Optimally securing interconnected information systems and assets

Vineet Kumar
Tepper School of Business, Carnegie Mellon University, vineetk@cmu.edu,

Rahul Telang
Heinz School of Public Policy and Management, Carnegie Mellon University, rtelang@andrew.cmu.edu,

Tridas Mukhopadhyay
Tepper School of Business, Carnegie Mellon University, tridas@cmu.edu,

Information security is a growing priority for organizations, many of which are struggling to decide the appropriate amounts of investments to counter threats to availability, confidentiality and integrity of information systems that put interlinked business processes at risk. The investments in security countermeasures usually have the characteristics of externalities since one entity's investment decision affects the utility of other entities that are connected to it. We specifically characterize the nature of the externalities that lead to divergence of interests between a central planner (or CIO) and division managers, where the former does not fully understand the nature of information systems deployed at different divisions within the enterprise. Also, most economics based prior research has considered information security to be a black box here we demonstrate the utility of opening up the black box by illustrating the differing strategic effects of availability and confidentiality losses in an interconnected enterprise and many of the findings turn out to be counterintuitive. Overall, we provide a rigorously derived framework to help firms design optimal mechanisms to deploy both protection and cryptographic countermeasures for availability and confidentiality losses. Several managerial implications that translate the research insights obtained in the paper to actionable recommendations are provided.

Key words: Information Security, Availability and Confidentiality Losses, Decision Rights, Optimal resource allocation

1. Introduction

Despite information security being a priority issue of many enterprises, the evaluation of investments in information security as well as how to determine firm policies is poorly understood. There are diverging views on whether decision rights for security be placed with the divisions, or with a central group responsible for security. This decision critically depends on the strategic nature

of countermeasures and the type of loss. We develop an analytical model that takes into account the heterogeneous information systems present in a multi-division enterprise, the various threats it faces to its information systems and assets, the kinds of losses these information assets can be targeted for, as well as the types of countermeasure technologies available to protect against different threats. We characterize how losses of confidentiality and availability are different and provide a rigorously derived framework to help firms design optimal mechanisms to deploy both protection and cryptographic countermeasures to combat threats.

The CSI/FBI 2005 security survey by Gordon et al. (2005) reports 13 different attacks types, ranging from web-site defacement to financial fraud to Internet worms and viruses. According to the survey, computer viruses caused the most damage in terms of dollar losses, with unauthorized access showing an increasing trend and coming in a close second. Effective countermeasures sometimes exist for many of these threats, but are often not optimally deployed. Much attention has been focused on detailing the kinds of *ad hoc* countermeasures that can protect against specific vulnerabilities, but very little attention is focused on the strategic nature of such decision making.

Within a multi-divisional firm, there are many information systems and assets, each of differing importance. For a software firm, ensuring the confidentiality of the source code to its products could be critical, whereas a financial service provider would be harmed if personal information of its customers is made public. We term these losses of confidentiality. On the other hand, firms also have data assets such as web sites or other information services that is clearly not meant to be confidential. These can face losses of availability when they are rendered unusable by a virus or denial of service attack. The type of countermeasures needed to protect against availability losses can be very different from confidentiality losses. One can commonly observe anti-virus software and firewalls that protect against worms and viruses that cause availability losses to systems, while cryptographic software and technologies like virtual private network (VPN), Secure Sockets Layer (SSL) and IPSEC protect against the loss of confidential data. Campbell and Zhou (2003) show that the change in stock market value of a breached company is higher if the breach involved confidential information as compared to a loss of availability. While measuring the extent of the

loss suffered by a firm with stock market data may have its drawbacks, it draws attention to market perceptions that certain breaches are more critical than others.

Indeed, this distinction is made in practice as detailed by Bill Boni, CISO, Motorola Inc. who says:

> *The key to security budgeting is this: Figure out what matters for the project or program in question. Is it availability of information? Confidentiality? Integrity? Take the answer to that question and couch your budget request in terms of how it will benefit the organization's specific objectives. If you speak to people about confidentiality as your lead element and the real issue they care about is availability, you're starting a losing conversation.*

1.1. Enterprise Systems - Heterogeneous and Interdependent

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 usually interconnected. For example, the engineering division in a firm may use UNIX workstations with specialized design software while the sales division may use Windows based PCs with sales software. In general, the central administrator is unlikely to know much about the systems used by the divisions, which leads us to employ an agency theoretic framework in this paper.

One distinguishing feature of our paper is to model the interdependence between heterogenous loss types (availability losses vs. confidentiality losses), heterogenous decision makers (CIO vs the divisions) and heterogenous countermeasures (protection countermeasures vs cryptographic countermeasures). We distinguish between availability losses from confidentiality losses by observing that the same threat source can cause repeated availability losses. For example, a firm will lose every time its website is brought down. On the other hand, confidentiality losses tend to have one time nature to it. Once the confidential information is lost (say customer credit card numbers) to a threat source (to hackers), losing that information again to the same hackers may not cause additional losses. We distinguish between the CIO and the divisions as decision makers. Some divisions may not value their information systems or assets as highly as others and may not be capable

of expertly deploying security measures. More importantly, the divisions may not care about the losses of the whole organization while CIO would. Finally, we distinguish between different countermeasures. Cryptographic countermeasures may not be able to stop availability related attacks while protection countermeasures may be able to stop both kind of attacks.

Depending on the enterprise network architecture, it could take just one insecure system to put the entire network at risk if internal security within the enterprise network is poor. So, each division's decisions on security deployment have an effect on the other divisions. This intuitive notion is formalized by the concept of externality in economics, where the decisions of an agent or user affect the utility of associated agents. In this paper, such externalities is the underlying cause for the divergence of interests of the CIO and the division managers.

When dealing with securing information systems, in order for there to be an externality effect in the allocation of countermeasures, the following requirements must be satisfied: (1) the systems must be interconnected (2) the threat must be capable of breaching a system and upon breaching a system, be able to attack others (3) countermeasures present at one system affect the expected breaches or damages faced by the other systems. We assume that the first requirement is satisfied - this is true in most current enterprises. Contagious threats as described below satisfy the second requirement.

1.2. Research Question

Information security has been treated like a black box by most academic researchers seeking to evaluate the effectiveness of countermeasures using a microeconomics-based approach. Our objective is to open the black-box sufficiently to be able to see the strategic effects in modeling the specifics of contagious threats in the context of different loss types (availability and confidentiality).

The principal research question we seek to answer is: How should a multi-divisional enterprise optimally deploy security countermeasures in the context of heterogeneous information systems possessed by its divisions and in response to different kind of damages that threats can cause to the enterprise's information systems and assets? We seek to explore these issues by developing

a theoretical model that accommodates these important characteristics. This work distinguishes itself from previous work along the following dimensions: (1) Heterogeneity of the information systems of the enterprise's divisions - each division is likely to have different systems and abilities to effectively deploy countermeasures (2) Different kinds of losses faced by information assets (availability and confidentiality) (3) Multiple types of countermeasures that are available to decision makers (protection and cryptography)

Our goal is to provide actionable recommendations that advance our understanding of how trade-off are affected by loss types, countermeasure types and who the decision makers are.

