Understanding and Influencing Attackers' Decisions: Implications for Security Investment Strategies

Marco Cremonini
Dept. of Information Technology
University of Milan, Italy
cremonini@dti.unimi.it

Dmitri Nizovtsev
School of Business
Washburn University, USA
dmitri.nizovtsev@washburn.edu

Abstract
We model economic behavior of attackers when they are able to obtain complete information about the security characteristics of targets and when such information is unavailable. We find that when attackers are able to distinguish targets by their security characteristics and switch between multiple alternative targets, the effect of a given security measure is stronger. That is due to the fact that attackers rationally put more effort into attacking systems with low security levels. Ignoring that effect would result in underinvestment in security or misallocation of security resources. We also find that systems with better levels of protection have stronger incentives to reveal their security characteristics to attackers than poorly protected systems. Those results have important implications for security practices and policy issues.

1. Introduction

The importance of developing quantitative models of computer security has been widely recognized in economics (Gordon&Loeb, 2005), computer science (Schechter, 2004, Liu et al., 2005), and dependable computing (Littlewood et al., 1993, Avizienis et al., 2004). More specifically, great attention has been paid to analyzing the interaction of computer systems with the operational environment, which, from the security standpoint, includes the behavior of attackers as one of the main components. Quantitative techniques for evaluating attackers' behavior have been fruitfully applied (Avizienis et al., 2000, Nicol et al., 2004) and several models of attackers' behavior have been proposed (Jonsson&Olovsson, 1997, Ortalo et al., 1999, McDermott, 2005). However, the attacker's behavior in the aforementioned works is modeled as exogenous and the principles guiding the attackers' behavior remain unclear.

The current paper examines attackers who are assumed to behave economically (that is, choose their actions optimally based on comparison of their costs to benefits). Viewing attackers as rational agents is consistent with several theoretical and empirical studies. Some prior work has recognized that attackers act strategically either by rationally selecting their targets or in response to targets' actions (Jajodia&Miller, 1993; NIST, 2002, Schechter and Smith, 2003). Leeson and Coyne (2006) make a distinction between fame-driven and profit-driven attackers,[1] with the former attracted by the possible

[1] We prefer the expression "economic behavior" to "profit-driven" as the more accurate and less restrictive one.

notoriety and the latter focused on gaining monetary rewards, and conclude that the two groups must be analyzed separately. In the past the stereotypical view of an attacker was mainly that of a fame-driven individual (Denning, 1990). Even more recently the "15 minutes of fame" was claimed to be one of the biggest motivation for attackers (Curry, 2002). While fame-driven attackers are certainly still numerous, ample evidence exists that economically-minded attackers are posing a much more serious threat to corporate information security. Recently, a pronounced shift toward financially motivated intrusions (Sieberg, 2005) has resulted in an increase in the average losses caused by unauthorized access to information and theft of proprietary information (Gordon et. al., 2005).

We follow the body of work (Anderson, 2001; Schechter and Smith, 2003) that addresses the issue of the optimal amount of security by studying attackers who rationally choose their course of action based on cost-benefit analysis. We specifically focus on the role of the opportunity cost of attacking a given target, which is represented by alternatives available to an attacker.

Our results support the important role the presence of alternative targets plays in attackers' decisions. More specifically, we show that in the presence of targets with heterogeneous security characteristics the amount of effort optimally spent by an attacker on a target decreases in the target's security level. Thus any given security measure affects the frequency of intrusions through two mechanisms. One is the increased ability of a target to withstand attacks of a given intensity. This direct technical effect is commonly recognized by practitioners dealing with security threats. The other effect contributing to a reduction in intrusions occurs through a change in attackers' perception of the target in question. In the presence of alternatives, a more secure target is less attractive for rational attackers, which eventually results in the decreased effort attackers put into attack attempts. Unfortunately, this behavioral component is largely neglected when security strategies are defined. This paper demonstrates that the behavioral effect can substantially exceed the direct effect of a security measure and discusses how taking both effects into account may help defenders choose better defense strategies.

The ability to signal one's security level to an attacker is also important for successful defense. The results from two alternative specifications of our model suggest that the absence of such signals makes systems with a low security level better off and those with a high security level worse off. Besides, lack of information attackers have about the security characteristics of each potential target reduces the incentives for individual firms to invest in security.

We use our findings to discuss various approaches to investments in security technology and make recommendations regarding security practices of individual firms as well as policy recommendations. In particular, we argue that the Annual Loss Expectancy (ALE) (Soo Hoo, 2000) and other widely adopted approaches to information security can severely underestimate positive effects of security investments, therefore leading to underinvestment in information security or misallocation of resources.

2. Related Work

Our research belongs to the field of economics of information security. Of all the issues within that broadly defined area, we are focusing on what economic research has to say about the best strategies for investing in security technologies. The advantage of the economic approach over traditional ones is that it recognizes and accounts for the presence of a strategic interaction between different parties involved.

The literature combining economic approach with information technology issues is vast. Clemons (1991) discusses the reasons why businesses have difficulty evaluating when to use information technology. Relevant to our research is his observation that some investments should be made to limit the possibility of future losses rather than to obtain long-term additional value. This notion applies perfectly to the case of security technology investments. When companies face environmental changes, a common scenario for information security, needed investments in information technology may be diverted if such changes are not foreseen. This is the effect Clemons called the "trap of the vanishing status quo".

Despite the amount of research making a case for wider use of economic approach to information security (Anderson, 2001, Gordon and Loeb, 2002a, Gordon et al., 2003, Rodewald, 2005, Schechter, 2005), little attention has been paid to those findings by security practitioners. An example in that regard is Gordon and Richardson (2004) which provides a comparative analysis of two traditional investment evaluation techniques, Return on Investment (ROI) and Net Present Value (NPV) and shows that the NPV approach is more applicable to computer security issues than ROI. Still , ROI is by far the most popular metric used, as documented by the 2005CSI/FBI Computer Crime and Security Survey (Gordon et al., 2005). This point is similar to one we make in our work, where we show how traditional investment evaluation techniques can greatly underestimate the effectiveness of a security solution by not considering the strategic nature of the problem and the interdependency between attackers' and defenders' actions.

Among other attempts to develop better techniques for evaluation investment in security, Geer (2005) introduces an alternative to traditional ROI formula suggesting to perform a cost-effective analysis, rather than a cost-benefit analysis, when costs and benefits are not commensurate. Purser (2005) proposes a modification to the ROI approach that would assign a monetary value to an increase or decrease in the risk resulting from an investment. According to that approach, a risk increase results in a lower ROI and vice versa. While the idea of considering a secondary effect of security investments resulting from a modified operational environment is similar to the one explored in our paper, our approach, based on game theory, is better suited to model such an interdependent behavior.

Other work in the game theory field that is related to ours includes Cavusoglu and Raghunathan (2004), which compares decision theory and game theory approaches in the context of the configuration of detection software. Although the subject of their paper is different from ours, the approach they follow is close to what we have done in terms of contrasting the results obtained by each of the two approaches.

Cavusoglu et al. (2004) and Cavusoglu et al. (2005) also propose game-theoretic models for evaluating security investments. However, their work focuses on specific security technologies, whereas we consider the more general problem of evaluating investments in computer security solutions.

The main purpose of investing in security is to defend against malicious attackers. Acquiring proper understanding of attackers' behavior is therefore a necessary step towards best security practices. Jonsson and Olovsson (1997) contributed to such understanding by performing an empirical study of attackers' behavior in a laboratory environment. While their work is descriptive in nature, we are able to use some aspects of their analysis as a starting point in setting up our model. In particular, they provide empirical evidence of several distinctive phases of an attack and hypothesize the presence of "behavioral" and "preventive" effects of security measures. We make a similar distinction in our model.

Schechter and Smith (2003) and Bier et al. (2005) model strategic defender-attacker interaction and show that a defender can influence the attacker's choice of targets by selecting its actions accordingly. An attack in their models is a one-time decision and the security characteristics of each target are assumed to be known. While the starting premise of our analysis is similar, we are able to move further in our analysis by studying the endogenous choice of effort made by attackers and examining the role of information about the target security level in the formulation of an optimal defense strategy.

