A Model for Evaluating IT Security Investments

Assessing the return on investment has always been a sticking point for technology investments. Similar to IT productivity paradox [1], Return on Security Investment (ROSI) has become a controversial topic due to immense growth of e-businesses. Defining the value of security investments is challenging. However, it is clear that "security consumers will need to understand the variables that define ROSI and endure the discomfort of assigning dollar values to quantities that currently are extremely ill-defined" [12].

While calculating ROSI seems taxing, increasing possibility and scope of IT security breaches due to increasing interconnectivity makes it imperative. As the number of security breaches increases exponentially according to the CERT (see Table 1) so does their cost. The 2003 CSI/FBI Computer Crime and Security Survey revealed that 56% of respondents detected security breaches. *Information Week* and PricewaterhouseCoopers LLP estimated that computer viruses and hacking took a $1.6 trillion toll on the worldwide economy and $266 billion in the U.S. [5]. Security breaches have a significant impact on the market values of firms too. We have estimated that compromised firms, on average, lost approximately 2.1% of their market values within two days surrounding security breaches [3]. This translates to an average loss of $1.65 billion in market capitalization per incident. Moitra and Konda [10] found that as investment in security increases the survivability of firms from security breaches increases rapidly at first and then more slowly at higher levels of investment. Undoubtedly these figures point to the importance of more studies on the economics and management of IT security investments.

Fear, uncertainty, and doubt (FUD) strategy has been used for years to sell investments in security [1]. However, according to Earthlink security experts Lisa Ekman and Lisa Hoyt, "Crying wolf may get the first firewall, but over the long run, you need a more well-rounded perspective" [12]. Since diverse security tech-

By Huseyin Cavusoglu, Birendra Mishra, and Srinivasan Raghunathan

nologies with different capabilities are available in the marketplace, and many of these technologies substitute or complement each other, a more rational methodology is required to analyze security investments.

The second approach is based on the cost of deploying security. For example, the approach based on cost effectiveness of investments asks, "What is the most I can get for $X, given that I am going to spend $X?" This analysis is tractable as it does not seek to quantify the benefits of security investment and assumes it simply as an overhead cost. The primary limitation of this approach is that it does not help a company decide how much to invest in IT security.

A third approach uses indirect estimation of dollar value of costs associated with security breaches such as the loss in market value associated with security breach announcements as a proxy for potential loss. As stated earlier, the loss in market value could be of the order of billions of dollars. While loss estimates can be a useful starting point in convincing firms to deploy security technologies, they are less useful to firms in deciding which technology to deploy or how much to invest.

Year	Number
1996	2,573
1997	2,134
1998	3,734
1999	9,859
2000	21,576
2001	52,658
2002	82,094
2003	137,529

Table 1. The number of security incidents reported to CERT over the years.

The fourth approach uses the traditional risk or decision analysis framework. The idea is to identify the potential risks, expected losses and their likelihoods, and compute the expected loss. Longstaff et al. [9] propose Hierarchical Holographic Model (HHM) to assess security risks of IT. Gordon and Loeb [6] develop an economic model to determine the optimal level of investment in information security. Hoo [7] provides a decision analytic framework to evaluate different policies for IT security. Though intuitive, the decision analysis approach for evaluating IT security investment treats security technology as a black box. This technique does not provide managers any insights into how the different variables of an IT security infrastructure affect the risk, expected loss, and likelihood. For example, it cannot answer questions such as: How does the firewall quality affect the likelihood of a security breach or the expected loss? What is the trade-off between preventive controls, such as a firewall, and detective controls, such as an Intrusion Detection System (IDS).

We propose a comprehensive analytical model to evaluate security investment decisions. The model offers several benefits. First, it captures the individual technologies used in a typical IT security infrastructure. Consequently, managers can evaluate the interaction among different technologies and jointly decide on investments in multiple technologies. The model also facilitates understanding the different drivers of return on IT security investment, enabling managers to conduct sensitivity analysis of return with respect to these drivers. Finally, the model is useful to consumers in selecting the optimal configuration of security technologies and to developers in the design and pricing of security systems. It is worthwhile to note here that even security technology developers have started to incorporate cost and benefit factors in algorithms used by the technology. For example, Lee et al. [8] studied the problem of building cost-sensitive IDSs. Here, we motivate our model framework by discussing a typical IT infrastructure.