1.3. Prior Research

This paper draws from two streams of research. The first explores investment in countermeasures for both stand-alone and interdependent systems, where the security of each entity depends on the decisions of the other entities in the system. In one of the earliest economic analyses of information security, Gordon and Loeb (2002) show that there is a trade-off in deploying security depending upon the marginal effectiveness of the countermeasure and the cost of the countermeasure - extremely low vulnerability and high vulnerability should not be protected against.

The interdependence between the security decisions of multiple entities has been explored by Kunreuther and Heal (2003) (K&H) in the context of airline security, where airlines that decide whether to invest in countermeasures, play a non-cooperative game. K&H consider a binary decision space, and conclude that there can be an equilibrium where everyone invests in protection measures. While an important early exercise in modeling the security of interdependent systems, the choice of a binary (invest/don't invest) decision space can mask divergent goals of the individual airlines in significant subsets of the parameter space. We allow the broadest possible space for both protection and cryptographic countermeasures by treating them as continuous variables. Ayres and Levitt (1998) empirically study the positive externalities due to the presence of a certain detection countermeasure(LoJack), and find that in cities with a higher presence of cars installed with Lojack, there are fewer proportional incidents of automobile theft, and show that all automobile owners benefit from this. Bier et al. (2004) identify how to optimally allocate resources to

protect interconnected systems, where there are two types of connections: series and parallel.[1]

The second stream deals with the allocation of decision rights as well as design of optimal mechanisms. Our modeling of decision structures and mechanisms is similar in spirit to the well known papers by Mendelson (1985) and Whang (1990) which look at decision structures for information systems, albeit in the context of pricing computer services using queueing theory.

1.4. Main results

We derive several analytical results in this paper regarding security deployment by the divisions, subsidies given by the CIO to the division managers, different countermeasure types and location of decision rights. For availability losses, we find that a constant subsidy for each unit will achieve first best levels of deployment - this result is valid for any functional form satisfying basic regularity condition. Also, we find that the subsidy is increasing with the level of internal vulnerability and as the internal vulnerability becomes worse, the CIO allocates both divisions the same level of protection measures even if their losses are very different. For confidentiality losses, the decision maker has two choices to make: the levels of protection and cryptographic countermeasures. With protection measures, we find that subsidies cannot achieve the first best unless the vulnerability function is exponential in protection measures. When only cryptographic measures are used, we find that there is no goal divergence between the CIO and the divisions; however, the level of countermeasures at each division depends on the susceptibility of the other division. This result implies that the CIO must ensure that this information is communicated to the decision maker. When both protection and cryptographic measures are used, we find that the divisions can invest less in protection measures while investing more in cryptographic measures than the CIO, and we explain how this occurs due to strategic considerations. We also derive the optimal mechanism that will result in the divisions deploying the first best levels of both protection and cryptographic security measures in the case of confidentiality losses.

[1] Series systems fail if any component fails, while parallel systems fail only if *all* components fail.

2. Threats, Losses and Countermeasures

2.1. Contagious Threats

The motivation for modeling contagious threats is driven by real-world examples of computer viruses and worms. A contagious threat source launches **T** attacks per time period, which is taken to be a random variable (since we assume risk-neutrality of the decision maker, we can replace it by its mean λ.). The threat source is external to the enterprise network, and attacks each division in the enterprise separately. The information systems of the divisions are potentially vulnerable to this (external) attack, with the probability of a successful breach of division j is given by $p_j(s) = v(\alpha_j, s_j)$ depending on the susceptibility (α_j) of the division's systems as well as the level of protection countermeasures deployed by the division (s_j). Suppose that a division's information system gets breached (and consequently infected) by this threat - it turns into a threat source and becomes an internal threat to the network, which is capable of further attacks and breaches leading to losses, with an internal attack leading to a new breach with probability η.

It should be noted that many threat sources may not be contagious, and do not spread from one system to another - a phishing attack is an example that comes to mind. We label these *independent* threats - they are subsumed in our model by setting the internal vulnerability η to 0.

2.2. Availability and confidentiality losses

In our setting, an external threat source can attack a division j directly or it can attack another division k and indirectly attack j. A key distinction for the availability losses we make is that the unit j thus can be attacked multiple times (twice). Consider a virus attack that takes down a company's web servers. This represents a loss of availability and can happen multiple times, once via a direct attack and again later via another infected division's systems. In fact, a Washington Post story (December 22, 2006) highlights how some of malwares can morph themselves and evade anti-virus software and can cause repeated attacks. We note that not every availability loss may occur repeatedly. In those instances, such losses would have the same flavor as confidentiality losses (see below). Thus, to make a useful distinction, in our model, availability losses have this features of being able occur directly as well as indirectly.

In addition, most enterprises have confidential data assets that need to be secured, and a breach (or unauthorized use) of these data assets could cause monetary, legal as well as reputation losses to them. Examples of these would include strategic plans, confidential customer data, employee information and medical data. There are several threats like Spyware and Trojan horses that are designed with data theft as an objective. For a given threat source, confidential losses can occur either directly or indirectly but not both. Of course, for a separate threat source, confidentiality losses can occur again but in our model it would simply imply adding a constant scaling.

2.3. Protection and Cryptographic countermeasures

There are many ways to deal with information security threats - the ones commonly mentioned in the computer science literature are protection, detection and reaction. Protection is the primary defense, especially against unknown attackers who may not even be traceable. Detection can be useful in situations where the identity of the attackers can be established and legal or other measures can be taken against them. Reaction measures the time taken to restore the systems back to normal operation after a successful breach has taken place. In this paper, we do not model detection or reaction - setting them exogenously to a fixed level does not affect our model. This is partly due to the fact that given the international nature of computer attackers and breaches, mere detection may not be useful and reaction has multiple levels and does not lend itself naturally to our model. Moreover, given the importance of information in a targeted attack (where detection would be helpful) - a very different model may be in order.

Protecting information systems against breaches involves measures like firewalls, antivirus software etc. which block the attack from succeeding - these can secure against both availability and confidentiality threats. Cryptographic measures, however, do not protect against availability losses - they do not stop the attackers from breaching the system. However, they do render the information unusable unless the attackers can decrypt the data. We include data access policies and restrictions broadly in the category of cryptographic measures. If an attacker can breach the protection measures deployed at a system, he can usually get access to the data housed on that system. An example of protection is a firewall while PGP is an instance of a cryptographic measure.

To protect against losses, firms use both protection countermeasures (like firewalls) as well as cryptographic countermeasures (like SSL). In order to access an encrypted confidential document one must perform the following operations in sequence: (1) breach the information system and obtain access to the document and (2) break the cryptographic security protecting the document to access its contents. The probability of (1) depends on the susceptibility α_j and the level of protection countermeasures (s_j), while the likelihood of (2) depends on the level of cryptographic countermeasures (χ_j).

