Enterprise Information Security: Who should mange it and how?

Vineet Kumar, Rahul Telang* and Tridas Mukhopadhyay
{vineetk, rtelang, tridas}@andrew.cmu.edu
Carnegie Mellon University, Pittsburgh, PA 15213

1 Introduction

In the recent past, Information Security has received a lot of attention in the business and trade press. Much of this can be attributed to an increase in security breaches leading to major losses to the affected enterprises.

The CSI/FBI 2005 [4] security survey reports 13 different attack types, ranging from web-site defacement to financial fraud to Internet worms and viruses. There are several reports in the business press that point to the increasing number of vulnerabilities in commonly used software as well as viruses and other threats that seek to exploit these vulnerabilities, and detail how it is becoming an increasing problem for enterprises. Effective countermeasures sometimes exist for many of these threats, but are often not correctly deployed due to the specific characteristics of the information systems in use, or the capabilities of the IT staff.

The economic analysis of information security has many dimensions to it as evidenced by the literature including risk management approaches, insurance, vulnerability analysis, information sharing etc. However, the role of decision making within the enterprise and the related issues of incentives and information asymmetry within a firm has not received much attention in the context of information security. It is our objective to specifically address how a multi-division should make optimal information security deployment decisions in the light of the above factors.

Much attention has been focused on detailing the operation of countermeasures (e.g. firewalls that protect against unauthorized traffic) but little attention is focused on who in the enterprise is making decisions regarding deployment of these measures and what policies are in place to deal with such decision-making.

When discussing Enterprise security, it is important to understand that enterprises are not homogeneous entities and their divisions often use varied information systems, which are commonly interconnected with each other as well as to the Internet. In addition, some divisions may not value their information assets as highly as others or may not be as capable of expertly deploying security

*Rahul Telang acknowledges generous support of National Science Foundation (NSF) through the NSF CAREER award CNS-0546009

measures. And, depending on the network architecture, it could take just one insecure network entity to put the entire network at risk. These factors imply that the allocation of decision rights regarding security among the divisions is particularly important in an interconnected enterprise because the security of each division could affect the overall expected loss of each of the other divisions.

The notion of information security and firm investments has been studied by [3] who derive an expression for optimal investments. However, intra-firm interdependencies was not their focus. The notion of interdependent security, where the security of each entity depends on the decisions of the other entities in the system has been explored by Kunreuther and Heal [6] in the context of airline security. The game is between airlines, who must decide whether to invest in security measures. This work considers a binary decision space, and concludes that there can be equilibrium where everyone invests in protection measures. However, the equilibrium in unstable in the sense that a single airline defecting would cause it to collapse. In addition, the probability of attack is exogenously specified.

The key goal of our paper is to understand the decision making within organization, when firm has multiple divisions. The incentives driving each division may diverge from the firm's objective, which would lead to the conclusion that the divisions being allocated decision rights would likely lead to sub-optimal deployment of security countermeasures. At the same time, a central IT team, which is commonly led by the CIO (and which we denote henceforth by the same abbreviation) need not necessarily understand the capabilities of each division's IT systems and staff. The CIO might attempt to set uniform security standards, which again could result in sub-optimal deployment of countermeasures. Indeed, there is much debate on who should be in charge of security in firms, with divergent accounts available at [7], [1] and [11]. This is a question we seek to answer in this paper, besides determining the optimal ways to invest in security countermeasures. In the context of this paper, an optimal deployment refers to security deployments which result in minimizing the expected losses for the enterprise.

The principal research question we seek to answer is: How should a multi-divisional enterprise optimally deploy security countermeasures in response to heterogeneous information systems, distinct threat situations, and different kind of damages that threats can cause to the enterprise's information systems and assets? The current paper is differentiated from previous work along the following dimensions:

- Different decision rights and incentive structures (Minimum Security Standards, Subsidies)
- Different kinds of countermeasures (divisional vs. perimeter security)
- Multiple threat types (contagious and independent).

We model the strategic interaction between divisions of the enterprise, while at the same time opening up the technology black-box sufficiently to evaluate the effectiveness of different technologies and decision structures. We draw from the following streams of research. Our stylized model of contagious threats is based on the research in computer science that details the propagation of viruses and Internet worms ([5], [9] and [12]). Our discussion of decision structures is inspired by Nadiminti et al [10] and Mendelson's work [8] which both look at decision structures in an Information Systems context, albeit with a different application in mind.

We derive several analytical results in this paper regarding security deployment by the divisions, subsidies, minimum security standards, and perimeter security. First, we find that the CIO should subsidize the divisions so that they deploy security measures optimally for contagious threats. Moreover, the subsidies are independent of the capabilities of the divisions. Second, imposing minimum security standards on the divisions results in sub-optimal deployment and could even increase the expected losses faced by the firm. Third, perimeter security plays an important role in enterprise security, and can even replace the security measures that divisions would ordinarily deploy, provided its cost is low enough. Otherwise, we find that there is an optimal level for deployment of perimeter security by the CIO.