IT Security Infrastructure

IT security infrastructure is the foundation of a secure environment. It provides a comprehensive plan that protects the confidentiality, integrity, and availability of information resources. An IT security plan is composed of risk assessment, technology architecture, and policies and procedures.

Security risk assessment determines the level of security risk that exists within the organization. Quantitative risk analysis attempts to assign numeric values to the components of the risk (likelihood and potential damage) and to security controls (cost and effectiveness). The purpose of IT security infrastructure is to mitigate the risk up to a point where the marginal cost of implementing controls is equal to the value of additional savings from security incidents. The risk must be systematically assessed to effectively manage it [9], even during the software development phase [4]. The technology architecture establishes security controls. Preventive controls seek to develop a defensive shield around IT systems. Detective controls aim to detect security violations when they occur and are increasingly deployed because there is no absolute security that will completely prevent intrusions. IT security policies and procedures address a firm's response to a suspected security breach, and others such as security training and password policies. Responses can be executed automatically or manually. Monitoring determines if the system is really under attack and if so, the type and the extent of the attack by analyzing the log files and audit trails.

IT Architecture for Security

We emphasize that no single technology can provide absolute security for an organization. While many of the security controls are automated, manual controls cannot be avoided. A typical IT security architecture is composed of multiple layers of controls with different capabilities. The value of a security control at any level depends on other controls surrounding it as they may substitute or complement the control. The security controls are imprecise. For example, a firewall may prevent a legal user from entering the system (false positive), or allow an illegal user to enter the system (false negative). The costs associated with false positives and false negatives affect the value of a security control. The cost structures associated with security mechanisms at different levels can be sharply different. These characteristics require a firm simultaneously design all layers of controls.

The strategic nature of the security problem is another dimension that must be considered. Hackers often determine the weaknesses of a new security technology as soon as it is developed. Hackers attack IT systems that are vulnerable and without appropriate controls. They also seek the challenge of breaking well-protected systems since financial gain is not the sole motivation for them. In essence, they strategically choose their victims and actions. To compete, organizations should also act strategically when choosing controls and their capabilities. A good illustration of the strategic game played by hackers and the security experts in a firm is provided at www.msnbc.com/modules/hack_attack/hach.htm.

Figure 1 shows the details of a typical IT security architecture in which firewalls comprise the outermost layer, IDSs comprise the middle layer, and detailed investigations, typically manual in nature, comprise the bottom layer.

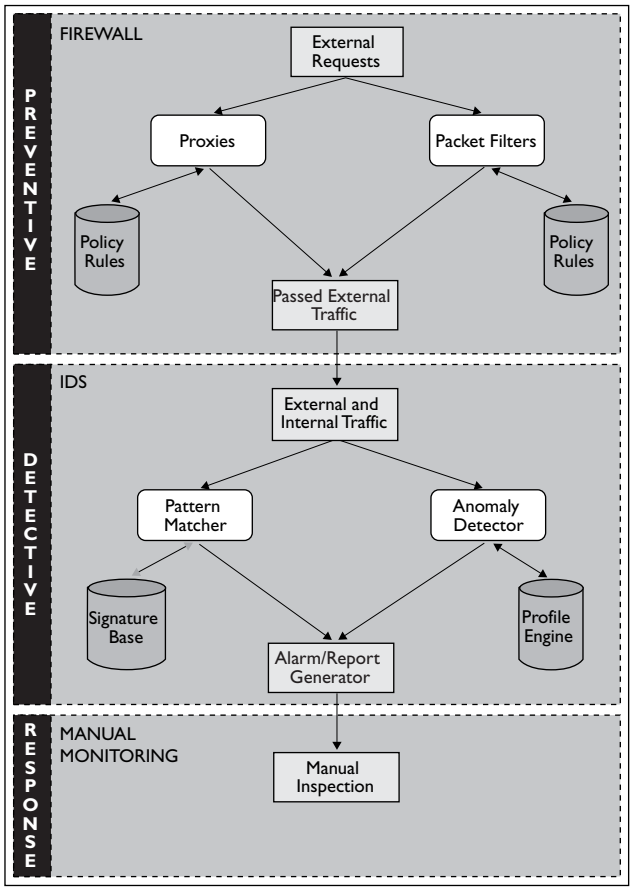

Figure 1. IT security infrastructure.