Enders and Sandler (2004) discuss the strategic interaction between defenders and attackers in the context of anti-terrorism policies. They identify two separate mechanisms, the income effect and the substitution effect, through which a defensive measure may mitigate a specific security threat. While their approach is conceptually similar to ours, our results are more directly applicable to information security.

Other works directly related to the issues we are interested in include Kuhnreuther and Heal (2003) and Kearns and Ortiz (2004), which introduce the concept of an interdependent security game and show that the ultimate safety of each participant may depend in a complex way on the actions of the entire population. In the present work we have used some elements of interdependent security models by discussing the effect of a security measure on the strategic behavior of attackers.

3. The model

The model consists of N corporate networks that serve as targets for malicious attacks and an unspecified number of attackers all of whom are identical. An attacker targets one network at a time and puts a certain amount of effort, x, into attacks. We are interested in modeling intrusions which require a series of breaches into systems that belong to the targeted network in order to achieve the ultimate goal (such as to access critical data stored on an internal database).

As intrusion progresses, it leaves behind a stream of evidence detectable by the victim. Intrusion detection literature (Ning et. al., 2004; Valeur et al., 2004; Lee and Xiang, 2001; Wespi et al., 2000)

recognizes that the more suspicious events are observed, the more likely is the data correlation to result in alerts, which improves the detection success rate. A successful detection of an attack may in turn lead to a punishment being imposed on the attacker. With that in mind, we assume the attacker's cost, *C*, to be increasing and concave in the amount of effort spent, $\frac{\partial C(x)}{\partial x} > 0$, $\frac{\partial^2 C(x)}{\partial x^2} > 0$.[2]

If successful, an attack results in an intrusion. The expected benefit from an attack is $E(B(x)) = \pi(x) \cdot G$, where $\pi(x)$ is the probability of success given the amount of effort put into attacking a given target and *G* is the one-time payoff the attacker receives in the case of a successful intrusion. The size of that payoff is assumed to be the same for all targets.

Attackers are maximizing their expected net payoff from attacks. They do so by deciding how much effort to spend on each target. The rule for optimally choosing the amount of that effort, further referred to as the optimal stopping rule, can be expressed in at least two alternative ways. It is in an attacker's best interest to stop attack attempts when the expected net benefit of his continued effort is no longer positive, $ENB(x) \le 0$, or when the marginal benefit, *MB*, of effort no longer exceeds its marginal cost, *MC*. Naturally, both decisions are equivalent.[3]

Denote $\hat{x}$ the amount of effort that solves $MB = MC$. The attacker's expected net payoff after he expended *x* units of effort but achieved no success is

$$ENB(x) = \int_x^{\hat{x}} \rho(\tau)(G - C(\tau) + C(x))d\tau - (1 - \pi(\hat{x} - x))(C(\hat{x}) - C(x)), \quad (1)$$

where $\pi(\hat{x}) = \int_0^{\hat{x}} \rho(\tau)d\tau$ and $1 - \pi(\hat{x}) = \int_{\hat{x}}^{\infty} \rho(\tau)d\tau$.

Here, $\rho(x)$ denotes the conditional probability distribution function given no success upon spending effort *x*.

We proceed by choosing specific functional forms for benefit and cost. The marginal cost of effort is assumed to be given by $MC = \alpha_0 + \alpha_1 x_i$. We also assume attack attempts result in an intrusion with constant probability per unit effort,[4] making the amount of effort that leads to an intrusion an exponential random variable with rate parameter λ. The probability of success of each attack is therefore given by $\pi(x) = 1 - e^{-\lambda x}$. We find it more convenient to characterize the aforementioned exponential distribution

[2] Following the approach of Jonsson and Olovsson (1997), we make no distinction between effort and time. Thus the intensity of effort, or how much effort is exerted in a unit of time, in our model is assumed to be exogenous and constant. A more complex setup is saved for later work.

[3] The literature on optimal search (Cozzolino, 1972) confirms that fact.

[4] This assumption is commonly used in the relevant literature (Littlewood et al, 1993; Cavusoglu and Raghunathan, 2004; and others). Jonsson and Olovsson (1997) confirmed the validity of that assumption empirically.

by the scale parameter $\mu = \lambda^{-1}$. Since $\frac{d\pi}{d\mu} < 0$, one can think of μ as the security level of the system subject to attacks, with greater values of μ corresponding to better protected systems.

Due to the memorylessness property of the exponential distribution, $\rho(x)$ in the expressions above equals $1/\mu$ for any *x*.

3.1 Scenario 1 – One target

We start with the simplest case in which there is only one specific target the attacker is interested in, *N*=1. Its security level, μ, is common knowledge. The attacker's expected net benefit from attacking that target is

$$ENB(x) = G - \mu\alpha_0 - \mu\alpha_1 x - \mu^2\alpha_1(1 - e^{x/\mu}e^{-(G/\mu-\alpha_0)/\mu\alpha_1}) \text{ for any } x \geq 0. \quad (2)$$

It is easy to show that $ENB(x) = 0$ is solved by $\hat{x} = \frac{G - \mu\alpha_0}{\mu\alpha_1}$. Thus, we have the following proposition:

Proposition 1. The amount of effort, $\hat{x}$, an attacker optimally puts into breaching a system increases in the size of the payoff he receives in the case of intrusion, ($\frac{\partial\hat{x}}{\partial G} > 0$), decreases in the target's security level ($\frac{\partial\hat{x}}{\partial\mu} < 0$), and decreases in the cost of performing an attack ($\frac{\partial\hat{x}}{\partial\alpha_0} < 0$, $\frac{\partial\hat{x}}{\partial\alpha_1} < 0$).

Here is a numerical example to help illustrate attacker's decisions in this case:

Let *G*=1000, $\mu = 50$, so that $\pi(x) = 1 - e^{-0.02x}$, and $\alpha_0 = 10, \alpha_1 = 2$.

From either $ENB(x) = 0$ or $MB = MC$, $\hat{x} = \frac{G - \mu\alpha_0}{\mu\alpha_1} = 5$.

The following graphs may clarify the decision making process and the equivalence of the two approaches to solving the optimal stopping problem. Figure 1 is trivial and shows the marginal benefit and marginal cost of effort. Clearly, stopping at $\hat{x} = 5$ is in the attacker's best interest.

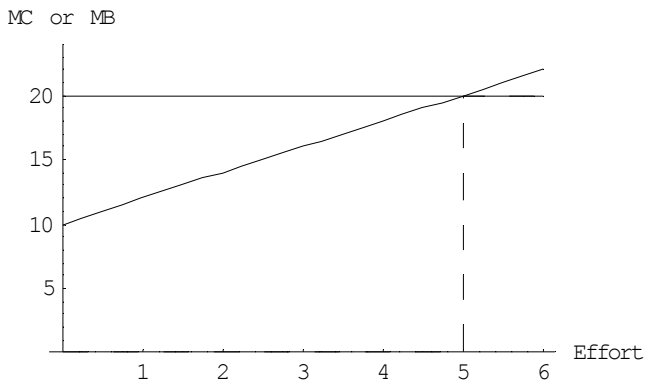

Figure 1. Marginal cost and marginal benefit of attacker's effort.

Figure 2 shows the overall expected net benefit that can be received from attacking the target as a function of effort that will be put in (as viewed before the start of attack attempts). Given the knowledge about the target's security parameter, μ, the attacker knows that the maximum value of the total expected net benefit is going to reach its maximum at $\hat{x} = 5$ ($ENB(\hat{x}) = 24.19$) and decrease afterwards. In other words, the attacker can make an advance commitment to putting in 5 units of effort (or less if he happens to succeed sooner).

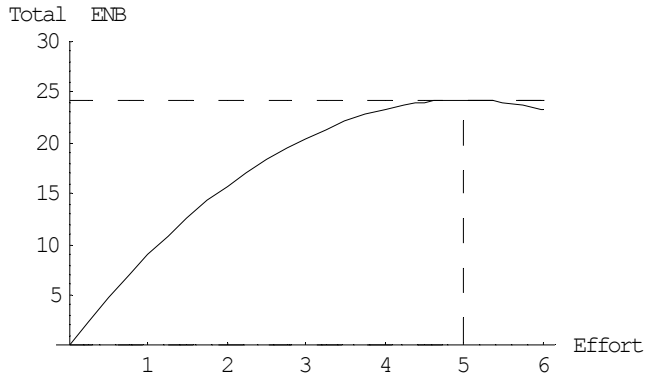

Figure 2. Expected net benefit from attacks as a function from future effort spent, as viewed before attacks start.

