}_1, \hat{\pi}_2, \hat{\pi}_3)$ is not an optimal contract. This is the contradiction implying that the (IR) constraint is binding.

■

Proof (Theorem 1). Proof by contradiction: Assume that $\hat{e} = 0$ and $t = 1$ and $\left\lceil \frac{(\alpha-\gamma c)(\exp\{\gamma l\}-1)}{\gamma c} \right\rceil > 1$. First we show that under these assumptions, $\hat{\pi}_1 = \hat{\pi}_2 = \frac{1}{\gamma}\ln(1-u^o) := w^o$. Because $\hat{e} = 0$ and $t = 1$, the optimization problem for finding $(\hat{\pi}_1, \hat{\pi}_2, \hat{\pi}_3)$ is as follows,

$$\begin{array}{l} \max_{\{\pi_1,\pi_2,\pi_3\}} \pi_1 + \pi_2 - 2l \\ s.t., \\ (IR)\ 1 - \exp\{\gamma\pi_1\} + 1 - \exp\{\gamma\pi_2\} = 2u^o \\ (IC)\ 0 = e^{in}(\pi_1, \pi_2, \pi_3) \end{array} \tag{24}$$

By (IR) constraint we have,

$$\frac{1}{\gamma}\ln(2 - 2u^o - \exp\{\gamma\pi_1\}) = \pi_2 \tag{25}$$

Therefore, we re-write the optimization problem (24) as follows,

$$\begin{array}{l} \max_{\{\pi_1,\pi_2,\pi_3\}} \pi_1 + \frac{1}{\gamma}\ln(2 - 2u^o - \exp\{\gamma\pi_1\}) - 2l \\ s.t., \\ (IC)\ 0 = e^{in}(\pi_1, \pi_2, \pi_3) \\ \qquad \frac{1}{\gamma}\ln(2 - 2u^o - \exp\{\gamma\pi_1\}) = \pi_2 \end{array} \tag{26}$$

Because π_3 does not appear in the objective function, we first find π_1 and π_2 such that they maximize the objective function. Then, we pick π_3 such that (IC) constraint is satisfied. By the first order optimality condition for the objective function, we have,

$$\hat{\pi}_1 = \hat{\pi}_2 = \frac{1}{\gamma}\ln(1 - u^o) \tag{27}$$

Without loss of generality, we set $\hat{\pi}_3 = \frac{1}{\gamma}\ln(\frac{\alpha-\gamma c}{\alpha}(1-u^o))$. By (9), $\hat{e} = 0$ (Notice that $\frac{\alpha}{\gamma c}\frac{\exp\{\gamma\hat{\pi}_2\}-\exp\{\gamma\hat{\pi}_3\}}{\exp\{\gamma\hat{\pi}_1\}} = 1$ and a slight decrease in $\hat{\pi}_3$, increases the agent's effort based on (9)).

Now we show that the decrease in $\hat{\pi}_3$ increases the insurer's payoff. Notice that a slight decrease in $\hat{\pi}_3$, increases the agent's effort (based on (9)) and improves agents' utility and the (IR) constraint is not violated. We write the insurer's objective function as a function of π_3. Therefore, we have (derivatives in the following equation are left derivatives),

$$\begin{aligned} h(\pi_3) &= \hat{\pi}_1 - p(e^{in}(\hat{\pi}_1, \hat{\pi}_2, \pi_3))(l - \hat{\pi}_2) + (1 - p(e^{in}(\hat{\pi}_1, \hat{\pi}_2, \pi_3)))\pi_3 - l \\ \frac{\partial h}{\partial \pi_3}|_{\pi_3=\hat{\pi}_3} &= \frac{\partial p(e^{in}(\hat{\pi}_1, \hat{\pi}_2, \pi_3))}{\partial \pi_3} \cdot (\hat{\pi}_2 - l) \\ &\quad - \frac{\partial p(e^{in}(\hat{\pi}_1, \hat{\pi}_2, \pi_3))}{\partial \pi_3} \cdot \pi_3 + (1 - p(e^{in}(\hat{\pi}_1, \hat{\pi}_2, \pi_3))) \\ &= \left(\frac{\partial p(e^{in}(\hat{\pi}_1, \hat{\pi}_2, \pi_3))}{\partial \pi_3}|_{\pi_3=\hat{\pi}_3} \cdot (-l + \hat{\pi}_2 - \hat{\pi}_3) - (1 - p(e_1^{in}(\hat{\pi}_1, \hat{\pi}_2, \hat{\pi}_3))) \right) \end{aligned}$$