3. Models

In this section, we describe in detail the parameters of our theoretical framework. We consider a firm, with two divisions, which seeks to implement security countermeasures to minimize its expected loss. Each division is assumed to have a manager or decision maker who has the goal of minimizing its expected losses, including expenditure on security countermeasures. The CIO on the other hand seeks to minimize the expected losses for the overall enterprise. Given the context for decision making, we detail below the parameters that characterize the threats, the information systems and security countermeasures as well as the decisions made by the CIO or the division managers. We list all the assumptions made in this paper in table 2, most of which are very general. It is important to note that each result only uses a subset of these assumptions - exactly which ones are used for a given result are specified along with the result.

Susceptibility of a division reflects the characteristics of its information systems as well as the capabilities of its IT staff in configuring and maintaining the system for example by closing back-doors or timely patching or disabling dangerous programs. Divisions with low susceptibility are less likely to be breached than those that have higher values for this parameter. For each division j, $\alpha_{\mathbf{j}}$ is a random variable with support $[\underline{\alpha}, \overline{\alpha}]$. The realization of $\alpha_{\mathbf{j}}$ is denoted by α_j and is known to the division manager but not to the CIO. In general, we denote $E[\alpha_{\mathbf{j}}] = \tilde{\alpha}$, $Var[\alpha_{\mathbf{j}}] = \sigma^2$ for $j = 1, 2$ and $cov(\alpha_{\mathbf{i}}, \alpha_{\mathbf{j}}) = \rho\sigma^2$ for $i \neq j$. We denote the marginal distributions of $\alpha_{\mathbf{j}}$ by $\phi(\cdot)$ and the joint distribution as $\phi(\cdot, \cdot)$. Unless explicitly mentioned, we assume $\alpha_{\mathbf{1}}$ and $\alpha_{\mathbf{2}}$ are independent.

Countermeasures include security software and hardware products that minimize the chances of a successful attack. This can be achieved by mitigating the vulnerability of systems or by reducing the likelihood of a breach of data with cryptographic measures. Each division's protection and cryptographic countermeasure deployment can be set independently of the other divisions. s_j and χ_j denote the levels of protection and cryptography at division j. We assume that s_j and χ_j are continuous with a support of $[0, \infty)$. There are monetary and non-monetary costs of implementing countermeasures and are denoted by $C(s, \chi) = \tilde{c} \cdot s + \tilde{c}_c \cdot \chi$. [2]. We normalize the cost with respect to the threat level λ so that we use $c = \frac{\tilde{c}}{\lambda}$ and $c_c = \frac{\tilde{c}_c}{\lambda}$ as our cost per threat level.

Vulnerability denotes the probability of a successful breach (given that an attack has occurred) of division j's information systems, given that its protection countermeasures are of level s_j. The vulnerability to an external attack is given by $p_j = v(\alpha_j, s_j)$. When the attack originates from within the network, the vulnerability is η and this is referred to as *internal vulnerability*. A loss of confidentiality only occurs when the encryption is also breached in addition to the information system breach. The conditional probability of breach of the cryptography given that the information system has been breached is denoted by the function $\psi(\chi_j)$ where χ_j refers to the level of cryptographic countermeasures deployed by division j.

Loss 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 two different kinds of losses to the *information systems and assets* of an enterprise: loss of *availability* (e.g. a DoS attack) denoted by l_j, and loss of *confidentiality* (*e.g.* theft of confidential data) denoted by L_j for division j.

We consider in table 3 the various events that may occur when the firm is attacked by a threat source. Briefly, we distinguish between losses of availability and confidentiality as well as what divisions are breached via external and internal attacks. A breach of division j via a direct attack

[2] These are generalizable: consider the cost structure $C(s) = c \cdot s^2$. Define $z \equiv s^2$ so that cost is linear in z and we could define a function $u(\cdot, \cdot)$ so that $v(\alpha, s) \equiv u(\alpha, z)$. Now, if $v(\cdot, \cdot)$ satisfies assumptions (A1)-(A3), so does $u(\cdot, \cdot)$ and so the results will apply for $u(\cdot, \cdot)$ as well - and we can think of $u(\alpha, z)$ as our breach probability function and z as the (redefined) security level.

Table 1 Notation

Symbol	Description
α_j	Division j's susceptibility or breach probability without countermeasures
s_j, χ_j	Level of protection and cryptographic countermeasures at division j
$p_j = v(\alpha_j, s)$	Breach probability or vulnerability to external threats when division j has susceptibility α_j and security countermeasures to level s
$v_s(\alpha, s)$ and $v_\alpha(\alpha, s)$	Partial derivative of $v(\cdot,\cdot)$ with respect to s and α respectively
η	Internal vulnerability
$h(s)$	Protection countermeasure response functions, respectively
$\psi(\chi)$	Probability of confidentiality breach of data asset given that the information system has been breached
c, c_c	Normalized cost of protection and cryptographic countermeasures
$s_j^{CIO}(\alpha_j, \alpha_k)^a$	Level of protection at division j when CIO is the DM and has *complete* information
s_j^{CIO-AI}	Level of protection at division j when CIO is the DM and has *incomplete* information
$s_j^{DIV}(\alpha_j)$	Level of protection at division j when division manager is the DM
$\tau_j(\hat{\alpha_j})$	Transfer from the CIO to division j when the division reports a susceptibility of α

[a] $\chi_j^{CIO}(\alpha_j, \alpha_k), \chi_j^{CIO-AI}$ and $\chi_j^{DIV}(\alpha_j)$ represent the corresponding levels of cryptographic measures and $p_j^{CIO}(\alpha_j, \alpha_k), p_j^{CIO-AI}$ and $p_j^{DIV}(\alpha_j)$ the corresponding breach probabilities for protection measures

Table 2 List of assumptions

#	Assumption	Implication
(A1)	$0 < v_\alpha(\alpha, s) < \infty$ and $v_s(\alpha, s) < 0 < v_{ss}(\alpha, s)$ $\forall \alpha \in [\underline{\alpha}, \overline{\alpha}], s \in [0, \infty)$	Regularity conditions for vulnerability
(A2)	$v(\alpha, 0) = \alpha$, $\lim_{s\to\infty} v(\alpha, s) = 0$	Perfect security only with infinite protection
(A3)	$\psi(\chi) > 0, \psi'(\chi) < 0$ and $\psi''(\chi) > 0$ $\forall \chi \in [0, \infty)$	Regularity conditions for cryptographic measures
(A4)	$\psi(0) = 1$, $\lim_{\chi\to\infty} \psi(\chi) = 0$	Unbreakable only with infinite cryptographic measures
(A5)	Convexity	Losses are convex in countermeasures (detailed in Appendix)
(A6)	$v(\alpha, s) = \alpha h(s)$	Simplified multiplicative separability

is listed as j_D while its breach by internal attack is denoted j_I. A "✓ " indicates that the division was breached while a "× " indicates that it was not.