Still another way to present the stopping decision is through the residual expected net benefit from attacks after some effort has been already spent. See Figure 3. The greater the amount of past effort, the greater is the marginal cost of effort and the smaller is the net benefit the attacker still expects to derive from future effort. In this version of the model the attacker keeps trying until $ENB(x) = 0$.

We find this representation the most insightful of the three and will utilize it in the analysis of more advanced versions of the model.

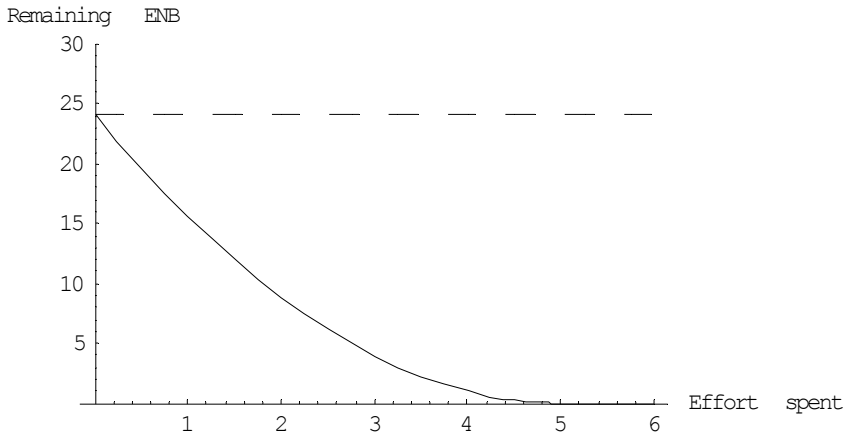

Figure 3. Expected net benefit of future effort as a function of effort spent.

3.2 Scenario 2 – Multiple identical targets

There are $N>1$ potential targets with the same security level, μ, which is common knowledge. Attackers are now able to stop working on one target and switch to another at any time. Switching to a different target involves some cost, which we interpret as the cost of effort put into the "learning phase" of an attack [5] in the context of the aforementioned Jonsson and Olovsson (1997). The size of the switching cost, C_S, is assumed to be the same for each target.

One major difference between this setup and the one discussed above is that there is now an outside opportunity present that has a certain value to an attacker. Therefore the attacker will make the decision to stop attacking one target and switch to another, randomly picked, target once his remaining expected net benefit from the current target gets smaller than the net benefit he expects to get if he switches. The optimal stopping rule for this case is therefore

$$ENB(x) = ENB(0) - C_S. \tag{3}$$

Proposition 2. The maximum amount of effort an attacker puts into attacking a target increases in the size of the switching cost, C_S. ($\frac{\partial \hat{x}}{\partial C_S} > 0$).

For $C_S = 5$ and the parameter values used above, (2) is solved by $\hat{x} = 0.5548$. See Figure 4 below.

[5] For example we may think of usual reconnaissance operations performed to gather information on potential targets like port scanning, OS and application fingerprinting, and so forth.

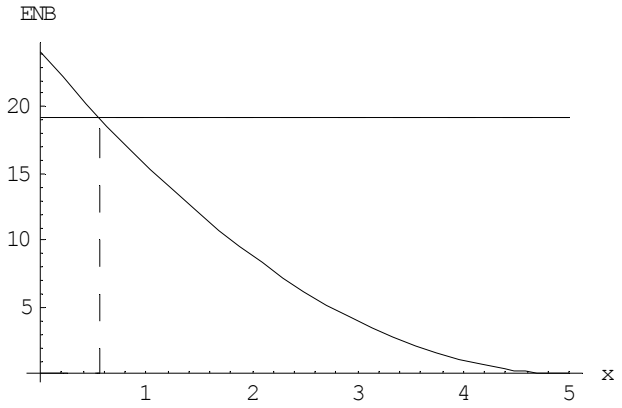

Figure 4. Optimal stopping decision in the presence of multiple targets and switching cost.

The downward sloping line on the graph above shows the expected net benefit from continuing attacks on the present target, $ENB(x)$. As discussed above, it decreases in the amount of effort already spent, $\frac{\partial ENB(x)}{\partial x} < 0$. The horizontal line represents $ENB(0) - C_S$. Once $ENB(x) \le ENB(0) - C_S$, the attacker is better off paying the one-time switching cost, C_S and switching to a different target.

Note that if $C_S = 0$, then $\hat{x} = 0$ trivially. Therefore in the absence of switching costs and presence of multiple identical targets attackers would be switching between targets all the time.

The analysis done insofar did not discuss decisions made by defenders. Those decisions are endogenized in the next modification of the model.

3.3 Scenario 3 – Heterogeneous targets

We now further advance our analysis by relaxing the assumption of target homogeneity. In the present setup, there are *(H+L)* targets present, *H* of which have a high security level and *L* have a low security level. We will further refer to such targets as being of *H-type* or *L-type*, respectively, where $\mu_H > \mu_L$. The type a target belongs to becomes known to attackers with certainty after the reconnaissance stage (hence upon paying the switching cost, C_S).

The expected net benefit received from a target of a known type is a modification of (2) obtained earlier:

$$ENB_i(x) = G - \mu_i\alpha_0 - \mu_i\alpha_1 x - \mu_i^{\,2}\alpha_1(1 - e^{x/\mu_i} e^{-(G/\mu_i - \alpha_0)/\mu_i\alpha_1}) \text{ for any } x \ge 0. \quad (4)$$

where *i = H, L*. Switching to a different, randomly chosen target involves cost C_S and gives the attacker an expected net benefit

$$\begin{aligned} &ENB_{random} = \eta ENB_H(0) + (1-\eta)ENB_L(0) = \\ &G - \eta\mu_H\alpha_0 - \eta\mu_H^2\alpha_1(1 - e^{-(G/\mu_H - \alpha_0)/\mu_H\alpha_1}) - (1-\eta)\mu_L\alpha_0 - (1-\eta)\mu_L^2\alpha_1(1 - e^{-(G/\mu_i - \alpha_0)/\mu_i\alpha_1}), \end{aligned} \quad (5)$$

where $\eta = \frac{H}{H+L}$ is the proportion of H-type systems in the population.

Applying the optimal stopping rule to this version of the model suggests that the attacker should continue putting effort into one target as long as $ENB_{random}(0) - ENB_i(x) \le C_S$ and switch to a different randomly chosen target when $ENB_{random}(0) - ENB_i(x) \ge C_S$, where *i = H, L* is the type of his present target. Thus, the effort after which it is optimal to switch to a different target is given by the solution to

$$ENB_{random}(0) - ENB_i(x) = C_S . \quad (6)$$

No explicit solution for (6) exists. However, using differentiating an implicit function we are able to get the result stated in the following proposition.

Proposition 3. Given the presence of targets of different security types and attackers' ability to determine the target type, the amount of effort optimally put by an attacker into a target decreases in its security level μ, $\frac{d\hat{x}}{d\mu} < 0$.

This important fact drives many results of our paper and has important implications for security practices. It suggests that every security solution affects the state of security through two distinct mechanisms. One is what we call the *direct* or *technical effect*, represented by the increased ability of a system to withstand intrusion attempts given the intensity of those attempts. The direct effect is commonly recognized by security practitioners. It can be shown that, when the probability of attack success is small, the direct effect is approximately proportional to the increase in the security parameter μ:

$$\frac{\partial \pi}{\partial \mu} \cdot \frac{\mu}{\pi} = \frac{x e^{-x/\mu}}{\mu(1-e^{-x/\mu})} \approx 1 .$$

Thus, according to the direct effect, a 10% increase in security level results in an approximately 10% reduction in the probability that each attack will result in a successful intrusion.

There is, however, another effect as well, which we call *indirect* or *behavioral effect*. Ceteris paribus, a more secure system is less appealing to attackers than a less secure one. Thus, a security enhancement performed at one system diverts attackers' effort away from it, and, since $\pi(\mu, x) = 1 - e^{-x/\mu}$, less effort on attackers' part translates into a lower probability of a security incident.