Firewalls can be configured to allow external traffic based on the service requested, the IP address, or user-ID. Firewalls use two types of mechanisms to filter the traffic. A packet-filtering mechanism performs filtering based on the set of rules in an access control list. The Application layer mechanism uses proxies. Each proxy checks both the header information and the service requested by the packet. The quality of a firewall is measured in terms of its effectiveness in stopping disallowed traffic and allowing valid traffic.

IDSs attempt to detect intrusions. An IDS generates an alarm when it detects something it considers suspicious or anomalous. IDSs use signature-based or anomaly detection approaches. Signature-based detection looks for events that match a predefined pattern of events, called a signature, associated with a known attack. Signature-based detectors are very effective in detecting common forms of attacks without generating an overwhelming number of false alarms. Anomaly detection identifies abnormal behavior using a "normal activity profile" that flags all system states that vary from the established profile in a statistically significant manner. Unfortunately, anomaly detection often produces a large number of false alarms. The quality of IDS is measured by its detection accuracy and false alarm rate.

A security analyst often examines the audit trails and log files to determine if there is a real attack and if so, the type and the extent of the attack. If an intrusion is confirmed by the investigation, appropriate corrective measures are undertaken to limit further damage and recover, if possible, the damage already incurred.

A Model for Evaluating IT Security Investments

Consider a firm interested in setting up its security infrastructure. First the firm estimates possible damages along with their likelihoods in the risk assess-

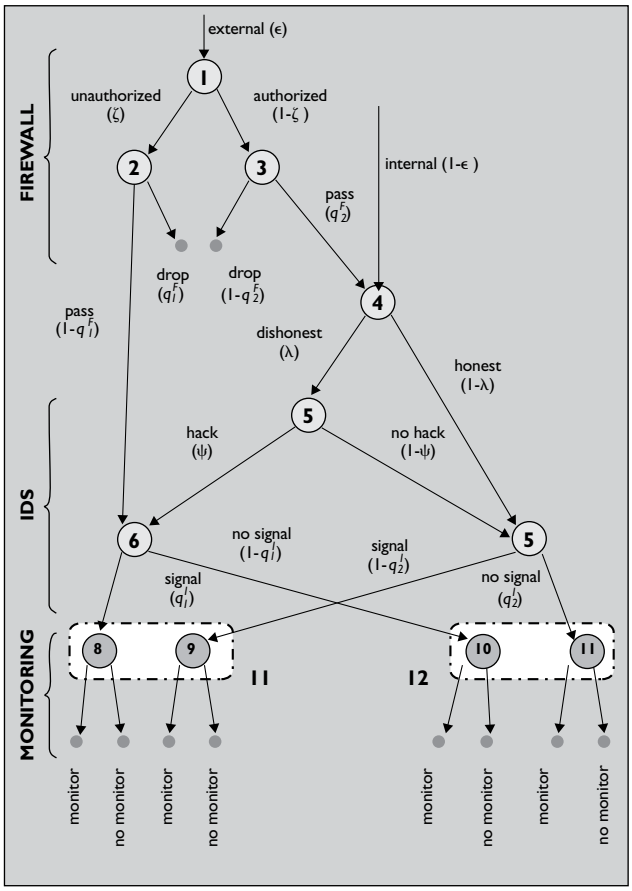

Figure 2. The game tree.

ment phase from which the firm calculates the expected damage d due to a security breach. Furthermore, let ε fraction of traffic come from external users and $(1 - \varepsilon)$ fraction originate from internal users. Assume ζ fraction of the external traffic comes from unauthorized users. These should be stopped at firewalls. We model the effectiveness of the firewall through two parameters q_1^F, the probability the firewall will stop an unauthorized user, and q_2^F, the probability that legitimate external traffic will pass through the firewall. If an authorized user is stopped at the firewall, the firm incurs a cost of σ.