Because $\left[\frac{(\alpha-\gamma c)(\exp\{\gamma l\}-1)}{\gamma c}\right] > 1$, (5) implies that e^o is not zero and $\hat{\pi}_2 = \frac{1}{\gamma}\ln(1-u^o) < l$. Moreover, $\frac{\partial p(e^{in}(\hat{\pi}_1,\hat{\pi}_2,\pi_3))}{\partial \pi_3}|_{\pi_3=\hat{\pi}_3} > 0$ implies that $\frac{\partial h}{\partial \pi_3}|_{\pi_3=\hat{\pi}_3} < 0$. Therefore, the decrease in $\hat{\pi}_3$ increases the insurer's payoff. This is a contradiction and the agent exerts non-zero effort in the optimal contract under given assumptions.

■

Proof (Theorem 2).

By (14), the agent exerts non-zero effort in a contract if $\beta = c$. If the discount factor $\beta = c$, then any positive number satisfies the (IC) constraint. Therefore, if $\beta = c$, then the desired effort maximizes the insurer's utility. By (14), we have,

$$\overline{e} = \arg\max_{e} w^o - ce - tl \exp\{-\alpha \cdot e\} - \gamma c^2 \sigma^2 \tag{28}$$

By the first order condition of optimality, the solution of above optimization problem is $\overline{e} = (\frac{1}{\alpha} \ln(\frac{\alpha \cdot t \cdot l}{c}))^+$. Moreover, if $\overline{e} > 0$, then the maximum insurer's profit using pre-screening (i.e., $\beta = c$) is given by,

$$\left\{ w^o - \frac{c}{\alpha} \ln(\frac{\alpha t l}{c}) - \frac{c}{\alpha} - \frac{\gamma c^2 \sigma^2}{2} \right\} \tag{29}$$

Without pre-screening (i.e., $\beta = 0$), the agent exerts zero effort and the insurer's profit is given by,

$$w^o - t \cdot l \tag{30}$$

Therefore, the insurer uses pre-screening if and only if,

$$\begin{array}{l} \frac{1}{\alpha} \ln(\frac{\alpha \cdot t \cdot l}{c}) > 0 \\ w^o - \frac{c}{\alpha} \ln(\frac{\alpha t l}{c}) - \frac{c}{\alpha} - \frac{\gamma c^2 \sigma^2}{2} \geq w^o - tl \end{array} \tag{31}$$

In other words, the insurer uses pre-screening and the agent exerts non-zero effort if and only if,

$$\begin{array}{l} \dfrac{\alpha \cdot t \cdot l}{c} > 1 \\ \sigma^2 \leq \dfrac{2}{\gamma c^2}(tl - \dfrac{c}{\alpha}(1 + \ln(\dfrac{\alpha t l}{c})) \end{array} \tag{32}$$

■

Proof (Theorem 3).

Assume $\sigma < \sigma'$.

Let $g(\beta, e, \sigma) = \left[w^o - ce - \frac{\gamma \beta^2 \sigma^2}{2} - p(e)l \right]$. It is easy to see that $g(\beta, e, \sigma') \leq g(\beta, e, \sigma)$. Therefore, we have,

$$\max_{\beta, e, IC\ constraint} g(\beta, e, \sigma') \leq \max_{\beta, e, IC\ constraint} g(\beta, e, \sigma)$$

Therefore, $V(\sigma') \leq V(\sigma)$.

■

Proof (Theorem 4).