This fact deserves to be summarized in another proposition.

Proposition 4. Given the heterogeneity of target types and the presence of rational attackers able to determine a target's type, any security improvement causes more than proportional reduction in the probability of success of each individual attack.

The result in Proposition 4 can be confirmed by differentiation by the chain rule:

$$\frac{d\pi}{d\lambda} = \frac{\partial \pi}{\partial \lambda} + \frac{\partial \pi}{\partial \hat{x}}\frac{\partial \hat{x}}{\partial \lambda} > \frac{\partial \pi}{\partial \lambda} > 0 \qquad (7)$$

(+) (+) (+)

While the probability with which an individual attack results in an intrusion is qualitatively important, it is not directly observed by defenders. Instead, defenders are primarily concerned with the frequency with which security incidents ("intrusions" according to our terminology) occur. Moreover, the loss incurred by defenders per unit of time (ALE being one such example) is directly related to the number of security incidents per unit of time. Therefore we chose to discuss security enhancement solutions in the context of their effect on the frequency of intrusions.

The frequency of intrusions is the product of the probability that an attack results in an intrusion, $\pi_i(x_i) = 1 - e^{-\hat{x}_i/\mu_i}$, and the rate of attackers' arrival at a target. The arrival rate is the same across targets. It is proportional to the overall number of attackers, N_A, and inversely proportional to the number of potential targets, N_T and the average length of an attacker's stay on each target, τ. To determine τ, one needs to realize that an attacker leaves one target and starts looking for another if he either has successfully breached the system or feels the current target is no longer worth his continued effort. Once we know the solution to the optimal stopping condition for each type of system, $\hat{x}_i$, $i = H, L$, we can determine the average, or "expected", amount of effort an attacker spends on a system:

$$\tau_i = \int_0^{\hat{x}_i} x f(x) dx + \hat{x}_i e^{-\hat{x}_i/\mu_i} = \mu_i (1 - e^{-\hat{x}_i/\mu_i}) . \qquad (8)$$

Keep in mind that if switching cost is interpreted as the opportunity cost of the reconnaissance effort, then it has to be included in the calculation of the length of stay as τ_S. A sufficiently close approximation for it is $\tau_S = C_S / \alpha_0$. Thus, an attacker spends an average of $(\tau_S + \tau_H)$ units of effort on an H-type system and $(\tau_S + \tau_L)$ units of effort on an L-type system. Since effort in our model is equivalent to time, attackers return to the pool and start probing another target at $\tau = (\tau_S + \eta\tau_H + (1-\eta)\tau_L)$ intervals. Finally, the frequency of intrusions equals

$$\nu_i = \frac{N_A \cdot \pi(\mu_i, \hat{x}_i)}{N_T(\tau_S + \eta\tau_H + (1-\eta)\tau_L)} = \frac{N_A(1-e^{-\hat{x}_i/\mu_i})}{N_T\left(\tau_S + \eta\mu_H(1-e^{-\hat{x}_H/\mu_H}) + (1-\eta)\mu_L(1-e^{-\hat{x}_L/\mu_L})\right)}, \quad (9)$$

where $i = H, L$.

We are now able to state the effect of a security enhancing solution on the frequency of intrusions and therefore on the annual loss expectancy.

Proposition 5. Given the heterogeneity of target types and presence of rationally behaving attackers who are able to determine a target's type, any security enhancement causes more than proportional reduction in the frequency of security incidents and in the expected annual loss from attacks (ALE). The extent of that reduction, $\xi = \left|\frac{\nu_L - \nu_H}{\mu_H - \mu_L}\right|$, *is inversely related to the size of the switching cost,* $\frac{d\xi}{dC_S} < 0$.

The specifics of the indirect behavioral effect may become clearer from the numerical simulation results. For simulations, we use the same parameters as before, $G = 1000$, $\alpha_0 = 10$, $\alpha_1 = 2$, $C_S = 5$. The security parameters of the two target types are $\mu_H = 55$ and $\mu_L = 50$. There is an equal proportion of systems of each type, $\eta = 0.5$. The number of attackers and the total number of targets are both normalized to 1.

Initially, an L-type defender with $\mu_L = 50$ suffers intrusions with frequency $\nu_L = 0.0197$. They result from attackers arriving at his system at the rate of 0.916 per unit of time, staying no more than $\hat{x}_L = 1.089$, and each attack leads to a success with probability $\pi_L = 0.0215$.

Next, a security enhancement is considered that would change the system security parameter to $\mu_H = 55$. If the rate of attackers' arrival and individual attacker's effort were assumed to remain the same, then it would be reasonable to expect the frequency of breaches to decrease to $\nu = 0.0178$. In reality, however, due to the amount of effort being substantially reduced (from $\hat{x}_L = 1.089$ to $\hat{x}_H = 0.094$), the frequency of intrusions also sees a significant drop to $\nu_H = 0.0017$.

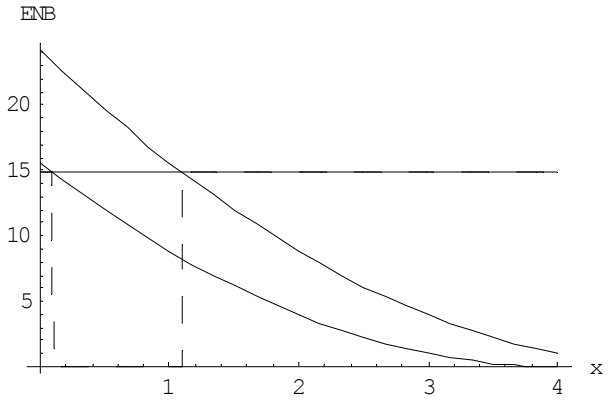

Figure 5. Difference in the optimal attacker's effort across target types.

Naturally, the outcome of a security enhancing measure depends on the distribution of target types and the size of the switching cost. Both effects are illustrated by the diagrams below.

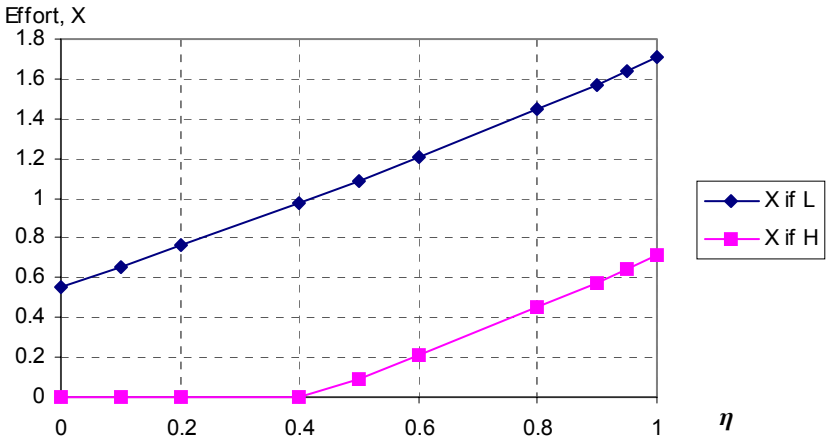

Figure 6. Optimal attacker's effort as a function of the type of the system and the composition of the population. $G = 1000$, $\alpha_0 = 10$, $\alpha_1 = 2$, $C_S = 5$, $\mu_H = 55$, $\mu_L = 50$. Parameter η denotes the proportion of H-type systems.