In the following sections, unless specified, we assume only (A1) and (A2) - that the vulnerability function satisfies regularity conditions. If any other assumptions are made, they are specified together with the result. We label the divisions j and k instead of 1 and 2 in order to minimize duplication of similar expressions.

Table 3 Scenarios for availability and confidentiality losses with contagious threats

	probability of event	1_D	2_D	1_I	2_I	Loss for specific scenario		
						Availability	Confidentiality	Confidentiality with cryptography
1	$\eta^2 p_1 p_2$	✓	✓	✓	✓	$2l_1 + 2l_2$	$L_1 + L_2$	$\psi(\chi_1)L_1 + \psi(\chi_2)L_2$
2	$p_1 p_2 (1-\eta)^2$	✓	✓	×	×	$l_1 + l_2$	$L_1 + L_2$	$\psi(\chi_1)L_1 + \psi(\chi_2)L_2$
3	$p_1 p_2 \eta(1-\eta)$	✓	✓	✓	×	$2l_1 + l_2$	$L_1 + L_2$	$\psi(\chi_1)L_1 + \psi(\chi_2)L_2$
4	$p_1 p_2 \eta(1-\eta)$	✓	✓	×	✓	$l_1 + 2l_2$	$L_1 + L_2$	$\psi(\chi_1)L_1 + \psi(\chi_2)L_2$
5	$p_1(1-p_2)\eta$	✓	×	×	✓	$l_1 + l_2$	$L_1 + L_2$	$\psi(\chi_1)L_1 + \psi(\chi_2)L_2$
6	$p_1(1-p_2)(1-\eta)$	✓	×	×	×	l_1	L_1	$\psi(\chi_1)L_1$
7	$(1-p_1)p_2\eta$	×	✓	✓	×	$l_1 + l_2$	$L_1 + L_2$	$\psi(\chi_1)L_1 + \psi(\chi_2)L_2$
8	$(1-p_1)p_2(1-\eta)$	×	✓	×	×	l_2	L_2	$\psi(\chi_2)L_2$
9	$(1-p_1)(1-p_2)$	×	×	×	×	0	0	0

4. Availability Losses

This section details how the CIO and the division managers respond differently to threats posed to the availability of information systems. We consider the impact of the commonly used practice of setting a common standard for heterogeneous systems. Availability losses are assumed to be unaffected by the presence of cryptographic countermeasures (which can protect data), and are not specifically modeled in this section. We explore whether there are simple methods of incentivizing the divisions to do what is optimal from the point of view of the overall enterprise. For this section, we assume (A1) and (A2) unless other assumptions are specifically mentioned in the results.

For availability losses, a direct external attack and internal attack cause additive losses. As we noted earlier, in availability losses, a division can be breached directly as well indirectly incuring the loss twice. In table 3, scenario 3 represents the case when division 1 suffers a direct attack as well as an indirect attack, thus facing a loss of $2 \cdot l_1$ while division 2 suffers only a direct attack which results in a loss of l_2. An indirect attack on division j can only happen if there has been a breach by a direct attack of division k.

The overall availability loss of the firm (we normalize everything wrt to λ) is given by the random variable $\mathcal{L}_{\mathcal{A}}$ (normalized loss) that sums up the losses under each scenario (as outlined in Table 3), weighted by its probability of occurring. It depends on the susceptibilities $\alpha_{\mathbf{1}}$ and $\alpha_{\mathbf{2}}$ as well as the level of protection countermeasures deployed at each of the divisions s_1 and s_2.

$$\mathcal{L}_{\mathcal{A}} = v(\alpha_{\mathbf{1}}, s_1)(l_1 + \eta l_2) + v(\alpha_{\mathbf{2}}, s_2)(l_2 + \eta l_1) + c\,[s_1 + s_2]$$

Note that the cost of countermeasures is included in the loss function to reflect the overall expenditure of the decision maker.

The loss faced by division j is given as:

$$\mathcal{L}_{\mathcal{A}}{}^{j} = v(\alpha_j, s_j)(l_j) + v(\alpha_k, s_k)(\eta l_j) + c[s_j]$$

Notice that the division only considers the losses it incurs (directly or indirectly) due to security breach and ignores the other division's losses.

PROPOSITION 1. *(i) The division managers always invest less in protection countermeasures as compared to when the CIO has decision rights and complete information and the divergence between* s_j^{CIO} *and* s_j^{DIV} *increases with internal vulnerability* η*.*

(ii)The investment in protection countermeasures are independent of other division's investment ($\frac{\partial s_i^{CIO}}{\partial s_j} = 0$ *and* $\frac{\partial s_i^{DIV}}{\partial s_j} = 0$*) and increasing in own susceptibility* $\frac{\partial s_j^{CIO}}{\partial \alpha_j}, \frac{\partial s_j^{DIV}}{\partial \alpha_j} > 0$*.*

All proofs are in the appendix. As we expect, since the division ignores the externality it imposes on the other division, the division manager always invest less than the CIO. More importantly, note that the protection deployment decision of each division is *strategically independent* from that of the other division- this holds whether the divisions or the CIO are vested with decision rights. In other words, division 1 will not consider the decision of division 2 while setting its security level even though each division's action affects the other division. However, since the internal vulnerability can not be reduced by more investments, each division invests its marginal dollar only on the direct losses it can stop.

As expected, the protection level of each division is increasing in its own susceptibility. More susceptible division invests more.

4.1. Optimality through subsidies

When one thinks about aligning the incentives of the division managers with the CIO (or the firm), the measure that immediately suggests itself is a subsidy (or tax which is a negative subsidy), because of its frequent use in problems with externalities - *e.g.* in the control of pollution. Here

we find that a subsidy can actually do much better than imposing a common standard - it can achieve the first best in the sense that the CIO can get the divisions to implement what is optimal for the firm.

PROPOSITION 2. *Subsidies can achieve the first best outcome. When the CIO subsidizes the cost of protection measures to the divisions then the divisions will implement exactly the level that is required to minimize the overall firm's expected loss. The subsidy per unit of protection countermeasure is given as:*

$$sub_j = \frac{c\eta l_k}{l_j + \eta l_k} \text{ where } j \neq k$$

What's interesting about this result is that it is independent of the functional form of $v(\cdot,\cdot)$ and only requires the very general assumptions (A1) and (A2) to be satisfied. As expected, as the internal vulnerability gets worse, the CIO subsidizes both divisions more.

5. Confidentiality losses

In this section, we consider the loss to a firm due to a breach of confidential data stored at its divisions, as was described in the introduction. In this section, we assume that when protection measures are used, the vulnerability and the cryptographic response function satisfy the regularity conditions given in table 2.

The key aspect of confidentiality losses is that for each attack, if a division'j' has been breached via a direct attack and has a loss of L_j, a further breach via an internal attack does *not* lead to an additional loss. This is because the attacker (spyware or worm/trojan horse) has already gained access to the sensitive information via the direct breach.