This is still a work-in-progress and we expect more results by the time workshop commences.

2 Models

In this section, we describe in detail the parameters of our theoretical framework. We consider a firm, with N divisions, which seeks to implement security countermeasures to minimize its expected loss.

- **Threat source** launches attacks that might cause damage to a system. These could include contagious threats like viruses or worms, phishing attacks, denial of service attacks(DoS), as well as hackers or insiders attempting to steal confidential information. We model three separate threat classes: independent, contagious and strategic. Each threat class is described and dealt with separately in the following sections.

  We attempt to derive optimal deployment of countermeasures for each of the threat classes. A threat source launches $T \sim Poisson(\lambda)$ attacks per time period implying that the *average* number of attacks from the threat source is λ. Some of these attacks will be successful, with the probability of success depending on the security countermeasures deployed, as detailed below.

- **Capability** The capability of a division reflects the characteristics of its information systems. Highly capable divisions will have a lower vulnerability than divisions that have a lower capability when both have the same level of security measures deployed. This could also be thought of as reflecting the heterogeneity among the divisions in effectively implementing the security countermeasures.

  $\alpha_i \in [\underline{\alpha}, \overline{\alpha}]$ denotes the security capability of a division i. Smaller values of α_i reflect divisions that have a better IT capability. The capabilities of the divisions are assumed to be drawn from a multivariate distribution. We denote $\gamma = E[\sum_{i=1} Nlog(\alpha_i)]$.

- **Vulnerability** is a measure of weakness in systems and policies with respect to a particular threat that may be exploited by an attack to cause damage.

  In our models, vulnerability is the probability of a successful breach (upon an attack) of division i's countermeasures which are set at level s. $p_i(s) = \alpha_i \cdot e^{-s}$ is the *external vulnerability*, or the probability of breach to attacks originating external to the enterprise network. When

the attack originates from within the network, the vulnerability is $\eta p_i(s_i)$, where $1 \leq \eta < \frac{1}{\alpha}$, and this is the *internal vulnerability*. This reflects the situation that internal and external traffic are treated often treated differently by enterprise networks, and information systems are *more* vulnerable to internal attacks. We call η the *internal vulnerability factor*.

Division i is vulnerable to level α_i when it has no security countermeasures deployed ($p_i(0) = \alpha_i$). Divisions can attain zero vulnerability only when they have an infinite level of security countermeasures deployed. This reflects the fact that perfect security is impossible to attain. Increasing the level of countermeasures deployed decreases vulnerability, which is also convex with respect to it. This implies decreasing marginal benefits of increased security countermeasures deployed.

- **Countermeasures** (also referred to as security measures): These include security software and hardware products as well as policies that are intended to minimize the exposure to threats that would result in a successful attack. In most cases, this is done by mitigating the vulnerability of systems or policies. Examples would include anti virus software, intrusion detection systems, firewalls, as well as policies to prevent unauthorized access to systems. Each countermeasure that an organization deploys has costs in monetary terms, whether it involves purchasing products or results from productivity losses caused by restrictive policies. The economic case to deploy countermeasures is due to lower expected losses. Each division's security deployment can be set independently of the other divisions. s_i denotes the level of security countermeasures deployed by division i. We assume that s_i is continuous with a support of $[0, \infty)$.

- **Loss**: This represents the monetary loss caused to the firm when it is subject to an attack that successfully exploits a vulnerability in its systems and/or policies. We consider three different kinds of losses to the *information systems and assets* of an enterprise. In addition there are often personnel costs to restore the IT system after a successful attack and loss due to disruption of business etc.

  L denotes, in dollar terms, the loss suffered by each division when a successful attack takes place.

- **Costs of security countermeasures**

  There are monetary costs of implementing both policy-based and product-based security countermeasures. In addition, firms need to consider the impact of such measures on end users. To illustrate an extreme example, one could have a low vulnerability and expected losses by disconnecting all information systems from the Internet, but this could lead to a large productivity loss. $C(s) = c \cdot s$ denotes the cost to a division of deploying countermeasures of a level s. We assume a linear cost structure here with the assumption being that a higher security level is achieved by stacking together security countermeasures.

We compare different decision structures, where the divisions or the CIO can deploy countermeasures under varied threat conditions as well as different countermeasure availability, and show how the firm must optimally deploy countermeasures and vest decision rights as the result of this

analysis. Optimality in our model implies security deployments and allocation of decision rights and countermeasure deployment that minimize expected losses for the firm.