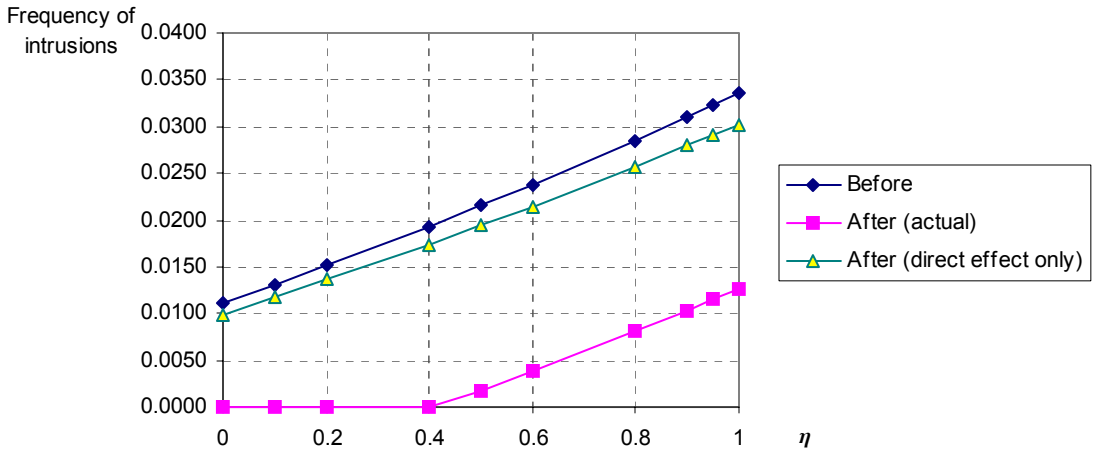

Figure 7. Direct only and overall effects of a security enhancement on the frequency of breaches, plotted against the proportion of H-type systems. $G = 1000$, $\alpha_0 = 10$, $\alpha_1 = 2$, $C_S = 5$, $\mu_H = 55$, $\mu_L = 50$.
The top curve, "Before", represents an L-type system before security enhancement. The middle one, "After (direct effect only)", shows the expected effect of upgrading to H-type, taking only the direct effect into account. The bottom line, "After (actual)", shows the frequency of intrusions that will actually occur as a result of an upgrade. It includes both the direct and the behavioral effect.

The size of direct and indirect effects of increased security on the frequency of intrusions and therefore the ALE can be seen from Figure 7. The top line represents the frequency of intrusion occurrences for an L-type system (before investment in security is made). The middle line shows the rate that is expected after its security level is raised from L to H if only the direct effect is accounted for. The bottom line is the intrusion rate after the security is raised, given the presence of both direct and behavioral effects. As the diagram indicates, the indirect effect a security enhancement may have on the frequency of breaches can substantially exceed the direct effect. It is also possible that an attacker will not find it worthwhile to spend any effort at all on an H-type target and will prefer to leave an H-type target immediately and look for an L-type target instead.

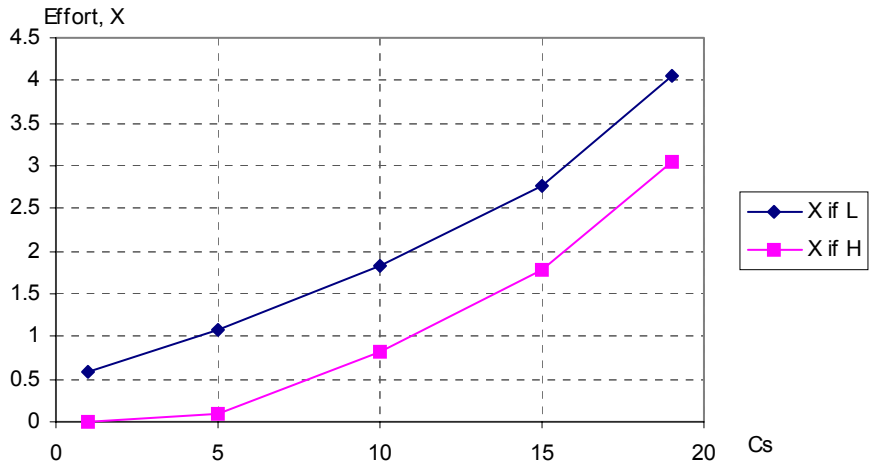

Figure 8. The effect of the switching cost C_s on the effort put by attackers into systems of each type. $G = 1000$, $\alpha_0 = 10$, $\alpha_1 = 2$, $\mu_H = 55$, $\mu_L = 50$, $\eta = 0.5$.

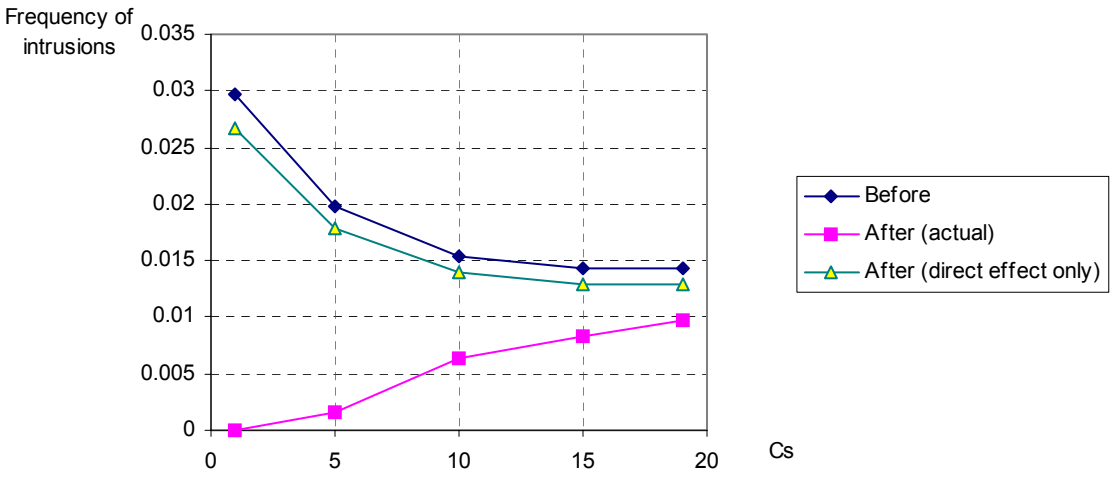

Figure 9. Frequency of intrusions before and after a security enhancement, plotted against the switching cost. $G = 1000$, $\alpha_0 = 10$, $\alpha_1 = 2$, $\mu_H = 55$, $\mu_L = 50$, $\eta = 0.5$.

The top curve, "Before", represents an L-type system before security enhancement. The middle one, "After (direct effect only)", shows the expected effect of upgrading to H-type, taking only the direct effect into account. The bottom line, "After (actual)", shows the frequency of intrusions that will actually occur as a result of an upgrade. It includes both the direct and the behavioral effect.

Figure 8 confirms the second result of Proposition 5, namely the positive correlation between the size of the switching cost and the amount of effort put into each attack on L-type and H-type systems, respectively.

Figure 9 allows us to elaborate on some specifics of the layered security approach and compartmentalized security architectures (Wells&Thrower, 2002). Those approaches to information security have gained popularity in the last several years and have proven to be superior to more traditional ones that rely on a security perimeter only. Still, existing security guidelines rarely make a distinction between investments into different security layers. If anything, there still seems to exist a tendency toward the "secure the perimeter" philosophy, according to which security investments should focus more on preventing intrusion attempts at early stages and therefore be concentrated on system areas that are closest to the external network. That means more attention is given to exterior layers of security than to interior ones. No clear consensus on the issue exists, however.

In the context of our model, the switching cost can represent the security of the exterior layer whereas the security parameter μ is a characteristic of the interior layer. As Figure 9 clearly indicates, strengthening outer echelons of defense may be less effective in reducing the frequency of intrusions than a combination of enhancing the security of the interior layer at the same time making that enhancement evident to attackers. It also shows that the higher the security level of a system, the more reason it has to signal its security level to attackers. As discussed above, such a signal may induce attackers to switch to other targets instead of continuing the intrusion attempts. This means that, at least for well-protected systems, it may be beneficial to have some means of implicit communication with attackers that would make them able to assess the target's security level. That is strikingly different from the aforementioned "secure the perimeter" approach and the traditional preference for opacity of protected networks that limit the amount of available information about deployed security measures.

Figure 9, however, has to be interpreted with care since the switching cost there is assumed to be the same across all systems. In reality, a security professional in charge of a specific system can control only the cost to an attacker of switching to his system but not to other systems. Second, as we try to translate this theoretical result into real world security practices, it is not completely clear what can serve as a credible signal of strong inner security and not undermine that security at the same time. Therefore, we are not ready to make any recommendations for security practices an individual firm may follow based on this result. A further exploration of this issue is among our priorities for future research.

The last result presents an interesting policy issue, however. It suggests that the same change in μ causes a bigger change in the frequency of intrusions when switching costs are smaller, that is, when it is easier for attackers to determine what type of a target they are dealing with. Therefore, incentives to invest in security are stronger when switching costs for all systems are small.