We proceed to compare the effects of vesting decision rights with the CIO or the division managers when they can only use protection countermeasures and when they use only cryptographic countermeasures, and finally the case when both types of countermeasures can be used together. We do this for two reasons: first, the focus on a single type of countermeasure gives us a benchmark and allows us to examine the strategic effects of each type separately and second, in some enterprises, a single type of countermeasure is predominantly used.

5.1. Exclusive use of protection countermeasures

We consider how a firm with access to protection countermeasures will deploy them to protect against losses of confidentiality under different decision rights structures. The loss is given by a random variable (from Table 3):

$$\mathcal{L}_{\mathcal{C}} = v(\alpha_{\mathbf{1}}, s_1)(L_1 + \eta L_2) + v(\alpha_{\mathbf{2}}, s_2)(L_2 + \eta L_1) - \eta v(\alpha_{\mathbf{1}}, s_1)v(\alpha_{\mathbf{2}}, s_2)(L_1 + L_2) + c\,[s_1 + s_2]$$

Again, the propensity of the divisions is to under-invest in security measures, as detailed by the following result.

PROPOSITION 3. *When facing threats to loss of confidentiality and using only protection measures to secure against these threats, we have the following:*

1. *division managers always invest less the the CIO in protection measures.*
2. *The investments in protection countermeasure by a division are strategic complement to the investment by the other division:* $\frac{\partial s_j}{\partial s_k} > 0$ *for both the CIO and the division*
3. *When the divisions are vested with decision rights,* $\frac{\partial s_j^{DIV}}{\partial \alpha_j} > 0$ *while* $\frac{\partial s_j^{DIV}}{\partial \alpha_k} = 0$
4. *When the CIO is vested with decision rights and knows the susceptibilities of the divisions, the sign of* $\frac{\partial s_j^{DIV}}{\partial \alpha_j}$ *cannot be determined without knowing the functional form of* $v(\alpha, s)$*. It can be positive or negative.*

There are several aspects of the above result that we want to expand on. First, notice that the security levels in the case of confidentiality losses are *strategic complements*, as opposed to the case in availability losses. This result implies that when division 1 has a higher level of security, it is optimal for division 2 to deploy a higher level of security. This is true whether the decision maker is the CIO or the division manager. Intuitively one might have expected the opposite - that the protection measures at each division substitute for each other. The intuition for why the protection levels are strategic complements follows: Consider when the protection at division 1 is reduced. Then, unit 1 has a higher probability of being breached by a direct attack. This implies that the probability that division 2 will be breached by an internal attack is also higher. Now, since division

2 cannot protect itself against an internal breach, and an external breach causes additional damage, it will find it optimal to reduce its level of protection, s_2.

Second, when the CIO knows the susceptibility of both divisions, he will choose to set the security level of each division contingent on the susceptibility of *both* divisions, not just the own susceptibility level. Contrast this with proposition 1, where the optimal action for the CIO involved setting the security levels of each division as a function of its own susceptibility only. This is because, there, the availability losses are strategically independent. In general, we expect that $\frac{\partial s_j}{\partial \alpha_j} > 0$, which is the formalization of the expectation that the more susceptible a division's system is, the higher should be the deployment of protection measures at the division. However, in addition to this intuitive direct effect, there is an indirect effect that works in an opposing direction to the direct effect: A higher susceptibility at division j will lead division k to reduce its protection level, s_k. Since the protection levels are strategic complements, this would lead to a lower optimal deployment s_j. Whether the direct (or own) effect dominates over the cross effect depends on the specific functional form as well as the region of the parameter space the problem is set in.

Without further restrictions on the functional forms of $v(\alpha, s)$, we cannot say whether the own effect or the cross effect will dominate. We provide two specific instances to show that either can reasonably be expected to occur.

Example 1 : If $v_{\alpha s} = 0$, we have $\frac{\partial s_j}{\partial \alpha_j} < 0$ or s_j decreases in α_j. So, as the susceptibility increases, divisions will actually decrease their security deployment. An example of this situation is when the vulnerability is additively separable: $v(\alpha, s) = h(\alpha) + g(s)$.

Example 2: $\eta \to 0$ When the internal vulnerability is low, then the second term is dominated by the first. This leads to the determinant being negative or $\frac{\partial s_j}{\partial \alpha_j} > 0$. As the susceptibility increases, protection deployment increases in this case.

5.1.1. Do subsidies get the first best for protection countermeasures? Subsidies achieve optimality when dealing with availability losses and it's important to verify whether they can achieve first best for confidentiality losses too. This would preempt the design of more complex

mechanisms and be of immense value to practitioners. Below, we show that strategic complementarity also means that subsidies in general can not achieve the first best - except for a very restrictive case of exponential vulnerability functions.

PROPOSITION 4. *When the firm faces threats to loss of confidentiality of its data assets and uses only protection countermeasures,*

1. *Subsidies do not achieve first best for the CIO in general.*

2. *Subsidies will be optimal only for an exponential vulnerability function* $v(\alpha_j, s_j) = \alpha_j e^{-bs_j}$ *where* $b > 0$.

When the vulnerability function is exponential or can be closely approximated by it, the CIO may choose to implement a subsidy to incentivize the divisions to implement higher protection measures. The (**can't we write the subsidy function without same L**) The subsidy increases with internal vulnerability and decreases with the loss level for each division (when the losses are equal), since with a higher level of loss, each division will deploy a higher level of protection measures. The maximum subsidy that each division receives is when the divisions are fully vulnerable to internal attacks ($\eta = 1$). In this case, the divisions receive a subsidy of $\frac{3}{4}c$, or the CIO subsidizes upto 75% of their protection countermeasures. This is more than the 50% maximum subsidy in the case of availability losses. We can see that the externality effect can be more harmful for confidentiality losses than for availability losses.

5.2. Exclusive use of cryptographic countermeasures

We now examine the impact of the cryptographic countermeasures. If divisions deploy only cryptographic countermeasures, then the probability of breach of their information systems is $v(\alpha_j, 0) \equiv \alpha_j$ (which is private information held by the division manager). We examine the levels of cryptographic measures deployed when the decision rights are vested with the CIO and with the division manager.

The loss faced by the enterprise depends on the susceptibilities as well as the level of cryptographic measures deployed at the divisions as is given as: (**shouldn't they be small l**)

$$\mathcal{L}_{\mathcal{C}} = (\alpha_{\mathbf{1}} + \eta\alpha_{\mathbf{2}} - \eta\alpha_{\mathbf{1}}\alpha_{\mathbf{2}})\,\psi(\chi_1)L_1 + (\alpha_{\mathbf{2}} + \eta\alpha_{\mathbf{1}} - \eta\alpha_{\mathbf{1}}\alpha_{\mathbf{2}})\,\psi(\chi_2)L_2 + c_c\,[\chi_1 + \chi_2]$$

$$\mathcal{L}_{\mathcal{C}}{}^{j} = (\alpha_{\mathbf{j}} + \eta\alpha_{\mathbf{k}} - \eta\alpha_{\mathbf{j}}\alpha_{\mathbf{k}})\,\psi(\chi_1)L_j + c_c\chi_j$$

The optimal levels of cryptographic measures are detailed in the following result.