All internal users are authorized users of the system. However, they as well as authorized external users can misuse the system by improperly accessing data or programs they are not authorized to use. In fact, many reports suggest the majority of hacking comes from legal users. A fraction λ of the legal users is assumed to be dishonest. Honest users never hack, but dishonest users may decide to hack. Previous studies have shown that incentives for intruders are usually not related only to financial gains. Hackers tend to be motivated by curiosity, self-esteem, vandalism, peer approval, public attention, technical prowess, and politics. We assume when the hacker breaks into the system, he or she gets a utility of μ.

IDSs are also not perfect. The parameter q_1^I denotes the probability that the IDS gives an alarm when an intrusion occurs. q_2^I is the probability that there is no signal when there is no intrusion.

The firm incurs a cost of c each time it monitors the audit trail of a user for a possible intrusion. Manual monitoring done by the firm may not detect intrusions with certainty. This imperfection of monitoring is captured by an effectiveness parameter α, the probability with which monitoring detects a true intrusion. If manual monitoring detects the intrusion, the firm recovers a fraction ϕ of the damage caused by the intruder without any additional cost. If an intrusion is detected, the hacker incurs a penalty. The penalty is composed of two components: A fixed penalty β and a variable penalty proportional to the expected amount of damage, γd. Hence, larger damage results in a larger penalty.

In order to evaluate investment in security, the firm determines the optimal frequencies of manual monitoring—when there is and is not a signal from the IDS—for the given firewall and IDS configurations. The firm can then determine its cost based on these optimal frequencies as well as the cost of deploying the firewall and IDS technologies. Similarly the firm can determine the optimal response and firm's cost when there is no IDS and firewall. The cost difference (that is, cost when there is and is not firewall and IDS) represents the value of firewall and IDS to the firm. Managers should then evaluate and choose the technology that maximizes the cost savings relative to the cost of the technology.

Strategic Investment Decisions Using Game Theory

Game theory [11] is used to analyze problems in which the payoffs to players depend on the interaction between players' strategies. For example, in the IT security investment problem, the firm and the hackers are players. The firm's payoff from security investment depends on the extent of hacking it is subjected to. The hacker's payoff from hacking depends on the likelihood he or she will be caught. Thus, the likelihood of the firm getting hacked depends on the likelihood the hacker will be caught, which, in turn, depends on the level of investment the firm makes in IT security. The first step in using game theory to analyze such strategic interactions among players is to develop a game tree that depicts the strategies of players. Figure 2 depicts the game tree for the IT security problem.

The game starts with nature selecting the type of

traffic to the system, which can be external (node 1) with probability ε or internal (node 4) with probability $(1 - \varepsilon)$. Node 2 represents the external hacker whereas node 3 is a legal external user. Node 5 characterizes the authorized access to an internal network. A dishonest user can take two actions: hack or do not hack. If the hacker decides to hack, the game moves to node 6, otherwise to node 7. IDS is captured by node 6. The firm makes decisions about whether to monitor or not based on the state (the signal or the no signal) it is in. The firm must make decisions without knowing exactly which node the game has reached. However, it can determine the probability of intrusion in the signal and no signal states using Bayes Rule as illustrated in Box 1.

Box 1. Bayes Rule.

$$\eta_1 = \text{P(intrusion}|\text{signal)} = \frac{\text{P(signal}|\text{intrusion)P(intrusion)}}{\text{P(signal}|\text{intrusion)P(intrusion)} + \text{P(signal}|\text{no-intrusion)P(no-intrusion)}}$$

and

$$\eta_2 = \text{P(intrusion}|\text{no-signal)} = \frac{\text{P(no-signal}|\text{intrusion)P(intrusion)}}{\text{P(no-signal}|\text{intrusion)P(intrusion)} + \text{P(no-signal}|\text{no-intrusion)P(no-intrusion)}}$$