3.4 Scenario 4 – Targets with unknown security level

In this final version of the model we consider the case when the defender knows its security type but attackers do not. Thus, in this case we are dealing with incomplete asymmetric information. In this case, attackers base their behavior on their beliefs about the security level of a particular target.

The attacker's belief that he is dealing with an L-type target after effort *x* has not resulted in a break-in, $P(i = L \mid x)$, can be determined from the Bayes' theorem,

$$P(i = L \mid x) = \frac{P(x \mid i = L) \cdot P(i = L)}{P(x)} = \frac{(1-\eta)e^{-x/\mu_L}}{(1-\eta)e^{-x/\mu_L} + \eta e^{-x/\mu_H}} \tag{10}$$

Here, $P(x \mid i = L)$ is the probability that an L-type target will remain intact after effort *x* has been spent on it. $P(i = L)$ is derived from the known prior distribution of systems within the population, and $P(x)$ is the probability that an attack on a randomly chosen target will not lead to success after effort *x* has been spent. It equals the weighted sum of corresponding probabilities for the two types.

It is easy to show that $\frac{\partial P(i = L|x)}{\partial x} < 0$, which implies that the more effort a target is able to withstand, the less likely the target is of the L-type. It also means that, unlike all the cases discussed so far, the marginal benefit of effort in this case is decreasing in effort. To see this is so, recall that in the deterministic cases discussed above $MB(x \mid i) = e^{-x/\mu_i} \cdot G/\mu_i$, where *i* is the system type. Given the Bayesian mechanism of forming beliefs, the marginal benefit of effort is given by

$$MB_{Bayes}(x) = P(i = L \mid x)MB(x \mid L) + P(i = H \mid x)MB(x \mid H) = \frac{G\left((1-\eta)\mu_H e^{-2x/\mu_L} + \eta\mu_L e^{-2x/\mu_H}\right)}{\mu_H\mu_L((1-\eta)e^{-x/\mu_L} + \eta e^{-x/\mu_H})} \tag{11}$$

and $\frac{\partial MB_{Bayes}(x)}{\partial x} < 0$. Even though the attacker's perception of his marginal benefit is constantly evolving, the optimal stopping rule can still be applied. Clearly, $MB_H(x) \le MB_{Bayes}(x) \le MB_L(x)$, which in turn implies $\hat{x}_H \le \hat{x}_{Bayes} \le \hat{x}_L$. See the graph below.

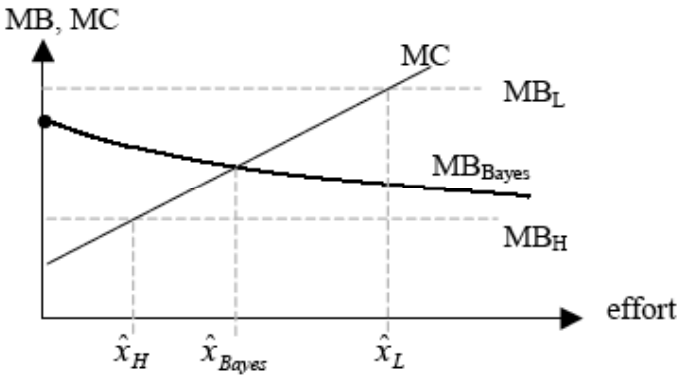

Figure 10. Marginal benefit and the solution to the attacker's optimal stopping problem for the complete information and the incomplete asymmetric information cases. The dot represents the prior probability of success from a randomly selected target. The expected marginal benefit of attacker's effort decreases because the more effort the attacker spends on a target with no success, the more he believes he is dealing with an H-type system.

In order to preserve consistency with the preceding analysis, the optimal stopping rule is applied by using the expected net benefit from future effort, which in this case equals

$$ENB_{Bayes}(x) = P(i = H \mid x)ENB_H(x) + P(i = L \mid x)ENB_L(x) =$$
$$\frac{\mu_H^2\alpha_1(z_H(x) - 1 + e^{-z_H(x)})\eta e^{-x/\mu_H} + \mu_L^2\alpha_1(z_L(x) - 1 + e^{-z_L(x)})(1-\eta)e^{-x/\mu_L}}{(1-\eta)e^{-x/\mu_L} + \eta e^{-x/\mu_H}} \quad (12)$$

where $z_i(x) = \frac{G/\mu_i - \alpha_0 - \alpha_1 x}{\mu_i \alpha_1}$, $i = H, L$. Once again, $ENB_H(x) \le ENB_{Bayes}(x) \le ENB_L(x)$ for any x and $\hat{x}_H \le \hat{x}_{Bayes} \le \hat{x}_L$.

The following two propositions summarize the results of our analysis of this case.

Proposition 6. When targets are heterogeneous and their type cannot be determined by attackers, the optimal amount of effort put forth by an attacker does not depend on the type of the system. For L-type targets, that amount is smaller than in the case when the target type is known to attackers (the complete information case) whereas for H-type targets it is greater than in the complete information case.

The expression for the frequency of intrusions at a given system in the presence of uncertainty about the target type is modified accordingly:

$$\nu_{i,uncert} = \frac{N_A(1 - e^{-\hat{x}_{i,Bayes}/\mu_i})}{N_T\left(\tau_S + \eta\mu_H(1 - e^{-\hat{x}_{i,Bayes}/\mu_H}) + (1-\eta)\mu_L(1 - e^{-\hat{x}_{i,Bayes}/\mu_L})\right)}. \quad (13)$$

It can be shown that $\nu_H < \nu_{i,uncert} < \nu_L$, where ν_H and ν_L are the frequencies of intrusions in the complete information case given by (9). As was pointed out earlier, the ALE at each system is directly related with the frequency with which intrusions occur. Therefore, we have the following proposition.

Proposition 7. Attackers' uncertainty about target types increases the annual loss expectancy of H-type systems and decreases it for L-type systems. The overall effect of the attackers' uncertainty about target types on the aggregate welfare is negative when the proportion of H-type systems in the population, η, is large and it is positive when η is small.

Numerical simulations were performed for the same parameters as before, $G = 1000$, $\alpha_0 = 10, \alpha_1 = 2$, $C_S = 5, \mu_H = 55$, and $\mu_L = 50$. Figure 11 confirms $\hat{x}_H \le \hat{x}_{Bayes} \le \hat{x}_L$.

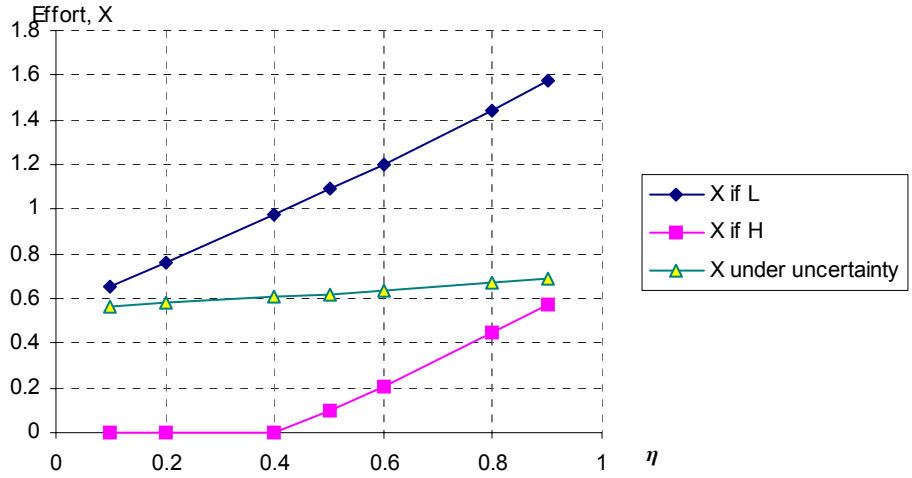

Figure 11. Effort optimally put by attackers into each system under different scenarios (as a function of the composition of the population). The top and the bottom lines represent $\hat{x}_L$ and $\hat{x}_H$, respectively, from the complete information case, as shown earlier in Figure 6.