PROPOSITION 5. *(i) Facing threats confidentiality losses, a firm that deploys only cryptographic countermeasures does not have any goal divergence between the division and the CIO making investment decisions.*

(ii) The investments by one division is strategically independent of the investments by the other division irrespective of who is vested with decision rights : $\frac{\partial \chi_j^{DIV}}{\partial \chi_k^{DIV}} = \frac{\partial \chi_j^{CIO}}{\partial \chi_k^{CIO}} = 0$ *(iii) Optimal Investments for a division when CIO is making the decision are increasing in the division's own susceptibility and the other division's susceptibility :* $\frac{\partial \chi_j^{CIO}}{\partial \alpha_j}, \frac{\partial \chi_j^{CIO}}{\partial \alpha_k} > 0$. *However, when a division is making decisions, investment decisions are independent of the other division's susceptibility :* $\frac{\partial \chi_j^{DIV}}{\partial \alpha_j} > 0, \frac{\partial \chi_j^{DIV}}{\partial \alpha_k} = 0,$

There are three observations that are important to note here: First, there is no goal divergence between the CIO and the division managers - indeed, if the CIO and the divisions had complete information they would make exactly the same decisions. This lack of strategic dependence happens because each division's cryptographic measures protect only its data assets from being breached, and have no externality effect on the other division's likelihood of breach. Also, one might be tempted to assume that the party with more information always makes the best decisions and allocate the decision rights to the divisions, since they know their own susceptibility while the CIO only knows the distribution. While this may hold in expectation, this will not hold for a subset of susceptibilities, i.e., there is a range of realizations of susceptibilities for which centralizing the decision will actually work better, but it is not known in advance whether the realized values of $\alpha_{\mathbf{1}}$ and $\alpha_{\mathbf{2}}$ fall in this region.

Second, the optimal level of cryptographic measures at each division depends on the susceptibilities of each division, much like protection countermeasures with confidentiality losses. So, there

must be effective communication channels either between the divisions or between the divisions and the CIO to transfer information about the susceptibilities. Since there are no incentive issues and all parties have the same preferences, we can use any suitable mechanism to communicate the susceptibility of each division to the other, with or without the intermediation of the CIO. The level of cryptographic measures at a division is increasing not only in its own susceptibility, but that of the other division too. Unlike protection measures, cryptographic measures secure against internal attacks as well since the internal threat can get access to a confidential file, but the encryption still prevents the data from being accessed.

Third, the optimal deployment of cryptographic at each division increases as the internal security (η) gets worse; this is exactly the opposite of protection measures. This happens because unlike protection countermeasures, which only reduce the likelihood of breach for an external attack, cryptographic measures ensure that only users with appropriate authorization can access the data assets. It prevents unauthorized access to the data assets no matter whether the attacker is internal or external to the enterprise network.

Now we are ready to explore the case when both countermeasures are used for confidentiality related losses. Recall that cryptographic measures are useless for availability losses.

5.3. Using both protection and cryptographic countermeasures

In many enterprises, multiple countermeasures are used to secure information assets against confidentiality losses. Protection measures protect against information systems being breached, while cryptographic measures protect unauthorized users from accessing sensitive data. While we have modeled the effect of using either protection or cryptographic measures above, in many enterprises, one often finds both types of countermeasures being used. In this case, the levels of each type of countermeasure is determined not only by the strategic interaction between the levels at each division, but the substitution effect between diffrerent types of countermeasure at a single division. The following result illustrates how the levels of protection and cryptographic measures interact with each other.

PROPOSITION 6. *When both protection and cryptographic measures are used to secure against confidentiality losses, we find that the security measures at each division can be strategic complements or substitutes depending on the type of measures. Specifically,* $\frac{\partial s_1}{\partial s_2} > 0$ *and* $\frac{\partial \chi_1}{\partial \chi_2} < 0$ *and this is independent of who is vested with decision rights.*

Recall that when protection measures were deployed exclusively, s_1 and s_2 were strategic complements. However, when only using cryptographic measures, we found that χ_1 and χ_2 were strategically independent. The above result shows that when both types of countermeasures are used, the strategic complementarity property of the protection measures is preserved. What's interesting is that cryptographic measures, which were strategically independent when used by themselves, become strategic substitutes in the presence of protection.

Next, we now look at whether the CIO always deploys a higher level of protection and the same level of cryptographic measures as compared to the division manager. One might expect it since this was the case when each type of countermeasure was used exclusively. The result below informs us that this may not hold.

5.4. Designing an optimal mechanism to secure against confidentiality losses

As we have observed above, protecting against confidentiality is a more complex problem for the CIO as well as division managers because of the strategic interdependence between their security decisions. Since we have a principal (CIO) and two agents (divisions) who possess private information, mechanism design is the appropriate methodology to implement the social choice function (first best countermeasure levels). We design an efficient mechanism that implements the exact decisions of the CIO using Bayesian implementation. This approach has been previously used in contexts such as pollution control by Baliga and Maskin (2003) and water resource management by Smith and Tsur (1997). Unlike the previous applications, in our problem with both types of countermeasures, we have two separate decisions for each report of susceptibility received from the divisions. So, the mechanism has to explicitly consider the strategic effect between both decisions for implementation.

By the revelation principle[3], we restrict ourselves to direct mechanisms in which the agents report their private information (susceptibility or $\hat{\alpha_j}$) to the CIO, who then allocates the division a protection security level $s_j(\hat{\alpha_j}, \hat{\alpha_k})$ and a cryptographic countermeasure level $\chi_j(\hat{\alpha_j}, \hat{\alpha_k})$ depending upon the reports of both divisions and a transfer $\tau(\hat{\alpha_j})$ that depends only on the division's own report.

We describe below the standard three step process by which this takes place:

1. The CIO offers a set of contracts to each division $(s_j(\hat{\alpha_j}, \hat{\alpha_k}), \chi_j(\hat{\alpha_j}, \hat{\alpha_k}), \tau_j(\hat{\alpha_j}))$ that depends on the reported susceptibility of each of the divisions

2. Each division j sends a report $\hat{\alpha_j}$ of its true susceptibility α_j

3. The CIO observes the reports of all the divisions, and allocates the security countermeasure level for each division and the transfer amounts based on the reports

If each division reports its susceptibility truthfully in equilibrium, then the (direct) mechanism is said to be truthful and is called a direct revelation mechanism.

The CIO wants to implement the protection and cryptographic countermeasure level at each division that results in minimizing the overall loss for the firm, as given in the proof of proposition 7 - we need to find the transfer scheme that will implement this. The following result gives us the solution.