– By (9), the agent exerts zero effort if $t_a \frac{\alpha}{\gamma c} \frac{\exp\{\gamma\pi_2\}-\exp\{\gamma\pi_3\}}{\exp\{\gamma\pi_1\}} \leq 1$. Because t_a goes to zero, $t_a \frac{\alpha}{\gamma c} \frac{\exp\{\gamma\pi_2\}-\exp\{\gamma\pi_3\}}{\exp\{\gamma\pi_1\}}$ also goes to zero. Therefore, the agent exerts zero effort under any insurance contract.
– Because the agent exerts zero effort inside the optimal contract, his utility is given by,

$$\begin{aligned} &U^{in}(0,\pi_1,\pi_2,\pi_3) = -\exp\{\gamma\pi\} - t_a \exp\{\gamma\pi_2\} - (1-t_a)\exp\{\gamma\pi_3\} \\ &\text{(IR) is binding and } t_a \to 0 \Rightarrow\ 1-\exp\{\gamma\pi_1\} + 1 - \exp\{\gamma\pi_3\} = 2u^o \end{aligned} \tag{33}$$

Therefore, the insurer's problem (8) can be written as follows,

$$\begin{aligned} &\max_{\pi_1,\pi_2,\pi_3} \pi_1 + \pi_3 - 2\cdot l_p \\ &s.t., \exp\{\gamma\pi_1\} + \exp\{\gamma\pi_3\} = 2 - 2u^o \end{aligned} \tag{34}$$

or

$$\max_{\pi_1} \pi_1 + \tfrac{1}{\gamma}\ln(2-2u^o-\exp\{\gamma\pi_1\}) - 2l_p \tag{35}$$

The optimal solution for the above optimization problem is $\pi_1 = \pi_3 = \frac{1}{\gamma}\ln(1-u^o)$ and also the value of π_2 does not affect insurer's or agent's utility and can be any positive value.

■

Proof (Theorem 5). The proof is similar to the proof of theorem 2 except that we should substitute l_p for $t \cdot l$.

■

Proof (Theorem 6). As the (IR) constraint is binding in (23), similar to (14) we can re-write optimization problem (23) as follows,

$$\begin{aligned} &R(\sigma) = \max_{\{\beta,e,\beta',e'\}} \left[w^o - ce + b(e-e') - \gamma\tfrac{(\beta-\beta')^2\sigma^2+(\beta')^2\sigma^2}{2} - p(e')l\right] \\ &s.t., (IC)(e,e') \in \arg\min_{(\tilde{e}\geq\tilde{e}')} \gamma(c-b+\beta'-\beta)\tilde{e} + \gamma(-\beta'+b)\tilde{e}' \end{aligned} \tag{36}$$

First we show that $\hat{e} = \hat{e}'$. Proof by contradiction. Assume $\hat{e} > \hat{e}' \geq 0$. Then, $\beta' - \beta = b - c$ since otherwise $\hat{e} = \infty$ or $\hat{e} = 0$. As $b \leq c$, then the objective function of (36) can be improved by decreasing $\hat{e}$ without violating (IC) constraint. This contradiction shows that $\hat{e} = \hat{e}'$.

By IC constraint, it is easy to see that if $\hat{e} = \hat{e}' > 0$, then $\beta = c$, and $\beta' = b$.

Let $\beta = \bar{\beta}, e = \bar{e}$ be the solution to (14). According to the IC constraint of (14), two cases can happen:

i) $\overline{\beta} = 0$ and $\overline{e} = 0$. Then, $(\beta = \beta' = e = e' = 0)$ satisfies the IC constraint in (36) and is a feasible point. We have,

$$\begin{aligned} &w^o - c\overline{e} - \frac{\gamma\overline{\beta}^2\sigma^2}{2} - p(\overline{e})l = \\ &w^o - c\overline{e} + b(\overline{e}-\overline{e}) - \gamma\frac{(\overline{\beta}-\overline{\beta}')^2 + (\overline{\beta}')^2}{2}\sigma^2 - p(\overline{e})l \end{aligned} \tag{37}$$

ii) $\overline{\beta} = c$. Then $(\beta = c, \beta' = b, e = e' = \overline{e})$ is a feasible point for (36) and satisfies the IC constraint. We have,

$$w^o - c\overline{e} - \frac{\gamma c^2 \sigma^2}{2} - p(\overline{e})l \leq$$
$$w^o - c \cdot \overline{e} + b(\overline{e} - \overline{e}) - \gamma \frac{(c-b)^2 + b^2}{2}\sigma^2 - p(\overline{e})l \tag{38}$$

Note that in this case $(\beta = c, \beta' = b, e = e' = \overline{e})$ is the solution to (36).

By (37) and (38) we have, $V(\sigma) \leq R(\sigma)$. Notice that if $b = c$, then (36) and (14) are equivalent and $V(\sigma) = R(\sigma)$ as $\hat{e} = \hat{e}'$.

■