Figure 12 shows the frequencies of intrusion in the incomplete asymmetric information case as a function of the relative proportion of H-type and L-type systems in the population. Since attackers now put the same amount of effort into attacking each system, the only difference in the frequency of intrusions across the two types (shown by the two lines in the middle) is attributed solely to the direct effect and is therefore proportional to the increase in the security parameter. This suggests that under incomplete asymmetric information the incentives to invest in security are substantially reduced.

The frequency of intrusions for each type in the certainty case is also provided for comparison. Clearly, attackers' uncertainty about target types benefits L-type systems and hurts H-type systems. The size of each type's welfare gain or loss from information asymmetry depends on the composition of the population. The fewer H-type systems there are, the more important it is for them to distinguish themselves from the rest of the population, therefore the loss in their welfare resulting from informational asymmetry will be greater, and vice versa.

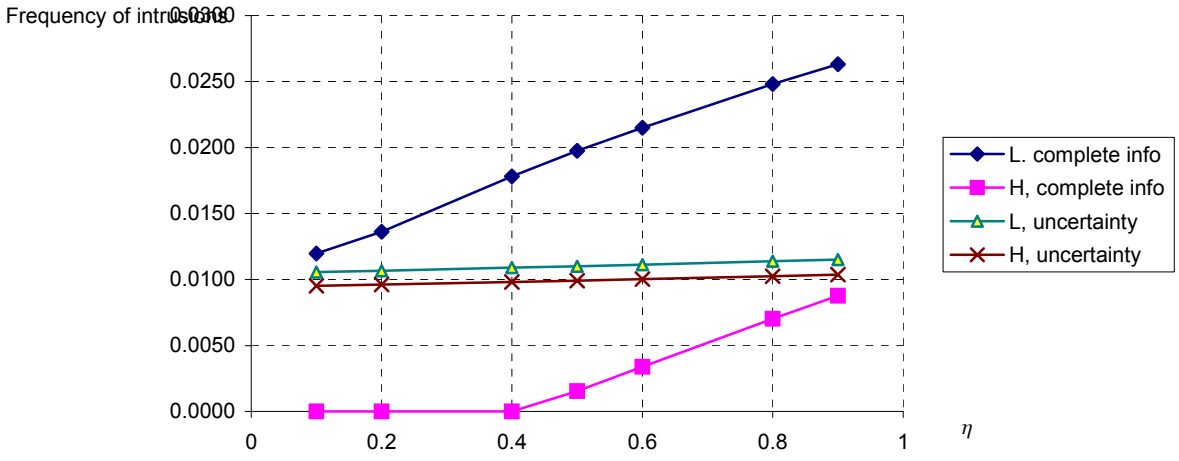

Figure 12. Frequency of intrusions under different scenarios (as a function of the composition of the population). The top and the bottom lines represent intrusion frequencies for L-type and H-type systems, respectively, from the complete information case, provided earlier in Figure 7. The two lines in the middle represent intrusion frequencies for the two types in the uncertain information case.

4. Discussion and Conclusion

Two cases were considered, one in which attackers were able to obtain information about each target's security level and the other in which the security level was known only to defenders. In both cases, attackers could choose from among multiple alternative targets. In the first, complete information case, attackers' optimal strategy is to put more effort into attacking systems with low security level than into systems with high security level. As a result, any increase in the defender's security level has two effects on the frequency of security incidents. One is the *direct effect* that is attributed to technical characteristics of a system and decreases the probability of success for a given attack effort. The second, indirect, or *behavioral effect* decreases the amount of effort an attacker puts into intrusion attempts, thus further decreasing the frequency of security incidents and the expected loss from attacks.

Traditional approaches to security practices tend to focus only on the direct effect of security solutions and overlook the behavioral one. Our analysis suggests that *the magnitude of the behavioral*

effect can greatly exceed that of the direct one. As a result, the benefit an individual system may receive from a security enhancement may be severely underestimated if the behavioral effect is not taken into consideration. This implies that some security investments worth making will not be made, leading to either underinvestment in security at the individual system level or, at the very least, to substantial misallocation of resources.

The magnitude of the behavioral effect depends on the security levels of available targets and the attackers ability to obtain information about potential targets' security characteristics and to rank those targets based on their attractiveness. In the complete information setup, attackers were able to determine a target's security level after some reconnaissance phase, with the effort an attacker had to spend to find out the target type represented by 'switching costs'. The analysis of that case suggests that the efficacy of security investment depends on the characteristic of the environment which we call "opacity". The greater the switching costs, the more "opaque" the environment is in the sense that it gets harder for attackers to determine the type of a target. As a result, the behavioral component of the overall effect gets weaker. If, on the contrary, the environment is "transparent" and determining the type is relatively easy, then the behavioral and the overall effects of a security investment will be stronger. From the practical perspective, this means that a given security solution will be more effective if potential attackers are aware of extra tools being deployed, and the incentives for firms to invest in security in that case will also be greater.

The above result was further confirmed by the second modification of our model, namely that of incomplete asymmetric information, in which attackers never know the target type with certainty. As a result, they treat every target the same, and the behavioral effect is not present. Once again, opacity penalizes better protected systems and favors those with weaker protection. As a result, systems whose security level is consistently low relative to the rest of the population will have preference for opacity over transparency since it gives them a better chance to disguise themselves as well-protected systems, thus reducing the amount of effort attackers put into attacking them. Well-protected systems, on the contrary, are better off in a transparent environment than in an opaque one and therefore have an incentive to signal their high security level to the attackers in order to separate themselves from less secure systems. This is consistent with existing theoretical research on economics of incomplete asymmetric information (Akerlof, 1970) that suggests that the ability to signal one's type (more transparency in the context of our model) benefit "high quality products" (well protected, or H-type systems) and penalize "low quality products" (poorly protected, or L-type systems).

The incomplete asymmetric information version of the model also allowed us to address the effect of informational issues and the distribution of system types on the expected welfare losses from attacks. When the proportion of H-type systems in the population is small, then the benefit each of them gets from

transparency (thus from identifying themselves as H-type systems) is substantial while L-type systems do not lose much since there are so many of them. When the proportion of H-type systems is large, the opposite is true. Interestingly, we did not find any effect of informational assumptions on the aggregate welfare since any benefit H-type systems as a group get from increased transparency was offset by a reduction in L-type systems welfare.

To a certain extent, the above discussion of opacity versus transparency is related to the debate surrounding the "security through obscurity" approach (Perens, 1998; Beale, 2000, Schneier, 2002, Swire, 2004). "Security through obscurity" is the term coined to denote technical security solutions the effectiveness of which is based on the secrecy of processes, protocols, or algorithms, which contrasts the basic rule of cryptography, the principles of the open source movement, and the open public disclosure of security vulnerabilities. In our model, a similar trade-off between disclosing information about the security level and keeping them secret exists. On both occasions, the resulting conclusions are controversial, although overall evidence seems to favor the rejection of the "security through obscurity" (or "security through opacity" in our case) approach. Today, when it comes to the information related to the security level of a corporate network, the practice is to keep it secret because its disclosure might favor attackers. The results of our study suggest, instead, that better protected organizations may be better off by letting their high level of security be evident to attackers. Such a strategy is the more justified the smaller is the share of those "H-type" organizations in the population of potential targets. Thus, our analysis further underscores the importance of an accurate and timely assessment an organization's security level relative to the rest of the population.

Our results also have practical implications for the case of layered security architecture. They suggest that poorly protected systems have more reason to invest in the exterior security layer, thus increasing opacity, whereas for well protected systems investment in the security of interior layers may be more beneficial, assuming they intend to maintain their advantage in protection level over the rest of the population.

As always, our analysis has some limitations. Most importantly, our model is static in the sense that it only examines the instantaneous effect of security enhancement assuming the distribution of system types and the rate of attackers' arrival at each target stayed the same. The outcome of any individual defender's decision will, however, also depend on what other defenders are doing at the same time. The interrelationship between the decisions of individual defenders and those of the rest of the population can be fully and properly understood only in a dynamic game-theoretic model. Therefore incorporating tools used in the analysis of interdependent security games is going to greatly enrich our understanding of security processes and practices.

4. References

Akerlof, G.A. (1970). The Market for 'Lemons': Quality Uncertainty and Market Mechanism. *Quarterly Journal of Economics*, 84(3):488-500.