PROPOSITION 7. *The direct revelation mechanism represented as:* $\left(s_j^{CIO}(\hat{\alpha_j}, \hat{\alpha_k}), \chi_j^{CIO}(\hat{\alpha_j}, \hat{\alpha_k}), s_k^{CIO}(\hat{\alpha_j}, \hat{\alpha_k}), \chi_k^{CIO}(\hat{\alpha_j}, \hat{\alpha_k}), \tau_j(\hat{\alpha_j}), \tau_k(\hat{\alpha_k})\right)$ *where* $\tau_j(\cdot)$ *is the solution to the differential equation below and* $s_j^{CIO}(\hat{\alpha_j}, \hat{\alpha_k})$ *and* $\chi_j^{CIO}(\hat{\alpha_j}, \hat{\alpha_k})$ *are derived from Lemma 3 above implements the first best security levels by the divisions.*

$$\tau_j^{'}(\alpha_j) = \frac{\partial}{\partial r_j}\Bigg[\int_{\underline{\alpha}}^{\overline{\alpha}} \Big\{ v\left(\alpha_j, s_j^{CIO}(r_j, \alpha_k)\right)\left(1 - \eta v\left(\alpha_k, s_k^{CIO}(r_j, \alpha_k)\right)\right)\chi_j^{CIO}(r_j, \alpha_k)L_j + c \cdot s_j^{CIO}(r_j, \alpha_k) + c_c \cdot \chi_j^{CIO}(r_j, \alpha_k)\Big\}\phi(\alpha_k)d\alpha_k\Bigg]\Bigg|_{r_j=\alpha_j}$$

[3] See Fudenberg and Tirole (1991) for a proof

The direction of the transfer is assumed to be from the CIO to the divisions and such a transfer is represented by a positive value. Transfers in the other direction are therefore negative. Since the transfers take place within an enterprise (between the CIO and the divisions), we assume there are no transaction costs for the transfers. Additionally, since the direction of the transfers can be made to be from the CIO to the divisions, this will likely hold true in practice as argued by Radner (1986). When we use this mechanism, we have the option of adding a fixed constant transfer of any amount which does not depend on the report of the divisions. This is a general feature of bayesian mechanisms, cf. Baliga and Maskin (2003), and such a constant transfer could be used to solve problems of individual rationality, in case divisions are reluctant to participate in the mechanism. However, given our context of a firm and its divisions, we can also exogenously impose the constraint that all divisions are required to participate in the mechanism as opposed to the case where they are incentivized to do so. This paper is agnostic with respect to the choice of which method to use in a practical situation and we do not deal with this matter further.

6. Effect of correlation on CIO's deployment decisions

We now look at the effect of correlation and variance of the susceptibilities of the two divisions on the optimal security deployments at each division as well as the expected loss of the firm. While availability losses were unaffected by correlation (or variance), for confidentiality losses we find that the expected loss can either increase or decrease with it as given by the following result.

PROPOSITION 8. *For threats involving availability losses, the overall expected loss is independent of correlation or variance of the susceptibilities of the two divisions (α_1 and α_2). However, when faced with threats to loss of confidentiality, the overall expected loss of the firm for any given level of protection and cryptographic measures at the divisions is decreasing (increasing) in ρ and σ^2 when $\rho > 0$ (< 0).*

For the correlation result, we have the following explanation: Attacks often target specific vulnerabilities in a software program or the operating system. So, the susceptibilities could be positively correlated because the systems of the divisions could have the same platform (e.g. Linux) or they

use similar software applications or software built from the same components (shared libraries or code bases). When this is so, the result above tells us that for any specified level of countermeasures, the expected loss falls as the correlation increases. This seems counterintuitive, but looking into the structure of confidentiality losses will help in understanding this result. Recall that, for confidentiality losses, when one division is breached, there is a positive likelihood of the other being breached via internal attack in addition to external attack. If division 1 has a high susceptibility, so that it is likely to be breached, a higher susceptibility for division 2 will not hurt marginally as much, as when division 1 had a low susceptibility. Therefore, correlation actually is preferable in the sense that if one division is less susceptible then it is more helpful for the other to be less susceptible to the same threat than if the first division were more susceptible. On the other hand, if a division is highly susceptible, it does not make much marginal difference if the other division's susceptibility were high as well. Also, it is counterintuitive that the expected loss is decreasing in variance when there is a positive correlation (σ^2 could be high because the number of programs installed is highly variable). However, the exact opposite effects result when correlation is negative. We assume the same attack rate on each division.

7. Managerial Implications

There are several points that practitioners can take away from our theoretical analysis and apply to real world situations. We must, however sound a word of caution that one must examine the assumptions underlying any model, and examine whether it holds for the situation at hand, as is true with all such attempts. Our assumptions are listed in table 2 and the reader will find that most of them are very general and not restrictive. We now discuss the main implications of our analysis from which CIOs and IT managers can derive actionable recommendations. The underlying theme here is that there are strategic issues in information security decision making and that the distortion due to incomplete knowledge of information systems by the CIO has to be weighed against incentive problems when division managers make decisions. Availability losses turn out to be relatively easier to deal with since the deployment decisions involving protection

countermeasures at each division are strategically independent. This implies that the CIO can set a constant subsidy and vest decision rights with the division managers, and this will be optimal for almost all forms of vulnerability functions. Availability losses are unaffected by correlation between the susceptibilities so this analysis applies irrespective of the heterogeneity of the information systems at the divisions, so the above recommendations hold whether all the systems run on a particular platform or whether they run different ones. The centralization of decision making for availability losses is unjustified if the CIO is less than fully informed about the systems in use at the divisions.

We illustrate the strategic effects of each countermeasure type and show that protection countermeasures are strategic complements and cryptographic measures are strategically independent when deployed exclusively. One might expect that in an interconnected network, if the protection measures at one division are reduced, the other division will in some sense try to make up by increasing its protection. We show that the exact opposite response is optimal - if protection is increased at one division, it is optimal to increase the deployment at the other division. This result holds no matter whether the decision rights are vested with the CIO or the division managers. It is natural for a manager to be curious whether a fixed subsidy per unit of protection countermeasures will solve the problem for confidentiality losses like it did for availability losses. Unfortunately, due to strategic interaction, this is not true in general and only holds for a restrictive class of exponential response vulnerability functions - for this class of vulnerability and identical losses for the divisions, the optimal subsidy for confidentiality losses ranges between 0% and 75% while the corresponding range for availability losses is 0%- 50%, and the subsidy increases with the internal vulnerability.

When only cryptographic countermeasures are used, there is no divergence of interests between the divisions and the CIO. However, the CIO must enable communication between the divisions so that each division knows the other's susceptibility in order to implement the first best level of cryptographic countermeasures. An alternative possibility is that units exchange such information between themselves, since there is no incentive to lie. When both protection and cryptographic

countermeasures can be used and the divisions are vested with decision rights, in general they invest less in protection but more in cryptographic countermeasures when compared with the case where the CIO with complete information makes the decision. Therefore, in the presence of multiple types of countermeasures, under-deployment of protection as well as over-deployment of cryptographic countermeasures by the divisions can simultaneously occur; prior analytical work has not indicated this since multiple types of countermeasures have not been modeled in information systems. We also design a direct revelation mechanism for confidentiality losses which the CIO can implement to achieve the first best. In such a mechanism, the divisions reveal their susceptibilities to the CIO who then sets the level of both protection and cryptographic countermeasures at each division. The mechanism includes a transfer payment that depends on the division's report.