Decision variables for the firm are the probability of monitoring given there is a signal from IDS, ρ_1, and the probability of monitoring given that there is no signal from IDS, ρ_2. The firm maximizes its expected payoff shown below with respect to these decision variables.

$$F = [P(signal)F_S + P(nosignal)F_N + P(drop)\sigma] \quad (1)$$

where $F_S = \{-\rho_1 c - \eta_1(1 - \rho_1)d - \eta_1\rho_1[(1 - \alpha)d + \alpha(1 - \phi)d]\}$ and

$F_N = \{-\rho_2 c - \eta_2(1 - \rho_2)d - \eta_2\rho_2[(1 - \alpha)d + \alpha(1 - \phi)d]\}$ are payoffs for signal and no signal states, respectively, each consisting of the loss from undetected intrusion, the loss from detected intrusion (such as the unrecoverable portion of the loss even if the intrusion is detected), and the monitoring cost.

The hacker's expected payoff includes the expected utility from the intrusion and expected cost if the intrusion is detected. The hacker, at the same time, maximizes its payoff function shown here with respect to his or her decision variable ψ.

$$H = P(hacking)(Benefit - Cost) = \psi\mu - \psi\alpha(\beta + \gamma d)[\rho_1 q_1^I + \rho_2(1 - q_1^I)] \quad (2)$$

For a given set of parameters, the solution of the game gives unique values for firm's and well as hacker's decision variables. We should note that the solution to the game involves maximization of a polynomial function. This is done by equating the first derivative of the function with respect to each decision variable to zero. At this point, firm knows how much manual monitoring it should implement in order to minimize the total cost of security. We illustrate this approach using a numerical example with parameters as follows: $\mu = 600$, $\alpha = 0.2$, $\beta = 5000$, $d = 1000000$, $\gamma = 0.01$, $q_1^I = 0.9$, $q_2^I = 0.9$, $q_1^F = 0.9$, $q_2^F = 0.9$, $\lambda = 0.5$, $\phi = 0.8$, $\varepsilon = 0.001$, $\sigma = 20$, $c = 5$. These numerical values are for a hypothetical scenario and are used for illustration purposes only. The firm that employs our model should estimate these parameters using any of the risk assessment techniques discussed earlier. The model predicts an optimal response of $\rho_1 = 0.56$ and $\rho_2 = 0$, meaning that the firm will monitor 56% of the cases for which it gets a signal from IDS and will not monitor no-signal cases at all.

Choice of Security Technology

Step 1. Select security technologies for deployment consideration.
Step 2. Collect data about quality parameters q_1^I, q_2^I, q_1^F, q_2^F, and costs of these technologies.
Step 3. Estimate hacker and firm specific parameters such as d, g, l, f, e, s, c, m, a, and b.
Step 4. Solve the model for each of the technology and determine the cost savings as (firm's cost with the technology—firm's cost without any security technology).
Step 5. Choose the technology that yields the maximum savings.

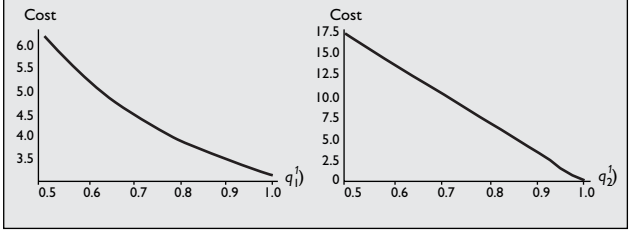

Figure 3. The effect of quality parameters of IDS on cost of security (a) q_1^I (b) q_2^I

Box 2. Choice of Security Technology

Step 1. Select security technologies for deployment consideration.
Step 2. Collect data about quality parameters q_1^I, q_2^I, q_1^F, q_2^F, and costs of these technologies.
Step 3. Estimate hacker and firm specific parameters such as d, γ, λ, ϕ, ε, σ, c, μ, α, and β.
Step 4. Solve the model for each of the technology and determine the cost savings as (firm's cost with the technology minus firm's cost without any security technology).
Step 5. Choose the technology that yields the maximum savings

Applications of the Model

A firm can apply this model in a variety of ways to evaluate investments in IT security. For example, the firm can find the procedure given in Box 2 useful in choosing a specific security technology.