Anderson, R.J. (2001). Why Information Security is Hard – an Economic Perspective. In *Proc. of the 17th Annual Computer Security Application Conference*, New Orleans, LA.

Avizienis, A., Laprie, J., and Randell, B. (2000). Fundamental Concepts of Dependability. In *Proc. of the Third Information Survivability Workshop*, Boston, MA.

Avizienis, A., Laprie, J., Randell, B., and Landwehr, C. (2004). Basic Concepts and Taxonomy of Dependable and Secure Computing. IEEE Transactions on Dependable and Secure Computing, 1(1):11-33.

Beale, J. (2000). "Security Through Obscurity" Ain't What They Think It Is. *Bastille Linux*. Available at http://www.bastille-linux.org/jay/obscurity-revisited.html.

Bier, V., Oliveros, S., and Samuelson, L. (2005). Choosing what to Protect: Strategic Defensive Allocation against an Unknown Attacker. *Journal of Public Economic Theory*, forthcoming.

Cavusoglu, H. and Raghunathan, S. (2004). Configuration of Detection Software: A Comparison of Decision and Game Theory Approaches. *Decision Analysis*, 1(3):131–148.

Cavusoglu, H., Mishra, B., and Raghunathan, S. (2004). A Model for Evaluating IT Security Investments. *Communications of the ACM*, 47(7):87-92.

Cavusoglu, H., Mishra, B., and Raghunathan, S. (2005). The Value of Intrusion Detection Systems in Information Technology Security Architecture. *Information Systems Research*, 16(1):28-46.

Clemons, E. K. (1991). Evaluation of Strategic Investments in Information Technology. *Communications of the ACM*, 34(l):22-36.

Cozzolino, J. M. (1972). Sequential search for an unknown number of objects of nonuniform size. *Operations research*. 20:293-308.

Curry, S. (2002). Bug Watch: Hacker Motivation. *Vnunet.com*. Available at http://www.vnunet.com/vnunet/news/2117147/bug-watch-hacker-motivation?vnu_lt=vnu_art_related_articles

Denning, D. E. (1990). Concerning Hackers Who Break into Computer Systems. *Phrack*, 3(32). Available at http://www.phrack.org/show.php?p=32&a=3.

Enders, W. and Sandler, T. (2004). What Do We Know About the Substitution Effect in Transnational Terrorism?. In Andrew Silke and G. Ilardi(eds.) *Researching Terrorism Trends, Achievments, Failures* (Frank Cass).

Geer, D. E. (2005). Making Choices to Show ROI. *Secure Business Quarterly*, 1(2). Available at http://sbq.com/sbq/rosi/sbq_rosi_making_choices.pdf.

Gordon, L. A. and Loeb, M. P. (2002a). The Economics of Information Security Investment. *ACM Transactions on Information and System Security*, 5(4):438-457.

Gordon, L. A., Loeb, M., and Lucyshyn, W. (2003). Information Security Expenditures and Real Options: A Wait-and-See Approach. *Computer Security Journal*, 19(2).

Gordon, L. A. and Loeb, M. (2005). *Managing Cybersecurity Resources: A Cost-Benefit Analysis*. McGraw-Hill.

Gordon, L. A., Loeb, M., Lucyshyn, W., and Richardson, R. (2005). *2005 CSI/FBI Computer Crime and Security Survey*. Computer Security Institute.

Gordon, L. A. and Richardson, R. (2004). The New Economics of Information Security. *Information Week*, 53-56. Available at http://www.banktech.com/aml/showArticle.jhtml?articleID=18901266.

Jajodia, S., J. Miller. (1993). Editor's preface. *Journal of Computer Security*. 16(4):43–53.

Jonsson, E. and Olovsson, T. (1997). A Quantitative Model of the Security Intrusion Process Based on Attacker Behavior. *IEEE Transactions on Software Engineering*, 23(4).

Kearns, M. and Ortiz, L. E. (2004). *Algorithms for Interdependent Security Games*. In S. Thrun, L. Saul, and B. Scholkopf, editors, *Advances in Neural Information Processing Systems* 16. MIT Press.

Kuhnreuther, H. and Heal, G. (2003). Interdependent Security. *The Journal of Risk and Uncertainty*, 26(2/3):231-249.

Lee, W. and Xiang, D. (2001). Information-theoretic measures for anomaly detection. In *Proc. of the IEEE Symposium on Security and Privacy*, Oakland, CA, USA.

Leeson P.T. and Coyne , C. J. (2006). The Economics of Computer Hacking. *Journal of Law, Economics and Policy*.

Littlewood, B., Brocklehurst, S., Fenton, N., Mellor, P., Page, S., Wright, D., Dobson, J., McDermid, J., and Gollmann, D. (1993). Towards Operational Measures of Computer Security. *Journal of Computer Security*, 2:211-229.

Liu, P., Zang, W. and Yu, M. (2005). Incentive-Based Modeling and Inference of Attacker Intent, Objectives, and

Strategies. *ACM Transactions on Information and System Security* , 8(1):78-118.
McDermott, J. (2005). Attack-Potential-Based Survivability Modeling for High-Consequence Systems. In *Proc. of the Third IEEE Int. Information Assurance Workshop*, Washington, DC.
Nicol, D. M., Sanders, W. H., and Trivedi, K. S. (2004). Model-Based Evaluation: From Dependability to Security. IEEE Transactions on Dependable and Secure Computing, 1(1):48-65.
Ning, P., Cui, Y., Reeves, D. S., and Xu, D. (2004). Techniques and Tools for Analyzing Intrusion Alerts. *ACM Transactions on Information and System Security*, (7)2:274–318.
NIST 800-30. (2002). Risk management guide for information technology systems. National Institute of Standards and Technology Special Publication, Gaithersburg, MD.
Ortalo, R., Deswarte, Y., and Kaâniche, M. (1999). Experiments with Quantitative Evaluation Tools for Monitoring Operational Security. *IEEE Transactions on Software Engineering*, 25(5):633-650.
Perens, B. (1998). Why Security-Through-Obscurity Won't Work. Slashdot. Available at http://slashdot.org/features/980720/0819202.shtml.
Purser, S. (2004). Improving the ROI of the Security Management Process. *Journal of Computers & Security*, 23 (7):542-546.
Rodewald, G. (2005). Aligning Information Security Investments with a Firm's Risk Tolerance. In *Proc. of the Information Security Curriculum Development (InfoSecCD) Conference '05*, Kennesaw, GA.
Schechter, S. E. and M. D. Smith (2003). How Much Security Is Enough To Stop a Thief? *The Seventh International Financial Cryptography Conference*, Gosier, Guadeloupe.
Schechter, S. E. (2004). *Computer Security Strength and Risk: A Quantitative Approach*. PhD thesis, Harvard University DEAS.
Schechter, S. E. (2005). Toward Econometric Models of the Security Risk from Remote Attack. *IEEE Security&Privacy*, 3(1):40-44.
Schneier, B. (2002). Secrecy, Security, and Obscurity. *Crypto-Gram*, Available at http://www.schneier.com/crypto-gram-0205.html#1.
Sieberg, D. (2005). Hackers shift focus to financial gain. CNN. Available at http://www.cnn.com/2005/TECH/internet/09/26/identity.hacker/
Soo Hoo, K. J. (2000). *How Much Is Enough? A Risk-Management Approach to Computer Security*. Doctoral dissertation, Stanford University School of Engineering.
Swire, P. P., (2004). A Model for When Disclosure Helps Security: What is Different About Computer and Network Security?. *Journal on Telecommunications and High Technology Law*, 2.
Valeur, F., Vigna, G., Kruegel, C. and Kemmerer, R. A. (2004). A Comprehensive Approach to Intrusion Detection Alert Correlation. *IEEE Transactions on Dependable and Secure Computing*, 1(3):146-169.
Wells, M., and Thrower, W. (2002). The Importance of Layered Security. *Symantec Corporation.* Available at http://enterprisesecurity.symantec.com/article.cfm?articleid=769&EID=0
Wespi, A., Debar, H., Dacier, M., and Nassehi, M. (2000). Fixed- vs. Variable-Length Patterns for Detecting Suspicious Process Behavior. *Journal of Computer Security*, 8(2/3).