3 Independent Threats

Each division of the multi-division firm independently faces a threat source, and is attacked $T \sim Poisson(\lambda)$ times in each period. Each attack will result in a successful breach with a probability $p_i(s_i)$ for division i. Then, the expected loss faced by a division i is given by $\tilde{L_i}$ where:

$$\tilde{L_i} = \left(\sum_{j=1}^{T} p_i(s_i) \right) L + c \cdot s_i$$

The firm's expected loss is simply the sum of the expected loss of each division:

$$E[\tilde{L}] + \sum_{i=1}^{N} c \cdot s_i = (\lambda L) \sum_{i=1}^{N} (p_i(s_i) + c \cdot s_i)$$

The divisions are assumed to know their capabilities exactly. Different informational situations allow for the CIO to either have full knowledge of the divisions' capabilities or partial knowledge, that the capabilities are drawn from a particular distribution.

3.1 Divisions with decision rights

Here we have the divisions deciding on the level of security countermeasures to deploy. Each division acts to minimize its total expected loss, which includes the cost of security countermeasures:

$$\min \qquad E[\tilde{L_i}] + c \cdot s_i = (\lambda L) p_i(s_i) + c \cdot s_i s_i$$

This gives the following FOCs:

$$\frac{1}{\lambda L} \frac{\partial E[\tilde{L_i}]}{\partial s_i} \equiv \frac{\partial p_i(s_i)}{\partial s_i} + \frac{c}{\lambda L} = 0 \ \forall \ i$$

3.2 CIO with decision rights

In this case the CIO is assumed to know the capability of each division. Then, the CIO's decision problem is to minimize the expected loss for the firm:

$$\min \qquad E[\tilde{L}] + \sum_{i=1}^{N} c \cdot s_i = (\lambda L) \sum_{i=1}^{N} (p_i(s_i) + c \cdot s_i)$$

$$s_1, \ldots, s_N$$

which gives the following FOCs:

$$\frac{1}{\lambda L} \frac{\partial E[\tilde{L}]}{\partial s_i} \equiv \frac{\partial p_i(s_i)}{\partial s_i} + \frac{c}{\lambda L} = 0 \ \forall \ i$$

Proposition 1. . *In the case of independent threats, the objective divisions and the CIO are perfectly aligned, and each would choose a division's security level to be:*

$$s_i = log\left(\frac{\alpha_i \lambda L}{c}\right)$$

We can see that the FOCs in both cases are the same. Therefore, the divisions have their interests aligned perfectly with the CIO in the case of independent threats. In case of any informational asymmetry between the CIO and the divisions where the CIO does not have full knowledge of the divisions capabilities, we note that the decision making is best left to the divisions. This is because the CIO will make suboptimal decisions due to lack of complete information about the divisions' capabilities, and this will cause the expected loss will be greater when the CIO makes the decisions.

4 Contagious Threats

In this section, we describe the idea of a contagious threat, and how it is modeled in our paper. The motivation for modeling this class of threat sources is derived from computer viruses and Internet worms, both of which are described as dangerous threats that enterprises spend a lot of resources protecting against [4]. The reader may like to consider these real-world examples in the discussion that follows. A good description of the detailed operating mechanisms of viruses and worms is given in [12].

We consider a specific model for contagious threats although there are many modes that threat sources may propagate by. Moreover, their varied interaction with countermeasures leads to a myriad number of ways by which their effects may be felt. Our objective is to model the threat behavior in enough detail to ensure derivation of accurate and representative results, while maintaining analytical tractability by electing to leave out esoteric capabilities (e.g. polymorphism, see [12] for an explanation) that some real-world threat sources may possess.

A contagious threat source launches T attacks per time period, similar to the case of independent threats. The specific characteristic of a contagious threat is that an attack, if successful, can lead to further propagation of the threat source within a network and cause losses before being stamped out, so to speak. Internet worms are known to attack and bring down hosts repeatedly, making them unavailable, even if such hosts are cleaned of such attacks, until they are no longer vulnerable to the attack (for example, by patching).

The threat source is initially external to the enterprise network, and attacks each division in it separately. The divisions are potentially vulnerable to this attack, with the probability of a successful breach, $p_i(s) = \alpha_i e^{-s}$ depending on the level of countermeasures deployed by the division. Now, consider that some divisions get breached (and consequently infected) by this threat source, which becomes an internal threat to the network. Each infected division has an internal threat source that is capable of further attacks and breaches leading to losses, with a breach probability $\eta \cdot p_i(s)$ for division i. Each incidence of breach costs a division L. Each division can suffer losses multiple times before it is patched. Patching is assumed to occur after a threat source

has successfully breached a division after an internal attack. The act of patching is exogenous to the model and we do not specifically model its effects.