When dealing with confidentiality losses with decision making by the CIO under incomplete information, we find that a positive correlation between the susceptibilities actually reduces the overall loss for the enterprise, and this is true for any level of security deployments at the divisions. This finding stands in stark contrast to prior work like Kataria et al. (2005) and O'Donnell and Sethu (2004) that have examined correlation and even concluded that it might increase expected losses. This is explained by the fact that we specifically model losses of confidentiality of data assets to reflect the fact that multiple breaches by the same attack (and hence, attacker) do not result in additive losses. If the CIO decides the optimal levels of security deployed at the divisions, he should choose deployments that *decrease* with correlation when the latter is positive. This result is robust and holds for both types for countermeasures - protection and cryptographic. When comparing two real world enterprises, one with the divisions having very similar IT systems (and similar configuration) the CIO must deploy a lower level of countermeasures than in an enterprise when the systems are very different (say Windows and Unix) *ceteris paribus*[4].

8. Conclusion and Future Research

Information security is a major IT priority for many firms, and spending on security products and services is ballooning according to the CSI/FBI report and others. However, the effectiveness of

[4] Our analysis assumes that the mean attack rate on each type of system is the same

this security spending and decision structures is a matter under debate and most current practice does not consider the incentives of decision makers in the enterprise. When interconnected systems are subject to attacks by contagious threats like Internet worms, this results in an externality effect if the decision rights for securing the different systems are vested with different agents. We have proposed a general framework for modeling different threat types, loss types as well as different countermeasures. This helps us evaluate the effects of vesting decision rights with the CIO or the division managers.

The analysis here is intended to lend insight to academics and practitioners on the different strategic implications of combinations of loss and countermeasure types along with decision rights. We have characterized the optimal levels of deployment of both protection and cryptographic countermeasures for the various permutations detailed previously. Additionally, we have derived optimal mechanisms to implement what is best from the overall enterprise's point of view. For example, for losses of availability, the levels of deployment of protection measures are strategically independent, even though there is an externality effect. This implies that a subsidy will result in the first best level of countermeasures at each division. In general, for losses of confidentiality we have a strategic dependence between the countermeasures at each division and consequently, a more complex mechanism is required for optimal implementation. For confidentiality losses, we have derived many non-obvious and counterintuitive results - for example, the optimal level of protection measures decrease with internal vulnerability, but the optimal level of cryptographic measures increases with it. This mechanism is required to consider the strategic effects of both protection and cryptographic measures. Throughout this analysis, we have attempted to use the most general functional forms that still give us specific insights into this important problem. The analysis is readily extensible on a number of dimensions, e.g. the attack rates for each division can be modeled to be different.

Our paper has examined what an enterprise must do to secure itself against multiple types of threats and losses. Further research efforts could consider explicitly the motivations of hackers and methods of deterrence and protection. There has been some work on this issue, notably by

Png et al. (2006), who consider the incentives for users to protect themselves against breaches by hackers by investing in precautions that mitigate losses. Also, modeling losses of integrity would be a valuable contribution to this area of research. Integrity could potentially be operationalized by considering n copies of a data asset, and changing a certain number of these could be considered a breach.

Another useful extension or possibly a separate research effort would be to look at the effect of minimum security standards. This is not unlike research in the economics of industrial organization which has examined the effect of imposing minimum quality standards in the setting of duopolistic competition. It turns out that the effects for consumers and firms are sensitive to the nature of competition - Bertrand or Cournot. We are uncertain whether the results extend to minimum security standards and it would be interesting to know either way since this is another commonly employed and practical method used by organizations today to secure against threats to information assets.

Acknowledgments

Vineet Kumar would like to express his appreciation to Carnegie Mellon's Cylab for a generous fellowship awarded for study of information security issues. Rahul Telang acknowledges generous financial support of NSF through the CAREER award, CNS-0546009

References

Ayres, Ian, Steven D. Levitt. 1998. Measuring positive externalities from unobservable victim precaution: An empirical analysis of lojack. *The Quarterly Journal of Economics* **113**(1) 43–77.

Baliga, Sandeep, Eric Maskin. 2003. Chapter 7 mechanism design for the environment. Karl-Goran Maler, Jeffrey R. Vincent, eds., *Environmental Degradation and Institutional Responses*, vol. Volume 1. Elsevier, 305–324.

Bier, V., A. Nagaraj, V. Abhichandani. 2004. Protection of simple series and parallel systems with components of different values. *Reliability Engineering and Systems Safety* **10(27)**.

Campbell, K., L. Zhou. 2003. The economic cost of publicly announced information security breaches: Empirical evidence from the stock market. *Journal of Computer Security* **11** 431–448.

Fudenberg, D., J. Tirole. 1991. *Game Theory*. MIT Press.

Gordon, Lawrence A., Martin P. Loeb. 2002. The economics of information security investment. *ACM Transactions on Information Systems Security* **5**(4) 438–457. doi:http://doi.acm.org/10.1145/581271.581274.

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

Kataria, Gaurav, Pei yu Chen, Ramayya Krishnan. 2005. Software diversity for information security. *Proceedings of the 2005 Workshop on the Economics of Information Security*.

Kunreuther, Howard, Geoffrey Heal. 2003. Interdependent security. *Journal of Risk and Uncertainty* **26**(2) 231–49.

Mendelson, Haim. 1985. Pricing computer services: queueing effects. *Commun. ACM* **28**(3) 312–321. doi:http://doi.acm.org/10.1145/3166.3171.

O'Donnell, Adam J., Harish Sethu. 2004. On achieving software diversity for improved network security using distributed coloring algorithms. *Proceedings of the 11th ACM conference on Computer and communications security* 121–131doi:http://doi.acm.org/10.1145/1030083.1030101.

Png, Ivan P.L., Candy Q. Tang, Qiu-Hong Wang. 2006. Information security: User precautions and hacker targeting. *Working Paper Series*. SSRN. URL `http://ssrn.com/abstract=912161`.

Radner, Roy. 1986. The internal economy of large firms. *The Economic Journal* **96** 1–22. URL `http://links.jstor.org/sici?sici=0013-0133%281986%2996%3C1%3ATIEOLF%3E2.0.CO%3B2-D`.

Smith, Rodney B. W., Yacov Tsur. 1997. Asymmetric information and the pricing of natural resources: The case of unmetered water. *Land Economics* **73**(3) 392–403.

Whang, S. 1990. Alternative mechanisms of allocating computer resources under queueing delays. *Information Systems Research* **1(1)** 71–88.