Our model is also useful for understanding how the different parameters affect the optimal investment as well as cost. That is, it can be used very effectively as a what-if analysis tool to explore different options. There are two reasons why the what-if analysis capability is required in evaluating security technologies. First, the parameters estimated in the risk assessment phase, such as the expected damage caused by a security breach, are simply estimates. The uncertainty associated with these estimates can be analyzed by solving the model with various estimates in order to understand how estimation errors affect a firm's decisions. Second, the security products can often be customized or configured by the firm. The configuration process changes the quality parameters. For example, a loose configuration for an IDS may result in a large number of undetected intrusions (that is, a low q_1^I) and also a low number of false alarms (that is, high q_2^I). A tight configuration may result in a high q_1^I and a low q_2^I.

Since the cost savings realized by the firm from a security technology depends on the value of the quality parameters, the firm must analyze the value of different configuration options. Our model can be used to analyze how the various quality parameters affect firm's cost. For instance, we ran our model for the previously mentioned numerical example using different values for q_1^I and q_2^I. Figure 3 shows how the firm's cost varies with the quality of IDS technology in this model. It illustrates that as the quality of IDS increases in either dimension, the cost of security decreases. By this way, a firm can assess the marginal effect of decrease or increase of one parameter on total cost.

Conclusion

IT security management is a demanding task. Assessing the value of security technologies is essential to manage IT security effectively. However, the lack of a comprehensive model that incorporates the specific features of IT security technologies has prevented firms from applying rigorous quantitative techniques to make security investment decisions. The current set of tools such as risk analysis and cost effectiveness analysis work with very high-level aggregate data, so these tools are of limited value in an IT security setting. We have proposed a comprehensive model to analyze IT security investment problems that overcome some of these limitations. We used this model to derive insights into the value of technologies. C

References

1. Berinato, S. Finally, a real return on security spending. *CIO Magazine* (Feb. 15, 2002).
2. Brynjofsson, E. The information technology and productivity paradox. *Commun. ACM 36,* 12 (Dec. 1993), 66–77.
3. Cavusoglu, H., Mishra, B., and Raghunathan, S. The effect of Internet security breach announcements on shareholder wealth: Capital market reactions for breached firms and Internet security developers. *International J. of Electronic Commerce.* Forthcoming.
4. Collofello, J. Software Development Risk Management, 2000; www.eas.asu.edu/~riskmgmt/
5. Denning, D. Reflections on cyberweapons controls. *Computer Security J. 16*, 4 (2000), 43–53.
6. Gordon, L. and Loeb, M. The economics of information security investment. *ACM Trans. IS Security 5*, 4 (Nov. 2002), 438–457.
7. Hoo, K.J.S. How much is enough? A risk management approach to computer security. Ph.D. Dissertation, Stanford University, 2000.
8. Lee, W., Fan, W., Miller, M., Stolfo, S., and Zadok, E. Toward cost-sensitive modeling for intrusion detection and response. *J. Computer Security* (2001).
9. Longstaff, T., Chittister, C., Pethia, R. and Haimes, Y. Are we forgetting the risk of information technology. *IEEE Computer* (Dec. 2000).
10. Moitra, S. and Konda, S. The survivability of network systems: An empirical analysis. Carnegie Mellon Software Engineering Institute. Technical Report, CMU/SEI-2000-TR-021.
11. Rasmusen, E. *Games and Information.* Blackwell Publishers, 1998.
12. *Secure Business Quarterly.* Issue on Return on Security Investment (Q4, 2001).

Huseyin Cavusoglu (huseyin@tulane.edu) is an assistant professor at Tulane University, New Orleans, LA.
Birendra Mishra (barry.mishra@ucr.edu) is an assistant professor at The University of California at Riverside.
Srinivasan Raghunathan (sraghu@utdallas.edu) is an associate professor of MIS at The University of Texas at Dallas, Richardson, TX.