We model the threats formally as follows:

$T \sim Poisson(\lambda)$ is the number of external threats faced by each of the firm's divisions

$T_k = \sum_{l=1}^{T} p_k(s_k)$ represents the number of successful breaches faced by each division

Each successful breach leads to an infection, and to the threat source propagating to the enterprise network, where it becomes an internal threat source. Consider such an internal threat source in division k. It will attack and potentially breach the other divisions'. Consider that another division j, is breached by this threat source, now present in division k. This happens with a probability p_j. Again, this would be described by a distribution $T_{j,k}$ which describes *successful* attacks on division j, resulting from a threat source present at division k.

$$T_{j,k} \sim Poisson(\lambda \eta p_k(s_k) p_j(s_j))$$

The loss to each division k is given by:

$$\tilde{L_k} = (\lambda) L \left[\sum_{l=1}^{T} p_k(s_k) + \sum_{l=1}^{T} \sum_{j \neq i} p_i(s_i) p_k(s_k) \right]$$

The loss to the firm is given by:

$$\tilde{L} = \sum_{i=1}^{N} \tilde{L}_i = L(\sum_{i=1}^{N} T_i + 2 \sum_{i=1}^{N} \sum_{j>i} T_{i,j}) + c \sum_{i=1}^{N} s_i$$

The expected loss to the firm is given as:

$$E[\tilde{L}] = \lambda \cdot L(\sum_{i=1}^{N} p_i + 2 \sum_{i=1}^{N} \sum_{j>i} \eta p_i \cdot p_j) + c \sum_{i=1}^{N} s_i$$

4.1 CIO with decision rights

The CIO's decision problem is to minimize the expected loss to the firm, given by:

$$\underset{s_1, \ldots, s_N}{min} \; E(\tilde{L}) = (\lambda L) \; E \left[(\sum_{i=1}^{N} p_i(s_i) + 2 \sum_{i=1}^{N} \sum_{j>i} \eta p_i(s_i) \cdot p_j(s_j)) \right] + c \sum_{i=1}^{N} s_i$$

This leads us to the first order conditions (FOCs) which are:

$$\frac{1}{\lambda L} \frac{\partial \tilde{L}}{\partial s_i} \equiv \frac{\partial p_i(s_i)}{\partial s_i} + 2\eta \frac{\partial p_i(s_i)}{\partial s_i} \sum_{j \neq i} p_j + \frac{c}{\lambda L} = 0 \;\; \forall \{i, j\}$$

We consider two different cases: one in which the CIO has full knowledge of the divisions' capabilities and another in which it only knows the distribution that the capabilities (which are random variables) are drawn from. In each case, we compute the expected loss, and compare the two cases.

4.1.1 CIO has complete information

Here, we evaluate the effectiveness of the CIO attempting to minimize the expected loss faced by the overall firm, solving the problem with perfect information. The solution to the FOCs given above gives us the optimal spending on security countermeasures by the divisions.

Proposition 2. . *Given that the CIO has complete information on the capabilities of the divisions, he will choose to set the security level of each division so that the security level deployed by division i is set to:*

$$s_i^* = log\left(\frac{4\alpha_i\eta(N-1)}{\sqrt{1+8\eta(N-1)\frac{c}{\lambda L}}-1}\right)$$

The expected loss is bounded as follows:

$$E[\tilde{L}^{CI}] = (cN)\left[\gamma + \frac{1}{2} + log\left(\frac{4\eta(N-1)}{\sqrt{1+8\eta(N-1)\frac{c}{\lambda L}}-1}\right)\right] + \frac{N\lambda L}{8\eta(N-1)}\left[\sqrt{1+\frac{8c\eta(N-1)}{\lambda L}}-1\right]$$

We note that the security level deployed increases in threat level as well as firm size. As the threat level becomes very large, the security level deployed tends to the same level as for the independent threats.

4.1.2 CIO has incomplete information

Often, the firm does not fully know the operational details of each of its divisions. In this section, we model that the firm has asymmetric information regarding its divisions, while each division is fully informed of the characteristics of its own IT systems.

In this case, we model the capabilities of the divisions to be uniformly distributed and correlated as given below:

$$\alpha_i \sim U[\mu-\delta, \mu+\delta] \,\forall\, i \text{ and } cov(\alpha_i, \alpha_j) = \rho\sigma^2 = \rho\frac{\delta^2}{3}$$

The CIO can only enforce the same level of security countermeasures throughout the firm and his problem is to find the single optimal level of security countermeasures to be deployed by all divisions. We observe that:

$$E[\alpha_i] = \mu, \forall i \text{ and } E[\alpha_i\alpha_j] = \mu^2 + \rho\sigma^2, \forall j \neq i$$

Given this, the FOCs reduce to minimizing the expected loss given below with respect to s:

$$E(\tilde{L}) = \lambda LN\left(\mu e^{-s} + \eta(N-1)e^{-2s}(\mu^2+\rho\sigma^2)\right) + Nc\cdot s$$

Solving this gives us the optimal security level to be deployed by all divisions.

Proposition 3. . *When the CIO knows that the capabilities of the divisions are distributed as given above, but cannot observe each division's capability, he chooses a common security standard for each division given by:*

$$s^* = log\left(\frac{4\eta(N-1)(\mu^2+\rho\sigma^2)}{\sqrt{\mu^2+8\frac{c}{\lambda L}\eta(N-1)(\mu^2+\rho\sigma^2)}-\mu}\right)$$

The expected loss in higher than in the case where the CIO has complete information and is given by:

$$E[\tilde{L}^{AI}] = Nc\left[\frac{1}{2}+log\left(\frac{4\eta(N-1)(\mu^2+\rho\sigma^2)}{\sqrt{\mu^2+\frac{8c\eta(N-1)}{\lambda L}(\mu^2+\rho\sigma^2)}-\mu}\right)\right] + \frac{\lambda LN}{8\eta(N-1)}\left[\frac{\sqrt{1+\frac{8c}{\lambda L}\eta(N-1)\left(1+\frac{\rho\sigma^2}{\mu^2}\right)}-1}{1+\frac{\rho\sigma^2}{\mu^2}}\right]$$

We observe that as the variation of capabilities of the divisions increases, even when the mean capability is the same, the CIO will choose to invest more in security measures for each division, and the expected losses for the firm are also higher.

If the CIO makes security deployment decisions, the expected loss to the firm is driven by the fact that a uniform security level has to be deployed by each division, and it cannot be tailored to the capabilities of the divisions. When compared to the case of the CIO having complete information, this factor is the cause of inefficiency leading to increased expected losses.

$$\begin{aligned}\Delta L^{AI} =& E[\tilde{L}^{AI}] - E[\tilde{L}^{CI}]\\ \geq& (Nc)log\left(\left(1+\frac{\rho\sigma^2}{\mu^2}\right)\frac{\sqrt{1+\frac{8c\eta(N-1)}{\lambda L}}-1}{\sqrt{1+\frac{8c\eta(N-1)}{\lambda L}\left(1+\frac{\rho\sigma^2}{\mu^2}\right)}-1}\right)\\ &+\frac{\lambda LN}{8\eta(N-1)}\left[\frac{\sqrt{1+\frac{8c}{\lambda L}\eta(N-1)\left(1+\frac{\rho\sigma^2}{\mu^2}\right)}-1}{1+\frac{\rho\sigma^2}{\mu^2}}-\sqrt{1+\frac{8c\eta(N-1)}{\lambda L}}+1\right]\end{aligned}$$

4.2 Divisions with decision rights

In this setting, each division makes its own decisions regarding the provision of investment in security countermeasures. Divisions know what their capabilities are, unlike the CIO who only has an imperfect idea. However, divisions do not know the capabilities of other divisions, whose decisions on countermeasure deployment exert an externality on them.

We model this as a game with incomplete information played between the divisions and the CIO does not have any role to play in it. Each division knows the distribution of the other divisions' capabilities, but does not know the actual realization of it.

The game is as follows: Each division observes its own capability perfectly. Divisions then make decision on the level of security countermeasure deployment. The concept of Bayesian Nash Equilibrium, gives us that in equilibrium, no division will find it optimal to deviate from its beliefs and strategies (which are consistent with each other), given that the other players do not deviate from their beliefs and strategies [2].

Proposition 4. . *When each division has to choose a security level, the equilibrium strategies for the divisions are given by:*

$$s_i^* = log\left(\frac{2\alpha_i\eta(N-1)}{\sqrt{1+\frac{4c\eta(N-1)}{\lambda L}}-1}\right)$$

From the firm's point of view, this is strictly lower than the optimal amount of security to be deployed by each division. This result holds whether each division knows only its own capability or fully knows the capabilities of all the other divisions as well. The expected loss for the firm is given by:

$$L^U = (cN)\left[1+\gamma+log\left(\frac{2\eta(N-1)}{\sqrt{1+\frac{4c\eta(N-1)}{\lambda L}}-1}\right)\right]$$

Each division, while deciding to deploy security countermeasures, will consider the effect of other division's decisions because of the network effect of security. Consider a division k that faces a decision: From the division's point of view, if all the other divisions' levels are held constant, it's security level minimizes its expected loss accounting for the network effect. However, each of the other divisions is also affected by division k's level of countermeasures, and this is not directly taken into account by division k, which acts to minimize its own expected loss. Looking at this from the firm's perspective, we see that this is clearly suboptimal.

It is interesting to note that the equilibrium strategy of each division is not affected even when it additionally knows that the capabilities of the divisions are draws from a multivariate distribution with correlation.

The firm's total expected loss under these conditions is:

$$L^U = (cN)\left[1+log\left(\frac{2\eta(N-1)}{\sqrt{1+\frac{4c\eta(N-1)}{\lambda L}}-1}\right)\right] + c\cdot E\left[\sum_{i=i}^{N} log(\alpha_i)\right]$$

The increase in expected loss between the CIO with complete information having decision rights and the divisions deciding is given by:

$$\Delta L^U = E[\tilde{L}^U] - E[\tilde{L}^{CI}]$$
$$= (cN)\left[\frac{1}{2} + \gamma + log\left(\frac{\sqrt{1+\frac{8c\eta(N-1)}{\lambda L}}-1}{2\left(\sqrt{1+\frac{4c\eta(N-1)}{\lambda L}}-1\right)}\right)\right] - \frac{N\lambda L}{8\eta(N-1)}\left[\sqrt{1+\frac{8c\eta(N-1)}{\lambda L}}-1\right]$$

It is surprising to see that the divisions would not change their strategies even when they know the capabilities of the other divisions. This information is not useful to the divisions because, in equilibrium, its strategies and beliefs on what other divisions' strategies are will not change.

Comparison

Here, we briefly compare whether a firm whose CIO does not know the capabilities of the division should keep security decision rights or whether such rights should be transferred to the divisions.

$$\Delta L = E[\tilde{L}^{AI}] - E[\tilde{L}^U]$$
$$= -\frac{cN}{2} + (cN)log\left[\frac{2\left(1+\frac{\rho\sigma^2}{\mu^2}\right)\sqrt{1+\frac{4c\eta(N-1)}{\lambda L}}-1}{\sqrt{1+\frac{8c\eta(N-1)}{\lambda L}\left(1+\frac{\rho\sigma^2}{\mu^2}\right)}-1}\right] + \frac{\lambda LN}{8\eta(N-1)}\left[\frac{\sqrt{1+\frac{8c}{\lambda L}\eta(N-1)\left(1+\frac{\rho\sigma^2}{\mu^2}\right)}-1}{1+\frac{\rho\sigma^2}{\mu^2}}\right]$$

Corollary 1. *The firm chooses to vest decision rights with the divisions whenever the following inequality holds and vests it with a CIO otherwise:*

$$\frac{1}{2} + log\left(\frac{1}{2\beta}\frac{\sqrt{1+8\nu\beta}-1}{\sqrt{1+4\nu}-1}\right) < \frac{\sqrt{1+8\nu\beta}-1}{8\nu\beta}$$

where

$$\nu = \frac{c(N-1)}{\lambda L}; \quad \beta = 1 + \frac{\rho\sigma^2}{\mu^2}$$

The underlying intuition here is as follows: Whenever there is lower uncertainty or variance (smaller values of β), the firm is more likely to vest decision rights with the CIO rather than the divisions, because the losses due to asymmetric information are smaller than the losses due to goal divergence. This is also the case for lower values of ν or higher threat levels (λ), because at higher threat levels, the incentives of the CIO and the divisions are more aligned.

4.3 Minimum security standards

When the divisions play the game of simultaneously deciding what security measures to deploy, we have previously derived the Bayesian Nash Equilibrium. Here, we consider the addition of a

minimum security standard by the CIO, so that all the divisions deploy at least that level of security countermeasures.

Prior to the imposition of a minimum standard, let the divisions be arranged in increasing order of security measures deployed:

$$s^*_{i_1} \leq s^*_{i_2} \leq \ldots \leq s^*_{i_N}$$

Now, let $\hat{s}$ be the minimum security standard level imposed. We assume that the divisions actually have complete information on capabilities. Even under this extreme condition, we show that using minimum security requirements is not optimal. We evaluate the effect of $\hat{s}$ under three cases:

(1) $\hat{s} < min(s^*_i)$: In this case, the minimum security requirement has no effect, and the divisions deploy the same levels that they did in the equilibrium in proposition 2

(2) $\hat{s} > max(s^*_i)$: This condition acts like a *mandatory* security requirement, so that each of the divisions deploys exactly $\hat{s}$.

(3) $s^*_{i_1} \leq s^*_{i_2} \leq \ldots \leq s^*_{i_M} \leq \hat{s} \leq s^*_{i_{M+1}} \leq \ldots \leq s^*_{i_N}$ The divisions $H = \{i_1, \ldots, i_M\}$ now deploy $\hat{s}$, which is more that they would have deployed under the equilibrium previously considered. However, the divisions $L = \{i_{M+1}, \ldots, i_N\}$ deploy lower level of security in equilibrium that previously. The divisions in H are high capability divisions, and they are forced to increase their security measures under the imposition of a minimum standard. The low capability divisions, in L, actually *decrease* their security level deployed under the minimum standard.

This leads us to the following proposition:

Proposition 5. . *When a minimum security deployment standard is imposed, then depending on the standard, one of the following will hold true: If the standard is set too low ($\hat{s} < s^l$), then all of the divisions will deploy the same amounts as in proposition 4. If it is set too high ($\hat{s} > s^u$), then all the divisions will deploy exactly the same as the standard. If it is intermediate, then the high capability divisions with $\alpha < \hat{\alpha}$ will increase their security deployments and the low capability divisions with $\alpha > \hat{\alpha}$ will lower their security deployments In either of the above cases, the CIO can never achieve first best using a minimum security requirement.*

4.4 Optimality through subsidies

In many firms, often firms purchase security countermeasures centrally for use by all the divisions. While there might be support for economies of scale in such decisions, an important reason for this is to subsidize the divisions for spending on (contagious) threats beyond a point that the divisions would themselves choose. We demonstrate this effect via the following proposition.

Proposition 6. . *When the decision rights are vested with the divisions, the CIO subsidizes the cost of security countermeasures to c' and this achieves the first best in terms of the security measures deployed by the divisions:*

$$c' = \frac{c}{2} + \frac{\lambda L}{8\eta(N-1)}\left(\sqrt{1 + \frac{8c\eta(N-1)}{\lambda L}} - 1\right)$$

We note that the subsidy is independent of the capability of the divisions and is increasing in firm size, which is to be expected because the interdependence of the divisions increases when there are more of them. It is also increasing in internal vulnerability, and the CIO has to incentivize the divisions more to counter this. A counterintuitive result is that the subsidy is decreasing in attack level. This is driven by the fact that when the attack level is higher, the incentives of the divisions are more aligned with that of the CIO and the firm, and hence the correction required for this goal divergence is lower. It is quite interesting to note that the CIO as well as the divisions who do not have complete information on the capabilities of the divisions can achieve first best by utilizing a subsidy mechanism.

4.5 Effect of Perimeter Security

Many firms often have some form of perimeter security that prevent certain kinds of access from entities that are external to the network. The idea is to block unauthorized traffic from accessing the enterprise network. We consider the effect of using such a security countermeasure, which protects each of the divisions that is situated within the *perimeter*. All access to the firm's network has to pass through the perimeter device, which permits or blocks requests.

We assume a poisson stopping process, with some of the attacks blocked out by perimeter security before they reach the firm's divisions. The parameters of the countermeasure are:

- Cost (c_p) of the perimeter countermeasure per unit deployed
- $p(t) = e^{-t}$, is the probability of breach, given a security deployment level of t.

The number of attacks seen successfully penetrating the outer perimeter is given by:

$$T^p = \sum_{i=1}^{N} p(t) \text{ where } N \sim Poisson(\lambda)$$

If the CIO deploys perimeter countermeasures, the divisions will observe that the number of attacks that they face has dropped, and will consequently adjust their security deployments. We also assume that the CIO subsidizes the divisions appropriately according to proposition 6, so that they will optimally deploy countermeasures from the firm's viewpoint. This is formalized as a two-stage game below played between the CIO and the divisions as given below:

First Stage The CIO determines the amount of perimeter security (t) to deploy as well as the subsidy that the divisions will receive

Second Stage The divisions observe the CIO's decisions regarding perimeter security and subsidy and decide to invest appropriate amounts in security deployments

We derive an equilibrium solution to this game by using backward induction. The second stage has already been solved previously for any value of λ that the divisions observe (In order to achieve the first-best for this stage, we assume that the CIO subsidizes the divisions optimally in order to achieve this).

We use the expected loss value from *proposition 2* appropriately modified for the effect described above. We find that in equilibrium, the CIO provisions perimeter security as given in the proposition below.

Proposition 7. . *The firm's optimal deployment of security countermeasures depends upon the relative costs of security countermeasures that the divisions deploy and perimeter security as follows:*

(1) ($c_p < \frac{cN}{2}$) The CIO will deploy perimeter security to the level $t^* = log\left(\frac{N\lambda L}{c_p}\right)$ and the divisions will not deploy any countermeasures

(2) ($\frac{cN}{2} < c_p < cN$) The CIO will deploy perimeter security of the level

$$t^* = log\left(\frac{\lambda\beta LN(cN - c_p)}{\eta(N-1)(2c_p - cN)^2}\right)$$

and the divisions will deploy countermeasures of the level:

$$s_i^* = log\left(\frac{4\alpha_i\eta(N-1)}{\sqrt{1 + \frac{8cN(cN-c_p)}{(2c_p-cN)^2}} - 1}\right)$$

The CIO subsidizes the countermeasures deployed by the divisions so that the cost to them is c' given by:

$$c' = \frac{c}{2} + \frac{(2c_p - cN)^2}{8N(cN - c_p)}\left(\sqrt{1 + \frac{8cN(cN - c_p)}{(2c_p - cN)^2}} - 1\right)$$

The perimeter security (t) and the countermeasures deployed by the divisions are substitute goods for the firm.

(3) ($c_p > cN$) The CIO will not deploy any perimeter security countermeasures, but will subsidize the countermeasures deployed by the divisions to ensure optimal deployment and the divisions will deploy countermeasures of the level given in proposition 2.

There are several notable things regarding this proposition. First, in case (2), as the internal security gets worse relative to external security, a *lower* level of perimeter security is deployed. This may appear counterintuitive because one would expect that the CIO would increase perimeter security to compensate for increased vulnerability. However, what happens is that in response to higher internal vulnerability the divisions increase their security deployments too. This increase is enough so that when the CIO behaves optimally, the level of perimeter security deployed will actually be *lower*. Second, in case (2) (and trivially in case (1)), the level of security countermeasures deployed by the divisions is independent of the threat level λ. The CIO essentially deploys perimeter security to ensure that the divisions will face a constant threat level, whether the threat level external to the enterprise network increases or decreases. Third, unlike in proposition 6, the subsidy in this case is independent of the internal vulnerability parameter, η. Fourth, as the cost of perimeter security increases relative to the cost of the countermeasures deployed by the divisions, the firm deploys less of the former and increase the deployments of the latter.

5 Conclusion and Proposed Research

This paper concentrates on looking at the effect of contagious and independent threats on heterogeneous multi-division firms and the countermeasures available to defend against them, as well as the organizational decision rights and incentives that firms can utilize to ensure an optimal outcome. We find that when the threat source is an independent one, the incentives for the divisions are aligned with those of the firm. For contagious threats, the firm can use a common subsidy to ensure that each division acts in the firm's best interests.

However, there are several areas that we intend to explore further, and we describe these briefly here:

(1) Risk Averse Firms: Clearly, there are many quantifiable losses that can be incorporated in our model. However, most firms and face losses of reputation and goodwill when they are breached. How should a risk-averse firm invest in countermeasures? This is especially important if a firm could face ruin upon a successful breach.

(2) Combination of threats: When faced with both independent and contagious threats, how should the firm respond optimally? This becomes complex if security countermeasures can counter multiple threat types.

(3) Reactive countermeasures: Here, we have considered mostly protective and detective countermeasures. However, the investment a firm places in recovering from an attack determines how successfully it can manage such threats. This becomes especially important if the attacks that are of new kinds, and patches are not available in a timely manner.

References

[1] Microsoft Strategic dialogue on Risk & Security. Risk and security: Who's in charge?

[2] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 1991.

[3] Lawrence A. Gordon and Martin P. Loeb. The economics of information security investment. *ACM Trans. Inf. Syst. Secur.*, 5(4):438–457, 2002.

[4] Lawrence A. Gordon, Martin P. Loeb, and William Lucyshyn. 2005 csi/fbi computer crime and security survey, 2005.

[5] Xinzhou Qin David Dagon Wenke Lee George Riley Guofei Gu, Monirul Sharif. Worm detection, early warning and response based on local victim information. 20th Annual Computer Security Applications Conference (ACSAC '04), Tucson, Arizona, USA, December 2004.

[6] Howard Kunreuther and Geoffrey Heal. Interdependent security. *Journal of Risk and Uncertainty*, 26(2):231–49, 2003. available at http://ideas.repec.org/a/kap/jrisku/v26y2003i2-3p231-49.html.

[7] CIO Magazine. How to staff up for security.

[8] Haim Mendelson. Pricing computer services: queueing effects. *Commun. ACM*, 28(3):312–321, 1985.

[9] David Moore, Colleen Shannon, and Jeffery Brown. Code-red: a case study on the spread and victims of an internet worm. In *Internet Measurement Workshop*, Marseille, France, September 2002.

[10] Raja Nadiminti, Tridas Mukhopadhyay, and Charles H. Kriebel. Research report: Intrafirm resource allocation with asymmetric information and negative externalities. *Info. Sys. Research*, 13(4):428–434, 2002.

[11] United States General Accounting Office. Executive guide: Information security management: Learning from leading organizations. Technical report, May 1998.

[12] Peter Szor. *The art of computer virus research and defense*. Addison-Wesley, 2005.